Uniform Pricing in US Retail Chains

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

November 16, 2017

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 the level of competition. Estimating a model of consumer demand reveals substantial within-chain variation in price elasticities and suggests that the average chain sacrifices seven percent of profits relative to a benchmark of flexible prices. In contrast, differences in average prices between chains broadly conform to the predictions of the model. As possible explanations for nearly-uniform pricing, we discuss advertising, tacit collusion, fairness concerns, and managerial fixed costs, and find the most support for the last explanation. We show that the uniform pricing we document significantly increases the prices paid by poorer households relative to the rich, likely dampens the overall response of prices to local economic shocks, and may also 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 and seminar participants at New York University, Stanford GSB, the University of Chicago (Department and Booth), UC Berkeley, UCLA, the University of Bonn, at the 2017 SITE Conference in Psychology and Economics and at the 2017 Berkeley-Paris conference in Organizational Economics for helpful comments. We thank Angie Acquatella, Sahil Chinoy, Bryan Chu, Johannes Hermle, Ammar Mahran, Akshay Rao, Sebastian Schaube, Avner Shlain, Patricia Sun, and Brian Wheaton for outstanding research assistance. Gentzkow acknowledges funding from the Stanford Institute for Economic Policy Research (SIEPR).

1 Introduction

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

In this paper, we show that most large US food, drugstore, and mass merchandise chains in fact set uniform or nearly-uniform prices across their stores. 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 held fixed across separate markets, rather than across multiple goods sold in a single market. We show that limiting price discrimination in this way costs firms significant short-term profits. We then show that the result of nearly-uniform pricing is a significant dampening of price adjustment, and that this has important implications for the pass-through of local shocks, the incidence of trade costs, and the extent of inequality.

Our analysis is based on store-level scanner data for 9,415 food stores, 9,977 drugstores, and 3,288 mass merchandise stores from the Nielsen-Kilts retail panel. In our baseline results, we focus on prices of ten widely available items. We consider larger sets of products in extensions and robustness analysis. We use the standard price measure in these data, which is defined to be the ratio of weekly revenue to weekly units sold.

Our first set of results documents the extent of uniform pricing. While we observe no cases in which the measured prices are the same for all products across stores, the variation in prices within chains is small in absolute terms and far smaller than the variation between stores in different chains. This is true despite the fact that consumer demographics and levels of competition vary significantly within chains: consumer income per capita ranges from $22,700 at the average 10th-percentile store to $40,900 at the average 90th-percentile store, and the number of competing stores within 10 kilometers varies from 0.6 at the 10th-percentile store to 8.3 at the 90th-percentile store.[1] Prices are highly similar within chains even if we focus on store pairs that face very different income levels, or that are in geographically separated markets. We can also look directly at the relationship between price and consumer income. Within chains, prices increase by 0.72 percent (s.e. 0.12)

[1]The within-chain heterogeneity is only slightly smaller than the cross-sectional heterogeneity across stores, with income at the 10th-percentile store of $19,300 and income at the 90th-percentile store of $41,700.

for each $10,000 increase in the income of local consumers. Between chains—that is, comparing chain-average prices to chain-average income—prices increase by 4.48 percent (s.e. 1.01).[2] Another way of looking at the same fact is to regress a store's log price on (i) the income of consumers in its own market and (ii) the average income of consumers in its chain; the coefficient on the former is an order of magnitude smaller than the coefficient on the latter (0.004 versus 0.040). All of these results remain similar for various alternative products, including store brands, lower-selling items, and high-priced items.

Next, we 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. This ratio can vary across stores not only because of variation in the prices but also because of variation in the quantity weights. In particular, we expect stores facing more elastic demand (e.g., lower income) to sell a larger share of units at relatively low prices, leading the weekly price measure to be lower in such stores, even if posted prices do not vary at all. Thus, the aggregation to weekly average prices can lead not only to excess variance in measured prices, but also to apparent correlation between measured prices and income.[3]

To assess the importance of this compositional bias, we turn to more detailed data from a major grocer studied in Gopinath et al. (2011) that allows us to see posted non-sale prices directly. These data suggest two reasons why consumers within a given week pay different prices. First, Nielsen's weeks run from Sunday to Saturday while this retailer typically changes prices mid-week. Second, consumers with loyalty cards pay lower prices than those without loyalty cards.[4] When we use the standard Nielsen price measure, this chain looks similar to other food chains in having a small but clearly non-zero price-income gradient.[5] When we adjust for the compositional bias and look directly at the posted prices, however, this relationship completely disappears.

For the large majority of the 73 chains in our data, measured prices vary very little across stores, and we suspect, based on our analysis of the major grocer, that their posted prices are in fact essentially uniform. For 11 food chains as well as the 2 major drugstore chains, prices vary at the level of large geographic zones, but vary much less within them.

Our second set of results uses a simple constant-elasticity model of demand to assess the extent

[2]The between-chain comparison includes just food stores, given that there are too few drugstore or mass merchandise chains for a meaningful between-chain comparison.

[3]The importance of the distinction between posted prices and average prices paid has been previously emphasized by Chevalier and Kashyap (2015) and Coibion, Gorodnichenko, and Hong (2015). They point out that average prices paid at the annual or market level may be responsive 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 (2015) and Coibion, Gorodnichenko, and Hong (2015) treat the weekly average price as equivalent to the posted price. Campbell and Eden (2014) and Cavallo (forthcoming) make a related point, noting that aggregation to the weekly level can lead observed data to overstate the frequency of price changes.

[4]Einav, Leibtag, and Nevo (2010) discuss further measurement error due to loyalty cards in Nielsen data.

[5]Give that this chain does zone pricing at the state level, we estimate the pricing gradient with state fixed effects.

to which uniform pricing represents a deviation from (short-run) optimal prices. 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 estimate of elasticity is both statistically precise and closely predicted by store-level measures of demographics and competition. Estimated elasticities vary widely within chains. Food stores whose elasticities fall at the 10th percentile within their chains have an average elasticity of -2.28. For those at the 90th percentile within their chains, the average elasticity is -2.98. This range is -1.94 to -2.65 for drugstores, and -2.92 to -3.67 for mass merchandise stores. Our model implies that the ratio of the optimal price to marginal cost for a store with elasticity η_s is $\eta_s/(1+\eta_s)$. Assuming no variation in marginal costs across stores, prices at stores with elasticities in the 90th percentile should be 18 percent higher than stores with elasticities in the 10th percentile in food stores, 29 percent higher in drugstores, and 11 percent higher in mass merchandise stores. However, observed prices are on average only 0.4 percent higher in food stores, 0.8 percent higher in drugstores, and 0.4 percent higher in mass merchandise stores.[6] To formally test the model's predictions, we regress log prices on the term $log\left[\eta_s/(1+\eta_s)\right]$, instrumenting this term with store income. This yields a between-chain coefficient for food chains of 0.94 (s.e. 0.22), very close to the value of 1 that the model would predict. The within-chain coefficient is an order of magnitude smaller, at 0.09 (s.e. 0.03), and the compositional issues discussed above suggest this is likely an over-estimate.

The model allows us to quantify the loss of profits from uniform pricing. The loss is highest for stores in high-income areas, where prices would be substantially higher under optimal pricing. We estimate that the average chain could increase profits by 6.9 percent under flexible pricing.

We consider a number of potential threats to the validity of our model. First, our model abstracts from variation in marginal costs across stores. Stroebel and Vavra (2015) 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. To the extent that marginal costs do vary, we would expect them to be positively correlated with income, meaning that our model if anything understates the gap between observed and optimal prices. Second, our baseline estimates assume that short-run week-to-week elasticities are equal to long-run elasticities. The long-run elasticities relevant to the store's problem could in fact be smaller (due to consumer stockpiling as in Hendel and Nevo, 2006) or larger (due to search). We repeat our analysis using prices and quantities aggregated to the quarterly level and find that the broad patterns are unchanged. We also show that the results are similar for storable and non-storable products. Third, our main analysis treats demand as separable across products. Cross-product substitution could lead us to overstate the relevant elasticities as consumers substitute among products, or understate them as consumers substitute on the store-choice margin

[6]For 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 first (i.e., the within-chain difference in average prices for stores near the 90th percentile and stores near the 10th percentile) and then take the average difference, weighting each retailer equally.

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. Finally, 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 suggests that constraints imposed by manufacturers are unlikely to be a key driver of our results.

The third section of our analysis considers potential explanations for uniform pricing. We argue that neither menu costs (Mankiw, 1985) nor price advertising provides plausible explanations. We see softening price competition and fairness concerns as potentially more plausible, but to the extent we are able to test them with our data we find limited support. The explanation we find the most support for is managerial decision-making costs (Bloom and Van Reenen, 2007).[7] Implementing more flexible pricing policies may impose chain-level fixed costs such as up-front managerial effort in pricing design, or a cost for inertial managers to deviate from the traditional pricing approach in the industry. Two pieces of evidence are consistent with such fixed costs. First, chains with more stores or facing more variable consumer income levels are more likely to deviate from uniform pricing. Second, among the drugstore and mass merchandise chains (though not for the food chains), the extent of uniform pricing has decreased over time, consistent with improvements in technology reducing managerial fixed costs.

In the final section of the paper, we turn to the implications of uniform pricing for the broader economy. We first show that uniform pricing exacerbates inequality, increasing prices posted to consumers in the poorest decile of zip codes by eight percent relative to the prices posted to consumers in the richest decile. We then show that 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. Finally, we show that uniform pricing may change the incidence of intra-national trade costs, benefiting more remote areas that otherwise would pay significantly higher prices, and that it can also bias estimates of these costs that use spatial price gaps as a key input. In several of these cases, we also note that the standard practice of treating average prices paid at weekly or greater time horizons as equivalent to posted prices can also be a source of bias in the results.

We are not the first to document uniform pricing policies in retailing.[8] Prior work has noted uniform pricing by European supermarkets and other major European retailers.[9] Early studies of

[7] A different version of this explanation is that managers are simply unaware of the income differences across their stores, or that they lack the information to recognize their implications for optimal prices. This seems unlikely to us.

[8] Our discussion of the literature focuses on private retail firms. Miravete, Seim, and Thurk (2014) offer a related analysis of the implications of a uniform markup regulation applied to government-run liquor stores in Pennsylvania. Their work parallels ours in estimating variation of demand elasticities and considering the distributional implications of uniform pricing.

