Retail Margins and Grocery Inflation

Aljoscha Janssen*

Singapore Management University, School of Economics

ajanssen@smu.edu.sg

19 November 2026

Abstract

I study whether retailers widen margins during inflationary periods. Matching supplier invoices to shelf prices for over 210,000 products in a major U.S. grocery group, I observe each margin through the 2021 to 2023 inflation. The margin rate rose, most on the cheapest products in each category. When a product's cost jumps, its neighbors' prices rise, but only where a category review falls near the shock. The review itself widens margins, raising prices whose costs did not move. Half of the price increases behind the rise came with no move in the product's own cost. Reviews recur a median of twenty weeks apart, and accumulated cost pressure brings one forward. Upward cost shocks spread across the category and downward ones much less. Raises come in round cash amounts and are a larger share of a cheap product's price. Households that traded down moved toward the products whose margins rose most.

JEL Codes: L11, L13, L81, E31, D12
Keywords: Grocery inflation, pass-through, margins, multiproduct pricing, distributional incidence

*I am grateful for comments received during seminars and conference presentations.

I. Introduction

U.S. grocery prices rose by roughly 19% from December 2020 to December 2022 (U.S. Bureau of Labor Statistics, 2024), and whether retailers profited from the rise became a public controversy. A Federal Trade Commission staff report found grocery retailer revenue rising further above cost than at any recent peak (Federal Trade Commission, 2024). Weber and Wasner (2023) argue that firms used supply bottlenecks as an occasion to raise prices and protect margins, and Bernanke and Blanchard (2025) attribute the episode mainly to commodity and sectoral price shocks. Conlon (2026) argues that evidence on profit margins and capital shares cannot by itself distinguish stronger demand from a change in competitive conduct. Each of these accounts treats a grocery price as a response to one product's own cost and demand.

How those prices are set is largely unobserved. A grocery retailer manages thousands of prices in each store, tens of thousands in the largest, and does not set each one in isolation (Bonomo et al., 2023). Prices are managed by category (Zenor, 1994; Basuroy, Mantrala, and Walters, 2001), changing them is costly and procedurally involved (Levy et al., 1997; Zbaracki et al., 2004), and they may be reset in category reviews, occasions on which many prices in a category are reconsidered together. Such a review could raise margins across a whole category rather than pass one cost through one price.

This paper asks how cost shocks change prices and margins in detail. I match the wholesale cost a major U.S. grocery group pays on each supplier invoice to the price it charges for the same product in the same store and month, across 551 stores from 2018 through 2025. The match covers more than 210,000 products, so the margin between price and cost is observed directly rather than inferred from a model of consumer demand (Villas-Boas, 2007). Managers set prices in dense categories, so a single cost shock has many neighboring prices within reach of the same review, and I measure how often those reviews occur and what brings them forward.[1]

I find that the gross margin rate, the wedge between the shelf price and the supplier's invoice as a share of the shelf price, rose from 23% in 2018–2019 to a peak of 28% in mid-2022 and stood more than two percentage points above its pre-pandemic level through 2024. Half of the price increases behind the rise fell in months in which

[1] In U.S. micro data, regular consumer prices change about every seven months (Klenow and Malin, 2010), and wholesale cost increases frequently trigger retail price increases (Anderson, Jaimovich, and Simester, 2015).

the product's own cost had not moved. The increase was not spread evenly across the shelf. Products in the cheapest tier of a category gained 2.6 percentage points more margin than the rest on the prices shoppers paid, and 1.5 points on posted prices. Their invoice costs do not explain the gap. The cheapest decile saw the largest cost growth of any, and its prices still outran its costs by 2.9 log points. At the premium end, prices fell 3.4 log points short of costs. Estimates of cost pass-through range from incomplete (Nakamura and Zerom, 2010; Alexander et al., 2024) to fast and complete for grocery price increases (Richards, Gómez, and Lee, 2014). Under either, a wedge that opens fastest on the cheapest products is not what price-by-price cost recovery would produce.

A cost shock to one product moves prices across the whole category around it. With cost shocks to individual products, isolated in time from other shocks to the same product, the shocked product's regular price stands 6.4% above its level of two months earlier in the month the new invoice cost is recorded. Its neighbors' prices, the products sharing its shelf in the same store (its subcategory), net of their own invoice costs, rise by 1.0% over the same two months and stand 0.7% higher six months on. Measured as slopes on the size of the cost change over six months, the neighbor response is about half the own-product response. A category review is a week in which the retailer reprices much of a category at once, detected from its weekly shelf prices. Neighbors' prices rise 1.8% where a review week falls within two weeks of the shock month and do not change where no review is held, whereas a shared cost or a shift in category demand would move neighbors whether or not a review happened to be due.

Regular prices, the posted shelf prices that hold outside promotions, do not drift between changes, and the changes that reset a category across most of the chain at once come in concentrated episodes. Those episodes recur a median of twenty weeks apart, with a wide spread, and most of the windows in which a review could have been held pass unused. Within a store's category and year, the weekly chance that a review is held rises by 2.2 percentage points on a base of 3.1 between the lower and upper quartiles of accumulated wholesale cost increases. When one is held, it moves the shelf rather than the product whose cost changed. Of the repricings made in review weeks, 59 percent had no cost increase of five percent or more of their own in the months around the change.

Pooled over the sample, an upward cost shock raises the share of neighbors repric-

ing upward and lowers the share repricing downward, while a downward shock does neither reliably. In 2021, the year costs rose fastest, an upward shock raised the share of neighbors repricing upward by 1.4 percentage points, while a downward shock did not raise downward repricing at all. In 2019, 2020 and 2023 the ordering runs the other way. The asymmetry runs upward where cost increases have accumulated and weakens or reverses where they have accumulated less. A sustained run-up therefore leaves more of the price structure raised than the recovery of 2023–2024 takes back.

Where the increase lands follows from the cash amounts in which those repricings are made. Raises made in a review, and raises made outside one, come from the same coarse menu of round amounts, so the same raise is a larger percentage change on a cheap product than on an expensive one. Repeated repricings accumulate that difference into the gap documented above. Comparing small and large packages of the same brand points to the cash amount rather than to demand, since a change in demand common to a brand would move both sizes together. A shift specific to one pack size cannot be excluded.

Switching products within the store gave households little relief. Priced at 2022 prices, the 2019 basket of the fifth of households that bought cheapest cost 24.9 percent more, against 20.2 percent for the fifth that bought least cheap, and of that increase $9.07 per $100 of 2019 spending would not have arisen had every product kept its 2019 invoice-margin rate, against $6.95. Every fifth of households saw the margin on that basket rise by between 4.3 and 5.6 percentage points, the most for the fifth buying cheapest. The products they moved to recovered less than a point of it, and households that traded down toward cheaper items moved toward the products whose margins rose most.

A final exercise reallocates each category's total margin across its price ladder so that each product's margin is the one its measured demand elasticity warrants, holding the category's total fixed. It covers the 319 categories where such a rule can be written down, those with a positive baseline margin and a complete price ladder, where the gap in margin growth between the cheapest and most expensive fifths is 0.94 percentage points against 2.44 across all categories. Under that rule, demand differences warrant only part of the gap. The reallocation closes the rest, 0.53 to 0.75 of the 0.94 points, between 22 and 31 percent of the 2.44-point gap across all categories. The exercise is an accounting decomposition that holds quantities fixed and predicts no equilibrium price.

These findings bear on two literatures. One asks how a multiproduct retailer passes costs through. Theory extends pass-through from the single product (Weyl and Fabinger, 2013) through differentiated-product competition (Griffith, Nesheim, and O'Connell, 2018) to the multiproduct portfolio, where it depends on the whole assortment and can reverse the sign of the price response (Friberg and Romahn, 2018; Armstrong and Vickers, 2023). Empirical work documents the frequency and rigidity of price changes (Nakamura and Steinsson, 2008; Bils and Klenow, 2004; Kehoe and Midrigan, 2015; Eichenbaum, Jaimovich, and Rebelo, 2011) and incomplete pass-through (Nakamura and Zerom, 2010; Goldberg and Hellerstein, 2013; Gopinath and Itskhoki, 2010; Hong and Li, 2017). It documents the response of national firms to cost shocks that arrive in only some of their local markets (Butters, Sacks, and Seo, 2022), markups in differentiated-product retail (Nevo, 2001; Hottman, Redding, and Weinstein, 2016; Stroebel and Vavra, 2019), and the synchronization of prices within a firm under multiproduct menu costs (Midrigan, 2011; Bhattarai and Schoenle, 2014; Alvarez and Lippi, 2014; Bonomo et al., 2023). Closest to this paper is the category management literature. It prices whole categories rather than single brands (Zenor, 1994; Sudhir, 2001; Chintagunta, 2002), finds cross-brand pass-through in scanner data of contested magnitude (Besanko, Dubé, and Gupta, 2005; Dubé and Gupta, 2008), and documents a directional asymmetry in own-cost pass-through (McShane et al., 2016). That literature studies prices; because I observe the invoice behind each price, I can follow the same repricing into *margins*, and show that a product's own pass-through is a poor guide to what a cost shock does to the shelf around it.

The other literature asks who pays for inflation. Cavallo and Kryvtsov (2024) document cheapflation across 91 retailers in 10 countries, separating the part due to temporary discounts from the part due to regular prices, without observing the pricing process behind it. Inflation differs systematically across the income distribution (Jaravel, 2019; Kaplan and Schulhofer-Wohl, 2017; Chen, Levell, and O'Connell, 2024), the cost of living varies spatially with income (Handbury, 2021), and households attribute inflation in part to firms raising prices to widen profits (Stantcheva, 2024), a belief this setting can speak to directly, since it observes what the firm did to each price and each margin. Mongey and Waugh (2025) propose a demand-side channel in which products serving less price-sensitive customer bases have higher markups. The debate over markups in the inflation reads aggregate profit and markup evidence (De Loecker, Eeckhout, and Unger, 2020; Glover, Mustre-del-Río, and von Ende-Becker, 2023; Con-

lon, 2026), macroeconomic decompositions (Bernanke and Blanchard, 2025), granular pre-pandemic markup estimates (Döpper et al., 2025), or aggregate and firm-level evidence together (Weber and Wasner, 2023). The rockets-and-feathers literature documents pervasive asymmetric pass-through (Peltzman, 2000; Borenstein, Cameron, and Gilbert, 1997) and rationalizes it through consumer search (Yang and Ye, 2008; Tappata, 2009) or, within the firm, through menu costs under moderate inflation (Ellingsen et al., 2026); related work studies coordination among sellers (Byrne and de Roos, 2019) and algorithmic pricing (Assad et al., 2024). This paper adds to both literatures the transmission of cost shocks across products. Cheapflation follows from repricing in round cash amounts, with no change in conduct required, and demand differences account for only part of the gap. The asymmetry arises inside a single firm's review calendar and spreads from the shocked product to the shelf around it, needing neither consumer search nor coordination nor algorithms. The redistribution the process produces within categories is large relative to the aggregate margin change it accompanies.

Section II describes the data, and Section III the margin expansion, its concentration at the cheap end of each category, and the three decisions behind a category's prices. Sections IV and V set out and report the designs built on isolated cost shocks, and Section VI reads the review calendar off weekly shelf prices. Section VII accounts for the rise in the margin rate, Section VIII draws out the implications, and Section IX concludes.

II. Data

My primary data come from a grocery scanner dataset (see Dubois, Griffith, and O'Connell, 2022, for a survey of scanner data in economics) covering 551 stores of a major U.S. grocery group, carrying roughly 270,000 distinct products, identified by their Universal Product Code (UPC), over the sample period.[2] For each store–UPC–month, I observe total dollar sales, unit sales, and the unit-value retail price.[3] Weekly

[2]The raw data contain over 900 stores, of which 689 have enough baseline coverage to be classified by their 2018–2019 assortment profile. Clustering on that profile separates 138 small-format express stores from the 551 supermarkets. The express stores are set aside because their thin assortment lacks the within-category density the analysis requires (Supplemental Appendix A.1). The event designs and the aggregate series use the 551 supermarkets; the few exhibits that use all of the group's stores say so in their notes (Supplemental Appendix A.3).

[3]Unit values are the average transaction price within a store–UPC–month and so include promotional discounts. Supplemental Appendix B.2 reports the promotional checks.

regular (posted, non-promotional) shelf prices from the same scanner system reveal the discrete repricing events that monthly unit values average over. Cumulating their persistent week-to-week changes, those the feed does not mark transient and whose absolute log size is at least 0.005, gives a monthly log regular price for each store–UPC. The paper therefore has two price objects. The posted regular price is the outcome of the event studies and the group-level regressions. The unit value enters wherever revenue does, in the margin rate, incidence and demand, and in the synchronization regressions, which count a repricing as a change in it. Supplemental Appendix B.3 sets out the construction and coverage of each, and Supplemental Appendix A.3 states the conventions behind every exhibit.

I also observe wholesale acquisition costs, the invoice-level prices the retailer pays to suppliers. I construct two cost measures: step-costs, the last observed invoice cost carried forward, and a seven-month trailing moving average ($C_{jt}^{wh,MA7}$).[4]

The invoice cost is the per-unit acquisition price from the supplier to the retailer's distribution center, the cost field the category manager observes.[5] It is net only of returned-goods credits. It does not net out off-invoice trade funding (promotional allowances, scanbacks and billbacks, and slotting fees), which the data do not record, nor the chain's own cost of moving goods from its warehouses to its stores. The measured wedge is therefore the invoice-based margin the category manager sees. It nets out neither trade funding nor operating costs (labor, occupancy, shrinkage) and so is not profit; Supplemental Appendix B.1 bounds what the omission can do.

