Pricing Off Accounting Costs

Michel Gutmann
Northwestern

Kunal Sangani
Columbia*

September 14, 2026

Preliminary draft—subject to change

Abstract

When costs change, firms often report waiting to update prices until inventory acquired at old prices is exhausted. This behavior is at odds with standard models, which treat historical purchase costs as sunk and predict that desired prices should follow current replacement costs. We propose that this behavior arises because *(i)* accounting costs are backward-looking measures that reflect prices actually paid for inputs and inventories, and *(ii)* firms largely set prices with reference to these accounting measures of cost rather than current replacement costs. We document firms pricing off accounting costs using airlines as a case study and using quarterly financial reports for a broader sample of public firms. Accounting cost pricing means that cost fluctuations take time to propagate along supply chains, as firms work through materials and inventories purchased at pre-shock prices. We show that the resulting price dynamics can resemble canonical sticky-price models, but with important differences: delays in price adjustment reflect inventory turnover rates rather than frequencies of price adjustment; prices are more responsive to transient shocks and slower to respond to persistent shocks; and anticipation of cost increases can further delay—rather than accelerate—pass-through due to forward buying. In a quantitative input–output model, accounting cost pricing generates substantial lags in the response of consumer prices to shocks and predicts price dynamics observed following the 2025–2026 tariffs.

*Emails: michelgutmann2030@u.northwestern.edu, kas2374@columbia.edu. We are grateful to Raphael Schoenle for fielding questions on a firm survey for us. We also thank Marios Angeletos, Gabriel Chodorow-Reich, Keshav Dogra, Martin Eichenbaum, Sebastian Heise, Diego Känzig, Andrea Prat, Matt Rognlie, and Andrei Shleifer for helpful conversations and suggestions.

1 Introduction

> *"A coffee roaster noted that while tariffs on coffee have largely been lifted, selling prices will only go down once the stock of inventory acquired at higher costs has been cleared."*
>
> — Federal Reserve Beige Book, January 2026

In standard models, a firm's selling price depends on the current and future replacement costs of its inputs. Costs that a firm previously incurred when acquiring its inputs or inventory are sunk at the time of a sale. Thus, given current and future replacement costs, the historical prices that firms faced when purchasing those inputs should generally not affect prices.

In practice, however, firms frequently report that the costs of previously acquired inputs and inventory influence pricing decisions. Table 1 lists examples from Federal Reserve interviews with business contacts about the price effects of tariffs over 2025–2026. Reports that pre-tariff inventory costs affect prices appear across Federal Reserve districts and industries, and the example of a coffee roaster waiting to pass along tariff exemptions suggests this behavior applies to both cost increases and decreases. Likewise, several executives describe pre-tariff input costs influencing either their own or competitors' prices on earnings calls. Appendix Table A1 catalogs examples from retailers like Wayfair, which cites a temporary price disadvantage relative to competitors carrying pre-tariff inventory, wholesale distributors like Global Industrial, which reports "implementing additional pricing actions as inventory affected by tariffs moves through [their] cost of sales," and manufacturers like Boeing, which expects to recoup duties paid on inputs when airplanes using those inputs are shipped.

Even beyond the 2025–2026 tariff episode, firm managers frequently cite inventory stocks in deciding how and when to pass through cost changes. Figure 1 reports responses from a survey of four hundred firm price-setters conducted by Schoenle (2026). Forty percent of respondents report that they would delay raising prices until they sell their existing inventory when faced with a cost increase.

While this behavior is at odds with standard models of price-setting, the idea that firms set prices with reference to the actual costs incurred for inputs and inventories relates to early discussions of price-setting in macroeconomics. For example, in his survey of price rigidity, Gordon (1981) notes, "[while] we have assumed that firms compute marginal costs on a replacement-cost basis, [....m]any firms, however, appear to base prices on actual costs paid in the past, so-called 'FIFO' (first-in-first-out) pricing practices." Likewise, Okun (1981) writes, "The reliance of firms on historical costs is evident in many empirical studies.... It is reasonably predictable that today's changes in cattle prices will show up on

Figure 1: Responses to "Suppose your input cost rises today, but you still have inventory bought at lower cost. When do you raise your sales price?"

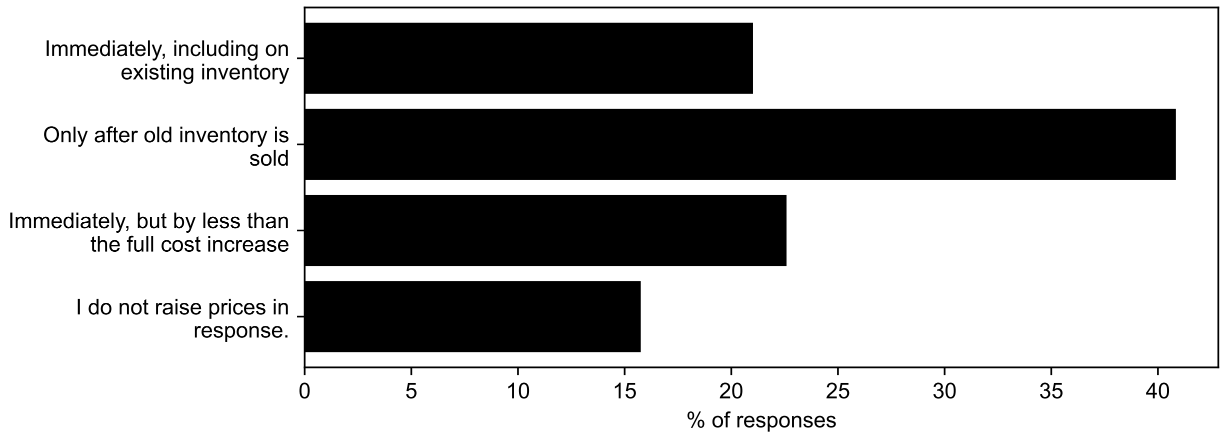

Note: Responses from a survey of 400 firm price-setters conducted by Schoenle (2026). Appendix Table A4 reports summary statistics about the surveyed individuals.

the supermarket counter only after a lag; the butcher does not keep changing price tags in pace with changes in some putative replacement cost."

In this paper, we advance two propositions. First, we argue that accounting measures of costs are largely backward-looking: they reflect the costs a firm actually incurred for inputs and inventories, rather than current or future replacement costs. Second, we argue that firms frequently set prices with reference to these accounting measures of cost rather than current replacement costs. We provide evidence for these propositions and explore their implications for macroeconomic price rigidity.

Section 2 uses airlines as a case study to demonstrate these two propositions. Airlines offer a useful empirical setting to explore accounting cost pricing because variation in fuel hedging practices across airlines and fluctuations in commodity prices generate differences in accounting costs of fuel reported across airlines, even while airlines face essentially identical fuel prices on the spot market. We first show that hedging creates a lag between changes in the spot price of fuel and airlines' accounting fuel costs: airlines that tend to hedge a larger share of their fuel needs experience a delayed pass-through of changes in jet fuel spot prices to their accounting fuel costs. We then show that hedging gains and losses incorporated into airlines' accounting costs of fuel are passed through to airlines' fares. Thus, the delayed pass-through of spot prices to accounting costs translates into delayed pass-through of spot fuel costs into ticket prices.

Section 3 then develops an approach to estimate the degree of accounting cost pricing using public firms' quarterly financial reports. We exploit the fact that changes in a firm's input costs are first capitalized in inventory before affecting the accounting cost of goods sold. Changes in the book value of inventory, alongside sales and accounting costs of goods sold, can therefore be used to estimate the degree to which firms price off replacement costs vs. accounting costs. We estimate that the extent of accounting cost pricing for public firms in Compustat ranges from 60 to 80 percent. The weight on accounting costs is lower for firms in industries with reference prices, firms that sell on organized exchanges, or firms with a large share of competitors using LIFO. In fact, once we zoom in on firms in industries with administered prices, the estimated extent of accounting cost pricing ranges from 85 to 100 percent.

In explaining why firms set prices this way, we emphasize the practical difficulties of setting prices according to replacement costs. John Cochrane's conversation about tariff pass-through with the CEO of a large retailer illustrates these challenges:[1]

> Retailer CEO: All tariff pauses just ended on August 8... we all brought in as much previously (lower) tariffed goods before that, and most fall merchandise was here in the spring and early summer... so we are still selling goods that didn't have the latest tariff burden. But we are all raising prices slowly and carefully as the tariffed inventory gets sold (remember, accounting wise, we have a cash hit paying tariffs upon receipt, but we have a P&L hit when we sell the tariffed goods). [....]
>
> Cochrane: It's interesting how the accounting treatment of bottom line matters. A pure economist would immediately mark up the current inventory to its opportunity cost[...], not its historical cost, book a windfall profit, and then mark any sales below that value as losses. Business[es] only raise prices when accounting losses force them to do so.

In effect, replacement cost pricing requires that firms value inputs and inventory on hand according to current cost conditions. The same concerns that prevent firms from marking inventories to market apply to tying prices to mark-to-market values: quotes of replacement costs may be costly to obtain, easy to manipulate, and difficult to verify or audit. Even if firms only mark to market internally for pricing decisions and not externally in financial reporting, fluctuations in replacement costs relative to accounting costs of goods sold would result in volatile performance metrics that would be difficult to justify to external observers.

Section 4 presents a two-period model that captures this intuition and explains why a firm may engage in *accounting cost pricing*. A principal delegates the operations of a firm

[1] `https://www.grumpy-economist.com/p/a-note-on-tariffs-from-the-real-world`

to a manager. The firm acquires inventory, and the manager observes the liquidation price for that inventory before choosing a price to set. Since the liquidation price determines the opportunity cost of a sale, the first-best price equates marginal revenue to this liquidation price. However, we posit that the manager may have a private preference for setting a higher or lower price than is optimal. This preference could reflect empire-building motives, a bias toward immediate earnings over longer-term value, a reluctance to engage in competition, or simply inattention to the optimal price (e.g., Baumol 1959; Bertrand and Mullainathan 2003; Graham et al. 2005). The private preference implies that, when left to his own devices, the manager may choose a price other than the one that maximizes firm value. Moreover, the presence of this private preference means that the manager cannot credibly communicate the liquidation price (which the manager can observe but the principal cannot independently verify) to the principal: regardless of the true liquidation price, the manager would always have an incentive to misreport it in order to justify his desired price.

The combination of the manager's private information about both the firm's true marginal cost and about his own preferences leads the principal to use accounting costs to discipline the manager's decision.[2] The model also generates comparative statics for the extent of accounting cost pricing that are consistent with the evidence. First, accounting cost pricing should be more prevalent when firms sell differentiated goods and face less elastic demand curves, and less prevalent when firms sell homogeneous, highly substitutable goods. Intuitively, a more elastic residual demand curve disciplines the extent to which the manager can use discretion in pricing, and the principal need not rely on high-powered incentives to achieve profits close to the first-best. In line with this prediction, we find in surveys that retail gas stations are the sole category of firms that report raising prices immediately when input costs go up, rather than waiting to raise prices until existing inventory is exhausted. Second, accounting cost pricing should diminish if there are external, accurate signals of current replacement costs. This explains why sellers of gold and silver jewelry often price to the current market price of precious metal inputs, rather than their historical purchase costs, even when they produce differentiated jewelry and face inelastic demand.[3]

[2]We do not address the question of how firms choose which accounting methods to use. Rogerson (2008, 2011) shows that if managers are incentivized using accounting measures of income, a measure based on replacement costs can restore the first best, and that this measure can be constructed from historical costs when the relationship between historical and future costs (for assets, the depreciation schedule) is known. Our model focuses on cases where replacement costs cannot be forecasted by the history of past costs.

[3]Okun (1981) (p. 156) discusses that firms appear to set prices using historical costs and arrives at an explanation remarkably close to ours: "The need for objective and systematic measurement forces the cost standard to be backward-looking—or at least sideward- rather than forward-looking. [...] Any attempt

We turn to the dynamics of prices implied by accounting cost pricing in Section 5. Compared to the two-period model, in which the book value of inventory is unambiguous, the accounting cost of goods sold in a dynamic setting depends on input prices that the firm faced in several prior periods and how the firm apportions costs across those prior purchases. We focus on the inventory valuation methods that firms use in practice when calculating cost of goods sold: most prominently average acquisition cost (AAC), first-in-first-out (FIFO), and last-in-first-out (LIFO). All three methods share the feature that accounting costs reflect realized historical costs rather than current replacement costs, but they differ in the weight they place on recent versus more distant purchases.

It turns out that there is a close connection between price dynamics under these inventory accounting methods and canonical sticky-price models. For example, price paths under average acquisition cost (AAC) pricing can be replicated by a model of Calvo (1983) rigidities, with an additional assumption that firms perceive future expected costs to equal current costs (a condition that Minton and Wheaton (2022) refer to as "myopic expectations"). Likewise, if firms make intermittent input purchases, price paths generated by LIFO pricing can coincide with a model of Taylor (1980) frictions in which firms are similarly myopic about future costs. Of course, the dependence on historical costs in the Calvo and Taylor models comes from the share of firms with stale prices, while in our model it comes from the lag between when inputs or inventories are purchased and when final goods using those inputs are sold.

Despite the close connection, price dynamics under accounting cost pricing differ from the sticky-price models in some key ways. Under accounting cost pricing, desired prices depend on historical input costs, in contrast to the forward-looking reset prices in sticky-price models.[4] Accounting cost pricing therefore predicts slower adjustment to persistent shocks and more responsiveness to transitory shocks, and further amplifies delays in price adjustment in supply chains. Moreover, while anticipation of future cost changes makes price adjustment more rapid in models with sticky prices, in our model, firms that anticipate future cost increases may stock up on inventories in advance, consequently delaying the pass-through of cost changes into prices. Thus, accounting cost pricing captures the role of forward buying in delaying the pass-through of tariffs in 2025–2026, as described in many of the quotes in Table 1.

to forecast input costs or to put them on a current replacement basis, which necessarily entails a forecast, would add to expenses, introduce a subjective element into cost calculations, and complicate the task of managerial control. Bygones cannot be mere bygones if they are the sole sources of systematic information."

[4]The dependence on historical costs does not imply that prices are completely independent of expectations. Our model predicts that accounting cost pricing arises from managerial incentives to target a certain accounting profit level. But this target profit level may adjust endogenously to information the principal has about its optimum, such as information about the rate of inflation.

The final section of the paper shows that accounting cost pricing can quantitatively generate substantial delays in aggregate price adjustment. Section 6 calibrates an input–output model of the U.S. economy with detailed data on industry inventories and simulates shocks to the costs of several industries. On average, cost shocks to goods-producing industries take 5–6 months to show up in consumer prices. These lags occur because commodity inputs take time to be used in production, and because we account for the time it takes for goods to move through the inventories of wholesale distributors and retailers. Our calibration suggests delays of five to twelve months in the effect of most commodity shocks on consumer prices.

