Retail Pricing and Organizational Structure *

Semyon Tabanakov
University of Chicago

Ali Goli
University of Rochester

Pradeep K. Chintagunta
University of Chicago

Abstract

We document the central role of organizational structure in determining pricing strategies in the U.S. food retail industry. Large parent companies that own multiple chains employ zone pricing aligned with their organizational divisions, reflecting delegation of pricing authority from central to divisional management. While chains typically belong to a single zone, explaining the uniform chain-level pricing commonly found in previous studies, parent-level pricing is more complex. Multiple chains can operate under coordinated pricing within a single zone, and some chains can span multiple zones with different prices. Price zones generally cover separate geographic areas, a design that mitigates coordination frictions under delegated pricing by preventing divisional managers from competing with one another within the same local market. When zones do overlap, they coincide with "discounter" and "traditional" chains operating at different price points, where limited direct competition softens such coordination frictions. Price zones are deeply rooted in organizational divisions and costly to adjust, as evidenced by a two-year restructuring effort following a large merger. Yet, parent companies can exhibit flexibility when required: the enactment of soda taxes demonstrates their ability to implement granular within-zone price adjustments in response to local economic shocks.

JEL Classification: L81, L11, L22, M31

*Email: stabanak@chicagobooth.edu, agoli@ur.rochester.edu, and Pradeep.Chintagunta@chicagobooth.edu.
We thank Jean-Pierre Dubé, Paul Ellickson, Camille Fredrickson, Matthew Gentzkow, Avi Goldfarb, Günter Hitsch, Ali Hortaçsu, Mitchell Lovett, Nadya Malenko, Olivia Natan, Elena Prager, David Reiley, Avner Strulov-Shlain, Chad Syverson, Joel Waldfogel, Thomas Wollman, and Tianli Xia for their helpful comments. We thank seminar participants at UC Riverside, Columbia, Nanyang Technological University, The University of Hong Kong, Dartmouth College, and conference participants at the International Industrial Organization conference and the Marketing Science conference. Researcher(s)' own analyses were 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 researcher(s) 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. Chintagunta thanks the Kilts Center for Marketing at the University of Chicago for financial support.

1 Introduction

Market concentration has increased in more than 75 percent of U.S. industries over the last three decades (Grullon et al., 2019). Such consolidation has enlarged firms, making them internally more complex. While theoretical microeconomists have long recognized the importance of the allocation of authority and decision rights within large organizations (Dessein, 2002), applied work in industrial organization and public policy often abstracts from these issues, treating the firm as a single cohesive entity making optimal choices. Yet if internal organization matters for firm behavior, we should observe its imprint in market outcomes like pricing decisions.

We study pricing in the U.S. food retail industry, a sector that accounts for a large share of household expenditures. Industry consolidation has created parent companies that manage portfolios of retail chains which may appear independent but in reality share common ownership. Recent research has documented that stores within the same chain typically charge uniform prices (DellaVigna and Gentzkow, 2019; Hitsch et al., 2021). However, because previous studies analyze pricing at the chain level, a fundamental question remains unanswered: at what level do prices actually vary when we account for common ownership across chains?

To answer this question, we use NielsenIQ Retail Measurement Service data covering approximately 54% of total U.S. food retail expenditures. We first infer common ownership by identifying chains that share similar private label assortments. We then apply a clustering algorithm to weekly prices of the highest-revenue products to uncover the structure of price variation within each parent company. This algorithm groups stores with similar pricing trajectories, relying exclusively on price data without incorporating information about geographical location, chain identity, or organizational structure.

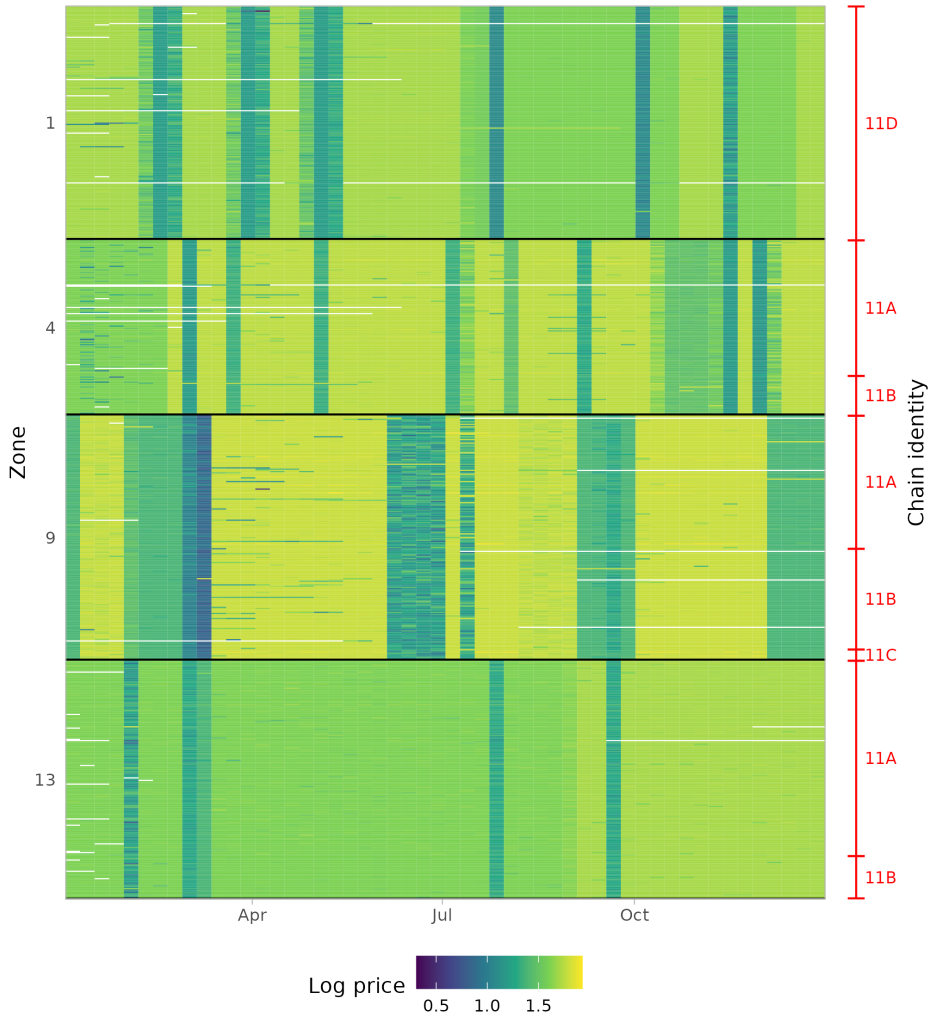
Our analysis reveals systematic price zones within parent companies. Using the fit of the clustering procedure, we classify parent companies into those practicing uniform, zone, or store-level pricing strategies. Although the majority of parents employ uniform pricing, these parent companies own only 33% of stores. In contrast, 61% of stores are operated by parents that use zone pricing. While chains typically belong to a single zone—explaining the uniform chain-level pricing documented in previous studies—multiple chains can operate under coordinated pricing within a

single zone, and some chains span multiple zones with different prices.

To illustrate the outcome of the clustering procedure and the relationship between price zones and chains, Figure 1 presents a heatmap of prices for a single cereal product sold by a large parent company ("parent 11"). Each pixel represents the price at a given store (rows) in a given week (columns), with darker colors indicating lower prices. The numbers on the left axis denote the stores' price zones, while the ticks on the right axis indicate chain identities (denoted 11A–11D). The parent has thirteen zones in total, and we select four to illustrate our key points. First, while prices are nearly uniform within each zone, they vary substantially across zones, validating our clustering approach. Second, the relationship between chains and price zones is complex. In some cases, a chain belongs entirely to a single price zone, as shown by chain 11D, which belongs exclusively to price zone 1. This represents an instance where uniform pricing at the chain level arises as a consequence of zone pricing at the parent level. However, price zones and chains do not exhibit a one-to-one mapping. Chains 11A and 11B operate in multiple price zones, indicating that zone pricing exists even at the chain level. Additionally, in all three price zones, multiple chains are present, resulting in uniform pricing *across* chains with common ownership. This coordination suggests that understanding retail pricing requires an analysis at the parent-company level rather than treating individual chains as independent units.

Although we demonstrate that most zone-pricing parents rely on geographically contiguous zones, we also observe some exceptions to this pattern. One large parent company, along with a few others, takes a different approach in certain regions where they operate multiple chains. Rather than assigning all these stores to a single price zone, they assign a separate zone to each chain. This occurs when one chain represents "traditional" stores and the other represents "discounter" stores offering lower prices. This pattern reveals how parent companies balance operational efficiency with market positioning. They tend to unify pricing when chains have similar formats but maintain separate pricing zones when formats differ significantly. Thus, analyzing pricing at the parent level allows us to document a novel instance of second-degree price discrimination, where the differentiation occurs across retail chains under common ownership. Such chains are vertically differentiated, maintaining distinct price points to attract different customer segments.

Figure 1: Heatmap of prices for a popular cereal product of parent 11 during the year 2019



Notes: Rows are stores, columns are weeks. Stores are ordered by their price zones indicated to the left and by horizontal black lines. Within the zone, stores are ordered by their chain identity, indicated by the red ticks on the right. The Figure displays four representative price zones out of thirteen total zones.

Do these price zones reflect formal organizational structure? For two large parents, which together account for more than a quarter of stores in our sample, we obtained independent data on their geographic organizational divisions. Remarkably, the price zones recovered purely from our price-based clustering algorithm closely align with these organizational divisions. This alignment suggests that price zones reflect the delegation of pricing authority from central to divisional management, with organizational structure directly shaping market conduct.

The fact that organizational divisions are endowed with pricing authority explains why their territories are typically contiguous and non-overlapping. When pricing authority is delegated to divisional managers, the parent can establish clear and separate geographic domains to prevent internal competition. If two divisions were to operate in the same local market, each would optimize for its own profits without internalizing the impact of its pricing on the other, generating price competition that erodes overall parent profitability. Geographic contiguity minimizes these inter-divisional externalities and limits profit cannibalization within the parent company. Overlaps could arise when divisions oversee chains with distinct formats such as "traditional" and "discounter" stores that target separate customer segments and therefore do not compete directly. Thus, inter-divisional externalities provide a natural explanation for the geographic contiguity of price zones, along with the exceptions we observe in the data.

Interviews with practitioners reveal that organizational divisions handle not only pricing but also advertising and workforce management. Tying pricing to these other functions makes price zones costly to change. This rigidity is evident in our case study of a large merger between two zone-pricing parents. By applying our clustering algorithm to each year of data, we track how the merged company's price zone structure evolved over time. Initially, the company inherited the price zones of the original parents, leading to geographical inconsistencies, with multiple price zones existing in the same markets. Over the following two years, these inconsistencies were resolved, resulting in a new geographically contiguous zone structure. Although we lack data on the corresponding evolution of organizational divisions, the slow adjustment of price zones is consistent with the high costs of overhauling the firm's internal organizational structure.

While prices are generally uniform within the zones we uncover, an important question is

whether some product categories can be priced at a more granular level. Put differently, can an organizational division respond to a local shock by adjusting prices only in a subset of its stores? Using the enactment of soda taxes as an example of a local cost shock, we find that one large parent increased prices only in the affected stores within the same price zone, demonstrating a targeted response to the tax. In contrast, another parent exhibited little divergence in prices between taxed and non-taxed stores, suggesting that they chose not to incorporate the tax into shelf price and instead applied it at the checkout. This case study shows that while overhauling the entire organizational structure may be costly, divisions can fine-tune prices for specific product categories to respond to local cost shocks.

Finally, we test whether divisions set prices in a manner consistent with profit-maximizing responses to local demand conditions. Using an isoelastic demand model, DellaVigna and Gentzkow (2019) demonstrate that optimal store-level pricing implies a coefficient of 1 in a regression of stores' prices on the term $\log\left(\frac{\eta_s}{1+\eta_s}\right)$, where η_s represents the demand elasticity. The authors find a coefficient of 0.09 (std. err. 0.02), which is far from the profit-maximizing benchmark. Focusing exclusively on zone-pricing parents and regressing zone-level prices on zone-level elasticities, we estimate a substantially larger coefficient of 0.71 (std. err. 0.054). This suggests that zone-pricing parents adjust prices in response to local demand elasticities, though not to the full extent implied by profit maximization. We find that zone pricing increases profits by 0.26%–6.81% and consumer surplus by 0.15%–4.88% relative to uniform pricing, with the modest profit gains reflecting organizational constraints on zone design rather than profit-maximizing zone boundaries.

In sum, our study demonstrates the central role of organizational structure in retail pricing. Price discrimination is implemented through delegating authority over pricing decisions from central to divisional management, thus turning a division into a price zone. These zones enable parent companies to respond to local market conditions, yet they are costly to modify due to their deep organizational roots. At the same time, divisions retain the flexibility to adjust prices more granularly when economic stakes are high. Our findings demonstrate that organizational structure has a direct effect on the firm's market behavior.

In what follows, Section 2 reviews related literature, Section 3 describes the data, and Section

4 outlines the methodology for identifying price zones. Section 5 classifies parent companies by pricing strategies. Sections 6 and 7 provide case studies examining the flexibility of zone pricing in response to global and local shocks, respectively. Section 8 evaluates the optimality and welfare implications of zone pricing. Section 9 concludes.

2 Related Literature

Our paper bridges and extends two distinct streams of research: organizational economics and retail pricing strategies, connecting internal firm design to observed market conduct. Within organizational economics, a central theme is the trade-off between centralized coordination and decentralized adaptation in multidivisional firms, where decision-making authority must be allocated across organizational units facing heterogeneous environments (Alonso et al., 2008; Rantakari, 2008). Empirical evidence on whether firms delegate economically significant decisions to local divisions remains limited, with most studies relying on survey-based measures of decentralization (Bloom et al., 2012; Dessein et al., 2022). A notable exception is Hortaçsu et al. (2023), who show that prices at a major airline emerge from a vertical structure in which departments pass inputs sequentially to a central pricing algorithm. In contrast, we study a horizontal structure: retailers divide operations by geography and delegate price-setting authority to local divisions.

Within the retail pricing literature, prior work has documented various regimes, including store-level pricing (Ellickson and Misra, 2008), zone pricing (Chintagunta et al., 2003; Adams and Williams, 2019; Eizenberg et al., 2021), and most recently, uniform pricing at the chain level (DellaVigna and Gentzkow, 2019; Hitsch et al., 2021). We contribute to this literature by providing systematic documentation of the prevalence of each pricing strategy across a broad set of parent companies in the U.S. food retail sector, showing that zone pricing is by far the dominant approach. The prevalence of uniform pricing documented at the chain level in prior work emerges because researchers have analyzed chains as independent entities rather than recognizing their common ownership. When chains operate within single parent-determined price zones, they naturally exhibit uniform pricing. In addition, we uncover a novel pattern of coordinated pricing across chains under common ownership: parents charge uniform prices across similar chains and

systematically differentiate prices across others, consistent with second-degree price discrimination by chain format.

3 Data

Our main data source is the NielsenIQ Retail Measurement Service (RMS) for the year 2019, which provides weekly sales and price information for each Universal Product Code (UPC) sold at various undisclosed retail establishments, covering approximately 54% of total U.S. food retail expenditures. Each record details the weekly sales and price of a specific UPC at an individual store, accompanied by a retailer code. These retailer codes allow us to group stores under their respective retail chains. In the following subsections, we first describe how we select stores and infer the parent of each chain. We then describe the selection process for the set of UPCs used in our analysis.

3.1 Store and Chain Selection

From the 114 food chains in the data, we select the 80 with the largest number of stores. These chains operate 11,089 stores, accounting for 98% of all food stores in the dataset. Multiple selected chains may belong to the same parent company. To the best of our knowledge, prior literature using NielsenIQ RMS data has relied on the field *parent_code* to identify a chain's ultimate parent. However, this field may not reflect true ownership when "*the company decentralizes their data release to NielsenIQ through multiple data centers representing each of their retail banners*" (Nielsen, 2021). As a result, chains with different *parent_code* values may still share common ownership.

To identify such cases, we group chains based on the similarity of their private label assortments across multiple product categories. Appendix A.1 provides a detailed description of this identification algorithm. Throughout the paper, we refer to each such group of chains as a "parent". We further exclude two chains whose stores appear to be operated by more than one parent, as their ownership cannot be clearly established. In the final sample, each parent owns one or more chains, and each chain belongs to a single parent. This yields 43 parents and 10,956 stores. Thus, the number of independent units of analysis falls from 80 chains to 43 parents once chains with common

ownership are combined. We refer to parents by randomly assigned numbers from 1 to 43.

Table 1: Summary statistics of chains and parents

	Mean	25th	Median	75th
Panel A: Chains				
Number of stores	140.5	34	72.5	157.2
Number of counties	42.1	10	20.5	41
Number of states	4.6	1	3	5
Panel B: Parents				
Number of stores	254.8	36.5	101	215.5
Number of counties	72.2	14	32	74
Number of states	7	2	3	6

Notes: Chain is defined by the *retailer_code*. Chain's parent is inferred as described in Section 3.1. Panel A has summary statistics at the chain level, and Panel B at the parent level.

Summary statistics are presented in Table 1. Panels A and B show the data across chains and parents, respectively. The average chain has 140 stores and is present in 42 counties. Accounting for the ownership structure increases the unit's size, with the average parent having 254 stores and being present in 72 counties. Median values do not change as drastically because only 8 out of 43 parents own more than one chain.

3.2 Products

We choose 40 product modules used in DellaVigna and Gentzkow (2019), to be consistent with previous work. As our analysis is conducted at the parent level, the set of analyzed products varies across parents. For each parent and product module, we select the 20 UPCs with the highest revenue to ensure that we include the most important products for the parent and reduce the potential for missing observations. Consequently, the maximum number of products considered for each parent is $40 \times 20 = 800$. However, since a parent may carry fewer than 20 UPCs for some categories, the final number of products could be smaller. Across parents, the average number of UPCs considered is 795, with a standard deviation of 14. We treat products of different pack sizes or flavors as separate rather than combining them into one, to avoid aggregation errors that could affect the accurate recovery of price zones.

3.3 Demographic and Competition Variables

To test whether demographic and competitive factors explain the zone boundaries we uncover, we collect store-level demographic and competition data. Demographic data comes from the 2015-2019 American Community Survey. When selecting the demographic variables likely to impact demand, we primarily rely on prior work by Hoch et al. (1995). We include income, house values, household size, the share of people with a bachelor's degree, the share of elderly people, and race.

To construct measures of local competition for each store, we separately count the number of grocery and mass merchandiser stores within a 5-kilometer radius. Similar measures were used by DellaVigna and Gentzkow (2019), and for a different product category by Ozturk et al. (2016). Other competition variables include the average distance to the nearest 10 grocery stores and 3 mass merchandiser stores. While NielsenIQ data might seem suitable for constructing these competition measures, it captures only a subset of U.S. stores. Therefore, we use the National Neighborhood Data Archive (NaNDA) (Finlay et al., 2020), which provides counts of grocery and dollar stores by NAICS codes from the National Establishment Time Series (NETS) database.[1]

3.4 Organizational divisions

For two large parent companies, parents 11 and 26, which together account for over a quarter of the stores in our dataset, we obtained external data that maps the geographic boundaries of their organizational divisions for the years 2015 and 2019, respectively. Using these maps, we manually matched the stores of these two parents in the RMS data to their corresponding divisions. This dataset is used only in Section 5.4 to interpret the results. It is not used in the clustering procedure or in the classification of parents by pricing strategies.

[1]We augment the data by manually adding big-box retailers that have different NAICS codes and therefore do not appear in the NaNDA data.

4 Identifying Zones

4.1 Clustering Description

To identify price zones, we cluster stores separately for each parent using the Partitioning Around Medoids (PAM) algorithm (Kaufman and Rousseeuw, 1990). The feature vector for store s consists of log prices p_{sjt}, where j indexes parent-specific products and t indexes weeks. The resulting feature vector includes over 40,000 elements, constructed from prices for 40 product modules, each with 20 UPCs, across 52 weeks. We chose PAM over alternatives like K-means because it scales better with the high dimensionality of our feature vectors and is more robust to missing data, noise, and outliers. The algorithm takes a matrix of pairwise dissimilarities as input, for which we use the ℓ_1 norm. It then iteratively selects representative stores (medoids) and assigns each store to its closest medoid, minimizing the total sum of dissimilarities within clusters. The output is a partition of stores into clusters, which we interpret as price zones. In this procedure, we only use price data and do not restrict zones to be geographically contiguous in any way.

4.2 Determining the Optimal Number of Clusters

To determine the optimal number of clusters, we primarily employ the Silhouette Coefficient (Rousseeuw, 1987), a standard metric that captures how well the clustering separates observations into coherent groups. More formally, it measures how well-matched each store is to its assigned cluster relative to the nearest neighboring cluster. For each observation i in cluster A, the Silhouette value $s(i)$ is defined as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}},$$

where $a(i)$ is the average dissimilarity of i to all other objects in A, and $b(i)$ is the average dissimilarity of i to all objects in its nearest neighboring cluster B. This value ranges from -1 to 1, with higher values indicating better clustering fit. The overall Silhouette Coefficient $SC(N)$ for N clusters is the average $s(i)$ across all observations:

$$SC(N) = \frac{1}{n} \sum_{i=1}^{n} s(i).$$

Choosing the number of clusters relying only on the Silhouette Coefficient would mean choosing N that maximizes $SC(N)$.

To complement the Silhouette Coefficient, we use t-SNE (van der Maaten and Hinton, 2008), a dimensionality reduction technique that transforms high-dimensional data into a two-dimensional representation while preserving local similarity relationships. Unlike methods such as Principal Component Analysis that focus on preserving global structure, t-SNE emphasizes maintaining the distances between nearby points, making it particularly effective for visualizing cluster structures in high-dimensional data. We apply t-SNE to the same dissimilarity matrix used for PAM clustering and visualize the results as scatterplots where each point represents a store.

This visualization approach is especially valuable when the Silhouette Coefficient fails to clearly identify the optimal number of clusters. Specifically, the Silhouette Coefficient can fail when dealing with hierarchical cluster structures, such as when there are two well-separated major clusters, each containing multiple meaningful subclusters. For example, this can occur when a parent operates both traditional and discounter format stores across several geographic regions. As demonstrated in our simulation in Appendix A.2, this produces multiple peaks in the Silhouette Coefficient: an initially high value capturing only the major separation, followed by a decrease, and then an increase once the clustering algorithm begins to identify individual subclusters.