Category assignments come from the retailer's own product master and span three levels: departments, 459 product categories, and 2,632 subcategories, counted over the 551 stores across 2018–2025. The matched panel spans 15 to 17 departments, about 410 categories and 2,200 to 2,300 subcategories in any one period (Table I). Coverage of the matched panel is, if anything, best at the cheap end of each category (Supplemental Appendix A.4).

My main outcome variable is the *gross margin rate*, the invoice-to-shelf wedge. Let j index a product (UPC), s a store, c a product category, and t a calendar month.

[4]The moving average serves the exhibits that compare margins across the price ladder; it spreads each invoice step across seven months and so cannot date a shock (Supplemental Appendix B.1). Each exhibit states which measure it uses.

[5]Supplemental Appendix B.1 checks that the neighbor response of Section V is the same for products invoiced often and rarely.

For each store–product–month, I define:

$$MR_{jst} = \frac{P_{jst} - C^{wh}_{jst}}{P_{jst}}, \tag{1}$$

where P_{jst} is the unit-value retail price and C^{wh}_{jst} the wholesale acquisition cost. I suppress the store subscript s when the context is unambiguous. Products are ranked within their category into price tiers, and because packages of very different sizes share a shelf the ranking normalizes package size. The *price-per-volume* (PPV) of a product is its sales-weighted unit value over 2018–2019 divided by its package size in common units (ounces, count, etc.) recorded in the product master. Tiers are cut on the 2018–2019 ranking of PPV within category, in deciles, in quintiles, or as a bottom-20% indicator ("cheap tier"), and each exhibit states the cut it uses. The exhibits that rank products by tier cover products with a recorded package size in their category's dominant unit, 79 percent of the sales with a matched cost (70 percent in the 2018 baseline year, because sizes come from the current master and items discontinued early lack them). The aggregate series include all products.

A store–product–month enters the matched panel when it records positive sales and carries an invoice cost for the same store and product, observed in that month or carried forward from an earlier one. Invoice costs are set for the chain, so what a store carries is a copy of the product's chain-level invoice. No minimum number of months is imposed, so products enter and leave over the sample, and the 2018–2019 window ranks products within their category rather than selecting them. Log unit values and log step-costs are winsorized once at the 0.5th and 99.5th percentiles, pooled over the panel; the regular shelf price is not winsorized, and the few exhibits that trim a change or a margin rate say so in their notes. The aggregate series that follow describe the matched assortment of each month rather than a fixed baseline cohort.[6] I organize the analysis around three subperiods, *baseline* (2018–2019), *run-up* (2021–2022) and *recovery* (2023–2024), omitting 2020 to avoid pandemic contamination; the monthly series and the event designs run through 2025. Table I summarizes the data.

Figure I shows the matched data for one store and category, salad dressings. Panel (a) follows one UPC's retail price and wholesale cost, and panel (b) the retail prices of the ranked products in the category, grouped by within-category price

[6]Table I reports period-specific sample sizes; Supplemental Appendix Table A1 traces observation counts through each analytical sample.

Table I: Summary statistics: matched retail–wholesale product-level panel.

	Baseline (2018–2019)		Run-Up (2021–2022)		Recovery (2023–2024)	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Prices and Costs						
Retail price ($)	5.37	4.85	6.41	5.49	7.01	5.79
Wholesale cost ($)	4.02	3.94	4.61	4.33	5.08	4.57
Margin rate (%)	23.4	26.1	26.7	24.2	26.3	24.2
Dollar markup ($)	1.35	2.02	1.80	2.56	1.93	2.71
Retail price vs. baseline (%)	–	–	19.4	–	30.5	–
Wholesale cost vs. baseline (%)	–	–	14.6	–	26.3	–
Panel B: Store–month quantities						
Distinct UPCs per store	16,516	3,408	16,882	4,310	18,192	3,100
Panel C: Product Coverage						
# UPCs in the matched panel	136,588		130,381		106,540	
of which private label	6,666 (5%)		6,022 (5%)		5,890 (6%)	
of which price-tier ranked	58,613		47,292		38,190	
PL share of dollar sales (%)	19.8		18.5		18.8	
# Stores	551		551		551	
# Departments	17		15		15	
# Product categories	408		409		401	
# Product subcategories	2,213		2,265		2,271	
Store–UPC–month obs.	∼218M		∼171M		∼176M	

Notes: Observation: store–UPC–month over the baseline (2018–2019), the run-up (2021–2022) and the recovery (2023–2024). Panel A: sales-weighted means and standard deviations across store–UPC–months carrying an observed wholesale cost; retail prices are monthly sales-weighted averages, the wholesale cost the last invoice cost carried forward, the gross margin rate $(P-C^{wh})/P$ in percent and the dollar markup $P - C^{wh}$; the change rows compare each period's mean with the baseline mean. Panel B counts distinct products with a matched cost within a store–month; the count rises because the share of the assortment with a matched invoice cost improves; assortments did not grow. Panel C counts distinct products with an observed wholesale cost in the 551 stores; the ranked row is the subset the tier exhibits can order, products with a pack size in their category's dominant unit, ranked once on 2018–2019 prices on the full costed stream, the monthly observations of all the group's stores that carry an observed invoice cost, aggregated over stores (Supplemental Appendix A.3). Those are 60,730 products, of which 58,613 sell with a matched cost in the 551 supermarkets in the baseline. The 551 stores are fixed once by the baseline classification; 550, 514 and 408 carry matched observations in the three periods (Supplemental Appendix Table A1).

per volume.

Figure I: Matched retail prices and wholesale costs in one store and category.



(a) Single UPC: retail price and wholesale cost

(b) Ranked UPCs, indexed within price bands

Notes: Salad dressings in one supermarket, 2018–2025: the category is carried in 550 of the 551 stores and holds 504 distinct products, and the store drawn is the one with the most matched observations, 432 products and 19,259 product-months. Panel (a): one product observed in all 88 months, its retail unit-value price (blue) and the wholesale cost the pricing process sees (red), the last invoice cost carried forward; the shaded gap is the observed dollar markup. Panel (b): sales-weighted retail price indices (January 2018 = 100) for the 206 of the store's products whose baseline price per volume is measurable in the category's dominant unit, grouped for legibility into four bands from cheapest to premium; the bands are not one of the tier partitions used in the analysis.

Prices outpaced costs. The average retail price rose 19.4% between the baseline and the run-up, against 14.6% for the average wholesale cost (Table I, Panel A).

Product turnover is substantial: 212,184 distinct UPCs carry an observed wholesale cost at some point in the sample, and the number in any one period falls from 136,588 in the baseline to 130,381 in the run-up and 106,540 in the recovery (Table I, Panel C). Among the products the tier exhibits rank the count falls from 58,613 to 38,190.[7]

Household panel. Household-level purchase records from the retailer's loyalty-card program carry no income or demographics and record purchases at this retailer only, so the household analysis compares baskets within households rather than across incomes. Matched to the product-level margin data, they give a balanced panel of 2.43 million households observed in both 2019 and 2022, covering purchases at every store of the group and only products with a matched cost sold in both years.[8]

[7]The recovery-period decline is amplified by the retailer's acquisition, after which a subset of stores converted to the acquirer's banner and exited the panel from August 2024 onward; no acquirer data enter the sample, so the recovery estimates describe the continuing stores.

[8]The panel yields 41.6 million household–category observations across 372 categories; the median household appears in 11 categories with $276 of annual spending at the retailer (Supplemental Appendix I).

III. Descriptive Evidence

Figure II, panel (a), plots the monthly gross margin rate from 2018 through 2025: it averaged 23.4% over the 2018–2019 baseline, peaked at 28.3% in June 2022, and fell back to 25.7% over 2024, still more than two percentage points above pre-pandemic levels (Table I gives the period averages).

Panel (b) draws the same rate by within-category price tier, on the seven-month cost and the wider store set its note describes. From the 2018–2019 baseline to the 2021–2022 run-up, margins in the cheapest quintile expanded by 5.37 pp, against 2.93 pp for the most expensive quintile, a gap of 2.4 percentage points (Supplemental Appendix K). The gap opened gradually over 2020–2022, consistent with accumulation over many repricings rather than a discrete repricing event.

Figure II: Monthly gross margin rate: aggregate and by price tier, 2018–2025.

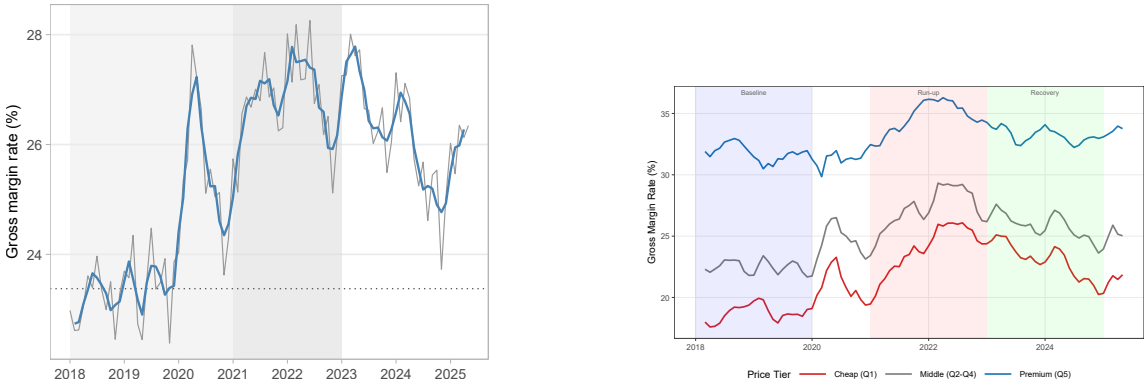

(a) Aggregate

(b) By within-category price tier

Notes: Panel (a): $MR_t = \sum_j w_{jt}(P_{jt} - C^{wh}_{jt})/P_{jt} / \sum_j w_{jt}$, weights dollar sales, C^{wh} the step cost of Section II, on the 551 supermarkets over the months in which a product carries an observed invoice cost. Blue: three-month moving average; grey: monthly values; shaded regions: pre-inflation, run-up and recovery; dotted line: the 2018–2019 average of 23.4%. Panel (b): the same rate on the seven-month moving-average cost, so the two panels do not share one cost object (Supplemental Appendix B.1), drawn on the full costed monthly stream of every store, the basis of Table II. Products are ranked once by 2018–2019 sales-weighted price per volume within category into quintiles, cheap (bottom, red), middle (second through fourth, grey) and premium (top, blue), each month's tier average taken over the ranked products observed that month. Three-month moving averages.

The expansion was broad but not universal. From 2018 to 2024, margin rates rose in about three fifths of the 336 categories with at least $100,000 of cost-matched sales in both years, and the rising categories hold seven tenths of 2018 sales. What matters for incidence is where within a category the wedge opened.

It opened away from the top. Table II ranks products into deciles of their 2018–

2019 price per volume within their category, the same ranking as Figure II, and compares the tiers' sales-weighted margin rates in 2018, 2022 and 2024. The bottom six deciles gained between 3.0 and 4.7 percentage points from 2018 to 2024, while the eighth gained 2.4 and the top two 1.6 points and essentially nothing. The seventh decile is the one interruption in an otherwise clean divide.

Table II: Margin changes by price tier, 2018 to 2024.

Price decile	MR 2018 (%)	MR 2022 (%)	MR 2024 (%)	ΔMR (pp)	$\Delta \ln C^{wh}$	$\Delta \ln P$
1 (Cheapest)	19.8	26.5	23.4	+3.53	+29.5%	+32.4%
2	17.0	24.1	20.3	+3.32	+22.6%	+24.4%
3	16.2	22.5	19.4	+3.15	+29.0%	+30.7%
4	22.9	29.1	25.9	+2.96	+15.5%	+20.5%
5	24.5	30.9	27.6	+3.07	+18.8%	+22.6%
6	25.0	31.5	29.6	+4.67	+9.5%	+15.7%
7	25.0	28.5	24.9	−0.06	+23.3%	+24.8%
8	24.8	30.0	27.2	+2.41	+14.9%	+17.9%
9	31.7	35.8	33.3	+1.60	+22.2%	+26.5%
10 (Premium)	32.8	34.8	32.6	−0.22	+24.4%	+21.0%

Notes: Observation: UPC. Products are ranked once, on 2018–2019 data, by their sales-weighted price per volume within their category, the ranking of Figure II, and cut into deciles (60,730 products; decile 1 has a median baseline price of $0.15 per unit of content, decile 10 of $2.72). The ranking universe is the full costed scanner stream among products with a recorded package size in their category's dominant unit, wider than the matched panel of Table I, which holds 58,613 of the 60,730 in the baseline. The margin rate is $(P - C^{wh})/P$ on the seven-month moving-average cost of Section II, and $\Delta \ln C^{wh}$ is on the same basis. The 2018, 2022 and 2024 columns are means over the year's monthly sales-weighted margin rates among the products observed in that year, so composition moves with turnover; ΔMR is the 2018–2024 change and $\Delta \ln C^{wh}$ and $\Delta \ln P$ the 2018–2024 log changes of the year-level sales-weighted average cost and price.

Costs alone do not explain the gap. They rose somewhat faster toward the bottom of the ladder, but at the two ends prices and costs parted in opposite directions. At the bottom prices rose about three log points ahead of costs, at the top three points behind. The growth columns are log changes of year-level averages over a changing assortment, so they do not line up with the margin change decile by decile. The seventh and ninth deciles show prices outgrowing costs with little margin gain. Because cheap products start from lower baseline margins, the same widening in percentage points is proportionally larger for them.

