

Uniform Pricing in US Retail Chains

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

May 2, 2019

Abstract

We show that most US food, drugstore, and mass merchandise chains charge nearly-uniform prices across stores, despite wide variation in consumer demographics and competition. Demand estimates reveal substantial within-chain variation in price elasticities and suggest that the median chain sacrifices \$16m of annual profit relative to a benchmark of optimal prices. In contrast, differences in average prices between chains are broadly consistent with the optimal benchmark. We discuss a range of explanations for nearly-uniform pricing, highlighting managerial inertia and brand-image concerns as mechanisms frequently mentioned by industry participants. Relative to our optimal benchmark, uniform pricing may significantly increase the prices paid by poorer households relative to the rich, dampen the response of prices to local economic shocks, alter the analysis of mergers in antitrust, and shift the incidence of intra-national trade costs.

*E-mail: sdellavi@berkeley.edu, gentzkow@stanford.edu. We thank Nicholas Bloom, Liran Einav, Benjamin Handel, Mitch Hoffman, Ali Hortacsu, Emir Kamenica, Kei Kawai, Carl Mela, Emi Nakamura, Peter Rossi, Stephen Seiler, Steven Tadelis, Sofia Villas-Boas. We thank seminar participants at Columbia University, MIT (Sloan), New York University, the San Francisco Federal Reserve, Stanford University (GSB), the University of Bonn, the University of Chicago (Department, Booth, Harris), UC Berkeley, UCLA, the University of Toronto, and conference participants at the NBER Summer Institute, the SITE Conference in Psychology and Economics, and the Berkeley-Paris conference in Organizational Economics for helpful comments. We thank Angie Acquatella, Sahil Chinoy, Bryan Chu, Johannes Hermlé, Christopher Lim, Ammar Mahran, Akshay Rao, Sebastian Schaubé, Avner Shlain, Patricia Sun, and Brian Wheaton for outstanding research assistance. Gentzkow acknowledges funding from the Stanford Institute for Economic Policy Research (SIEPR). Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilt Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

The adjustment of local retail prices to local economic conditions is central to a range of economic policy questions. Differences in local retail prices across poor and rich areas may exacerbate or moderate real income inequality (Allcott et al., 2019a). The response of local prices to consumer demand is a key input to understanding business cycles (Stroebel and Vavra, forthcoming). Standard antitrust analysis of retail mergers depends critically on the way local prices respond to changes in local competitive conditions (Federal Trade Commission, 2010). Variation in prices as a function of intra-national trade costs determines the distribution of gains from globalization (Atkin and Donaldson 2015). Analysis in all these areas typically starts 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 food, drugstore, and mass merchandise chains in fact set uniform or nearly-uniform prices across their stores, resulting in a significant loss of profits. This fact echoes uniform pricing “puzzles” in markets such as soft drinks (McMillan, 2007), movie tickets (Orbach and Einav, 2007), rental cars (Cho and Rust, 2010), and online music (Shiller and Waldfogel, 2011), but is distinct in that prices are fixed across separate markets, rather than across multiple goods in one market. We also show that limited within-chain variation applies not just to prices but to product assortment as well. These findings have important implications for the policy areas mentioned above.

Our main analysis is based on store-level scanner data from the Nielsen-Kilts retail panel for sales of 1,365 products in 9,415 food stores. In addition, we extend some results to a sample of 9,977 drugstores and 3,288 mass merchandise stores from the same data source, and also consider expanded product sets with as many as 40,000 products. We use the standard price measure in these data, defined to be the ratio of weekly revenue to weekly units sold.

Our first set of results in Section 3 highlights the extent of uniform pricing. The price variation within chains is small in absolute terms and far smaller than the variation between stores in different chains. This is true despite substantial variation in consumer demographics and levels of competition within chains. Consumer income per capita ranges from \$23,100 at the average chain’s 10th-percentile store to \$41,600 at the average chain’s 90th-percentile store, and the number of competing stores within 10 km varies from 0.7 at the 10th-percentile store to 8.9 at the 90th-percentile store. Prices are highly similar within chains even for store pairs that face very different income levels and are in geographically separated markets. A corollary is that the relationship between prices and consumer income within chains is limited: prices increase by only 0.47 log points (≈ 0.47 percent) for each \$10,000 increase in the income of local consumers. In contrast, the price-income relationship between chains is an order of magnitude larger. Chain average prices

increase instead by 4.2 log points (≈ 4.2 percent) per \$10,000 of chain average income. These results are similar for various alternative sets of products, including store brands, high-revenue products, and low-revenue products. We find similar evidence of nearly-uniform within-chain pricing for drug and mass-merchandise chains.

We then show that the way prices are measured in the Nielsen data means that the degree of uniform pricing is likely even greater than these results would suggest. If not all consumers pay the same price within a given week, the weekly ratio of revenue to units sold will yield the quantity-weighted average price. For example, even if posted prices are uniform across stores on any given day of the week, stores facing more elastic demand (e.g., lower income) sell a larger share of units on days with relatively low prices, leading the weekly average price to be lower in such stores. Thus, compositional differences can create apparent correlation between measured prices and income, even if posted prices are actually uniform.¹ We assess the importance of this compositional effect using data from a major grocer studied in Gopinath et al. (2011). These data provide additional information that allows us to adjust for compositional effects, and once we do so the price-income relationship is essentially zero.

Our second set of results shows that the pattern of chain-level uniformity extends beyond pricing to other dimensions including product assortment. We define an assortment index based on the average national-level prices of the products sold by a store. This is high when the store carries relatively expensive products but does not depend on the prices charged by the store itself. We show that both overall variation in this index and its correlation with income are minimal within chains and much larger between chains. We find similar, though less extreme, patterns for the share of products that are organic, the share of generic products, and the number of products carried.

Our third set of results compares the prices we observe to an optimal benchmark based on a simple constant-elasticity model of demand. We address price endogeneity using a novel instrumental variables strategy broadly related to those in Hausman (1996) and Nevo (2001). The model fits the data well, with an observed relationship between weekly log quantity and weekly log price very close to linear. The store-level estimates of elasticity vary in an intuitive way with store-level demographics and competition.

Our model implies that the ratio of optimal price to marginal cost for a store with elasticity η_s is $\eta_s / (1 + \eta_s)$. Assuming no variation in marginal costs across stores, prices at stores with elasticities

¹Chevalier and Kashyap (2019) and Coibion et al. (2015) point out that average prices paid at the annual or market level may respond to macroeconomic shocks even when posted prices are constant. To the best of our knowledge, we are the first to emphasize that the same force affects the weekly average prices commonly available in scanner data. Indeed, both Chevalier and Kashyap (2019) and Coibion et al. (2015) treat the weekly average price as equivalent to the posted price. Campbell and Eden (2014) and Cavallo (2018b) make a related point, noting that aggregation to the weekly level can cause the frequency of price changes to be overstated.

in the 90th percentile within the typical chain ($\eta_s = -2.81$) should be 17 percent higher than at stores with elasticities in the 10th percentile ($\eta_s = -2.22$). However, observed prices are on average only 0.4 percent higher.² To formally test the model’s predictions, we regress log prices on the term $\log[\eta_s/(1 + \eta_s)]$, instrumented with store income. This yields a between-chain coefficient for food chains of 0.83 (s.e. 0.23), close to the value of 1 predicted by the model. The within-chain coefficient is an order of magnitude smaller, at 0.09 (s.e. 0.02), and the compositional issues discussed above suggest this is likely an overestimate. Our estimated model implies that the median chain could increase annual profits by \$16.1m (1.6 percent of revenue) by moving to optimal flexible pricing.

We consider a number of potential threats to the validity of our model. First, it abstracts from variation in marginal costs across stores. Stroebel and Vavra (forthcoming) present a range of evidence suggesting that such variation is likely to be small, and this is supported by our analysis of the major grocer’s data. We argue that any remaining cost variation is likely to work against what we find. Second, our main analysis treats demand as separable across products. Cross-product substitution could lead us to overstate the relevant elasticities as consumers substitute among products, or to understate them as consumers substitute on the store-choice margin as in Thomassen et al. (2017). To partially address this concern, we show that estimated elasticities are similar when we aggregate prices and quantities to the product category level. Third, prices and promotions are often determined jointly by retailers and manufacturers (Anderson et al., 2017). The fact that our results are similar for store brands provides evidence against the view that constraints imposed by manufacturers are a key driver of our results. Fourth, the products in our main analysis account for a minority of stores’ total sales. While our robustness and availability analyses expand the set of products along a number of dimensions, the results could be different for the products we exclude, as well as for other retail chains not included in our analysis.

Finally, and perhaps most importantly, our baseline results focus on short-run weekly elasticities that need not be equal to the longer-run elasticities relevant to setting a store’s average price level. Long-run elasticities could be smaller (due to consumer stockpiling as in Hendel and Nevo, 2006) or larger (due to search frictions). We address this in two ways. First, we show that the results are similar using prices and quantities aggregated to the quarterly level. Second, we present an event-study analysis of long-run price changes driven by chain mergers. We consider 114 stores that changed ownership from one chain to another in our sample. We show that these stores switch sharply from tracking the prices of their former chain to tracking the prices of their new chain. This provides a validation of our uniform pricing result, as well as a natural source of long-run

²For observed prices, we calculate this by selecting stores that have elasticities within 0.05 of the 10th and 90th elasticity percentiles in each chain. We compute the within-chain difference in average prices for stores near the 90th percentile versus stores near the 10th percentile, and then take an equal-weighted average across retailers.

price variation. Products whose prices were higher in the acquiring chain persistently increase; products whose prices were lower persistently fall. We show that long-run elasticities computed using these changes are highly correlated with our estimated short-run elasticities for the same stores. Nevertheless, we cannot rule out meaningful differences between short-run and long-run elasticities, nor can we rule out other sources of bias in our elasticity estimates; any such bias could have a significant impact on our conclusions about the optimality of uniform pricing.

The next section discusses potential explanations for uniform pricing. Conclusively identifying the correct explanation(s) is beyond the scope of this paper. However, discussions and interviews with chain managers, consultants, and industry analysts suggest two leading explanations. The first is *managerial inertia*, encompassing both agency frictions and behavioral factors that prevent firms from implementing optimal policies, even though the benefits exceed the economic costs traditionally defined. The second is *brand image concerns*, encompassing various mechanisms by which varying prices across stores would lead to negative reactions from consumers that could depress demand for a chain in the long run. We discuss several reasons that lead us to suspect that the former may be more important than the latter. We also consider the possible roles played by tacit collusion, menu costs, and engineering costs.

The final section turns to the broader economic implications of uniform pricing. First, uniform pricing may exacerbate inequality. In a calibrated example, we find that uniform pricing increases the prices faced by consumers in the poorest decile of zipcodes by ten percent relative to the prices faced by consumers in the richest decile. Second, uniform pricing is likely to substantially dampen 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. Third, uniform pricing dramatically changes the standard analysis of retail mergers. Finally, uniform pricing alters the incidence of intra-national trade costs and can lead to severe biases in standard estimates of trade costs.

We are not the first to document uniform pricing policies in retailing. Online Appendix Table 1 provides an overview of relevant papers.³ Prior work by Nakamura (2008) and contemporaneous work by Hitsch et al. (2017) find that chain effects explain a substantial share of price variation in large samples of stores and products. Cavallo (2017, 2018a) documents uniform pricing using data from retailer websites in the context of comparing online to offline prices. Other papers have focused on zone pricing for select retail chains, including early studies of the Dominicks food chain (e.g., Hoch et al., 1995, Chintagunta et al., 2003) and more recent evidence on three home improvement

³Our discussion focuses on private retail firms. Miravete et al. (forthcoming) analyze the implications of a uniform markup regulation for government-run liquor stores in Pennsylvania. Their work parallels ours in estimating variation of demand elasticities and considering the distributional implications of uniform pricing.

chains including Home Depot (Adams and Williams, 2019).⁴ Our paper differs from most prior work focusing directly on the extent of uniform pricing in that it compares observed pricing to an optimal benchmark⁵ and highlights broader economic implications. As a separate contribution, the novel instrumental variables strategy we develop in Section 4.2 may be applicable to demand estimation in other settings where multi-establishment firms set partly or fully uniform prices.

Our paper also relates to a large body of work on the extent and implications of local retail price responses to economic shocks or incentives. Work by Gagnon and López-Salido (forthcoming), Cawley et al. (2018), and Leung (2018) provide direct evidence supporting our claim that uniform pricing dampens responses to local shocks. Our work also speaks to the literature tracing out the implications of retail firms’ price setting for macroeconomic outcomes, including influential early work using BLS data by Bils and Klenow (2004) and Nakamura and Steinsson (2008), and recent contributions using scanner data such as Anderson et al. (2017).⁶

Finally, our paper relates to work in behavioral industrial organization. For a review, see Heidhues and Koszegi, 2018. Most of the work in this area has focused on firms optimally responding to behavioral consumers (e.g., DellaVigna and Malmendier, 2004). Our paper is instead part of a smaller literature on behavioral *firms* — that is, cases in which firms deviate from simple benchmarks of profit maximization (Bloom and Van Reenen, 2007; Hortaçsu and Puller, 2008; Goldfarb and Xiao, 2011; Massey and Thaler, 2013; Hanna et al., 2014).

2 Data

Our primary data sources are the Nielsen Retail Scanner (“RMS”) and Nielsen Consumer Panel (“Homescan”) data provided by the Kilts Center at the University of Chicago. RMS records the average weekly revenue and quantity sold for over 35,000 stores and roughly four million unique products (UPCs). We use RMS data for the period 2006 to 2014. Homescan tracks the purchases of roughly 115,000 individual consumers across different stores. We use Homescan data to define the demographic composition of each store’s consumers. We provide additional details on the sample construction in the Online Appendix.

Stores. We focus our main analysis on food stores, and show supplemental analysis for drug and mass-merchandise stores. Table I Panel A shows that the initial Nielsen sample includes 38,539

⁴Other related papers summarized in Online Appendix Table 1 are Hwang et al. (2010), Cavallo et al. (2014), Kaplan and Menzio (2015), Gagnon and López-Salido (forthcoming), and Dobson and Waterson (2008).

⁵Adams and Williams (2019) is the only prior study that computes a benchmark for optimal pricing.

⁶A large body of prior work has considered the broader question of how responses to local and aggregate shocks may differ. See, for example, Nakamura and Steinsson (2008), Stroebel and Vavra (forthcoming), and Beraja et al. (2019). Our contribution is to show that uniform pricing can change predictions of local responses and to provide an additional mechanism for different responses to shocks at different geographic levels.

stores with total average yearly revenue of \$224 billion recorded in RMS.⁷

We define a chain to be a unique combination of two identifiers in the Nielsen data: *parent_code* and *retailer_code*. The former generally indicates the parent company that owns a chain and the latter indicates the chain itself. Nielsen does not disclose the names of the chains in the data, but a general example would be the Albertson’s LLC parent company which owns chains including Albertson’s, Shaw’s, and Jewel-Osco. Sometimes, a single *retailer_code* appears under multiple *parent_codes*, possibly for reasons related to mergers.

To define our main analysis sample, we exclude stores that switch chains over time,⁸ stores in the sample for fewer than 104 weeks, and stores without any consumer purchases in the Homescan data. We next introduce restrictions at the chain level. First, we require that the chains are present in the sample for at least 8 of the 9 years. Second, in cases where the same *retailer_code* appears for stores with different *parent_codes*, we only keep the *parent_code* associated with the majority of such stores, and we further exclude cases in which this *retailer_code-parent_code* combination accounts for less than 80% of the stores with a given *retailer_code*. Third, we exclude chains in which 60% or more of stores belonging to a *retailer_code-parent_code* combination change either *parent_code* or *retailer_code* in our sample.

As Table I Panel A shows, these restrictions result in a final sample of 22,680 stores from 73 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). We focus our main analysis on food stores and present results for drug and mass merchandise stores in supplemental analyses.

To define store demographics such as income and education, we use the characteristics of consumers in the Homescan panel. The median food store has 27 Homescan panelists ever purchasing at the store, for a total of 1,067 trips (Table I Panel B). We associate each consumer with the demographics of their zipcode of residence as measured in the 2008-12 American Community Survey (ACS). We then define the demographics of a given store to be the average of these zipcode demographics across its consumers, weighted by the number of trips they take to the store.⁹ We let Y_s denote the resulting measure of income for store s . This is equal to \$27,420 for the median store, \$22,770 for the 25th percentile store, and \$34,300 for the 75th percentile store.

Turning to chain-level summary statistics, Panel C of Table I shows that the median food chain has 66 stores spanning 4 Designated Market Areas (DMAs) and 2.5 states. Online Appendix Figure

⁷This figure omits revenue from products not in the RMS data, including prescription drugs and most produce.

⁸We return to such switchers in an event-study analysis in Section 5.

⁹This demographic information is more granular than a measure computed from the store location recorded in the RMS data, which is either a county or a 3-digit zipcode.

1 shows the locations of these stores. Online Appendix Table 2a shows that the number of stores and chains in the sample is roughly constant over time.

Products. We focus on a sample of products that are frequently sold and available across many of the chains. This guarantees clean comparisons both within and between chains. It also minimizes the frequency of missing observations due to prices not being observed in weeks with zero sales. Such missing values can cause systematic bias since zero sales are more likely to occur in weeks when prices are relatively high.

