Retail Margins and Grocery Inflation

Aljoscha Janssen*

Abstract

I match wholesale costs to scanner prices for 46,000 grocery products between 2018 and 2025. Gross margins rose from 23% to 29% during the post-COVID inflation episode, loading disproportionately on cheap products. The mechanism is category-level repricing: a cost shock to one product triggers a category review that raises prices on unshocked neighbors. These spillovers account for up to three-quarters of within-category price impact, and upward shocks trigger repricing 14 times more readily than downward shocks. A profit-neutral counterfactual attributes 44–61% of the cheap–premium margin gap to pricing constraints rather than optimal discrimination.

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

---

*Singapore Management University, School of Economics. Email: ajanssen@smu.edu.sg.

1 Introduction

A grocery retailer manages tens of thousands of prices but does not set them one at a time (Ellickson and Misra, 2008; Bonomo et al., 2023). When a wholesale cost shock hits a product, the pricing team reviews the entire category, examining costs, margins, and competitor prices for dozens of related products, and adjusts multiple prices in a coordinated review. This organizational reality suggests that inflation may not pass through product by product; it can propagate across the shelf through category-level repricing. The 2021–2022 inflation episode, which raised U.S. grocery prices by roughly 25% (Bureau of Labor Statistics, 2024), provides a laboratory for studying this process. Public debate framed the episode as a conduct question, asking whether firms raised margins during inflation (Weber and Wasner, 2023; Blanchard and Bernanke, 2023), but recent work argues that aggregate accounting data cannot distinguish demand shifts from conduct changes (Conlon, 2026). This article evaluates three mechanical questions. First, *how* does the pricing process transmit cost shocks across the shelf? Second, can it raise margins even absent any change in firm conduct? Third, does it generate distributional consequences? I open the black box by matching wholesale costs to scanner prices inside a multi-banner grocery group.

I study a major U.S. grocery group operating three retail chains with 551 stores, matching invoice-level wholesale costs to scanner prices at the store–UPC–month level from 2018 through 2025.[1] The match yields direct observation of the gross price–cost wedge for over 46,000 products: observed costs, not model-inferred markups (Villas-Boas, 2007; Alexander et al., 2024). Grocery retail is informative because managers set prices in dense categories with recurring category reviews: audits in which the pricing team adjusts multiple prices in a coordinated review, typically within one to three months of a triggering cost shock (Levy et al., 1997; Zbaracki et al., 2004).[2] I organize the analysis around four sets of results.

First, descriptive facts. Four patterns emerge from the data before any causal analysis. (i) Aggregate margin expansion: the gross margin rate rose from 23% in 2018–2019 to a peak of 29% in early 2022, then partially retraced but remained 3 percentage points above pre-pandemic levels through 2024. (ii) The expansion was not confined to a few categories: nearly all 142 product categories show margin increases, with substantial heterogeneity. (iii) Cheap-tier incidence: the cheapest decile gained 5.8 percentage points of margin versus 2.8 for premium items. Wholesale costs rose *more* for cheap products (14.2% vs. 8.4%), yet their margins expanded more, a puzzle since larger cost shocks should compress margins

---

[1] Sub-periods: *baseline* (2018–2019), *run-up* (2021–2022), *recovery* (2023–2024), with 2020 omitted. Unless labeled as an extended window (2019–2022), "run-up" refers to 2021–2022.

[2] Reviews typically occur one to two times per year per category (Klenow and Malin, 2011). Wholesale cost increases are the primary trigger (Anderson, Jaimovich, and Simester, 2015).

under standard product-level pass-through (Weyl and Fabinger, 2013). (iv) Within-category synchronization: when at least one product in a category receives a cost shock, 51.2% of unshocked neighbors also reprice, compared with 33.5% in unshocked categories in the same store and month, pointing to category-level repricing as the key organizational mechanism.

These descriptive patterns motivate the empirical analysis. The causal challenge is to establish that repricing spillovers are genuine, not an artifact of correlated costs or common demand shifts.

Second, the central causal finding addresses this directly: a cost shock to one product leads unshocked neighbors, other products in the same store and category whose own costs did not change, to reprice. The panel event study (the paper's main specification) shows that own prices jump 5.9% at the shock month while neighboring products' prices, after controlling for their own wholesale costs, remain flat at impact and drift upward over subsequent months. Each neighbor's price response is about one-quarter of the own-product response over a six-month horizon, consistent with pricing that operates at the category level, not product by product. A complementary design based on 212,640 group-level natural experiments confirms the pattern through dose-response variation in shock size.

Third, mechanism and incidence. Spillovers are larger for closer substitutes (consistent with demand-linked transmission) and much stronger for upward than downward cost shocks, contributing to upward margin drift over a cost cycle. During the run-up, upward shocks drove a 1.51 percentage point increase in neighbors' probability of raising prices, while downward shocks produced only 0.11 percentage points—a 14-to-1 asymmetry. By 2023, the pattern reversed as accumulated cost decreases breached the review threshold. Cheap-tier incidence builds through repeated category reviews: within each review, coarse adjustments that do not fully differentiate across the price ladder load larger percentage price changes on cheaper products, so repeated reviews accumulate larger margin increases on cheap items. A within-brand pack-size test (comparing small and large packages of the same product) confirms this supply-side channel, ruling out demand-side explanations. An inverted elasticity gradient, where more elastic products gain more margin, is inconsistent with independent product pricing (Lerner, 1934).

Fourth, household outcomes and magnitudes. Using a loyalty panel of 2.85 million households, I document that overall product switching provides only modest relief from margin exposure, and that within-category trade-down partially offsets this relief: consumers substitute toward exactly the products with the largest margin increases. Two profit-neutral counterfactuals, reallocating markups according to the Lerner inverse-elasticity rule (Lerner, 1934) and the Ramsey inverse semi-elasticity rule (Ramsey, 1927) while keeping total category profit fixed, suggest that 44–61% of the observed gap during the run-up reflects category-level pricing constraints rather than optimal price discrimination. The conservative Lerner estimate (44%) recognizes that cheap products have slightly lower percentage elasticities and thus warrant somewhat higher margins; the higher Ramsey estimate (61%) additionally captures the price-level arithmetic. The constraint is consistent with organizational costs of product-by-product optimization; the retailer could achieve at least the same profit with a less regressive allocation. If the cheap-tier loading operates similarly at other retailers, programs indexed to low-price baskets (such as SNAP) would lose purchasing power faster than headline inflation (Hastings and Shapiro, 2018).

These findings speak to four literatures. First, the paper provides a supply-side mechanism for cheapflation and inflation inequality. Cavallo and Kryvtsov (2024) documents cheapflation across 91 retailers in 10 countries but leaves the mechanism open. Jaravel (2019) traces differential price growth to product innovation; Kaplan and Schulhofer-Wohl (2017) show that inflation rates vary substantially across households; Handbury (2021) demonstrates that non-homothetic demand generates income-specific price indices; Chen, Levell, and O'Connell (2024) measure cost-of-living inequality during the recent surge; Griffith, O'Connell, and Smith (2015) show that commodity price shocks pass through to relative food prices with distributional consequences for household nutrition. On the mechanism side, Mongey and Waugh (2025) propose a demand-side channel: households with less elastic demand face higher markups. No prior paper directly observes the supply-side transmission—from wholesale cost shocks through category reviews to margin redistribution—inside a retailer using matched cost-price data. I show that category-level repricing with imperfectly differentiated adjustments can generate cheapflation mechanically, without requiring changed firm conduct or differential demand elasticities.

Second, I document how wholesale cost shocks propagate across products within a retailer's categories. Weyl and Fabinger (2013) develop the theory of pass-through under imperfect competition; Griffith, Nesheim, and O'Connell (2018) model tax pass-through in differentiated-product oligopoly with income effects and distributional welfare consequences; Armstrong and Vickers (2023) extend this to multiproduct firms, showing that cost pass-through depends on the full product portfolio. Empirically, Nakamura and Steinsson (2008) document five facts about the frequency and size of price changes; Bils and Klenow (2004) provide foundational evidence on price rigidity; Kehoe and Midrigan (2015) show that prices are stickier than raw frequency suggests once temporary sales are excluded; Eichenbaum, Jaimovich, and Rebelo (2011) distinguish reference prices from temporary sales using scanner data; Nakamura and Zerom (2010) estimate pass-through for coffee; Goldberg and Hellerstein (2013) use wholesale-to-retail data to structurally identify sources of incomplete pass-through; Gopinath and Itskhoki (2010) establish that the frequency of price adjustment shapes pass-through rates; Hong and Li (2017) study how market structure shapes retail pass-through; Nevo (2001) quantifies multiproduct pricing in grocery using structural demand estimation; Hottman, Redding, and Weinstein (2016) show that variable markups and product scope drive firm heterogeneity; and Stroebel and Vavra (2019) document how local demand conditions shape retail markups. In macro, Midrigan (2011) shows that economies of scope in menu costs generate within-category synchronization; Bhattarai and Schoenle (2014) confirm that within-firm synchronization far exceeds within-industry synchronization; Alvarez and Lippi (2014) derive optimal pricing rules for multiproduct firms; and Bonomo et al. (2023) document that large retailers adjust multiple products simultaneously. I provide large-scale evidence from natural upstream cost variation (212,640 quasi-experiments and a panel of 1.6 million balanced neighbor UPCs across 551 stores), showing that product-level pass-through is not a sufficient statistic in multiproduct retail.

Third, I connect to the category management literature on cross-brand pass-through. That literature studies how category-level pricing affects retail *prices*; this paper studies *margins* and their redistribution across the shelf. Besanko, Dubé, and Gupta (2005) estimate own-brand and cross-brand pass-through for 78 products in 11 categories, finding significant cross-brand effects; Dubé and Gupta (2008) confirm with longer data. McShane et al. (2016) find that wholesale price increases are passed through 70% of the time but decreases only 9%, a directional asymmetry similar to the rockets-and-feathers pattern I document. The category management institutional context is established by Zenor (1994); Sudhir (2001); Chintagunta (2002); Ellickson and Misra (2008); Anderson, Malin, Nakamura, Simester, and Steinsson (2017). Relative to this literature, I advance along three dimensions: persistent upstream cost shocks (versus temporary trade promotions (Besanko, Dubé, and Gupta, 2005; Dubé and Gupta, 2008)), scale (46,000 products across 142 categories versus 78 in 11), and a welfare analysis linking cross-brand pass-through to cheapflation and household margin exposure.

Fourth, I shift the greedflation debate from aggregate markups to within-category redistribution. Weber and Wasner (2023) argue that firms exploited inflation expectations to raise markups; Blanchard and Bernanke (2023) counter with standard pass-through; Conlon (2026) argues that most evidence cited for the profits-inflation hypothesis cannot distinguish demand shifts from conduct changes. The debate focuses on aggregate markups (De Loecker, Eeckhout, and Unger, 2020; Glover, Mustre-del-Río, and von Ende-Becker, 2023), and even the long-run markup trend may partly reflect changing consumer preferences rather than market power (Dopper et al., 2025), but neither side examines how the pricing process *redistributes* margins across products within categories. I show that pricing architecture generates large within-category redistribution even with modest aggregate changes. The directional asymmetry connects to the rockets-and-feathers literature (Peltzman, 2000; Tappata, 2009; Borenstein, Cameron, and Gilbert, 1997), where recent work traces the pattern to learned coordination (Byrne and de Roos, 2019) and algorithmic pricing (Assad et al., 2024). My finding adds a within-firm channel—category-level review costs—and shows that the asymmetry is regime-dependent, reversing during recovery as accumulated cost decreases breach the review threshold.

2 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 a major U.S. grocery group comprising three retail chains and 551 stores, carrying roughly 40,000–50,000 distinct UPCs.[3] For each store–UPC–month, I observe total dollar sales, unit sales, and the unit-value retail price.[4] I also observe wholesale acquisition costs: invoice-level prices the retailer pays to suppliers. I construct two cost measures: a seven-month trailing moving average ($C_{jt}^{wh,MA7}$) for descriptive statistics,[5] and step-costs (the raw invoice cost carried forward) for the causal analyses. These are *observed* costs, not imputed (Villas-Boas, 2007; Alexander et al., 2024). The invoice cost reflects the per-unit acquisition price from the supplier to the retailer's distribution center—the cost field the category manager observes when evaluating margins and triggering reviews.[6] Category assignments span three levels: 30 super-categories, 142 product categories (my primary unit), and approximately 450 product modules.

My main outcome variable is the *gross margin rate*.[7] Let j index a product (UPC), s a store, c a product category, and t a calendar month. For each store–product–month, I define:

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

where P_{jst} is the unit-value retail price and C_{jst}^{wh} the wholesale acquisition cost. I suppress

---

[3] The full raw data contain over 900 stores. I restrict the analysis sample to full-size supermarkets ($\geq$5,000 distinct UPCs) with complete baseline coverage (2018–2019), yielding 551 stores. Small-format and express stores are excluded because their thin assortment lacks the within-category density required for the spillover analysis; Online Appendix A.1 reports format-specific results.

[4] Unit values reflect the average transaction price within a store–UPC–month, incorporating promotional discounts. Results are robust to excluding promotional observations (Online Appendix B.5).

[5] Smoothing compresses individual shock magnitudes ($r = 0.27$; Online Appendix B.1) but modestly amplifies the cross-sectional cheap-tier gradient.

[6] Online Appendices B.2–B.4 report extensive robustness checks on the cost measure, including sensitivity to cost dating, momentum controls, and invoice frequency.

[7] Throughout, "margin" refers to the gross price–cost wedge $(P - C^{wh})/P$, excluding operating costs (labor, occupancy, shrinkage). This wedge is a direct measure of the retailer's transmission of upstream costs to downstream prices. It should not be interpreted as profit.

the store subscript s when the context is unambiguous. To define within-category price tiers, I compute the *price-per-volume* (PPV) for each product using 2018–2019 data:

$$PPV_{jc} = \frac{\bar{P}_{j,\text{pre}}}{\text{Volume}_j}, \tag{2}$$

where Volume_j is the package size in common units (ounces, count, etc.). I rank products within each category by PPV and assign them to tiers: deciles for descriptive gradients, a binary bottom-20% indicator ("cheap tier") for regressions, and quintiles for the household and counterfactual analyses. Results are robust across all definitions.

