Quantitative Marketing and Economics (2021) 19:289–368
https://doi.org/10.1007/s11129-021-09238-x

Prices and promotions in U.S. retail markets

Günter J. Hitsch[1] · Ali Hortaçsu[2] · Xiliang Lin[3]

Received: 18 August 2019 / Accepted: 11 August 2021 / Published online: 29 October 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We provide generalizable results on the price and promotion tactics employed in the U.S. retail grocery industry. First, we document a large degree of price dispersion for UPCs and brands across stores, both nationally and at the local market level. Base price differences across stores and price promotions contribute to the overall price variance, and we show how to decompose the price variance into base price and promotion components. Second, we document that a large percentage of the variation in prices and promotion tactics across stores can be explained by retail chain and especially market/chain factors, whereas market factors explain only smaller percentage of the variation. Third, we show that the chain-level price and promotions similarity can be explained by similarity in demand. In particular, a large percentage of the variance in price elasticities and promotion effects can be explained by retail chain and especially market/retail chain factors. Further, price elasticities and promotion effects across stores of the same chain are hard to distinguish from the chain-market-level mean, and cross-price elasticities are typically imprecisely estimated. These findings suggest that retail managers may plausibly consider price discrimination across stores to be infeasible.

Keywords Price dispersion · Pricing · Promotions · Retail industry

JEL Classification D22 · L1 · L81 · M31

Researchers' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researchers and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

✉ Günter J. Hitsch
guenter.hitsch@chicagobooth.edu

Extended author information available on the last page of the article.

1 Introduction

What are the important features of the pricing and promotion tactics used in the U.S. retail grocery industry? We study this question using Nielsen RMS scanner data that record weekly store-level prices and quantities for nearly 50,000 products in 17,184 stores that belong to 81 different retail chains, including grocery stores, drug stores, and mass merchandisers. This comprehensive data set allows us to provide *generalizable insights* that are not specific to a small number of products, categories, stores, or retailers as in most of the extant literature. Similar (or the same) data have been widely used for academic research in marketing, industrial organization, and macroeconomics as well as by managers and analysts in the industry. Our goal is to provide a range of generalizable insights on pricing, promotion, and demand patterns that we hope will spur future research. Further, we intend these insights to inform industry practitioners to evaluate and improve current pricing and promotion tactics.

The first part of the paper documents the degree of price dispersion for (almost) identical products across stores at a given moment in time. We define products as UPCs,[1] which are identical across stores, and also as brands, which are "almost identical" aggregates of UPCs that differ in form or pack size but contain the exact same product content. The overall degree of price dispersion is large even at narrowly defined geographic levels. For instance, at the 3-digit ZIP code level, the ratio of the 95th to 5th percentile of prices is 1.294 for the median UPC and 1.433 for the median brand. Further, there is substantial heterogeneity in the degree of price dispersion across products.

The overall price dispersion may be due to variation in base prices, i.e. the regular shelf-price of a product, or due to price promotions when a product is sold at a temporary discount. To distinguish between these two sources of price dispersion we develop a new algorithm that classifies prices as base or promoted prices. We document a substantial degree of base price dispersion that is only modestly smaller than the overall degree of price dispersion, both nationally and at the local market level. Further, most products are frequently promoted. For instance, the median product is promoted once in 6.8 weeks, and the median promotional discount is 19.5%. Also, there is much heterogeneity in the promotion frequency and depth across stores.

To quantify the contribution of the different sources to the overall price dispersion, we decompose the overall variance in prices across stores and weeks in a year into separate components. Within markets, persistent UPC base price differences across stores account for the largest share of the price variance (46.5%), whereas the within-store variance of base prices during a year only accounts for a substantially smaller share (18%). Despite extensive promotional activity, the contribution of promotions to the market-level price variance is modest (35%). This is due to co-existing EDLP (every-day low price) vs. Hi-Lo pricing patterns, whereby stores with systematically high base prices offer deeper or more frequent price discounts than stores with systematically low base prices. Thus, the EDLP vs. Hi-Lo pattern compresses the overall price dispersion across stores.

[1] UPC is the acronym for "universal product code".

In the second part of the paper we study if differences across stores in prices and promotion tactics, i.e. promotion frequency and depth, are related to market versus retail chain-specific factors. We find that market (3-digit ZIP code) factors can explain a significant percentage of the national price variance (46.5% for the median product). However, a substantially larger percentage of the price variance is explained by retail chain factors (69.9%), and local market/retail chain factors in particular (88.1% for the median product). An analogous analysis for differences in promotion frequency and depth across stores reveals almost identical results. Further, it is not just that the overall frequency and depth of promotions are similar at the chain and especially at the market/chain level, but individual promotion events are coordinated, too. In particular, we document that the incidence of a store-level promotion is strongly associated with the percentage of other stores in the same retail chain that promote the product in the same week, both a the local market and at the national level.

The third part of the paper documents that the similarity in prices and promotions among stores in the same retail chain can be explained by similarity in demand. We estimate store-level price elasticities and promotion effects for 2,000 brands using similar data and a statistical approach (a Bayesian hierarchical model) that is employed by sophisticated analysts in the industry. Hence, our estimates are similar to the estimated effects used by retailers to make price and promotion decisions.

Mirroring our findings for price and promotion tactics, we find that a large percentage of the overall variance in price elasticities and promotion effects is explained by retail chain and especially local market/retail chain factors. In particular, using the estimates that most closely emulate the information available to industry analysts, 51.0% of the overall variation in price elasticities is explained by chain factors and 70.5% is explained by market/chain factors. This finding raises the question if retail managers can actually distinguish among price elasticities and promotion effects across stores. To address this question, we predict credible intervals for the estimates and document the percentage of store-level price and promotion effects with credible intervals that exclude the local chain-level average of the estimated effects. Using this approach, a large majority of the estimates is indistinguishable from the market/chain-level price and promotion effects. Further, our results show that store-level cross-price elasticities in particular are typically imprecisely estimated.

The difficulty to distinguish among local store-level price and promotion effects and to obtain precise cross-price elasticity estimates suggests that retail managers may plausibly consider price discrimination across stores as infeasible. This hypothesis is also consistent with the anecdotal evidence from conversations with retail chain managers, who frequently indicated that local price discrimination ("store-specific marketing") is challenging to implement in practice. Our explanation for price and promotion similarity is different, although not mutually exclusive, from other explanations that have been proposed in the literature, including managerial inertia (DellaVigna & Gentzkow, 2019), brand image and fairness concerns (Ater & Rigbi, 2020), and the softening of competition if firms can commit to less granular pricing (Adams & Williams, 2019; Dobson & Waterson, 2005). Industry insiders have also proposed that price discrimination across stores is infeasible in practice due to the institutional constraints of feature advertising, i.e. retailer advertising that highlights price promotions in local markets. Our analysis rejects this last hypothesis, however, because not only promoted but also base prices are similar across stores.

The paper is organized as follows. We discuss the related literature in Section 2, and provide an overview of the data sources in Section 3. Section 4 documents the basic facts of price dispersion, Section 5 separately documents the degree of base price dispersion and the prevalence of price promotions, and Section 6 provides a decomposition of the overall price variance into base price and promotion components. Section 7 shows how price and promotion differences across stores are related to market versus retail chain-specific factors, and Section 8 documents chain-level price similarity and promotion coordination. In Section 9 we examine how the chain-level similarity in prices and promotions is related to similarity in demand, and Section 10 examines if retailers can distinguish among store-level price elasticities and promotion effects. We discuss our proposed explanation for price similarity in Section 11 and relate it to alternative hypotheses. Section 12 concludes.

2 Literature review

Several studies have documented price dispersion in retail grocery stores, including (Lach, 2002; Eden, 2014; Dubois & Perrone, 2015; Eizenberg et al., 2021), although with less extensive product, market, or retail chain coverage. Most closely related to our analysis of price dispersion in Section 4 is Kaplan and Menzio (2015), which studies price dispersion using the Nielsen Homescan household panel data. The Homescan panel data only contains the prices of products *purchased*. Hence, because demand decreases in price, the panel data do not represent the full distribution of prices at which products were available for purchase. This problem affects our analysis to a substantially lesser degree, because we observe the price of a product whenever at least one unit was sold in a given store/week.[2] The size of the Homescan panel only allows to systematically capture the prices of a small number of products. Related, due to the limited sample size, Kaplan and Menzio (2015) present results for 54 Scantrack markets, compared to the 840 3-digit ZIP codes that we use as the most granular market definition. However, an advantage of the Homescan panel is that it contains purchase data from all retail chains, unlike the Nielsen RMS data used in our paper.[3]

Our paper documents a large degree of similarity in prices and promotions among stores in the same retail chain. Similar results are reported in Nakamura (2008) and DellaVigna and Gentzkow (2019). Our main explanation for this similarity is based on similarity in demand, and in particular the difficulty to distinguish among and to obtain precise price and promotion effect estimates at the store level. We will discuss this explanation and the relationship to other explanations that have been provided

[2] As explained in Section 3.1), we impute prices in weeks when a product did not sell using the most recent base (non-promoted) price.

[3] In work that has a different focus than our study, Kaplan et al. (2019) analyze the extent to which product-level price dispersion is due to persistent price-level differences across *stores*, based on a sample of 1,000 UPCs from the Nielsen RMS data.

by (Adams & Williams, 2019; DellaVigna & Gentzkow, 2019), and Ater and Rigbi (2020) in Section 11.

Related to our work, Adams and Williams (2019) provide evidence on retail price similarity in the home improvement industry. Similar to our findings, they document that prices are not store-specific, but cluster within pricing zones.

Our analysis, which focuses on the dispersion of prices at a given moment in time, is related to work that provides generalizable evidence on the time-series variation in prices (Bronnenberg et al., 2006). A related literature documents the frequency of price adjustments and price rigidity, which has important implications for macroeconomics (Nakamura & Steinsson, 2008, 2013). Our work is also related to research on assortments across stores. For a sample of products in four categories, Hwang et al. (2010) find that stores that belong to the same retail chain in a market (state) carry similar assortments. This finding mirrors our results on price and promotion similarities within retail chains.

Our analysis of store-level price elasticities and promotion effects in Section 9 is related to meta-analyses of price elasticities in the marketing literature, in particular Tellis (1988) and Bijmolt et al. (2005). Other studies relate brand or category elasticities and promotion effects to market characteristics (Bolton, 1989), demographic and competitor information (Hoch et al., 1995; Boatwright et al., 2004), and category characteristics (Narasimhan et al., 1996).

3 Data description

Our analysis primarily uses the Nielsen RMS (Retail Measurement Services) retail scanner data. The data are available for academic research through a partnership between the Nielsen Company and the James M. Kilts Center for Marketing at the University of Chicago Booth School of Business.[4] We also use the Nielsen Homescan consumer panel data to select the products in our sample.

3.1 Nielsen RMS retail scanner data

The RMS scanner data record sales units and prices at the week-level, separately for all stores and UPCs (universal product code). The data available from the Kilts Center for Marketing covers close to 40,000 stores across various channels, including grocery stores, mass merchandisers, drug stores, convenience stores, and gas stations. Although the data cover cover many stores and retailers, the subset of the RMS data available from the Kilts Center for Marketing is neither a census nor a randomly selected sample. However, the data have broad geographic coverage and, on average, account for between 50 and 60% of all market-level spending in grocery and drug stores and for one third of all spending at mass merchandisers.[5] The data allow us to

[4] https://www.chicagobooth.edu/research/kilts/datasets/nielsen

[5] See the *Retail Scanner Dataset Manual* provided by the Kilts Center for Marketing for the Scantrack market-level data indicating the coverage of spending for the three main retail channels.

identify all stores that belong to the same retailer. However, the exact identity (name) of a retail chain is concealed.

A week in the RMS data is a seven-day period that ends on a Saturday. If the shelf price of a UPC changes during this period the quantity-weighted average over the different shelf prices in the week is recorded.

The RMS data only contain records for weeks when at least one unit of the product was sold in a given store. Hence, in particular for small (in terms of revenue) products the incidence of weeks without price observations can be high. For such products the observed sample is more likely to include low, promoted prices than high, regular prices, and hence the sample will not represent the true price distribution. To ameliorate this problem, we impute the missing prices using on an algorithm that first classifies the observed prices as either base (regular) prices or promoted prices. The algorithm distinguishes between regular and promoted prices based on the frequently observed saw-tooth pattern in a store-level time series of prices whereby prices alternate between periods with (almost) constant regular price levels and shorter periods with temporarily reduced price levels. We perform this classification separately for each store-UPC pair. We assume that weeks without sales are non-promoted weeks, which is justified by the large sales spikes that are frequently observed in promoted weeks. Hence, we impute the missing prices using the predictions of the current regular (non-promoted) price levels based on the price classification algorithm. Two examples of observed store-level price series and the corresponding predicted base prices are given in the Appendix in Figs. 15 and 16.

3.2 Nielsen Homescan household panel data

We use the Nielsen Homescan household panel to select the product sample for our analysis. The participating households scan all purchased items after each shopping trip, and thus Homescan provides comprehensive data on the UPCs purchased and the corresponding prices paid. During the sample period in this paper, 2008-2010, the Homescan panel includes more than 60,000 households. Nielsen provides sampling weights ("projection factors") to make summary statistics from the data, such as total spending in retail stores, representative of the US population at large.

3.3 Sample selection

The 2008-2010 RMS data include information on almost one million UPCs. We intend our analysis to be as comprehensive as possible. However, including all products is challenging, in particular because "small" products rarely sell. For such products, store-level price observations are only available for a small percentage of weeks, and hence our price imputation algorithm (Section 3.1) is likely to yield noisy results.

Therefore, we use a subset of all products in our analysis. We select the UPCs that are observed in both the RMS scanner data and the Homescan household panel data, and we then choose the top 50,000 products based on total Homescan expenditure. We use the Homescan expenditure data to select products that are representative of the overall buying-behavior of US households. As discussed above, Homescan is

Table 1 Product sample descriptive statistics

Year	Observed prices (million)	Observed and imputed prices (million)	Revenue ($ million)
Panel A: Products (UPC's)			
2008	6100.0	8669.9	140156.7
2009	6352.7	9226.3	146830.2
2010	6355.4	9192.4	146608.4
All	18808.1	27088.6	433595.3
Percent imputed prices 30.57			
Panel B: Brands			
2008	2302.9	2958.2	140156.7
2009	2401.8	3131.4	146830.2
2010	2409.5	3146.7	146608.4
All	7114.2	9236.3	433595.3
Percent imputed prices 22.97			
Panel C: Percent private label			
Products	10.8	9.3	15.4
Brands	8.6	7.6	15.4

Note: The first column indicates the number of price observations obtained from the RMS data. The second column also includes the imputed prices for weeks with zero sales. The observation numbers are expressed in millions, and the revenue data are expressed in millions of dollars. The table also indicates the percentage of imputed prices among all observed and imputed prices, and the percentage of private label observations among all product or brand price observations and revenue

intended to be representative, whereas the stores and retailers in the RMS data need not be representative of the US population at large.

The 50,000 chosen products account for 73% of the total Homescan expenditure and 79% of the revenue in the RMS data.[6] Excluding some infrequently sold UPCs, our final sample contains 47,355 products. Table 1 provides summary statistics for the product/store/week level price observations. In each year we observe more than 6 billion prices. In total, there are 18.81 billion price observations corresponding to 434 billion dollars in revenue. Including the imputed prices the number of observations is 27.09 billion; 30.6% of the prices are imputed.

3.4 UPCs versus brands

We will frequently compare results that use UPCs as product definition with results for products defined as a brand. We obtain brand-level data by aggregating UPCs that share a common brand name. For example, all UPCs with the brand description

[6]See Appendix A.2 for details on the product size (revenue) distribution in our data.

"COCA-COLA CLASSIC R" or "COCA-COLA R" belong to the same brand, Coca-Cola, while products with the description "COCA-COLA DT" belong to a different brand, Diet Coke. The sales volume of a brand is measured in equivalent units, such as ounces or counts, and the brand price is measured as the average price per equivalent unit (e.g. 12 cents per ounce). We calculate *weighted* average prices, using the total product-level revenue over all stores and weeks as weights. Thus, differences in brand prices are entirely due to differences in the underlying UPC prices, not due to differences in the aggregation weights.

Using this aggregation process we obtain 11,279 brands. Summary statistics for the brand sample are shown in Table 1.

3.5 Chain and store coverage

Our sample includes data from 17,184 stores that belong to 81 different retail chains, including grocery stores, drug stores, and mass merchandisers. These stores represent more than 90% of the total revenue in the Kilts-Nielsen RMS data. For stores with a large incidence of weeks when the products in our sample did not sell, the predictions of the price imputation algorithm will be unreliable. We hence exclude these predominantly small stores, especially convenience stores and gas stations, from the analysis.

Table 8 in the Appendix summarizes the number of retail chains and stores in our sample at the DMA (designated market area) and ZIP+3 level.[7] Three-digit ZIP codes and counties are the smallest geographic areas accessible to researchers in the Kilts-Nielsen RMS data. Our sample contains 205 DMAs and 840 3-digit ZIP codes with at least one store. At the DMA level, the median number of chains is 6 and the median number of stores is 32. At the ZIP+3 level, the corresponding numbers are 4 and 10, respectively. Hence, even at the smallest geographic level there are multiple retailers and stores, and thus the measured price dispersion is unlikely to be limited due to a small number of store or retail chain observations.

3.6 Product assortments

The dispersion of prices will also be limited if many products are sold only in few stores. We provide a detailed analysis of UPC and brand availability across stores and retail chains in Appendix A.3. We find that most "large" (as measured by total revenue) products are widely available, but there are also many predominantly "small" UPCs and brands that are available only in few stores. We account for product size differences by analyzing the revenue-weighted distribution of prices, which places less weight on small products that are not widely available.

[7] Table 8 also provides chain-level summary statistics on the geographic coverage and the total number of stores of different retail chains.

4 Price dispersion: The basic facts

We start our analysis by documenting the price dispersion of identical or almost identical products across stores at any given moment in time. We present the results separately for products defined as UPCs or brands. UPCs are identical across stores. Brand prices are calculated as a weighted average over the prices of the individual UPCs that share the same brand name (Section 3.4). These UPCs typically differ along pack size (15 oz, 20 oz, etc.) or form factor (bottles, cans, etc.). Hence, the brand aggregates need not be exactly identical across stores. Apart from the packaging, however, the main product is *de facto* physically identical across UPCs, and consumers are likely to *perceive* the product content as identical. Hence, the comparison of weighted average brand prices across stores is meaningful, and we refer to brands as "almost identical products."

4.1 Price dispersion measures

For each product j we measure the dispersion of prices in week t using two statistics. Both statistics are calculated using the sample of store-level prices, $\mathcal{P}_{jt} = \{p_{jst} : s \in \mathcal{S}_{jt}\}$, where $\mathcal{S}_{jt}$ is the set of all stores that sell product j in week t. The first statistic is the standard deviation of the log of prices,

$$\sigma_{jt} = \sqrt{\frac{1}{N_{jt}-1} \sum_{s \in \mathcal{S}_{jt}} \left(\log(p_{jst}) - \overline{\log(p_{jt})}\right)^2}.$$

σ_{jt} measures dispersion based on the percentage price differences from the geometric mean across stores.[8] The second statistic, $r_{jt}(0.05)$, is the ratio of the 95th to the 5th percentile of the price observations $\mathcal{P}_{jt}$. We calculate the statistics for each week in 2010, and then take the mean over all weeks to report the average dispersion statistics σ_j and $r_j(0.05)$.

4.2 Price dispersion: UPCs

Figure 1 and Table 2 summarize the distribution of the price dispersion statistics across all 47,355 UPCs. To account for differences in the "importance" of each product we summarize the weighted distribution of the dispersion statistics using total product-level revenue across all stores and weeks in 2010 as weights.

The top row of Fig. 1 displays the weighted distribution of the price dispersion statistics at the national level. Overall, the degree of price dispersion for identical products across stores at any given moment in time is large. The log-price standard deviation, σ_j, for the median product is 0.161, which roughly indicates that 95% of prices vary over a range from 32% below to 32% above the average national price

[8] N_{jt} is the number of stores in $\mathcal{S}_{jt}$.

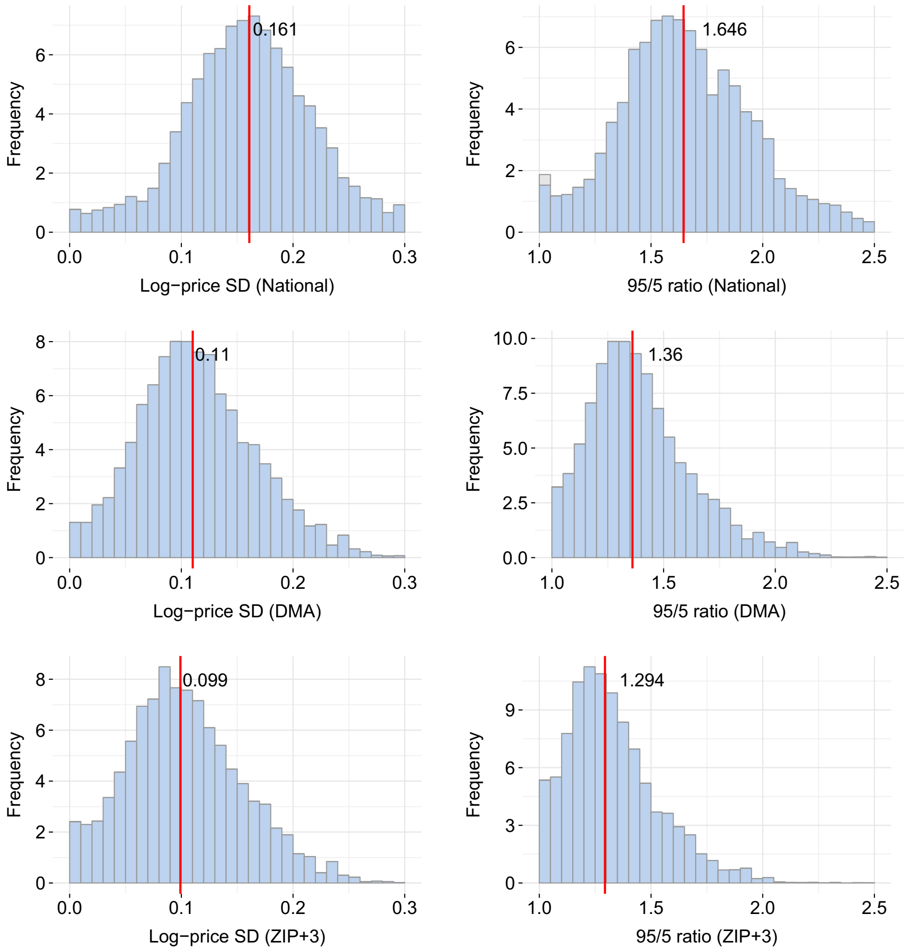

Fig. 1 Price dispersion statistics: UPC prices

of the median product.[9] The ratio of the 95th to 5th percentile of prices, $r_j(0.05)$, is 1.646 for the median product.

The large degree of price dispersion at the national level may simply reflect systematic differences in price levels across regions. Our main focus, however, is on documenting the price dispersion of identical products in local markets, where consumers could at least in principle buy the same product from any of the local stores. To account for systematic regional price differences, we calculate the dispersion statistics separately for each market m using the price observations $\mathcal{P}_{jmt} = \{p_{jts} : s \in \mathcal{S}_{jmt}\}$.[10] We then take the weighted average over all market-level dispersion

[9] The median product is the median of the revenue-weighted distribution of σ_j.
[10] $\mathcal{S}_{jmt}$ is the set of all stores that sell product j in market m in week t.