[9] Reports from UK regulators show that roughly half of UK supermarket chains charge uniform prices across stores (Competition Commission, 2003, 2005; Dobson and Waterson, 2008) as do the main UK electronics retailers (MMC, 1997a,b). Cavallo, Nieman, and Rigobon (2014) show that Apple, IKEA, H&M, and Zara charge nearly uniform prices across the Euro zone in their online stores, though they charge different (real) prices across countries with different currencies. Eizenberg, Lach, and Yiftach (2016) study price variation across supermarkets in Jerusalem.

the Dominicks chain in the Chicago market showed that Dominicks varied prices between pricing zones but kept prices constant within zone, and the same is true for the widely-studied major US grocer from Gopinath et al. (2011).[10] The more comprehensive Nielsen data set shows that, while a minority of US food chains engage in zone pricing, nearly-uniform pricing is the industry norm.

Two recent papers are particularly related. Adams and Williams (2017) show that the Home Depot and Lowe's US hardware chains use a zone pricing strategy, with different degrees of price flexibility for different products. They estimate a structural model of demand and oligopoly pricing for a single product, drywall, and use it to evaluate how profits would change under more flexible pricing for this product. Contemporaneous work by Hitsch, Hortaçsu, and Lin (2017) uses the same Nielsen data we do to decompose price variation for close to 50,000 products. Though their main focus is separating the roles of regular price variation and promotions, they also note that prices vary more between chains than within chains, and they also estimate elasticities for a large set of stores. Our paper differs in more sharply characterizing the extent of uniform pricing, in comparing observed pricing to an optimal benchmark, and in addressing broader economic implications.

More broadly, our paper relates to a large body of work on the extent and implications of local retail price responses to economic shocks or incentives. Contemporaneous work by Leung (2017) and Kryvtsov and Vincent (2017) provide direct evidence supporting our claim that uniform pricing dampens responses to local shocks. Other relevant work includes Broda and Weinstein (2008), Gopinath et al. (2011), Fitzgerald and Nicolini (2014), Coibion, Gorodnichenko, and Hong (2015), Handbury and Weinstein (2015), Kaplan and Menzio (2015), and Dubé, Hitsch, and Rossi (forthcoming). Our work also speaks to the literature tracing out the implications of retail firms' price setting for macroeconomic outcomes, including influential early work using scanner data by Bils and Klenow (2004) and Nakamura and Steinsson (2008), and recent contributions such as Anderson et al. (2017).

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 (DellaVigna and Malmendier, 2004; Gabaix and Laibson, 2006). Our paper is part of a smaller literature which considers behavioral *firms* instead, that is, cases in which firms deviate from simple benchmarks of profit maximization (Romer, 2006; Bloom and Van Reenen, 2007; Hortaçsu and Puller, 2008; Goldfarb and Xiao, 2011; Massey and Thaler, 2013; Hanna, Mullainathan, and Schwartzstein, 2014; Hortaçsu et al., 2017; Ellison, Snyder, and Zhang, 2016).

[10] See Hoch et al. (1995), Montgomery (1997), and Chintagunta, Dubé, and Singh (2003). It may not be a coincidence that the chains which have been most likely to partner with researchers are also the ones that implement the most sophisticated pricing policies. The focus of research on these chains may also explain why the full extent of uniform pricing has been under-appreciated.

2 Data

Our primary data sources are the Nielsen Retail Scanner (RMS) and Consumer Panel (HMS) data provided by the Kilts Center at the University of Chicago.[11] The retailer scanner panel records the average weekly revenue and quantity sold for over 35,000 stores in the US over the 2006-2014 period, covering about a million different unique products (UPCs). We use this data set to extract the information on weekly price and quantity. We also use some information from the consumer panel which is based on following the purchase of more than 60,000 consumers across different stores. We present the main information in this section, with additional detail in the Appendix.

Stores. We focus the analysis on three store types, or channels: food (i.e., grocery), drug, and mass merchandise. Table 1, Panel A shows that the initial Nielsen sample includes 38,539 stores for a total average yearly revenue (as recorded in the RMS data) of $224 billion.[12]

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 company that owns a store and the latter indicates the chain itself.[13] 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. We introduce additional restrictions detailed below to exclude cases like these where chain identity is unclear.

We first introduce restrictions at the store level. We exclude stores that switch chains over time,[14] stores that are in the sample for fewer than two years,[15] and stores without any consumer purchases in the HMS data. This reduces the sample to 22,985 stores in 113 chains.

We next introduce restrictions at the chain level. We require that the chains are present in the sample for at least 8 of the 9 years. This eliminates a few chains with typically only a small number of stores each with inconsistent presence in the data. Next, we resolve cases where the mapping of stores to chains is not sufficiently clear. A first concern occurs when the same retailer_code identifier appears for stores with different parent_codes. It is unclear whether the use of the same retailer_code in this case indicates that these stores belong to one chain, or perhaps they belong to a subchain that changed owner, or something else. Thus, for each retailer_code, we only keep the parent_code associated with the majority of its stores, and then further exclude cases in which this retailer_code-

[11] The data are collected by Nielsen and made available through the Marketing Data Center at the University of Chicago Booth School of Business. Information on availability and access to the data can be found at https://research.chicagobooth.edu/kilts/marketing-databases/nielsen.

[12] This figure omits revenue from prescription drugs, most kinds of produce, and a number of other products that are not in our data.

[13] However, some parent companies report their data to Nielsen in a decentralized manner representing each of their retail chains, so we do not observe all ownership relationships.

[14] In some cases, we can validate the ownership change with significant observed changes in prices in the switching stores. However, some of the pricing changes occur up to two years in advance, or two years after, the change, suggesting a possible inaccurate record of the timing of ownership changes. We therefore exclude such switches.

[15] This sample restriction is especially important for our estimation of price elasticities, since our elasticity estimates include controls for 52 week-of-year indicators, requiring multiple observations per week-of-year.

parent_code combination accounts for less than 80% of the stores with a given retailer_code. A second concern is for chains in which a number of stores switch chain, given that this may indicate a change in ownership of the entire chain. We thus exclude chains in which 60% or more of stores belonging to the retailer_code-parent_code change either parent_code or retailer_code in our sample.

To define the demographics of the stores, we use the HMS data, which includes all shopping trips for the consumers in the panel. The median store has 21 Nielsen consumers ever purchasing at the store, for a total of 502 trips (Panel B of Table 1). We use variables like income and education from the 2008-2012 5-year ACS for the 5-digit zip code of residence of the consumers shopping in each store, and then compute the weighted average across the consumers, weighting by the number of trips that they take to the store.[16] We let Y_s denote this measure of income for store s.

Table 1 provides summary statistics for our main sample. Panel A summarizes the sample restrictions, which 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). Panel B summarizes these stores' demographics. The median store has an average per-capita income of $26,900, with sizable variation across stores; for example, the 75th percentile is at $33,450. Panel C-E provide chain-level summary statistics. The median food chain (Panel C) has 66 stores, and has locations in 4 DMAs (Designated Market Area) and 2.5 states. Drugstore and mass merchandise chains (Panel D and E) are significantly larger and span more states, with the vast majority of stores in both cases belonging to 2 chains. Given this high concentration, our between-chain analysis below is limited to food chains.

Our store sample covers the entire continental US (see map in Appendix Figure 1), with a number of stores and chains that is fairly constant between 2007 and 2013 (Online Appendix Table 1).

Products. We focus most of our analysis on a set of products that are both frequently sold and widely available. This guarantees clean comparisons both within and between chains, and avoids the problem that the price measure is missing in weeks with zero sales. This is an issue especially because the price is not missing at random—a week with no purchases is less likely to occur when a product is on a price reduction—thus introducing a potential bias in the price measure.

For food stores, we focus on one UPC from each of ten product categories ("modules"): canned soup, cat food, chocolate, coffee, cookies, soda, bleach, toilet paper, yogurt and orange juice.[17] We define the first eight to be storable products and the last two to be non-storable products. These modules together account for an average yearly revenue of $13.6 billion across the 9,415 stores in our food store sample, or 10.0 percent of total revenue. For drugstores and mass merchandise stores,

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

[17] These modules have a large overlap with ones used in previous analysis, e.g., Hoch et al. (1995).

we focus on the subset of these modules in which some UPC is available in at least 90 percent of stores: soda and chocolate for drugstores, and soda, chocolate, cookies, bleach, and toilet paper for mass merchandise stores. Within each module, we choose a specific UPC to maximize sales and availability across our sample of stores. In most cases, these UPCs remain the same across years; in all cases, the UPC is the same across stores within a channel and year. Examples of products in our sample (Table 2) include a 12-can package of Coca-Cola, a single 10.75 oz can of Campbell's Cream of Mushroom Soup, and a 59 oz. bottle of pulp-free Simply Orange juice.

In our robustness analysis, we consider larger sets of products. To avoid availability issues, we construct these larger product samples only for food stores. They include less commonly sold items (the 20th highest-availability product across chains for each module), high-quality items (chosen to have a high-unit-price), the top-selling generic product within each chain, a subset of generic products comparable across chains, and a large basket of products we use to construct module-level price and quantity indices.[18]

To define the module-level baskets, we include all UPCs in a module such that the average share of weeks with non-zero sales for that UPC is at least 95 percent, where the average is taken across stores.[19] For some modules such as soda and orange juice, products meeting this criterion cover 50-60 percent of the total module revenue, while for other modules like chocolate or coffee, they cover just 15-20 percent (see Online Appendix Table 1, Panel B). Summing over the 10 modules, these products cover an average annual revenue of $6.3 billion.

Prices. As is standard in the literature, 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 unit is sold in a UPC-store-week. We let p_{sjt} denote the standardized log price, defined as $\log(P_{sjt})$ minus the average of $\log(P_{sjt})$ across stores and weeks within each year and store format. To define the average log price of store s, $\overline{p}_s$, we first average p_{sjt} within years to produce a mean $\overline{p}_{sjy}$; we then define $\overline{p}_s$ to be the simple average of $\overline{p}_{sjy}$ across products j and years y.

