Shrinkflation and Consumer Demand*

Aljoscha Janssen† Johannes Kasinger‡

July 22, 2024

Abstract

This paper investigates shrinkflation—the phenomenon of reducing product size while maintaining or slightly changing prices—in the U.S. retail grocery market. We analyze a decade of retail scanner data to assess the prevalence and patterns of product size changes across various product categories. Our findings show that around 1.8% of products have experienced product downsizing, and it is at least five times more prevalent than product size increases. Product downsizing typically occurs without a corresponding decrease in price and is widespread across product categories. Consequently, consumers end up paying more per unit volume. The study also reveals that shrinkflation is a common strategy not just among manufacturers but also at the retail level. We further find that consumers are more responsive to price adjustments than to changes in product size. This finding suggests that reducing product sizes is an effective strategy for retailers and manufacturers to increase margins or respond to cost pressures, offering valuable implications for retailers and policymakers.

Keywords: Shrinkflation, Product Size, Package-downsizing, Inattention
JEL Codes: D12, L10, M31

---

*Researcher's own analyses are calculated (or derived) based in part on data from NielsenIQ 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 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.

†Singapore Management University, ajanssen@smu.edu.sg

‡Tilburg School of Economics and Management and Leibniz Institute for Financial Research SAFE, j.kasinger@tilburguniversity.edu

1 Introduction

Retail grocery pricing is traditionally characterized by frequent changes, where price adjustments are the primary tool for retailers and manufacturers to increase their margins. However, current discussions have shifted to exploring another aspect: product size changes as a means of increasing prices per volume. This trend, widely referred to as 'shrinkflation', has been increasingly highlighted in global news, suggesting a tendency among firms to subtly raise effective prices by reducing product sizes (Lempert, 2023; Barrett and Rachwani, 2023; BBC News, 2023). Shrinkflation has caught the attention of consumer protection agencies and political circles (CNBC, 2024; Toeniskoetter, 2022; European Parliament, 2022), prompting some supermarkets to introduce warning labels to underscore these practices (Vidalon, 2023).[1]

For retailers and manufacturers, product downsizing emerges as an additional strategy to inflate profit margins, potentially contributing to the broader trend of rising retail margins and profits (Döpper et al., 2022). From a behavioral economics perspective, product downsizing can incur high costs for consumers who are inattentive to product sizes or find product comparison costly. Shrinkflation is similar to other behavioral pricing practices that exploit consumer underreactions by hiding or obfuscating price increases. Examples of such practices include complex add-on pricing schemes (e.g., Ellison, 2005; Gabaix and Laibson, 2006) and hidden shipping and handling costs (e.g., Hossain and Morgan, 2006; Brown et al., 2010). Moreover, shrinkflation could disproportionately affect low-income households and their ability to afford groceries (Davidson, 2023; Danley, 2022). Consequently, shrinkflation has important policy implications for consumer protection. Yet, the extent and consumer response to shrinkflation remain underexplored.

This paper investigates trends, characteristics, and implications of product size adjustments within the U.S. retail grocery market over the past decade. We document the prevalence of product size decreases compared to increases across different product groups and examine associated price adjustments. Additionally, the study examines whether these size changes are implemented uniformly by manufacturers at a national level or are more localized decisions made by retailers. We further examine consumer sensitivity to size changes relative to price adjustments and how these two factors interact.

In the first part of this article, we use the full sample of product offerings in the U.S. retail

[1] One reason for the heightened media focus on shrinkflation may be that consumers perceive product downsizing as more inequitable than price increases (Evangelidis, 2024).

grocery market covered by store-level scanner data from the Nielsen-Kilts retail panel. Our results indicate that product downsizing is not confined to recent years but a consistent trend across all observed years. We find that reductions in product size occur at least five times more often than increases.[2] Among all observations (weeks of non-zero sales for a specific store-product pair), 1.81% of the observations pertain to products that experienced a size decrease in a specific store during our observation period, while only 0.33% of observations relate to size increases. The proportions are very similar when considering the total sales of affected products. Our findings show that size decreases are widespread across various product categories. They are particularly prominent in categories such as candy and snacks but also extend commonly to categories like detergents and hair care products.

Next, our research shows that size decreases usually occur with little or no accompanying price changes. This price stickiness is not apparent for product size increases, which are typically associated with price hikes. A closer analysis of the absolute changes in price per volume highlights that product size decreases generally result in consumers ending up paying more for the same volume. On average, the price per volume in the year after a product size reduction in a store is around 12% higher than in the year before the size reduction. This pattern is observable across all product categories except liquor. Personal care products, particularly cosmetics, exhibit the highest price per volume increases, averaging over 40%. Conversely, product size increases reduce the price per volume for consumers in most product categories. On average, the price per volume decreases by around 2% following a product size increase.

Examining absolute purchase volume one year before and after a product size reduction, we find that consumers, on average, reduce their purchased volume by about 5%, with substantial heterogeneity across product categories. However, revenues for sellers still significantly increase in most product categories, generating, on average, around 6% more revenues in the year after a size reduction than in the year before. For product size increases, purchase volumes tend to rise by a greater extent than the price per volume decreases, yielding a rise in total revenues.

Moreover, we show that product size changes are typically not national occurrences affecting all markets simultaneously but rather long-running developments that vary across individual stores and retail chains. Our analysis reveals a compelling pattern: while manufacturers occasionally initiate size reductions by entirely replacing products with smaller versions, they are not

[2] When identifying instances of product size changes, we aim to identify each occurrence within a store where a product is permanently replaced by an equivalent product of a different size. This new product must share the same brand, brand description, and product description, and its size must fall within a specified range of the original—33% larger for increases and 25% smaller for decreases.

the predominant driver. Size reductions more commonly occur at the retail level, with stores choosing to substitute a manufacturer's product with a smaller version while the larger product continues to be available in other stores. Thus, we conclude that both retailers and manufacturers contribute to the widespread phenomenon of shrinkflation within the U.S. retail grocery market, which contradicts reports that primarily blame manufacturers.

As a final step, we empirically analyze consumer demand responses to product price and size variations. We estimate the elasticities associated with price and size by using data on weekly sales at the store-product level, alongside corresponding product sizes and prices. Our findings indicate a significantly higher sensitivity of the average consumer to price fluctuations as opposed to changes in product size. Our model predicts that a 1 percent hike in price leads to a 1.14 percent decline in sales on average, whereas a similar rise in product size tends to increase sales by approximately 0.05 percent. To tackle potential endogeneity concerns of prices and package size, we use an instrumental variable approach similar to DellaVigna and Gentzkow (2019). We instrument the product's price and size with the corresponding average values observed in different stores of the same chain but located outside the focal market area. In this analysis, the price elasticity is -1.83, while the estimate of the package size elasticity is 0.19. The observed disparity in price and size elasticities offers important insights for retailers, manufacturers, and policymakers. For the supply side, they imply that reducing package sizes may be a more effective strategy than raising prices to increase margins or address cost pressures. Conversely, assuming that preferences for specific product sizes do not drive consumers' responses, the results suggest that consumers tend to underreact to product size changes, which provides an argument for regulating excessive shrinkflation practices to protect consumers.

To further investigate the underlying driving factors of the disparities between size and price elasticities, we analyze whether price and product size elasticities differ for downsized and upsized products. Using a general least square and an instrumental variable approach, we find that price elasticities between downsized and upsized products are not statistically different from each other. Second, for downsized products, the size elasticity is close to zero, while price and size elasticities closely align for upsized products. The results are consistent with the idea that retailers and manufacturers strategically hide product size decreases while making size increases more salient. This finding echoes numerous studies showing that consumers tend to underreact to non-salient attributes of goods (DellaVigna, 2009; Chetty et al., 2009) and that firms exploit this tendency (Hossain and Morgan, 2006; Ellison and Ellison, 2009; Brown et al., 2010).

Our study is not the first that investigate shrinkflation and consumer responses to changes in package size. In one of the earliest studies on downsizing, Çakır and Balagtas (2014) show that consumers in the Chicago ice cream market are more sensitive to price in comparison to product size changes. Çakır et al. (2013) presents evidence of the prevalence and demand effects of shrinkflation in the peanut butter and shelf-stable tuna categories, and Çakır (2022) documents that price pass-through of wholesale costs in the ice cream market is strongly associated with product downsizing. Kim (2024) examines the South Korean milk market and finds a consumer preference for package downsizing. Meeker (2021) develops a structural model illustrating consumer inattention to product size decreases. Yonezawa and Richards (2016) study the effects of product downsizing on price competition in the cereal industry. Our study contributes to the literature as the first comprehensive analysis of the prevalence of size changes, the differences between size increases and decreases, and consumer reactions across various product groups.[3]

Additionally, there are several consumer behavior studies that examines consumers' ability to perceive changes in product sizes in experimental settings. Ordabayeva and Chandon (2013) find that consumers tend to underestimate increases in product sizes, whereas Chandon and Ordabayeva (2017) contend that consumers are acutely aware of product downsizing, successfully recognizing product shrinkage. Further, Yao et al. (2020) conduct field and laboratory experiments, uncovering that consumer responses to decreases and increases in product sizes are analogous when price information is accessible. Evangelidis (2024) investigates when and why shrinkflation behavior is perceived as unfair. Our research augments the existing body of work in consumer behavior by employing retail scanner data. Distinctively, our findings reveal significant disparities in consumer reactions to product size expansions and reductions. We observe that size expansions frequently coincide with price hikes, while reductions in size typically occur without any change in price. Moreover, our analysis suggests that consumers exhibit a markedly higher sensitivity to variations in price than to alterations in size, highlighting a critical dimension of consumer decision-making in the retail context.

Our analysis parallels a research note from the U.S. Bureau of Labor Statistics (McNair, 2023) that investigates occurrences of product upsizing and downsizing and its effect on the Consumer Price Index (CPI). Rojas et al. (2024) assesses product downsizing and its impact on inflation measures.[4] Their findings reveal that significant product downsizing influences inflation

[3]Recent studies further underline these findings, examining the effects of shrinkflation and consumer habits in the context of canned tuna. See Harris-Lagoudakis et al. (2023) and Webb et al. (2022) for detailed analyses.

[4]Related work by Ochirova (2017) illustrates the effect of product downsizing on inflation in the United Kingdom.

indices. We corroborate these findings, documenting a significant number of product reductions across different categories. In addition to those studies, we provide details on the characteristics of product shrinkage and related consumer responses.

More broadly, we document an empirical pattern within the U.S. retail grocery market. This research adds to influential papers that demonstrate uniform pricing behavior of various goods by retail chains across different geographies (DellaVigna and Gentzkow, 2019), highlight the heterogeneous impact of advertisement on consumer demand across products (Shapiro et al., 2021), document price dispersion and similarities in promotions across stores in the U.S. (Hitsch et al., 2021), and reveal the significance of pink-taxes across products (Moshary et al., 2023). Similar to those studies, we exploit rich scanner data that allow us to make broad conclusions across product categories, stores, and the U.S. retail sector at large. Our findings underscore the importance of size shrinking as a noteworthy phenomenon in the U.S. retail grocery market, with important implications for policy-makers and suppliers.

2 Data

The backbone of our analysis is the Nielsen Retail Scanner data provided by the Kilts Center at the University of Chicago, covering the years 2010 to 2020. We exclude observations prior to 2010 as product size changes are not recorded. The Nielsen Retail scanner documents weekly quantities and prices of approximately 4 million products sold across a spectrum of up to 50,000 retail establishments, including grocery stores, drugstores, mass merchandisers, and other types of retail outlets.

We limit our sample to products that are sold in a specific unit size. Thus, we exclude product groups that are not suitable due to non-standard sizes. Examples are deli products, fresh vegetables, or fruits. Our analysis requires the identification of products a retailer offers within a week. Given that the scanner data only records instances of product sales per week, we have constrained our dataset to encompass primarily leading brands within each of the 1,100 product modules. We operationalize 'leading brands' as those constituting 80 percent of the sales revenue within their respective module. The reduction permits a more accurate estimation of retailer stock assortments on a weekly basis. In Appendix A we show detailed summary statistics of the data set and the impact of the restrictions.

Given that also the price of a product is only observed in weeks with positive sales, we impute missing pricing following the approaches of Hitsch et al., 2021; Moshary et al., 2023 or Shapiro et al. (2021). For a discussion on the imputation approach and implication, see Appendix B. For the estimation of price and size elasticities of demand, we further only consider food, drugstores, and mass merchandise chains. The final sample and summary statistics are illustrated in Appendix A.

2.1 Identification of Product Size Change

A key challenge in our research framework is the accurate identification of products undergoing product size reduction or increases. We approach this process from the perspective of a regular consumer who visits a store and may not notice when a familiar product has been replaced by a similar item with a smaller or larger size. Our goal is to systematically identify each instance within a store where a product is permanently substituted by an equivalent product of a different size.

The substitution is essential since product assortment expansions have been prevalent over the past decade (Neiman and Vavra, 2023). When identifying products that have decreased or increased in size, we maintain a neutral stance regarding the underlying cause. These variations may result from manufacturers altering products—introducing new sizes and discontinuing old ones—or from retailers modifying their assortment to include different sizes from the same manufacturer. We explore these different scenarios in detail in Section 4.

Our identification process is as follows. First, we classify products as equivalent if they share the same brand, brand description, and UPC (Universal Product Code) description, requiring that these attributes be consistently reported within the same store. Equivalent products that change in size only differ in the UPC itself, in addition to the reported size. A detailed justification for using UPC changes as indicators of product size changes, and why ignoring size changes of products with the same UPCs does not undermine our analysis, is provided in Online Appendix D. Second, we establish that the two products are in substitution; that is, when the old product is removed from the store's inventory, it is replaced by a new product of a smaller or larger size. This substitution is not temporary but permanent. We anticipate some heterogeneity among stores—some may allow an overlap period to sell off old stock, while others may experience a temporary gap between the discontinuation of the old product and the introduction of the new one. Therefore, our third criterion is a maximum of either an eight-week overlap or an eight-week gap between the discontinuation of the old product and the introduction of the new one. Finally, we also require that the size changes are in a specific range, with an increase of at most 33.33% and a decrease of at most 25%.[5]

Our method could mistakenly classify a product as equivalent when size changes coincide with rebranding or package redesigns that do not alter the brand or UPC description. To reduce such errors and validate our methodology, we employ a UPC lookup database (`https://www.upcitemdb.com/`) for a visual inspection to accurately identify packaging and product details before and after size changes. Despite limitations in tracking older, discontinued products, the inspection allowed us to validate over 800,000 out of 1,300,000 instances of size reduction in inspected groups as identical products, supporting our analysis's reliability. The similarities in price adjustments and size reduction patterns between the overall data and the subset with confirmed identical products post-reduction further reassure us of our study's validity. Appendix C.1 details the visual inspection approach and results. Additionally, we use Nielsen's household scanner data to assess whether households consistently purchase the same products before and after a change in package size (for details, see Appendix C.2). We document a consistent buying pattern among panelists, which reassures us that products experiencing size changes are, in most instances, identical in every aspect apart from size.

Table 1 displays summary statistics for all products that underwent size changes, considering all sales before and after changes in stores where the products were substituted with newly-sized products. In comparison, Column 1 shows the summary stats considering all sales in the sample.

The data shows that decreases in product size are significantly more common than increases. We recorded 1,88 billion of observations (weeks of non-zero sales for a specific store-product pair) that refer to products that were downsized in a specific store during our observation period, generating about \$37.94 billion in sales.[6] In comparison, size increases occurred 0.34 billion times, totaling \$7.36 billion in sales. The average price of products that decreased in size was about \$3.49, while the average price for products that increased in size was \$4.62. This implies that the divergence in prevalence between product size decreases and increases becomes more evident when considering total units sold—14.49 billion for decreased versus 2.32 billion for increased.

Downsized products represent a considerable fraction of all transactions. Approximately

[5] We adjust the percentage decrease and increase to maintain equal relative changes between the maximally shrunk product and the old product, and between the old product and the maximally enlarged product. If we consider a 25%-change in both directions, the number of shrunk compared to increased products rises. The other results are stable.

[6] Note that unique UPCs affected by size changes, as reported in Table 1, provide an inflated measure of affected products because the same product has different UPCs before and after size changes occur, according to our definition.

1.86% of total sales are attributed to products downsized to smaller versions, with these products accounting for about 2.28% of total units sold. The proportion of unique store-product pairs affected by size reduction stands at 1.16%. In Appendix A, we extend the comparison to include only food stores, mass merchandisers, and drug stores—the sample used in the demand estimation. The descriptive statistics are comparable, except for variables referring to unique stores. For example, the ratio of unique store-product pairs affected by size reduction increases to 1.81%.

Table 1: Summary Statistics of Products with Size Changes

	All Products	Downsized Products	Upsized Products
Total Observations	103.84 Bn.	1.88 Bn.	0.34 Bn.
Unique UPCs	540,182	24,779	15,893
Unique Stores	59,345	58,634	55,554
Unique Retailers	141	141	141
Unique UPC-Store	1,554.92 Mn.	18.04 Mn.	4.09 Mn.
Total Sales (in $)	2,044.99 Bn.	37.94 Bn.	7.36 Bn.
Total Units	635.40 Bn.	14.49 Bn.	2.32 Bn.
Average Price (in $)	4.71	3.49	4.62
	(4.10)	(2.01)	(2.50)
Average Product Size	30.34	21.10	30.63
	(53.54)	(16.55)	(20.5)

Notes: This table compares summary statistics of all products across stores and those products that decreased or increased in size between 2010 and 2020. Note that all products solely include the top 80%-percentiles brands within each module. Column 1 shows the summary stats considering all sales in the sample. Columns 2 and 3 show statistics for downsized and upsized products, considering all sales before and after the product size changes in stores that changed the product size. Average prices and sizes are calculated annually by product module and averaged across years and modules, weighted by the respective shares of total observations. Similarly, standard deviations (in parentheses) are computed annually by product module and averaged across years and modules

In Figure 1, we present an overview of product size variations across different product groups between 2010 and 2020. Subfigure 1a illustrates the count of unique products undergoing at least one size change, while Subfigure 1b details the frequency of these size alterations, counting each instance of increase or decrease at the store level. We observe a size reduction being more common than size expansion across almost all product groups.

The product categories of candies and snacks are notable for their frequent size variations, both in terms of decreases and increases. Additionally, it is worth mentioning that size changes are also frequently observed in product groups such as hair care and cosmetics, where changes in product size tend not to impact consumer usage patterns directly. Finally, there is substantial heterogeneity across product groups—with some product categories, such as milk or wine, where size changes rarely occur.

Additionally, we present a timeline analysis of product size changes from 2010 to 2020 in Figure 2, showing weekly observations in both subfigures. Subfigure 2a illustrates the weekly count of unique products that either increase or decrease in size. In contrast, Subfigure 2b details the weekly total of size change occurrences across all stores.

A consistent observation is that product size decreases, both at the individual product level and in terms of occurrences, are more frequent than increases. Throughout the decade from 2010 to 2020, product size changes have been a continuous trend, indicating that these changes are not a recent phenomenon but have been occurring since the early 2010s.

We also note a slight seasonal pattern, with product size changes being more common at the start of each year. This pattern is particularly apparent for the number of occasions across stores and shows some irregularities across years. Additionally, we observe a correlation between product size decreases and increases over time. This relationship may be linked to the infrequent adjustments in stores' product assortments, indicating that periods marked by significant product size reductions also tend to experience more occurrences of size increases, though to a lesser extent.