Table 2 Price and base price dispersion statistics

		Median	Mean	Percentiles							
				0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
Prices											
Product definition: UPC											
Log-price SD	National	0.161	0.163	0.015	0.059	0.090	0.124	0.200	0.238	0.268	0.318
	DMA	0.110	0.114	0.006	0.032	0.050	0.078	0.147	0.185	0.208	0.249
	ZIP+3	0.099	0.103	0.002	0.021	0.039	0.066	0.136	0.174	0.196	0.236
95/5 percentile ratio	National	1.646	1.690	1.018	1.171	1.300	1.460	1.869	2.103	2.331	2.836
	DMA	1.360	1.398	1.011	1.075	1.129	1.232	1.524	1.721	1.831	2.071
	ZIP+3	1.294	1.329	1.004	1.045	1.094	1.182	1.438	1.614	1.713	1.946
Product definition: Brand											
Log-price SD	National	0.175	0.185	0.074	0.098	0.114	0.141	0.216	0.260	0.320	0.427
	DMA	0.138	0.146	0.045	0.067	0.084	0.105	0.174	0.214	0.250	0.356
	ZIP+3	0.129	0.137	0.035	0.060	0.072	0.097	0.165	0.202	0.239	0.341
95/5 percentile ratio	National	1.728	1.841	1.232	1.340	1.414	1.537	1.981	2.289	2.753	4.002
	DMA	1.508	1.583	1.113	1.208	1.264	1.367	1.697	1.929	2.146	3.042
	ZIP+3	1.433	1.492	1.081	1.170	1.211	1.303	1.593	1.786	1.978	2.666
Base prices											

Table 2 (continued)

		Median	Mean	Percentiles							
				0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
Product definition: UPC											
Log-price SD	National	0.133	0.138	0.006	0.046	0.073	0.103	0.168	0.205	0.240	0.320
	DMA	0.088	0.093	0.002	0.023	0.039	0.062	0.118	0.151	0.177	0.239
	ZIP+3	0.078	0.083	0.001	0.014	0.029	0.052	0.107	0.140	0.166	0.225
95/5 percentile ratio	National	1.516	1.564	1.001	1.130	1.230	1.373	1.702	1.933	2.131	2.710
	DMA	1.276	1.310	1.003	1.050	1.096	1.177	1.399	1.556	1.680	1.983
	ZIP+3	1.220	1.250	1.001	1.028	1.066	1.134	1.328	1.464	1.569	1.832
Product definition: Brand											
Log-price SD	National	0.162	0.174	0.062	0.091	0.101	0.127	0.206	0.253	0.310	0.425
	DMA	0.126	0.137	0.034	0.062	0.071	0.093	0.166	0.209	0.244	0.357
	ZIP+3	0.118	0.127	0.025	0.055	0.064	0.086	0.157	0.197	0.231	0.345
95/5 percentile ratio	National	1.654	1.780	1.187	1.297	1.350	1.478	1.912	2.229	2.671	3.857
	DMA	1.460	1.541	1.081	1.192	1.225	1.315	1.659	1.893	2.133	2.918
	ZIP+3	1.389	1.453	1.054	1.154	1.185	1.257	1.564	1.759	1.953	2.645

statistics using the number of observations in each market as weights. We use two market definitions: DMAs (designated market areas) and 3-digit ZIP codes.[11]

The market-average price dispersion statistics are shown in the middle and lower rows of Fig. 1 (see Table 2 for detailed numbers). The log-price standard deviation for the median product is 0.110 at the DMA level and 0.099 at the 3-digit ZIP code level, and the corresponding $r_j(0.05)$ values are 1.360 and 1.294 respectively. Hence, even in local markets there is a large degree of price dispersion for identical products.

While the *overall* degree of price dispersion of identical products at any given moment in time is large, Figure 1 and Table 2 also reveal another, equally important fact: There is substantial heterogeneity in the dispersion statistics across products. At the market (3-digit ZIP code) level, the standard deviation of log prices ranges from from 0.021 at the 5th percentile to 0.196 at the 95th percentile of the dispersion statistics. Similarly, the 95th to 5th percentile ratio of prices ranges from 1.045 to 1.713 when comparing the 5th and 95th percentile values.

4.3 Price dispersion: Brands

As discussed in Section 3.4, differences in brand prices across stores with the same brand-level assortments of UPCs are entirely due to differences in the UPC prices. However, brand prices may differ across stores if the assortments differ, even if the underlying UPC prices are identical.

Figure 2 and Table 2 summarize the distributions of the brand-price dispersion statistics. Generally, the degree of price dispersion is larger at the brand level compared to the UPC level. Nationally, the standard deviation of log brand prices is 0.175, compared to 0.161 for UPC prices. At the 3-digit ZIP code level the corresponding numbers are 0.129 (brands) and 0.099 (UPCs), respectively.

The large degree of heterogeneity in price dispersion that we documented for UPCs is also evident for brands. For example, at the 3-digit ZIP code level the standard deviation of log brand prices is 0.060 at the 5th percentile, compared to 0.239 at the 95th percentile.

We refer the reader to Section B in the Appendix for some additional results, in particular a sensitivity analysis that uses two alternative measures of price dispersion. Section B also compares our results to the results in Kaplan and Menzio (2015).

5 Base prices and promotions

For many products, prices alternate between periods when the product is sold at the base price, i.e. the regular or every-day shelf price, and periods when the product is promoted, i.e. offered at a discount. Base prices change only infrequently. The large degree of price dispersion documented in the previous section could be due to

[11] A small number of markets (2 DMAs and 45 3-digit ZIP codes) contain only one store. We exclude these markets from the analysis, and we also exclude markets where only one store carries product j.

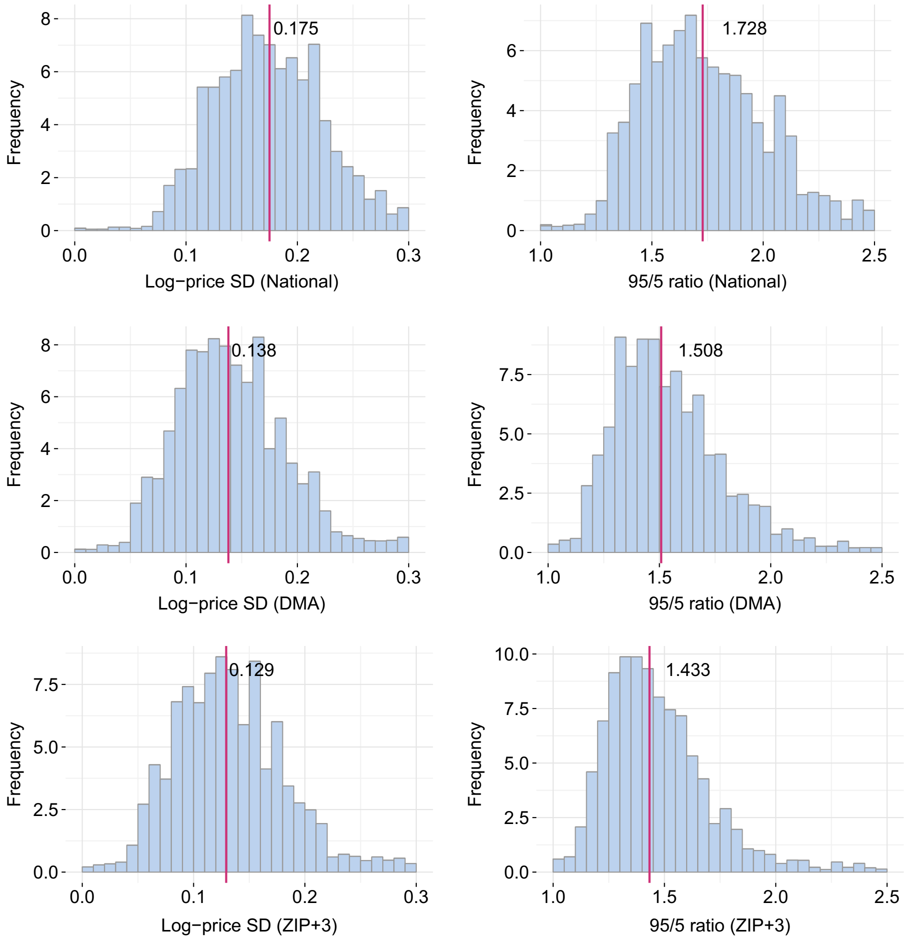

Fig. 2 Price dispersion statistics: Brand-level prices

differences in base prices across stores, reflecting a relatively persistent component in the dispersion of prices, or due to price promotions that are idiosyncratic to stores. In this section we document the dispersion of base prices and the promotion policies for the products in our data. We also document the "importance" of price promotions based on the percentage of the total product volume sold on promotion.

5.1 Base price dispersion

We measure the dispersion of base prices, $\mathcal{B}_{jt} = \{b_{jst} : s \in \mathcal{S}_{jt}\}$, using the same statistics as before, the standard deviation of the log of base prices and the ratio of the 95th to the 5th percentile of base prices across stores. We show the results for UPCs in Fig. 3 and the results for brands in Fig. 19 in the Appendix. Detailed summary statistics are in Table 2.

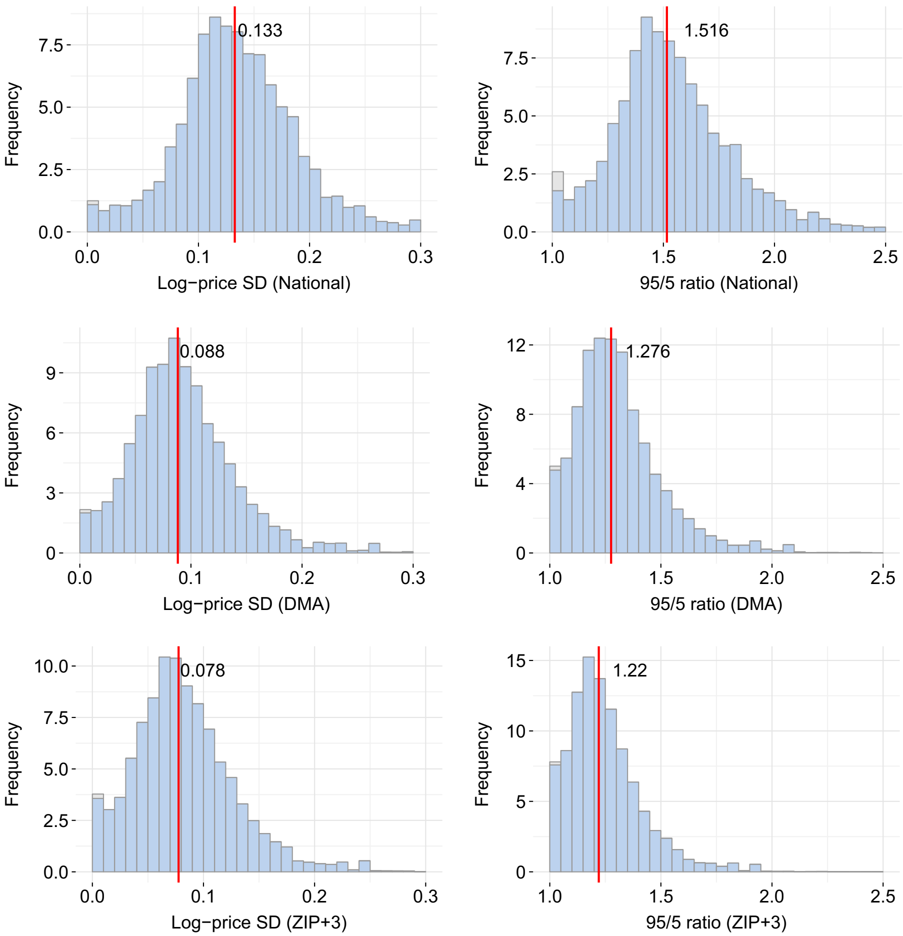

Fig. 3 Base prices dispersion statistics: UPC base prices

Generally, the degree of base price dispersion is substantial. The standard deviation of the log of base prices is 0.133 for the median product at the national level and 0.078 at the 3-digit ZIP code market level. The degree of base price dispersion is not much smaller than the degree of price dispersion. For example, the standard deviation of the log of prices at the 3-digit ZIP code level is 0.099, compared to 0.078 for the log of base prices. For brands, the difference is even smaller. Similar to the results in Section 4, there is much heterogeneity in the base price dispersion statistics across products.

5.2 Price promotions

The Nielsen RMS data do not directly indicate if a product was promoted. Instead, we infer a price promotion from the difference between the imputed base price and the realized price. We define the percentage price discount, or promotion depth, as follows:

$$\delta_{jst} = \frac{b_{jst} - p_{jst}}{b_{jst}}.$$

We then classify the product as promoted if the percentage price discount exceeds a threshold $\bar{\delta}$. The variable $D_{jst} = \mathbb{I}\{\delta_{jst} \geq \bar{\delta}\}$ captures promotional events, where $D_{jst} = 1$ indicates a promotion. We use a promotion threshold of $\bar{\delta} = 0.05$ in our analysis, a choice that we justify in detail in Appendix C.1. In short, small percentage price discounts where δ_{jst} is close to 0 are unlikely to represent a planned price promotion, but are likely due to measurement error in prices.

The price dispersion analysis so far was based on data from 2010. Here we extend the sample period to 2008-2010 to reduce measurement error, in particular in the promotion frequency statistic discussed below.

5.2.1 Promotion frequency

We measure the promotion frequency of product j in store s as

$$\pi_{js} = \frac{1}{N_{js}} \sum_{t \in \mathcal{T}_{js}} D_{jst},$$

where $\mathcal{T}_{js}$ includes the weeks in 2008-2010 when product j was sold in store s, and N_{js} is the corresponding number of observations. We calculate the product-level promotion frequency, π_j, by taking a weighted average of π_{js} across all stores using the number of store-level observations as weights.[12] We measure the heterogeneity in the promotion frequency across stores using the difference between the 95th and the 5th percentile of the π_{js} observations (weighted by N_{js}).

The left panel in the middle row of Fig. 4 displays the weighted distribution (using total product revenue) of the promotion frequency across products (see Table 3 for key summary statistics). The average promotion frequency for the median product is 0.147, implying that the product is promoted about once in 6.8 weeks. The promotion frequency varies strongly across products, ranging from 0.011 (once in 91 weeks) at the 5th percentile to 0.370 (once in 2.7 weeks) at the 95th percentile level. The left panel in the middle row of Fig. 4 shows the corresponding differences between the 95th and the 5th percentile of π_{js} across stores s. For the median product this difference is 0.314, compared to the median average promotion frequency of 0.147. Hence, there are large differences in the promotion frequency of a given product across stores.

[12]This is equivalent to calculating π_j based on all D_{jst} observations, pooled across stores and weeks.

Table 3 Price promotion statistics

	Median	Mean	Percentiles							
			1%	5%	10%	**25%**	75%	90%	95%	99%
Pooled across products, stores, and weeks										
Percentage price discount	0.177	0.207	0.013	0.032	0.051	0.101	0.287	0.402	0.486	0.595
Product-level statistics										
Promotion frequency	0.147	0.163	0.001	0.011	0.023	0.074	0.231	0.317	0.370	0.502
Promotion depth	0.195	0.201	0.083	0.102	0.119	0.149	0.246	0.290	0.315	0.390
Percentage volume on promotion	0.287	0.298	0.002	0.018	0.046	0.157	0.426	0.544	0.614	0.755
Promotion multiplier	3.043	3.783	1.054	1.335	1.610	2.169	4.458	6.790	9.220	23.297
Product-level statistics: Differences between 95th and 5th percentiles across stores										
Promotion frequency difference	0.314	0.316	0.007	0.046	0.083	0.184	0.436	0.545	0.591	0.667
Promotion depth difference	0.180	0.190	0.023	0.070	0.096	0.135	0.233	0.292	0.342	0.444
Percentage promoted volume difference	0.541	0.510	0.013	0.085	0.163	0.381	0.658	0.758	0.821	0.933
Promotion multiplier difference	5.531	7.854	0.603	1.316	1.897	3.225	9.096	16.289	25.518	65.013

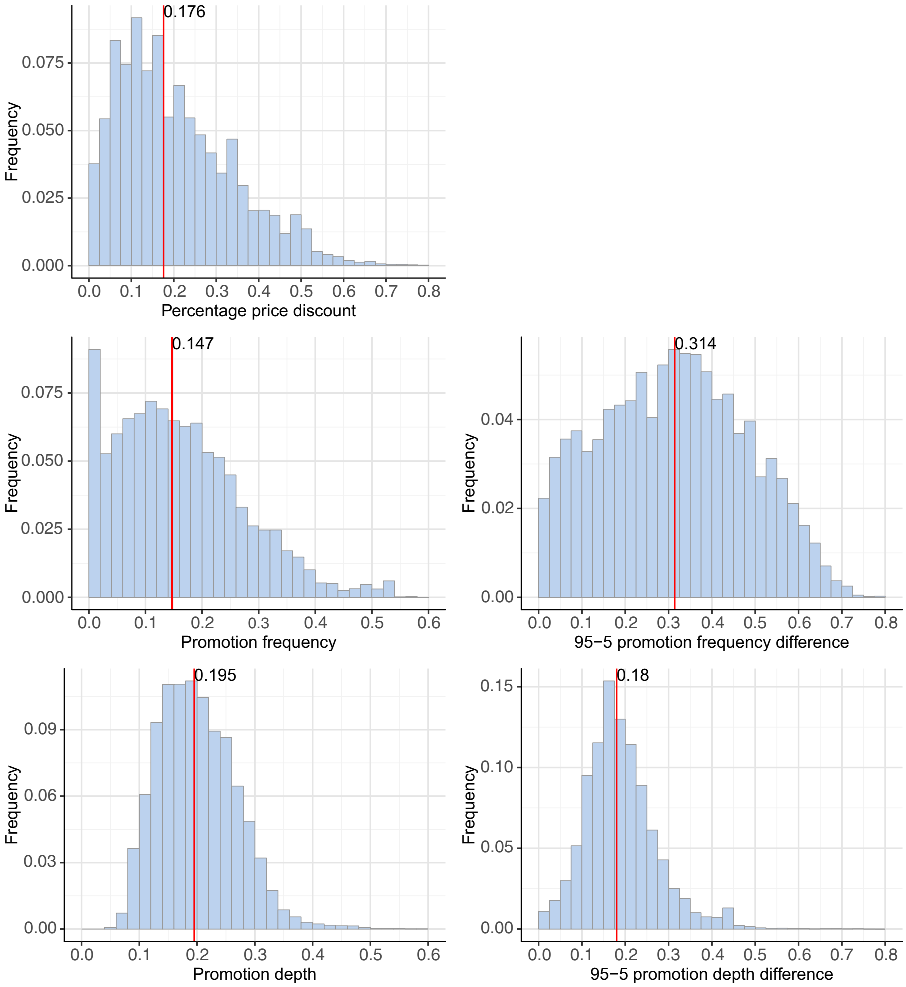

Fig. 4 Promotion depth, frequency, and differences in promotion depth and frequency across stores. Note: The top-left panel displays the distribution of the percentage price discounts, δ_{jst}, pooled across all products j, stores s, and weeks t when the price is strictly less than the base price, $p_{jst} < b_{jst}$. The middle and bottom-left panels display the weighted distribution of promotion frequency and promotion depth across products, j. Here, promotion depth is measured conditional on the product being promoted at a 5 percent promotion threshold. The middle and bottom-right panels summarize across-store differences in promotion frequency and promotion depth for all products. In particular, for each product j the differences are based on the 95th and 5th percentile of promotion frequency and promotion depth across all stores where the product is sold

5.2.2 Promotion depth

The promotion depth for product j in store s, i.e. the average promotional discount across all weeks when the product was promoted, is given by

$$\delta_{js} = \frac{1}{N_{js}^D} \sum_{t \in \mathcal{T}_{js}, D_{jst}=1} \delta_{jst},$$

where N_{js}^D is the number of promotion events. The product-level promotion depth statistic, δ_j, is the weighted average of δ_{js} across all stores using N_{js}^D as weights. As for the promotion frequency, we measure the heterogeneity in the promotion depth across stores using the difference between the 95th and the 5th percentile of δ_{js} across stores (the distribution of δ_{js} is weighted using N_{js}^D).

The bottom row of Fig. 4 shows the distributions of the average promotion depth, δ_j, and the across-store heterogeneity in δ_{js}. The promotion depth for the median product is 19.5%, and the distribution across products ranges from 10.2% at the 5th percentile to 31.5% at the 95th percentile. As for the promotion frequency, there is a much heterogeneity in the promotion depth across stores.

5.3 Volume sold on promotion and promotion multipliers

We measure the "importance" of promotions using the percentage of product volume that is sold during a promotional period. We also document the ratio of the average product volume sold during a promotion relative to the average product volume when the product was not promoted. In the retail industry and in brand management, this ratio is called a *lift factor* or *promotion multiplier*.

The top left panel in Fig. 5 displays the weighted distribution of the percentage volume sold on promotion across products,[13] and Table 3 provide detailed summary statistics. The median percentage of volume sold on promotion is 28.7%, and ranges from 1.8% to 61.4% at the 5th and 95th percentiles. The bottom left panel displays the corresponding distribution of the promotion multipliers, with a median of 3.04 and a range from 1.34 to 9.22. Hence, as expected, the volume sold on promotion is disproportionately high (relative to the overall incidence of promotions), and units sales spike relative to the non-promoted volume when a product is promoted. The right panels in Fig. 5 reveal a large degree of heterogeneity in these two statistics across stores for most products.

6 Price variance decomposition

We documented both a substantial degree of base price dispersion and that many products are frequently promoted. To quantify the contribution of these factors to the

[13] The weights are given by total product revenue.

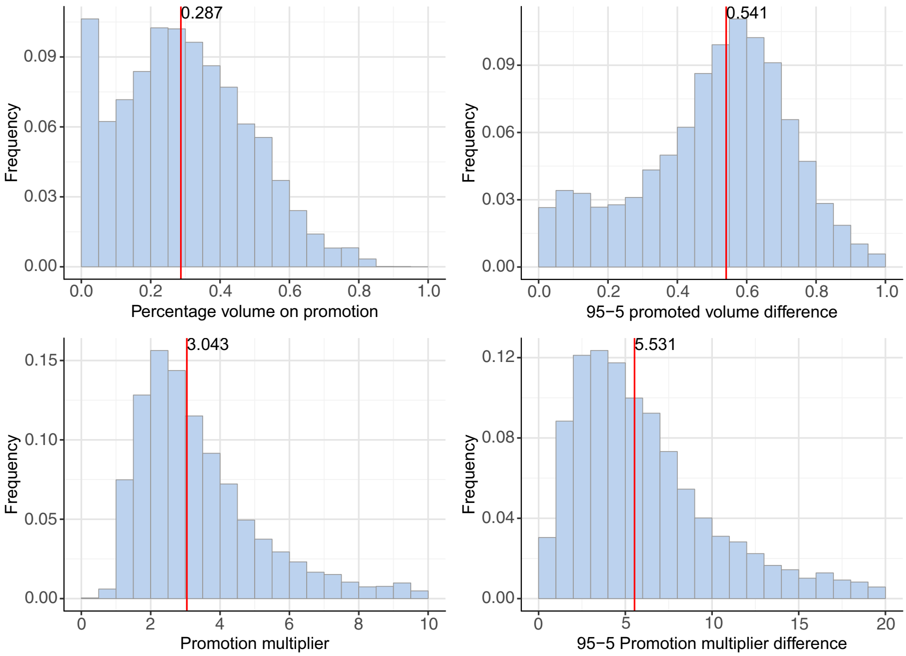

Fig. 5 Promotion volume and multiplier, and differences in volume and multipliers across stores

overall price dispersion, we decompose the price variance of a product into components that capture (i) price differences across markets, (ii) persistent price or base price differences across stores within markets, (iii) within-store price or base price variation over time, and (iv) price variation due to promotions.

We perform the variance decomposition separately for each product (UPC or brand). To simplify the notation, we drop the product subscript j. $\mathcal{M}$ is the set of all markets, and $\mathcal{S}$ is the set of all stores. For each store s we observe prices in periods $t \in \mathcal{T}_s$. N_s is the number of observations for stores s, N_m is the total number of observations (across stores and time periods) in market m, and N is the total number of observations across all markets. $\bar{p}$ is the overall (national) average price, $\bar{p}_m$ is the average price in market m, and $\bar{p}_s$ is the average price in store s. The overall price variance of a product is

$$\text{var}(p_{st}) = \frac{1}{N} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}_s} (p_{st} - \bar{p})^2.$$

6.1 Basic decomposition

The basic decomposition quantifies the contribution of price variation across markets, within-market price variation across stores, and the variation of prices within

Table 4 Price variance decomposition

	UPC	Brand	Within-Market UPC	Within-Market Brand
Basic decomposition				
Across-market	32.7	29.7		
Across-store	27.0	42.3	40.2	60.1
Within-store	40.3	28.0	59.8	39.9
Decomposition into base prices and promotions				
Across-market	32.7	29.7		
Across-store mean base price variance	31.3	49.9	46.5	70.9
Within-store base price variance	12.3	13.4	18.3	19.0
Total contribution of promotions	23.7	7.0	35.2	10.0
Promotional price discounts	36.0	29.9	53.5	42.6
EDLP vs. Hi-Lo adjustment	−12.3	−22.9	−18.3	−32.6

stores:[14]

$$\begin{aligned} \text{var}(p_{st}) &= \text{var}(\bar{p}_m) && \text{(across-market)} \\ &+ \tfrac{1}{N} \textstyle\sum_{m \in \mathcal{M}} N_m \text{var}(\bar{p}_s | m) && \text{(across-store)} \\ &+ \tfrac{1}{N} \textstyle\sum_{s \in \mathcal{S}} N_s \text{var}(p_{st} | s). && \text{(within-store)} \end{aligned} \quad (1)$$

$\text{var}(\bar{p}_m)$ is the variance of average market-level prices across markets, $\text{var}(\bar{p}_s|m)$ is the within-market variance of average store-level prices, and $\text{var}(p_{st}|s)$ is the within-store variance of prices over time.[15]

We report the revenue-weighted mean of the variance components in Table 4 using the 2010 data. Price-level differences across markets (ZIP+3 areas) account for 32.7% of the overall price variance for UPCs and 29.7% for brands. Hence, more than two-thirds of the national price variance is due to price variation within markets. Price-level differences across stores within markets account for 27% of the overall variance for UPCs and 42.3% for brands. Furthermore, the contribution of within-store price variation is 40.3% for UPCs and 28% for brands.

[14] All results are derived in detail in Appendix D.

[15] $\text{var}(p_m)$ and $\text{var}(\bar{p}_s|m)$ are calculated as weighted averages using the number of observations in each market and the number of observations for each store as weights (see Appendix D).