We consider 40 product categories or “modules” (Online Appendix Table 2b), ranging from canned soup to soda to yogurt.¹⁰ These account for an average yearly revenue of \$49.3 billion across the 9,415 food stores, or 36.2 percent of total revenue in the Nielsen data for these stores. Within a module, we keep all the products (UPCs) which satisfy an availability restriction: pooling across chains, a product must have positive revenue in at least 80% of store-weeks. This leaves us with 1,365 products covering \$10.8 billion in yearly revenue for the food stores, corresponding to 20 percent of the \$49.3 billion total revenue in the 40 modules.¹¹ Examples include a 12-can package of Coca-Cola and a 59 oz. bottle of pulp-free Simply Orange juice. Within this sample, for robustness we consider products in the top decile of revenue as well as products in the bottom decile of revenue; the patterns for this latter group allow us to conjecture pricing patterns for products which we do not cover due to inconsistent availability. We apply similar restrictions for drug and mass merchandise stores, covering the modules listed in Online Appendix Table 2c.

We also consider different, and often larger, sets of products for additional robustness. We create a sample of store-brand products, as indicated by Nielsen. We also define baskets of products to construct module-level indices. For a given chain and module, we include all products such that the average across stores of the share of weeks with non-zero sales for that UPC is at least 95 percent. For some modules such as soda and orange juice, products meeting this criterion cover 55-65 percent of the total module revenue, while for other modules like chocolate or coffee, they cover just 17-30 percent. Finally, for our analysis of availability we consider an even broader set of 43,269 products, covering \$44.6 billion, that is, 90.5% of the revenue in the 40 modules considered.

Prices. We define the price P_{sjt} in store s of product j in week t to be the ratio of the weekly revenue to weekly units sold. The price is not defined if no units are sold in a store-product-week. As Table I Panel D shows, the products in our main sample have positive sales in 90.7% of weeks, with an average price of \$2.86. We let p_{sjt} denote the demeaned log price, defined as the residuals

¹⁰These modules have a large overlap with ones used in previous analyses, e.g., Hoch et al. (1995). Of the remaining 1148 modules in the data set, many have small total revenues or include products that are very differentiated across chains, and thus would not satisfy our restrictions.

¹¹The remaining revenue is accounted by products that fail our criteria because they sell relatively infrequently, or because they sell in some chains but are absent in many chains, e.g., regional products like dairy.

from a regression of $\log(P_{sjt})$ on product-year fixed effects. We define the store-product average price \bar{p}_{sj} to be the simple average of p_{sjt} across t , and the store average price \bar{p}_s to be the simple average of \bar{p}_{sj} across j .

We also define price and quantity indices at the module level. We start from the weekly log price $\log(P_{sjt})$ and weekly log units sold $\log(Q_{sjt})$, then average across all products j included in the basket defined above for that module-chain-year, weighting by the total quantity sold for product j in the chain-year. If a product j has no sales in a particular store s and week t , we omit product j from the store-week cell and scale the weights on other products up accordingly.¹²

Major Grocer’s Data. We use supplemental scanner data from a major grocer (*parent_code*) studied in Gopinath et al. (2011). These data contain the same variables as the Nielsen data, plus gross revenue (defined to be the total revenue had all transactions occurred at the non-sale posted price), wholesale prices paid, and gross profits. The definition of weeks in these data differs from Nielsen, and is aligned with the timing of the retailer’s weekly price changes. The data cover 250 stores belonging to 12 retailers (*retailer_codes*) from 2004 until mid-2007. We focus on the largest retailer, which has 134 stores, and match 132 stores to stores in our main sample.

3 Descriptive Evidence

3.1 Example

We begin with a visualization of pricing by a chain that is representative of the typical patterns in our data. Figure I Panel A shows the prices of one orange-juice product. The rows correspond to the stores in the chain, and are sorted by income. The columns correspond to weeks from January 2006 to December 2014. The color of each store-week indicates the demeaned log price p_{sjt} . Darker colors correspond to higher prices and white indicates missing values due to zero sales.

The figure shows substantial variation of prices across weeks, with frequent sales, but virtually no variation across stores within a week. Variation across stores is not visibly correlated with store income, i.e., the vertical position of stores in the chart, even though income ranges from only \$13,000 at the bottom of the chart to roughly \$50,000 at the top. Figure I Panel B shows a similar pattern for five other products in the yogurt, chocolate, soda, cookies, and cat food modules respectively. Here we display 50 of the stores shown in Panel A. We see variation across products in the depth and timing of sales, but again no systematic variation in prices across stores. The

¹²We use the same weights for price and quantity so that, under the assumption that all products in a module have a constant-elasticity demand with the same elasticity η , we can recover η by regressing the module index quantity on the module index price. We use quantity weights to mirror a geometric modified Laspeyres Index, similar to the index of Beraja et al. (2019) and consistent with the BLS category-level price indices. Our index is not exactly a geometric Laspeyres Index because the weights are not week 1 weights but instead average quantities in year y .

pricing patterns of this chain are representative of the large majority of chains in our sample. We present two other examples in Online Appendix Figure 2 Panels A and B.

While patterns like these are typical of the majority of chains, a few other chains follow a pattern of zone pricing: within a zone (often a state or group of states), the patterns will look like the ones documented above, with larger price differences across zones.

3.2 Measures of Pricing Similarity

To describe chain pricing patterns more systematically, we introduce three measures of the extent of uniform pricing. Each is defined separately for each pair of stores s and s' and product j .

The first measure is the *quarterly absolute log price difference*. For each pair of stores s and s' , product j , and quarter v , we first compute the average of $\log(P_{s jt})$ and $\log(P_{s' jt})$ respectively across weeks with non-missing prices in both store s and s' . We then compute the absolute difference between the average for store s and the average for store s' . Finally, we average these quarterly absolute differences across quarters.

The second measure is the *weekly correlation of log prices*. For each pair of stores s and s' and product j , we first compute the residuals from regressions of $\log(P_{s jt})$ and $\log(P_{s' jt})$ on store-product-year fixed effects. We then compute for a product j the correlation between the store s residuals and the store s' residuals, including all weeks t which are non-missing in both store s and s' for product j , provided there is a minimum of 26 such weeks.

These two measures capture two orthogonal dimensions: cross-sectional differences in average prices, and correlation of price changes over time. Two stores with the same timing and depth of sales, but different regular prices would have high weekly correlation but also a high quarterly difference. Conversely, two stores with similar average prices at the quarterly level, but different timing of sales would have a low quarterly difference but also a low weekly correlation.

The third measure is the *share of (nearly) identical prices*, defined as the share of observations across weeks t for which $|\log(P_{s jt}) - \log(P_{s' jt})| < 0.01$ for a product j . We require a minimum of 26 weeks of non-missing observations for a pair-product.

To compute within-chain measures of similarity for product j and chain c , we sample up to 200 pairs of stores (s, s') in chain c , and we average the measures of similarity across all such pairs, using the same pairs for all products. To compute between-chain measures of similarity for product j and chain c , we follow a similar procedure but draw 200 pairs composed of a store s in chain c and a store s' belonging to a different chain c' .

In Figure II Panels A, B, and C we plot the distribution of these measures, with each product, chain forming one observation. Prices for within-chain pairs (solid bars) are far more similar than

for between-chain pairs (hollow bars) on all three measures. The absolute log price difference (Panel A) is typically below 5 log points for the within-chain pairs, and typically above 8 log points for the latter. The weekly correlation (Panel B) is nearly always above 0.5 for within-chain pairs but almost always below 0.3 for between-chain pairs. The share of identical prices (Panel C) is typically above 0.3 for within-chain pairs, but is rarely above 0.3 for between-chain pairs.

Table II summarizes variants of these measures. Compared to the baseline price similarity in these pairs (Panel A), Panel B shows that the patterns are essentially unchanged if we restrict attention to cases where stores s and s' are in the same geographic market (DMA). Panel C shows the same for cases where s and s' are in different DMAs and also face very different income levels (with one store in the pair in the top third of the income distribution and the other in the bottom third). Online Appendix Figure 3 displays the distribution of distance between pairs for these samples. These results provide initial evidence against the possibility that the observed similarity reflects same-chain store pairs serving more homogeneous consumers in terms of either geography or demographics. The results also suggest that the observed similarity does not result from constraints specific to stores operating in the same geographic market, for example, because price advertising is determined at the newspaper or television market level. The results are similar if we focus just on food stores (Table II Panel D), drug stores (Panel E), or mass-merchandise stores (Panel F). The results are also similar for products in the top and in the bottom decile of revenue, and for generic products in food stores (Online Appendix Table 3).¹³

Figure III summarizes pricing similarity at the chain level. In Panel A, quarterly absolute log price difference is on the horizontal axis, weekly correlation is on the vertical axis, and each point indicates the similarity measure for a chain, averaged across all products. The majority of chains cluster in the upper-left of the figure: of the 73 chains, 59 have both an average correlation of weekly prices above 0.7 and an absolute quarterly distance in prices below 5 log points. These two measures of pricing similarity are highly correlated: chains that are similar in one dimension are also similar in the other dimension. This correlation is not mechanical. One might have seen, for example, highly correlated sales but varying levels of regular prices across stores.¹⁴

Panel B returns to the phenomenon of zone pricing. We decompose the pricing similarity into similarity for pairs of stores within a state, versus across state boundaries. Focusing on chains that operate at least three stores in two or more states, we plot the within-state log price difference on

¹³We cannot compute the between-chain similarity in pricing for generic products, which are not obviously comparable across chains. Nielsen assigns the same “masked” UPC to generic products that it deems similar across chains, but we cannot verify that these different generic products are of similar quality.

¹⁴The measure of similarity based on quarterly log prices differences is correlated with the measure based on the share of (near-)identical prices (Online Appendix Figure 4 Panel A). Also, chains with uniform prices for our benchmark products also tend to have uniform prices for generic products (Online Appendix Figure 4 Panel C).

the horizontal axis, and the between-state log price difference on the vertical axis. Zone pricing following state boundaries should show up in this figure as larger differences between state than within state, i.e., points lying above the 45-degree line. For the majority of chains, the within-state and between-state differences are similar; these chains do not appear to determine prices by state. A minority of chains do have clearer zone pricing patterns, notably chain 9.¹⁵

An alternative measure of price uniformity, motivated by Nakamura (2008) and Hitsch et al. (2017), is the share of variation in prices explained by chain fixed effects. For each product j , we run a regression of the store-product prices \bar{p}_{sj} on chain and DMA fixed effects. If pricing is near-uniform within chains, chain fixed effects should explain most of the variation. If local factors such as competition or demand shifters are the key drivers, DMA fixed effects should explain most of the variation. Online Appendix Figure 5 shows that the median R^2 with DMA fixed effects is 0.486, while the median R^2 with chain fixed effects is 0.824. The R^2 with chain fixed effects rises to 0.872 when we drop the chains identified in Figure III Panel B as engaging in zone pricing.

3.3 Price Response to Consumer Income

We now turn to the relationship between prices and income, focusing on food stores. All else equal, we would expect stores in high-income areas to have less elastic consumers, a prediction we confirm below; in this case, we expect these stores to charge higher prices.¹⁶ Though we argue that variation in marginal costs across stores is likely to be small, any such variation might well be positively correlated with income, and so tend to strengthen this relationship.

Figure IV shows the relationship between income and log price within and between chains. For the within-chain relationship, we regress both store average log price \bar{p}_s and store income Y_s on chain fixed effects, and show a binned scatterplot of the residuals in Panel A. The relationship is positive and highly significant, but the magnitude is very small: an increase in per-capita income of \$10,000, equivalent to a move from the 30th to the 75th percentile, is associated with a price increase of only 0.47 log points (0.47 percent). For the between-chain relationship, we show a scatterplot of chain averages in Panel B. This relationship is also highly significant, but an order of magnitude larger: a \$10,000 increase is associated with a price increase of 4.2 log points.

We view this sharp contrast between the within- and between-chain results as one of our key findings. It suggests that chains are either varying their prices far too little across stores in response to income, or varying their prices far too much at the overall chain level. Our model below separates

¹⁵Online Appendix Figure 4 Panel B shows a parallel figure using correlation in weekly prices.

¹⁶That elasticities vary inversely with income is a standard prediction of theory and the motivation for the common practice of interacting price with income in empirical demand specifications. Berry et al. (1995), for example, derive such a relationship from a simple Cobb-Douglas utility specification and impose it in estimation.

these two hypotheses, providing strong support for the former.

Panel C shows how these key relationships vary across products. For each of the 1,365 products, we estimate the slope of the within-chain relationship (as in Panel A) and of the between-chain relationship (as in Panel B). We then plot the distribution of the two estimated slopes. The within-chain price-income slope is remarkably consistent, with almost all estimates between 0 and 0.015. In contrast, the between-chain price-income slope ranges more broadly, mostly between 0 and 0.10.

In the Online Appendix, we extend this analysis focusing on zone pricing. In Online Appendix Figure 6 Panel A, we re-estimate the within-chain relationship, but now plot residuals after taking out chain-state fixed effects. This reduces the price-income slope to 0.38 log points per \$10,000 of income, but it remains statistically significant. In Online Appendix Figure 6 Panel B, we show the complementary plot of chain-state mean prices after subtracting the chain mean. Across states within a chain, a \$10,000 income increase is associated with an increase in prices of 1.24 log points, a slope about a quarter the size of the between-chain relationship. This relationship is stronger for the 12 food chains that we identify as zone pricers based on Figure III Panel B.

Next, we consider an extensive set of robustness checks. Estimating within-chain coefficients separately by chain (Online Appendix Figure 7) shows that all but three chains have coefficients below 0.01. The pattern of small within-chain response and large between-chain response is clear in nearly every module (Online Appendix Figure 8). The results are similar replacing income with the fraction of college graduates (Online Appendix Figure 9 Panels A and B), or median home prices (Online Appendix Figure 9 Panels C and D). We also obtain similar results for store-brand, high-revenue, and low-revenue products, for prices weighted by product revenue, and for module-level price indices (Online Appendix Figure 10). Further, Online Appendix Table 4 shows that the within-chain results are stable to (i) running the specifications at the store-product level instead of at the store level, (ii) including only products with valid elasticities, and (iii) including chain-state fixed effects instead of just chain fixed effects. The between-chain results are similarly robust to (i) running the specification at the chain-product level instead of at the chain level, (ii) running an unweighted regression, and (iii) dropping two chains that appear as outliers in Figure IV Panel B. Moving beyond demographics, a simple measure of store-level competition has a similarly small, but statistically significant, impact on prices within chain (Online Appendix Table 5). The within-chain evidence for mass merchandise and drugstore chains is similar to the evidence for food stores, with a larger zone pricing relationship for drug stores (Online Appendix Figure 11).

Table III presents an alternative view of the price-income relationship. We run a store-level regression of average log price \bar{p}_s on both store income Y_s and the average income of stores in the chain to which s belongs. In some specifications, we also include the average income in s 's chain-

state. We separate food stores (Panel A) from drugstores and mass merchandise stores (Panel B and C), as we can only do a meaningful between-chain comparison for the food stores. The first column presents the regression including only own-store income as a benchmark. The second column adds chain average income for food stores. Consistent with the evidence in Figure IV, a store’s response to its own consumers’ income is an order of magnitude smaller than its response to the average income served by its chain. This result remains when we add county fixed effects (third column). Thus, if we look at two stores in the same county both attracting consumers of the same income, one from a mainly high-income chain and one from a mainly low-income chain, the former will tend to charge higher prices than the latter.¹⁷ The fourth and fifth columns add chain-state average income as a regressor. This response is larger than the own-store-income response but smaller than the response to overall chain average income, consistent with our zone pricing results. In Online Appendix Table 6 we show that the results are parallel for alternative sets of products.

3.4 Composition Effects

The within-chain pricing in Figure IV Panel A poses a puzzle. Why would chains vary prices in a systematic way with consumer income, but then do so with a tiny magnitude far smaller than the one with which they respond to income at the chain level, and far smaller than the analysis below suggests would be profit maximizing? We show that this small price-income relationship is likely to be mainly an artifact of compositional differences, due to the fact that the standard Nielsen price measure is the weekly average price paid rather than the price the store posts on any given day.

If all consumers of a store in a given week paid the same prices, weekly average price paid and posted price would be equal. For them to diverge, prices paid must vary within a week. There are two main reasons why they are likely to do so. First, Nielsen’s weekly revenue and units sold are based on a week that runs from Sunday through Saturday. Although most retailers change prices at a weekly frequency, their price changes may occur on a different day of the week. For example, if a retailer changes prices on Wednesdays, consumers who buy in the first half of Nielsen’s week will pay a different price from those who buy in the second half. Second, some consumers may use store cards or obtain other discounts that lead them to pay lower prices.