For most parents, we select the optimal number of clusters using the Silhouette Coefficient, which generally aligns with the visually distinct clusters identified by t-SNE. However, for parents 26 and 22, t-SNE clearly reveals multiple distinct clusters that differ markedly from what the Silhouette Coefficient would suggest. In both of these cases, when we run PAM with the number of clusters suggested by the visible separations in the t-SNE plot (rather than the number that maximizes the Silhouette Coefficient), the resulting PAM clusters align well with the visually distinct groups of points in the t-SNE visualization. Thus, for these two parents, we choose the number of clusters based on t-SNE. In what follows, we denote N_r as the chosen number of zones for the parent r.

4.3 Clustering Fit

To categorize parents as using uniform, zone, or store-level pricing strategies, we quantify how well the zone structure explains variation in product prices. Fix a parent company and the number of clusters (zones). Let $\mathcal{S}_{zjt}$ denote the set of stores in the zone z that sell product j in week t, and $M_{zjt} \equiv |\mathcal{S}_{zjt}|$ the number of such stores. With p_{sjt} denoting the log price in store s as before, define the following dispersion measure:

$$\sigma_{zjt} \equiv \frac{1}{M_{zjt}^2} \sum_{s \in \mathcal{S}_{zjt}} \sum_{s' \in \mathcal{S}_{zjt}} |p_{sjt} - p_{s'jt}|.$$

This statistic measures the mean absolute difference in log prices between two randomly chosen stores with replacement. For example, $\sigma_{zjt} = 0.05$ implies that, within week t, two stores randomly drawn from zone z will differ on average in price for product j by approximately 5%. We adopt this measure for its intuitive interpretation and to be consistent with DellaVigna and Gentzkow (2019).

Next, we aggregate σ_{zjt} across zones, products, and weeks to get a single measure of price dispersion for the given parent r and number of zones N:

$$V_r(N) \equiv \frac{\sum_{j,t,z \in \{1,\cdots,N\}} M_{zjt} \cdot \sigma_{zjt}}{\sum_{j,t,z \in \{1,\cdots,N\}} M_{zjt}}.$$

This is a simple weighted average of the similarity measure σ_{zjt} with weights equal to the number of stores in each zone-week-product bin. The interpretation of $V_r(N)$ is as follows: take a random week, a random product, and a store selling that product. Then $V_r(N)$ is the expected absolute log price difference between that store and another store from the same zone, which approximately equals a percentage price difference.

As the number of clusters increases, stores within the clusters will have more similar prices resulting in the decrease of V_r. However, even if prices are perfectly uniform within a zone, V_r can be larger than 0. NielsenIQ reports weekly average prices, calculated as total revenue divided by total quantity sold. For instance, if a store sells q_h units at the regular price p_h, and q_l units at a discount p_l, the recorded price is a weighted average: $p_l \frac{q_l}{q_l+q_h} + p_h \frac{q_h}{q_l+q_h}$. This can make two stores with identical shelf prices appear to have different prices in the data when their realized quantity

weights differ.

4.4 Clustering Outcomes for Two Large Parents

To illustrate our clustering procedure and its implications, we present detailed clustering outcomes for two large parent companies (accounting for over a quarter of stores in our data), denoted as parent 11 and parent 26, in Figure 2.

Panel A plots the Silhouette Coefficient against the number of clusters for each parent. For parent 11, the coefficient is maximized when the number of clusters is 13, which we select as the optimal number. For parent 26, however, the Silhouette Coefficient has multiple peaks: it peaks at 2 clusters, decreases until 4 clusters, and then increases again before dropping sharply beyond 23 clusters. This pattern suggests a hierarchical structure where two major clusters contain multiple well-defined subclusters, precisely the scenario where t-SNE visualization becomes valuable.

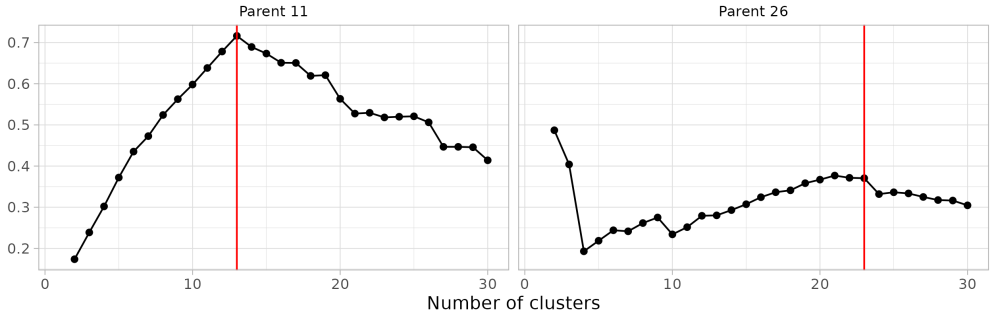
To resolve this ambiguity for parent 26, we examine the t-SNE visualization in Panel B. Each point in this visualization represents a store, with coordinates determined by t-SNE applied to the same dissimilarity matrix used in PAM clustering. The color of each point indicates the store's cluster assignment from PAM clustering, not from t-SNE itself, making any visual alignment between colors and spatial groupings a meaningful validation of our clustering approach. For parent 11 with 13 clusters, we observe strong agreement between t-SNE and PAM, as points with the same color form visually distinct groups. This alignment confirms that 13 is an appropriate number of clusters for this parent.

For parent 26, however, t-SNE reveals many more distinct clusters than the two suggested by maximizing the Silhouette Coefficient. When we set the number of clusters to 23, just before the sharp drop in the Silhouette Coefficient, we observe a close alignment between these clusters and the natural groupings revealed by t-SNE. This visual evidence suggests that 23 is a more appropriate number of clusters for capturing the true pricing zone structure of parent 26.[2] Similar reasoning applies to parent 22, where we adjust the number of clusters from 17 to 12 based on t-SNE visualizations (see Figure A.4 for its clustering results).

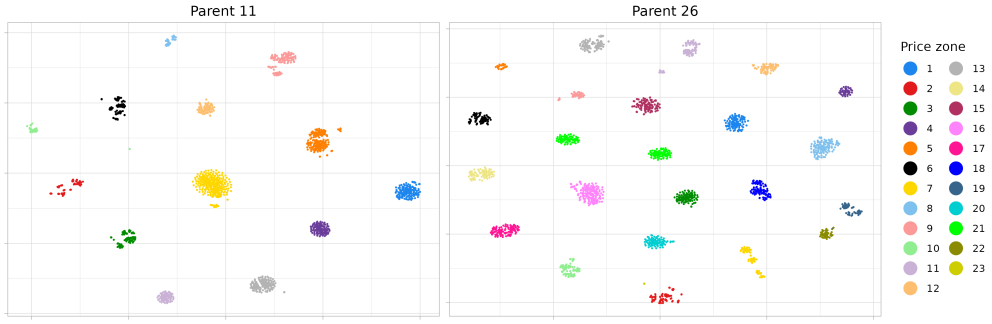
[2]We further validate our selection of 23 clusters for parent 26 through a temporal consistency check, see Appendix A.3

Figure 2: Clustering fit for two large parents

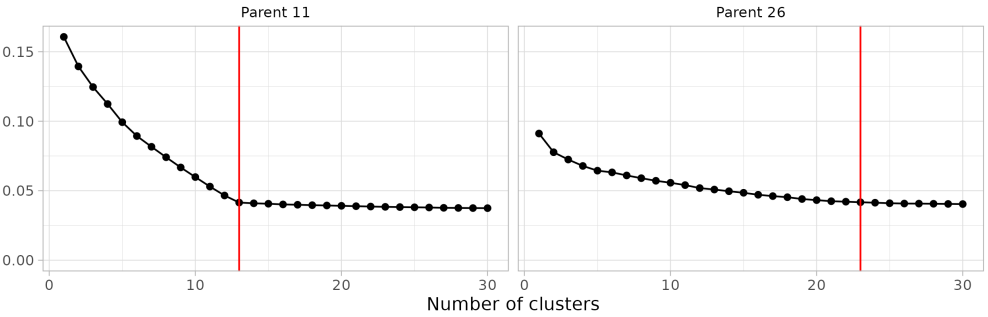
Panel A: Silhouette Coefficient



Panel B: t-SNE dimensionality reduction



Panel C: Mean absolute log price difference



Notes: Panel A shows Silhouette Coefficients for different numbers of clusters, the red line indicates the chosen number of clusters. Panel B displays the results of t-SNE, which reduces each store's high-dimensional price vector to two coordinates for visualization. Colors represent the store's cluster when running PAM on the original price data, with 13 clusters for parent 11 and 23 clusters for parent 26. The close alignment between visually separated clusters in the t-SNE plot and the PAM-based colorings suggests that two distinct methods capture the same underlying structure, indicating that the clustering is robust. Panel C plots $V_r(N)$, the absolute log price difference between two random stores within a zone, as a function of the number of zones N.

While the Silhouette Coefficients and t-SNE outputs help determine the appropriate number of clusters, they do not illustrate the economic significance of zone pricing. Well-defined clusters could still have negligible price differences across zones. Panel C addresses this concern by plotting $V_r(N)$ against different numbers of clusters N.

The results show substantial price variation reduction. Note the high values of $V_r(1)$ for both parents: 0.16 for parent 11 and 0.09 for parent 26. At the optimal number of clusters, these values drop to 0.04 for both parents. In words, it means that for parent 11 the prices between two random stores differed by approximately 16%, but between two random stores *within the same zone* by 4%. For parent 26, the price differences drop from 9% to 4%. Overall, this panel demonstrates that accounting for zone pricing substantially increases the share of explained price variation for these two large parents.

In the following sections, when presenting results for all parents, we use $V_r(1)$ as a measure of price uniformity, and $V_r(N_r)$ as a measure of how well the zone structure explains the parent's prices. The increase in the explanatory power due to clustering is denoted as $\Delta_r \equiv V_r(1) - V_r(N_r)$. Definitions of all variables introduced in this section are listed in Panel A of Table 2.

Table 2: Variable definitions

Variable	Definition
Panel A: Clustering fit	
N_r	Optimal number of clusters
$V_r(N)$	Explained variation in prices using N clusters
Δ_r	$V_r(1) - V_r(N_r)$. Increase in explained variation from clustering relative to uniform pricing baseline
Panel B: Accuracy of predicting zones using store's ...	
chain	Chain affiliation
geo	Geography - latitude and longitude
chain_geo	Chain affiliation and geography. Only for parents owning more than 1 chain
dem	Demographic variables (see Section 3.1)
comp	Competition variables (see Section 3.1)

Notes: r indexes parents. Variables from Panel B correspond to column names in Table 3

4.5 Explaining Zones Using Observable Characteristics

To systematically analyze what drives zone structures, we develop measures that quantify how well different store characteristics explain the observed price zones. For each group of variables (chain affiliation, geographic location, demographics, and competition), we train a random forest to predict the zone label of each store. Prediction accuracy is measured using the out-of-bag error.[3] For a recent example of using machine learning accuracy as a descriptive tool, see Bertrand and Kamenica (2023).

The first measure, *chain*, reflects how accurately zone membership can be predicted from chain identity alone. If a parent sets different prices across chains but maintains uniform prices within each chain, then chain identity will determine zone assignment, and the prediction accuracy will be close to one. Conversely, if a parent sometimes splits a single chain into multiple price zones, the accuracy will be lower.

The second measure, *geo*, captures how accurately zone membership can be predicted using the store's latitude and longitude. If zones are geographically contiguous, then location will strongly determine zone assignment, and the *geo* score will be close to one. In contrast, if price zones cover overlapping geographical areas, prediction accuracy will fall. For multi-chain parents, we also compute a combined *chain_geo* score using both chain identity and geographic coordinates as predictors, to see whether parents employ both chain-based and geography-based price discrimination. We create two additional scores, *dem* and *comp*, that assess how well demographic and competition variables defined in Section 3.3 explain zones. Definitions of all five scores are reported in Panel B of Table 2.

These metrics provide a systematic framework for identifying which factors influence zone formation across parent companies. In the next section, we use the clustering results together with these metrics to classify parents into distinct pricing strategy groups and examine whether their zones are determined by geography, chains, or other factors.

[3]Out-of-bag (OOB) error provides an accuracy estimate in random forests by evaluating each tree on observations not included in the bootstrap sample used to train that tree. This method yields a reliable measure of out-of-sample prediction accuracy without requiring a separate test set.

5 Results

Having described our data, clustering procedure, and the metrics for evaluating zone structures, we now turn to the main results. For each parent company, we implement the clustering, determine the optimal number of zones, and compute how well these zones explain price variation. We then classify parents into distinct pricing strategy groups and examine what determines the observed zone boundaries.

5.1 Classification of Parents

Table 3 presents the results for each parent company, including their optimal number of clusters (N_r), baseline price uniformity measure ($V_r(1)$), clustering fit ($V_r(N_r)$), and increase in explained variation from clustering (Δ_r). The table also reports how well different store characteristics predict zone membership. These predictors include chain affiliation, geographic location, the combination of chain and geography (for multi-chain parents), local demographics, and competition measures.

Given the substantial heterogeneity in clustering results across parents, we classify them into distinct pricing strategy groups. We define a parent as following uniform pricing if the difference in log prices between two random stores is less than 0.04 ($V_r(1) < 0.04$).[4] While a 4% price difference may seem small at first glance, there are two reasons to consider this cutoff economically meaningful. First, DellaVigna and Gentzkow (2019) report a mean absolute log price difference of 3% across all chains, which is interpreted as evidence of uniform pricing at the chain level. Our 4% threshold is therefore more conservative, erring on the side of classifying pricing as uniform. Second, price changes have disproportionately large effects on retailer profit margins. For example, with a gross margin of 25%, a 4% price increase (holding costs and quantity fixed) raises the gross margin to approximately 28%, a 12% proportional increase. If we assume a typical net margin of 2% (FMI, 2024), the same price increase raises net margin to 5.7%, a 185% relative increase. Thus, a 4% price difference is economically substantial, and when observed, is likely the result of a deliberate pricing strategy rather than noise. Under this classification, 26 parents (60% of the

[4]We report full results in Table 3 so that the reader can examine the underlying measures and apply alternative groupings if desired.

Table 3: Projection of price zones on observables

Parent	N_r	$V_r(1)$	$V_r(N_r)$	Δ_r	*chain*	*geo*	*chain_geo*	*dem*	*comp*
Zone pricing: geography and/or chains									
11	13	16.1	4.1	11.9	57	99	99	59	43
20	3	9.6	5.4	4.2	100	99	100	78	66
22	12	4.1	1.5	2.6	12	95		34	20
23	2	10	5.6	4.4	100	61	100	71	69
26	23	9.1	4.2	4.9	52	91	99	51	36
31	2	5.5	2.6	3	87	78	97	86	83
36	2	5	2.4	2.6	65	100		75	71
42	3	9.8	3.7	6.1	88	74	95	54	50
Zone pricing: idiosyncratic									
28	28	6.1	1.4	4.6	11	62		25	15
30	2	4.6	2.3	2.3	89	95		95	95
43	2	5.6	2.5	3	90	86		88	83
Store-level pricing									
5	30	7.9	6	1.9	12	69		31	23
6	3	4.2	2.9	1.2	55	76		55	38
12	2	5.6	4.3	1.3	51	70		70	89
18	2	5	4.6	0.3	81	84		82	81
34	2	4.1	2.5	1.6	91	88		82	91
38	3	11.4	10	1.4	80	80	80	71	78
Uniform pricing									
1	2	1.1	1	0.1	96	88		100	100
2	8	0.6	0.3	0.3	31	19		31	19
3	2	2.7	2.4	0.3	76	98		87	79
4	3	3.8	3.4	0.3	80	98		84	79
7	3	3.8	3.1	0.7	52	48		48	48
8	5	0.8	0.7	0.1	44	48		48	33
9	3	3.9	2.8	1.1	81	98		84	81
10	2	2.4	2	0.3	51	79		63	49
13	4	2.8	2	0.8	58	76	79	69	45
14	12	3.7	1.6	2.1	28	77		44	45
15	18	2.1	0.8	1.3	8	14		11	8
16	2	1.5	1.2	0.3	96	100		92	92
17	3	2.9	2.4	0.4	67	96		63	63
19	2	1.9	1.8	0.1	95	94		94	94
21	21	3.1	1.5	1.6	9	84		54	39
24	3	1.4	0.4	1	78	96		91	70
25	7	2.4	0.9	1.5	46	96		63	51
27	4	3	1.9	1.1	47	82		52	40
29	3	2.6	2.3	0.4	90	90		90	85
32	6	2	1.4	0.7	75	76		74	69
33	2	3.8	2.9	0.8	68	76		71	69
35	2	3.6	3.2	0.4	75	86		86	93
37	2	3.2	2.5	0.7	94	97		97	90
39	8	2.6	0.6	1.9	23	94		50	38
40	13	2.2	1.6	0.6	17	80		21	15
41	2	3.2	2.9	0.3	69	99		90	79

Notes: The table is best understood by using the first row as an example. Parent 11 has thirteen price zones, the mean price difference between two random stores of approximately 16.1%, between two stores in the same zone of 4.1%, implying a reduction of 11.9 p.p. Zone prediction accuracy is 57% using chain identity, 99% using geography, 99% using both, 59% using demographics, and 43% using competition variables.

sample) employ uniform pricing. However, these parents tend to be smaller, collectively operating 33% of stores in the sample.

For uniform-pricing parents, interpreting the clustering results requires careful consideration. One possibility is that the clusters simply reflect measurement error from price aggregation in our data, making them uninformative. However, since our classification is based on the mean price variation across products, a parent might maintain uniform pricing for most products while employing zone pricing for some products, and still be classified as a uniform pricer. While we suspect measurement error drives clustering for most uniform-pricing parents, there are clear cases of limited zone pricing. Parent 41 provides an illustrative example. Our algorithm identifies two clusters containing 69% and 31% of stores, respectively, with a *geo* score of 0.99 indicating highly contiguous geographic zones. Such clear geographic patterns are unlikely to arise from pure measurement error, suggesting that parent 41 employs zone pricing for some products despite maintaining uniform pricing overall. This example highlights a broader point: our classification, like any such categorization, captures general pricing patterns but may miss product-specific variation. Individual products or even entire categories might deviate from a parent's predominant pricing strategy, a phenomenon we explore in detail through our analysis of soda taxes in Section 7.

Turning to parents that are not classified as uniform pricers, we first identify those that exhibit what we term store-level pricing. For these parents, imposing a zone structure does not significantly reduce price variation. We quantify this as parents with $\Delta_r < 0.02$, meaning that applying the clustering structure reduces average within-zone price differences by less than 2 percentage points. This threshold can be compared to the one used by DellaVigna and Gentzkow (2019), who classify chains as practicing geographic zone pricing if restricting comparisons to within-state store pairs reduces the average quarterly absolute log price difference by at least 0.01. As before, our threshold is more conservative, erring on the side of classifying pricing as store-level rather than zone-level. Under this criterion, 6 parents (14% of the sample) exhibit store-level pricing, accounting for 5% of stores.

The remaining parents are classified as using zone-pricing. A natural question, given the previous literature, is whether such parents construct zones based on their geography. On one hand,

Dominick's provides an example where zones are not fully determined by geography (Chintagunta et al., 2003), with stores of different zones often located very close to each other. On the other hand, Adams and Williams (2019) document clear geographic zones for large oligopolists. A separate question, not previously addressed in the literature, is whether a parent with several chains in a geographical market price discriminates across them. We find that almost all remaining parents construct their zones based on geography, chains, or both. There are a few exceptions — two parents (30, 43) have close to 90% of stores belonging to one zone, rendering the zone prediction exercise using random forest irrelevant due to the small number of stores in another zone. Another exception is parent 28, for which we obtain price zones that are not explained by geography. We label these parents as having idiosyncratic zone pricing. The zones of all other remaining parents are determined by geography, chains, or both. This is evident from the very high *geo* scores for parents with a single chain, and *chain_geo* scores for parents with more than one chain. This group consists of 8 parents (19%) that own 61% of stores. Since zones are defined by geography and chains, stores with different demographic and competitive conditions can belong to the same zone. This overlap leads to zones being poorly explained by demographic and competition variables, as evidenced by low *dem* and *comp* scores.

Even among zone-pricing parents, some price variation remains at the store level, as indicated by values of $V_r(N_r)$ above 0. This within-zone variation may reflect limited pricing discretion by store managers, who could adjust prices in response to factors such as local inventory levels or nearby competitors (Ellickson and Misra, 2008). However, because we cannot distinguish true managerial discretion from measurement error in the price data, we interpret this residual variation as an upper bound on store-level pricing autonomy. For example, for parent 11, this residual accounts for at most $\frac{V_r(N_r)}{V_r(1)} = 25\%$ of the total price variation.

In summary, we find two broad patterns in retail pricing strategies, with size being a key differentiating factor. The majority of parents employ uniform pricing, but these parents are typically smaller retailers accounting for less than a third of all stores in our sample. In contrast, most large parents employ zone pricing, with zones determined by geography, chain affiliation, or both, collectively operating 61% of stores. A further novel contribution is that we systematically

describe how price zones are structured in practice. Prior work has assumed that zones correspond to administrative boundaries such as states, whereas we use a data-driven approach to recover price zones and analyze the determinants of their boundaries. The following sections continue to use parents 11 and 26 as running examples to illustrate how zone pricing can be based on either geography or chain affiliation.

5.2 Discrimination by Geography

We define geographical discrimination as cases where price zones are geographically contiguous. Note that our clustering algorithm did not use any information on the stores' geographical location, relying solely on weekly prices to uncover price zones. Therefore, there can be no mechanical relationship between the discovered price zones and the store locations.

Figure 3a plots the locations of parent 11's stores, colored by price zone. To preserve confidentiality, we plot a linearly transformed version of the original store locations.[5] The Figure demonstrates that the zones are remarkably geographically contiguous, aligning with the *geo* score of 0.99 from Table 3. The *chain* score of 0.57, being less than 1, indicates that stores within a single retail chain can be split into several price zones. To illustrate this fact, we zoom in on the region outlined by the dashed rectangle, and present price zones and chain affiliations in Figures 3b and 3c respectively. These zoomed-in views clearly show that chains 11B and 11C operate in both price zones 7 and 9, demonstrating how individual chains can span multiple price zones.