A first clue to how that happened is that prices within a category move together. When at least one product in a category takes a cost shock, a wholesale cost change of at least two percent, more than half of the *unshocked* products in that category reprice in the same store and month, more than in categories where nothing happened

to costs.[9] Something beyond their own costs moved the neighbors' prices. The answer has to lie in three decisions the retailer takes for a category: when its prices are reconsidered, which products are reconsidered together, and by how much each price moves. A price change records each decision in a different place. Its date records when, the set of products changing with it records which, and its cash amount records how much.

The third decision alone can open a gap of the kind Table II records. Take two products in one category, at \$2 and \$10, with the same margin rate, and let neither cost change. Raise both by 30 cents. The raise is 15 percent of the first price and 3 percent of the second. With costs unchanged, the first product's margin rate rises about four and a half times as much, because the same cents are a larger share of a smaller price.[10] A margin rate rises only when the price's percentage change exceeds the cost's; each raise adds another such increment, and a cut in the same cash amounts undoes what a raise did. Section VIII.A asks whether a downward cost shock moves the shelf around it as readily as an upward one. The raises in the data vary in cash amount across the shelf, and outside review weeks cheap products are also repriced more often. Section VIII.B measures how much of the cheapest fifth's extra price growth comes from the size of its raises and how much from their number. Section IV builds the designs that identify the second decision, and Section VI reads the first off weekly shelf prices, when a review is held and what brings one forward.

IV. Empirical Strategy

The designs must answer whether a cost shock to one product moves the prices of products around it, net of their own costs, and if so, when. A retailer that reprices continuously and optimally would already move a neighbor's price when a product's cost changes, because it owns the substitutes and internalizes the diversion between them. That benchmark delivers a magnitude, since the neighbor's optimal response runs through cross-price demand, and a location, since the response should be concen-

[9] Across the three sample periods, 56 to 58 percent of a category's unshocked products reprice in store–category–months with a shocked product, against 44 to 52 percent in those without (Supplemental Appendix D.2).

[10] Supplemental Appendix C states the formal versions. Proposition 1 gives what a cost shock does to a neighbor's price under joint and under product-by-product pricing, and Proposition 2 gives when a review is held and shows that one resets every price in the category. Proposition 3 gives which direction of shock brings one forward, and Proposition 4 the common-increment result the example illustrates.

trated among close substitutes. The designs put both to the data. Supplemental Appendix C states the pricing problem formally: under joint optimization the neighbor response is zero under linear demand and negative under logit and constant-elasticity demand, and under a product-by-product rule, the benchmark most generous to the demand channel, it is graded by how close a substitute the neighbor is, with none of it common to the shelf.

A dose-response design asks whether *larger* cost shocks produce larger neighbor responses, comparing event groups, each a store, category and month in which some product's cost moved in isolation. A panel event study traces the *timing* of the response product by product, controlling for each neighbor's own wholesale cost. The two handle the obvious threat, that neighbors' costs moved too, in different ways: the first excludes neighbors whose costs changed and the second controls for them directly. Neither removes a store-wide change that coincides with the shock, so the causal reading rests on their agreement together with the review-calendar test of Section VI. A third, cross-sectional design asks where the resulting margin expansion landed across the price ladder (Section VIII.B).

The two designs take different events. The dose-response design counts a cost change of at least two percent, isolated from other changes to the same product by six months either side. The panel event study takes the first isolated change of at least ten percent for a given product in a given store, a stricter bar that buys large, discrete shocks whose timing can be read.

Group-level dose-response. For each event group I construct four estimates, each the average log price change over the six months after the event. They are the shocked product's own pass-through Y_g^{own}; the neighbor spillover Y_g^{nb}, the average change of the unshocked neighbors; the low-price differential $Y_g^{cheapdiff}$ between the lowest-priced fifth of the category, cut on the price level in the first twelve months of the period rather than on the 2018–2019 ranking of Section II, and the rest; and the price-proximity difference $Y_g^{closediff}$ between neighbors priced within twenty percent of the shocked product at baseline and the rest. Each is estimated from a simple regression on the event sample:

$$Y_g^m = a_m + \lambda_{\text{period}(g)} + b_m \, Shock_g + u_{gm}, \tag{2}$$

with standard errors clustered two-way by category and event month. Each observation is one event group g, a store–category–month triplet. When more than one product's cost qualifies in the same cell, the group's shock is the average of their log cost changes, the own-price response averages over the shocked products, and the neighbor set is the products below the two percent threshold. The regression pools 3,509,123 such groups on the 551-supermarket sample, behind which stand 1,314,864 distinct chain-level (product-month) cost shocks. $Shock_g$ is the focal product's wholesale cost change in logs, and b_m is the log-point price change per log point of cost change. Period fixed effects (run-up and recovery) absorb shifts in the level of repricing activity across periods. The slope requires that the size of a group's shock be unrelated to whatever else moved that shelf in that period. Comparing groups within a category, or within a category and period, gives the same slopes (Table III, note).

Panel event study. The panel event study operates at the store×UPC×month level and follows each product's price around the shock month; the full neighbor panel of Figure III, one of the two neighbor samples defined below, contains 96.9 million observations (Supplemental Appendix Table A1). The neighbors are the other products on the focal product's own shelf, those sharing its subcategory in the same store, the finest grouping the product master provides and the most conservative reading of what a review reprices. The product category and a coarser grouping of the product master were estimated as well, under a rule stated in Supplemental Appendix F.3 that admits a construction only if its leads are flat and its response peaks in the month of the shock or one of the two following it. The full neighbor panel's third-month response lies between 0.64 and 0.81 percent across every construction, while the cost-stable panel's runs from 0.06 to 0.86 (Supplemental Appendix F.3).

I define a *panel event* as the first temporally isolated *persistent* wholesale cost increase of at least 10% for a given store–UPC pair.[11] Three filters ensure clean identification. *Persistence*: over the following six months the cost must be observed and must neither decline below its post-shock level nor rise more than one further percentage point, so products leaving the cost data do not become events. *Isolation*: no other shock within ±6 months. *First event*: each store–UPC pair contributes one event. Isolation is a property of the shocked product: 41 percent of balanced events, those whose thirteen-month window lies inside the sample period, have another first

[11]The synchronization analysis, which measures the share of a category's unshocked products that reprice in a given store and period, uses a lower 2% threshold.

event in the same store–subcategory inside the window. On store–subcategories with a single event the response of cost-stable neighbors, a sample defined below, is 0.53 (0.08) percent in the shock month (standard errors in parentheses throughout) and 0.41 (0.11) at the third month, against 0.88 and 0.86 on all events (Supplemental Appendix Table F3). Shelves on which several invoices move are repriced more, as a review that fires on accumulated cost pressure implies and as a cost shock shared across the category would also produce. The own-product specification is:

$$\ln R_{jst} = \alpha_{js} + \gamma_t + \sum_{\tau \neq -2} \beta_\tau^{own} \mathbf{1}[\tau_{jst} = \tau] + \varepsilon_{jst}, \tag{3}$$

where $\ln R_{jst}$ is the log regular shelf price for UPC j in store s at month t, the posted price of Section II rather than the unit value P_{jst} of the margin measure. Store×UPC fixed effects α_{js} absorb permanent level differences, year-month fixed effects γ_t absorb ambient inflation, and τ_{jst} is event time relative to the shock month, binned at ± 7 (months seven or more from the shock pooled into two end bins). The reference period is $\tau = -2$.

The neighbor specification adds the neighbor's own wholesale cost:

$$\ln R_{jst} = \alpha_{js} + \gamma_t + \sum_{\tau \neq -2} \beta_\tau^{nb} \mathbf{1}[\tau_{jst} = \tau] + \delta \ln C_{jst}^{wh} + \varepsilon_{jst}, \tag{4}$$

so that $\hat{\beta}_\tau^{nb}$ is the neighbor's price change at each event time net of its own cost movements. The control assumes that what the invoice misses, trade funding above all, does not arrive with a lag that lands in event time. In 82 percent of consecutive month pairs the invoice cost is unchanged, so the control is active only where the invoice moves at all.

Identifying assumption. The identifying assumption for equation (4) is that, conditional on the neighbor's own wholesale cost, the store×UPC effect and the year-month effect, the timing of the focal shock carries no information about what the neighbor's price would otherwise have done: $\mathbb{E}[\varepsilon_{jst} \mid \tau_{jst}, \ln C_{jst}^{wh}, \alpha_{js}, \gamma_t] = 0$. A violation requires manufacturers to time idiosyncratic invoice increases on one product to coincide with unobserved retailer-side determinants of other products' prices, net of those products' own invoice costs, and to do so only where the category is being repriced anyway and asymmetrically for increases versus decreases. Event time zero is the month in

which the new invoice cost first appears in the retailer's cost record. In the month before the record the shocked product's shelf price already stands 3.2 percent above its level of the month before that, and the neighbor panels about 0.5 percent, so every coefficient is read against the month two before the record, and the month before it counts as part of the response. The pricing response of products whose own invoices took no such step identifies the spillover.

I estimate this on two samples. The *full neighbor sample* includes every product on the shelf that never has a qualifying event of its own anywhere in the panel, whatever smaller or non-persistent cost changes it takes, and controls for each neighbor's own cost. The *cost-stable neighbor sample* restricts to products whose recorded invoice cost never moves by ten percent or more in any month of the 2018 to 2025 panel, so a contemporaneous invoice shock of their own cannot drive their response. The rule conditions on a product's whole cost history, and a change in trade funding that never reaches the invoice is excluded by neither sample. Each neighbor UPC is assigned to one event, the first balanced event in its store–subcategory, so no product enters twice and no event is stacked.

Supplemental Appendix F.6 counts what each filter keeps, from 16.9 million candidate cost increases to the 1.64 million first events of the panel and the 24,361 distinct focal products at balanced events.

The design's main validation is that neighbor prices show no pre-trend. Before the shock their coefficients are small, unsystematic in sign, and do not build toward the event over the four tested leads, in both neighbor samples, though both rise in the month before the invoice moves.[12]

Threats to identification. Two threats matter; the review calendar answers the first and the dose-response the second. If neighbors were simply passing through cost increases of their own, the response would not depend on whether the retailer happened to be repricing the category. It does, and among neighbors whose invoice cost never moves by ten percent in eight years it is if anything larger. If instead the events were picking up category demand shifts, the size of one product's invoice change would not predict how far its neighbors move. What limits that reading is the dose-response of

[12]Supplemental Appendix Table F1 reports every pre-period coefficient with its standard error for all three panels, and Supplemental Appendix F.1 reports a placebo test and a worst-case bias bound in the spirit of Rambachan and Roth (2023). With samples running to tens of millions of observations, formal tests reject equality with zero for deviations far smaller than the effect, so the relevant question is magnitude rather than significance.

Table III, and future shock size does not predict past price changes (Supplemental Appendix F.1).

Timing of the shocks. Wholesale prices are negotiated and the review calendar of Section VI is predictable, so a supplier could time an increase to a scheduled review, in which case the spillover would be the review's echo. Three facts argue against that reading (Supplemental Appendix F.4). Shocks and reviews are linked in time, but loosely. A shock month contains a detected review week, a week in which much of a category reprices at once, 28 percent of the time, against 18 percent when the calendar is shifted by a quarter. For 55 percent of shocks no review week falls within two weeks of their month at all. The nearest review week precedes the shock month as often as it follows it, 51 percent against 49, which is not the one-sided alignment that placing an increase just ahead of a known review date would produce. And the calendar is not a fixed appointment. Among the intervals between reviews still open eight weeks after the last review, a cost increase of at least ten percent on any product of the category in weeks three to eight raises the probability that a review has been held by week fifteen from 18 to 27 percent. Once accumulated cost pressure, defined in Section VI, enters the regression for the weekly chance that a review is held, the indicator for a single shock turns negative, but that shock is itself one of the increases the pressure measure counts. The timing of the next review lines up with the stock of increases since the last one rather than with any one of them. A review that would have fired regardless also cannot explain why the size of the focal shock predicts how far the neighbors move.

Inference. Invoice costs are set for the chain, so the store copies of one product's cost shock are not independent experiments. The 1.64 million event store–products of the panel correspond to 199,997 distinct product-months, and the 405,137 store–subcategories with a balanced event to 24,967 subcategory-neighborhood cohorts, one per subcategory and event month. Supplemental Appendix F.2 recomputes the standard errors of every headline coefficient with clusters at the store-neighborhood, the neighborhood being the subcategory cell of the design; at the neighborhood-cohort and at the neighborhood, each alone and two-way with the calendar month; and at the product two-way with the calendar month. The own-product response is unaffected. The neighbor response keeps its point estimate at every level and loses precision as the clusters coarsen, but not its significance. The third-month spillover of 0.71 percent has a standard error of 0.02 percentage points under store-by-neighborhood cluster-

ing, 0.18 at the level of the subcategory neighborhood and event month, the scheme Figure III uses, and 0.21 at the coarsest, the subcategory neighborhood itself.[13]

The dose-response design fares better, because it uses the size of 1.31 million distinct chain-level shocks rather than the timing of each store copy. With standard errors clustered by category and event month, its own pass-through of 0.045 has a standard error of 0.006 and its neighbor spillover of 0.023 a standard error of 0.004, so Table III reports those. The text reads the panel figure as evidence on timing and takes the magnitude from the designs that pool the store copies or use clean controls (Supplemental Appendix F.5).

V. Results

In the short run the retailer absorbs most of any individual cost shock. Only 4.5 percent of it passes into the shelf price over the six months that follow, 0.45 percent for a 10 percent cost change (Table III). Unshocked neighbors reprice too, by 2.3 percent of the shock, about half of the own-product slope, even though their own wholesale costs did not move. In 96.5 percent of the neighbor-months the cost change is exactly zero.[14] That is what pricing by category implies.