Table 2 summarizes prices and availability for the products in our main sample. The average price varies from $0.49 for cat food in food stores to $8.60 for toilet paper in mass merchandise stores (column 3). The products have at least one recorded sale in the large majority of store-week, for example in 99.7% of store-week-UPC observations for chocolate in food stores (column 4). Cat food, coffee, and toilet paper have somewhat lower availability in food stores, as do most of the products sold in drugstores and mass merchandise stores, but are still in the range around or above 95%. We also compute the average yearly revenue per store that these products generate, with the highest number associated with the soda product in food stores, $34,100.

To compute the module-level price and quantity index for store s, we start from the weekly log price p_{sjt} and weekly log units sold q_{sjt}, then average across all products j included in the basket for

[18] See Appendix Section A.1.4 for more details

[19] We omit weeks from this calculation in which the store has zero recorded sales in all ten modules.

that module-chain-year. As weights, we use the total quantity sold for product j in a chain-year. If a product j has no sales in a particular store s and week t, product j is omitted for that store-week cell, and the other weights are scaled up accordingly.[20]

Major Grocer's Data. We use supplemental scanner data from a single major grocer (parent_code) studied in Gopinath et al. (2011).[21] 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 also differs from Nielsen, and is aligned with the timing of the retailer's weekly price changes. The data cover 250 stores belonging to 12 chains (retailer_codes) beginning in 2004 and ending in mid-2007. We focus on the largest retailer, which has 134 stores. We match 132 of these 134 stores to stores in our main sample (see Appendix Section A.1.7).

3 Descriptive Evidence

3.1 Example

We begin with a visualization of pricing by a single chain (chain 128), which we choose to be representative of the typical patterns observed in our data. Figure 1a shows the prices of the orange juice product. The 108 rows in the figure correspond to the 108 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 standardized 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 of up to 30 log points, but virtually no variation across stores within a week. To the extent that prices vary across stores, this variation is uncorrelated with store per capita income, i.e., the vertical position of stores in the chart. This is despite the fact that store income ranges from about \$13,000 at the bottom of the chart to about \$50,000 at the top. Figure 1b shows a similar pattern for five other products: yogurt, chocolate, soda, cookies, and cat food. Here we display just 50 of the 108 stores shown in Figure 1a, with the same 50 stores shown for the 5 products, and still ordered by income. We see variation across products in the depth and frequency of sales, but again no systematic variation of prices across stores. The pricing patterns of this chain are representative of the large majority of chains in our sample. Two additional examples are in Online Appendix Figure 1a and 1b.

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

[21] The data-sharing agreement between this retailer and the research community is managed through the SIEPR-Giannini data center (http://are.berkeley.edu/SGDC).

While patterns like these are typical of the majority of chains, a few other chains follow a different pattern, which we will call zone pricing. Figure 2 displays an example for chain 130 (showing a random sample of 250 stores), returning to the orange juice product. Figure 2 follows the same design as Figure 1a, except that we group stores geographically by sorting them by 3-digit zip codes within states. This chain operates in 12 different states. Prices are essentially uniform within horizontal bands, but then differ for different bands. For example, stores in Georgia and Kentucky share the same pricing patterns, which are different in Illinois and most of Indiana. Note that the pricing zones are strongly correlated with state borders but do not follow them perfectly. Online Appendix Figure 2a and 2b show other examples of zone pricing.

3.2 Measures of Pricing Similarity

To describe chain pricing patterns more systematically, we introduce three measures of the extent of uniform pricing. Each measure defines the similarity of prices of a pair of stores s and s'. To compute a chain-level measure of similarity, we average the raw measure across pairs within a chain.

The first measure is the *quarterly absolute log price difference*. We denote the average of p_{sjt} across weeks in quarter v by $\overline{p}_{sjv}$. We compute for each pair of stores s and s' the absolute difference in the average quarterly log price, and average this difference across quarters and products: $a_{s,s'} = \frac{1}{N_{vj}} \sum_{v,j} |\overline{p}_{sjv} - \overline{p}_{s'jv}|$, where N_{vj} denotes the number of valid product-quarter observations.

The second measure is the *weekly correlation in prices*. We first demean the log price p_{sjt} at the store-year-product level to obtain $\widetilde{p}_{sjt}$. Then we compute the correlation of $\tilde{p}_{sjt}$ and $\tilde{p}_{s'jt}$, including all weeks t and all products j which are non-missing in both store s and store s'.

These two measures capture, by design, two orthogonal aspects of similarity: differences in average prices, and the 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*. This is defined as price differences smaller than one percent, i.e., the share of observations across products j and weeks t for which $|P_{sjt} - P_{s'jt}|/((P_{sjt} + P_{s'jt})/2) < 0.01$.

Figure 3 displays the distribution of these measures for store pairs in the same chain (solid blue bars) and pairs that belong to different chains (hollow red bars). To form the within-chain pairs, we keep all stores for chains with fewer than 200 stores, and a random sample of 200 stores from larger chains, and then compute similarity for all pairs within the resulting set of stores. Prices for same-chain pairs are far more similar on all three measures than for different-chain pairs. The absolute log price difference (Figure 3a) is typically below 5 log points for the former, and typically above 10 log points for the latter. The weekly correlation (Figure 3b) is typically above 0.8 for the

former and below 0.2 for the latter. The share of identical prices (Figure 3c) is often as high as 0.5 or 0.6 for the former, but is rarely above 0.2 for the latter.

Table 3 summarizes a number of variants of these measures. The first row summarizes the same information shown in Figure 3, reporting the mean and standard deviation of the three similarity measures for our full set of same-chain and different-chain store pairs. The second row 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). The third row 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).[22] These results provide initial evidence against the possibility that the observed similarity just reflects same-chain store pairs serving more homogeneous consumers in terms of either geography or demographics. They also suggest that it does not result from constraints specific to pairs of stores operating in the same geographic market, for example, because price advertising is determined at the newspaper or television market level. The remaining rows of Table 3 show that these pricing patterns are not an artifact of focusing on our set of widely-available products. Focusing on food stores, the table shows similar patterns for: (i) the 20th-highest-selling product within a category; (ii) the top-selling store-brand product, and (iii) products with high unit-prices.[23]

Figure 4 summarizes pricing similarity at the chain level. In Figure 4a, weekly correlation is on the vertical axis, absolute log price difference is on the horizontal axis, and each point indicates the average value for a single chain. The vast majority of chains cluster in the upper-left of the figure, with low price differences and high correlation. Out of 73 chains, 61 have both an average correlation of weekly prices above 70 percent and an absolute quarterly distance in prices below 5 percent. These two measures of pricing similarity are also highly correlated: chains that are similar in one dimension are also similar in the other dimension. One might have *ex ante* expected to see deviations from this - for example, highly correlated sales but varying levels of regular prices across stores - but these deviations do not appear to a substantial degree in the data.

Figure 4b returns to the phenomenon of zone pricing. We decompose the measures of 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 each of two or more states, we plot the within-state log price difference on the horizontal axis, and the between-state log price difference on the vertical axis. To the extent that zone pricing follows state boundaries, it should show up in this figure as smaller differences within and larger differences between, i.e., points above the 45-degree line in the figure. Between differences are indeed larger in almost all cases, but for the majority of chains only slightly so; these chains do not appear to determine prices by state to a significant extent. A minority of

[22] Online Appendix Figure 3 displays the distribution of distance between pairs for these samples.

[23] In Online Appendix Figure 5, we show that chains with uniform prices for our benchmark products also tend to have uniform prices for these alternative products.

chains do have clearer zone pricing patterns, however. Most notable is chain 9, which has an average within-state difference of roughly 2 log points but an average between-state difference of more than 9 log points.[24]

3.3 Price Response to Consumer Income

We now turn to the relationship between prices and income. Stores in high-income areas have less elastic consumers (as we confirm below), so all else equal, we expect these stores to charge higher prices. Though we argue that variation in marginal costs across stores is likely to be small, any such variation would likely be positively correlated with income, and so tend to strengthen this relationship.

Figure 5 shows the relationship between log income and price within and between chains.[25] We first regress both store average log price $\overline{p}_s$ and store income Y_s on chain fixed effects. Figure 5a shows a binned scatterplot of the residuals, illustrating the within-chain relationship. The relationship is positive and highly significant, but the magnitude is very small economically: an increase in per-capita income of $10,000, equivalent to a move from the 30th to the 75th percentile, is associated with an increase in prices of only 0.72 percent. Figure 5b shows a scatterplot of the chain averages, illustrating the between-chain relationship.[26] This is also highly significant, and its magnitude is more than five times larger: a $10,000 increase is associated with a price increase of 4.5 percent.

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 these two hypotheses, providing strong support for the former.

If we estimate the within-chain coefficients separately by chain (Online Appendix Figure 6a-c), the majority of chains have small, positive coefficients in the range between 0 and 0.01, with 27 coefficients positive and significantly different from zero. Only five chains have coefficients above 0.01. The pattern of tiny within-chain response and large between-chain response is robust to dropping the two outlier chains 98 and 124 (Online Appendix Table 2). It also holds for lower-selling and high price products, as well as for all but one module separately (Online Appendix Figure 7).

Figure 5c-d examine the role of zone pricing. As we documented in Figure 2 and 4b, some chains vary prices more between states than within states. In Figure 5c, we re-estimate the within-chain relationship, but now plot residuals after taking out chain*state fixed effects. This further reduces the price-income slope to 0.56 percent per $10,000 of income, but it remains statistically significant.

[24] Online Appendix Figure 4 shows a parallel figure using correlation in weekly prices, instead of the log price difference. The classification of chains as zone pricers based on this alternative figure is overall similar.

[25] The results are similar replacing income with the fraction of college graduates (Online Appendix Figure 9).

[26] Here we omit drugstore and mass merchandise chains, since comparing across formats may be less informative and the number of chains for these channels is small. Online Appendix Figure 8 shows the plot including these chains. The overall pattern remains unchanged.

In Figure 5d, 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 2.16 percent, a slope about half the size of the between-chain relationship.

Figure 6a-d zoom into this zone pricing result for different groups of chains. For the 54 food chains that are not identified as doing zone pricing based on Figure 4b, there is no relationship between chain-state income and chain-state prices (Figure 6a). There is instead a clear relationship for the 11 food chains identified as zone pricing (Figure 6b) and for the 2 major drugstore chains (Figure 6c). The mass merchandise chains (Figure 6d) exhibit modest evidence of zone pricing.