I restrict to products observed in both retail and wholesale datasets for at least six months in the baseline period (2018–2019, used to assign price tiers). Products need not survive to all subsequent periods; Table 1 reports period-specific sample sizes. I winsorize prices and costs at the 0.5th and 99.5th percentiles. The final sample includes over 46,000 UPCs and over 12 million store–UPC–month observations in the baseline alone (Table 1). Online Appendix Table A3 traces observation counts through each analytical sample used in the paper.

I organize the analysis around three sub-periods: *baseline* (2018–2019), *run-up* (2021–2022), and *recovery* (2023–2024), omitting 2020 to avoid pandemic contamination. Table 1 summarizes the data.

Figure 1 illustrates the data architecture for a single store and category (condiments). Panel (a) tracks a single UPC's retail price and wholesale cost; the price–cost wedge is directly observed each month. Panel (b) plots all UPCs in the category by PPV quartile, showing heterogeneous cost dynamics across the price ladder.

Prices outpaced costs: retail prices rose by 17% versus 12% for wholesale costs, and the aggregate margin rate expanded from 22.8% to 26.6% during the run-up before partially retracing to 25.7%. Private-label products carried higher baseline margins (39% vs. 19%) but experienced smaller margin expansion (+1.8 vs. +4.6 pp). Product turnover is substantial: the number of distinct UPCs falls from 46,971 in the baseline to 43,258 during the run-up and 38,445 in recovery (Table 1, Panel C), reflecting normal product churn.[8]

Household panel. I also observe household-level purchase records linked through the retailer's loyalty-card program, covering 96% of transactions. The loyalty data do not include household demographics, limiting the welfare analysis to within-household comparisons. Matching to the product-level margin data, I construct a balanced panel of 2.85 mil-

---

[8] Because entering products embed current-period pricing, composition changes could affect cross-sectional comparisons. A balanced-panel check (Online Appendix G.2) confirms that cheap-tier incidence partially attenuates on continuing products while the elasticity gradient is fully robust.

Table 1: 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 ($)	3.82	3.15	4.47	3.68	4.89	4.02
Wholesale cost ($)	2.95	2.62	3.29	2.91	3.63	3.24
Margin rate (%)	22.8	14.2	26.6	14.8	25.7	14.5
Dollar markup ($)	0.87	1.48	1.18	1.72	1.26	1.78
$\Delta \ln P$ (%, ann.)	—	—	10.2	18.4	4.6	15.7
$\Delta \ln C^{wh}$ (%, ann.)	—	—	7.8	22.1	5.3	19.8
Panel B: Store–month quantities						
Units sold (000s)	238	142	226	159	198	147
Dollar sales ($000s)	718	402	820	536	816	556
Distinct UPCs per store	14,842	—	13,768	—	13,387	—
Panel C: Product Coverage						
# UPCs in sample	46,971		43,258		38,445	
of which private label	7,017 (15%)		5,288 (12%)		5,065 (13%)	
PL share of dollar sales (%)	20.0		18.4		18.4	
# Stores	551		551		551	
# Super-categories	30		30		30	
# Product categories	142		142		142	
# Product modules	∼450		∼450		∼450	
Store–UPC–month obs.	∼12M		∼11M		∼10M	

Notes: Observation: store–UPC–month. Statistics reported for three sub-periods: baseline (January 2018 to December 2019), inflation run-up (January 2021 to December 2022), and recovery (January 2023 to December 2024). Panel A reports sales-weighted means across all store–UPC–month observations. Retail prices are monthly sales-weighted averages (total sales divided by total units for a given store–UPC–month); wholesale costs are smoothed using a seven-month trailing moving average of invoice-level data (see text). The gross margin rate is $(P - C^{wh})/P$ in percent; the dollar markup is $P - C^{wh}$ in dollars; log changes are annualized. Panel B aggregates across all UPCs within a store–month: a typical store sells approximately 238,000 units and generates $718,000 in revenue per month during the baseline period, spanning roughly 15,000 distinct UPCs. Units and sales are in thousands. Panel C reports counts of distinct products (including the private-label subset), private-label share of total dollar sales, stores, and the three-level category hierarchy. Private-label products are identified from the product master file. Prices and costs are winsorized at the 0.5th and 99.5th percentiles within each UPC.

Figure 1: Data architecture: matched retail prices and wholesale costs.

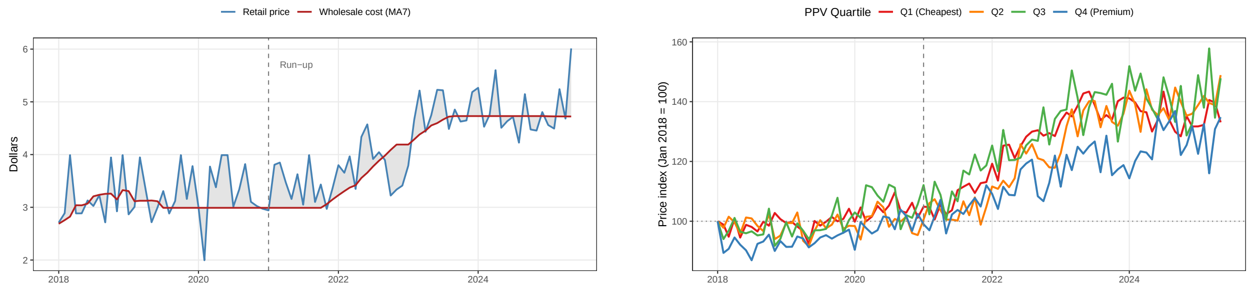

(a) Single UPC: retail price and wholesale cost

(b) All UPCs by within-category price quartile

Notes: Illustration of the matched retail–wholesale data for a single store and category (condiments), 2018–2025. Panel (a): one representative UPC showing the retail unit-value price (blue) and seven-month trailing moving-average wholesale cost (red); the vertical gap is the observed gross margin. Panel (b): all UPCs in the category, colored by within-category PPV quartile (Q1 = cheapest, Q4 = premium). The layered structure shows heterogeneous price and cost dynamics across the price ladder: cheaper products (Q1–Q2) experienced proportionally larger cost shocks during 2021–2022 than premium products (Q3–Q4). Data: matched scanner prices and invoice-level wholesale costs for over 46,000 products across 551 stores.

lion households observed in both 2019 and 2022, yielding 37.7 million household–category observations across 142 categories. Table 2 summarizes the household panel.

Table 2: Summary statistics: household panel.

Panel A: Panel construction		*Panel B: Household spending*	
Balanced households	2,847,224	Annual spend 2019, median ($)	410
Product categories	142	Annual spend 2019, mean ($)	960
HH × category obs.	37,710,629	Annual spend 2022, median ($)	452
Loyalty-card coverage (%)	96.1	Annual spend 2022, mean ($)	1,081
Categories per HH (med./mean)	10 / 13.2	Cheap-tier share, 2019 (%)	22.3
		Cheap-tier share, 2022 (%)	20.9

Notes: The balanced panel requires a household to purchase in at least three product categories with a minimum of $5 per category per period in both 2019 (pre-inflation) and 2022 (run-up). "Cheap tier" denotes products in the bottom 20% of within-category price-per-volume (PPV), assigned using 2018–2019 baseline prices so that tier assignments are fixed over time. Cheap-tier expenditure share declined from 22.3% to 20.9% between 2019 and 2022. Loyalty-card coverage reflects transaction coverage, not population coverage; the 4% of unlinked transactions may come from infrequent or non-enrolled shoppers.

3 Descriptive Evidence

Using the matched retail–wholesale panel, I now document four patterns that motivate the empirical analysis.

Pattern 1: Aggregate margin expansion. Figure 2 tracks the monthly gross margin rate from 2018 through 2025. Panel (a) shows the aggregate series: the margin was roughly 23% in the 2018–2019 baseline, peaked at 29% in a single month (early 2022), and partially retraced to 25–26% by 2024. Over the two-year run-up window (2021–2022), the sales-weighted average margin rate was 27.4%, an increase of 4.4 percentage points above baseline. Even after retracement, margins remain three percentage points above pre-pandemic levels.

Panel (b) decomposes by price tier. Cheap-tier margins expanded by 4.9 pp during the run-up, nearly double the 2.5 pp for premium.[9] The gap opened gradually over 2020–2022, consistent with accumulation through repeated category reviews rather than a discrete repricing event.

Figure 2: Monthly gross margin rate: aggregate and by price tier, 2018–2025.

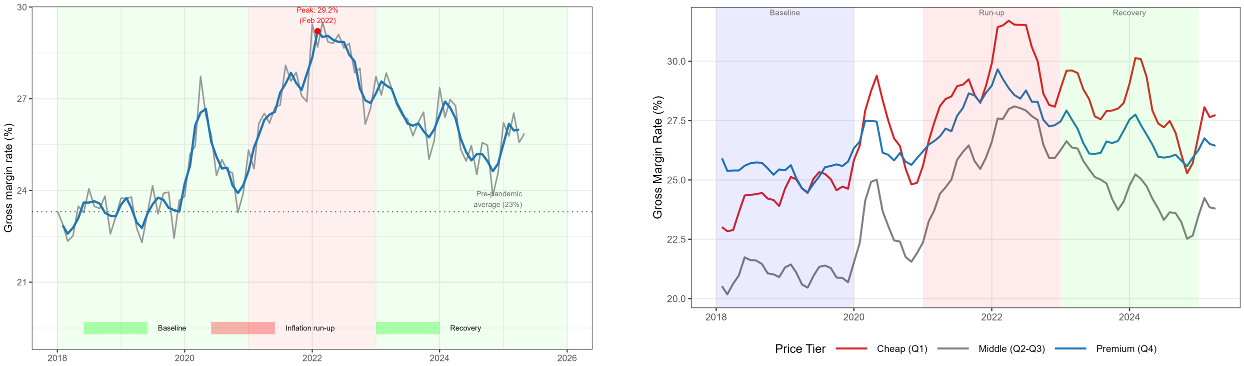

(a) Aggregate

(b) By within-category price tier

Notes: Panel (a): the margin rate is $MR_t = \sum_j w_{jt}(P_{jt} - C_{jt}^{wh})/P_{jt}/\sum_j w_{jt}$, where weights are dollar sales. Blue line: three-month moving average; grey: monthly values. Shaded regions: pre-inflation (2018–2020), run-up (2021–2022), recovery (2023+). Dotted line: pre-pandemic average (23%). Panel (b): products assigned to tiers based on 2018–2019 average price within each category: cheap (bottom 20%, red), middle (interquartile range, grey), premium (top quartile, blue). Lines show three-month moving averages. Cheap-tier margins expanded by 4.9 pp above baseline versus 2.5 for premium. The sample is restricted to products observed in both baseline and comparison months.

Pattern 2: Cross-category heterogeneity. The margin expansion was not confined to a few categories. Plotting the change in gross margin rate against wholesale cost growth across all 142 categories reveals substantial heterogeneity: some categories with modest cost growth saw large margin gains, while others with large cost shocks saw margins compress (Online Appendix Figure A1). This heterogeneity motivates the within-category analysis that follows.

---

[9] The top quartile pools Q4 and Q5 from the structural counterfactual; the Q5 figure alone is 1.2 pp (Online Appendix Table K1).

Pattern 3: Cheap-tier incidence. The most striking feature of the data is *where* the margin wedge loaded within categories. Table 3 reports margin changes by within-category price-per-volume (PPV) decile, comparing 2018 to 2024, spanning the full sample period. The gradient is steep: the cheapest decile saw a margin increase of 5.8 percentage points, roughly double the 2.8 percentage points for the most expensive decile, though it is not perfectly monotonic across the middle deciles (in particular decile 6, where category composition shifts across tiers compress the gradient).

Table 3: Margin changes by within-category price tier, 2018 to 2024 ($N \approx 47{,}000$ UPCs).

PPV Decile	MR 2018 (%)	MR 2022 (%)	MR 2024 (%)	ΔMR (pp)	$\Delta \ln C^{wh}$	$\Delta \ln P$
1 (Cheapest)	12.6	**20.5**	18.4	**+5.78**	+14.2%	+21.8%
2	12.3	19.5	15.8	+3.45	+13.8%	+18.9%
3	12.0	19.4	15.7	+3.70	+14.1%	+19.5%
4	14.2	21.7	17.7	+3.51	+13.5%	+18.7%
5	16.3	23.1	20.9	+4.53	+12.9%	+19.4%
6	21.3	26.0	21.9	+0.61	+11.8%	+12.6%
7	21.7	27.6	24.8	+3.08	+11.2%	+15.6%
8	23.8	28.9	25.9	+2.15	+10.5%	+13.5%
9	28.0	31.6	29.5	+1.49	+9.8%	+11.9%
10 (Premium)	29.7	34.5	32.5	+2.83	+8.4%	+12.1%

Notes: Observation: UPC (product). Products are assigned to within-category price-per-volume (PPV) deciles based on their 2018–2019 average PPV; decile 1 products have a median retail price of approximately \$2, while decile 10 products have a median of approximately \$10. The margin rate is $(P - C^{wh})/P$. Columns report the 2018, 2022 (peak), and 2024 sales-weighted margin rates, the 2018–2024 change in percentage points (ΔMR), and cumulative log changes in wholesale costs ($\Delta \ln C^{wh}$) and retail prices ($\Delta \ln P$). All statistics are sales-weighted within each decile across 142 product categories and 551 stores. A formal test of the implementation-frictions mechanism (Proposition 4), which predicts that margin loading appears in percentage terms but is attenuated in dollar terms, is reported in Section 6.2.

