Markups Along the Supply Chain*

Santiago Alvarez-Blaser
Bank of Spain

Alberto Cavallo
Harvard University

Alexander MacKay
University of Virginia

Paolo Mengano
Esade Business School

October 2026
[REVISION IN PROGRESS - PLEASE DO NOT CIRCULATE OR POST]

Abstract

We use a novel dataset of production costs, wholesale prices, and retail prices from a large global manufacturer to study markups and pricing behavior along the supply chain. We document new facts about the distribution and dynamics of markups, and we propose a model of supply-chain pricing behavior that rationalizes key patterns in our data. We find substantial dispersion in markups across products and negative correlations between manufacturer and retail markups in the cross section and over time. Despite time-series variation in markups at each stage, total markups—reflecting the relationship between retail prices and production costs—are stable over time, even as prices rose during the 2022 inflationary episode in the United States. This evidence indicates frequent *margin reallocation* along the supply chain. We provide a test for vertical conduct that allows us to distinguish coordinated pricing from sequential pricing. Our test implies that retail prices are coordinated along the supply chain, and we interpret wholesale prices as reflecting bargaining weights. We use the model to quantify predictors of relative bargaining power between the manufacturer and retailers. Using counterfactuals, we show the implications of coordinated vertical pricing for policy debates on tariffs and trade.

Keywords: Markups, Supply Chain, Vertical Relationships, Pass-through, Inflation
JEL Codes: D22, D40, E3, L11, L81

*This research project was conducted at the HBS Pricing Lab and funded by the Digital, Data, and Design Institute at Harvard. We thank Chris Conlon, Silke Forbes, Sharat Ganapati, Basile Grassi, Luigi Iovino, Nir Jaimovich, Sarah M. Lein, Igor Letina, Isabelle Mejean, Nate Miller, Ivan Werning, and Brian Wheaton for their valuable feedback, as well as the participants at MIT Economics, DC IO Day, Esade Macro Meetings, Salento Macro Meetings, SED 2024, CEBRA Annual Meeting, and ASSA Annual Meeting. We also acknowledge Tomas Pacheco for his excellent research assistance. We are grateful to the research team at the unnamed global manufacturing company for generously providing the manufacturing data used in this paper. Alberto Cavallo is a co-founder of PriceStats LLC, a private company that contributed proprietary price data without compensation and without imposing any conditions to review the findings. Paolo Mengano acknowledges financial support from the Spanish Ministry of Science and Innovation with the project PID2024-159188NB-I00 financed by MICIU/AEI/10.13039/501100011033 and FEDER, UE. Contact Information: Alvarez-Blaser: s.alvarez.blaser@gmail.com; Cavallo: acavallo@hbs.edu; MacKay: mackay@virginia.edu; Mengano: paolo.mengano@esade.edu.

1 Introduction

Markups are central to understanding firm behavior, market power, and macroeconomic dynamics. A large literature documents substantial variation in markups across firms and over time, with important implications for resource allocation, inflation, and welfare (e.g., De Loecker et al., 2020; Döpper et al., 2025). Yet markups are rarely observed directly. Standard approaches infer them from demand systems or production functions that abstract from supply chain interactions. This obscures fundamental questions about how surplus is divided between upstream manufacturers and downstream retailers and how this division shapes prices and cost pass-through. Ultimately, understanding supply chain dynamics can better inform economic questions related to equilibrium prices, firm profits, and inflation—issues that are particularly relevant amid renewed trade tensions and tariff uncertainty.

In this paper, we use a novel dataset that allows us to observe markups directly at multiple stages of the supply chain. Our data include production costs, wholesale prices, and retail prices for approximately 2,000 products from a large global manufacturer of nondurable household goods. The data span four countries and five years, from July 2018 through June 2023. This window into the vertical structure of prices reveals several new patterns. Markups vary substantially across products at each stage, but they exhibit negative correlation along the supply chain: a product with a high manufacturer markup typically has a low retail markup, and vice versa. Over time, when manufacturer markups rise, retail markups tend to fall. As a result, *total markups*—reflecting the relationship between retail prices and production costs—exhibit less variation across products and over time.

These facts suggest that the supply chain often experiences *margin reallocation* rather than uniform margin expansion or contraction: profits shift between manufacturers and retailers while the overall wedge between retail prices and production costs remains relatively stable. This has direct implications for policy debates about firm conduct. During the 2022 inflationary episode, retail prices in our data increased substantially. At the same time, retail markups increased, but manufacturer markups fell, and total markups remained essentially unchanged. This finding is inconsistent with the "greedflation" hypothesis—which holds that firms uniformly exploited inflation to increase markups—and it suggests that sector-level markup estimates may obscure relevant offsetting dynamics along the supply chain.

To interpret these patterns, we develop a model of pricing along the supply chain in which prices are determined in two stages. In the first stage, the manufacturer sets retail prices to maximize total variable profits for its products. In the second stage, the manufacturer and retailer negotiate wholesale prices through Nash bargaining, splitting total variable profits according to their bargaining weights. The model brings into focus an empirical object we term the *margin share*—the share of supply-chain variable profits accruing to the manufacturer. Changes in the margin share produce negatively correlated movements in upstream and downstream markups

and can explain stable total markups even when stage-specific markups are volatile. We invert the model to recover observation-level bargaining weights and use variation across products, countries, and time to quantify which economic forces shift bargaining power between manufacturers and retailers.

Our empirical analysis yields new facts about total markups, as well as the markups for each segment of the supply chain. Average total markups are similar across the United States, the United Kingdom, Canada, and Mexico, ranging from 0.60 to 0.67 (measured as the Lerner index).[1] In the United States, the average manufacturer markup is 0.54 and the average retail markup is 0.28. Averaging across the two sectors (or stages), we obtain a value of 0.41. In magnitudes, these markups are consistent with structural estimates in the literature (Döpper et al., 2025; De Loecker et al., 2020; De Loecker and Eeckhout, 2018).[2]

In the cross section, we document substantial product-level heterogeneity in total markups and markups at each stage of the supply chain. One factor that drives heterogeneity in markups is product quality. Manufacturer margins and markups are positively correlated with product quality, but the relationship is weaker for retailers. Because retailers pay more for higher-quality products, retail markups are negatively correlated with quality.

Due to the negative correlation between manufacturer and retailer markups, total markups tend to be less dispersed than sector-level markups. In all four countries, the coefficient of variation is lower for total markups than for markups at either the manufacturer or retail level. Over time, total markups are less volatile than manufacturer or retail markups. These patterns indicate that, to accurately capture the behavior in our sample, (i) models of demand should be sufficiently flexible to allow for variable markups, even for a single firm, and (ii) vertical conduct should account for or endogenously generate the margin reallocation that we observe along the supply chain. This would rule out demand systems that yield constant firm-level markups, such as constant elasticity of substitution (CES) demand, and conduct assumptions that impose perfect competition in the upstream or downstream sector.

The offsetting movements in manufacturer and retailer markups motivate a model in which retail prices are determined separately from how variable profits are split along the supply chain. Using a generalized retail pricing first order condition, we derive an empirical test of vertical conduct that nests coordinated supply chain pricing and sequential pricing, where the manufacturer first sets wholesale prices, then the retailer sets prices in response. The conduct parameter captures whether the manufactuer's margin is relevant marginal cost for retail pricing. We estimate near-zero values for this parameter, implying coordinated pricing.

Our model allows us to interpret margin reallocation through the lens of bargaining power.

[1] The Lerner index for price p and cost c is given by $\frac{p-c}{p}$, and typically ranges from 0 to 1. An alternative measure, p/c, has also been used in the literature. For a given value $L = \frac{p-c}{p}$, the price-over-cost markup p/c is equal to $\frac{1}{1-L}$.

[2] Structural estimates recover markups by estimating consumer demand or firm production functions, then inverting optimality conditions. See De Loecker and Scott (2022) for an overview and a comparison of the two methods.

Manufacturer bargaining weights—and the corresponding margin shares—are greater than retailer bargaining weights in the United States, but this relationship is reversed in the other three countries. As a step toward understanding the determinants of bargaining power between manufacturers and retailers, we present a descriptive analysis that relates bargaining weights to features of the product and market. Our analysis indicates that manufacturer bargaining power is higher in markets where the manufacturer has greater market penetration and in higher-income environments. Retailers capture a larger share of margins when production costs are higher and during periods of elevated inflation—the latter of which may reflect overall consumer demand.[3]

Vertical conduct matters for price levels, profits, and predicted policy impacts of tariffs and other taxes. In the baseline counterfactual, sequential pricing raises average retail prices by about 54 percent and reduces sales of the manufacturer's products by 57 percent. Combined profits fall, with the manufacturer bearing most of the decline. We also use an illustrative production-cost increase to show how the pricing arrangement changes the distribution of losses between the manufacturer and retailers.

Overall, our analysis shows that the pass-through of production costs to prices cannot be treated as two independent reduced-form relationships along the supply chain, that is, one linking costs to wholesale prices and another linking wholesale prices to retail prices. Focusing on a single of these links would also be misleading. Instead, the dynamic response of wholesale prices reflects bargaining, and gradual changes in bargaining power can mechanically produce gradual pass-through into retail prices. Our estimated dynamics provide a new channel through which the nature of a shock affects both the timing of pass-through and the incidence of costs across supply chain stages.

Our paper contributes to three distinct areas. First, we bring direct evidence to the empirical literature on markups. Existing work typically relies on structural approaches to recover unobserved markups from demand or production data (see De Loecker and Scott, 2022, for an overview). We complement this literature by providing product-level measurements at each stage of the supply chain using the same data the firm itself uses to evaluate its products. This allows us to characterize how margins are split between manufacturers and retailers—information that is rarely available.[4] The markups we recover are consistent with structural estimates, offering a degree of cross-validation between the two approaches.

Regarding time-series patterns, we document stable total markups in recent years, adding to evidence against the view that firms increased markups during the recent inflation episode (Leduc et al., 2024; Bilyk et al., 2023). We also reconcile conflicting evidence on cyclical

[3] In Appendix D, we consider alternative assumptions with a fixed bargaining weight for each country and where variation in margin shares is attributed to variation in outside options.

[4] Villas-Boas (2007) infer vertical relations using retail price movements and find stronger retailer bargaining power, a pattern reversed in our micro data. Related work on vertical relationships includes Alviarez et al. (2023), Demirer and Rubens (2025), Avignon et al. (2025), and Toxvaerd (2024).

markup behavior: Anderson et al. (2025) document mildly procyclical retail markups, whereas Bils et al. (2018) find countercyclical patterns in manufacturing and at the aggregate level.[5] By observing both margins simultaneously, we show that manufacturer and retailer markups move asymmetrically, so aggregate measures conceal important stage-specific dynamics.[6]

Second, we contribute to the literature on cost pass-through. A large body of work documents incomplete pass-through in response to aggregate shocks such as exchange rates, commodity prices, or tariffs (Amiti et al., 2019; Auer et al., 2018, 2021; Bonadio et al., 2020). Supply-chain analyses of pass-through remain rare, largely due to data constraints. Nakamura (2008) finds limited influence of manufacturing shocks on retail pricing; Nakamura and Zerom (2010) documents complete retailer pass-through of commodity shocks.[7] More closely related to our setting, Flaaen et al. (2025) use confidential US transaction data across multiple distribution stages to study tariff pass-through on European wines. We build on this work by documenting how pass-through varies with the nature of the shock—aggregate versus idiosyncratic, expected versus unexpected—and by linking pass-through dynamics to changes in bargaining power along the supply chain.

Third, we inform methodological debates on markup estimation and modeling choices in both microeconomics and macroeconomics. Structural approaches typically assume perfect competition or passive cost pass-through in either the upstream or downstream sector—assumptions that are inconsistent with the coordinated markup adjustments we observe at both stages. Our insights on cross- and within-product markup variation and the empirical patterns we document, including the negative correlation in markups along the chain and the offsetting adjustments over time, can discipline modeling choices in future work. We propose a test for vertical conduct and derive a vertical conduct parameter from two standard assumptions about pricing. This testing approach complements the literature on testing for horizontal conduct (e.g., Miller and Weinberg, 2017; Duarte et al., 2024) and provides an empirical approach to distinguish among different models, following Villas-Boas (2007). Observing both production costs and wholesale prices lets us estimate how much of the manufacturer's margin enters the marginal cost relevant for retail pricing.

Our analysis is subject to limitations. The data come from a single firm, which limits our ability to capture the full range of firm behavior in the economy. Nonetheless, our findings reveal important economic aspects of supply chain behavior. To the extent that other large manufacturers of consumer products behave similarly, our representation of pricing behavior—and the implications of margin reallocation—may generalize.[8] The patterns we document may

[5]Other contributions include Stroebel and Vavra (2019), Kryvtsov and Midrigan (2013), Rotemberg and Woodford (1999), Bils (1987), Galeotti and Schiantarelli (1998), Bils and Kahn (2000), Vaona (2016), and Nekarda and Ramey (2020).

[6]Our evidence complements Burstein et al. (2020) on differences between firm- and sector-level markups.

[7]Related studies include Koujianou Goldberg and Hellerstein (2013) for beer markets, Minton and Wheaton (2023) on supply chain networks, and Hahn (2024) for the European car industry.

[8]The evidence suggests that large manufacturers behave similarly: posting list prices, tracking con-

differ for manufacturers with less bargaining power, for durable goods, or in supply chains with a different structure. We abstract from downstream heterogeneity and may not capture the full range of retailer conduct. We note, however, that the manufacturer and retail markups in our data are in line with estimates from the existing literature, providing some reassurance regarding external validity and the relevance of our findings.

The paper proceeds as follows. Section 2 describes the data and defines our markup measures. Section 3 documents time-series and cross-sectional patterns in markups, including the negative correlation along the supply chain and the relationship with product quality. Section 4 presents our model of supply chain pricing and tests its retail-pricing condition against sequential pricing using an estimated demand system. Section 5 recovers bargaining weights from the model, examines their determinants, and studies how production-cost changes are passed through to wholesale and retail prices and reallocate margins along the supply chain. Section 6 uses the estimated model to quantify the cost of double marginalization and the incidence of a cost increase under alternative vertical conduct. Section 7 concludes.

2 Data

2.1 Prices, Costs, and Quantities

Our analysis relies on two distinct data sources that together allow us to measure markups and examine pricing behavior along the value chain. The first source comprises detailed product-level information from a global manufacturer. The second consists of a vast collection of retail prices provided by PriceStats, a private company related to the Billion Prices Project (Cavallo and Rigobon, 2016). This combined dataset offers a unique opportunity to study price formation and pass-through mechanisms along the entire supply chain, from production to retail.[9]

The first dataset comes from a large global manufacturer of nondurable household products that operates multiple brands in each of the countries in our analysis. The dataset includes monthly country-SKU-level records of revenues, quantities, and costs from Canada, Mexico, the United Kingdom, and the United States. A product is defined as a combination of a brand, product form, package size, and variant (such as its fragrance). When multiple SKUs are assigned to identical products (for accounting or logistical reasons), we aggregate to the product level. The manufacturer provides information about per-product monthly expenses on raw materials, packaging, manufacturing operating expenses, transportation, and warehousing. We sum up these components (which exclude fixed and overhead costs) to construct a single measure of

sumer behavior, and using high-frequency negotiations with retailers. For an illustration of the focus on downstream consumer behavior, see, e.g., coverage of earnings calls in which large consumer product firms discuss their responses to shifts in consumer sentiment (`https://consumergoods.com/pg-kdp-nestle-colgate-palmolive-unilever-respond-dive-consumer-sentiment`).

[9]Construction of this dataset was supported by the Pricing Lab at Harvard Business School.

variable costs at the product-month level. We calculate unit prices and unit costs by dividing revenues and costs by the quantity sold in each month.[10]

The manufacturer dataset features two additional features that are useful for our purposes: measures of product quality and a breakdown of realized costs into expected and unexpected costs. Quality measures are derived from consumer surveys in which respondents rate perceived quality and rank the uniqueness of the product relative to competitors. The resulting variables, *perceived quality* and *product differentiation*, are available only for products sold in the United States. We standardize these measures for our analysis and focus on perceived quality in the main text. As for the cost breakdown, the manufacturer relies on an in-house team of experts to forecast costs for each product, which are recorded as *expected* costs. Any discrepancies between these forecasts and the actual realized costs are recorded as *unexpected* costs. These features of the data allow us to examine how pricing behaviors vary with quality and in response to different types of cost shocks.

The second dataset includes retail prices and product details, including category and sale status. It is provided by PriceStats, a private firm that collects prices from online retailers using web-scraping techniques and uses them to provide insights into daily price changes. Although the products in our data are also sold through brick-and-mortar channels, we view online prices as a reasonable measure for our purposes.[11] We aggregate daily observations to monthly average prices for each product-retailer. We have a single production cost and wholesale price for each product, but typically multiple retail prices. We construct an average retail price per product by netting out retailer fixed effects.[12] Once we have a single retail price per product, we manually match the prices to the manufacturer's production data. We drop outliers and then winsorize all variables at the 1% level.[13]

The resulting dataset, which we refer to as our main sample, includes 1,918 matched products across 16 brands and covers July 2018 to June 2023. Appendix Table A.1 shows the number of observations, products, brands, and retailers available for each country. We were able to match more than 85% of the manufacturer's sales to the retail data in the United States,

[10]Revenues are recorded net of all volume- or quantity-based rebates that the manufacturer grants to individual retailers. The resulting figure therefore reflects an average transaction price per product. We use this average price throughout, while noting that marginal prices could diverge if the manufacturer applies additional ad-hoc discounts.

[11]Cavallo (2017) shows that online prices are very similar to offline ones, even identical in 72% of the cases. Other papers using this data include Cavallo (2013), Cavallo et al. (2024); Cavallo and Kryvtsov (2023); Cavallo (2018), and Alvarez et al. (2022).

[12]Specifically, within each country, we run the following specification:

$$\tilde{p}^R_{ijt} = b + \delta_i + \delta_t^{s(i)} + \delta_j + \varepsilon_{ijt}, \tag{1}$$

where δ_i, $\delta_t^{s(i)}$, and δ_j are product, time, and retailer fixed effects, respectively. Time fixed effects are calculated separately by product category, $s(i)$. In the second step, we subtract the retailer fixed effect (δ_j) from the raw retail price ($\tilde{p}^R_{ijt}$) and average within product and period to compute retail prices, $p^R_{it} = \frac{1}{|J_{it}|}\sum_{j\in J_{it}}(\tilde{p}^R_{ijt} - \delta_j)$.

[13]We classify observations as outliers if (a) products have negative revenues, costs, or volumes or (b) a product's log absolute change in volume from one month to the next is greater than 1 and the log absolute change in price or cost also exceeds 1.

and around 70% considering all countries. Appendix Table A.2 reports summary statistics for the main sample.

2.2 Product Markups and Margin Shares

A key object of interest for our study is the supply chain *total markup*. For product i at time t, we compute this value as the Lerner index in terms of retail price (p_{it}^R) and production cost (c_{it}), $\mu_{it}^{TOT} = \frac{p_{it}^R - c_{it}}{p_{it}^R}$. This markup reflects the wedge between the price that consumers pay and the production costs, and it takes on values between 0 (price equals cost) and 1 (prices substantially greater than costs).

Our data also allow us to construct markups at the manufacturer and retailer level separately for individual products. We compute manufacturer markups as the Lerner index in terms of wholesale prices (p_{it}^M) and production costs, $\mu_{it}^M = \frac{p_{it}^M - c_{it}}{p_{it}^M}$. Similarly, we compute retail markups with retail and wholesale prices, $\mu_{it}^R = \frac{p_{it}^R - p_{it}^M}{p_{it}^R}$, following other papers that study retail markups (e.g., Aguirregabiria, 1999; Eichenbaum et al., 2011; Anderson et al., 2025). The rationale for this approach is that the marginal cost of a product for a retailer is its replacement cost, with other retailing costs fixed over short horizons (Gopinath et al., 2011).

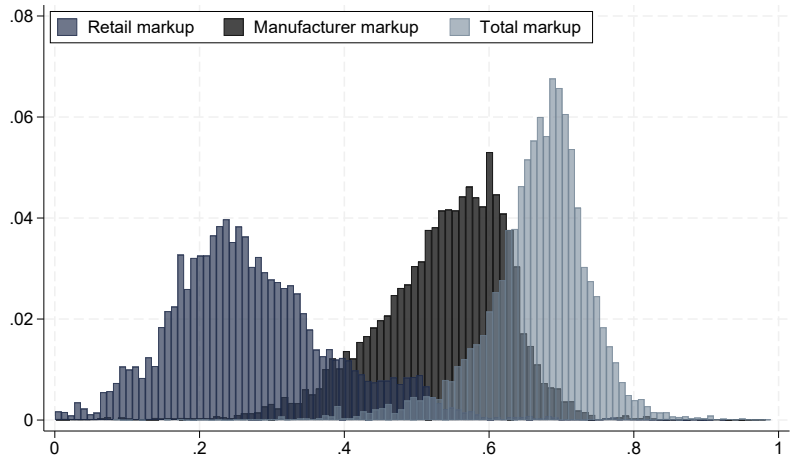
We note that our price and cost measures are constructed as within-month averages of revenues and variable costs, which may not always correspond to revenues and costs for the marginal unit. However, given the nature of the products and the fact that these data are used by the manufacturer to assess margins, we believe average variable costs and revenues are reasonable approximations to marginal costs and marginal prices in our setting. A separate consideration is that, to the extent that additional variable retailing costs (such as home delivery costs) exist, our measure of retail markups will be biased upward. In that case, our retail markups may be interpreted as upper bounds.[14] Similarly, our measure of total markups omits these additional retailer marginal costs and may also be biased upward. In our regressions, we employ fixed effects to control for certain types of unobserved costs that are common across products.

We report key results using an alternative price-over-cost markup measure in Section C. Since there is a monotone relationship between the Lerner index and price-over-cost markups, changing the measure yields very similar patterns. We maintain the Lerner index in the main text because it is bounded, easier to interpret in levels, and less sensitive to outliers.

Finally, we construct the manufacturer and retailer *margin share*—the share of total markups (or dollar margins) obtained by each stage of the supply chain. Margin shares are calculated as $\frac{p_{it}^M - c_{it}}{p_{it}^R - c_{it}}$ and $\frac{p_{it}^R - p_{it}^M}{p_{it}^R - c_{it}}$. To illustrate the difference between markups and margin shares, note that an increase in cost while holding fixed wholesale and retail prices will increase the retail

[14] Sangani (2022) compares retail markups measured using replacement costs with estimates resulting from structural methods based on demand and production models, thereby incorporating a richer notion of marginal costs, and finds similar markup levels.

Figure 1: Markups Along the Supply Chain



Notes: This figure shows the sales-weighted frequency distribution of markups along the supply chain in the United States. Retail markups ($\frac{p^R - p^M}{p^R}$) are shown in dark blue, manufacturer markups ($\frac{p^M - c}{p^M}$) in dark grey, and total markups ($\frac{p^R - c}{p^R}$) in light blue.

margin share but leave retail markups unchanged. Also, because the wholesale price enters the denominator of the retail markup, a retailer may have a lower markup than the manufacturer but a higher margin share. Our measures of margin shares have a direct connection to the model we present in Section 4.

3 Supply Chain Markups

In this section, we analyze markups and pricing behavior in four countries: the United States, the United Kingdom, Canada, and Mexico. We begin by discussing the distribution of markups and time-series variation for the largest market in our sample, the US market. Subsequently, we expand the analysis by comparing it with the UK, Canadian, and Mexican markets.

3.1 Markup Levels and Heterogeneity

In Figure 1, we plot the distribution of markups for each stage of the supply chain, along with total markups, in the US.[15] Manufacturer markups are substantially larger than retailer markups. The mean manufacturer markup is 0.54, nearly twice the mean retailer markup of 0.28. This indicates an average product markup of 0.41, while the average total markup, reflecting the entire supply chain, is 0.67.

Retail markups, manufacturer markups, and total markups exhibit substantial heterogeneity. The 10th and 90th percentiles for retail markups and manufacturer markups are (0.15, 0.43) and (0.42, 0.63), respectively. The 10th and 90th percentiles for total markups are (0.58, 0.74).

[15] Appendix Figure C.1 presents the same distributions with markups expressed as price over marginal cost.

This degree of heterogeneity is particularly striking because the products in our study come from a single manufacturer in a limited set of product categories. Though Figure 1 displays markups for all years in our data, we observe similar levels of dispersion in the cross section (see Appendix Figure A.1 for separate plots by year). Summary statistics for the distribution of markups and margin shares are presented in Appendix Table A.4.

These findings have important implications for firm behavior and consumer demand. First, to accurately capture product-level cost and price variation, models of supply and demand should allow markups for a single firm to vary meaningfully across products. Second, the differences in markups at each segment of the supply chain suggest that manufacturers and retailers are not symmetric in terms of, e.g., competitive pressures, bargaining positions, or conduct. The differences in markups between the two segments are not simply due to the fact that retailer markup is calculated with a larger denominator. The margin share—i.e., the share of the variable profits in dollar terms—is much higher for manufacturers (0.59).