Related literature. The idea that firms set prices with reference to realized costs, rather than current and anticipated replacement costs, dates back to Gordon (1981) and Okun (1981). Blinder et al. (1998) find that this theory of cost-based pricing resonates with firms, and Bewley (2025) stresses that the way firms measure unit costs—including allocated overhead and other accounting standards—appears to influence pricing for many surveyed firms. Most treatments of cost-based pricing refer to the role of a firm's most recent cost observations in shaping price—often because firms are inattentive to or imperfectly observe information that would predict future costs (e.g., Gordon 1990; Mankiw and Reis 2002; Woodford 2003; Minton and Wheaton 2022; Afrouzi et al. 2024)—rather than the idea that more distant historical costs, which affect accounting costs due to production and inventory lags, shape prices.

Our empirical results on airlines relate to prior work by Hortaçsu et al. (2024) and Shi et al. (2024). Using data from a large airline, Hortaçsu et al. (2024) show that pricing decisions are suboptimal due to both the pricing heuristic used and conflicting incentives of separate departments that supply its inputs. Our model focuses on how pricing can be distorted by managerial incentives and information frictions between manager and principal, complementing their findings on how organizational structure affects pricing. Shi et al. (2024) argue that sunk costs affect airlines' pricing decisions using evidence that cross-sectional variation in fuel costs across airlines at a point in time predicts variation in fares. We build on Shi et al. (2024) by documenting that (i) cross-sectional variation in fuel costs comes from differences in hedging and the backward-looking accounting cost standard, (ii) variation in fuel costs specifically from the backward-looking cost standard affects fares, and (iii) the delayed pass-through of jet fuel spot prices to fares for hedgers is robust to controlling for differences in advance ticket purchases.

Finally, an alternative explanation for the behaviors we observe is that firms set prices based on their realized costs out of fairness concerns (Okun 1981; Kahneman et al. 1986;

Table 1: Inventories and prices: Examples from the Federal Reserve Beige Book.

District	Date	Quote
Chicago	Apr 2025	Dealers expected strong demand to continue until existing inventories run down and tariffs start to affect pricing.
St. Louis	Apr 2025	A pharmacy reported they were starting to see price increases from manufacturers, but their costs had not been impacted yet as they had built up inventory.
Boston	June 2025	Manufacturers' prices increased moderately on average, driven by a partial pass-through of tariff-related cost increases. However, some manufacturers held prices fixed, either because they were not exposed to tariffs or because they stockpiled inventories of foreign inputs in advance of tariffs.
Atlanta	June 2025	Some retailers noted they were still working through previously acquired inventories and thus had not raised prices.
Atlanta	July 2025	Several contacts noted they were still working through pre-tariff inventories, thus delaying price adjustments. [. . . .] Many merchants who were still selling off pre-tariff inventories expect demand to erode if, or when, tariff-related price increases are implemented.
Philadelphia	Sep 2025	One contact reported that pre-tariff inventory is almost sold out and that tariffs have begun to cause an increase in new car prices.
San Francisco	Sep 2025	Several contacts across different retail sectors phased in price increases as old inventories were sold and replaced with goods subject to tariffs. Retailers expect additional price increases as this phase-in continues.
Atlanta	Sep 2025	Some contacts reported recent price increases on final products as pre-tariff inventories were depleted. [....] Some retailers expressed concern over further declines in demand in the coming months, as prices may increase after pre-tariff inventories are depleted.
Richmond	Oct 2025	A couple of firms said they expected to raise prices once they worked through their pre-tariff inventory.
National summary	Jan 2026	Several contacts that initially absorbed tariff-related costs were beginning to pass them on to customers as pre-tariff inventories became depleted or as pressures to preserve margins grew more acute.
New York	Jan 2026	A coffee roaster noted that while tariffs on coffee have largely been lifted, selling prices will only go down once the stock of inventory acquired at higher costs has been cleared.

Rotemberg 2005, 2011; Eyster et al. 2021).[5] Fairness and accounting cost pricing share many predictions because the underlying mechanisms are in fact quite similar: in both cases, stakeholders with incomplete information outside the firm aim to assess whether the price chosen by the firm is justified. Our model explains why these pricing behaviors arise even when consumers cannot realistically audit firms' choices—for example, by verifying the inventory a firm holds when costs change.

2 Airlines: A Case Study

We begin by using airlines as a laboratory to explore the extent to which firms' prices reflect accounting measures of cost. Airlines offer a useful empirical setting for three reasons. First, prices of jet fuel—which constitutes 20–30 percent of airlines' operating expenditures—exhibit substantial variation over time. Second, airline carriers differ in their fuel hedging strategies: some carriers like Southwest and Delta have historically hedged a substantial portion of their fuel needs, while other carriers like United and American tend to engage in minimal hedging. In the same way that two firms may face the same replacement cost but differ in accounting costs because one has existing inventory in stock and the other does not, hedging gains and losses create variation in accounting costs of fuel across carriers even while they face the same spot price of jet fuel on the margin. Finally, the Bureau of Transportation Statistics collects detailed data on airlines' costs of fuel and ticket fares, allowing us to measure the extent to which accounting costs of fuel influence prices.

Our analysis builds on Shi et al. (2024), who show that deviations in airlines' average fuel costs affect prices. We further show that deviations in airlines' accounting costs are linked to historical spot prices through hedging, and that the gap between an airline's average realized fuel cost and the jet fuel spot price passes through to ticket prices at a rate similar to that of the spot price itself, consistent with airlines pricing off average realized costs of fuel rather than marginal fuel costs.

Section 2.1 describes the data on airline accounting costs of fuel and fares. In Section 2.2, we show that hedging causes airlines' accounting costs to reflect prior jet fuel prices and thus respond sluggishly to changes in spot prices. Finally, Section 2.3 shows that airlines' fares are influenced by accounting costs of fuel.

[5] Kahneman et al. (1986) present participants with the following scenario, which bears a striking resemblance to our motivating evidence: "A grocery store has several months supply of peanut butter in stock which it has on the shelves and in the storeroom. The owner hears that the wholesale price of peanut butter has increased and immediately raises the price on the current stock of peanut butter." 79 percent of participants judged this behavior as unfair, compared to 21 percent who found it acceptable.

2.1 Data on Airline Ticket Prices and Fuel Costs

We use ticket fare data from the Airline Origin and Destination Survey (DB1B), a 10 percent sample of airline tickets collected by the Bureau of Transportation Statistics (BTS) from 1990 to 2025. We restrict the sample to nonstop itineraries and keep only tickets for which the operating carrier and ticketing carrier are the same to ensure that fares can be matched to a single carrier's costs.

Data on jet fuel spot prices and airlines' fuel accounting costs come from the Energy Information Administration (EIA) and the BTS, respectively. We use the spot price of kerosene-type jet fuel from the EIA.[6] Airlines' realized fuel costs come from the BTS Air Carrier Financial Reports Schedule P-12(a). We focus on fuel costs for scheduled domestic operations, since these operations correspond most closely to the tickets in the DB1B sample; nonscheduled operations refer to cargo charters and other irregular services. For each carrier i in quarter t, we compute the average fuel cost per gallon as

$$\text{Fuel Cost}_{it} = \frac{\text{Total Scheduled Domestic Fuel Cost}_{it}}{\text{Total Scheduled Domestic Fuel Consumption}_{it}}.$$

BTS Directive 324 instructs airlines to use an average acquisition cost method when computing fuel costs for the Schedule P-12(a) submission, so these accounting measures of fuel cost reflect actual prices paid for fuel inclusive of hedging gains and losses.

Our analysis uses data on ten airlines—American, Alaska, Allegiant, Delta, Frontier, Hawaiian, JetBlue, Southwest, Spirit, and United—over the period from 1990 to 2025. Over this period, these ten airlines account for two-thirds of all fuel consumption among airlines reporting to the BTS. We merge data on these airlines' fuel costs with quarterly financial data from Compustat.

2.2 Delayed Adjustment of Accounting Costs to Spot Price

We first show that hedging leads airlines' accounting measures of cost to respond sluggishly to changes in the spot price of jet fuel. We measure the pass-through of changes in the spot price of jet fuel to each airline's accounting cost of fuel using the specification

$$\Delta\text{Fuel Cost}_{it} = \alpha_i + \sum_{k=0}^{12} \beta_{ik} \Delta\text{Spot Fuel Price}_{t-k} + \varepsilon_{it}, \tag{1}$$

[6] This spot price is strongly correlated with accounting fuel costs for airlines that do not hedge, suggesting that it is a good measure of spot prices faced by airlines. For example, regressions of United's and Frontier's quarterly accounting fuel costs per gallon on the spot price yield coefficients of 1.02 and 1.01, with $R^2 = 0.992$ and 0.991 respectively.

Figure 2: Pass-through of changes in jet fuel spot price to airlines' accounting costs of fuel.

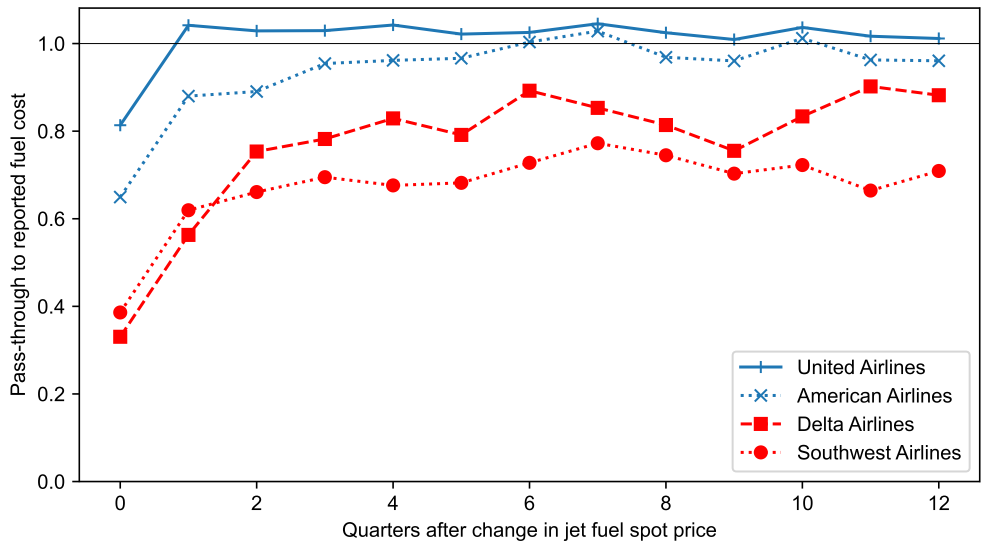

Note: The plotted points show the cumulative pass-through $\sum_{k=0}^{h} \beta_{ik}$ for each horizon $h = 0, ..., 12$, estimated using Equation (1) for the "Big Four" airline carriers.

where ΔFuel Cost$_{it}$ is the change in airline i's accounting cost of fuel from quarter $t-1$ to t and ΔSpot Fuel Price$_{t-k}$ is the change in the spot price of jet fuel from quarter $t-k-1$ to $t-k$.

This distributed lag specification is used in the pass-through literature to trace out the speed of adjustment of prices to costs (e.g., Goldberg and Hellerstein 2008; Nakamura and Zerom 2010). We instead use it to show sluggish adjustment of the accounting measure of fuel costs to the spot price. For any horizon h, the sum of the coefficients $\beta_{i0} + ... + \beta_{ih}$ reflects the portion of the change in the spot fuel price in quarter t that is captured in the accounting cost measure by quarter $t+h$.

Figure 2 shows the pass-through of changes in the spot price to airlines' accounting fuel costs for the "Big Four" airlines. For United and American Airlines, which engage in limited hedging, the pass-through of changes in the jet fuel spot price to accounting costs is rapid and converges to full pass-through within a few quarters. On the other hand, Delta and Southwest Airlines, which engage in substantial hedging, exhibit delayed pass-through of changes in the spot fuel price to their accounting costs. Intuitively, Delta and Southwest have locked in prices for part of their fuel needs at spot prices prevailing in the past, so that their accounting costs will depend on past prices of fuel. Four quarters after a change in the jet fuel spot price, the accounting cost of fuel for Southwest reflects 71% of the change in the spot price, and the accounting cost of Delta reflects 80% of the change in the spot price, compared to 97% and 104% for American and United, respectively.

These empirical results on delayed adjustment of accounting fuel costs to spot prices for hedging carriers are supported by managers' explanations in financial reports and earnings calls. For example, when fuel prices tumbled in late 2008, Delta Air Lines noted in its 2008 year-end annual report, "[l]osses on our derivative contracts that relate to jet fuel purchases in 2009 are deferred on our Consolidated Balance Sheet for 2008 and will be recognized in 2009 when the hedged jet fuel is purchased and consumed."[7] In its earnings call on April 21, 2009, Delta described to shareholders how hedges locked in during the high fuel prices in early 2008 resulted in fuel costs being elevated by \$0.34 per gallon relative to spot prices throughout the 2009 calendar year (see Appendix Figure A1).

2.3 Pricing Off Accounting Costs of Fuel

We estimate the extent to which accounting costs of fuel influence airlines' prices using the specification

$$\text{Log Fare}_{irt} = \gamma \text{Log Spot Fuel Price}_t + \varphi \text{Fuel Cost Diff}_{it} + \delta_{ir} + \varepsilon_{irt}, \tag{2}$$

where Fare_{irt} is the mean fare for tickets operated by carrier i on route r in quarter t, Spot Fuel Price_t is the spot price of jet fuel per gallon in quarter t, and $\text{Fuel Cost Diff}_{it} = \text{Log Fuel Cost}_{it} - \text{Log Spot Fuel Price}_t$ is the log difference between carrier i's accounting cost of jet fuel in quarter t and the spot price. The inclusion of carrier–route fixed effects δ_{ir} means that the coefficients φ and γ estimate how a carrier's fare on a given route fluctuates with the spot price of fuel and the carrier's accounting cost of fuel.

Table 2 reports the results from specification (2). Column 1 finds that a 10 percent increase in spot fuel prices is associated with 3.5 percent higher ticket fares charged by a carrier on a given route (i.e., $\gamma \approx 0.35$).[8] The log difference between the carrier's accounting cost of fuel and the spot price also significantly affects fares, and in fact is passed through at a rate similar to that of the spot price ($\varphi \approx 0.52$).