Four features stand out. First, margins peaked in 2022 and partially reversed, but the cheap tier remains 5.8 pp above its 2018 level. Second, wholesale costs rose *more* for cheap products (14.2% vs. 8.4%); if costs were the sole driver, cheap margins should have compressed. Third, retail prices rose even more for cheap products (+21.8% vs. +12.1%), with the price–cost gap wider at the cheap end. Fourth, cheap products start with lower baseline margins, so the same pp increase is proportionally larger.

Why do cheap products gain disproportionately more margin despite facing larger cost increases? Three forces could generate this: *implementation frictions* (category-wide repricing that does not fully differentiate across the price ladder loads larger margin wedges on cheaper items, an effect compounded by the bigger cost shocks cheap products face), *trade-down* (consumers shift toward cheap products, raising their marginal profitability), and *differential product turnover*: new entrants carry current-period margins while exiting products carry pre-inflation ones. The cheap-tier gradient is robust among continuing products (Section 6.2), though turnover amplifies the cross-sectional magnitude.

Pattern 4: Within-category synchronization. When at least one product in a category experiences a cost shock, the synchronization rate among *unshocked* products averages 51.2%, versus 33.5% in unshocked categories in the same store and month.[10] Because the comparison is within store and month, store-level demand or aggregate trends cannot explain the gap.

Taken together, the four patterns point to category-level repricing as the transmission mechanism. The next section formalizes this and generates additional predictions, including two not apparent from the descriptive evidence: spillovers should exhibit a closeness gradient, and cheap-tier loading should appear in log-price changes but not dollar changes if pricing frictions dominate.

4 A Model of Multiproduct Pricing with Category Reviews

I develop a reduced-form model of multiproduct retail pricing that rationalizes the descriptive patterns and generates testable predictions. The model is an organizing framework, not a structural estimation exercise. It has three ingredients: (i) a multiproduct retailer that prices categories as systems, (ii) within-category substitution, and (iii) category-level review costs that make repricing infrequent. I build the model sequentially and extract predictions at each step.[11]

Consider a category with J products. The retailer chooses prices $p_t = (p_{1t}, \dots, p_{Jt})$ to maximize:

$$\Pi_t(p_t) = \sum_{j=1}^{J} (p_{jt} - c_{jt})\, q_{jt}(p_t; \theta_t), \tag{3}$$

where demand q_{jt} depends on the *entire* price vector (reflecting within-category substitution) and an exogenous demand state $\theta_t \geq 0$ that captures trade-down (consumers shifting toward cheaper products when θ_t rises; formally defined in Online Appendix C). The first-order condition for product k equates the direct margin loss from a price increase with the

---

[10] These rates are from a within-store-month paired comparison that pools all periods; Table G4 reports period-specific unpaired averages that differ slightly.

[11] Key terms: *Category-level repricing*: coordinated adjustment of multiple products' prices during a category review triggered by a cost change to one or more products. *Synchronization*: the share of products in a store–category–month that reprice. *Spillover*: the price response of unshocked products to a neighbor's cost shock. *Cheap-tier*: bottom 20% of within-category price-per-volume.

diversion gain across all substitutes $j \neq k$. Because the retailer internalizes within-category substitution, the optimal price for each product depends on costs and demand conditions across the entire category. This yields the first prediction:

Proposition 1 (Spillover pass-through)**.** *If there exist substitutes* $i \neq j$ *with* $\partial q_{jt}/\partial p_{it} \neq 0$*, then a cost change for SKU* i *generically changes the optimal price for other SKUs* $j \neq i$*, even holding their own costs fixed.*

Intuitively: if a cheap cereal's cost rises, the retailer raises its price, diverting consumers toward the premium brand. Because the retailer internalizes this substitution, it also raises the premium cereal's price, even though its cost has not changed. The spillover is larger for close substitutes (e.g., two brands of the same type) than for distant ones (e.g., oatmeal vs. granola), generating a testable *closeness gradient*.

In practice, retailers do not continuously re-optimize prices. I introduce a fixed cost $K > 0$ of conducting a category review: the costly process of examining margins, costs, and competitor prices for an entire category (Levy et al., 1997; Alvarez and Lippi, 2014). Economies of scope in this cost generate within-category synchronization (Midrigan, 2011; Bhattarai and Schoenle, 2014). The retailer reviews a category only when the gain from re-optimization exceeds K: $\Pi_t(p_t^*) - \Pi_t(p_{t-1}) \geq K$.

Proposition 2 (Synchronization under review costs)**.** *With a category-level review cost, sufficiently large shocks to a subset of costs increase the probability of a price review. Conditional on a review, multiple prices within the category adjust together, generating broad within-category comovement.*

Once a review is triggered, many products adjust, not just the shocked product. During widespread cost increases, the threshold is breached more frequently. The review-cost framework also predicts an asymmetry:

Proposition 3 (Directional asymmetry in review triggering)**.** *Under review costs, upward cost shocks trigger category reviews more readily than downward cost shocks of the same magnitude, generating a rockets-and-feathers pattern in synchronization.*[12]

The intuition is that rising costs compress margins, creating urgency to review, while falling costs raise margins at old prices, weakening the incentive to cut—a ratchet that pushes margins upward over a cost cycle (Peltzman, 2000; Tappata, 2009). Ellingsen et al.

---

[12] Recent work identifies distinct sources of asymmetric price adjustment: learned coordination among competitors (Byrne and de Roos, 2019), algorithmic pricing (Assad et al., 2024), and menu costs with inflation (Ellingsen et al., 2026). The category-review mechanism here is complementary: it amplifies single-product asymmetry across the full shelf.

(2026) formalize this insight, showing that a monopolist facing menu costs and stochastic costs optimally adjusts prices faster upward than downward—providing a micro-foundation for the rockets-and-feathers pattern that the category-level repricing mechanism amplifies across products. Even conditional on a review, two channels can generate cheap-tier incidence. *Repricing frictions*: adjustments that do not fully scale with price levels load larger percentage changes on cheaper products (a \$0.25 increase is 12.5% on a \$2 item but only 4.2% on a \$6 one).[13] *Demand*: if consumers trade down toward cheaper products (higher θ_t), those products become more profitable to reprice upward.

Proposition 4 (Cheap-tier incidence)**.** *Two channels generate larger* cumulative *margin-rate increases for cheap-tier products across repeated category reviews. (i) Under common-increment repricing frictions (where price adjustments do not fully scale with the price level), each category review loads a larger percentage price change and margin-rate increase on cheaper items (Online Appendix C). (ii) Under trade-down, increased demand for cheap products raises their shadow value and amplifies repricing. The diagnostic signature differs: common-increment frictions predict cheap-tier loading in log-price changes but not dollar changes; trade-down predicts loading in both.*

These channels are complementary. I use the logs-versus-cents distinction to evaluate their relative importance in Section 6.2.

A final diagnostic exploits the relationship between demand elasticity and margin changes. Under standard single-product pricing, the Lerner condition $m = 1/|\eta|$ governs margin levels, and under the empirically typical case of log-concave demand, margins *fall* after cost increases (Weyl and Fabinger, 2013). More generally, the curvature of demand—Weyl and Fabinger's curvature parameter ρ—determines whether pass-through exceeds or falls short of unity for a single product. Category-level repricing short-circuits this mechanism because the review adjusts many prices simultaneously without fully differentiating by product-level demand curvature.

Proposition 5 (Inverted elasticity gradient)**.** *Under category-level repricing, own-price elasticity is not a sufficient statistic for pass-through because it depends on the full demand system: cross-elasticities, category scope, and the review-cost threshold. An inverted gradient (more elastic products gaining more margin) is inconsistent with independent-pricing benchmarks under any common demand specification.*

---

[13] The Online Appendix (Section C) derives this under the polar case of common-dollar repricing. The logs-vs-cents diagnostic (Table 5) confirms that cheap-tier loading appears in log-price changes and not in dollar changes, consistent with pricing frictions that do not fully differentiate across the price ladder.

Online Appendix C provides formal derivations and maps each proposition to its empirical test (Table C1). The six hypotheses jointly are difficult to generate from alternative models; the central causal claim rests on the spillover test (Proposition 1).

5 Empirical Strategy

The model generates testable predictions; this section describes how I take them to data. The central empirical challenge is isolating the causal effect of a cost shock to one product on the prices of other products in the same category. The analysis operates at two distinct observation levels, each suited to different questions:

- *Event groups* (one observation per store×category×month with an isolated shock; $N \approx$ 212,640). This level tests whether and how much spillovers exist via dose-response variation in shock *size*.
- *Store× UPC×month panel* ($N \approx 79$ million for the neighbor sample). This level (the main specification) traces the *timing* of price responses, controlling directly for each neighbor's own wholesale cost.

The two designs address the correlated-cost threat differently: the group-level estimands exclude neighbors with concurrent cost changes; the panel event study absorbs correlated cost pass-through through an explicit cost control. Because both comparisons are within store–category–month, store-wide demand shocks and macroeconomic conditions cancel out. Agreement across the two designs strengthens the causal interpretation. Long-difference regressions (one observation per UPC, $N = 32{,}971$) test distributional incidence; a synchronization analysis (Online Appendix D.1) confirms that the micro-level spillovers aggregate into broad within-category repricing.

Identification is strongest for Propositions 1 (spillovers) and 3 (directional asymmetry), both identified from idiosyncratic within-category shocks. Proposition 2 (synchronization) is tested in Online Appendix D.1 using the same variation but is more exposed to serial persistence. Proposition 4 (cheap-tier incidence) provides distributional implications that combine cross-sectional evidence with quasi-experimental support, and the elasticity diagnostic provides additional discrimination. Online Appendix Table J3 maps each prediction to its identifying variation and primary threat.

The two designs use different event definitions. For *group-level estimands*, an event is a store–category–month in which at least one product experiences a wholesale cost change of $\geq$2%, with temporal isolation (±6 months), yielding 212,640 events. For the *panel event study*, an event is the first temporally isolated cost change of $\geq$10% for a store–UPC pair; the stricter threshold ensures large, discrete shocks suitable for tracing dynamic price paths.

The identifying assumption, that wholesale cost variation reflects upstream supply factors (commodity prices, packaging, manufacturer decisions) rather than local demand conditions, is stated formally below (equation 6) and tested in the threats discussion. Temporal isolation ensures I study discrete events; excluding neighbors with concurrent cost changes (group-level) or controlling for each neighbor's own wholesale cost (panel) purges correlated cost pass-through.

Group-level dose-response. The first design asks: does a *larger* cost shock cause a *larger* neighbor price response? This dose-response test uses only cross-event variation in shock size, requiring no timing assumptions. For each event, I construct four estimands:

1. **Own pass-through** (Y_g^{own}): cumulative log price change of the shocked product over the post-event horizon ($H = 6$).
2. **Neighbor spillover** (Y_g^{nb}): average log price change of unshocked neighbors, the core test of category-level repricing.
3. **Cheap-tier differential** ($Y_g^{cheapdiff}$): difference in neighbor response between cheap and non-cheap products.
4. **Closeness gradient** ($Y_g^{closediff}$): difference in response between close substitutes (same product module within the event category) and distant ones (different module within the same category).

Each estimand is estimated from a simple regression on the event sample:

$$Y_g^m = a_m + \lambda_{\text{period}(g)} + b_m \, Shock_g + u_{gm}, \tag{4}$$

with standard errors clustered at the event-group level. Each observation is one event group g, a store–category–month triplet. The regression pools 212,640 such groups. $Shock_g$ is the focal product's wholesale cost change (in logs); the coefficient b_m is a unitless pass-through elasticity (percentage-point price response per unit of log cost shock). I include period fixed effects (run-up vs. recovery) to absorb regime shifts in the average level of repricing activity across periods; identification remains cross-event variation in shock size within the event sample.

Panel event study. The group-level estimands test *whether* spillovers exist via dose-response, but they cannot trace the *timing* of the response or control for neighbors' own cost movements at the product level. The panel event study addresses both. It operates at the store×UPC×month level (roughly 79 million observations for the full neighbor sample before balancing; the balanced estimation sample for Figure 3 contains approximately 11.5 million observations; see Online Appendix Table A3 for the full observation-count mapping) and traces the dynamic path of each product's price around the shock month. Crucially, the neighbor specification includes each neighbor's own wholesale cost as a control, so the coefficients $\hat{\beta}_\tau^{nb}$ measure price changes *net of* any correlated cost pass-through.

I define a *panel event* as the first temporally isolated *persistent* wholesale cost increase of at least 10% for a given store–UPC pair.[14] Three filters ensure clean identification: *persistence* (cost must not decline below its post-shock level over six months), *isolation* (no other shock within ±6 months), and *first event* (each store–UPC pair contributes one event). The own-product specification is:

$$\ln P_{jst} = \alpha_{js} + \gamma_t + \sum_{\tau \neq -2} \beta_\tau^{own} \, \mathbf{1}[\tau_{jst} = \tau] + \varepsilon_{jst}, \tag{5}$$

where $\ln P_{jst}$ is the log retail price for UPC j in store s at month t, α_{js} are store×UPC fixed effects absorbing permanent level differences, γ_t are year-month fixed effects absorbing ambient inflation, and τ_{jst} is event time relative to the shock month (binned at ±7 to absorb outside-window variation). The reference period is $\tau = -2$. This specification contains no cost control because the cost shock *is* the treatment.

The neighbor specification adds the critical wholesale cost control:

$$\ln P_{jst} = \alpha_{js} + \gamma_t + \sum_{\tau \neq -2} \beta_\tau^{nb} \, \mathbf{1}[\tau_{jst} = \tau] + \delta \ln C_{jst}^{wh} + \varepsilon_{jst}, \tag{6}$$

where $\ln C_{jst}^{wh}$ is the neighbor's own wholesale cost. The control δ absorbs any correlated cost pass-through, so the remaining $\hat{\beta}_\tau^{nb}$ isolate *pure* pricing spillovers: price changes in neighbor products at each event time, net of their own cost movements.