Our direct measurements are consistent with estimates from the literature that use indirect (econometric) approaches to recover markups. Using the production approach, De Loecker et al. (2020) estimate an average firm-level markup (across manufacturers, retailers, and other sectors) of 0.38 in 2016, which is close to our value of 0.41. Papers that use the demand approach to estimate markups for non-durable household products have employed assumptions about vertical conduct between manufacturers and retailers that implicitly recover the total markup. With this approach, Döpper et al. (2025) estimates a median markup of 0.63 in 2019, comparing favorably to our estimate of 0.67. Brand (2021) and Atalay et al. (2025) report slightly lower median markups and 75th percentile markups of 0.69 and 0.64. Using direct measures of prices and wholesale costs, Sangani (2022) finds a retail markup of 0.24 in 2007, close to our estimate of 0.28, albeit from a period over 10 years prior.[16]

3.2 Markups Over Time

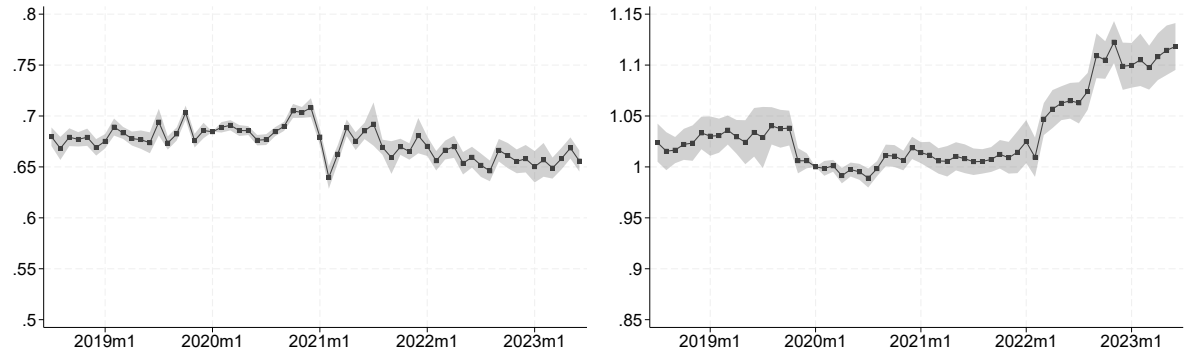
We now consider time-series variation in markups. We first examine the pattern in total markups from July 2018 through June 2023. We estimate time series trends by regressing our outcomes of interest on product and period (month-year) fixed effects, and then reporting the coefficients from the period fixed effects.[17]

Figure 2 displays the results for total markups and retail prices. Panel (a) shows that the average total markup in the US was fairly stable over time. Monthly values remained close to the average value of 0.67, with slightly higher values in 2020 and a slight decline starting in 2021. The degree of stability is striking, especially with the inflationary period occurring in the

[16] De Loecker et al. (2020) and Sangani (2022) report price-over-cost markups of 1.6 and 1.32, respectively, which we translate to the Lerner index values.

[17] We weight observations by sales. The inclusion of product fixed effects controls for changes in the composition of products over time.

Figure 2: Time Series of Prices and Total Markups



(a) Total Markups (b) Retail Prices

Notes: This figure shows the time series trends in (a) total markups and (b) retail prices in the United States. Total markups are defined as the Lerner index using retail prices and manufacturing costs. We report time period coefficients from regressions on period and product fixed effects using sales weights, with 2020m1 as the base period. In panel (a), the base period is indexed to its (sales-weighted) mean value. In panel (b), the dependent variable is log price and the base period is indexed to a value of 1. 95 percent confidence intervals are displayed in gray computed using product-level clustered standard errors.

US toward the end of our sample.[18] Panel (b) shows the time series for retail prices, which increased approximately 10 percent in 2022. Despite the sharp rise in retail prices, we see no corresponding change in total markups. Thus, for the products in our sample, retail price changes coincided with increases in production costs. This finding indicates that, during this period, there was no discrete change in pricing practices across the supply chain, as would be implied by narratives for which inflation was driven primarily by higher markups, such as the "greedflation" hypothesis.

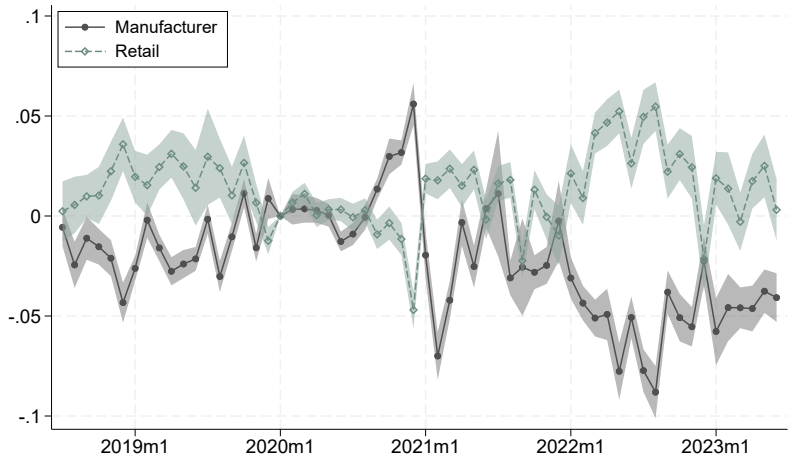
In contrast to total markups, there are meaningful fluctuations in markups at individual stages in the supply chain. In Figure 3, we plot the time series of manufacturer and retailer markups separately, normalizing January 2020 to 0. Manufacturer markups increased sharply at the end of 2020, returned to previous levels in 2021, then declined in 2022. Retailer markups followed an inverse pattern, declining at the end of 2020 and increasing in 2022. Even at a monthly frequency, we observe that changes in one sector's markups are often offset by inverse changes in the other, contributing to stable total markups.[19] This establishes evidence for margin reallocation in the time series. In our data, shifts in the split of surplus occur at higher frequency—and greater magnitudes—than supply chain margin expansions or contractions.

These time series patterns yield two main implications for supply chain behavior. First, vertical conduct along the supply chain may stabilize total markups, even in the face of significant shocks to production costs or wholesale prices that generate variation in manufacturer or

[18] Appendix Figure C.2 presents the same time-series patterns with markups expressed as price over marginal cost. The stability over time remains virtually identical.

[19] Appendix Figure C.3 shows the same time-series relationship using markups expressed as price over marginal cost. The negative co-movement between manufacturer and retailer markups is similar.

Figure 3: Markup Dynamics



Notes: This figure shows the negative dynamic correlation of markups along the supply chain in the United States. Changes in average markups and 95% confidence intervals computed using product-level clustered standard errors resulting from a regression on period and product fixed effects using sales weights.

retailer markups. Second, the empirical relationship between firm profitability and consumer prices can depend on where the firm is in the supply chain. During the period of increased consumer prices starting in 2022, variable profits, as captured by the markup, were relatively higher for retailers and relatively lower for the manufacturer. Thus, restricting attention to a single part of the supply chain can be misleading; credible inference requires either multi-stage data or an explicit model of supply-chain conduct and pass-through.

3.3 Exploring Cross-Sectional Markup Variation

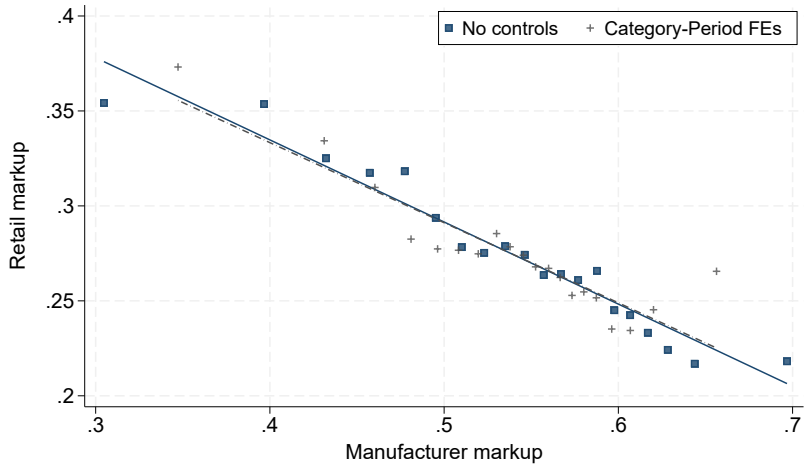
The distributions displayed in Figure 1 show that the variance in total markups is smaller than the variation in manufacturer and retail markups. This is due in part to the fact that manufacturer and retail markups are negatively correlated in the cross section. To demonstrate this more formally, we plot a binscatter of manufacturer versus retailer markups in Figure 4. The plot displays two sets of values: the raw data and the residual markups after controlling for category-period fixed effects. In both specifications, we find statistically significant negative slope coefficients.[20] This negative relationship is corroborating evidence for the presence of margin reallocation along the supply chain.

We further investigate these correlations by decomposing markups into time, product, and residual components using the regression:

$$\mu_{it}^{d} = \delta_{t}^{d} + \delta_{i}^{d} + \varepsilon_{it}^{d} \tag{2}$$

[20] Appendix Figure C.4 presents the same relationship using markups expressed as price over marginal cost.

Figure 4: Markup Correlation Along the Supply Chain



Notes: This figure shows the negative correlation of markups along the supply chain in the United States. Bins include sales-weighted values and values residualized on category-period fixed effects.

where $d \in \{\text{Manufacturer, Retail}\}$, δ_t^d denotes time fixed effects, δ_i^d product fixed effects, and ε_{it}^d the residual component. The correlation coefficients are then computed for the aggregate markup series and for each of these components in the two stages of the supply chain. Results are shown in Appendix Table A.5.[21] The aggregate correlation between manufacturer and retail markups is -0.34. For the US, the correlation is negative for all components. There is a strong negative correlation over time (-0.70; Figure 3), along with a cross-sectional correlation across products of -0.42 and a residual correlation of -0.12. Our proposed model in Section 4, which invokes bargaining between the manufacturer and retailers, provides a way to rationalize the negative correlation in markups across products.

Finally, to explore one potential determinant of these patterns, we document the correlation between product quality and margins and markups in Table 1. We find that higher product quality is associated with higher total margins and markups, consistent with consumers having a less elastic demand for high-quality products. However, examining the patterns separately for the manufacturer and the retailer reveals stark differences in pricing behavior. The manufacturer achieves higher margins and markups with higher quality products. A one standard deviation increase in quality is associated with an 18 percent increase in manufacturer margins and a 5.7 percent increase in manufacturer markups. In contrast, retail margins do not exhibit a statistically significant relationship with quality,[22] which leads to a negative correlation between retail markups and product quality. Product quality explains a larger share of the variance in margins than in markups for all stages of the supply chain, reflecting the fact that higher-quality prod-

[21] We use per-period sales weights in both the regression and the correlations.

[22] Uniform margins within a category can be rationalized by a multi-product firm facing a logit demand (de Roos, 2011).

Table 1: Margins and Quality

	Margins			Markup		
	Total	Manufacturer	Retail	Total	Manufacturer	Retail
Perceived Quality	9.948***	17.869***	1.705	1.176***	5.678***	-8.182***
	(1.127)	(1.047)	(2.078)	(0.358)	(0.619)	(1.153)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.79	0.81	0.56	0.38	0.35	0.30
Observations	16,476	16,476	16,476	16,476	16,476	16,476

Notes: This table reports the linear relation between margins or markups and (standardized) product quality estimated from the following linear model: $y_{it} = a + \alpha q_i + \beta X_{it} + \phi_{it} + \varepsilon_{it}$, with y_{it} being the variable of interest, margins or markups, for product i at time t, q_i (standardized) product quality, X_{it} controls including linear and quadratic package size to account for quantity discounts, ϕ_{it} fixed effects including product category interacted with time, and ε_{it} a mean zero error. Coefficients and standard errors reported as percent relative to the mean of the dependent variable. Product-level clustered standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ucts tend to be more costly to produce. We find similar patterns with our alternative measure of quality (product differentiation), which we report in Appendix Table A.6.[23]

These results further highlight the importance of considering separate stages of the supply chain for questions related to profitability, markups, and price levels. Our results indicate that product quality matters more to the manufacturer than it does to the retailer (in equilibrium). We proposed a two-stage pricing-bargaining framework in Section 4 to rationalize these patterns.

3.4 The UK, Canada, and Mexico

We now examine markup patterns across the four countries in our data. Figure 5a shows total markups in Canada, Mexico, the United Kingdom, and the United States. The distribution of total markups is remarkably similar across countries, with average values of 0.64 in the UK, 0.65 in Canada, and 0.60 in Mexico, compared to 0.67 in the US. All four countries exhibit meaningful dispersion in total markups (see Appendix Table A.4 for summary statistics).[24]

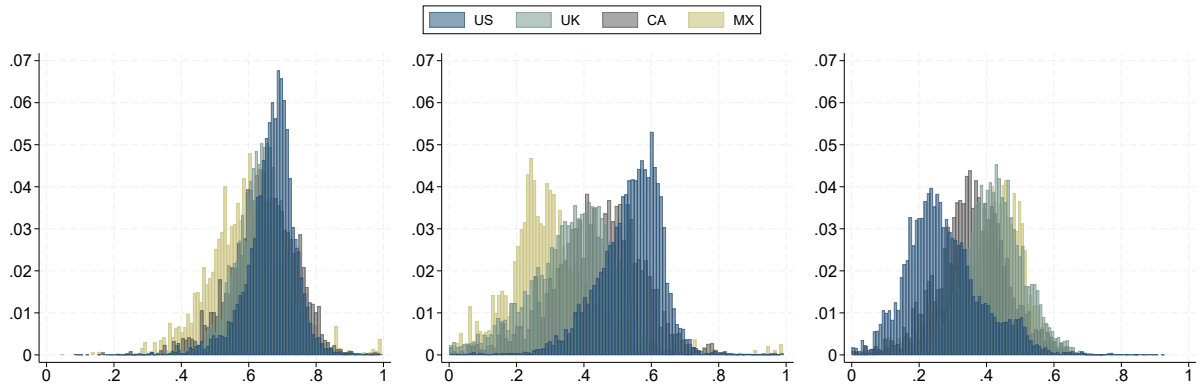
Despite comparable total markup levels, we find substantial heterogeneity in the distribution of markups along the supply chain between the US and other countries (Figure 5b and Figure 5c). In Canada, manufacturer and retailer markups are more evenly distributed, though the manufacturer has slightly higher markups on average. In the UK, retail markups are systematically higher than manufacturer markups. Finally, Mexico has the highest retailer markups and the lowest manufacturer markups in the sample. These findings stand in stark contrast to the US market structure, as retailers in these three countries capture a substantially larger proportion of total markups.

To see this more directly, Figure 6 presents the variation in product-level margin shares

[23] Appendix Figure A.2 and A.3 show binscatter plots along the entire distribution of the two quality indicators.

[24] Appendix Figure A.4 shows total markup distributions using all raw retail prices instead of our constructed average per-product retail prices.

Figure 5: Markups Distributions



(a) Total (b) Manufacturer (c) Retail

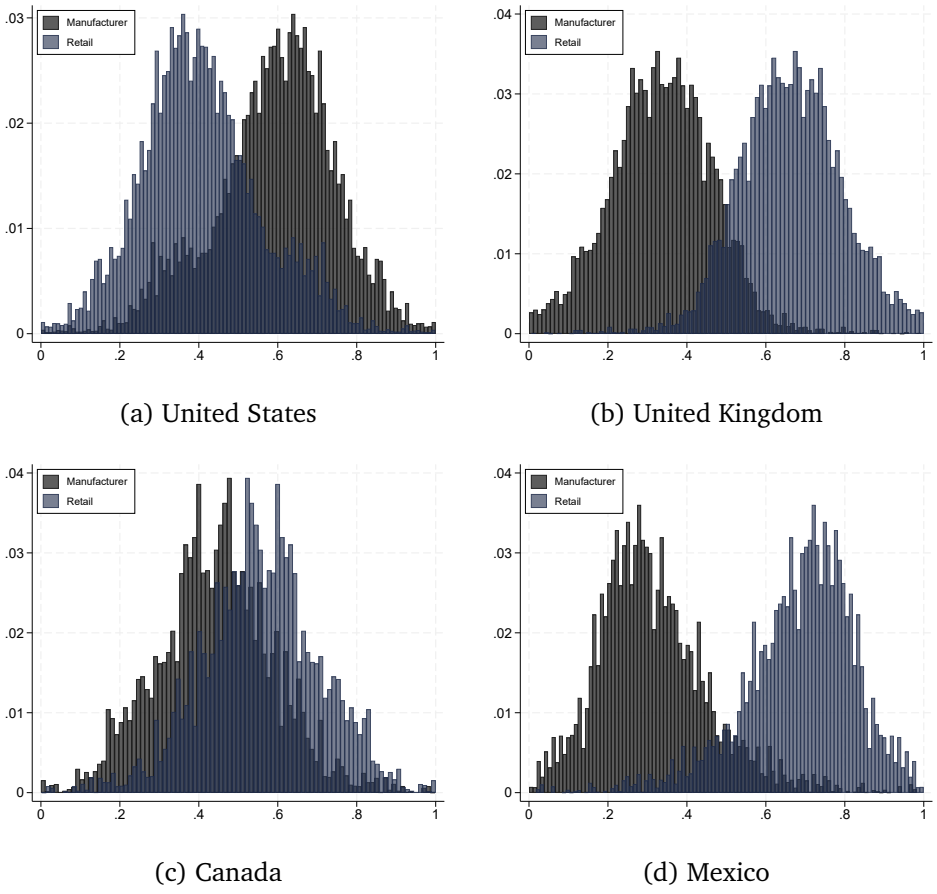
Notes: This figure shows the sales-weighted frequency distribution of markups along the supply chain. The total markup is defined as the Lerner index of retail prices and production costs, $(p^R - c)/p^R$, the manufacturer markup as the Lerner index of wholesale prices and production costs, $(p^M - c)/p^M$, while the retail markup as the Lerner index of retail prices and wholesale costs, $(p^R - p^M)/p^R$.

across countries. The gray bars illustrate the distribution of the manufacturer margin shares, which are highest in the US. Canada, the UK, and Mexico all have higher average margin shares for the retail sector. Another striking feature about the figure is the heterogeneity in margin shares across observations. The standard deviation ranges from 0.13 to 0.16 across countries, with all countries realizing manufacturer margin shares below 0.1 and greater than 0.9. This indicates a great deal of flexibility in how surplus is split among the supply chain at the product level. This is reflected in a negative correlation between manufacturer and retailer markups, which we show for each country in Appendix Figure A.5.

Total markups remain relatively stable over time in all countries. This stability reflects the negative correlation between manufacturing and retail markups over time, confirming the pattern documented for the US. Appendix Figure A.6 displays this dynamic relationship, where increases in manufacturer markups correspond to decreases in retailer markups, and vice versa. This compensatory mechanism helps maintain overall markup stability by balancing fluctuations between different stages of the supply chain. Appendix Table A.5 reports the corresponding correlation coefficients based on the decomposition in equation (2), highlighting cross-country heterogeneity. The UK exhibits a negative correlation in both aggregate markups and all underlying components. By contrast, Canada and Mexico have a negative correlation in two components and positive correlation in the third. Canada does not show a negative correlation across product fixed effects. Mexico stands out as the main exception, with weak negative correlations in the cross section, a positive correlation in the time series, and a positive overall correlation in manufacturer and retail markups.

Table 2 summarizes our findings across countries. We report means of our measures of markups, as well as the manufacturer and retailer margin shares. The margin share patterns are

Figure 6: Manufacturer and Retail Margin Shares



(a) United States (b) United Kingdom

(c) Canada (d) Mexico

Notes: This figure shows the sales-weighted frequency distribution of margin shares along the supply chain for all years separately by country. Manufacturer margin share is defined as the ratio of the manufacturer margin to the total margin $\frac{p^M - c}{p^R - c}$. Retailer margin is share defined as the ratio of the retail margin to the total margin $\frac{p^R - p^M}{p^R - c}$.

similar to the markup patterns, with the exception of Canada: retailers have a lower markup than manufacturers but a higher margin share.[25] Our markup estimates at the sector level (averaging across the manufacturer and retail markup) are consistent with estimates reported in the literature.[26]

[25] As an example, consider $P^R = \$2.75$, $P^M = \$1.75$, and $c = \$1.00$. The retailer has a higher margin share ($1.00 out of the $1.75 total margin) but a lower markup.

[26] In the UK, De Loecker and Eeckhout (2018) and Díez et al. (2021) obtain values of 0.40 and 0.23, close to our estimate of 0.39. In Mexico, De Loecker and Eeckhout (2018) obtain an estimate of 0.35, compared to 0.37 in our data. In Canada, De Loecker and Eeckhout (2018) obtain an estimate of 0.35, which is slightly lower than our estimate of 0.40.

Table 2: International Comparison

	Countries			
	US	UK	Canada	Mexico
Average Total Markup	0.67	0.64	0.65	0.60
Average Manufacturer Markup	0.54	0.36	0.46	0.32
Average Retail Markup	0.28	0.42	0.35	0.42
Average Sector Markup	0.41	0.39	0.40	0.37
Average Manufacturer Margin Share	0.59	0.33	0.46	0.31
Average Retail Margin Share	0.41	0.67	0.54	0.69
Negative Time Series Corr.	Yes	Yes	Yes	No
Negative Product Corr.	Yes	Yes	No	Yes

Notes: This table summarizes the main results across the countries included in our sample. Margin shares are computed as manufacturer or retail margin over total margins, i.e., $\tilde{M}_t^M = M_t^M/(M_t^M + M_t^R)$ and $\tilde{M}_t^R = M_t^R/(M_t^M + M_t^R)$.

4 Empirical Model of Vertical Conduct

We propose a model of supply-chain pricing behavior to rationalize our key findings. The model provides a framework that can generate the negative correlation between manufacturer and retailer markups—a pattern not readily explained by standard vertical pricing models. It yields a tractable estimable equation that maps observed prices and costs to an economically interpretable parameter: the manufacturer's bargaining power.

Our data also allow us to test between our model and alternative forms of vertical conduct. We show a more general retail pricing condition can be expressed as a function of the relevant marginal cost to the retailer. We express this relevant marginal cost as a weighted average of production costs and wholesale prices, and show that the data indicate that production costs are the relevant marginal cost for the retail prices, providing support for our assumption about supply chain pricing.

4.1 Supply Chain Pricing

The framework considers a manufacturer that supplies a set of products, $\mathcal{I}$, to multiple retailers. The supplier-buyer network is fixed, and pricing occurs in two stages. First, the manufacturer determines the optimal retail prices (p_i^R) that maximize joint manufacturer-retailer profits for the set of products $\mathcal{I}$. Second, the manufacturer and the retailer negotiate wholesale prices (p_i^M) through a bargaining process for each product separately. A key advantage of this two-stage pricing formulation is that retail prices, quantities, and joint profits are determined in stage 1, leaving only the division of profits to be negotiated in stage 2.

In the first stage, the manufacturer determines the retail prices that maximize supply-chain

profits for its own products:

$$\max_{p_i^R} \sum_{i \in \mathcal{I}} \left(p_i^R - c_i\right) q_i \ , \tag{3}$$

where p_i^R is the retail price, c_i is the production cost, and q_i is the total retail quantity for product i. Without loss of generality, we write the model as if there is a single downstream retailer.

In practice, manufacturers often avoid directly setting retail prices through contracts. Instead, they incentivize retailers to adhere to preferred prices using a combination of "list prices" at the retail level (e.g., MSRP) and "promotional" rebates to the retailer. Moreover, manufacturers tend to focus on the behavior of the end consumer and spend extensive resources analyzing them. For example, executives often focus on end-consumer retail price elasticities rather than wholesale price elasticities.[27] Our specification is consistent with this behavior.

At the retail pricing stage, the manufacturer internalizes downstream consumer substitution patterns across products in $\mathcal{I}$, products of other manufacturers sold by the retailer, and the outside option. Downstream prices $\{p_i^R\}$ determine quantities $\{q_i\}$, which reflect quantities demanded at the retail prices by end consumers and are equal to the quantities produced by the manufacturer. Retail prices, production costs, and quantities determine supply chain profits.

In the second stage, the manufacturer and the retailer bargain over the wholesale price for each product. Total supply-chain profits are fixed in the first stage, so only the split of those profits—implemented through the wholesale price—remains to be determined.[28] Let $\pi_i^M(p_i^M, q_i)$ and $\pi_i^R(p_i^M, q_i)$ denote manufacturer and retailer profits on product i at quantity q_i, and let $\bar{\pi}_i^M$ and $\bar{\pi}_i^R$ be their outside options. The Nash problem is:

$$\max_{p_i^M} \left[\pi_i^M(p_i^M, q_i) - \bar{\pi}_i^M\right]^{\tau_i} \left[\pi_i^R(p_i^M, q_i) - \bar{\pi}_i^R\right]^{1-\tau_i}, \qquad \forall i \in \mathcal{I}. \tag{4}$$

Here, τ_i and $(1 - \tau_i)$ represent the bargaining weights of the manufacturer and retailer, respectively, reflecting their relative bargaining power.