Consider a retailer that charges identical prices in all stores and that changes prices on Wednesdays. In a particular week, they cut the price from P^{high} to P^{low} . The weekly average price in the Nielsen data for store s will be $P_s^{RMS} = \theta_s P^{high} + (1 - \theta_s) P^{low}$, where θ_s is the share of purchases in the first half of the week in store s . If the share θ_s varies across stores, this will obscure the

¹⁷An alternative explanation is that the chain-level income measure may capture the elasticity of consumers in a store better than the own-store income measure due to measurement error in the latter. Counter to this explanation, the store elasticity is predicted by own-store income but not by chain-level income (Online Appendix Table 8).

fact that the chain is charging uniform prices. In fact, the share θ_s will vary systematically: for stores facing less elastic consumers, fewer will shift purchases to the low price, and θ_s will be higher. Measured prices P_s^{RMS} will thus be higher for stores facing higher income or otherwise less elastic consumers, even with uniform prices. A similar effect arises if the share of consumers who use store cards or other discounts is greater among more price elastic consumers.

We illustrate this with an example calibrated to data from the major grocer described in Section 2, which indeed changes prices on Wednesdays. Suppose that the income of store s is \$10,000 greater than the income of store s' , and that, consistent with the estimated income-elasticity relationship (Table V, Column 2), this translates into price elasticities among their respective consumers of $\eta_s = -2.5$ and $\eta_{s'} = -2.63$. Suppose that P^{low} is about 30 percent lower than P^{high} , consistent with the average depth of a sale for this grocer. Consistent with our model below, assume a constant-elasticity demand function $Q_s = kP_s^{\eta_s}$. Then $(\theta_s, \theta_{s'}) = (0.298, 0.289)$, and the difference in average weekly log prices would be 0.0034.¹⁸ For this grocer, products go on sale and off of sale in about 30% of weeks, which implies a slope of 0.0010 in Figure IV Panel A. Thus, under reasonable assumptions, the week misalignment could explain one fifth of the observed slope of 0.0047. These calibrations do not account for heterogeneity in the share of consumers using store cards. To the extent that this share is correlated with consumer income, this second compositional effect could account for another large part, or potentially all, of the within-chain price-income slope.¹⁹

We use the data from the major grocer to provide direct evidence on compositional effects. We expect to see the effect arising from mid-week price changes in the Nielsen data, where the data is reported on a Sunday-to-Saturday basis, but not in the grocer's data, where it is reported on a Wednesday-to-Tuesday basis. Figure V Panel A shows a binned scatterplot of the within-chain relationship using the Nielsen price measure for the 132 stores in both data sets. This grocer uses geographic pricing zones, so we focus on the within-chain-state relationship, yielding a slope of 0.0032, similar to the one for all stores in Online Appendix Figure 6 Panel A. Using weekly prices from the grocer's data yields a slope of 0.0018 (Figure V Panel B). The difference between these slopes ($0.0032 - 0.0018 = 0.0014$), which is likely due to the week-offset effect, lines up nicely with

¹⁸ Assume prices change in the middle of the week and note that $\theta_s = \frac{q_s^{high}}{q_s^{high} + q_s^{low}} = \frac{(P^{high})^{\eta_s}}{(P^{high})^{\eta_s} + (P^{low})^{\eta_s}}$. Plugging in the values for P^{high} , P^{low} , η_s , and $\eta_{s'}$ yield the values of $(\theta_s, \theta_{s'})$ which in turn yield values of P_s^{RMS} and $P_{s'}^{RMS}$.

¹⁹ Assume that (i) the regular price is unaffected by the use of store cards; (ii) in case of sale, a share r of consumers without a store card pays the full price, instead of the sale price; and (iii) the share r is a linear function of consumer income: $r = R + \rho Y$. The log price in case of sale will be $p^{sale} = \log(r(Y)P^{high} + (1 - r(Y))P^{low})$, and thus $\partial p^{sale} / \partial Y = \rho(P^{high} - P^{low}) / p^{sale}$. For the major grocer, the average depth of a sale is 30 percent, which implies $(P^{high} - P^{low}) / p^{sale} \approx 3/7$. Given that items are on sale about 30 percent of weeks, the average log price p would respond to consumer income as $\partial p / \partial Y = 0.3 * \rho * 3/7 \approx 0.13\rho$. Thus, if a \$10,000 increase in income were associated with a 3 percentage point increase in the share of non-store-card users (that is, $\rho = 0.03$), this second bias could explain the remaining observed slope.

our calibration above.²⁰ Figure V Panel C shows the same slope using the posted non-sale price, which we observe in the grocer’s data. This is not the object we would ideally like—if stores vary the frequency or depth of their sales, we would consider this real variation in posted prices—but it provides a benchmark. Here nearly all remaining slope disappears, with an insignificant point estimate of 0.0010. This chain sets near uniform non-sale prices with respect to income.

These compositional effects impact also the analysis of price similarity in Section 3.2, as we illustrate in Online Appendix Table 7 for within-DMA pairs of stores for the major grocer. In the RMS data, the price distance measures are similar to the ones for other retailers: 0.026 quarterly log price distance, 0.874 correlation, and 47.5% of identical prices within chain. In the major grocer data, we can control not only for the week offset, but also for the differential coupon use across stores by excluding pairs with prices that are not exact to the cent, likely indicating that different consumers paid different prices. When we do so, the quarterly log price distance decreases to 0.013, the correlation increases to 0.949, and the share of identical prices nearly doubles to 81.4%. Once one controls for both compositional effects, the pricing for this grocer is nearly perfectly uniform.

Our conclusion is that the within-chain slope shown in Figure IV Panel A is at least partly an artifact of composition effects. We suspect that a large majority of chains are in fact charging the same prices in all of their stores, or in all stores within geographic zones.

3.5 Product Assortment

Firms choose not just prices, but also other store-level variables including which brands and how many product to carry, and whether to sell organic products. These dimensions may provide alternative ways to cater to consumer demand. For example, stores in lower-income areas might carry a cheaper olive oil while stores in higher-income areas carry an expensive extra-virgin version.

To investigate this margin, we define an assortment index based on the average national-level prices of the products sold by a store. The index is higher when the store carries relatively expensive products but does not depend on the prices charged by the store itself. We define a per-unit constant price for product j to be the average log price for product j in year y across all stores s that carry it, divided by the unit size (e.g., 40 oz). We then average this measure over products carried by store s to create an assortment price index for a store s , year y , within a sub-module b . Sub-modules are defined as subsets of modules with similar product size. To create the final assortment index for store s , we demean the index by sub-module-year, average over the years and over the sub-modules. We provide additional detail in the online appendix in Section A.1.9. We emphasize that, as Table

²⁰The fact that the differences in slope, 0.0014, lines up with the calibrated prediction, 0.0010, provides an indirect validation for our elasticity estimates, which are used for the calibration. The difference in slopes is similar if we do not account for geographic pricing zones (Online Appendix Figure 12).

1 shows, the sample of products covered in the assortment decision is an order of magnitude larger than for our main analysis, and covers 90.5% of the revenue in the 40 modules considered.

Figure VI displays the within-chain and between-chain relationship of this assortment index with store-level income. The patterns are strikingly parallel to the ones for pricing: the within-chain slope of the assortment index with respect to income (0.0058, Panel A) is an order of magnitude smaller than the between-chain slope (0.0489, Panel B). The distribution of the within-chain and between-chain slopes across sub-modules of products (Panel C) reinforces this pattern.²¹

Column 1 of Table IV shows that the within-chain difference in assortment between pairs of stores is about four times smaller than the between-chain difference in assortment (Panel A). Furthermore, the assortment index responds an order of magnitude more to chain-level income than it responds to store-level income variation (Panel B). In Columns 2-4 we document similar patterns for additional measures of assortment: the fraction of products within a sub-module that are in the top 10% of unit prices (Column 2), the fraction of organic products for the relevant modules (Column 3), and the fraction of generic products (Column 4).

Finally, in Column 5 we consider a different decision: how many products to carry in a store. As with assortment, the stores are more similar within a chain than across chains, and respond more to chain-level income than to store-level income.²²

4 Demand Estimation and Optimal Prices

4.1 Model

We introduce a simple demand model to gauge the degree to which we would expect prices to vary within and between chains. The model makes strong assumptions, and we do not necessarily take deviations from the model predictions to imply a failure of profit maximization.

A monopolistically competitive chain r chooses a price P_{sj} in each product-store to maximize total profits. Each store's residual demand for product j takes the constant elasticity form $Q_{sj} = k_{sj}P_{sj}^{\eta_{sj}}$, where Q_{sj} is the number of units sold, k_{sj} is a scale term, and η_{sj} is the store's price elasticity for product j . Total cost consists of a marginal cost c_{rj} that may vary by chain but does not vary by store within a chain, as well as a store-level fixed cost C_{sj} . The chain solves

$$\max_{\{P_{sj}\}} \sum_{s(r),j} (P_{sj} - c_{rj}) Q_{sj}(P_{sj}) - \sum_{s(r)} C_{sj}. \quad (1)$$

²¹The zone pricing assortment patterns also parallel the results for pricing (Online Appendix Figure 13).

²²Many relevant managerial decisions, such as the store size, opening hours, and shelf space allocations, are not observed in the data. We did attempt to examine the incidence of featured and displayed products, but the information is not recorded with enough consistency to allow informative within- and between-chain comparisons.

The first order conditions imply the optimal price P_{sj}^* satisfies

$$\log(P_{sj}^*) = \log\left(\frac{\eta_{sj}}{1 + \eta_{sj}}\right) + \log(c_{rj}). \quad (2)$$

Thus, under optimal pricing a regression of log prices on $\log\left(\frac{\eta_{sj}}{1 + \eta_{sj}}\right)$ with chain-product fixed effects should yield a coefficient of one.

While the assumption of constant marginal costs within a chain is strong, several pieces of evidence suggest that it may be a reasonable approximation. Stroebel and Vavra (forthcoming) use data on wholesale costs for a retailer to show that geographic variation in these costs is minimal. Since wholesale costs account for three-quarters of total costs, this limits the scope for cost variation. They then present further evidence suggesting that neither variation in labor costs nor variation in retail rents plays a significant role. If marginal costs vary positively with factors like labor costs or retail rents this would introduce upward bias in the price-income relationship, and so strengthen our core finding that this relationship is surprisingly small. We confirm Stroebel and Vavra (forthcoming)’s findings for wholesale costs in our large grocer’s data, finding no relationship between wholesale costs and store income (Online Appendix Figure 14).

4.2 Elasticity Estimates

To evaluate the extent to which observed pricing matches the prediction of equation (2), we require estimates of the elasticities of demand η_{sj} for each store-product pair. As our benchmark measure, we estimate the response of weekly log quantity to the weekly log price. We allow for store-product-year fixed effects α_{sjy} and store-product-week-of-year fixed effects γ_{sjw} and estimate separately for each store-product sj :

$$\log(Q_{sjt}) = \eta_{sj} \log(P_{sjt}) + \alpha_{sjy} + \gamma_{sjw} + \epsilon_{sjt}. \quad (3)$$

To address potential endogeneity of prices, we instrument $\log(P_{sjt})$ with the average of $\log(P_{sjt})$ across other stores in s ’s chain that are located outside s ’s DMA. This is a version of the instrumenting approach introduced by Hausman (1996) and applied by Nevo (2001), where prices of a product in other markets serve as instruments. The existence of uniform pricing within chains combined with frequent sales makes this approach particularly convincing. The key assumption—that the timing of chain-level sales is unrelated to local demand shocks for a store (after controlling for seasonality via γ_{sjw})—is in our view compelling, given the jagged and idiosyncratic pattern of sales illustrated in Figure I and the heterogeneous demand conditions faced by stores within chains. As

further evidence, we discuss below robustness analysis in which we allow for DMA-product-week fixed effects, which will soak up any local market shocks. The first-stages of our regressions are very strong, with coefficients close to 1 (Online Appendix Figure 15 Panel C).²³

We believe this instrumental variables strategy may be applicable in other settings with uniform pricing. Average chain decisions may provide a valid instrument for price differences both over time (as here) and across products. It is also possible to extend the strategy to instrument for a store’s prices with the average *demographics* of other stores in the chain, following a logic similar to George and Waldfogel (2003). Allcott et al. (2019b) and Allcott et al. (2019a) are examples of papers that have already built on this strategy.

Our estimated price elasticity $\hat{\eta}_{sj}$ is the coefficient on $\log(P_{sjt})$ from this IV regression. We cluster standard errors by two-month periods. We require that valid $\hat{\eta}_{sj}$ have at least 104 weeks of data as well as standard errors within (0.01, 1.25). To adjust for sampling error, we use an Empirical Bayes (EB) procedure for each chain-product that shrinks elasticities to their chain-product mean.²⁴ We denote the shrunk elasticities by $\tilde{\eta}_{sj}$ and winsorize both $\hat{\eta}_{sj}$ and $\tilde{\eta}_{sj}$ at -7 and -1.2 . We also define store-level price elasticities $\hat{\eta}_s$ (respectively, $\tilde{\eta}_s$) by first subtracting from $\hat{\eta}_{sj}$ ($\tilde{\eta}_{sj}$) its mean across s and then averaging the demeaned values across products j .

Figure VII Panel A provides evidence on the fit of the demand model. The figure shows a binned scatterplot of residuals of $\log(Q_{sjt})$ against residuals of $\log(P_{sjt})$ after partialing out the fixed effects α_{sjy} and γ_{sjw} , for a random sample of 25 products. To illustrate variation in elasticities by income, we choose the 50 stores nearest to the \$20,000 income level (11th percentile) and the 50 stores nearest to the \$60,000 income level (98th percentile). The model assumes a linear relationship, and the figures shows that this assumption holds to a remarkable degree. The plot also shows that demand is less elastic in the higher-income stores and more elastic in the lower-income stores.

Panel B shows the distribution of the raw and EB-adjusted store-level elasticities $\hat{\eta}_{sj}$ and $\tilde{\eta}_{sj}$.²⁵ These are well-behaved, with all but a handful of values smaller than the theoretical maximum of -1 , and most of the mass falling between -1.5 and -3 .

Alternative Elasticities. A concern is geographically correlated demand shocks that affect many stores within a chain and thus could bias even the instrumented elasticity estimates. To address this concern, we re-estimate specification (3) at the product-DMA level including an additional fixed effect ζ_{djt} for the DMA-product-week, capturing any time-series component that is common across all stores in a DMA. Online Appendix Figure 16 Panel B shows that this has a

²³The estimated elasticities with this IV procedure are highly correlated with elasticities estimated from specification (3) with OLS, as Online Appendix Figure 16 Panel A shows.

²⁴Online Appendix Figure 15 Panel A shows the distribution of the shrinkage parameter. See detail on the shrinkage procedure in the online appendix in Section A.1.11.

²⁵Online Appendix Figure 15 Panel B shows the distribution of the elasticities at store-product level.

very limited impact on the elasticities. A different omitted variable could be chain-level variation in promotions or product displays. For some products and stores, the Nielsen data set carries information on feature and display promotion. Adding these controls once again has a very limited impact (Online Appendix Figure 16 Panel C).

Our short-run elasticities may differ from the longer-run elasticities that are relevant to the pricing problem. Longer-run elasticities could be smaller due to stockpiling, or larger if it takes consumers time to adjust to price changes. As a step toward addressing these concerns, we estimate quarterly elasticities. We average the weekly log price and log units sold across all weeks in a quarter, and re-estimate a modified version of equation (3).²⁶ As Online Appendix Figure 17 shows, the quarterly elasticities are smaller (in absolute value) than the benchmark ones, but the two measures are highly correlated. As with the benchmark elasticities, the log-log specification is approximately linear and the elasticities are highly correlated with local income. We return to these elasticities in the next section, and we present alternative evidence on long-run elasticities in Section 5.

To illustrate the role of stock-piling, we run a regression of $\log(Q_{sjt})$ on $\log(P_{sjt})$ including leads and lags of prices (Hendel and Nevo, 2006). The specification includes the same set of fixed effects as (3), but restricts the coefficient on the log price variables to be the same across products and stores. As Online Appendix Figure 18 shows, the coefficients on lagged prices variables, while statistically significant and in line with the predictions of a stockpiling model, are an order of magnitude smaller than the coefficients on price in week t , even for the storable products.

Our model also ignores substitution between products. If some of the response in our benchmark elasticities reflects within-store substitution, the optimal price response could be smaller than our model predicts. To partially address this issue, we re-estimate our elasticities using the module-level price and quantity indices described in Section 2. As Online Appendix Figure 19 shows, the module-level elasticities are again smaller in absolute value than the benchmark elasticities, but the two are highly correlated. We return to these elasticities as well in the next section.

Correlates. Returning to our benchmark elasticities, Table V relates the raw store-level elasticities $\hat{\eta}_s$ to demographic and competition measures. The first two columns show a robust relationship between elasticity and income.²⁷ In column 3, we add the share of college graduates, the median home price, controls for the fraction of urban area, and indicators for the number of competitor stores within 10 kilometers. Income remains the most significant determinant, and

²⁶We estimate quarterly elasticities in groups of up to 20 products within each module, forcing the same elasticity in a module-group. Specifically, we estimate $\log(Q_{sjv}) = \eta_{sg} \log(P_{sjv}) + \alpha_{sjy} + \gamma_{sjq} + \epsilon_{sjv}$ for groups of products j in each module-group g . The fixed effect γ_{sjq} controls for store-product-quarter fixed effects.