I estimate this on two samples. The *full neighbor sample* includes all non-shocked products and controls for each neighbor's own cost. The *clean neighbor sample* restricts to products that never experience a ≥10% cost change (products whose costs are stable throughout), so any price response must reflect pure pricing spillovers. Each neighbor UPC is assigned to one event (the first balanced event in its store–category); standard errors are clustered at the store×category level. Results are robust to 5% and 20% thresholds (Online Appendix F.1).

---

[14] The synchronization analysis uses a lower 5% threshold; Online Appendix Table D1 reports results across a 4×4 threshold grid.

Distributional incidence. Throughout the paper, I use several tier definitions: deciles for descriptive gradients, a binary bottom-20% indicator ("cheap tier") for regressions, quintiles for the counterfactual, and a three-tier split (bottom 20%, middle, top quartile) for time-series figures. Results are robust across all definitions.

The causal analysis establishes that spillovers exist, aggregate broadly, and ratchet upward. The remaining question is: where does the resulting margin expansion land across the price ladder? The event-study framework is not suited for this question because cheap-tier loading accumulates slowly through many small adjustments, not through differential per-event responses. A cross-sectional long-difference approach directly asks: which products gained most margin between baseline and run-up? Each observation is a UPC j in category c ($N =$ 32,971 UPCs observed in both baseline and run-up). I estimate:

$$\Delta MR_{jc} = \alpha_c + \beta \cdot Cheap_{jc} + \gamma \, \Delta \ln C^{wh}_{jc} + \varepsilon_{jc}, \tag{7}$$

where ΔMR_{jc} is the change in margin rate for product j in category c between baseline and run-up, $Cheap_{jc}$ indicates cheap-tier status, and α_c are category fixed effects. Table 6 additionally reports a specification that includes a $Cheap_{jc} \times \Delta \ln C^{wh}_{jc}$ interaction.

A sharper test replaces $Cheap_{jc}$ with the product's pre-period own-price elasticity (Proposition 5). Under independent product pricing (single-product Lerner), margins should rise most on *inelastic* products. An inverted gradient (more elastic products gaining more margin) rules out independent pricing and points directly to category-level repricing. I estimate:

$$\Delta MR_{jc} = \alpha + \delta \cdot z_{\beta,j} + \gamma \, \Delta \ln C^{wh}_{jc} + \varepsilon_{jc}, \tag{8}$$

where $z_{\beta,j}$ is the standardized pre-period elasticity.[15]

Two further specifications link the causal and distributional results. First, I augment the long-difference regression (equation 7) with an interaction between cheap-tier status and category-level synchronization intensity; a positive interaction coefficient means categories with more spillover-driven repricing exhibit more cheap-tier loading. Second, I interact shock size with pre-period elasticity at the neighbor-product level within the event-study framework; a negative coefficient means more elastic neighbors exhibit larger price responses, confirming the inverted gradient within the causal identification.

The primary validation is the neighbor panel's pre-period coefficients: $\hat{\beta}^{nb}_{\tau} \approx 0$ for $\tau < 0$, confirming no differential pre-trends. A placebo test confirms that future shock size does

---

[15] I estimate product-specific own-price elasticities from $\ln q_{sjt} = \alpha_{sj} + \gamma_t + \beta_j \ln p_{sjt} + \varepsilon_{sjt}$ on 2018–2019 data, with store×UPC and month fixed effects. This yields a product-specific own-price elasticity using within-store-UPC price variation. I standardize the elasticity estimates and trim outliers; Online Appendix E.1 reports details and robustness to a Hausman instrument.

not predict past price changes, and a worst-case bias bound (Rambachan and Roth, 2023) generates at most 12% bias for the single-month neighbor coefficient and less than 1% for the cumulative own pass-through estimate (Online Appendix F.1).

Threats to identification. Five concerns deserve discussion; detailed robustness is in Online Appendix F.1.

Correlated wholesale costs and staggered shocks. Products in the same category may share commodity inputs, so neighbor price responses could reflect correlated own-cost pass-through rather than spillovers. The panel event study controls for each neighbor's own cost; the contamination filter restricts to groups where $\geq$70% of neighbors have $<$5% cost changes; richer cost controls (levels, momentum, acceleration) leave the spillover unchanged (Online Appendix Table B3). Crucially, the *timing pattern* provides further identification: correlated costs predict immediate neighbor co-movement, but the data show neighbors flat at impact and drifting upward only over subsequent months. The closeness gradient supports demand-linked transmission, since close and distant substitutes often face similar commodity costs. A manufacturer-boundary decomposition (Online Appendix F.2) confirms that the spillover crosses brand boundaries: after controlling for each neighbor's cost, same-brand and different-brand neighbors converge by $\tau = +3$ ($z = 0.14$). A cross-category placebo confirms category specificity: cost shocks in *other* categories within the same store do not generate positive synchronization ($\hat{\beta} = -0.27$ pp; Online Appendix D.1), ruling out store-wide operating cost inflation as a confound.

Competitive response. Competitor price increases could drive the measured spillover. Three features argue against this: the delayed timing ($\tau = +3$, not $\tau = 0$); private-label products, with no cross-retailer equivalent, also show significant spillovers (Online Appendix F.3); and the within-store-category-month design nets out store-wide competitive pressure.

Smoothing artifacts. The seven-month trailing moving average is used only for descriptive statistics; all causal analyses use unsmoothed step-costs (raw invoice cost carried forward from the last shipment). The MA(7) modestly amplifies the descriptive cheap-tier gradient (2.85 pp vs. 2.42 pp under raw costs), but the pattern is robust across smoothing windows (Online Appendix B.1).

Invoice timing and cost measurement. The step-cost series carries forward the most recent invoice cost until a new shipment arrives, which may not align perfectly with the retailer's cost basis for pricing decisions. If the measured cost systematically lags the decision-relevant cost, the delayed neighbor drift could partly reflect cost measurement lag. Three features mitigate this: cost-dating sensitivity checks (±1 month shift of the cost series) produce stable estimates (Online Appendix B.2); the 14-to-1 directional asymmetry is inconsistent with symmetric invoice-timing noise; and unobserved cost pressures would predict immediate co-movement at $\tau = 0$, but neighbors are flat at impact and drift only over subsequent months.

External validity of isolated shocks. The event sample spans 141 of 142 categories with a tier distribution nearly identical to the full shelf. A 4×4 sensitivity grid (Online Appendix Table F4) confirms stable spillover estimates across shock thresholds and contamination cutoffs.

Unit values versus regular prices. Excluding promotional observations yields nearly identical spillover coefficients, and six-month persistence rules out temporary promotions (Online Appendix B.5). Online Appendix B.6 decomposes unit-value changes into regular-price adjustments and promotional-mix shifts using the retailer's monthly promotional calendar, confirming that the delayed neighbor drift reflects regular-price repricing rather than promotional cycling.

Demand-side confounds. If events capture category-level demand shifts rather than upstream cost shocks, neighbor quantities should trend upward before the event. A panel event study of neighbor log quantities (Online Appendix F.6) shows flat pre-trends; the two significant pre-period coefficients are negative, the opposite of the demand-confound prediction. Post-shock, neighbor quantities decline cumulatively (-0.5% by $\tau = +6$), consistent with consumer response to repricing rather than a positive demand shift. A within-brand pack-size test (Section 6.2) provides further evidence: products within the same brand line share the same demand environment, so any differential margin loading across pack sizes reflects the pricing process rather than demand shifts.

6 Results

All results address a single chain of reasoning: category-level repricing transmits cost shocks in a way that widens margins on cheap products (see Online Appendix Figure F1 for a schematic overview). I present them in four cumulative steps. First, I establish the causal core: cost shocks to one product spill over to unshocked neighbors (existence). Second, I show that synchronization and directional asymmetry aggregate this mechanism into a margin ratchet (scope). Third, I trace where the resulting expansion lands across the price ladder (incidence). Fourth, household-level evidence shows whether consumers can escape (welfare). The hypothesized mechanism operates cumulatively: spillovers aggregate into synchronization, the asymmetric ratchet contributes to margin expansion, and cheap-tier loading generates distributional consequences.

6.1 Category-Level Repricing and Spillovers

In the short run, the retailer absorbs over 90% of any individual cost shock, passing through only 8.1% to the retail price within six months (Table 4). This is the group-level dose-response elasticity, measuring cumulative six-month price impact per unit of log cost shock across 212,640 natural experiments; the panel event-study coefficient at $\tau = 0$ (+5.9%; Figure 3) captures the instantaneous price jump in a balanced panel absorbing ambient inflation and is not directly comparable.[16] But unshocked neighbors also reprice: the neighbor spillover coefficient is 0.022 (Table 4), about one-quarter of the 0.081 own-product pass-through rate, even though the neighbors' wholesale costs did not change. This is consistent with category-level repricing (Proposition 1). The result is stable across shock thresholds and contamination cutoffs (all $t > 14$; Online Appendix Table F4), and the low own-product pass-through rate is not an artifact of selecting large shocks: own pass-through is similarly modest at the 2% and 20% thresholds. Cumulative pass-through over repeated reviews is substantially higher (Section 6.2).

Table 4: Event-level estimand regressions: own pass-through, neighbor spillover, and closeness gradient.

Estimand	Coeff.	Std. Err.	t-stat	p-value	N groups	95% CI
Y^{own} (Own pass-through)	**0.081**	0.003	27.75	<0.001	212,640	[0.075, 0.086]
Y^{nb} (Neighbor spillover)	**0.022**	0.001	18.08	<0.001	210,249	[0.020, 0.025]
$Y^{cheapdiff}$ (Cheap-tier loading)	0.000	0.001	0.22	0.826	192,441	[−0.002, 0.002]
$Y^{closediff}$ (Closeness gradient)	**0.003**	0.001	3.45	<0.001	175,229	[0.001, 0.005]

Notes: Regression: $Y_g^m = a_m + \lambda_{\text{period}(g)} + b_m\, Shock_g + u_{gm}$ (equation 4; Proposition 1). *Observation:* event group g = (store, category, month); N = 175,229–212,640 depending on estimand. *RHS:* $Shock_g = \Delta \ln C^{wh}_{i_g,t}$. *Dependent variables:* Y^{own} = log price change over a 6-month horizon (post minus pre average) of the shocked SKU; Y^{nb} = sales-weighted average log price change of unshocked neighbors; $Y^{cheapdiff}$ = difference in neighbor response between cheap-tier and non-cheap-tier products; $Y^{closediff}$ = difference between close substitutes (same product module) and distant substitutes (same category, different module). Events span the full sample period (2018–2025) with period fixed effects for run-up (2021–2022) and recovery (2023–2024). Standard errors clustered at the event-group level.

The closeness gradient is small (0.003, an additional 14% on top of the 0.022 average) but precisely estimated: if spillovers were purely mechanical, close and distant substitutes would respond identically.[17] The cheap-tier differential is zero, meaning cheap-tier loading accumulates through repeated reviews rather than per-event differentials.[18]

---

[16] The apparent tension with aggregate pass-through exceeding unity (retail prices rose 17% while wholesale costs rose 12%) reflects cumulative compounding across repeated category reviews, augmented by neighbor spillovers and the rockets-and-feathers ratchet.

[17] National-brand neighbors drive the spillover; private-label neighbors respond negatively, consistent with the retailer lowering PL prices to capture trade-down.

[18] Excluding promotional observations or restricting to low-promotion UPCs yields nearly identical spillover

Panel event study. The group-level estimands test whether spillovers exist; the panel event study (Figure 3) traces their *dynamics* at the store×UPC×month level using persistent cost shocks and two complementary neighbor samples (equations 5–6).

Panel (a) shows own-product pass-through: the log price jumps sharply at $\tau = 0$ (+5.9%) and remains elevated through $\tau = +6$ (+4.8%). The gradual post-impact decline reflects the year-month fixed effect absorbing ambient inflation rather than cost mean-reversion.[19]

Panels (b) and (c), the critical tests, examine neighbor price responses. Panel (b) uses the *full neighbor sample* (controlling for each neighbor's own cost): prices are flat at impact but drift upward, reaching +0.74% by $\tau = +3$ ($t = 5.6$). Panel (c) uses the *clean neighbor sample* (products that never experience a $\geq$10% cost change): the pattern is similar, rising to +0.91% ($t = 5.7$).[20] The convergence of both samples (one absorbing correlated costs via regression control, the other eliminating them by sample restriction) establishes that the delayed response reflects genuine pricing spillovers. Alternative thresholds (5% and 20%) yield similar patterns (Online Appendix F.1).

Ruling out correlated upstream costs. A manufacturer-boundary decomposition (Online Appendix Figure F3) confirms that the delayed neighbor spillover crosses brand boundaries: same-brand neighbors show a larger immediate response (correlated upstream costs), but after controlling for each neighbor's own wholesale cost the gap vanishes by $\tau = +3$ ($z = 0.14$), establishing retailer-initiated repricing as the operative channel (details in the threats discussion above).

Timing: trigger-then-review. The timing pattern (own prices jumping at $\tau = 0$ while neighbors drift upward over 1–3 months) is consistent with a two-step process. The own-product adjustment is a narrow, product-specific response (updating a shelf tag to reflect a new invoice cost) that occurs immediately. The broader category review, examining margins, competitor prices, and promotional calendars across the shelf, is triggered by the cost shock but requires additional time to execute, generating the 1–3 month delay visible in Figure 3. Monthly aggregation further compresses timing: if the cost shock arrives mid-month and the review adjusts neighbor prices late in the same month, neighbor adjustments may straddle the month boundary and load partly onto $\tau = +1$. This "trigger-then-review" interpretation

---

coefficients (Online Appendix Table B4); the six-month persistence rules out temporary promotions.

[19] The +5.9% is the panel event-study coefficient at a single horizon; the group-level 8.1% (Table 4) measures cumulative pass-through over six months in a different sample; see that table's footnote for a full reconciliation.