Figure 1: Overview of Product Size Changes

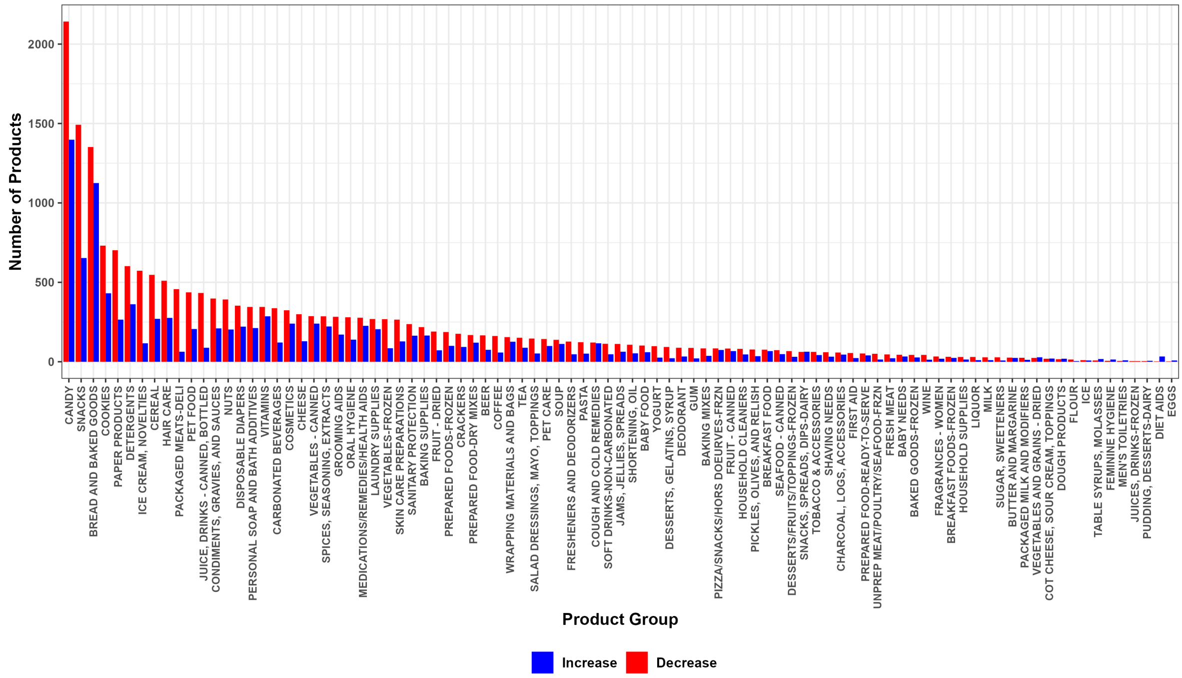

(a) Number of Products

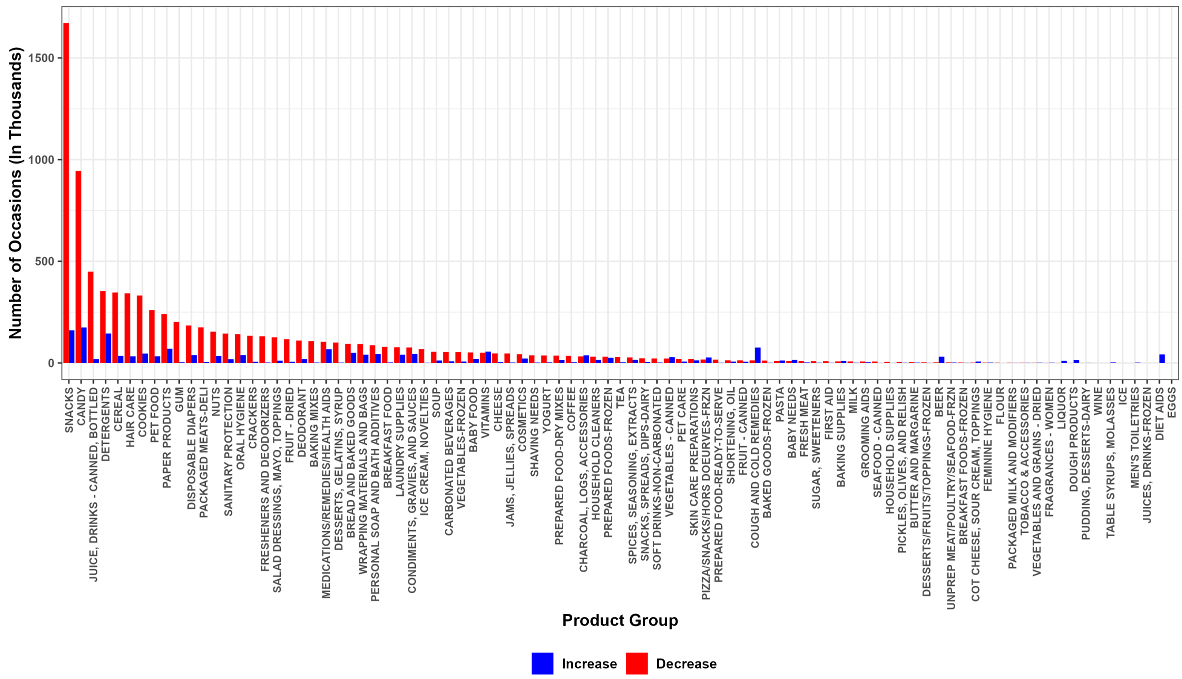

(b) Number of Occasions Across Stores

Notes: This figures illustrates the trends in product size changes across various product groups between 2010 and 2020. In Subfigure 1a, we display the total number of product size changes, counting each product once regardless of the number of stores affected. This contrasts with Subfigure 1b, which considers each instance of a size change within an individual store as a separate occurrence. In both subfigures, blue bars represent an increase in product size, while red bars denote decreases. Note that for a size change to be included, it must occur in an individual store where the new product size replaces the old one permanently, not just temporarily. Additionally, the size changes fall within specific parameters: Increases are capped at 33.33% and decreases at a maximum of 25%.

Figure 2: Timeline of Product Size Changes

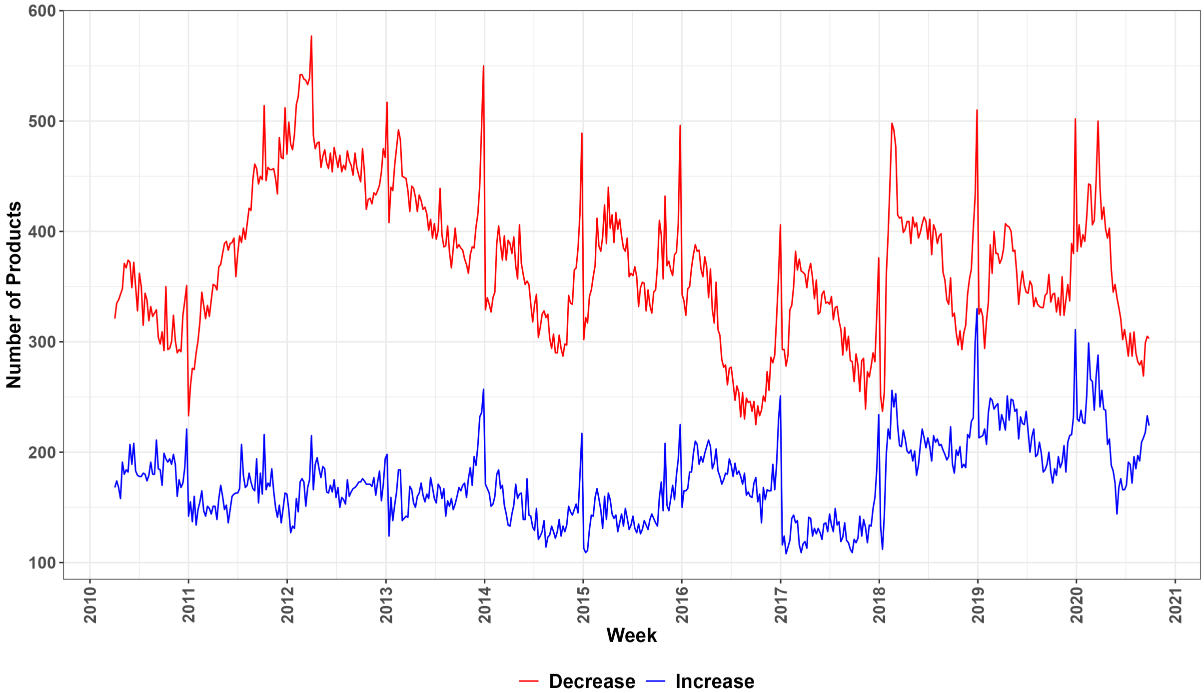

(a) Number of Products

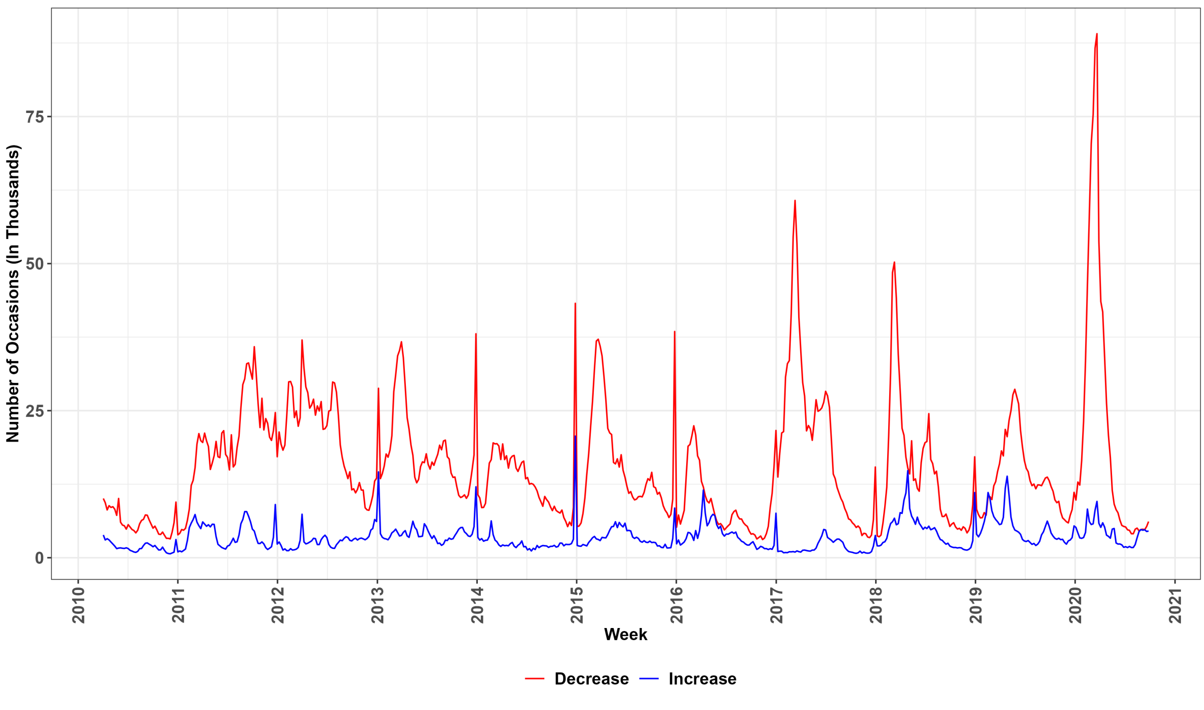

(b) Number of Occasions Across Stores

Notes: The figures provide a detailed timeline of product size changes from 2010 to 2020. In Subfigure 2a, we depict the trends in weekly product size changes, where each product is counted once, irrespective of the number of stores experiencing the change. In contrast, Subfigure 2b offers a different perspective by accounting for each size change event as a distinct occurrence within individual stores. Throughout both subfigures, blue lines represent increases in product size, and red lines indicate decreases. This dual approach allows for a comprehensive understanding of the frequency and distribution of product size alterations over the decade.

3 Insights into Product Size Changes

In this section, we analyze product size changes, examining how these variations differ across product categories, over time, and in their effects on prices, price per volume, and sales. When analyzing these dimensions, we distinguish between the characteristics of product size increases and decreases.

We begin our analysis by examining the temporal trends in product size changes across various product categories from 2011 to 2019, as illustrated in Figure 3. Subfigure 3a represents each product category in a separate row. The color intensity indicates the proportion of sales impacted by size reductions to total sales within each category in any given week across all stores. We categorize all sales in a specific store a year before and after a product size decrease as impacted by size reductions. The proportion is then equal to the ratio of all impacted sales and the rolling sum of sales for all products in that category within the same two-year span. Consequently, we exclude 2010 and 2020 due to the inability to calculate a two-year rolling sum. A darker shade signifies a larger percentage of category sales impacted by size reductions. Similarly, Subfigure 3b depicts the trends of product size increases, with darker shades of red denoting a higher fraction of sales affected by size increases.

It is noteworthy that the overall impact of size changes on sales is relatively modest. For product size decreases, the maximum proportion of affected sales in a week is 5.2%, with a sales-weighted average of 0.005% across all categories and time. In contrast, for size increases, the peak is 2.5%, and the weighted average is merely 0.001%. Two key observations emerge from this analysis. Firstly, there is significant variation in the frequency and timing of product size adjustments across different categories. These adjustments do not transpire within a brief interval; instead, we observe ongoing size reductions in various categories and stores. These reductions in individual product categories are not synchronized across stores but occur progressively over several months. Secondly, the modifications in product size, encompassing both increases and decreases, are not restricted to a specific time frame. Our research suggests that, although infrequent, these fluctuations are evident across most product categories and persist throughout the entire study period.

In our next analysis, depicted in Figure 4, we explore the relationship between product price changes and product size alterations. Subfigure 4a focuses on cases of product size reductions, while Subfigure 4b examines instances of size increases. In both cases, we consolidate data from individual product modules into broader product group categories. In the aggregation, we

Figure 3: Timing of Product Size Changes in Product Groups

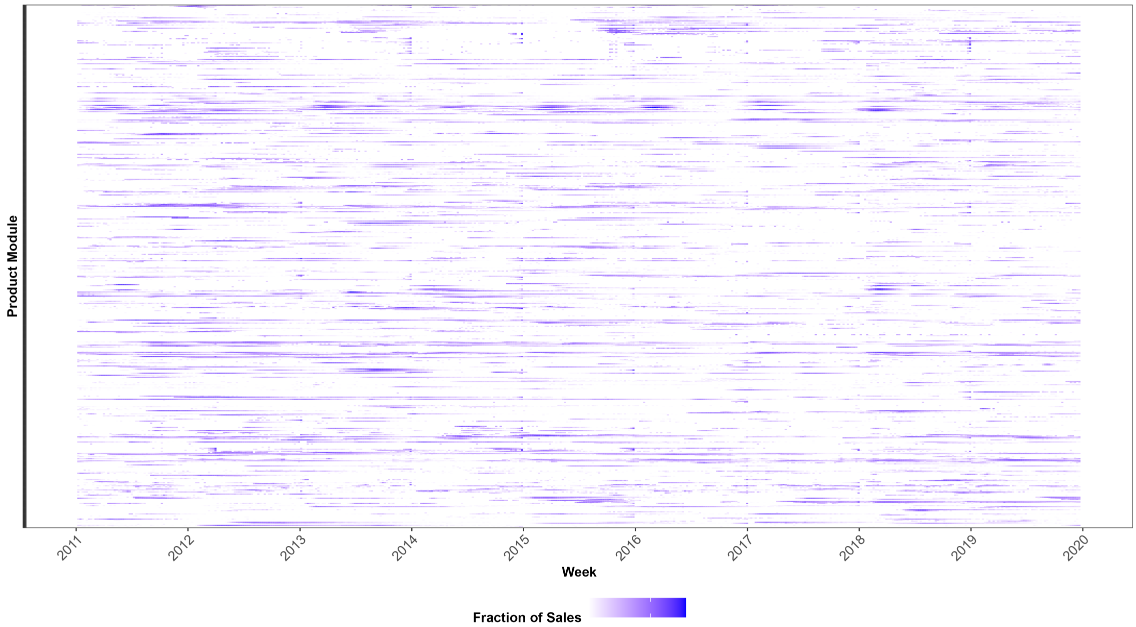

(a) Product Size Decreases

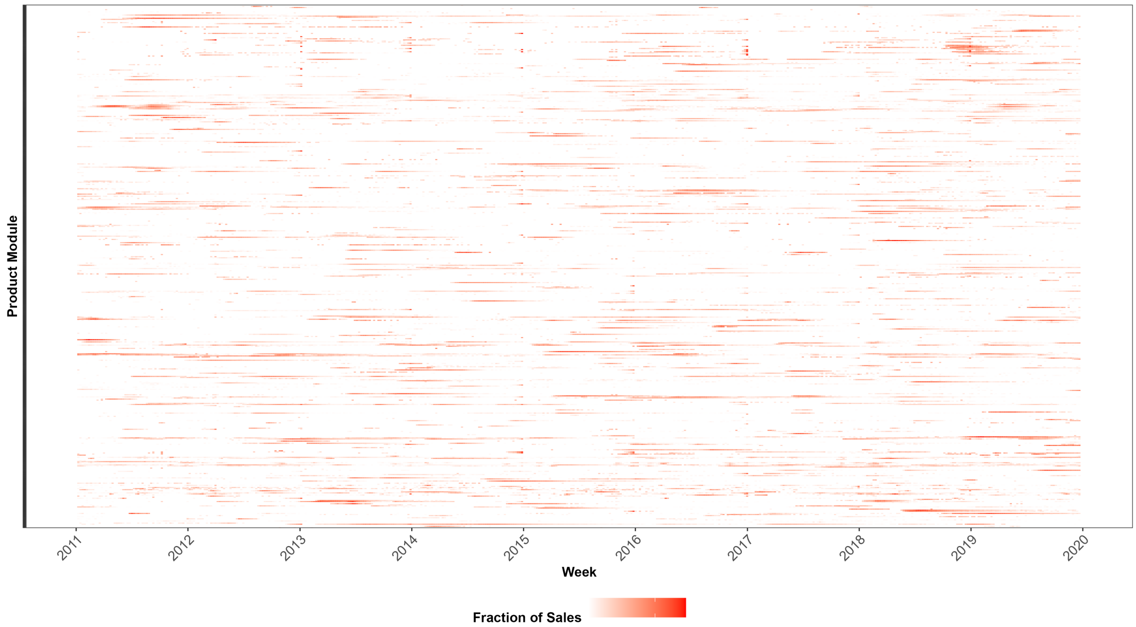

(b) Product Size Increases

Notes: The figures presented herein offer insights into the timing of product size variations across different product modules. Specifically, Subfigure 3a illustrates the instances of product size reductions, whereas Subfigure 3b focuses on the occurrences of product size expansions. In these visualizations, each row is dedicated to a distinct product module. The intensity of blue or red hues within these rows represents the proportion of sales impacted by changes in product size, either through decrease or increase, respectively. This proportion, termed as the 'Fraction of Sales' is calculated using a moving average approach. It encompasses the sales of products that underwent size changes from one year prior to one year post-change. This figure is then compared to the aggregate sales of all products within the same module over the two-year period, across all stores. Note that due to the specifics of this moving average calculation, data from the years 2010 and 2021 are excluded from the analysis.

compute the weighted average of price changes and size alterations, using the sales volume of each module within a group as the weight. We measure price changes at the store level, assessing the year before and after any given price alteration.

In these subfigures, the blue line represents a linear weighted regression of product groups, with the weighting based on the sales volume of each group across all years. The red line denotes a 45-degree reference line. This line represents an equitable price adjustment, where the size change is proportional to the price change. Points above this line indicate scenarios where the price per unit volume has increased, while those below imply a decrease in price per volume.

Figure 4 reveals how product size changes interact with price adjustments, particularly highlighting the differences between size increases and decreases. Price changes tend to be more modest when product sizes decrease compared to when they increase. Most product groups exhibit minimal changes in price following a size reduction, hovering near zero. Conversely, size increases often coincide with price changes. As these price hikes are typically lower than accompanying size increases, this pattern benefits consumers on average through lower prices per volume. This beneficial pattern is not observable for cases of size decreases. Figure 5 shows that for all product groups except liquor, the average price per volume increases when product sizes decrease. Conversely, changes in price per volume following a product size increase exhibit more heterogeneity across various product groups; however, prices per volume slightly decrease on average.

In sum, our analysis suggests that retailers or manufacturers often reduce product sizes while maintaining similar pricing levels, effectively increasing the price per unit volume. Conversely, when product size increases occur, prices tend to rise. However, this price hike is typically not enough to offset the size change, resulting in a lower price per volume.