We make the following assumptions regarding the disagreement payoffs. First, we abstract from strategic considerations across products or spillover from individual negotiations to other products. Because the manufacturer and the retailer negotiate over thousands of products, and the median item accounts for less than one percent of volumes and sales within its category, we assume that products are sufficiently small and do not meaningfully affect the others. Second, if negotiations break down, we assume the retailer can immediately replace the product with a

[27] As one typical example, in a March 2023 earnings call, the executives of General Mills (a manufacturer) discussed pricing and elasticities several times, referring to retail consumer behavior. `https://seekingalpha.com/article/4589620-general-mills-inc-gis-q3-2023-earnings-call-transcript`

[28] The split can be represented as a wholesale price or a lump sum transfer; either will yield equivalent results. Consistent with our assumptions, the data used by the manufacturer does not distinguish between the two.

close substitute that yields zero profits. Finally, we assume negotiations occur at high frequency. Therefore, this implies that the production process of the manufacturer is already set and that fixed costs do not affect the bargaining result. Under these assumptions, the disagreement payoffs satisfy $\bar{\pi}_i^M = \bar{\pi}_i^R = 0$, i.e., the outside options are effectively zero for both sides.

Under these conditions, Equation (4) simplifies to:

$$\max_{p_i^M} \left[\left(p_i^M - c_i\right) q_i\right]^{\tau_i} \left[\left(p_i^R - p_i^M\right) q_i\right]^{1-\tau_i}, \qquad \forall i \in \mathcal{I}. \tag{5}$$

The first component represents the profits of the manufacturer, expressed as manufacturer margin $(p_i^M - c_i)$ times quantity. The second component is retailer profits, expressed as retail margins $(p_i^R - p_i^M)$ times quantity.

Under our timing assumptions, q_i is fixed from the first stage, thus the second-stage problem reduces to:

$$\max_{p_i^M} \left(p_i^M - c_i\right)^{\tau_i} \left(p_i^R - p_i^M\right)^{1-\tau_i}, \qquad \forall i \in \mathcal{I}. \tag{6}$$

In equilibrium, the solution to the bargaining problem yields a share of margins for the manufacturer that is equal to the manufacturer bargaining weight, τ_i. This can be expressed as

$$\frac{p_i^M - c_i}{p_i^R - c_i} = \tau_i. \tag{7}$$

Using our data on retail prices, wholesale prices, and production costs, we can recover these bargaining weights directly from the data.

With this approach, we interpret variation in profit split as arising from time-series and cross-sectional variation in bargaining power. In Section D, we present two alternative approaches that interpret this variation as arising from time-series and cross-sectional variation in disagreement payoffs, while bargaining weights are fixed at the country level. Both approaches indicate that the manufacturer has more bargaining power in the US than other countries, and, despite the differences in interpretation, they feature similar time-series variation as our baseline approach.

4.2 Framework for Testing Vertical Conduct

The first stage of the model has a testable implication. If the manufacturer sets retail prices to maximize joint profits, retail prices are set as a margin over production costs. If instead the manufacturer sets wholesale prices and retailers set retail prices on top of them, as in the standard sequential-pricing model, the wholesale price is the retailer's marginal cost, and retail prices are set as a margin over wholesale prices. Let m_i denote the implied dollar margin. Then,

for each product, the dollar margin reflects either

$$\text{Coordinated pricing:} \quad p_i^R - c_i = m_i^D, \qquad \text{Sequential pricing:} \quad p_i^R - p_i^M = m_i^D.$$

They differ only in the cost that the retail price setter treats as marginal, and because wholesale prices move relative to production costs across products and over time, they have different implications. Following this logic, we develop a test of vertical conduct that asks whether observed retail prices are better rationalized by production costs, wholesale prices, or any combination of the two.

The test builds on the multiproduct retail-pricing problem, whose first-order condition yields a margin formula that is agnostic about which cost is marginal. Regardless of which cost is used, the margins depend on consumer demand through its price elasticities. We recover them by estimating demand and then ask which of the two conditions the observed margins satisfy. Let $\boldsymbol{s}_t(\boldsymbol{p}_t^R)$ denote the vector of product shares generated by consumer demand in period t, $\boldsymbol{m}_t^D$ denote dollar margins, and Ω_t denote the ownership matrix, whose (j,k) entry equals one if optimal prices takes into account the effect of the price of product j on the profits earned from product k, and zero otherwise.[29] The first-order condition of the problem is

$$\boldsymbol{m}_t^D = \boldsymbol{p}_t^R - \widetilde{\boldsymbol{c}}_t(\psi) = -\left(\Omega_t \circ \left[\frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{p}_t^R}\right]'\right)^{-1} \boldsymbol{s}_t, \tag{8}$$

where $\boldsymbol{m}_t^D$ is the vector of price-cost margins implied by consumer demand and the pricing portfolio, and $\widetilde{\boldsymbol{c}}_t(\psi)$ is the relevant marginal cost for retail prices. We express the latter as

$$\widetilde{c}_{it}(\psi) = (1-\psi)c_{it} + \psi p_{it}^M = c_{it} + \psi\left(p_{it}^M - c_{it}\right), \qquad \psi \in [0,1]. \tag{9}$$

When $\psi = 0$, that cost is the production cost, as under coordinated pricing. When $\psi = 1$, it is the wholesale price, as under sequential pricing. For intermediate values, ψ is a reduced-form index of vertical conduct: it measures the fraction of the wholesale wedge that the price setter treats as a cost.

Thus, an econometric estimate of ψ can help inform the relevant assumptions about vertical conduct. In our setting, we obtain measures of production costs and wholesale prices directly, facilitating this exercise. Our empirical strategy is to solve for the value of ψ that sets predicted margins close to observed margins, using equation (8).

[29] We assume that retail prices internalize the profits for all the manufacturer's products within a product category. Rival products enter consumer demand, but their profits are not internalized.

4.3 Demand

To bring Equation (8) to the data, we require a demand model that allows margins to differ across products of the same manufacturer, as the evidence in Section 3 requires, and that remains tractable. Standard demand systems do not deliver such heterogeneity, CES, logit, and even a random coefficient on price alone fit average margins but imply nearly constant Lerner indices across products (Section B). We therefore specify a random-coefficients logit model in which a single unobserved consumer type governs both price sensitivity and the taste for individual products, as in Döpper et al. (2025), and estimate it separately for each of the two broad product categories in the United States. Let a market be a broad product category–month, and the choice set consist of the manufacturer's products, one aggregate rival for each narrow product category, and an outside option. Consumer h's utility from product j in market t is

$$u_{hjt} = \delta_{jt} + \beta_j(y_h - 0.5) - (\alpha_0 + \alpha_1 y_h)p^R_{jt} + \varepsilon_{hjt}, \qquad \alpha_0, \alpha_1 > 0, \tag{10}$$

where δ_{jt} is the non-price component of mean utility, and ε_{hjt} is an idiosyncratic taste shock with a Type-I extreme-value distribution. Consumer type is represented by y_h, which is uniformly distributed on $[0, 1]$. Consumers with higher values of y_h are more sensitive to prices, and $\beta_j(y_h - 0.5)$ lets them sort across products: a product with a positive β_j attracts price-sensitive consumers and faces more elastic demand. Thus, our model allows the consumer mix, and with it the margin, differs across products. We normalize the average β_j across the manufacturer's products within each broad category to zero. Market shares are obtained by integrating the resulting choice probabilities over consumer types,

$$s_{jt}(\boldsymbol{p}^R_t; \theta) = \int_0^1 \sigma_{jt}(y; \boldsymbol{p}^R_t, \theta)\, \mathrm{d}y, \tag{11}$$

where $\sigma_{jt}(y; \boldsymbol{p}^R_t, \theta)$ is the choice probability conditional on consumer type y, and θ collects the demand parameters.

4.4 Estimation

We construct empirical market shares within a broad product category by combining observed manufacturer quantities with measures of total category purchases and potential market size. Let Q^{obs}_{gt} denote the observed quantity of the manufacturer's products in narrow product category g and market t. Because rival quantities are not observed in our data, we introduce a coverage parameter ρ_g as the share of total quantity in narrow product category g accounted

for by the manufacturer's observed products and recover total category quantity as

$$Q_{gt}^{\text{all}} = \frac{Q_{gt}^{\text{obs}}}{\rho_g}. \tag{12}$$

We assign the remaining quantity, $Q_{gt}^{\text{all}} - Q_{gt}^{\text{obs}}$, to the aggregate rival and estimate the ρ_g jointly with the demand and conduct parameter.

We estimate the model by combining demand moments with the pricing condition. Each β_j is pinned down by the average margin for a product j; β_j determines selection to each product by consumer type, which allows the distribution of price sensitivities to vary flexibly across products. We use cost-shifters as instruments for α_0 and α_1. The instruments are total production cost and its available components: materials, manufacturing, transportation and warehousing, and other recorded costs. We residualize each instrument on the product and month fixed effects, discard instruments with no remaining variation, and standardize the remaining instruments to have unit variance. Identification through these instruments requires that, conditional on the fixed effects, the cost shifters be orthogonal to unobserved demand shocks. Cost shifters move retail prices under either conduct hypothesis, so they identify price sensitivity without assuming a specific vertical conduct.

On the supply side, conditional on the demand and coverage parameters, we choose ψ to minimize the weighted distance between the demand-implied margins $\widehat{m}_{it}^{D}(\theta, \boldsymbol{\rho})$ and the margins observed in the data $m_{it}(\psi)$:

$$\widehat{\psi}(\theta, \boldsymbol{\rho}) = \arg\min_{\psi \in [0,1]} \sum_{i,t} w_{it} \left[m_{it}(\psi) - \widehat{m}_{it}^{D}(\theta, \boldsymbol{\rho}) \right]^2, \qquad w_{it} \propto q_{it}. \tag{13}$$

We then choose the demand and coverage parameters to minimize the sum of the demand criterion and the supply criterion evaluated at $\widehat{\psi}(\theta, \boldsymbol{\rho})$. Because wholesale prices are observed and conditioned on, the estimation imposes neither a first-order condition for manufacturer pricing nor any restriction from the bargaining model; the retail-pricing condition is the only supply-side restriction. Section E provides additional details on market construction, the instruments, and numerical implementation.

Two features of this estimation strategy determine what identifies ψ. The cost instruments tie price sensitivity to within-product variation in retail prices over time, so the estimated elasticities must be consistent with how quantities respond to price changes and not only with the level of margins. And with one β_j per product, the model can match each product's average margin regardless of the value of ψ, but it cannot match how margins move from month to month unless it uses the right cost. The conduct parameter is therefore identified from within-product variation in the two candidate margins relative to the demand-implied markup.

Table 3: Demand and Vertical-Conduct Estimates

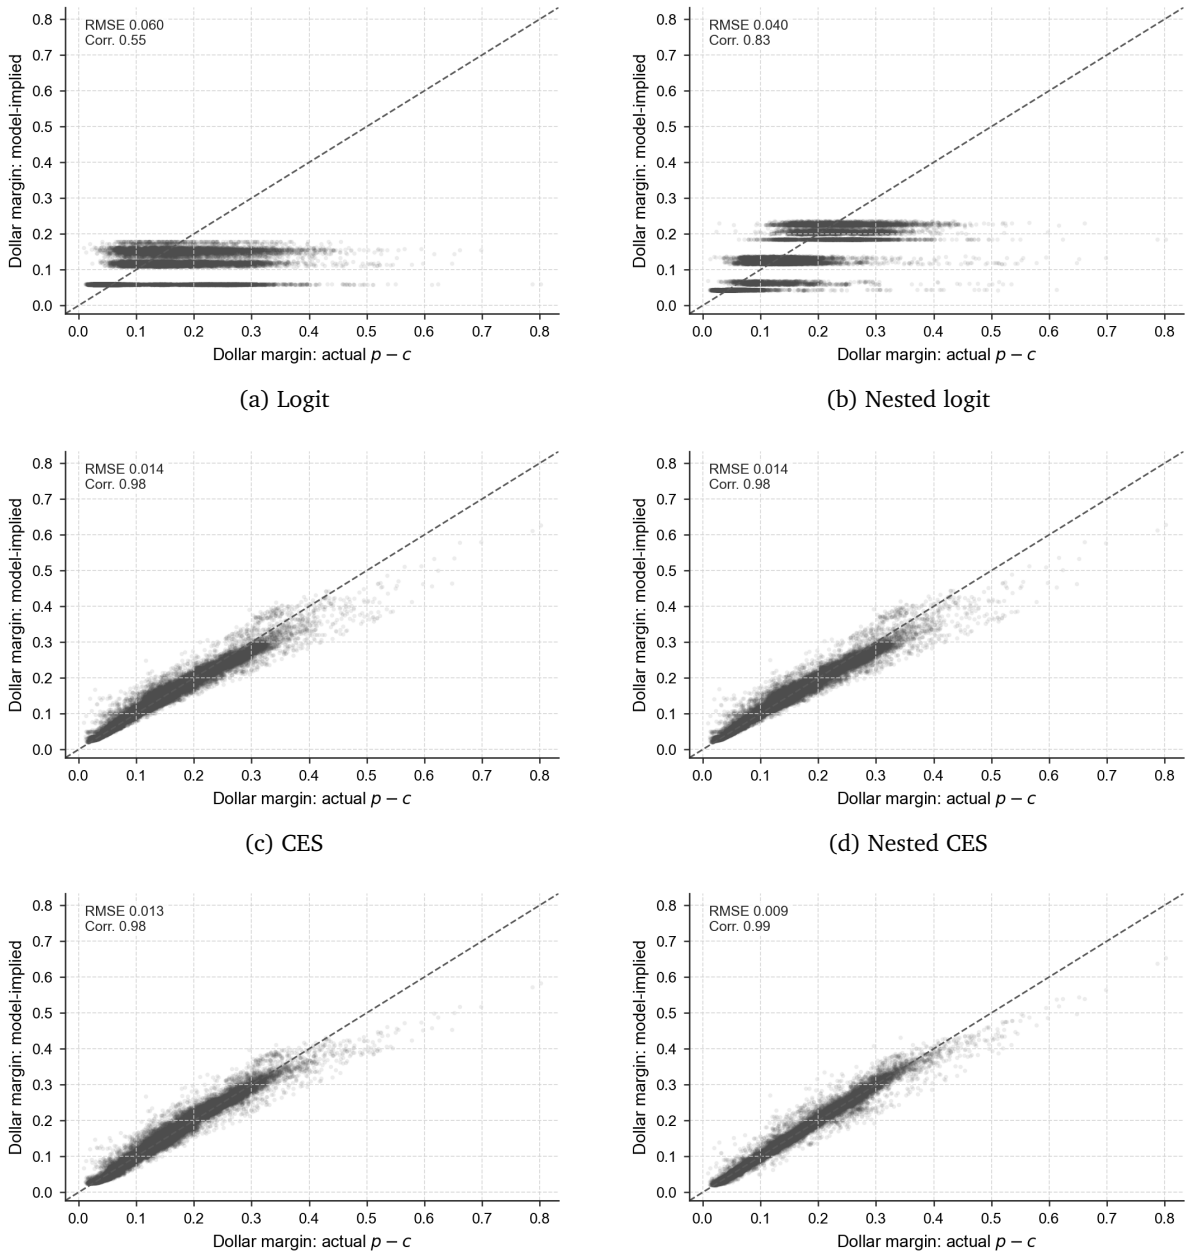
	First Category	Second Category
Demand parameters		
α_0	0.34	1.35
	[0.15, 0.28]	[0.71, 2.22]
α_1	100.90	202.16
	[101.48, 113.27]	[193.86, 308.27]
β_j range	-5.96–8.59	-8.17–16.20
Conduct parameter		
$\widehat{\psi}$	0.000	0.001
	[-0.0002, 0.0092]	[-0.0002, 0.0516]
Statistics		
Observations	10,732	13,506
Margin RMSE	0.0060	0.0112
10th percentile own-price elasticity	-1.74	-9.91
Median own-price elasticity	-1.46	-6.57
90th percentile own-price elasticity	-1.31	-3.69
Diversion to own products	0.04	0.82

Notes: The demand and conduct models are estimated separately by product category. Brackets report 95% percentile confidence intervals from 200 product-level bootstrap draws per category. Each draw assigns an exponential weight to each product and re-estimates all parameters. The interval for ψ uses the unrestricted estimate before imposing the $[0, 1]$ bound. Additional optimization in the bootstrap shifts the first-category price parameters enough that the reported point estimates fall outside their intervals. Margin RMSE is calculated using the quantity-weighted supply-side criterion. Diversion to own products is the share of diversion from a product that is recaptured by the manufacturer's other products.

4.5 Results

Parameter estimates are reported in Table 3. The estimated of ψ, the vertical conduct parameter, is 0.000 in the first category and 0.001 in the second (Table 3). The test results therefore strongly supports coordinated pricing over sequential pricing: the margins recovered from consumer demand align with the total supply-chain margin, rather than with the retailer margin, implying that production costs are the relevant marginal costs for the retail-pricing problem. Figure A.7 shows how the fit changes as ψ moves from zero to one, holding the demand estimates fixed. The estimation criterion is minimized at zero in the first category and very near zero in the second, while the fit is substantially worse under sequential pricing.

Table 3 reports the estimated demand parameters and shows that the two categories reach similar margins through different demand structures. (Note that the current point estimate lies outside of the boostrapped confidence intervals; this is because the estimates in the current draft use different convergence criteria. This will be updated.) In the first category, own-price elasticities are low (a median of -1.46) and only 4 percent of the demand that a product loses when it raises its price goes to the manufacturer's other products: each product faces a very inelastic demand of its own. In the second category, own-price elasticities are high (a median of -6.57), but 82 percent of the lost demand is recaptured by the manufacturer's other products: the portfolio, not the individual product, faces inelastic demand, and this is internalized in the



(a) Logit

(b) Nested logit

(c) CES

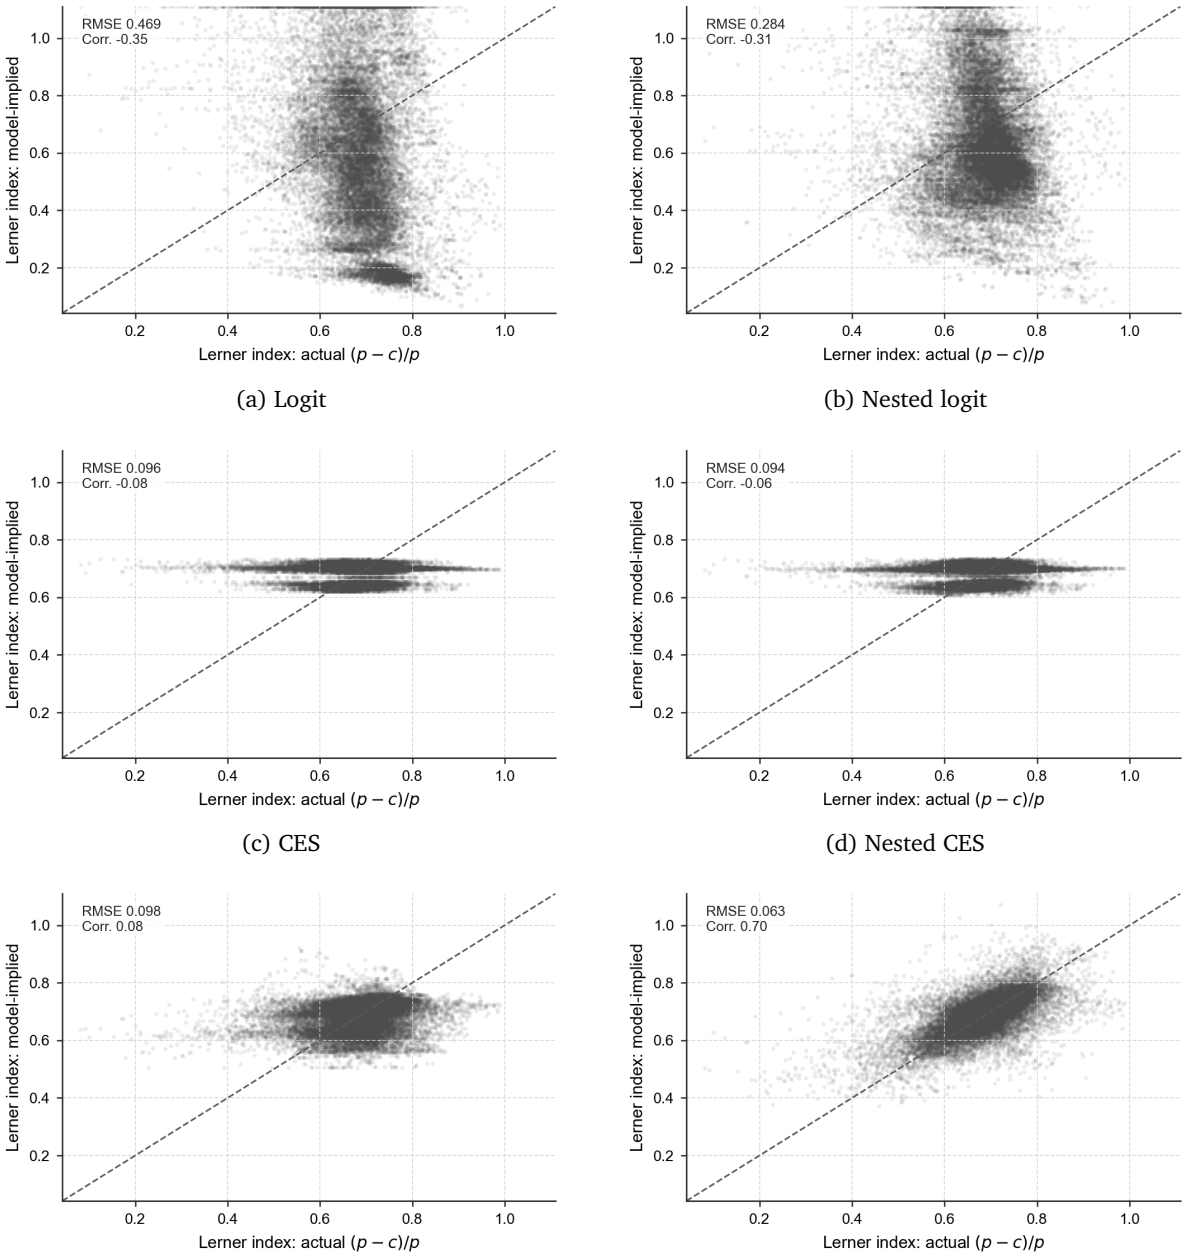
(d) Nested CES

(e) Random coefficient on price

(f) Baseline (uniform y)

Figure 7: Actual and model-implied dollar margins by demand system

Notes: Each point is a product-month, both categories pooled. Actual dollar margin $p^R - c$ against the model-implied joint-pricing markup ($\psi = 0$). All panels share the same axes. The dashed line is the 45-degree line. RMSE and correlation are quantity-weighted.



(a) Logit

(b) Nested logit

(c) CES

(d) Nested CES

(e) Random coefficient on price

(f) Baseline (uniform y)

Figure 8: Actual and model-implied Lerner indices by demand system

Notes: Each point is a product-month, both categories pooled. Actual Lerner index $(p^R - c)/p^R$ against the one implied by the joint-pricing markup ($\psi = 0$). All panels share the same axes; implied values outside the axis range are drawn at the edge. The dashed line is the 45-degree line. RMSE and correlation are quantity-weighted.

price optimization. As an additional diagnostic, Figure A.8 shows that the margins implied by estimated demand closely track observed dollar margins.

Taken together, the test results strongly support the coordinated-pricing framework over the sequential-pricing framework in all product categories. Under the bargaining assumptions detailed above, and maintaining the coordinated-pricing framework across countries, we interpret the manufacturer's margin share as its bargaining weight. In the next section, we use this measure to examine variation in surplus division across products, countries, and time.

Figures 7 and 8 illustrate the importance of flexibility in demand to match the dollar margins and markups in our data. To illustrate, we provide estimates from alternative demand systems, along with our baseline. Logit and nested logit yield strong restrictions on dollar margins and Lerner index markups that are negatively correlated with those we observed in our data, as shown in panels (a) and (b) of the figures. Panels (c) and (d) show that CES and nested CES demand provide more flexibility on dollar margins but strong restrictions on markups that are uncorrelated with those in our data. A random coefficient on price (panel (e)) does better than these alternatives but provides only a weak positive correlation with markups. Our baseline model that allows for non-price characteristics to drive selection does substantially better in fitting product-level markups, as shown in panel (f).