[20] Exact coefficients: full sample $\hat{\beta}_0^{nb} = +0.16\%$ ($t = 1.80$), $\hat{\delta} = 0.24$ ($t = 27.3$); clean sample $\hat{\beta}_0^{nb} = -0.15\%$ ($t = -1.25$).

is consistent with industry practice, where category reviews are initiated by cost-change alerts but executed on a review calendar (Levy et al., 1997; Zbaracki et al., 2004).[21]

Time variation. Spillovers were amplified during the inflationary period (Figure 4): larger cost shocks generate larger neighbor responses, and spillovers were dramatically larger during the run-up than recovery. Net of ambient inflation, run-up spillovers are 2–2.5% for positive shocks, versus near-zero during recovery.[22]

Synchronization and directional asymmetry (Propositions 2 and 3). The micro-level spillovers aggregate into broad within-category repricing. A synchronization analysis (Online Appendix D.1) confirms that the presence of cost shocks in a category raises the probability that unshocked neighbors reprice, with a dose–response pattern, threshold nonlinearity, and chain-wide dominance consistent with centralized category reviews.

Critically, this repricing is directionally asymmetric. Upward cost shocks drive a 1.51 pp increase in neighbors' probability of raising prices ($t = 26.4$), while downward shocks produce only 0.11 pp ($t = 2.0$), a 14-to-1 asymmetry that ratchets margins upward over the cost cycle.[23] The asymmetry is regime-specific (Figure 5): in 2018, up and down shocks were roughly symmetric; in 2021, upward shocks drove +2.44 pp while downward shocks *reduced* repricing (−0.53 pp); by 2023, the pattern reversed as accumulated cost decreases breached the review threshold (+1.38 pp for downward shocks, zero for upward). Online Appendix Figure F6 confirms this within the panel event-study framework.

6.2 Distributional Incidence

Category-level repricing with directional asymmetry provides a mechanism for margin expansion. The remaining question is where this expansion lands across the price ladder.

Implementation frictions: the logs-vs-cents diagnostic (Proposition 4). Table 5 provides indirect evidence on the implementation-frictions mechanism. The prediction is not that all products receive identical dollar adjustments unconditionally (premium products naturally have larger absolute price changes because they have higher base prices and face larger

---

[21] The delay cannot reflect inventory turnover: most grocery products turn over in one to three weeks, far shorter than the one-to-three-month repricing lag. Moreover, the clean neighbor sample (products whose costs never change) shows the same delayed drift, ruling out inventory depletion as the mechanism.

[22] Table 4 reports a regression slope ($\hat{\beta} = 0.022$), while Figure 4 reports conditional means including ambient inflation; the dose-response pattern is consistent across both.

[23] Roughly one-third of the synchronization coefficient reflects serial persistence (Online Appendix Table D4); the directional asymmetry is robust to serial-correlation purging.

Figure 3: Panel event study: own vs. neighbor price response to persistent cost shocks.

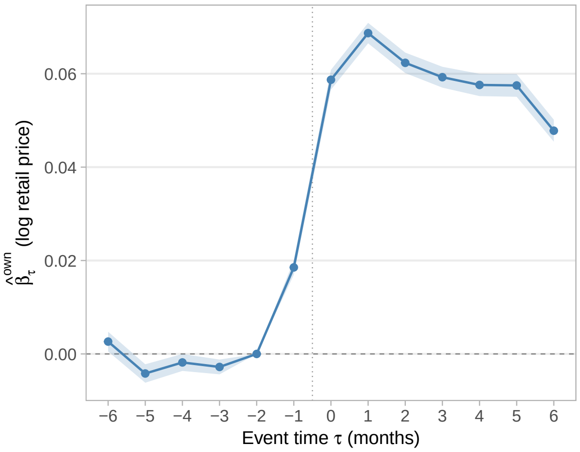

(a) Own product

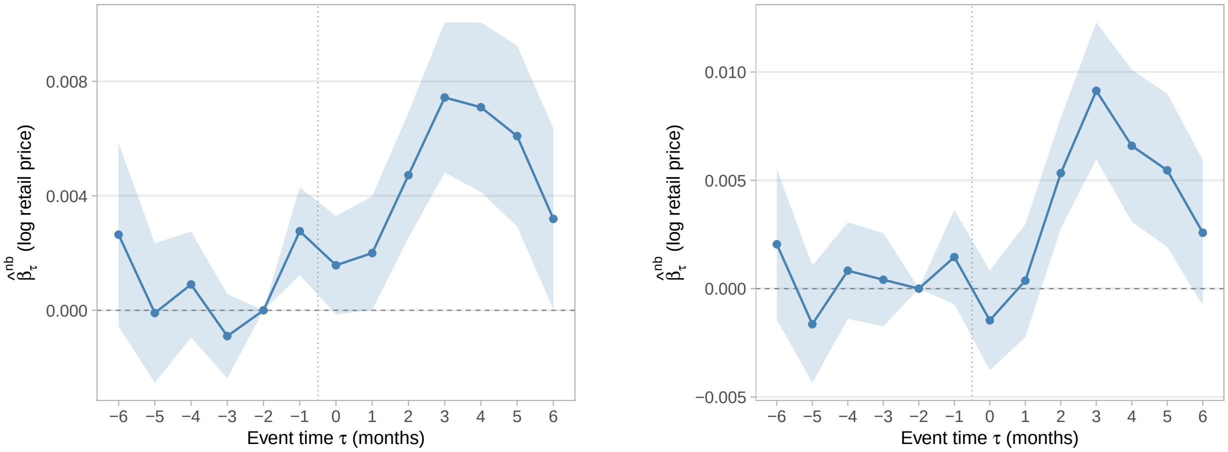

(b) Neighbor products: full sample (cost-controlled)

(c) Neighbor products: cost-stable sample

Notes: Coefficients $\hat{\beta}_\tau$ from two-way fixed-effects regressions at the store×UPC×month level. Events are defined as the first persistent, isolated wholesale cost increase of at least 10% for a given store–UPC pair: the cost must not decline below its post-shock level and must not rise more than one additional percentage point over the following six months, with no other shock of this magnitude within ±6 months. Panel (a) estimates equation (5) on products experiencing the cost shock (57,777 balanced UPCs, 3.83 million observations). Panel (b) estimates equation (6) on the *full neighbor sample* (all non-shocked products in the same store–category), controlling for each neighbor's own wholesale cost ($\hat{\delta} = 0.24$, $t = 27.3$; 194,432 balanced UPCs, 11.5 million observations). Panel (c) estimates equation (5) on the *clean neighbor sample*, products that never experience a wholesale cost change exceeding 10% at any point in the sample (85,631 balanced UPCs, 4.1 million observations); no cost control is needed because these products' costs are stable by construction. Both neighbor samples show flat pre-trends (0 of 4 pre-event coefficients significant) and a delayed upward drift reaching approximately +0.7–0.9% by $\tau = +3$, consistent with delayed category reviews rather than staggered cost arrivals. Events span 2018–2025. Store×UPC and year-month fixed effects; reference period: $\tau = -2$; standard errors clustered at the store×category level. Shaded bands: 95% confidence intervals.

Figure 4: Neighbor price response by shock size bin and period.

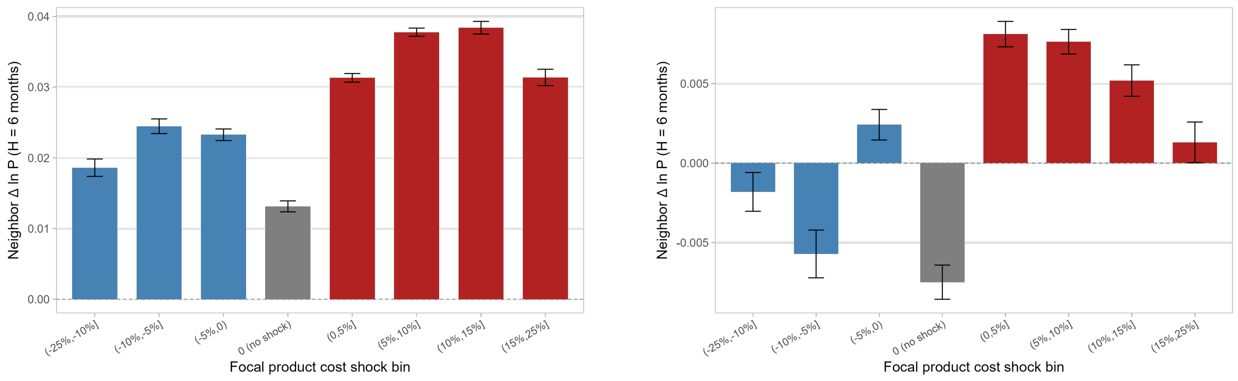

(a) Extended run-up window (2019–2022)

(b) Recovery (2023–2025)

Notes: Each bar shows the conditional mean of the log price change over a 6-month horizon ($\Delta \ln P^{H=6}$, post-event minus pre-event average) for products with $|\Delta \ln C| < 2\%$ in the same store–category–month. Observation: store–category–month group $g = (s, c, t)$, including months with no cost shock; bars are sales-weighted averages within each shock-size bin. Panel (a): extended run-up window (2019–2022); Panel (b): recovery (2023–2025). The run-up gradient (1.3% for no-shock months rising to 3.8% for large positive shocks) is substantially steeper than the recovery gradient. Error bars show 95% confidence intervals. Corresponds to Online Appendix Table F2.

Figure 5: Directional asymmetry in neighbor repricing by year.

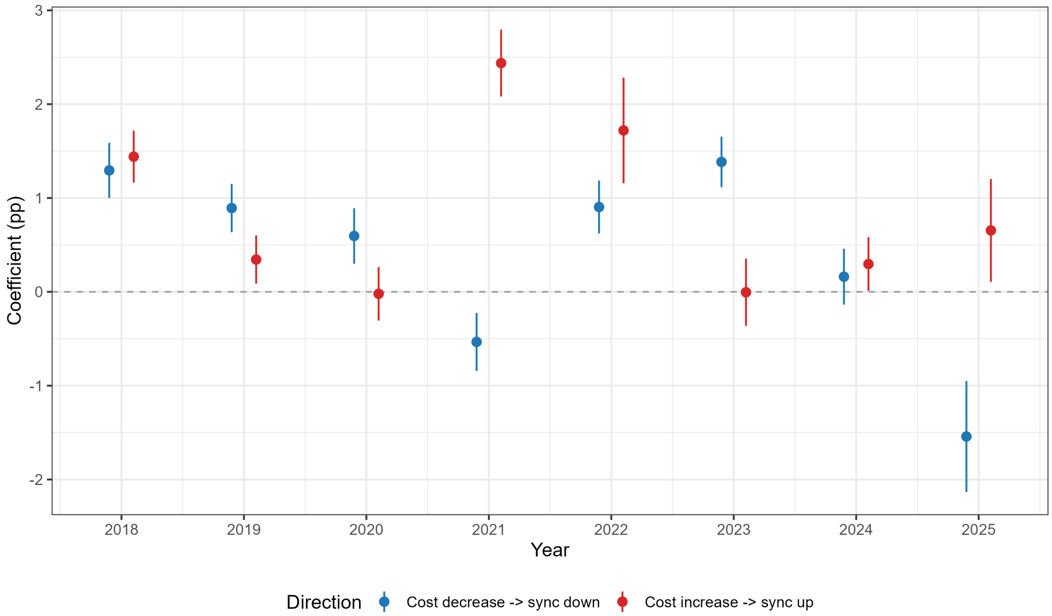

Notes: Each bar shows the coefficient from a synchronization regression of the share of unshocked neighbors with an upward (red) or downward (blue) price change on an indicator for at least one product with a $\geq$5% cost increase or decrease, estimated separately by year. Observation: store–category–month ($N \approx 467{,}000$). Fixed effects: store×category, year-month. Standard errors clustered at store×category. In 2021 (inflation onset), upward shocks dominate; in 2023 (recovery), the pattern reverses. Whiskers show 95% confidence intervals. Online Appendix D.1 reports the full specification and robustness.

absolute cost movements) but rather that, *conditional on a category review*, adjustments are more uniform in dollar terms than in percentage terms. The unconditional means confirm the predicted divergence between percentage and dollar gradients: cheap and premium products show comparable percentage-change magnitudes (+0.27% vs. +0.77%/month), but the dollar-change gap is dramatically wider (−0.56 vs. +11.51 cents/month), driven by the five-fold difference in base prices.[24] A histogram of individual price changes (Online Appendix Figure G1) makes the pattern vivid: percentage-change distributions are nearly identical across tiers (median: 9.4% vs. 9.1%), while dollar-change distributions differ dramatically (median: 20 vs. 75 cents).

The sharpest test of the friction mechanism is the within-brand pack-size test, which compares different sizes of the *same product* within the same brand (e.g., 12-oz vs. 24-oz cereal), controlling for Brand×Module fixed effects. The cheaper pack size gained 1.65 pp more margin (Online Appendix Table G1), accounting for 93% of the price-tier gradient. This test directly rules out demand-side explanations: products within the same brand line share the same manufacturer, the same shelf placement, and the same demand environment, so any differential demand shift (trade-down, elasticity changes, or compositional reallocation) would affect both pack sizes equally. The margin differential can only reflect the mechanical interaction between roughly uniform repricing adjustments and heterogeneous base prices. The recovery period shows the same pattern in reverse.

Cheap-tier incidence (Proposition 4). Table 6 reports that cheap-tier products gained a statistically significant 3.5 pp ($t = 4.49$) more margin than non-cheap products, controlling for cost changes and category fixed effects. The insignificant interaction with cost growth means this is a level shift, not differential pass-through, consistent with the friction mechanism: common-increment repricing frictions predict a *level* effect (cheap products gain more margin per review regardless of cost-shock magnitude), not differential pass-through scaling with cost growth.

The result is robust unweighted (+1.62, $t = 4.01$), sales-weighted (+2.10, $t = 3.72$), and strengthens with a baseline margin-rate control (+1.97, $t = 5.20$; Online Appendix Table G2).[25] The within-brand pack-size result (−1.65 pp; see the implementation-frictions

---

[24] The negative mean for the cheapest decile reflects within-decile composition: items range from roughly \$0.50 to \$2.50, and a few large dollar decreases on higher-priced items within the bin pull the mean negative. The median is positive (+\$0.03/month), confirming this is a composition artifact, not a data error.