How that spillover is distributed across the shelf is as telling as its size. Were the retailer reoptimizing against demand, the response would concentrate on close substitutes. Products priced far from the shocked item instead reprice nearly as much as those priced within twenty percent of it. The extra response among the price-proximate is an eighth of the common one (Table III), and Section VIII.D makes the comparison on measured cross-price elasticities.[15] The spillover is common to the category, which is what a review of the whole shelf produces. Within a single event the extra response of the lowest-priced fifth is small and at best marginal (Table III); the tier gap of Table II accumulates across repeated repricings (Section VIII.B).

[13]Supplemental Appendix Table F2 reports the pre-period tests under all six schemes. The fine store-by-neighborhood clustering rejects in all three panels, as tests on tens of millions of observations will, and under every coarser scheme the neighbor panels do not reject. The own panel rejects under every scheme but one, since its leads three to five months before the shock sit 0.31 to 0.64 percentage points below the reference month (Supplemental Appendix Table F1).

[14]Both are cumulative six-month responses per unit of log cost shock. The event-study coefficient of +6.4% reported below is an instantaneous jump in a balanced panel that absorbs ambient inflation, so the two are not the same quantity. Supplemental Appendix H.1 reports the three-month slopes, 0.035 and 0.014.

[15]Excluding promotional observations leaves the coefficients almost unchanged (Supplemental Appendix Table B1), and the six-month persistence rules out temporary promotions.

Table III: Event-level regressions: own pass-through, neighbor spillover, and the low-price and price-proximity differences.

Estimate	Coeff.	Std. Err.	N groups	95% CI
Y^{own} (Own pass-through)	0.045***	0.006	3,509,123	[0.034, 0.057]
Y^{nb} (Neighbor spillover)	0.023***	0.004	3,459,578	[0.015, 0.030]
$Y^{cheapdiff}$ (Lowest-priced fifth)	0.001	0.001	2,903,443	[−0.001, 0.003]
$Y^{closediff}$ (Price proximity)	0.003***	0.001	2,594,639	[0.002, 0.004]

Notes: Regression: $Y_g^m = a_m + \lambda_{\text{period}(g)} + b_m\, Shock_g + u_{gm}$ (equation 2); coefficients are log-point price changes per log point of cost change, averaged over the six months after the event month, so 0.045 is a price response of 0.45 percent to a cost change of 10 percent. *Observation:* event group $g = (\text{store}, \text{category}, \text{month})$. *RHS:* $Shock_g$ is the average log change in the step cost of the cell's qualifying products, those whose invoice moved by at least two percent in either direction and were isolated by six months; a fifth of the groups have a net decrease. *Dependent variables:* Y^{own} is the sales-weighted change in the shocked products' log regular price, the mean over the six months after the event month less the mean over the six before it, and Y^{nb} the same object for the group's unshocked neighbors, weighted by pre-event sales. $Y^{cheapdiff}$ is the difference in neighbor response between the lowest-priced fifth of the category, cut on the unit-value level over the first twelve months of the period, and the rest, and $Y^{closediff}$ the difference between price-proximate neighbors (baseline price within twenty percent of the shocked product's) and the rest. A group enters the neighbor row only with an unshocked neighbor whose six-month change is computed, which 49,545 groups lack, and the contrast rows only with at least three such neighbors on each side. Events span 2019–2025 on the 551-supermarket sample, with period fixed effects for the run-up and the recovery. Standard errors are clustered two-way by category and event month. Comparing groups within a category gives the same slopes, category fixed effects moving the own and neighbor slopes to 0.043 and 0.022, category-by-period effects to 0.041 and 0.020 and store-by-category effects to 0.043 and 0.021. Calendar-month effects, which absorb whatever moved every group's prices in a month, cut the own slope to 0.033 (0.004) and the neighbor slope to 0.008 (0.002); the design's variation is across months by construction. $^{***}p < 0.01$, $^{**}p < 0.05$.

Panel event study. Figure III follows prices month by month around a persistent cost shock, and the timing points to where the spillover comes from.

The shocked product responds immediately, standing 6.4% above its pre-shock level in the month the cost lands and 4.7% six months later.[16] Its neighbors respond at the same time and by far less. In the full neighbor sample, which controls for each neighbor's own cost, prices rise 0.96% in the shock month, about a seventh of the own-product jump, and stand 0.71% and 0.69% above their pre-shock level three and six months later. Among neighbors that never take a cost shock of their own the profile is the same: 0.88% at impact, 0.86% at three months and 0.85% at six. In neither neighbor panel do the four tested leads, $\tau = -6$ to -3, build toward the shock, the largest being 0.13 percentage points in the full sample and 0.05 in the cost-stable

[16]The slow post-impact decline reflects the year-month effects absorbing ambient inflation rather than costs mean-reverting.

one.

The contrast identifies the process, because prices change through two channels. More than half of persistent regular-price changes, 56 percent, fall outside the review weeks that Section VI identifies, which hold the other 44 percent. A product whose own invoice jumps ten percent is repriced individually without waiting for the calendar. A neighbor has no invoice change of its own to answer, so its repricing comes only with one of those episodes.

The argument rests on the calendar. A confound, whether an unmeasured common cost or a shift in category demand, moves prices when it arrives. To produce what the data show it would have to move neighbor prices by 1.8 percent where a repricing episode falls near the shock and by nothing where none is detected over the following six months, and leave the shocked product's own response present in both (Section VI). Only a confound implemented at those episodes would reproduce that pattern, and the cost and review-timing checks narrow that reading.

The response also depends on the period. On a 184-store subsample, two of the six blocks into which the data are divided, events dated 2020–2022 carry a third-month neighbor response of 1.40 percent against 0.38 for events dated 2018–2019 and 0.89 for 2023–2024 (Supplemental Appendix F.5).

The size and the persistence of the step follow from the same process: a review moves a median 18 percent of the category at once (Section VI), so the coefficient in panel (b) of Figure III averages repriced and unrepriced products, and the level then holds because regular prices do not drift between episodes.

The spillover reaches well beyond the products whose costs moved: the shocked product moves more than any single neighbor, but the median event group has 46 unshocked neighbors.[17]

Four qualifications. First, the invoice month is not the onset. At $\tau = -1$ the shocked product's price stands 3.23 percent above the reference month, against 6.35 at impact, and the neighbor panels 0.49 and 0.45 against 0.96 and 0.88 (Supplemental Appendix Table F1). A price that moves before the invoice records the change is what one would expect where the increase was contracted in advance, though nothing in these

[17]The count is the median over event groups of the unshocked neighbors that enter the neighbor outcome (mean 73, interquartile range 20 to 97); 84 and 75 percent of those groups supply the three neighbors on each side that the two contrast rows of Table III require. The average group has about five shocked products rather than one, so I attach no share of the category's total price movement to either channel.

Figure III: Panel event study: own vs. neighbor price response to persistent cost shocks.

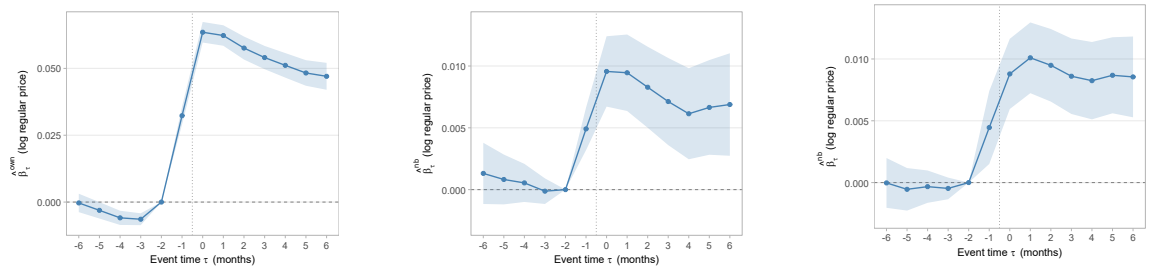

(a) Own product

(b) Neighbor products: full sample

(c) Neighbor products: cost-stable sample

Notes: Coefficients $\hat{\beta}_\tau$ from two-way fixed-effects regressions at the store×UPC×month level. The outcome is the log regular shelf price, the posted price that holds outside promotions, read from the retailer's weekly price feed; Supplemental Appendix B.3 defines it and reports the same estimates on the monthly unit value. Events are the first persistent, isolated wholesale cost increase of at least 10% for a store–UPC pair (Section IV). Of the 1,640,564 such events in the 551 supermarkets over 2018–2025, 405,137 are the first balanced event in their store–subcategory, 259,319 enter the full neighbor panel and 148,556 the cost-stable panel. Panel (a) estimates equation (3) on the shocked products (45.0 million observations); panel (b) equation (4) on the *full neighbor sample* (96.9 million), controlling for each neighbor's own log step cost; panel (c) the same without the cost control on the *cost-stable neighbor sample* (28.7 million). The largest of the four tested leads is 0.13 percentage points in panel (b) and 0.05 in panel (c), and neither panel's leads build toward the shock (Supplemental Appendix Table F1). Event time zero is the month the new invoice cost is first recorded; panels (b) and (c) share a vertical axis. Store×UPC and year-month fixed effects; reference period $\tau = -2$; standard errors clustered on the neighborhood by event-month cell, the cohort that pools the store copies of one shock (Section IV) (24,967 cells in panel (b)); shaded bands are 95% confidence intervals. Supplemental Appendix F.2 reports five further clustering schemes, Supplemental Appendix F.3 alternative neighborhoods and Supplemental Appendix Table F3 alternative cost-stable definitions.

data dates the contract. Second, the third-month neighbor response depends on the comparison group. On the same 184 stores, estimators that compare shocked shelves only with never-shocked ones (Gardner, 2022; Sun and Abraham, 2021; Cengiz et al., 2019) give 0.38 (0.21) to 1.54 (0.46) percent in the full sample against 0.72 under the panel design on the same stores, and 0.64 to 1.70 among cost-stable neighbors against 0.89. On a further eight-store subset, on which two of the estimators run, the full-sample range is 0.08 (0.27) to 0.79 (0.17) (Supplemental Appendix F.5). Third, the cost-stable rule, no invoice move of ten percent or more in any month, selects on eight years of costs. Definitions that screen only the months around the event fail their pre-period tests, and controlling for the neighbor's own cost halves the cost-stable panel's third-month response from 0.86 to 0.40 percent (Supplemental Appendix Table F3). Fourth, the review trigger of Section VI, that accumulated

cost pressure brings a review forward, is identified within the store–category. Defined at the chain by category and week, where the invoice actually moves, the stock of accumulated cost increases has a coefficient of 0.026 with a standard error of 0.014 on 977 chain-level events, and a calendar shifted forward by thirteen weeks returns a larger coefficient still (Supplemental Appendix F.4).

VI. Category Reviews

The results point at a process that reprices a whole category at once. The neighbors move with the shocked product net of their own invoices, and by roughly the same amount whichever neighbor one looks at, which is what a decision taken for the category rather than for the product implies. The repricing episodes are directly observable in the weekly data, so the process can be read off them rather than inferred from timing. Monthly averages turn a series of discrete steps into a gentle drift; weekly regular shelf prices make the steps visible. I call a store–category–week a *review week* when far more of the category's products change regular price together than that cell's own baseline repricing rate allows.[18]

Review weeks are 3.2 percent of the store–category weeks the detection can score, yet they contain 44 percent of the persistent regular-price changes made in the weeks the detection can score. The median review reprices 18 percent of the category's products across three distinct brands. Reviews are chain-level category events. When a category is reviewed, one calendar week holds the reviews of at least half of the stores reviewing it in 86 percent of category-years, against 7 percent when the same repricing is scattered at random. Reviews recur without keeping to a fixed clock. Adjacent review weeks are collapsed into one review, an episode that can span several weeks. The median gap between consecutive reviews of a store–category, from the end of one episode to the start of the next, is 20 weeks; nearly a fifth of gaps are shorter than nine weeks and a sixth longer than a year. Gaps of eleven to seventeen weeks are about forty-five percent more common than a permutation of each category's episode dates produces, and the rhythm runs from the last review rather than from the calendar,

[18]The rule is an exact-binomial test of a store–category–week's repricing share against the cell's own rate over the other weeks of its calendar quarter. The week is flagged when the tail probability is below one in ten thousand and the share is at least three times that baseline and five percentage points above it. Cells with fewer than twenty actively priced products, or fewer than fifty active product-weeks in the rest of the quarter, are never flagged and are excluded from the shares that follow. Moving the tail probability by an order of magnitude either way changes the share of weeks flagged only from 3.2 to 3.8 and 2.7 percent (Supplemental Appendix D.1).

since episode starts are no more common in the weeks that turn a calendar quarter than in any other week (Supplemental Appendix D.1).

Define a *slot* as the window 12 to 15 weeks after the category's last review, repeated every thirteen weeks until the next review, a grid the analysis imposes rather than one the retailer is observed to keep. Define *cost pressure* as the share of the category's cost-observed products that have received a wholesale cost increase of at least five percent since that review. The chance that a slot is used rises with pressure; I say a slot fires when a review starts in it. Figure IV shows the two objects in three store–category histories. In breakfast cereal, reviewed often, pressure of 34 percent or more fires the next slot every time it is reached, and the one slot reached at 19 percent passes. In cocoa and milk modifiers, reviewed half as often, the slots reached at 28 to 55 percent fire and those reached at 11 to 18 percent pass. In oral care nothing fires for a year while pressure creeps from 4 to 19 percent, and the fifth slot fires once pressure reaches 41 percent. Across all store–categories, roughly six in seven slots pass with no review (Figure V, panel (a)).

The probability that a slot fires rises from just under 10 percent when little has happened to costs to 19 percent at high pressure (Figure V), so across the range of pressure the chance that a scheduled opportunity is taken roughly doubles. That comparison pools slots across years while pressure is itself concentrated in the run-up; the controlled version holds the store–category and the year fixed at a given number of weeks since the last review. There the weekly chance of a new review rises by 2.2 percentage points on a base of 3.1 between the lower and upper quartiles of pressure (Supplemental Appendix F.4).