5 Understanding Wholesale Prices

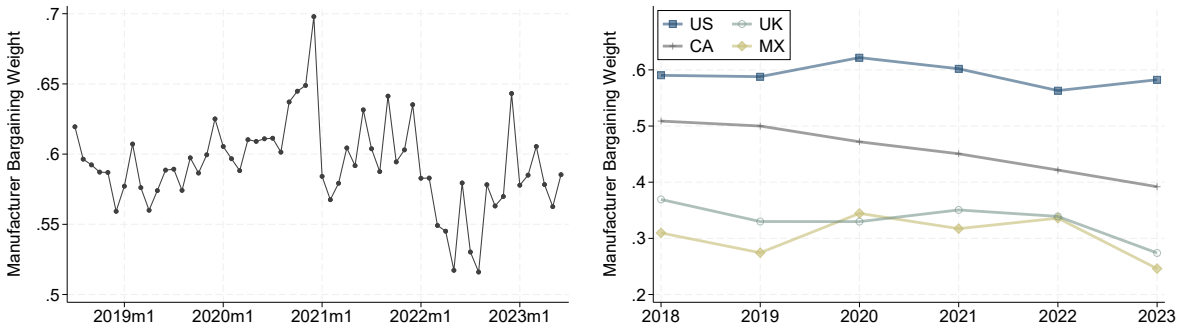
Under the coordinated-pricing framework developed in Section 4, the wholesale price divides the total supply-chain margin between the manufacturer and the retailer. Given the assumptions introduced there, the manufacturer's share of this margin recovers its bargaining weight. In this section, we first document how the manufacturer's bargaining weight varies across products, countries, and time and examine its relationship with several economic factors. We then study how changes in production costs are reflected in wholesale and retail prices and in the division of margins between the manufacturer and retailers.

5.1 Margin Allocation and Bargaining Power

We begin by documenting the relative shifts in markups between the manufacturer and the retailer and changes in bargaining power. Using the expression derived in Section 4, we calculate the manufacturer's bargaining weight for each product and period. We then construct sales-weighted averages at the monthly level for the United States and at the annual level for cross-country comparisons.

Figure 9 plots the time series of manufacturer bargaining weights. In the US (Figure 9a), the manufacturer consistently held higher bargaining power than retailers, averaging 0.62 over the observed period. Moreover, the bargaining weight remained relatively stable until late 2020, when it experienced a sharp increase and subsequent decline due to changes in production

Figure 9: Manufacturer Bargaining Power



(a) The United States

(b) All countries

Notes: This figure shows the time series of the manufacturer's bargaining weights in the US (a) and across all countries (b). Bargaining weights are calculated using equation (7) and aggregated at the country-month level, weighted by sales.

costs. By 2022, bargaining power decreased significantly, coinciding with rising production costs.

The manufacturer exhibits lower bargaining power outside of the US (Figure 9b) when negotiating wholesale prices. The average bargaining weight is 0.47 in Canada, 0.33 in Mexico, and 0.36 in the United Kingdom. The evolution over time also varies across countries: in Canada, bargaining power has steadily declined since late 2019; in the United Kingdom, it remained unchanged over the period until a significant drop in 2023; while in Mexico, the manufacturer's bargaining power increased substantially in 2020 but fell again in 2023.

The differences documented in Figure 9 show substantial variation in margin allocation and bargaining power across countries and over time. We next examine how this variation is related to several economic factors to identify potential determinants of the manufacturer's bargaining power. Factors considered include production costs (which reflect supply shocks), total margins (which reflect market structure and competition), product-level sales per capita (market penetration), GDP per capita (income levels)[30], and the Consumer Price Index (CPI) as a proxy for overall demand (once production costs and margins are controlled for). Using variation across products, time, and countries, this analysis aims to provide qualitative insights into the predictors of bargaining power.

The results in Table 4 shed light on the economic factors that influence the manufacturer's bargaining power and their relative importance in different contexts. All variables except quality (which is standardized) are in logs, allowing for the interpretation of coefficients as elasticities. The first column shows that production costs are negatively correlated with bargaining weights. This suggests that higher production costs, such as those observed during the post-pandemic surge in input costs, are linked to lower manufacturer bargaining power. A 10 percent increase in production costs is associated with a 4.96 percent decline in margin share.

[30]We adjust GDP per capita for PPP to compare across countries.

Table 4: Bargaining Weight Predictors

	(1)	(2)	(3)	(4)	(5)
Production Cost	-0.496***	-0.496***	-0.481***	-0.471***	-0.480***
	(0.014)	(0.014)	(0.017)	(0.026)	(0.036)
Total Margin	0.571***	0.559***	0.611***	0.486***	0.476***
	(0.013)	(0.014)	(0.014)	(0.029)	(0.045)
Sales per Capita	0.049***	0.051***	0.064***	0.048***	0.047***
	(0.004)	(0.004)	(0.006)	(0.005)	(0.005)
GDP per Capita	0.412***	-0.101	0.556***	0.351***	0.419***
	(0.029)	(0.071)	(0.055)	(0.104)	(0.139)
CPI	-0.168***	0.028	-0.233***	-0.134	-0.086
	(0.058)	(0.062)	(0.084)	(0.093)	(0.112)
Quality					0.027**
					(0.011)
Country FEs		X			
Non US			X		
US Only				X	X
Observations	42,458	42,458	18,220	24,238	16,476
R^2	0.672	0.675	0.688	0.477	0.477

Notes: This table reports the outcome of regressing the (log) manufacturer bargaining weights on (log) production costs, (log) total margins, (log) product sales, (log) GDP per capita, (log) CPI, and (standardized) quality. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Conversely, total margins show a positive relationship with bargaining weights. Because higher margins persist for products with fewer close substitutes, this suggests that lower competition in the downstream sector strengthens the manufacturer's position.

The coefficients of product-level sales per capita and GDP per capita show that higher market penetration and higher income are also associated with higher manufacturer bargaining power. This suggests that products with a larger presence in the market and wealthier consumer bases provide greater leverage in negotiations for this manufacturer. Finally, the coefficient of CPI shows that higher demand, as captured by aggregate price levels after accounting for costs and product-specific margins, is associated with lower manufacturer bargaining power. This suggests that positive aggregate demand shocks may lead to disproportionate gains for retailers. Together, these factors explain more than 67 percent of the variation in bargaining weights.

In columns (2) through (4), we examine alternative specifications by introducing country fixed effects, restricting the sample to non-US countries, and focusing solely on the US. Across all specifications, the estimated coefficients are remarkably stable. In column (2), the inclusion of country fixed effects emphasizes within-country variation. Relative to our baseline specification, the coefficients on GDP per capita and CPI become statistically insignificant, with almost no change in R^2. The stability in the other coefficients and R^2 suggests that our chosen variables effectively capture persistent cross-country differences in bargaining power. In column (3), we limit the analysis to Canada, Mexico, and the UK, while column (4) narrows the focus

to the US. The coefficients are very similar across these specifications, underscoring the robustness of our results and suggesting the existence of a stable framework to explain bargaining weights across countries.

Finally, in column (5) we introduce our measure of product quality, which we have only for the US. Our earlier evidence indicated that higher product quality is associated with higher manufacturer markups but lower retailer markups (Section 3). Our analysis here provides corroborating evidence in terms of bargaining power: higher quality is associated with higher manufacturer bargaining power. Specifically, the estimated coefficient implies that a one standard deviation increase in product quality is associated with a 2.7 percent increase in the manufacturer's bargaining power.

These results characterize the economic factors associated with differences in margin allocation across products, countries, and time. We next turn to margin reallocation over time and examine how the manufacturer's bargaining weight evolves following expected and unexpected changes in production costs.

5.2 Margin Reallocation and Cost Pass-Through

The estimates in Table 4 indicate that higher production costs are associated with a smaller manufacturer share of the total supply-chain margin. We now examine how wholesale and retail prices adjust as production costs change and what these adjustments imply for the split of the margin. In the model of Section 4, production costs enter retail prices, while wholesale prices determine how the resulting margin is divided. The response of retail prices to costs therefore tells us how the total margin changes, and the response of wholesale prices tells us how the manufacturer's share changes.

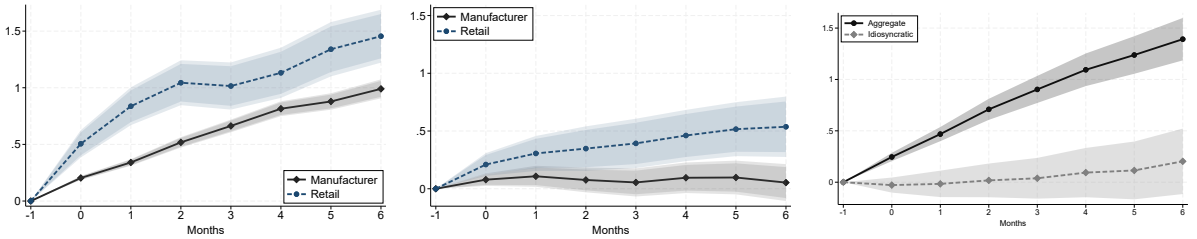
Let τ_0 denote the manufacturer's initial margin share. Following a positive production-cost change Δc, let $\rho_h^M = \Delta_h p^M / \Delta c$ and $\rho_h^T = \Delta_h p^R / \Delta c$ denote the cumulative responses of wholesale and retail prices at horizon h. The manufacturer margin changes by $(\rho_h^M - 1)\Delta c$, while the total margin changes by $(\rho_h^T - 1)\Delta c$. Provided that the total margin remains positive, the manufacturer's share declines if and only if

$$\rho_h^M - 1 < \tau_0 \left(\rho_h^T - 1\right), \tag{14}$$

that is, whenever the manufacturer's margin grows proportionally less than the total margin. Preserving the initial share requires wholesale-price pass-through of $1 + \tau_0(\rho_h^T - 1)$: if retail prices rise by more than costs, the wholesale price must rise by more than costs as well, by a fraction τ_0 of the excess. Wholesale prices can therefore rise dollar for dollar with production costs and still leave the manufacturer with a smaller share of the margin.

We begin with evidence from the United States. Because we observe product-specific production costs, we can estimate pass-through in levels and express price responses in dollars per

Figure 10: Cost Pass-Through Along the Supply Chain



(a) Aggregate Components (b) Idiosyncratic Components (c) Production Costs to Retail Prices

Notes: This figure reports cumulative price responses in the United States over six months. In panels (a) and (b), the manufacturer series reports the response of wholesale prices to a permanent $1 increase in the corresponding component of production costs, while the retail series reports the response of retail prices to a permanent $1 increase in the corresponding component of wholesale prices. Panel (c) reports the directly estimated response of retail prices to aggregate and idiosyncratic production-cost increases. All specifications include product fixed effects and contemporaneous and six monthly lagged values of both components. Cumulative responses sum the coefficients through each horizon. Shaded areas show 90 and 95 percent confidence intervals calculated using product-level clustered standard errors.

unit. We decompose production costs into a category-month average, which we refer to as the aggregate component, and the product-specific deviation from that average, which we refer to as the idiosyncratic component. The distinction matters economically: a category-wide cost increase raises the costs of a product's closest substitutes as well, whereas a product-specific increase leaves them unchanged, and prices should respond differently to the two. Our specifications include product fixed effects and the contemporaneous value and six monthly lags of both components. At each horizon, we sum the contemporaneous and lagged coefficients to obtain the cumulative response to a permanent $1 increase in the corresponding cost component, and we cluster standard errors at the product level. We focus on the first six months to limit the loss of observations associated with longer lag structures.

We estimate three sets of responses: wholesale prices to production costs, retail prices to wholesale prices, and retail prices to production costs directly, decomposing wholesale prices into aggregate and idiosyncratic components in the same way. The direct response is the one we use in Equation (14). Under coordinated pricing, the wholesale price is set as a share of the retail margin, so the response of retail prices to wholesale prices mixes the response of the total margin with changes in its split, and combining the two stage-specific responses would not recover the response of retail prices to costs. Section F reports the coefficients for the direct response and additional results separating expected and unexpected costs.

Figure 10 reports the cumulative responses: the stage-specific relationships for the aggregate and idiosyncratic components in panels (a) and (b), and the direct response of retail prices to production costs in panel (c).

Following a $1 increase in aggregate production costs, the manufacturer raises wholesale prices by $0.20 on impact and by $0.99 after six months, while retail prices rise by $0.25 on impact and by $1.39 after six months. The manufacturer's margin is compressed at first and

gradually recovers its initial level after six months. Retail prices, however, rise by more than costs. This is consistent with consumers differing in their price sensitivity, as in the demand system of Section 4.2: a category-wide price increase pushes the most price-sensitive consumers to the outside option, the consumers who remain are less elastic, and the profit-maximizing markup on them is higher.[31] The total margin expands by about \$0.39 per dollar of cost, and the entire expansion accrues to retailers. The manufacturer keeps its margin in dollars and loses share of a larger total: with an initial share of 0.62, keeping that share would require wholesale prices to rise by \$1.24 per dollar of cost, well above the \$0.99 we estimate.[32]

Product-specific cost increases are absorbed upstream. After six months, wholesale prices have risen by \$0.05 per dollar of idiosyncratic cost and retail prices by \$0.20, and neither response is statistically distinguishable from zero. The retail price of a single product is set against substitutes whose costs have not moved, so raising it would mainly divert sales to those substitutes. Passing the increase into the wholesale price would instead shift the loss to the retailer, whose margin on an unchanged retail price would shrink, and the estimates indicate that the retailer largely resists this: in the point estimates, the manufacturer's margin falls by \$0.95 per dollar of cost after six months, while the retailer's margin rises by \$0.15. The manufacturer's share falls here as well, although the total margin contracts rather than expands. Keeping the share would require wholesale-price pass-through of \$0.51, far above the \$0.05 we estimate.

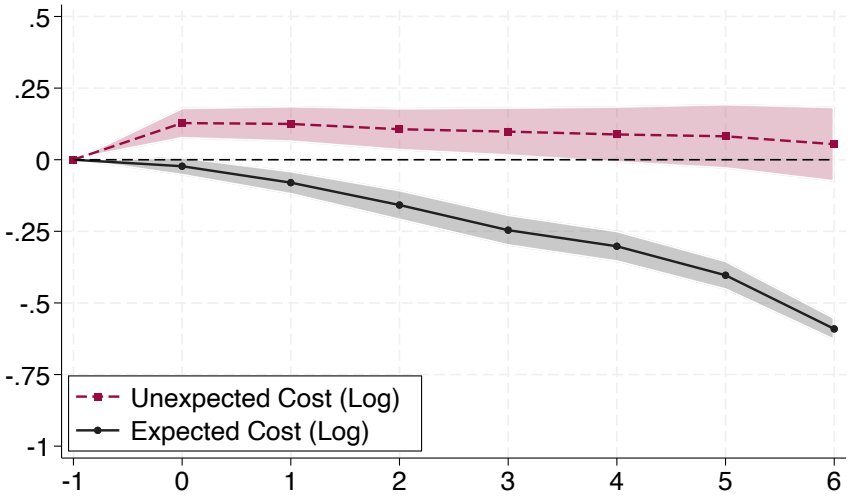
Retail prices nevertheless move with wholesale prices in both panels of Figure 10: a \$1 increase in aggregate wholesale prices is followed by a retail-price increase of \$0.50 on impact and \$1.45 after six months, and a \$1 increase in idiosyncratic wholesale prices by \$0.21 and \$0.54. Under coordinated pricing, these coefficients do not measure the transmission of costs to consumers. The wholesale price is $p^M = c + \tau(p^R - c)$, so a product whose retail price rises for demand reasons, with its cost unchanged, has a higher wholesale price as well, and a regression of retail on wholesale prices records the comovement as pass-through. This is why retail prices can move with idiosyncratic wholesale prices while barely responding to idiosyncratic costs, and why we rely on the direct response of retail prices to production costs for the consumer-price response.

We next ask whether the timing of this reallocation differs between cost changes the manufacturer expected and cost changes that surprised it. The distinction is separate from the one between aggregate and idiosyncratic components: expected and unexpected changes can each be shared across products or specific to one. Figures F.1 and F.2 show that wholesale and retail prices rise more on impact after expected aggregate cost increases than after unexpected ones.

[31] We return to the size of this response in Section 6, where the estimated model delivers a retail-price response of similar magnitude.

[32] The wholesale-price regressions include observations for which retail prices are unavailable. Comparisons combining these coefficients therefore provide an illustrative application of the margin-allocation condition rather than an exact decomposition on a common sample.

Figure 11: Cumulative Responses of Manufacturer Bargaining Weight



Notes: This figure reports cumulative responses of the log manufacturer bargaining weight to expected and unexpected production costs, using observations from the United States, Canada, Mexico, and the United Kingdom. Expected costs enter in logs, and unexpected costs are measured as the log ratio of realized to expected costs. Each response sums the contemporaneous and lagged coefficients through the indicated horizon. The specification controls for total margins, sales per capita, GDP per capita, and CPI. Shaded areas report 95 percent confidence intervals calculated using product-level clustered standard errors. Appendix Table A.7 reports the underlying regression coefficients.

By month six, the estimated retail-price response to unexpected costs overtakes the response to expected costs, while the wholesale-price response remains larger for expected costs.

To examine the margin-share dynamics, we extend the descriptive regressions in Table 4, using observations from all four countries. We include expected and unexpected costs separately, together with six monthly lags, and retain the controls for total margins, sales per capita, GDP per capita, and CPI.[33] Figure 11 reports the resulting cumulative responses of the manufacturer's bargaining weight.

Once lagged cost changes are included, the contemporaneous response to expected costs is nearly zero, while the lagged coefficients are negative and statistically significant. The cumulative coefficient reaches -0.59 after six months, close to the coefficient on expected costs in the specification without lags, which indicates that the contemporaneous estimate in the baseline regression captures the relationship between bargaining weights and both current and recent cost changes. Expected cost increases reallocate margins from the manufacturer to retailers gradually, over several months.

Unexpected cost changes generate a different pattern. Their contemporaneous relationship with the manufacturer's bargaining weight is positive, but this initial response quickly dissipates over the following months: after six months, the cumulative coefficient is -0.01. A surprise thus raises the manufacturer's margin share temporarily, whereas an expected change lowers

[33] Expected costs enter in logs. Unexpected costs are measured as $\log(c/c^e)$, where c^e denotes expected production costs. Appendix Table A.7 reports the regression coefficients.

it persistently. One possible explanation is that negotiations are conducted on the basis of expected costs. An increase that the manufacturer anticipated is already part of the negotiation when it materializes, so it does not move the split on impact; as it is transmitted into retail prices and the total margin over the following months, retailers gradually capture a larger part of the margin. A surprise arrives after retail prices have been set and after the retailer's own plans are in place, and the manufacturer can point to it to obtain a temporary concession in the split, which unwinds once the retailer has adjusted its own prices or renegotiated. Under the assumptions in Section 4, these movements in margin shares correspond to changes in bargaining weights.

These results distinguish changes in the total supply-chain margin from changes in its division between the manufacturer and retailers. Wholesale-price adjustment reflects both, so a stage-specific pass-through coefficient cannot be interpreted independently of the mechanism determining wholesale prices. Under coordinated pricing, movements in wholesale prices and bargaining weights describe the allocation of margins without determining the retail-pricing decision. The direct production-cost-to-retail estimates describe the consumer-price response, while the wholesale-price and bargaining-share estimates show how price adjustment is accompanied by margin reallocation.

6 Implications

6.1 Alternative Conduct

Our estimates in Section 4.2 indicate that the supply chain coordinates retail pricing: production cost, rather than the wholesale price, is the marginal cost in the retail-pricing problem. Under coordinated pricing, the wholesale price divides the total margin between the manufacturer and the retailer but does not affect retail prices or quantities. We examine the welfare consequences of this vertical conduct for consumers and firms by asking how outcomes would change under sequential pricing, when the wholesale price instead enters the retailer's marginal cost, thus having an allocative role, and the manufacturer and the retailer add margins at different stages of the supply chain. The counterfactual quantifies the cost of the double marginalization that the supply chain we observe avoids, in terms of prices, quantities, consumer surplus, and profits.

To implement the counterfactual, we replace the two-stage coordinated-pricing model with a Stackelberg pricing game. The manufacturer moves first and chooses wholesale prices to maximize its profits, anticipating the retailer's pricing response. After observing the wholesale prices, the retailer chooses retail prices, treating the wholesale prices as its marginal costs. Wholesale prices are therefore determined by manufacturer profit maximization rather than bargaining. We solve the game by backward induction, first characterizing the retailer's pricing response and then solving for the manufacturer's optimal wholesale prices.

We first compute the coordinated-pricing residual from observed retail prices, production costs, and the demand-implied margins evaluated at observed prices:

$$\widehat{\boldsymbol{\zeta}}_t^J \equiv \boldsymbol{p}_t^R - \boldsymbol{c}_t - \widehat{\boldsymbol{m}}_t^D\left(\boldsymbol{p}_t^R; \widehat{\theta}\right). \tag{15}$$

We hold this residual fixed across the counterfactuals. Let $\boldsymbol{r}_t(\boldsymbol{p}_t^M)$ denote the retailer's pricing response to a vector of wholesale prices. The retailer prices the manufacturer's products within a category jointly, as the price setter does in Section 4.2, and its response satisfies

$$\boldsymbol{r}_t(\boldsymbol{p}_t^M) - \boldsymbol{p}_t^M = \widehat{\boldsymbol{m}}_t^D\left(\boldsymbol{r}_t(\boldsymbol{p}_t^M); \widehat{\theta}\right) + \widehat{\boldsymbol{\zeta}}_t^J, \tag{16}$$

where $\widehat{\boldsymbol{\zeta}}_t^J$ is the residual from Equation (15).

Anticipating the retailer's response, the manufacturer chooses wholesale prices according to

$$\boldsymbol{p}_t^{M,S} = \arg\max_{\boldsymbol{p}_t^M}\left(\boldsymbol{p}_t^M - \boldsymbol{c}_t\right)' \boldsymbol{q}_t\left(\boldsymbol{r}_t(\boldsymbol{p}_t^M)\right), \qquad \boldsymbol{p}_t^{R,S} = \boldsymbol{r}_t\left(\boldsymbol{p}_t^{M,S}\right). \tag{17}$$

The manufacturer accounts for how each wholesale price affects retail prices and, through consumer demand, the quantities of all products in its portfolio. Relative to coordinated pricing, the counterfactual changes both the cost that enters the retailer's pricing problem and the way wholesale prices are set.

We solve the counterfactual separately for each product category-month in the two US categories used for demand estimation. We hold fixed production costs, the estimated demand system, market sizes, product availability, rival prices, and pricing residuals. Consumers can substitute toward rival products and the outside option, but rival prices do not respond. We aggregate prices using quantities in each pricing framework and calculate manufacturer, retailer, and supply-chain profits as $(p_{it}^M - c_{it})q_{it}$, $(p_{it}^R - p_{it}^M)q_{it}$, and $(p_{it}^R - c_{it})q_{it}$, respectively.

The results in Table 5 imply large costs from double marginalization. Moving from coordinated to sequential pricing increases the average wholesale price from 0.105 to 0.130, or by 24%, and the average retail price from 0.149 to 0.228, or by 53.7%. The retail-price increase reflects more than the increase in wholesale prices: the wholesale margin itself now enters the retailer's marginal cost, and the retailer adds its own margin on top, so consumers pay two markups instead of one. The resulting increase in retail prices reduces the quantity sold by the manufacturer by 57%.

The two stages of the supply chain bear this loss in sales differently. Total supply-chain profits fall by 23.8%: the higher prices obtained through separate upstream and downstream pricing do not compensate for the contraction in demand. Retailer profits decline by 3.8%, while manufacturer profits decline by 39.6%. The difference reflects what each stage prices. Under coordinated pricing, the retailer receives a share of the joint margin. Under sequential pricing, it sets a full markup over the wholesale price, and its margin per unit more than

Table 5: The Cost of Double Marginalization

	Data ($\hat{\psi} \approx 0$)	Sequential ($\psi = 1$)	% change
Average Prices			
Retail price p^R	0.149	0.228	53.7
Wholesale price p^M	0.105	0.130	24.0
Scale-masked outcomes			
Quantity	–	–	-57.0
Retailer profit	–	–	-3.8
Manufacturer profit	–	–	-39.6
Supply-chain profit	–	–	-23.8
Consumer surplus: total	–	–	-3.8
Consumer surplus: manufacturer products	–	–	-53.0
Welfare: manufacturer products	–	–	-41.9

Notes: This table compares observed outcomes under coordinated pricing with counterfactual outcomes under sequential pricing in the two US categories used for demand estimation. Prices are weighted by quantities in the corresponding pricing framework. Supply-chain profits sum manufacturer and retailer profits from the manufacturer's products. Total consumer surplus is calculated over the full choice set. Consumer surplus attributable to the manufacturer's products is the difference between total surplus and surplus when only rival products and the outside option are available. Welfare associated with the manufacturer's products sums this incremental consumer surplus and supply-chain profits; it excludes changes in rival profits. The counterfactual holds fixed demand, production costs, rival prices, market sizes, product availability, and pricing residuals. Pricing residuals are not subtracted from profits. Absolute levels of scale-sensitive outcomes are omitted to preserve confidentiality.

doubles on the reported averages, from 0.044 to 0.098, which nearly offsets its loss in volume. The manufacturer's margin per unit rises by much less, because it must anticipate the retailer's markup when choosing wholesale prices, so it bears most of the lost volume. The retailer therefore retains a larger share of the resulting smaller supply-chain profits.