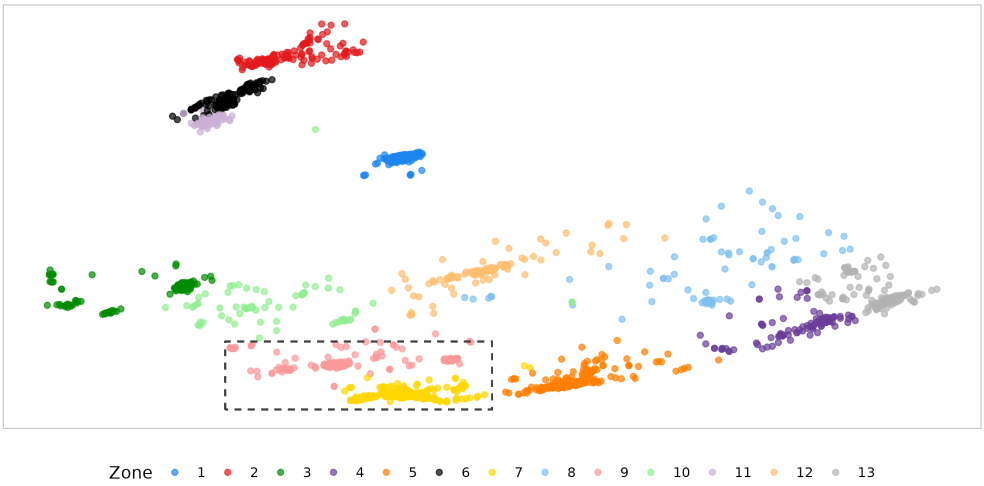
The most interesting consequence of purely geographical discrimination is that different chains owned by the same parent and operating in the same market belong to the same price zone. As Figures 3b and 3c demonstrate, price zone 9 includes stores from chains 11A, 11B, and 11C. Similarly, zone 7 includes stores from chains 11B, 11C, and 11E. Thus, when a parent uses purely geographic price discrimination, in markets where several chains overlap, these chains will be pooled together into a single price zone, resulting in uniform pricing across these chains.

To visualize this finding, Figure 4 plots the weekly prices of a cereal product within zones that

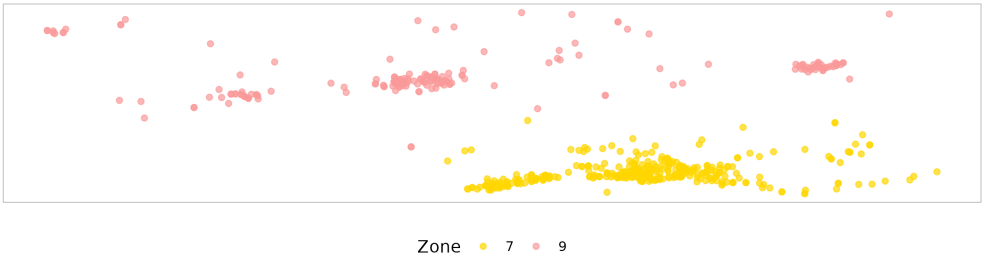
[5]Note that the relative positions of stores are preserved under a linear transformation. That is, if point C is on a line between A and B, it will remain on the line between A and B. Therefore, the relative position of the zones/stores remains the same.

Figure 3: Price zones and chains of parent 11

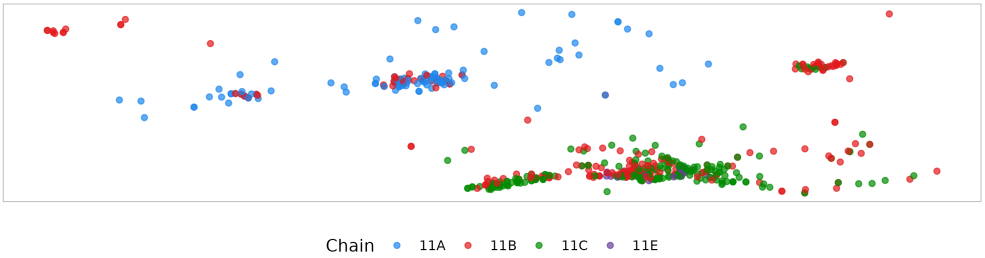
(a) Price zones



(b) Price zones (inset map)



(c) Chains (inset map)



Notes: Panel (a) shows store locations for parent 11, with each dot representing a store and colors indicating its assigned price zone from the clustering algorithm. The map is a linearly transformed version of the United States. The transformation preserves relative positions but distorts absolute coordinates to protect confidentiality. Panel (b) zooms in on the area marked by the dashed rectangle in (a). Panel (c) shows the same zoomed-in region, but colors stores by their retail chain.

have several chains separately for each chain. Apart from zone 2, prices are identical across chains as long as the stores belong to the same price zone. This pattern of identical pricing and synchronized promotions across retail chains could easily be misinterpreted as coordination between independent retailers, possibly due to shared wholesale costs or promotional schedules, or more concerningly, as evidence of collusion. In reality, these similarities result from centralized decision-making by a single parent company across its chains. This underscores the importance of analyzing pricing at the parent rather than the chain level, as failing to account for common ownership risks misattributing these patterns to inter-firm coordination or collusion rather than to decisions by a unified corporate entity.

Figure 4: Prices for a cereal product of parent 11 by zones and chains



Notes: Each panel represents a zone. Zone numbers map to those from Figure 3a. Colors represent different chains within a given price zone. Lines show weekly prices averaged across stores of a given chain in a given zone.

There is an interesting price dynamic in zone 2, where despite belonging to the same price zone, chains have different base prices but correlated promotions. The fact that our clustering algorithm

pools these two chains together into the same zone suggests that such base price differences are small compared to the total price variation. However, for other parents, this type of discrimination across chains located in the same geographical market is more widespread and captured by the clustering algorithm. We explore discrimination by chains in the next section.

5.3 Discrimination by Chains

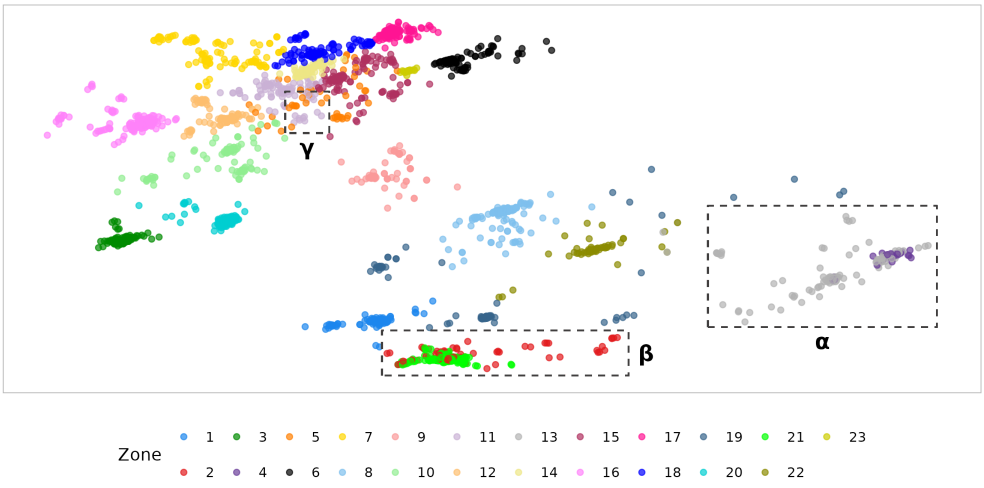
We define chain-based discrimination as cases where the parent company sets different prices across its retail chains operating in the same geographical area. This strategy can be identified by comparing the *chain_geo* and *geo* scores in Table 3; a substantially higher *chain_geo* score indicates that chain affiliation improves explanatory power beyond geography alone. This pattern appears for parents 23, 26, 31, and 42. We examine parent 26 to illustrate this pricing strategy in detail.

Figure 5a maps the stores of parent 26, showing their price zones. While some aspects of parent 26's pricing strategy mirror that of parent 11, there are important differences. Like parent 11, some price zones are geographically contiguous, as is evident from the left part of the Figure. Additionally, there exist chains that are split into multiple price zones, as is evidenced by the low *chain* score of 0.52.

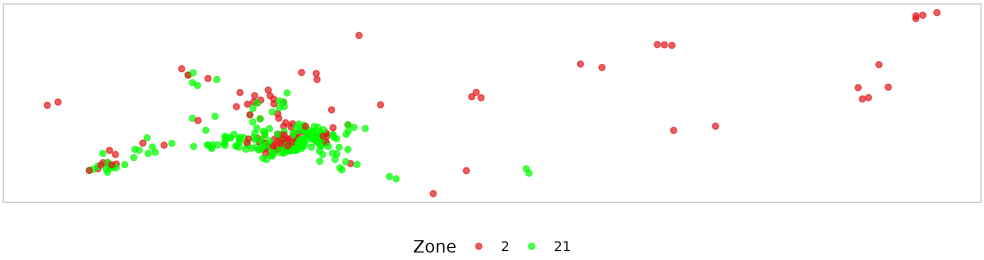
The key difference lies in how parent 26 handles multiple chains operating in the same geographic area. Unlike parent 11, which maintains uniform pricing across chains in the same area, parent 26 may assign different price zones to different chains operating in the same region. This is evident from the *geo* score of 0.91, which, while high, is notably lower than parent 11's value of 0.99. The high *chain_geo* score of 0.99 confirms that price zones are determined by both geography and chain affiliation. The regions where the discrimination by chains takes place (that is, where price zones overlap) are indicated by dashed rectangles and denoted as Markets α, β, and γ. Figures 5b and 5c zoom in on the Market β presenting price zones and chain affiliations, respectively. On the left side of the Figure, we see that chain 26A belongs to the price zone 21, while chain 26B belongs to the price zone 2. Thus, parent 26 is able to price discriminate across different chains with overlapping markets. Importantly, this does not always happen. Parent 26 can still pool different chains into the same price zone, as is evidenced by price zone 2 encompassing both chains 26B and 26C.

Figure 5: Price zones and chains of parent 26

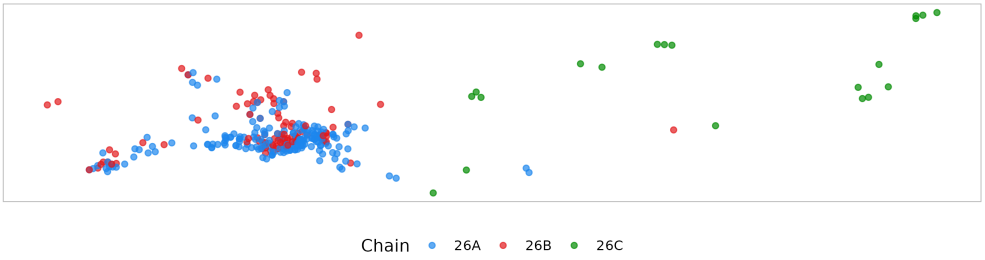
(a) Price zones



(b) Price zones (inset map for Market β)



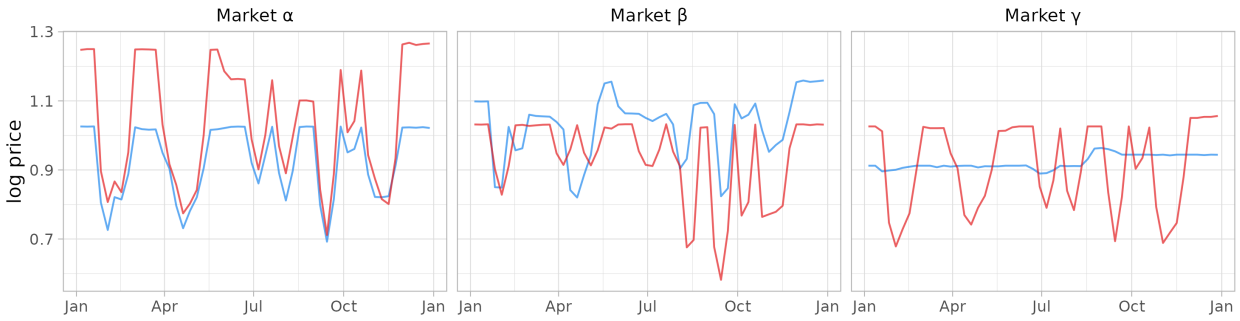
(c) Chains (inset map for Market β)



Notes: Panel (a) shows store locations for parent 26, with each dot representing a store and colors indicating its assigned price zone from the clustering algorithm. The map is a linearly transformed version of the United States. The transformation preserves relative positions but distorts absolute coordinates to protect confidentiality. Panel (b) zooms in on the area marked by the dashed rectangle in (a). Panel (c) shows the same zoomed-in region, but colors stores by their retail chain.

This finding raises an interesting question: why would a parent want to discriminate between closely located stores with different chains? We hypothesize that this occurs when chains cater to different market segments, with one chain acting as a discounter store, and another offering a more traditional shopping experience. An example of such price differentiation is the Canadian parent Loblaws Companies Limited, which operates in Toronto both the traditional chain Loblaws and the discounter chain No Frills. These chains can have different operating costs and attract different types of customers, thereby motivating the parent company to engage in price discrimination. Another natural way to differentiate would be through assortments with discounter chains carrying fewer premium products. While assortment differences are certainly relevant, in this paper our focus is specifically on pricing strategies.

Figure 6: Prices for a cereal product of parent 26 by chains within geographical markets



Notes: Each panel represents a market with overlapping price zones. Colors represent chains that form these zones. Lines show weekly prices averaged across stores of a given chain in a given market.

To visualize the parent 26's discrimination by chains, Figure 6 plots prices for a frozen entree product within three markets where discrimination by chains takes place. Within Market α base prices are different across chains, however, promotions seem to occur during the same weeks. The other two markets show that chains differ in their base prices as well, though they exhibit less coordination in price promotions. This visual evidence suggests that discrimination by chains aligns with their differing price images. Below, we provide more comprehensive evidence supporting this phenomenon.

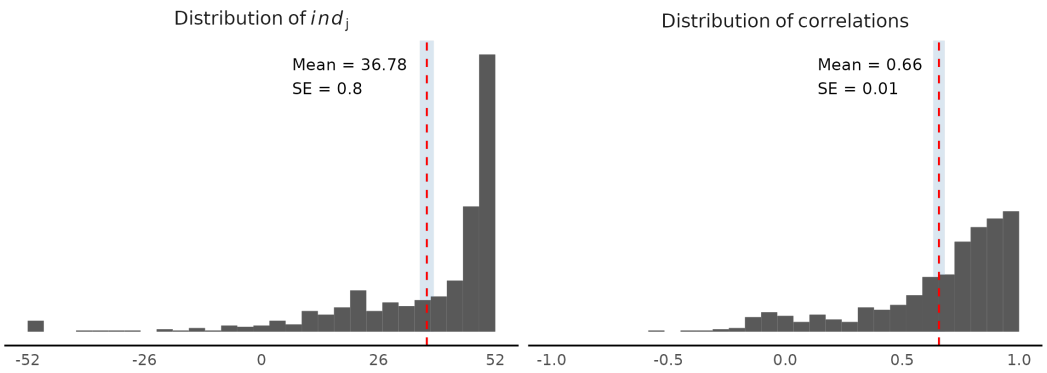
Discrimination by Traditional and Discounter Chains

For the three Markets α, β, and γ where parent 26 exhibits discrimination by chains, we now examine the extent to which chains belonging to different price zones convey different price images and whether the parent is indeed setting different price levels across the chains. For each pair of chains operating in the same market that are assigned to different price zones, consider some product j sold in both chains during all 52 weeks. Let p_{1jt} be the mean log price in stores of the first chain at week t, and similarly define p_{2jt}. Define the following variable as the difference between two indicators: $ind_{jt} \equiv \mathbb{1}(p_{2jt} - p_{1jt} > 0.01) - \mathbb{1}(p_{2jt} - p_{1jt} < -0.01)$. Its value is 0 when prices in the two chains differ by less than 0.01 log points, 1 when the price in the second chain is sufficiently higher, and -1 otherwise. The tolerance level of 0.01 is chosen to account for measurement error and is the same as that used by DellaVigna and Gentzkow (2019). Define $ind_j = \sum_{t=1}^{52} ind_{jt}$. If $ind_j = 52$, then the second chain has higher prices than the first chain for product j every week; if $ind_j = -52$, it is the other way around. If $ind_j = 0$, then the number of weeks when the product was priced higher by the second chain is the same as when it was priced lower. If there are no substantial differences in the price levels of the two chains and all differences in pricing are idiosyncratic, then we would expect the distribution of ind_j across products j to be centered around 0 and the mean of ind_j to be 0. If the second chain is a traditional grocery store and the first one is a discounter offering lower prices, then the distribution of ind_j would be skewed to the right with the mean greater than 0. Motivated by the coordinated promotions visible in the left panel of Figure 6, we also examine the distribution of correlations between weekly prices. Even though two chains belong to different price zones, if promotions are in part funded by the manufacturer and therefore determined at the parent level, promotion schedules can be the same in both chains. If this is the case, we would expect high correlations in weekly prices.

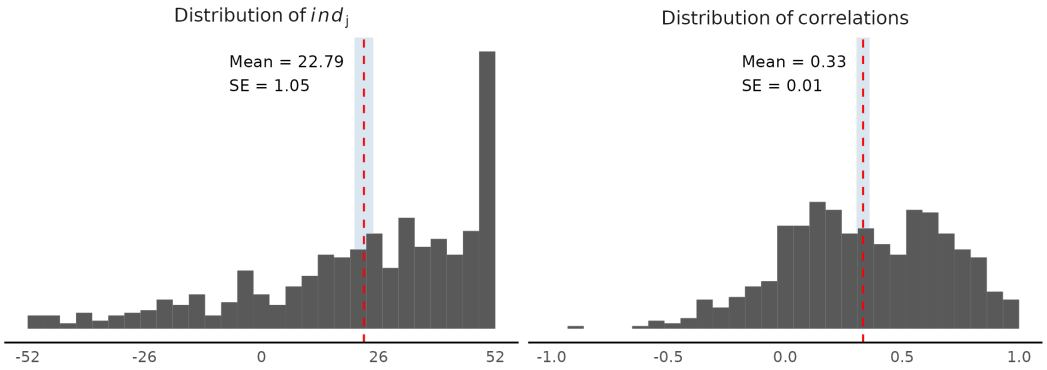
Figure 7 presents distributions of ind_j and correlations for three geographical markets where parent 26 discriminates across chains. The red dashed line shows the mean, and the blue shaded region shows the 95% confidence interval for the mean. Perhaps the starkest distinction between chains is in the Market α. The mean ind_j is 37, statistically and economically greater than 0, indicating that most products are more expensive in the stores of one chain compared to another

Figure 7: Discrimination by chains of parent 26

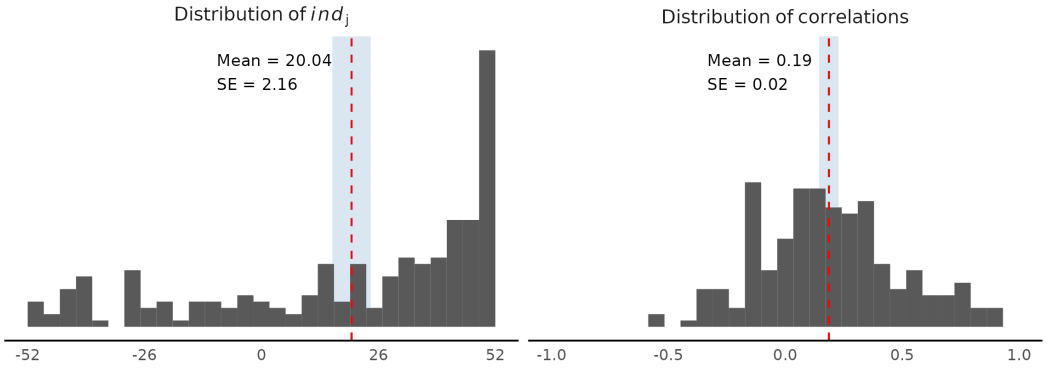
(a) Market α



(b) Market β



(c) Market γ



Notes: Each panel corresponds to a geographical market indicated by a dashed rectangle in Figure 5, where parent 26 practices price discrimination across chains. Panels display the distributions of ind_j and the correlations of weekly prices across products j. Vertical lines indicate means, and shaded blue regions denote 95% confidence intervals for the means.

for most of the year, suggesting that one chain has a traditional format and another is a discounter. Prices in the two chains are heavily correlated—the median correlation is 0.76, indicating coordination of promotions and generalizing the visual evidence from Figure 6. The distributions of ind_j in Markets β and γ are also skewed to the right, suggesting that one chain is a discounter, and the other is a traditional one. Correlations in weekly prices are positive but smaller, indicating modest coordination of promotions. Overall, the results suggest that parent 26 discriminates across chains in a geographical market if these chains have different formats.

Figure A.5 shows results for the remaining three parents. For two of them, we continue to find heavily skewed distributions indicating discrimination between traditional and discounter chains. For the remaining parent, the distribution is centered at 0 with the mean not statistically different from 0, suggesting that there is no clear vertical differentiation of chains by price images. Instead, these chains could be horizontally differentiating by targeting consumers with different tastes but similar price sensitivities.

5.4 Linking price zones to organizational divisions

We have identified two forms of zone pricing. For most parent companies, zones are determined primarily by geography, while for some they are determined by both geography and chain membership. This raises a fundamental question: how do parent companies implement such systematic price discrimination across thousands of products and stores? We hypothesize that zones reflect the parent's internal organizational structure.

Large parent companies typically follow a multi-divisional organizational structure, with divisions defined primarily by geography. To examine whether price zones align with this structure, we used an external data source independent of Nielsen to obtain maps of the organizational divisions for large parents 11 and 26, which together account for more than a quarter of stores in our data. Figures A.7 and A.8 display the geographic boundaries of these divisions alongside the price zones identified by our clustering algorithm. The visual alignment is striking given that our clustering algorithm relied exclusively on price data, without incorporating any information about divisions.

To quantify this alignment, we compute the Adjusted Rand index (Hubert and Arabie, 1985),

which equals 1 when two partitions are identical and 0 when they are unrelated.[6] The index equals 0.832 for parent 11 and 0.83 for parent 26, indicating a high degree of alignment between price zones and organizational divisions. We were unable to obtain data on organizational divisions for other zone-pricing parents. However, because price zones coincide with divisions in the two cases where such data are available, we suspect that this correspondence may hold more broadly across other large zone-pricing parent companies.

Thus, what we refer to as a "price zone" is, in practice, an operating division with authority over pricing decisions. These divisions also manage assortment selection, promotions, distribution, and other operational functions (Springer, 2013).[7] Each division has a dedicated managerial hierarchy and a regional head. Central management reviews operating results and financial statements for each division separately. In short, divisions function as distinct business units within the larger company.