A review moves the shelf rather than the individual product whose cost changed. Of the 7.0 million regular-price changes in detected review weeks across all of the group's stores, on products whose invoice cost is observed around the change, 59 percent involved no wholesale cost increase of five percent or more for the product being repriced in the month of the change or the month on either side.

How far a price then moves depends much less on the size of the product's own cost increase. I match each product raised in a big review week, a detected review week in which at least a quarter of the store–category's products change regular price, to the increase in its own wholesale cost per unit over the review month and the two before it, and compare products of the same kind at the same price. Splitting them by whether that cost increase was under 50 cents or 50 cents or more, a median cost

increase of 12 cents in the first group and one of 82 cents in the second are followed by the same median raise of 30 cents, and by average raises of 33.8 and 36.5 cents (Figure V, panel (b)). Close to a sevenfold difference in the median cost increase moves the average raise by under three cents.

The two groups draw their raises from the same menu of amounts: close to nine in ten land on a multiple of ten cents in either group, and four amounts account for roughly two thirds of them. The response is coarse rather than absent, since a raise above a dollar is 1.7 times as likely under the heavy cost increases. The firing probability responds strongly to category-wide pressure and the raise barely to the product's own cost increase. The weight of pass-through therefore lies on how often a category is reviewed rather than on the size of the raise a review makes.

The spillover of Section V also appears only where the calendar places a detected review. A shock is adjacent to a review when a detected review week of its own store–category falls within two weeks of the shock month, and without one when the category has no detected review week from six weeks before the shock month to twenty-seven weeks after its first day. Estimated separately on the two groups, which do not exhaust the shocks, the neighbor response differs sharply (Figure VI). Adjacent to a review, neighbor prices are flat through the pre-period and step up with the shock, standing 1.8 log points above their pre-shock level in the shock month and 1.4 six months later. Where no review is detected, the impact coefficient is 0.01 log points, the share repriced in the shock month is 0.070 against 0.295 in the adjacent group, and the level then drifts down, because the year-month effects follow shelves that are inflating while these stand still. The shocked product is repriced in both groups, by 7.0 and 5.5 percent, so the spillover varies with the review and the pass-through much less (Supplemental Appendix F.4).

The split uses the calendar the shocks realized, so whether it is drawn from the outcome it explains is worth checking. I split the same shocks by a quantity fixed before them, the number of weeks since the category's last review. The share that turn out to be review-adjacent rises from 0.33 where the last review fell three to five weeks earlier to 0.50 at twelve to fourteen weeks, and falls to 0.29 beyond twenty. The neighbor response follows it, 1.64 log points (0.81) in the shock month where the last review fell nine to fourteen weeks earlier against 0.62 (0.34) where it fell three to eight weeks earlier. Three of the six bins fail their pre-period test (Supplemental Appendix F.4).

This is the process the designs above were reading: prices are chosen for a category, at a moment, on a calendar whose firing rate rises with accumulated costs. A cost shock changes the shocked product's own optimal price and brings forward the reconsideration of the prices around it; the neighbor response measures the second.

Figure IV: Cost pressure, review slots, and review firing: three store–category histories.

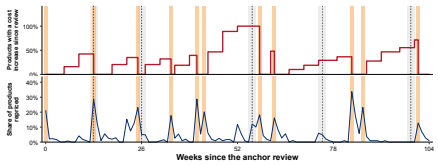

(a) Breakfast cereal, one store: ten reviews in two years

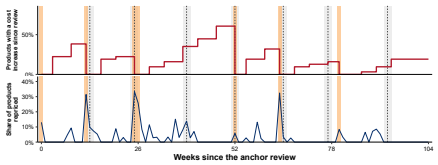

(b) Cocoa and milk modifiers, one store: five reviews in two years

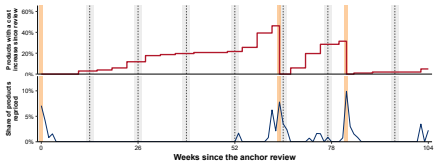

(c) Oral care, one store: a year of slots passing, then a review

Notes: Each panel shows one store–category over two years of event time, weeks since an anchor review. Orange bands: detected review weeks. Dashed vertical lines: slots, every thirteen weeks after the last review, the clock restarting after every review; the slot windows of the text run from the week before each line to two weeks after it. Top track (red staircase): cost pressure, the cumulative number of the category's products with a wholesale cost increase of at least 5% since the last review, divided by the store–category's average number of cost-observed products, reset to zero when a review fires. Bottom track: the share of the category's products changing regular price each week; humps without an orange band do not clear the detection rule. The histories were chosen by rule from the 341,575 two-year windows of the 551 supermarkets that begin at a review, tobacco and pet excluded. Panels (a) and (b) come from the 553 windows in which every slot above a high-pressure cutoff (20 or 30 percent) fires and none below a low cutoff (5 or 10 percent) does. Panel (c) comes from the 361 windows whose first review comes at a slot between weeks 26 and 70 after two or more slots have passed below 20 percent. Observation: store×category×week.

VII. Accounting for the Margin Increase

The decomposition. The margin rate rose from 23.4 percent in the baseline to 26.7 percent in the run-up, a rise of 3.3 points (Table I). I split that rise in three steps, product by product within each store. First, I set aside a composition term, the products that entered or left a store's assortment and the shift of sales among the continuing products, and keep the products sold in the same store in both periods. Second, for each continuing product, weighted by its baseline sales, I split the change in its margin rate exactly into a price effect, the price changes alone with cost held

Figure V: Cost pressure, the probability a review fires, and the size of the raise.

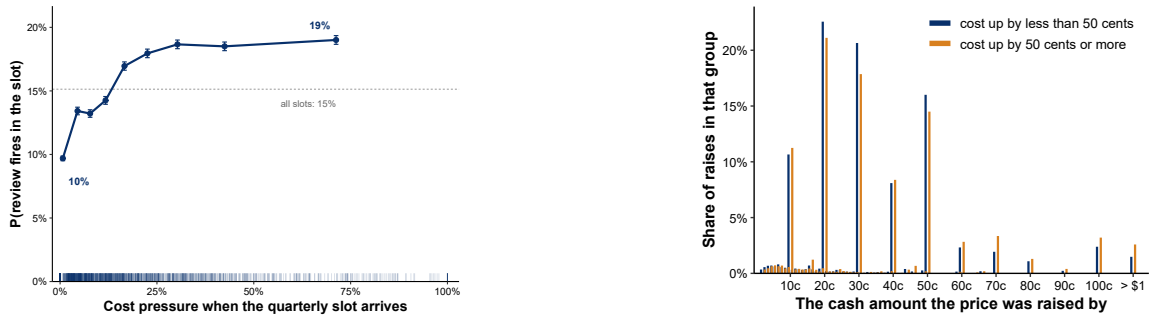

(a) Probability the slot fires

(b) Size of the raise

Notes: Panel (a). Horizontal axis: cost pressure when the slot arrives, the share of the category's cost-observed products with a wholesale cost increase of at least 5% accumulated since the last review, in deciles, the two lowest pooled by ties at zero. Vertical axis: the probability that a new review starts in the first slot after the last review, weeks 12–15 since it. Observation: store×category intervals between reviews that reach their first slot, one per interval, on the 914 stores whose weekly price feed has usable cost coverage (Supplemental Appendix A.3), January 2018 to May 2025 (N = 503,201); the dashed line is the unconditional rate of 15%, and the within-store–category and within-year slope is in Supplemental Appendix F.4. Panel (b). The cash amount by which a regular price was raised in a big review week, for products whose wholesale cost per unit had risen by less than 50 cents (navy) and by 50 cents or more (orange). The cost increase is the largest month-on-month rise in the product's own invoice cost per unit over the review month and the two before it. Prices are held to \$2 to \$6 and the two groups matched cell by cell on category and 25-cent price point, the orange group reweighted to the navy group's composition, and the final bar collects every raise above a dollar, 1.7 times as likely in the orange group. The median cost increase is 12 cents in the navy group against 82 in the orange, at a median price of \$3.79 in both; the median raise is 30 cents in both and the averages 33.8 and 36.5 cents; close to nine in ten raises (87% and 86%) land on a multiple of ten cents, and the four commonest amounts, 10, 20, 30 and 50 cents, account for 70% and 65% of them. Observation: a product raised in a big review week, 1 January 2018 to 22 April 2024, in the 550 supermarkets with invoice cost levels, matched to a rise in its own invoice cost (54 percent of the 3,865,560 raises of Figure VIII). Prices lie between \$2 and \$6, in one of the 1,503 category by price-point cells that hold both groups (N = 1,239,222).

at the baseline, and a cost effect, the cost changes alone with price held there. Third, I sort each product's monthly changes by whether its own invoice cost moved by at least two percent in that month, the threshold of the dose-response design, and split both effects by that rule. Made month by month, the split leaves a term of 0.5 points that belongs to neither side. Table IV reports the parts.[19]

Prices, costs and composition. On the continuing products, prices alone would have raised the rate by 6.1 points and costs alone would have lowered it by 5.0 (Table IV,

[19]Supplemental Appendix H.2 states the identity and the checks behind it, and Supplemental Appendix Table H1 gives every part to three decimals, including the alternative look-back windows and cost cutoffs reported below; the split by the week of the change is in the same section.

Figure VI: The neighbor response comes with the review, and only with it.

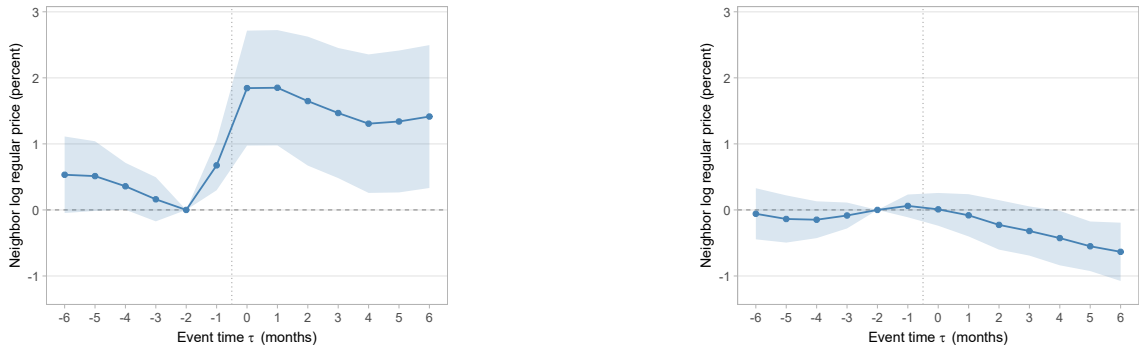

(a) Review week adjacent to the shock

(b) No review week for six months

Notes: The full neighbor panel of Figure III, re-estimated separately on two groups of shocks, both restricted to events whose six-month horizon lies inside the clean weekly review calendar. Panel (a) keeps shocks for which a detected review week of the shocked product's own store–category falls within fourteen days of the shock month: 107,585 balanced store–subcategory events and 46.0 million neighbor product-months. Panel (b) keeps shocks whose store–category has no detected review week from six weeks before the shock month to twenty-seven weeks after its first day: 61,783 events and 17.1 million product-months. In panel (a) neighbors stand 1.84 log points above their pre-shock level in the shock month (s.e. 0.44), 1.47 at three months (0.50) and 1.41 at six (0.55). In panel (b) the impact coefficient is 0.01 (0.13) and the path drifts to −0.63 by the sixth month, since no detected episode reaches these shelves while the year-month effects follow a shelf that is inflating. The pre-period coefficients are jointly indistinguishable from zero in both panels, and no tested lead exceeds 0.53 log points in panel (a) or 0.15 in panel (b). The share of neighbors whose regular price changes in the shock month is 0.295 in panel (a) against 0.070 in panel (b), on pre-period baselines of about 0.10 and 0.08. The two groups are not a partition. Specification, reference period and event definition as in Figure III; standard errors clustered by category and calendar month; the two panels share a vertical axis; shaded bands are 95% confidence intervals.

right-hand column). With the 0.5-point term the three give 1.6 points on those products, and composition gives the other 1.6, of which the shift of sales among continuing products added 2.3 points, products entering a store's assortment 4.9, and products leaving it removed 5.6. The price effect exceeds the whole 3.3-point rise because cost increases took back 5.0 of its 6.1 points.

Sorting the price effect by the product's own cost. The first two columns sort the same parts by the two percent rule. Months in which the product's own cost had not moved contributed 3.3 of the 6.1 points of price effect, half of the price increases behind the rise, and months in which it had moved the other 2.9. The cost effect falls almost entirely in the second group, 4.9 of the 5.0 points. Of the 3.3 points of price effect in months without an own cost move, 3.1 therefore remained as margin, more than the 1.6 points continuing products gained in all; the months with one lowered

Table IV: Accounting for the rise in the margin rate, 2018–2019 to 2021–2022.