We then calculate total consumer surplus over the full estimated choice set. We also report the consumer surplus attributable to the manufacturer's products, defined as the difference between total consumer surplus and the surplus consumers would obtain if only rival products and the outside option were available. This second measure captures the additional value consumers obtain from the availability of the manufacturer's products.

Total consumer surplus falls by 3.8%. This loss amounts to 53% of the baseline consumer surplus attributable to the availability of the manufacturer's products. The two figures express the same dollar loss relative to different bases: consumers switch toward rival products and the outside option, whose prices and utilities are unchanged, so the loss is small relative to all the surplus they obtain in the category and large relative to the surplus that the manufacturer's products add to it. The sum of consumer surplus attributable to the manufacturer's products and supply-chain profits falls by 41.9%. This measure excludes changes in rival profits and should not be interpreted as total market welfare.

Taken together, the results show that coordinated pricing benefits consumers and raises joint supply-chain profits relative to sequential pricing. These gains come from avoiding an additional vertical pricing distortion, not from eliminating market power in the product market.

6.2 An Import Tariff Increase

The tariff increases implemented during our sample motivate a second application of the model. In September 2018, the United States introduced additional tariffs of 10% on approximately $200 billion of imports from China, raising the rate to 25% in May 2019. A cost increase of this kind raises the question of who pays for it, consumers, the manufacturer, or retailers, and whether the answer depends on how the supply chain prices. Under coordinated pricing, production costs determine retail prices while wholesale prices divide the resulting profits, so the two parts of the question can be answered separately. We use the model to examine both following an illustrative 10% increase in production costs.[34]

We increase each product's production cost to $c'_{it} = 1.10c_{it}$ and solve for prices separately in each product category-month. We hold fixed the estimated demand system, rival prices, market sizes, product availability, and pricing residuals. Under coordinated pricing, we first determine retail prices using the joint-pricing condition with the higher production costs. Retail prices rise because the joint-profit-maximizing markup is now set over a higher cost, and they rise by more than the cost increase in dollar terms, for the reason discussed in Section 5: the increase drives the most price-sensitive consumers toward rival products and the outside option, and the markup on those who remain is higher. We then recover wholesale prices from the division of the resulting margins between the manufacturer and the retailer.

We allow the manufacturer's bargaining weight to adjust using the relationship between margin shares, production costs, and total margins reported in column (1) of Table 4.[35] This empirical rule determines how the cost increase reallocates profits between the two stages of the supply chain. With a fixed bargaining weight, manufacturer and retailer profits would both fall by roughly the same percentage as supply-chain profits; the rule is what moves the split. It does not affect retail prices or quantities under coordinated pricing, since wholesale prices do not enter the retail-pricing condition.

We then construct the sequential pricing counterfactual by solving the manufacturer's wholesale-pricing problem and the retailer's pricing response with the higher production costs. The manufacturer chooses wholesale prices anticipating how the retailer will adjust its prices and how consumers will respond. No bargaining adjustment is imposed. We compare outcomes after the cost increase with the no-shock baseline of the same pricing framework. The sequential-pricing baseline is therefore the counterfactual studied in the previous subsection, rather than the observed allocation.

Under coordinated pricing, the cost increase raises the average retail price by 5.2%, to 0.156,

[34] We do not calibrate the cost increase to the manufacturer's exposure to particular tariffs. A statutory tariff applies to the value of covered imports and need not generate an equivalent percentage increase in total production costs. The exercise also excludes tariff revenue and therefore does not provide a complete welfare assessment of trade policy.

[35] The adjustment uses $\Delta \log \tau_{it} = -0.496\Delta \log c_{it} + 0.571\Delta \log(p^R_{it} - c_{it})$. We apply this change to each product's initial margin share and bound the resulting weight within $(0, 1)$. The adjustment uses only the production-cost and total-margin coefficients, the contributions of the other predictors are unchanged.

Table 6: A 10% Cost Increase Under Different Pricing Frameworks

	Joint ($\psi = 0$)		Sequential ($\psi = 1$)	
	Level	% change	Level	% change
Average Prices				
Retail price p^R	0.156	5.2	0.238	4.2
Wholesale price p^M	0.109	4.1	0.136	5.1
Outcomes: % change vs relevant baseline				
Quantity	–	-7.8	–	-8.0
Retailer profit	–	-0.6	–	-5.2
Manufacturer profit	–	-6.8	–	-5.6
Supply-chain profit	–	-4.1	–	-5.4
Consumer surplus: total	–	-0.3	–	-0.3
Consumer surplus: manufacturer products	–	-4.8	–	-8.9
Welfare: manufacturer products	–	-4.5	–	-7.1

Notes: This table reports outcomes following a 10% increase in every product's production cost under coordinated and sequential pricing. Percentage changes are relative to each framework's own no-shock baseline. Prices are weighted by quantities in the corresponding scenario. Under coordinated pricing, the manufacturer's bargaining weight adjusts using the production-cost and total-margin coefficients in column (1) of Table 4. Under sequential pricing, wholesale prices follow from manufacturer profit maximization. Demand, rival prices, market sizes, product availability, and pricing residuals are held fixed. Profits use the higher production costs without subtracting pricing residuals. Consumer surplus and welfare associated with the manufacturer's products are defined as in Table 5; the latter excludes changes in rival profits and tariff revenue.

and the average wholesale price by 4.1%, to 0.109. The higher retail prices reduce the quantity sold by the manufacturer by 7.8%, and total supply-chain profits decline by 4.1%. The losses are unevenly distributed: manufacturer profits fall by 6.8%, while retailer profits decline by only 0.6%. The cost increase therefore falls on the manufacturer, whose share of a smaller margin shrinks under the estimated bargaining response, while retailers are almost untouched. Consumers pay the same retail prices as they would under any other split of the margin.

Under sequential pricing, the average wholesale price increases by 5.1%, to 0.136, and the average retail price increases by 4.2%, to 0.238. The percentage increase in retail prices is smaller than under coordinated pricing, but it applies to a substantially higher initial price: on the reported averages, retail prices rise by about 0.010 per unit under sequential pricing and by about 0.007 under coordinated pricing, so consumers face higher prices after the shock under sequential pricing. A smaller percentage response to costs does not, by itself, imply lower prices or better consumer outcomes. Quantities fall by 8.0%, and supply-chain profits decline by 5.4%. The profit losses are more evenly distributed across stages, because each stage re-optimizes its own markup over the higher cost: retailer profits fall by 5.2% and manufacturer profits by 5.6%.

Consumers can respond to the higher prices by switching toward rival products whose prices remain unchanged. Total consumer surplus declines by approximately 0.3% under either pricing framework. Consumer surplus attributable to the manufacturer's products falls by 4.8% under coordinated pricing and 8.9% under sequential pricing. Within each framework, these measures express the same absolute consumer-surplus loss relative to different baseline levels. The sum of consumer surplus attributable to the manufacturer's products and supply-chain profits declines

by 4.5% and 7.1%, respectively. Because each framework starts from a different allocation, these percentage changes should not be interpreted as a comparison of absolute welfare losses.

We compare the model's retail-price response with the six-month reduced-form response to aggregate production costs estimated in Section 5.2. This comparison uses the direct pass-through estimate of production costs into retail prices, rather than combining the stage-specific pass-through estimates. Let C_{gt} denote the average production cost in narrow category g and month t. We consider a 10% increase in this aggregate component, holding each product's deviation from the category average fixed. The estimated six-month aggregate pass-through coefficient is 1.392 (column (7) of Table F.1), implying a predicted retail-price change of $\widehat{\Delta_6 p_{it}^R} = 1.392 \times 0.10 C_{gt}$. Aggregating these changes using baseline quantities gives a 4.89% increase in the average retail price, with a 95% confidence interval of $[4.02\%, 5.75\%]$.[36] A dollar pass-through above one and a percentage response of about 5% are consistent because the category-average component of production costs is about 35% of the retail price in our sample: a coefficient of 1.39 corresponds to an elasticity of retail prices with respect to costs of about 0.49.

The model-implied retail-price increase of 5.2% is close to this six-month response and lies within its confidence interval.[37] We read this as a non-targeted empirical validation on the size of the model's price response since the reduced-form estimate was not used in estimation.

The two exercises make different points about vertical conduct. Conduct matters especially for the level of retail prices, which sequential pricing would raise by more than half. It matters less for the percentage pass-through of a cost increase, which is similar under the two frameworks, but it crucially determines who within the supply chain is affected: the manufacturer under coordinated pricing with the estimated bargaining response, and both stages in similar proportion under sequential pricing.

7 Conclusion

This study analyzes markups and pricing strategies along the supply chain in the United States, the United Kingdom, Canada, and Mexico. We assembled a unique dataset combining detailed product price and cost information from a major global manufacturer with corresponding retail prices collected online. With data from around 2,000 products, we measure markups for both the manufacturer and retailers without relying on strong assumptions.

Our findings reveal that total markups are similar in magnitude across countries, averaging around 0.65, but there is significant heterogeneity along the supply chain. This indicates

[36] The percentage response is calculated as $100 \sum_{i,t} q_{it} \widehat{\Delta_6 p_{it}^R} / \sum_{i,t} q_{it} p_{it}^R$. The confidence interval uses the product-clustered standard error of the cumulative aggregate-cost coefficient and conditions on baseline costs and quantities.

[37] Using the twelve-month coefficient of 1.681 from column (8) of Table F.1 gives a 5.9% increase, with a confidence interval of approximately $[4.6\%, 7.2\%]$.

substantial market power at both the manufacturer and the retailer levels and suggests that economic models that assume a single sector or perfect competition in upstream or downstream sectors may not be well-suited to capture the richness of markup patterns in the real world.

Total markups are stable, and the dollar margins are shared between the manufacturer and retailers. We do not find evidence that inflation is driven by higher markups; in the time series, higher prices are driven by higher costs and total markups are fairly stable. These findings are not consistent with the "greedflation" hypothesis, which suggests that in recent years inflation was driven by firms increasing their markups. Total markup stability over time is due to the negative correlation along the supply chain, where adjustments in the manufacturing and retail sectors are asymmetric and partially offset each other.

We develop a two-stage supply chain pricing model to rationalize these findings and provide a testing framework for vertical conduct. Our model fits key facts in our data that are not readily explained by standard models in the industrial organization and macroeconomic literature. This model can be adapted for future research on how supply chain relationships influence various economic phenomena.

Our model allows us to quantify the relative bargaining power between the manufacturer and retailers. The analysis reveals that the manufacturer holds greater bargaining power than retailers in the US, while retailers have greater bargaining power than the manufacturer in other countries. Over time, the manufacturer's bargaining power has declined in the United States and Canada, remained stable in the United Kingdom, and increased significantly in Mexico in 2020.

Using variation across products, countries, and time, we analyze different predictors of bargaining power. We find that the manufacturer bargaining weights increase with income and market penetration and decrease with production costs and inflation. Manufacturer bargaining weights also increase with the total margin of a product, indicating the manufacturer captures a greater share when fewer close substitutes are available.

Finally, we show that the manufacturer and the retailers adjust their prices at different rates and magnitudes. Differences in the pass-through along the supply chain have significant implications for the transmission of various shocks. Idiosyncratic shocks generate little price adjustment: the manufacturer does not respond, and retailers adjust only partially and slowly. As a result, product-specific cost changes are transmitted only partially to final prices. In contrast, aggregate shocks are passed gradually, but long-run pass-through is larger as both the manufacturer and retailers eventually fully pass on these costs. Our model of supply chain pricing adds nuance to the dynamics of pass-through, since changes in wholesale prices—which enter both retailer and manufacturer pass-through regressions—may reflect dynamic changes in bargaining power.

From a policy perspective, these findings underscore the importance of considering supply chain dynamics when addressing inflation and market power. The differing speeds at which cost

changes propagate through the supply chain can delay consumer price responses to supply-side shocks. Furthermore, shifts in economic conditions can reshape the bargaining positions of suppliers and buyers, affecting price setting and market outcomes. Accounting for these factors is crucial to developing effective strategies that promote price stability and competitive markets.

One clear topical application of our findings is trade policy, especially in light of ongoing US trade tensions and proposed tariffs on a range of imports. Our results suggest that the impact of tariffs on consumer prices may depend on how predictably and broadly they are applied. When the tariffs affect entire product categories and are more expected, the manufacturer and retailers tend to adjust prices more quickly, leading to faster—and complete—pass-through. In contrast, if tariffs are product-specific or more uncertain, price adjustments may tend to be slower. Markups may shift across stages of the supply chain over time: initially falling for manufacturers as import costs rise, and later for retailers as wholesale and retail prices gradually adjust.

References

Aguirregabiria, V. (1999). The dynamics of markups and inventories in retailing firms. *The Review of Economic Studies*, 66(2):275–308.

Alvarez, F., Lippi, F., and Oskolkov, A. (2022). The macroeconomics of sticky prices with generalized hazard functions. *The Quarterly Journal of Economics*, 137(2):989–1038.

Alviarez, V. I., Fioretti, M., Kikkawa, K., and Morlacco, M. (2023). Two-sided market power in firm-to-firm trade. Technical report, National Bureau of Economic Research.

Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402.

Anderson, E., Rebelo, S., and Wong, A. (2025). Markups across space and time. *Review of Economic Studies*.

Atalay, E., Frost, E., Sorensen, A. T., Sullivan, C. J., and Zhu, W. (2025). Scalable demand and markups. *Journal of Political Economy*.

Auer, R., Burstein, A., and Lein, S. M. (2021). Exchange rates and prices: evidence from the 2015 swiss franc appreciation. *American Economic Review*, 111(2):652–686.

Auer, R. A., Chaney, T., and Sauré, P. (2018). Quality pricing-to-market. *Journal of International Economics*, 110:87–102.

Avignon, R., Chambolle, C., Guigue, E., and Molina, H. (2025). Markups, markdowns, and bargaining in a vertical supply chain.

Bils, M. (1987). The cyclical behavior of marginal cost and price. *The American Economic Review*, pages 838–855.

Bils, M. and Kahn, J. A. (2000). What inventory behavior tells us about business cycles. *American Economic Review*, 90(3):458–481.

Bils, M., Klenow, P. J., and Malin, B. A. (2018). Resurrecting the role of the product market wedge in recessions. *American Economic Review*, 108(4-5):1118–1146.

Bilyk, O., Grieder, T., and Khan, M. (2023). Markups and inflation during the covid-19 pandemic. Technical report, Bank of Canada.

Bonadio, B., Fischer, A. M., and Sauré, P. (2020). The speed of exchange rate pass-through. *Journal of the European Economic Association*, 18(1):506–538.

Brand, J. (2021). Differences in differentiation: Rising variety and markups in retail food stores. *Available at SSRN 3712513*.

Burstein, A., Carvalho, V. M., and Grassi, B. (2020). Bottom-up markup fluctuations. Technical report, National Bureau of Economic Research.

Cavallo, A. (2013). Online and official price indexes: Measuring argentina's inflation. *Journal of Monetary Economics*, 60(2):152–165.

Cavallo, A. (2017). Are online and offline prices similar? evidence from large multi-channel retailers. *American Economic Review*, 107(1):283–303.

Cavallo, A. (2018). Scraped data and sticky prices. *Review of Economics and Statistics*, 100(1):105–119.

Cavallo, A. and Kryvtsov, O. (2023). What can stockouts tell us about inflation? evidence from online micro data. *Journal of International Economics*, page 103769.

Cavallo, A., Lippi, F., and Miyahara, K. (2024). Large shocks travel fast. *American Economic Review: Insights*.

Cavallo, A. and Rigobon, R. (2016). The billion prices project: Using online prices for measurement and research. *Journal of Economic Perspectives*, 30(2):151–178.

De Loecker, J. and Eeckhout, J. (2018). Global market power. Technical report, National Bureau of Economic Research.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

De Loecker, J. and Scott, P. (2022). Markup estimation using production and demand data. an application to the us brewing industry. Technical report, Working paper.

de Roos, N. (2011). Markup uniformity under the multi-nomial logit. *Available at SSRN 1804735*.

Demirer, M. and Rubens, M. (2025). Welfare effects of buyer and seller power. Technical report, National Bureau of Economic Research.

Díez, F. J., Fan, J., and Villegas-Sánchez, C. (2021). Global declining competition? *Journal of International Economics*, 132:103492.

Döpper, H., MacKay, A., Miller, N., and Stiebale, J. (2025). Rising markups and the role of consumer preferences. *Journal of Political Economy*.

Duarte, M., Magnolfi, L., Sølvsten, M., and Sullivan, C. (2024). Testing firm conduct. *Quantitative Economics*, 15(3):571–606.

Eichenbaum, M., Jaimovich, N., and Rebelo, S. (2011). Reference prices, costs, and nominal rigidities. *American Economic Review*, 101(1):234–262.

Flaaen, A. B., Hortaçsu, A., Tintelnot, F., Urdaneta, N., and Xu, D. (2025). Who pays for tariffs along the supply chain? evidence from european wine tariffs. Technical report, National Bureau of Economic Research.

Galeotti, M. and Schiantarelli, F. (1998). The cyclicality of markups in a model with adjustment costs: econometric evidence for us industry. *Oxford Bulletin of Economics and Statistics*, 60(2):121–142.

Gopinath, G., Gourinchas, P.-O., Hsieh, C.-T., and Li, N. (2011). International prices, costs, and markup differences. *American Economic Review*, 101(6):2450–2486.

Hahn, N. (2024). Who is in the driver's seat? markups, markdowns, and profit sharing in the car industry. *ZEW-Centre for European Economic Research Discussion Paper*, (24-047).

Koujianou Goldberg, P. and Hellerstein, R. (2013). A structural approach to identifying the sources of local currency price stability. *Review of Economic Studies*, 80(1):175–210.

Kryvtsov, O. and Midrigan, V. (2013). Inventories, markups, and real rigidities in menu cost models. *Review of Economic Studies*, 80(1):249–276.

Leduc, S., Li, H., and Liu, Z. (2024). Are markups driving the ups and downs of inflation? *FRBSF Economic Letter*, 2024(12):1–5.

Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the MillerCoors joint venture. *Econometrica*, 85(6):1763–1791.

Minton, R. and Wheaton, B. (2023). Delayed inflation in supply chains: Theory and evidence. *Available at SSRN 4470302*.

Nakamura, E. (2008). Pass-through in retail and wholesale. *American Economic Review*, 98(2):430–437.

Nakamura, E. and Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, 77(3):1192–1230.

Nekarda, C. J. and Ramey, V. A. (2020). The cyclical behavior of the price-cost markup. *Journal of Money, Credit and Banking*, 52(S2):319–353.

Rotemberg, J. J. and Woodford, M. (1999). The cyclical behavior of prices and costs. *Handbook of macroeconomics*, 1:1051–1135.

Sangani, K. (2022). Markups across the income distribution: Measurement and implications. *Available at SSRN 4092068*.

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

Toxvaerd, F. (2024). Bilateral monopoly revisited: Price formation, efficiency and countervailing powers. Technical report, CEPR Discussion Papers.

Vaona, A. (2016). A nonparametric panel data approach to the cyclical dynamics of price-cost margins in the fourth kondratieff wave. *Eurasian Business Review*, 6:155–170.

Villas-Boas, S. B. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies*, 74(2):625–652.

Appendix

A Additional Tables and Figures

Table A.1: The Matched Manufacturer-Retailer Dataset

	Main Sample						Expanded Sample	
	Observations	% Sales	% Products	Products	Brands	Retailers	Manufacturer	Retail
US	24,238	0.85	0.64	971	8	16	40,603	24,238
UK	7,713	0.84	0.59	406	5	7	12,275	7,997
CA	6,618	0.55	0.34	309	6	12	21,987	6,943
MX	3,889	0.60	0.47	232	9	5	9,122	4,273
All	42,458	0.70	0.53	1,918	16	40	83,987	43,451

Notes: This table reports summary statistics for the main sample and for the expanded sample used in the pass-through analysis in Section 5.2. For each country and for the pooled sample, we report the number of matched product-period observations, the share of the manufacturer's sales covered by matched products, the share of products with a retail match, and the number of unique products, brands, and retailers. The final columns refer to the expanded sample and report the number of product-periods with available manufacturer data and the corresponding number with available retail prices.

Table A.2: Summary Statistics of Costs, Prices and Quality

	Mean	p25	p75	SD	CV_i	N
Retail Price	0.214	0.116	0.295	0.11	0.08	42,458
Manufacturer Price	0.133	0.068	0.183	0.07	0.08	42,458
Cost	0.068	0.039	0.093	0.04	0.13	42,458
Cost Breakdown						
Expected	0.066	0.037	0.090	0.03	0.09	39,123
Unexpected	0.002	-0.001	0.006	0.01	0.10	39,123
Quality Indicators						
Differentiation	68.23	65.84	71.23	3.13	-	668
Perceived Quality	86.14	85.20	87.44	1.92	-	668

Notes: This table reports summary statistics for costs, prices, and quality in the matched sample from July 2018 to June 2023 (with cost breakdown data available starting in January 2019). All prices and costs are put in standard category-level units. N denotes the number of product-month observations for prices and costs, or the number of products with available quality information. All prices and costs are converted to January 2019 US dollars. "SD" reports the standard deviation across all observations. As a measure of time-series variation, we compute the coefficient of variation for each product across time and report the average across products as CV_i. Since unexpected costs have a mean close to zero, we report coefficients of variation for both expected and unexpected costs relative to total costs to aid interpretation. Summary statistics for expected and unexpected costs are based on fewer observations due to their later start date. All variables are winsorized at the 1% level. Quality information is available only for a subset of products sold in the US

Table A.3: Imputed Data

	Manufacturer Data				Retail Data			
	Observations		Regression Obs. (6M)		Observations		Regression Obs. (6M)	
	No Fill	Fill	No Fill	Fill	No Fill	Fill	No Fill	Fill
UK	11,549	12,275	6,712	8,352	7,713	7,997	4,868	5,522
CA	19,817	21,987	9,693	15,586	6,618	6,943	3,581	4,607
MX	7,788	9,122	3,102	5,563	3,889	4,273	1,752	2,537
All	79,757	83,987	46,831	56,825	42,458	43,451	26,541	29,006

Notes: This table reports the number of product-period observations for all products ("Observations") and for those continuously observed over a six-month period ("Regression Obs. (6M)"). For each group, we show the raw number of observations and the total after imputing missing values by carrying forward production costs and wholesale prices to fill gaps of up to 90 days in a product's time series. Row "All" includes also the observations from the US, where no observations were filled in the manufacturer dataset.

Table A.4: Summary Statistics For Markups and Margin Shares

Variable	Mean	SD	p10	p25	p50	p75	p90
Country: United States							
Total Markup	0.67	0.07	0.58	0.64	0.68	0.71	0.74
Manufacturer Markup	0.54	0.09	0.42	0.49	0.55	0.60	0.63
Retail Markup	0.28	0.11	0.15	0.20	0.26	0.34	0.43
Product Markup	0.41	0.06	0.34	0.37	0.41	0.44	0.48
Manufacturer Margin Share	0.59	0.15	0.37	0.51	0.60	0.69	0.77
Retail Margin Share	0.41	0.15	0.23	0.31	0.40	0.49	0.63
Country: Canada							
Total Markup	0.65	0.11	0.51	0.59	0.65	0.72	0.77
Manufacturer Markup	0.46	0.14	0.28	0.37	0.46	0.54	0.62
Retail Markup	0.35	0.11	0.21	0.29	0.35	0.42	0.49
Product Markup	0.40	0.09	0.29	0.35	0.40	0.46	0.51
Manufacturer Margin Share	0.46	0.16	0.25	0.36	0.46	0.56	0.65
Retail Margin Share	0.54	0.16	0.35	0.44	0.54	0.64	0.75
Country: Mexico							
Total Markup	0.60	0.12	0.46	0.53	0.60	0.67	0.75
Manufacturer Markup	0.32	0.16	0.15	0.23	0.30	0.40	0.52
Retail Markup	0.42	0.10	0.27	0.36	0.43	0.49	0.53
Product Markup	0.37	0.10	0.26	0.30	0.36	0.42	0.49
Manufacturer Margin Share	0.31	0.14	0.14	0.22	0.30	0.39	0.49
Retail Margin Share	0.69	0.14	0.51	0.61	0.70	0.78	0.86
Country: United Kingdom							
Total Markup	0.64	0.07	0.54	0.59	0.64	0.68	0.72
Manufacturer Markup	0.36	0.12	0.20	0.29	0.38	0.45	0.50
Retail Markup	0.42	0.09	0.31	0.37	0.42	0.48	0.54
Product Markup	0.39	0.06	0.31	0.35	0.39	0.43	0.47
Manufacturer Margin Share	0.33	0.13	0.17	0.25	0.34	0.42	0.49
Retail Margin Share	0.67	0.13	0.51	0.58	0.66	0.75	0.83

Notes: This table reports summary statistics of the sales-weighted markup and margin share distributions in each country. Markups are computed as the Lerner index using retail prices and production costs for the total indicator ($\frac{p^R-c}{p^R}$), wholesale prices and production costs for the manufacturer indicator ($\frac{p^M-c}{p^M}$), and retail prices and wholesale prices for the retail indicators ($\frac{p^R-p^M}{p^R}$). Product markup is computed as the average between the manufacturer markup and the retail markup. Margin shares indicate the percentage of the total $ margin available per product that the manufacturer ($\frac{p^M-c}{p^R-c}$) and the retailer ($\frac{p^R-p^M}{p^R-c}$) retain.