Research in organizational economics (Alonso et al., 2008; Rantakari, 2008; Malenko, 2024) highlights the trade-off between coordination and adaptation that multi-divisional firms face. Pricing is a key decision for retailers and is subject to this same trade-off. Differences in local demand and competition require adaptation, while scale efficiencies and brand consistency call for coordination across stores. Delegating pricing authority to geographically based divisions serves as a middle ground, capturing economies of scale at the division level while allowing for price discrimination that exploits local market differences.

The coordination–adaptation trade-off also helps explain why parent companies pool chains of similar format into a single division, while assigning chains with different formats to separate divisions. When chains have similar formats and compete for the same customer base, the need for coordination is high. If such chains were managed by separate divisions, they could engage in aggressive price competition without internalizing the cross-divisional externalities, lowering

[6] To define the Rand index, consider all pairs of stores. Partitions are said to "agree" on a pair if both classify the stores to the same group or to different groups. The Rand index is the proportion of agreeing pairs out of all possible pairs. The Adjusted Rand index adjusts for chance, ensuring that the expected value is zero under random partitions.

[7] SEC filings confirm the substantial autonomy of regional divisions in the food retail industry, please see `https://web.archive.org/web/20250128051840/https://www.sec.gov/Archives/edgar/data/56873/000155837023004767/kr-20230128x10k.htm`, `https://web.archive.org/web/20250925023648/https://www.sec.gov/Archives/edgar/data/1646972/000164697218000017/acify1710-k.htm`.

overall parent profits. To avoid such price wars, the parent pools these chains into a single division. This results in geographically separated divisions with minimal overlap, reducing inter-divisional spillovers. The need for adaptation is relatively low in this case, as similar-format chains in the same markets face highly correlated demand and competitive conditions. In contrast, when a parent operates both discounter and traditional chains targeting different segments, the need for coordination is lower due to limited direct competition. At the same time, the need for adaptation is higher, since the chains face distinct market conditions and compete with different sets of rivals. In Appendix A.6, we document limited cross-shopping between such chains, consistent with this logic. These forces help explain why parents assign different-format chains to separate divisions.

To conclude, we do not attempt to definitively identify the factors that determine the observed organizational structure or assess which design is optimal. Rather, our goal is to demonstrate that the company's pricing decisions are fundamentally shaped by its current organizational structure.

6 Case Study of Zone Pricing Pre- and Post-Merger

The relationship between price zones and organizational structure has important implications. First, if firms design organizational structure to maximize profits, this structure should adapt when circumstances make the existing configuration suboptimal. Second, altering the organizational structure of a large company is a complex and costly process that constrains how quickly such adjustments can occur.[8] We examine these predictions using a natural experiment: a merger between two large zone-pricing parent companies. Because they operated in overlapping geographic markets, forming the merged firm's organizational structure by simply pooling the divisions of the two original parents would be suboptimal. In overlapping markets, divisions from the two parents would compete with and cannibalize one another's profits. We therefore expect the merged company to modify overlapping divisions, and thus price zones, while largely preserving non-overlapping ones. Importantly, we also expect these adjustments to occur gradually, potentially spanning multiple

[8] Major consulting firms such as McKinsey, BCG, and PwC maintain dedicated post-merger integration practices, and industry reports indicate that full integration may require 12 to 24 months (McKinsey & Company, 2010; Boston Consulting Group, 2021; PwC, 2023). The existence of specialized practices devoted to post-merger integration underscores both the demand for and the difficulty of reorganizing large firms after mergers.

years.

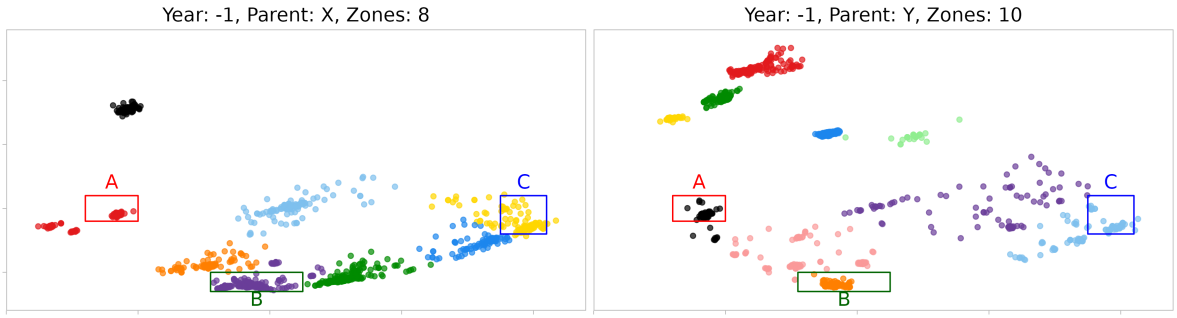
Using similarities in private label assortments as in our main analysis, we identify two parent companies, X and Y, that merged to form parent Z. We normalize the year right before the merger to -1, so the first year after the merger is 0. We applied the same clustering procedure to the pre-merger data for parents X and Y separately for the year -1 and to the post-merger data for the new parent company Z for years 0 to 4. The number of clusters was selected separately for each clustering using the same method as in the main analysis. For Silhouette Coefficient plots, t-SNE results, and clustering fit, see Figures A.9–A.11.

Next, we describe the evolution of price zones following the merger. Figure 8 presents maps of the stores colored according to their corresponding price zones, visually illustrating their evolution. To complement these visual patterns with quantitative evidence, we employ two measures. First, we measure the extent of changes in price zones between consecutive years using the Adjusted Rand index, plotted in Figure 9a. Higher values of this index indicate fewer changes in price zones from one year to another. Comparing years from 0 to 4 is straightforward; however, comparing years -1 to 0 is more complex, as we clustered parents separately in the year -1. To make this comparison, we pool all stores from both parent companies together, keep their original clusters (separate for each parent), and calculate the index relative to clusters from year 0. Second, we measure the geographic contiguity of price zones using the method described in Section 4.5, which is the accuracy of predicting the clustering label from a random forest trained on stores' latitude and longitude, as shown in Figure 9b. This measure ranges from 0 to 1, with higher values indicating more geographically contiguous zones. Using Figures 8 and 9, we can now describe how the zones changed following the merger.

In year -1, before the merger, both parent companies had perfectly geographically contiguous price zones, as evident from Figure 8 and the contiguity score close to 1 for both parents in Figure 9b. Parent X had 8 zones, while parent Y had 10 zones. Immediately after the merger in year 0, the newly formed parent company Z had 20 geographically non-contiguous zones. Consider regions A, B, and C highlighted on the maps. In year -1, both parent companies had stores in these regions. After the merger, these stores came under common ownership but continued to belong to

Figure 8: Price zones before and after the merger

(a) Before merger



(b) After merger

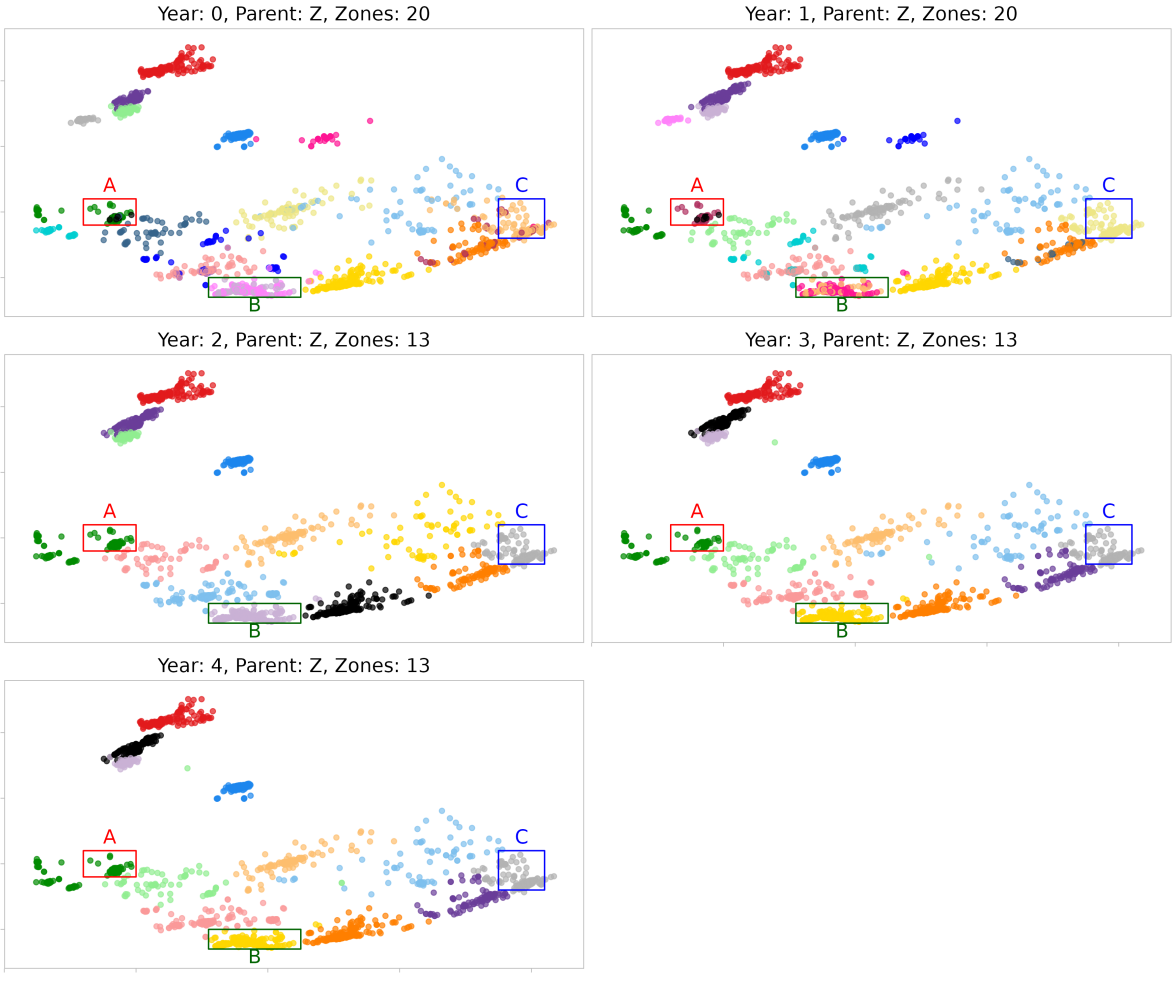
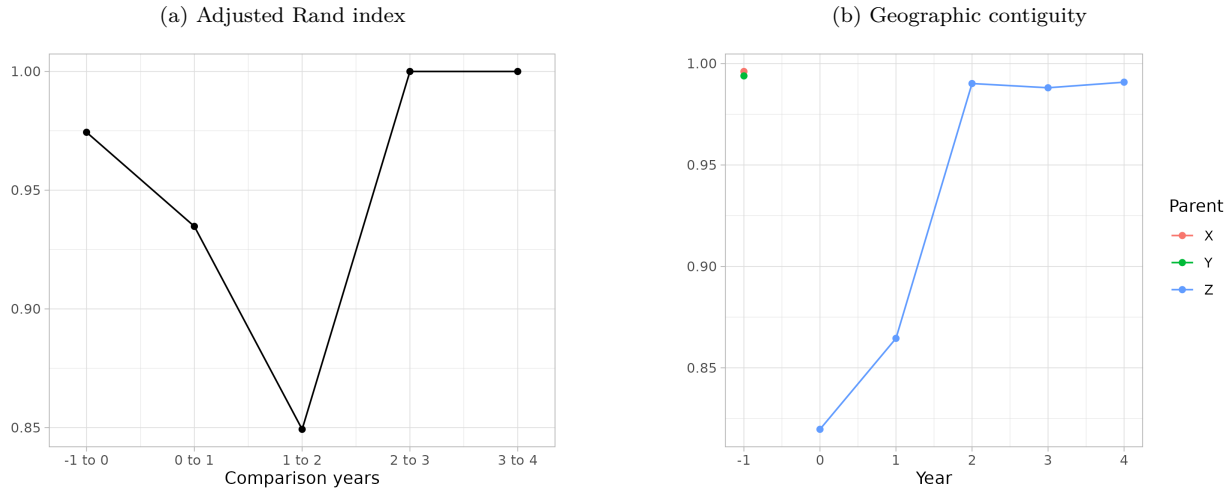


Figure 9: Description of price zones following the merger



Notes: Panel (a) shows the Adjusted Rand index comparing price zones in consecutive years. The index ranges from 0 for unrelated partitions to 1 for identical ones. Thus, lower index values indicate greater changes in the price zone structure between year t and year $t+1$. Before the merger (year -1), clustering was performed separately for the two independent parent companies. To compare year -1 to 0, we compute the Adjusted Rand Index between the pooled zones from the two separate clusterings in year -1 and the clustering of the merged parent in year 0. Panel (b) reports the geographic contiguity of price zones, measured as the out-of-bag prediction accuracy of a random forest trained on store latitude and longitude (see Section 4.5 for details). Contiguity is computed separately for the two parents before the merger and for the merged entity after.

separate price zones. This is visually evident from the maps, where the zones overlap in regions A-C, and formally from the adjusted Rand index of 0.975 in Figure 9a when comparing years -1 and 0. This high value indicates that the clustering structure barely changed after the merger, meaning that the zones of the newly formed parent company were simply pooled from the original two parent companies.

In year 1, the first notable changes in price zones occurred. In several areas where the zones overlapped in year 0, such as region C, the overlap disappeared, and all stores in that region began belonging to the same zone. Zones became more geographically contiguous compared to year 0, as indicated by the increase in score from Figure 9b. However, there were still regions with overlapping zones, such as A and B. By year 2, there were no more such regions—zones became perfectly contiguous, as visually evident from Figure 8 and formally from Figure 9b showing the score close to 1. The number of zones decreased from 20 to 13 as stores within the same region that had belonged to different zones were merged into a common zone. In years 3 and 4, the zones did not change, as evidenced by the Rand index being almost exactly equal to 1 when comparing years 2 and 3, as well as 3 and 4.

To summarize, it took two years to establish new price zones after the merger. Immediately after the merger, the zones were simply inherited from the original parent companies. Over the next two years, all geographical inconsistencies were smoothed out, resulting in a new, persistent zone structure. Importantly, price zones in non-overlapping markets remained largely unchanged throughout this period. Both the slow adjustment in overlapping areas and the stability in non-overlapping areas reflect the substantial costs of reorganizing price zones. Given our previous finding that zones are tied to geographical divisions, we believe this two-year adjustment period reflects the time needed to reorganize fundamental aspects of the merged company's organizational structure. Such large shocks to parent companies are rare, and in their absence price zones remain stable. Clustering analysis for all zone-pricing parents over 2017–2020 shows consistently high Adjusted Rand indices between consecutive years (see Appendix Table A.1), consistent with organizational divisions being costly to adjust.

7 Case Study of Soda Taxes

We have shown that large parent companies *generally* use zone pricing, meaning prices are set at the level of defined zones rather than at the individual store level. Does that mean firms must *always* adjust prices uniformly across the entire zone in response to local shocks? Since these price zones typically align with the boundaries of organizational divisions, this raises the question of whether divisional managers with pricing authority are able to deviate from zone-level uniformity in response to store-specific shocks. This speaks directly to how flexible delegated pricing authority is within the organizational structure. We examine this question in the context of a policy shock that creates sharp variation in marginal costs within zones.

In particular, we study the implementation of soda taxes (Seiler et al., 2021; Keller et al., 2024), which introduce substantial and localized increases in marginal costs. These taxes create a clear divide within price zones: some stores become subject to the tax, while others in the same zone remain untaxed. We investigate how zone-pricing parents respond—whether they maintain zone-level uniformity or allow store-level adjustments. Among all zone-pricing parents in our data, only two (denoted A and B) have multiple stores affected by these taxes, enabling meaningful analysis. One additional parent has only a single affected store, while the remaining zone-pricing parents are either unaffected or lack data for the relevant years (2016 and 2017) needed for pre-post comparisons.

The paper most closely related to this section is Butters et al. (2022). Using tax changes on beer, liquor, cigarettes, and sodas, the authors show that chains can adjust prices only in the affected stores. Our analysis differs in two key ways. First, we present the results separately for the two parent companies and each tax regime. This allows us to uncover heterogeneity in price responses across both parents and markets. Second, we focus on whether the parent adjusts prices within a price zone and make comparisons exclusively among stores within the same zone. This approach enables us to isolate the within-zone price response, whereas the results of Butters et al. (2022) may reflect both within-zone and across-zone responses.

Table 4 presents information on all excise soda taxes included in our study. The first four rows show taxes implemented in different geographical markets. As these taxes affected stores belonging

to different price zones, we analyze each scenario independently. Within each affected price zone, we designate stores subject to the tax as "taxed" and the remaining stores in the same zone as "untaxed". The last soda tax implementation we consider is in California, where multiple nearby cities implemented taxes within a short time frame. In total, we analyze five different soda tax instances: Philadelphia, Boulder, Cook County, Seattle, and the California case described above.

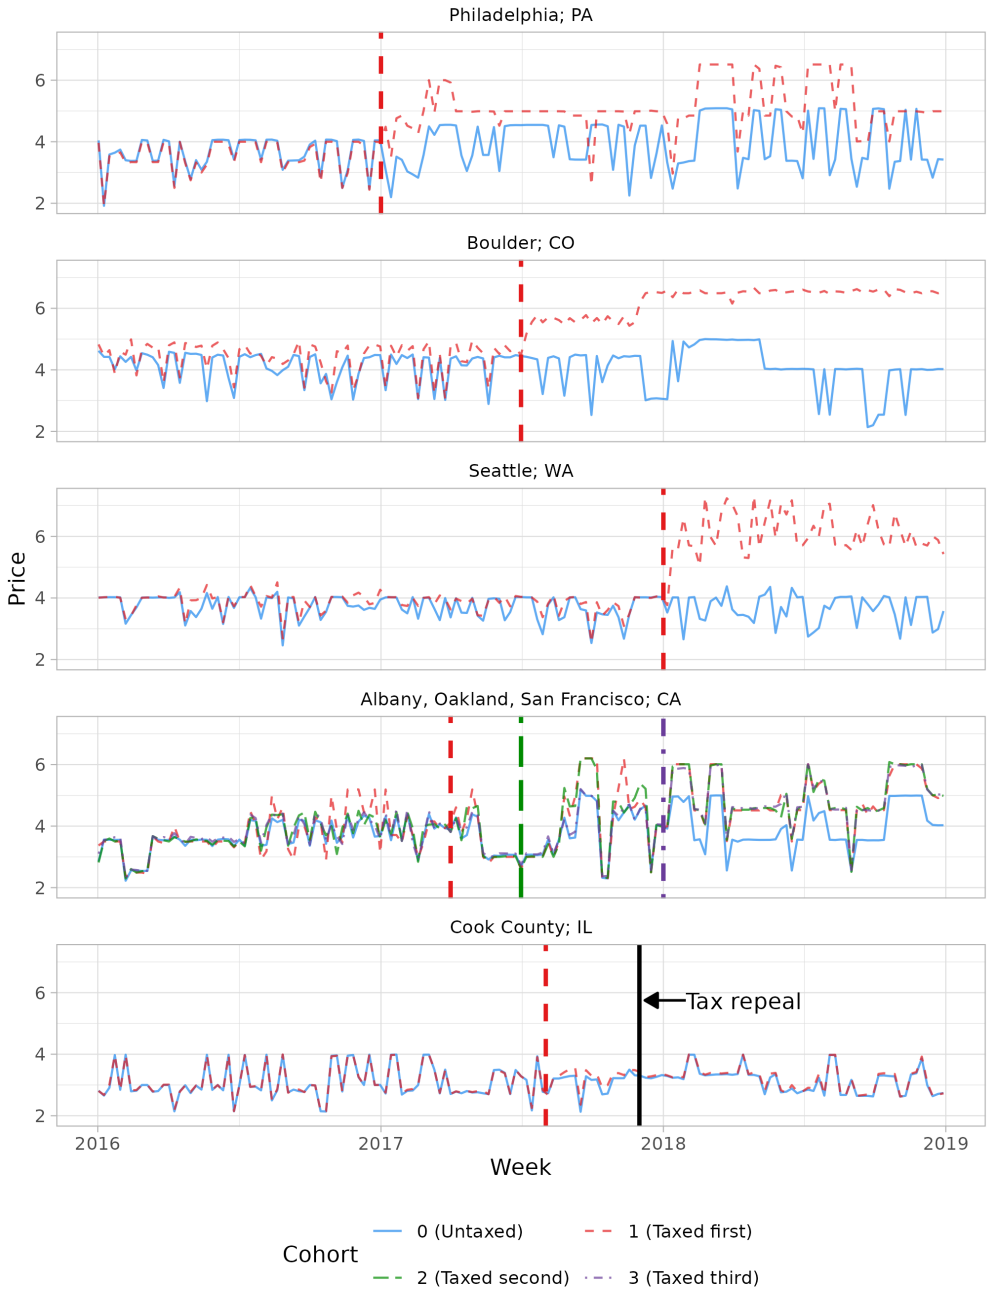
Table 4: Soda tax implementation details

Market	Region, State	Start Date	End Date	Tax Size (cents/ounce)
1	Philadelphia, PA	January 1, 2017	-	1.5
2	Boulder, CO	July 1, 2017	-	2
3	Cook County, IL	August 2, 2017	December 1, 2017	1
4	Seattle, WA	January 1, 2018	-	1.75
California taxes:				
5	Albany, CA	April 1, 2017	-	1
5	Oakland, CA	July 1, 2017	-	1
5	San Francisco, CA	January 1, 2018	-	1

Consider the null hypothesis that the parent responded to the tax at the zone level. If this hypothesis holds, prices in taxed and untaxed stores should remain similar after the enactment of the tax, as they were before. However, if the parent responded at a more granular level by differentially changing prices in taxed and untaxed stores, then the prices in these two groups should diverge after the tax enactment. It is tempting to call taxed stores as "treated" units, and untaxed stores as "control" units, but strictly speaking, we are not in the classical treatment effects setup. The Stable Unit Treatment Value Assumption (SUTVA) is violated under the null hypothesis of a zone-level response since untaxed stores are affected by the tax through the parent's zone-wide response. Therefore, we do not use the language of treatment effects and do not interpret our analysis as estimating the treatment effect of a tax policy. What we are investigating is whether there is a significant divergence in pricing between taxed and untaxed stores after the tax enactment relative to the pre-tax period. Evidence of such divergence would indicate a granular response on the parent's part, as opposed to a zone-level response.