	Months in which the product's own cost		
	did not move	moved	All months
Price changes, costs fixed (the price effect)	3.3	2.9	6.1
Cost changes, prices fixed (the cost effect)	−0.2	−4.9	−5.0
Term from the month by month split			0.5
Continuing products, price and cost effects combined	3.1	−2.0	1.6
Composition (reallocation 2.3, entry 4.9, exit −5.6)			1.6
Change in the margin rate, 23.4 to 26.7 percent			3.3
Share of the price effect	0.53	0.47	1.00

Notes: Observation and sample: the store–UPC–months of Table I in the 551 supermarkets, 2018–2019 against 2021–2022; continuing products are store–UPCs with matched-cost sales in both periods, entering at baseline sales weights. Entries are percentage points of the sales-weighted margin rate, $100\,(P - C^{wh})/P$ on the monthly unit value and step cost, and the last row shares of the price effect. Within a continuing product the split into a price effect and a cost effect is exact by the logarithmic-mean identity of Supplemental Appendix H.2. Made month by month against the product's baseline average price and cost it leaves a further term of 0.5 points, which belongs to neither column, so the combined row reads $3.1 - 2.0 + 0.5 = 1.6$ and the last column $6.1 - 5.0 + 0.5 = 1.6$. A month counts as one in which the product's own cost moved when its invoice cost changed by at least two percent in absolute value that month. Parts are rounded to one decimal and shares computed from the three-decimal parts; that the 3.3 points in the first column equal the 3.3-point rise in the last row is a coincidence. Supplemental Appendix Table H1 gives every part to three decimals.

their margin rate by 2.0 points. If a cost move in the two or six preceding months also counts as the product's own, the share of the price effect in months without one drops from a half to a third or a quarter. A cutoff of 1, 5 or 10 percent in place of 2 puts it at 51, 66 and 81 percent. Across these rules the share runs from a quarter to four fifths.

When the increases were made. The weekly regular-price feed dates each price change to a week, so the price changes can also be sorted by when they were made. Applying the split to each product's average price and cost in each period leaves no third term: the price effect is 6.3 points and the two sides sum to 1.6. Two fifths of the price effect came from price increases made in detected review weeks, a little more than two fifths from increases made outside them, and about a sixth is movement in the transaction price that the weekly feed does not record. This is an accounting of where the observed increase was booked, not of what the margin rate would have been had those prices not been raised or had no review been held.

VIII. Implications

VIII.A. Directional Asymmetry

Synchronization and directional asymmetry. The micro-level spillovers aggregate into broad within-category repricing. In a panel of routine store–category–months rather than selected events (Supplemental Appendix D.2), the share of a category's unshocked products that reprice rises with the number of its products taking a cost shock during the run-up, by 0.4 percentage points with one shocked product and by 1.1 with four or more, relative to store–category–months with none. In the recovery the same shares are lower where shocks are present, so the pooled coefficient is small, and only cost moves of at least two percent raise repricing (Supplemental Appendix D.2). What this panel adds is this gradient inside the run-up and the direction of the response by year.

This repricing is directionally asymmetric, and the asymmetry is not a fixed feature of the review process. In 2021, the year costs rose fastest, an upward shock raised neighbors' probability of repricing upward by 1.41 pp while a downward shock did not raise downward repricing at all (the point estimate is -0.13 pp). The ordering runs the other way in 2019, 2020 and 2023, and upward dominance holds in 2021 and returns in 2024 and 2025 (Figure VII), so the one-sidedness is a property of the year rather than of the process.

Directional asymmetry and the year's cost environment. Each pair in Figure VII is estimated within one year on the same store–category–months, so the year's cost environment is held fixed inside the comparison. On the posted price the reversals do not appear, so the year-by-year pattern belongs to the price shoppers paid. What holds on both price objects and over the whole sample is narrower (Supplemental Appendix D.2): an upward cost shock raises upward repricing by 1.15 percentage points (0.18) and lowers downward repricing by 1.20 (0.24), while a downward shock does neither reliably. The mirror case, a calendar that fires on accumulated decreases as readily as on accumulated increases, is what the model of Supplemental Appendix C predicts and what these data cannot test. Cost decreases are a minority of the cost changes of at least two percent in every year of the sample, 44 percent at their highest in 2020, net accumulated cost movement is upward in every year, and the pressure measure of Section VI counts increases only. The distribution of price changes shifts markedly over time in other data as well (Berger and Vavra, 2018).

This locates the rockets-and-feathers asymmetry, prices that rise faster than

Figure VII: Directional asymmetry in neighbor repricing by year.

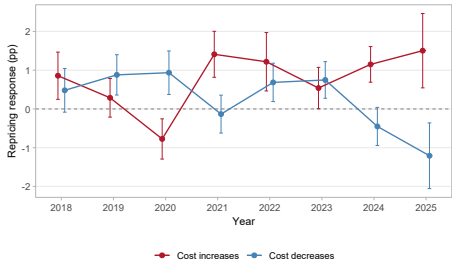

Notes: Observation: store–category–month, over all of the group's stores (Supplemental Appendix A.3) and on the monthly unit value. Each point is the coefficient, in percentage points, from a regression of the share of unshocked neighbors whose monthly unit value rises (red, on an indicator for at least one product in the category taking a cost increase of at least 2%) or falls (blue, on an indicator for a cost decrease of at least 2%), estimated separately by year. Store×category and year-month fixed effects; standard errors clustered on the category by year-month cell, the level at which invoice costs change. Whiskers are 95% confidence intervals, and twelve of the sixteen coefficients are distinguishable from zero. Estimating the difference between the two responses directly, upward less downward, with standard errors clustered two way on the category and the month, gives +0.38 (0.85) in 2018, −0.59 (0.34) in 2019, −1.71 (0.53) in 2020 and +1.54 (0.62) in 2021. It gives +0.53 (0.52) in 2022, −0.21 (0.34) in 2023, +1.60 (0.56) in 2024 and +2.71 (0.92) in 2025. A joint test that the eight differences are equal is rejected. Supplemental Appendix D.2 reports the specification.

they fall, (Peltzman, 2000; Borenstein, Cameron, and Gilbert, 1997) inside a single firm's pricing process rather than across market stages, and it shows the asymmetry Ellingsen et al. (2026) derive from a menu-cost model with moderate inflation to be state-dependent.

The asymmetry requires no error or bias on the retailer's part. During a run-up the pressure that has accumulated by the time a review is held is mostly upward, so reviews predominantly raise the category's price structure; where increases have accumulated less, 2020 being the year in which the two directions came closest to balance, the ordering is weaker or runs the other way. Anticipation would produce the opposite of what the data show: if a focal shock signaled correlated future cost increases, neighbors would move ahead of their realized costs, yet among neighbors away from any review week prices do not move in the month the shock lands (Figure VI, panel (b)).

VIII.B. Distributional Incidence

Because the extra margin on cheap products accumulates over many adjustments rather than within one event, a cross-sectional long difference asks which products gained the most margin between 2018 and 2022, on the 34,680 UPCs observed in

both years. I estimate:

$$\Delta MR_{jc} = \alpha_c + \beta \cdot Cheap_{jc} + \gamma \, \Delta \ln C_{jc}^{wh} + \varepsilon_{jc}, \tag{5}$$

where ΔMR_{jc} is the change in the margin rate of product j in category c between 2018 and 2022, $Cheap_{jc}$ indicates cheap-tier status, and α_c are category fixed effects. Table V reports this regression augmented with a $Cheap_{jc} \times \Delta \ln C_{jc}^{wh}$ interaction, centered at the sample's average cost growth, so the $Cheap_{jc}$ coefficient there is the cheap-tier differential at average cost growth. The common-increment mechanism of Section III, the same cash raise applied across the shelf, predicts a positive main effect with a null interaction, because a near-uniform repricing increment loads proportionally more on cheap items regardless of the size of any single shock.

The increment lands at the bottom of the ladder. Cheap-tier products gained 2.6 pp more margin than the rest at the sample's average cost growth, holding cost changes and category fixed effects constant (Table V). That is the margin on the price shoppers paid, with the tiers ranked on the same price. Ranking and measuring on the posted regular price throughout gives 1.54 points (0.64) on the products the weekly feed covers, and the two rankings agree on 88.5 percent of the cheap fifth (Supplemental Appendix G.2). The difference is promotional. Between the baseline and the run-up the gap between the posted price and the price paid narrowed by 1.75 points, by 1.12 points at the cheap end of the ladder against a widening of 0.15 at the premium end (Supplemental Appendix B.2). A large part of what the unit-value margin records is therefore promotional depth withdrawn rather than a higher posted price. The tier differential survives inside the retailer's own brand, at 2.34 points (0.94) on 1,162 products in 32 categories, where a third party's trade funding cannot produce it.[20] The gain looks like a level shift rather than differential pass-through: the interaction with cost growth is imprecise, and evaluated at the mean cost growth it could move the differential by about a point in either direction.

Interacting cheap-tier status with the category's synchronization share, the share of its unshocked products that reprice when a shock lands, gives +1.67 percentage points (0.41) of extra cheap-tier margin at full synchronization against none, so the cheap tier's extra margin arrives where repricing was most synchronized (Supplemen-

[20] On the products present in both 2019 and 2022, the differential stays positive and significant unweighted, sales-weighted, or with a baseline margin-rate control, at +1.57, +2.13 and +0.99 (Supplemental Appendix G.2).

Table V: Long-difference margin regressions on cheap-tier status and wholesale cost changes.

	Estimate	Std. Err.
$Cheap_{jc}$	$+2.59^{***}$	0.54
$\Delta \ln C_{jc}^{wh}$	-30.05^{***}	2.91
$Cheap_{jc} \times \Delta \ln C_{jc}^{wh}$	-1.22	3.26
Category FE	Yes	
N UPCs	34,680	

Notes: Regression: $\Delta MR_{jc} = \alpha_c + \beta\, Cheap_{jc} + \gamma\, \Delta \ln C_{jc}^{wh} + \delta\, Cheap_{jc} \times (\Delta \ln C_{jc}^{wh} - \overline{\Delta \ln C^{wh}}) + \varepsilon_{jc}$, equation (5) with the interaction added, where $\overline{\Delta \ln C^{wh}}$ is the sales-weighted mean cost growth, 0.143 log points, so the $Cheap_{jc}$ coefficient is the differential at average cost growth. *Observation:* UPC j within category c, on year-level sales-weighted prices, seven-month moving-average costs and margin rates for 2018 and 2022, built from the objects of Section II; products whose margin rate falls outside $(-1, 1)$ in either year are dropped, and the regression is weighted by 2018 sales. *Dependent variable:* the change in the gross margin rate, percentage points, 2018 to 2022. *RHS:* $Cheap_{jc}$ indicates the bottom 20% of 2018–2019 sales-weighted price per volume within category, the ranking used throughout; $\Delta \ln C_{jc}^{wh}$ is the cumulative log wholesale cost change. *Fixed effects:* category. Standard errors clustered at the category level ($N = 34{,}680$ UPCs across 232 categories with at least twenty ranked products). The interaction is imprecise, with a 95 percent interval from -7.6 to $+5.2$ per log point of cost growth, which at the mean cost growth moves the differential by between -1.1 and $+0.7$ points. The same regression on the regular shelf price and on the products the weekly feed covers is in Supplemental Appendix G.2. $^{***}p < 0.01$, $^{**}p < 0.05$.

tal Appendix D.2). The gap accumulates through both the number and the size of a cheap product's repricings, and the two work in different places: in big review weeks the cheapest fifth is repriced almost as often as the most expensive, 36.8 percent against 34.0. Outside review weeks it is repriced two and a half times as often, 2.7 percent against 1.1 (Supplemental Appendix G.3).

Why the increment lands on cheap products. Prices are adjusted in cash, in reviews and outside them alike, and the cash amounts rise more slowly than the prices they are applied to. Ranking every product by the shelf price at which it enters the review week, within its own store–category, the average raise in a big review runs from 22 cents on the cheapest tenth to 83 cents on the dearest, a factor of nearly four. The prices themselves span a factor of five and a half, \$1.96 against \$10.85 (Figure VIII, panel (a)). Divided by the price it lands on, the average raise is 12.1 percent on the cheapest tenth of a shelf against 8.4 percent on the dearest, and it falls at five of the nine steps in between.

The tilt is a property of the menu rather than of any decision about cheap products: close to nine in ten raises in big review weeks land on a multiple of ten cents, and the amounts in between are almost never used (Figure V, panel (b)). No single review

creates the whole gap, since one cost shock produces only a small, at best marginal extra response of the lowest-priced fifth (Table III). A product repriced outside a review draws from the same menu, so the calendar times these adjustments rather than concentrating the gap in review weeks (Supplemental Appendix G.3).

The gap accumulates over the inflation episode: each pass that repriced a cheap product moved its price by a larger fraction, and the asymmetry kept the increases while letting few decreases through.

What the cheapest fifth's extra growth is made of. The menu of round cash amounts hands the larger percentage raise to a cheap product. I ask what the gap in regular-price growth between the cheapest fifth and the rest would have been had the retailer used one percentage for everything it repriced in a week instead. I keep the dates on which prices moved, the set of products that moved, and the invoice costs, and give every product repriced in a store and category in a given week the average percentage move of the products repriced alongside it. The category's own inflation that week is unchanged; only its spread across the shelf changes. Figure VIII, panel (b), draws the gap week by week from 2018 to 2022, as it happened and under three versions of that rule: equal moves in detected review weeks only, equal moves in every week, and equal moves with equal frequency. As made, the cheapest fifth's regular price ends 2022 5.9 log points above the rest's. Applying the rule to the moves made in detected review weeks alone brings that to 4.6, and applying it in every week brings it to 4.1.

The rule works on the size of each move and leaves a difference in their number untouched. Of the 5.9 points, 1.5 are made in detected review weeks and 4.3 outside them, the parts rounded to one decimal. In a big review week the cheapest fifth is raised by 9.5 percent against 8.4 for the rest, and equalizing size takes that part of the gap from 1.5 points to 0.3. Outside review weeks the two are raised by almost the same percentage, 8.4 against 8.5, and the same rule takes the remaining 4.3 only to 3.9. What differs outside a review is how often a price moves, 2.7 percent of the cheapest fifth's active product-weeks against 1.4 percent for the rest. The equal-frequency line gives every product in a category that category's own average move each month, repriced or not, and the gap ends at 0.1 log points. That last step is four of the 5.9 points, and it is a difference in how often the cheapest fifth is repriced, since the size of each raise is already equal along the every-week line.