²⁷We consider two alternative income measures: (i) the income of the county where the store is located, and (ii) the self-reported income of store consumers in the Homescan data, averaged by trip. Online Appendix Table 9 shows that our benchmark income measure is the best predictor of elasticity and of log elasticity.

elasticities increase with the degree of competition. We find similar results in columns 4-6 with the log elasticity term $\log\left(\frac{\hat{\eta}_{sj}}{1+\hat{\eta}_{sj}}\right)$ as the dependent variable.²⁸

4.3 Comparing Observed and Optimal Prices

In this section, we test the model and estimate the empirical analogue of equation (2),

$$\bar{p}_{sj} = \alpha_{rj} + \beta \hat{\lambda}_{sj} + \varepsilon_{sj}, \quad (4)$$

where α_{rj} are chain-product fixed effects and $\hat{\lambda}_{sj} = \log\left(\frac{\hat{\eta}_{sj}}{1+\hat{\eta}_{sj}}\right)$. Under the assumptions of the model, the coefficient β on the log elasticity term equals 1.²⁹ If chains under-respond to elasticity variation, we will instead observe $\beta < 1$. We instrument $\hat{\lambda}_{sj}$ with store-level income to address measurement error in the estimates $\hat{\eta}_{sj}$. The standard errors are block bootstrapped by *parent_code*.

We then compare the within-chain price variation captured by equation (4) to the analogous variation between chains. We regress chain average prices \bar{p}_{rj} on product fixed effects α_j and the simple average of $\hat{\lambda}_{sj}$ across stores s within chain r , instrumenting for the latter with chain average income. The model of equation (2) makes no direct prediction about the coefficient in this regression, since it allows marginal costs c_{rj} to vary across chains. However, provided that such marginal cost variation is small relative to demand-side variation, we expect a coefficient of approximately one.

Finally, to facilitate data visualization and simplify computation, we estimate versions of both the within-chain and between-chain specifications aggregated to the store level. We define the store-level log elasticity $\hat{\lambda}_s$ to be the store average of the residuals from a regression of $\hat{\lambda}_{sj}$ on product fixed effects. For the within-chain specification, we regress the store average price \bar{p}_s on $\hat{\lambda}_s$ and chain fixed effects, instrumenting for $\hat{\lambda}_s$ with store average income. For the between-chain specification, we regress the chain mean of \bar{p}_s on the chain mean of $\hat{\lambda}_s$ instrumenting with chain average income.

Figure VIII displays the store-level first stage relationship between the log elasticity term and income, both within chains (Panel A) and between chains (Panel B). It also shows the distribution of estimates of the product-by-product first-stages (Panel C). Consistent with expectations, the relationship of elasticity to income is of similar magnitude within and between chains.

²⁸While the finding that stores in lower-income areas face more elastic demand is intuitive and consistent with prior work, it need not be the case as a matter of theory. For example, it could be that high-income retail markets are more competitive and so the residual demand of stores in these areas is actually more elastic. Table V shows that this competition force is present but that it is small relative to the main effect of income on price sensitivity.

²⁹Recall that \bar{p}_{sj} is the average residual from a regression of $\log(P_{sjt})$ on product-year fixed effects. Since the model implies that equation (2) holds period-by-period, it also holds when averaging across t .

Table VI presents the main results, beginning with the within-chain estimates. The first column shows estimates of equation (4), the most direct test of our model. In this specification, we allow the first-stage coefficient on store income to vary by product. The estimated coefficient is $\hat{\beta} = 0.0482$ (*s.e.* = 0.0137). In the second column, we simplify the model by constraining the first-stage coefficient to be the same across products; the first stage for this case is in Column 4 of Table V. The estimated coefficient increases somewhat to $\hat{\beta} = 0.0852$ (*s.e.* = 0.0210). The third column shows that the estimates are essentially identical when we aggregate to the store level and regress \bar{p}_s on $\hat{\lambda}_s$. The first stage regression for this case is Column 5 of Table V. In all these cases, the within-chain variation in prices is an order of magnitude smaller than the model prediction of $\hat{\beta} = 1$.

The patterns are quite different for the between-chain estimates, which we run at the chain-product level (Column 4) and at the chain level (Column VI 5). To maximize power, we fix the first stage coefficient for these regressions at the values from Columns 4 and 5 of Table V. The coefficients are $\hat{\beta} = 0.7609$ (*s.e.* = 0.2158) and $\hat{\beta} = 0.8339$ (*s.e.* = 0.2315) respectively, an order of magnitude larger than the within-chain versions and statistically indistinguishable from $\hat{\beta} = 1$.

Robustness. Table VII presents a series of robustness checks. To simplify computation, we perform these checks using the store-level aggregated specification of Table VI, Columns 3 and 5.

Our baseline results use per-capita income of the zipcodes of residence as an instrument for the log elasticity $\hat{\lambda}_{sj}$. Row 1 shows that the results are similar if we use as instruments the full set of demographic and competition variables in column 6 of Table V. In row 2 we use the self-reported income of Homescan panelists rather than their zipcode incomes to construct the instrument. The first stage is not as strong, but the within and between results are similar.

Row 3 reports results using the quarterly elasticity (see Section 4.2). The within-chain coefficient is even smaller than the benchmark value. This reflects two offsetting forces: (i) the quarterly elasticities vary less across stores, which would reduce the optimal price variation, but, (ii) the elasticities are lower in magnitude, which makes the log elasticity term $\log\left(\frac{\eta_s}{1+\eta_s}\right)$ more responsive to a given change in elasticity. The smaller coefficient indicates that the second effect dominates. The between-chain relationship is still an order of magnitude larger than the within-chain response.

Row 4 reports results using module-level price and elasticity measures introduced in Section 4.2. The within-chain coefficient is similar to the baseline. We do not estimate a between-chain specification because our module-level price indices are chain-specific.

Rows 5-9 show that the results are also similar (i) for store-brand products, thus at least partially addressing the concern that the uniform pricing is driven by negotiation with the manufacturers of branded products; (ii) for high-revenue products; (iii) for low-revenue products, (iv) weighting prices and elasticities by product revenue; and (v) winsorizing elasticities at -1.5 (instead of -1.2).

In rows 10 and 11 we show that for the drug and mass merchandise chains the within-chain IV slope is larger than for the food chains but still much smaller than one ($\hat{\beta} = 0.2286$ and $\hat{\beta} = 0.2013$).

Finally, in Online Appendix Table 10 and Online Appendix Figure 20 we present OLS estimates not instrumenting the log elasticity term with income. The OLS within-chain price-elasticity relationship is an order of magnitude too flat to be consistent with the model. The between-chain results provide weaker (and not statistically significant) evidence of response to chain-level income, possibly because the elasticity estimates are less comparable across chains than income levels.

Average Prices Paid. Our main question is how the pricing decisions of firms compare to the benchmark of optimal pricing, and we are therefore interested mainly in the prices firms choose to post. With regards to the welfare effects on consumers, however, we also want to consider the average price paid over a longer time horizon. The presence of sales works as a kind of “automatic” price discrimination, guaranteeing that more elastic consumers pay lower prices over the year, even in presence of uniform pricing. This is closely related to the way sales affect responsiveness of average prices to macroeconomic shocks (Chevalier and Kashyap, 2019 and Coibion et al., 2015).

In Online Appendix Table 11, we estimate the model with the average price paid over a year. The within-chain coefficient ($\hat{\beta} = 0.1786$) is larger than the benchmark estimate, but is still five times smaller than under optimal pricing. The between-chain coefficient is similar to the benchmark.

4.4 Loss of Profits

The model allows us to compute the profits lost as a result of nearly-uniform pricing. Under uniform pricing, we assume that a chain r sets a constant price \bar{P}_{rj} across all stores s to maximize its profit

$$\max_{\{\bar{P}_{rj}\}} \sum_{s(r),j} \left[\bar{P}_{rj} k_{sj} \bar{P}_{rj}^{\eta_{sj}} - c_{rj} k_{sj} \bar{P}_{rj}^{\eta_{sj}} \right] - \sum_{s(r)} C_s,$$

where $s(r)$ is the set of stores s belonging to chain r . This leads to the first order condition

$$\sum_{s(r)} k_{sj} \left[(1 + \eta_{sj}) \bar{P}_{rj}^{\eta_{sj}} - c_{rj} \eta_{sj} \bar{P}_{rj}^{\eta_{sj}-1} \right] = 0. \quad (5)$$

Each chain r sets their average price \bar{P}_{rj} equal to the solution to (5). We estimate η_{sj} using the EB-adjusted elasticity estimates $\tilde{\eta}_{sj}$, and the scale factor k_{sj} using the average \hat{k}_{sj} over the weeks t of $Q_{sjt}/P_{sjt}^{\tilde{\eta}_{sj}}$. We estimate marginal costs c_{rj} by plugging in the observed average price \bar{P}_{rj} into $\hat{c}_{rj} = \left[\sum_s \hat{k}_{sj} (1 + \tilde{\eta}_{sj}) \bar{P}_{rj}^{\tilde{\eta}_{sj}} \right] / \left[\sum_s \hat{k}_{sj} \tilde{\eta}_{sj} \bar{P}_{rj}^{\tilde{\eta}_{sj}-1} \right]$, which, by equation (5), is a consistent estimator.

Given these assumptions, we compute the profits for store s and product j under uniform pricing, $\bar{\Pi}_{sj}$, and under optimal pricing, Π_{sj}^* , which we obtain by setting $P_{sj}^* = \hat{c}_{rj} \cdot \tilde{\eta}_{sj} / (1 + \tilde{\eta}_{sj})$.

We then aggregate across all 982 products j for which we have estimated elasticities and scale up this amount to account for the products we do not cover, to estimate the yearly store-level profit variables, $\bar{\Pi}_s$ and Π_s^* . We report the estimated profit loss in dollars per year, $\Pi_s^* - \bar{\Pi}_s$, as well as in fraction of revenue, $(\Pi_s^* - \bar{\Pi}_s)/R_s$, where R_s is the annual revenue for store s in the Nielsen data.³⁰

The results are reported in Panel A of Table VIII. In the median store, the annual profit loss from uniform relative to optimal pricing is \$239,000, or 1.79 percent of revenue, with the losses as high as 3.7 percent of revenue at the 90th-percentile store. Next, we compute the approximate loss from the actual within-chain price-elasticity slope we observe, rather than assuming fully uniform pricing. We use the prices implied by our benchmark IV specification: $P_{sj} = A_{rj}[\tilde{\eta}_{sj}/(1 + \tilde{\eta}_{sj})]^{\beta^{IV}}$, where β^{IV} is the estimate in column 3 of Table VI and A_{rj} is a constant that guarantees that the average price P_{sj} across all stores s in chain is equal to the uniform price \bar{P}_{rj} .³¹ This tempers the losses to a median profit loss of 1.47 percent of revenue. We also compute the loss of profits for state-zone optimal pricing, where prices are set optimally, but are uniform at the state level. This leads to losses comparable to those estimated using the actual price-elasticity slope.

In Panel B of Table VIII, we aggregate the profits for all stores s in a chain r . The annual chain-level profit loss from uniform pricing is \$19.6 million for the median chain and \$111.5 million for the 90th percentile chain. The annual loss from varying prices according to the observed price-elasticity slope is \$16.1m for the median chain and \$91.6m for the 90th percentile chain.

Online Appendix Table 12 shows that the losses of profits computed using the quarterly elasticity and the module-level index elasticity are somewhat smaller than the benchmark estimate. The losses for drug and mass merchandise chains are if anything larger than for food stores.

5 Evidence from Mergers

In this section, we study how prices and quantities change when stores move from one chain to another as part of a merger. The extent to which stores adopt the pricing scheme of the acquiring chain provides a direct test of uniform pricing, as the demand and competitive conditions the store faces are held essentially constant while chain identity changes. It also provides a powerful experiment to estimate long-run price elasticities.

³⁰We report the loss of profit as a fraction of revenue as it does not require us to estimate the fixed costs C_s . To compute the implied loss of profit as a fraction of profits, as we reported in earlier drafts, we follow Montgomery (1997), who estimates gross profit margins of 25 percent and operating profit margins of 3 percent; we therefore assume that $\hat{C}_s = (1 - (3/25)) * \Pi_s^*$. Under these assumptions, the estimated profit loss as fraction of profits is 10.2% for the median store and 11.7% for the median chain.

³¹If we were to use the actual observed prices rather than the prices predicted from the estimated price-elasticity coefficient β^{IV} , the estimated loss in profits would be much larger, since there is a large component of variation in the actual prices that is not explained by the model. See Online Appendix Table 12.

We identify as potential switchers all food stores that change *parent_code* at some point. Recall that these stores are excluded from our main sample. We refer to the first *parent_code* associated with a store as its “old” chain and the *parent_code* it switches to as its “new” chain. While the switch of *parent_code* is recorded, the precise timing is not. To identify the timing, we use a combination of criteria including store closures and sudden changes in the products carried (product assortment). These criteria, which do *not* include pricing information, identify 8 episodes in which 2 or more stores switch chains, for a total of 114 switching stores. In all 8 episodes the assortment of the switching stores mirrors the assortment in the old chain before the acquisition, but switches promptly to the assortment of the new chain afterwards (Online Appendix Figure 21).³²

As an example of the way prices adjust, Figure IX Panel A shows prices of a single product around one merger. The prices in the 15 switching stores during the pre-acquisition period follow exactly the prices of stores in the old chain, and are not synchronized to prices in the acquiring chain. After the acquisition, following a short closure, the switching stores’ prices track precisely the prices in the acquiring chain. While this plot shows the behavior for just one product, it is representative of the pricing for products in this chain.

As more systematic evidence, Panel B averages across the 8 events and all products, and plots the average absolute difference in log prices between the switching stores and (i) stores in the old chain (ii) stores in the new chain. Before the event, the absolute difference is close to zero for the old chain and higher for the new chain, suggesting that the switching stores followed the prices of the old chain before acquisition. After the event, the pattern reverses, suggesting that the switching stores jump to following the prices of the new chain. Online Appendix Figure 22 shows that this pattern holds separately for each of the eight acquisitions. Online Appendix Figure 23 shows that the demographics of customers shopping at switching stores does not change systematically post-acquisition relative to non-switching stores.

Given the sharp changes in pricing, these mergers provide a unique opportunity to estimate long-run price elasticities. Consider the single product whose price is shown in Figure IX Panel A. The new chain generally charges lower prices for this product than the old chain; thus, it is likely to become cheaper in the switching stores after the acquisition. For other products, the new chain charges higher prices; these items are likely to become more expensive in the switching stores. These are long-term, persistent price changes, which quite closely mimic the kind of experiment we would see if a chain switched some of its stores from uniform to flexible pricing.

We use these changes to estimate long-run elasticities. In a first step, we estimate separately

³²In some cases, all the stores in the “old” chain switch, in which case we cannot measure the similarity to the “old” chain post acquisition. See the online appendix in Section A.1.12 for additional detail.

for each acquisition episode and each product j

$$\log(\text{outcome}_{sjt}) = \beta_{sj}^y d_t^{Post} d_s^{Switcher} + \xi_{sj} + \zeta_{jt} + \epsilon_{sjt}, \quad (6)$$

where outcome_{sjt} may be either price P_{sjt} or quantity Q_{sjt} , the stores s include the switching stores and stores in the old and new chain, the weeks t include up to 52 weeks pre-acquisition and up to 52 weeks post-acquisition, the indicator d_t^{Post} denotes weeks post acquisition, and the indicator $d_s^{Switcher}$ denotes stores that switch ownership. In each product-level regression, we control for invariant product-store differences with the fixed effects ξ_{sj} and for common time series changes with product-week fixed effects ζ_{jt} . The regression with price as the outcome yields the estimated event-study change in log prices due to acquisition, $\hat{\beta}_{sj}^P$. For products that are cheaper in the new chain, as in the case of Figure IX Panel A, $\hat{\beta}_{sj}^P$ will be negative, and vice versa for products that are more expensive in the new chain. The regression with quantity as the outcome yields the estimated event-study change in log quantity due to acquisition, $\hat{\beta}_{sj}^Q$.

In the second step, we interpret the regression slope of $\hat{\beta}_{sj}^Q$ on $\hat{\beta}_{sj}^P$ for all the switching stores and all products as the relevant elasticity. As the bin scatter in Figure X Panel A shows, the event-study elasticity is -1.4171 (s.e. 0.0207), smaller (in absolute value) than the average weekly elasticity for the pre-acquisition period for these stores of -1.9905 and similar to the quarterly elasticity -1.6313 . Online Appendix Figure 24 shows a well-behaved pattern of elasticities for each of the 8 acquisitions.

While Figure X Panel A shows the overall event-study elasticity, we estimate a store-specific event-study elasticity η_s^{LR} by running a univariate regression of $\hat{\beta}_{sj}^Q$ on $\hat{\beta}_{sj}^P$ for a store s and all the relevant products j . This allows us to compare in Panel B the store-level long-run elasticities with our benchmark store-level short-run elasticities. (The elasticities are demeaned by acquisition episode.) The two elasticities are clearly correlated, with a slope of 0.80 (s.e. 0.14), not significantly different from one. Further, the long-run elasticities are also highly correlated with the quarterly elasticities and with store-level income (Online Appendix Figures 25 Panels A and B). Finally, the event-study elasticities are similar if we estimate them using as controls (i) only the stores in the old chain or (ii) only the stores in the new chain (Online Appendix Figure 25 Panel C).