The effect of carriers' accounting costs of fuel on fares is surprising vis-à-vis standard price-setting theory because even frequent hedgers like Southwest and Delta do not hedge the entirety of their fuel needs. Thus, all carriers face the spot price for jet fuel purchases on the margin. Lump-sum gains and losses on hedges should not affect the marginal costs relevant for pricing decisions.

[7] See `https://www.sec.gov/Archives/edgar/data/27904/000119312509042726/d10k.htm`.

[8] Sangani (2025) points out that the log pass-through will equal the cost share of an input in revenues if firms pass through cost changes on a dollars-and-cents basis. Since we don't observe fuel needs for each route, we are unable to measure the pass-through in levels in this setting, but a log pass-through of 30–35 percent is broadly consistent with airlines' reported expenditures on fuel as a share of revenues.

Table 2: Effect of airlines' accounting costs of fuel on ticket prices.

	Log Fare$_{irt}$			*Fuel Cost Diff$_{it}$*
	OLS	RF	IV	FS
	(1)	(2)	(3)	(4)
Log Spot Fuel Price$_t$	0.348**	0.332**	0.354**	
	(0.028)	(0.030)	(0.030)	
Fuel Cost Diff$_{it}$	0.522**		0.613**	
	(0.106)		(0.184)	
$\Delta_{t,t-4}$ Log Spot Price × Hedge Ratio$_i$		−0.593**		−1.018**
		(0.192)		(0.127)
Carrier × Route FEs	Yes	Yes	Yes	Yes
N	1765168	1765536	1765168	1765168
R^2	0.58	0.58	0.58	0.37

Note: The outcome variable, Log Fare$_{irt}$, is the log of the average fare for carrier i on route r in quarter t in the DB1B sample. Spot Fuel Price$_t$ is the spot price of jet fuel from the EIA, and Fuel Cost Diff$_{it}$ = Log Fuel Cost$_{it}$ − Log Spot Fuel Price$_t$ is the log difference between carrier i's accounting cost of jet fuel in quarter t and the spot price. $\Delta_{t,t-4}$Log Spot Price is the change in the log spot price from quarter $t-4$ to quarter t, and Hedge Ratio$_i$ is one minus the pass-through of spot prices to accounting costs at four quarters estimated in Figure 2. Standard errors are two-way clustered by route and quarter. * indicates significance at the 10% level, ** at 5%.

A concern is that differences in accounting costs of fuel across airlines could also reflect genuine differences in fuel prices that carriers face on the spot market. The remaining columns of Table 2 show that variation in accounting costs of fuel specifically caused by carriers' hedging practices affects ticket fares. To do so, we use the change in the spot price of fuel over the past year interacted with an airline's hedge ratio as an instrument for the deviation in an airline's accounting cost of fuel from the spot price. We calculate each carrier's hedge ratio using the historical dependence of accounting costs on spot prices at one year, i.e., HedgeRatio$_i = 1 - \sum_{k=0}^{3} \beta_{ik}$ from (1).[9] Intuitively, a carrier's hedge ratio reflects the degree to which its accounting cost of fuel will not fully reflect recent changes in spot prices, and so the interaction of the hedge ratio with the change in the spot price over the past year predicts how the carrier's accounting fuel cost will differ from the spot price.

Columns 2–4 use this instrument to isolate how variation in accounting costs induced by hedging affects carriers' fares. As expected, when spot fuel prices have risen over the previous year, carriers that hedge more have accounting fuel costs that lie further

[9] Airlines sporadically report the percent of fuel needs one year out that they have hedged. We collect these data from Bloomberg and validate that the hedge ratios estimated using (1) are correlated with these reports.

below the prevailing spot price. The coefficient of -1.02 in column 4 shows that this relationship is roughly one-for-one: a carrier that hedges an additional 10 percent of its fuel needs experiences roughly a 1 percent larger decline in its accounting cost relative to the spot price following a 10 percent increase in spot prices. The reduced-form regression in column 2 shows that a higher hedge ratio is associated with less pass-through of recent changes in the spot fuel price to fares. Finally, column 3 instruments Fuel Cost Diff$_{it}$ with the hedging-induced variation in accounting fuel costs. The IV estimate ($\varphi \approx 0.61$) is similar in magnitude to the OLS estimate of 0.52 in column 1, showing that variation in the accounting cost of fuel generated by carriers' hedging practices is passed through to ticket fares.

Taken together with the pass-through of spot prices to accounting fuel costs in Figure 2, these results show that (i) accounting measures of fuel cost reflect historical spot prices for carriers that hedge, and (ii) since ticket fares appear to be influenced by carriers' accounting costs, the delayed pass-through of spot prices to accounting costs for hedgers translates into delayed pass-through of spot prices to fares.

Robustness to early ticket purchases. A threat to identification in this setting is that airlines choose how much to hedge, and characteristics that shape hedging decisions may also directly affect how airlines' fares respond to fuel prices. While the carrier–route fixed effects in (2) absorb persistent differences in fare levels across airlines, they do not absorb differences in airlines' desired pass-through of spot fuel prices coming from unobserved differences in cost structure, financial position, or customer base.

An especially salient version of this concern involves early ticket purchases. Airlines whose customers book flights further ahead may have stronger incentives to hedge fuel costs because they commit to prices in advance. For example, Southwest's larger share of leisure travelers—who tend to buy tickets early—could explain both its greater propensity to hedge fuel costs and the stronger relationship between its observed fares and historical spot fuel prices. In this case, the apparent dependence of Southwest's fares on accounting costs arises because both its ticket revenues and its fuel costs are locked in early, rather than because it is pricing off accounting costs.

We address this concern using the new DB1C data collected by the BTS starting in July 2025, which include indicators for whether each ticket itinerary was purchased more than 21 days or 90 days in advance.[10] Appendix Figure A2 plots the share of tickets purchased

[10]The DB1C database includes a 40 percent sample of all itineraries, compared to the 10 percent DB1B data it succeeds. The disclosure of three broad buckets—tickets purchased within 21 days of departure, between 21 and 90 days before departure, and more than 90 days before departure—was the outcome of negotiations between the Department of Transportation and industry trade associations, since airlines

within 21 days or within 90 days of departure against each airline's pass-through of spot fuel prices to accounting costs. These are the complements of the early-purchase shares and hedge ratios, respectively. The positive correlations therefore indicate that airlines with more advance purchases tend to hedge more. However, there remains substantial variation in hedging practices across airlines with similar early-purchase shares. For example, Southwest, JetBlue, and Alaska Airlines all have a similar share of customers that buy tickets more than 21 days and more than 90 days in advance; however, of the three airlines, only Southwest appears to engage in extensive fuel hedging. Likewise, Delta, American, and United Airlines all face similar early-purchase shares, but Delta has historically engaged in much more hedging than the other two. We view the variation in hedging across airlines that have similar customer purchase patterns as largely reflecting differences in managerial strategy and outlook.[11]

Appendix Table A2 shows that while higher early-purchase shares indeed predict lower pass-through of recent changes in the spot fuel price, the estimated effect of carriers' accounting fuel costs on fares remains significant and similar in magnitude to our baseline results. In other words, airlines' accounting fuel costs appear to shape ticket prices even after controlling for the share of tickets that are sold in advance.

Evidence from accounting margins. Complementary evidence that airlines price off accounting costs comes from data on airlines' operating margins. If airlines set fares based on prevailing spot fuel prices, hedging gains and losses should generate fluctuations in reported profits. For example, when spot fuel prices rise, fares rise with them while reported fuel expenses remain low, temporarily raising accounting margins. Revaluing fuel expenses at the spot price should remove this source of volatility. Conversely, if airlines set fares to track accounting fuel costs, revenues and reported expenses move together, and replacing accounting fuel costs with costs valued at the spot price can make margins more volatile, especially for airlines that hedge extensively.

We compare the volatility of airlines' reported operating margins—which use accounting costs of fuel—with the volatility of operating margins calculated after revaluing

argued that exact ticket purchase and departure dates were commercially sensitive information.

[11]For example, Delta partially hedges fuel costs by owning the Pennsylvania Trainer Refinery, which it purchased in 2012. In 2018, Forbes wrote, "Plenty of experts in both the airline and the oil industries said at the time that Delta's purchase [...] was a bad idea, and have continued to say so ever since. But only now are Delta's own executives, who steadfastly have kept on defending the six-year-old mistake, beginning to admit it" (*Forbes*, "Delta Belatedly Is Facing Up To Its One Big Mistake"). Likewise, when Southwest decided to close its hedging portfolio in 2025 (ironically, just a year before tensions in the Strait of Hormuz in March 2026 caused jet fuel prices to double), a senior equity analyst at JPMorgan Chase commented, "We still don't understand why Southwest clung to fuel hedging for this long" (*Wall Street Journal*, "Why Southwest Airlines Finally Soured On Jet Fuel Hedging").

consumed fuel at the prevailing spot price. Appendix Table A3 reports the volatility of the unadjusted and adjusted operating margins for each airline using public quarterly financial reports over 1990–2026, extending beyond the 1990–2025 fare sample. (We exclude the COVID-19 period from 2020–2021, which saw operating margins plummet to negative values for all airlines.) Revaluing fuel consumption expenses at the spot price sharply increases the volatility of operating margins for hedgers: for Southwest and Delta, the adjustment raises the volatility of operating margins by 31 and 10 percent, respectively. This pattern corroborates the ticket-level evidence that revenues track accounting fuel costs rather than spot prices of fuel.

Other extensions. Appendix Table A2 also explores asymmetries and strategic complementarities. We find that accounting costs of fuel both above and below the spot price are passed through to fares—that is, the influence of accounting costs on prices does not appear to be limited to periods when accounting fuel costs exceed the spot price. We also find evidence of strategic complementarities: the pass-through of a carrier's accounting cost of fuel to fares is more pronounced on routes where the airline has a higher market share, and fares appear to be affected by both a carrier's own fuel costs and the average fuel cost of other airlines servicing the same route.

3 Evidence from Quarterly Financial Reports

In this section, we develop an approach to estimate the extent of accounting cost pricing using firms' quarterly financial reports. Our approach exploits the fact that, while a firm's cost of goods sold will reflect the price that the firm paid for inputs in the past, the book value of inventory captures information about input prices that the firm currently faces. The advantage of this approach is that we can estimate the extent of accounting cost pricing solely using firms' public, quarterly financial reports, without microdata on prices or detailed data on input cost shocks, as we used in the case study of airlines.

3.1 Estimation Approach

The model in Section 4 predicts that firms' prices are a blended average of prices based on replacement costs and accounting costs, i.e., $p_t \approx \mu(\phi c_t + (1-\phi)h_t)$. Taking the ratio of sales to cost of goods sold therefore yields

$$\frac{\text{Sales}_t}{\text{COGS}_t} = \frac{p_t D_t}{h_t D_t} \approx \frac{\mu(\phi c_t + (1-\phi)h_t)D_t}{h_t D_t} = \mu\left(\phi \frac{c_t}{h_t} + (1-\phi)\right). \tag{3}$$

Intuitively, if firms put a positive weight on replacement costs ($\phi > 0$), then the ratio of sales to cost of goods sold should be systematically higher when the current replacement cost is greater than the accounting unit cost. Conversely, if firms price entirely off accounting costs ($\phi = 0$), the ratio of sales to cost of goods sold should be insensitive to fluctuations in replacement costs relative to accounting costs.

How can we observe fluctuations in current replacement costs relative to historical input purchase costs? Notice that changes in the book value of inventory reflect purchases of new inputs, made at the current cost c_t, minus the cost of goods sold, valued at the accounting unit cost h_t:

$$\text{BookValueInventory}_t = \text{BookValueInventory}_{t-1} - \underbrace{\text{COGS}_t}_{h_t D_t} + \underbrace{\text{Purchases}_t}_{c_t Y_t}.$$

Thus, the ratio of purchases to cost of goods sold can be calculated using changes in the book value of inventory from a firm's balance sheet and cost of goods sold from the firm's income statement, and is equal to

$$\frac{\text{Purchases}_t}{\text{COGS}_t} = \frac{\Delta\text{BookValueInventory}_t + \text{COGS}_t}{\text{COGS}_t} = \frac{c_t Y_t}{h_t D_t}. \tag{4}$$

The ratio of purchases to cost of goods sold contains two pieces of information. First, it reflects how current replacement costs compare to accounting costs (c_t/h_t). Second, it reflects how the quantity of purchases compares to the quantity of shipments (Y_t/D_t). The former ratio of current and historical accounting costs is precisely the ratio we are interested in to understand the degree of accounting cost pricing; the latter is a fluctuation in relative quantities that we will need to treat as an unobserved, potential confounder in our estimation strategy.

Log-linearizing (3) and (4) around a steady state with constant costs, sales, and purchases, we find

$$\left(\widehat{\frac{\text{Sales}}{\text{COGS}}}\right)_t = \phi\left(\widehat{\frac{\text{Purchases}}{\text{COGS}}}\right)_t - \phi\left(\hat{Y}_t - \hat{D}_t\right). \tag{5}$$

Thus, absent changes in markups and changes in the relative quantities of purchases and shipments, the extent of replacement cost pricing ϕ is identified by the sensitivity of the sales–COGS ratio to the purchases–COGS ratio. Our estimation strategy for ϕ will exploit this relationship between the ratio of purchases to cost of goods sold and the fluctuations in replacement costs relative to accounting costs.

Does the purchases–COGS ratio capture variation in c_t/h_t? Before turning to the estimation, it is useful to verify whether the ratio of purchases to cost of goods sold indeed captures meaningful variation in the gap between replacement costs and accounting costs. We do so using two firms whose input costs are closely tied to a single, publicly traded commodity: Starbucks and Gildan Activewear. Starbucks, which carries roughly one month of inventory on average, purchases large quantities of Arabica coffee beans to produce its coffee products. Gildan, which is one of the world's largest manufacturers of basic apparel, holds about six months of inventory on average, much of which reflects the cost of cotton and other textile inputs. Because both firms use commodity inputs to produce differentiated, branded products, we expect their purchases–COGS ratios to track fluctuations in the respective commodity prices.

Figure 3 compares the ratio of purchases to cost of goods sold for Starbucks and Gildan to commodity prices for coffee and cotton, respectively. In each figure, we construct the ratio of purchases to cost of goods sold using the sum of purchases and cost of goods sold over the trailing four quarters, to minimize the influence of seasonal variation.[12] We plot this ratio alongside the year-over-year inflation rate for the matched commodity series, which we lag by two quarters to approximately match the timing of the four-quarter trailing sums of purchases and cost of goods sold.