We first focus on parent A and present the results for parent B later in this section. Figure

Figure 10: Average weekly prices for a 6-pack (16.9 oz) soda in five markets affected by soda taxes.



Notes: This figure compares average prices across parent A's taxed and untaxed stores. Each panel represents the affected market. Within each market, all selected stores belong to a single price zone. Dashed vertical lines show the tax enactment date for the respective cohort. California has three taxed cohorts since three cities implemented taxes. The solid vertical line marks the date when the tax in Cook County was repealed.

10 presents weekly prices for a popular brand's 6-pack of 16.9 oz soda bottles across five distinct markets in stores of parent A. For each market, we plot the prices averaged across taxed and untaxed stores. Pretax prices in both groups are nearly identical since in each market we selected stores belonging to a single price zone. Consider taxes in Philadelphia, Boulder, and Seattle. After the tax was implemented, there was a sharp increase in prices in the taxed stores, which was not mirrored by the untaxed stores. Thus, for this specific product, the parent company responded to the tax not at the zone level but changed pricing specifically for the taxed stores. The price patterns are similar in California, though stores in Albany and Oakland showed a delayed response to the tax enactment. In 2018, average prices in all of the taxed stores were higher than in the untaxed stores, again reflecting a targeted response to the local shock.

Let us now examine Cook County, Illinois, which presents a unique scenario in our analysis of soda taxes. Unlike other instances we have discussed, this tax was reflected in the prices at the register rather than the shelf prices (Leider et al., 2022). Note that NielsenIQ RMS captures shelf prices rather than prices at the register. As a result, the tax itself is not included in prices from the Nielsen RMS dataset, much like other sales taxes applied at the register. Therefore, what appears in our data as a null effect is actually a full pass-through of the tax to consumers in the taxed stores. The parent company did not have to adjust the shelf price labels to "respond" to the shock at a more granular level, it simply applied the tax at the register. Had the tax not been repealed, retailers would have been required to include taxes in the shelf price starting January 1, 2018, (Cook County, Illinois Department of Revenue, 2017), and then we would likely have observed an increase in prices in our data. However, as it stands, the parent continued to set zone-level prices and fully passed the tax to the consumer at the register.

After presenting visual evidence for a single product, we turn to a more comprehensive analysis. We focus on major soda products that were consistently among the top sellers in our sample period. For each year from 2016 to 2018, we identify the top 20 revenue-generating sodas subject to tax; of these, 17 products appear in the top 20 list across all three years. Let $\tilde{p}_{sjt}$ denote prices in cents per ounce in store s for product j on week t. We consider the following time periods. For each tax regime except California, we examine one year before and after the tax enactment (for Cook

County, one year before enactment and after repeal). Given California's staggered implementation, we analyze one year prior to its earliest tax and one year after its latest tax. We estimate the following OLS regression specification:

$$\tilde{p}_{sjt} = \beta \cdot D_{st} + \gamma_{jt} + \delta_{sj} + \epsilon_{sjt}, \tag{1}$$

where D_{st} is an indicator variable equal to 1 if store s is subject to soda taxes in week t and 0 otherwise. The coefficient β measures the differential price response of taxed stores relative to untaxed stores. The specification includes product-week fixed effects γ_{jt} and store-product fixed effects δ_{sj}. To interpret this regression, consider allowing the coefficient on D_{st} to vary by product (β_j). This would amount to estimating separate two-way fixed effects regressions for each product. Instead, by constraining β to be common across products, we obtain a weighted average of the product-specific coefficients β_j, making the results more easily interpretable.[9] As a robustness check, we also estimate a two-way fixed effects regression at the store-week level using a price index that averages prices across products. These results, reported in Appendix A.11, are very similar to our main specification.

We estimate equation (1) and report the results with standard errors clustered at the store level in Table 5. Under the null hypothesis of a zone-level response to the tax, we would expect $\beta = 0$. For all tax regimes except for Cook County, the estimates are statistically and economically significant. As for Cook County, recall that the tax is not included in the prices, so the real pass-through rate is 107%. Across markets, pass-through rates vary from 86% in Boulder to 110.9% in Seattle. Such high levels indicate that the parent did respond to soda taxes not at the zone level, but at the level of stores affected by the taxes. We conclude that when faced with a large cost shock affecting only specific stores within a zone, a parent company can deviate from its zone pricing strategy. The parent may be willing to bear the administrative costs of implementing store-specific pricing to avoid the larger losses that would result from maintaining uniform zone-level prices.

We now discuss the response to soda taxes of parent B. Appendix Table A.4 presents results analogous to Table 5. There is very little divergence in prices between taxed versus untaxed stores.

[9] Formally, β is the weighted average of β_j with non-negative weights. For the proof, see Appendix A.10.

Table 5: Parent A's response to soda taxes. Estimation results

Dependent Variable:	Price (cents / ounce)				
Tax regime	Philadelphia	Boulder	Cook County	Seattle	Albany, Oakland, SF
D_{st}	1.458***	1.721***	0.069***	1.941***	0.804***
	(0.0114)	(0.0114)	(0.0035)	(0.1227)	(0.0176)
Store × product FE	Yes	Yes	Yes	Yes	Yes
Week × product FE	Yes	Yes	Yes	Yes	Yes
Observations	267,285	191,881	357,039	309,838	557,985
Tax rate	1.5	2	1	1.75	1
Implied pass-through rate	97.2%	86%	106.9%	110.9%	80.4%
	(0.76%)	(0.57%)	(0.35%)	(7%)	(1.76%)

Notes: Clustered (store) standard errors in parentheses. Signif. Codes: ***: 0.001, **: 0.01, *: 0.05.

Additionally, there was no economically significant increase in prices in the untaxed stores after the tax was implemented. The average price per ounce during the 6 months after the tax is only 0.075 cents higher than during the same period immediately before the tax. Taken together, these facts suggest that parent B did not include taxes in the shelf prices across all markets, similar to what parent A did in Cook County. The practice of not including the tax in the shelf price was documented in markets other than Cook County, for example, in Seattle.[10] One rationale for excluding the tax from the shelf price is that doing so reduces tax salience, leading consumers to underreact to the resulting price increase (Chetty et al., 2009). While this demand-side channel is relevant, we discuss two additional advantages of this strategy that stem from organizational frictions.

Firstly, not including the tax in the shelf price reduces menu costs. Appendix Figure A.12 shows a price tag that simply adds a fine print message "*Plus imposed sugar expense*" rather than displaying different prices for each product. This approach simplifies both physical and online price management, as it allows the parent to maintain uniform zone-level pricing while adding a single tax notification in affected stores.

Secondly, by not including the tax in the shelf price, the parent does not have to adjust its

[10] `https://sph.washington.edu/news-events/news/seattles-soda-tax-almost-entirely-passed-consumers-researchers-find`

feature flyers. As we show in Appendix A.14, organizational divisions are responsible not only for pricing but also for feature promotions. Since feature ads usually list prices, having varying within-zone shelf prices would prevent the parent from having a common feature ad for soft drinks at the zone level. There are three ways of dealing with this complication. The first way is to stop featuring soda drinks altogether, which may be prohibitively costly if features on soda drinks drive substantial traffic to the store. The second way is to create store-specific feature fliers by having one ad for the taxed stores and another ad for the untaxed stores. Appendix Figure A.13 shows an example of the ad for the taxed stores with the price before tax made very salient, and the after-tax price written in fine print. The third way is to use quantity discounts, an example is in the Appendix Figure A.13c. Quantity discounts do not advertise prices, so the feature fliers can be the same across taxed and untaxed stores. By not including taxes in shelf prices, the parent avoids having to alter its existing feature promotion strategy.

To conclude, we find that when faced with a large local cost shock, such as soda taxes, the parent can break its zone-pricing policy to granularly adjust prices in the affected stores. However, to avoid the operational costs associated with changing price tags and feature promotions, a parent may instead preserve its zone pricing policy by passing the tax through at the register.

8 Do Parents Vary Prices with Demand Elasticities?

We have shown that price zones are tied to organizational divisions and are costly to adjust, as evidenced by the two-year restructuring following the merger. However, while the zone structure itself may be sticky, divisional managers have authority over setting prices within their zones. A natural question is whether they use this authority effectively: do zone-level prices respond to zone-level demand elasticities in the manner predicted by profit-maximizing behavior?

Answering this question helps us understand the sources of deviations from optimal pricing. DellaVigna and Gentzkow (2019) document that store-level prices within chains respond weakly to local demand elasticities, and propose two main explanations. Behavioral frictions arise when managers have limited attention, fail to accurately estimate demand elasticities, or lack understanding of how to apply optimal pricing rules. Agency frictions arise when misaligned incentives prevent

managers from implementing optimal policies even when they understand them. Our analysis at the zone level allows us to distinguish between these explanations. Divisional managers in our setting already practice price discrimination, setting different prices across zones that serve largely separate geographic markets with distinct demand conditions. This demonstrates that managers possess both the authority to vary prices and face incentives to respond to local demand. Therefore, any remaining deviation from optimal zone-level pricing is more plausibly attributed to behavioral factors rather than agency frictions. Managers may not fully optimize even when they have both the authority and incentive to do so.

We focus on the eight parent companies from the first group of Table 3, which use price zones based on geography and/or chains. We closely follow the model of DellaVigna and Gentzkow (2019) to ensure that our results are comparable with established findings. Parent company r faces demand for product j in store s at week t of the form $Q_{sjt}(P) = k_{sjt}P^{\eta_{z(s)j}}$, where $z(s)$ represents the zone of store s and $\eta_{z(s)j}$ denotes the zone-level demand elasticity. For tractability, we assume that elasticities are the same across stores belonging to the same zone in our model. This assumption lets us estimate a single zone-level price elasticity using the log-log demand model and pooling all stores within the zone.[11]

Let c_{rjt} denote marginal costs for parent company r selling product j in week t. Within a parent company, costs can vary over time but are assumed to be constant across stores in any given week. This assumption is critical for testing the parent's optimal pricing, and we adopt it to maintain consistency with the model of DellaVigna and Gentzkow (2019). When we later assess the welfare implications of zone pricing, we discuss why this assumption may be questionable, but we adopt it here to remain consistent with the previous literature.

For every zone-product-week combination, the parent company chooses the price to maximize profits:

$$\max_{P_{zjt}} \sum_{s:z(s)=z} (P_{zjt} - c_{rjt})Q_{sjt}(P_{zjt}). \tag{2}$$

[11] Hitsch et al. (2021) show that store-level price elasticities are noisily estimated. Since prices are set uniformly at the zone level, the firm can increase precision by estimating a single zone-level elasticity, pooling data from stores within the same zone. Our analysis mirrors this approach, assuming that the price elasticities are the same across stores within the same zone. This methodology allows for a more accurate estimation of price elasticities at the zone level, potentially aligning our results with how firms might operate in practice.

The optimal price is then derived from the Lerner index:

$$\log(P^*_{zjt}) = \log\left(\frac{\eta_{zj}}{1+\eta_{zj}}\right) + \log(c_{rjt}). \tag{3}$$

Averaging equation (3) across t and setting $\lambda_{zj} \equiv \log\left(\frac{\eta_{zj}}{1+\eta_{zj}}\right)$ yields:

$$\overline{\log(P^*_{zj})} = \lambda_{zj} + \overline{\log(c_{rj})}. \tag{4}$$

This equation suggests that if the parent company prices according to the model above, then regressing average zone-level prices on λ_{zj} and product-parent fixed effects, absorbing marginal costs, would yield a coefficient of 1. To test this, we need zone-level demand elasticities. However, we do not observe zone-level elasticities directly and therefore have to estimate them. Taking logs of the demand equation gives:

$$\log(Q_{sjt}) = \eta_{z(s)j} \log(P_{sjt}) + \log(k_{sjt}). \tag{5}$$

The demand shock $\log(k_{sjt})$ is a store-specific error term that is potentially correlated with prices. For each zone z and product j, we estimate the following specification:

$$\log(Q_{st}) = \hat{\eta} \cdot \log(P_{st}) + \tau \cdot F_{st} + \delta_s + \epsilon_{st}, \tag{6}$$

where s indexes stores belonging to the price zone z, δ_s is the store fixed effect, and F_{st} is the dummy for the product being featured. NielsenIQ reports feature data only for a small subsample of stores. Because we find that features are determined at the level of price zones, we impute missing feature observations at the zone level (see Appendix A.14).

We instrument $\log(P_{st})$ with the average log price of the same product across stores outside of store s's price zone, following the approach in Hausman (1996). This instrument exploits the fact that zone-level pricing decisions are correlated due to common supply-side factors, such as changes in wholesale costs, while prices in other zones are unlikely to respond directly to demand shocks in the focal zone. When using simple OLS, we continue to find a strong relationship between demand elasticities and prices, reinforcing the evidence for systematic price discrimination across zones (see Table A.6).

The regressions yield estimated zone-level price elasticities $\hat{\eta}_{zj}$. Some of them are noisily estimated: 16% of the estimates are not statistically different from 0. Among those that are significantly different from 0, 7.2% are positive. In what follows, we select elasticities that lie in the interval $(-7, -1.2)$ and have a standard error of less than 0.3. Results are robust both to the choice of the interval and of the tolerated standard error (see Table A.6). From $\hat{\eta}_{zj}$, we construct the empirical analog of λ_{zj}, which is $\hat{\lambda}_{zj} = \log(\frac{\hat{\eta}_{zj}}{1+\hat{\eta}_{zj}})$. From equation (4), it follows that if all zone-pricing parents behave according to the outlined model, then regressing average log prices of product j in zone z on $\hat{\lambda}_{zj}$ and parent-product fixed effects should yield a slope of 1. That is, the following regression:

$$\overline{\log(P_{zj})} = \beta \cdot \hat{\lambda}_{zj} + \delta_{rj} + \epsilon_{zj},$$

should yield a coefficient of $\beta = 1$. A coefficient of 1 would indicate that parents are setting prices exactly as predicted by (3), fully exploiting differences in demand elasticities across zones to maximize profits. The issue is that $\hat{\lambda}_{zj}$ itself is estimated with error, leading to attenuation bias and β being biased towards 0. To address this, we instrument for $\hat{\lambda}_{zj}$ using the zone-level consumer income (an approach similar to DellaVigna and Gentzkow [2019]), which is simply store-level income averaged across stores in the zone. Table 6 presents the results for the two regression specifications. The first is OLS, while the second uses income as an instrument.

Table 6: Testing for optimal pricing of zone-pricing parents

Dependent Variable:	Log price	
Model:	OLS	IV
$\hat{\lambda}_{zj}$	0.0485***	0.7084***
	(0.0041)	(0.0541)
Parent × product FE	Yes	Yes
F-stat 1st stage		509.9
Observations	18,641	18,641

Notes: Clustered (product×parent) standard errors in parentheses. Signif. Codes: ***: 0.001, **: 0.01, *: 0.05.

The estimated coefficient is close to 0 in the OLS specification, which is consistent with a high degree of measurement error. When we use instrumental variables, the coefficient increases to 0.71.

While this is distinct from the value of 1 implied by the model, it is substantially higher than the 0.085 obtained by DellaVigna and Gentzkow (2019). The driver of this difference is that we narrow our attention to parent companies that we identified as zone-pricers. Therefore, for parent companies that engage in zone pricing, we find that prices vary systematically as a function of zone-level elasticity.

To summarize, for zone-pricing parents, we find a strong positive relationship between demand elasticities and prices. Although this relationship is not as strong as the optimal benchmark would predict for most parents, it is considerably stronger than previously reported in the literature, where all chains were pooled together. By accounting for ownership structure, focusing solely on zone-pricing parents as opposed to all parents, and on the relevant unit of decision-making (a price zone), we document that prices vary with local elasticities to a considerable extent.

Welfare implications of zone pricing

Finally, we assess how the observed zone pricing policy affects profits and consumer surplus. This analysis requires knowledge of marginal costs, which are not observed. To recover them, we assume that the observed zone prices are optimal and back out store-level marginal costs by inverting the first-order condition for profit maximization: $\hat{c}_{sjt} = P_{sjt}\left(1 + \frac{1}{\hat{\eta}_{z(s)j}}\right)$, where P_{sjt} is the observed price and $\hat{\eta}_{z(s)j}$ is the estimated elasticity. Because prices vary at the zone level, this assumption implies that marginal costs also vary across zones. This differs from the assumption made earlier and in the prior literature, which treats marginal costs as constant across stores. In what follows, we explain why, in our setting of large and geographically dispersed zone-pricing parent companies, marginal costs can vary across zones.

Historically, manufacturers sought to offer region-specific wholesale price discounts via off-invoice allowances. However, these efforts were undermined by diversion, where retailers resold products acquired in low-cost regions to higher-cost regions, and by forward-buying, where retailers overstocked during promotional periods (Buzzell and Quelch, 1990; Zwiebach, 1996). To mitigate these moral hazard problems, manufacturers increasingly adopted scanbacks, which tie discounts to actual sales rather than shipments, enabling more effective region-specific promotions

(Drèze and Bell, 2003). Additionally, Anderson and Fox (2019) document that most trade promotions are communicated bilaterally between manufacturers and retailers. This structure plausibly allows manufacturers to vary promotional terms across regions, depending on manufacturer performance, objectives for the region, and retailer characteristics. Trade press and consulting reports also indicate that manufacturers can tailor trade promotions to regional market conditions.[12] In Appendix A.16, we document cases in which the same product is promoted in one zone but not in others within the same week, consistent with zone-specific negotiations of trade promotions between parents and manufacturers.

While existing academic literature finds limited geographic variation in wholesale prices, we believe this is a consequence of data limitations. Besanko et al. (2005) show that wholesale prices in the Dominick's chain were uniform across stores. However, Dominick's operated in the relatively compact Chicago market, whereas the large zone-pricing chains we study span multiple states and have stronger incentives to engage in region-specific trade promotions. Stroebel and Vavra (2019) find little geographical variation in wholesale prices using data from independent wholesalers, who typically serve smaller, unaffiliated retailers. In contrast, the large parent companies in our setting likely negotiate directly with manufacturers and may have greater flexibility to secure region-specific promotional terms, potentially resulting in geographic variation in effective wholesale costs.

We define the parent's optimal uniform price as the solution to:

$$\max_{P_{jt}} \sum_{s} k_{sjt} P_{jt}^{\eta_{z(s)j}} (P_{jt} - c_{sjt}).$$

We compute the optimal profits by substituting the estimated elasticities $\hat{\eta}_{z(s)j}$, demand shocks as $\hat{k}_{sjt} = \frac{Q_{sjt}}{P_{sjt}^{\hat{\eta}_{z(s)j}}}$, and marginal costs $\hat{c}_{sjt}$. For each parent, we then calculate the percentage change in profits and consumer surplus from moving from the optimal uniform pricing to the observed zone pricing. Results are shown in Table 7. Across parents, zone pricing increases profits by 0.26%–6.81%, consumer surplus by 0.15%–4.88%, and total welfare by 0.18%–5.49%.

These modest profit gains reflect organizational constraints on zone design. Price zones align

[12] See Eversight, "What Diversification Means for Trade Promotion Effectiveness", accessed June 3, 2025; and McKinsey & Company, "How Analytics Can Drive Growth in Consumer-Packaged-Goods Trade Promotions", accessed June 3, 2025.

Table 7: Welfare change of zone pricing compared to uniform pricing

Parent	11	20	22	23	26	31	36	42
Δ Profit (%)	6.81	2.82	0.26	4.19	2.04	1.62	0.85	4.02
Δ Consumer Surplus (%)	4.88	2.24	0.15	3.05	1.43	1.32	0.35	2.09
Δ Welfare (%)	5.49	2.44	0.18	3.41	1.63	1.42	0.52	2.71

with organizational divisions and are determined primarily by geography rather than local demand and competitive conditions. Since divisions span broad geographic areas, they aggregate stores facing heterogeneous market conditions that would ideally be priced differently. This organizational constraint limits potential gains from zone pricing relative to a structure designed purely for profit maximization. While we find positive effects of zone pricing on total welfare, the welfare impact of third-degree price discrimination depends on the curvature of demand (Miravete, 2024). We adopt the constant-elasticity form for consistency with our prior analysis of pricing optimality, but recognize that the welfare results should be interpreted with this limitation in mind.

9 Conclusion

How do large multi-unit firms organize decision-making? This paper uses retail pricing to study this question. We show that identifying actual decision-making units requires tracing retail chains to their parent companies. Grouping commonly owned chains reveals that large retailers systematically divide operations into geographic price zones. For two major parents where we can verify organizational structure, these zones align closely with organizational divisions, indicating that pricing authority resides at the divisional level.

We then examine how this organizational structure shapes firms' responses to changing market conditions. Following a merger, the reorganization of price zones took nearly two years, revealing the high costs of overhauling internal divisions. While changing divisional boundaries is costly, divisional managers can deviate from uniform pricing when facing local shocks, as evidenced by our analysis of soda taxes. Finally, we show that zone prices respond to local demand elasticities, though imperfectly, suggesting that delegated authority enables adaptation while falling short of

full optimality.

These results challenge standard models that treat firms as unitary actors costlessly optimizing in response to market conditions. Our evidence shows organizational structure creates real constraints on firm behavior. Even sophisticated firms in competitive markets cannot freely adapt to local conditions. Internal boundaries determine who makes decisions, what information gets used, and how quickly firms respond to shocks. The two-year reorganization period demonstrates these constraints bind not just for pricing but for any strategic dimension coordinated within organizational units. Understanding firm responses to policy changes, competition, or market entry requires modeling not just market incentives but the internal architecture through which firms act. Any large firm delegating decisions across geographic or product markets faces similar tradeoffs: banks setting lending terms across branches, manufacturers coordinating production across plants, hospital systems determining treatment protocols across facilities. For empirical work, this implies the need to identify where decisions are actually made within firms, since true decision boundaries may diverge substantially from nominal corporate lines.