Our analysis proceeds with a descriptive analysis of the impact of product size changes on purchased volumes and sales. In Subfigures 6a and 6b, we illustrate changes in purchased absolute volumes before and after product size adjustments, for decreases and increases respectively. Notably, there is a discernible heterogeneity across product groups in response to size decreases. While some groups exhibit an uptick in purchased volume, others show a decline. On average, however, a reduction in product size is associated with a slight decrease in purchased volume. Conversely, when product sizes increase, the purchased absolute volume rises across the vast majority of product groups, indicating that consumers do not perfectly adjust their purchasing behavior to maintain the same level of absolute consumption in both scenarios of size decreases

Figure 4: Price and Size Changes, Overview

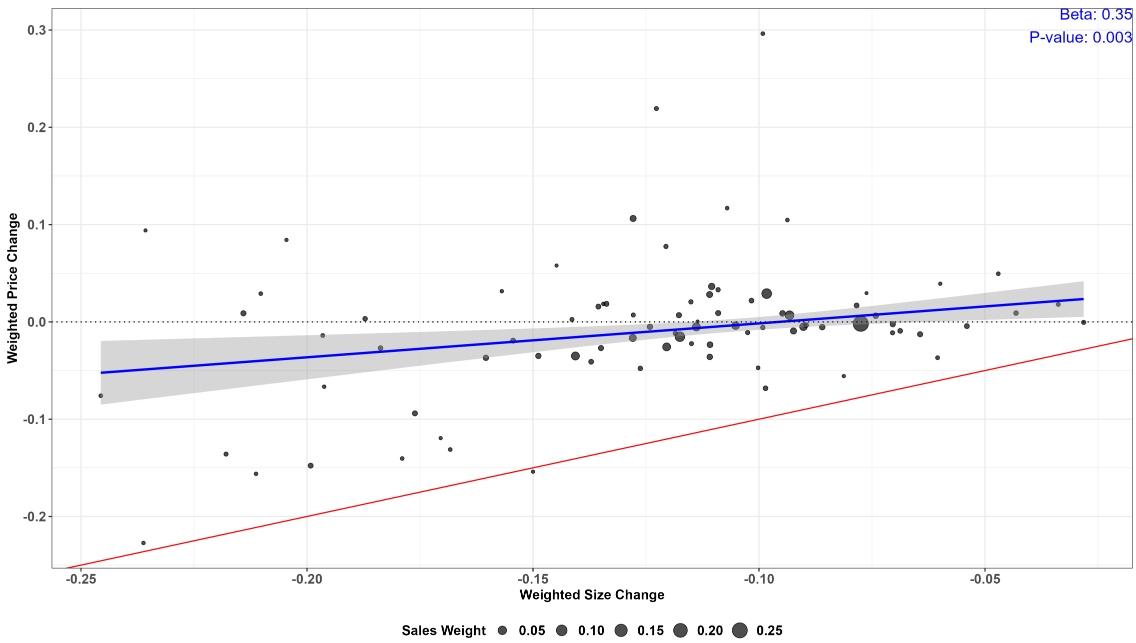

(a) Product Size Decreases

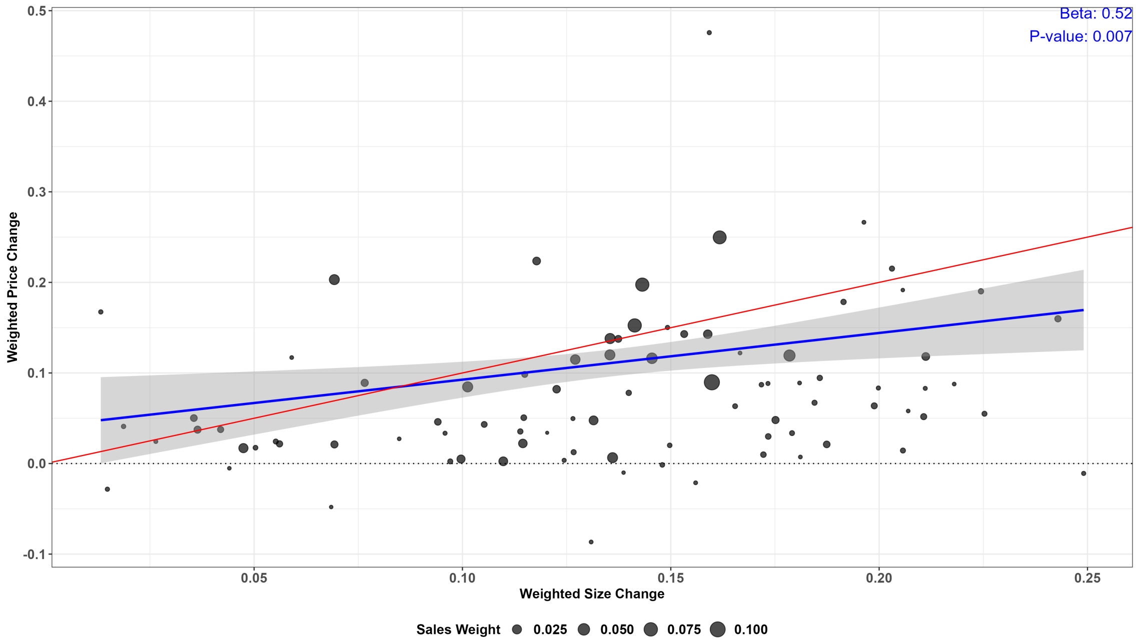

(b) Product Size Increases

Notes: The figures present an examination of how product size changes correlate with price adjustments. The analysis is divided into two parts: Subfigure 4a investigates product size reductions, and Subfigure 4b assesses size increases. For both scenarios, the data is aggregated from individual modules to product groups. Weighted average price and size changes are calculated using sales volume of the products as a weighting factor. Weighted average price changes are analyzed at the store level by comparing prices a year before and after the change. The figure features a blue line that depicts a linear weighted regression, where weights are based on sales of products that underwent size changes within their respective groups. The coefficients and p-values of these regressions are displayed in the upper right corner of each subfigure. Additionally, a red line is included to provide a baseline for proportional price-size adjustments.

Figure 5: Price per Volume, Overview

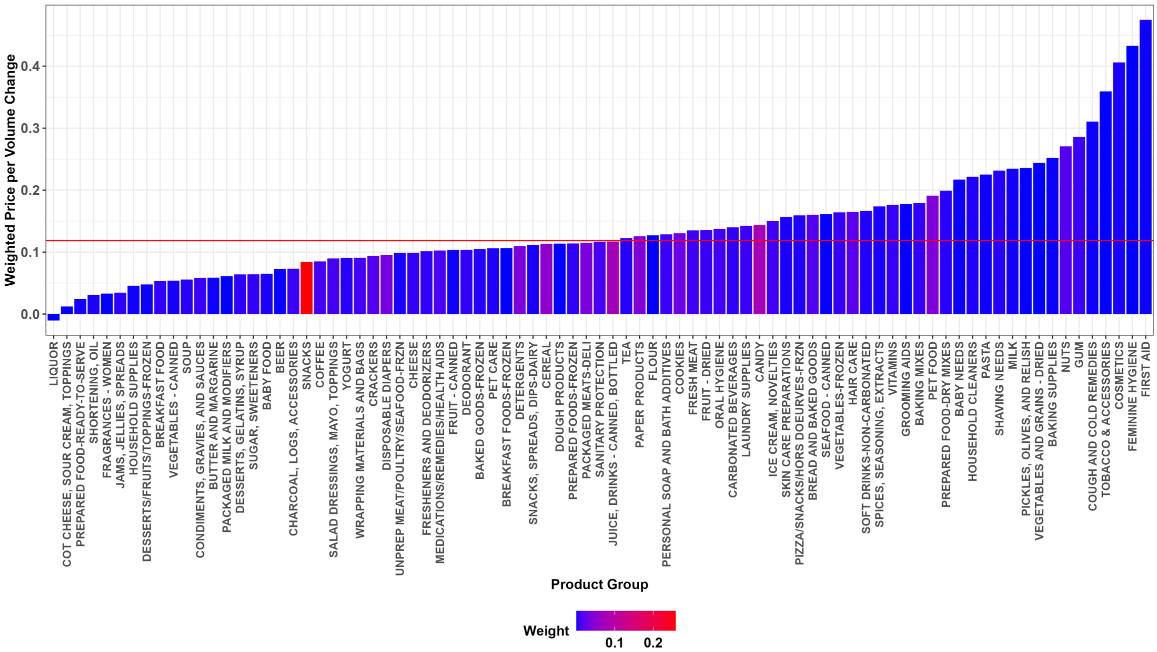

(a) Price per Volume, Product Size Decreases

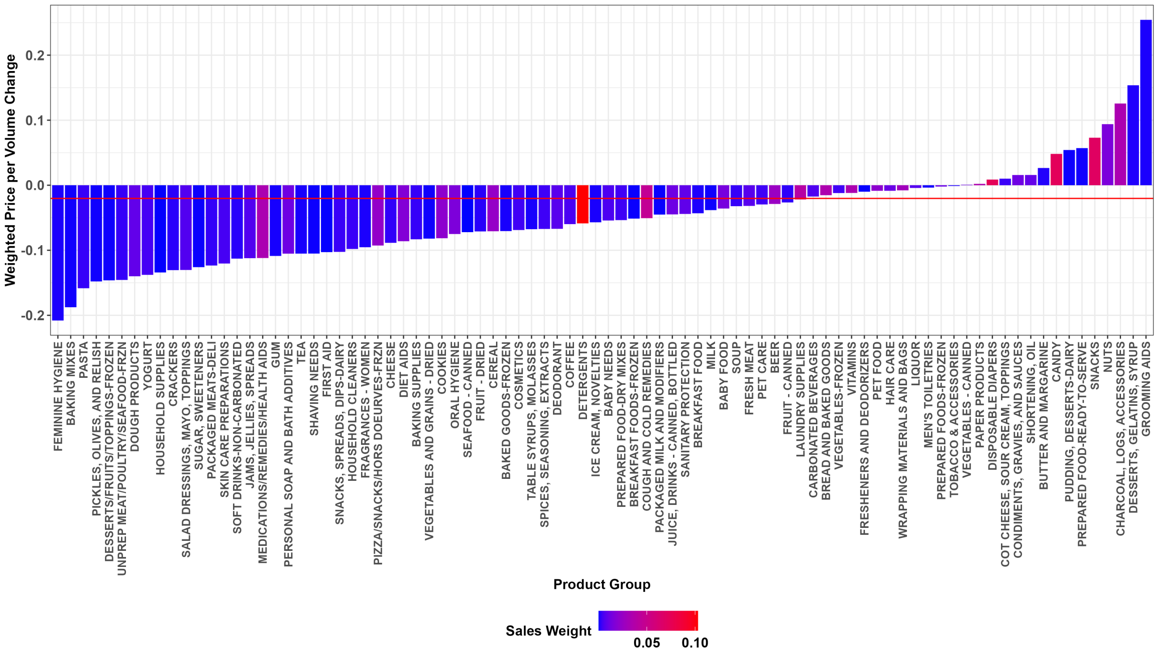

(b) Product Size Decreases, Product Size Increases

Notes: The figure presents an examination of how product size changes affect prices per volume. The analysis is divided into two parts: Subfigure 5a investigates product size reductions, and Subfigure 5b assesses size increases. For both scenarios, the data is aggregated from individual modules to product groups. Weighted average price per volume changes are calculated using sales volume of the products as a weighting factor. Note that prices per volume are prices divided by the size of a product using product group-specific measurements. The analysis compares these metrics from the year preceding to the year following a product size change. The color intensity of each bar in the figure corresponds to the sales weight of the product group, with deeper shades of red indicating higher sales volumes within that group. The red line shows the weighted average price per volume changes across all product groups.

and increases.

Subfigures 6c and 6d further explore the effects on stores' sales. Here, the trend is somewhat consistent for both size decreases and increases: most product categories experience an uptick in sales. These findings suggest that, on average, consumers are spending more on these products and that retailers and/or manufacturers successfully leverage product-downsizing and upsizing strategies to boost revenue.

4 Manufacturer or Store-Level Size Changes

In this section, we explore who is primarily accountable for the observed changes in product sizes. Generally, either retailers or manufacturers might instigate size reductions or expansions. Recent media focus has largely been on manufacturer-led size alterations. Retailers, notably, have been pointing fingers at manufacturers for supplying smaller products, as highlighted in various reports (The Associated Press, 2022). Some retailers in other countries have even taken steps to transparently inform consumers about size reductions in products (BBC News, 2023; The Korea Times, 2023).

However, it is also conceivable that retailers play a role in size changes. Manufacturers often provide several size variations of the same product. Retailers may choose to stock different sizes or select a particular size for their stores. Consequently, product size shifts at the store level could reflect retail decisions rather than manufacturing changes.

We will next examine both manufacturer- and retailer-induced size changes. Generally, we argue that a size change observed nationally is attributed to manufacturers, while those observed only at the store level are retailer-induced. As explained in section 2.1, we define size changes based on exits and entries on the store level. The definition of exits and entries on the store level and the corresponding definitions of size changes remain unchanged. Next, we identify if a store-level size change is likely due to a manufacturer or a retailer.

Identifying manufacture-induced shrinking requires multiple assumptions. We never observe a coordinated size change across all stores for two reasons. First, we observe residual sales of all products even when the product is unavailable in nearly all US stores. Second, we often observe size changes on a store level preceding a more coordinated change in size nationwide. To tackle this challenge, we follow a two-step procedure: First, we define a manufacture-induced exit date as the date when the sales of the older version of a size-changing product drop permanently to

Figure 6: Volume Purchased and Sales across Product Groups

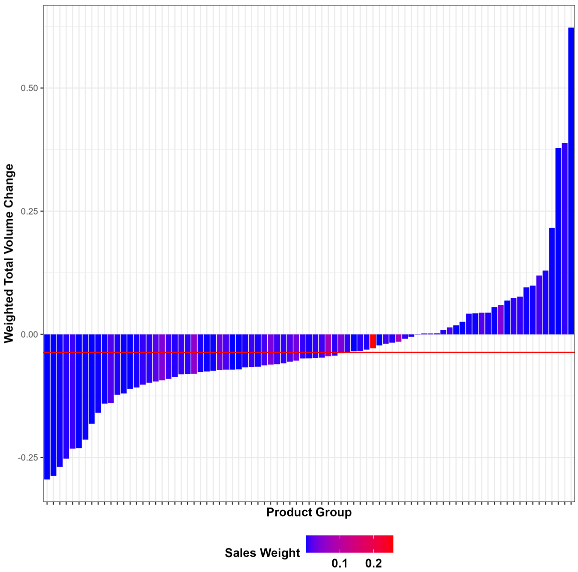

(a) Purchased Volume, Product Size Decrease

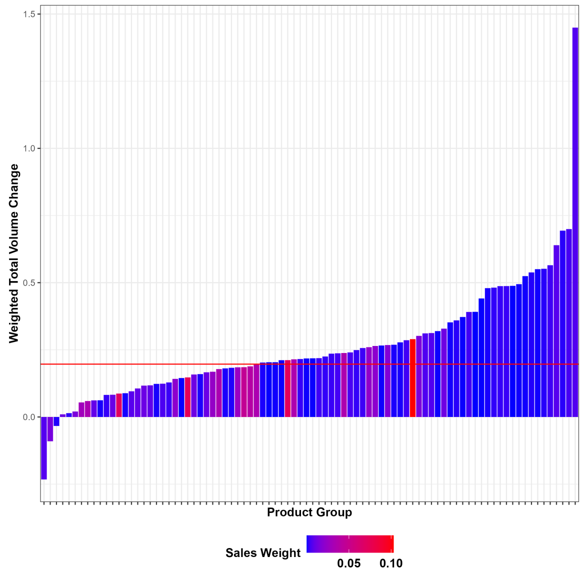

(b) Purchased Volume, Product Size Increase

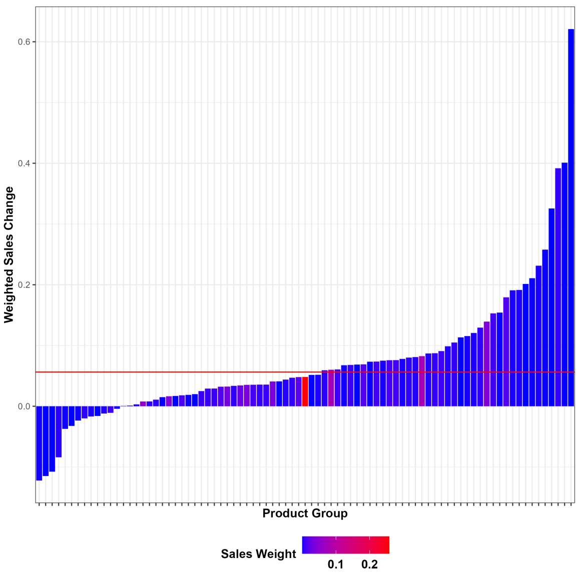

(c) Sales, Product Size Decrease

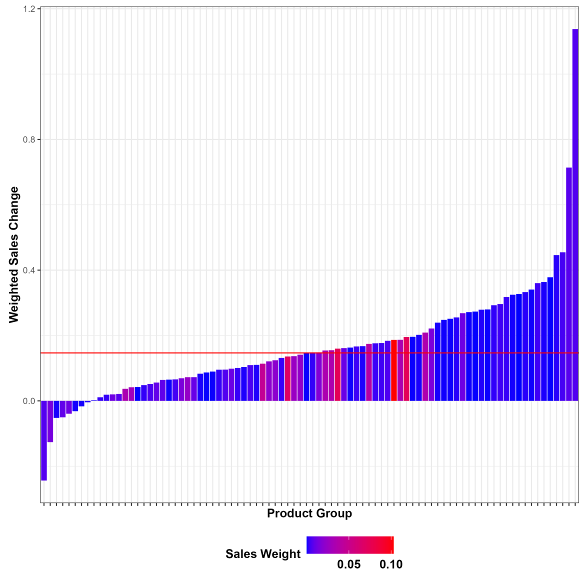

(d) Sales, Product Size Increase

Notes: The figures show the changes of purchased volume and sales of products that decrease or increase in size across product groups. Each bar corresponds to one product group. Subfigure 6a considers the dynamics of purchased volumes following product size reductions, while Subfigure 6b focuses on the effects of size increases. Changes are measured on the product level and compare volumes measured in module-specific units (e.g., ounces) from the year before to the year after the product size change, The analysis employs an aggregation approach, consolidating data from individual products into product groups, considering their weights in sales. Similarly, Subfigure 6c and 6d analyze the corresponding changes in sales, adopting the same methodological framework. In these subfigures, the intensity of color in each bar graphically represents the sales weight of the product group, where richer red hues signify higher sales volumes of a group. A distinct red line across the figures denotes the aggregated weighted changes in volume or sales across all product groups.

less than 5 percent of the weekly national average of sales.[7] Second, as retailers may shrink their product actively way before the manufacturer-induced exit rate, we solely characterize store-level size changes as manufacturer-induced if the size change happens within 52 weeks before the estimated manufacturer-induced exit date.

On the other hand, a retailer-induced size change is identified if sales exceed the 5th percentile in the remaining stores in at least the following 52 weeks. This approach is designed to be robust against small retailers continuing to sell old stock (thereby inflating the prevalence of retailer-induced size changes) while avoiding the overestimation of manufacturer-induced changes by attributing every store-level variation to a decision by the manufacturer.

To clarify our methodology for distinguishing between retailer- and manufacturer-induced size changes, we present a visual illustration in Figure 7. Specifically, we focus on the top 100 products experiencing size reductions based on sales and observe a manufacturer-induced exit rate within our sample period.[8] In the figure, these products are arranged by their sales volume, beginning with the highest. We highlight the initial observed size reduction event in red, followed by the estimated exit point – the moment when sales fall below 5 percent of the national average across all weeks in our data on a permanent basis.

The example of top-selling products that decrease in size reveals considerable variability in the time span between the first size reduction and the eventual national exit, with some products showing a shorter duration and others spanning several years. To discern whether the size reductions for each product are predominantly manufacturer-induced, we use a light blue highlight for products where the majority of size reductions occurred within a year prior to the national exit. Notably, our findings suggest that only about 10 percent of the products predominantly exhibit manufacturer-induced size reductions.

In Figure 8, we present an analysis comparing the sales volume associated with retailer versus manufacturer-induced size changes. Subfigure 8a illustrates the sales volume of products experiencing size decreases, while Subfigure 8b focuses on size increases. The sales size for both manufacturer and retailer-induced changes is represented as a two-year moving sum.

Consistent with the patterns identified in Figure 2, we find that product size decreases are more prevalent in terms of both frequency and sales volume. While manufacturer-induced size changes generally show lower sales volumes compared to those initiated by retailers, there

[7] In detail, we require four consecutive weeks of national sales that are lower than 5 percent of the national sales level.

[8] A similar analysis is conducted for size increases, yielding comparable results across all aspects.

Figure 7: Example Retailer- vs. Manufacturer-Induced Size Decrease

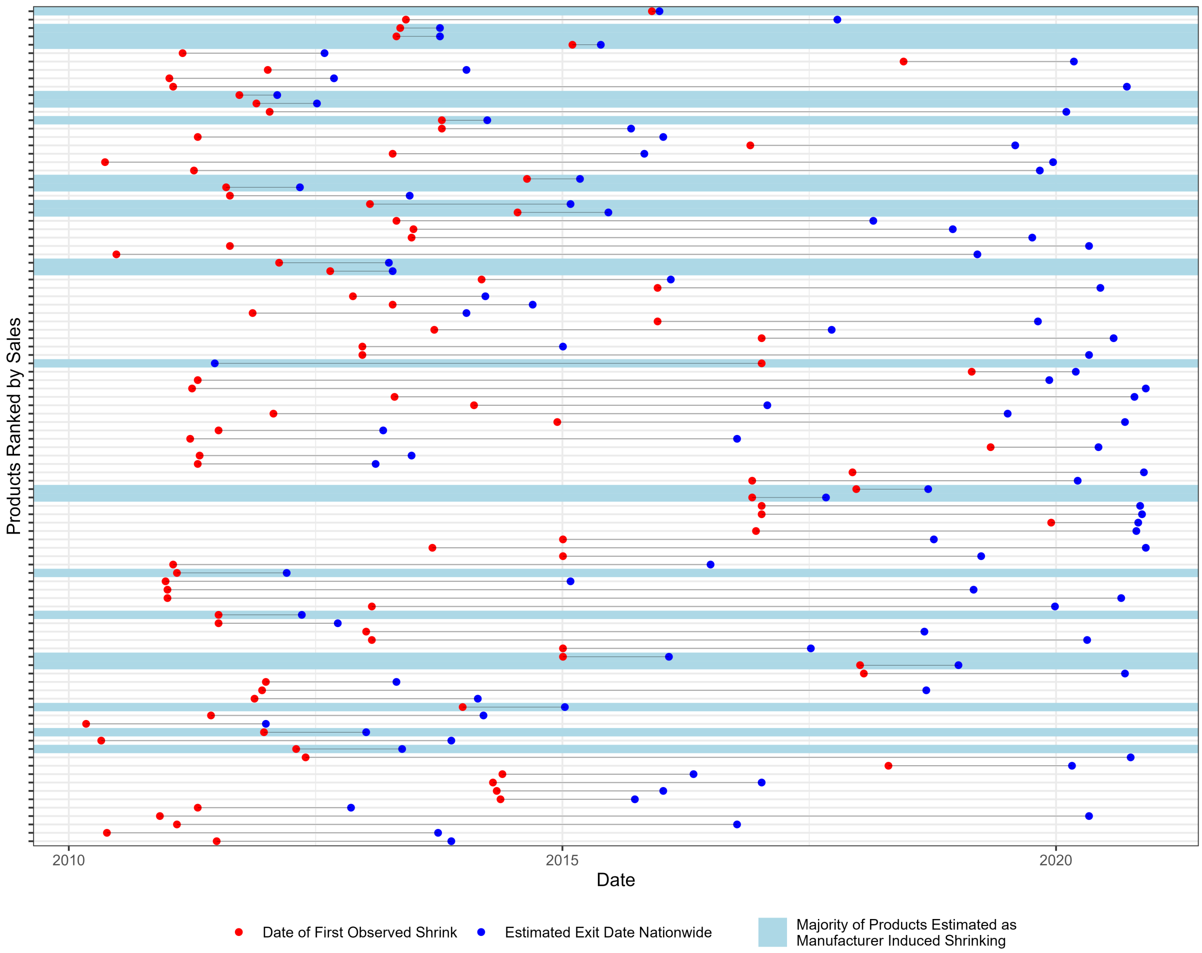

Notes: This figure visually represents the leading 100 products by sales volume that have undergone size reductions on the store level, and there exists a date we define as manufacturer-induced exit date, meaning the sales of the older version of a size-changing product drop permanently to less than 5 percent of the weekly national average of sales. Products are ordered according to their sales, starting with the highest. Key events are marked: initial size reduction occurrences are indicated in red, and the estimated exit point—defined as the moment when a product's sales consistently fall below 5 percent of the national average across the full sample period—is shown in blue. Products where the majority of size reductions are identified as manufacturer-induced, occurring predominantly within a year prior to the national exit, are highlighted in light blue.

appears to be a notable correlation between the two.