Table A.5: Markup Correlation Along the Supply Chain

	Aggregate	Time FEs	Product FEs	Residuals
US	-0.34	-0.70	-0.41	-0.12
CA	-0.09	-0.29	0.02	-0.13
MX	0.05	0.33	-0.01	-0.07
UK	-0.30	-0.24	-0.27	-0.35

Notes: This table reports correlation coefficients between manufacturer and retail markups. Column (1) presents the aggregate (raw) correlation, while Columns (2)-(4) reports the correlation between the different markup components resulting from equation (2) across the two stages of the supply chain. Regressions and correlations are weighted by per-period sales.

Table A.6: Margins and Alternative Quality Measure (Product Differentiation)

	Margins			Markup		
	Total	Manufacturer	Retail	Total	Manufacturer	Retail
Differentiation	11.181***	18.577***	3.484	0.345	3.528***	-7.088***
	(1.095)	(1.122)	(2.130)	(0.342)	(0.667)	(1.226)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.79	0.81	0.56	0.37	0.31	0.29
Observations	16,476	16,476	16,476	16,476	16,476	16,476

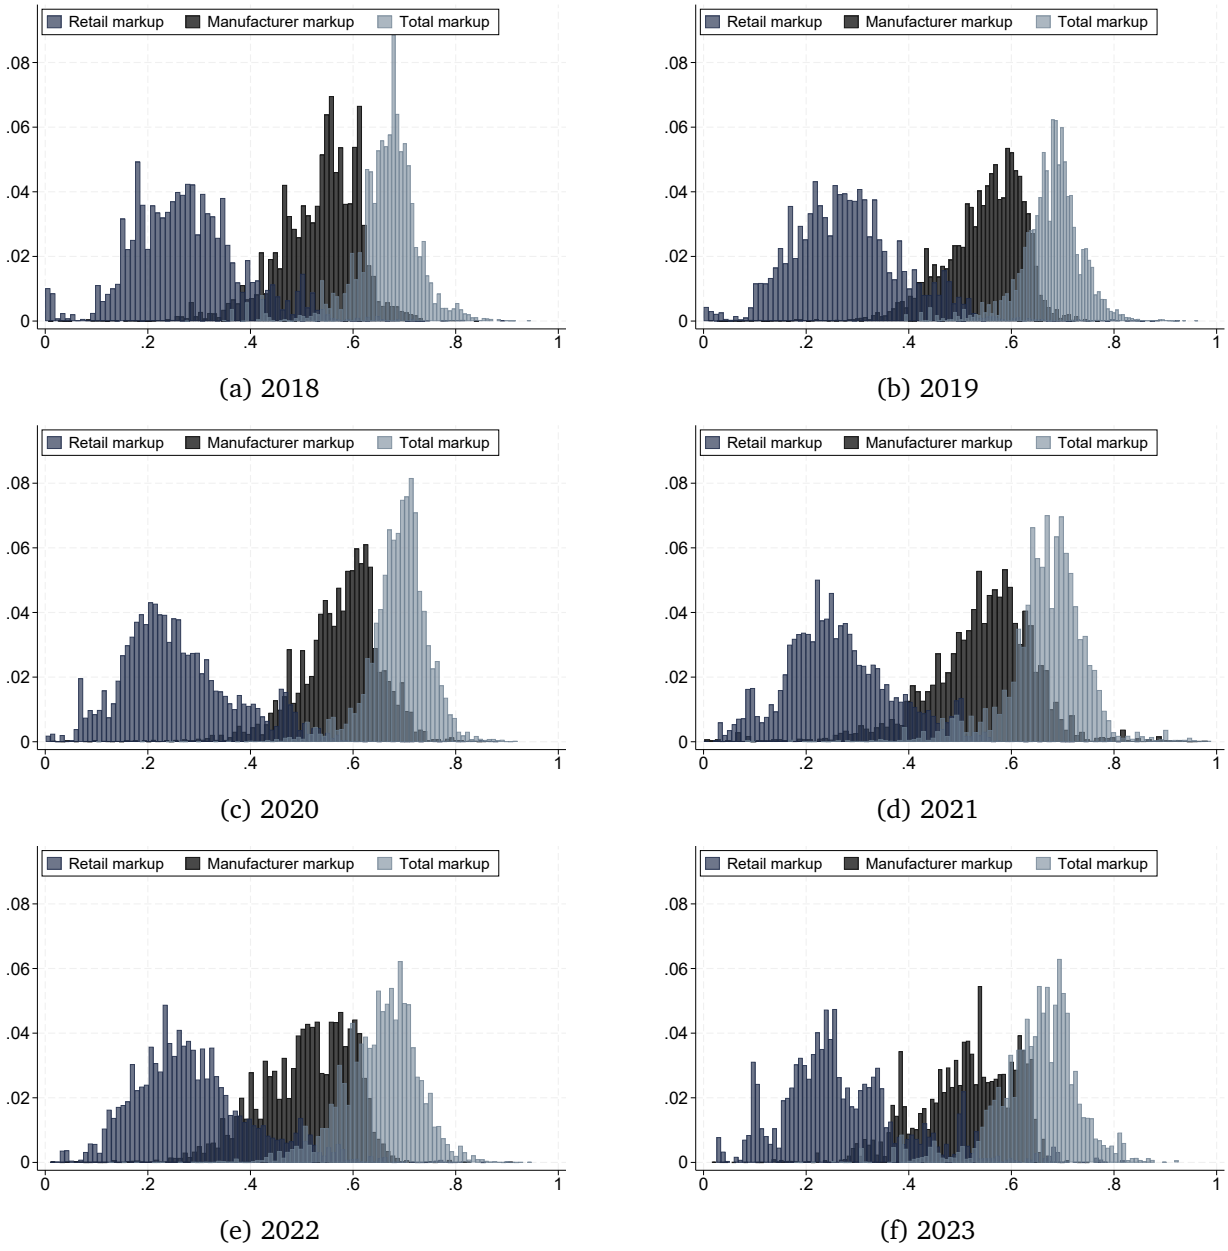
Notes: This table reports the linear relation between margins or markups and (standardized) product quality estimated from the following linear model: $y_{it} = a + \alpha q_i + \beta X_{it} + \phi_{it} + \varepsilon_{it}$, with y_{it} being the variable of interest, margins or markups, for product i at time t, q_i (standardized) product quality, X_{it} controls including linear and quadratic package size to account for quantity discounts, ϕ_{it} fixed effects including product category interacted with time, and ε_{it} a mean zero error. Coefficients and standard errors reported as percent relative to the mean of the dependent variable. Product-level clustered standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Bargaining Weight Predictors with Dynamics

	(1)	(2)	(3)	(4)
Production Cost	-0.496***			
	(0.014)			
Total Margin	0.571***	0.611***	0.622***	0.629***
	(0.013)	(0.015)	(0.015)	(0.017)
Sales per Capita	0.049***	0.048***	0.047***	0.051***
	(0.004)	(0.004)	(0.004)	(0.004)
GDP per Capita	0.412***	0.339***	0.282***	0.235***
	(0.029)	(0.031)	(0.032)	(0.036)
CPI	-0.168***	-0.388***	-0.402***	-0.547***
	(0.058)	(0.058)	(0.062)	(0.079)
Expected Cost$_t$		-0.554***	-0.157***	-0.023
		(0.016)	(0.015)	(0.016)
Expected Cost$_{t-1}$			-0.168***	-0.057***
			(0.013)	(0.013)
Expected Cost$_{t-2}$			-0.252***	-0.078***
			(0.016)	(0.013)
Expected Cost$_{t-3}$				-0.088***
				(0.011)
Expected Cost$_{t-4}$				-0.056***
				(0.009)
Expected Cost$_{t-5}$				-0.101***
				(0.012)
Expected Cost$_{t-6}$				-0.187***
				(0.018)
Unexpected Cost$_t$		-0.088***	0.018	0.064***
		(0.024)	(0.016)	(0.014)
Unexpected Cost$_{t-1}$			-0.024*	-0.003
			(0.014)	(0.009)
Unexpected Cost$_{t-2}$			-0.053***	-0.018*
			(0.015)	(0.010)
Unexpected Cost$_{t-3}$				-0.008
				(0.009)
Unexpected Cost$_{t-4}$				-0.010
				(0.010)
Unexpected Cost$_{t-5}$				-0.007
				(0.011)
Unexpected Cost$_{t-6}$				-0.028**
				(0.013)
Observations	42458	39072	32678	23941
R^2	0.672	0.689	0.691	0.689

Notes: This table reports the outcome of regressing the (log) manufacturer bargaining weight on (log) total margins, (log) product sales, (log) GDP per capita, (log) CPI, present and lagged (log) production costs. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: Markups Along the Supply Chain by Year (United States)



(a) 2018

(b) 2019

(c) 2020

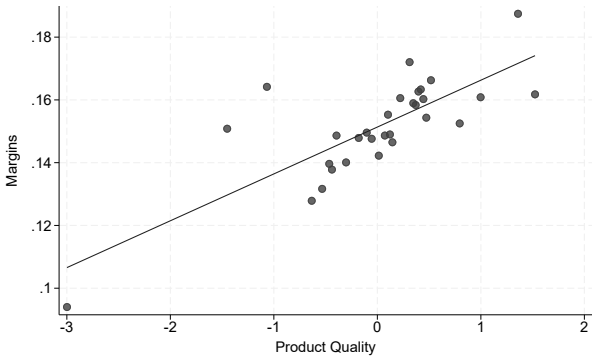
(d) 2021

(e) 2022

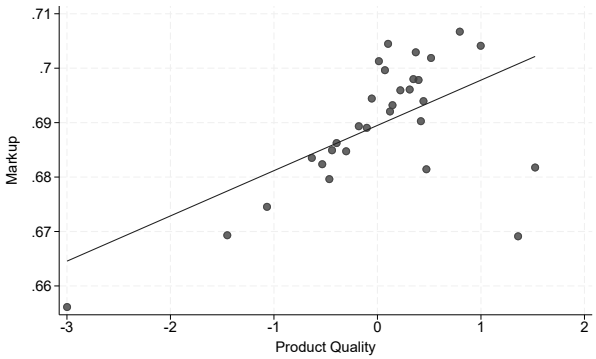
(f) 2023

Notes: This figure shows the sales-weighted frequency distribution of markups along the supply chain in the United States by year. Retail markups ($\frac{p^R - p^M}{p^R}$) are shown in dark blue, manufacturer markups ($\frac{p^M - c}{p^M}$) in dark grey, and total markups ($\frac{p^R - c}{p^R}$) in light blue.

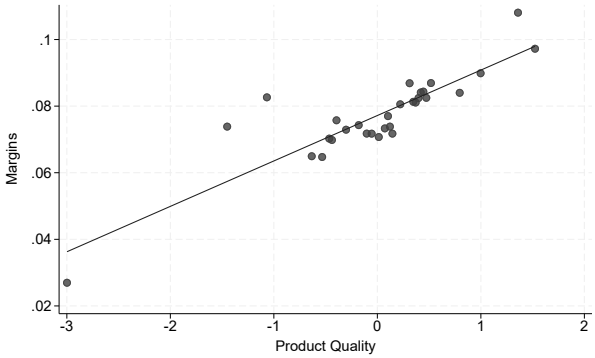
Figure A.2: Quality Differentiation Along the Supply Chain (Perceived Quality)



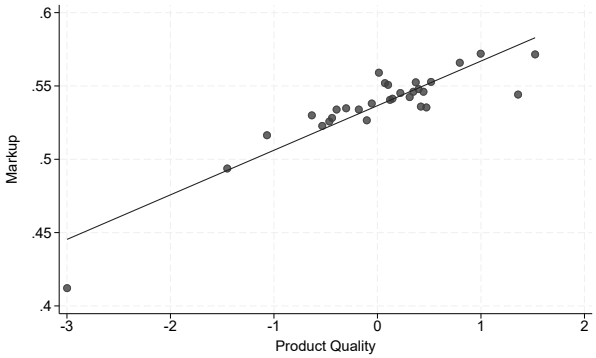
(a) Total Margins



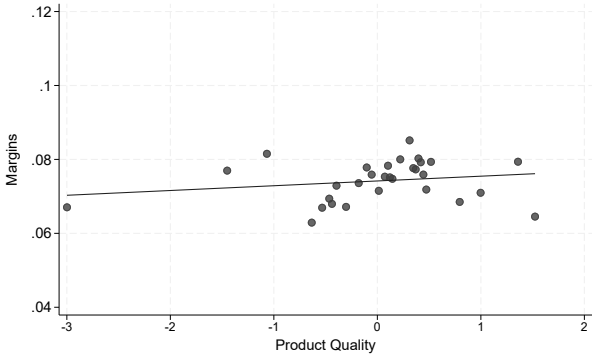
(b) Total Markups



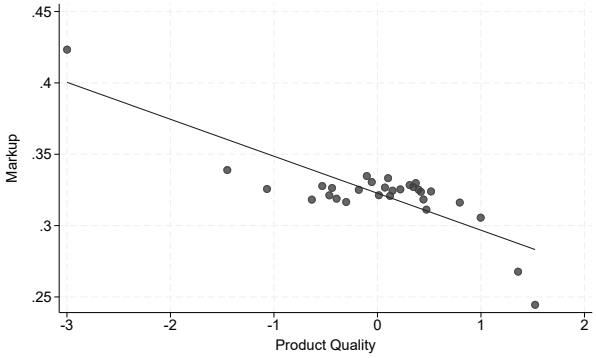
(c) Manufacturer Margins



(d) Manufacturer Markups



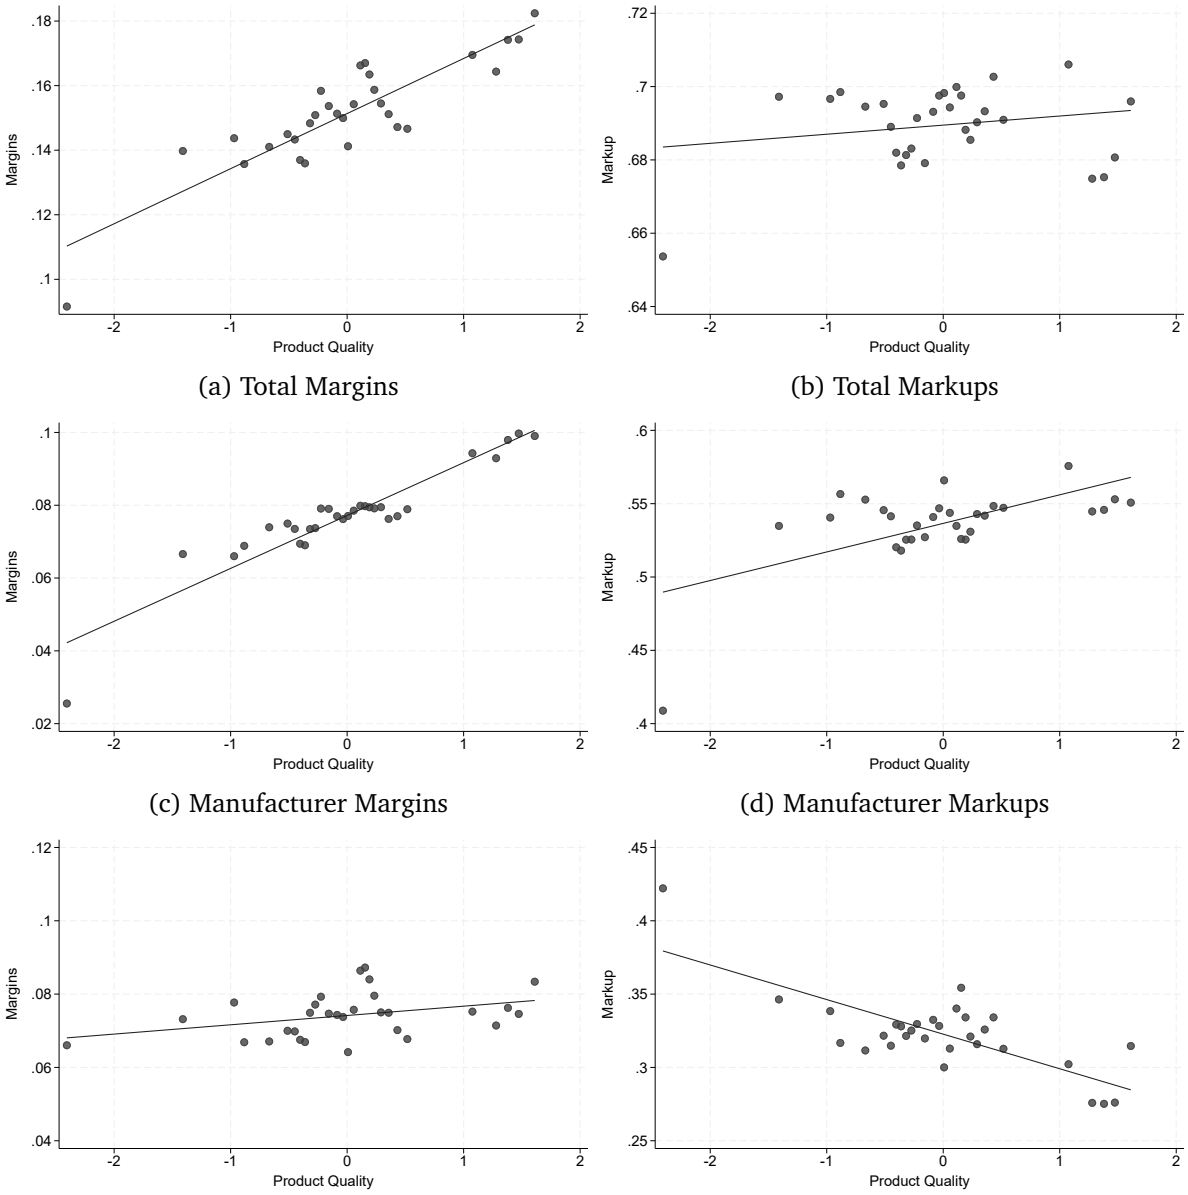
(e) Retail Margins



(f) Retail Markups

Notes: This figure shows the relation of product (a) total margins, (b) total markups, (c) manufacturer margins, (d) manufacturer markups, (e) retail margins, and (f) retail markups to a (standardized) measure of perceived quality. Values are residualized on linear and quadratic package size, to account for quantity discounts, and on a set of fixed effects including product category interacted with periods.

Figure A.3: Quality Differentiation Along the Supply Chain (Product Differentiation)



(a) Total Margins

(b) Total Markups

(c) Manufacturer Margins

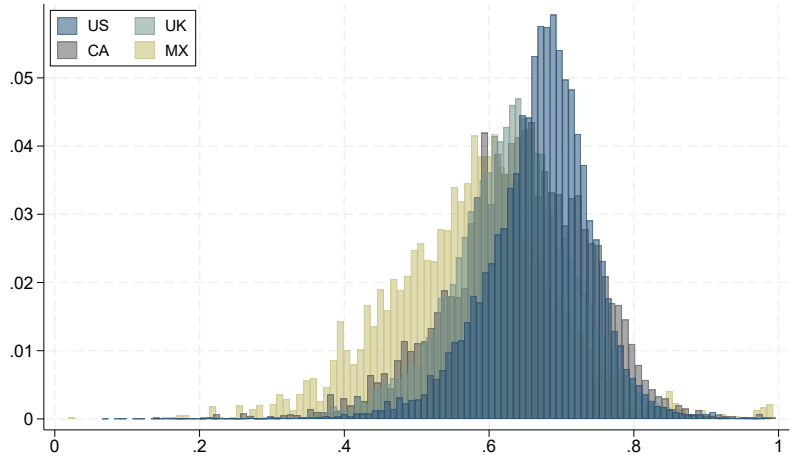
(d) Manufacturer Markups

(e) Retail Margins

(f) Retail Markups

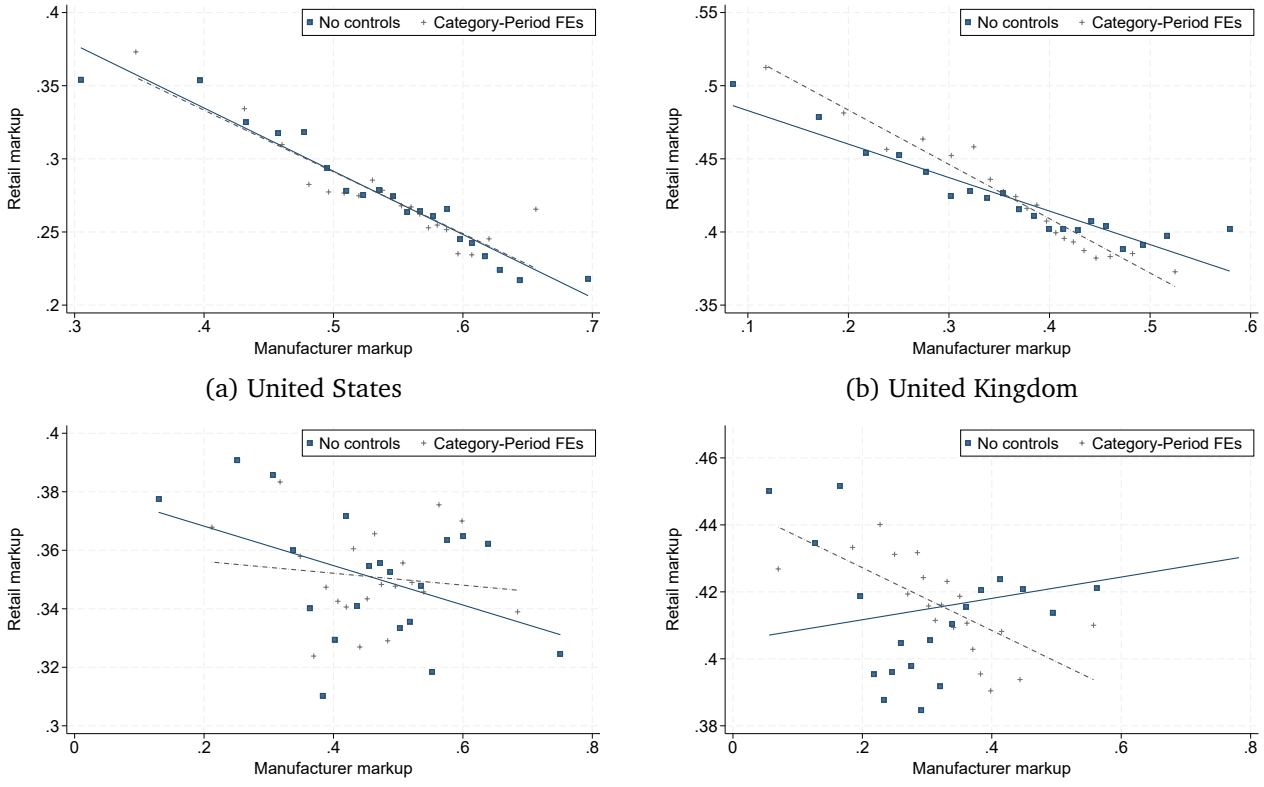
Notes: This figure shows the relation of product (a) total margins, (b) total markups, (c) manufacturer margins, (d) manufacturer markups, (e) retail margins, and (f) retail markups to a (standardized) measure of product differentiation. Values are residualized on linear and quadratic product size, to account for quantity discounts, and product category interacted with periods.

Figure A.4: Total Markups Without Retail Price Adjustment



Notes: This figure shows the sales-weighted frequency distribution of total markups ($\frac{p^R - c}{p^R}$) along the supply chain for each country, using multiple retail prices for each product when observed at different retailers.