A Appendix

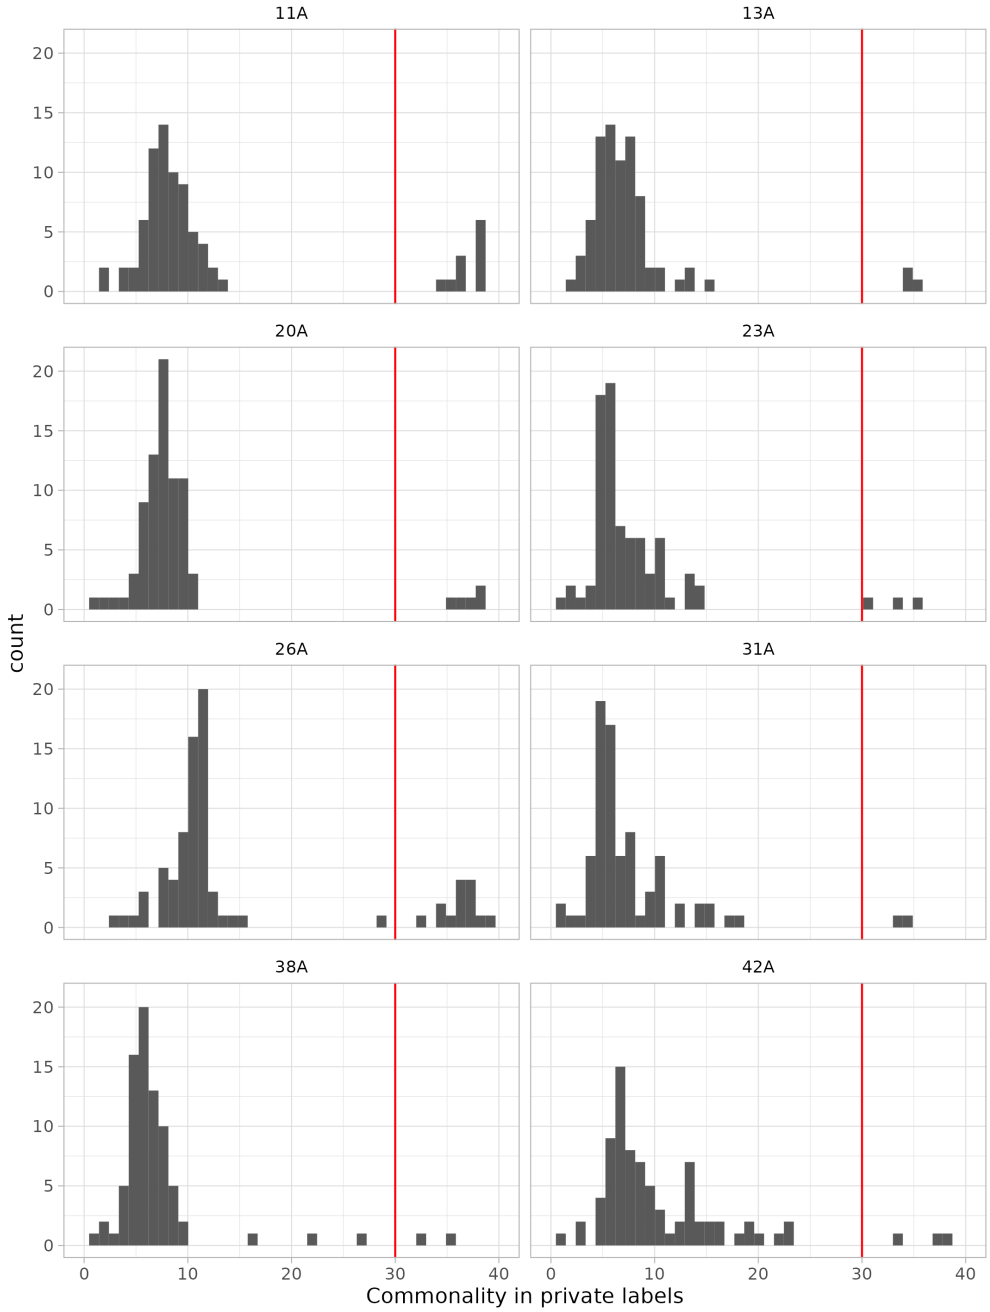
A.1 Inferring chain's ultimate parent

The strategy of determining the true parent is based on the commonalities of private labels across different chains. The exact algorithm is as follows. Fix a product module, one chain A, and construct the set of private labels UPCs sold at that chain. Then for every other chain B calculate the share of revenues from selling private label products which comes from selling the products that appeared in chain A. Summing this share across all modules we obtain a score between 0 to 40, which shows how much of the private labels revenues in the chain B are coming from the products that appeared in the chain A. If the value is high, then chains A and B likely belong to the same parent since their private label assortments are similar.

Figure A.1 illustrates how we formed eight parents with multiple chains. Each panel represents the largest chain in the parent company, with the distribution of scores across all other retail chains. The bimodal distributions indicate groups of chains very similar to the focal one in their assortments of private labels. We chose 30 as the cutoff for assigning a chain to the parent company, indicated by a red vertical line in the Figure. Out of eight parents uncovered this way, only three companies have common values in the field *parent_code* across their stores. These are parents 42, 31, and 38. The stores of the other five parent companies have multiple values in the *parent_code* field, and thus can't be recovered by simply relying on this field. However, the fact that the few chains whose stores share the same *parent_code* value are also recovered by our algorithm provides reassurance that it captures the underlying organizational structure accurately.

We believe this procedure likely underestimates the size of the true parent companies, as a parent may not necessarily expand its private label products to all chains under its ownership, resulting in some chains being missed in this analysis. Another potential challenge to our strategy is the presence of large wholesalers that supply the same private label brands to independent retailers. These wholesalers typically serve many small independent chains that wish to carry inexpensive private label products but lack the scale to develop their own branding. We identified one group of chains with common private label products, where several chains operate fewer than five stores.

Figure A.1: Construction of parents that own multiple chains



Notes: In our analysis, there are eight parents that own more than one chain. Each panel represents the largest chain of the parent company, showing the distribution of similarities in product label assortments of all other chains relative to the focal chain. The parent company is formed by combining all chains with a score larger than 30 (to the right of the red line) with the focal chain.

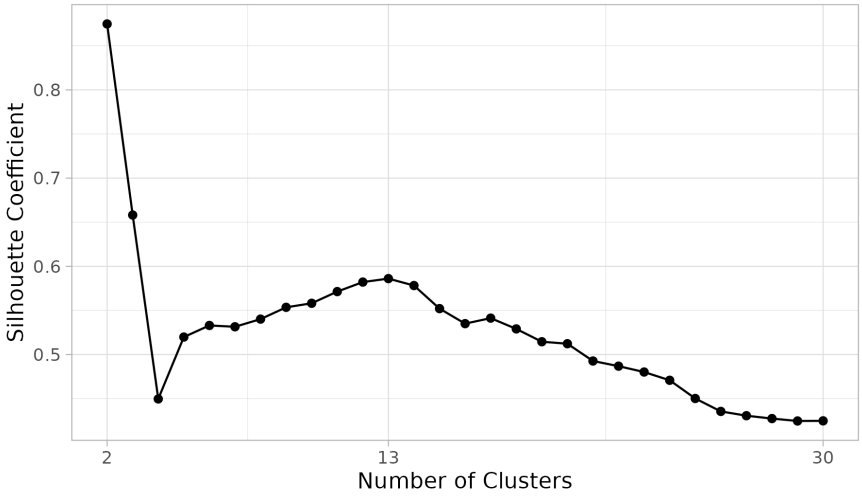
We suspect that these chains may share a common private label supplier but otherwise operate independently. We treat all these chains as independent entities in our analysis. If anything, this conservative treatment again leads us to underestimate the size of parent companies, biasing our classification toward uniform pricing rather than zone pricing.

A.2 Simulation with multi-peaked Silhouette Coefficient

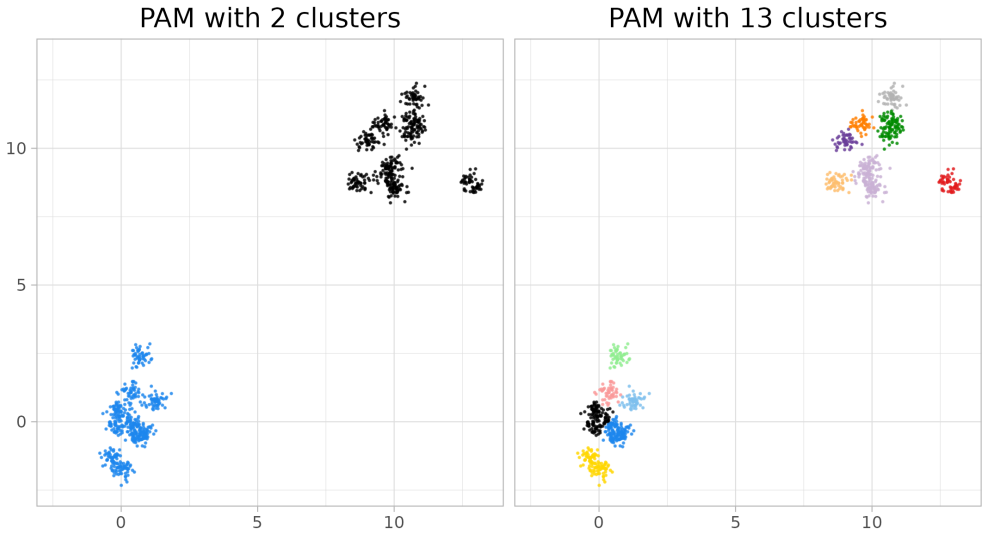
We suspect that the multi-peak behavior of the Silhouette Coefficient exhibited for parent 26 in the Figure 2 is a result of two well-separated clusters with multiple subclusters within them. Let N be the number of clusters used for PAM. Intuitively, the coefficient is high for $N = 2$ because the distance between clusters is much larger than the distances within them. When N increases, the coefficient immediately drops because major clusters are now divided into a few poorly separated groups of subclusters. However, when N increases even further, PAM starts picking up individual subclusters, increasing the coefficient. To illustrate this intuition, we simulated 1000 points with two well-separated clusters and multiple subclusters. Figure A.2a shows the dynamics of the Silhouette Coefficient, which is very similar to the one of parent 26: high value at $N = 2$, a decrease up to $N = 4$, and then an increase until reaching a local maximum at $N = 13$. Figure A.2b shows that at $N = 2$ the algorithm identifies two large clusters, while at $N = 13$ the algorithm successfully identifies subclusters.

Figure A.2: Simulation results: clustering evaluation and partitioning

(a) Silhouette Coefficient for the simulated dataset



(b) PAM clustering results for the simulated dataset



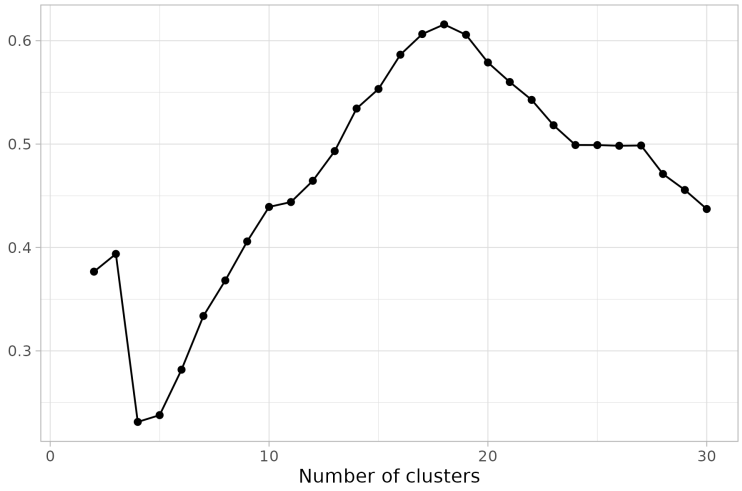
Notes: See Section A.2 for details of the simulation setup.

A.3 Temporal consistency check for parent 26

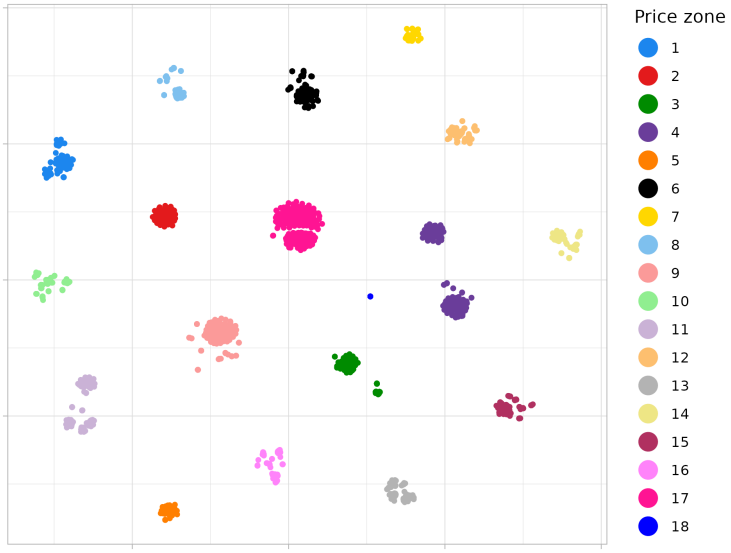
To further verify our selection of the number of clusters for parent 26, we perform a temporal consistency check by repeating the clustering on 2010 data and comparing the results across years. Appendix Figure A.3 shows that, when analyzing 2010 data, the Silhouette Coefficient suggests 18 clusters, which is much closer to our selection of 23 than to 2. The t-SNE plot likewise indicates that the correct number of clusters in 2010 is around 18. Moreover, the clusters identified from 2010 data exhibit high congruence with those from 2019, with an Adjusted Rand index of 0.9 (for a detailed explanation of this metric, see footnote 6). Since the correct number of clusters in 2010 appears to be high, and these clusters align closely with those from 2019, we feel confident in preferring 23 clusters over 2 clusters in 2019.

Figure A.3: Choosing optimal number of clusters for parent 26 in the year 2010

(a) Silhouette Coefficient



(b) t-SNE

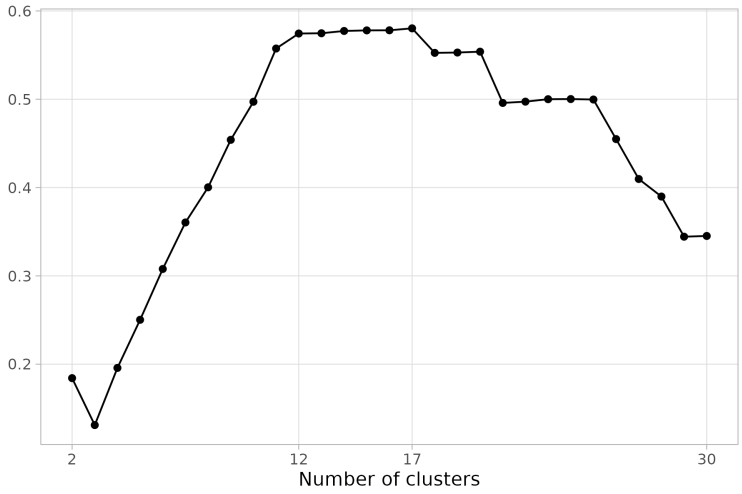


Notes: The first panel plots the Silhouette Coefficient for parent 26 in the year 2010, which is maximized when the number of clusters is 18. The second panel plots the output of t-SNE, with each point being a store. Colors represent the store's cluster when running PAM on the original price data with 18 clusters. Thus, colors *do not* rely on t-SNE in any way. As the Figure shows, with 18 clusters, the results of PAM and t-SNE agree.

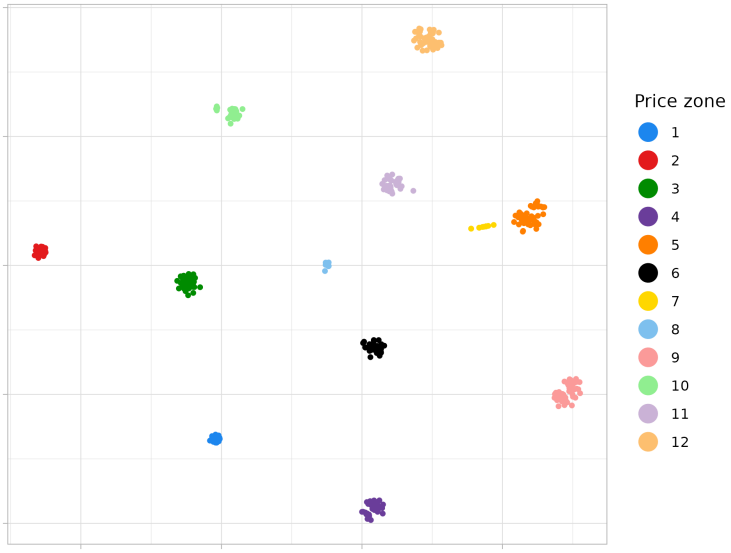
A.4 Choosing optimal number of clusters for parent 22

Figure A.4: Choosing optimal number of clusters for parent 22

(a) Silhouette Coefficient



(b) t-SNE

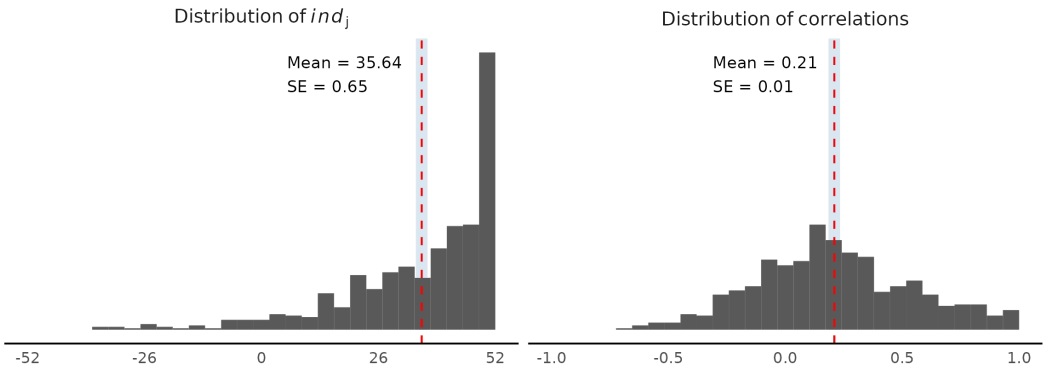


Notes: The first panel plots the Silhouette Coefficient for parent 22, which is maximized when the number of clusters is 17. The second panel plots the output of t-SNE, with each point being a store. Colors represent the store's cluster when running PAM on the original price data with 12 clusters. Thus, colors *do not* rely on t-SNE in any way. As the Figure shows, with 12 clusters, the results of PAM and t-SNE agree, so we select 12 to be the number of clusters.

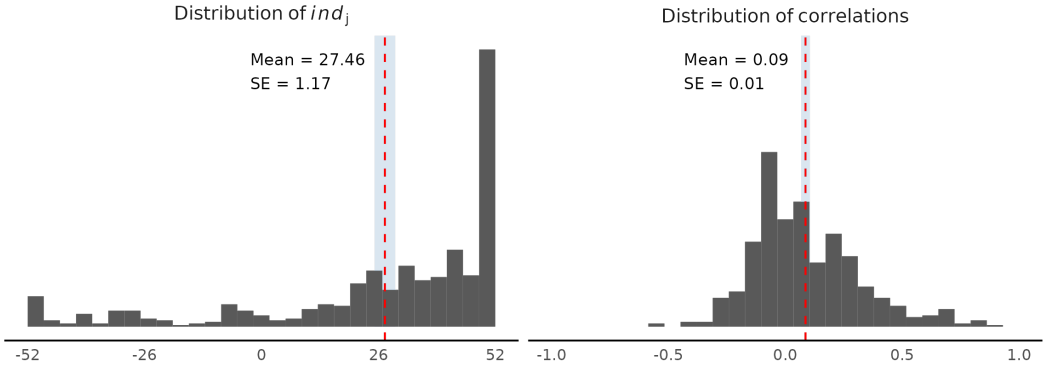
A.5 Discrimination by chains of other parents

Figure A.5: Discrimination by chains of other parents

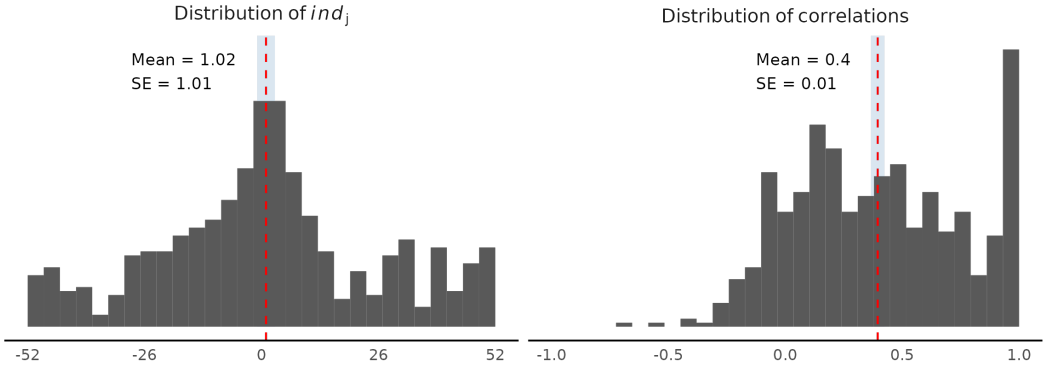
(a) Parent 23



(b) Parent 42



(c) Parent 31



Notes: Panels show the distributions of ind_j and correlations of weekly prices across products j. Vertical lines indicate means. Shaded blue regions indicate 95% confidence intervals for the means.

A.6 Spillovers between traditional and discounter chains

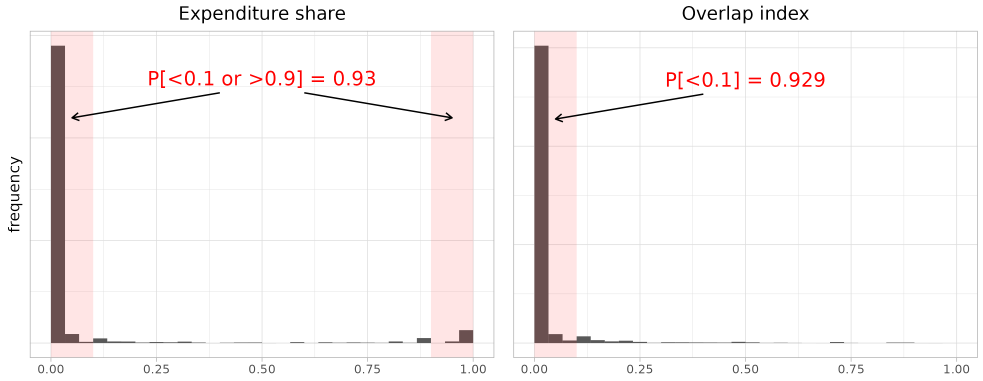
We investigate whether the parent company should account for the spillovers between its traditional and discounter chains when setting prices. Consider parent 26, pick one of the Markets from α, β, and γ. Let X and Y denote chains in the chosen market that belong to different price zones. Our objective is to measure the strength of competition between X and Y. One way to assess this is by examining whether consumers frequently shop at both chains or mainly shop only in one of them. Formally, let i denote a consumer in the HMS panel who shopped at least once in either chain. Let X_i and Y_i represent the consumer's yearly expenditures at the respective chains. We calculate two measures of cross-shopping behavior. First, we compute the expenditure share $\frac{\mathrm{X}_i}{\mathrm{Y}_i+\mathrm{X}_i}$. Shares close to 0 or 1 suggest limited cross-shopping. Second, we construct the overlap index, defined as $\frac{\min(\mathrm{X}_i,\mathrm{Y}_i)}{\max(\mathrm{X}_i,\mathrm{Y}_i)}$. A value near 0 similarly indicates weak cross-shopping behavior.