6.2 Decomposition into base prices and promotions

We now provide a more detailed decomposition that distinguishes between the contribution of base price differences and price promotions to the overall price variance

$$
\begin{aligned}
\mathrm{var}(p_{st}) &= \mathrm{var}(\bar{p}_m) && \text{(across-market)}\\
&+ \tfrac{1}{N} \textstyle\sum_{m\in\mathcal{M}} N_m \mathrm{var}(\bar{b}_s|m) && \text{(across-store base price var.)}\\
&+ \tfrac{1}{N} \textstyle\sum_{s\in\mathcal{S}} N_s \mathrm{var}(b_{st}|s) && \text{(within-store base price var.)}\\
&+ \tfrac{1}{N} \textstyle\sum_{m\in\mathcal{M}} N_m \mathrm{var}(b_{st}-p_{st}|m) && \text{(promotional discount var.)}\\
&- 2\tfrac{1}{N} \textstyle\sum_{m\in\mathcal{M}} N_m \mathrm{cov}(b_{st}-p_{st}, b_{st}|m) && \text{(EDLP vs. Hi-Lo adjustment)}
\end{aligned}
\tag{2}
$$

The first component captures the variance of price levels across markets. The second component is the within-market variance of base prices, indicating persistent base price differences across stores, whereas the third component is the within-store variance in base prices over time.

The last two terms in Eq. 2 capture the contribution of price promotions to the overall price variance. The fourth term is the variance of promotional price discounts, $b_{st} - p_{st}$, which is zero in the absence of price promotions. The last term is negative if the weighted average of the covariances between promotional price discounts and base prices is positive. A positive correlation indicates an EDLP (everyday low price) vs. Hi-Lo pricing pattern at the product level: Stores with above average base prices offer larger promotional price discounts than stores with below average base prices. Correspondingly, we call the last term in the decomposition (2) the "EDLP vs. Hi-Lo adjustment." In the absence of an EDLP vs. Hi-Lo pricing pattern, price promotions increase the overall price variance. However, if there is EDLP vs. Hi-Lo pricing, the adjustment term will be negative and the overall variance in prices will be reduced. Intuitively, EDLP vs. Hi-Lo pricing compresses the price distribution and thus reduces the variance in prices. In Appendix D.3 we present an example that shows that price promotions may even *decrease* the overall price variance in the presence of an EDLP vs. Hi-Lo pricing pattern.

The results in Table 4 show that the within-market variance of average store-level base prices accounts for 31.3% of the overall price variance for UPCs and for 49.9% of the price variance for brands. The within-store base price variation over the course of a year accounts for a much smaller percentage of the overall price variance, 12.3% for UPCs and 13.4% for brands. Hence, store-level base prices are fairly persistent over the course of a year.

Of particular interest if the role of price promotions. The promotional price discount component in Eq. 2 is large and positive—36.0% for UPCs and 29.9% for brands. However, the EDLP vs. Hi-Lo adjustment term is negative, -12.3% for UPCs and -22.9% for brands. Hence, there is strong evidence for EDLP vs. Hi-Lo pricing at the product-level, which reduces the contribution of price promotions to the overall price variance. Table 4 also shows the total contribution of promotions to the overall price variance, 23.7% for UPCs and 7.0% for brands.[16]

[16]The total contribution is the sum of the last two components in Eq. 2.

The last two columns in Table 4 express the contributions of the across-store and within-store base price variances and the contribution of price promotions as a percentage of the *within-market* price variance. The results highlight the importance of persistent base price differences across stores (46.5% for UPCs and 70.9% for brands) relative to the total contribution of price promotions (35.2% for UPCs and 10% for brands).

More detailed statistics, including the key percentiles of the price variance components, are provided in Table 9 in the Appendix. In particular, Table 9 reports the percentage of products for which price promotions *decrease* the overall price variance: 3.6% for UPCs and 30.8% for brands.

A key result from the price variance decomposition is that the large differences in base (non-promoted) prices across stores in a local market persist over time. This finding implies the presence of potentially large consumer search costs (e.g. Sorensen, 2000, or Honka et al., 2019 for a general survey) or rational inattention (e.g. Joo, 2020).

7 Price dispersion at the market and retail chain level

The analysis so far has shown that some of the overall, national price variance is due to market-specific factors. We now extend this analysis and compare how much of the overall variance in prices and differences in promotion tactics across stores can be attributed to market versus retail chain-specific factors.

To measure the percentage of the price variance that can be explained by market or chain factors, we regress the price of a UPC, p_{jst}, on indicator variables for (i) all markets (3-digit ZIP codes), (ii) all retail chains, or (iii) all market/retail chain combinations. Specifically, we estimate three *separate* regressions corresponding to one of the three sets of indicator variables. The regressions are performed separately for each product j and week t in 2010, and we take the average over all weeks to obtain a single R^2 value for each product.

The revenue-weighted distribution of the R^2 values is shown in the top row of Fig. 6 (Table 10 in the Appendix contains detailed numbers). For the median product, 46.5% of the overall price variance is explained by local market factors.[17] Chain-specific factors explain 69.9% of the price variance for the median product, and 88.1% of the price variance is explained by market/chain factors. Hence, prices are substantially more homogenous within the 81 different retail chains than within the 840 different 3-digit ZIP code areas in our data, and prices are particularly homogenous in retail chains at the local market level. To further illustrate the large difference in price homogeneity at the market versus market/chain level, note that for 90% of

[17] The R^2 values from the market indicator regressions are comparable to the across-market price variance component in the variance decomposition (1). The values are not identical, because the variance decomposition in Section 6 is performed using all weeks in 2010, whereas the product-level R^2 values in this section are obtained by averaging over the R^2 values from separate regressions for each week.

products the R^2 is at least 26.6% based on market factors and at least 70.7% based on market/chain factors.

We perform a similar analysis for the store-level promotion frequency and promotion depth, π_{js} and δ_{js}. The results are shown in the middle and bottom rows of Fig. 6. The results mirror the previous findings. For the median product, 36.1% of the variation in the promotion frequency across stores is explained by market factors, 62.3% is explained by retail chain factors, and 80.0% is explained by market/chain factors. The corresponding findings for promotion depth are similar.

We visualize the similarity of prices at the retail chain and market level for the case of Tide HE Liquid Laundry Detergent (100 oz) in Figs. 7 and 8. The figures display a two-dimensional representation of the store-level time series of prices using

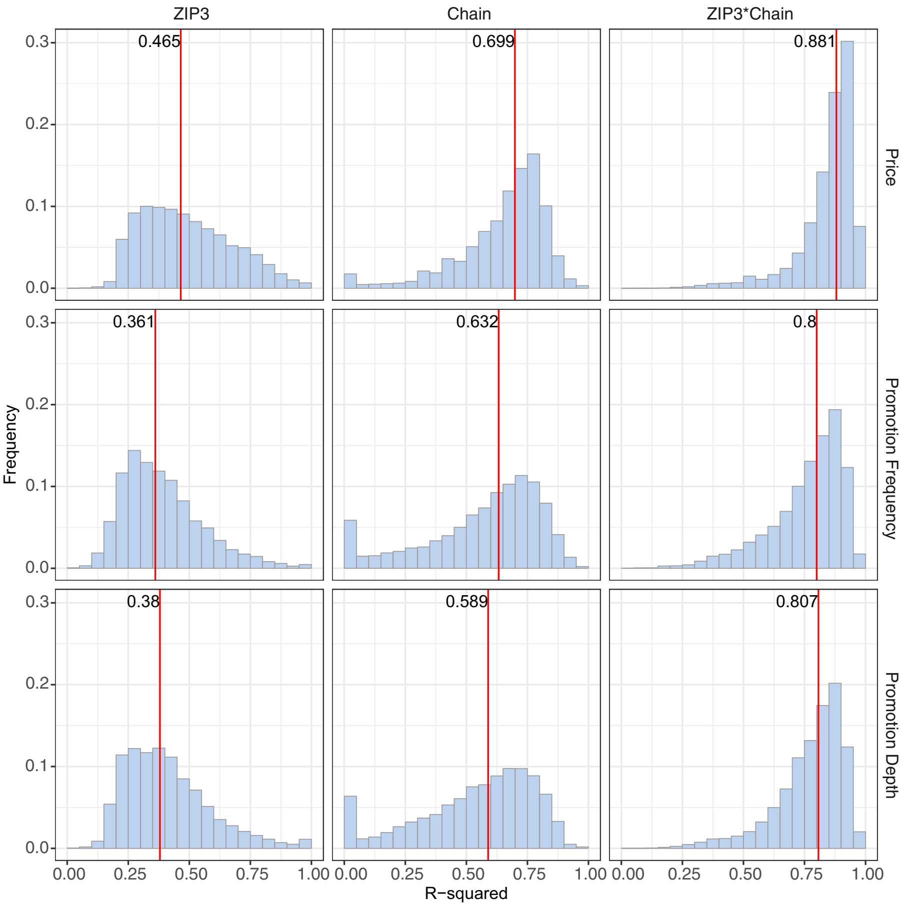

Fig. 6 Percentage of variance of prices, promotion frequency, and promotion depth explained by market and chain factors

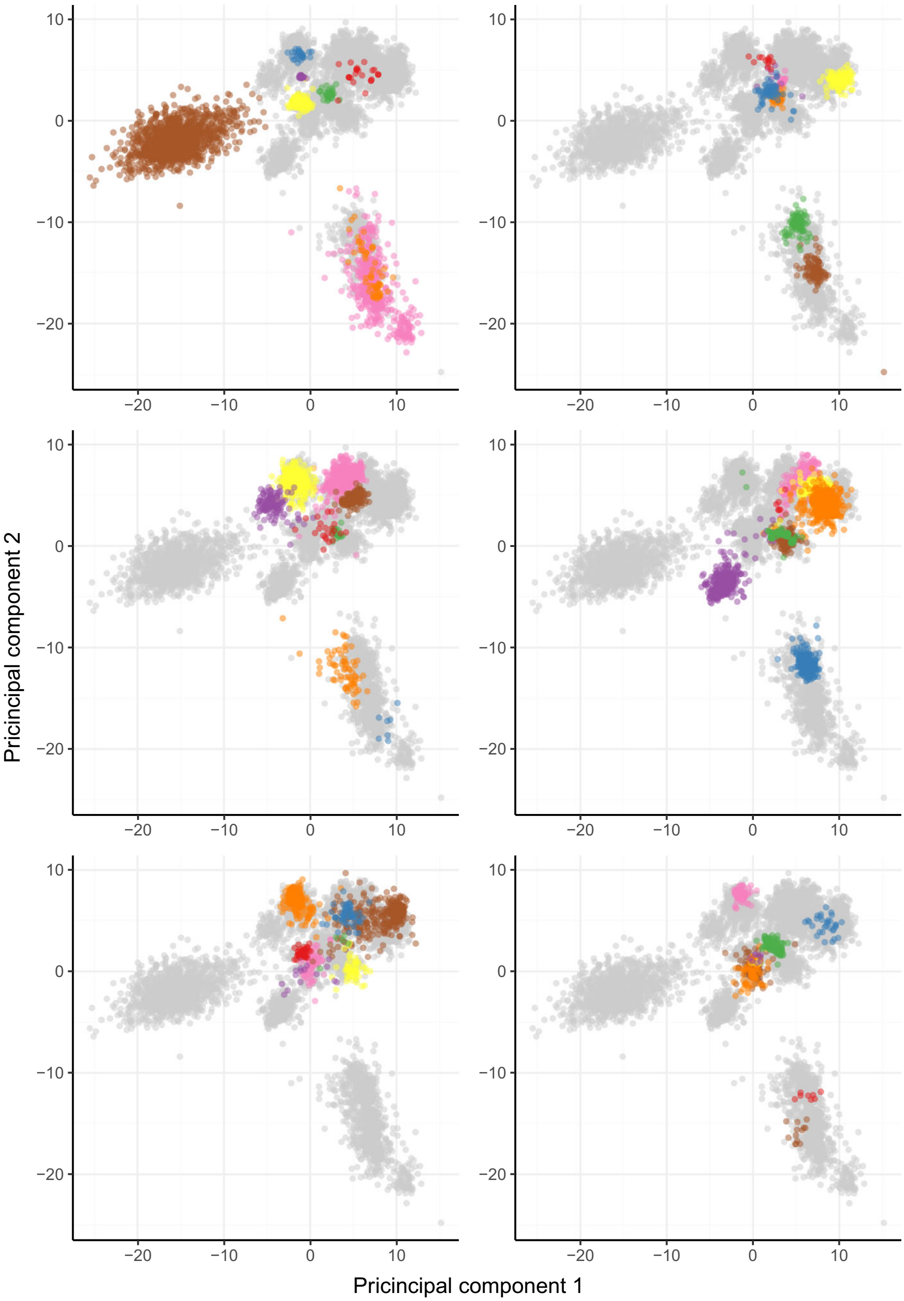

Fig. 7 Projected store-level price vectors colored by retail chain

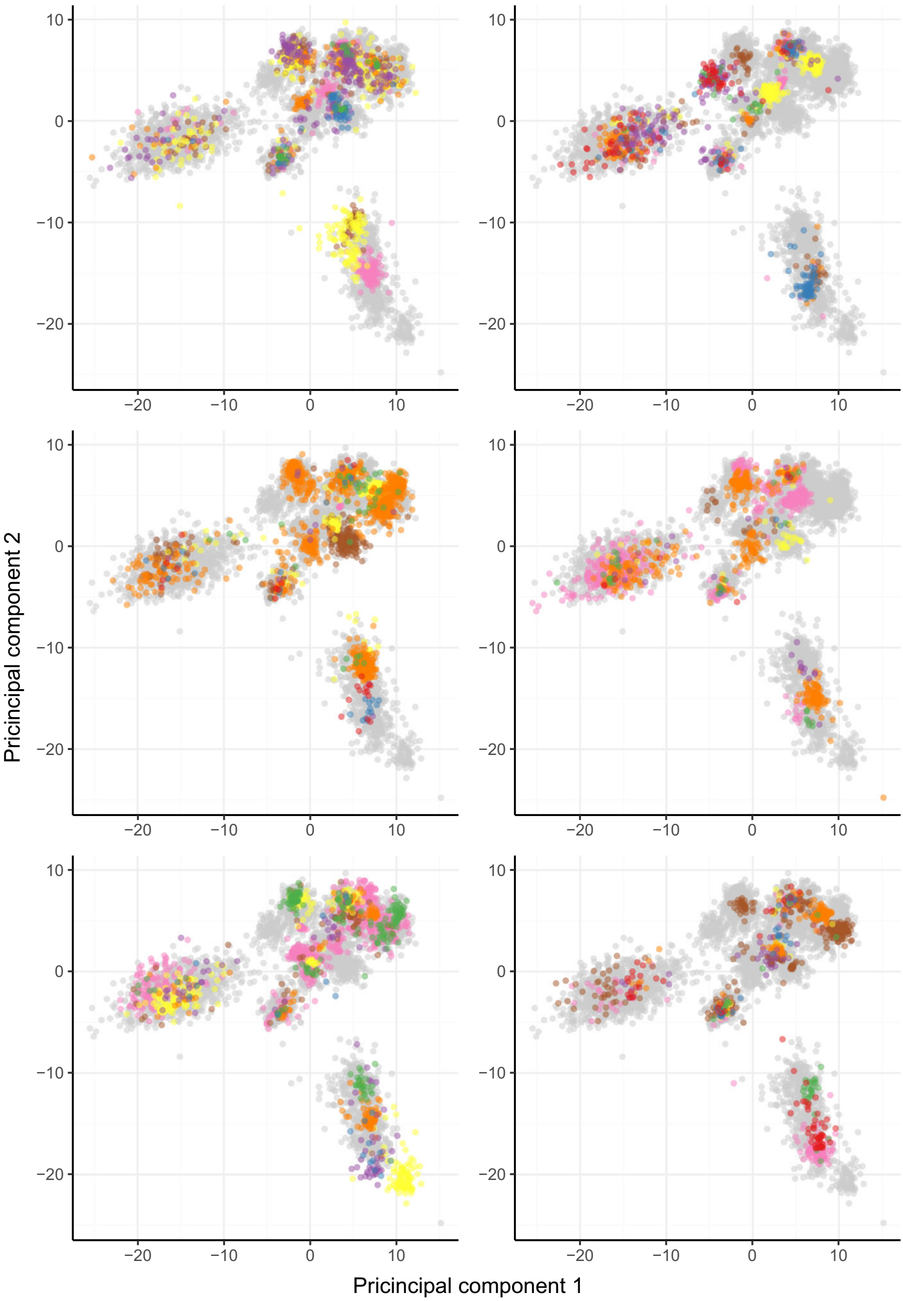

Fig. 8 Projected store-level price vectors colored by DMA

the projections onto the first two principal components.[18] Figure 7 is split into six panels that contain identical gray dots representing all projected store-level prices. In each of the panels some of the dots are colored according to the retail chain that the prices belong to. All stores that belong to the same retail chain appear in exactly one panel.[19] Figure 7 shows that the projected price vectors that belong to the same retail chain cluster and exhibit much less variance compared to the overall variance in projected prices. In Fig. 8 we color the projected store-level price vectors according to the market (DMA) that a store belongs to.[20] There is some clustering of prices also at the market level, but the similarity of prices within a market appears much smaller than the similarity of prices within a retail chain.

8 Chain-level price similarity and promotion coordination

We now provide a more detailed analysis of the similarity in product pricing and promotions at the retail chain level.

8.1 Price dispersion at the retail chain level: Summary statistics

We quantify the chain-level price dispersion using the approach used to measure price dispersion at the national and market level in Section 4. Table 5 shows the results for UPCs. The overall log-price standard deviation at the chain level is 0.079 for the median product. At the market level, the chain-level price dispersion, 0.039, is substantially smaller, both compared to the overall chain-level price dispersion and the market-level price dispersion, which is 0.099 at the 3-digit ZIP code level. The smaller degree of chain-level price dispersion at the market level compared to the national level is evidence for *zone pricing* (Adams & Williams, 2019).

We discussed in Section 5.2 that the discrepancy between a retailer's promotion calendar and the Nielsen RMS week definition creates measurement error that may exaggerate the variance in chain-level prices. However, this measurement error predominantly affects promoted prices, not base prices. For base prices, the log-price standard deviation is 0.078 at the chain level versus 0.027 at the chain/market (3-digit ZIP code) level. Hence, even if we account for measurement error, the local chain-level prices are not exactly identical, but similar across stores.[21]

[18]The first two principal components explain 33% of the price variance for the median product. See Appendix E for detailed explanations and more empirical examples.

[19]The color labels are not mutually exclusive across the panels. For example, red dots in two different panels represent the projected prices for stores that belong to two different retail chains.

[20]We use DMAs instead of 3-digit ZIP codes as markets because the large number of 3-digit ZIP codes is hard to visualize.

[21]Table 5 shows analogous patterns based on the ratio of the 95th to 5th percentile of prices and base prices.

Table 5 UPC price and base price dispersion statistics: Market vs. retail chain-level

		Median	Mean	Percentiles							
				1%	5%	10%	25%	75%	90%	95%	99%
Prices											
Product definition: UPC											
Log-price SD	DMA	0.110	0.114	0.006	0.032	0.050	0.078	0.147	0.185	0.208	0.249
	ZIP3	0.099	0.103	0.002	0.021	0.039	0.066	0.136	0.174	0.196	0.236
	Chain	0.079	0.083	0.008	0.032	0.043	0.061	0.101	0.125	0.145	0.220
	Chain/DMA	0.050	0.054	0.004	0.018	0.025	0.036	0.067	0.085	0.103	0.163
	Chain/ZIP3	0.039	0.043	0.002	0.012	0.018	0.027	0.054	0.070	0.087	0.151
95/5 percentile ratio	DMA	1.360	1.398	1.011	1.075	1.129	1.232	1.524	1.721	1.831	2.071
	ZIP3	1.294	1.329	1.004	1.045	1.094	1.182	1.438	1.614	1.713	1.946
	Chain	1.254	1.281	1.006	1.070	1.112	1.183	1.346	1.471	1.570	1.960
	Chain/DMA	1.131	1.148	1.004	1.035	1.054	1.086	1.186	1.253	1.317	1.560
	Chain/ZIP3	1.072	1.085	1.001	1.011	1.025	1.044	1.109	1.153	1.196	1.351

Table 5 (continued)

		Median	Mean	Percentiles							
				1%	5%	10%	25%	75%	90%	95%	99%
Base Prices											
Product definition: UPC											
Log-price SD	DMA	0.088	0.093	0.002	0.023	0.039	0.062	0.118	0.151	0.177	0.239
	ZIP3	0.078	0.083	0.001	0.014	0.029	0.052	0.107	0.140	0.166	0.225
	Chain	0.064	0.067	0.002	0.020	0.032	0.048	0.081	0.104	0.125	0.188
	Chain/DMA	0.037	0.040	0.000	0.010	0.017	0.026	0.050	0.065	0.077	0.126
	Chain/ZIP3	0.027	0.031	0.000	0.006	0.011	0.019	0.038	0.052	0.063	0.111
95/5 percentile ratio	DMA	1.276	1.310	1.003	1.050	1.096	1.177	1.399	1.556	1.680	1.983
	ZIP3	1.220	1.250	1.001	1.028	1.066	1.134	1.328	1.464	1.569	1.832
	Chain	1.197	1.219	1.000	1.037	1.080	1.137	1.267	1.368	1.484	1.776
	Chain/DMA	1.091	1.103	1.000	1.017	1.035	1.060	1.129	1.177	1.213	1.404
	Chain/ZIP3	1.045	1.054	1.000	1.005	1.014	1.028	1.068	1.098	1.123	1.245

8.2 Promotion coordination

Next, we focus on promotions. In particular, we examine if price promotions are coordinated, in the sense that the same product is systematically promoted at the same time among the stores in a retail chain or in a market.

Appendix F provides an overview of promotion coordination, and analyzes the distribution of chain/market level promotion percentages, i.e. the fraction of local stores in a retail chain that simultaneously promote a product.

Here, we focus on a more formal statistical analysis to estimate the dependence in the incidence of promotions across stores. The promotion incidence is captured using the indicator $D_{jst} \in \{0, 1\}$, where $D_{jst} = 1$ if product j is promoted in store s in week t. For store s in market m we define the *inside promotion percentage*, the mean promotion incidence in week t among all other stores in the local retain chain except s:

$$I_{jst} = \frac{1}{|\mathcal{S}_{sm}|} \sum_{r \in \mathcal{S}_{sm},} D_{jrt}.$$

$\mathcal{S}_{sm}$ is the set of all stores in market m that belong to the same retail chain as store s, not including store s itself.[22] Vice versa, the *outside promotion percentage* is the mean promotion incidence in week t among all stores in other local retain chains:

$$O_{jst} = \frac{1}{|\bar{\mathcal{S}}_{sm}|} \sum_{r \in \bar{\mathcal{S}}_{sm}} D_{jrt}.$$

$\bar{\mathcal{S}}_{sm}$ is the set of all stores in market m that belong to a retail chain other than the chain that store s belongs to. We define markets as DMAs to ensure that the local retail chains have sufficiently many stores. In the extreme case when a chain had only one local store, the inside promotion percentage would not be defined.

To test for promotion coordination, we estimate the statistical association between the promotion indicator D_{jst} and the inside and outside promotion percentages, $x_{jst} = (I_{jst}, O_{jst})$, separately for each product j and store s:

$$\mathbb{E}[D_{jst}|x_{jst}] = \Pr\{D_{jst}|x_{jst}\} = \alpha_{js} + \beta_{js} I_{jst} + \gamma_{js} O_{jst}. \tag{3}$$

If the promotions in store s are set independently of the same-chain and other-chain promotions in market m, then $\beta_{js} = \gamma_{js} = 0$.

The estimates are summarized in Table 6 (Fig. 23 displays corresponding histograms).[23] We provide the results separately for the full model (3) and for a restricted version that only includes O_{jst} as independent variable. We first focus on the DMA-level results for the full model in the top panel. The median across all inside percentage coefficients, β_{js}, is 1.005, and 97.7% of the estimates are positive. Furthermore, we reject the null hypothesis that the inside percentage coefficient is not positive, $\beta_{js} \leq 0$, for 95.8% of all estimates, and 52.8% of the estimates are not statistically different from 1 at the 5% level. These results provide clear evidence that

[22] $|\mathcal{S}_{sm}|$ is the number of stores in $\mathcal{S}_{sm}$.

[23] The distributions are weighted using total product revenue weights.

Table 6 Promotion coordination: Promotion dependence regressions

		Median	Mean	NR = 0	NR = 1	R ≥ 0	Est. > 0	Percentiles							
								0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
DMA															
(a)	Inside Percentage	1.005	0.975		0.528	0.958	0.977	−0.147	0.297	0.553	0.861	1.086	1.233	1.401	2.247
	Outside Percentage	−0.001	0.008	0.897		0.036	0.490	−0.768	−0.272	−0.168	−0.064	0.064	0.183	0.309	0.916
	Intercept	0.000	0.008	0.872			0.497	−0.125	−0.056	−0.037	−0.015	0.021	0.060	0.097	0.235
(b)	Outside Percentage	0.082	0.173	0.777		0.252	0.631	−1.036	−0.363	−0.222	−0.067	0.308	0.646	0.957	2.106
	Intercept	0.147	0.186	0.289			0.956	−0.044	0.003	0.017	0.060	0.272	0.419	0.513	0.686
National															
(a)	Inside Percentage	1.017	1.036		0.503	0.933	0.972	−0.294	0.183	0.458	0.825	1.140	1.369	1.629	3.296
	Outside Percentage	0.000	0.027	0.876		0.044	0.498	−1.374	−0.560	−0.360	−0.144	0.164	0.429	0.685	1.666
	Intercept	−0.002	0.006	0.873			0.476	−0.267	−0.115	−0.074	−0.029	0.030	0.091	0.154	0.383
(b)	Outside Percentage	0.363	0.510	0.649		0.394	0.754	−1.299	−0.499	−0.268	0.006	0.875	1.495	1.961	3.553
	Intercept	0.077	0.114	0.635			0.794	−0.255	−0.101	−0.047	0.010	0.190	0.338	0.449	0.693

Note: NR stands for "Not Rejected", R stands for "Rejected", and "Est." stands for "Estimate" the price promotions for most products are coordinated across stores within a retail chain. On the other hand, the median of the outside percentage coefficients, γ_{js}, is -0.001, and 89.7% of the estimates are not statistically different from 0. Hence, conditional on the inside promotion percentage, I_{jst}, information on the contemporaneous promotion incidence in other retail chains in the local market is typically not predictive of D_{jst}. The estimates of the intercept are small and mostly not distinguishable from 0, indicating that the promotion probability in a store is 0 if none of the other stores in the chain promote the product. This is further evidence that promotions are coordinated at the local chain level.