Figure A.5: Manufacturer and Retailer Markup Relation



(a) United States

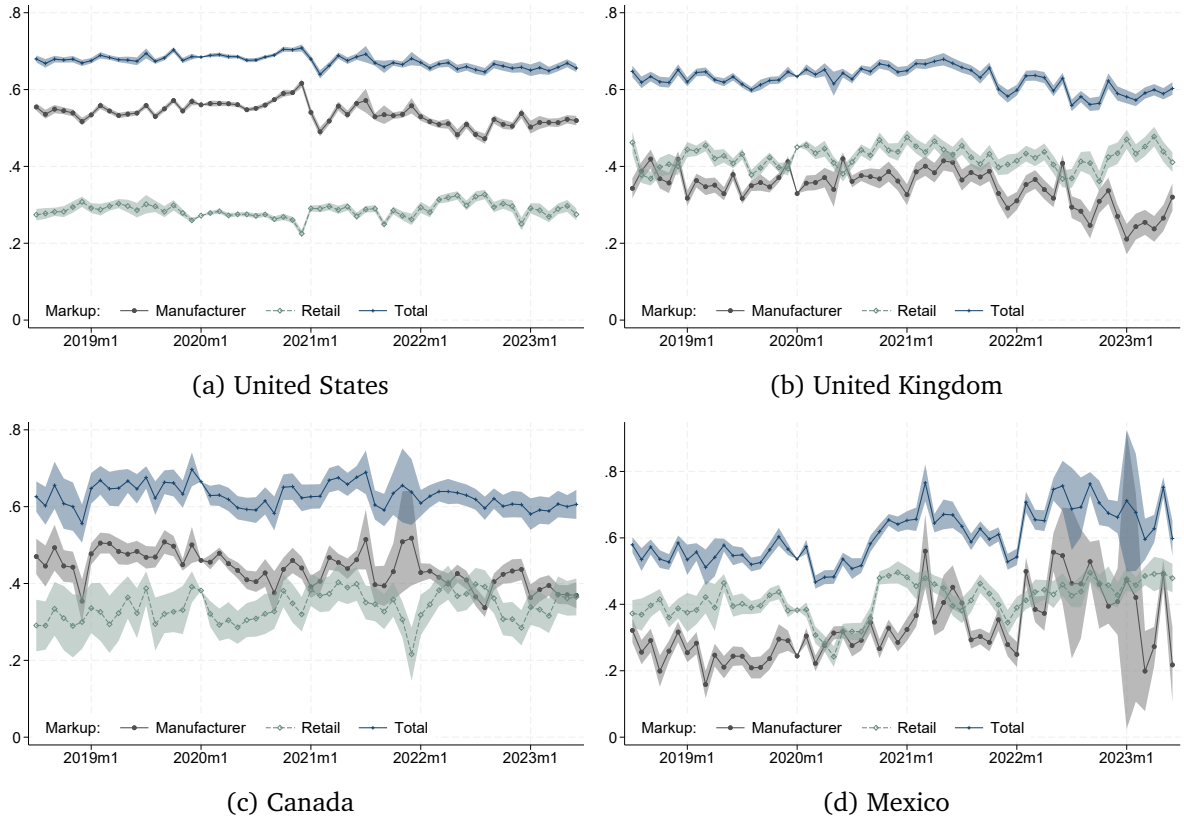
(b) United Kingdom

(c) Canada

(d) Mexico

Notes: This figure shows the correlation of markups along the supply chain in every country. Bins include sales-weighted values and values residualized on category-period fixed effects.

Figure A.6: Markup Dynamics

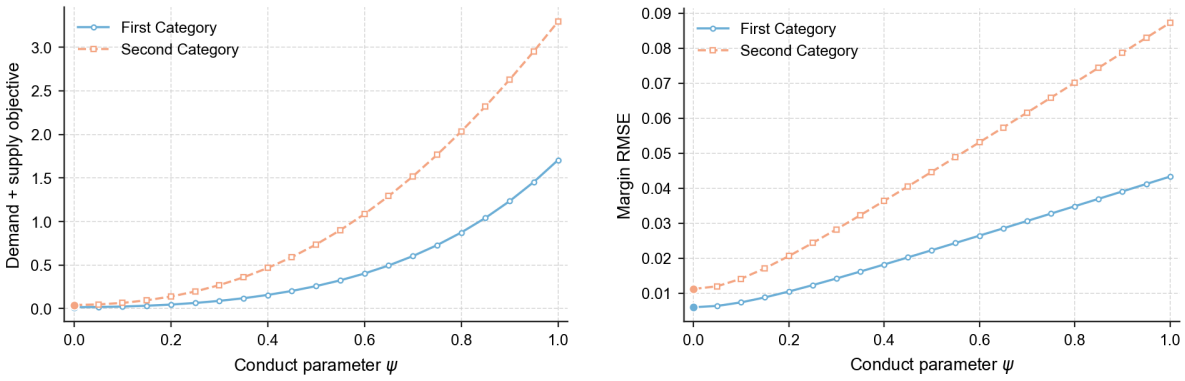


(a) United States (b) United Kingdom

(c) Canada (d) Mexico

Notes: This figure shows markup averages and 95% confidence intervals computed using product-level clustered standard errors in the United States, United Kingdom, Canada, and Mexico, resulting from a regression on period and product fixed effects using sales weights. The total markup is defined as the Lerner index of retail prices and production costs, $\frac{p^R - c}{p^R}$, the manufacturer markup as the Lerner index of wholesale prices and production costs, $\frac{p^M - c}{p^M}$, while the retail markup as the Lerner index of retail prices and wholesale costs, $\frac{p^R - p^M}{p^R}$.

Figure A.7: Profiles of the Vertical-Conduct Parameter

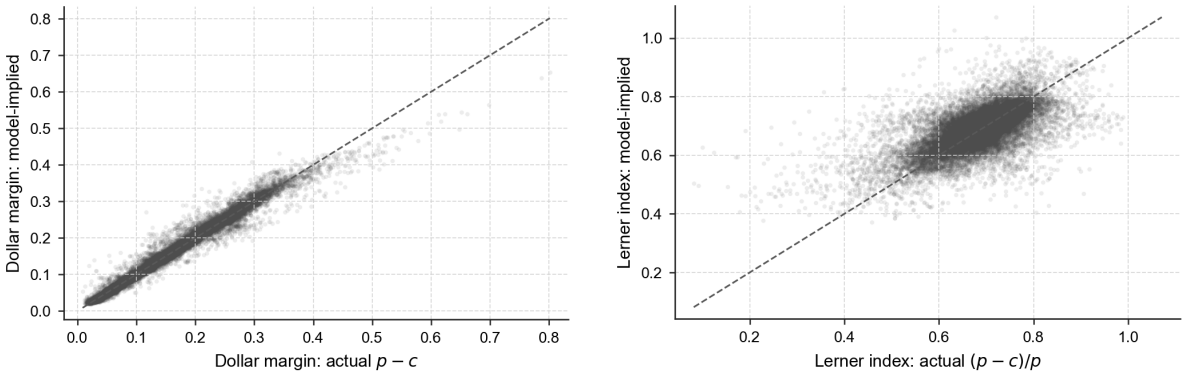


(a) Total estimation objective

(b) Margin RMSE

Notes: This figure shows the test fit at different values of the conduct parameter, ψ. The left panel reports the total estimation objective, and the right panel reports the supply-side margin RMSE over $\psi \in [0, 1]$, separately by product category. Both panels use quantity weights.

Figure A.8: Model Fit for Dollar Margins and Lerner Indices



(a) Dollar margins

(b) Lerner indices

Notes: The horizontal axes report margins observed in the data, and the vertical axes report the corresponding margins implied by estimated demand. The dashed line is the 45-degree line.

B Alternative Demand Systems

This appendix documents the demand systems we estimated, how they differ, and how well each one fits the observed margins. Every model is estimated separately for each broad category on the U.S. sample (60 months, July 2018 to June 2023; 10,732 product-months and 419 products in the First Category, 13,506 product-months and 552 products in the Second Category).

Common structure. All models share the same choice set, moments and supply side; they differ only in the demand kernel. A market is a category-month. Consumers choose among the manufacturer's products, one aggregate rival alternative for each narrow category g, and an outside good. Rival quantities are recovered from coverage parameters ρ_g, the manufacturer's share of narrow category g, which are estimated. The market size M is set so that the largest inside share across months equals 0.6. Given the nonlinear parameters, mean utilities δ_{jt} are recovered by the BLP contraction, and the structural error ξ_{jt} is the residual of δ_{jt} after product and month fixed effects. The demand moments are $E[z_{jt}\xi_{jt}] = 0$, where z_{jt} contains total unit cost and its four components (materials, manufacturing, transport and warehousing, other), residualized on the same fixed effects.

On the supply side, let μ_{jt} be the markup implied by joint pricing of the manufacturer's portfolio, $\mu_t = -(\Delta_t')^{-1}s_t$, where Δ_t is the matrix of share derivatives with respect to the retail price among the manufacturer's products. The conduct margin is

$$m_{jt}(\psi) = (p^R_{jt} - c_{jt}) - \psi\,(p^M_{jt} - c_{jt}),$$

so that $\psi = 0$ corresponds to joint pricing (the whole margin $p^R - c$ is set against demand) and $\psi = 1$ to sequential pricing (only the retail margin $p^R - p^M$ is). The criterion is $Q = Q_D + Q_S$, where Q_D is the mean squared demand moment scaled by the variance of ξ, and Q_S is the quantity-weighted mean squared difference between $m_{jt}(\psi)$ and μ_{jt}, scaled by the variance of $m_{jt}(\psi)$. Given demand, ψ has a closed form (the weighted projection of $p^R - c - \mu$ on $p^M - c$), which is concentrated out and restricted to $[0, 1]$. Q is minimized with Adam from three starting values per category, keeping the start with the lowest Q.

Random-coefficient models. Each of the four random-coefficient models has a single latent consumer type y_i that shifts both price sensitivity and the taste for each product through a product-specific loading β_j (normalized to mean zero):

1. *Two types.* Two types with equal mass. Type L has utility $\delta_{jt} - \alpha^L p^R_{jt}$; type H has $\delta_{jt} - \alpha^H p^R_{jt} + x_j$, with $\alpha^L > \alpha^H$ and a free product taste x_j.

2. *Continuous y (log-normal price coefficient).* $y_i \sim N(0,1)$, utility $\delta_{jt} + \beta_j y_i - \alpha e^{-\sigma_y y_i} p_{jt}^R$; consumers with high y are less price sensitive.

3. *Normal y (normal price coefficient).* $y_i \sim N(0,1)$, utility $\delta_{jt} + \beta_j y_i - (\alpha - \sigma_y y_i) p_{jt}^R$, with a penalty that keeps $\alpha \geq 3\sigma_y$ so that almost no consumer has a positive price coefficient.

4. *Uniform y (baseline).* $y_i \sim U(0,1)$, utility $\delta_{jt} + \beta_j (y_i - \frac{1}{2}) - (\alpha_0 + \alpha_1 y_i) p_{jt}^R$.

Benchmark models. To show what the product-specific heterogeneity buys, we also estimate five standard demand systems with the same choice set, instruments and criterion: logit; nested logit with one nest per narrow category; CES in expenditure shares; nested CES with one nest per narrow category; and a random coefficient on price only (the baseline with $\beta_j = 0$ for every product). The nested CES is the nested logit in log prices and expenditure shares, $\ln(e_{jt}/e_{0t}) = \delta_{jt} - (\eta - 1) \ln p_{jt}^R + \sigma_n \ln e_{jt|g}$, where e denotes expenditure shares. The elasticity of substitution is η between nests and $\sigma_w = 1 + (\eta - 1)/(1 - \sigma_n) \geq \eta$ within a nest; with $\sigma_n = 0$ it reduces to CES with elasticity η. For these benchmarks ρ_g and M are held at the baseline estimates, so all models are fit to the same shares, and ψ is fixed at 0, so the comparison isolates the demand kernel.

Results. Table B.1 reports the criterion, the conduct parameter, the fit of the margins and the median own-price elasticity for all models. Table B.2 reports the price parameters of the random-coefficient models. Three results stand out.

First, the estimated conduct parameter is essentially zero in every random-coefficient model: at most 0.002 in the First Category and 0.028 in the Second Category. Joint pricing is selected regardless of the demand kernel. Among the benchmarks, only the logit would choose $\psi = 1$ if ψ were left free. The logit implies nearly identical markups for all products in a market and cannot match margins under either conduct (dollar margin RMSE of 0.057 and 0.063 against 0.006 and 0.011 for the baseline). The nested logit would choose $\psi = 0.44$ in the Second Category; CES, nested CES and the random coefficient on price would all choose $\psi = 0$.

Second, the uniform-y model attains the lowest criterion and the best fit of margins in both categories. Figures 7 and 8 show why the product-specific loadings matter. The logit and nested logit cannot generate dispersion in markups within a market. CES and the random coefficient on price fit dollar margins well, because dollar margins scale with prices, but they imply nearly constant Lerner indices: the correlation between actual and implied Lerner indices is -0.22 and 0.01 in the First Category. Nesting does not change this. In the First Category the estimated nesting parameter is close to zero ($\sigma_n = 0.03$) and the nested CES collapses to CES. In the Second Category products within a narrow category are closer substitutes ($\sigma_w = 2.16$ against $\eta = 1.82$), so the implied Lerner index differs across narrow categories, but it remains nearly constant within each of them, and the correlation with actual Lerner indices stays at -0.06. The

baseline reproduces the cross-sectional pattern of Lerner indices, with a correlation of 0.70 in both categories.

Third, the models agree on the First Category, where the median own-price elasticity lies between -1.4 and -1.5 in all four random-coefficient models, but not on the Second Category, where it ranges from -1.6 to -6.6.

Table B.1: Fit and conduct across demand systems

	First Category					Second Category				
	Q	Q_D	$\hat{\psi}$	RMSE(m)	Med. ϵ	Q	Q_D	$\hat{\psi}$	RMSE(m)	Med. ϵ
Random-coefficient models (ψ estimated)										
Two types	0.043	0.005	0.000	0.0111	-1.44	0.060	0.009	0.000	0.0136	-1.79
Continuous y	0.020	0.005	0.000	0.0068	-1.42	0.052	0.009	0.000	0.0125	-2.60
Normal y	0.025	0.005	0.002	0.0080	-1.52	0.085	0.010	0.028	0.0164	-1.64
Uniform y (baseline)	0.016	0.005	0.000	0.0060	-1.46	0.041	0.007	0.001	0.0112	-6.57
Benchmark models ($\psi = 0$; ρ_g, M at baseline)										
Logit	1.004	0.005	[1.00]	0.0567	-1.57	1.105	0.007	[1.00]	0.0632	-3.12
Nested logit	0.326	0.048	[0.00]	0.0299	-2.21	0.670	0.019	[0.44]	0.0487	-21.63
CES	0.034	0.005	[0.00]	0.0095	-1.44	0.089	0.009	[0.00]	0.0171	-1.85
Nested CES	0.033	0.006	[0.00]	0.0095	-1.45	0.088	0.010	[0.00]	0.0169	-2.15
RC on price only	0.035	0.005	[0.00]	0.0097	-1.49	0.072	0.010	[0.00]	0.0151	-6.22

Notes: $Q = Q_D + Q_S$ is the estimation criterion and Q_D its demand part. RMSE(m) is the root mean squared difference between the conduct margin $m(\hat{\psi})$ and the model-implied markup, in dollars. Med. ϵ is the median own-price elasticity. For the benchmarks, ψ is fixed at 0; the number in brackets is the value of ψ, restricted to $[0, 1]$, that the benchmark would choose if ψ were estimated.

Table B.2: Price parameters of the random-coefficient models

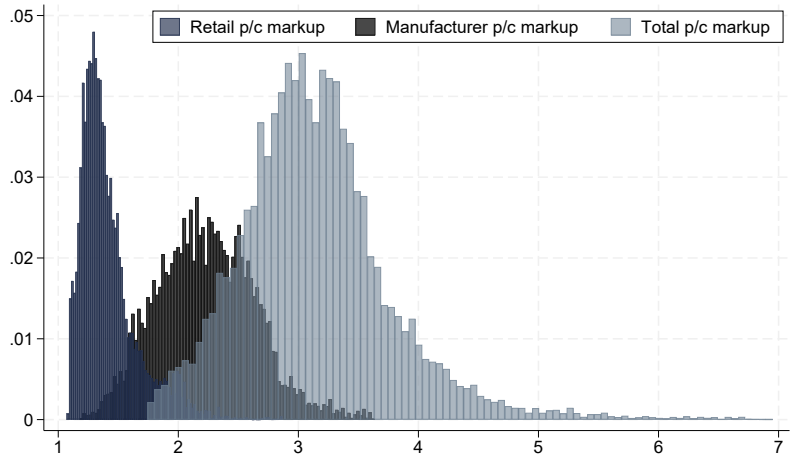
Model	Parameter	First Category	Second Category
Two types	α^L, α^H	36.48, 3.88	19.92, 2.99
Continuous y	α, σ_y	53.19, 1.47	30.88, 1.15
Normal y	α, σ_y	28.45, 8.62	10.55, 3.34
Uniform y (baseline)	α_0, α_1	0.34, 100.90	1.35, 202.16
Nested logit	σ	0.88	0.95
CES	σ	1.44	1.85
Nested CES	η, σ_w (σ_n)	1.44, 1.45 (0.03)	1.82, 2.16 (0.29)
RC on price only	α_0, α_1	0.52, 86.67	1.78, 172.78

C Alternative Markup Definition

In this section, we present our main results using price-over-cost (p/c) markups rather than the Lerner index. Consistent with our previous notation, we define total, retail, and manufacturer price-over-cost markups as $\frac{p_{it}^R}{c_{it}}$, $\frac{p_{it}^M}{c_{it}}$, and $\frac{p_{it}^R}{p_{it}^M}$, respectively.

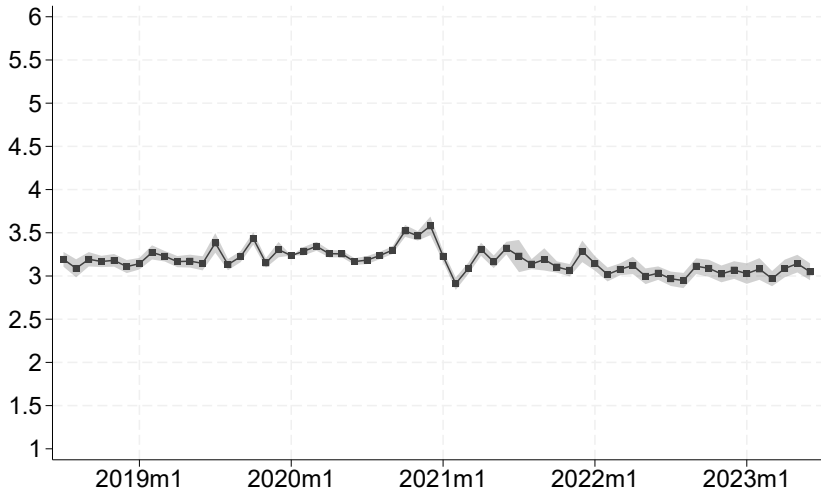
Figures C.1, C.2, C.3 and C.4 display the equivalent results to Figures 1, 2a, 3 and 4, respectively, using price-over-cost markups. Figure C.3 is constructed using an index of manufacturer and retail price-over-cost markups due to the different average levels of price-over-cost markups for the manufacturer and the retailer.

Figure C.1: Price-Over-Cost Markups Along the Supply Chain



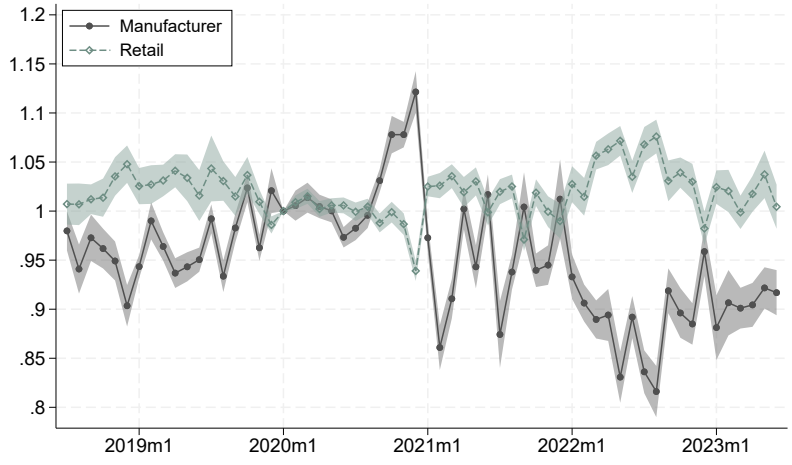
Notes: This figure shows the sales-weighted frequency distribution of price-over-cost (p/c) markups along the supply chain in the United States. Retail markups over costs ($\frac{p^R}{p^M}$) are shown in dark blue, manufacturer markups ($\frac{p^M}{c}$) in dark grey, and total markups ($\frac{p^R}{c}$) in light blue. Top and bottom percent of price-over-cost markups not included in the figure.

Figure C.2: Time Series of Total Price-Over-Cost Markups Along the Supply Chain



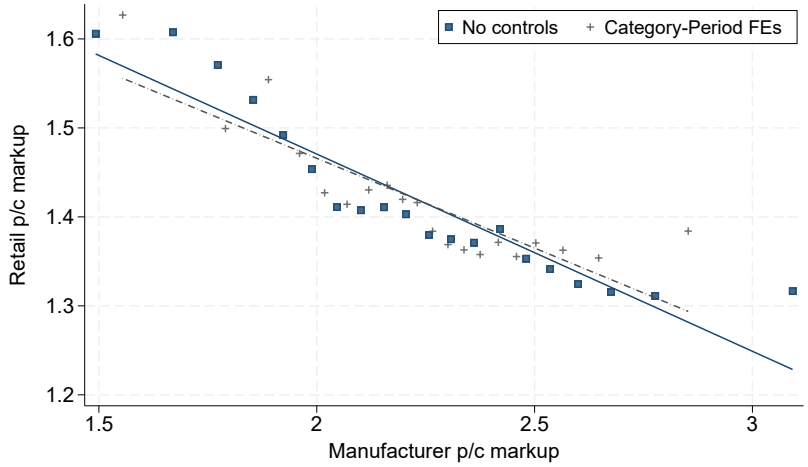
Notes: This figure shows the time series trends in total price-over-cost (p/c) markups. Total price-over-cost markups are defined as retail prices divided by manufacturing costs. We report time period coefficients from regressions on period and product fixed effects using sales weights, with 2020m1 as the base period. The base period is indexed to its (sales-weighted) mean value. 95 percent confidence intervals computed using product-level clustered standard errors are displayed in gray. Top and bottom percent of price-over-cost markups not included in the estimation of the index.

Figure C.3: Price-Over-Cost Markup Dynamics



Notes: This figure shows the negative dynamic correlation of price-over-cost (p/c) markups along the supply chain in the United States. We report time period coefficients from regressions of manufacturer or retail price-over-cost log markups on period and product fixed effects using sales weights, with 2020m1 as the base period indexed to a value of 1. 95% (robust) confidence intervals computed using product-level clustered standard errors resulting from a regression on period and product fixed effects using sales weights. Top and bottom percent of price-over-cost markups not included in the estimation.

Figure C.4: Price-Over-Cost Markup Correlation Along the Supply Chain



Notes: This figure shows the negative correlation of price-over-cost markups along the supply chain in the United States. Bins include sales-weighted values and values residualized on category-period fixed effects.

D Alternative Bargaining Assumptions

We propose an alternative version of the theoretical framework introduced in Section 4, which differs from the baseline in one key respect: it assumes fixed bargaining weights for both parties and introduces a time-varying outside option $\bar{\pi}_i^M \geq 0$ for the manufacturer and $\bar{\pi}_i^R \geq 0$ for the retailer. Product pricing remains unchanged: it follows a two-stage process in which the manufacturer first sets retail prices to maximize joint supply chain profits, and then the manufacturer and retailer negotiate the division of product-level margins.

Product margins are split via a Nash bargaining protocol with fixed weights $\tau_i \in [0, 1]$ for the manufacturer:

$$\max_{p_i^M} \left(p_i^M - c_i - \bar{\pi}_i^M\right)^{\tau_i} \left(p_i^R - p_i^M - \bar{\pi}_i^R\right)^{1-\tau_i}, \ \forall i \in \mathcal{I}. \tag{D.1}$$

Taking the first-order condition yields:

$$\tau_i(p_i^R - c_i - \bar{\pi}_i^R - \bar{\pi}_i^M) = p_i^M - c_i - \bar{\pi}_i^M, \tag{D.2}$$

where the bargaining weight corresponds to the share of surplus the manufacturer captures in the negotiation. Unlike Equation (7) in the baseline model, this expression cannot be directly mapped to the data, as it depends on two unobserved outside options.