Two of the three decisions of Section III are at work here, and they are complements: how often a price is reconsidered sets how many increments accumulate, and

the size of each adjustment sets what an increment is worth. On the price the size rule leaves 4.1 of the 5.9 points standing; on the margin it takes the cheap-tier coefficient of the Table V regression, re-estimated on the 13,474 products whose weekly feed is balanced across the window, from 2.96 as made to 1.59, close to half. Removing the difference in frequency as well leaves that coefficient at −1.37 (1.20), indistinguishable from zero; the point estimate below zero is not a reversal, since that benchmark also removes the price growth each product's own cost change required. This is an accounting benchmark on the observed dates and quantities, not a prediction of what the retailer would have charged.

Figure VIII: The size of a review's price raises across the shelf, and where the cheapest fifth's extra regular-price growth comes from.

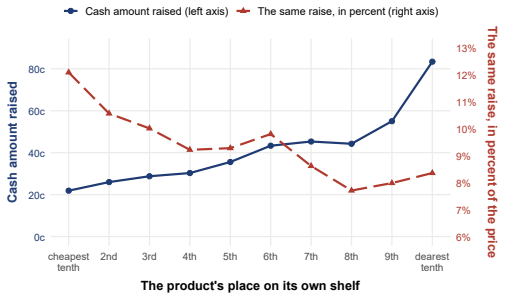

(a) Raises in big review weeks, in cash and in percent

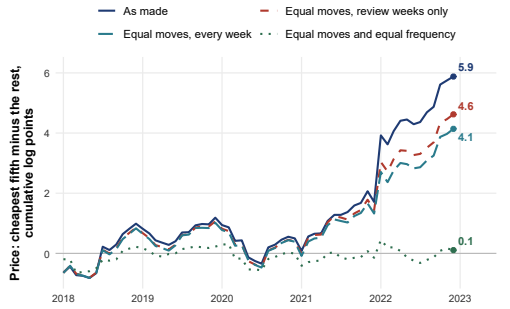

(b) The cheapest fifth's extra growth under equal moves

Notes: Panel (a). *Observation:* a product whose regular price was raised in a big review week, meaning a week in which at least a quarter of its store–category was repriced; 3,865,560 raises across the 551 supermarkets, 2018 to April 2024. The raise is the regular price two weeks after the review week minus the price the week before. Products are ranked into tenths by the package price entering the review week, within their own store–category–week, so every comparison is between products on one shelf; this is a shelf-price ranking rather than the price-per-volume ranking behind the tiers of Table V, and the tenths run from an average price of $1.96 to $10.85. Solid navy: the average cash amount raised, left axis; dashed red: the same raises divided by the price they land on, right axis. Nothing is excluded, with no cap on the price or on the size of the raise. Medians are lower and flatter, 9.1 percent against 7.1, because the median sits on whichever cash amount is modal for that part of the shelf. Panel (b). *Observation:* store–product–week. Each line is the cumulative difference between the cheapest fifth's and the rest of the ladder's log regular price, in log points, from the first week of 2018 to the last week of 2022, on the balanced store copies of the products in Table V that the weekly feed covers throughout, weighted by baseline sales. *As made* uses the price changes the retailer made. *Equal moves, review weeks only* gives every product repriced in a store, category and week of a detected review the average move of the products repriced alongside it, increases and decreases kept apart, and leaves the changes made outside those weeks as they are. *Equal moves, every week* applies the same rule in every week, and *equal moves and equal frequency* gives every product in a category that category's own average move each month, repriced or not, so that within a category no product moves more often or by more than another, the category's own monthly inflation unchanged.

Two sizes of the same product. The sharpest test of that mechanism compares different sizes of the *same product* within the same brand, a 12-ounce against a 24-ounce box of the same cereal. Within a brand and product type the lower-priced packs gained 1.81 pp more margin than the rest between 2018 and 2022, the window of Table V, about 70 percent of the gap across the full ladder, and 1.04 (0.28) comparing the 2018–2019 baseline with the 2021–2022 run-up. Two sizes of one product share a manufacturer, a shelf, and a demand environment, so trade-down, a change in elasticity, or a compositional shift would move both together. What separates them is that the same cash adjustment is a larger proportional move on the smaller pack, though this comparison alone cannot exclude a demand change specific to a pack size. The same gap is present in the recovery, at 1.28 points on the two-year windows (Supplemental Appendix G.1).

VIII.C. Household Incidence

I rank the 2.43 million households observed in both years into fifths by the cheap-tier share of their 2019 basket, and price each household's 2019 basket at 2019 and at 2022 prices. The 2019 basket of the fifth that bought cheapest cost 24.9 percent more in 2022, against 20.2 percent for the fifth that bought least cheap (Figure IX). Of that increase, \$9.07 per \$100 of 2019 spending would not have arisen had every product kept its 2019 invoice-margin rate, against \$6.95 for the least cheap fifth. The ordering holds in every cut of the sample, though two categories, soft drinks and water, account for more than all of the \$2.12 (Figure IX, note). Since 55 percent of the sales-weighted category price increases of the run-up came from products with no contemporaneous cost shock of their own (Supplemental Appendix H.1), a shopper might hope to avoid those products.[21] Across the fifths, the retailer's margin on the basket each bought before the inflation rose by between 4.3 and 5.6 percentage points, the most for the fifth buying cheapest. The products households switched to recovered between 0.24 and 0.88 points of it (Supplemental Appendix I). Households traded down more in the categories whose costs rose most, and trading down moves a basket toward the products whose margins rose most, though the fifth already buying cheapest moved up rather than down and recovered the most. Clustered on the category, the conservative choice for a shock that varies across 372 categories, the trade-down response to cost, the change in the spending-weighted price tier of a

[21] That is a share of category price increases, a different object from the share of the price increases behind the margin rise in Section VII.

household's purchases in a category as that category's cost rises, −0.043, is imprecisely estimated (standard error 0.034), while the link between trading down and margin exposure, the average margin-rate change of the products a household buys, −0.011 (0.0043), survives. A household's pre-inflation cheap-tier share interacted with the cost shock has a coefficient of 3.50 with a category-clustered standard error of 0.50, so this evidence says which households traded down further and not how far the average household did (Supplemental Appendix I.1). This is margin exposure inside one retailer, measured on continuing shoppers and continuing products, so substitution to other retailers, exit from a category and changes in quantity lie outside it.

Figure IX: What the 2019 basket cost in 2022, by how much of the basket is cheap-tier.

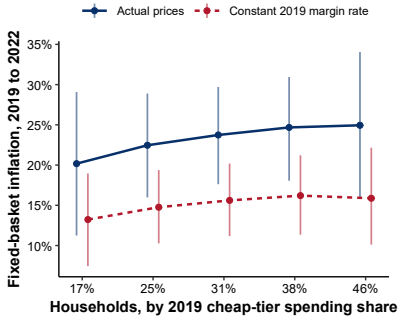

Notes: Observation: household. The 2.43 million households observed in both 2019 and 2022 are ranked into fifths by the share of their 2019 spending that went to cheap-tier products, from 16.6 percent in the lowest fifth to 45.8 in the highest, which labels the horizontal axis. Each household's 2019 basket is held fixed and priced at both years' prices, and the vertical axis is the change in what that basket costs between 2019 and 2022, in percent, averaged over the fifth's households. The solid navy line prices it at the prices the retailer charged, the dashed red line holds every product's 2019 invoice-margin rate fixed so that each price moves with its own invoice cost alone, and the distance between them is the margin component, in dollars per $100 of 2019 spending. Vertical bars span the 10th to the 90th percentile of households within a fifth, the two series offset slightly on the horizontal axis. Prices and costs are each product's 2019 and 2022 sales-weighted unit value and seven-month-average invoice cost across every store in the monthly data (Supplemental Appendix I). Each household's 2019 spending on a category and tier is priced by that cell's 2019-sales-weighted price relative, a Laspeyres at the category-by-tier level that ignores substitution inside a tier, on the 39,251 products sold in both years, which cover 55.6 percent of the 2019 purchase rows and 55.2 percent of the 2022 rows. The difference of $2.12 between the extreme fifths runs from $1.73 to $2.27 across six cuts of the sample. Treating the 367 of the panel's 372 categories in which either extreme fifth spent in 2019 as the independent draws gives the $2.12 a standard error of $1.98, so it is not distinguishable from zero. Soft drinks and mixers at 1.33 and water at 1.07 together exceed the whole difference, so the ordering holds in every cut and its size rests on those categories.

VIII.D. Alternative Explanations, Scope, and Policy

A demand-system account of neighbor repricing. One reading of the spillover runs through the demand system: when products are substitutes, a cost shock to one good changes the optimal prices of related goods. The demand channel fails on the *cross-sectional pattern* of the repricing. Measured on the impact window, where the focal product's transaction price has moved by 4.3 percent and the neighbors' transaction price by less than 0.05 percent, though their posted price does rise in that month (Figure III), the pooled cross-price elasticity is 0.084, with a standard error of 0.025.[22] Same-subcategory neighbors would need a cross-price elasticity of 0.38 to justify what they do and measure 0.40, so for the products a demand model calls substitutes demand does account for the repricing. Neighbors in the same category but a different subcategory, two thirds of the category-level neighbor sample, measure a *negative* cross-price elasticity, so no benchmark implies that they should reprice at all. Pooled over all neighbors and on the posted price, they move 0.71 percent against a focal 5.41 at the third month, and 0.86 percent among cost-stable neighbors, 1.4 to 1.7 times the response the measured cross-elasticity delivers (Supplemental Appendix E.2); that is what a category-level rule applied across the shelf produces and what product-by-product reoptimization does not.

The same account extends to the distributional result: if consumers moved toward cheap products, demand for them may have become less elastic than for premium ones, which would justify raising their margins by more with no pricing friction at all. Estimated separately by tier and period, cheap products are the more elastic at baseline and both tiers become less elastic during the run-up, the cheap tier by more (Supplemental Appendix E.1), a rotation in the direction the demand account requires.

The static single-product Lerner rule (Lerner, 1934) ties the margin level to con-

[22]The required elasticity is the value that would make the demand-implied neighbor response equal the observed one, under two demand benchmarks: a calibrated logit and, more generously, an envelope that allows the largest response any single-product pricing rule can justify (Supplemental Appendix E.2). Against the raw five-month change in the neighbor's transaction price, the logit puts the pooled requirement at 155 and the envelope at 0.32, against a measured 0.084 and a median near 0.10 across more than 7,000 published cross-price elasticities (Auer and Papies, 2020; Nevo, 2001). Once store–product and year-month effects absorb the ambient inflation of the window, both requirements change sign, which is why the level test is not decisive. On the posted price the envelope requirement is 0.119 in the full neighbor sample and 0.143 among cost-stable neighbors, and for same-subcategory products it is met, 0.379 required against 0.399 measured (Supplemental Appendix E.2).

temporaneous elasticity. Read on changes, and if whatever moved elasticities or the conduct term moved them by a common factor along the price ladder, it predicts that margins rise most on the least elastic products; under a shift concentrated on the elastic end they need not (Supplemental Appendix Hypothesis 5). I estimate:

$$\Delta MR_{jc} = \alpha + \delta \cdot z_{\beta,j} + \gamma \, \Delta \ln C^{wh}_{jc} + \varepsilon_{jc}, \tag{6}$$

where $z_{\beta,j}$ is the standardized pre-period elasticity, estimated for each UPC from a log-log regression of quantity on price on 2018–2019 data, $\ln q_{sjt} = \alpha_{sj} + \gamma_t + \beta_j \ln p_{sjt} + \varepsilon_{sjt}$, with store×UPC fixed effects and month effects.[23] The most demand-sensitive products gained more margin than the least, by 0.65 to 1.35 percentage points per standard deviation of the elasticity across the specifications of Supplemental Appendix E.1, the ordering a shelf-wide increment produces with no change in demand at all. Because the cheap tier became less elastic by more than the premium tier, product-by-product pricing can produce the same ordering, so the ranking on its own separates neither account; the timing evidence and the cross-price bound above do.

A calibrated version of that rule puts a number on the distance. Reallocating each category's total gross margin across its price ladder by the static inverse-elasticity condition, holding the category's gross margin and its quantities fixed, closes 56 to 80 percent of the cheap–premium gap on the 319 categories where such a rule can be written down. Across all categories it closes 22 to 31 percent of the 2.4-point gap of Section III, the 0.53 to 0.75 points it removes from the 0.94-point gap between the cheapest and most expensive fifths in the categories it can price (Supplemental Appendix K).

External validity. The data come from a single multi-banner grocery group, so what matters is whether the *mechanism* extends to other settings. It rests on three features: dense within-category assortments, so that a cost shock to one product has many unshocked neighbors to reprice; centralized, periodic category reviews, through which the adjustment propagates; and recurring idiosyncratic upstream cost shocks to trigger them. The first and third are general to large-format grocery. Grocery prices are managed by category (Zenor, 1994; Basuroy, Mantrala, and Walters, 2001)

[23] Retail prices may respond to product-level demand shocks, so these elasticities are descriptive and serve only to rank products by demand sensitivity. I restrict to products with $\beta \leq 0$, trim the margin change and the cost change at the 1st and 99th percentiles, and standardize the estimates on that sample (Supplemental Appendix E.1).

and nearly uniform across stores within a chain (DellaVigna and Gentzkow, 2019), with temporary sales planned substantially in advance (Anderson et al., 2017). The extra margin on cheap products follows from coarse adjustments applied to a heterogeneous price ladder, so it should arise wherever adjustments are coarse and the ladder is heterogeneous, and its outcome, larger price increases on cheaper products, is documented across 91 retailers in 10 countries (Cavallo and Kryvtsov, 2024). In the framework of Al-Ubaydli and List (2015) and List (2020), category management and chain-level pricing are constraints observed to match elsewhere; the review calendar and the adjustment menu are, so far, this retailer's.