[25] The estimand regression (Table 4) finds no cheap-tier loading per event ($Y^{cheapdiff} = 0.000$, $t = 0.22$); the cross-sectional +3.5 pp accumulates through two channels. First, cheap products face more and larger cost shocks, triggering more reviews; repeated rounds of repricing cumulate to +1.0 pp on continuing products (Online Appendix Table G2). Second, product turnover accounts for roughly half the gap: new cheap-tier entrants carry current-period margins. Categories with more synchronization exhibit more cheap-tier loading

Table 5: Monthly price changes in percentage and dollar terms by price tier and review status.

PPV Decile	Base Price ($)	$\Delta \ln P$ (%/mo)	ΔP (cents/mo)	Period
1 (Cheapest)	1.92	+0.27	−0.56	Run-up (2021–22)
5	5.47	−0.17	−1.07	Run-up (2021–22)
10 (Premium)	15.07	+0.77	+11.51	Run-up (2021–22)
1 (Cheapest)	1.87	−1.15	−3.63	Post (2023–24)
5	5.66	−0.70	−3.13	Post (2023–24)
10 (Premium)	15.26	−0.51	−5.50	Post (2023–24)

Notes: Observation: UPC–store–month, averaged to PPV decile level via two-stage sales-weighted aggregation (first within category×decile, then across categories). Each row reports the sales-weighted average monthly log price change ($\Delta \ln P$, percent) and average monthly absolute price change (ΔP, cents) for products in the indicated PPV decile. Base Price is the average pre-period retail price in dollars. Compositional artifacts cannot drive this pattern: cheap products are actually *less* likely to reprice in any individual review (−1.21 pp; Online Appendix D.1), so the extensive margin works against the finding, not in its favor. The cheapest decile's opposite signs ($\Delta \ln P > 0$ but $\Delta P < 0$) reflect within-decile price heterogeneity: items range from roughly $0.50 to $2.50, so a few large dollar decreases on the higher-priced items within this tile pull the cents mean negative while their moderate log decreases leave the log mean positive. Medians are positive in both metrics (+1.6%, +3.0 cents), confirming the mean divergence is driven by composition, not a coding artifact. Under implementation frictions (Proposition 4), the log ratio (cheap-to-premium) should be closer to parity than the cents ratio; under trade-down (Proposition 4), both should show cheap-tier loading. The sample covers ∼47,000 UPCs across 551 stores.

discussion above) confirms that the gradient operates within brand lines, ruling out demand-side trade-down (Online Appendix G.1).

The inverted elasticity gradient (Proposition 5). The most demand-sensitive products gained 6.2 pp of margin versus 0.5 pp for the least sensitive (Online Appendix Table E3), the opposite of what independent product pricing predicts. Under standard Lerner pricing, margins should rise most on *inelastic* products. Table 7 confirms the inversion: a one-standard-deviation increase in elasticity is associated with a 1.66 pp ($t = -3.51$ with category clustering) *larger* margin increase.

The result survives price-decile controls, category fixed effects, and a Hausman IV specification (Online Appendix E.2). Because elasticities are estimated on 2018–2019 data and related to *subsequent* margin changes, measurement error attenuates rather than inflates the gradient. A pass-through diagnostic confirms: products with low pass-through rates saw the largest margin expansion (Online Appendix Table E4). A demand-shift explanation would predict the same sign, but the gradient is driven primarily by the price level (5:1 variation in dollar semi-elasticity vs. 1.2:1 in percentage elasticity), so demand rotation would need

---

(Online Appendix Table D8).

Table 6: Long-difference margin regressions on cheap-tier status and wholesale cost changes.

	Estimate	Std. Err.	t-stat	p-value
$Cheap_{jc}$	$\mathbf{+3.51}$	0.78	4.49	<0.001
$\Delta \ln C_{jc}^{wh}$	-26.58	4.21	-6.31	<0.001
$Cheap_{jc} \times \Delta \ln C_{jc}^{wh}$	-3.14	3.99	-0.79	0.433
Category FE	Yes			
N UPCs	32,971			
R^2 (within)	0.14			

Notes: Regression: $\Delta MR_{jc} = \alpha_c + \beta\, Cheap_{jc} + \gamma\, \Delta \ln C_{jc}^{wh} + \delta\, Cheap_{jc} \times \Delta \ln C_{jc}^{wh} + \varepsilon_{jc}$, augmenting equation (7) with the interaction term $Cheap_{jc} \times \Delta \ln C_{jc}^{wh}$. *Observation:* UPC j within category c. *Dependent variable:* change in gross margin rate (ΔMR_{jc}, percentage points, 2018 to 2022). *RHS:* $Cheap_{jc}$ = indicator for bottom 20% of within-category PPV; $\Delta \ln C_{jc}^{wh}$ = cumulative log wholesale cost change. *Fixed effects:* category (α_c). Standard errors clustered at the category level (N = 32,971 UPCs across 142 categories). The insignificant interaction ($\hat{\delta}$) indicates that the cheap-tier margin premium is a level shift rather than differential pass-through.

to be implausibly large. The within-brand pack-size test (Online Appendix G.1) further disciplines this concern.

Table 7: Margin changes regressed on pre-period own-price elasticity.

Variable	Estimate	Std. Error	t-stat	p-value
Intercept	5.61	0.068	82.1	<0.001
z_β (std. elasticity)	$\mathbf{-1.66}$	0.069	-24.1	<0.001
$\Delta \ln C^{wh}$	$\mathbf{-12.56}$	0.362	-34.7	<0.001
Sample	UPCs with $\beta \leq 0$, trimmed 1–99%			
N	39,387			
R^2	0.05			

Notes: Regression: equation (8), $\Delta MR_{jc} = \alpha + \delta\, z_{\beta,j} + \gamma\, \Delta \ln C_{jc}^{wh} + \varepsilon_{jc}$ (elasticity diagnostic). *Observation:* UPC j (N = 39,387). *Dependent variable:* ΔMR_{jc} (percentage points, 2019 to 2022). *RHS:* z_β = standardized pre-period own-price elasticity estimated via log-log demand regressions (see footnote 15) with store×UPC and month FE on 2018–2019 data (more negative = more elastic); $\Delta \ln C_{jc}^{wh}$ = cumulative log wholesale cost change. Standard errors are heteroskedasticity-robust. The negative $\hat{\delta}$ means more elastic products saw larger margin increases, the opposite of Lerner pricing. With category FE and category-clustered SEs: $z_\beta = -1.59$ ($t = -3.51$); with price-decile FE and category clustering: $z_\beta = -1.81$ ($t = -4.09$).

A causal test within the event-study framework confirms the gradient (Online Appendix Table F6). The inversion reflects a mechanical correlation between the price level and dollar semi-elasticity $\alpha \propto |\eta|/p$: $|\eta|$ varies only 1.2-to-1 while α varies 5-to-1, so the cost gradient dominates (Online Appendix Figure E1).

6.3 Putting It Together: Magnitudes

The category-repricing mechanism generated \$154–320 million in additional gross margin above a constant-markup benchmark during the two-year run-up, on \$7.2 billion of total sales. The dollar-sales-weighted margin rate rose from approximately 23% to 27.4% (peaking near 29% in early 2022), an increase of 4.4 percentage points. The bottom-up figure sums product-level margin changes across matched UPCs; the top-down imputation (0.044×\$7.2B) includes unmatched products.[26] Both figures represent gross margin (before accounting for rising operating costs such as labor, energy, and shrinkage) and should not be interpreted as additional profit.

A back-of-the-envelope illustration anchors the micro estimates. Consider a cereal category with 30 products, one experiencing a 10% cost increase. The shocked product's price rises by 0.81% (0.081 × 10%). But the cost shock also triggers a category review, causing seven unshocked neighbors (the median number of affected products per event) to reprice by 0.22% each (0.022 × 10%). The cumulative neighbor repricing (7 × 0.22% = 1.54 percentage points of product-level price changes) exceeds the own-product response by a factor of 1.9, so neighbor repricing accounts for roughly two-thirds of total within-category repricing activity. The per-event own pass-through rate of 8.1% appears low relative to long-run estimates (Nakamura and Zerom, 2010), but cumulative pass-through exceeds unity: retail prices rose 17% while wholesale costs rose 12% (Table 1), consistent with partial adjustments compounding through repeated reviews.[27]

Micro-to-macro decomposition. Five channels contribute to the aggregate expansion: (i) own pass-through (8% of each cost shock transmitted to the retail price); (ii) spillovers, accounting for 56–74% of within-category price impact (Online Appendix Table H1); (iii) the rockets-and-feathers ratchet (14:1 asymmetry during the run-up); (iv) accumulation through repeated reviews (no per-event cheap-tier loading, but +3.5 pp cross-sectionally); and (v) product turnover, accounting for roughly half the cheap-tier loading. A Shapley decomposition attributes the margin expansion primarily to retail price increases outpacing cost growth; the wholesale cost contribution is negative (cost increases partially offset the

---

[26] The gap between the two figures reflects product turnover: entering and exiting UPCs that cannot be matched across periods account for 14% of run-up revenue, and the top-down calculation includes their contribution while the bottom-up sum does not. The 4.4 pp increase uses the dollar-sales-weighted margin rate over the two-year run-up window (2021–2022), weighting both across products and across months by dollar sales. The 3.8 pp difference implied by Table 1 (26.6% minus 22.8%) uses sales-weighted aggregation within each month but equal-weights across months within the period; the gap between 3.8 and 4.4 pp reflects the additional cross-month weighting.

[27] The panel event study's $\tau = 0$ coefficient ($\approx 5.9\%$) is lower than the group-level 8.1% because it captures instantaneous impact absorbing ambient inflation; Online Appendix Table H4 reconciles the estimands.

margin expansion), so the retail-price contribution exceeds 100% of the net change (Online Appendix Table G5). This decomposition is accounting, not structural; the components interact and the spillover share is a local estimate from isolated shocks.

Among continuing UPCs, the bottom quintile accounts for 16.5% of run-up sales and 16.9% of markup expansion (roughly proportional); the premium tier contributed only 18.6% from 24.7% of sales. Product turnover amplifies the pattern: new cheap-tier entrants carry higher margins (Online Appendix Table G3). For a consumer spending \$200/month, the margin-driven price increase is roughly \$12/year larger for budget-oriented than premium-oriented baskets.

Spillover share: definition and aggregation. The spillover share measures how much of within-category repricing activity is attributable to unshocked neighbors rather than the directly shocked product. I define it through two complementary approaches. The *bridge based on local average treatment effects (LATEs)* computes the per-event share as $\hat{b}_{nb}\bar{N}/(\hat{b}_{own}+\hat{b}_{nb}\bar{N})$, where $\hat{b}_{own} = 0.081$ is own pass-through, $\hat{b}_{nb} = 0.022$ is the per-neighbor spillover, and $\bar{N}$ is the median number of unshocked neighbors. At $\bar{N} = 7$, the share is 66%; the plausible range across $\bar{N} \in [5, 10]$ and ±1 SE on $\hat{b}_{nb}$ is 56–74% (Online Appendix Table H1). Because these LATEs come from temporally isolated shocks, the per-neighbor response may be overstated during high-density periods: the per-product spillover is concave in shock density (Online Appendix Table H3), though total category-level spillover continues to rise. A complementary *accounting decomposition* partitions each store–category–month into shocked and unshocked products: during the run-up, unshocked products account for 54.5% of sales-weighted price increases (Online Appendix Table H2). The accounting share is lower than the LATE estimate because it captures all repricing of unshocked products, including lagged cost pass-through. Both approaches bracket the spillover share in the range of roughly half to three-quarters, which should be interpreted as an order-of-magnitude characterization (Online Appendix H.2).

6.4 Household-Level Evidence: Product Switching and Margin Exposure

I compute *margin exposure* (spending-weighted average margin change) for 2.85 million households observed in both 2019 and 2022. *Counterfactual exposure* holds the 2019 basket fixed; *actual exposure* uses the 2022 basket.

Margin exposure is large and positive for all household types (Figure 6): 5.0–5.6 pp counterfactual, 4.7–5.0 pp actual. Overall product switching, including cross-category substitution and within-tier adjustments, modestly reduces exposure (-0.17 to -0.62 pp), consistent with households partially adjusting away from the most-affected products. However, this net relief masks a countervailing within-category channel: cost shocks predict within-category trade-down toward cheaper products ($\hat{\beta} = -0.10$, $p < 0.001$), and such trade-down *increases* margin exposure because consumers substitute toward exactly the products with the largest margin increases. Trade-down mediates roughly half the cost-shock–to–margin-exposure relationship (Online Appendix I).

Figure 6: Margin exposure by pre-period cheap-tier spending quintile.