While we cannot compute these long-run elasticities for the stores in our main sample, this suggests that the short-run elasticities we use are informative of the more relevant long-term elasticity.

6 What Explains Uniform Pricing?

The results above present a puzzle. The retail chains in our sample include some of the largest consumer-facing firms in the country. Most of them have decades of experience. They manage complex supply chains with tens of thousands of products, and they implement elaborate patterns of sales and promotions whose timing, depth, and duration differ dramatically across products and across time. Yet even though they are changing thousands of prices every week, their prices are essentially uniform across stores—at most varying at the level of large geographic zones. Why do chains not vary prices more and pick up the large gains in profit we predict would result?

Conclusively answering this question is beyond the scope of this paper. However, discussions and interviews with chain managers, industry consultants and analysts have given us at least a qualitative sense of the factors that may be important. Two explanations come up most frequently.

The first is what we will call *managerial inertia*. This encompasses both agency frictions and behavioral factors that prevent firms from implementing optimal policies even though the benefits of doing so exceed the economic costs traditionally defined. The combination of sophisticated sales patterns with uniform pricing across stores seems a close cousin of other cases in which firms devote significant resources to optimizing some variables, while ignoring or devoting limited resources to others. This pattern has been observed in small firms such as a bagel delivery service (Levitt, 2006) and seaweed farmers (Hanna et al., 2014). It has also been observed in much higher stakes contexts, such as rental car companies failing to vary prices with car mileage (Cho and Rust, 2010), baseball teams ignoring statistics like on base percentage (Hakes and Sauer, 2006), fracking operations failing to optimize quantities of sand and water (Covert, 2015), and manufacturing firms failing to implement simple routines that would dramatically improve productivity (Bloom et al., 2013). Hanna et al. (2014) develop a model of this phenomenon (building on Schwartzstein 2014) in which managers devote scarce attention to a subset of choice variables. Such inertia may be exacerbated by herding (Scharfstein and Stein, 1990), as managers who recognize the value of experimenting along some previously unexplored dimension may not internalize all the potential benefits if the experiment succeeds, and may pay a disproportionate share of the costs if it fails.

The industry participants we spoke to frequently pointed to such inertia as an explanation for uniform pricing. One leading consultant suggested that when information systems made sophisticated pricing possible beginning in the 1990s, firms faced a choice between optimizing price levels and sales patterns or optimizing variation across stores and opted to focus on the former. Other participants noted that pricing teams within retailers often have limited sophistication, and that they face a range of organizational barriers—from opposition by teams that would pay costs of introducing flexible pricing but reap few benefits, to constraints imposed by budgeting and report-

ing structures, to herding incentives—that make a major change in pricing difficult. The lack of experimentation with varying prices across stores also means that there is limited awareness of how large the potential gains could be. As one former executive at a large national retailer told us:

Getting buy-in from managers on incremental improvements [to pricing strategy] is hard. You need to show how it will integrate with markdown budgets, with gross margin optimization, with regional profit targets. The manager will say, “Your science does not solve my job. And if it doesn’t solve my job I don’t want your science.” It’s probably true that not varying prices across stores is a missed opportunity. But there are many other missed opportunities and this may be one of the smaller ones.³³

We suspect that managerial inertia may be the most important explanation for uniform pricing.

The second explanation is what we will call *brand image concerns*. This encompasses various mechanisms by which varying prices across stores would lead to negative reactions from consumers that could depress demand for a chain in the long run. This could arise if consumers who observe different prices for the same item in multiple stores perceive this as unfair or a breach of an implicit contract. In a report on UK grocery pricing, the Competition Commission (2003) 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.

Such concerns could be reinforced if prices are explicitly advertised in local or national media, and consumers thus expect all stores’ prices to match what was advertised. A similar mechanism could be at play if chains are posting prices for individual products online, a force shown to be important in other settings by Cavallo (2018b) and Ater and Rigbi (2017). Our conversations confirm that these concerns are frequently cited by managers as a potential risk if chains were to vary prices.

While brand image concerns may be important, there are several reasons why we suspect they may not be the key causal factor in our setting. First, most of the gains from flexible pricing would come from varying prices between stores that are separated geographically, making price comparisons less likely. Very few consumers would know that a store in Akron, OH is charging different prices from a store in Milwaukee, WI. Second, the optimal pricing implied by our model would likely amount to giving discounts to poorer consumers and raising prices on wealthier consumers. This seems less likely to cause a public relations outcry than the reverse. Third, virtually all price advertising is local, and so should not constrain pricing across markets, and most of the retailers we

³³Personal telephone interview, February 8, 2019.

study did not post prices of individual items online during our sample period.³⁴ Finally, retailers in many other sectors do vary prices across stores without appearing to provoke customer outcry, including Starbucks (Luna, 2017), McDonalds, and Burger King (Thomadsen, 2005). Gasoline prices vary widely across locations within a chain, and even grocery chains that sell gasoline vary gasoline prices across their stores.³⁵ None of this precludes some managers viewing brand image risk as a reason to oppose varying prices. Still, we suspect such concerns may be more important as an *ex post* justification rather than the main driver of uniform pricing.

Several other explanations could contribute to uniform pricing but were mentioned less often if at all in our interviews. A first is that committing to uniform pricing may allow chains to soften price competition (Dobson and Waterson, 2008 and Adams and Williams, 2019). We cannot rule out this being an important factor, though in evidence discussed below we fail to find evidence that the extent of uniform pricing varies with the extent of competition. A second is traditional menu costs. We expect such costs to be a small factor given the frequency with which stores change prices over time. A third is IT or engineering costs of varying prices across stores. Though sophisticated versions of non-uniform pricing could be costly to optimize, our results suggest large profit gains to even coarse schemes that adjust all prices in a store by a single factor based on local income. The engineering costs of such a scheme are likely small relative to the predicted gains.

While we cannot empirically distinguish the various explanations, looking at the correlates of price uniformity across chains offers some suggestive evidence. In Table IX we regress the chain-by-chain coefficients on the log elasticity term in equation (4) — a measure of the extent to which each chain deviates from uniform pricing — on chain characteristics.³⁶ Chains with larger gains from varying prices—larger chains and chains with stores serving more diverse income levels—vary prices more (Columns 1-2). When we combine such factors into a measure of profit loss from uniform pricing (Column 3), chains with greater loss vary prices more. This effect loads on the total (log) profit loss rather than the percentage loss. These results are consistent with the managerial inertia mechanism to the extent that inertia functions like a chain-level fixed cost of implementing non-uniform pricing. It is also consistent with some role for IT or engineering costs. Finally, neither the share of a chain’s stores with nearby competitors nor the share with nearby same-chain stores are significant predictors of uniformity (Column 4).³⁷ In Online Appendix Table 14 we show that the

³⁴If online posting were a key driver of uniform prices, we might expect to see chains moving to uniformity as they start posting online. As Online Appendix Table 14 shows, there is no significant increase in the extent of uniform pricing over our sample.

³⁵As an illustrative example, we examined the gasoline prices for three Safeway stores in different cities using gasbuddy.com. For these stores, we found regular gasoline prices of \$2.94, \$3.21, and \$3.07.

³⁶For each chain, we regress the store-level log price on the store-level income, yielding the coefficients in Online Appendix Figure 7. We then divide these coefficients by the first-stage coefficient in Table V, column 5.

³⁷In Online Appendix Table 13, we document similar results using the chain-level quarterly absolute log price

price-log elasticity relationship does not differ for stores with no competing stores nearby. These facts provide some evidence against either tacit collusion or brand image concerns (where the key issue is consumers directly observing prices across stores) playing a driving role.

7 Implications

In this section, we consider four broader economic implications of uniform pricing.

7.1 Inequality

Allcott et al. (2019a) and Jaravel (2019) both highlight the way variation in prices can potentially exacerbate inequality. Uniform pricing by chains may have a substantial effect in this sense, since optimal prices co-vary positively with income and thus uniform pricing will tend to lower the prices paid by the rich and raise the prices paid by the poor.

To quantify this effect, Figure XI shows the predicted relationship between local income and prices for a hypothetical representative product sold by every chain with the same marginal cost c under (i) uniform pricing and (ii) flexible pricing. We also present the relationship between local income and the average yearly price paid, to incorporate variation in prices paid due to endogenous substitution by consumers as discussed in Section 4.3.

We compute the uniform price, $p_r^{Uniform}$, as in equation (5), taking the predicted elasticity based on the income first stage. We compute optimal flexible prices as $p_s^* = \log(c) + \log(\frac{\eta_s}{1 + \eta_s})$, where the log elasticity term is also replaced with the predicted value from the first stage. We compute the yearly price paid using $p_s^{Yearly} = p_r^{Uniform} + \beta^{Yearly}(\lambda_s - \bar{\lambda}_r)$, where β^{Yearly} is the estimated slope using price paid in Online Appendix Table 11, column 1 and $\bar{\lambda}_r$ is the average of λ_s within chain r . We present additional details in Section A.1.13 in the online appendix.

Figure XI shows that uniform prices are somewhat increasing in store income because low-income and high-income areas are served by distinct chains. In the flexible pricing counterfactual, prices are much more responsive to local income. The slope of yearly average prices is only slightly steeper than under uniform pricing, suggesting that consumer substitution does not substantially change the implications of uniform pricing. These patterns are similar if we use the quarterly or index elasticities instead of the weekly elasticities (Online Appendix Figure 26).

This calibration suggests important implications for inequality. Consumers of stores in the lowest income decile pay about 3.4 log points higher prices under the yearly price paid benchmark than under flexible pricing, while consumers of stores in the top income decile pay about 6.5 log difference, as in Figure III Panel A.

points lower prices. We emphasize that this is not a complete analysis, and these patterns could be moderated or strengthened by unmodeled competitive responses including entry and exit.

7.2 Response to Local Shocks

A second implication of our findings relates to the response of prices to local shocks. Benchmark models assume that when a negative shock to income or wealth hits consumers in an area, the impact on welfare will be offset to some extent by reductions in local retail prices. Similarly, any shocks that increase local costs would tend to be reflected in higher prices. Such responses will be dampened by uniform pricing, especially if the geographic area affected by the shock is small.

As an illustration, in Table X we consider stores facing a negative shock to local consumer income of \$2,000. We translate this into a change in the log elasticity term using the coefficient from Table V, column 5. We consider three cases: (i) a nationwide shock that affects all stores in the country; (ii) a shock that affects all stores in a given state; (iii) a shock that affects all stores in a given county. For each shock, we compute the response of log prices under flexible pricing (column 1), under uniform pricing (column 2), and under the average yearly price paid scenario (column 3), following the same approach as in Section 7.1. For (ii), we compute the results separately for each of the 48 states and then average the results. For (iii), we average analogously across counties.

The dampening effect of uniform pricing on local price responses is dramatic. Under flexible pricing, the \$2,000 income shock leads to 1.01 log points lower prices regardless of the geographic level. When the shock is national, the response is similar (though not identical) under uniform pricing, since it is optimal for all stores in a chain to adjust their prices by roughly the same amount. For state and county-level shocks, however, the responses are far smaller. Under uniform pricing, the average state-level shock reduces prices by only 0.32 log points, and the average county-level shock by only 0.04 log points. The responses of yearly average prices to local shocks is larger, but still highly attenuated relative to flexible pricing.³⁸

Panels B and C explore this further and consider a shock affecting California or Nevada. A negative shock in California leads to a price decrease in the California stores of 0.72 log points under uniform pricing, still dampened but much closer to the response under optimal pricing, given that the chains operating in California have a majority of their stores in the state. In contrast, an equal-sized shock in Nevada lowers prices in Nevada stores by only 0.16 log points under uniform pricing, since many Nevada chains operate most of their stores elsewhere. This example also illustrates potential price spillovers. Under uniform pricing, a shock in California causes prices in Nevada to

³⁸The patterns for local price response to shocks are similar if we use the quarterly or price index elasticities instead of the benchmark elasticity, with a larger overall response to the negative shock (Online Appendix Table 15).

decrease by 0.38 log points, as some Nevada stores are part of the same chains as the California stores. In contrast, the impact of a shock in Nevada on California is negligible at 0.04 log points.

Some recent research provides evidence consistent with these simulations. Gagnon and López-Salido (forthcoming) show that large localized demand shocks due to labor conflicts, population displacement, and weather events translate into minimal changes in local supermarket prices. Cawley et al. (2018) show that pass-through of a Philadelphia soda tax into supermarket prices was smaller at chain stores than at independent retailers. Leung (2018) shows that the pass-through of a local minimum wage change into prices is moderate for supermarkets and negligible for drug and mass merchandise stores, citing the uniform pricing documented here as a likely explanation. In contrast, Stroebel and Vavra (forthcoming) document relatively large effects of house prices on local retail prices; one explanation is that house price changes are correlated at the regional level and so translate into larger impacts than more localized shocks, consistent with our predictions.³⁹

7.3 Merger Enforcement

A third implication relates to the analysis of horizontal mergers by competition authorities. The standard approach to analyzing such mergers in the US proceeds in three steps (Balto, 2001): (1) the territory in which the merging firms operate is divided into separate markets; (2) the change in concentration and other relevant competitive conditions is analyzed separately in each market; (3) if the merger would reduce competition in one or more markets to a sufficient degree, the merger is blocked or the parties are required to divest stores in the affected markets. This analysis is economically coherent under the assumption that chains set prices independently in each market.

Uniform pricing changes this picture dramatically. If firms charge uniform prices, separate markets can no longer be analyzed in isolation. The impact of a merger in a market will depend not only on the change in concentration in that market but also on the change in concentration in all other markets where the firms operate. Even markets where the firms operate but do not compete will be affected, and will also alter the impact of the merger in the markets where they do compete. To be clear, we do not establish here how common uniform pricing is in the cases the competition authorities consider; we only argue that the analysis of mergers needs to account for uniform pricing in cases where it does occur.

When the existence or lack of uniform pricing is noted in merger cases, it is often used as a *test* of the extent of competition. For example, in evaluating a proposed merger between Whole

³⁹There are several other reasons why past work may have found larger responses to local shocks beyond broad geographic scope. First, some chains do vary prices by geographic zone so prices are not perfectly uniform. Second, some prior work looks at price indices that also incorporate variation in the set of available products. Third, average prices do vary within chain as a result of compositional effects.

Foods and Wild Oats supermarkets (FTC v. Whole Foods, 2007), the District Court cited the fact that “Whole Foods prices are essentially the same at all of its stores in a region, regardless of whether there is a Wild Oats store nearby” as evidence that Whole Foods and Wild Oats are not each others’ primary competitors and that the relevant market must include other supermarkets (Varner and Cooper, 2007). Such conclusions are invalid if firms charge uniform prices.

7.4 Incidence of Trade Costs

A fourth implication relates to the incidence of trade costs. A large literature estimates trade costs by examining differences in the prices of specific products at geographically separated stores (see surveys 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 for the source locations of the products and allowing for spatially varying 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. The adjustment for markups would also be affected by uniform pricing.⁴⁰

Uniform pricing also affects the real incidence of trade costs. Just as it tends to lower prices in high-income areas and raise them in low-income areas, it will tend to raise prices in locations close to where products are produced and lower them in remote locations. It thus leads to a relative reduction in the benefits of trade for those close to ports or origin locations and a relative increase in these benefits for those far away.

8 Conclusion

In this paper, we show that most large US food, drugstore, and mass merchandise chains set uniform or nearly-uniform prices across their stores, and that limiting price discrimination in this way costs firms significant short-term profits. The result of nearly-uniform pricing is a significant dampening of price adjustment, which has important implications for the extent of inequality, the pass-through of local shocks, the analysis of mergers, and the incidence of trade costs.

⁴⁰Atkin and Donaldson (2015) infer the extent of market power from the observed pass-through 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 so they use the origin retail price as a proxy. Uniform pricing will tend to increase the estimated pass-through, as it increases the correlation between changes in retail prices in origin and destination markets that are served by stores from the same chains. It will therefore tend to reduce the level of estimated markups, while (correctly) implying less variation in markups across space.

References


- Adams, Brian and Kevin R. Williams. 2019. "Zone Pricing in Regional Oligopoly." *American Economic Journal: Microeconomics*, 11(1): 124-56.
- Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dube, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. 2019a. "Food Deserts and the Causes of Nutritional Inequality." Working Paper.
- Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky. 2019b. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." Working paper.
- 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*, 90(C): 64-83.
- Anderson, James E. and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature*, 42(3): 691-751.
- Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters. 2017. "Leveraging Lotteries For School Value-Added: Testing and Estimation." *Quarterly Journal of Economics*, 132(2): 871-919.
- Ater, Itai and Oren Rigbi. 2018. "The Effects of Mandatory Disclosure of Supermarket Prices." Working paper.
- Atkin, David and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." *NBER Working Paper No. 21439*
- Baker, Jonathan B. 1999. "Econometric Analysis in *FTC v. Staples*." *Journal of Public Policy & Marketing*, 18(1):11-21.
- Balto, David A. 2001. "Supermarket Merger Enforcement." *Journal of Public Policy and Marketing*, 20(1):38-50.
- Beraja, Martin, Erik Hurst, and Juan Ospina. 2019. "The Aggregate Implications of Regional Business Cycles." Working paper.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4):841-90.
- Bils, Mark and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy*, 112(5): 947-985.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, John Roberts. 2013. "Does Management Matter? Evidence from India." *Quarterly Journal of Economics*, 128(1): 1-51.
- Bloom, Nicholas and John Van Reenen. 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *Quarterly Journal of Economics*, 122(4): 1351-1408.
- Campbell, Jeffrey R. and Benjamin Eden. 2014. "Rigid Prices: Evidence from U.S. Scanner Data." *International Economic Review*, 55(2): 423-442.
- Cavallo, Alberto. 2017. "Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers." *American Economic Review*, 107(1), 283-303.