Figure A.6 shows the distribution of these measures for all three Markets α, β, and γ. Each consumer is weighted by the product of the projection factor provided by NielsenIQ (to correct for non-random sampling in the HMS panel) and their total expenditures at the two chains (to account for their relevance to the parent company). As shown in the figure, expenditure shares are highly concentrated at 0 and 1, suggesting minimal cross-shopping across all three markets. Additionally, overlap indices cluster near 0, further indicating weak cross-shopping behavior. Together, we interpret these findings as evidence of limited spillover effects between the parent's chains, suggesting that incorporating these effects into pricing decisions is not of primary importance.

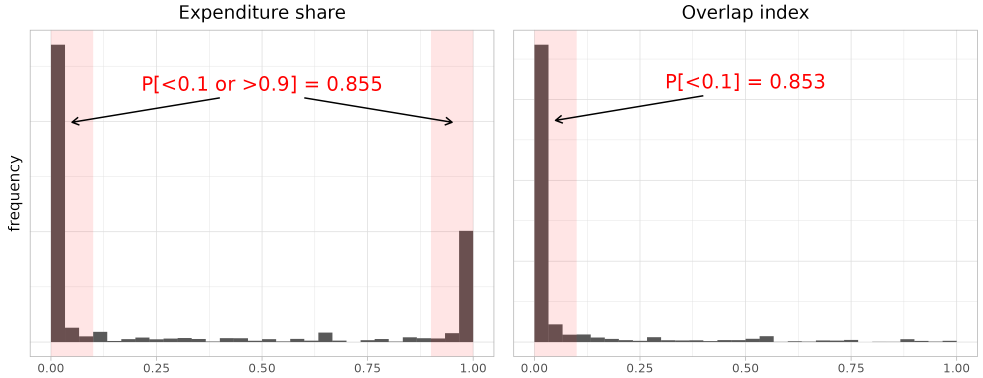
There are at least two possible explanations for the low levels of cross-shopping. First, these chains differ in format and target consumers with distinct price sensitivities and quality preferences. Consequently, consumers may be more likely to substitute for a competitor's chain of the same format rather than the parent's chain of a different format. Second, this analysis conditions on the observed store locations. When determining store placement, the parent company likely accounts for potential cannibalization effects and avoids opening stores too close to one another. As a result, stores from both chains may be strategically located at sufficient distances to minimize direct competition.

Figure A.6: Spillovers between traditional and discounter chains of parent 26

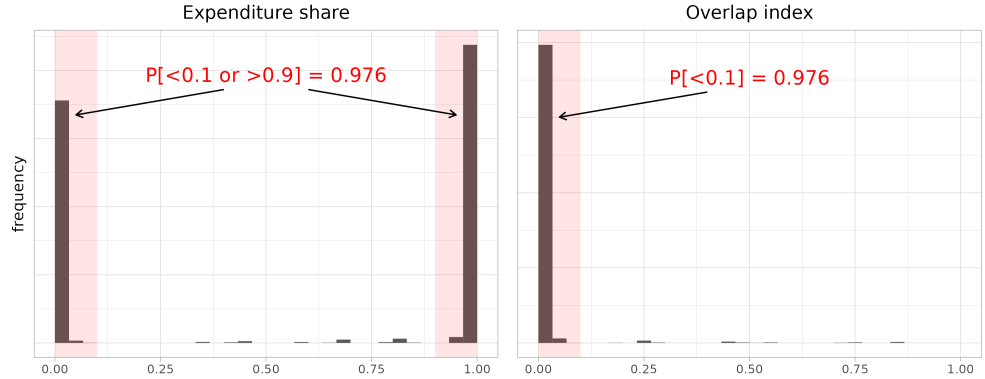
(a) Market α



(b) Market β



(c) Market γ

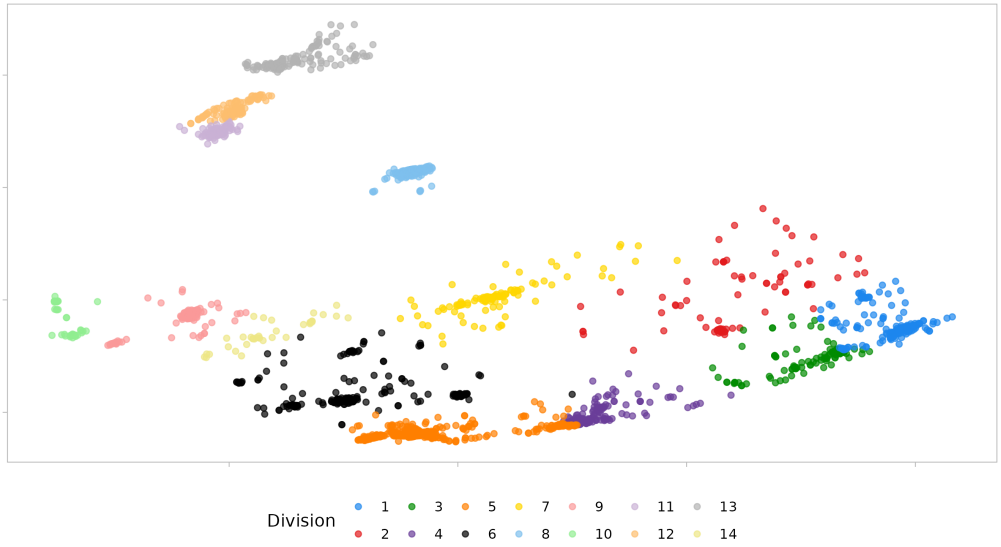


Notes: Panels show distributions of expenditure shares and overlap indices across consumers in the HMS panel. Each consumer is weighted by the product of the projection factor supplied by NielsenIQ and their total expenditures in these chains. Red text shows the probability of share lying in the region $[0, 0.1] \cup [0.9, 1]$, and the overlap index being less than 0.1.

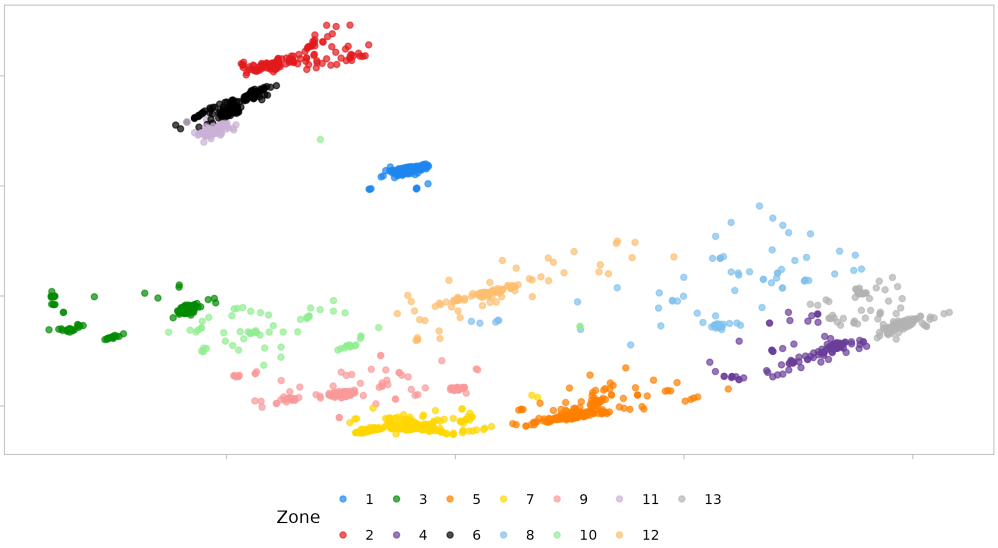
A.7 Comparing price zones to organizational divisions

Figure A.7: Organizational divisions and price zones of parent 11

(a) Organizational divisions



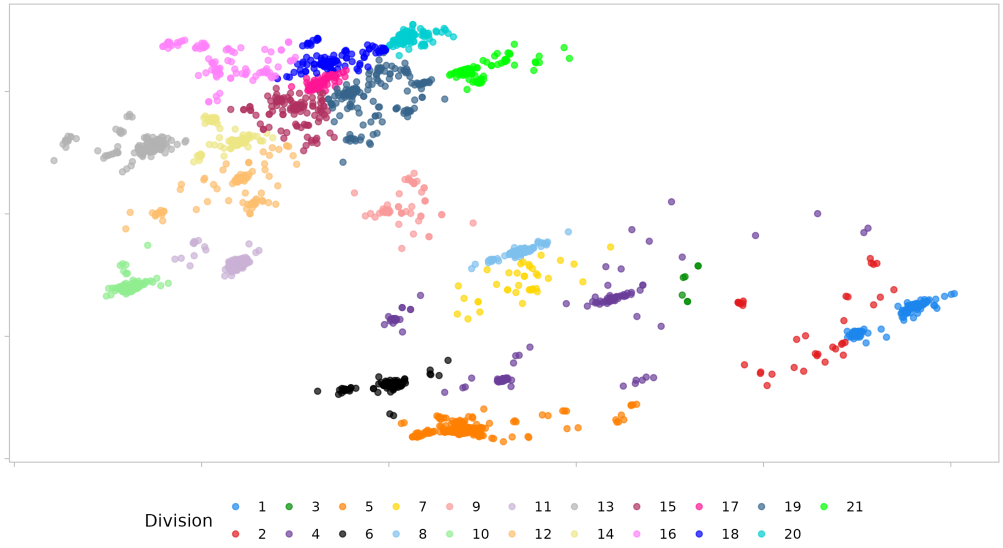
(b) Price zones



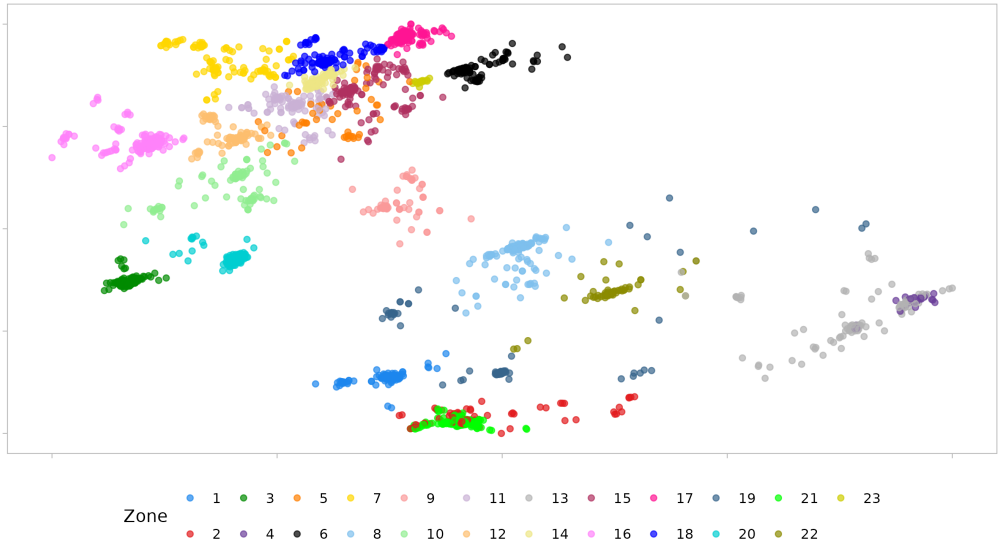
Notes: Each panel shows store locations for parent 11. In Panel (a) dots represent stores colored by organizational division, and in Panel (b) by price zone. The map is a linearly transformed version of the United States. Division data come from an external source independent of Nielsen.

Figure A.8: Organizational divisions and price zones of parent 26

(a) Organizational divisions



(b) Price zones

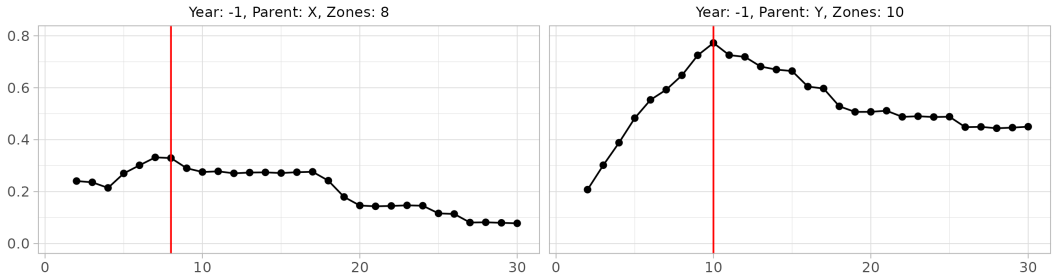


Notes: Each panel shows store locations for parent 26. In Panel (a) dots represent stores colored by organizational division, and in Panel (b) by price zone. The map is a linearly transformed version of the United States. Division data come from an external source independent of Nielsen.

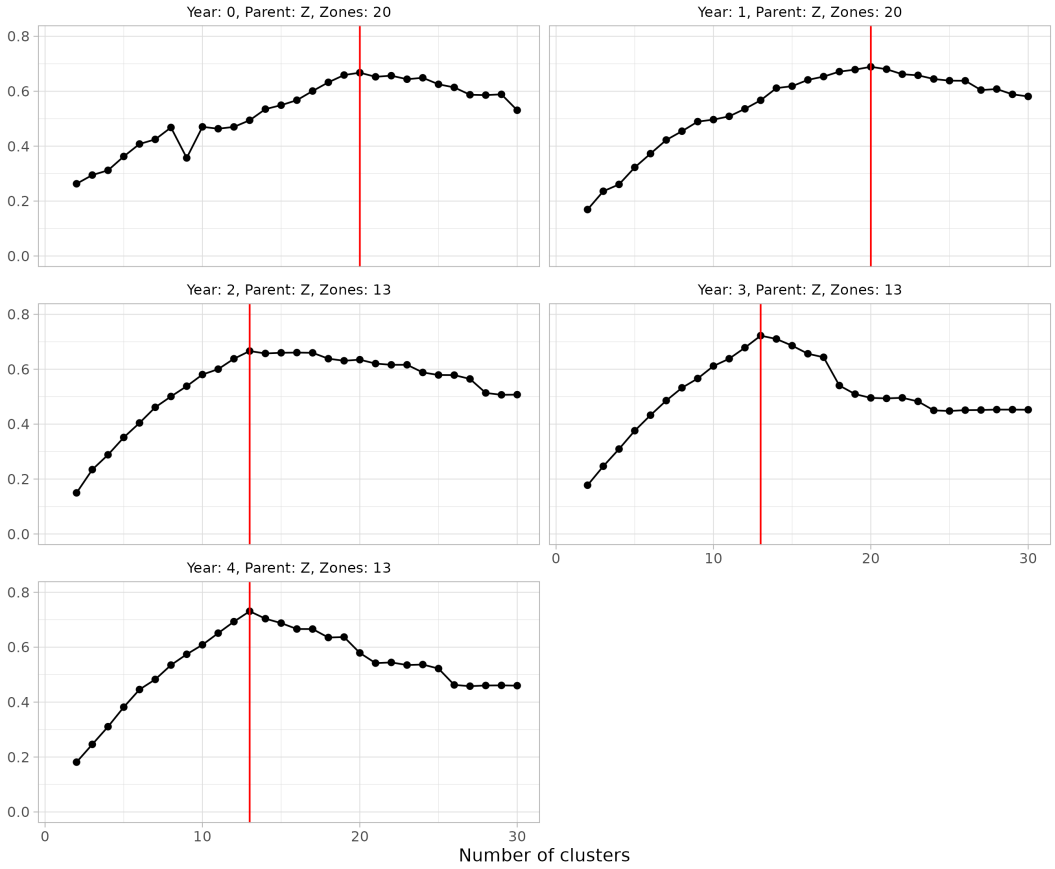
A.8 Clustering fit for merger analysis

Figure A.9: Silhouette Coefficients for the merger analysis

(a) Before merger



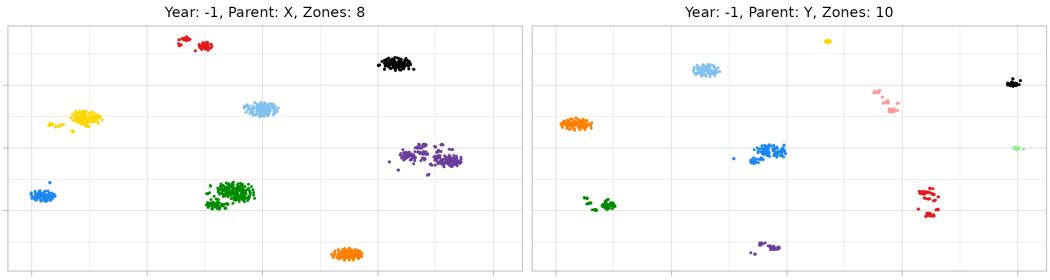
(b) After merger



Notes: The red line represents the chosen number of clusters. It maximizes the Silhouette Coefficient, except for the parent X, where we chose 8, which has a slightly lower coefficient compared to the argmax 7, to be consistent with the clear visual evidence from t-SNE.

Figure A.10: t-SNE for the merger analysis

(a) Before merger



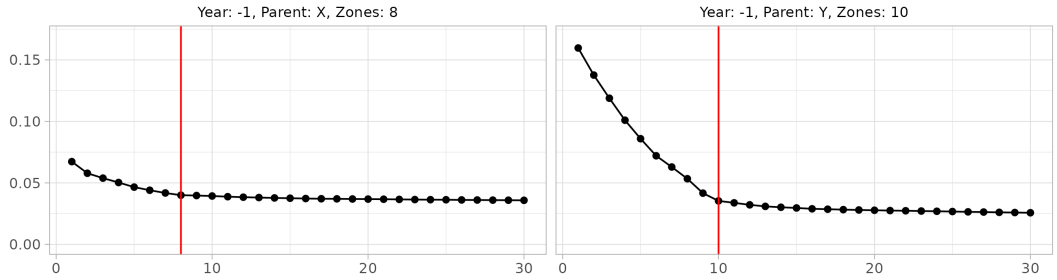
(b) After merger



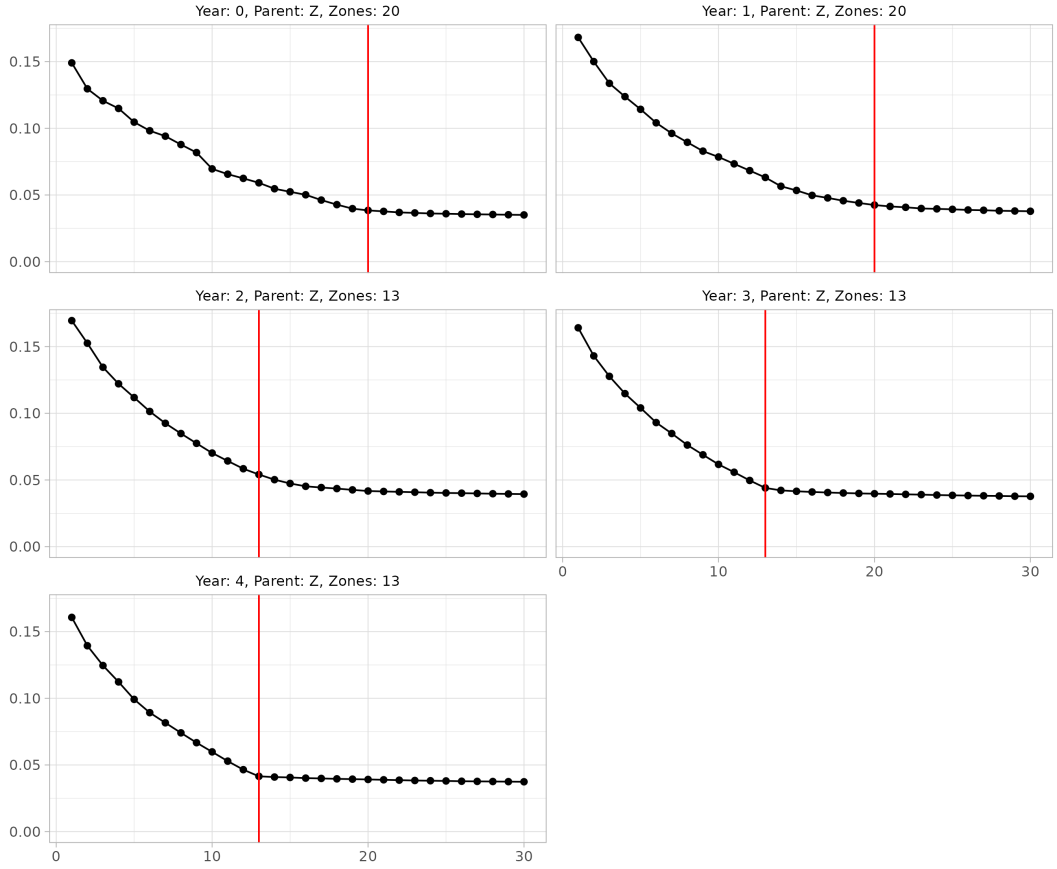
Notes: Each panel displays the results of t-SNE, which reduces each store's high-dimensional price vector to two coordinates for visualization. Colors represent the store's cluster when running PAM on the original price data. The close alignment between visually separated clusters in the t-SNE plot and the PAM-based colorings suggests that two distinct methods capture the same underlying structure, indicating that the clustering is robust.

Figure A.11: Mean absolute log price difference for the merger analysis

(a) Before merger



(b) After merger



Notes: The panels plot $V_r(N)$ against the number of clusters N for the two parents before the merger, and for the newly formed parent after the merger. The red line represents the chosen number of clusters.