5 An Empirical Model of Consumer Responses

5.1 Empirical Strategy

In this section, we move to evaluating consumer responses to product price and size changes. In order to evaluate whether consumers react differently to unit price changes than to product size changes, we estimate elasticities of demand for each store-product pair. We distinguish between price elasticity η_p and product size elasticity η_l. As our benchmark, we estimate the response of weekly log quantity, $log(x(p_{t,s,i}))$, to the weekly log unit price and the weekly log product size of a product within a store s, allowing for store-product fixed effects, $\alpha_{s,i}$, and week fixed effects, γ_t:

$$log(x(p_{t,s,i})) = \eta_p log(p_{t,s,i}) + \eta_l log(l_{t,s,i}) + \alpha_{s,i} + \gamma_t + \epsilon_{t,s,i}, \tag{1}$$

where $p_{t,s,i}$ and $l_{t,s,i}$ denote the weekly unit price and weekly package size of product i in store s. We define a product as in previous sections, i.e., having identical brand, UPC description and brand description. Our primary interest lies in the comparative analysis of elasticities. The implicit assumption allowing for meaningful comparison of the two elasticities in our benchmark model is that product size only affects demand through its effect on price per volume. If only the price per volume matters to consumers, we would expect that the two elasticities equate for fully attentive consumers, i.e., consumers' demand reacts to a percentage change in the unit price in the same way as to a percentage change in price per volume due to product size changes.

In our estimations, we use scanner the data with imputed prices for zero-sales weeks using approaches from Hitsch et al., 2021 (see Appendix B for a discussion) and focus on food, drugstores, and mass merchandise chains. Moreover, we selectively focus on products that experienced a variation in product size across at least one store within the U.S. across the observation period. By focusing on these products, we ensure a comparable set of products and computational feasibility. Yet, we still have a diverse set of observations, not subject to product size changes, that allow for meaningful time-fixed effects. The final sample and summary statistics are detailed in Appendix Table A.1. We estimate equation 1 separately for each product group.[9]

[9] Due to computational constraints, we treat product modules "CANDY - CHOCOLATE", "CANDY - NON-CHOCOLATE" (both within product group "CANDY"), and "SNACKS - POTATO CHIPS" (within product group "SNACKS") as separate product groups. The estimates for product groups "CANDY" and "SNACKS" include all other product modules within these groups, excluding the separately estimated product modules.

Figure 8: Sales Volume of Retailer- vs. Manufacturer-Induced Size Changes

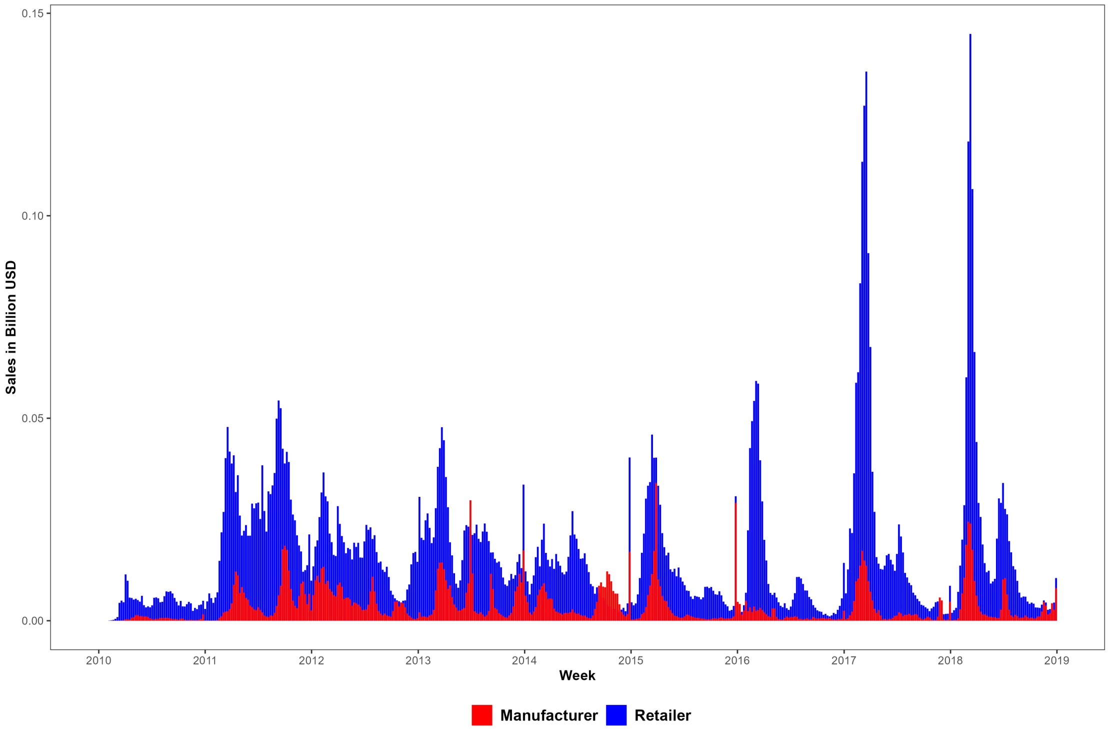

(a) Product Size Decreases

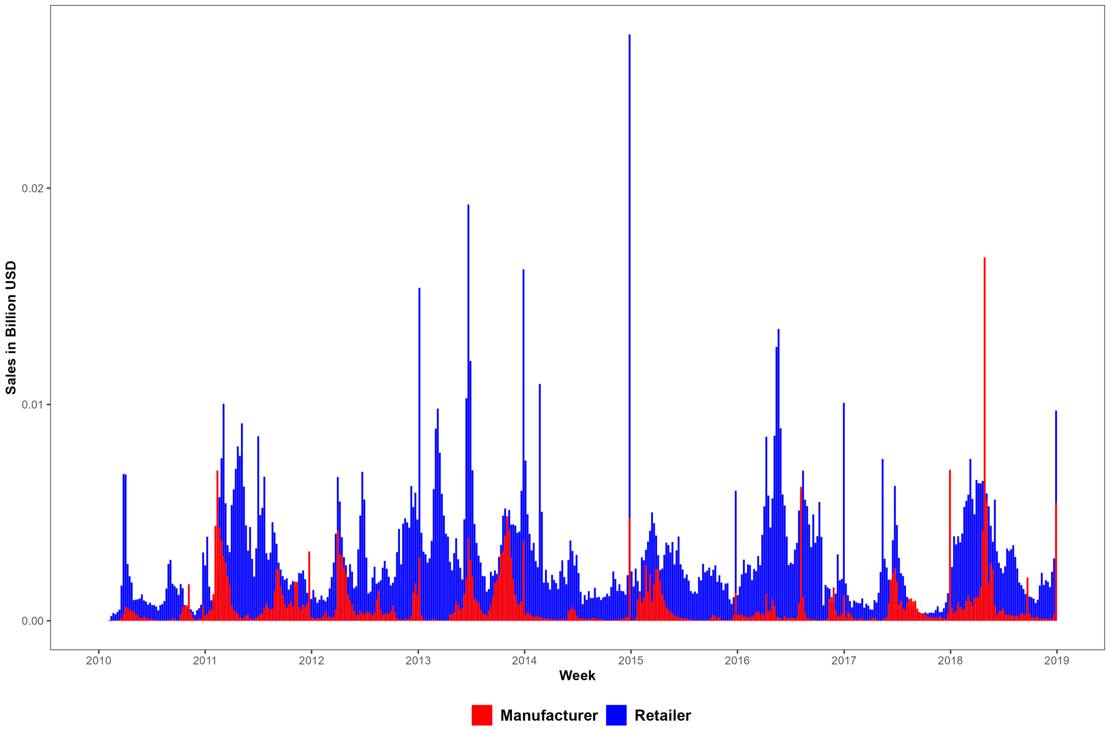

(b) Product Size Increases

Notes: This figure delineates the scale of product size alterations instigated by both manufacturers and retailers. Subfigure 8a focuses on product size decreases, while Subfigure 8b examines instances of product size increases. Sales volumes are calculated at the product level for each week, represented as a cumulative sum over two years. Note that data from the year 2020 is excluded from this analysis. Our current methodology for differentiating between manufacturer and retailer-induced size changes is not applicable to periods shorter than 52 weeks prior to the last date in the sample.

A concern challenging the validity of our estimation approach is that prices and sizes are correlated to unobserved local demand shocks. To tackle endogeneity concerns, we follow DellaVigna and Gentzkow (2019) and instrument the weekly price and the weekly product size in-store s with the average of prices and product sizes across other stores in s's chain that is located outside the respective store's Designated Market Area (DMA). This instrument approach where prices in other markets are used as instruments is common in the literature of industrial organization, following Hausman (1996) and Nevo (2001). The key assumption here is that the timing of chain-level sales is unrelated to local demand shocks after controlling for time-fixed effects. DellaVigna and Gentzkow (2019) shows that prices and stores' assortments are very similar within chains, but different between chains. These findings imply that chains do not account for local demand conditions when setting prices or product assortment, making our instrument approach convincing with respect to both price and product size.[10] Moreover, product size changes are a long-term strategy that is not directly influenced by local demand shocks, which suggests that our fixed effects adequately capture relevant local demand variations.

While the estimation of model 1 allows us to estimate price and quantity elasticities of demand, we do not evaluate how the sensitivity of consumers differs when products are de- or increasing in size. Thus, we extend our main model to estimate differential elasticities between products that decrease and increase in size. In detail, we consider the following model which is a direct extension of model 1:

$$\begin{aligned} log(x(p_{t,s,i})) = &\eta_p log(p_{t,s,i}) + \eta_l log(l_{t,s,i}) + \\ &\beta_p log(p_{t,s,i}) \times \mathbf{I}(Shrink_{i,s} = 1) + \beta_l log(l_{t,s,i}) \times \mathbf{I}(Shrink_{i,s} = 1) + \\ &\alpha_{s,i} + \gamma_{t,\mathbf{I}(Shrink_{i,s}=1)} + \epsilon_{t,s,i}, \end{aligned} \tag{2}$$

where the key differences are the interaction terms of $log(p_{t,s,i})$ and $log(l_{t,s,i})$ with $\mathbf{I}(Shrink_{i,s} = 1)$, which is a dummy variable that takes the value one if the specific product has been decreasing in size. Additionally, we also interact the dummy variable with the week-fixed effects ($\gamma_{t,\mathbf{I}(Shrink_{i,s}=1)}$) to allow for diverging time trends for both subgroups. As a result, β_p indicates how the price elasticities of downsized products differ from those of non-downsized products, encompassing both upsized products and products in stores where no size changes occur. Similarly,

[10] In contrast, to our benchmark model, DellaVigna and Gentzkow (2019) use store-product-week-of-year and store-product-year fixed effects. The qualitative results remain robust when using these alternative fixed effects (results not shown).

β_l shows how the size elasticity differs when considering downsized products. To corroborate our findings, we use the fitted values from the first-stage of our instrument variable approach presented in Equation 1 and interact the fitted values with the $\mathbf{I}(Shrink_{i,s} = 1)$ as in Equation 2. To ensure consistency between the first and second stages, we employ non-interacted time fixed effects (γ_t) in the second stage of this specification.[11]

5.2 Results

Figure 9 presents the estimates by product group. The response to unit price changes appears more pronounced than to product size changes. The average estimated unit price elasticity, weighted by total sales across all product groups, which is illustrated by the solid blue line, is equal to -1.14. The average elasticity for product size, illustrated by the solid red line, is -0.05. Thus, on average, a 1 percent increase in price would reduce sales by 1.14 percent, while the regression results suggest, on average, a zero effect of size on purchases. The difference between price and quantity elasticities suggests consumers are less responsive to the price-per-volume changes resulting from product size variations than to unit price changes. Such a tendency offers profit-maximizing firms an opportunity to capitalize on potential consumer underreactions to these changes by decreasing product sizes.

Second, Figure 10 presents the results for the instrument variable regression. The first stages of the IV regressions are strong in most cases. However, we find that for a few smaller product groups, this is not always the case. Therefore, we exclude product groups where the ratio of F-statistics to the number of observations in the first stage is lower than 0.05 to avoid biases due to outliers. In Appendix E.1, we provide detailed information on the first stage and discuss the sensitivity of our results to product groups with weak first stages. After excluding outliers, the weighted average estimate for price elasticity is -1.83, and the estimate for package size elasticity is 0.19. In contrast to the results from Figure 9, the package size elasticity is clearly different from zero when aggregating all product groups. While estimated elasticities increase in absolute size, the IV regression results produce a consistent picture. Consumers tend to underreact to price changes due to package size changes relative to unit price changes.

Our findings also reveal significant heterogeneity in elasticities across different product groups. The variation is more pronounced for product size elasticities, which can be partly attributable to considerably smaller variations in product size compared to unit price. Large

[11] Results do not change considerably if we use interacted fixed effects as in Equation 2 instead.

Figure 9: Estimated Unit Price and Package Size Elasticities for Different Product Groups

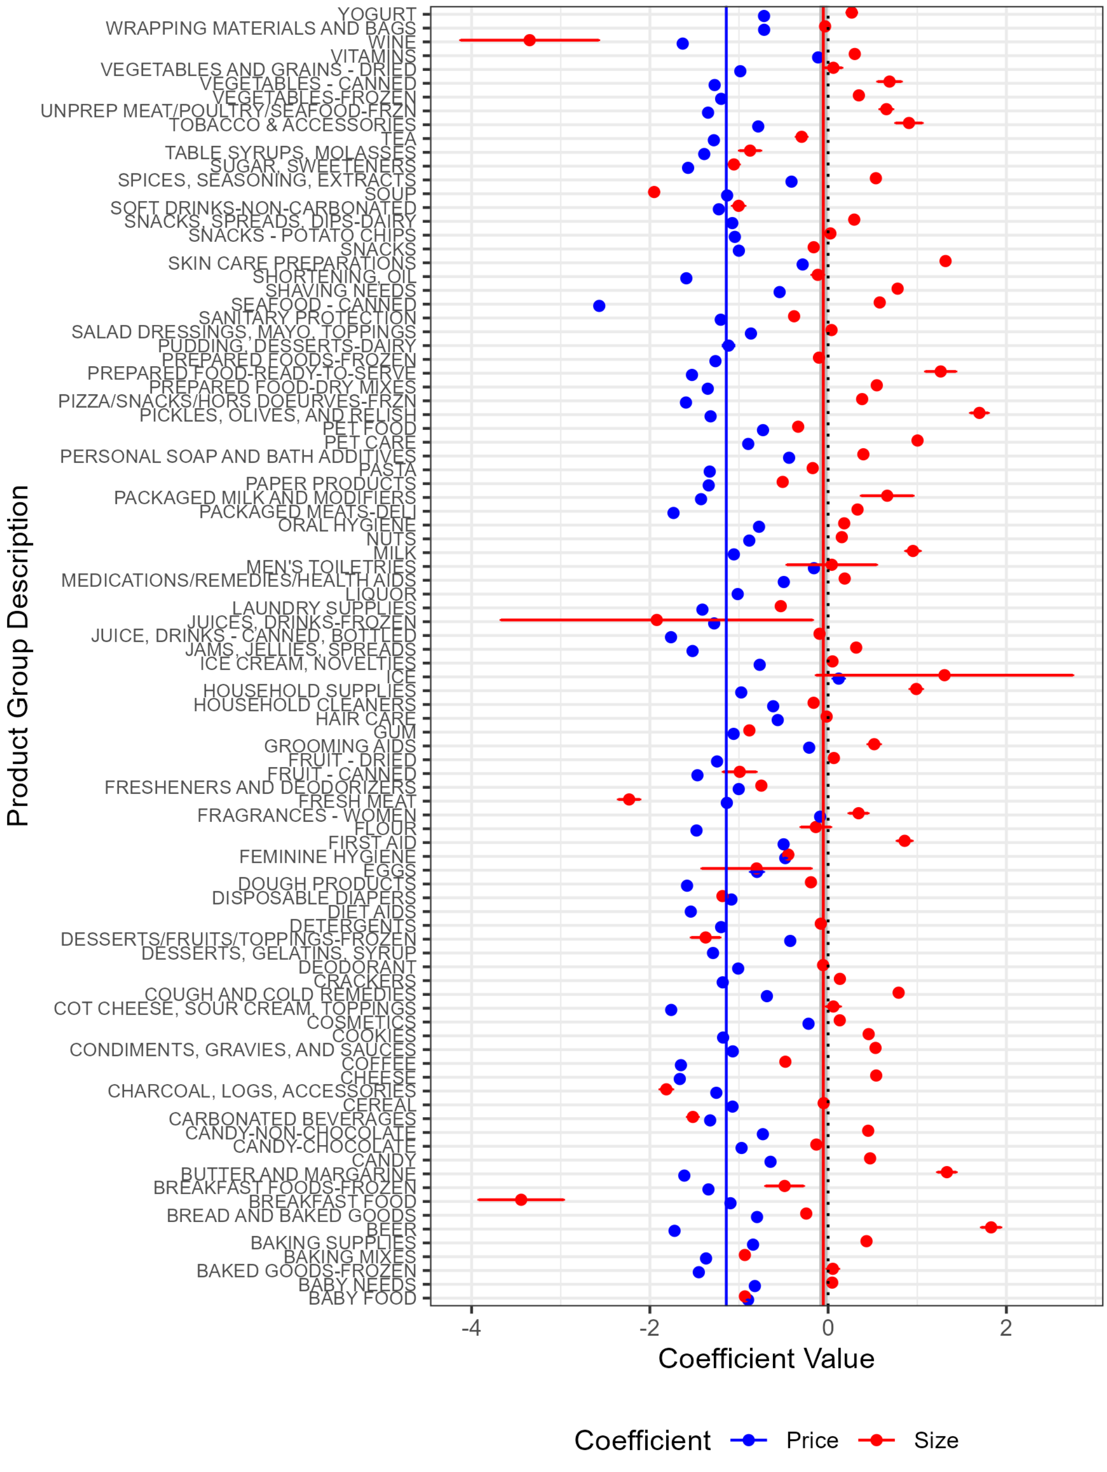

Notes: This graph presents the estimated unit price elasticities (in blue) and package size elasticities (in red) and corresponding 95% confidence intervals in different product groups, according to equation 1. The coefficients, therefore, refer to $\hat{\eta}_p$ and $\hat{\eta}_l$ in model 1. The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. Weights are based on the sales within product groups. Error bars denote the 95% confidence intervals of the weighted coefficients. For visibility, we exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5. However, those product groups are part of the weighted average calculations.