Both examples suggest that movements in the ratio of purchases to cost of goods sold at least in part reflect fluctuations in replacement costs relative to accounting costs. In both cases, rising input costs lead to a rise in the firm's ratio of purchases to accounting costs. For example, the purchases–COGS ratio for Starbucks shows pronounced peaks in 1995, 1997, 2011, and 2022, coinciding with increases in coffee commodity prices. Likewise, the purchases–COGS ratio for Gildan peaks in 2001, 2011, and 2023, alongside sharp increases in cotton commodity prices. These examples suggest that the ratio of purchases to cost of goods sold indeed captures variation in the gap between replacement costs and accounting costs necessary to estimate ϕ.

3.2 Estimation and Results

Data and specification. We use quarterly financial statements for publicly traded firms from Compustat. Appendix Table A5 enumerates filters that we use for our estimation sample. We exclude firms with missing NAICS codes and firms in the information (NAICS 51) and finance, insurance, and real estate (NAICS 52–53) sectors. We require that

[12] Figure 3 truncates the time series for Gildan before the Gildan–Hanes merger in 2025Q4, which led to large changes in the firm's reported assets and inventories unrelated to cost fluctuations. We apply filters for mergers and spin-offs in our estimation sample, described below.

Figure 3: Comparing Purchases / COGS to commodity input prices: Two examples.

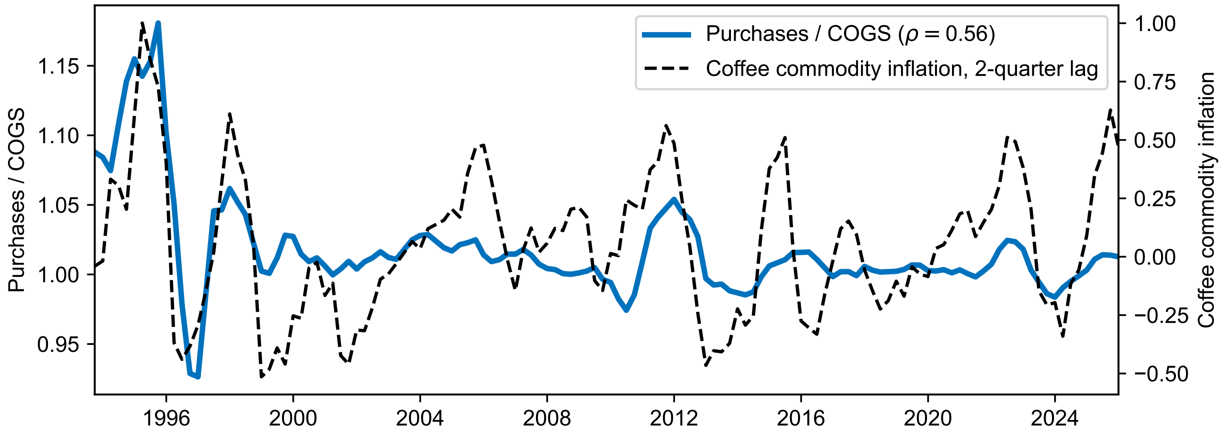

(a) Starbucks and coffee commodity prices.

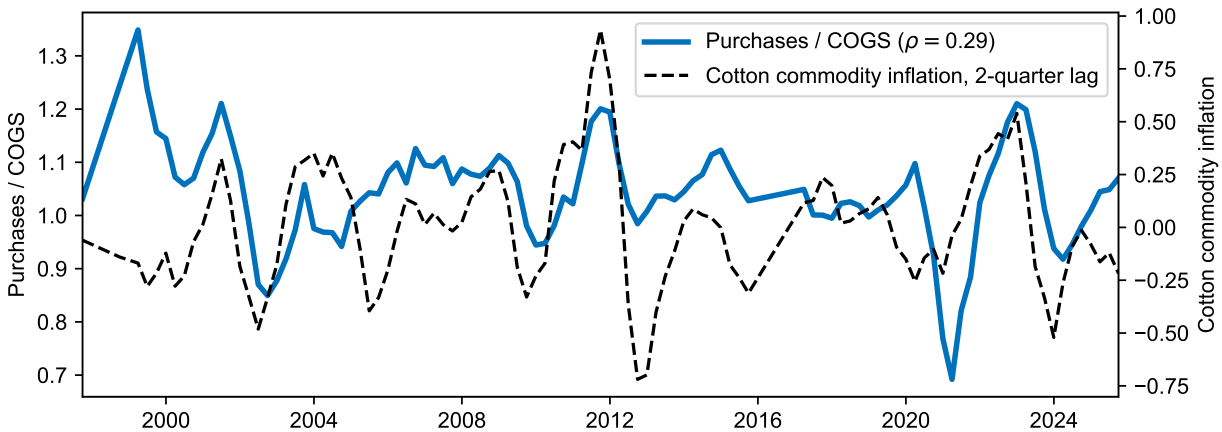

(b) Gildan and cotton commodity prices.

observations have strictly positive values for inventory stocks, sales, and costs of goods sold. Since we construct purchases using the change in the book value of inventory, we also exclude observations where changes in inventory stocks are likely to be driven by mergers or spin-offs. We do so by excluding observations where Compustat deems the observation not comparable to previous quarters (recorded in the field `COMPSTQ`) and where the year-over-year change in the firm's total assets is greater than 30 percent. Finally, we restrict the sample to firms that report using FIFO or AAC accounting as their inventory valuation method. This is because, under LIFO, costs of goods sold already reflect the firm's most recent input purchase prices, and so $h_t \approx c_t$ by construction. The final sample includes about 15,000 firms, with an average of 29 quarterly observations per firm.

Using these data, we estimate the specification

$$\log(\text{Sales}/\text{COGS})_{it} = \phi \log(\text{Purchases}/\text{COGS})_{it} + \alpha_{iq} + \varepsilon_{it}, \tag{6}$$

where $(\text{Sales}/\text{COGS})_{it}$ and $(\text{Purchases}/\text{COGS})_{it}$ are firm i's ratios of sales and purchases, respectively, to cost of goods sold in quarter t, α_{iq} are firm $\times$ quarter-of-year fixed effects, and ε_{it} is a mean-zero error.

Comparing (6) to (5) shows that the term $\phi\left(\hat{Y}_t - \hat{D}_t\right)$, which depends on the ratio of the quantity of purchases to shipments, is an omitted variable in (6). This omission introduces two biases, which can be seen in the expression below.

$$\hat{\phi} \to \phi\left[1 - \frac{Var\left(\hat{Y}_t - \hat{D}_t\right)}{Var\left(\widehat{\frac{\text{Purchases}}{\text{COGS}}}\right)} + \frac{Cov\left(\hat{h}_t - \hat{c}_t, \hat{Y}_t - \hat{D}_t\right)}{Var\left(\widehat{\frac{\text{Purchases}}{\text{COGS}}}\right)}\right].$$

First, fluctuations in the ratio of the quantity of purchases to shipments mean that the ratio of purchases to cost of goods sold is a noisy measure of c_t/h_t, therefore introducing an attenuation bias from measurement error. Second, quantities of purchases and shipments may move systematically with cost fluctuations. For example, if firms tend to reduce purchases when input costs rise and stock up when input costs fall below their historical levels (i.e., $Cov\left(\hat{c}_t - \hat{h}_t, \hat{Y}_t - \hat{D}_t\right) < 0$), then the omitted variable will bias our estimate of ϕ upward.

We take a few steps to mitigate this omitted variable bias. First, one of the main reasons why quantities of purchases and shipments may fluctuate over time is seasonality: many firms build up inventory during one part of the year and have higher sales volumes during other parts of the year. The firm $\times$ quarter-of-year fixed effects α_{iq} in (6) absorb such firm-specific seasonal patterns.

Second, to account for possible firm-specific trends in purchases and shipments over time, we also estimate the first-differences specification

$$\begin{aligned}\log(\text{Sales}/\text{COGS})_{it} &- \log(\text{Sales}/\text{COGS})_{it-4} \\ &= \phi\left[\log(\text{Purchases}/\text{COGS})_{it} - \log(\text{Purchases}/\text{COGS})_{it-4}\right] + \alpha_{iq} + \varepsilon_{it}.\end{aligned} \tag{7}$$

The first-differences specification uses year-over-year changes in both the dependent and independent variables, to ensure our results are not driven by trends in firms' shipment–purchase ratios, or by secular trends in inventory holdings or markups more broadly.

Finally, we also estimate both (6) and (7) using oil supply shocks from Känzig (2021)

Table 3: Estimating ϕ, the extent of replacement cost pricing.

	Log(Sales / COGS)			
	Levels		First Differences	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Log(Purchases / COGS)	0.147**	0.380**	0.126**	0.351**
	(0.015)	(0.104)	(0.014)	(0.079)
Firm × Quarter-of-Year FEs	Yes	Yes	Yes	Yes
N	382 211	369 469	345 468	333 935
R^2	0.86	0.86	0.09	0.04

Note: Columns 1–2 report results from (6), and columns 3–4 report results from (7). We exclude observations in the bottom and top 1 percent of values for the dependent and independent variables. Columns 2 and 4 use oil supply shocks from Känzig (2021) interacted with NAICS-4 fixed effects as an instrument for Log(Purchases / COGS). Regressions are weighted by sales. Driscoll-Kraay standard errors with eight lags in parentheses. * indicates significance at 10%, ** at 5%.

interacted with NAICS-4 industry fixed effects as an instrument for log(Purchases/COGS). Note that while the instrument should remove the attenuation bias from measurement error, it does not address the bias that results from how quantities of purchases and shipments respond to cost changes. If purchase quantities tend to fall when input costs rise (which strikes us as the plausible case), our IV estimates will tend to overestimate the extent of replacement cost pricing and thus underestimate the extent of accounting cost pricing.

Baseline results. Table 3 columns 1 and 3 report the results from estimating specifications (6) and (7). We find $\phi \in [0.12, 0.15]$, suggesting a substantial extent of accounting cost pricing. Columns 2 and 4 report IV estimates that place $\phi \in [0.35, 0.40]$.

We find similar point estimates for ϕ under a number of alternate specifications in Appendix Table A6. We extend specifications (6) and (7) to include time fixed effects, thus exploiting relative differences in cost fluctuations facing firms, and find a similar extent of accounting cost pricing for both firms using AAC and FIFO accounting. We further include a control for the ratio of purchases to cost of goods sold for other firms in the same NAICS-4 industry using LIFO accounting. The idea is that, under LIFO accounting, $h_t \approx c_t$, and thus changes in the ratio of purchases to cost of goods sold reflect only changes in quantities of purchases and shipments. To the extent that purchase and shipment quantities are correlated across firms in an industry, we can then use these LIFO peers to control for changes in purchase and shipment quantities. Across these

Table 4: Heterogeneity in the extent of replacement cost pricing by industry.

	Log(Sales / COGS)					
	Levels		First Differences			
	OLS	OLS	OLS	OLS	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Purchases / COGS)	0.122**	0.125**	0.101**	0.103**	−0.268	−0.136
	(0.016)	(0.019)	(0.017)	(0.016)	(0.376)	(0.265)
× Reference prices (r)	0.079**		0.074*		1.006*	
	(0.029)		(0.042)		(0.599)	
× Organized exchange (w)	0.254**		0.168**		0.720**	
	(0.062)		(0.051)		(0.329)	
× Industry LIFO share		0.185**		0.144**		0.744**
		(0.090)		(0.068)		(0.353)
Firm × Quarter-of-Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	251 379	276 804	227 964	251 033	225 546	242 882
R^2	0.88	0.88	0.09	0.09	0.08	0.01

Note: Industries with reference prices (r) and organized exchanges (w) are identified using the Rauch (1999) liberal classification, which we map to NAICS codes using a SITC-NAICS concordance. The industry LIFO share is constructed using data from the 2017 Economic Census. Columns 5–6 use oil supply shocks from Känzig (2021) interacted with NAICS-4 fixed effects as an instrument for Log(Purchases / COGS). Regressions are weighted by sales. Driscoll-Kraay standard errors with eight lags in parentheses. * indicates significance at 10%, ** at 5%.

specifications, we find similar estimates of ϕ to the baseline estimates in Table 3.

Heterogeneity by industry characteristics. Section 4 predicts that the extent of replacement cost pricing should rise in competitive industries with homogeneous output goods and in industries with readily observable costs. We test this prediction in Table 4 by interacting the firm's ratio of purchases to cost of goods sold with an indicator for whether the firm's NAICS industry has reference prices or an organized exchange, using categorizations by Rauch (1999). We find a higher extent of replacement cost pricing in both industries with reference prices and organized exchange markets. In fact, IV estimates suggest that firms in these industries largely account for the positive weight on replacement cost prices in the sample. Focusing exclusively on firms in industries with administered prices, the OLS estimates of the extent of replacement cost pricing are $\phi \in [0.10, 0.13]$, and IV estimates are not significantly different from zero.

We also expect that firms should exhibit greater replacement cost pricing if their competitors' prices are more closely correlated with recent costs rather than historical input costs. Table 4 interacts the firm's ratio of purchases to cost of goods sold with the share of inventories in a firm's NAICS industry that are valued using LIFO accounting. In line with our predictions, the extent of replacement cost pricing increases with the share of a firm's competitors that use LIFO. For firms in industries where LIFO is rare, our OLS estimates of the extent of replacement cost pricing are again in the range $\phi \in [0.10, 0.13]$, and IV estimates are not significantly different from zero.

Thus, Table 4 suggests two conclusions. First, the comparative statics predicted by the theory—that the extent of replacement cost pricing should be higher for firms in industries with competitive, homogeneous outputs and with observable input costs—appear to borne out in the data. Second, the degree to which public firms using FIFO and AAC price off replacement costs appears to be largely driven by firms in those industries. For firms in industries with administered prices and few LIFO competitors, the estimated extent of accounting cost pricing ranges from 85% to complete.

4 A Two-Period Model of Accounting Cost Pricing

This section develops a two-period model of accounting cost pricing. The purpose of the model is twofold. First, the model generates comparative statics for when accounting cost pricing or replacement cost pricing should be more prevalent, which we can verify in the data. Second, the model explains how a pricing rule that depends on backward-looking cost measures can be the outcome of agents optimizing with rational expectations. The latter is important to reconcile pricing off accounting costs with the effects of trend inflation or inflation expectations.

The model we develop in this section builds on the justification for backward-looking cost measures described by Okun (1981),

> The need for objective and systematic measurement forces the cost standard to be backward-looking—or at least sideward- rather than forward-looking. [...] Any attempt to forecast input costs or to put them on a current replacement basis, which necessarily entails a forecast, would add to expenses, introduce a subjective element into cost calculations, and complicate the task of managerial control. Bygones cannot be mere bygones if they are the sole sources of systematic information.