Finally, Table 4 presents an alternative view of the price-income relationship. We run a store-level regression of average log price $\overline{p}_s$ on both store income Y_s and the average income of stores in the chain to which s belongs. In some specifications, we include separately the average income in s's chain-state. We separate food stores (Panel A) from the drugstores and mass merchandise stores (Panel B and C), since it is only for the food stores that we can do a meaningful between-chain comparison. The first column presents the regression including only own-store income as a benchmark.[27] The second column adds chain average income for food stores. Consistent with the evidence in Figure 5, 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 of which is from a mainly high-income chain and one of which is from a mainly low-income chain, the former will tend to charge much higher prices than the latter.[28]

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 (for food stores) than the response to overall chain average income, consistent with our other zone pricing results. Online Appendix Table 4 shows that the results are similar using generic products, lower-selling items, and higher-quality items.

3.4 Composition Bias

The pattern of within-chain pricing in Figure 5a-c poses a puzzle. Why would chains exert the effort to vary their prices in a highly systematic way with consumer income, but then do so with an economically 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 here that this small price-income relationship is likely to be mainly an artifact of composition bias,

[27] For mass merchandise stores, there is a negative relationship between prices and income when not including chain fixed effects because among the largest two mass merchandise chains, the one operating in, on average, higher income areas has lower prices (see Online Appendix Figure 8b).

[28] One could worry that this result stems from the chain-level income measure capturing the elasticity of consumers in a store better than the own-store income measure due to measurement error in the latter. In Online Appendix Table 3, we show that the elasticity in a store is predicted by own-store income but not predicted by chain-level income, counter to this explanation.

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 do not change prices at more than weekly frequency, they may institute their price changes on a different day of the week. If, for example, they change 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 but not all consumers may use store cards or obtain other discounts that lead them to pay lower prices.

For an example of the bias that can arise, consider a retailer that charges identical prices in all stores and that changes prices on Wednesdays. Suppose in a particular week, they cut the price from P^{high} to P^{low}. The measured 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 made in the first half of the week in store s. If the share θ_s varies across stores for any reason, this will obscure the 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 if posted prices do not vary at all. A similar bias arises if the share of consumers who use store cards or other discounts is greater among consumers who are most price elastic.

This bias can explain a price-income gradient broadly consistent with what we observe. We illustrate this with calibrated magnitudes for the case of a major grocer (see also Section 2), which indeed changes prices on Wednesday. Suppose that the income of store s is \$10,000 greater than the income of store s', and that, consistent with our estimates of the income-elasticity relationship below, this translates into price elasticities among their respective consumers of $\eta_s = -2.5$ and $\eta_{s'} = -2.65$. Suppose that P^{low} is 40 percent lower than P^{high}, consistent with the average depth of a sale for this grocer. Consistent with our model below, assume store s faces a constant-elasticity demand function $Q_s = kP_s^{\eta_s}$. Then it is straightforward to show that $(\theta_s, \theta_{s'}) = (0.218, 0.205)$, and the difference in log prices is .0075.[29] For this grocer, products go on sale and off of sale about twice every five weeks, which implies a bias in 40 percent of the weeks and thus a slope of 0.0030 in the analogue of Figure 5c. Thus, under reasonable assumptions based on data from a grocer, this bias could explain half of the observed slope across all stores (0.0056 in Figure 5c) and all of the slope for food stores (0.0029, Online Appendix Table 2). This calibrated bias would be lower in stores with less offset (say, stores changing prices on Monday), and would be likely higher in presence of

[29] To see this, assume prices change in the middle of the week and note that $\theta_s = \frac{q_s^{high}}{q_s^{high}+q_s^{low}} = \frac{\left(P^{high}\right)^{\eta_s}}{\left(P^{high}\right)^{\eta_s}+\left(P^{low}\right)^{\eta_s}}$. 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 $P_{s'}^{RMS}$.

heterogeneity in the share of consumers using store cards.

To provide direct evidence on the magnitude of this bias, we use the major grocer's data described in Section 2. We expect the bias arising from mid-week price changes to be present 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 7a 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 as in Figure 5c. The slope of 0.0027 is similar to the one in Figure 5c and marginally significant. Figure 7b reproduces the same estimate, but using weekly prices from the grocer's data. The estimated slope falls to 0.0008 and is no longer significantly different from zero. We thus cannot reject the view that posted prices for this retailer do not vary at all with income, and that all of the within-chain-state slope for this retailer is an artifact of the weekly offset. Figure 7c 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 to measure—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 all remaining slope disappears: this chain sets uniform non-sale prices with respect to income. We observe further that within a state, 74 percent of store pairs charge identical non-sale prices (i.e., within one percent as defined above).

As an additional check about the importance of this bias, we use an algorithm described in Appendix Section A.1.8 to estimate non-sale prices in the Nielsen data and repeat the analysis of Figure 5. Online Appendix Figure 10a-d show that this flattens the within-chain price-income relationship but not the between-chain relationship.

Our conclusion is that a large part of the within-chain and within-chain-state slopes shown in Figure 5 is likely an artifact of composition bias. We suspect that a large majority of chains are in fact charging exactly the same prices in all of their stores, or in all stores within geographic zones. Moreover, we note that this bias will not only affect the cross-sectional relationship between prices and income, but also the apparent response of prices to income shocks observed in panel data. We discuss the implications of this bias for the literature on local price responses in Section 6 below.

3.5 Variation in Marginal Cost

Stroebel and Vavra (2015) provide a detailed analysis suggesting that variation in marginal costs among retailers such as those in our data is likely to be negligible. They first use novel data on wholesale costs to show that geographic variation in these costs is minimal. Since wholesale costs account for three-quarters of total costs and a presumably much larger share of marginal costs, this significantly 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.

We confirm the findings for wholesale costs in our large grocer's data. A binned scatterplot of

the wholesale cost variable for this grocer against store income (Figure 7d) displays no evidence of a positive relationship between the two variables. Given this evidence, we will assume for the remainder of our analysis that marginal costs are constant across stores within a chain.

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. It provides a valuable benchmark, however, on the extent to which short-run pricing incentives vary across stores.

A monopolistically competitive chain chooses a price P_{sj} for each product in each of its stores to maximize total profits. Each store's residual demand for product j takes the constant elasticity form $Q_{sj} = k_{sjw}P_{sj}^{\eta_s}$, where Q_{sj} is the number of units sold, k_{sjw} is a scale term that may vary seasonally by week of year w, and η_s is the store's price elasticity. Total cost $c_j Q_{sj} + C_s$ for a store consists of a constant marginal cost c_j and a store-level fixed cost C_s. The chain solves

$$max_{\{P_{sj}\}} \sum_{s(c),j} \left(P_{sj} - c_j\right) Q_{sj}\left(P_{sj}\right) - \sum_{s(c)} C_s. \tag{1}$$

The first order conditions yields, for all j,

$$P_{sj}^* = \frac{\eta_s}{1+\eta_s} c_j \tag{2}$$

or in log terms

$$p_{sj}^* = log\left(\frac{\eta_s}{1+\eta_s}\right) + log\left(c_j\right). \tag{3}$$

There is thus a simple relationship between elasticities and optimal prices, and under optimal pricing a regression of log prices on $log\left(\frac{\eta_s}{1+\eta_s}\right)$ within chains should yield a coefficient of one.

4.2 Elasticity Estimates

The model above requires estimates of the store-level elasticity of demand η_s. As our benchmark measure, we estimate the response of weekly log quantity to the weekly log price product-by-product, for each store s. More precisely, letting $q_{sjt} = log\left(Q_{sjt}\right)$, we estimate separately for each store s,

$$q_{sjt} = \eta_s p_{sjt} + \alpha_{sjy} + \gamma_{sjw} + \epsilon_{sjt}, \tag{4}$$

where α_{sjy} is product*year fixed effects, γ_{sjw} is product*week-of-year fixed effects, and ϵ_{sjt} is an error term. The former controls for the fact that the exact UPC associated with a product varies in some cases across years, and the latter captures seasonal variation in k_{sjw}. The coefficient on the log price is the estimated price elasticity, $\widehat{\eta}_s$. We use price variation for all 9 years and all 10 products in order to maximize precision. We cluster the standard errors by bi-monthly period, thus allowing for correlation across products, as well as over time within a 2-month period.

This stylized demand structure abstracts away from two important margins: intertemporal substitution and cross-product substitution. The former could lead quantities in week t to vary with prices in prior weeks. The latter could lead to demand for one product to depend on the price of other products. We revisit these assumptions below.

To adjust for sampling error in the elasticity estimates, we use a simple empirical shrinkage procedure. We re-estimate the elasticity using just the first 26 weeks of each year and again using the next 26 weeks of each year; label these elasticity estimates $\widehat{\eta}_{1,s}$ and $\widehat{\eta}_{2,s}$. We choose a shrinkage parameter ρ to minimize the mean squared difference between $(1-\rho)\,\widehat{\eta}_{1,s}+\rho\overline{\eta}_1$ and $\widehat{\eta}_{2,s}$, where $\overline{\eta}_1$ is the overall mean of $\widehat{\eta}_{1,s}$ across stores.[30] The estimated optimal shrinkage is just $\hat{\rho}=0.104$ for food stores, though it is slightly larger at $\hat{\rho}=0.305$ for drugstores and $\hat{\rho}=0.408$ for mass merchandise stores. We then adjust our overall estimates $\hat{\eta}_s$ as $\tilde{\eta}_s=(1-\hat{\rho})\,\hat{\eta}_s+\hat{\rho}\overline{\eta}$, where $\overline{\eta}$ is the mean of $\hat{\eta}_s$.

Figure 8a shows the distribution of the resulting elasticity estimates $\tilde{\eta}_s$ for food, drug, and mass merchandise stores. All are well-behaved, with all but a handful of values smaller than the theoretical maximum of -1, and most of the mass falling between -2 and -4. Figure 8b shows the distribution of associated standard errors, which are mostly between 0.05 and 0.2 for food stores and between 0.2 and 0.4 for drugstores and mass merchandise stores. The lower precision for the latter is expected given the smaller number of products in the drug and mass merchandise samples.