Figure 10: Estimated Unit Price and Package Size Elasticities for Different Product Groups, IV Regression

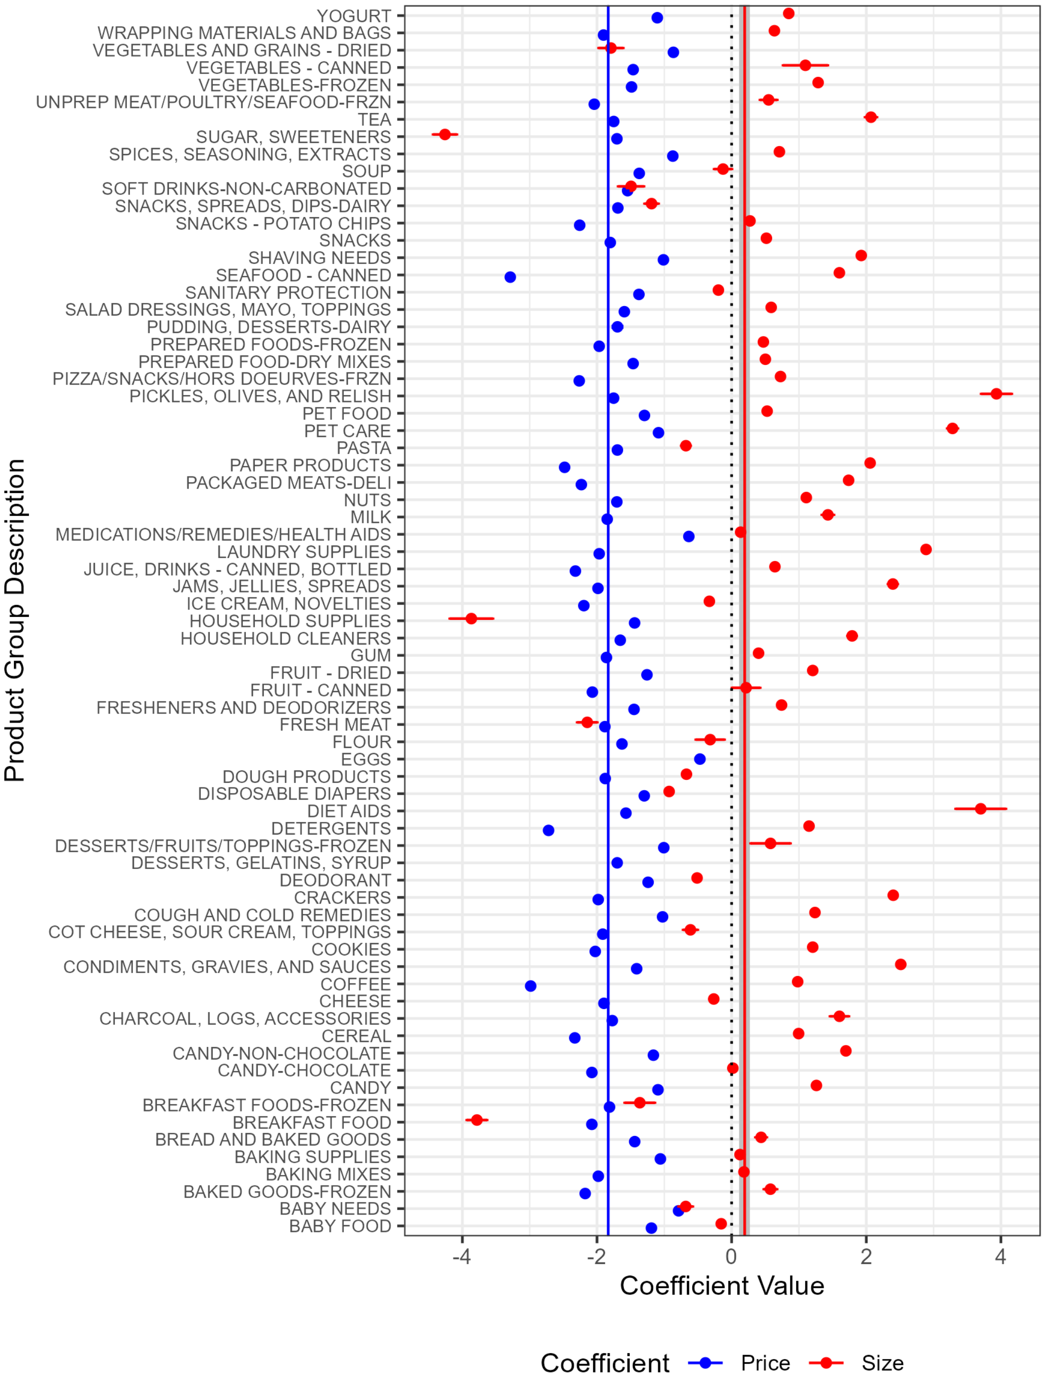

Notes: This graph presents the estimated unit price elasticities (in blue) and package size elasticities (in red) and corresponding 95% confidence intervals in different product groups, using an instrument variable approach. The coefficients refer to $\hat{\eta}_p$ and $\hat{\eta}_l$ in model 1. We instrument the weekly price and the weekly package size in store s with the average of prices and package sizes across other stores in s's chain that are located outside the respective store's DMA. The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. Weights are based on the sales within product groups. We exclude product groups where the adjusted F-statistics—the ratio of F-statistics to the number of observations in the first stage—is lower than 0.05. We explain details and show the robustness of this decision in Appendix E. For visibility, we exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5. However, product groups solely excluded for visibility are part of the weighted average calculations.

negative product size elasticities, indicating that demand is increasing for smaller product sizes, are more prevalent for food and drink product groups, such as fresh meat, carbonated beverages or soup. This tendency could be related to consumers exhibiting a preference for product size, e.g., because of stockpiling constraints and healthier eating habits for some product categories. We expect this to be less relevant for the size elasticity in product groups that are more durable and where "one" product is typically connected to multiple uses, such as deodorants, hair care, cosmetics, but also cereals. In most of these groups, size elasticities are close to zero, which is in line with the idea that consumers are not attentive to product size changes. In the subsequent analyses, we further try to disentangle preference from price effects related to product size changes.

We start by investigating if elasticities vary between downsized products and those that have increased in size. For this purpose, we interact the logarithm of prices and product size with a dummy variable indicating whether a product decreased in size at the store-product level. Figure 11 illustrates the elasticities for non-downsized products, and Figure 12 presents the interaction term's coefficients across product groups.

Estimating the price and product size elasticities for non-downsized products gives weighted average estimates of -0.29 and 0.15, respectively. Thus, for non-downsized products, the reaction to price changes is about twice as strong than for size changes. Moving to Figure 12, the average coefficient of the price interaction term is close to zero (-0.02), indicating no significant difference in unit price elasticities between products that decrease in size and those that do not. Secondly, the interaction effect for size changes is -0.14, suggesting that consumers react less to size changes when related to product size reductions. The weighted average coefficient of the size interaction term is of similar magnitude to the size elasticity of nonshrinking products, implying that the size elasticity of products size decreases is close to zero. The higher size elasticities for increased products aligns with the idea that retailers make size increases more salient, leading to more pronounced consumer reactions. Conversely, the near-zero average size elasticity for shrunk products suggests that product size decreases tend to be hidden.

Next, we turn to the results of the instrumental variable specification of Equation 2. In some product groups, the instrument variable specification yields very noisy and extreme size and size-interaction coefficients. This is likely because the identification of coefficients for downsized products in these groups is based on a small number of products with infrequent size changes. The fitted values often have a large, sometimes even perfect, overlap with non-downsized prod-

Figure 11: Estimated Unit Price and Package Size Elasticities for Different Product Groups with interaction term

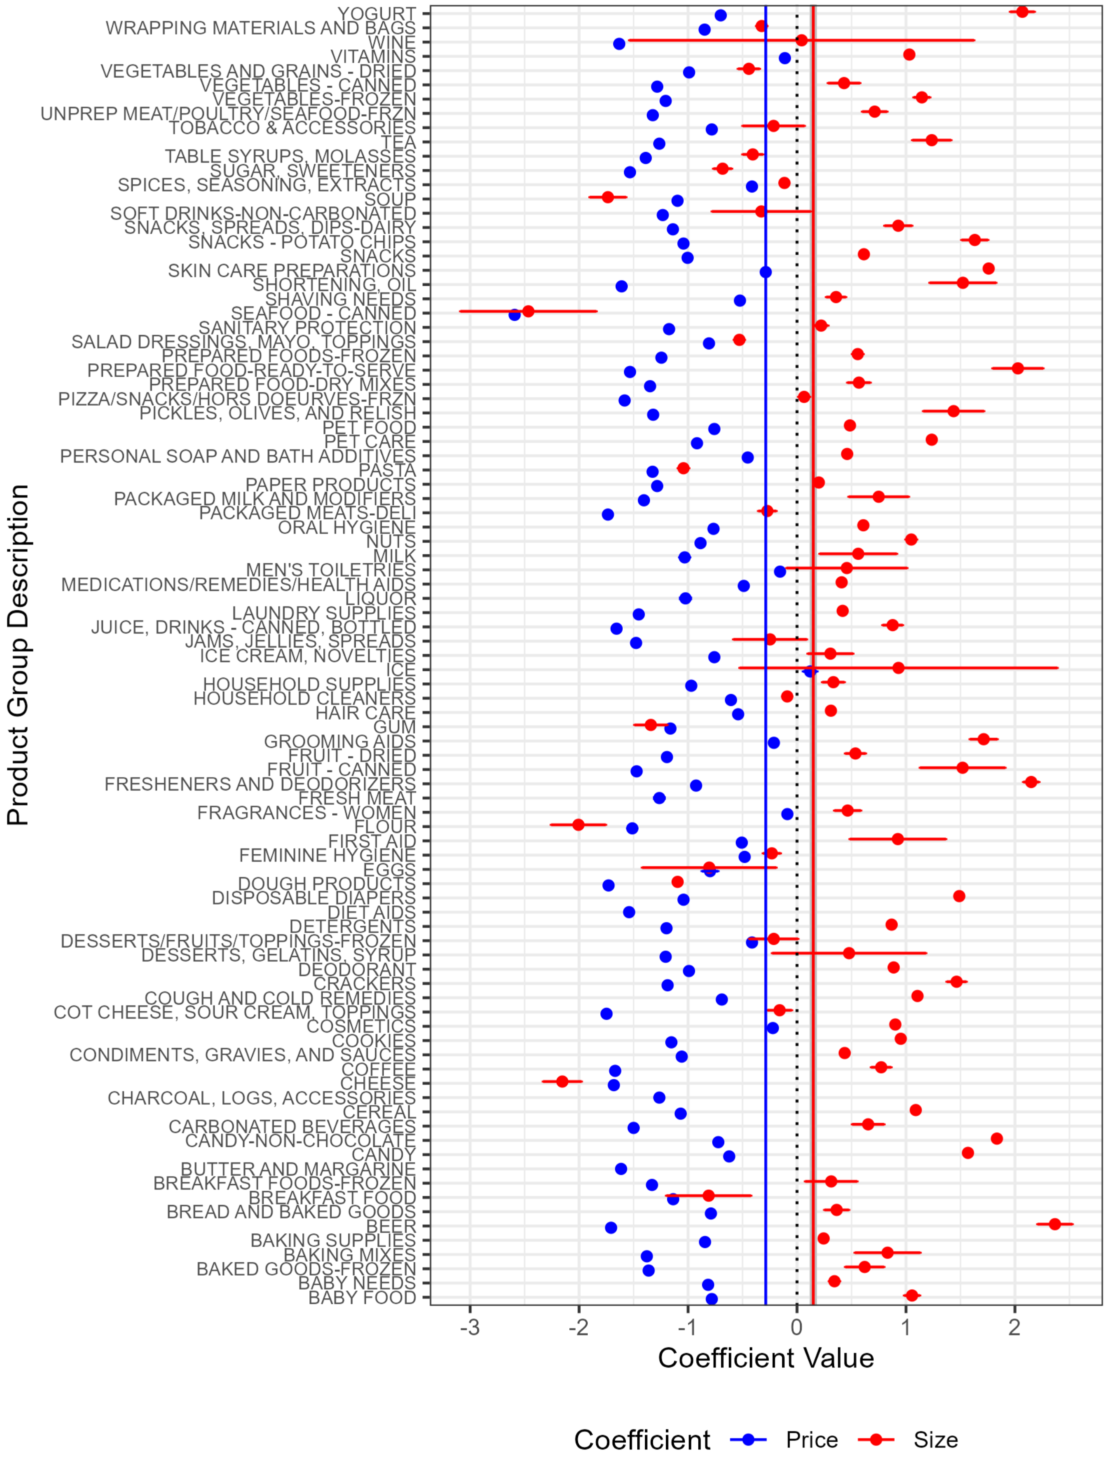

Notes: This graph presents the estimated unit price elasticities (in blue) and package size elasticities (in red) and corresponding 95% confidence intervals in different product groups for products that do not decrease in size. The coefficients, therefore, refer to $\hat{\eta}_p$ and $\hat{\eta}_l$ in model 2.The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. Error bars denote the 95% confidence intervals of the weighted coefficients. Weights are based on the sales within product groups. For visibility, we exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5. However, those product groups are part of the weighted average calculations.

Figure 12: Estimated Interaction Terms for Shrinking Products Across Different Product Groups

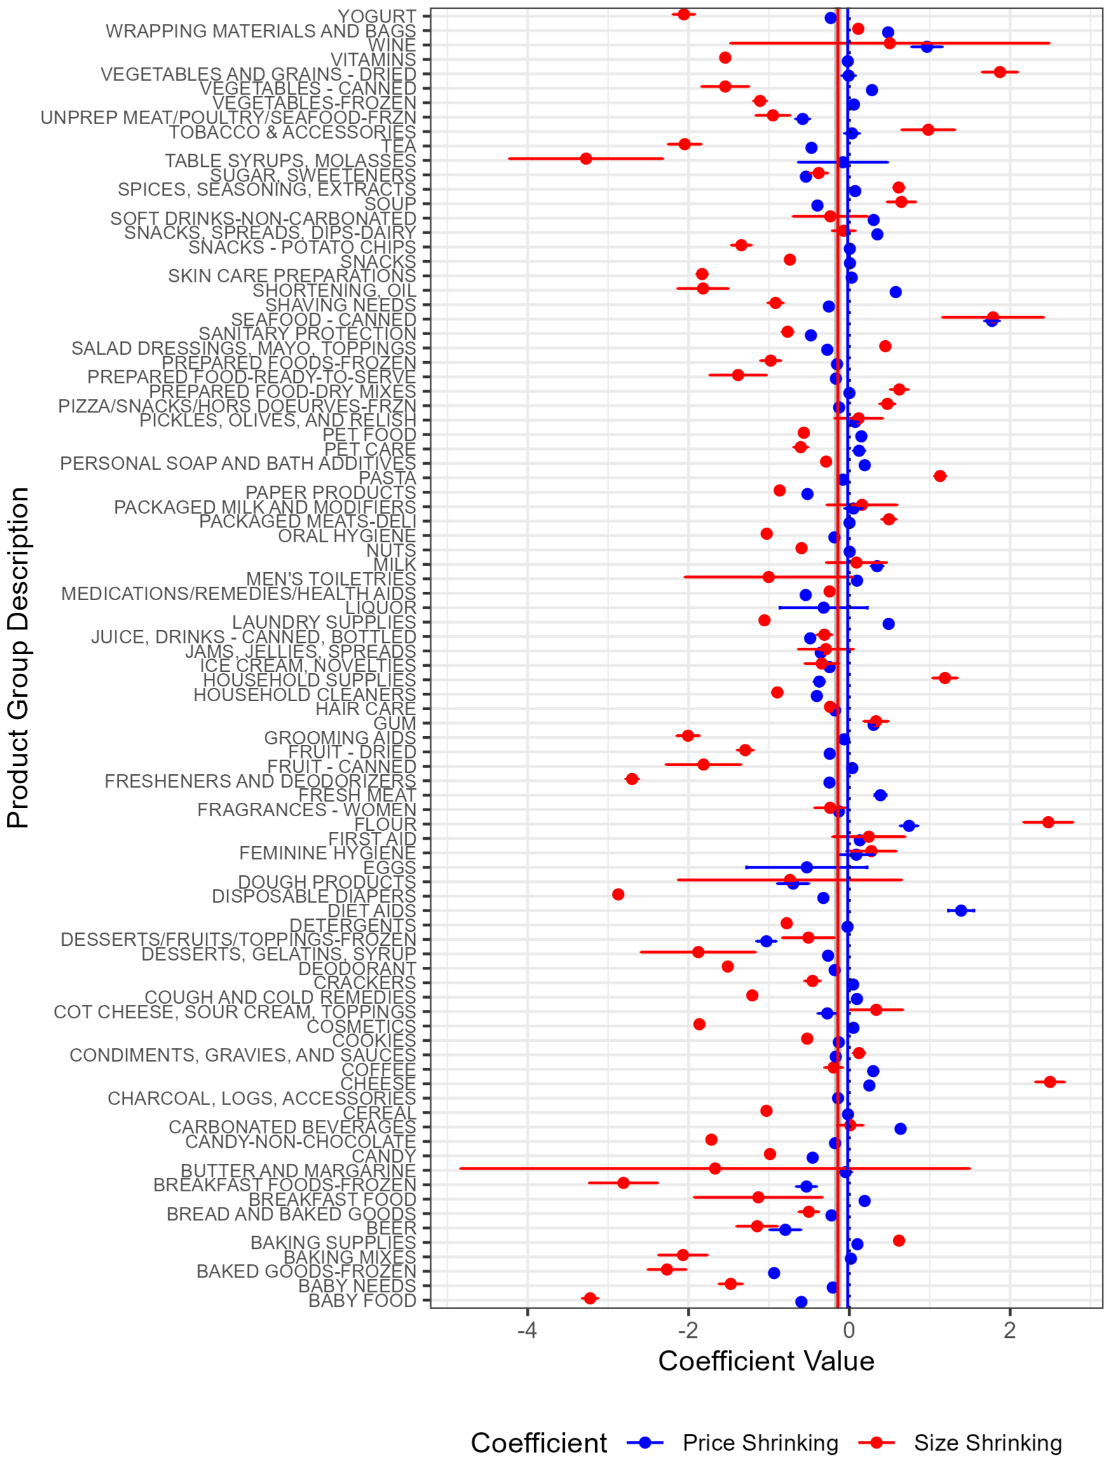

Notes: This graph presents the estimated interaction terms of equation 2. In detail, we show the differences for unit price elasticities (in blue), package size elasticities (in red), and corresponding 95% confidence intervals in different product groups for products that decrease in size. The coefficients, therefore, refer to $\hat{\beta}_p$ and $\hat{\beta}_l$ in model 2. The coefficients in each group refer to the sales-weighted average of estimated elasticities across all product modules within one product group. The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. For visibility, we exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5. However, those product groups are part of the weighted average calculations.

ucts, which significantly limits the power of our instrumental variable approach. Therefore, we exclude groups with extreme standard errors and coefficients to avoid biased results (more details in Appendix E.2).

Appendix Figure E.2 presents the results for the general unit price and product size elasticities, and Appendix Figure E.3 presents the estimated interaction terms across different product groups after excluding outliers. First, the product group-specific coefficients show higher variability; however, the general results align with the IV estimation of model 1 and the OLS regression of 2.

The price and product size elasticities for non-downsized products yield weighted average estimates of -1.80 and 1.66, respectively. Thus, for non-downsized products, we observe a strong positive size coefficient and a negative price coefficient. As in the OLS regression, we observe a small negative price-interaction coefficient (-0.12), indicating no strong difference in price elasticities between downsized and non-downsized products. In contrast, the absolute magnitude of the size-interaction coefficients is high, as detailed in Appendix Figure E.3. The average coefficient is -1.46, about the same absolute size as the size elasticity for non-downsized products. Thus, our findings from above are confirmed: consumers tend to be nearly unresponsive to size changes for shrinking products, while for non-downsized products, size elasticity is strong.

6 Discussion

This study shows that shrinkflation is widespread in the U.S. retail market, subtly increasing the price per volume for various products. Product size decreases occur significantly more often than increases, and product downsizing usually happens without equivalent price reductions. In comparison, product size increases are correlated with increases in prices. Our findings reveal that both retailers and manufacturers strategically use product size changes to boost their revenues. Using a regression framework and instrumental variable approach, we show that consumers exhibit more sensitivity to price adjustments than to changes in product size. Additionally, the sensitivity to size changes is lower when product size decreases compared to product size increases.

The diverging responsiveness to price and size changes aligns with the idea that consumers are inattentive to non-salient or hidden attributes, as numerous studies demonstrate (DellaVigna, 2009). For instance, in a seminal paper on tax salience, Chetty et al. (2009) show that consumers react less to non-salient taxes not included in posted prices than to similar salient price changes.[12] This reasoning implies our results may be driven by inattentive consumers who react less to changes in price per volume attributable to size changes as they are typically less salient than direct changes to posted unit prices.

The underreaction to non-salient price attributes incentivizes profit-maximizing firms to hide or obfuscate relevant price attributes. Numerous examples exist of firms strategically hiding or obfuscating relevant prices to exploit consumers' underreactions. Examples of such practices include add-on pricing (e.g., Ellison, 2005; Gabaix and Laibson, 2006) and hidden shipping and handling costs (e.g., Hossain and Morgan, 2006; Brown et al., 2010). According to models exploring optimal obfuscation behavior and market outcomes (e.g. Gabaix and Laibson, 2006; Ellison and Wolitzky, 2012; Janssen and Kasinger, 2024), price obfuscation decreases the price sensitivity of consumers, allowing firms to limit competitive pressure and boost their profits by obfuscating prices. Following this logic, our findings suggest shrinkflation as an effective strategy for retailers and manufacturers to increase prices non-saliently. This idea is supported by the result that for non-shrunk products, where firms typically have no incentive to hide those price changes as they benefit consumers, size and price elasticities are similar.