In the model, a principal delegates the operation of a firm to a manager. The model features two key frictions to formalize Okun's (1981) intuitions. First, the principal can observe and verify accounting profits, but cannot precisely observe the replacement costs that determine the firm's optimal price. Second, the manager may have private preferences that are not aligned with maximizing the firm's value. The risk of managerial discretion leads the principal to penalize the manager for deviating too sharply from a target accounting markup.

4.1 Model Setup

Figure 4 shows the timing of the two-period model. At time $t = 0$, the firm manager purchases I units of inventory at price h. At time $t = 1$, the manager observes the price c at which he can liquidate inventory in the future ($t = 2$). After observing c, the manager chooses a price p to set at $t = 1$ and sells $D(p)$ units of the inventory. In the final period, the manager liquidates the leftover inventory, $I - D(p)$, at the liquidation price c.

Figure 4: Two-period model timing.

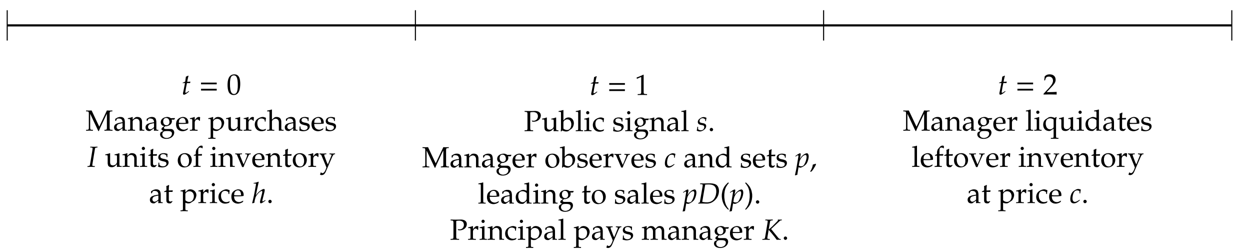

For simplicity, we assume that the inventory stock I purchased in the first period is sufficiently large that $D(p) < I$ for all feasible values of p, so that the optimal price is not affected by the risk of stockouts. We further assume no discounting. Thus, the first-best value of the firm is

$$V^{\text{FB}} = \max_p \left\{-hI + pD(p) + c(I - D(p))\right\}. \tag{8}$$

Notice that the optimal price does not depend on the initial price of inventory h. Instead, the value of the firm is maximized by choosing $p^* = \arg\max_p (p - c)D(p)$, i.e., the price that maximizes profits using the liquidation price c as the measure of the marginal cost. Intuitively, the liquidation price c reflects the opportunity cost of selling one more unit of inventory, analogous to the replacement cost in a dynamic model.

We now introduce two frictions. First, we assume that the principal can observe the firm's sales, $pD(p)$, and the book value of the goods sold, $hD(p)$, but cannot observe the liquidation price c. Instead, the liquidation price c is the private information of the manager. Along with sales and the book value of goods sold, we allow the principal to observe a noisy signal s of how the liquidation price compares to the acquisition price at the beginning of period $t = 1$. The signal s represents information, such as trend inflation or market prices of related inputs, that can be used to estimate how the liquidation price c relates to the acquisition price h.

Second, we assume that delegating the pricing decision to the manager can be costly because the manager's goals may not be aligned with the principal. The manager's value function is

$$V^{\text{manager}} = \max_{p} \Big\{ \underbrace{-hI + pD(p) + c(I - D(p))}_{\text{Firm value}} + \underbrace{(\alpha - 1)p}_{\text{Pricing wedge}} + \underbrace{K(pD(p), hD(p), s)}_{\text{Incentive payment}} \Big\}. \tag{9}$$

The manager's value function includes the true value of the firm, which is the first bracketed term. But it also includes a private benefit that the manager enjoys from setting the price above or below the optimum, which is the second term. This private benefit could reflect a manager's preference for prices that are lower than the value-maximizing level ($\alpha < 1$), for example due to empire-building motives (Baumol 1959), or a preference for higher-than-optimal prices ($\alpha > 1$), for example to avoid intense competition with rivals (Bertrand and Mullainathan 2003). More broadly, the wedge introduced by $\alpha \neq 1$ stands in for any way in which managerial discretion, inattention, or slack may distort pricing at the expense of the firm's value. We assume that α is the private information of the manager and cannot be observed by the principal.

The third term is a payment K that the principal can make to the manager. Crucially, the principal is limited in what information she can use to determine the payment. She observes accounting quantities, but cannot observe the true opportunity cost of a sale, c. Moreover, since the manager's pricing wedge α is private information, the manager also cannot credibly communicate the liquidation price c to the principal: regardless of the true value of c, the manager has an incentive to misreport it in order to justify his desired price. The principal must therefore condition the payment solely on observable accounting quantities—sales $pD(p)$ and the book value of goods sold $hD(p)$—and the public signal s. We assume that the payment schedule $K(\cdot)$ chosen by the principal is known by the manager.

A natural question is why the principal cannot wait until $t = 2$ and condition the payment on the liquidation proceeds, which reveal c to both parties. Our assumption that the contract is tied to short-term accounting performance, rather than the long-run value of decisions, reflects the need to monitor and incentivize managers in a timely manner. Conditioning on liquidation proceeds at $t = 2$ requires the manager to commit to a longer-term contract and attribute distant outcomes to current decisions, both of which are likely difficult in practice.[13]

Given the manager's pricing choice, the principal's problem is to design the payment terms, $K(\cdot)$, to maximize expected firm value:[14]

$$V^{\text{principal}} = \max_{K(\cdot)} \mathbb{E}\left[-hI + pD\left(p\right) + c\left(I - D\left(p\right)\right)\right]. \tag{10}$$

The principal faces a trade-off when choosing $K(\cdot)$: higher-powered incentives limit the extent to which the manager will allow his idiosyncratic preferences, captured by α, to erode firm value, but they also lead the manager to disregard his valuable private information about the liquidation price c. In this sense, the model captures the intuitions described by Okun (1981): the manager has a systematic informational edge in determining the value-maximizing price, but because of the misaligned incentives between manager and principal, any attempt to convey replacement costs to the principal would be colored by the manager's subjective outlook on the best price, and the principal is unwilling to let 'bygones be bygones.'

4.2 Assumptions and Solution Method

We assume that demand is isoelastic, $D(p) = p^{-\eta}$, where the price elasticity $\eta > 1$. We assume that the liquidation price, public signal, and managerial preference take the form

$$c = h \cdot \exp(\delta C), \qquad s = (c/h)\exp(\delta S), \qquad \alpha = \exp(\delta A),$$

where C, S, and A are independently distributed according to

$$C \sim N(0, \sigma_c^2), \qquad S \sim N(0, \tau^{-1}), \qquad A \sim N(0, \sigma_\alpha^2),$$

and where δ parameterizes the degree of frictions.

We characterize the principal's optimal contract and the manager's pricing rule locally around $h = c = \alpha = s = 1$. Specifically, we let $\log h = O(\delta)$ and take $\delta \to 0$, holding the distributions of C, S, and A fixed. For any variable x, let $\hat{x}$ denote the log-deviation of x from its value when $\delta = 0$. For example, $\hat{p} = \log p - \log(\eta/(\eta - 1))$. We approximate the manager's and principal's payoffs to second order in δ and report the resulting first-order

[13]Equity compensation can link managerial rewards to subsequent liquidation proceeds if the manager must retain the equity until those proceeds are realized. However, if the manager's payoff depends on the equity's market value at $t = 1$, and c is unobserved by market participants, the value of the equity is likewise affected by what the market can infer about c and α from observed accounting metrics.

[14]We abstract from the resource cost of providing these incentives and study the payment schedule chosen by the principal that maximizes expected firm value before managerial compensation.

pricing rules.

4.3 Pricing With and Without Incentives

Proposition 1 characterizes the price that the manager would set in the absence of incentives from the principal.

Proposition 1 (Pricing without incentives)**.** *If $K(\cdot) = 0$, then the price set by the manager is*

$$\hat{p}^{\text{no incent}} = \hat{c} + \psi_0 \hat{\alpha}, \qquad \textit{where} \qquad \psi_0 = \frac{1}{\eta - 1}\left(\frac{\eta}{\eta - 1}\right)^{\eta}.$$

The elasticity of the price to the managerial preference, ψ_0, is decreasing in the demand elasticity η.

In the absence of incentives, the price set by the manager fully incorporates information about the opportunity cost of selling inventory: changes in the liquidation value c pass through one-for-one to prices. However, the manager's idiosyncratic preferences α also distort prices. The degree of distortion falls as the elasticity of demand η increases. Intuitively, when the elasticity of substitution is high, deviating from the optimal price leads to a more rapid fall-off in firm value. Thus, market forces penalize the manager for diverging too far from the optimum. This mechanism captures the idea that managerial slack is more difficult to sustain in competitive industries (Hart 1983; Giroud and Mueller 2010, 2011).

The manager's objective can be written to second order in δ (omitting terms independent of the chosen price) as

$$\max_{\hat{p}} \left\{ -\frac{\eta}{2}\left(\frac{\eta}{\eta - 1}\right)^{-\eta} \left(\hat{p} - \hat{p}^{\text{no incent}}\right)^2 + K(pD(p), hD(p), s) + o(\delta^2) \right\}.$$

Conditional on the acquisition price h and the public signal s, the manager's private information (c, α) only affects his choice of price through $\hat{p}^{\text{no incent}}$. Moreover, sales and accounting costs are functions of the manager's chosen price. This implies that there is no contract on observables at $t = 1$ that the principal can use to fully separate the effects of c from α on the manager's chosen price; that is, the principal cannot choose a contract $K(\cdot)$ that always leads to the first-best price. Furthermore, it means that any contract chosen by the principal will induce a pricing rule that depends on $\hat{p}^{\text{no incent}}$, along with the observables h and s.

We proceed as follows. First, we derive the best possible pricing rule that the principal could induce the manager to follow, ignoring implementability constraints. Then, we show that the pricing rule can be implemented by a simple contract rule that penalizes the manager for deviating too far from a target accounting markup.

Lemma 1 (Optimal pricing rule)**.** *Given observables, the principal can do no better than a contract that induces the pricing rule*

$$\hat{p}^*(\hat{p}^{no\ incent}, \hat{h}, \hat{s}) = \mathbb{E}\left[\hat{c} \mid \hat{p}^{no\ incent}, \hat{h}, \hat{s}\right] = \lambda \hat{p}^{no\ incent} + (1-\lambda)\mathbb{E}[\hat{c}|\hat{h}, \hat{s}],$$

where $\lambda = 1/(1 + \psi_0^2(\sigma_\alpha^2/\sigma_c^2)(1 + \tau\sigma_c^2))$.

Given isoelastic demand and the liquidation price c, the firm's value is maximized by setting $\hat{p} = \hat{c}$. Thus, an optimal contract would induce the manager to set the price equal to the expected value of $\hat{c}$ given the information observed by the principal and the manager's private information summarized by $\hat{p}^{\text{no incent}}$. The expected value of $\hat{c}$ is given by the expression in Lemma 1, with the coefficient λ reflecting the relative informativeness of the principal's observables and the manager's no-incentive price for the liquidation cost $\hat{c}$.

This optimal pricing rule can be implemented by a simple contract that penalizes deviations from a target accounting markup. Proposition 2 characterizes the parameters of this contract.

Proposition 2 (Optimal contract parameters)**.** *The optimal pricing rule can be implemented with a contract payment that takes the form*

$$K = \bar{K}(s) - \frac{\kappa^*}{2}\left(\frac{pD(p)}{hD(p)} - \mu^*(s)\right)^2,$$

where $\bar{K}(s)$ *is an arbitrary function depending on* s*, and*

$$\underbrace{\mu^*(s) = \frac{\eta}{\eta-1}\exp\left(\frac{\tau\sigma_c^2}{1+\tau\sigma_c^2}\hat{s}\right),}_{\text{Target accounting markup}} \qquad \underbrace{\kappa^* = \frac{1}{\eta}\left(\frac{\eta}{\eta-1}\right)^\eta \frac{\sigma_\alpha^2}{\sigma_c^2}\left(1+\tau\sigma_c^2\right).}_{\text{Strength of incentives}}$$

The strength of incentives κ^* *is decreasing in the firm's demand elasticity* η*, increasing in the variance of managerial preferences* σ_α^2*, and decreasing in the variance of the cost shock* σ_c^2*.*

The contract in Proposition 2 has two parameters of interest: $\mu^*(s)$, which is the target accounting markup, and κ^*, which determines the strength of the incentives.

The target accounting markup includes the standard monopoly markup, $\eta/(\eta-1)$, and an adjustment to the book value of inventory based on the public signal s. Intuitively, the public signal effectively allows the principal to "mark-to-market" the firm's inventory, given the principal's estimate of the value of the firm's inventory stock relative to its book value.

The degree to which the principal's payment to the manager falls as the manager deviates from the target accounting markup is captured by κ^*. This parameter is increasing in the ratio of the variance of the manager's pricing wedge to the variance of cost shocks, $\sigma_\alpha^2/\sigma_c^2$. Intuitively, a greater risk of managerial discretion warrants stronger incentives, but incentives should be tempered when costs are volatile and the manager's private information about c is especially valuable. A more competitive market (i.e., a higher elasticity of substitution η) leads to lower-powered incentives, because market forces limit the degree to which the manager prices according to his own idiosyncratic preferences. Finally, greater visibility into the real-time value of inventory—an increase in the signal precision τ—also allows the principal to set higher-powered incentives.

Proposition 3 presents the price set by the manager under the optimal contract.

Proposition 3 (Pricing with optimal contract)**.** *Under the optimal contract, the manager's expected log price, averaging over s, is*

$$\mathbb{E}[\hat{p}|c, \alpha, h] = \underbrace{\phi\,[\hat{c} + \psi\hat{\alpha}]}_{\substack{\text{Replacement} \\ \text{cost pricing}}} + \underbrace{(1-\phi)\,\hat{h}}_{\substack{\text{Accounting} \\ \text{cost pricing}}},$$

where

$$1-\phi = \frac{\sigma_\alpha^2/\sigma_c^2}{\sigma_\alpha^2/\sigma_c^2 + \psi_0^{-2} + \tau\sigma_\alpha^2}, \qquad \text{and} \qquad \psi = \frac{1}{1+\tau\sigma_\alpha^2\psi_0^2}\psi_0.$$

The degree of accounting cost pricing $(1-\phi)$ *is increasing in the variance of managerial preferences* σ_α^2*, decreasing in the variance of the cost shock* σ_c^2*, decreasing in the demand elasticity* η*, and decreasing in the signal precision* τ*.*

Under the optimal contract, the manager's price is a weighted average of two components: one that depends on the true liquidation price as well as the manager's preference, and another that depends on the accounting cost at which inventory was acquired. We refer to the former as *replacement cost pricing* and to the latter as *accounting cost pricing*.