Figure 8c provides evidence on the fit of the constant elasticity demand model. The figure shows a binned scatterplot of residuals of q_{sjt} against residuals of p_{sjt} after partialing out the fixed effects α_{sjy} and γ_{sjw}. The model assumes a linear relationship, and this test proves that it is true to a remarkable degree. This plot aggregates across all products and tens of thousands of stores of all types. Visual inspection of this relationship by product and store-by-store generally yields similarly well-behaved linear relationships (with different slopes as expected, given the different mean elasticity estimates for each store type and module); some examples are in Online Appendix Figure 12a-b.

We provide two additional pieces of evidence validating the elasticity estimates for food stores in Online Appendix Figure 12c-d. First, the log price variable explains about half of the remaining variation (in terms of R^2) after controlling for the fixed effects. Second, we run a module-by-module regression that pools across stores and augments equation (4) by including also the prices charged in weeks $t-2$ and $t-4$, as well as in week $t+4$ (and using store*year fixed effects μ_{sy} in place of the

[30] Online Appendix Figure 11 shows the mean squared error as function of the shrinkage parameter.

product*year fixed effects α_{sjy}). The coefficients on these variables, while statistically significant and in line with the predictions of a model of stockpiling, are an order of magnitude or more smaller than the coefficients on price in week t. Furthermore, they are not systematically larger for storable products, like toilet paper and canned soup, than for non-storables, like orange juice and yogurt.

Finally, we examine the correlates of our estimated elasticities. Figure 8d shows that the estimated elasticities vary monotonically (and in fact linearly) with store income. Table 5 presents regressions of the estimated elasticities $\tilde{\eta}_s$ on a broader set of demographic and competition measures. The results confirm the robust relationship between elasticity and income: an increase of \$10,000 is associated with an increase of the elasticity of 0.140 (s.e. 0.014). This estimate remains similar with the addition of chain fixed effects. In columns 3 and 4, we add as determinants the share of college graduates, the median home price, controls for the fraction of urban area, as well as a simple measure of competition with other stores: indicators for the number of competitor stores within 10 kilometers of the store. The coefficients have the expected signs, with income as the strongest determinant, and the expected effect of the competition proxies.[31]

Next, we replace the estimated elasticity $\tilde{\eta}_s$ as the dependent variable with the log elasticity term $log\left(\frac{\tilde{\eta}_s}{1+\tilde{\eta}_s}\right)$ suggested by equation (3). The regressions in columns 6-8 are the first-stage of the instrumental variables regressions we estimate below. Income is a strong predictor of log elasticity, with a coefficient that is relatively similar across food, drug, and mass merchandise stores.

4.3 Comparing Observed and Optimal Prices

In this section, we test the model's prediction for optimal pricing (equation 3) directly. We estimate

$$\overline{p}_s = \alpha + \beta log\left(\frac{\tilde{\eta}_s}{1+\tilde{\eta}_s}\right) + \epsilon_s. \tag{5}$$

Equation (5) follows from averaging equation (3) across products. Under the assumptions of the model, the coefficient β on the log elasticity term is equal to 1. If the chains under-respond to the elasticity variation, instead, we will observe $\beta < 1$. For our benchmark specification, we instrument the log elasticity term with the store-level income to address remaining measurement error in these estimates.[32] The standard errors are block bootstrapped, clustering by parent_code in food stores to allow for any within-chain correlation in errors. For the drugstores and mass merchandise stores, given that there are only, respectively, 4 and 5 chains, we block bootstrap by parent_code*state.[33] Figure 9 displays the first stage relationship between the log elasticity term and income within

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

[32] For this specification we winsorize the store elasticity $\tilde{\eta}_s$ at -1.2. This happens very rarely in the case of our benchmark weekly elasticity estimates (nine food stores and one drugstore) but more frequently for weekly index (564 food stores) and quarterly top-product estimates (796 food stores).

[33] We cluster by parent_code instead of by chain (retailer_code) to allow for correlation between chains that fall under the same ultimate owner (parent_code).

chains, between chains, and within and between chain-states. Unlike the analogous plots of the price-income relationship in Figure 5, the within and between relationships are remarkably similar, with a first stage coefficient varying between 0.034 and 0.055.

We estimate variants of this IV regression within and between chains and chain states. To compute within-chain estimates, we replace α with chain fixed effects. To compute within-chain-state estimates we replace α with chain*state fixed effects. To compute between-chain-state estimates we average $\overline{p}_s$ and the log elasticity term within chain-states, and run the regression including chain fixed effects. To compute between-chain estimates we average both terms within chains and then run the regression with no fixed effects. In all cases, we fix the first-stage coefficients at the values from the within-chain first stage shown in the final three columns of Table 5. This is motivated by the similarity of the coefficients in Figure 9, and prevents us from losing first-stage power in the specifications where we aggregate the data.

Table 6 presents our main IV estimates, starting with the food stores. The first two columns show within-chain and within-chain-state results respectively. Both coefficients are statistically significant, but an order of magnitude smaller than the model prediction of $\hat{\beta} = 1$. Thus, the observed nearly-uniform pricing within chains occurs despite significant incentives to vary prices to achieve profit maximization. The third column shows between-chain-state estimates. The results imply a substantial response of price to the elasticity term, though smaller than predicted by the model, $\hat{\beta} = 0.351$ ($s.e.\,0.193$). Finally, the fourth column shows the between-chain estimates. The estimated coefficient on the log elasticity term in this regression, $\hat{\beta} = 0.944$ ($s.e.$ 0.220), indicates that average pricing at the chain level is consistent with the model: we cannot reject a slope $\beta = 1$.

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

We present OLS estimates in Online Appendix Table 6 and Online Appendix Figure 13a-d. The results are qualitatively similar, but the point estimates are 3 times or more smaller. This may reflect measurement error in our estimated elasticities or a larger local average treatment effect in the IV regressions due to variation in income being more salient to chains than variation in other determinants of elasticities. The OLS results reinforce the conclusion that the within-chain price-elasticity relationship is an order of magnitude too flat to be consistent with the model.

4.4 Robustness

Table 7 presents a series of robustness checks, focusing on the results for food stores.

Quarterly Elasticities. Our short-run elasticities may differ from the longer-run elasticities that are relevant to the chains' pricing problem. Longer-run elasticities could be smaller due to

stockpiling, or larger if it takes consumers time to adjust to price changes—for example, if price increases induce gradual substitution to other stores. As a step toward addressing these concerns, we estimate quarterly elasticities. That is, we average the weekly log price and log units sold across all weeks in a quarter, and then re-estimate our main equation (4). The controls include product*year fixed effects as well as product*quarter-of-year fixed effects. The estimated quarterly elasticities are smaller (in absolute value) than the benchmark ones (Online Appendix Figure 15a), but the two measures are highly correlated (Online Appendix Figure 15b). Importantly, the quarterly elasticity measure passes the same validation exercises as our benchmark measure, as Online Appendix Figure 15c-e document: (i) the log-log specification is approximately linear; (ii) the standard errors of the estimated elasticity are still relatively small (at 0.15 to 0.4), even if clearly larger than in the benchmark elasticities,[34] and (iii) the measure is still highly correlated with local income.

Panel A of Table 7 reports our main IV results using the quarterly elasticity measure, with the first stage reported in Online Appendix Table 5. The finding that within-chain responses are an order of magnitude too small is robust to using the quarterly measure. In fact, the within-chain $\hat{\beta}$ coefficients become substantially smaller. This reflects two offsetting forces: the quarterly elasticities vary less across stores, which would tend to reduce the optimal price variation, but the fact that they are lower in magnitude 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, and the consequently larger gap between actual and predicted price variation, means that the second effect dominates. The between-chain relationship is still an order of magnitude larger than the within-chain response, but is now significantly smaller than the model predicts ($\hat{\beta} = 0.409$).

Module-Level Indices. Another important limitation of our model is the fact that it ignores substitution between products. Firms care about the profits they earn from all products. If some of the response in our benchmark elasticities reflects within-store substitution, the optimal price response could be smaller than our model would predict. To partially address this issue, we re-estimate our elasticities using the module-level price and quantity indices described in Section 2. As in the case of the quarterly elasticities, the module-level price elasticities are smaller (in absolute value) than the benchmark ones (Online Appendix Figure 16a), as expected, but are highly correlated with the benchmark elasticity (Online Appendix Figure 16b). The index elasticity also passes the same validation tests (Online Appendix Figure 16c-e).

In Panel B of Table 7, we regress the price index on the index elasticities, instrumented with income.[35] As with the quarterly elasticities, this yields even smaller within-chain coefficients: our main finding that chains vary prices too little from the model's perspective is robust. Here we cannot re-estimate the between-chain specifications given that the price indices are not comparable across

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

[35] The first stage is in Online Appendix Table 5.

chains. In Online Appendix Table 7 Panel A, we show that the within-chain results are similar if we use the index elasticity while keeping the benchmark prices as the dependent variable.

Additional Robustness. So far we have used per-capita income as an instrument for the log elasticity term. The results are very similar (Panel C) if instead we use the full set of demographic and competition variables in columns (3) and (4) of Table 5. In Panel D, we consider the role of branding by considering the price of a top-selling generic that is common to many chains within a subset of the modules considered. The within-chain relationship between price and income or elasticity remains very similar to the one for the top-selling branded good. In Online Appendix Table 7, we show that these patterns persist when we use the 20th most-available good in each of the 10 product categories as dependent variable (Panel B), generic top sellers within chains (Panel C), and some high-quality (high-price) products (Panel D).[36]

Yearly Average Price Paid. Our main question of interest 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. If we are interested in the welfare effects on consumers, however, we may also want to consider the average price paid over a longer time horizon. When prices vary over time, more elastic consumers could end up paying substantially lower prices on average than less elastic consumers even if the posted prices they face are the same.

To assess how large this force is, in Table 8, we compare our benchmark results using the log of the weekly average prices (Panel A) with the results using the log of the yearly average prices (Panel B). The within-chain coefficient ($\hat{\beta} = 0.223$) for the yearly average prices, while larger than the benchmark estimate in Panel A as expected, is still 5 times smaller than the model prediction under optimal pricing ($\beta = 1$). Thus, even taking into account this margin of adjustment does not bring the level of prices up to what is expected in light of the model. Still, it is interesting that the presence of sales works as a kind of "automatic" price discrimination, guaranteeing that consumers in stores facing more elastic demand pay lower prices over the year, even in presence of uniform pricing.[37]

4.5 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 c sets a constant price $\bar{P}_{jc}$ across all stores s to maximize

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

[36] Online Appendix Figure 14 presents the within-chain and between-chain plots for these alternative products.