To investigate if promotions are also unconditionally independent of promotions in other retail chains, we estimate a restricted model that only includes the outside percentage as independent variable. The median of the coefficient estimates is 0.082, and the distribution of the estimates is skewed to the right. Hence, there is evidence that for some products promotions are unconditionally dependent on the promotion incidence in other local retail chains. This dependence is likely due to promotional allowances—trade deals that are offered by the product manufacturers to multiple or all retail chains. Another explanation is seasonality in demand, although seasonality is unlikely to account for the large documented degree in promotion coordination.

We also test if promotions are coordinated across markets at the national level. We define the *national inside promotion percentage*, the mean promotion incidence in week t among all other stores of the chain in different markets:

$$I'_{jst} = \frac{1}{|\mathcal{S}_{s,-m}|} \sum_{r \in \mathcal{S}_{s,-m}} D_{jrt}.$$

$\mathcal{S}_{s,-m}$ includes all stores in the same chain at the national level, excluding the market that store s belongs to. We similarly define the national outside promotion percentage, O'_{jst}, the mean promotion incidence in other chains outside the market. The estimates in Table 6 show that promotions are also strongly coordinated at the national level. The median of the national inside promotion percentage estimates is 1.017, and 97.2% of the estimates are positive. Further, mirroring the market-level results, promotions in store s are conditionally independent of the promotions in other retails chains outside the market. However, in the restricted regression, the median of the national outside percentage coefficients is 0.363. Thus, the national-level estimates provide stronger evidence for unconditional promotion dependence across retail chains than the market-level results.

8.3 Is price discrimination constrained by feature advertising?

Retailers use feature advertising to inform households of specific products and their prices. Feature ads are distributed in the form of circulars (print or digital) or newspaper inserts. Because feature ads typically apply to all stores in a market, they constrain the degree of price discrimination that is feasible for a retail chain. Hence, we analyze if the similarity of prices at the chain-level is due to feature advertising.

Our analysis relies on the fact that features are typically used to advertise promoted prices but not base prices. We are able to measure the association between feature

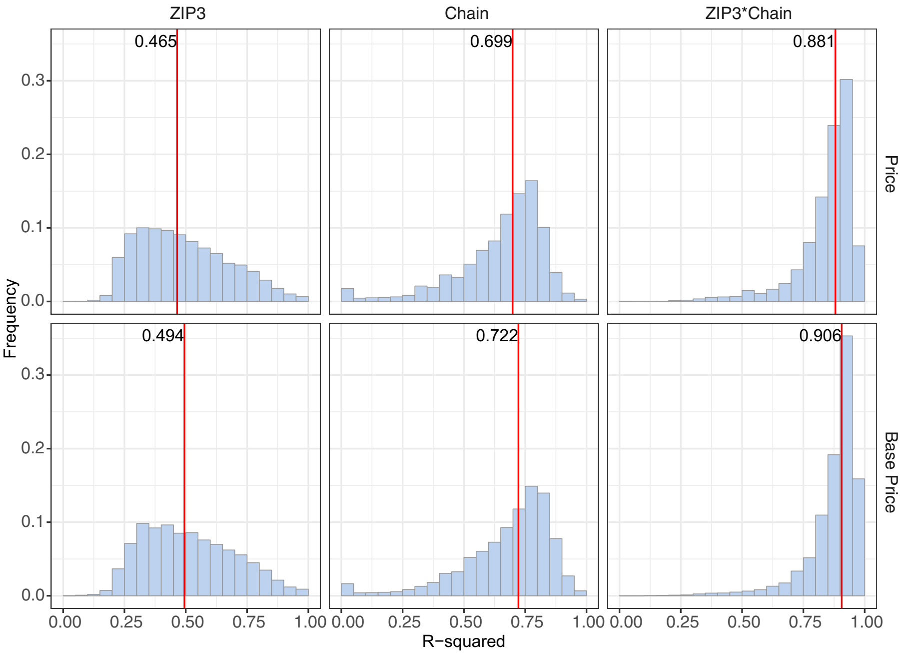

Fig. 9 Percentage of variance of prices and base prices explained by market and chain factors

ads and promotions because data on feature advertising are available in the Nielsen RMS scanner data for a sub-sample of 17% of stores.[24] We find that in weeks when a product is not promoted, the probability that a product is featured is only 0.03. Hence, while feature advertising may constrain price discrimination in promoted prices, it cannot constrain base prices.

We repeat the analysis in Section 7, where we measure the percentage of the price variance that can be attributed to market or retail chain-specific factors, for base prices only. If price similarity were due to feature advertising, chain and market/chain factors should explain a smaller percentage of the base price variance compared to the overall price variance. Figure 9 shows the revenue-weighted distributions of the R^2 values from regressions of base prices on market (3-digit ZIP code), chain, and market/chain indicators. For comparison, we also show the corresponding distributions discussed in Section 7, where we do not differentiate between promoted and base prices. For the median product, the R^2 values are somewhat larger for the base price regressions. For example, the median R^2 for the base price regressions using market/chain indicators is 90.6%, compared to 88.1% for the analogous price regressions.

[24]All but four retail chains have stores that are in this sub-sample. Among the covered retailers, feature ads are recorded for about 20% of stores, and in these stores feature advertising is measured consistently for most products and weeks. Among the covered stores, feature advertising is measured for 99% of all non-imputed product/week observations and for almost 90% of all products.

The results indicate that the degree of price similarity at the chain and market/chain level is somewhat higher for base prices, which are typically not featured. We conclude that the observed price similarity is not primarily due to feature advertising as an institutional constraint on price discrimination.

We also compared the degree of price dispersion across weeks with and without feature advertising. This analysis is inconclusive, however, because there is more measurement error for promoted prices than for base prices (Section 5.2), which introduces more artificial price dispersion in featured than in non-featured weeks.

8.4 Discussion

The similarity in prices and promotion strategies in our data differs from the heterogeneity in pricing strategies within the same retail chain that is documented in Ellickson and Misra (2008). The analysis in Ellickson and Misra (2008) uses data from the 1998 *Trade Dimensions Supermarkets Plus Database*, which provides information on store-level pricing strategies based on surveys of retail chain managers. Hence, the data are not directly comparable and cover a different time period than our work. In particular, the managers surveyed in the *Trade Dimensions* data classify store-level pricing policies as EDLP (everyday low price), promotional/Hi-Lo, or as a hybrid of EDLP and Hi-Lo. These qualitative responses may be consistent with the residual variation in pricing and promotion policies after accounting for market/chain dummies as shown in Fig. 6.

Also related to our work, Arcidiacono et al. (2020) find that the price similarity pattern remains unchanged even after the entry of a strong competitor, a Walmart Supercenter, in the local market.[25]

9 Does demand similarity explain price similarity?

As shown in the previous two sections, prices are more similar at the chain level than at the market level, and prices are especially similar within chains at the local market level. Furthermore, promotions are highly coordinated across stores that belong to the same retail chain.

Without context, the economic implications of these findings are hard to assess. In particular, if there is a loss in profits because price discrimination is not employed depends on whether price elasticities and promotion effects differ significantly across the stores of a retailer. To provide this necessary context, we estimate store-level demand for the top 2,000 brands in our data, and we compare the similarity in prices and promotions to the corresponding similarity in price elasticities and promotion effects.

[25] In particular, the entry of a Walmart Supercenter leads to a 16% drop in the revenue of the nearby retailers, but to no corresponding change in the prices offered by the incumbents.

9.1 Demand model

We estimate demand for each product at the store level. We define products as brands. Estimating demand at the UPC level is difficult because of the large number of UPCs in most categories and because stores carry different assortments of UPCs under the same brand name. Hence, demand estimation in industrial organization and marketing is typically performed at the brand level (e.g. Hoch et la., 1995 and Nevo, 2001).

We use a log-linear demand model for brand j in store s and week t:

$$\log(1 + q_{jst}) = \alpha_{js} + \sum_{k \in \mathcal{J}_{js}} \beta_{jks} \log(p_{kst}) + \sum_{k \in \mathcal{J}_{js}} \gamma_{jks} D_{kst} + \tau_j(s,t) + \epsilon_{jst}. \quad (4)$$

We add 1 to the sales quantity, q_{jst}, to ensure that the demand model is valid for the substantial number of observations with 0 sales in the data. p_{kst} is the price of brand $k \in \mathcal{J}_{js}$, where $\mathcal{J}_{js}$ is a set of products in store s that are in the same category that brand j belongs to, including j itself. D_{kst} is a promotion indicator. The demand model includes brand-store fixed effects, α_{js}, and time fixed effects, $\tau_j(s,t)$. $\tau_j(s,t)$ is identical for all stores in a local market, defined based on the 3-digit ZIP code. Our main estimates are obtained with $\tau_j(s,t)$ defined as month fixed effects.[26]

Using the demand model (4), the predicted price elasticity for brand j with respect to the price of product k is given by

$$e_{jk} = \frac{\partial q_{jst}}{\partial p_{kst}} \frac{p_{kst}}{q_{jst}} = \beta_{jks} \frac{q_{jst} + 1}{q_{jst}}.$$

The price coefficient β_{jks} approximates the price elasticity e_{jk}, and for simplicity we will refer to β_{jks} as a price elasticity from now on.

9.2 Sample selection

We estimate demand for the top 2,000 brands (based on total revenue) during the 2008-2010 period. We focus on these large brands to avoid measurement error in prices. This measurement error occurs because the price of a small brand frequently needs to be imputed for weeks with no sales, $q_{jst} = 0$. Additionally, to avoid measurement error, for each brand we only estimate demand for stores where prices are observed in at least 80% of weeks. In many product categories it is not feasible to include the prices and promotions of all competing brands in the demand model. Hence, we only include the brands that account for at least 80% of the category revenue, with a maximum of 5. In total, we estimate 27.2 million brand-store demand models.

[26] If s and s' are two stores in the same 3-digit ZIP code, and if t and t' are two weeks in the same year and month, then $\tau_j(s,t) = \tau_j(s',t')$.

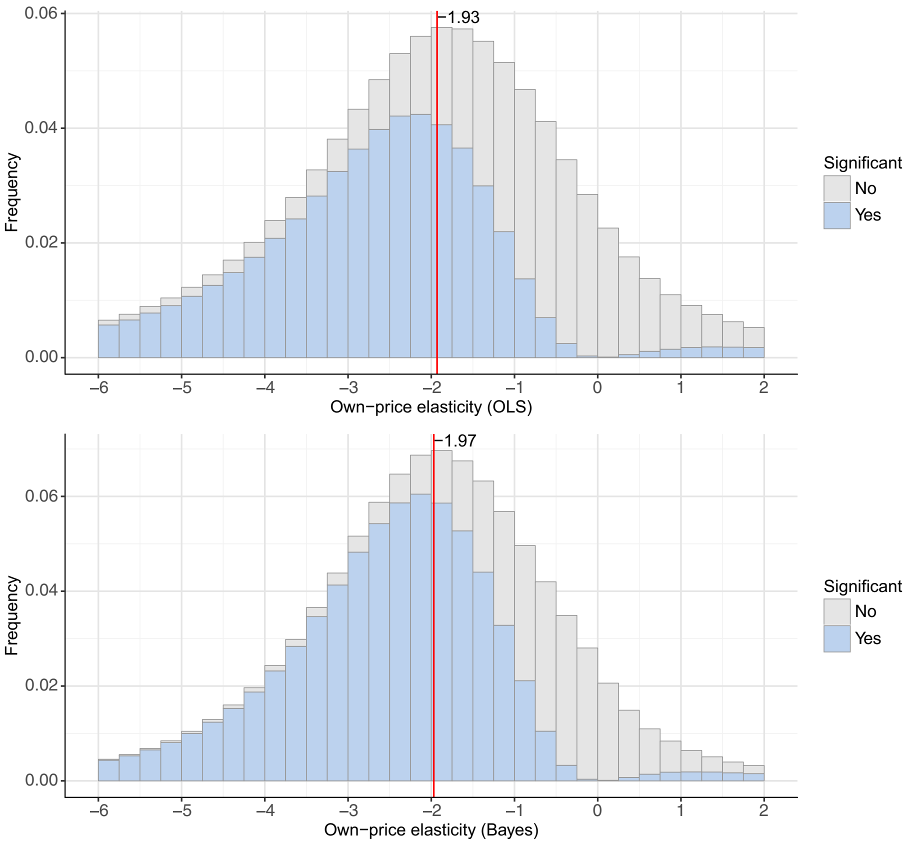

Fig. 10 Own-price elasticity estimates

9.3 Estimation results

9.3.1 Main results

The distribution of the own-price elasticities, pooled across all brand-store estimates, is shown in the top panel of Fig. 10 (see Table 12 in the Appendix for detailed summary statistics).[27] The estimates are obtained using local (3-digit ZIP code) month fixed effects. We color estimates that are not statistically different from 0 at the 5% level in gray and all other estimates in blue.

The median of the brand-store price elasticity estimates is -1.93. There is a large degree of heterogeneity, with the estimates ranging from -6.647 at the 5th percentile to 2.025 at the 95th percentile of the distribution. 85.6% of the own-price elasticity estimates are negative, but only a small percentage, 4.1%, of the estimates is positive and statistically different from 0. That many parameters are not precisely estimated

[27] The distribution is weighted using total brand revenue. The weights are brand, not brand/store-specific.

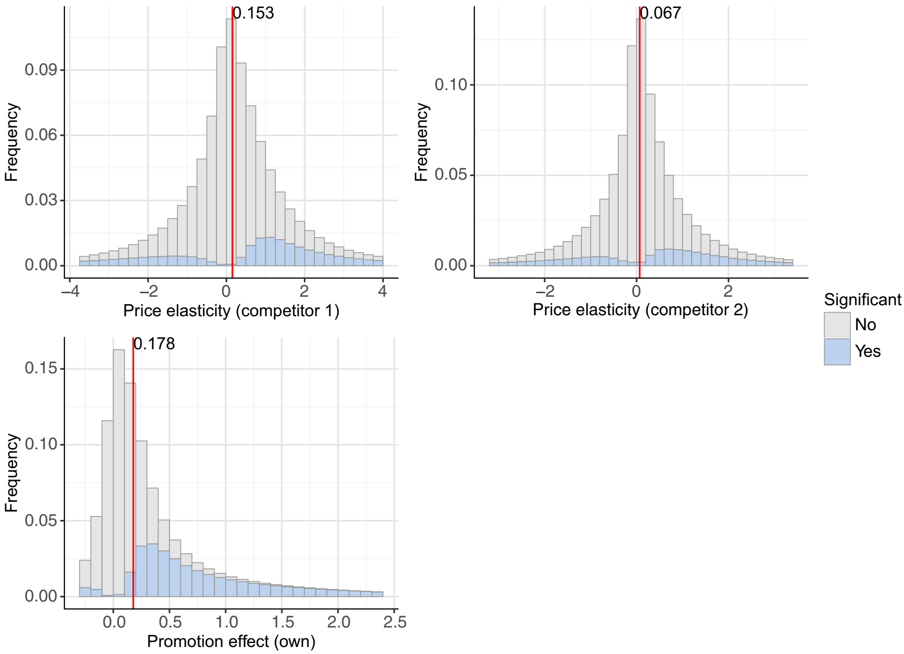

Fig. 11 Cross-price and promotion effect estimates

or have the "wrong" sign is expected because the estimates are brand/store-specific and obtained using observations for at most 156 weeks during the 2008-2010 period. 70.6% of the estimates indicate elastic demand, $\beta_{jjs} < -1$.

The top panels in Fig. 11 display the distributions of the estimated cross-price elasticities with respect to the two largest competitors in the product category. The medians of the cross-price elasticities are positive, but a large percentage of the estimates, 42.8% for the largest and 44.7% for the second largest competitor, is negative. The majority of the estimates is not statistically different from zero. This evidence indicates that it is particularly challenging to obtain precise cross-price elasticity estimates at the brand/store level using a 156 time series of weekly observations. To improve on these estimates we would have to impose parameter restrictions or estimate a demand model that relies on a smaller number of parameters, such as a logit or random coefficients logit demand system (Berry et al., 1994, 1995).

The own-promotion effect estimates are shown in the bottom panel of Fig. 11. Among these estimates a larger percentage have the expected sign: 77.5% of all estimates and 94.8% of the estimates that are statistically different from zero are positive.

9.3.2 Bayesian hierarchical model results

As an alternative to the OLS estimates we use a Bayesian hierarchical model to obtain the posterior distribution of the demand parameters for each brand and store, $\theta_{js} =$

$(\alpha_{js}, \beta_{js}, \gamma_{js})$.[28] In the Bayesian hierarchical model specification, θ_{js} is assumed to be distributed according to the population distribution or first-stage prior $p(\theta)$, which we assume to be normal, $N(\bar{\theta}_j, V_j)$. The posterior distribution of the store-level and population parameters is obtained using MCMC sampling. We obtain store-level estimates of the demand parameters based on the posterior means of θ_{js}. A detailed summary of the model specification and sampling approach is provided in Appendix G.

There are two reasons that motivate us to provide these additional estimates. First, the posterior means in the Bayesian hierarchical model are shrinkage estimators by which imprecise store-level estimates are shrunk to the population mean. This shrinkage property provides a form of regularization to guard against noisy parameter estimates, which is particularly important given the goal to obtain a large number of brand/store-level demand estimates. Second, Bayesian hierarchical models are widely used in the industry by in-house analysts and analytics companies that provide demand estimates for retail chains and brand manufacturers. Hence, the pricing and promotion decisions in our data are often made using estimates from a Bayesian hierarchical model.

The bottom panel in Fig. 10 displays the distribution of the own-price elasticity estimates, i.e. posterior means, from the Bayesian hierarchical model. The median is almost identical to the median of the OLS estimates. However, the distribution of the Bayesian hierarchical model estimates has thinner tails than the distribution of the OLS estimates, which is expected due to the shrinkage property of the Bayesian hierarchical model (see Table 12 for detailed results). Also, the percentage of negative elasticities is larger for the Bayesian hierarchical model compared to the OLS estimates, 90.3% versus 85.6%, and 74.9% versus 70.6% of the OLS estimates indicate elastic demand. In this sense, the Bayesian hierarchical model estimates conform more to expectations, although the overall difference with respect to the OLS estimates is only moderate.

9.3.3 Causal price and promotion effects?

In the presence of endogeneity or confounding the price and promotion coefficients do not have a causal interpretation. Our strategy to adjust for confounding relies on the store and market/time fixed effects that are included in the demand model. If $\tau_j(s, t)$ captures all time-varying demand components that are associated with the prices, and if the residual variation in prices, conditional on the fixed effects, is due to factors such as costs that do not directly affect demand, then we can interpret the estimated price and promotion coefficients as causal. This strategy is discussed further in Appendix H, and we show that the coefficient estimates are largely insensitive to the choice of less granular (quarterly) or more granular (weekly) time fixed effects. The robustness to the exact choice of the fixed effects suggests that confounding is of little concern, but we cannot conclusively rule out some remaining endogeneity.

[28] β_{js} is a vector that includes the own and cross-price elasticities, β_{jks}, and the promotion parameters, γ_{jks}.

We could have alternatively pursued an instrumental variables strategy (e.g. Berry et al., 1994, 1995). However, for the case of retail pricing, good instruments are hard to find and often weak, as emphasized by Rossi (2014). Ultimately, if we *should* use instruments or not is besides the point, for an important reason that we explain next.

In particular, instrumental variables strategies have not been employed widely in the industry for the purpose of price and promotion effect estimation. Hence, the estimates that we obtain are similar to the estimates available to sophisticated retailers and brand manufacturers. In our empirical analysis below we investigate if the price elasticities and promotion effects that are available to managers are similar or discernibly different across stores at the local retail chain level. For this purpose, what matters is an analysis of the estimates that managers use to set prices and promotions, whereas any potential bias due to confounding is of little relevance.

9.4 Similarity of price and promotion effects at the market and retail chain level

In Section 7 we measured the percentage of the variance in prices, promotion frequency, and promotion depth at the market and retail chain level. Analogous to this analysis, we now document what percentage of the variance of the own-price and promotion effects can be attributed to market (3-digit ZIP code), retail chain, and market/retail chain factors.

Figure 12 displays the distribution of the revenue-weighted R^2 values from regressions of $\hat{\beta}_{jjs}$ and $\hat{\gamma}_{jjs}$ on market, chain, and market/chain indicators (for detailed results see Table 13 in the Appendix).[29] The top row shows the results for the OLS price elasticity estimates. For the median product, market factors explain 14.6% of the price elasticity variance across stores. Retail chain factors explain 17.2%, and in particular, market/chain factors explain almost half, 47.3% of the variance across stores. The results indicate that market factors account for a modest percentage of the differences in own-price elasticities across stores. However, within local markets, the elasticities are much more similar for stores that belong to the same retail chain.

Although similar, the price elasticities are not identical at the local retail chain level—slightly more than half of the variation is not captured by market/chain fixed effects. This remaining variation in the elasticities might represent an unexploited opportunity to price discriminate across stores, or it may simply reflect measurement error in the *estimated* own-price elasticities. Hence, we compare the R^2 values for the OLS own-price elasticity estimates to the corresponding R^2 values for the posterior means of the elasticities from the Bayesian hierarchical demand model, which reduces measurement error by design. The results in the second row of Fig. 12 are consistent with less measurement error in the estimates from the Bayesian hierarchical demand model, as the R^2 values are uniformly higher. Based on these alternative estimates, local market factors explain 20.7% of the overall variance, whereas chain factors explain 23.5% and market/chain factors explain more than half, 52.3% of the variation in own-price elasticities.

[29] The own-price elasticity and promotion effect estimates are from the main model specification that includes 3-digit ZIP code/month fixed effects.

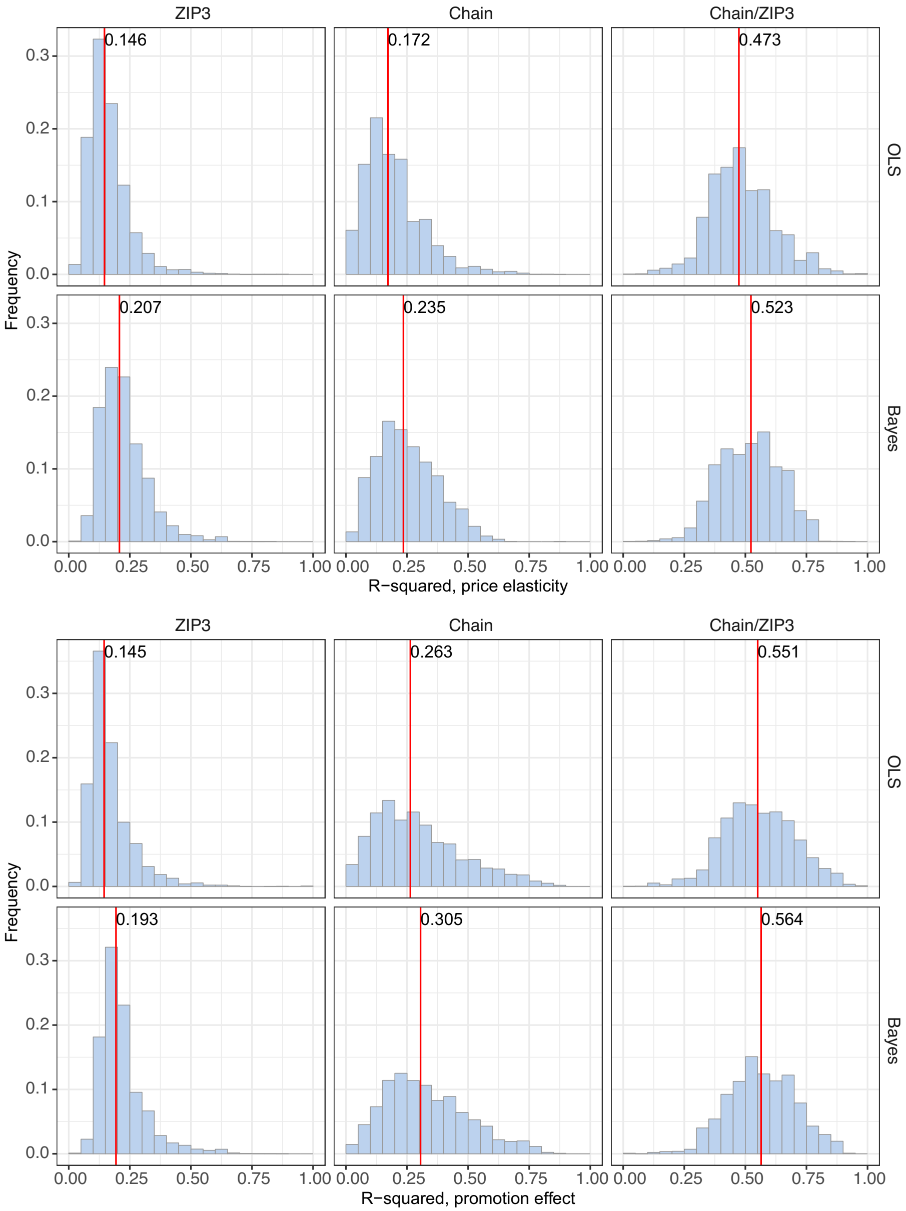

Fig. 12 Percentage of variance of own-price elasticities and own-promotion effects explained by market and chain factors

We perform an analogous analysis for the estimated own-promotion effects. The results are shown in the bottom two rows of Fig. 12. Based on the estimates using the Bayesian hierarchical model, market factors explain 19.3%, chain factors explain

30.5%, and market/chain factors explain 56.4% of the variance in the estimated promotion effects. Hence, compared to the results for the own-price elasticities, chain factors explain a larger percentage of the variation in the estimated promotion effects.