- Cavallo, Alberto. 2018a. "More Amazon Effects: Online Competition and Pricing Behaviors." *Jackson Hole Economic Symposium Conference Proceedings (Federal Reserve Bank of Kansas City)*.
- Cavallo, Alberto. 2018b. "Scraped Data and Sticky Prices." *Review of Economics and Statistics* 100(1): 105-119.
- Cavallo, Alberto, Brent Nieman, and Roberto Rigobon. 2014. "Currency Unions, Product Introductions, and the Real Exchange Rate." *Quarterly Journal of Economics*, 129(2): 529-595.
- Cawley, John, David Frisvold, Anna Hill, and David Jones. 2018. "The Impact of the Philadelphia Beverage Tax on Prices and Product Availability." *NBER Working Paper No. 24990*.
- Chevalier, Judith and Anil Kashyap. 2019. "Best Prices: Price Discrimination and Consumer Substitution." *American Economic Journal: Economic Policy*, 11(1): 126-159.
- 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.
- Cho, Sungjin and John Rust. 2010. "The Flat Rental Puzzle." *The Review of Economic Studies*, 77(2): 560-594.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong. 2015. "The Cyclicity of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review*, 105(3): 993-1029.
- Covert, Thomas R. 2015. "Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale." Working paper.
- Competition Commission. 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.
- DellaVigna, Stefano and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence" *Quarterly Journal of Economics*, 119(2), 353-402.
- Dobson, Paul W. and Michael Waterson. 2008. "Chain-Store Competition: Customized vs. Uniform Pricing." *Warwick Economic Research Papers No. 840*.
- Fackler, Paul L. and Barry K. Goodwin. 2001. "Spatial Price Analysis." *Handbook of Agricultural Economics, vol. 1*: 971-1024.
- Federal Trade Commission. 2010. *Horizontal Merger Guidelines*.
- United States District Court, District of Columbia Circuit. 2007. "Federal Trade Commission v. Whole Foods Market, Inc." 502 F.Supp.2d 1 (2007), No. 07-1021(PLF).
- United States District Court, District of Columbia Circuit. 1997. "FTC v. Staples, Inc. and Office Depot, Inc." *FTC File No. 9710008. Civil Action No. 197CB00701*.
- Gagnon, Etienne and David J. López-Salido. Forthcoming. "Small Price Responses to Large Demand Shocks." *Journal of the European Economic Association*.

- George, Lisa and Joel Waldfogel. 2003. "Who Affects Whom in Daily Newspaper Markets?" *Journal of Political Economy*, 111: 765–784.
- 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.
- 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.
- Hakes, Jahn, K., and Raymond D. Sauer. 2006. "An Economic Evaluation of the *Moneyball* Hypothesis." *Journal of Economic Perspectives*, 20 (3): 173-186.
- 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.
- Hausman Jerry A. 1996. "Valuation of new goods under perfect and imperfect competition." In *The Economics of New Goods*, ed. TF Bresnahan, RJ Gordon, pp. 209–48. Chicago: Univ. Chicago Press
- Heidhues, Paul and Botond Koszegi. 2018. "Behavioral Industrial Organization," in *Handbook of Behavioral Economics*, Vol. 1, ed. D Bernheim, S DellaVigna, and D Laibson, Elsevier, 517-612.
- Hendel, Igal and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica*, 74(6): 1637-1673.
- Hitsch, Guenter, Ali Hortaçsu, and Xiliang Lin. 2017. "Prices and Promotions in U.S. Retail Markets: Evidence from Big Data." *Booth Working paper 145*.
- 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.
- 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.
- Hwang, Minha, Bart J. Bronnenberg, Raphael Thomadsen. 2010. "An Empirical Analysis of Assortment Similarities Across U.S. Supermarkets." *Marketing Science*, 29(5): 858-79.
- Jacob, Brian A., and Larse Lefgren. 2008. "Can Principals Identify Effective Teachers? Evidence on Subjective Performance Evaluation in Education." *Journal of Labor Economics*, 2008(26): 101-36.
- Jaravel, Xavier. 2019. "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector." *Quarterly Journal of Economics*. 134(2):715-83.
- Kane, Thomas J., and Douglas O. Staiger. 2008. "Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation." *NBER Working Paper No. 14607*.
- Kaplan, Greg and Guido Menzio. 2015. "The Morphology of Price Dispersion." *International Economic Review*, 56(4): 1165-1206.
- Leung, Justin. 2018. "Minimum Wage and Real Wage Inequality: Evidence from Pass-Through to Retail Prices." Working paper.

- Levitt, Steven D. 2006. "An Economist Sells Bagels: A Case Study in Profit Maximization." *NBER Working Paper No. 12152*.
- Luna, Nancy. 2017. "Starbucks quietly raises Southern California prices up to 30 cents." *Orange County Register*.
- 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.
- McMillan, Robert Stanton. 2007. "Different Flavor, Same Price: The Puzzle of Uniform Pricing for Differentiated Products." Working paper.
- Miravete, Eugenio J., Katja Seim, and Jeff Thurk. Forthcoming. "One Markup to Rule Them All: Taxation by Liquor Pricing Regulation" *American Economic Journal: Microeconomics*.
- Montgomery, Alan L. 1997. "Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data." *Marketing Science*, 16(4): 315-337.
- Nakamura, Emi. 2008. "Pass-Through in Retail and Wholesale." *American Economic Review Papers and Proceedings*, 98(2): 430-437.
- Nakamura, Emi and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics*, 123(4): 1415-1464.
- Nevo Aviv. 2001. "Measuring market power in the ready-to-eat cereal industry." *Econometrica*, 69: 307-42.
- 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.
- Scharfstein, David S., and Jeremy C. Stein. 1990. "Herd Behavior and Investment." *American Economic Review*, 80: 465-479.
- Schwartzstein, Joshua. 2014. "Selective Attention and Learning." *Journal of the European Economic Association* 12(6): 1423-52.
- 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.
- Stroebel, Johannes and Joseph Vavra. Forthcoming. "House Prices, Local Demand, and Retail Prices." *Journal of Political Economy*.
- Thomadsen, Raphael. 2005. "The Effects of Ownership Structure on Prices in Geographically Differentiated Industries." *RAND Journal of Economics*, 36(4): 908-929.
- Thomassen, Øyvind, Howard Smith, Stephan Seiler, and Pasquale Schiraldi. 2017. "Multi-Category Competition and Market Power: A Model of Supermarket Pricing." *American Economic Review*, 107(8): 2308-2351.
- Varner, Carlton, and Heather Cooper. 2007. "Product Markets in Merger Cases: The *Whole Foods* Decision." *The Antitrust Source*, October 2007.

Panel A. Single Chain, Prices of a Single Product in Orange Juice Category



Panel B. Single Chain, Prices of Products in Five Categories

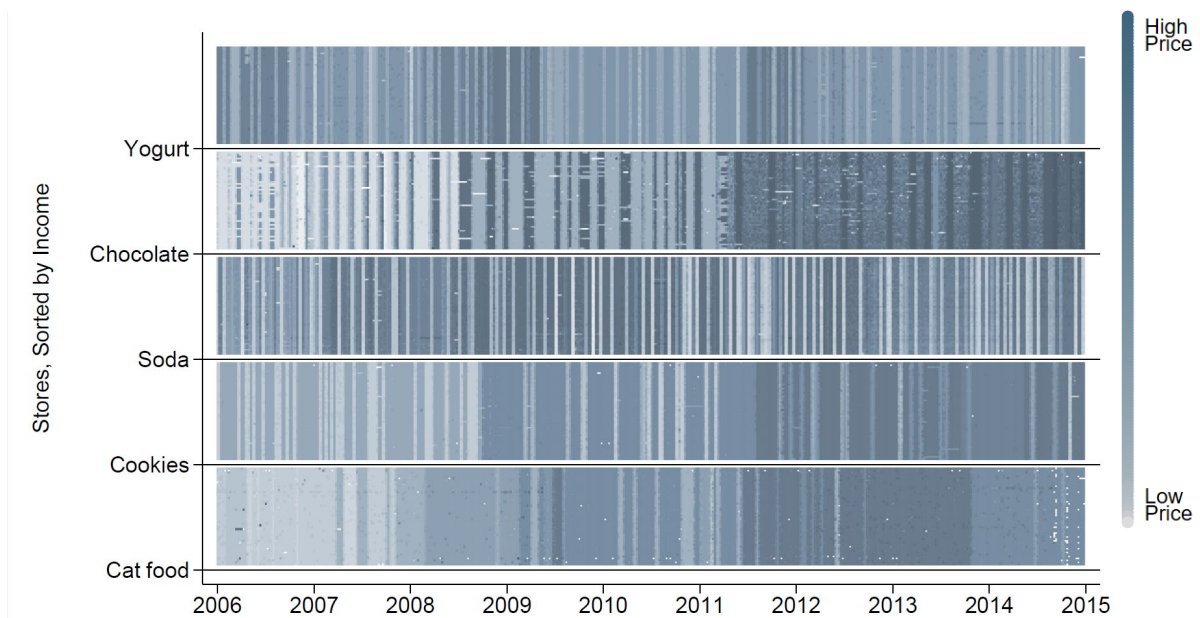



Figure I

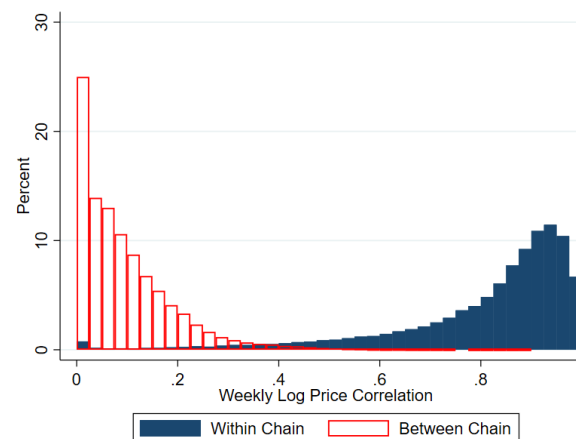
Examples of Uniform Pricing

Notes. Figures depict log price in store s and week t for a particular product j . To facilitate comparison across products, we standardize prices by demeaning the log price by the average log price across all stores s in all chains. Darker colors indicate higher price and the figure is blank if price is missing. Each column is a week. Each row is a store, and stores are sorted by store-level income per capita. In Panel B, the same 50 stores appear for each product.

Panel A. Quarterly Absolute Log Price Difference



Panel B. Weekly Log Price Correlation



Panel C. Weekly Share of Identical Prices

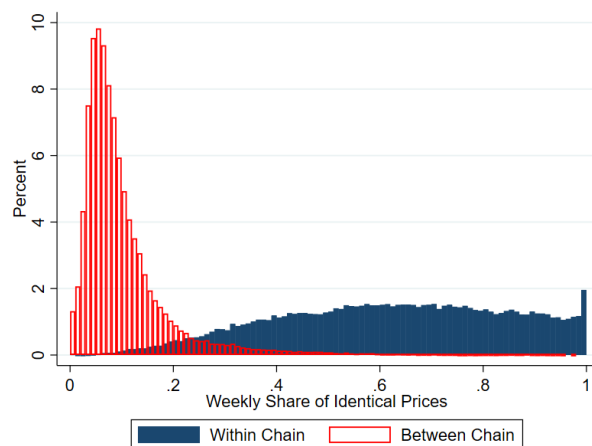



Figure II

Similarity in Pricing Across Stores: Same-Chain Comparisons vs. Different-Chain Comparisons

Notes. Each observation in the histograms is a chain-UPC representing the average relationship between up to 200 store-pairs belonging to each chain. The “same chain” pairs are formed from stores belonging to the same chain; the “different chain” pairs are formed from stores in different chains, requiring in addition that the two stores do not belong to the same *parent_code*. Panel A displays the distribution of the average absolute difference in log quarterly prices between two stores in a pair, winsorized at 0.3. Panel B displays the distribution of the correlation in the weekly (demeaned) log prices between two stores, winsorized at 0. Panel C displays the share of prices in a pair of stores that are within 1 log point of each other.

Panel A. Quarterly Similarity in Pricing versus Weekly Correlation of Prices, by Chain



Panel B. Within-State versus Between-State Quarterly Absolute Log Price Difference, by Chain






Figure III
Similarity in Pricing, Chain-Level Measure

Notes. Each observation in Panels A and B is a chain, with circles representing food stores, diamonds representing mass merchandise stores, and squares representing drug stores. In Panel A, for each chain, we plot the average across all within-chain pairs of the quarterly absolute difference in log price (as in Figure II Panel A) and of the weekly correlation in log price (Figure II Panel B). We compute the averages using up to 200 pairs of stores within a chain. In Panel B, each observation is a chain that operates at least three stores in each of at least two states. Chains that differentiate pricing geographically—with difference between across-state and within-state quarterly absolute price difference greater than 0.01—are denoted with solid markers.

Panel A. Within Chain



Panel B. Between Chains



Panel C. Within-Chain and Between-Chain Coefficients by UPC

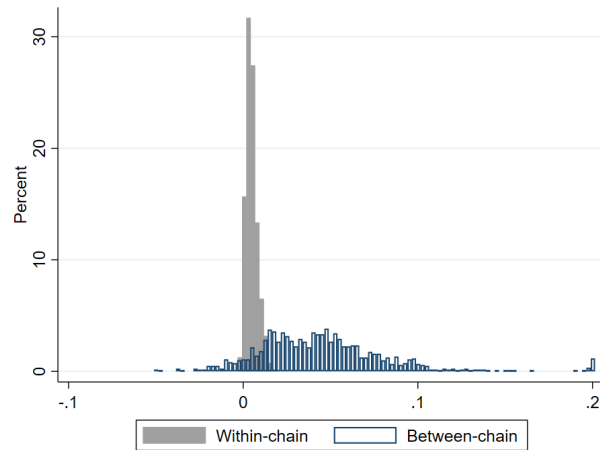




Figure IV
Price vs. Income

Notes. Panel A is a binned scatterplot with 50 bins of the residual of log price in store s on the residual of income in store s . Residuals are after removing chain fixed effects. Panel B is a scatterplot of average price on average income at the chain level for the food stores, with the labels indicating a chain identifier. Panel C shows the distribution of the coefficient, within and between chains, UPC-by-UPC. The figures report the coefficients of the relevant regressions, with standard errors clustered by *parent_code*. Analytic weights equal to the number of stores in each aggregation unit are used for the regression in Panel B.

Panel A. Data from Nielsen: Average Weekly Log Price



Panel B. Data from Major Grocer: Average Weekly Log Price



Panel C. Data from Major Grocer: Nonsale Log Price





Figure V


Price vs. Store-Level Income Within State: Investigation Using Major Grocer's Data

Notes. Panels A-C are binned scatterplots of the residuals of log price in store s on the residuals of income in store s . Residuals are after removing chain-state fixed effects. The figures report the coefficients and robust standard errors of the relevant regressions. Axes ranges have been chosen to make the slopes visually comparable. Prices are demeaned by UPC. Products were selected using the following criteria: products must be sold for at least 40 weeks in 99% of Major Grocer stores.

Panel A. Within Chain



Panel B. Between Chain



Panel C. Within-Chain and Between-Chain Coefficients by Submodule





Figure VI

Assortment Price Index vs. Store-Level Income (Food Stores)

Notes. These figures show the relationship of which products stores carry, a non-price measure of store-level decision making, vs. income. The assortment price index is constructed as follows: first, for each product that is in the top 20% of national units sold, we calculate the average national log unit price. Second, we divide each module into up to five sub-modules based on product package size. Third, we calculate the average log national unit price for each store-sub-module-year. We collapse this measure to a store level Assortment Price Index (API) by averaging over the sub-module-years for each store. Panel A is a binned scatterplot with 50 bins of the residual of API in store s on the residual of income in store s . Residuals are after removing chain fixed effects. Panel B is a scatterplot of average API on average income at the chain level for the food stores, with the labels indicating a chain identifier. Panel C shows the distribution of the coefficient, within and between chains, submodule-by-submodule for a total of 388 submodules sorted by package size. The figures report the coefficients of the relevant regressions, with standard errors clustered by *parent_code*. Analytic weights equal to the number of stores in each aggregation unit are used for the regression in Panel B.