[37] Online Appendix Figure 17a-d reproduce the key findings in Figure 5a-d, comparing the yearly average price to the weekly average price. The yearly average price is more responsive to within-chain differences in income than the weekly average price (Online Appendix Figure 17a and 17c), with a slope that is about twice as steep. Similarly, the between-state zone-pricing relationship is also stronger using the yearly average price (Online Appendix Figure 17d). The between-chain relationship for food stores, instead, is not much affected (Online Appendix Figure 17b).

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

$$\sum_{s(c)} k_{sj} \left[(1+\eta_s) \bar{P}_{cj}^{\eta_s} - c_{cj}\eta_s \bar{P}_{cj}^{\eta_s - 1} \right] = 0. \tag{6}$$

We assume that each chain c sets their average price $\bar{P}_{cj}$ equal to the solution to (6). This is a weaker assumption than the standard assumption of profit maximization, and is consistent with the between-chain results observed above. To compute the implied profits, we need estimates for the scale factor k_{sj}, the marginal costs c_{cj}, and the fixed cost C_s, in addition to the estimates $\tilde{\eta}_s$ of the elasticities computed above.[38] To estimate the quantity term, we take $\hat{k}_{sj}$ to be the average over the weeks t of $Q_{sjt}/P_{sjt}^{\tilde{\eta}_s}$. To estimate the marginal costs, we define $\hat{c}_{cj} = \left[\sum_s \hat{k}_{sj} (1+\tilde{\eta}_s) \bar{P}_{cj}^{\tilde{\eta}_s}\right] / \left[\sum_s \hat{k}_{sj}\tilde{\eta}_s \bar{P}_{cj}^{\tilde{\eta}_s - 1}\right]$, which, by equation (6), is a consistent estimator. Finally, to estimate the fixed costs, we follow Montgomery (1997), who estimates gross profit margins for supermarkets of 25 percent and operating profit margins of 3 percent; thus we posit $\hat{C}_s = (1 - (3/25)) * \Pi_s^g$, where Π_s^g is the (optimal) gross profit.

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 $P_{sj}^* = \hat{c}_{cj} * \tilde{\eta}_s/(1+\tilde{\eta}_s)$.[39] We then aggregate across products j to obtain $\bar{\Pi}_s$ and Π_s^*, and compute the percent of profit loss due to uniform pricing as $\left(\Pi_s^* - \bar{\Pi}_s\right)/\Pi_s^*$. Panel A of Table 9 shows that, in the average store, the profit loss is 8.9 percent of variable profits, with the losses as high as 23 percent at the 90th-percentile store. The highest losses occur for stores in high-income areas (Online Appendix Figure 18a), where the lower elasticities (in absolute value) would call for significantly higher prices under flexible pricing.

While uniform pricing captures well the observed pricing in most food chains, we do estimate a small, but statistically significant, response of prices to elasticity within chains. In the second row of Table 9, we compute the percent of profit loss for the actual profits Π_s, relative to optimal pricing. To compute the actual profits, we use the prices implied by our benchmark IV specification: $P_{sj} = A_{cj} * [\tilde{\eta}_s/(1+\tilde{\eta}_s)]^{\beta_{IV}}$, where β_{IV} is the estimate in column 1 of Table 6 and A_{cj} is a constant that guarantees that the average price P_{sj} across all stores s in chain is equal to the uniform price $\bar{P}_{cj}$.[40] We then compute the profits Π_{js} and aggregate across products j. The second row in Panel A of Table 9 shows that this tempers the losses to an average profit loss of 6.5 percent. In the final row, we 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 similar losses as using the actual price-elasticity slope.

In Panel B of Table 9, we aggregate the profits for all stores s in a chain c, and compute the

[38] For the computation of the lost profits, we use the modules with constant products throughout the 9 years so as to avoid estimating year-specific marginal cost and quantity terms.

[39] For the stores with elasticity $\tilde{\eta} > -1.2$, the elasticity is winsorized at -1.2.

[40] If we instead use the observed prices, the estimated lost profits would be larger, as there is significant idiosyncratic variation in these prices, which is both unrelated to income and sub-optimal from the perspective of our model; see Online Appendix Table 8.

chain-level profit loss $\left(\Pi_c^* - \bar{\Pi}_c\right)/\Pi_c^*$. The mean loss from uniform pricing at the chain level is 8.7 percent, and the mean loss using the actual price-elasticity slope is 6.9 percent. For the average food chain (Online Appendix Table 8), the profit loss from uniform pricing is about 8.8 percent, and 7.2 percent under the actual pricing. For the drugstore chains, the loss from uniform pricing is larger, though this larger loss is mitigated when taking into account the more sizable price-elasticity slope. Finally, the losses are smaller for the mass merchandise chains.

The loss of profits computed so far uses the benchmark weekly elasticity. In Panels B and C of Online Appendix Table 8 we compute the loss of profits using the quarterly elasticity and the module-level index elasticity in food stores. We follow the same procedure outlined above to infer the marginal costs, but using the relevant elasticity in place of $\tilde{\eta}_s$. The implied loss of profits is larger under these alternative elasticities, with mean loss of 34.4 percent for the quarterly elasticity and 17.2 percent under the index elasticity.

5 Determinants of Pricing and Explanations

We now discuss chain-level factors that predict flexible pricing, then turn to possible explanations.

Determinants of Flexible Pricing. Table 10 shows the relationship between the degree of uniform pricing and chain characteristics. The dependent variable is a chain's estimated response to the log elasticity term in equation (5), based on chain-by-chain estimate of the IV specification in column 1 of Table 6, where we pool chains in the first stage.[41]

We first relate this measure of flexible pricing to measures of chain size: the log number of stores, the log number of states the chain operates in, and the log of average revenue per store. As column 1 of Table 10 shows, both the number of stores and the number of states positively predict flexible pricing. Just these three variables achieve an $R^2 = 0.548$, implying that chain size is an important determinant of the uniformity decision.

Next, we add a reduced-form measure of the returns from flexible pricing: the standard deviation of store-level income across the stores in the chain. As column 2 shows, even controlling for chain size, heterogeneity in income across stores is a significant predictor of flexible pricing.

We then consider two model-based variables of the return to tailored pricing: the percentage gains from optimal pricing, $\left(\Pi_c^* - \bar{\Pi}_c\right)/\Pi_c^*$, as in Table 9, and the log absolute gain from optimal pricing, $log(\Pi_c^* - \bar{\Pi}_c)$. To compute the second variable, we take the store-level loss from uniform pricing in dollar terms, $\Pi_s^* - \bar{\Pi}_s$, and scale it up by the share of revenue in that store due to the selected UPCs; we then sum the dollar losses across all stores in a chain, and take the log.

The percent profit loss is likely the most relevant variable if the chain is comparing the gain

[41] That is, for each chain, we regress the store-level log price on the store-level income, yielding the coefficients in Online Appendix Figure 6. We then divide these coefficients by the first-stage coefficient in Table 5, columns 6-8 (depending on chain type). For example, for food chain 32, the coefficient is $0.0098/0.0474 = 0.207$.

from optimal pricing to a store-by-store cost of optimal pricing, such as, for example, the possible backlash from consumers who may perceive optimal pricing to be unfair. The profit loss in dollars is more relevant if the chain is comparing the gain from optimal pricing to a chain-wide fixed cost, such as a chain-wide managerial cost. As column 3 shows, of the two variables, only the (log of) the dollar profit losses is significant, consistent with the results in columns 1 and 2, which suggest that the chain size matters. While the log dollar profit loss is clearly correlated with the extent of tailored pricing (see Online Appendix Figure 18b for a scatterplot at the chain level), the explanatory power of these two variables ($R^2 = 0.3$) is about half the one for column 2.

Next, in column 4, we consider two variables that one would expect to play a role according to some of the possible explanations for the findings. As a measure of competition, we take the share of stores in chain c that have at least a competitor store within ten kilometers. Stronger competition could lead one to expect a stronger incentive to charge optimal prices. Conversely, though, as discussed below, tacit collusion incentives could lead one to expect the opposite. We also consider a measure of density, the share of stores which have a second store of the same chain within ten kilometers. If consumers compare prices across stores in the same chain and dislike price variation due to fairness, varying prices will be less attractive for contiguous stores. We find that neither variable matters significantly.

In column 5 we consider all determinants jointly. A key size variable, the number of states, remains a strong determinant, as does a variable indicating the dispersion of consumer demand, the income heterogeneity. The model-based measures are not statistically significant, though they are highly correlated with the other determinants. In Online Appendix Table 9, we document similar results using the chain-level quarterly absolute log price difference, as in Figure 4a.

Explanations. In light of this evidence, we now consider explanations for the documented prevalence of uniform pricing. Some traditional explanations do not appear to apply to this setting. One is **menu costs** (Mankiw, 1985). Food stores change prices regularly to implement sales. Thus, it is implausible that a menu cost limits the ability to set different prices at the store level. A behavioral explanation that also appears unlikely in this setting is **limited attention** on the part of managers (e.g., Gabaix and Laibson 2006). It is hard to imagine that managers are not aware, or are optimally inattentive, with regards to the local income in their stores, especially given that we consider local income averaged over several years.

We also think it is unlikely that uniform pricing is due to constraints imposed by price **advertising** or price posting online. Advertising could certainly create an incentive to price uniformly within a geographic market: if a firm advertises a particular price in a newspaper, for example, it might want to honor that price in all stores in that newspaper's circulation area. Chains rarely advertise prices nationally, however, and this constraint would only apply at the level of a newspaper market, or at most a television market (DMA), which is typically larger. In Online Appendix Figure 19a-b,

we compare the zone pricing at the state level to the zone pricing at the DMA level (after taking residuals for chain*state fixed effects). We find *less* evidence of zone pricing at the DMA level than at the state level, and we see many chains that charge nearly uniform prices across DMAs. A related possibility is that chains post product prices online and that this constrains them to honor these prices nationally.[42] This may be relevant to some cases in our data, but it is unlikely to be a primary explanation as typical food chains do not post a single set of product prices online.