10 Can retailers distinguish among store-level price elasticities and promotion effects?

The documented chain-level similarity in own-price and promotion effects raises the question if retail managers are able to distinguish among the corresponding estimates that are available to them. Hence, unlike in the previous section, we now focus on the estimates used by retail managers who employ marketing analytics to make price and promotion decisions. In the industry, Bayesian hierarchical models of demand that are similar to our specification have been used since the early 2000s.[30] To emulate the empirical analysis available to a sophisticated retailer, we now estimate Bayesian hierarchical demand models *separately* for each retailer. Using only data for their own stores is a practical necessity for retailers, because scanner data for their competitors are typically not available to them. Different from the prior analysis, the population distribution (first-stage prior) of the demand parameters is now specified at the chain level, not at the national level. Hence, the store-level estimates will be shrunk to the chain-level mean, and the price and promotion effects that are visible to the managers will likely appear more similar compared to the previous estimates that were obtained using a national first-stage prior.

10.1 Similarity of price and promotion effects visible to managers

We first replicate the analysis in Section 9.4 using the store-level own-price elasticities and promotion effect estimates that are available to retailers, and document the percentage of the variance of the estimates that can be explained by market, retail chain, and market/chain factors. The results are shown in Fig. 13, with the estimates using a national first-stage prior displayed in the first row and the results using chain-level first-stage priors in the second row of each panel.

The variation in price elasticities and promotion effects explained by market factors is almost identical across the two specifications. However, chain and market/chain factors explain a substantially larger fraction of the variation in the estimates that are obtained using chain-level first-stage priors. Whereas 23.5% and 52.3% of the variation in price elasticities is explained by chain and market/chain factors when using national first-stage priors, 51.0% and 70.5% of the variation is explained by these factors when using using chain-level first-stage priors. Similarly, focusing on the promotion effects, the variation explained across the specifications is 30.5% versus 56.0% for chain factors and 56.4% versus 73.5% for market/chain factors.

[30]For example, DemandTec, which was founded in 1999 and later acquired by IBM, offered analytic services to its retail clients using such demand models.

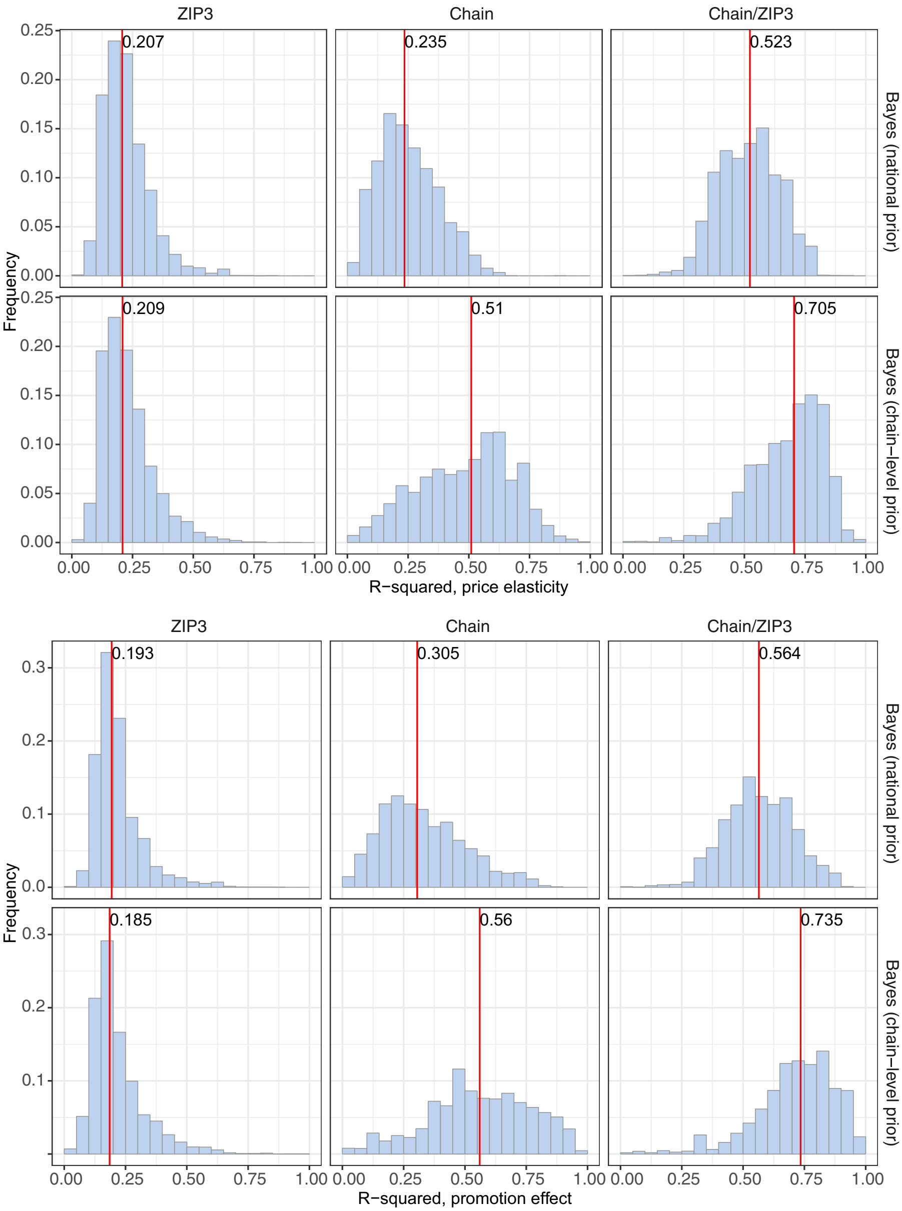

Fig. 13 Percentage of variance of own-price elasticities and own-promotion effects explained by market and chain factors: Chain level vs. national priors

10.2 Statistical distinguishability of the store-level estimates

The results in the previous sub-section indicate a large degree of similarity in the price and promotion effect estimates that are visible to sophisticated retailers who

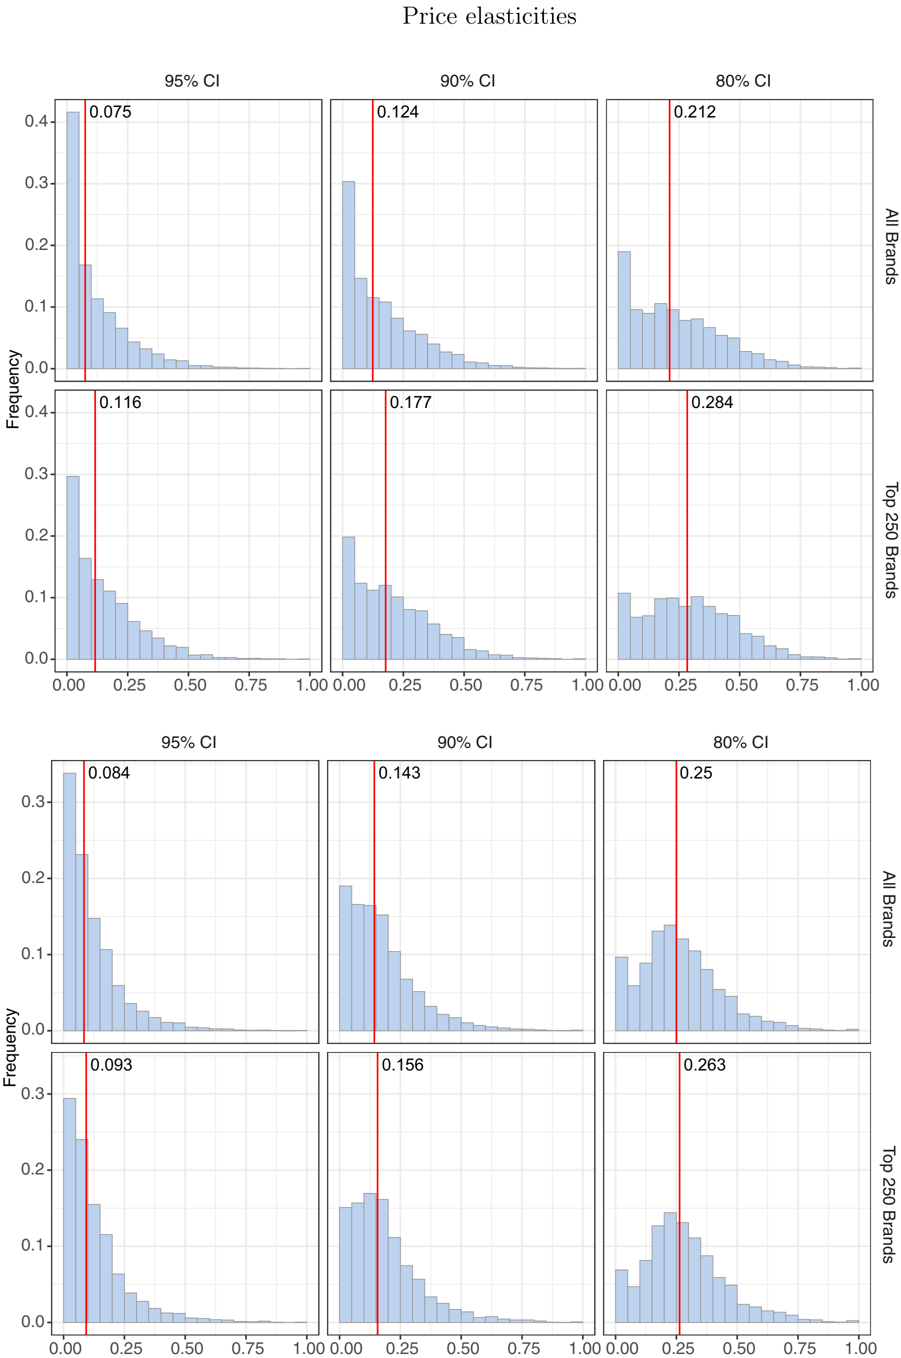

Fig. 14 Percentage of estimated price elasticities and promotion effects different from market/retail chain level mean

make data-driven pricing and promotion decisions. The similarity is especially stark considering that the store-level price and promotion effects are obtained using at most 156 weekly observations and hence likely affected by substantial sampling error. Hence, although the posterior means of the price and promotion effects are different, managers may consider price discrimination across stores to be impractical or infeasible due to the statistical uncertainty around the estimates.

To assess the statistical distinguishability of the estimates we again emulate the demand analysis that sophisticated retailers use in practice, and we construct credible intervals for the store-level price elasticity and promotion effect estimates. For each brand/market/retail chain combination, we then record the percentage of store-level price and promotion effects with credible intervals that exclude the mean of the local chain-level estimates. This approach to assess the difference in the estimated effects is not exactly identical to a frequentist hypothesis test, but plausibly corresponds to how differences between price and promotion affects are assessed in the industry practice.

Figure 14 displays histogram of the percentages of price and promotion effects that can be distinguished from the local chain-level mean. The observations in the distributions are at the brand/market/retail chain level. The results are shown separately using 95%, 90%, and 80% percent credible intervals. In the top row of each panel we show the distributions for all brands, whereas the histograms in the bottom row show the distributions for the 250 largest brands.[31]

Focusing on the results for all brands, the median percentages of statistically distinguishable price elasticities are 7.5%, 12.4%, and 21.2% when using 95%, 90%, and 80% credible intervals. For the 250 largest brands, the respective median percentages are 11.6%, 17.7%, and 28.4%. The somewhat larger fraction of statistically distinguishable price effects reflects the higher precision of the estimates for the largest products. However, even for the largest brands, retail managers who conduct a similar analysis will find that between 71.6% and 88.4% of all price elasticity estimates are not distinguishable from the chain/market-level average. A similar pattern holds for the estimated promotion effects, shown in the bottom panel of Fig. 14.

Note that the Bayesian hierarchical model estimates were obtained after residualizing the data to account for the common time/market fixed effects (Appendix G). Our credible intervals do not account for the uncertainty in the fixed effect estimates, and hence they understate the uncertainty about the price and promotion effects. Therefore, our numbers should be interpreted as upper bounds on the true statistical distinguishability of the store-level estimates.

11 Discussion: Explanations for price similarity

The analysis in the previous section provides one explanation for the observed lack of local price discrimination by retailers. In particular, the lack in precision and the corresponding difficulty to distinguish between store-level price and promotion effects

[31] Brand size is measured by total brand revenue.

may lead managers to conclude that price discrimination across stores is infeasible in practice. This explanation also conforms with the anecdotal evidence that we gathered in informal discussions with retail chain managers.

A related explanation is that cross-price elasticities are hard to estimate. Both valid own and cross-price elasticity estimates are required to predict and optimize category profits. However, as discussed in Section 9.3.1, the brand/store cross-price elasticity estimates in our data are imprecise, which may be expected given the limited sample size of at most 156 weekly observations for each brand/store. Our data, sample size, and estimation approach are similar to what is used in the industry.[32] Hence, it is unlikely that managers have access to better estimates. Further, pricing is part of what is called *category management* in the industry[33], and managers are aware that price changes for one product may also affect demand for other products in a category (or even lead to substitution across categories or stores). Hence, given the two related challenges to distinguish among price elasticities across stores and to obtain precise cross-price elasticities, retail managers may plausibly consider price discrimination across stores as infeasible.

We do not attempt to predict optimal, profit-maximizing prices in this paper, largely because of the challenges discussed above when estimating store-level price elasticities. Closely related to our work, DellaVigna and Gentzkow (2019) also document price similarity at the retail chain level. They estimate log-linear demand models that are similar to our specification using UPCs as product definition. Using the elasticity estimates they predict optimal product-level prices separately for each store within a chain. The corresponding predicted annual profit loss at the observed prices compared to optimal pricing for all products is $239,000 (1.79% of revenue) for the median store in their sample. To predict the optimal prices, DellaVigna and Gentzkow (2019) maximize profits for each UPC separately only with respect to its own price. In particular, competitor prices are not included in the estimated demand models. Thus, by abstracting away from substitution to other UPCs sold under the same brand name or to other brands in the category, they avoid the problems that we discussed above. DellaVigna and Gentzkow (2019) propose that managerial inertia, including agency frictions and behavioral factors, are the reason for price similarity and the deviation of actual prices from their prediction of optimal prices. This hypothesis is different from our proposed explanation for price similarity, which neither requires agency issues nor behavioral factors, although the two alternative explanations are not mutually exclusive.

We ruled out that price similarity is due to feature advertising (Section 8.3), in particular because not just promoted prices but also base prices exhibit a high degree of price similarity. However, consumers may prefer retailers with predictable, consistent

[32] We could potentially have obtained more precise demand estimates using data that covered a larger number of years than the three years used in our analysis. In practice, however, demand analyses performed for manufacturers and retailers have typically been based on at most two years of data. Hence, the demand estimates that we analyze are likely to overstate, not understate the precision of the estimates available in the industry practice.

[33] See, for example, the discussion of pricing and promotion tactics in *Consumer-Centric Category Management* by ACNielsen (2005).

pricing patterns over chains that sell products at different price points across stores, and hence the increase in profits from store-level price discrimination may be offset by the loss of consumers who substitute to competing retail chains. This aversion of consumers to price discrimination may be due to fairness concerns, as proposed by Kahneman et al. (1986). We are unable to test this explanation with our data. However, fairness concerns are consistent with the findings in Ater and Rigbi (2020), who provide evidence that a food price transparency regulation in Israel caused an increase in price similarity in grocery chains. As corroborating evidence, Ater and Rigbi (2020) argue that fairness concerns were an important part of the public debate that resulted in the legislation that included the price transparency regulation.

Adams and Williams (2019) and Dobson and Waterson (2005) argue that uniform pricing may soften competition and increase retailer profits. To achieve uniform pricing as the outcome in the Dobson and Waterson (2005) model, one competitor needs to be able to pre-commit to uniform pricing. In an empirical analysis of the home-improvement industry, Adams and Williams (2019) find that one retailer in the home-improvement industry would achieve higher profits if both the retailer and its main competitor used uniform pricing. Even with pre-commitment, however, uniform pricing is not an equilibrium outcome in the empirical example.

12 Conclusions

We analyze patterns in retail (grocery) price and promotion strategies and in the relationship between pricing and promotion strategies and product demand. We emphasize generalizable results that are based on a large sample of products, representing almost 80% of retail revenue and sold across a large number of stores.

We document a large degree of price dispersion at a given moment in time for identical products (UPCs) and almost identical products (brands) across U.S. retail stores, both nationally and at the local market level. The degree of price dispersion strongly varies across products. Most products are frequently promoted, but the overall price dispersion is not only due to promotions but also due to a large degree of dispersion in the non-promoted base prices. Decomposing the overall yearly price variance into base price and promotion components, we find that within markets, persistent base price differences across stores account for the largest share of the price variance, whereas the contribution of the within-store variance of base prices is relatively small. Despite the high degree of promotional activity, promotions account for only a moderate share of the overall price variance. This counterintuitive finding is due to an EDLP vs. Hi-Lo pricing pattern, which compresses the overall price dispersion across stores.

A key finding is the substantial similarity in prices, promotion frequency, and promotion depth at the retail chain and especially at the local market/retail chain level. Individual promotion events are also strongly coordinated within chains. This similarity in prices and promotions may appear like a missed opportunity to price discriminate, and it contradicts what has been the conventional wisdom in the academic marketing community, that there was a trend towards “store-specific marketing” in the grocery retail industry. However, mirroring the observed patterns for prices and promotions, we find that also store-level estimates of price elasticities and promotion effects are similar at the retail chain level, and even more so at local market/chain level. Further, the price and promotion effects across stores are hard to statistically distinguish from the local average effects, and cross-price elasticity estimates in particular are highly imprecise. We are confident that managers do not have access to more precise information, because our demand model, estimation method, and data closely emulate what sophisticated analysts use in practice. Therefore, managers may consider store-by-store pricing and promotion decisions to be practically infeasible.

Our novel explanation for price similarity differs from related work that has attributed price similarity to managerial inertia (DellaVigna & Gentzkow, 2019), brand image or fairness concerns (Ater & Rigbi, 2020), and competitive considerations (Adams & Williams, 2019). However, although different, the explanations are not mutually exclusive.

We emphasize, however, that our analysis is not *normative*. We argued that retail managers or analysts may consider price discrimination across stores to be infeasible if they attempt to statistically distinguish between price and promotion effects using an approach similar to Section 10.2. Our argument provides an explanation for the observed price and promotion similarity, but does not imply that retailers should not try to price discriminate based on store-level price and promotion effects. In particular, a strategy that sets store-level prices and promotions based on the expectation of profits with respect to the posterior distribution of the store-level demand estimates may increase profits despite the large statistical uncertainty. An empirical investigation into whether such a price discrimination strategy can improve retailer profits, and if it is feasible given the potential brand image and fairness concerns, would be of great practical value.

Acknowledgements We thank Itai Ater, Susan Athey, Pierre Dubois, Wes Hartmann, Kirthi Kalyanam, Carl Mela, Helena Perrone, Stephan Seiler, Steve Tadelis, Raphael Thomadsen, and especially Paul Ellickson for their helpful comments and suggestions. We also benefitted from the comments of seminar participants at the 2017 QME Conference at Goethe University Frankfurt and the 2017 Columbia Business School Marketing Analytics and Big Data Conference. Jacob Dorn, George Gui, Jihong Song, and Ningyin Xu provided outstanding research assistance. Part of this research was funded by the Initiative on Global Markets (IGM) at the University of Chicago Booth School of Business and the Becker Friedman Institute at the University of Chicago. Xiliang Lin worked on the paper and had access to the data when he was a Ph.D. student at the University of Chicago.

Appendix

A Data description: Details

A.1 UPCs and UPC versions

The RMS scanner data record sales units and prices at the week-level, separately for all stores and UPCs. Over time, a UPC can be reassigned to a different product. Therefore, the Kilts Center for Marketing also provides a version code (`upc_ver_uc`) such that the combination of the UPC and UPC version code uniquely identifies a product.[34] A change of the brand name (description) is one of the reasons why a new UPC version is created. Sometimes, a new UPC version reflects a different spelling or abbreviation of the brand name, for example "MOUNTAIN DEW R" versus "MTN DEW R." We attempt to identify and correct all such instances.

In the paper we only refer to UPCs, with the understanding that at the most disaggregated level a product is characterized by a unique combination of a UPC and UPC version code.

A.2 Product size (revenue) distribution

Between 2008 and 2010 the RMS data include information on 967,832 products (UPCs).[35] A large percentage of the total sales revenue is concentrated among a relatively small number of products. To illustrate, in Fig. 17 we rank all products based on total revenue between 2008 and 2010 and plot the cumulative revenue of the top N products on the y-axis. For example, the top 1,000 products account for 20.7 percent, the top 10,000 products account for 56.5, and the top 50,000 products account for 89.3 percent of the total revenue in the 2008-2010 data, respectively.

A.3 Product assortments

We document the distribution of product and brand availability across stores and retail chains in our sample. We classify a product (brand) as available in a specific store or retail chain if it was sold in the store or chain at least once during 2010.

Figure 18 displays the distribution of store availability for brands (column one) and products (column two). The histograms are shown separately for the top 100 (based on total revenue), top 1,000, top 10,000, and—at the bottom of the figure—for all brands and products included in the analysis. The median product in the top 100 group is sold in 12,771 stores, whereas the corresponding median brand is sold in 15,985 stores, representing 93% of all 17,184 stores. Hence, the top products and in particular the top brands are widely available. However, even the top 100 products and brands are not consistently available across all stores, indicating differences in store-level assortment choices. Also, the top brands are more consistently available across stores than products, implying assortment differences whereby stores that carry the same brand offer the brand in different pack sizes or forms (e.g. cans versus bottles). Whereas the top 100 and also top 1,000 products and brands are widely available, the corresponding distributions for the top 10,000 and all products and brands indicate a smaller degree of availability across stores. For example, the median product among all products in the sample is available at 3,854 stores, and the median brand is available at 5,281 stores, i.e. 31% of all stores. Overall, we find that

[34] A new UPC version is created when one or more of the "core" UPC attributes change. The core attributes include the product module (category) code, brand code, pack size (volume), and a multi-pack variable indicating the number of product units bundled together.

[35] If we define a product as a combination of UPC and UPC version (the variable `upc_ver_uc`) the number is 967,863.

assortments across stores tend to be somewhat specialized, with the exception of top-selling products and brands that are typically available at a vast majority of all stores and chains.

The distributions of brand and product availability across retail chains, shown in columns three and four of Fig. 18, are similar to the corresponding distributions across stores. In particular, whereas the top-selling brand and products are widely available (for example, the median brand among the top 100 is sold in 76 out of 81 retail chains), availability is much more limited among the top 10,000 and among all brands and products in the sample. Compared to the store availability distributions, however, the differences across the top and bottom groups are less pronounced. For example, the median brand among all brands in the sample is still available in the majority of retail chains (45 out of 81). In particular, the brand availability distribution for all brands exhibits a pronounced bi-modal shape, indicating a mass of brands available at most retailers and a mass of brands available at a very limited number of retail chains.

A.4 Private label products

The Nielsen data contain both national brand and private label products. However, the brand description of private label products is always "CTL BR" (control brand), and hence we do not know the brand name under which the product is sold. Also, we cannot infer the brand name based on the store where the product was sold because the name of the retail chain that the store belongs to is not revealed. However, we know the product (UPC) description of a product, such as "CTL BR RS BRAN RTE" for a private label Raisin Bran product in the ready-to-eat breakfast cereal category.

In our analysis we treat all private label UPCs as the same product if they share the same product description and contain the same volume. In particular, we treat such UPCs as the same product even if the UPCs are different. The UPCs are typically different because the product is sold by different retail chains, whereas the product itself of often physically identical because it is produced by the same manufacturer that supplies multiple retailers. Even if the product is identical the packaging and specific brand name (e.g. "Kroger Raisin Bran") will differ across retailers. Hence, treating different private label UPCs as the same product is not entirely innocuous, but it is the best we can do to compare the price dispersion of national brands to the price dispersion of private label products across retail chains.

Table 1 shows the percentage of observations accounted for by private label products. Private label products account for 10.8 percent of all price observations and 15.4 percent of total revenue.