The density constraint predicts a smaller spillover where a category has few products to reprice. The dose-response spillover is present in full supermarkets, which carry between 16,500 and 18,200 distinct products in a store–month (Table I, Panel B), and absent in the thin-assortment express stores set aside from the analysis sample. Those stores show no own pass-through either and carry a different assortment at a different price level, so the contrast is consistent with the density constraint without isolating it (Supplemental Appendix A.1). The group's prices follow national food-at-home inflation closely through 2022, though by May 2025 its price index stands 11 index points above the food-at-home CPI (Supplemental Appendix J.2).

Centralized pricing. A store's cost record updates when a shipment to it is invoiced, so the same chain-level cost change reaches some stores' records in one month and others' later. Where it is recorded in some stores but not others in a month, neighbor price changes in the stores recording it run 0.072 percentage points (standard error 0.003) below those in the other stores six months after the shock, and 0.069 points sales-weighted. The comparison covers 1,096,543 such differential events with at least five stores on either side, in 550 of the 551 supermarkets. That stands against a spillover of 0.69% in the full neighbor sample and 0.85% among cost-stable neighbors six months on. A differential a tenth that size, pointing the wrong way for a store-level channel, indicates that the category review is centralized (Supplemental Appendix J.1), consistent with evidence that U.S. grocery prices are nearly uniform at the chain or zone level (Hitsch, Hortaçsu, and Lin, 2021; DellaVigna and Gentzkow, 2019). Companion work on the same retailer examines how store-level cost differences interact with a single chainwide price (Janssen, 2026).

Policy implications. First, a regulatory focus on firm-level markups misses the redistribution within categories, which turns on how margin is allocated across products.

Second, a benefit or an allowance indexed to an average basket understates what a household buying at the cheap end faced, since cheap-tier retail *prices* rose by more than premium prices in these data (Table II); Section VIII.C measures the amount for one retailer's shoppers.

IX. Conclusion

How a retailer organizes its price changes shapes the distributional incidence of inflation. The prices of a shocked product's neighbors, net of their own invoice costs, rise with it where a category-wide repricing episode is detected around the shock and show no average response where none is. The resulting margin expansion concentrates on cheap products (Table V) and falls hardest on the baskets of the households that buy cheapest at this retailer; had every product repriced in a category-week moved by that cell's average percentage, the cheapest fifth's extra margin would have been about half as large (Section VIII.B). An average basket therefore understates what a cheap-oriented basket paid (Section VIII.C). Because the gross-to-net cost wedge (trade funding, labor, occupancy, shrinkage) may itself differ across tiers, the gross-margin gap need not equal the change in economic markup.

References

Al-Ubaydli, Omar, and John A. List. 2015. "On the Generalizability of Experimental Results in Economics." In *Handbook of Experimental Economic Methodology*, ed. Guillaume R. Fréchette and Andrew Schotter, 420–462. New York: Oxford University Press.

Alexander, Patrick, Lu Han, Oleksiy Kryvtsov, and Ben Tomlin. 2024. "Markups and Inflation in Oligopolistic Markets: Evidence from Wholesale Price Data." Bank of Canada Staff Working Paper No. 2024-20.

Alvarez, Fernando and Francesco Lippi. 2014. "Price Setting With Menu Cost for Multiproduct Firms." *Econometrica*, 82(1): 89–135.

Anderson, Eric, Nir Jaimovich, and Duncan Simester. 2015. "Price Stickiness: Empirical Evidence of the Menu Cost Channel." *Review of Economics and Statistics*, 97(4): 813–826.

Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jón Steinsson. 2017. "Informational Rigidities and the Stickiness of Temporary Sales." *Journal of Monetary Economics*, 90: 64–83.

Armstrong, Mark and John Vickers. 2023. "Multiproduct Cost Pass-Through: Edgeworth's Paradox Revisited." *Journal of Political Economy*, 131(10): 2645–2665.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2024. "Algorithmic

Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy*, 132(3): 723–771.

Auer, Johannes and Dominik Papies. 2020. "Cross-Price Elasticities and Their Determinants: A Meta-Analysis and New Empirical Generalizations." *Journal of the Academy of Marketing Science*, 48(3): 584–605.

Basuroy, Suman, Murali K. Mantrala, and Rockney G. Walters. 2001. "The Impact of Category Management on Retailer Prices and Performance: Theory and Evidence." *Journal of Marketing*, 65(4): 16–32.

Berger, David and Joseph Vavra. 2018. "Dynamics of the U.S. Price Distribution." *European Economic Review*, 103: 60–82.

Bernanke, Ben S. and Olivier J. Blanchard. 2025. "What Caused the U.S. Pandemic-Era Inflation?" *American Economic Journal: Macroeconomics*, 17(3): 1–35.

Besanko, David, Jean-Pierre Dubé, and Sachin Gupta. 2005. "Own-Brand and Cross-Brand Retail Pass-Through." *Marketing Science*, 24(1): 123–137.

Bhattarai, Saroj and Raphael Schoenle. 2014. "Multiproduct Firms and Price-Setting: Theory and Evidence from U.S. Producer Prices." *Journal of Monetary Economics*, 66: 178–192.

Bils, Mark and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy*, 112(5): 947–985.

Bonomo, Marco, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, and Rodolfo Rigato. 2023. "Multi-Product Pricing: Theory and Evidence from Large Retailers." *The Economic Journal*, 133(651): 905–927.

Borenstein, Severin, A. Colin Cameron, and Richard Gilbert. 1997. "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, 112(1): 305–339.

U.S. Bureau of Labor Statistics. 2024. "Consumer Price Index for All Urban Consumers: Food at Home in U.S. City Average." Series CUSR0000SAF11 (seasonally adjusted), retrieved from FRED, Federal Reserve Bank of St. Louis, `https://fred.stlouisfed.org/series/CUSR0000SAF11`.

Butters, R. Andrew, Daniel W. Sacks, and Boyoung Seo. 2022. "How Do National Firms Respond to Local Cost Shocks?" *American Economic Review*, 112(5): 1737–1772.

Byrne, David P. and Nicolas de Roos. 2019. "Learning to Coordinate: A Study in Retail Gasoline." *American Economic Review*, 109(2): 591–619.

Cavallo, Alberto and Oleksiy Kryvtsov. 2024. "Price Discounts and Cheapflation during the Post-Pandemic Inflation Surge." *Journal of Monetary Economics*, 148 (Supplement): 103644.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics*, 134(3): 1405–1454.

Chen, Tao, Peter Levell, and Martin O'Connell. 2024. "Measuring Cost of Living Inequality during an Inflation Surge." CEPR Discussion Paper No. 19388.

Chintagunta, Pradeep K. 2002. "Investigating Category Pricing Behavior at a Retail Chain." *Journal of Marketing Research*, 39(2): 141–154.

Conlon, Christopher. 2026. "Did Profits Cause Inflation?" *International Journal of Industrial Organization*, 105: 103257.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics*, 135(2): 561–644.

DellaVigna, Stefano and Matthew Gentzkow. 2019. "Uniform Pricing in U.S. Retail Chains." *Quarterly Journal of Economics*, 134(4): 2011–2084.

Döpper, Hendrik, Alexander MacKay, Nathan H. Miller, and Joel Stiebale. 2025. "Rising Markups and the Role of Consumer Preferences." *Journal of Political Economy*, 133(8): 2462–2505.

Dubé, Jean-Pierre and Sachin Gupta. 2008. "Cross-Brand Pass-Through in Supermarket Pricing." *Marketing Science*, 27(3): 324–333.

Dubois, Pierre, Rachel Griffith, and Martin O'Connell. 2022. "The Use of Scanner Data for Economics Research." *Annual Review of Economics*, 14: 723–745.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference Prices, Costs, and Nominal Rigidities." *American Economic Review*, 101(1): 234–262.

Ellingsen, Tore, Richard Friberg, and John Hassler. 2026. "Menu Costs and Asymmetric Price Adjustment." *International Journal of Industrial Organization*, 105: 103259.

Federal Trade Commission. 2024. "Feeding America in a Time of Crisis: FTC Staff Report on the United States Grocery Supply Chain and the COVID-19 Pandemic." Staff Report, Bureau of Competition and Bureau of Economics, Washington, DC, March.

Friberg, Richard and André Romahn. 2018. "Pass-Through by Multi-Product Firms." *International Journal of the Economics of Business*, 25(2): 265–295.

Gardner, John. 2022. "Two-Stage Differences in Differences." arXiv:2207.05943.

Glover, Andrew, José Mustre-del-Río, and Alice von Ende-Becker. 2023. "How Much Have Record Corporate Profits Contributed to Recent Inflation?" *Federal Reserve Bank of Kansas City Economic Review*, 108(1): 23–35.

Goldberg, Pinelopi K. and Rebecca Hellerstein. 2013. "A Structural Approach to Identifying the Sources of Local Currency Price Stability." *Review of Economic Studies*, 80(1): 175–210.

Gopinath, Gita and Oleg Itskhoki. 2010. "Frequency of Price Adjustment and Pass-Through." *Quarterly Journal of Economics*, 125(2): 675–727.

Griffith, Rachel, Lars Nesheim, and Martin O'Connell. 2018. "Income Effects and the Welfare Consequences of Tax in Differentiated Product Oligopoly." *Quantitative Economics*, 9(1): 305–341.

Handbury, Jessie. 2021. "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities." *Econometrica*, 89(6): 2679–2715.

Hitsch, Günter J., Ali Hortaçsu, and Xiliang Lin. 2021. "Prices and Promotions in U.S. Retail Markets." *Quantitative Marketing and Economics*, 19(3–4): 289–368.

Hong, Gee Hee and Nicholas Li. 2017. "Market Structure and Cost Pass-Through in Retail." *Review of Economics and Statistics*, 99(1): 151–166.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein. 2016. "Quantifying the Sources of Firm Heterogeneity." *Quarterly Journal of Economics*, 131(3): 1291–1364.

Janssen, Aljoscha. 2026. "The Cost Side of Uniform Pricing: Cost Differences and Store Margins." Unpublished working paper, Singapore Management University; manuscript on file with the author.

Jaravel, Xavier. 2019. "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector." *Quarterly Journal of Economics*, 134(2): 715–783.

Kaplan, Greg and Sam Schulhofer-Wohl. 2017. "Inflation at the Household Level." *Journal of Monetary Economics*, 91: 19–38.

Kehoe, Patrick J. and Virgiliu Midrigan. 2015. "Prices Are Sticky After All." *Journal of Monetary Economics*, 75: 35–53.

Klenow, Peter J. and Benjamin A. Malin. 2010. "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics*, Vol. 3A, ed. Benjamin M. Friedman and Michael Woodford, 231–284. Elsevier.

Lerner, Abba P. 1934. "The Concept of Monopoly and the Measurement of Monopoly Power." *Review of Economic Studies*, 1(3): 157–175.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains." *Quarterly Journal of Economics*, 112(3): 791–824.

List, John A. 2020. "Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial." NBER Working Paper 27535.

McShane, Blakeley B., Chaoqun Chen, Eric T. Anderson, and Duncan I. Simester. 2016. "Decision Stages and Asymmetries in Regular Retail Price Pass-Through." *Marketing Science*, 35(4): 619–639.

Midrigan, Virgiliu. 2011. "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations." *Econometrica*, 79(4): 1139–1180.

Mongey, Simon and Michael E. Waugh. 2025. "Pricing Inequality." NBER Working Paper No. 33399.

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

Nakamura, Emi and Dawit Zerom. 2010. "Accounting for Incomplete Pass-Through." *Review of Economic Studies*, 77(3): 1192–1230.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69(2): 307–342.

Peltzman, Sam. 2000. "Prices Rise Faster than They Fall." *Journal of Political Economy*, 108(3): 466–502.

Rambachan, Ashesh and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies*, 90(5): 2555–2591.

Richards, Timothy J., Miguel I. Gómez, and Jun Lee. 2014. "Pass-Through and Consumer Search: An Empirical Analysis." *American Journal of Agricultural Economics*, 96(4): 1049–1069.

Stantcheva, Stefanie. 2024. "Why Do We Dislike Inflation?" *Brookings Papers on Economic Activity*, Spring 2024: 1–46.

Stroebel, Johannes and Joseph Vavra. 2019. "House Prices, Local Demand, and Retail Prices." *Journal of Political Economy*, 127(3): 1391–1436.

Sudhir, K. 2001. "Structural Analysis of Manufacturer Pricing in the Presence of a Strategic Retailer." *Marketing Science*, 20(3): 244–264.

Sun, Liyang and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*, 225(2): 175–199.

Tappata, Mariano. 2009. "Rockets and Feathers: Understanding Asymmetric Pricing." *RAND Journal of Economics*, 40(4): 673–687.

Villas-Boas, Sofia B. 2007. "Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data." *Review of Economic Studies*, 74(2): 625–652.

Weber, Isabella M. and Evan Wasner. 2023. "Sellers' Inflation, Profits and Conflict: Why Can Large Firms Hike Prices in an Emergency?" *Review of Keynesian Economics*, 11(2): 183–213.

Weyl, E. Glen and Michal Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121(3): 528–583.

Yang, Huanxing and Lixin Ye. 2008. "Search with Learning: Understanding Asymmetric Price Adjustments." *RAND Journal of Economics*, 39(2): 547–564.

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics*, 86(2): 514–533.

Zenor, Michael J. 1994. "The Profit Benefits of Category Management." *Journal of Marketing Research*, 31(2): 202–213.