We see three other explanations as more plausible. First, committing to uniform or zone pricing may allow chains to **soften price competition**. Suppose chains start out in a high-price equilibrium, and that they would be tempted to cut prices in certain markets to capture market share from their competitors. Under flexible pricing, they could do so while leaving their other stores' prices unchanged. Under uniform pricing, such price cuts become more costly, as the chain would be forced to cut prices in all stores. Dobson and Waterson (2008) present a model along these lines, and Adams and Williams (2017) estimate a richer model in which this incentive exists and find mixed support for it as an explanation using data from the hardware industry.

Such models predict that the value of committing to uniform pricing would be greater for chains that face significant competition than for chains that do not. We do not observe any chains that are literally monopolists in all their markets, but there is substantial variation in the extent of competition chains face. Yet, the share of stores in a chain facing competitors nearby is not a predictor of the extent of uniform pricing (Table 10). Further, in Online Appendix Table 10, we ask whether the price-log elasticity relationship differs for stores with no competing stores nearby. There is no evidence for food and drug stores. These exercises are not a definitive test – the incentive to soften competition need not be related in a simple way to our competition proxies – but we see them as evidence suggesting that softening price competition may not be the driving force.

Second, managers may avoid varying prices because consumers could perceive charging different prices in different stores as unfair (**fairness**). If varying prices would damage the reputation of the chain, this might outweigh the gains from flexible pricing. Anecdotal evidence from store managers themselves provides some support for this explanation. In a report on the UK grocery pricing, the UK Competition Commission writes "*Asda said that it would be commercial suicide for it to move away from its highly publicized national EDLP pricing strategy and a breach of its relationship of trust with its customers, and it would cause damage to its brand image, which was closely associated with a pricing policy that assured the lowest prices always*" and "*Morrisons stated that adopting a policy of local prices would be contrary to its long-standing marketing and pricing policy, it would damage its brand and reputation built up over many years and would adversely affect customer goodwill, as well as being costly to implement and manage*" (Competition Commission, 2003).

These quotes notwithstanding, fairness may be a less compelling explanation for uniform pricing

[42] Consistent with this, Ater and Rigbi (2017) document a sharp decline in overall price dispersion after the introduction of a law requiring mandatory online disclosure of prices for food stores in Israel.

in our setting than in others. Few consumers visit multiple stores from a chain in geographically separated markets, so if chains did choose to price discriminate across these stores, few consumers would observe this directly. In addition, if our chains were to price discriminate they would be raising prices on the rich and cutting prices for the poor (as we discuss in Section 6.1 below), not an obviously objectionable practice from an ethical point of view.

To provide direct evidence, we examine if there is more price similarity when detecting price variation is easier for consumers. As we show in Table 10, there is no evidence that chains with a higher fraction of isolated stores are more likely to do targeted pricing. Further, in Online Appendix Table 10, we consider if stores with no same-chain stores nearby are more likely to price flexibly. We find no evidence for food store or mass merchandise stores, with some evidence for drugstores. Overall, we find only limited support for the fairness explanation.

A final explanation is **managerial decision-making costs**. Implementing more flexible pricing policies may impose costs such as up-front managerial effort in policy design or investments in more sophisticated information technology. More generally, the costs of implementing sophisticated pricing policies may exhibit increasing returns to scale. Bloom and Van Reenen (2007) offer possible support for this, showing that larger chains have better management practices generally. To the extent that uniform pricing has been the prevalent policy in the past, change may impose additional costs, as managers may not be well incentivized to make the change, while fearing the cost if a change backfires.

The key prediction of any such model is that chains should invest more in flexible pricing when the returns to doing so are high. In particular, we would expect to see more such investment for large chains in terms of number of stores and revenue, and in chains where the variation of optimal prices across stores is large. The findings in Table 10 support these predictions.

Another prediction is that pricing would become more flexible over time, as information technology costs fall and inertial managers turn over. In Online Appendix Table 10, we re-estimate our IV specification forming the log price either using the first years (2006-08) or the most recent years (2012-14). Both drugstore and mass merchandise chains appear to have doubled the flexibility of their pricing over these years, though we do not see a comparable trend for food chains.

Overall, though none of this evidence is definitive, we see the most support for managerial decision-making costs or inertia, consistent with prior work on management and behavioral firms (e.g., Bloom and Van Reenen, 2007; Hanna, Mullainathan, and Schwartzstein, 2014).

6 Implications

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

6.1 Inequality

Jaravel (2016) among others brings attention to the role of store pricing for the rise of inequality in the past decades, and shows that the introduction of novel products reduced the relative prices faced by high-income consumers. Uniform pricing by chains will have a similar effect, since optimal prices covary positively with income.

To quantify this effect, we compare the relationship between average local prices and income to the counterfactual one we would observe if firms priced flexibly as in our benchmark model. We compute the observed level of prices at a particular income level by simply taking the average price charged by stores in all chains with local income in that range. For the counterfactual, we compute the optimal price under flexible pricing as detailed in Section 4.5: $p^*_{sj} = log(\hat{c}_{cj}) + log(\tilde{\eta}_s/(1+\tilde{\eta}_s))$. Since our observed price measure is standardized, we standardize the optimal prices as well by subtracting the average observed log price in a year-product cell. We then average this standardized optimal price across products j and further across all stores s at a particular income level. Note that the extent to which the price-income gradient increases in the counterfactual depends both on how much steeper the within-chain pricing slope becomes and on the extent to which low-income and high-income areas are served by distinct chains.

Figure 10a shows the result for food stores. Areas with higher income have higher actual prices: an extra $10,000 of local income increases prices on average by about 2 percent. This relationship is consistent with our between-chain relationship (e.g., Table 6 column 4): chains operating in higher average income areas charge higher prices. Yet, this price-income slope is much flatter than it would be if firms were setting prices optimally. Under flexible pricing, the price increase associated with $10,000 higher local income would be about 5 percent. The pattern is similar for drugstores (Figure 10b). For mass merchandise stores, the observed price-income relationship is in fact negatively sloped, due to the fact that of the two major chains, the one operating in higher income areas has lower prices. Even so, the counterfactual price has a positive slope with respect to income.

These patterns have quantitatively important implications for inequality. For the food stores, consumers of stores in the lowest income decile pay about 0.7 percent higher prices than they would pay under flexible pricing, but consumers of stores in the top income decile pay about 9.0 percent lower prices than under flexible pricing. Consolidation between retailers could further strengthen this pattern. These quantitative implications are even larger if we use the quarterly or index elasticities instead of the benchmark weekly elasticities (Online Appendix Figure 20).

6.2 Response to Local Shocks

A second implication of our findings relates to the response of prices to local economic shocks (Beraja, Hurst, and Ospina, 2016; Stroebel and Vavra, 2015; Leung, 2017; Kryvtsov and Vincent, 2017). 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, and this will be more true the smaller the geographic area affected by the shock.

The magnitude of this effect depends on the geographic distribution of stores, as well as the degree of uniform pricing. We provide an illustrative calibration in Table 11. For simplicity, we focus on our benchmark orange juice product. We assume that some group of stores faces a negative \$2,000 income shock, translating into an elasticity shift as in Table 5, column 6. We consider a set of such shocks that differ in the set of stores affected. These are (i) a nationwide shock that affects all stores in the country; (ii) a set of 48 state-level shocks, each of which affects all stores in a given state; (iii) a set of analogous county-level shocks. For each shock, we compute the response of prices under optimal pricing (column 1), under uniform pricing (column 2), and under our approximation of actual pricing (column 3), similar to what we did in Section 4.5. For the local shocks in groups (ii) and (iii), we report the mean response across shocks in each group.

The dampening effect of uniform pricing on local price responses is dramatic. Under flexible pricing, our \$2,000 income shock leads to roughly 0.9 percent lower prices. As expected, the response under uniform and actual pricing is similar when the shock is nationwide.[43] For state and county-level shocks, however, the responses are far smaller. Under uniform pricing, the average state-level shock reduces prices by 0.29 percent, and the average county-level shock reduces prices by 0.02 percent. These responses are somewhat higher under actual pricing (which allows some flexibility across stores), but still far smaller than under flexible pricing: the average state-level shock reduces prices by 0.35 percent and the average county-level shock by 0.10 percent.[44]

Panel B and C explore these responses further by considering an example of a cost shock affecting California or Nevada. Given that California is a large state, a negative shock in California leads to a decrease in prices in the California stores of 0.71 percent, still dampened but much closer to the response under optimal pricing. In contrast, an equal-sized shock in Nevada lowers the prices in Nevada stores by only 0.23 percent. This example also illustrates potential price spillovers. Under flexible pricing, a shock in California leaves prices unaffected in Nevada. Under uniform pricing, instead, it causes prices in Nevada to decrease by 0.36 percent, as some stores in Nevada are part of the same chains as stores in California. Uniform pricing thus causes not just under-response to local shocks, but also spillover in pricing beyond the region affected by the shock.

[43] The response to a nationwide shock under uniform and actual pricing is not identical to the flexbile pricing response because elasticities enter equation (6) non-linearly.

[44] The 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 11).

6.3 Incidence of Trade Costs

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

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

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

Both of these points relate to the estimation of trade costs. Uniform pricing also affects the true incidence of these costs. Just as we noted above that uniform pricing tends to raise prices in high-income areas and lower them in low-income areas, so here it will tend to raise prices in locations close to where products are produced and lower them in remote locations. It thus shifts the incidence of trade costs away from those located far from the place of origin to those located closer to it.

7 Conclusion

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. We show that limiting price discrimination in this way costs firms significant short-term profits. We find managerial costs to be the most plausible explanation for this pattern. We show that the result of nearly-uniform pricing is a significant dampening of price adjustment, and that this has important implications for the extent of inequality, the pass-through of local shocks, and the incidence of trade costs.

References

Adams, Brian and Kevin R. Williams. 2017. "Zone Pricing in Regional Oligopoly." *Cowles Foundation Discussion Paper No. 2079*

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

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

Ater, Itai and Oren Rigbi. 2017. "The Effects of Mandatory Disclosure of Supermarket Prices." *CEPR Discussion Paper No. DP12381.*

Atkin, David and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." *NBER Working Paper No. 21439*

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

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

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

Campbell, Jeffrey R. and Benjamin Eden. 2014. "Rigid Prices: Evidence from U.S. Scanner Data." *International Economic Review*, 55(2): 423-442.