B Additional price dispersion results

B.1 Price dispersion: Sensitivity analysis

We calculate two alternative dispersion statistics that are related to the standard deviation of log-prices. First, the distribution of percentage price differences can be measured using the standard deviation of prices normalized relative to the mean price (nationally or at the market level), $p_{jst}/\bar{p}_{jmt}$, which is the approach used in Kaplan and Menzio (2015). Second, we can report the square root of the variance of log-prices calculated using the following approach:

$$\text{var}(\log(p_{jst})|m) = \frac{1}{N_{jmt}} \sum_{s \in \mathcal{S}_{jmt}} (\log(p_{jst}) - \overline{\log(p_{jmt})})^2,$$
$$\text{var}(\log(p_{jst})) = \frac{1}{N_{jt}} \sum_{m \in \mathcal{M}} N_{jmt} \text{var}(\log(p_{jst})|m). \tag{5}$$

Note that we do not use Bessel's correction in these two variance formulas. This approach is equivalent to demeaning each $\log(p_{jst})$ observation with respect to the average log price in market m, and then calculating the variance over all observations. We include this approach because it is more closely related to the variance decomposition in Section 6.

Summary statistics for these two alternative approach are shown in Table 7, separately for products defined as UPCs and brands. As expected the difference between the dispersion statistics based on the standard deviation of log prices and the standard deviation of normalized prices is negligible. On the other hand, the standard deviation calculated as the square root of Eq. 5 is slightly larger at the DMA and 3-digit ZIP code level compared to the standard deviation of the log of prices. Overall, our main conclusions are unchanged using these two alternative dispersion statistics.

B.2 Comparison to Kaplan and Menzio 2015

Our results are not directly comparable to Kaplan and Menzio (2015) because their work is based on different data and a substantially smaller number of products, as we already discussed in Section 2. Also, they measure price dispersion using the standard deviation of prices normalized relative to the market-average price level, $p_{jst}/\bar{p}_{jmt}$, whereas our main price dispersion measure is the standard deviation of log prices. However, as expected, the different dispersion statistics yield almost identical dispersion measures (see the sensitivity analysis in Appendix B.1) and hence they are not a source of differences in the results.

Kaplan and Menzio (2015) report that the standard deviation of normalized prices for the mean UPC at the Scantrack/quarter level is 0.19. In our data, the corresponding standard deviation is 0.10 for the median UPC at the 3-digit ZIP code/week level and 0.12 at the the Scantrack/week level. Hence, the price dispersion of identical products at a given moment in time is substantially smaller than the dispersion level that Kaplan and Menzio report at the quarter level for a small product sample. Comparable brand-level results are not reported in Kaplan and Menzio (2015).[36]

[36]In Kaplan and Menzio (2015) a brand aggregate is obtained using a "set of products that share the same features and the same size, but may have different brands and different UPCs."

C Base prices and promotions: Details

C.1 Choice of promotion threshold $\bar{\delta}$

Assuming that every event when the price of a product is strictly less than the base price is a price promotion, we could define the promotion indicator as $D_{jst} = \mathbb{I}\{\delta_{jst} > 0\}$. However, it is unlikely that any brand or category manager designs a price promotion that offers only a negligible price discount. Hence, to find a suitable threshold value $\bar{\delta}$ to use in our analysis we examine the distribution of the percentage price discounts, δ_{jst}, pooled across all products, stores, and weeks when $p_{jst} < b_{jst} \Leftrightarrow \delta_{jst} > 0$. This distribution is shown in the top right panel of Fig. 4 and summarized in Table 3. The median percentage price discount across all events is 17.7%. There are instances of small percentage price discounts, but the overall incidence of such events is small. For example, in only slightly less than 10% of all events the price discount is less than 5%, $0 < \delta_{jst} < 0.05$. It is implausible that these observations represent a planned price promotion. Rather, such observations are likely due to measurement error in either the price or base price. Such measurement error can arise due to differences between the promotional calendar in a store and the Nielsen RMS definition of a week. For example, suppose a product was offered at a 20 percent price discount during a two week period starting on a Monday and ending on Sunday, May 30. Because a week in the RMS data ends on a Saturday, the RMS week that begins on May 30 and ends on Saturday, June 4 will include one day when the product was offered at the 20 percent price discount and six days when the product was sold at the regular (base) price. The data report the average price over these seven days, which is an average over the promoted and non-promoted prices. The inferred percentage price discount, δ_{jst}, is likely to be small in this example, and it will not accurately represent the promotional price discount. In order to ameliorate measurement error we use the threshold of $\bar{\delta} = 0.05$, as discussed in Section 5.2.

C.2 Calculation of percent volume sold on promotion and lift factors

The percentage of product volume that is sold during a promotional period is given by

$$\nu_{js} = \frac{\sum_{t \in \mathcal{T}_{js}, D_{jst}=1} q_{jst}}{\sum_{t \in \mathcal{T}_{js}} q_{jst}}.$$

Here, q_{jst} is the number of product j units sold in store s in week t. To calculate a corresponding product-level statistic, ν_j, we take a weighted average of ν_{js} over all stores s, with weights N_{js}, the number of observations for store s.

The lift factor (promotion multiplier) is given by

$$L_{js} = \frac{\frac{\sum_{t \in \mathcal{T}_{js}, D_{jst}=1} q_{jst}}{N_{js}^D}}{\frac{\sum_{t \in \mathcal{T}_{js}, D_{jst}=0} q_{jst}}{N_{js} - N_{js}^D}}.$$

Alternatively, we could calculate L_{js} using the predicted volume in the absence of a promotion in the denominator, based on a demand model or a weighted average of the observed non-promoted volume. To obtain a product-level lift factor L_j we aggregate over L_{js} using the same approach that we use to aggregate the volume percentages.

D Derivation of price variance decompositions

All decompositions are performed at the product level, hence we drop the subscript j. $\mathcal{M}$ is the set of all markets, $\mathcal{S}_m$ is the set of all stores in market $m \in \mathcal{M}$, and $\mathcal{S} = \cup_{m\in\mathcal{M}}\mathcal{S}_m$ is the set of all stores. For each store s we observe prices in periods $t \in \mathcal{T}_s$. Correspondingly, S_m is the number of stores in market m, $S = \sum_{m\in\mathcal{M}} S_m$ is the number of all stores, and N_s is the number of observations for stores s. Then the total number of observations is $N = \sum_{s\in\mathcal{S}} N_s$, and the number of observations in market m is $N_m = \sum_{s\in\mathcal{S}_m} N_s$.

Define the overall (national) average price, the average price in market m, and the average price in stores s:

$$\begin{aligned}
\bar{p} &= \frac{1}{N}\sum_{s\in\mathcal{S}}\sum_{t\in\mathcal{T}_s} p_{st}, \\
\bar{p}_m &= \frac{1}{N_m}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s} p_{st}, \\
\bar{p}_s &= \frac{1}{N_s}\sum_{t\in\mathcal{T}_s} p_{st}.
\end{aligned}$$

Similarly, define the average base price in market m and the average base price in store s :

$$\begin{aligned}
\bar{b}_m &= \frac{1}{N_m}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s} b_{st}, \\
\bar{b}_s &= \frac{1}{N_s}\sum_{t\in\mathcal{T}_s} b_{st}.
\end{aligned}$$

Our goal is to provide a decomposition for the overall variance of prices,

$$\text{var}(p_{st}) = \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st} - \bar{p})^2.$$

D.1 Basic decomposition

Define

$$\begin{aligned}\mathrm{var}(\bar{p}_m) &= \frac{1}{N}\sum_{m\in\mathcal{M}} N_m(\bar{p}_m-\bar{p})^2,\\ \mathrm{var}(\bar{p}_s|m) &= \frac{1}{N_m}\sum_{s\in\mathcal{S}_m} N_s(\bar{p}_s-\bar{p}_m)^2,\\ \mathrm{var}(p_{st}|s) &= \frac{1}{N_s}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_s)^2.\end{aligned}$$

$\mathrm{var}(p_m)$ is the variance of average market-level prices across markets, $\mathrm{var}(\bar{p}_s|m)$ is the within-market variance of average store-level prices, and $\mathrm{var}(p_{st}|s)$ is the within-store variance of prices over time. Note that $\mathrm{var}(p_m)$ and $\mathrm{var}(\bar{p}_s|m)$ are calculated as weighted averages, using the number of observations in each market and the number of observations for each store as weights.

We first decompose the overall variance of prices, $\mathrm{var}(p_{st})$, into an across-market and a within-market term:

$$\begin{aligned}\mathrm{var}(p_{st}) &= \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p})^2\\ &= \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m+\bar{p}_m-\bar{p})^2\\ &= \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2+\frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(\bar{p}_m-\bar{p})^2\\ &= \mathrm{var}(\bar{p}_m)+\frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2. \qquad (6)\end{aligned}$$

Note that the third line in this formula follows because

$$\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)(\bar{p}_m-\bar{p}) = \sum_{m\in\mathcal{M}}(\bar{p}_m-\bar{p})\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)=0.$$

To further decompose the within-market term, note that

$$\begin{aligned}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2 &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_s+\bar{p}_s-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_s)^2+\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(\bar{p}_s-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}N_s\mathrm{var}(p_{st}|s)+N_m\mathrm{var}(\bar{p}_s|m). \qquad (7)\end{aligned}$$

Here, to derive the second line we used

$$\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_s)(\bar{p}_s-\bar{p}_m) = \sum_{s\in\mathcal{S}_m}(\bar{p}_s-\bar{p}_m)\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_s)=0.$$

Substituting (7) in Eq. 6, we obtain the desired decomposition of the overall price variance into the variance of average market-level prices, the weighted average of the within-market variances of average store-level prices, and the weighted average of the within-store variances of prices:

$$\begin{aligned}\mathrm{var}(p_{st}) &= \mathrm{var}(\bar{p}_m) + \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2\\ &= \mathrm{var}(\bar{p}_m) + \frac{1}{N}\sum_{m\in\mathcal{M}} N_m \mathrm{var}(\bar{p}_s|m) + \frac{1}{N}\sum_{m\in\mathcal{M}}\sum_{s\in\mathcal{S}_m} N_s \mathrm{var}(p_{st}|s)\\ &= \mathrm{var}(\bar{p}_m) + \frac{1}{N}\sum_{m\in\mathcal{M}} N_m \mathrm{var}(\bar{p}_s|m) + \frac{1}{N}\sum_{s\in\mathcal{S}} N_s \mathrm{var}(p_{st}|s). \qquad (8)\end{aligned}$$

D.2 Decomposition into base price and promotion components

We start with an alternative decomposition of the within-market term (7):

$$\begin{aligned}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2 &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-b_{st}+b_{st}-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-b_{st})^2 + \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{p}_m)^2\\ &\quad +2\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-b_{st})(b_{st}-\bar{p}_m). \qquad (9)\end{aligned}$$

Note that

$$\begin{aligned}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{p}_m)^2 &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{b}_m+\bar{b}_m-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{b}_m)^2 + N_m(\bar{b}_m-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{b}_s+\bar{b}_s-\bar{b}_m)^2 + N_m(\bar{b}_m-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-\bar{b}_s)^2 + \sum_{s\in\mathcal{S}_m} N_s(\bar{b}_s-\bar{b}_m)^2 + N_m(\bar{b}_m-\bar{p}_m)^2\\ &= \sum_{s\in\mathcal{S}_m} N_s \mathrm{var}(b_{st}|s) + N_m \mathrm{var}(\bar{b}_s|m) + N_m(\bar{b}_m-\bar{p}_m)^2. \qquad (10)\end{aligned}$$

Substituting (10) in (9) and rearranging terms, we obtain

$$\begin{aligned}\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(p_{st}-\bar{p}_m)^2 &= N_m \mathrm{var}(\bar{b}_s|m) + \sum_{s\in\mathcal{S}_m} N_s \mathrm{var}(b_{st}|s)\\ &\quad + N_m(\bar{b}_m-\bar{p}_m)^2 + \sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-p_{st})^2\\ &\quad - 2\sum_{s\in\mathcal{S}_m}\sum_{t\in\mathcal{T}_s}(b_{st}-p_{st})(b_{st}-\bar{p}_m). \qquad (11)\end{aligned}$$

Define

$$\text{var}(b_{st} - p_{st}|m) = \frac{1}{N_m} \sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} \big((b_{st} - p_{st}) - (\bar{b}_m - \bar{p}_m)\big)^2$$
$$= \frac{1}{N_m} \sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} (b_{st} - p_{st})^2 - (\bar{b}_m - \bar{p}_m)^2. \quad (12)$$

Rearranging (12) and substituting in Eq. 11, we obtain

$$\sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} (p_{st} - \bar{p}_m)^2 = N_m \text{var}(\bar{b}_s|m) + \sum_{s \in \mathcal{S}_m} N_s \text{var}(b_{st}|s)$$
$$+ N_m \text{var}(b_{st} - p_{st}|m) + 2N_m(\bar{b}_m - \bar{p}_m)^2$$
$$- 2 \sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} (b_{st} - p_{st})(b_{st} - \bar{p}_m). \quad (13)$$

Define the within-market covariance between the promotional price discounts, $b_{st} - p_{st}$, and the difference between the store-level base price and the average market price, $b_{st} - \bar{p}$:

$$\text{cov}(b_{st} - p_{st}, b_{st} - \bar{p}_m|m) = \frac{1}{N_m} \sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} \big((b_{st} - p_{st}) - (\bar{b}_m - \bar{p}_m)\big)$$
$$\times \big((b_{st} - \bar{p}_m) - (\bar{b}_m - \bar{p}_m)\big)$$
$$= \frac{1}{N_m} \sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} (b_{st} - p_{st})(b_{st} - \bar{p}_m) - (\bar{b}_m - \bar{p}_m)^2. \quad (14)$$

Rearranging and substituting (14) in Eq. 13, we then obtain

$$\sum_{s \in \mathcal{S}_m} \sum_{t \in \mathcal{T}_s} (p_{st} - \bar{p}_m)^2 = N_m \text{var}(\bar{b}_s|m) + \sum_{s \in \mathcal{S}_m} N_s \text{var}(b_{st}|s)$$
$$+ N_m \text{var}(b_{st} - p_{st}|m) + 2N_m(\bar{b}_m - \bar{p}_m)^2$$
$$- 2N_m \text{cov}(b_{st} - p_{st}, b_{st} - \bar{p}_m|m) - 2N_m(\bar{b}_m - \bar{p}_m)^2$$
$$= N_m \text{var}(\bar{b}_s|m) + \sum_{s \in \mathcal{S}_m} N_s \text{var}(b_{st}|s)$$
$$+ N_m \text{var}(b_{st} - p_{st}|m) - 2N_m \text{cov}(b_{st} - p_{st}, b_{st} - \bar{p}_m|m). \quad (15)$$

Finally, we substitute (15) in Eq. 6 and note that $\text{cov}(b_{st} - p_{st}, b_{st} - \bar{p}_m|m) = \text{cov}(b_{st} - p_{st}, b_{st}|m)$ to obtain the variance decomposition:

$$\text{var}(p_{st}) = \text{var}(\bar{p}_m)$$
$$+ \frac{1}{N} \sum_{m \in \mathcal{M}} N_m \text{var}(\bar{b}_s|m) + \frac{1}{N} \sum_{s \in \mathcal{S}} N_s \text{var}(b_{st}|s)$$
$$+ \frac{1}{N} \sum_{m \in \mathcal{M}} N_m \text{var}(b_{st} - p_{st}|m) - 2\frac{1}{N} \sum_{m \in \mathcal{M}} N_m \text{cov}(b_{st} - p_{st}, b_{st}|m). \quad (16)$$

D.3 Example: Price promotions may decrease the overall price variance

We focus only on price variation within one market. Assume that base prices in each store s are constant over time, $b_{st} \equiv \bar{b}_s$, and uniformly distributed around the mean base price $\bar{b}$ on the interval $[\bar{b} - \nu, \bar{b} + \nu]$. Suppose that only stores with above average base prices, $b_{st} = \bar{b}_s > \bar{b}$, promote the product, and that the promoted price is always $p_{st} = \bar{b}$. All stores with base prices $b_{st} = \bar{b}_s \leq \bar{b}$ always sell the product at the base price, $p_{st} = b_{st}$. Define the promotion indicator $D_{st} = \mathbb{I}\{p_{st} < b_{st}\}$ (this corresponds to the promotion definition in Section 5.2 with a threshold $\bar{\delta} = 0$). The incidence of promotions is constant, $\pi \equiv \Pr\{D_{st} = 1 | b_{st} > \bar{b}\}$. There is a continuum of stores with mass 1, and promotions are independent across stores and across time periods.

The mean price is given by:

$$\begin{aligned} \mathbb{E}[p_{st}] &= \frac{1}{2}\mathbb{E}[p_{st}|b_{st} \leq \bar{b}] + \frac{1}{2}\left((1-\pi)\mathbb{E}[p_{st}|D_{st}=0, b_{st} > \bar{b}] + \pi\mathbb{E}[p_{st}|D_{st}=1, b_{st} > \bar{b}]\right) \\ &= \frac{1}{2}\left(\bar{b} - \frac{\nu}{2}\right) + \frac{1}{2}\left((1-\pi)\left(\bar{b} + \frac{\nu}{2}\right) + \pi\bar{b}\right) \\ &= \bar{b} - \frac{\pi\nu}{4}. \end{aligned}$$

The across-store variance of base prices is given by the variance of a uniform distribution,

$$\text{var}(\bar{b}_s) = \frac{\nu^2}{3}.$$

To derive the variance of the promotional price discounts we first calculate

$$\begin{aligned} \mathbb{E}\left[(b_{st} - p_{st})^2\right] &= \frac{1}{2}\pi\mathbb{E}\left[(b_{st} - p_{st})^2 | D_{st} = 1, b_{st} > \bar{b}\right] \\ &= \frac{1}{2}\pi\mathbb{E}\left[(b_{st} - \bar{b})^2 | D_{st} = 1, b_{st} > \bar{b}\right] \\ &= \frac{1}{2}\pi\int_{\bar{b}}^{\bar{b}+\nu}(x - \bar{b})^2\frac{1}{\nu}dx \\ &= \frac{\pi\nu^2}{6}. \end{aligned}$$

Similarly, to derive the covariance between the promotional price discounts and the base prices we use the expression for $\mathbb{E}\left[(b_{st} - p_{st})^2\right]$ above to obtain

$$\begin{aligned} \mathbb{E}\left[(b_{st} - p_{st})(b_{st} - \bar{b})\right] &= \frac{1}{2}\pi\mathbb{E}\left[(b_{st} - p_{st})(b_{st} - \bar{b}) | D_{st} = 1, b_{st} > \bar{b}\right] \\ &= \frac{1}{2}\pi\mathbb{E}\left[(b_{st} - \bar{b})^2 | D_{st} = 1, b_{st} > \bar{b}\right] \\ &= \frac{\pi\nu^2}{6}. \end{aligned}$$

Also, the squared difference between the mean base price and shelf price is:

$$(\bar{b} - \bar{p})^2 = \frac{\pi^2\nu^2}{16}.$$

Hence,

$$\begin{aligned}\mathrm{var}(b_{st} - p_{st}) &= \mathbb{E}\left[(b_{st} - p_{st})^2\right] - (\bar{b} - \bar{p})^2 \\ &= \frac{\pi \nu^2}{6} - \frac{\pi^2 \nu^2}{16}.\end{aligned}$$

Also,

$$\begin{aligned}\mathrm{cov}(b_{st} - p_{st}, b_{st}) &= \mathrm{cov}(b_{st} - p_{st}, b_{st} - \bar{b}) \\ &= \mathbb{E}\left[(b_{st} - p_{st})(b_{st} - \bar{b})\right] \\ &= \frac{\pi \nu^2}{6}.\end{aligned}$$

Combining the three components we obtain the variance of prices,

$$\begin{aligned}\mathrm{var}(p_{st}) &= \mathrm{var}(\bar{b}_s) + \mathrm{var}(b_{st} - p_{st}) - 2\mathrm{cov}(b_{st} - p_{st}, b_{st}) \\ &= \frac{\nu^2}{3} + \frac{\pi \nu^2}{6} - \frac{\pi^2 \nu^2}{16} - 2\frac{\pi \nu^2}{6} \\ &= \nu^2 \left(\frac{1}{3} - \frac{\pi}{6} - \frac{\pi^2}{16}\right).\end{aligned}$$

The EDLP vs. Hi-Lo adjustment factor is positive if $\pi > 0$ and strictly increasing in π, and the variance of prices is strictly decreasing in the frequency of promotions, π:

$$\mathrm{var}(p_{st}) = v^2 \left(\frac{1}{3} - \frac{\pi}{6} - \frac{\pi^2}{16}\right).$$

E Visualization of price similarity using principal component analysis

We examine the similarity of pricing patterns across stores based on the whole time series of store-level prices. For product j, we observe the vector of prices $\boldsymbol{p}_s = (p_{s1}, \ldots, p_{sT})$ for each store $s \in \mathcal{S}$ and over the prior 2008-2010 (we suppress the index j for notational simplicity). The sample of prices for product j then consists of $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_S$. Our goal is to visualize the price vectors $\boldsymbol{p}_s$, which is not directly feasible given the dimensionality of $\boldsymbol{p}_s$. Instead, we conduct a principal components analysis (PCA) of the store-level price vectors. PCA is an unsupervised dimensionality reduction technique that allows us to represent each $\boldsymbol{p}_s$ in a low-dimensional space while maintaining as much of the original information (variance) contained in $\boldsymbol{p}_s$ as possible.[37]

We perform a PCA for the top 1,000 products (UPCs) in our sample, based on total revenue rank. We only choose these top products because we need to be able to consistently observe the weekly prices, p_{st}, across stores for the analysis to be feasible. For smaller products there is a larger incidence of missing values.

[37] See, for example, Hastie et al. (2009) for a thorough introduction to principal components analysis.

The top panel in Fig. 20 displays box plots of the percentage of the price variance that is explained by the first twenty principal components. Each box plot shows the distribution (weighted by total product revenue) of these percentages across the products in our sample. The first principal component explains 20% of the price variance for the median product, and all the first five principal components explain at least 5% of the price variance. In the bottom panel of Fig. 20 we display box plots of the cumulative percentage of the price variance explained by the top principal components. The top five principal components explain 53% and the top ten principal components explain 68% of the price variance for the median product. Hence, a large percentage of the information in the original store-level price vectors over the 2008-2010 period can be explained by a small number of principal components. A representation of the original, high-dimensional price data in a low-dimensional space is therefore meaningful.

Following the case of Tide HE Liquid Laundry Detergent (100 oz) in Figs. 7 and 8, we present additional examples in Fig. 21, including Prilosec (42 count), Pepsi (12 oz cans 12 pack), and private-label milk (2 percent, 1 gallon). In this graph we only include the store-level price vectors for a subset of all retail chains. For Prilosec and Pepsi we find a pattern that is similar to the case of Tide laundry detergent—a large degree of price similarity within retail chains and a significantly smaller degree of price similarity at the market level. The case of private-label milk is quite different, however, as there is much heterogeneity in prices both at the chain and the market level.

F Overview of promotion coordination

To provide an overview of promotion coordination we summarize the percentage of all stores in a retail chain that promote a specific product during week t.

Promotions are captured using the indicator $D_{jst} \in \{0, 1\}$, such that $D_{jst} = 1$ if product j is promoted in store s in week t. We calculate the chain/market level *promotion percentage* for product j:

$$\phi_{jcmt} = \frac{\sum_{s \in \mathcal{S}_{jcmt}} D_{jst}}{|\mathcal{S}_{jcmt}|}.$$

$\mathcal{S}_{jcmt}$ includes the stores that belong to retail chain c in market m and carry product j in week t, and $|\mathcal{S}_{jcmt}|$ is the corresponding number of stores.

The graphs in the top row of Fig. 22 display histograms of the promotion percentages ϕ_{jcmt}, pooled over all products, chains, markets, and time periods between 2008 and 2010 (Table 11 contains detailed summary statistics). The distributions are weighted using total product revenue weights. We display the promotion percentage distributions conditional on $\phi_{jcmt} > 0$, i.e. weeks when at least one store in chain $c$and market m promotes the product, to avoid that the histograms are dominated by large mass points at 0. The percentage of observations when none of the stores promoted the product, $\phi_{jcmt} = 0$, is indicated separately at the bottom of each graph.

We define markets as DMAs to ensure that the chains have a sufficiently large number of stores in a local market.[38] To see why this is important consider the case when a chain has only one local store. Then the promotion percentage is always 0 or 1, indicating perfect promotion coordination. For the same reason, we summarize the promotion percentage distributions only for observations when the retail chain c carries the product in at least five stores in the local market.

The top row in Fig. 22 shows the promotion percentage distributions for all products (left panel) and the top 1,000 products as measured by annual, national product revenue (right panel). Both histograms reveal a large mass point at 1, indicating perfect promotion coordination. The promotion percentages, conditional on $\phi_{jcmt} > 0$, are overall larger among the top 1,000 products, with a median of 0.548 compared to a median of 0.41 for all products. This indicates a larger degree of promotion coordination among the top-selling products, although the percentage of $\phi_{jcmt} = 0$ observations is smaller for the top 1,000 products: 55.1% versus 62.2% among all products. However, the latter finding may simply reflect an overall higher promotion frequency among the top 1,000 products, not a smaller degree of coordination on weeks when none of the stores in the chain promote the product.