The extent of accounting cost pricing, $1-\phi$, is increasing in the relative importance of the agency problem, which is captured by the ratio of variances $\sigma_\alpha^2/\sigma_c^2$. When managerial discretion is the dominant concern ($\sigma_\alpha^2/\sigma_c^2 \to \infty$) and the principal cannot verify the liquidation price (i.e., $\tau = 0$), the principal forces the manager to price entirely off the historical acquisition cost h. Conversely, when the manager's private information about replacement costs is especially valuable relative to agency concerns ($\sigma_\alpha^2/\sigma_c^2 \to 0$ and $\tau = 0$), incentives are weak and the manager prices off the true marginal cost as in Proposition 1.

The extent of accounting cost pricing also falls when the elasticity of demand η is high: market competition prevents the manager from incorporating his own preferences into prices, obviating the need for a strict contract. Thus, the model predicts that accounting cost pricing should be less prevalent in industries where firms sell undifferentiated, highly substitutable goods. In fact, Okun (1981) comments on this case:

> The reliance of firms on historical costs is evident in many empirical studies of price-cost relations. [....] The exceptions that occur to me arise when processed raw materials are sold on auction markets, as in the case of flour or soybean meal. In those cases, the processed products promptly reflect changes in the (auction-market) prices of the primary products (Okun 1981, pp. 156–157).

Finally, the extent of accounting cost pricing falls when the principal can obtain a precise signal of replacement costs. This can explain why sellers of gold and silver jewelry often set prices using current gold and silver market prices, rather than the cost they incurred for those inputs. In fact, from an accounting rule perspective, FASB ARB 43 explicitly allows inventories to be marked to market in such cases,

> "Only in exceptional cases may inventories properly be stated above cost. For example, precious metals having a fixed monetary value with no substantial cost of marketing may be stated at such monetary value; any other exceptions must be justifiable by inability to determine appropriate approximate costs, immediate marketability at quoted market price, and the characteristic of unit interchangeability."

4.4 Extensions: Average Cost Pricing and Demand Shocks

The model shows that principal–agent frictions can lead managers to price off accounting costs—even when those accounting costs reflect information that is irrelevant to the first-best price—and to ignore valuable information that is relevant to the first-best price. These same forces can lead firms to pass through changes in average cost that do not affect marginal costs (e.g., changes in overhead costs) and to neglect changing market conditions, such as demand shocks, that affect the optimal price but cannot be externally verified. We briefly sketch these extensions.

Average cost pricing. Surveys dating back to Hall and Hitch (1939) find that managers often set prices by adding allowances for overhead expenses and profits to average unit costs. Bewley (2025) stresses that the measures of costs that firms use to set prices often include allocated overhead expenses.

Our model can be extended to explain why firms pass through changes in overhead costs. Suppose that selling units of the good at $t = 1$ requires paying a stochastic overhead cost $F = \bar{F} + \nu$ and stochastic variable cost per unit sold w. The realizations of the overhead cost F and variable cost w can be observed by the firm manager at $t = 1$ but not by the principal. Thus, the principal observes total accounting costs at $t = 1$, which are equal to $(h + w)\, D(p) + F$, but cannot disentangle the portion of costs that is due to variable costs from the portion due to fixed costs. (The presence of the stochastic variable cost w also prevents the principal from using other information, such as the change in the book value of inventory, to back out the realization of the fixed cost.) In this case, the principal penalizes the manager for deviating from a target accounting markup that reflects the principal's best guess of both the variable accounting cost and the overhead cost. Indeed, Appendix Figure A4 shows that most firm price-setters surveyed by Schoenle (2026) report passing through changes in overhead costs at the same rate as changes in variable costs.

Pass-through of demand *vs.* cost shocks. Several empirical studies suggest that firms' prices underreact to demand shocks relative to similarly sized cost shocks. Cagan (1979) reports that "Empirical studies have long found that short-run shifts in demand have small and often insignificant effects [on prices]." More recently, the asymmetric response of prices to costs relative to demand is documented both in surveys (e.g., Fabiani et al. 2006; Kohler et al. 2026) and empirically (e.g., Bils and Chang 2000). Studies exploiting local demand and cost shocks likewise find small responses of prices to large local demand shocks (e.g., Arcidiacono et al. 2020; Gagnon and Lopez-Salido 2020) but large responses to local cost shocks (Butters et al. 2022).

Our model can be extended to reflect this asymmetric response to demand *vs.* cost shocks. In the same way that the optimal contract leads the manager to partially neglect valuable information about replacement costs, so too can these incentives lead the manager to ignore valuable information about the state of demand that cannot be externally verified. Suppose, for example, that the elasticity of demand η is stochastic. Low realizations of η mean that consumer demand is unusually inelastic, leading to a higher first-best price, while high realizations of η lower the first-best price. The principal's inability to verify the state of demand leads the manager to incompletely pass through changes in the optimal markup (in the same manner that the true replacement cost c is incompletely passed through).[15]

5 Price Dynamics Under Accounting Cost Pricing

This section explores how accounting cost pricing affects price dynamics. We focus on the limiting case where firms price solely off accounting costs (i.e., $\kappa^* \to \infty$ and $\tau = 0$). Since the prior section shows that managers choose a weighted average of accounting cost- and replacement cost-based prices, price dynamics for intermediate values of κ would (to a first-order approximation) follow a weighted average of the price dynamics we derive here and the current replacement cost.

We first discuss how we model the evolution of accounting costs and inventory stocks. Then, we characterize the dynamics of prices based on accounting costs and compare the results to canonical sticky-price models.

5.1 Inventory Valuation Methods

To characterize price dynamics, we need to take a stance on how firms calculate the accounting cost of goods sold. In the two-period model, the firm purchases inventory once at price h, so the accounting cost per unit is h regardless of whether the firm uses first-in-first-out (FIFO), average acquisition cost (AAC), or last-in-first-out (LIFO) accounting.[16] In a dynamic setting, however, the firm purchases inputs repeatedly, and the valuation method determines how it apportions costs for units that are sold. Under FIFO, the cost of goods sold reflects the price of the oldest inventory in stock; under LIFO, the cost of goods sold reflects the most recent purchase price; and under AAC, the cost of goods sold depends on a blended average of prior acquisition prices. All three methods share the feature that accounting costs reflect realized historical costs rather than current replacement costs, but they differ in the weight they place on recent versus prior purchases.

We focus largely on AAC and FIFO accounting, because these are the predominant methods used by firms that are likely to price off accounting costs. Figure 5a shows that AAC and FIFO are unconditionally the most common inventory valuation methods used by public firms, together accounting for over 80 percent of inventory-carrying firms in Compustat from 2010–2025. AAC and FIFO are also especially prevalent in industries

[15] This result requires that we specify the model so that the principal cannot back the demand elasticity out by observing both the price level p and sales $pD(p)$.

[16] The average acquisition cost method is often referred to as "average cost." We use "average acquisition cost" throughout to distinguish between this inventory valuation method and the full cost pricing theory that firms set prices according to average costs (inclusive of fixed costs) rather than marginal costs.

Figure 5: Prevalence of inventory valuation methods.

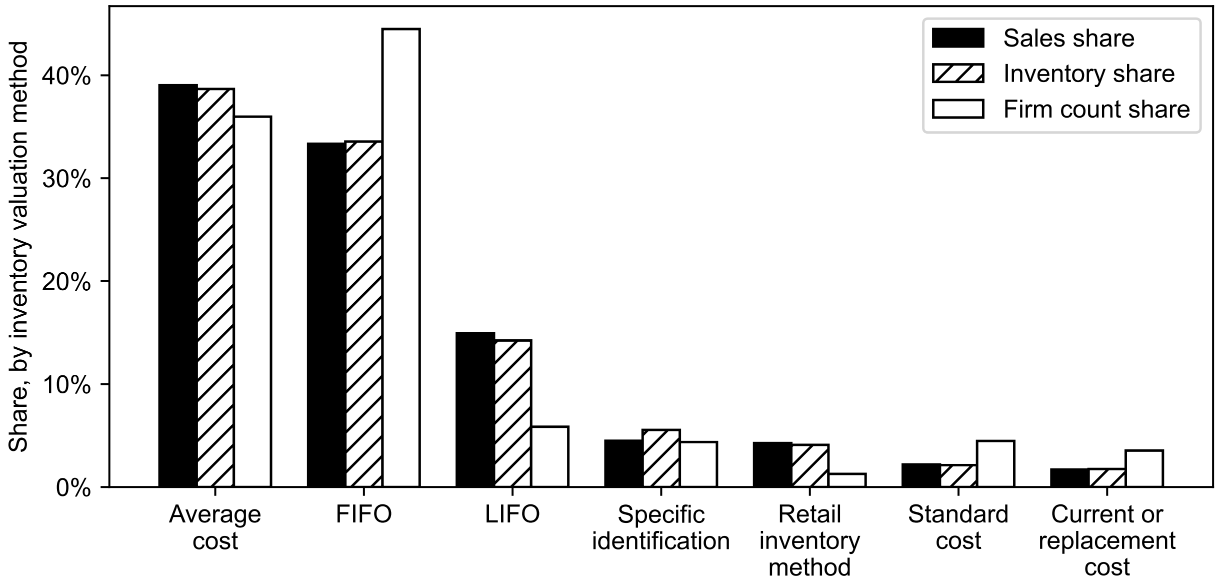

(a) Inventory valuation methods for public firms, 2010–2025.

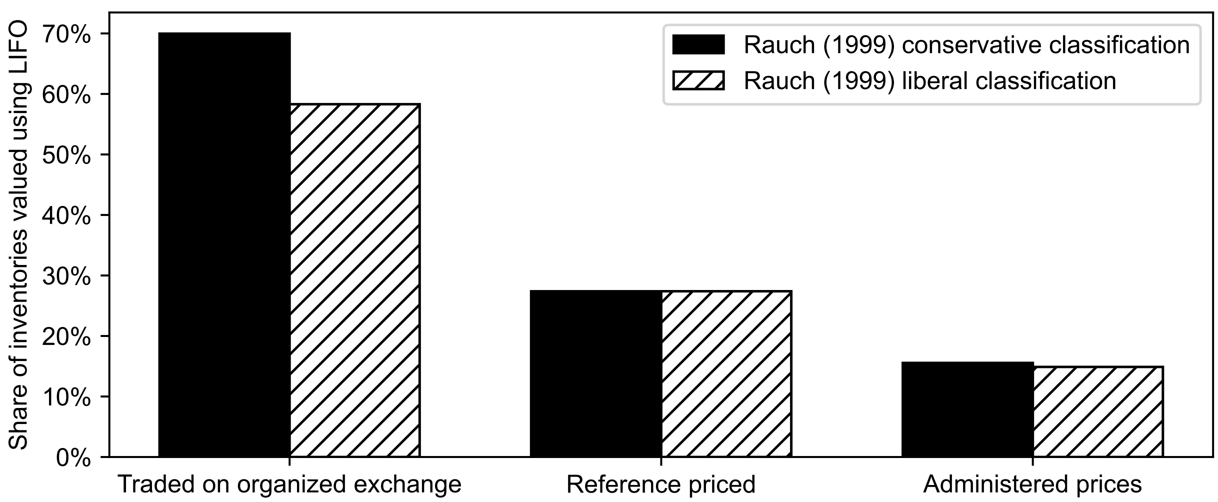

(b) Share of inventories in 2017 Economic Census valued using LIFO, by Rauch (1999) classification.

Note: The top panel shows the share of Compustat firms with each inventory valuation method (`INVVAL`) over the period 2010–2025. We exclude firms whose inventory valuation methods are recorded as "No inventory" (44.1% of observations), as "Not reported" (5.5% of observations), or are missing (<0.1% of observations). The bottom panel shows the fraction of inventories in the 2017 Economic Census priced using LIFO accounting, split by the most common Rauch (1999) categorization of products within the NAICS-6 industry code. The 2017 Economic Census covers inventory valuation methods for the mining (NAICS 21), manufacturing (NAICS 31–33), and wholesale trade (NAICS 42) sectors. We match Rauch (1999) categorizations to 258 NAICS-6 industries covering 41.6 percent of inventories in these sectors.

with administered prices, as opposed to LIFO, which is more common in industries whose outputs are primarily traded on organized exchanges or whose products are priced relative to reference prices in trade publications (see Figure 5b). Both AAC and FIFO put more weight on older acquisition costs than LIFO, which instead computes costs using the price of the most recently acquired inventory and thus more closely tracks current replacement costs.[17]

5.2 Price Dynamics Under Average Acquisition Cost Accounting

Average acquisition cost accounting yields especially tractable price dynamics because, unlike FIFO or LIFO, it does not require keeping track of inventory vintages. The law of motion for inventories is

$$I_t = I_{t-1} - D_t + Y_t,$$

where I_t is the inventory stock at the end of period t, D_t is the number of units sold at t, and Y_t are purchases of new units of inventory at t.

Under AAC accounting, the accounting value of inventories at the end of period t is calculated as

$$h_t I_t = (I_{t-1} - D_t)\, h_{t-1} + Y_t c_t. \tag{11}$$

That is, the accounting cost per unit of inventory h_t is a weighted average of the prior period's accounting cost h_{t-1} and the current replacement cost c_t. The weights are the shares of leftover inventory from the prior period and new purchases in the final inventory stock, respectively. We maintain the assumption of isoelastic demand from Section 4, so that under AAC pricing, the output price is $p_t = (\eta/(\eta-1))\, h_t$.

We log-linearize around a steady state where the cost of inventory is fixed at c^{ss}. We denote the steady-state ratio of per-period demand to inventories by $\delta = D^{ss}/I^{ss}$. We refer to δ as the *shipment–inventory ratio* to differentiate it from sales–inventory ratios reported by statistical agencies; the sales–inventory ratio measures the ratio of sales to the book value of inventory, so in steady state it is instead $\delta\eta/(\eta-1)$. We use $\hat{x}_t$ to denote the log-deviation of any variable x_t from its steady-state value, $\log x_t - \log x^{ss}$.

[17] The fact that LIFO pricing more closely approximates replacement cost pricing might make it puzzling why so few firms use LIFO, especially given that LIFO also confers a tax advantage when input costs trend upward. Part of the reason is institutional: LIFO is prohibited under International Financial Reporting Standards (see IAS 2 paragraph 25). A further reason is that LIFO can be administratively costly, requiring firms to classify inputs and determine when newly received shipments are sufficiently similar to existing inventory. Using variation in tax benefits from LIFO over time, Ling and Tinkelman (2024) estimate that the administrative and compliance costs of LIFO accounting average more than 0.35% of costs of goods sold for public firms.