Panel A. Specification Check: Linearity of Log Quantity and Log Price for Two Income Levels



Panel B. Distribution of Elasticity Estimates: Store-Level Average

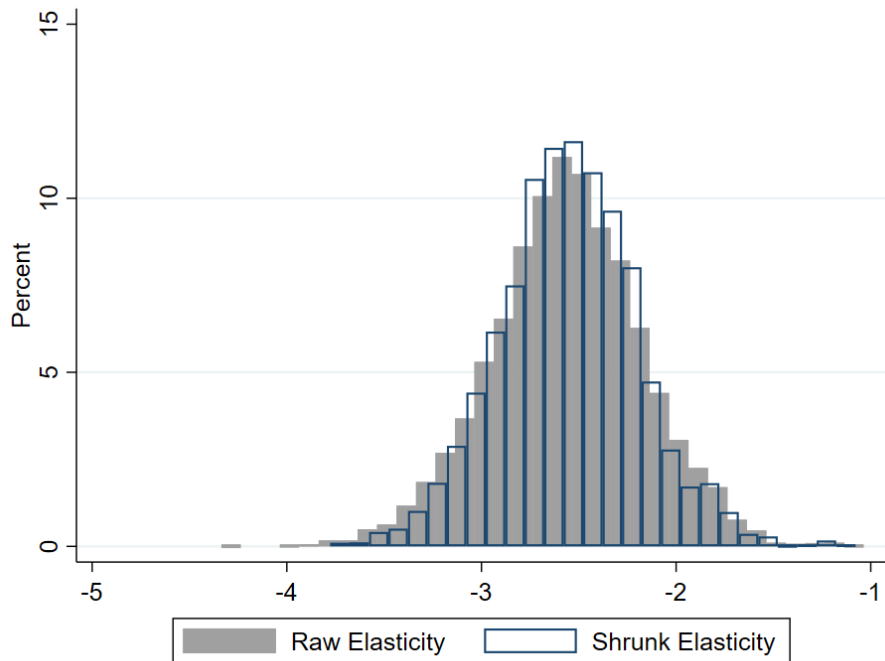




Figure VII
Elasticity Estimates and Validation

Notes. Panel A is a binscatter of log P and log Q for 25 randomly sampled UPCs after residualizing for week-of-year and year fixed effects. One set of observations includes the 50 stores nearest to the \$20,000 income level; the second set includes 50 stores nearest to the \$60,000 income level. The slope of each line is the elasticity for that group of stores and products. Panel B plots the distribution of the store-level elasticity, shrunk and raw. Store-level elasticities are mean-zero. We have recentered these elasticities using the average of the store-UPC-level elasticities to show level.

Panel A. Within Chain



Panel B. Between Chain



Panel C. Distribution of Within Chain vs. Between Chain Coefficients, UPC-by-UPC

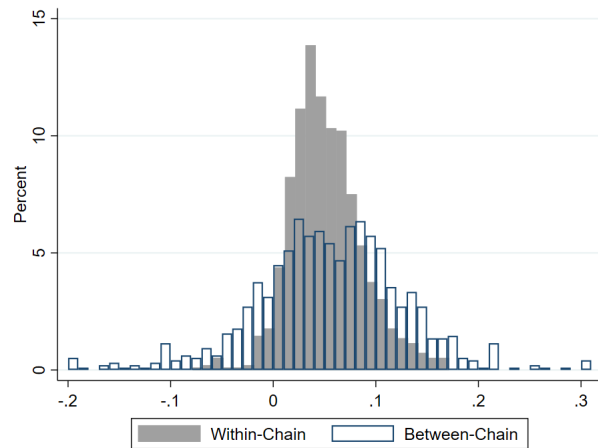



Figure VIII

Log Elasticity versus Store-Level Income

Notes. Panel A is a binned scatterplot with 50 bins of the residual of $\log\left(\frac{\eta}{1+\eta}\right)$ store s on the residual of income in store s . Residuals are after removing chain fixed effects. Panel B is a scatterplot of average $\log\left(\frac{\eta}{1+\eta}\right)$ on average income at the chain level, with the labels indicating a chain identifier. Panel C shows the distribution of the coefficients, within-and-between-chain, UPC-by-UPC.

Panel A. Quarterly Log Prices for Merger 5, Selected Product



Panel B. Event Study of Weekly Price Changes for All Switching Stores


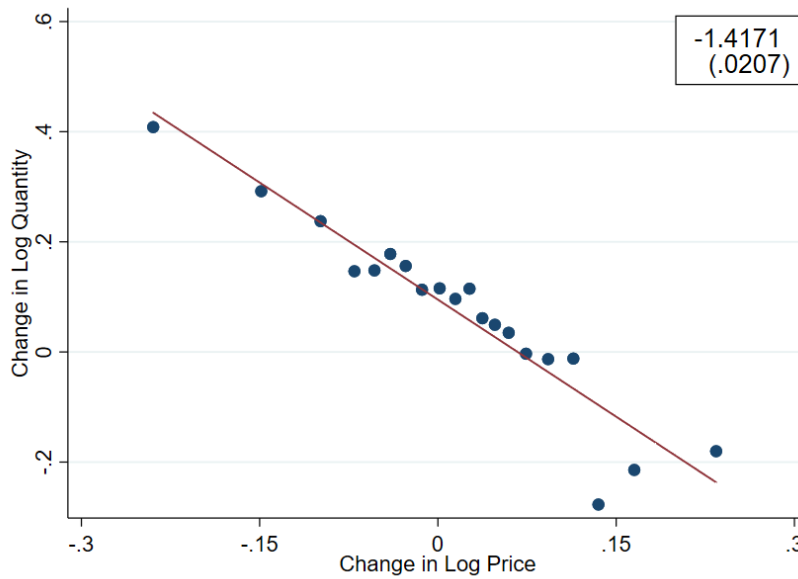


Figure IX

Event-Study Graph of Pricing in Stores that Change Owner

Notes. Panel A is a quarterly log price series of one UPC in Merger #5. Old and new chain price series are leave-out means of stores in the new and old *parent_code* that do not switch, respectively. The vertical lines show the window where the switch has occurred. The y-axis scale is redacted to prevent product identification. Panel B shows the absolute log price difference of all switching stores from their old and new chains. We standardize the weeks according to the switch time to make different switches comparable. The increased jaggedness of the solid line over time is caused by some old chains closing down and changes in product assortment.

Panel A. Acquisition-Induced Change in Quantity vs. Change in Price



Panel B. Store-Level Estimated Long-run Elasticity Vs. Weekly Elasticity

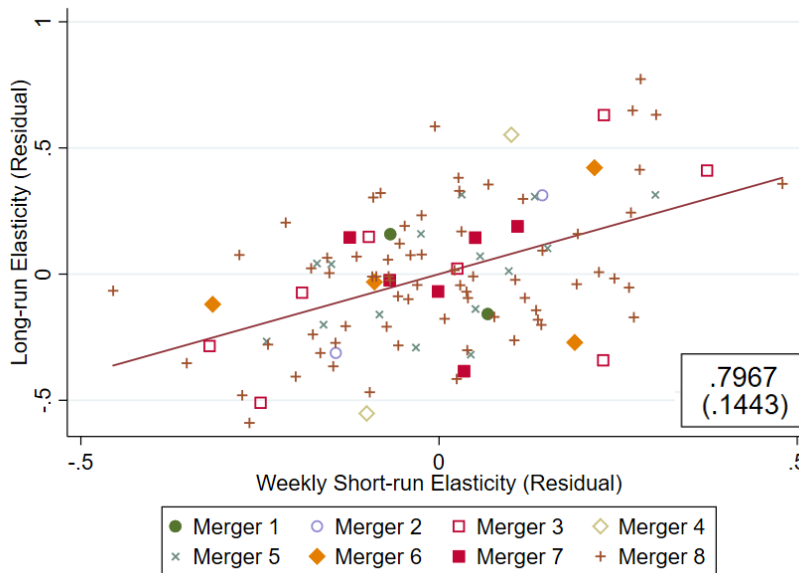


Figure X

Long-Term Elasticities from Store Acquisitions

Notes. Panel A shows a binscatter plot of changes in log quantity versus changes in log price after a store has been acquired by a different chain for the universe of products. Formally, for a given product, changes in log quantity and price are obtained as OLS coefficients on an interaction of a dummy for the post-acquisition period with the particular store indicator, controlling for store and date fixed effects. Control stores include up to 200 stores in the original and acquiring chain which exhibit at least 50% availability for the particular product before and after the acquisition. The long-run elasticity is obtained as the coefficient in an OLS regression of changes in log quantity on changes in log price for all stores and products. Panel B shows store-level long-run elasticities versus (weekly) short-run elasticities after demeaning both variables on the merger-level. For each store, the long-run elasticity is computed as in Panel A, restricting the observations to the particular store. Short-run elasticities are estimated using data prior to the merger only.

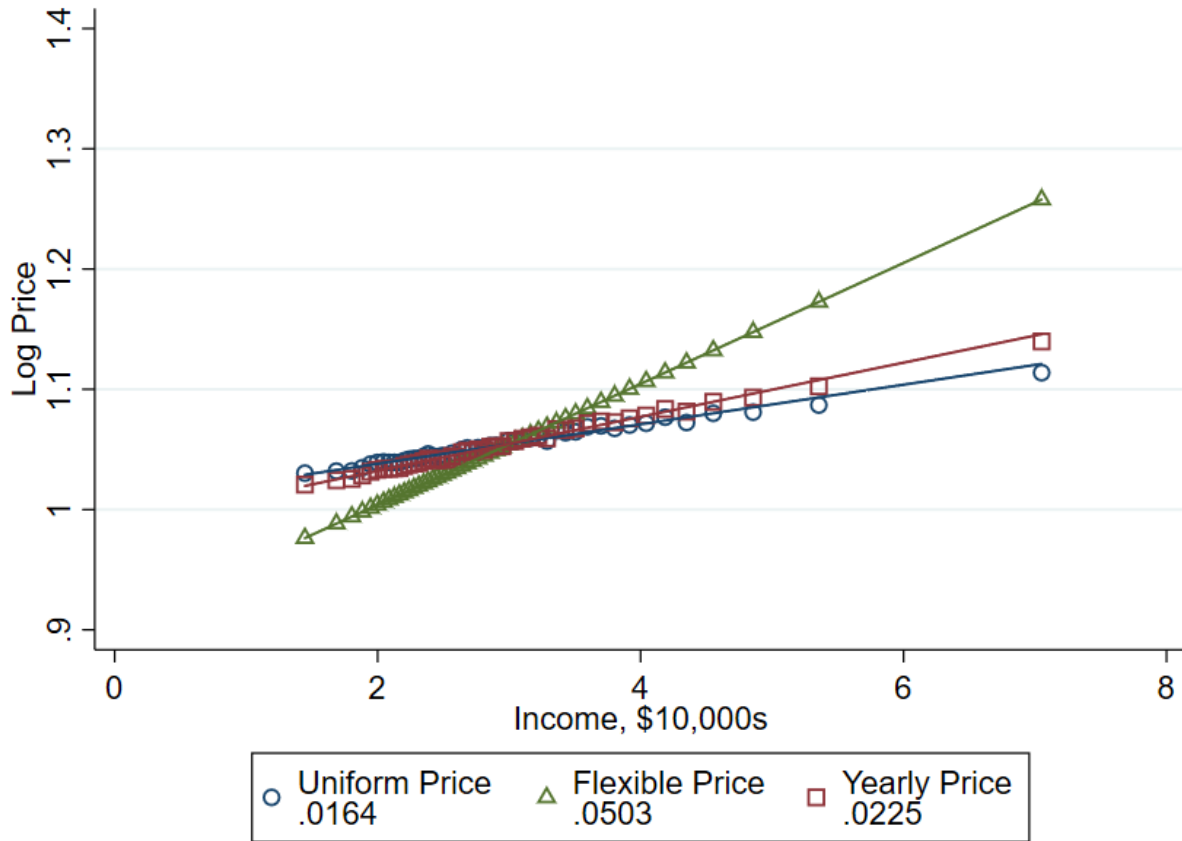


Figure XI
Price Rigidity and Inequality: Prices in Areas with Different Income.

Notes. In this figure, using a representative product, we plot binned scatterplots with 50 bins of store-level uniform price and counterfactual log prices (under flexible pricing) versus store-level income. The elasticity of each store is according to the predicted elasticity using the coefficients of Table 5 Column 5. The counterfactual price assumes flexible pricing, that is $\log P^* = \log\left(\frac{\eta}{1+\eta}\right) + \log(c)$ using the estimated elasticity for each store s and a constant marginal cost for all chains. The uniform price is the profit maximizing uniform price set for each chain. Yearly Price perturbs the optimal uniform price within each chain by the yearly IV coefficient of prices on elasticity (Table 8 Column 1) which is meant to include the “automatic stabilizer” effect of intertemporal substitution due to sales. Chain average Price Paid and chain average Uniform Price are equal. We use the median of the estimated marginal costs as the marginal cost for the representative product.

TABLE I
SAMPLE FORMATION AND SUMMARY STATISTICS

Panel A: Sample Formation				
	No. of Stores	No. of Chains	No. of States	Total Yearly Revenue
Initial Sample of Stores	38,539	326	48+DC	\$224bn
Store Restriction 1. Stores do not Switch Chain, >= 104 weeks	24,489	119	48+DC	\$193bn
Store Restriction 2. Store in HMS dataset	22,985	113	48+DC	\$192bn
Chain Restriction 1. Chain Present for >= 8 years	22,771	83	48+DC	\$191bn
Chain Restriction 2. Valid Chain	22,680	73	48+DC	\$191bn
Final Sample, Food Stores	9,415	64	48+DC	\$136bn
Final Sample, Drug Stores	9,977	4	48+DC	\$21bn
Final Sample, Merchandise Stores	3,288	5	48+DC	\$34bn

Panel B: Store Characteristics, Food Stores				
	Mean	25th	Median	75th
Average per-capita Income	\$29,770	\$22,770	\$27,420	\$34,300
Percent with at least Bachelor Degree	10.2%	5.0%	8.3%	13.3%
Number of HMS Households	31.3	16	27	42
Number of Trips of HMS Households	1382	522	1067	1860
Number of Competitors within 10 km	3.6	0	1	5

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: Product Characteristics, Food Stores				
	No. of Products	Yearly Revenue Covered and Share of Total	Weekly Average Price	Weekly Availability
Benchmark Sample (40 Modules)	1,365	\$10.8bn 22.0%	\$2.86 (0.134)	90.70% (6.39)
Sample with Estimated Elasticities (38 Modules)	982	\$9.20bn 18.7%	\$2.65 (0.152)	94.04% (4.26)
Sample for Assortment Analysis (40 Modules)	43,279	\$44.5bn 90.5%		

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 Yearly Revenue is the yearly average total revenue for products in that sample in the Nielsen RMS dataset. To get yearly revenues, we take total revenue and divide by the number of years a store appears in the sample. The Share of Total is computed as the share of the Yearly Revenue figure to the yearly revenue for all products in the 40 Nielsen modules we consider (\$49.2bn in revenue per year). For comparison, across all products in all Nielsen modules, there is an average of \$136bn in revenue per year. The average price is calculated by summing both the number of units sold and revenue for a product in a store. The ratio is taken and then averaged to obtain a store-level average price. Weekly availability is calculated first at the yearly level by counting the number of weeks a product is sold divided by the number of weeks a store is open. Then, we average across years for store-level availability. Since the assortment analysis does not rely on weekly prices or availability, we omit these figures for the assortment analysis sample.

TABLE II
SIMILARITY IN PRICING ACROSS STORES, WITHIN-CHAIN VS. BETWEEN-CHAIN

Measure of Similarity:	Absolute Difference in Quarterly Log Prices		Correlation in (Demeaned) Weekly Log Prices		Share of Weekly <i>Log</i> Prices within One Log Point	
	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	0.031	0.137	0.806	0.091	0.622	0.104
Standard Deviation	(0.031)	(0.062)	(0.190)	(0.101)	(0.222)	(0.094)
Number of Chain-UPCs	86,426	86,197	84,875	83,930	85,140	86,001
Panel B: Benchmark UPCs, Store Pairs Within a DMA						
Mean	0.025	0.126	0.846	0.106	0.679	0.131
Standard Deviation	(0.025)	(0.067)	(0.161)	(0.173)	(0.201)	(0.141)
Number of Chain-UPCs	86,344	83,483	84,804	73,498	85,057	83,198
Panel C: Benchmark UPCs, Store Pairs Across DMAs, Top 33% income vs Bottom 33% Income Only						
Mean	0.041	0.141	0.759	0.087	0.544	0.102
Standard Deviation	(0.039)	(0.065)	(0.225)	(0.098)	(0.245)	(0.095)
Number of Chain-UPCs	61,167	80,344	59,542	79,052	59,755	85,715
Panel D: Food Stores, All Store Pairs						
Mean	0.031	0.137	0.809	0.089	0.624	0.104
Standard Deviation	(0.031)	(0.062)	(0.189)	(0.096)	(0.222)	(0.094)
Number of Chain-UPCs	83,793	83,558	82,282	81,312	82,546	83,382
Panel E: Drugstores, All Store Pairs						
Mean	0.052	0.136	0.651	0.041	0.498	0.095
Standard Deviation	(0.031)	(0.062)	(0.254)	(0.100)	(0.286)	(0.079)
Number of Chain-UPCs	551	552	528	546	529	547
Panel F: Mass Merchandise Stores, All Store Pairs						
Mean	0.032	0.130	0.761	0.195	0.569	0.081
Standard Deviation	(0.020)	(0.082)	(0.178)	(0.191)	(0.186)	(0.083)
Number of Chain-UPCs	2,082	2,087	2,065	2,072	2,065	2,072

Notes. This table presents measures of similarity of pricing for pairs of stores both within a chain, and across chains. To form the pairs we select a maximum of 200 pairs per chain (within the appropriate channel only) that correspond to the comparison criteria we impose (see text for additional details). In Panel A, we sample from the universe of store pairs within each channel. In Panel B, we sample from store pairs where both stores are located in the same DMA only, while in Panel C we compare only pairs of stores in different DMAs and such that one store in the pair is in the bottom third of the income measure, while the other store is in the top third. Panels D, E, and F are the means per channel.