Although suggestive of promotion coordination, the overall extent of promotion coordination conveyed by Fig. 22 is difficult to assess without a comparison to a baseline where promotions are not coordinated. To provide such a baseline we simulate data assuming that promotions are chosen independently across stores. For each product j and stores s we calculate the promotion frequency π_{js} using the 2008-2010 data, as in Section 5.2. For each store s and week t we then draw a promotion indicator $\tilde{D}_{jst}$ from a Bernoulli distribution with success probability π_{js}. The distributions of promotion percentages for the simulated data are shown in the second row of Fig. 22. The histograms indicate a much smaller degree of promotion coordination compared to the observed promotion percentages. The median promotion percentage among all products in the simulated data is 0.200, compared to 0.41 in the actual data, and the corresponding percentages among the top 1,000 products are 0.279 in the simulated data versus 0.548 in the observed data. Indeed, the displayed distributions understate the difference between the simulated and the original data because they are conditional on $\phi_{jcmt} > 0$, and hence do not reveal the large difference in observations where none of the stores in a chain promotes a product. In the observed data, $\phi_{jcmt} = 0$ for 62.2% of all observations, compared to 28.3% of all observations in the simulated data.

The distributions in the top rows of Fig. 22 are based on observations at the chain/market level in a given week, conditional on at least one store in the chain promoting product j. This leads to an asymmetry between observations with highly coordinated promotions and observations with a small number of stores promoting a product. For example, suppose that all promotions were perfectly coordinated within a retail chain, such that $\phi_{jcmt} = 1$ in the week when the promotion is held. Suppose there were some small differences in the timing of the end date of the promotion across stores in the chain, such that a small number of stores would still offer the

[38] See Table 8 for summary statistics on the number of stores per retail chain at the DMA and ZIP+3 level.

promotion for a few days in week $t+1$. Then each perfectly coordinated promotion observation, $\phi_{jcmt}=1$, would have an associated observation with a small, positive $\phi_{jcm,t+1}$ value, suggesting that perfectly coordinated promotion events were as frequent as almost completely uncoordinated promotion events. Hence, as an alternative summary of promotion coordination we associate each store-level promotion event, $D_{jst}=1$, with the corresponding promotion percentage, ϕ_{jcmt}, and display the distribution of the promotion percentages based on all store level observations such that $D_{jst}=1$. This is equivalent to displaying the distribution of ϕ_{jcmt} weighted by the number of stores that promote product j in chain c and market m in week t. The results are shown in the third row of Fig. 22 (Table 11 contains detailed numbers), and strongly indicate that store-level promotions are coordinated at the chain-market level. Furthermore, the differences between the actual and simulated data, shown in the bottom row of Fig. 22, are large.

G Bayesian hierarchical demand model

This section provides an overview of the Bayesian hierarchical linear regression model to estimate the brand/store-level demand parameters. See Rossi et al. (2005), Chapter 3.7, for a detailed exposition of the model and the MCMC sampling approach.

The goal is to obtain the posterior distribution of the demand parameters for brand j and store s in model (4). The store-level parameter vector $\theta_{js}=(\alpha_{js},\beta_{js},\gamma_{js})$ includes the intercept, α_{js}, the own and cross-price elasticities, β_{jks}, and the promotion parameters, γ_{jks}. We do not estimate the large number of 3-digit ZIP code/time fixed effects, $\tau_j(s,t)$, as part of the Bayesian hierarchical model. Instead, we first project all variables in the demand model, $\log(1+q_{jst})$, $\log(p_{kst})$, and D_{kst}, on the fixed effects. We then use the residuals from this projection to estimate the store-level demand parameters in the model:

$$\widetilde{\log(1+q_{jst})}=\alpha_{js}+\sum_{k\in\mathcal{J}_{js}}\beta_{jks}\widetilde{\log(p_{kst})}+\sum_{k\in\mathcal{J}_{js}}\gamma_{jks}\widetilde{D_{kst}}+\epsilon_{jst}, \qquad (17)$$

$$\epsilon_{jst}\sim N(0,\sigma^2_{js}).$$

ϵ_{jst} is i.i.d. across stores and time. From now on we drop the brand index j to simplify the notation.

We specify a normal first-stage prior or population distribution for the the store-level demand parameters, θ_s:

$$\theta_s\sim N(\mu,V_\theta).$$

The parameter vectors for different stores are conditionally independent, given μ and V_θ. More flexible priors are possible, such as a mixture of normal distributions, which has been used in the literature (Rossi et al., 2005).

We further specify the second-stage prior distribution of V_θ and μ:

$$V_\theta\sim \mathrm{IW}(\nu,V),$$
$$\mu|V_\theta\sim N(\bar{\mu},V_\theta\otimes A^{-1}).$$

IW denotes an inverse Wishart distribution. The error term variances, σ_s^2, are independent draws from an inverse chi-squared distribution,

$$\sigma_s^2 \sim \frac{\nu_\epsilon r_s^2}{\chi^2_{\nu_\epsilon}}.$$

Here, ν_ϵ denotes the degrees of freedom and r_s^2 is a scale parameter.

The MCMC algorithm to obtain the posterior distribution of the model parameters is performed using Peter Rossi's `bayesm` package[39] in R. We run the algorithm using the default settings for the hyper-parameters in the `bayesm` package. These settings allow for a diffuse prior:

$$\begin{aligned} v &= 3 + n, \\ V &= v I_n, \\ \bar{\mu} &= 0, \\ A &= 0.01, \\ \nu_\epsilon &= 3, \\ r_s^2 &= \text{var}(\widetilde{\log(1 + q_{st})}). \end{aligned}$$

Here, n is the dimension of the parameter vector θ_s.

We choose a chain length of 20,000 (after 2,000 initial burn-in draws) and keep every 10th draw to calculate the posterior means and the 95% credible intervals of the parameters. A visual inspection of the trace plots for a large number of randomly selected parameters (across brands and stores) indicates convergence of the chain.

H Causal price and promotion effects: Sensitivity analysis

Price and promotion endogeneity occurs if the retail chains set prices or promotions based on demand shocks that are observed to them but not to us. To avoid endogeneity bias we include time-invariant store fixed effects α_{js} in the demand model (4). The store fixed effects account for systematic demand differences across stores that are associated with systematic differences in prices and promotions. Further, we include the market/time fixed effects $\tau_j(s, t)$ to account for demand shocks and brand-specific trends in demand at a narrowly defined geographic level.

We can interpret the price and promotion estimates as causal if (i) $\tau_j(s, t)$ captures all time-varying demand components that may be correlated with p_{kst} and D_{kst}, (ii) there is variation in the price and promotion *changes* over time across stores, and (iii) the difference in price and promotion changes across stores reflects store or chain-specific changes in costs, wholesale prices, markups, or other factors that affect prices and promotion but not directly demand. These assumptions are not directly testable, but we can perform an analysis to indicate if our estimates are sensitive to the inclusion and exact specification of the fixed effects. Thus, we first estimate

[39] https://cran.r-project.org/web/packages/bayesm/index.html

Table 7 Additional price and base price dispersion statistics

		Median	Mean	Percentiles							
				0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
Prices											
Product definition: UPC											
Normalized price SD	National	0.157	0.160	0.014	0.058	0.088	0.122	0.197	0.233	0.267	0.327
	DMA	0.107	0.112	0.006	0.031	0.049	0.076	0.142	0.181	0.205	0.258
	ZIP+3	0.097	0.101	0.002	0.021	0.038	0.066	0.132	0.170	0.192	0.247
Demeaned log-price SD	National	0.161	0.163	0.014	0.058	0.090	0.124	0.200	0.238	0.268	0.318
	DMA	0.121	0.125	0.010	0.040	0.060	0.089	0.158	0.195	0.218	0.265
	ZIP+3	0.113	0.116	0.002	0.026	0.048	0.080	0.151	0.187	0.209	0.257
Product definition: Brand											
Normalized price SD	National	0.174	0.187	0.073	0.097	0.115	0.140	0.217	0.263	0.330	0.450
	DMA	0.135	0.145	0.045	0.066	0.082	0.104	0.173	0.216	0.251	0.388
	ZIP+3	0.127	0.135	0.035	0.059	0.071	0.096	0.162	0.203	0.237	0.362
Demeaned log-price SD	National	0.175	0.185	0.074	0.098	0.114	0.141	0.216	0.260	0.319	0.427
	DMA	0.146	0.154	0.055	0.073	0.092	0.114	0.181	0.223	0.261	0.364
	ZIP+3	0.139	0.145	0.044	0.067	0.081	0.107	0.172	0.209	0.244	0.332

Table 8 Retail chain and store descriptive statistics

	Mean	SD	Min	Percentiles					Max
				10	25	50	75	90	
Panel A: At National level									
No. of retail chains	81								
No. of stores	17184								
Panel B: At DMA level									
No. of retail chains	6.0	2.6	1	3	4	6	7	9	18
No. of stores	83.8	143.7	1	6	13	32	82	224	1061
Panel C: At ZIP+3 level									
No. of retail chains	4.4	2.2	1	2	3	4	6	8	11
No. of stores	20.5	24.9	1	2	5	10	27	54	170
Panel D: Markets covered by retail chains									
No. of DMA's	15.1	34.8	1	1	2	5	9	31	192
No. of ZIP+3	45.9	105.7	1	4	6	13	32	112	572
Panel E: Stores per retail chain									
National	212.1	443.8	2	10	25	77	171	505	3007
DMA	14.0	26.7	1	1	2	5	13	36	320
ZIP+3	4.6	6.6	1	1	1	2	5	11	109

demand without fixed effects, and then including $\tau_j(s, t)$ defined at the 3-digit ZIP code/quarter, 3-digit ZIP code/month, and 3-digit ZIP code/week levels.

The estimated distributions of the price effects are shown in Fig. 24 and summarized in Table 12. The median elasticity estimate is -1.767 in the model without fixed effects, -1.924 in the model with 3-digit ZIP code/quarter fixed effects, -1.93 with 3-digit ZIP code/month fixed effects, and -1.859 with 3-digit ZIP code/week fixed effects. Hence, controlling for time fixed effects at the local market level moderately changes the distribution of the estimates, and the direction of this change is consistent with price endogeneity if positive demand shocks are correlated with higher prices. However, the elasticity estimates are not particularly sensitive to the exact choice of fixed effects, and the direction of the change in the estimated elasticities when we use year-month or year-week fixed effects instead of year-quarter fixed effects is not indicative of a price endogeneity problem. For a severe price endogeneity problem to exist it would have to be true that there are high-frequency demand shocks that occur at level that is more local than a 3-digit ZIP code area, and that the store or chain managers are able to predict these shocks and correspondingly change prices. This seems a priori implausible, and in particular such localized price-setting is inconsistent with the strong similarity in price and promotion patterns at the retail chain level that we document.

Table 9 Details of price variance decompositions

	Median	Mean	% > 0	Percentiles							
				1%	5%	10%	25%	75%	90%	95%	99%
UPC's											
Basic decomposition											
Across-market	0.269	0.327		0.009	0.068	0.100	0.164	0.444	0.662	0.774	0.980
Across-store	0.251	0.270		0.000	0.028	0.076	0.158	0.373	0.486	0.557	0.670
Within-store	0.396	0.403		0.000	0.033	0.075	0.214	0.578	0.719	0.802	0.977
Decomposition into base prices and promotions											
Across-market	0.269	0.327		0.009	0.068	0.100	0.164	0.444	0.662	0.774	0.980
Across-store mean base price variance	0.303	0.313		0.000	0.028	0.084	0.182	0.432	0.545	0.612	0.760
Within-store base price variance	0.096	0.123		0.000	0.008	0.022	0.051	0.160	0.244	0.323	0.622
Total contribution of promotions	0.188	0.237	0.964	−0.096	0.000	0.007	0.070	0.359	0.537	0.653	0.938
Promotional price discounts	0.329	0.360		0.000	0.005	0.029	0.143	0.546	0.735	0.828	0.975
EDLP vs. Hi-Lo adjustment	−0.084	−0.123	0.103	−0.772	−0.375	−0.286	−0.183	−0.017	0.000	0.004	0.040
Covariance between price discounts and store base price level											
	0.042	0.062	0.920	−0.020	−0.002	0.000	0.009	0.091	0.143	0.188	0.386

Table 9 (continued)

	Median	Mean	% > 0	Percentiles 1%	5%	10%	25%	75%	90%	95%	99%
Brands											
Basic decomposition											
Across-market	0.242	0.297		0.066	0.107	0.133	0.175	0.372	0.564	0.674	0.794
Across-store	0.420	0.423		0.052	0.157	0.201	0.291	0.555	0.652	0.710	0.793
Within-store	0.251	0.280		0.020	0.043	0.062	0.139	0.394	0.528	0.616	0.776
Decomposition into base prices and promotions											
Across-market	0.242	0.297		0.066	0.107	0.133	0.175	0.372	0.564	0.674	0.794
Across-store mean base price variance	0.513	0.499		0.053	0.167	0.234	0.351	0.642	0.752	0.812	0.881
Within-store base price variance	0.094	0.134		0.008	0.023	0.033	0.053	0.159	0.272	0.328	0.755
Total contribution of promotions	0.059	0.070	0.692	−0.613	−0.255	−0.101	−0.012	0.174	0.312	0.411	0.635
Promotional price discounts	0.242	0.299		0.006	0.027	0.053	0.126	0.426	0.621	0.722	0.990
EDLP vs. Hi-Lo adjustment	−0.151	−0.229	0.028	−1.466	−0.785	−0.457	−0.260	−0.067	−0.021	−0.008	0.017
Covariance between price discounts and store base price level											
	0.076	0.115	0.974	−0.008	0.004	0.010	0.033	0.130	0.229	0.393	0.733

Table 10 Percentage of variance of prices, promotion frequency, and promotion depth explained by market and chain factors

	Median	Mean	Percentiles							
			0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
Price										
Market (ZIP+3)	0.465	0.490	0.198	0.241	0.266	0.341	0.620	0.757	0.820	0.932
Chain	0.699	0.652	0.001	0.299	0.410	0.573	0.774	0.826	0.855	0.917
Market/chain	0.881	0.845	0.374	0.589	0.707	0.812	0.921	0.946	0.957	0.978
Promotion frequency										
Market (ZIP+3)	0.361	0.389	0.125	0.182	0.211	0.270	0.477	0.611	0.708	0.879
Chain	0.632	0.587	0.000	0.019	0.177	0.444	0.747	0.822	0.856	0.910
Market/chain	0.800	0.759	0.287	0.443	0.541	0.688	0.873	0.910	0.929	0.963
Promotion depth										
Market (ZIP+3)	0.380	0.407	0.148	0.191	0.215	0.279	0.505	0.649	0.753	0.965
Chain	0.589	0.562	0.000	0.005	0.177	0.402	0.724	0.804	0.840	0.894
Market/chain	0.807	0.772	0.305	0.483	0.589	0.711	0.875	0.913	0.930	0.980

Table 11 Promotion coordination: Summary statistics for promotion percentages

		Median	Mean	% = 0	Percentiles							
					1%	5%	10%	25%	75%	90%	95%	99%
DMA/chain promotion percentages												
Top 1,000	Data	0.548	0.541	55.060	0.011	0.026	0.048	0.143	0.991	1.000	1.000	1.000
	Simulated	0.279	0.313	22.451	0.025	0.060	0.091	0.167	0.429	0.600	0.667	0.833
All	Data	0.410	0.484	62.184	0.010	0.023	0.039	0.111	0.886	1.000	1.000	1.000
	Simulated	0.200	0.255	28.274	0.021	0.048	0.071	0.125	0.353	0.500	0.600	0.800
DMA/chain promotion percentages, store/promotion-weighted												
Top 1,000	Data	0.957	0.822		0.047	0.188	0.379	0.745	1.000	1.000	1.000	1.000
	Simulated	0.416	0.419		0.060	0.125	0.174	0.274	0.556	0.667	0.732	0.833
All	Data	0.901	0.779		0.038	0.158	0.323	0.650	0.991	1.000	1.000	1.000
	Simulated	0.333	0.355		0.045	0.094	0.132	0.212	0.482	0.606	0.681	0.800

Note: Results are based observations when a product is carried in at least 5 stores by the retailer in the DMA

Table 12 Own-price coefficient estimates

	Median	Mean	% > 0	% < 0	% < -1	% significant	Percentiles							
							1%	5%	10%	25%	75%	90%	95%	99%
OLS														
Own-price, no FE's	−1.767	−1.941		0.839	0.670	0.564	−15.182	−6.803	−4.907	−3.078	−0.585	0.689	2.093	9.095
Year-quarter/ZIP+3	−1.924	−2.046		0.857	0.704	0.595	−13.849	−6.608	−4.906	−3.204	−0.753	0.534	1.958	8.851
Year-month/ ZIP+3	−1.930	−2.050		0.856	0.706	0.601	−13.948	−6.647	−4.929	−3.216	−0.763	0.557	2.025	9.116
Year-week/ ZIP+3	−1.859	−1.972		0.847	0.692	0.592	−14.153	−6.658	−4.878	−3.137	−0.691	0.669	2.198	9.568
Cross-price 1	0.153	0.157	0.572			0.203	−8.767	−3.251	−1.855	−0.538	0.925	2.144	3.424	8.338
Cross-price 2	0.067	0.080	0.553			0.180	−6.733	−2.508	−1.401	−0.410	0.611	1.578	2.587	6.482
Cross-price 3	0.048	0.049	0.540			0.163	−6.827	−2.499	−1.395	−0.411	0.547	1.474	2.518	6.657
Own-promotion effect	0.178	0.373	0.775			0.347	−0.610	−0.218	−0.112	0.016	0.494	1.216	1.913	3.579
Bayesian hierarchical														
Own-price	−1.970	−2.026		0.903	0.749	0.708	−8.215	−5.116	−4.163	−2.989	−0.994	−0.033	0.754	4.070
Cross-price 1	0.159	0.193	0.626			0.191	−3.511	−1.364	−0.768	−0.195	0.615	1.205	1.734	3.899
Cross-price 2	0.089	0.084	0.592			0.177	−3.253	−1.226	−0.676	−0.196	0.405	0.871	1.300	2.751
Cross-price 3	0.070	0.063	0.584			0.148	−2.392	−1.029	−0.614	−0.184	0.336	0.718	1.080	2.389
Own-promotion effect	0.184	0.299	0.844			0.452	−0.271	−0.104	−0.042	0.052	0.408	0.835	1.247	2.353

Table 13 Percentage of variance in estimated own-price elasticities and own-promotion effects explained by market and chain factors

	Median	Mean	Percentiles							
			0.01	0.05	0.1	0.25	0.75	0.9	0.95	0.99
OLS estimates										
Own-price elasticity										
Market (ZIP+3)	0.146	0.163	0.047	0.072	0.082	0.108	0.197	0.263	0.311	0.476
Chain	0.172	0.195	0.016	0.042	0.068	0.110	0.250	0.353	0.412	0.601
Market/chain	0.473	0.486	0.186	0.292	0.341	0.394	0.563	0.668	0.732	0.827
Own-promotion effect										
Market (ZIP+3)	0.145	0.170	0.057	0.077	0.089	0.111	0.198	0.283	0.346	0.519
Chain	0.263	0.297	0.013	0.067	0.087	0.158	0.409	0.564	0.653	0.772
Market/chain	0.551	0.558	0.209	0.338	0.381	0.455	0.659	0.749	0.807	0.876
Bayesian hierarchical model estimates										
Own-price elasticity										
Market (ZIP+3)	0.207	0.224	0.082	0.106	0.124	0.158	0.269	0.340	0.400	0.563
Chain	0.235	0.252	0.039	0.072	0.099	0.161	0.334	0.421	0.479	0.563
Market/chain	0.523	0.520	0.239	0.323	0.356	0.424	0.608	0.681	0.724	0.793
Own-promotion effect										
Market (ZIP+3)	0.193	0.217	0.087	0.109	0.126	0.157	0.247	0.328	0.407	0.597
Chain	0.305	0.332	0.044	0.087	0.130	0.202	0.439	0.568	0.657	0.773
Market/chain	0.564	0.567	0.250	0.347	0.392	0.471	0.667	0.744	0.801	0.867

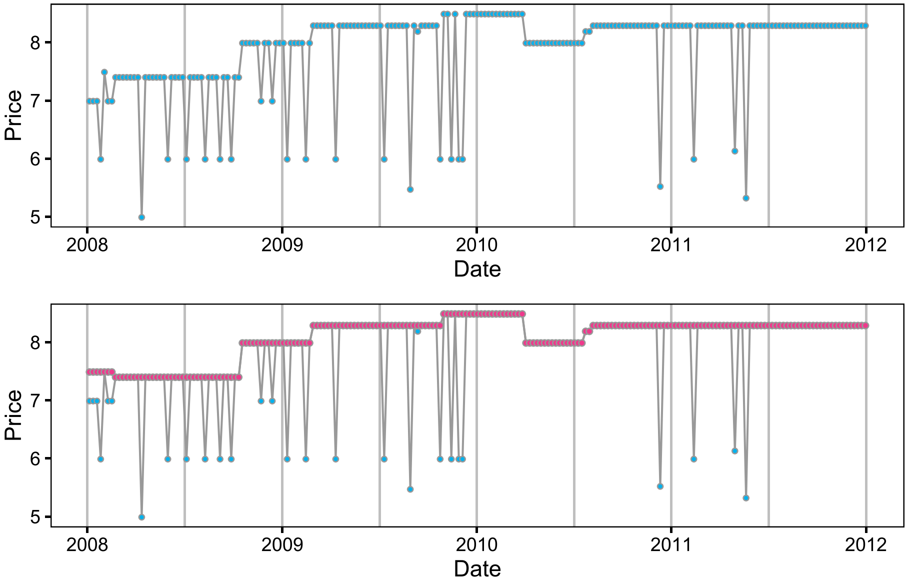

Fig. 15 Predicted base prices, Tide liquid laundry detergent (70 oz)

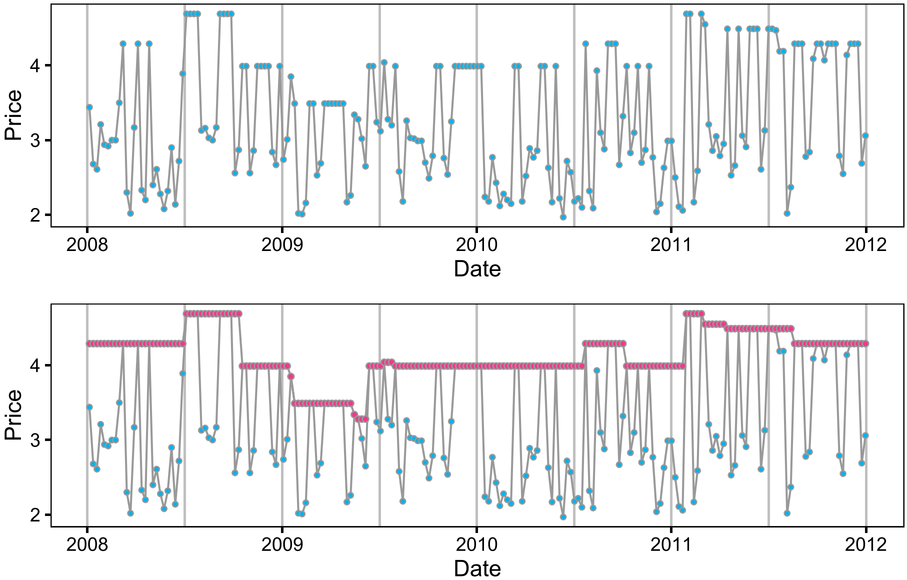

Fig. 16 Predicted base prices, Kellogg's Raisin Bran (20 oz)

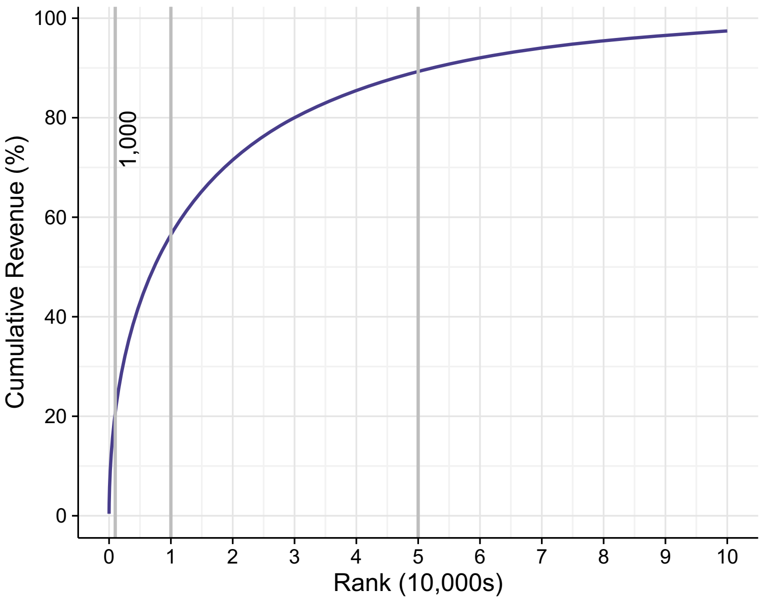

Fig. 17 Cumulative revenue for top 100,000 products

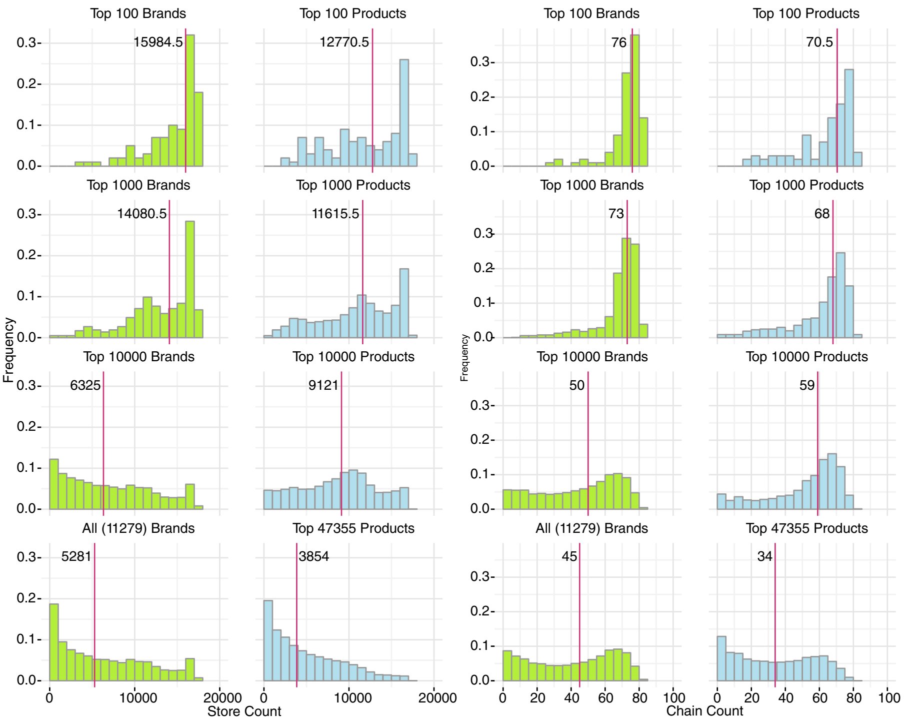

Fig. 18 Product and brand availability across stores and retail chains. Note: The graph displays the distribution of the number of stores (the two columns to the left) and chains (the two columns to the right) at which a product or brand is available. The median of each distribution is indicated using a vertical line. A product (brand) is classified as available if it was sold at least once in a store or chain in 2010. The sample includes 17,184 stores and 81 retail chains