A.9 Temporal stability of price zones

Table A.1: Comparing price zones across years using Adjusted Rand index

	11	42	36	20	26	23	22	31
Years 2017–2018	0.99	—	0.15	—	0.86	1.00	—	0.81
Years 2018–2019	1.00	0.94	1.00	1.00	1.00	1.00	1.00	1.00
Years 2019–2020	1.00	0.95	1.00	1.00	0.95	1.00	0.89	1.00

Notes: For each parent company listed in Table 3 as a zone pricer, we perform the clustering procedure separately by year from 2017 to 2020 to examine the temporal stability of price zones. The number of clusters in each year is chosen using the Silhouette Coefficient, except for parent 26, where it is fixed at 23 for reasons outlined in Section 4.2. We then compute the Adjusted Rand index between the zone assignments in consecutive years (2017–2018, 2018–2019, and 2019–2020). The index ranges from 0 to 1, with higher values indicating greater similarity of price zones across years, 0 meaning no alignment, and 1 indicating perfect alignment. Data for 2017 are unavailable for three parent companies, resulting in missing values for those comparisons. Parent 36 practiced uniform pricing in 2017 ($V_r(1) = 0.025$), so the low Rand index in that year reflects the absence of meaningful zone structure rather than temporal instability.

A.10 Proof that β is the weighted average of β_j

Lemma. Consider scalar random variables Y, X, and a discrete random variable W. Let α_W denote the slope of the best linear predictor (BLP) of Y on X conditional on W:

$$\alpha_W \equiv \frac{\mathbb{E}[YX|W]}{\mathbb{E}[X^2|W]}$$

Let α denote the unconditional slope of BLP of Y on X:

$$\alpha \equiv \frac{\mathbb{E}[XY]}{\mathbb{E}[X^2]}$$

Then α is the weighted average of α_W with non-negative weights, with the expectation of one. That is, there exists a function ν such that:

$$\alpha = \int \alpha_w \nu(w) dF_W(w)$$

$$\int \nu(w) dF_W(w) = 1$$

$$\forall w \in \text{supp}(W) : \nu(w) \geq 0$$

Proof. Define the residual $U \equiv Y - \alpha_W X$. It is uncorrelated with X conditional on W: $\mathbb{E}[UX|W] = 0$, and therefore unconditional as well: $\mathbb{E}[UX] = 0$. Substitute $Y = \alpha_W X + U$ into the definition of α to get:

$$
\begin{aligned}
\alpha &= \frac{\mathbb{E}[X(\alpha_W X + U)]}{\mathbb{E}[X^2]} \\
&= \frac{\mathbb{E}[\alpha_W X^2] + \overbrace{\mathbb{E}[XU]}^{=0}}{\mathbb{E}[X^2]} \\
&= \frac{\mathbb{E}[\alpha_W X^2]}{\mathbb{E}[X^2]} \\
&= \frac{\mathbb{E}[\mathbb{E}[\alpha_W X^2 | W]]}{\mathbb{E}[X^2]} \\
&= \frac{\mathbb{E}[\alpha_W \mathbb{E}[X^2 | W]]}{\mathbb{E}[X^2]} \\
&= \mathbb{E}\left[\alpha_W \underbrace{\frac{\mathbb{E}[X^2|W]}{\mathbb{E}[X^2]}}_{\equiv \nu(W)}\right] \\
&= \mathbb{E}[\alpha_W \nu(W)]
\end{aligned}
$$

So, we got that α is a weighted average of α_W. The weights are non-negative and integrate to 1 by the law of iterated expectations. Thus, the populational slope of BLP is the weighted average of conditional slopes. □

Now back to our regression:

$$\tilde{p}_{sjt} = \beta \cdot D_{st} + \gamma_{jt} + \delta_{sj} + \epsilon_{sjt}.$$

Let $\bar{D}_{sjt}$ be the residual from the regression of D_{st} on time-product and store-product fixed effects. Then by Frisch-Waugh-Lovell theorem β is the slope of the regression of $\tilde{p}_{sjt}$ on $\bar{D}_{sjt}$. By the same theorem β_j is the slope from regression of $\tilde{p}_{sjt}$ on $\bar{D}_{sjt}$ conditional on product j. Using the lemma above, we get the desired result that β is the weighted average of β_j with non-negative weights.

Using the formula of weights in the lemma, we can say a bit more. If there are no missing observations for any product (i.e., all stores at every week sold at least one item), then $\bar{D}_{sjt}$ does not vary across products, and the weights would be the same for every product. And since there are no missing observations, each product would have an equal number of observations. Thus, the weighting would simply recover the mean of β_j. In our case, however, there are some missing

observations, and this result does not hold exactly. But since the number of missing observations is low, β is still very close to the mean β_j as shown in Table A.2.

Table A.2: Comparison of β and mean of β_j for parent A

Tax regime	Philadelphia	Boulder	Cook County	Seattle	Albany, Oakland, SF
Mean of β_j	1.462	1.721	0.069	1.916	0.805
β	1.458	1.721	0.069	1.941	0.804

A.11 Alternative specification for the soda tax

Let $\bar{\tilde{p}}_{st} = \frac{1}{|J|} \sum_{j=1}^{J} \tilde{p}_{st}$ be the mean price per ounce across products for store s and week t. The regression is:

$$\bar{\tilde{p}}_{st} = \beta \cdot D_{st} + \gamma_t + \delta_s + \epsilon_{st}$$

Estimation results are in Table A.3 and are similar to our main specification.

Table A.3: Response of parent A to soda taxes. Estimation results of the pooled regression.

Dependent Variable:	Price (cents / ounce)				
Tax regime	Philadelphia	Boulder	Cook County	Seattle	Albany, Oakland, SF
D_{st}	1.710***	1.752***	0.0693***	2.242***	0.8686***
	(0.0262)	(0.0097)	(0.0036)	(0.1359)	(0.0186)
Store FE	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes
Observations	16,065	12,810	22,326	18,304	34,992

Notes: Clustered (store) standard errors in parentheses. Signif. Codes: ***: 0.001, **: 0.01, *: 0.05.

A.12 Soda tax results for parent B

Table A.4: Response of parent B to soda taxes. Estimation results

Dependent Variable:	Price (cents / ounce)				
Tax regime	Boulder	Cook County	Seattle 1	Seattle 2	Albany, Oakland, SF
D_{st}	0.1609***	-0.0170*	0.0763**	0.0470***	-0.0642
	(0.0126)	(0.0070)	(0.0269)	(0.0104)	(0.0418)
Store × product FE	Yes	Yes	Yes	Yes	Yes
Week × product FE	Yes	Yes	Yes	Yes	Yes
Observations	238,755	25,796	188,831	89,099	261,157
Tax rate	2	1	1.75	1.75	1
Implied pass-through rate	108.0%	98.3%	104.4%	102.7%	93.6%
	(0.63%)	(0.70%)	(1.54%)	(0.59%)	(4.18%)

Notes: This is the analog of Table 5 for parent B. There are no stores in Philadelphia, so there are no results for that market. In Seattle, there are two zones, and each is analyzed separately as if they were different tax markets. The implied pass-through rate is calculated under the assumption that the tax is added at the register and not included in the shelf price. Standard errors in parentheses are clustered at the store level. Signif. Codes: ***: 0.001, **: 0.01, *: 0.05

A.13 Examples of soda taxes on shelf tags and feature fliers

Figure A.12: An example of tax not being included in the shelf price

(a) Taxed product



(b) Non-taxed product



Notes: Panels (a) and (b) show price tags for regular and diet soda, respectively, from a parent company that does not include soda taxes in the shelf price. The only difference between the price tags is the fine print "*Plus imposed sugar expense*" under the price in Panel (a), indicating that the soda tax is added at checkout. Both products have the same shelf price and would be recorded as having the same price in the NielsenIQ RMS data.

Figure A.13: Examples of taxed products on feature fliers

(a) Soda prices in taxed stores



(b) Soda prices in non-taxed stores



(c) Featured soda quantity discounts



Notes: This figure illustrates how a parent company that includes the soda tax in the shelf price presents taxed products on feature fliers. Panels (a) and (b) show the same promotional flyer with and without a disclaimer indicating that the promoted price does not include the soda tax. The disclaimer appears only in stores subject to the tax. Panel (c) shows an alternative strategy using quantity discounts that avoid listing a specific price, allowing the same flyer to be used in both taxed and non-taxed stores.

A.14 Uniformity of features within price zones

Along with price information, NielsenIQ RMS data records whether the product was featured in a given store. These include ads in local newspapers, inserts, free-standing circulars, and ads from the retailer's website (Nielsen, 2021). The issue is that the feature data is recorded only for a subset of "audited" stores. This subset appears to be small. For example, 88% of stores of parent 26 seem to have never been audited in 2019. These stores, in all weeks for all 800 products, have NAs in the feature field. There are two questions regarding features. Firstly, since features increase demand and are often paired with price discounts, they are a possible source of omitted variable bias in demand estimation. The missing data problem makes it hard to directly include indicators for features in the regressions, so we should have a way of imputing missing observations. Secondly, the coordination of features may be of interest in itself. To our knowledge, there was no research describing at what level the decisions about product features are made.

Given our finding that price zones coincide with organizational divisions, it is natural to hypothesize that feature promotions are also determined at the division level. Denote NielsenIQ's data about whether the product was promoted as F_{sjt}. It can be either NA, 0, or 1. If promotions are coordinated at the division level, then at every week-zone pair, we should see that either all or none of the stores feature the product. Denote $s(z, j, t)$ as the set of stores that have non-missing feature data in zone z for product j in week t and define:

$$G_{zjt} = \frac{1}{|s(z,j,t)|} \sum_{s \in s(z,j,t)} F_{sjt}.$$

If our hypothesis is correct, the distribution of G_{zjt} must have high mass on 0 and 1. Note that if $|s(z, j, t)| = 1$, then mechanically G_{zjt} can only be 0 or 1, so to mitigate this concern we limit our attention to observations with $|s(z, j, t)| \geq 5$. We also experimented with setting this threshold to 1 (i.e., no restrictions) and 10, the results are virtually identical. Table A.5 presents the probability of G_{zjt} being equal to 0, 1, and the sum of these two probabilities for parents that engage in zone pricing.

The first two rows show that the distribution of G_{zjt} is not degenerate at 0 or 1, which would make our analysis nonsensical. Parents feature products rather often, apart from parents 22 and

Table A.5: Probability mass of G_{zjt} at 0 and 1 across parents under different sample restrictions

	11	20	22	23	26	31	36	42
Restricting to sample with $\lvert s(z,j,t)\rvert \geq 5$:								
$P(G_{zjt}=0)$	0.781	0.816	0.984	0.899	0.839	0.878	0.857	0.845
$P(G_{zjt}=1)$	0.185	0.169	0.016	0.083	0.148	0.116	0.135	0.105
Total	0.966	0.985	0.999	0.981	0.987	0.995	0.991	0.951
Restricting to sample with $\lvert s(z,j,t)\rvert \geq 1$:								
$P(G_{zjt}=0)$	0.785	0.817	0.986	0.906	0.846	0.897	0.859	0.859
$P(G_{zjt}=1)$	0.182	0.168	0.014	0.077	0.141	0.099	0.132	0.094
Total	0.966	0.985	0.999	0.983	0.987	0.995	0.992	0.953
Restricting to sample with $\lvert s(z,j,t)\rvert \geq 10$:								
$P(G_{zjt}=0)$	0.772	0.818	0.980	0.822	0.829	0.823	0.848	0.836
$P(G_{zjt}=1)$	0.193	0.170	0.019	0.135	0.158	0.177	0.144	0.129
Total	0.965	0.988	0.999	0.958	0.986	1.000	0.992	0.965

Notes: Column names correspond to parent IDs.

25. The third row presents the check of our hypothesis that features are determined at the division level. The distribution of G_{zjt} has a high mass on 0 and 1, ranging from 0.951 to 1 across all zone-pricing parents. Note that we get zones from our clustering algorithm, which may not perfectly align with the true price zones and divisional structure. This estimation error causes some G_{zjt} not to be exactly equal to 0 or 1 even if features are perfectly coordinated at the divisional level. Despite this, we still find a remarkable degree of coordination on features within the uncovered price zones. This leads us to two conclusions.

Firstly, a useful strategy for imputing missing feature observations is to do so at the zone level. That is, if F_{sjt} is missing, impute it with the mode of non-missing $F_{s'jt}$ where s' are stores belonging to the price zone of store s. It is feasible as long as there is at least one such store s'. For parent 26, this imputation reduces the share of missing observations from 88% to 0.3% and for parent 11 from 84% to 0.1%. The remaining small share of missing observations is assumed to be 0, i.e., no feature.

Secondly, decisions about product features seem to be made at the same level as general pricing decisions—at the organizational division level. One explanation for why this is so is cost synergies.

Since features often come with price discounts, decisions about product features and pricing are connected, and having them made at the same level may be cost-effective.

A.15 Robustness of pricing optimality

Table A.6: Robustness checks for $\hat{\lambda}_{zj}$

Specification	OLS	IV	F-stat	N
OLS elasticities	0.0411***	0.4760***	596.3	24,191
	(0.0050)	(0.0274)		
Admissible elasticity values:				
$\hat{\eta}_{zj} \in (-10, -1.2)$	0.0500***	0.7133***	362.2	18,817
	(0.0041)	(0.0541)		
$\hat{\eta}_{zj} \in (-4, -1.2)$	0.0430***	0.6896***	294.4	16,308
	(0.0040)	(0.0582)		
Admissible elasticity standard error:				
SE < 1	0.0406***	0.7521***	314	22,897
	(0.0032)	(0.0610)		
SE < 0.1	0.0831***	0.6518***	467.9	9,530
	(0.0075)	(0.0469)		

Notes: This table reports robustness checks for the results in Table 6. The first row uses elasticities estimated with the OLS specification. The second group varies the admissible range of elasticity values, and the third group varies the maximum admissible standard error of elasticities. Across specifications, the IV estimates remain similar to our main result. All regressions include parent × product fixed effects. Clustered (parent × product) standard errors in parentheses. Signif. Codes: ***: 0.001, **: 0.01, *: 0.05.

A.16 Can wholesale costs vary across markets?

Anderson and Fox (2019) argue that retail promotions are usually a consequence of trade promotions. One way to assess whether trade promotions, and thus marginal costs, can vary across stores is to examine whether the same product is discounted in one price zone but sold at a regular price in another zone during the same week. This pattern would be consistent with the parent negotiating a trade promotion, such as a scanback, in the promoted zone that lowers its marginal cost, while the other zone receives no such allowance facing a regular wholesale cost.

We proceed as follows. For a given parent and product, we define a zone-week as promotional if the weekly average price is at least 30 percent lower than the product's median price in that zone. For each week, we record whether the product is on promotion in at least one zone and not on promotion in at least one other zone. We then count the number of such weeks for each product and summarize the distribution across products.

Table A.7: Assessing variation in wholesale costs using price promotions

	11	20	22	23	26	31	36	42
Share with $\geq$ 1 week	0.78	0.26	0.48	0.25	0.65	0.38	0.25	0.43
Share with $\geq$ 5 weeks	0.61	0.10	0.26	0.09	0.34	0.17	0.09	0.21
Share with $\geq$ 10 weeks	0.49	0.05	0.10	0.04	0.18	0.07	0.01	0.11

Notes: Column names correspond to parent IDs. For each product, we count the number of weeks when it is on promotion in at least one zone and not on promotion in at least one other zone. The table reports the share of products with at least 1, 5, or 10 such weeks.

Table A.7 reports the share of products that exhibit at least one, five, or ten weeks in which they are on promotion in some zones but not others within the same parent. There is substantial heterogeneity across parent companies: the share of products with at least one such week ranges from 0.25 to 0.78; with at least five such weeks, from 0.09 to 0.61; and with at least ten such weeks, from 0.01 to 0.49. We interpret these patterns as evidence that some parents engage in zone-specific trade promotions with manufacturers, resulting in varying effective wholesale costs across parent's stores.

References

Adams, Brian and Kevin R Williams, "Zone pricing in retail oligopoly," *American Economic Journal: Microeconomics*, 2019, *11* (1), 124–156.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek, "When does coordination require centralization?," *American Economic Review*, 2008, *98* (1), 145–179.

Anderson, Eric T and Edward J Fox, "How price promotions work: A review of practice and theory," *Handbook of the Economics of Marketing*, 2019, *1*, 497–552.

Bertrand, Marianne and Emir Kamenica, "Coming Apart? Cultural Distances in the United States over Time," *American Economic Journal: Applied Economics*, 10 2023, *15* (4), 100–141.

Besanko, David, Jean-Pierre Dubé, and Sachin Gupta, "Own-brand and cross-brand retail pass-through," *Marketing Science*, 2005, *24* (1), 123–137.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, "Americans do IT better: US multinationals and the productivity miracle," *American Economic Review*, 2012, *102* (1), 167–201.

Boston Consulting Group, "Successful Merger Integration: How BCG Helps Maximize the Value from a Deal (available at `https://web.archive.org/web/20250327192708/http://media-publications.bcg.com/flash/PDF/PMI.pdf`)," 2021.

Butters, R Andrew, Daniel W Sacks, and Boyoung Seo, "How do national firms respond to local cost shocks?," *American Economic Review*, 2022, *112* (5), 1737–1772.

Buzzell, Robert D and John A Quelch, "The costly bargain of trade promotion.," *Harvard Business Review*, 1990, *68* (2), 141–149.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and taxation: Theory and evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Chintagunta, Pradeep K, Jean-Pierre Dubé, and Vishal Singh, "Balancing profitability and customer welfare in a supermarket chain," *Quantitative Marketing and Economics*, 2003, *1*, 111–147.

Cook County, Illinois Department of Revenue, "Sweetened Beverage Tax Regulation 2017-2: Inclusion of Tax in the Sale Price," 2017.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in us retail chains," *The Quarterly Journal of Economics*, 2019, *134* (4), 2011–2084.

Dessein, Wouter, "Authority and communication in organizations," *The Review of Economic Studies*, 2002, *69* (4), 811–838.

_ , Desmond Lo, and Chieko Minami, "Coordination and organization design: Theory and micro-evidence," *American Economic Journal: Microeconomics*, 2022, *14* (4), 804–843.

Drèze, Xavier and David R Bell, "Creating win–win trade promotions: Theory and empirical analysis of scan-back trade deals," *Marketing Science*, 2003, *22* (1), 16–39.

Eizenberg, Alon, Saul Lach, and Merav Oren-Yiftach, "Retail Prices in a City," *American Economic Journal: Economic Policy*, 5 2021, *13* (2), 175–206.

Ellickson, Paul B and Sanjog Misra, "Supermarket pricing strategies," *Marketing Science*, 2008, *27* (5), 811–828.

Finlay, Jessica, Mao Li, Michael Esposito, Iris Gomez-Lopez, Anam Khan, Philippa Clarke, and Megan Chenoweth, "National Neighborhood Data Archive (NaNDA): Grocery Stores by Census Tract, United States, 2003-2017," 2020.

FMI, "The Food Retailing Industry Speaks 2024," Technical Report, FMI – The Food Industry Association, Arlington, VA 2024.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely, "Are US industries becoming more concentrated?," *Review of Finance*, 2019, *23* (4), 697–743.

Hausman, Jerry A, "Valuation of New Goods under Perfect and Imperfect Competition," in "The Economics of New Goods," University of Chicago Press, 1996, pp. 207–248.

Hitsch, Günter J, Ali Hortacsu, and Xiliang Lin, "Prices and promotions in US retail markets," *Quantitative Marketing and Economics*, 2021, pp. 1–80.

Hoch, Stephen J, Byung-Do Kim, Alan L Montgomery, and Peter E Rossi, "Determinants of store-level price elasticity," *Journal of Marketing Research*, 1995, *32* (1), 17–29.

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams, "Organizational Structure and Pricing: Evidence from a Large U.S. Airline," *The Quarterly Journal of Economics*, 09 2023, p. qjad051.

Hubert, Lawrence and Phipps Arabie, "Comparing partitions," *Journal of Classification*, 1985, *2*, 193–218.

Kaufman, Leonard and Peter J. Rousseeuw, "Partitioning Around Medoids (Program PAM)," in "Finding Groups in Data: An Introduction to Cluster Analysis," New York: Wiley, 1990, chapter 2, pp. 68–125.

Keller, Kristopher O, Jonne Y Guyt, and Rajdeep Grewal, "Soda taxes and marketing conduct," *Journal of Marketing Research*, 2024, *61* (3), 393–410.

Leider, Julien, Andrea Pipito, and Lisa M Powell, "The impact of the Cook County, Illinois, sweetened beverage tax on prices, 2017," 2022.

Malenko, Nadya, "Information flows, organizational structure, and corporate governance," in "Handbook of Corporate Finance," Edward Elgar Publishing, 2024, pp. 511–546.

McKinsey & Company, "Perspectives on Merger Integration (available at `https://web.archive.org/web/20250817221622/https://www.mckinsey.com/client_service/organization/latest_thinking/~/media/1002A11EEA4045899124B917EAC7404C.ashx`)," June 2010.

Miravete, Eugenio J, "On Anticompetitive Third-Degree Price Discrimination," 2024.

Nielsen, "Retail Scanner Dataset Manual," 2021.

Ozturk, O. Cem, Sriram Venkataraman, and Pradeep K. Chintagunta, "Price Reactions to Rivals' Local Channel Exits," *Marketing Science*, 2016, *35* (4), 588–604.

PwC, "Transact to Transform: PwC's 2023 M&A Integration Survey (available at `https://web.archive.org/web/20250723134448/https://www.pwc.com/us/en/services/consulting/deals/library/ma-integration-survey.html`)," 2023.

Rantakari, Heikki, "Governing adaptation," *The Review of Economic Studies*, 2008, *75* (4), 1257–1285.

Rousseeuw, Peter J, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, 1987, *20*, 53–65.

Seiler, Stephan, Anna Tuchman, and Song Yao, "The impact of soda taxes: Pass-through, tax avoidance, and nutritional effects," *Journal of Marketing Research*, 2021, *58* (1), 22–49.

Springer, Jon, "Kroger Adds New Divisions in South," *Supermarket News*, July 2013. [Accessed January 15, 2025].

Stroebel, Johannes and Joseph Vavra, "House prices, local demand, and retail prices," *Journal of Political Economy*, 2019, *127* (3), 1391–1436.

van der Maaten, Laurens and Geoffrey Hinton, "Visualizing data using t-SNE," *Journal of Machine Learning Research*, 2008, *9* (11).

Zwiebach, Elliot, "Diverting Paths," *Supermarket News*, November 18 1996. [Accessed May 26, 2025].