TABLE III
DETERMINANTS OF PRICING

	Log Price				
	(1)	(2)	(3)	(4)	(5)
Panel A: Food Stores					
Own Store Income (in \$10,000s)	0.0168*** (0.0041)	0.0047*** (0.0009)	0.0042*** (0.0009)	0.0038*** (0.0003)	0.0038*** (0.0003)
Chain Average Income (in \$10,000s)		0.0372*** (0.0096)	0.0320*** (0.0108)	0.0296** (0.0114)	
Chain-State Average Income (in \$10,000s)				0.0086** (0.0042)	0.0086** (0.0042)
Fixed Effect for County	No	No	Yes	No	No
Fixed Effect for Chain	No	No	No	No	Yes
Observation Level	Store	Store	Store	Store	Store
Observations	9,415	9,415	9,415	9,415	9,415
R-squared	0.128	0.266	0.702	0.268	0.933
Panel B: Drug Stores					
Own Store Income (in \$10,000s)	0.0077*** (0.0011)			0.0054*** (0.0006)	0.0054*** (0.0006)
Chain-State Average Income (in \$10,000s)				0.0247*** (0.0068)	0.0242*** (0.0049)
Fixed Effect for Chain	No			No	Yes
Observation Level	Store			Store	Store
Observations	9,977			9,977	9,977
R-squared	0.100			0.185	0.394
Panel C: Mass Merchandise Stores					
Own Store Income (in \$10,000s)	-0.0117*** (0.0029)			0.0033*** (0.0010)	0.0033*** (0.0010)
Chain-State Average Income (in \$10,000s)				-0.0679*** (0.0092)	0.0071*** (0.0016)
Fixed Effect for Chain	No			No	Yes
Observation Level	Store			Store	Store
Observations	3,288			3,288	3,288
R-squared	0.041			0.276	0.932

Notes . In Panel A, standard errors are clustered by *parent_code* . In Panels B and C, standard errors are clustered by *parent_code* -state. In Panels B and C we do not report the specifications with chain-average income given that there are only 4 drug chains and only 5 mass merchandise chains.

*** p<0.01, ** p<0.05, * p<0.1

TABLE IV
PRODUCT ASSORTMENT AND OTHER FIRM DECISIONS IN FOOD STORES

Measures:	Product Assortment Price Index	Share Products in Top Decile by Unit Price	Fraction of Organic Products	Fraction of Generic Products	Log Number of UPCs Carried
	(1)	(2)	(3)	(4)	(5)
Panel A: Within/Between Chain Relationships: Absolute Log Difference Store Pairs					
Within-Chain Pairs Mean	0.012	0.007	0.004	0.011	0.126
Between-Chain Pairs Mean	0.052	0.019	0.007	0.055	0.273
Panel B: Relationship with Income					
Store Income, \$10,000s	0.0058*** (0.0013)	0.0119*** (0.0017)	0.0125*** (0.0011)	-0.0062*** (0.0014)	0.0281*** (0.0077)
Chain Income, \$10,000s	0.0425*** (0.0107)	0.0331*** (0.0084)	0.0200** (0.0076)	-0.0080 (0.0110)	0.1178** (0.0494)
County Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Modules	40	40	8	37	40
Products Considered	Top 20%	All	All	All	All
Sample Mean	0.000	0.219	0.085	0.231	5.495
Sample Standard Deviation	0.042	0.042	0.042	0.048	0.222
Observation Level	Store	Store	Store	Store	Store
Observations	9,415	9,415	9,415	9,415	9,415
R-squared	0.527	0.539	0.315	0.150	0.211

Notes. In Panel A, we use the same sample of store-pairs as in our price analysis in Table II. Each observation is the average measure in a store-sub-module-year, and we report the average over the pair-sub-module-years for within-chain pairs and between-chain pairs. The Product Assortment Price Index is a measure of the average log unit price for nationwide top-selling products that are carried by each store. Specifically, the products must be in the top 20% of units sold among all products belonging to the module. See notes to Figure VI for additional details of how the Product Assortment Price Index is calculated. Fraction of Products that are in the Top 10% Unit Price is the share of products carried at the store-level that have national average unit prices in the top 10% for each sub-module. Fraction of Organic Products and Fraction of Generic Products require that at least 1 percent of products belong to that category in order to ensure that variation across stores is possible (for example, organic toilet paper does not exist). In Panel B, standard errors are clustered by *parent_code*.

*** p<0.01, ** p<0.05, * p<0.1

TABLE V
DETERMINANTS OF STORE-LEVEL ELASTICITY IN FOOD STORES

Dependent Variable:	Elasticity			Log $\frac{\text{Elasticity}}{1 + \text{Elasticity}}$		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Demographic Controls</i>						
Income Per Capita (in \$10,000s)	0.1395*** (0.0226)	0.1329*** (0.0146)	0.1170*** (0.0187)	0.0527*** (0.0051)	0.0503*** (0.0051)	0.0501*** (0.0084)
Fraction with College Degree (or higher)			0.1833* (0.1081)			0.0364 (0.0345)
Median Home Price (in \$100,000s)			0.0064 (0.0113)			-0.0014 (0.0052)
Controls for Urban Share			X			X
<i>Controls for Number of Competitors Within 10km</i>						
1 Other Store			-0.0186 (0.0149)			-0.0110** (0.0053)
2 Other Stores			-0.0319*** (0.0106)			-0.0144*** (0.0047)
3+ Other Stores			-0.0624*** (0.0151)			-0.0236*** (0.0053)
Fixed Effect for Chain	No	Yes	Yes	Yes	Yes	Yes
Fixed Effect for UPC	No	No	No	Yes	No	No
Observation Level	Store	Store	Store	Store-UPC	Store	Store
Observations	9,415	9,415	9,415	6,593,513	9,415	9,415
R-squared	0.138	0.666	0.671	0.398	0.688	0.692

Notes. Standard errors are clustered by *parent_code*. All independent variables are our estimates of store-level demographics at the zip-code level based on Nielsen Homescan (HMS) panelists' residences. Demographics are from 2012 ACS 5-year estimates. 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]. Columns 4 - 6 represent the first stage we use in our IV specification (Table VI and Table VII). Columns 4 and 5 are the first stages of the IV regressions in Table VI, corresponding to the disaggregated and aggregated specifications respectively. Column 6 is the first stage of Table VII row 1. The first stage allowing for the log elasticity-income relationship to vary by UPC is not shown here.

*** p<0.01, ** p<0.05, * p<0.1

TABLE VI
DETERMINANTS OF STORE-LEVEL ELASTICITY IN FOOD STORES

Dependent Variable:	Within-Chain, IV			Between-Chain, IV	
	Log Price			Average Log Price	
	(1)	(2)	(3)	(4)	(5)
Log $\frac{\text{Elasticity}}{1 + \text{Elasticity}}$	0.0482*** (0.0137)	0.0852*** (0.0210)	0.0934*** (0.0220)		
Average Log $\frac{\text{Elasticity}}{1 + \text{Elasticity}}$				0.7609*** (0.2158)	0.8339*** (0.2315)
First Stage Varies by UPC?	Yes	No	N/A	No	N/A
Fixed Effects	Chain x UPC	Chain, UPC	Chain	UPC	
Observation Level	Store-UPC	Store-UPC	Store	Chain-UPC	Chain
Observations	6,593,513	6,593,513	9,415	54,364	64

Notes. This table reports the results of instrumental variable regressions, in which the log elasticity term is instrumented with store-level income as in Table V columns 4 and 5 for disaggregated and aggregated specifications, respectively. The standard errors are block bootstrapped by *parent_code*. Columns 1 and 2 are done at the disaggregated store-UPC level. In column 1, we allow the log elasticity-income relationship to vary by UPC. In columns 2 - 5 we pool this relationship. Column 4 is done at the chain-UPC level. Elasticities are winsorized at -1.2 to -7. Column 5 reports the mean average log elasticity term (not log of average elasticity).

*** p<0.01, ** p<0.05, * p<0.1

TABLE VII
LOG PRICES AND STORE-LEVEL LOG ELASTICITY, ROBUSTNESS

Dependent Variable:	Within-Chain, IV	Between-Chain, IV	First Stage
	Log Price	Average Log Price	Log $\frac{\text{Elasticity}}{1 + \text{Elasticity}}$
	(1)	(2)	(3)
Benchmark Coefficient	0.0934*** (0.0220)	0.8339*** (0.2315)	0.0503*** (0.0051)
1. Benchmark Products, IV with all variables	0.0986*** (0.0186)	0.8274*** (0.2121)	0.0501*** (0.0084)
2. Homescan Income as Instrument	0.1112*** (0.0313)	1.8686** (0.7408)	0.0232*** (0.0040)
3. Quarterly Elasticity	0.0626*** (0.0148)	0.5579*** (0.1405)	0.0749*** (0.0059)
4. Module Index Price and Elasticity	0.0716*** (0.0200)		0.0678*** (0.0061)
5. Generics	0.0686*** (0.0230)		0.0407*** (0.0039)
6. Top-Decile Products by Revenue	0.0947*** (0.0180)	0.6195*** (0.2318)	0.0466*** (0.0036)
7. Bottom-Decile Products by Revenue	0.0868*** (0.0235)	0.9799*** (0.2469)	0.0558*** (0.0070)
8. All Products, Weighted by Revenue	0.0829*** (0.0157)	0.6526*** (0.2105)	0.0511*** (0.0042)
9. Elasticities Winsorized at -1.5	0.1338*** (0.0320)	1.1942*** (0.3160)	0.0351*** (0.0033)
10. Drug Stores	0.2286*** (0.0394)		0.0330*** (0.0050)
11. Mass Merchandise Stores	0.2013*** (0.0504)		0.0220*** (0.0038)

Notes. Each row and column represents a different specification. This table reports the results for various instrumental variable regressions, in which the log elasticity term is instrumented with store-level income. The standard errors are block bootstrapped by *parent_code* in rows 1 through 9. Standard errors are block bootstrapped by *parent_code-state* in rows 10 and 11. The benchmark coefficients are the corresponding coefficients from Table VI columns 3 and 5. Row 1 uses a richer set of regressors for the first stage and the benchmark weekly price; see Table V column 6. In row 2, we instrument using Homescan per capita income. Row 3 replaces our benchmark elasticities with quarterly elasticities. In row 4, we replace both price and elasticity with module index prices and elasticities. In row 5, we reestimate elasticities for our within-chain generic products and replace both price and elasticity with equivalent generic prices and elasticities using a total of 12,423 generic product. Row 6 keeps the top decile of products by average yearly revenue, taking into account the number of years a product is in the sample. Row 7 keeps the bottom decile of products by average yearly revenue, also taking into account the number of years a product is in the sample. Row 8 keeps all benchmark products but weights by store-UPC revenue. Row 9 winsorizes elasticities above at -1.5 rather than -1.2, still winsorizing below at -7. Rows 10 and 11 show the within results for Drug and Mass Merchandise stores. Between chain comparisons are not possible. In column 3, we show the corresponding first stage of the IV regression.

*** p<0.01, ** p<0.05, * p<0.1

TABLE VIII
ESTIMATED LOSS OF PROFITS IN FOOD STORES

	10th	25th	Median	75th	90th
Panel A: Store-Level (N=9,415)					
Comparing Flexible Pricing to Uniform Pricing					
Yearly Dollars	\$112.0K	\$158.0K	\$239.0K	\$374.1K	\$602.3K
Percent of Revenue	1.11%	1.39%	1.79%	2.45%	3.71%
Comparing Flexible Pricing to IV Price-Elasticity Slope					
Yearly Dollars	\$92.2K	\$129.4K	\$196.1K	\$307.5K	\$494.1K
Percent of Revenue	0.90%	1.14%	1.47%	2.01%	3.04%
Comparing Flexible Pricing to State-Zone Optimal Pricing					
Yearly Dollars	\$97.9K	\$138.6K	\$212.0K	\$329.5K	\$530.2K
Percent of Revenue	0.98%	1.22%	1.56%	2.14%	3.27%
Panel B. Chain-Level (N=64)					
Comparing Flexible Pricing to Uniform Pricing					
Yearly Dollars	\$2.12M	\$7.07M	\$19.57M	\$50.81M	\$111.45M
Percent of Revenue	1.18%	1.43%	1.93%	2.26%	2.82%
Comparing Flexible Pricing to IV Price-Elasticity Slope					
Yearly Dollars	\$1.73M	\$5.80M	\$16.09M	\$42.19M	\$91.57M
Percent of Revenue	0.97%	1.17%	1.59%	1.84%	2.33%
Comparing Flexible Pricing to State-Zone Optimal Pricing					
Yearly Dollars	\$1.78M	\$6.41M	\$18.53M	\$48.88M	\$97.02M
Percent of Revenue	0.89%	1.35%	1.61%	1.95%	2.66%

Notes. This table reports the difference between the profits computed under optimal pricing and the profits under alternative scenarios, divided by the store-level revenue. Optimal pricing is assuming the monopolistic competition model and thus deriving optimal prices using $\log(P) = \lambda + \log(c)$, where λ is log elasticity. Elasticities are winsorized at -1.2 and -7. Uniform pricing assumes that each chain sets the optimal uniform price across its stores. Pricing according to the IV price-elasticity slope assumes that chains set prices according to β the IV estimate in Table VI column 3. State-Zone Optimal Pricing assumes that the chain charges a uniform price within each state, with the price set optimally in the chain-state. In Panel A each observation is a store. In Panel B we aggregate to the chain level. The percentiles do not indicate the same store. For example, the median store in terms of dollars lost is not the median store in terms of losses as a percent of revenue.

TABLE IX
ESTIMATED LOSS OF PROFITS IN FOOD STORES

	Price-Elasticity Relationship (IV)			
	(1)	(2)	(3)	(4)
Log (No. of Stores)	0.0178*** (0.0053)	0.0083 (0.0089)		
Log (No. of States)	-0.0047 (0.0124)	0.0054 (0.0127)		
Log (Average Yearly Store Sales)	-0.0151 (0.0169)	-0.0189 (0.0130)		
Standard Deviation of Store-level Per-capita Income, \$10,000s		0.0685** (0.0278)		
Log Dollar Profit Loss from Uniform Pricing			0.0153** (0.0059)	
Percent Profit Loss from Uniform Pricing			-0.0130 (0.0145)	
Share of Stores with Competitor Stores within 10 km				0.0071 (0.0325)
Share of Store with Same-Chain Stores within 10 km				-0.0057 (0.0635)
Analytic Weights	Yes	Yes	Yes	Yes
Observation Level	Chain	Chain	Chain	Chain
Observations	64	64	64	64
R-squared	0.186	0.314	0.156	0.001

Notes: The dependent variable is the chain-by-chain estimate of the IV specification, as in Table VI column 3, computing the first stage using all chains. Standard errors are clustered by *parent_code*. Analytic weights equal to the inverse standard error squared of the reduced form chain-level regression of price on income are used. The chain-level percent profit loss from uniform pricing is as in Table VIII, Panel B, row 1. The log dollar profit loss from uniform pricing is computed taking the store-level loss from uniform pricing in dollar terms, and scaling it up by the share of revenue in that store due to the selected UPCs; we then sum the dollar losses across stores in a chain, and take the log.

*** p<0.01, ** p<0.05, * p<0.1

TABLE X
RESPONSE TO LOCAL SHOCKS USING A REPRESENTATIVE PRODUCT IN FOOD STORES

Outcome:	Estimated Log Point Change in Prices: \$2,000 Decrease in Income		
Assumed Price Setting:	Flexible Pricing	Uniform Pricing	Yearly Pricing
	(1)	(2)	(3)
Panel A: \$2000 Negative Shock in All Stores			
National Shock, Impact on All Stores	-1.01	-1.00	-1.00
State-Level Shock, Impact on Same-State Stores	-1.01	-0.32	-0.45
County-Level Shock, Impact on Same-County Stores	-1.01	-0.04	-0.21
Panel B: \$2000 Negative Shock in California			
Impact on California Stores	-1.01	-0.72	-0.77
Impact on Nevada Stores	0.00	-0.38	-0.31
Panel C: \$2000 Negative Shock in Nevada			
Impact on Nevada Stores	-1.01	-0.16	-0.31
Impact on California Stores	0.00	-0.04	-0.03

Notes. Displayed are the estimated log point response to a permanent \$2,000 decrease in income using a representative product with predicted elasticities using the coefficients in Table V column 5. Additionally, we assume that the income shock translates into a change of the log elasticity in terms of that same column. In Panel A, the averages are the mean response for stores in each locality, weighting each locality equally. Uniform Pricing assumes that chains set one uniform price across all stores with a constant marginal cost for all chains, and that all stores are the same size. Yearly pricing takes into account consumer substitution by adjusting using the coefficient in Online Appendix Table 11 column 1. For more detail, see section 7.1 Inequality.