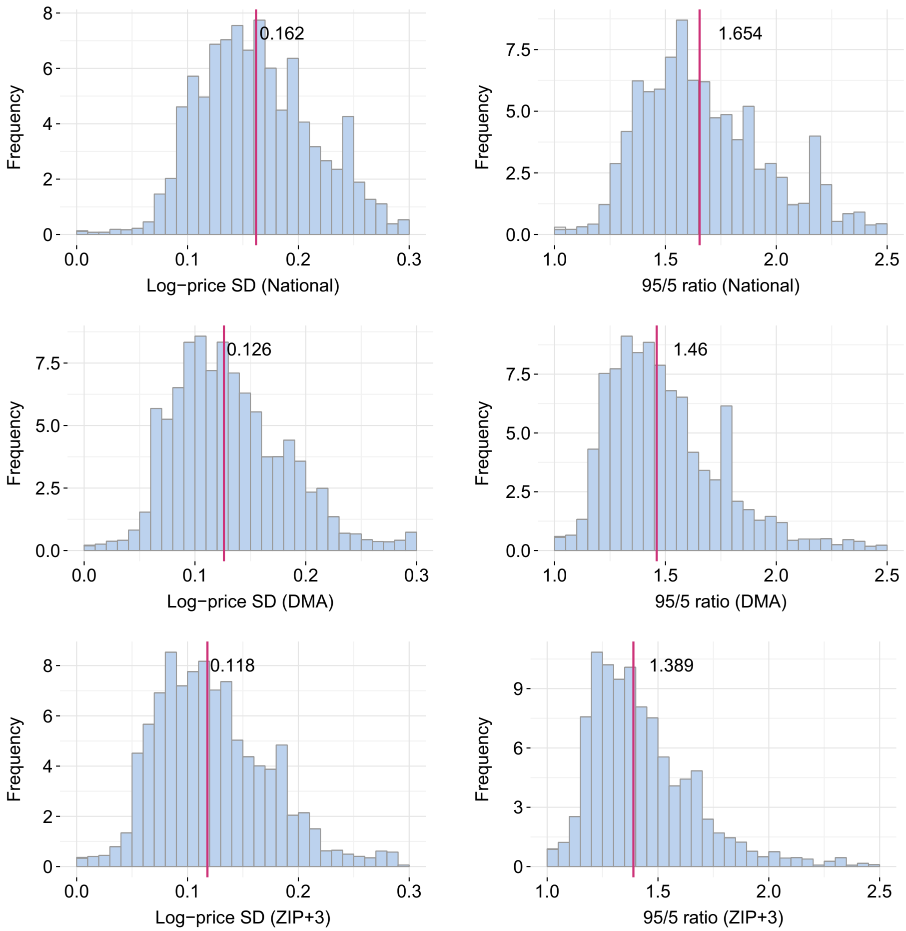

Fig. 19 Base prices dispersion statistics: Brand-level base prices

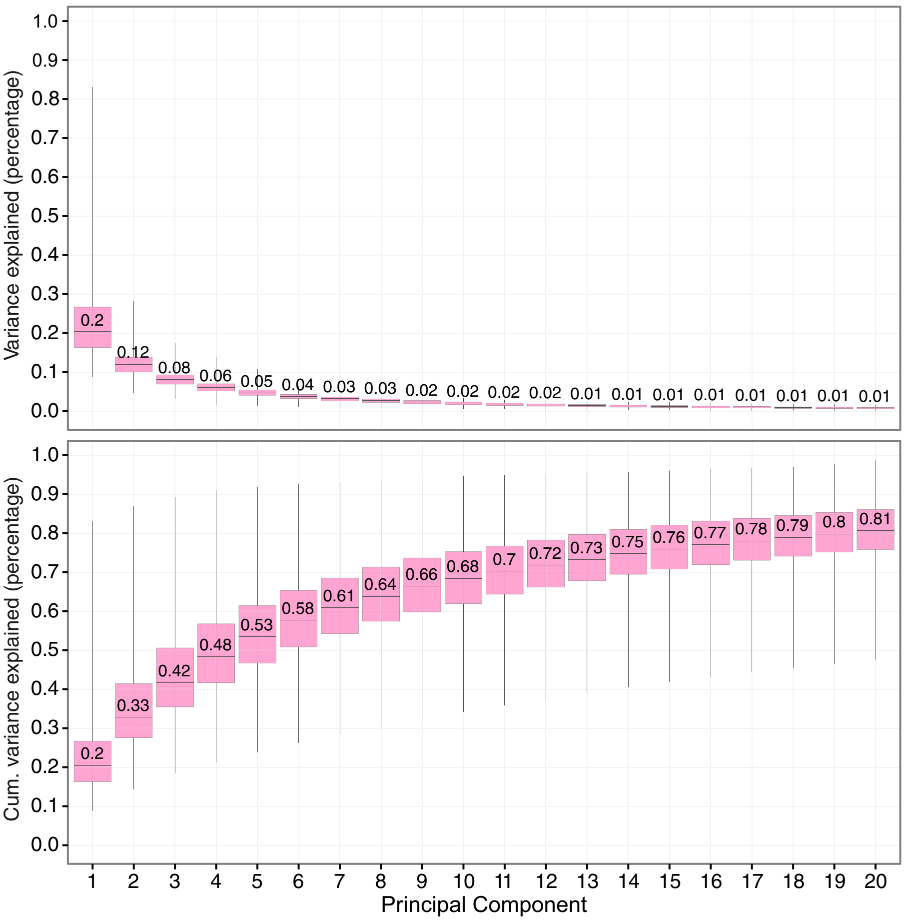

Fig. 20 Percentage and cumulate percentage of price variance explained by principal component

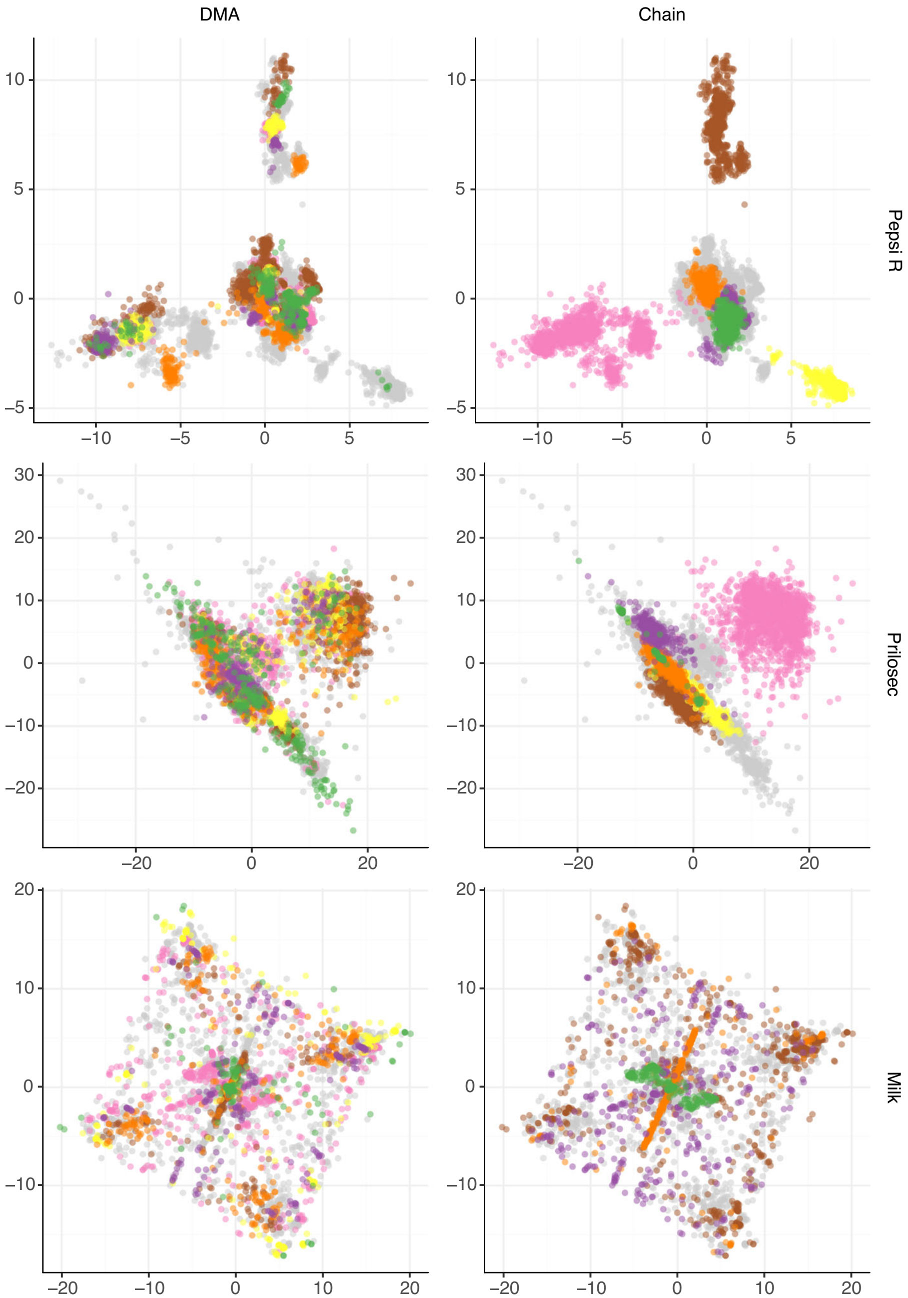

Fig. 21 Projected store-level price vectors colored by retail chain and DMA

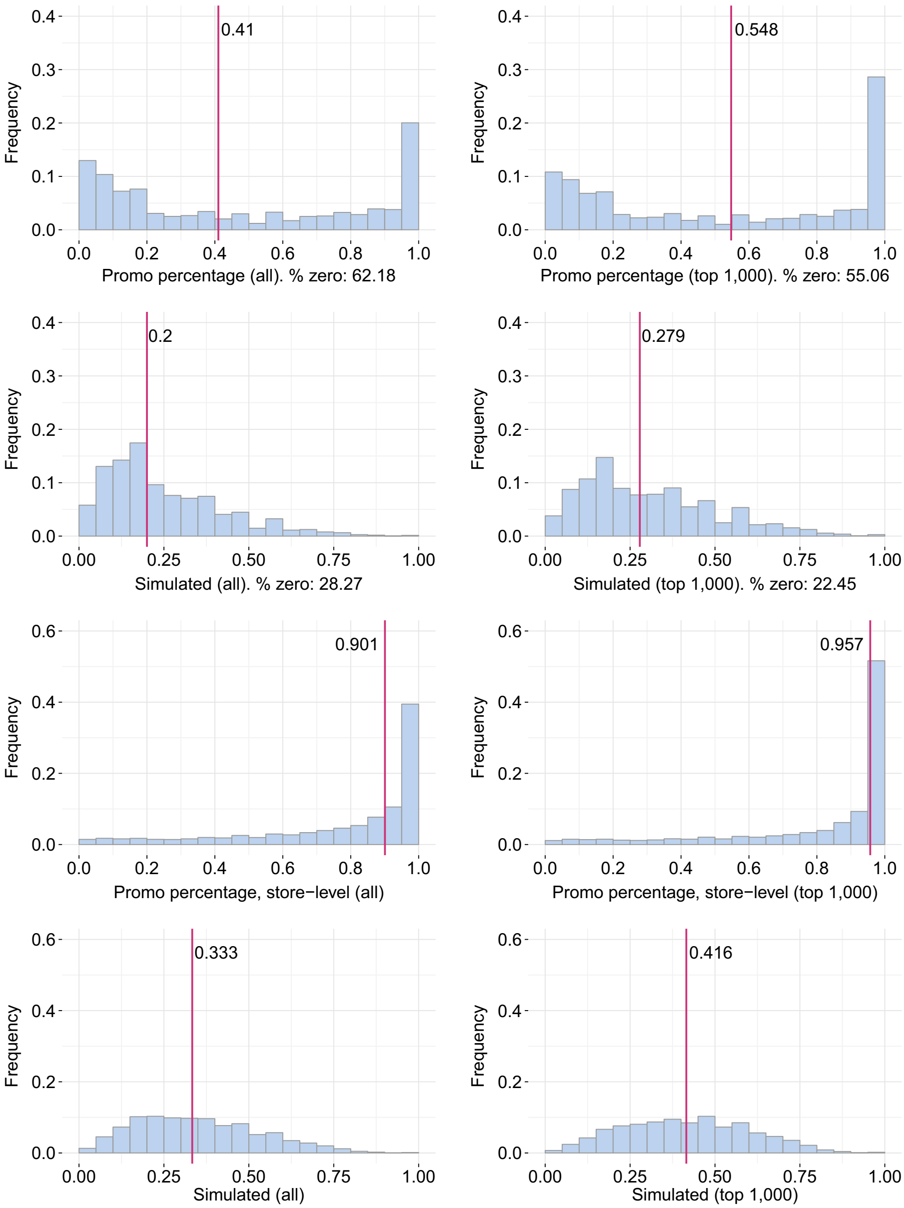

Fig. 22 Distribution of chain/DMA promotion percentages

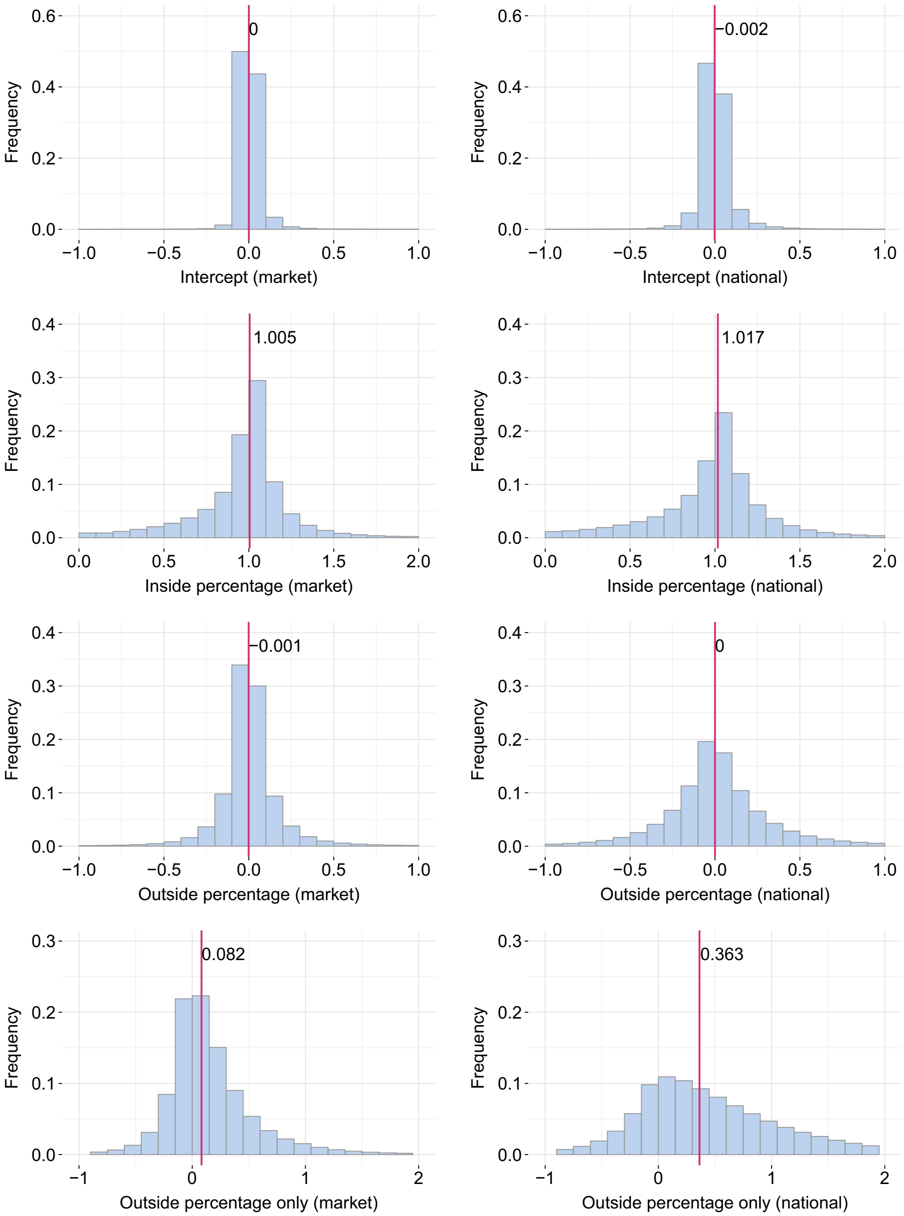

Fig. 23 Promotion coordination regression results: Promotion incidence and inside and outside promotion percentages

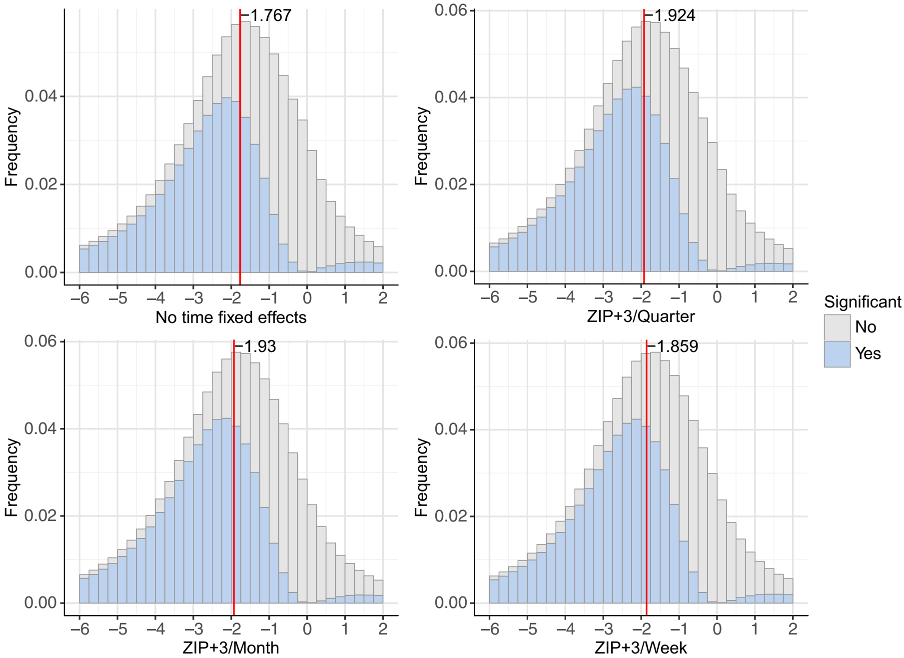

Fig. 24 Own-price elasticity estimates for different fixed effects definitions

References

ACNielsen (2005). *Consumer-Centric Category management: How to Increase Profits by Managing Categories Based on Consumer Needs*. Wiley.

Adams, B., & Williams, K. R. (2019). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics*, *11*, 124–156.

Arcidiacono, P., Ellickson, P. B., Mela, C. F., & Singleton, J. D. (2020). The competitive effects of entry: Evidence from supercenter expansion. *American Economic Journal: Applied Economics*, *12*, 175–206.

Ater, I., & Rigbi, O. (2020). *Price Transparency, Media and Informative Advertising*. manuscript.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, *63*, 841–890.

Berry, S. T. (1994). Estimating Discrete-Choice models of product differentiation. *Rand Journal of Economics*, *25*, 242–262.

Bijmolt, T. H. A., van Heerde, H. J., & Pieters, R. G. M. (2005). New empirical generalizations on the determinants of price elasticity. *Journal of Marketing Research*, *42*, 141–156.

Boatwright, P., Dhar, S., & Rossi, P. E. (2004). The role of retail competition, demographics and account retail strategy as drivers of promotional sensitivity. *Quantitative Marketing and Economics*, *2*, 169–190.

Bolton, R. N. (1989). The relationship between market characteristics and promotional price elasticities. *Marketing Science*, *8*, 153–169.

Bronnenberg, B. J., Mela, C. F., & Boulding, W. (2006). The periodicity of pricing. *Journal of Marketing Research*, *43*, 477–493.

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

Dobson, P. W., & Waterson, M. (2005). Chain-Store Pricing across local markets. *Journal of Economics and Management Strategy*, *14*, 93–119.

Dubois, P., & Perrone, H. (2015). *Price Dispersion and Informational Frictions: Evidence from Supermarket Purchases*. manuscript.

Eden, B. (2014). *Price Dispersion and Demand Uncertainty: Evidence from US Scanner Data*. manuscript.

Eizenberg, A., Lach, S., & Oren-Yiftach, M. (2021). Retail prices in a city. *American Economic Journal: Economic Policy*, *13*, 175–206.

Ellickson, P. B., & Misra, S. (2008). Supermarket pricing strategies. *Marketing Science*, *27*, 811–828.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning*, second ed. Berlin: Springer.

Hoch, S. J., Kim, B.-D., Montgomery, A. L., & Rossi, P. E. (1995). Determinants of Store-Level price elasticity. *Journal of Marketing Research*, *32*, 17–29.

Honka, E., Hortaçsu, A., & Wildenbeest, M. (2019). Empirical search and consideration sets. In J.-P. Dubé, & P. E. Rossi (Eds.) *Handbook of the Economics of Marketing, chap. 4*, (Vol. 1 pp. 193–257). B.V. Elsevier.

Hwang, M., Bronnenberg, B. J., & Thomadsen, R. (2010). An empirical analysis of assortment similarities across U.S. *Supermarkets. Marketing Science*, *29*, 858–879.

Joo, J. (2020). *Rational Inattention as an Empirical Framework: Application to the Welfare Effects of New Product Introduction*. manuscript.

Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, *76*, 728–741.

Kaplan, G., & Menzio, G. (2015). The morphology of price dispersion. *International Economic Review*, *56*, 1165–1205.

Kaplan, G., Menzio, G., Rudanko, L., & Trachter, N. (2019). Relative price dispersion: Theory and evidence. *American Economic Journal: Microeconomics*, *11*, 68–124.

Lach, S. (2002). Existence and persistence of price dispersion: an empirical analysis. *The Review of Economics and Statistics*, *84*, 433–444.

Nakamura, E. (2008). Pass-Through in Retail and Wholesale. NBER Working Paper 13965.

Nakamura, E., & Steinsson, J. (2008). Five facts about prices: a reevaluation of menu cost models. *Quarterly Journal of Economics*, *123*, 1415–1464.

Nakamura, E. (2013). Price rigidity: Microeconomic evidence and macroeconomic implications. *Annual Review of Economics*, *5*, 133–163.

Narasimhan, C., Neslin, S. A., & Sen, S. K. (1996). Promotional elasticities and category characteristics. *Journal of Marketing*, *60*, 17–30.

Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, *69*, 307–342.

Rossi, P., Allenby, G., & McCulloch, R. (2005). *Bayesian Statistics and Marketing*. Wiley.

Rossi, P. E. (2014). Even the rich can make themselves poor: The dangers of IV and related methods in marketing applications. *Marketing Science*, *33*, 655–672.

Sorensen, A. T. (2000). Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy*, *108*, 833–850.

Tellis, G. J. (1988). The price elasticity of selective demand: a Meta-Analysis of econometric models of sales. *Journal of Marketing Research*, *25*, 331–341.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Günter J. Hitsch[1] · Ali Hortaçsu[2] · Xiliang Lin[3]

Ali Hortaçsu
hortacsu@uchicago.edu

Xiliang Lin
xilianglin@gmail.com

1 University of Chicago Booth School of Business, Chicago, IL, USA

2 University of Chicago and NBER, Chicago, IL, USA

3 JD.com, Beijing, China