Cavallo, Alberto. Forthcoming. "Scraped Data and Sticky Prices." *Review of Economics and Statistics.*

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

Chevalier, Judith and Anil Kashyap. 2015. "Best Prices: Price Discrimination and Consumer Substitution." *NBER Working Paper No. 20768.*

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 Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications." *American Economic Review*, 105(3): 993-1029.

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

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

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

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

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. Forthcoming. "Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession." *Marketing Science.*

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2010. "Recording Discrepancies in Nielsen Homescan data: Are they present and do they matter?" *Quantitative Marketing and Economics*, 8(2): 207-239.

Eizenberg, Alon, Saul Lach, and Merav Yiftach. 2016. "Retail Prices in a City." *CEPR Discussion Paper No. DP11590*

Ellison, Sara, Christopher M. Snyder, and Hongkai Zhang. 2016. "Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market." *CESifo Working Paper Series No. 6285.*

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

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

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

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

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

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

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.

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

Hendel, Igal and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica,* 74(6): 1637-1673.

Hitsch, Guenter, Ali Hortacsu, 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, Fernando Luco, Steven L. Puller, and Dongni Zhu. 2017. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." *NBER Working Paper No. 23526*

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.

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

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

Kryvtsov, Oleksiy and Nicolas Vincent. 2017. "The Cyclicality of Sales and Aggregate Price Flexibility."

Leung, Justin. 2017. "Minimum Wage and Real Wage Inequality: Evidence from Pass-Through to Retail Prices."

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

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

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

Miravete, Eugenio J., Katja Seim, and Jeff Thurk. 2014. "Complexity, Efficiency, and Fairness of Multi-Product Monopoly Pricing." *CEPR Discussion Paper No. DP9641.*

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

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

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

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

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.

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

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

Stroebel, Johannes and Joseph Vavra. 2015. "House Prices, Local Demand, and Retail Prices." *CESifo Working Paper Series No. 5607.*

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.

A Appendix

A.1 Data

A.1.1 Store Selection

In the RMS data, Nielsen provides a basic categorization of stores into five "Channel Codes": Convenience, Food, Drug, Mass Merchandise, and Liquor. Of these, we select Food, Drug, and Mass Merchandise chains since the Convenience and Liquor stores are typically not covered in the HMS data and thus would not be included in our final sample. In the HMS data, there are more detailed "Retailer Channel Codes" and each store is assigned to one of 66 mutually exclusive categories such as Department Store, Grocery, Fruit Stand, Sporting Goods, and Warehouse Club. Our starting sample of food stores includes all stores that are categorized as "Food" stores in the RMS data. All the food stores selected in the final sample fall into the "Grocery" store category in the HMS channel code categorization[45], the drugstores are all "Drug Stores," and the Mass Merchandise stores are all "Discounters."

Some stores change DMA or FIPS code over the time that they are in the sample. Since Nielsen identifies store by the physical location of the store, this occurs because DMA regions or county lines are redefined over the nine years we observe. In other words, the stores themselves are not changing physical locations. For stores that switch, we use the modal DMA and FIPS code. This does not affect how we aggregate store-level demographics for our main analysis.

A.1.2 Demographics

All demographics are zip code level data from the 2008-2012 5-year ACS. We aggregate this zip code level demographics into store-level demographics as explained in Section 2. There are two special cases: (i) for one store with missing median home price data, we impute this value by regressing median home price on the other demographics (income, fraction with a bachelor's degree, race, and fraction of urban area); (ii) three drugstores are only visited by one household each, and these households provide a PO Box zip code as its zip code, making it impossible to use our usual procedure; we use county-level demographics for these three stores.

A.1.3 Competition Measures

We use the HMS panel data to help us construct a measure of competition based on geodesic distance. To compute the location of each store, we use the more detailed location information in the HMS data.[46] First, we assume that each HMS household lives at the center of its zip code. For each of the stores in the HMS dataset, we use a trip-weighted average of the coordinates of each household in order to arrive at an imputed location for the store. For our measure of competition for store s, we then count the number of stores of the same channel (e.g., food stores, counting only other food stores) within various distances (e.g., 5 or 10 km) of store s by geodesic distance.[47]

A.1.4 Product Selection

We select 10 modules (product categories) based on commonly available and highly-sold products in food stores. These products include five that belong to product groups used in Hoch et al. (1995) (soup, cookies, orange juice, soda, and toilet paper) and in Montgomery (1997) (orange juice).

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

[46] Recall that the location of the store in the Nielsen publicly available data is only recorded up to the 3-digit zip code or county.

[47] i.e., distance as the crow flies

Within a module (e.g., soda), we select a high-selling product (e.g., 12-pack cans of Coke). The product choice aims to ensure that (i) the product is available across as many chains and stores as possible (to ensure comparability across stores and across chains), and that (ii) within a store, it is sold in as many weeks of the year as possible (since otherwise the price is not recorded). We select the product within a module-year with the highest number of week-store observations with positive sales. This procedure typically selects, for a given module, different products in different years. We refine this initial selection with an eye to identifying a constant product across all 9 years, when possible. We search for a single product (UPC) that is present in all nine years and with availability in each module-year at most 10 percentage points below that of the top product. In food stores, this leads us to identify a consistent product for 7 out of 10 modules; for the remaining 3 modules, the selected product varies across the years. Because of low availability, we identify two constant products in drugstores and five products in mass merchandise stores, two of which are constant.[48]

For robustness analysis, just for the food stores, we also identify an additional set of four products. (i) (*Less Commonly Sold Items*) For each module, we also select the 20th-highest availability seller (for example, 6 oz Chobani Black Cherry); this product is chosen to be the same across all chains. (ii) (*High-quality products*) Since high-quality, high-price products are not well represented in our benchmark products, for three modules (chocolate, cookies, and coffee) we identify products that are in the top 2 quintiles of average price (conditional on package size) are high in availability, as defined above. We select three coffee products (for example, Starbucks 12 oz French Roast), 3 cookie products (Pepperage Farm Chessmen), and 2 chocolate product (Lindt 5.1 oz bar). (iii) (*Chain-specific generic product*) We identify within each module a generic (store-brand) product as the product with highest availability within each chain among the UPCs with the Nielsen identifier "CTL BR" (which identifies store-brand products). This generic product is not comparable across chains. (ii) (*Generic product across chains*) We choose a different set of generic products to enable between-chain comparisons. We take advantage of the fact that Nielsen assigns the same (masked) UPC to products it deems similar across chains, e.g. "Yogurt 32 oz." To minimize the chance that we identify non-comparable products across chains, we require that the average price for each store-product is within 20% of each other for stores in the same DMA. On this sample of products, we apply a similar product selection procedure as for our benchmark products, leading to generic products that are comparable across (most) chains for 4 modules (soup, cookies, soda, and yogurt).[49]

A.1.5 Prices

As described in the text, we compute the weekly price P_{sjt} as the ratio of weekly revenue and weekly units sold for that store-product. We apply the following filters: (i) Following the Nielsen manual, we divide the weekly units sold by the variable 'prmult' (price multiplier); (ii) We drop all prices $\leq$ \$0.10 since almost surely these represent cases of measurement error. This affects very few observations: 1,298 store-product-weeks (0.003% of observations) in food stores, 1,613 observations (0.017%) in drugstores, and 459 observations (0.006%) in mass merchandise stores.

A.1.6 Pairs Dataset for the Analysis of Store Pricing Similarity

For the measure of similarity in pricing across stores, we create a data set of pairs of stores. For the within-chain comparisons, we select up to 200 stores per chain, out of which we form all possible pairs. For the chains with more than 200 stores, we select 200 stores with consistent presence in the sample, since we could not compare the similarity in pricing between stores operating in different sample years. (Recall that the condition for inclusion of a store is data availability for at least 104

[48] In drugstores, we replace the top-availability soda UPC with the fifth-highest-availability soda UPC as the top four products go on temporary price reductions extremely rarely and thus make it difficult to identify the demand elasticity.

[49] For each store, we keep only the generic modules with at least 80% availability over the 9 years. Some stores have pricing information for only a subset of the 4 generic products.

weeks). Specifically, we select stores with non-missing price observations for at least 60% of the 468 sample weeks in at least 5 out of 10 modules in food stores, 1 out of 2 modules for drugstores, and 3 out of 5 modules for mass merchandise stores. For the between-chain comparisons, given that some of the comparisons condition on DMA, we sample one store per chain-DMA if there are multiple chains in a DMA and two stores per chain-DMA otherwise.

For the between-chain sample, we begin with the within-chain sample (limited to a maximum of 200 stores per chain) and select the maximum of 1 store or 7.5% of stores (rounded down) in that sample, giving preference to stores with consistent presence in the sample as described above.

A.1.7 Major Grocer's Data

As additional data, we use the scanner data for 250 stores from a major grocer as in Gopinath et al. (2011). Since we want to compare the results using the Nielsen price measure versus the price measures in this major grocer's data, it is important to identify the stores in the Nielsen data which correspond to stores in this additional data set. Since the dataset in Gopinath et al. (2011) covers 2004 to mid-2007, we focus on the 1.5 years of overlap covering all of 2006 and part of 2007. We match the two data sets using the 3-digit zip code and (with a fuzzy match) using the sum of units sold in 2006 for our benchmark products in the 10 modules. This results in 132 matches to stores in our main sample, all of which belong to a single Nielsen retailer_code. We validate the correctness of the matches using data on price and with an alternative matching algorithm.

A.1.8 Imputing Nonsale Prices from Nielsen RMS Data

The major grocer's data provides information on non-sale prices. While we do not have the same information in the Nielsen RMS dataset, we infer non-sale prices using the following procedure. We keep only prices that are above the 80th percentile of prices in a given store-year–module. Futher, we require that a given price is charged for three weeks in a row (two weeks in a row for the first and last weeks of each year). This yields $P_{sjt}^{nonsale}$ for the weeks t where this variable is defined. We then calculate the price level as detailed in Section 2. For the 132 stores that are both in the Nielsen and in the major grocer's data, we can provide a validation of this procedure comparing the non-sale price measure computed this way with the non-sale price variable in the major grocer's data. The two data series match perfectly for all but 13 store-product-weeks.