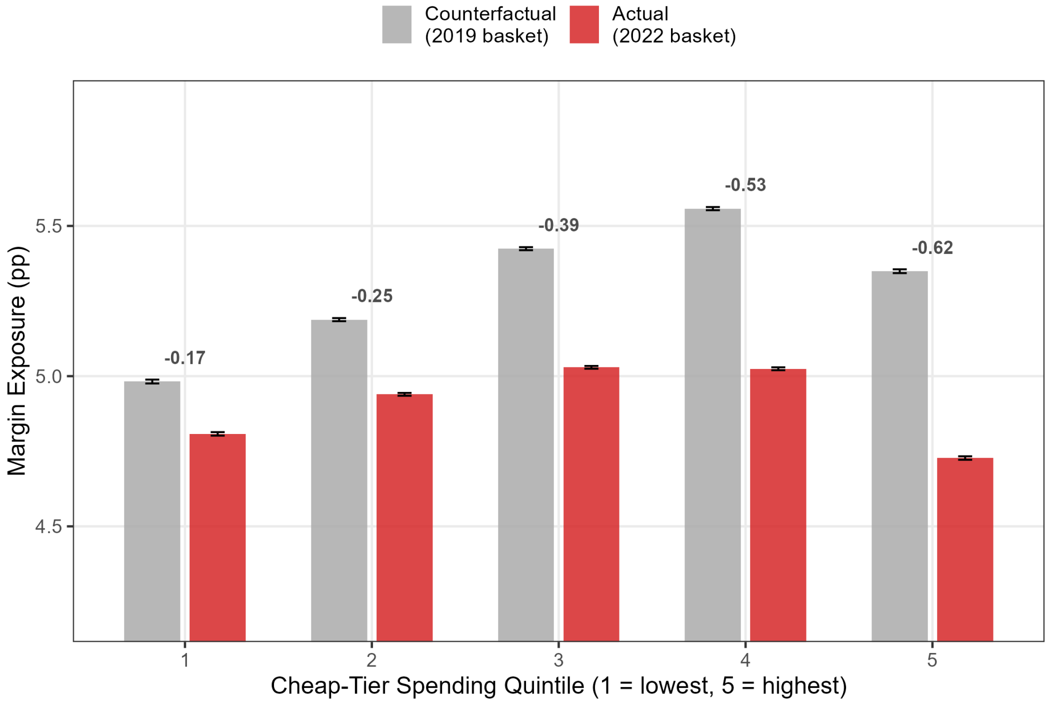

Notes: Observation: household, averaged to quintile (N = 2.85 million balanced households observed in both 2019 and 2022 across 142 categories). Q1 = households spending least on cheap-tier products in 2019; Q5 = most. Margin exposure is the spending-weighted average ΔMR (2019–2022, percentage points) of products purchased. Grey bars: counterfactual exposure (holding the 2019 product basket fixed); red bars: actual exposure (using the 2022 product basket). Numbers above each pair: reallocation effect (actual minus counterfactual). All quintiles face 4.7–5.6 pp of margin expansion; product switching reduces exposure only modestly (-0.17 to -0.62 pp). Quintile assignment is based on the household's 2019 share of spending on bottom-20% PPV products within each category. Confidence intervals are omitted because these are population means over 2.85 million households (standard errors are negligible).

6.5 Counterfactual: Category vs. Product-Level Pricing

The logs-vs-cents diagnostic and within-brand test establish that pricing frictions are the dominant channel for cheap-tier incidence. How much of the regressive outcome is attributable to the category-pricing constraint? I answer with two gross-margin-neutral counterfactuals in the spirit of DellaVigna and Gentzkow (2019): holding total category gross margin fixed, I reallocate markups using (i) a Lerner rule based on percentage elasticities and (ii) a Ramsey rule based on semi-elasticities, and compare the resulting gradients to the data.

Approach. The counterfactuals require demand elasticities for each price quintile. While the standard *percentage* elasticity varies modestly across the price spectrum ($|\hat{\eta}|$ ranges from 1.92 for Q1 to 2.35 for Q5; Table 8, Panel A), the *semi-elasticity* α_q—which measures how demand responds to a one-*dollar* price change—varies dramatically because $\alpha_q = \lambda_c \cdot |\eta_q|/[p_q(1-s_q)]$, where s_q is the market share of quintile q and λ_c is a category-specific scaling factor (median 4.3; Appendix K). The proportionality $\alpha_q \propto |\eta_q|/p_q$ means that the same dollar price increase represents a much larger percentage shock for a cheap product than for a premium one. After calibrating a multinomial logit demand system for each of 139 product categories (three categories are dropped for insufficient product coverage; Appendix K), the median semi-elasticity ranges from 4.6 for Q1 (cheapest) to 0.9 for Q5 (premium), a five-to-one gradient driven primarily by the price level rather than by the elasticity gradient itself. Elasticities are estimated from OLS using within-store-UPC variation (see footnote 15). A Hausman IV specification—instrumenting each store's price with the leave-one-out mean across other stores (first-stage $F > 10^5$; Appendix E.2)—yields elasticities approximately 40% larger in magnitude but highly correlated with OLS ($r > 0.7$). The Ramsey attribution is virtually unchanged: 57% under OLS-only versus 61% under IV-only (Table K1). Even setting a completely uniform elasticity across all quintiles yields 58%, confirming that the result is driven by the price-level gradient ($1/p_q$), not by heterogeneous demand parameters. Re-estimating elasticities on 2021–2022 data yields a Spearman rank correlation of 0.28 with the baseline estimates ($N = 35{,}299$ UPCs in both periods; Online Appendix E.1), suggesting that the broad rank ordering is partially preserved despite level shifts. The results should therefore be interpreted as approximate but robust to the elasticity source.

Given the demand parameters, both counterfactuals are closed-form. For each category c in month t, I hold total observed profit $M_{ct} = \sum_q m_q Q_q$ fixed and reallocate markups using two rules:

$$\text{Ramsey:} \quad m_q^{\mathrm{R}} = \frac{\kappa_{ct}^{R}}{\alpha_q}, \qquad \text{Lerner:} \quad m_q^{\mathrm{L}} = \frac{\kappa_{ct}^{L} \cdot p_q}{|\eta_q|}, \tag{9}$$

where $m_q = p_q - c_q$ is the observed dollar markup, Q_q is observed units sold, and the scaling constants κ_{ct}^{R} and κ_{ct}^{L} are chosen so that $\sum_q m_q^{\mathrm{new}} Q_q = M_{ct}$ by construction (gross-margin-neutral).[28] Quantities are held fixed at observed levels (a first-order approximation).

The Lerner rule sets percentage margin rates proportional to $1/|\eta_q|$, the standard inverse-elasticity pricing condition (Lerner, 1934). Because $|\eta|$ varies only 1.2-to-1 across quintiles, this rule preserves most of the observed margin-rate gradient: cheap products (lower $|\eta|$) *should* have higher percentage margins even under product-level optimization, so only the excess regressivity beyond what elasticities justify is attributed to the pricing constraint. The Ramsey rule operates on dollar markups via the semi-elasticity $\alpha_q \propto |\eta_q|/p_q$, which additionally scales markups with the price level. The Ramsey rule therefore removes both the elasticity-based and the price-level components of the gradient, providing an upper bound on the constraint attribution.

The gross-margin-neutral constraint eliminates the revealed-preference concern: the retailer earns the same total gross margin under all three allocations, so the question is how the *distribution* differs.

Results. Table 8 reports the main results. Panel A shows the demand parameters. The percentage elasticity gradient across quintiles is modest ($|\eta|$ varies 1.2-to-1), while the semi-elasticity gradient is steep (α varies 5-to-1) because a dollar change looms much larger for a \$2 product than for a \$10 one.

Panel B reports margin rate changes from baseline (2018–2019) under observed category pricing and under both gross-margin-neutral counterfactuals. The Lerner counterfactual—the more conservative benchmark—narrows the Q1–Q5 gap from +3.7 percentage points to +2.1 during the run-up, a 44% reduction. The Ramsey counterfactual narrows it further to +1.5, a 61% reduction. During recovery (2023–2024), the reductions are 38% (Lerner) and 80% (Ramsey).

Figure 7 visualizes the result. The Lerner counterfactual (light bars) compresses the gradient partially—Q1's expansion falls from 4.9 to 4.6 pp while Q5's rises from 1.1 to 2.5 pp—but a substantial gap remains because some regressivity is optimal under product-level pricing (cheap products have lower $|\eta|$ and thus higher optimal margin rates). The Ramsey counterfactual (medium bars) compresses more aggressively, nearly equalizing the quintiles.

Dynamic path. Figure 8 traces the Q1–Q5 gap month by month. The observed gap rises sharply during 2021, peaks above +5 pp, and remains elevated. Both counterfactual gaps track below, with the Lerner gap closer to observed and the Ramsey gap closer to zero. The

---

[28] Explicitly: $\kappa^R_{ct} = M_{ct}/\left(\sum_k Q_k/\alpha_k\right)$ and $\kappa^L_{ct} = M_{ct}/\sum_k p_k Q_k/|\eta_k|$. The κ notation is distinct from the contamination threshold in Appendix D.2.

Table 8: Gross-margin-neutral counterfactual: margin changes by price quintile under observed vs. counterfactual pricing.

	Q1 (Cheapest)	Q2	Q3	Q4	Q5 (Premium)	Q1–Q5 gap
Panel A: Demand parameters						
Elasticity $\lvert\hat{\eta}\rvert$	1.92	2.00	2.20	2.22	2.35	
Semi-elasticity α_q	4.58	2.99	1.99	1.42	0.92	
Lerner weight $1/\lvert\eta_q\rvert$	0.52	0.50	0.45	0.45	0.43	
Ramsey weight $1/\alpha_q$	0.22	0.33	0.50	0.70	1.09	
Panel B: Margin rate change from baseline (pp)						
Run-up (2021–2022)						
Observed (category)	+4.9	+3.6	+4.3	+5.3	+1.1	+**3.7**
Counterfactual (Lerner)	+4.6	+4.2	+4.0	+4.8	+2.5	+**2.1**
Counterfactual (Ramsey)	+4.2	+4.2	+4.3	+4.4	+2.7	+**1.5**
Recovery (2023–2024)						
Observed (category)	+4.1	+1.4	+2.9	+3.0	+0.3	+**3.8**
Counterfactual (Lerner)	+3.7	+2.0	+2.7	+3.0	+1.4	+**2.4**
Counterfactual (Ramsey)	+2.8	+1.9	+3.0	+3.1	+2.1	+**0.7**
Panel C: Regressivity attribution						
Run-up						
Lerner (percentage elasticity)						44%
Ramsey (semi-elasticity, upper bound)						61%
Recovery						
Lerner (percentage elasticity)						38%
Ramsey (semi-elasticity, upper bound)						80%

Notes: Two gross-margin-neutral counterfactual reallocations of markups across within-category price quintiles ($N = 139$ product categories). *Panel A:* Demand parameters calibrated from a multinomial logit model (Appendix K). Elasticities $\lvert\hat{\eta}\rvert$ are sales-weighted means of OLS UPC-level estimates on 2018–2019 data (see footnote 15; Appendix E.1); a Hausman IV check yields similar results (attribution: 57% OLS-only vs. 61% IV-only; 58% under uniform elasticities). Lerner weights $1/\lvert\eta_q\rvert$ govern the Lerner allocation of percentage margin rates; Ramsey weights $1/\alpha_q$ govern the Ramsey allocation of dollar markups. *Panel B:* Margin rate = $(p-c)/p$. "Observed" rows use actual prices. "Lerner" rows set percentage margin rates $\propto 1/\lvert\eta_q\rvert$; "Ramsey" rows set dollar markups $\propto 1/\alpha_q$; both hold total category gross margin fixed. Statistics are sales-weighted means across categories. The Q1–Q5 gap measures the difference in margin-rate expansion; positive values indicate regressive patterns. *Panel C:* Attribution = 1 − (counterfactual gap / observed gap). The Lerner counterfactual uses percentage elasticities only; its attribution (44%) reflects the excess regressivity beyond what product-level inverse-elasticity pricing would predict. The Ramsey counterfactual additionally removes the price-level scaling embedded in the semi-elasticity; its higher attribution (61%) is an upper bound. The outside good share s_0 (the fraction of potential demand captured by products outside the category) is set at 0.30; results are robust to $s_0 \in [0.15, 0.50]$ (Table K1). Attributions in Panel C are computed from unrounded values and may not match hand calculations from the rounded figures in Panel B.

Figure 7: Margin rate change from baseline: observed vs. counterfactual pricing by quintile.

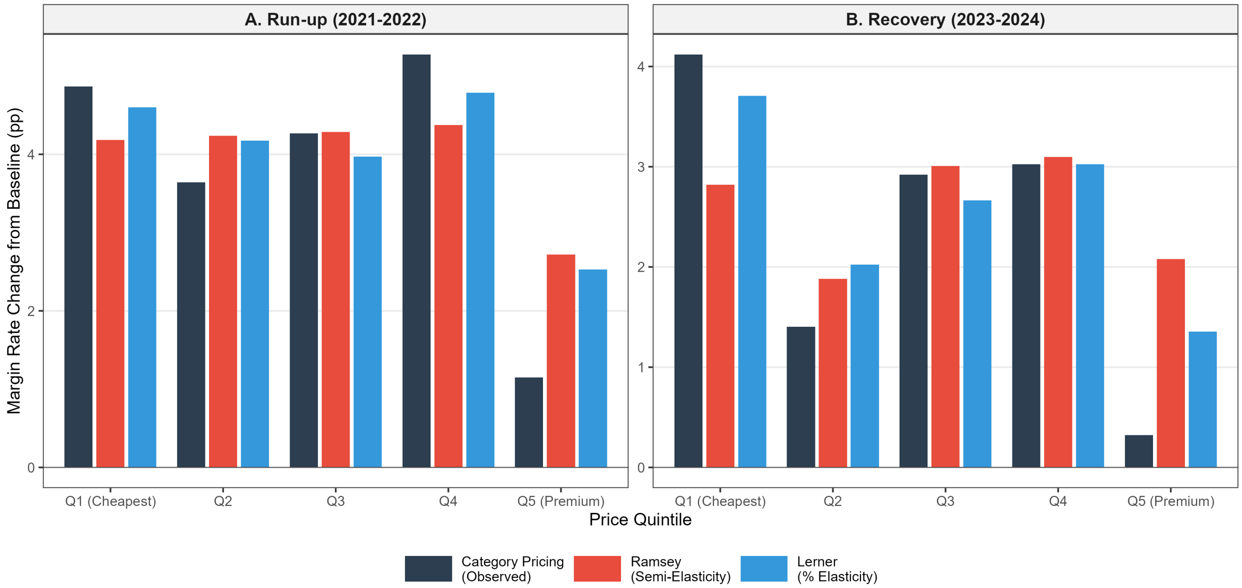

Ramsey: m_q = kappa / alpha_q (semi-elasticity rule). Lerner: MR_q = kappa' / |eta_q| (percentage elasticity rule). Both profit-neutral.