Log-linearizing (11) and using $\hat{p}_t = \hat{h}_t$ yields the recursive equation

$$\hat{p}_t = (1-\delta)\hat{p}_{t-1} + \delta\hat{c}_t. \tag{12}$$

The weights on the current replacement cost $\hat{c}_t$ and the last-period price $\hat{p}_{t-1}$ depend on the rate at which inventory is depleted. When turnover is high, the selling price largely reflects the current replacement cost. When turnover is low, the average unit acquired spends a long time in inventory before being sold, and prices are more reflective of historical costs.

The recursive formulation in (12) is reminiscent of the law of motion for prices under the Calvo (1983) model of sticky prices. We formalize the relationship between the price dynamics under AAC pricing and Calvo in Proposition 4.

Proposition 4 (AAC pricing *vs.* Calvo)**.** *Given a path of costs $\{\hat{c}_t\}_{t=-\infty}^{\infty}$, the path of prices under AAC pricing and replacement cost pricing with Calvo frictions are*

$$\hat{p}_t^{AAC} = \delta \sum_{k=0}^{\infty} (1-\delta)^k \hat{c}_{t-k},$$
$$\hat{p}_t^{Calvo} = \theta \sum_{k=0}^{\infty} (1-\theta)^k \hat{c}_{t-k} + \underbrace{\theta\left[1-\beta(1-\theta)\right] \sum_{k=0}^{\infty}\sum_{s=0}^{\infty} \beta^s (1-\theta)^{k+s}\, \mathbb{E}_{t-k}\left[\hat{c}_{t+s-k} - \hat{c}_{t-k}\right]}_{\textit{Due to forward-looking reset prices}},$$

where δ is the steady-state shipment–inventory ratio, θ is the fraction of firms able to reset their prices each period under Calvo, and β is the discount factor. The paths of prices under both models coincide if $\delta = \theta$ and either $\beta = 0$ or $\mathbb{E}_t[\hat{c}_{t+s}] = \hat{c}_t$ for all periods t and forecast horizons s.

Both AAC pricing and Calvo frictions predict prices with geometrically decaying weights on historical costs, but the forces determining the speed of decay are different. Under Calvo frictions, the decay rate is determined by the share of firms $1-\theta$ that are unable to change their prices each period and are left with stale prices reflecting historical costs. Under AAC pricing, prices are flexible but depend on historical costs because the average unit spends $1/\delta$ periods in inventory before being sold.

Compared to AAC pricing, prices under Calvo frictions have an extra term that reflects how firms expect future replacement costs to differ from current costs. This forward-looking term drives a wedge between the Calvo and AAC price paths unless firms do not care about the future (i.e., $\beta = 0$) or firms' expectations of future costs are equal to current costs. The latter condition is satisfied if costs follow a martingale or if firms have what Minton and Wheaton (2022) term myopic expectations. Minton and Wheaton (2022)

Figure 6: Price response to immediate *vs.* slow-moving cost shocks.

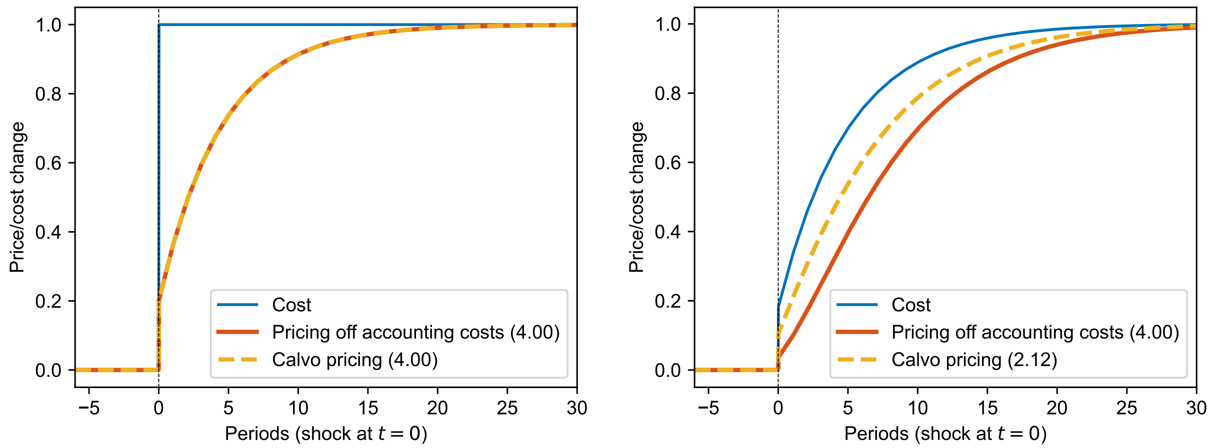

(a) Immediate shock. **(b)** Slow-moving shock.

Note: The Calvo model uses $\beta = (0.96)^{1/12}$ and $\theta = 0.2$, and the pricing off accounting costs model uses weighted average acquisition costs with $\delta = 0.2$ (i.e., an inventory–shipments ratio of five). The left panel shows the response to an unanticipated, permanent cost shock ($\rho = 0$) while the right panel shows the response to an unanticipated, slow-moving shock ($\rho = e^{-0.2}$). The numbers in parentheses reflect the sluggish response of prices to costs, $\sum_{t=0}^{\infty}(\hat{c}_t - \hat{p}_t)$.

estimate that "complete myopia" is necessary to match the dynamics of pass-through for most commodities through the input–output network.

We use two illustrative examples to highlight how AAC pricing leads to different predictions than Calvo frictions.

Example 1 (Sluggish response to slow-moving shocks)**.** Calvo prices respond more quickly to a cost shock that unfolds over time, because firms increase their prices in anticipation of future cost changes, unlike firms with AAC pricing. To illustrate this, suppose there is an unanticipated shock to costs at time $t = 0$, with $\hat{c}_t = 1 - \rho^{t+1}$ for $t \geq 0$. The parameter ρ describes how slow-moving the shock is: when $\rho = 0$, costs jump to $\hat{c}_0 = 1$ immediately, while when $\rho > 0$ the cost shock gradually increases to one over time. Applying Proposition 4 yields the following price paths for $t \geq 0$,

$$\hat{p}_t^{\text{AAC}} = 1 - (1-\delta)^{t+1} - \delta\rho\left[\frac{(1-\delta)^{t+1} - \rho^{t+1}}{(1-\delta) - \rho}\right],$$

$$\hat{p}_t^{\text{Calvo}} = 1 - (1-\theta)^{t+1} - \theta\rho\left[\frac{(1-\theta)^{t+1} - \rho^{t+1}}{(1-\theta) - \rho}\right]\left[1 - \frac{(1-\rho)\beta(1-\theta)}{1-\rho\beta(1-\theta)}\right],$$

When $\delta = \theta$, $\hat{p}_t^{\text{Calvo}} \geq \hat{p}_t^{\text{AAC}}$ for all $t \geq 0$, with the inequality holding strictly when $\rho > 0$

and $\beta > 0$. In words, firms raise their prices more quickly under Calvo frictions, because firms anticipate that future costs are higher than present costs.

Figure 6 plots these price responses. The left panel shows that the responses of prices under Calvo frictions and pricing off accounting costs coincide for an immediate, permanent shock. However, the coincidence fails when anticipated future costs differ from current costs, as in the case of the slow-moving shock shown in the right panel. In that case, prices respond more sluggishly when firms price off accounting costs as higher-priced inventory slowly makes its way through firms' inventory stocks. We calculate the cumulative gap between costs and prices over the path of the shock (embedded in a simple general equilibrium model, this gap would be proportional to the cumulative impulse response of the cost shock on output). In the right panel of Figure 6, AAC pricing increases this statistic by over 80 percent relative to Calvo pricing.

Example 2 (Greater responsiveness to transitory shocks)**.** The two preceding examples show that prices respond more sluggishly to costs under accounting cost pricing than Calvo. This is because in both examples, current cost increases are correlated with higher future costs. If the opposite is true—for example, in the case of transitory shocks—then accounting cost pricing can instead generate *more* responsiveness than Calvo.

The left panel of Figure 7 shows how prices respond to an unanticipated, transitory cost shock that arrives at $t = 0$. Prices under both Calvo frictions and AAC accounting costs rise on impact, but the response is stronger for AAC pricing, because the firms with Calvo frictions incorporate the fact that they expect costs to revert. In fact, in this example, the price response under AAC pricing is always larger than the pricing response under Calvo, though both responses peak below the size of the initial cost shock.

For an anticipated cost shock, such as the one shown in the right panel of Figure 7, the picture is more complicated. Firms with Calvo pricing are initially more responsive before the shock materializes, because they increase their prices in anticipation of the expected cost increase. But once the shock materializes, the prices under AAC pricing jump and ultimately reach a higher peak response. In fact, while price responses for unanticipated versus anticipated cost shocks differ under Calvo frictions, the price response under AAC pricing is the same whether or not firms anticipate the cost increase; they mechanically increase their price in accordance with the average cost of inventory.

5.3 Price Dynamics Under FIFO and LIFO

Unlike AAC pricing, both FIFO and LIFO require keeping track of which inventory was purchased when, since costs of goods sold are calculated by tallying up unit costs for

Figure 7: Increased responsiveness to transitory shocks.

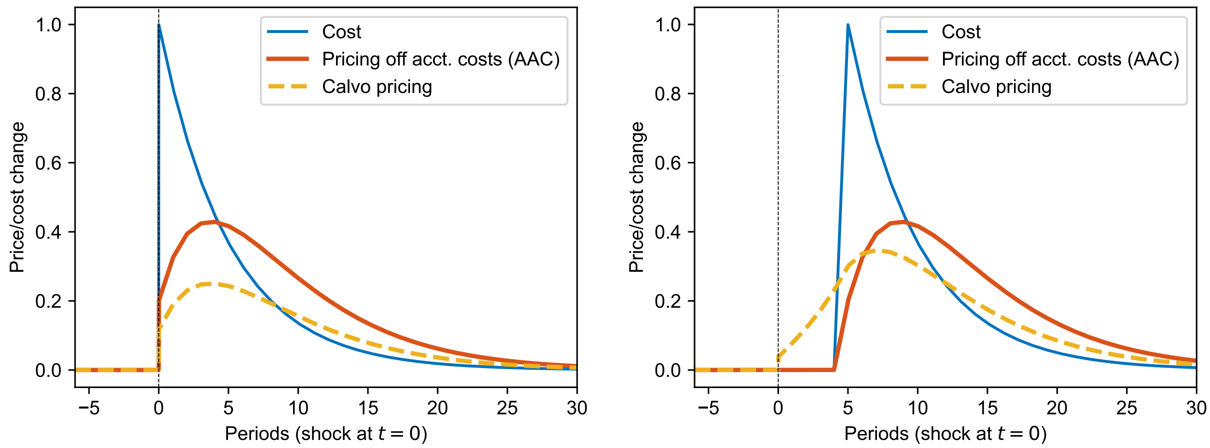

(a) Unanticipated cost increase at $t = 0$. **(b)** News at $t = 0$ about cost increase at $t = 5$.

Note: The Calvo model uses $\beta = (0.96)^{1/12}$ and $\theta = 0.2$, and the pricing off accounting costs model uses weighted average acquisition costs with $\delta = 0.2$ (i.e., an inventory–shipments ratio of five). In the left panel, the path of costs is $\hat{c}_t = e^{-0.2t}$ for $t \geq 0$. In the right panel, the path of costs is $\hat{c}_t = e^{-0.2(t-5)}$ for $t \geq 5$.

items in the inventory stock that were either purchased first or most recently. However, if the firm always has enough inventory on hand to satisfy current demand (in steady state, this corresponds to $\delta < 1$) and if $1/\delta$ is an integer, then

$$\hat{p}_t^{\text{FIFO}} = \hat{h}_t^{\text{FIFO}} = \hat{c}_{t-1/\delta}$$
$$\hat{p}_t^{\text{LIFO}} = \hat{h}_t^{\text{LIFO}} = \hat{c}_t.$$

That is, prices and accounting costs under FIFO depend on the replacement cost lagged by $1/\delta$ periods, and prices and accounting costs under LIFO are given by current replacement costs. In this respect, FIFO pricing introduces a lag exactly equal to the time that units spend in inventory before being sold, and LIFO pricing replicates the path of prices without nominal rigidities.

The fact that pricing under LIFO replicates replacement cost pricing is a byproduct of the assumption that the firm buys new inventory in each period. Okun (1981) notes that, in reality, "[firms] are not in the market for all their inputs at any time" (Okun, 1981, p. 156). One reduced-form way of modeling this idea is to consider an industry that consists of a set of producers, each of whom buys inputs once every $1/\delta$ periods and slowly exhausts those inputs over $1/\delta$ periods.[18] Since firms are only in the market for their inputs once

[18]Our use of δ here is a slight abuse of notation, since the ratio of shipments to the average inventory stock in this setting is instead 2δ.

Figure 8: Response to immediate *vs.* slow-moving cost shocks: FIFO, LIFO, and Taylor.

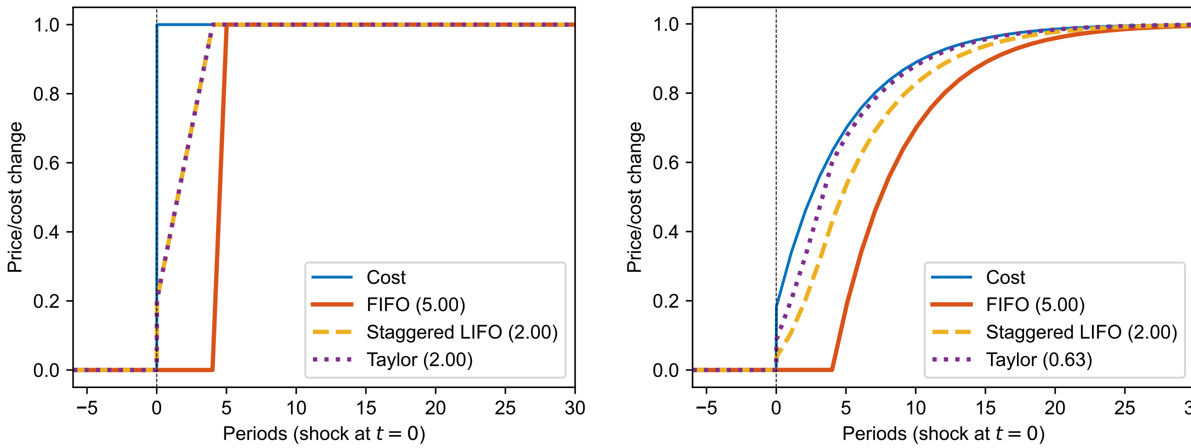

(a) Immediate shock.