We proceed in two complementary ways to bring Equation (D.2) to the data. First, we impose a parametric restriction by assuming that the expected value of each party's outside option is a constant fraction $A \in [0, 1)$ of their *average* margin:

$$\mathbb{E}[\bar{\pi}_i^M] = A\ \mathbb{E}[p_i^M - c_i], \qquad \mathbb{E}[\bar{\pi}_i^R] = A\,\mathbb{E}[p_i^R - p_i^M].$$

Taking expectations of Equation (D.2) and substituting these moment conditions yields:

$$\tau = \frac{\mathbb{E}[p_i^M - c_i]}{\mathbb{E}[p_i^R - c_i]}. \tag{D.3}$$

Since A cancels out, the ratio of sample means identifies τ for any assumed value of A.

Second, we provide a more flexible alternative by estimating the manufacturer's bargaining weight directly. Rearranging Equation (D.2) yields the following linear regression specification:

$$p_i^M - c_i = \tau_i(p_i^R - c_i) + (1 - \tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\,, \tag{D.4}$$

which we estimate via OLS using period sales weights. The estimate of τ_i is unbiased under the assumption that outside options are mean-independent of the total margin.

Table D.1 reports the results. Table D.1a shows bargaining weights derived from the calibration exercise with parametric restrictions, while Table D.1b reports country-level OLS esti-

Table D.1: Bargaining Weight with Time-Varying Outside Options

(a) Calibration Exercise

	(1) US	(2) UK	(3) CA	(4) MX
Bargaining Weight	0.59	0.33	0.42	0.31

(b) Estimation Exercise

	(1) US	(2) UK	(3) CA	(4) MX
Bargaining Weight	0.399^{***}	0.281^{***}	0.383^{***}	0.282^{***}
	(0.032)	(0.011)	(0.019)	(0.024)
Constant	0.023^{***}	0.007^{***}	0.008^{***}	0.002
	(0.004)	(0.002)	(0.003)	(0.003)
R^2	0.61	0.52	0.67	0.43
N	24,238	7,713	6,618	3,889

Notes: This table reports estimates for bargaining weights with time-varying outside option. Panel a) shows weighted averages from Equation (D.3). Panel b) shows OLS coefficients of Equation (D.4) using per period sales as weights. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

mates.[38] Both approaches yield the same picture: the manufacturer has the strongest bargaining power, meaning it retains the largest share of surplus, in the United States, while its bargaining power is weaker in the UK, Canada, and Mexico. The calibrated weights (Table D.1a) closely match the average of the time-varying bargaining weights recovered in Section 4. In contrast, the OLS estimates (Table D.1b) are lower. One reason for this is that our OLS estimates attribute more of the manufacturer's bargaining power to greater relative outside options, as we now discuss.

While we cannot separately identify $\bar{\pi}^M$ and $\bar{\pi}^R$ for each product without further assumptions, we can recover the combined term $\left[(1-\tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\right]$, which reflects a weighted difference between the manufacturer's and retailer's outside options. Using the estimated bargaining weights, we can recover the combined outside option term from equation (D.4) for both specifications. For the OLS specification, the regression constant identifies the (sales-weighted) average of this combined term, and the product-specific value is recovered by adding the regression residuals to the constant.[39] Estimated constants are reported in Table D.1b, and are positive and significant in the US, UK, and Canada, but not significantly different from zero in Mexico.[40] In the US, we estimate a greater relative outside option, which explains the greater decline in the bargaining weight (compared to other countries) from the calibration exercise.

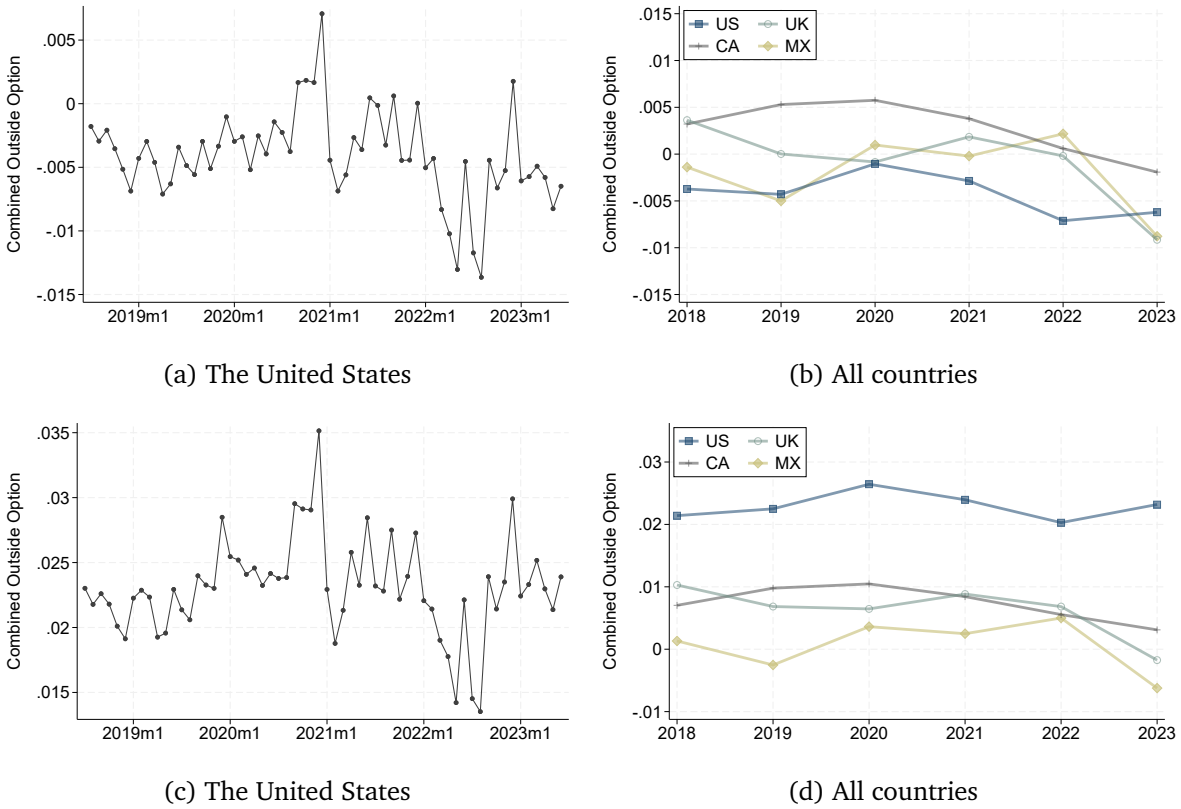
We explore the time-series dynamics of the combined outside option term in Figure D.1. The top panels show results from the calibration exercise, and the bottom panels from the estimation

[38] All moments are computed using period sales weights, consistent with Section 4.

[39] We trim the resulting estimates of the combined outside option term at the 1% level.

[40] A regression without a constant—effectively assuming this term equals zero on average—yields higher τ estimates in all countries: 0.53 in the US, 0.33 in the UK, 0.43 in Canada, and 0.31 in Mexico. Median coefficients from quantile regressions with a constant are also close to those from the OLS without constant.

Figure D.1: Combined Outside Option Over Time



(a) The United States

(b) All countries

(c) The United States

(d) All countries

Notes: This figure shows the time series of the estimates of the combined outside options term $\left[(1-\tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\right]$. Panel (a) and (b) show estimates from the calibration exercise, while Panel (c) and (d) estimates from the estimation exercise.

exercise. Figure D.1a displays monthly variation in the US under the calibration approach. The combined term fluctuates around zero by construction, with notable volatility: it spikes at the end of 2020 and falls sharply in early 2022, before returning to trend. Figure D.1c shows similar high-frequency variation under the estimation approach, though the series now centers around a positive mean of approximately 0.02.

Figure D.1b plots annual averages across countries from the calibration exercise. The series remains relatively stable for the US and Canada, while Mexico and the UK exhibit a pronounced decline in 2023. A comparable pattern emerges from the estimation-based series in Figure D.1d, with the US, UK, and Canada consistently showing positive averages.

Finally, we examine the correlation between estimated combined term and a set of economic variables to shed light on the potential determinants of outside options in favor of one or the other party involved in the negotiation. This exercise mirrors the analysis presented in the main text, where we explored how bargaining power covaries with market structure and supply chain characteristics.

Since a non-negligible share of the combined outside option terms are negative, we cannot

Table D.2: Bargaining Weight Predictors

	(1)	(2)	(3)	(4)	(5)
Production Cost	-0.742***	-0.759***	-0.607***	-0.921***	-0.964***
	(0.023)	(0.023)	(0.026)	(0.040)	(0.051)
Total Margin	0.883***	0.806***	0.783***	0.910***	0.811***
	(0.016)	(0.017)	(0.016)	(0.034)	(0.046)
Sales per Capita	0.127***	0.133***	0.098***	0.146***	0.132***
	(0.009)	(0.009)	(0.011)	(0.012)	(0.013)
GDP per Capita	1.179***	-0.260*	0.909***	0.919***	1.111***
	(0.067)	(0.143)	(0.103)	(0.243)	(0.313)
CPI	-0.820***	-0.242*	-0.553***	-0.758***	-0.628**
	(0.133)	(0.131)	(0.165)	(0.221)	(0.266)
Quality					0.200***
					(0.026)
Country FEs		X			
Non US			X		
US Only				X	X
Observations	42,458	42,458	18,220	24,238	16,476
R^2	0.563	0.582	0.507	0.396	0.432

Notes: This table reports the outcome of regressing standardized manufacturer bargaining weights on (log) production costs, (log) total margins, (log) product sales, (log) GDP per capita, (log) CPI, and (standardized) quality. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

apply a log transformation. Instead, we standardize it and regress it on the same set of predictors. For completeness, we also replicate Table 4 using the standardized level of bargaining power (recovered from Equation (7)) as the dependent variable. The results, presented in Table D.2, remain consistent with the baseline specification, though the R^2 values are lower than in our baseline specification.

We now examine how the combined outside option term correlates with product and market characteristics, noting that an increase in the combined term $\left[(1-\tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\right]$ implies an increase in the manufacturer's outside option relative to retailer's. Table D.3 reports the estimated coefficients for the calibration exercise, and Table D.4 reports the same analysis using the combined outside option terms resulting from the estimation exercise.

Some patterns are consistent across all three tables. Product sales per capita remain a strong and stable predictor, with large positive coefficients throughout. This suggests that products with broader market penetration are associated with stronger fallback positions for the manufacturer, possibly due to demand scale or distribution leverage. Inflation, measured as the consumer price index (CPI) is negatively and significantly associated with the combined term across all specifications.

However, we do see changes in sign for some key variables across specifications and within specifications across countries. Overall, the predictive power of these regressions is substantially lower than our baseline specification. R^2 values for column (1) are 0.20 and 0.17 for the

Table D.3: Combined Outside Option Predictors

	(1)	(2)	(3)	(4)	(5)
Production Cost	-0.228***	-0.029	-0.106***	0.209***	0.112*
	(0.027)	(0.032)	(0.032)	(0.049)	(0.062)
Total Margin	-0.306***	-0.479***	-0.126***	-0.834***	-0.969***
	(0.033)	(0.042)	(0.043)	(0.062)	(0.082)
Sales per Capita	0.166***	0.159***	0.093***	0.202***	0.159***
	(0.013)	(0.013)	(0.014)	(0.018)	(0.020)
GDP per Capita	-0.572***	-0.306*	-0.067	0.872**	1.137***
	(0.092)	(0.161)	(0.114)	(0.340)	(0.424)
CPI	-0.217	-0.497***	-0.618***	-1.257***	-1.264***
	(0.177)	(0.175)	(0.223)	(0.294)	(0.331)
Quality					0.362***
					(0.034)
Country FEs		X			
Non US			X		
US Only				X	X
Observations	41,612	41,612	17,858	23,754	16,094
R^2	0.202	0.224	0.055	0.280	0.367

Notes: This table reports the outcome of regressing the combined outside options term $\left[(1-\tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\right]$ on (log) production costs, (log) total margins, (log) product sales, (log) GDP per capita, (log) CPI, and (standardized) quality. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

combined outside options specifications, compared to 0.56 for the bargaining weight predictions. Our preferred specification is the one used in the main text.

Table D.4: Combined Outside Option Predictors

	(1)	(2)	(3)	(4)	(5)
Production Cost	0.207***	0.005	-0.148***	0.151***	0.116*
	(0.026)	(0.035)	(0.035)	(0.054)	(0.064)
Total Margin	0.081**	0.257***	0.110**	0.240***	0.050
	(0.033)	(0.042)	(0.045)	(0.063)	(0.080)
Sales per Capita	0.160***	0.171***	0.100***	0.192***	0.146***
	(0.013)	(0.013)	(0.015)	(0.018)	(0.020)
GDP per Capita	1.626***	-0.244	-0.053	0.614*	0.760*
	(0.093)	(0.174)	(0.124)	(0.337)	(0.409)
CPI	-1.343***	-0.483***	-0.698***	-0.741**	-0.832**
	(0.180)	(0.180)	(0.240)	(0.293)	(0.326)
Quality					0.382***
					(0.033)
Country FEs		X			
Non US			X		
US Only				X	X
Observations	41,609	41,609	17,856	23,753	16,145
R^2	0.165	0.188	0.036	0.169	0.281

Notes: This table reports the outcome of regressing the combined outside options term $\left[(1-\tau_i)\bar{\pi}_i^M - \tau_i\bar{\pi}_i^R\right]$ on (log) production costs, (log) total margins, (log) product sales, (log) GDP per capita, (log) CPI, and (standardized) quality. Standard errors (in parentheses) are clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Demand Estimation

E.1 Market Construction

We provide additional details on the construction of the markets and market shares used to estimate the model in Section 4.

Within each broad product category c, let $\mathcal{G}_c$ denote the set of narrow product categories and let $\mathcal{I}_{gt}$ denote the manufacturer's products observed in narrow category g and month t. We include one aggregate rival for each narrow category in which the manufacturer has positive sales. Its quantity is the difference between the estimated total quantity in the narrow category and the quantity observed for the manufacturer. In particular,

$$Q_{gt}^{\text{obs}} = \sum_{j \in \mathcal{I}_{gt}} q_{jt}, \qquad Q_{gt}^{\text{rival}} = Q_{gt}^{\text{obs}} \left(\frac{1}{\rho_g} - 1 \right), \qquad \rho_g \in (0, 1). \tag{E.1}$$

We estimate ρ_g jointly with the demand and conduct parameters. Because we do not observe rival prices, we assign the aggregate rival the quantity-weighted average retail price of the manufacturer's products in the same narrow category and month:

$$p_{gt}^{\text{rival}} = \frac{\sum_{j \in \mathcal{I}_{gt}} q_{jt} p_{jt}^{R}}{Q_{gt}^{\text{obs}}}. \tag{E.2}$$

We use a common potential market size for all months within a broad product category. Letting $Q_{gt}^{\text{all}} = Q_{gt}^{\text{obs}} / \rho_g$, we set

$$M_c = \frac{5}{3} \max_t \left\{ \sum_{g \in \mathcal{G}_c} Q_{gt}^{\text{all}} \right\}. \tag{E.3}$$

The manufacturer's product, rival, and outside option shares are therefore

$$s_{jt} = \frac{q_{jt}}{M_c}, \qquad s_{gt}^{\text{rival}} = \frac{Q_{gt}^{\text{rival}}}{M_c}, \qquad s_{0t} = 1 - \sum_{g \in \mathcal{G}_c} \left(\sum_{j \in \mathcal{I}_{gt}} s_{jt} + s_{gt}^{\text{rival}} \right). \tag{E.4}$$

E.2 Details on Estimation and Identification

Let $\boldsymbol{z}_{jt}$ collect the L instruments and let N denote the number of manufacturer product-month observations. For each candidate value of $(\theta, \boldsymbol{\rho})$, we recover mean utilities and the corresponding demand residuals as described in the main text. The demand moments and their contribu-

tion to the estimation criterion are

$$\boldsymbol{g}_D(\theta, \boldsymbol{\rho}) = \frac{1}{N} \sum_{j,t} \boldsymbol{z}_{jt} \xi_{jt}, \qquad Q_D(\theta, \boldsymbol{\rho}) = \frac{\boldsymbol{g}_D(\theta, \boldsymbol{\rho})' \boldsymbol{g}_D(\theta, \boldsymbol{\rho})}{L V_\xi}, \tag{E.5}$$

where $V_\xi = N^{-1} \sum_{j,t} (\xi_{jt} - \bar{\xi})^2$ is the variance of the recovered demand residuals. These moments use the manufacturer's products, for which costs are observed.

For each candidate value of $(\theta, \boldsymbol{\rho})$, we choose $\widehat{\psi}(\theta, \boldsymbol{\rho})$ using the quantity-weighted minimum-distance criterion in Equation (13). The supply component of the estimation criterion is

$$Q_S(\theta, \boldsymbol{\rho}, \psi) = \frac{\sum_{j,t} q_{jt} \left[m_{jt}(\psi) - m_{jt}^D(\theta, \boldsymbol{\rho}) \right]^2}{\sum_{j,t} q_{jt} \left[m_{jt}(\psi) - \bar{m}_q(\psi) \right]^2}, \tag{E.6}$$

where $\bar{m}_q(\psi) = \sum_{j,t} q_{jt} m_{jt}(\psi) / \sum_{j,t} q_{jt}$ is the quantity-weighted mean of the data-implied margins. The denominator normalizes the pricing discrepancies by the weighted variance of these margins.

We estimate the demand and coverage parameters by minimizing the sum of the two components:

$$(\widehat{\theta}, \widehat{\boldsymbol{\rho}}) = \arg\min_{\theta, \boldsymbol{\rho}} \left\{ Q_D(\theta, \boldsymbol{\rho}) + Q_S\left(\theta, \boldsymbol{\rho}, \widehat{\psi}(\theta, \boldsymbol{\rho})\right) \right\}. \tag{E.7}$$

The conduct estimate is then $\widehat{\psi} = \widehat{\psi}(\widehat{\theta}, \widehat{\boldsymbol{\rho}})$. This procedure uses both the cost instruments and the fit of the pricing condition to estimate demand, without imposing either coordinated or sequential pricing in advance.

The profiles in Figure A.7 hold the demand and coverage parameters at their estimated values and vary ψ. They show how the fit of the pricing condition changes with vertical conduct, conditional on the estimated demand system.

F Additional Pass-Through Results

Table F.1 reports the US estimates of direct pass-through from production costs to retail prices. Each specification includes product fixed effects and jointly estimates the effects of aggregate and idiosyncratic costs. The six- and twelve-lag specifications provide the cumulative coefficients used in Section 6.

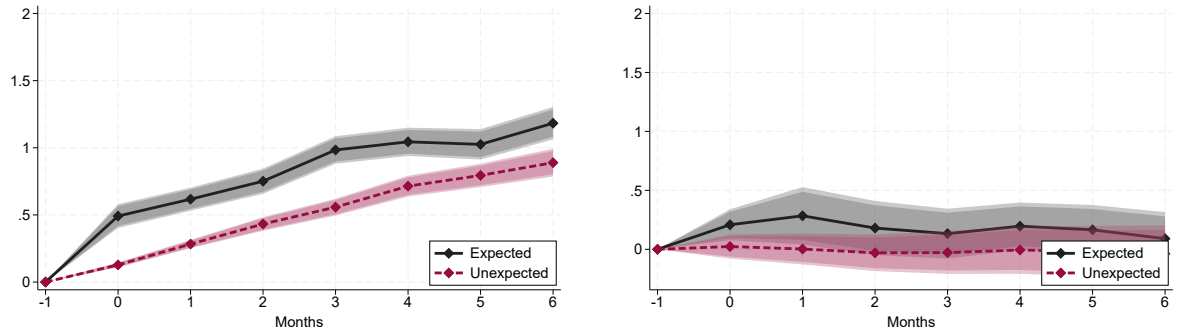
To compare expected and unexpected costs, we decompose each into a category-month average and a product-specific deviation. All four components enter jointly with their contemporaneous values and six monthly lags. We report the resulting cumulative wholesale- and retail-price responses in Figures F.1 and F.2.

Table F.1: Direct Pass-Through from Production Costs to Retail Prices

	Idiosyncratic					Aggregate			
	(1)	(2)	(3)	(4)		(5)	(6)	(7)	(8)
$\hat{c}_t$	0.163*	0.003	-0.028	-0.036	C_t	0.340***	0.262***	0.246***	0.244***
	(0.099)	(0.045)	(0.045)	(0.054)		(0.033)	(0.026)	(0.025)	(0.028)
$\hat{c}_{t-1}$		0.075*	0.012	0.047	C_{t-1}		0.262***	0.223***	0.240***
		(0.045)	(0.039)	(0.044)			(0.023)	(0.021)	(0.021)
$\hat{c}_{t-2}$		0.084	0.034	0.048	C_{t-2}		0.223***	0.241***	0.247***
		(0.058)	(0.037)	(0.046)			(0.022)	(0.023)	(0.026)
$\hat{c}_{t-3}$		0.108	0.021	0.001	C_{t-3}		0.216***	0.194***	0.218***
		(0.077)	(0.037)	(0.044)			(0.025)	(0.023)	(0.025)
$\hat{c}_{t-4}$			0.055	-0.008	C_{t-4}			0.191***	0.178***
			(0.045)	(0.050)				(0.021)	(0.021)
$\hat{c}_{t-5}$			0.021	-0.007	C_{t-5}			0.144***	0.177***
			(0.044)	(0.049)				(0.020)	(0.025)
$\hat{c}_{t-6}$			0.089*	0.011	C_{t-6}			0.154***	0.158***
			(0.053)	(0.039)				(0.023)	(0.024)
$\hat{c}_{t-7}$				0.055	C_{t-7}				0.166***
				(0.058)					(0.026)
$\hat{c}_{t-8}$				0.042	C_{t-8}				0.107***
				(0.045)					(0.030)
$\hat{c}_{t-9}$				0.066	C_{t-9}				-0.020
				(0.052)					(0.027)
$\hat{c}_{t-10}$				0.170**	C_{t-10}				-0.025
				(0.068)					(0.034)
$\hat{c}_{t-11}$				-0.024	C_{t-11}				-0.013
				(0.058)					(0.028)
$\hat{c}_{t-12}$				0.005	C_{t-12}				0.002
				(0.057)					(0.035)
Long-term PT	0.163	0.270	0.203	0.369		0.340	0.964	1.392	1.681
s.e.	0.099	0.185	0.194	0.336		0.033	0.085	0.125	0.189
Product FE	Yes	Yes	Yes	Yes		Yes	Yes	Yes	Yes
Date FE	No	No	No	No		No	No	No	No
N	24,233	19,677	16,340	11,346		24,233	19,677	16,340	11,346

Notes: Estimates use US observations with matched retail prices. Each pair of columns reports the idiosyncratic and aggregate coefficients from the same specification, with zero, three, six, or twelve monthly lags. Long-term pass-through is the sum of the contemporaneous and lagged coefficients. Standard errors are clustered by product. Standard errors for cumulative responses use the delta method. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

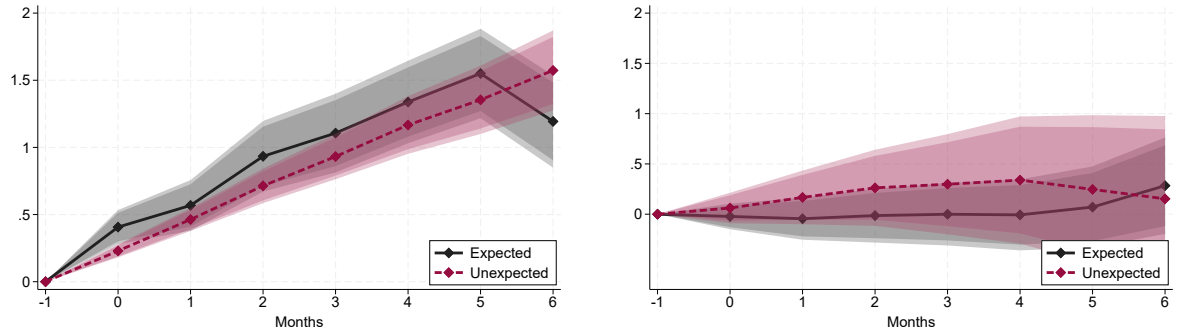
Figure F.1: Wholesale Price Responses to Expected and Unexpected Costs



(a) Aggregate Cost (b) Idiosyncratic Cost

Notes: The figure reports cumulative wholesale-price responses to expected and unexpected production costs in the United States. Shading shows 90 and 95 percent confidence intervals calculated using the delta method with product-level clustered standard errors.

Figure F.2: Retail Price Responses to Expected and Unexpected Costs



(a) Aggregate Cost (b) Idiosyncratic Cost

Notes: The figure reports cumulative retail-price responses estimated directly from expected and unexpected production costs in the United States. Shading shows 90 and 95 percent confidence intervals calculated using the delta method with product-level clustered standard errors.