Notes: Dark bars: observed category pricing. Light bars: Lerner counterfactual (percentage margin rates $\propto 1/|\eta_q|$). Medium bars: Ramsey counterfactual (dollar markups $\propto 1/\alpha_q$). Panel A: inflationary run-up (2021–2022). Panel B: recovery (2023–2024). Both counterfactuals hold total category gross margin fixed. The Lerner counterfactual narrows the Q1–Q5 gap by 44% (run-up) and 38% (recovery); the Ramsey by 61% and 80%. The gap between the two attributions reflects the price-level scaling in $\alpha_q \propto |\eta_q|/p_q$.

two counterfactuals diverge noticeably during the recovery: the Lerner gap tracks close to observed (attribution falls to 38%), while the Ramsey gap drops toward zero (attribution rises to 80%). This divergence reflects the asymmetric return to baseline across the price spectrum. Premium-tier margins nearly revert to pre-inflation levels (+0.3 pp), while cheap-tier margins remain elevated (+4.1 pp). The Ramsey rule—which weights dollar markups by $1/\alpha_q \propto p_q$—assigns a large share of total markup to premium products, so when their observed margins are low, the Ramsey correction on both ends is large. The Lerner rule operates only on percentage elasticities, which vary modestly, and therefore preserves most of the observed gradient regardless of the price level.[29]

Figure 8: Monthly Q1–Q5 margin-rate gap under observed category pricing vs. gross-margin-neutral counterfactuals.

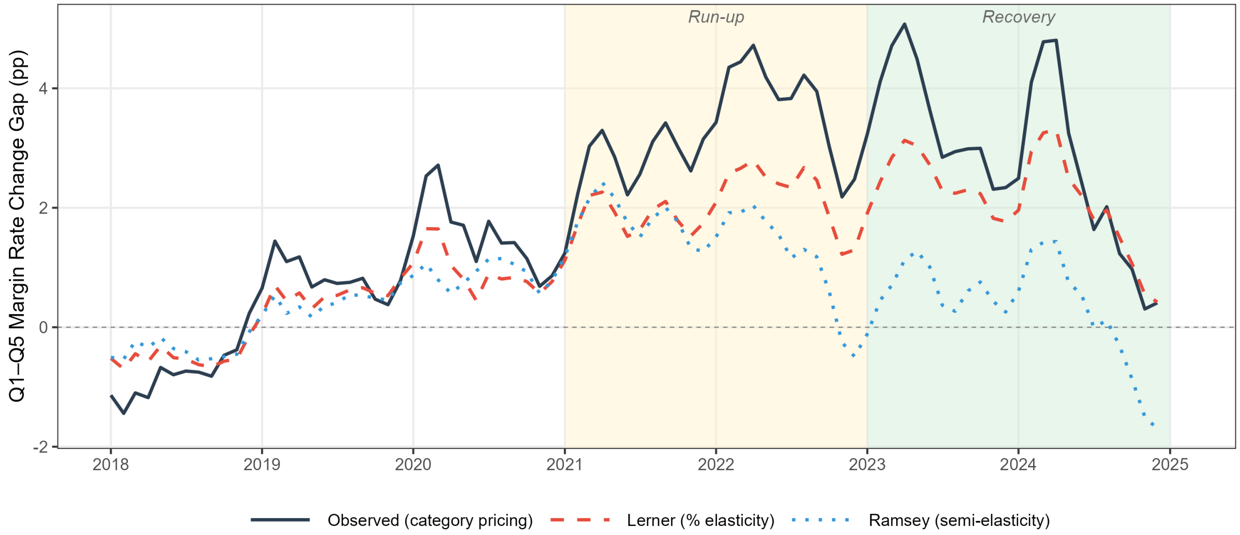

Notes: Solid line: Q1–Q5 gap in margin-rate changes from baseline (2018–2019) under observed category pricing. Dashed red line: same gap under the gross-margin-neutral Lerner counterfactual (percentage margin rates $\propto 1/|\eta_q|$). Dotted blue line: same gap under the gross-margin-neutral Ramsey counterfactual (dollar markups $\propto 1/\alpha_q$). All series are seasonally adjusted (subtracting the 2018–2019 monthly pattern) and smoothed with a three-month centered moving average. Both counterfactuals hold total category gross margin fixed in each month (equation 9). Shaded regions indicate the inflationary run-up (2021–2022) and recovery (2023–2024). The observed gap exceeds both counterfactuals throughout the inflation episode. The Lerner gap tracks closer to observed because it preserves the modest elasticity gradient; the Ramsey gap is compressed further by additionally removing the price-level scaling. Full sample: 139 product categories. Confidence bands are not shown because all three series are deterministic given the calibrated elasticities; uncertainty in α_q is addressed via sensitivity analysis (Table K1).

Interpretation. The two counterfactuals bracket the attribution. The Lerner rule attributes 44% of the run-up regressivity to the category constraint: even under optimal

---

[29] Monthly attributions differ from period averages due to equal vs. sales-weighted aggregation.

product-level pricing using percentage elasticities, 44% of the Q1–Q5 gap would disappear. This estimate is conservative because it recognizes that some regressivity is consistent with the modest elasticity gradient (cheap products face lower $|\eta|$ and thus warrant higher percentage margins under the Lerner condition). The Ramsey rule attributes 61%, reflecting the additional price-level scaling—because dollar adjustments are roughly uniform rather than proportional to price, they generate larger percentage-margin increases for cheaper products. This additional 17 percentage points captures the arithmetic channel documented in the logs-vs-cents diagnostic (Table 5).

The exercise is best interpreted as a calibrated accounting decomposition rather than a prediction of equilibrium prices. Three caveats apply. First, holding quantities fixed is a first-order approximation that likely overstates the scope for reallocation, since demand responses would partially offset the markup compression. However, an iterated equilibrium that allows quantities to adjust via $\Delta \ln q = \eta \cdot \Delta \ln p$ confirms the correction is small (+2.4 pp; quantity adjustments below 0.4% per quintile; Appendix K). Second, the Ramsey attribution ranges from 57–61% across elasticity sources and is 58% even under uniform elasticities (Table K1). Third, the counterfactual captures both organizational and competitive constraints without distinguishing between them.[30]

7 Discussion

The results provide strong evidence that pricing architecture, the organizational process by which retailers set prices, shapes the distributional incidence of inflation. Several implications follow.

Product-level pass-through is not a sufficient statistic. The standard approach in pass-through analysis treats each product as an independent pricing unit (Weyl and Fabinger, 2013; Nakamura and Zerom, 2010). The results show this misses the dominant channel: category-level repricing generates cross-product spillovers that account for roughly half to three-quarters of within-category price impact. Own-product pass-through captures only about one-third of total price adjustment once neighbor responses are aggregated. This finding echoes the theoretical prediction that multiproduct firms' pass-through depends on the full product portfolio (Armstrong and Vickers, 2023), and provides large-scale empirical confirmation using 212,640 natural upstream cost shocks.

---

[30] This connects to DellaVigna and Gentzkow (2019), who study uniform pricing across *stores*. My exercise operates within stores, across the quality spectrum.

Pricing frictions generate regressive incidence without changed conduct. The public debate framed inflation-era margin expansion as a conduct question, asking whether firms exploited market power (Weber and Wasner, 2023; Conlon, 2026). This paper identifies a distinct, mechanical channel: category-wide repricing with imperfectly differentiated adjustments loads larger percentage price changes on cheaper products, because the same absolute adjustment represents a larger proportional change for a $2 product than a $10 one. The within-brand pack-size test isolates this supply-side friction from demand-side explanations, and the inverted elasticity gradient (more elastic products gaining more margin) directly contradicts independent product-level optimization (Lerner, 1934). These findings complement Mongey and Waugh (2025), who propose a demand-side channel through heterogeneous elasticities; here, the supply-side pricing process generates regressive outcomes even with modest elasticity differences, without requiring differential demand elasticities as the primary driver. The gross-margin-neutral counterfactual attributes 44–61% of the cheap–premium margin gap to category-level pricing constraints (Section 6.5). Since low-income households allocate a larger budget share to cheap-tier grocery items (Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019), the incidence is likely regressive, though the data lack household demographics to quantify the welfare effect directly.

Directional asymmetry ratchets margins upward. Across the sample, upward cost shocks trigger category reviews 14 times more readily than downward shocks. This asymmetry is concentrated in the run-up and is not a fixed institutional feature: it reverses during recovery as accumulated cost decreases breach the review threshold. The regime-dependent pattern extends the rockets-and-feathers literature (Peltzman, 2000; Borenstein, Cameron, and Gilbert, 1997; Ellingsen et al., 2026) by showing that asymmetric pass-through can operate *within* a firm's pricing process, not only across market stages. Ellingsen et al. (2026) derive this asymmetry from a menu-cost model with moderate inflation; my results confirm the prediction at scale and reveal that the asymmetry is state-dependent rather than structural. The mechanism here differs from inter-firm channels: Byrne and de Roos (2019) attribute asymmetric adjustment to learned coordination among competitors, and Assad et al. (2024) trace it to algorithmic pricing adoption. Category-level repricing generates asymmetry *within* a single firm's pricing process, without requiring algorithmic sophistication or tacit coordination. Household-level trade-down provides limited relief: the net reallocation effect on margin exposure ranges from -0.17 to -0.62 pp, small relative to the 5+ pp margin expansion on cheap-tier products.

Heterogeneity across store formats. The spillover mechanism requires within-category density. Full supermarkets (551 stores, ∼15,000 UPCs) show strong dose-response; express/convenience stores (138 stores, ∼1,500 UPCs) show null effects, consistent with their thin assortment lacking the product density needed for category-level repricing (Online Appendix A.1). This format dependence suggests that the mechanism is most relevant for large-format grocery, which accounts for the majority of U.S. food-at-home expenditure.

Centralized pricing. A direct test of whether spillovers operate at the store or chain level finds no store-level amplification: when the same product receives a cost shock in some stores but not others, neighbor price changes are 0.11 percentage points *smaller* in shocked stores ($t = -3.75$; 227,994 differential events). The category review is centralized, consistent with evidence that U.S. grocery chains coordinate pricing at the chain or zone level rather than store by store (Hitsch, Hortacsu, and Lin, 2021; DellaVigna and Gentzkow, 2019) (Online Appendix J.1).

Policy implications. Three policy-relevant implications emerge. First, regulatory focus on firm-level markups may miss the within-category redistribution that drives distributional incidence; the relevant margin is not the aggregate but the product-level allocation. Second, SNAP benefit levels are tied to the Thrifty Food Plan (U.S. Department of Agriculture, 2021), a low-cost food basket that overweights precisely the cheap-tier products whose margins expanded most. If the mechanism documented here operates similarly at other retailers, and the cross-country evidence of cheapflation across 91 retailers in 10 countries is consistent (Cavallo and Kryvtsov, 2024), then SNAP purchasing power erodes faster than headline CPI food-at-home would suggest.

8 Conclusion

Pricing architecture shapes the distributional incidence of inflation. Category-level repricing transmits upstream costs in a way that product-level pass-through models miss. The core finding is causal: unshocked neighbors' prices drift upward over subsequent months while remaining flat at impact. A calibration exercise suggests that neighbor spillovers account for roughly half to three-quarters of within-category price impact, and the resulting margin expansion is regressive: 3.5 pp more on cheap-tier products, with 44–61% attributable to the category-level pricing constraint in a calibrated accounting exercise (Section 6.5).

For measurement, the results imply that price indices constructed from retail scanner data may understate the distributional burden of inflation episodes: households that concentrate spending on cheap products face a systematically larger margin-driven price increase than headline food-at-home CPI suggests. Transfer programs indexed to average food baskets—notably SNAP—may therefore undercompensate precisely the households most exposed to category-level repricing.

Three questions remain open. First, how do operational decisions (shelf-space allocation, promotion timing, private-label sourcing) interact with category reviews to amplify or dampen the cheap-tier gradient? Second, does competition across retailers attenuate spillovers—if rivals do not follow a category review, the initiating retailer may lose price-sensitive shoppers—or reinforce them through parallel category management practices? Third, can the mechanism be detected in real time, using wholesale cost data or review-frequency signals, to improve the responsiveness of transfer-program indexation to distributional price pressures?

References

Bureau of Labor Statistics. 2024. "Consumer Price Index for All Urban Consumers: Food at Home in U.S. City Average." Series CUSR0000SAF11, retrieved from FRED, Federal Reserve Bank of St. Louis.

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.

Blanchard, Olivier J. and Ben S. Bernanke. 2023. "What Caused the US Pandemic-Era Inflation?" NBER Working Paper No. 31417.

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

Conlon, Christopher T. 2026. "Did Profits Cause Inflation?" *International Journal of Industrial Organization*, 105: 103257. https://doi.org/10.1016/j.ijindorg.2026.103257.

Dopper, 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.

Glover, Andrew, Jose 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.

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.

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.

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.

Hastings, Justine S. and Jesse M. Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review*, 108(12): 3493–3540.

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

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

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

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

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.

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.

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.

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

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

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

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.

Griffith, Rachel, Martin O'Connell, and Kate Smith. 2015. "Relative Prices, Consumer Preferences, and the Demand for Food." *Oxford Review of Economic Policy*, 31(1): 116–130.

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.

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

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

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

Klenow, Peter J. and Benjamin A. Malin. 2011. "Microeconomic Evidence on Price-Setting." In *Handbook of Monetary Economics*, Vol. 3A, ed. Benjamin M. Friedman and Michael Woodford, 231–284. Elsevier.

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

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

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

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.

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.

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): 791–825.

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

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.

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.

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

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

Ramsey, Frank P. 1927. “A Contribution to the Theory of Taxation.” *Economic Journal*, 37(145): 47–61.

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

Cavallo, Alberto and Oleksiy Kryvtsov. 2024. “Price Discounts and Cheapflation during the Post-Pandemic Inflation Surge.” *Journal of Monetary Economics*, 148(S1): 103644.

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

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.

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

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

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.

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

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

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

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.

U.S. Department of Agriculture. 2021. *Thrifty Food Plan, 2021.* FNS-916, Food and Nutrition Service, Washington, DC.