While beneficial for firms, the resulting increase in profits and revenues often comes at the expense of (inattentive) consumers. Consequently, policymakers may want to regulate shrinkflation behavior to protect consumers. Banning product size changes is, however, impractical and may limit beneficial size changes. A more viable policy could involve making size changes more salient and transparent—for instance, mandating that retailers display prices per volume or clearly communicate any changes in product sizes to consumers. The regulation would, therefore, follow the warning labels that we observe in French supermarkets (Vidalon, 2023). Importantly, the policy should not be solely tailored to manufacturer-induced product size decreases, given our findings that retailer-driven size changes are also common.

Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

[12]Other examples on the tax salience effect include Finkelstein (2009), Sallee (2011), Goldin and Homonoff (2013), Feldman and Ruffle (2015) and Taubinsky and Rees-Jones (2018).

References

Jonathan Barrett and Mostafa Rachwani. The curse of shrinkflation: how food is being sneakily downsized – but prices aren't. The Guardian - Australia, 2023. https://www.theguardian.com/australia-news/2023/aug/02/shrinkflation-groceries-packaging-size-decrease-brands-. Retrieved 11 October 2023.

BBC News. France's carrefour puts up 'shrinkflation' warning signs. https://www.bbc.com/news/business-66809188, 2023. Retrieved 3 January 2024.

Jennifer Brown, Tanjim Hossain, and John Morgan. Shrouded attributes and information suppression: Evidence from the field. *Quarterly Journal of Economics*, 125(2):859–876, 05 2010.

Metin Çakır. Retail pass-through of package downsizing. *Agribusiness*, 38(2):259–278, 2022.

Metin Çakır and Joseph V Balagtas. Consumer response to package downsizing: Evidence from the chicago ice cream market. *Journal of Retailing*, 90(1):1–12, 2014.

Metin Çakır, Joseph Valdes Balagtas, and Abigail M Okrent. The effects of package downsizing on food consumption. https://ageconsearch.umn.edu/record/150680/?v=pdf, 2013. Retrieved 3 March 2024.

Pierre Chandon and Nailya Ordabayeva. The accuracy of less: Natural bounds explain why quantity decreases are estimated more accurately than quantity increases. *Journal of Experimental Psychology: General*, 146(2):250, 2017.

Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77, 2009.

CNBC. What is shrinkflation? here's why consumers may be getting less for their money. https://www.cnbc.com/2024/03/17/shrinkflation-why-consumers-might-be-getting-less-for-their-money.html, 2024. Retrieved 3 March 2024.

Sam Danley. How lower-income households are responding to grocery inflation. *Supermarket Perimeter*, Aug 2022.

Paul Davidson. Inflation is battering lower-income households most as food, housing costs soar, fed study says. *USA TODAY*, Jan 2023. https://www.usatoday.com/story/money/2023/01/18/inflation-hurts-lower-income-households-food-housing-costs-rise/11074945002/. Retrieved 2 January 2024.

Stefano DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–72, 2009.

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

Hendrik Döpper, Alexander MacKay, Nathan Miller, and Joel Stiebale. Rising markups and the role of consumer preferences. *Harvard Business School Strategy Unit Working Paper*, (22-025), 2022.

Glenn Ellison. A model of add-on pricing. *Quarterly Journal of Economics*, 120(2):585–637, 2005.

Glenn Ellison and Sara Fisher Ellison. Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2):427–452, 2009.

Glenn Ellison and Alexander Wolitzky. A search cost model of obfuscation. *RAND Journal of Economics*, 43(3):417–441, 2012.

European Parliament. 'Shrinkflation' – a deceptive marketing practice. Commission protection for consumers against the phenomenon of product downsizing. 2022. https://www.europarl.europa.eu/doceo/document/E-9-2022-002635_EN.html. Retrieved 11 October 2023.

Ioannis Evangelidis. Frontiers: Shrinkflation aversion: When and why product size decreases are seen as more unfair than equivalent price increases. *Marketing Science*, 43(2):280–288, 2024.

Naomi E. Feldman and Bradley J. Ruffle. The impact of including, adding, and subtracting a tax on demand. *American Economic Journal: Economic Policy*, 7(1):95–118, 2015.

Amy Finkelstein. E-ztax: Tax salience and tax rates. *The Quarterly Journal of Economics*, 124 (3):969–1010, 2009.

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121(2):505–540, 05 2006.

Jacob Goldin and Tatiana Homonoff. Smoke gets in your eyes: Cigarette tax salience and regressivity. *American Economic Journal: Economic Policy*, 5(1):302–336, 2013.

GS1US. Gtin management guidelines, net content. https://www.gs1.org/1/gtinrules/en/rule/266/declared-net-content, 2023. Retrieved 3 January 2024.

Katherine Harris-Lagoudakis, John Crespi, and Xibo Wan. Shrinkflation, collusion & unit price disclosure: Evidence from canned tuna. *Working paper*, December 2023.

Jerry A Hausman. Valuation of new goods under perfect and imperfect competition. In *The Economics of New Goods*, pages 207–248. University of Chicago Press, 1996.

Günter J Hitsch, Ali Hortacsu, and Xiliang Lin. Prices and promotions in us retail markets. *Quantitative Marketing and Economics*, 19(3):289–368, 2021.

Tanjim Hossain and John Morgan. ... plus shipping and handling: Revenue (non) equivalence in field experiments on ebay. *Advances in Economic Analysis & Policy*, 5(2):1–27, 2006.

Aljoscha Janssen and Johannes Kasinger. Obfuscation and rational inattention. *The Journal of Industrial Economics*, 72(1):390–428, 2024.

In Kyung Kim. Consumers' preference for downsizing over package price increases. *Journal of Economics & Management Strategy*, 33(1):25–52, 2024.

Phil Lempert. Shoppers speak up! #shrinkflation and #foodcrimes in their supermarkets. Forbes, 2023. https://www.forbes.com/sites/phillempert/2023/02/27/shoppers-speak-up-shrinkflation-and-foodcrimes-in-their-supermarkets/. Retrieved 11 October 2023.

Kari McNair. Getting less for the same price? explore how the cpi measures "shrinkflation" and its impact on inflation. Beyond the Numbers: Prices & Spending, vol. 12, no. 2 (U.S. Bureau of Labor Statistics), 2023. https://www.bls.gov/opub/btn/volume-12/measuring-shrinkflation-and-its-impact-on-inflation.htm. Retrieved 11 October 2023.

Ian Meeker. Does peter piper pick a package of pepper inattentively? the consumer response to product size changes. *The Consumer Response to Product Size Changes (September 14, 2021)*, 2021.

Sarah Moshary, Anna Tuchman, and Natasha Vajravelu. Gender-based pricing in consumer packaged goods: A pink tax? *Marketing Science*, 2023.

Brent Neiman and Joseph Vavra. The rise of niche consumption. *American Economic Journal: Macroeconomics*, 15(3):224–264, 2023.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2): 307–342, 2001.

Nadezhda Ochirova. The impact of shrinkflation on cpih, uk: January 2012 to june 2017. 2017. URL https://www.ons.gov.uk/economy/inflationandpriceindices/articles/theimpactofshrinkflationoncpihuk/january2012tojune2017.

Nailya Ordabayeva and Pierre Chandon. Predicting and managing consumers' package size impressions. *Journal of Marketing*, 77(5):123–137, 2013.

Christian Rojas, Edward C Jaenicke, and Elina Page. Shrinkflation? quantifying the impact of changes in package size on food inflation. *Quantifying the Impact of Changes in Package Size on Food Inflation (April 23, 2024)*, 2024.

James M Sallee. The surprising incidence of tax credits for the toyota prius. *American Economic Journal: Economic Policy*, 3(2):189–219, 2011.

Bradley T Shapiro, Günter J Hitsch, and Anna E Tuchman. TV advertising effectiveness and profitability: Generalizable results from 288 brands. *Econometrica*, 89(4):1855–1879, 2021.

Dmitry Taubinsky and Alex Rees-Jones. Attention variation and welfare: Theory and evidence from a tax salience experiment. *The Review of Economic Studies*, 85(4):2462–2496, 2018.

The Associated Press. 'shrinkflation' accelerates globally as manufacturers quietly shrink package sizes. https://www.npr.org/2022/06/08/1103766334/shrinkflation-globally-manufacturers-shrink-package-sizes, 2022. Retrieved 3 January 2024.

The Korea Times. Korea pushing to require shrinkflation signs on products for downsizing. https://www.koreatimes.co.kr/www/biz/2023/12/602_365009.html, 2023. Retrieved 3 January 2024.

Clare Toeniskoetter. Meet the man on a mission to expose sneaky price increases. New York Times, 2022. https://www.nytimes.com/2022/11/26/climate/fighting-shrinkflation-consumer-products.html. Retrieved 11 October 2023.

Dominique Vidalon. Carrefour sticks price warnings on food to shame suppliers. Reuters, 2023. https://www.reuters.com/business/retail-consumer/carrefour-sticks-price-warnings-food-shame-suppliers-2023-09-14/. Retrieved 11 October 2023.

Ryan Webb, Clarice Zhao, Matthew Osborne, Peter Landry, and Colin Camerer. A neuro-autopilot theory of habit: Evidence from canned tuna. *Rotman School of Management Working Paper*, (4130496), June 2022.

Jun Yao, Harmen Oppewal, and Di Wang. Cheaper and smaller or more expensive and larger: How consumers respond to unit price increase tactics that simultaneously change product price and package size. *Journal of the Academy of Marketing Science*, 48(6):1075–1094, 2020.

Koichi Yonezawa and Timothy J Richards. Competitive package size decisions. *Journal of Retailing*, 92(4):445–469, 2016.

Online Appendix

A Additional Summary Statistics

Tables A.1 provide summary statistics for various configurations of the raw dataset, showing how the dataset changes as we impose additional restrictions. Descriptive statistics in each column are calculated using the underlying data set of the preceding column imposing one additional restriction. It displays summary statistics for (1) the full sample; (2) the top 80%-percentile brands within each product module; (3) only considering food, mass merchandise, and drug stores; (4) only considering products with size changes, and (5) the final imputed data used for our regression analysis.

Table A.2 focuses on products that experienced size changes, comparing downsized and upsized products against the full sample for food, mass merchandise, and drug stores. This table is equivalent to Table 1 in the main text, but it only considers food, mass merchandise, and drug stores.

B Price Imputation

In the RMS retail scanner data, price information is only available for weeks when a product is sold. This leads to missing price observations, especially for smaller products with rare sales. We use the price imputation algorithm developed by Hitsch et al. (2021) to overcome this issue. The algorithm works by identifying base prices and promoted prices. It does so by examining the typical price pattern at the store level, where base prices are stable for some time, followed by shorter periods of promoted prices. The algorithm can fill in missing price data, ensuring a more accurate representation of the full price spectrum for these products. It assumes that weeks with no sales are periods when prices are non-promoted. The algorithm imputes prices if, for the last observed (base) price, i.e. the reference price, there is another observed price (close enough to the reference price) within a 12-week window. If this is the case, prices that are missing between these two observed prices will be set equal to the reference price, assuming that the reference price is correct and stable within the specified window. As a result, the price in a respective period can still be missing if there is no observed price within 12 weeks from the last observed price. For more details see the online appendix of Hitsch et al. (2021)

We only consider products, as defined by their UPC, that experienced either a product size increase or decrease in at least one store within the U.S. across the observation period. Columns 4 and 5 in Table A.1 compare aggregated summary statistics of the original movement files with those of the imputed dataset. Through the price imputation algorithm, we increase the number of observations from near 16.6 billion to 25.4 billion observations, implying that around 34% of the observations in the final data set are imputed. A ratio that is similar to Hitsch et al. (2021). Unsurprisingly, the average units and value of purchase are both lower in the imputed dataset as it now includes zero quantity observations. Similarly, it is sensible that the average price is slightly higher in the imputed data set, given that lower prices are positively associated with quantity. Average sizes slightly increase.[13]

C Are Products Identical?

C.1 Visual Inspection

In our primary analysis, we focus on identifying products that exhibit decreasing or increasing trends. This is achieved by examining products from the same brands and analyzing UPC descriptions that are substituted at the store level with either a smaller or larger equivalent. However, this approach is not without its risks, particularly in terms of potentially mislabeling certain products. For instance, it's possible that a product is replaced by another version that is not only different in size but also distinct in its marketing approach. Consider a scenario where a product undergoes a rebranding or packaging overhaul. While our experience suggests that such concurrent changes are common, we acknowledge this limitation and address it by implementing several robustness checks in the subsequent sections of our study.

In the first robustness check, we examine each individual UPC associated with products that have demonstrated either a decrease or increase in size. This involves a detailed comparison of packaging to ascertain whether the products in question are truly identical. Given the sheer volume of products across all modules, manual verification for each is impractical. Therefore, we have chosen to concentrate on specific modules where changes in product size are most commonly observed: Candy, Snacks, Detergents, Hair Care, and Cereals.

For each product within these modules, we utilize a UPC lookup database (accessible at https://www.upcitemdb.com/) to accurately determine the packaging of each product both

[13] The slight differences in total sales, total units, and unique store-UPCs between Columns 4 and 5 of Table A.1 can be attributed to instances where products share the same (reused) UPC within a year—making the products distinguishable only by the UPC's version code. When this occurs, we consider only the product with the highest sales within that year.

prior to and following any changes in size. However, it's important to note that this database has its limitations, particularly concerning products that were phased out more than a few years ago and are no longer on the market. Consequently, we rely on a subset of products with confirmed size changes as a representative sample to demonstrate the validity of our analysis. The remaining products may still be valid product size changes, but we are not able to confirm the changes visually.

In this analysis, we present a detailed comparison between the complete collection of products that have undergone size reduction and those specifically verified as identical despite this decrease. The relevant statistics for the product group of snacks are detailed in Columns 1 and 2 of Table C.1.[14] Our focus centers on key aspects such as variations in product size, price dynamics, and sales trends before and after the size reduction event.

In over 800,000 instances, we have been able to confirm that the size reduction corresponds to identical products. This figure stands in comparison to the 1,320,041 size reduction events reported in the main study. Each observation here represents a single instance of size reduction within a store, suggesting that a large proportion of these reductions are indeed with the same product.

It is also important to note that the unconfirmed observations do not necessarily imply inaccuracies or different products. In these cases, the issue often lies in the insufficient clarity to definitively confirm the product's identity.

When analyzing both sets of products against our selected metrics, only slight differences emerge. Crucially, trends in size reduction, price adjustments, and sales fluctuations are remarkably similar across the general range of products and those where the identity post-reduction has been confirmed. This uniformity leads us to conclude that, apart from the change in size, the products remain essentially identical in other respects. This insight offers a valuable perspective into the market dynamics associated with changes in product size.

[14] We also show corresponding statistics for the product group of hair-care, detergent, and candy in Table C.1. Results are comparable.

Table A.1: Summary Statistics by Restrictions

	(1) Full data	(2) Top 80%-pctl. brands	(3) Food, Mass merch. & Drug stores	(4) Shrunk/increased UPCs	(5) Imputed Data
Total Observations	143.99 Bn	103.84 Bn	101.19 Bn	16.62 Bn	25.39 Bn
Unique UPCs	1,552,445	540,182	533,239	30,019	30,019
Unique Stores	59,352	59,345	44,537	44,534	44,534
Unique Retailers	141	141	121	121	121
Unique UPC-Store	2,551.12 Mn	1,554.92 Mn	1,514.60 Mn	181.57 Mn	181.46 Mn
Total Sales	2,643.62 Bn	2,044.99 Bn	1,974.14 Bn	349.66 Bn	348.71 Bn
Total Units	818.76 Bn	635.40 Bn	616.87 Bn	128.47 Bn	127.98 Bn
Avg. Price	4.67	4.71	4.72	4.01	4.50
Pooled SD Price	4.12	4.10	4.01	2.65	2.97
Avg. Product Size	29.89	30.34	30.09	24.73	25.55
Pooled SD Product Size	51.39	53.54	52.64	26.32	28.32

Notes: This table displays summary statistics for various configurations of our dataset. Column 1 shows statistics for the entire dataset. Column 2 restricts the dataset to brands comprising at least 80% of sales within a product module. Column 3 includes only food stores, mass merchandisers, and drug stores. Column 4 further limits the data to UPCs that underwent a product size change in at least one store during our study period, as defined in our methodology. Column 5 shows the statistics for the imputed data set based on the dataset in column 4. Average prices and sizes are calculated annually by product module and averaged across years and modules, weighted by the respective shares of total observations. Similarly, standard deviations are computed annually by product module and averaged across years and modules using the formulas: $SD_p = \sqrt{\sum_t (n_t - 1) \times SD_t^2 / \sum_t (n_t - 1)}$ *and* $SD_{all} = \sqrt{\sum_p (n_p - 1) \times SD_p^2 / \sum_p (n_p - 1)}$.

Table A.2: Summary Statistics of Product with Size Changes, Restricted Stores

	Downsized products	Upsized Products	All products
Total Observations	1.83 Bn	0.34 Bn	101.19 Bn
Unique UPCs	24,449	15,600	533,239
Unique Stores	44,268	44,132	44,537
Unique Retailers	120	119	121
Unique UPC-Store	17.34 Mn.	4.01 Mn.	1,514.60 Mn.
Total Sales	37.18 Bn	7.27 Bn	1,974.14 Bn
Total Units	14.09 Bn	2.27 Bn	616.87 Bn
Average Price (in $)	3.53	4.66	4.72
	(2.03)	(2.51)	(4.01)
Average Product Size	21.45	30.20	30.09
	(16.76)	(20.68)	(52.64)

Notes: This table compares summary statistics of downsized and upsized products to those of all products. Only food, mass merchandise, and drug stores, as well as top 80%-percentiles brands within each module, are considered. Column 1 and 2 shows statistics for shrunk and increased products, considering all sales before and after the products size changes in stores that changed the product size. Column 3 shows the summary stats considering all sales in the sample. Average prices and sizes are calculated annually by product module and averaged across years and modules, weighted by the respective shares of total observations. Similarly, standard deviations (in parentheses) are computed annually by product module and averaged across years and modules using the formulas: $SD_p = \sqrt{\sum_t (n_t - 1) \times SD_t^2 / \sum_t (n_t - 1)}$ *and* $SD_{all} = \sqrt{\sum_p (n_p - 1) \times SD_p^2 / \sum_p (n_p - 1)}$.

Table C.1: Comparison of Products That Are Confirmed by Visual Inspection in Different Product Groups

	Snacks		Candy		Detergents		Hair Care	
	All	Confirmed	All	Confirmed	All	Confirmed	All	Confirmed
Number of Observations	1,320,041	800,658	910,657	459,837	344,696	223,040	325,286	283,199
Avg Size Before Decrease	7.59	7.79	13.13	14.7	36.75	33.84	16.53	16.25
	(3.98)	(3.65)	(10.23)	(11.47)	(32.93)	(30.28)	(7.42)	(7.14)
Avg Size After Decrease	7.06	7.29	11.79	13.31	32.25	29.69	14.62	14.39
	(3.73)	(3.45)	(9.19)	(10.24)	(28.9)	(25.83)	(6.33)	(6.15)
Avg Price Before Decrease	2.57	2.64	3.75	4.59	4.54	4.25	4.12	3.97
	(1.14)	(1.14)	(2.89)	(3.09)	(3.84)	(3.47)	(1.63)	(1.41)
Avg Price After Decrease	2.56	2.63	3.8	4.67	4.35	4.06	4.23	4.13
	(1.14)	(1.15)	(2.92)	(3.11)	(3.62)	(3.23)	(1.67)	(1.57)
Avg Price per Volume Before Decrease	3.84	3.96	5.37	4.69	12.81	13.36	6.03	6.22
	(6.68)	(6.95)	(11.04)	(10.79)	(33.42)	(35.68)	(14.7)	(15.02)
Avg Price per Volume After Decrease	3.51	3.6	4.39	3.6	11.3	11.54	7.03	7.35
	(5.95)	(6.16)	(7.42)	(6.22)	(25.36)	(27.22)	(20.62)	(21.36)
Avg Weekly Units in Store Sold Before Decrease	9.63	10.06	3.61	3.64	4.86	5.35	1.81	1.8
	(13.97)	(14.92)	(5.15)	(5.85)	(6.38)	(6.96)	(1.22)	(1.14)
Avg Weekly Units in Store Sold After Decrease	9.67	10.12	3.56	3.58	4.76	5.22	1.79	1.77
	(14.34)	(15.23)	(5.01)	(5.63)	(6.1)	(6.6)	(1.28)	(1.23)
Avg Weekly Sales in Store Sold Before Decrease	22.73	24.35	10.59	12.08	15.59	16.17	7.2	6.78
	(34.08)	(37.14)	(14.37)	(12.51)	(31.17)	(19.51)	(7.18)	(5.54)
Avg Weekly Sales in Store Sold After Decrease	23.11	24.63	10.73	12.31	14.95	15.29	7.25	6.92
	(35.57)	(38.28)	(14.2)	(12.58)	(25.64)	(17.63)	(7.28)	(6.26)

Notes: This table presents a comparative analysis of products within various subsegments (snacks, candy, detergents, hair care), focusing on those that underwent a size reduction. It differentiates between products confirmed to be identical despite the size decrease and their counterparts in terms of pricing adjustments. The confirmation is based on individual UPC lookups and visual inspection of the package. Key metrics such as changes in size, price fluctuations, and sales variations before and after the size reduction event are examined. The dataset encompasses the entire range of products, offering insights into the typical market behavior surrounding size and price modifications in these categories. Standard deviations are indicated in parentheses to provide a statistical perspective on the data's variability.

C.2 Household Data

In our second robustness check, we employ Nielsen's household scanner data to examine whether households consistently purchase identical products before and after a change in package size. Although there might be variations in consumer behavior, this pattern generally supports the accurate identification of products undergoing size changes. However, it is important to note that the Nielsen household scanner data is not perfectly suited for this analysis. Product size changes at the store-level are infrequently observed, and their occurrence is relatively rare compared to the number of households in the panel purchasing products within a specific segment. Consequently, our analysis is conducted on a small sample of observed households within a specific store.

We investigate each product size reduction recorded over the 10 years within our sample, focusing on households that purchased a product which decreased in size at least four times in independent shopping trips within one year before and after the size reduction. This criterion ensures that the analysis includes only households that regularly purchase the product, rather than those trying it incidentally without awareness of previous price changes, thereby precluding their purchase decisions from being influenced by size alterations.

Figure C.1 presents the findings, depicting the number of households that exclusively purchased the product either before, after, or both before and after the product size decrease. The data reveal that repeated purchases both before and after the size reduction are the most common. The frequencies of households that exclusively purchase the product either before or after the size decrease are relatively similar, with a marginal tendency towards more purchases before the size reduction.

While this analysis does not provide in-depth insights into the household choices regarding product size changes—due to infrequent purchases and the limited number of households buying specific products that decrease in size—the summary statistics suggest a higher likelihood of households purchasing the product both before and after a size increase. This observation further corroborates our belief that the products undergoing size decreases are identical in all other dimensions besides size.

Figure C.1: Repeated Purchases By Households

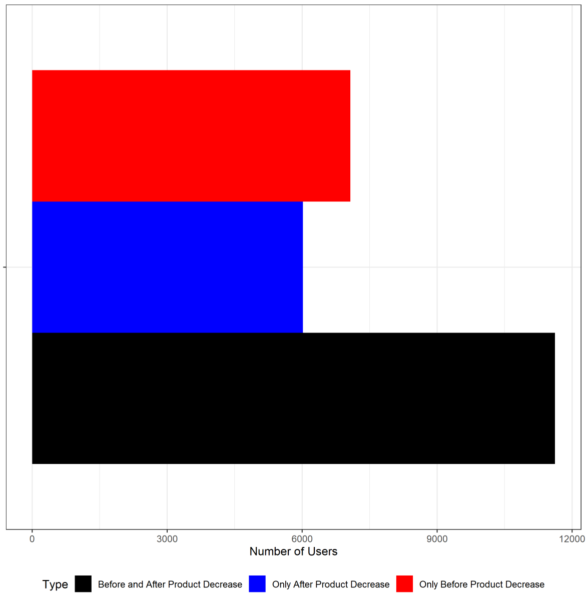

Note: This graph shows the number of households purchasing products that decrease in size. Specifically, it includes households with at least four purchase occasions of products experiencing size decreases in the year prior and following the event of a decrease. The graph categorizes the households into three groups: those who only purchased the product before the size decrease (red), after the size decrease (blue), and both before and after the size decrease (black).

D Shrinking a Product without changing the UPC

We primarily identify a product that is reducing in size at the store level using the following criteria: we notice a new UPC from the same brand entering the system. This new UPC, while having an identical brand and description to the previous UPC, is up to 25% smaller. Concurrently, the old UPC exits the system. However, our data source does not enable us to detect product shrinkage when a smaller product is substituted under the same UPC code. Since product details for a specific UPC are only updated at the start of the year, identifying the exact time of substitution for products with unchanged UPC codes is not feasible.

In this section, we contend that neglecting size reductions of same UPC products does not compromise the integrity of our analysis for two primary reasons. First, it's important to recognize the prevalent practices among retailers and manufacturers. Most adhere to guidelines set forth by the Uniform Code Council, now recognized as GS1 US, the originators of the UPC system. The UPC is predominantly designed to identify individual consumer products at retail point-of-sale, serving as a distinct barcode predominantly utilized in North America for retail commodities.[15] As per the management standards stipulated by GS1, any alteration in net

[15]It's worth noting that the UPC is frequently used interchangeably with the GTIN (Global Shipment Identification Number). The GTIN varies in length, but the UPC-A barcode, the version most prevalent at points of

content mandates the introduction of a new GTIN, and consequently, a new UPC (GS1US, 2023). To verify the adherence of manufacturers to this practice, we explored popular consumer forums like https://www.reddit.com/r/shrinkflation. Here, consumers voluntarily report recent instances of product shrinkage. In the majority of these cases where UPC information was accessible, we observed modifications to the UPCs. However, it's crucial to highlight that there's no stringent legal mandate enforcing UPC alterations for every size change, which implies that some manufacturers might retain an unchanged UPC despite modifying product dimensions.

Secondly, we approach this from an empirical perspective. While our current methods do not permit us to pinpoint instances where products under the same UPC experience size reductions (given our inability to temporally and geographically trace the introduction of an identical UPC at the store-level), we can juxtapose the overall count of such UPCs exhibiting varied sizes against the tally of multiple UPCs manifesting similar size variations nationwide from 2010 to 2020. By contrasting the aggregate of products with unchanged versus altered UPCs that display size differences within a 25% range, we gain preliminary insight into the frequency of occurrences where products retain their original UPC despite size adjustments.

In Figure D.1, we present the proportion of products within each category that exhibit multiple size variations within a 25% range, while maintaining identical brand, brand description, and UPC description. We distinguish between products with differing UPCs (depicted in blue) and those with identical UPCs (depicted in red). On average, 9.7% of products display size variations alongside different UPCs, despite being marketed under the same brand with identical descriptions. By contrast, a mere 0.4% of products show size differences under the same UPC with the same brand and description. This marked discrepancy suggests that instances of product size changes without corresponding UPC updates are relatively rare. It reinforces our belief that the vast majority of manufacturers and retailers likely adhere to GS1's management standards when implementing product changes such as size reductions or substitutions.

E Instrumental Variables

E.1 First Stage

In our demand analysis, we utilize two primary instruments for regression equations 1 and 2: averages of weekly in-store prices and product sizes across other stores in the same chain but located outside the respective store's DMA. When considering the size instrument running these

sale and within our dataset, employs data from the GTIN-12 (GS1US, 2023).

Figure D.1: Different Sized Products, Same and Different UPCs

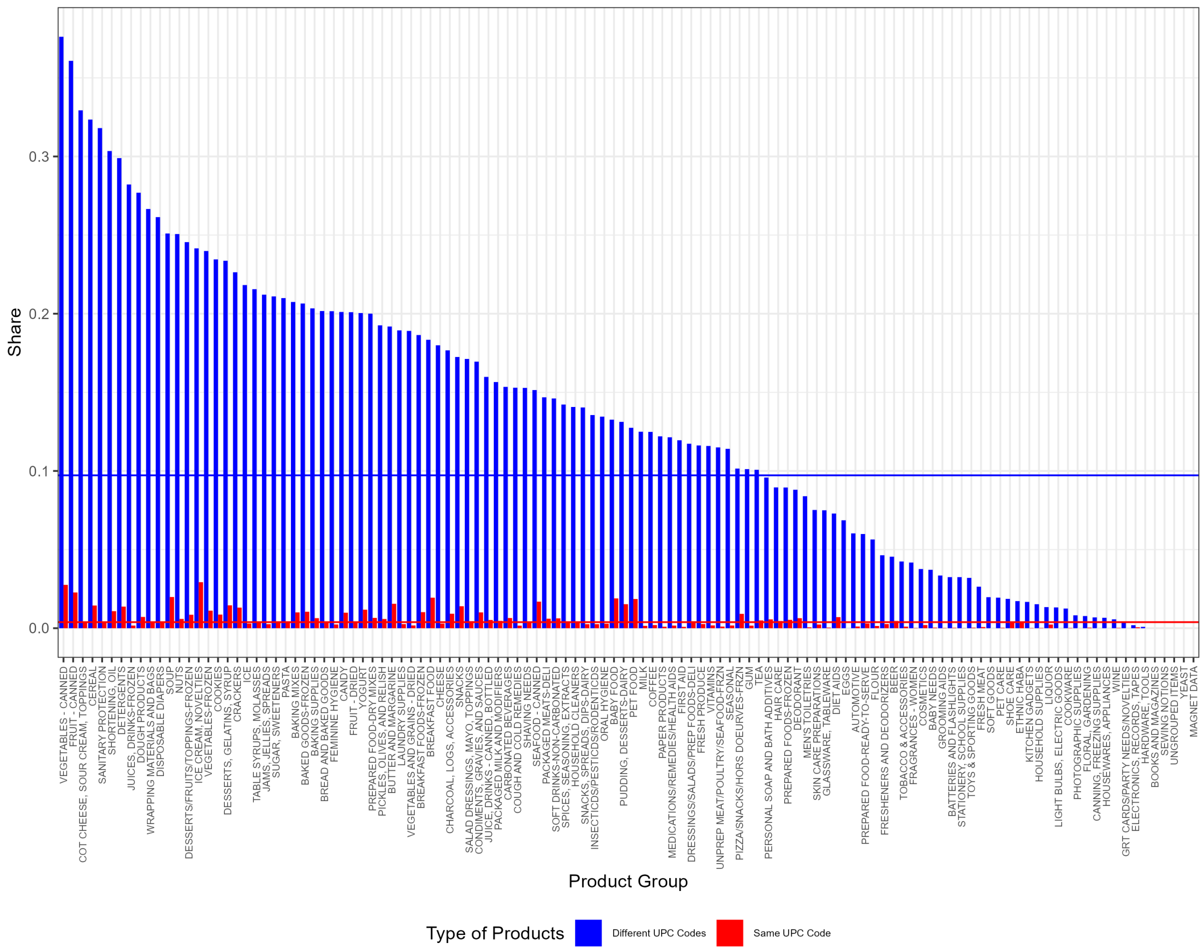

Notes: This graph illustrates the proportion of products that feature size variations within a 25% range, while retaining identical brand attributes and UPC descriptions. Each bar represents a distinct product category. The data is differentiated by products with variant UPCs (in blue) and those with consistent UPCs (in red). The solid blue and red lines indicate the average percentage of products with multiple sizes across all categories, corresponding to different and identical UPC scenarios, respectively.

regressions at the product group level can sometimes lead to weak instruments, especially after accounting for store-product and week-fixed effects. This is due to the potentially insufficient correlation between product sizes in different stores of the same chain. Although we expect retailer-induced changes in product sizes to occur at the chain level, independent decisions at individual stores may result in limited correlation in the first stage of our instrumental variable regression.

Our analysis begins by highlighting a few product groups with comparatively low F-statistics of the first-stage size instruments, though such cases are not prevalent. We then examine the robustness of our results through various exclusions of product groups with lower F-statistics. Specifically, Figure E.1a illustrates the correlation between the first-stage F-statistic for the size instrument and the second-stage size elasticity estimate. Each point on the graph represents a product group, with point sizes proportional to sales. Here, we note several extreme outliers for product groups with relatively low F-statistics. Due to the large dataset and the frequent observation of unchanging product sizes within a chain's store, F-statistics are often inflated. To address this, we make a minor adjustment in Figure E.1b by dividing the F-statistics by the number of observations, also showing how outliers of size elasticity estimates correlate with a low ratio of these adjusted F-statistics.

Figures E.1c and E.1d explore how second-stage elasticity estimates vary when excluding product groups with low (adjusted) F-statistics. Specifically, we omit groups with F-statistics lower than the x-value shown, analyzing size and price instruments separately. For the remaining groups, we compute the weighted average of size or price elasticity and display these values. The approach in Figure E.1d follows the same logic but applies restrictions based on the adjusted F-statistics.

It's important to note that the number of product groups excluded based on these first-stage restrictions is minimal, with over 90% of groups unaffected. Excluding product groups with weaker first stages significantly enhances the precision of size elasticity estimates. Outliers in size elasticity tend to be large, negative, and associated with weak F-statistics, but these estimates stabilize after excluding groups with small F-statistics that previously showed large point estimates in the second stage. Conversely, similar exclusions in the first stage of the price instrument regression do not significantly alter the price elasticity estimates or their precision.

In conclusion, our analysis supports the exclusion of product groups with low F-statistics. Specifically, we exclude groups where the adjusted F-statistics—the ratio of F-statistics to the number of observations in the first stage—is lower than 0.05. This restriction markedly improves the precision of point estimates and eliminates implausible outliers in size elasticity.

E.2 Instrument variables specification with interaction terms

Similar to our main analysis (Equation 1), endogeneity concerns may also arise in the specification with interactions terms (Equation 2). To address these concerns, we use the fitted values from the first stage of our instrumental variable approach and additionally interact them with the $\mathbf{I}(Shrink_{i,s} = 1)$ as in Equation 2.[16]

Figure E.2 shows estimates for the general unit price and product size elasticities across different product groups, and Figure E.3 presents the corresponding interaction terms. In some product groups, the identification of coefficients for downsized products is based on a small number of products with little variation in size. The fitted values often have a large, sometimes even perfect, overlap with non-downsized products, significantly limiting the power of our instrumental variable approach. Consequently, the instrumental variable specification can yield very noisy and extreme size and size-interaction coefficients for these groups.

To prevent our results from being driven by extreme values, we exclude groups with very large standard errors and coefficients. Specifically, we exclude product groups with standard errors and coefficients that lead to rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5 at a 95%-confidence level. This restriction balances the need to avoid biases from outliers while maintaining the vast majority of product groups in the sample. Overall, our findings are robust: consumers show little responsiveness to size changes for downsized products, whereas their size elasticity for non-downsized products remains strong.

F Additional Figures

[16] The only difference to Equation 2 is that we employ non-interacted time fixed effects (γ_t) in the second stage to ensure consistency in fixed effects between the first and second stages.

Figure E.1: Evaluating First Stage of Instrumental Variable Regression

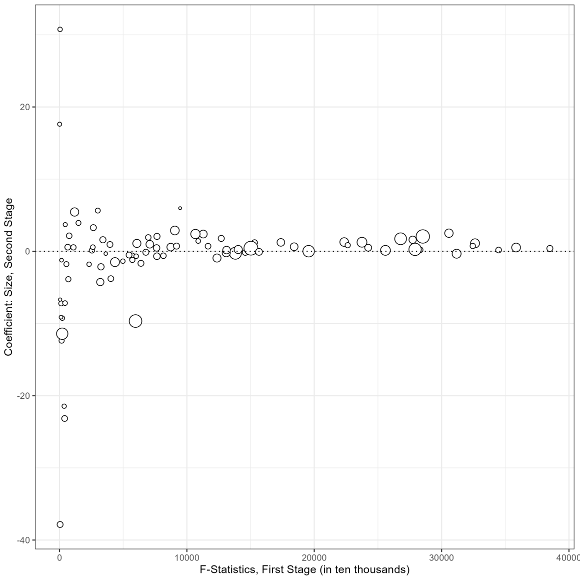

(a) F-statistics and Size Coefficient

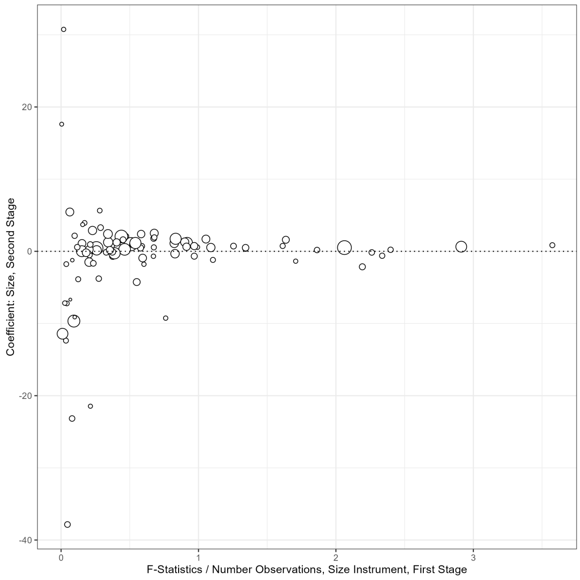

(b) Observation Adjusted F-statistics and Size Coefficients

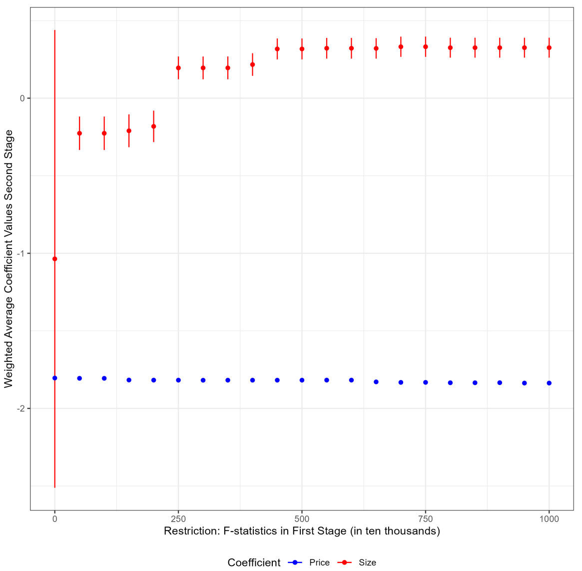

(c) Different F-statistic Restrictions

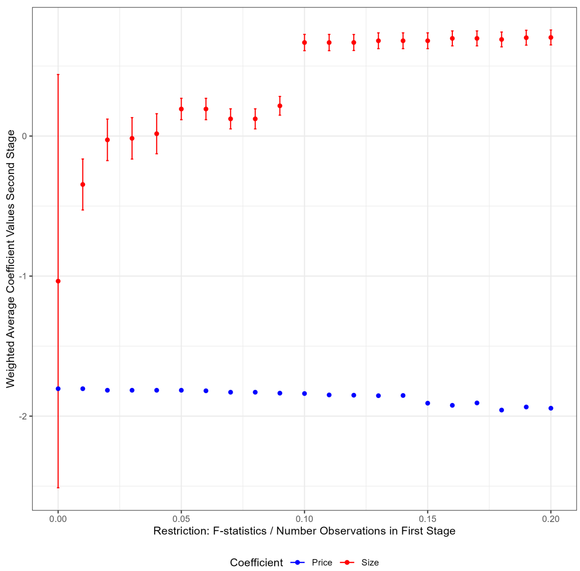

(d) Different Observation Adjusted F-statistics Restrictions

Notes: The figures show an analysis of the first stage in our instrumental variable approach of equation 1. Subfigures E.1a and E.1b show the correlation between F-statistics or F-statistics divided by the number of observations in the first stage, i.e. the regression of sizes of products in a store on the average sizes of the same product in other stores of the same chain in other markets, and the coefficient estimates of the size elasticities in the second stage. Each point corresponds to a product group, and the size of the point is based on the sales. For visibility, we exclude those product groups with extremely high F-statistics (all of the corresponding size elasticity estimates are small positives and not outliers). The graph also excludes one outlier with a second-stage estimate of smaller than -3000 and a very small F-statistic. In Subfigures E.1c and E.1d, we exclude the product group according to restrictions on the first stage that are displayed on the x-axis. We then calculate corresponding weighted averages with sales as weights of the estimated size (shown in red) and price (shown in blue) elasticities. In Subfigures E.1c, the exclusion is only based on F statistics in the first stage. For each of the instrumental variables, we exclude those product groups where the first stage F-statistic falls below a certain value displayed on the x-axis. In Subfigures E.1d, we solely consider groups where the ratio of the F-statistics and the number of observations falls below the value displayed on the x axis. The error bars correspond to the 95% confidence interval.

Figure E.2: Estimated Unit Price and Package Size Elasticities for Different Product Groups, IV Regression

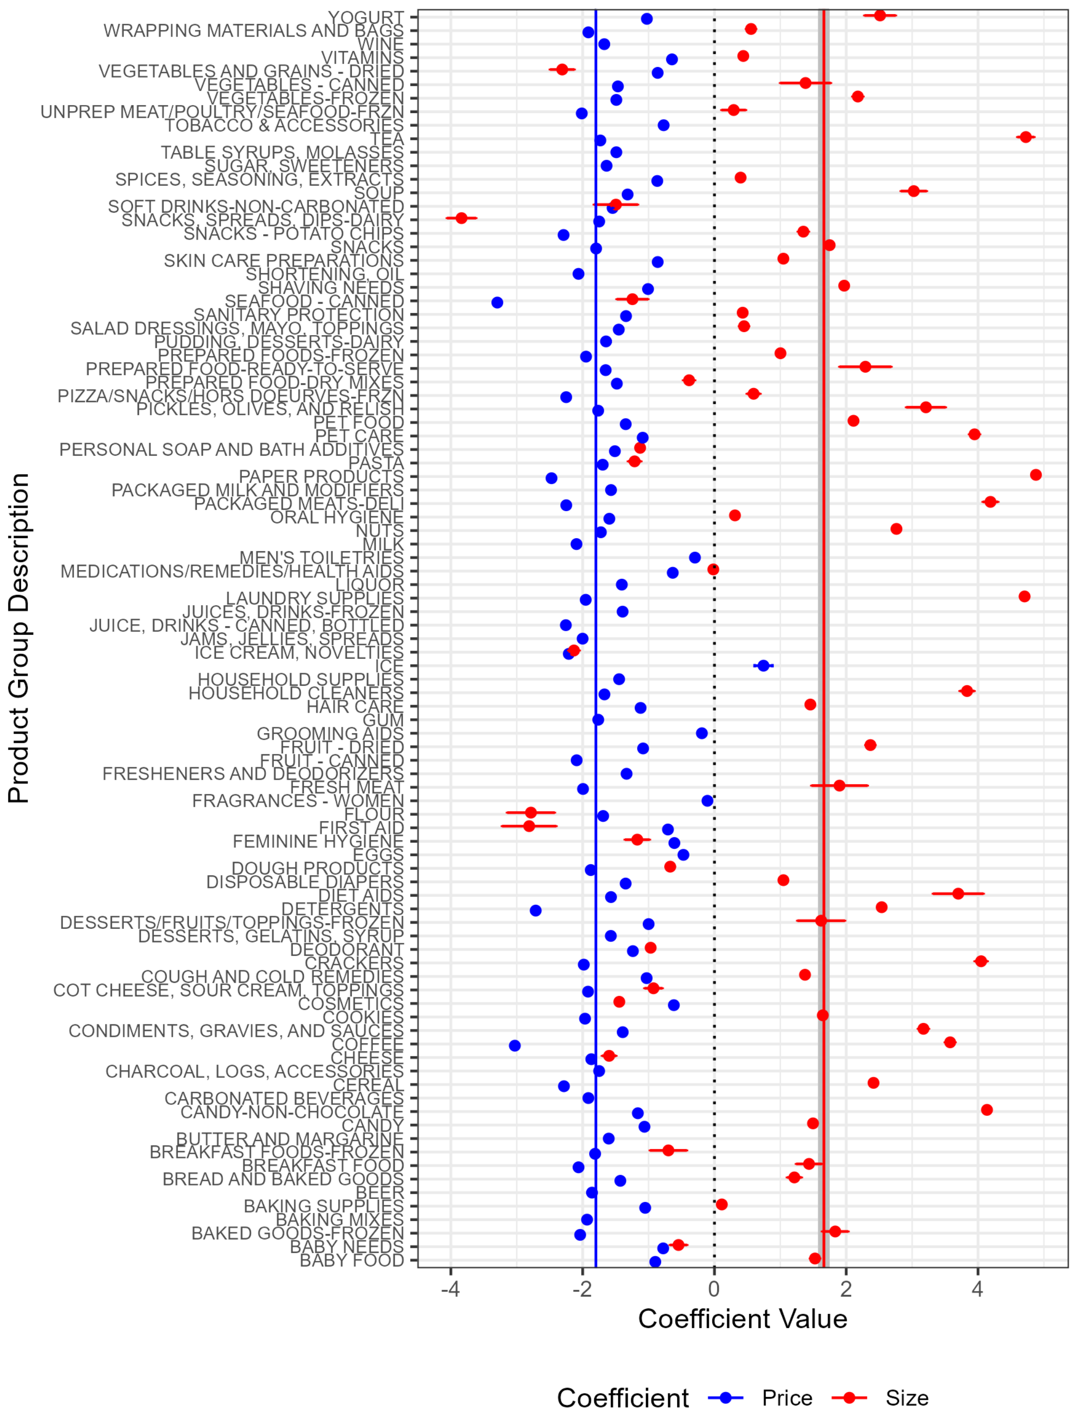

Notes: This graph presents the estimated unit price elasticities (in blue) and package size elasticities (in red) and corresponding 95% confidence intervals in different product groups for products that do not decrease in size. The coefficients, therefore, refer to $\hat{\eta}_p$ and $\hat{\eta}_l$ in model 2.The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. We instrument the weekly price and the weekly package size in store s with the average of prices and package sizes across other stores in s's chain that are located outside the respective store's DMA. Error bars denote the 95% confidence intervals of the weighted coefficients. Weights are based on the sales within product groups. We exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of ≤ -5 or ≥ 5.

Figure E.3: Estimated Interaction Terms for Shrinking Products Across Different Product Groups, IV Regression

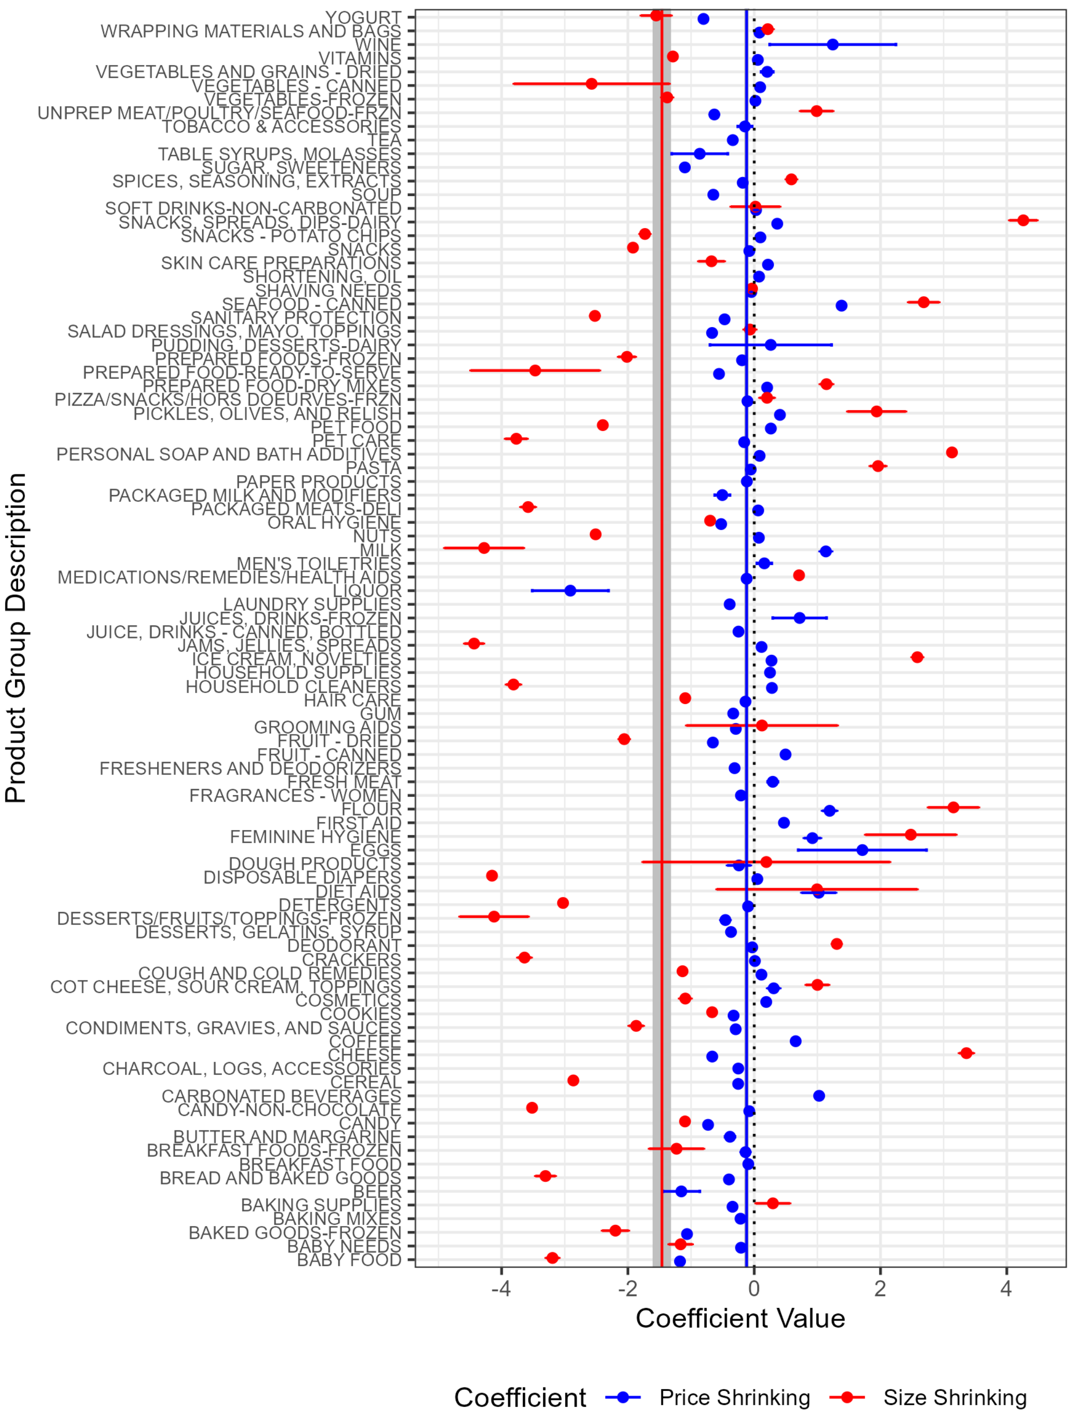

Notes: This graph presents the estimated interaction terms of equation 2. In detail, we show the differences for unit price elasticities (in blue), package size elasticities (in red), and corresponding 95% confidence intervals in different product groups for products that decrease in size. The coefficients, therefore, refer to $\hat{\beta}_p$ *and* $\hat{\beta}_l$ *in model 2. We instrument the weekly price and the weekly package size in store s with the average of prices and package sizes across other stores in s's chain that are located outside the respective store's DMA. The coefficients in each group refer to the sales-weighted average of estimated elasticities across all product modules within one product group. The solid blue line (red line) illustrates the sales-weighted average unit price (package size) elasticity across all product groups. For visibility, we exclude product groups with extremely large standard errors that result in rejecting the hypothesis that the point estimates are within the range of* ≤ -5 *or* ≥ 5*. However, those product groups are part of the weighted average calculations.*

Figure F.1: Changes of Volume Purchased across Product Groups

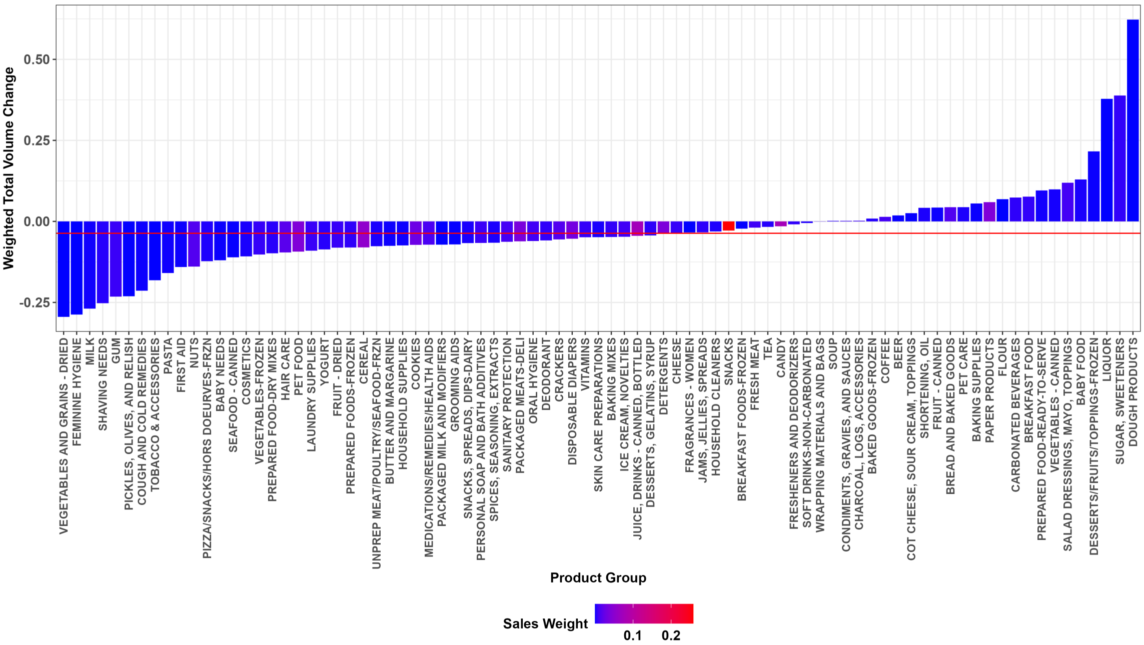

(a) Purchased Volume, Product Size Decrease

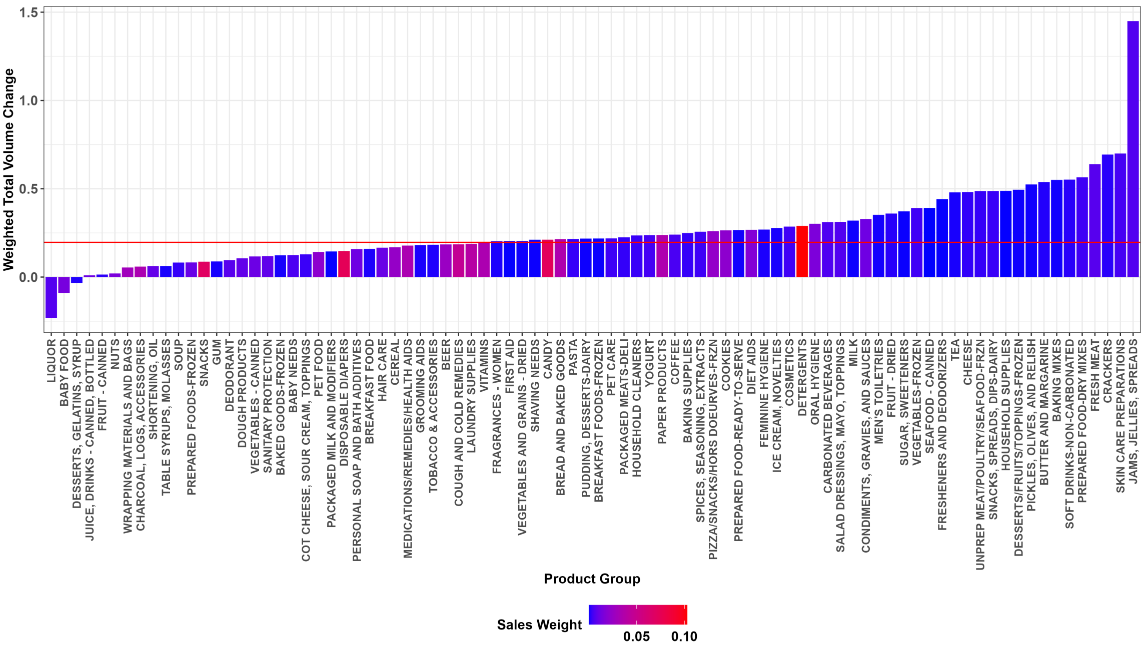

(b) Purchased Volume, Product Size Increase

Notes: This graph illustrates the purchased volume changes for product groups for product size decreases and product size increases. Subfigure F.1a investigates product size reductions, and Subfigure F.1b assesses size increases. For both scenarios, the data is aggregated from individual modules to product groups. Weighted average price per volume changes are calculated using sales volume of the products as a weighting factor. Changes in volume are measured on the product module level and considering the year before and after the product size change. The color intensity of each bar in the figure corresponds to the sales weight of the product group, with deeper shades of red indicating higher sales volumes within that group. The red line corresponds to the weighted average across all product modules.

Figure F.2: Changes of Sales across Product Groups

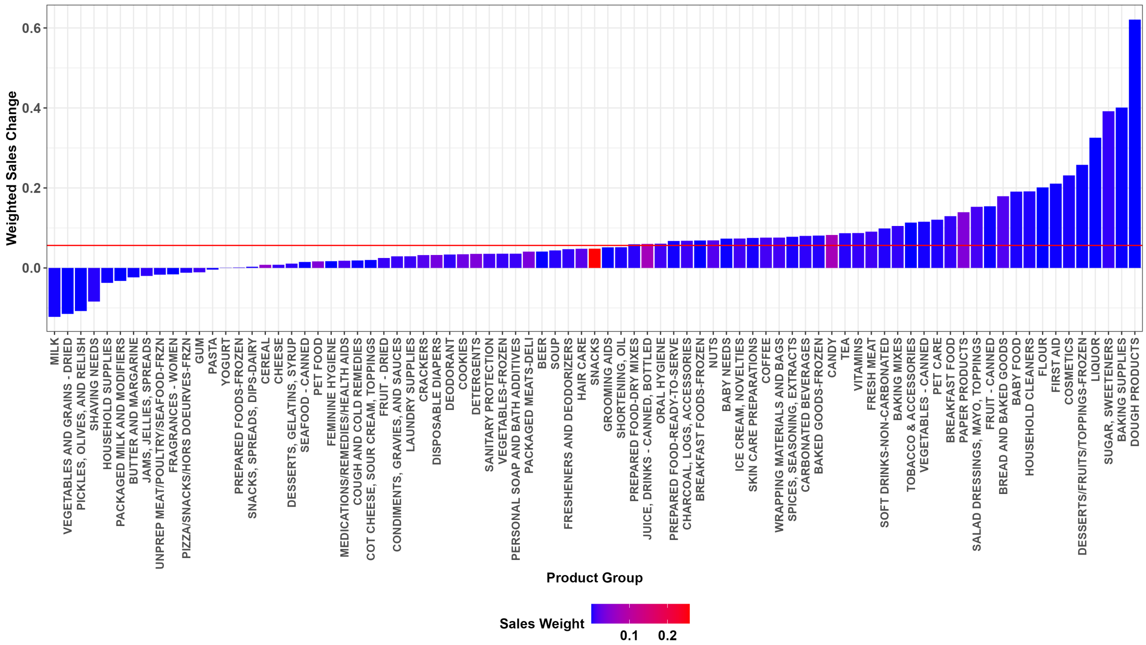

(a) Sales, Product Size Decrease

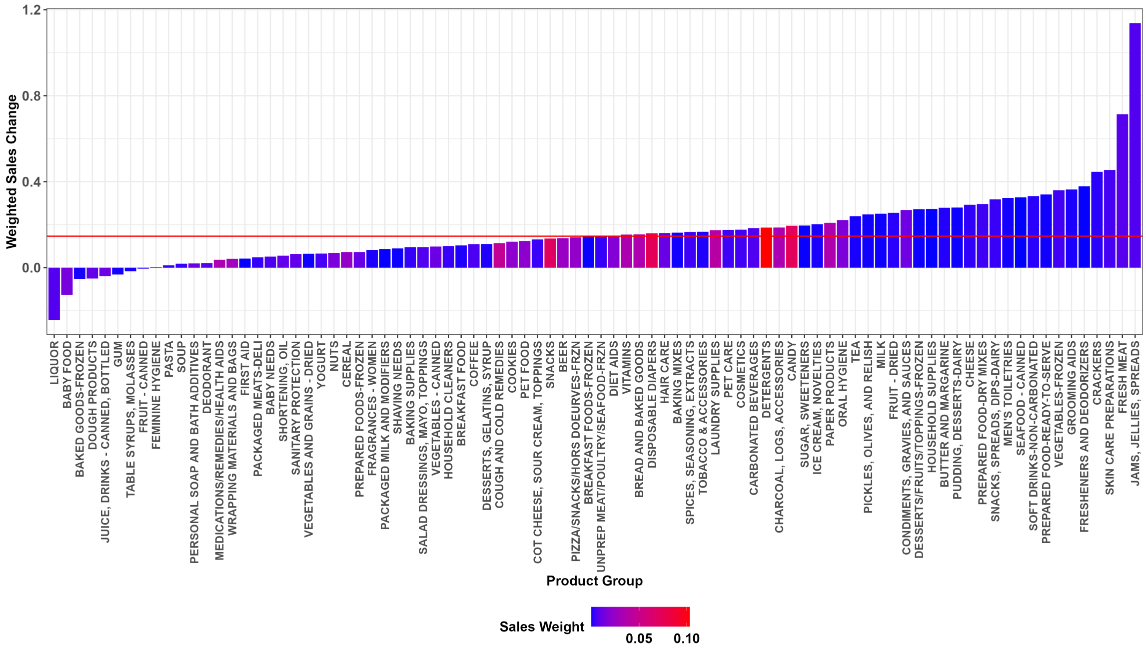

(b) Sales, Product Size Increase

Notes: This graph illustrates the sales changes for product groups for product size decreases and product size increases. Subfigure F.2b investigates product size reductions, and Subfigure F.2b assesses size increases. For both scenarios, the data is aggregated from individual modules to product groups. Weighted average price per volume changes are calculated using sales volume of the products as a weighting factor. Changes in sales are measured on the product module level and considering the year before and after the product size change. The color intensity of each bar in the figure corresponds to the sales weight of the product group, with deeper shades of red indicating higher sales volumes within that group. The red line corresponds to the weighted average across all product modules