(b) Slow-moving shock.

Note: The Taylor model uses $\beta = (0.96)^{1/12}$ and $N = 5$, and all models with pricing off accounting costs use $\delta = 0.2$. The cost shock sequences are the same as in Figure 6. In both panels, the cost shock is unanticipated before $t = 0$. The numbers in parentheses reflect the sluggish response of prices to costs, $\sum_{t=0}^{\infty}(\hat{c}_t - \hat{p}_t)$.

every $1/\delta$ periods, the average price across producers in this "staggered" setting reflects historical costs from the last purchase date. Just as Proposition 4 drew a parallel between AAC pricing and Calvo frictions, Proposition 5 shows a connection between price paths under LIFO pricing with intermittent purchases and Taylor (1980) pricing frictions.

Proposition 5 (LIFO pricing with intermittent purchases *vs.* Taylor)**.** *Given a path of costs* $\{\hat{c}_t\}_{t=-\infty}^{\infty}$*, the paths of average prices under LIFO pricing when producers buy inputs every* $1/\delta$ *periods and under replacement cost pricing with Taylor frictions are*

$$\hat{p}_t^{\textit{Staggered LIFO}} = \delta \sum_{k=0}^{1/\delta - 1} \hat{c}_{t-k},$$

$$\hat{p}_t^{\textit{Taylor}} = \frac{1}{N}\sum_{k=0}^{N-1} \hat{c}_{t-k} + \underbrace{\frac{1}{N}\frac{1-\beta}{1-\beta^N}\sum_{k=0}^{N-1}\sum_{s=0}^{N-1}\beta^s \mathbb{E}_{t-k}\left[\hat{c}_{t+s-k} - \hat{c}_{t-k}\right]}_{\textit{Due to forward-looking reset prices}},$$

where firms under Taylor pricing reset prices every N *periods and* β *is the discount factor. The paths of prices under both models coincide if* $\delta = 1/N$ *and either* $\beta = 0$ *or* $\mathbb{E}_t[\hat{c}_{t+s}] = \hat{c}_t$ *for all periods* t *and forecast horizons* s*.*

Figure 8 illustrates Proposition 5 in a quantitative example. For an immediate, permanent shock, the paths of average prices under LIFO with intermittent purchases and

Taylor pricing coincide, because current costs exactly equal future costs. This is not true when the cost shock unfolds slowly over time: in that case, Taylor frictions predict a faster response of prices, because firms' reset prices incorporate the anticipated higher future costs. Compared to these models, price adjustment is always more delayed under FIFO—the price path lags the cost shock by $1/\delta$ periods—and prices under LIFO pricing with continuous purchases (not shown) coincide with the path of replacement costs.

5.4 Comparing Pricing Off Accounting Costs to Sticky-Price Models

We highlight three implications of how price dynamics under accounting cost pricing differ from those under sticky-price models, even when the models are calibrated to the same parameters. First, we show that pricing off accounting costs can amplify delays in the pass-through of cost shocks along supply chains. Second, we show that anticipation of cost increases can further delay pass-through (rather than accelerate pass-through, as it would in sticky-price models) if firms stock up on inventories. Finally, we show that the response of accounting markups to cost shocks can be misleading about the response of true markups to shocks.

Amplified delays in supply chains. Both sticky prices and pricing off accounting costs translate immediate cost shocks into gradual price changes. In supply chains, this means that one-time, permanent cost shocks affecting an upstream producer turn into slow-moving cost shocks for their customers, and even slower-moving cost shocks for producers further downstream. As we saw in Example 1, firms pricing off accounting costs exhibit more delayed pass-through of slow-moving shocks. Thus, the combination of pricing off accounting costs and supply chains can further amplify delays in price adjustment.

Figure 9 illustrates the response of prices in a three-firm vertical supply chain to an immediate, permanent cost increase in the most upstream firm's costs. Whether firms price off accounting costs or with Calvo frictions, the first firm in the chain has the same response to the cost shock (assuming $\delta = \theta$). However, the slow-moving cost shock faced by the middle firm in the chain leads to a slower path of price increases when firms price off accounting costs. The same pattern repeats further downstream. The result is that, even though both supply chains face the same cost shock, the cumulative gap between the upstream cost and downstream prices over the path of the shock is more than 50 percent larger under accounting cost pricing.

Figure 9: Amplification of sluggish price response in supply chains.

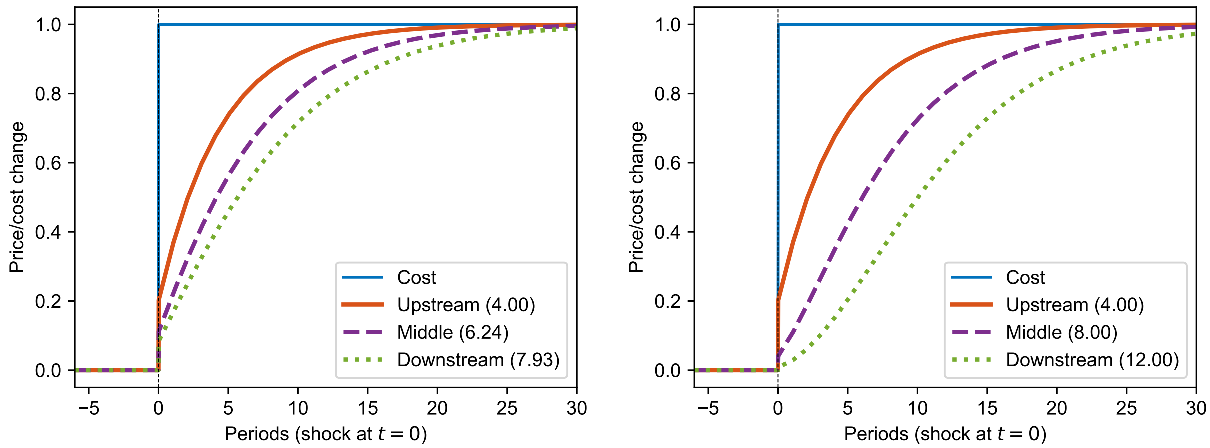

(a) Calvo. **(b)** Pricing off accounting costs (AAC).

Note: The Calvo model uses $\beta = (0.96)^{1/12}$ and $\theta = 0.2$, and the pricing off accounting costs model uses weighted average acquisition costs with $\delta = 0.2$ (i.e., an inventory–shipments ratio of five). In both panels, the shock to upstream costs is unanticipated before $t = 0$. The numbers in parentheses reflect the sluggish response of prices to costs, $\sum_{t=0}^{\infty}(\hat{c}_t - \hat{p}_t)$.

Anticipation and pass-through. Firms may respond to anticipated cost changes by stocking up on inventory when costs are expected to rise (often referred to as "forward buying") or slowing purchases when costs are expected to fall. Carreras-Valle and Lee (2026) document forward buying behavior in advance of the 2018–2019 tariffs on steel, aluminum, and Chinese imports, and news coverage around changing tariff rates in 2025–2026 suggests that firms engaged in substantial forward buying when they expected tariffs to increase.[19]

The AAC pricing equation (12) shows that firm decisions to stock up on inventories or adjust purchase decisions do not affect prices to a first-order approximation. But a second-order approximation of the average acquisition cost yields

$$\hat{h}_t = \underbrace{(1-\delta)\,\hat{h}_{t-1} + \delta\hat{c}_t}_{\text{As in eq. (12)}} + \underbrace{\frac{1}{2}\delta\,(1-\delta)\left(\hat{c}_t - \hat{h}_{t-1}\right)^2}_{\text{Second-order correction}} + \underbrace{\delta\left(\hat{c}_t - \hat{h}_{t-1}\right)\left(\hat{Y}_t - \hat{I}_t\right)}_{\text{Effect of inventories / purchases}} \,. \tag{13}$$

[19]For example, PIIE notes a sharp increase in U.S. imports from many trading partners in early 2025, which they attribute to anticipatory purchases (`https://www.piie.com/blogs/realtime-economics/2025/despite-tariffs-us-merchandise-imports-increased-and-exports-held`). In the aftermath of the Supreme Court's ruling on the illegal tariffs, The Wall Street Journal likewise reported, "The Race to Seize on Lower Tariffs Has Businesses Rushing In Shipments" (`https://www.wsj.com/economy/trade/tariffs-are-lower-and-businesses-are-racing-to-take-advantage-977bd276`).

Figure 10: Forward buying of inventories delays pass-through of cost increases.

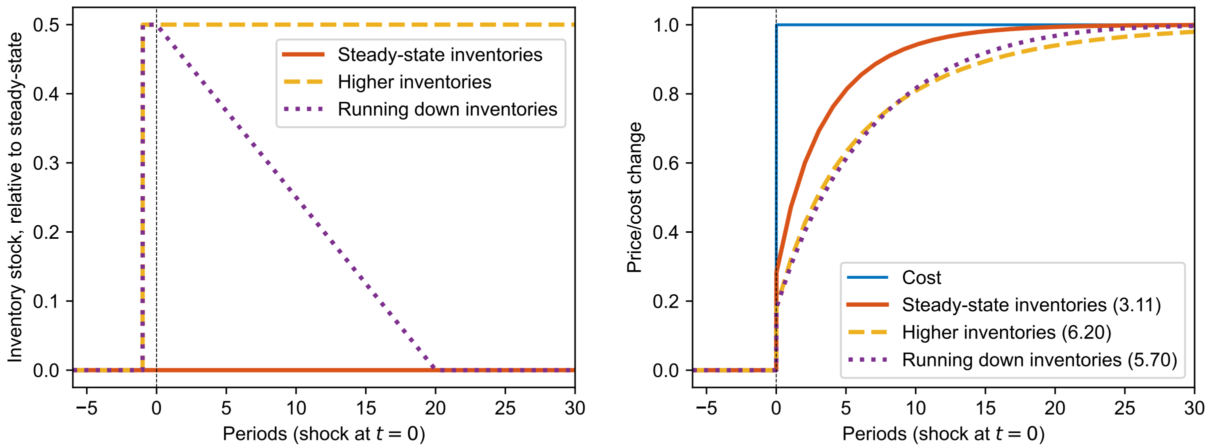

(a) Inventory paths. **(b)** Pricing off accounting costs (AAC).

Note: The pricing off accounting costs model uses average acquisition costs with $\delta = 0.2$ (i.e., an inventory–shipments ratio of five). The left panel shows three paths for inventories relative to steady state. The law of motion for purchases is $\hat{Y}_t = (1/\delta)(\hat{I}_t - \hat{I}_{t-1}) + \hat{D}_t$, and we assume that $\hat{D}_t = 0$ for all periods so that the inventory paths come from changes in purchases relative to steady state. The right panel shows the resulting path of prices. The numbers in parentheses reflect the sluggish response of prices to costs, $\sum_{t=0}^{\infty}(\hat{c}_t - \hat{p}_t)$.

The final term reflects the effect of inventories and purchases on accounting costs and thus on prices. If a firm stocks up on inventories before costs rise, then when the cost increase arrives, the combination of $\hat{I}_t > 0$ and $\hat{h}_{t-1} < \hat{c}_t$ depresses the current accounting cost $\hat{h}_t$ and the firm's price. (Reducing purchases $\hat{Y}_t < 0$ and depleting existing inventories has a similar effect.) This is precisely the mechanism that many of the quotes in Table 1 and Appendix Table A1 suggest played out during the onset of the tariffs in 2025.

We illustrate the effects of forward buying on the timing of price changes using this second-order approximation in Figure 10. In the figure, costs increase permanently starting at $t = 0$. We consider three scenarios for how inventories evolve in response to the shock. The first scenario assumes inventories remain fixed at the steady-state level. The second scenario assumes the firm stocks up on inventory at $t = -1$, increasing its inventory stock by 50 percent, and keeps the inventory stock permanently higher after the shock. The third scenario again assumes that the firm stocks up on inventory at $t = -1$, but that it gradually runs down the inventory by cutting back on purchases until the inventory stock returns to its steady-state level. We model each of these scenarios, shown graphically in the left panel of Figure 10, by assuming a path of purchases $\hat{Y}_t$ that yields the desired path of inventories.

The right panel of Figure 10 shows that forward buying substantially delays the pass-through of the cost shock to prices. In the scenario where the firm permanently increases its inventory stock, the delayed pass-through occurs because the firm holds more inventory purchased at pre-shock prices, and thus the weighted average acquisition cost converges more slowly to the new cost. In the scenario where the firm builds up and then gradually runs down its inventory stock, the average acquisition cost of inventory is slower to rise because the firm reduces purchases of new inventory immediately after the shock. Despite the different mechanisms, in both scenarios with forward buying, the cumulative gap between costs and prices over the path of the shock is about twice as large as the gap under the first scenario when the inventory stock is held fixed.

This result highlights a contrast between pricing off accounting costs and sticky prices. In models with sticky prices, anticipation speeds up the pass-through of persistent cost shocks, because firms account for future cost changes when they have an opportunity to reset prices. Minton and Wheaton (2022) find that firm "myopia"—i.e., inattention to predictable future cost changes—is needed to have these models match the empirical responses of prices to many commodity cost shocks. Under pricing off accounting costs, however, anticipation can in fact delay the pass-through of cost increases. Forward-looking firms that expect costs to rise stock up on inventory, and this forward buying slows down the convergence of accounting costs to current replacement costs. In other words, under accounting cost pricing, forward-looking behavior can reinforce delayed price adjustment rather than erasing it.

Response of accounting *vs.* true markups to shocks. Nekarda and Ramey (2020) summarize a long-standing debate over whether the response of markups to shocks conforms with the predictions of the New Keynesian model. Recently, Anderson et al. (2025) and Burstein et al. (2025) revisit these questions using firm-level data on markups constructed using firms' reported costs.

Backward-looking accounting measures of costs can make accounting markups a poor proxy for true markups over replacement costs, which are the relevant object for the New Keynesian model. Figure 11 illustrates this possibility with two cases. In the first case, the firm sets its price as a fixed markup over accounting costs. In this case, the accounting markup appears unresponsive to cost shocks *by construction*, even as the cost increase causes the markup over replacement costs to fall. This decline in "true" markups is disguised by the lag between changes in replacement costs and accounting costs.

The problem is exacerbated if the firm's price reflects a combination of both accounting costs and replacement costs. The right panel of Figure 11 shows that a cost increase then leads accounting markups to slightly rise, even though the true markup over replace-

Figure 11: Response of measured vs. true markups to a cost shock.

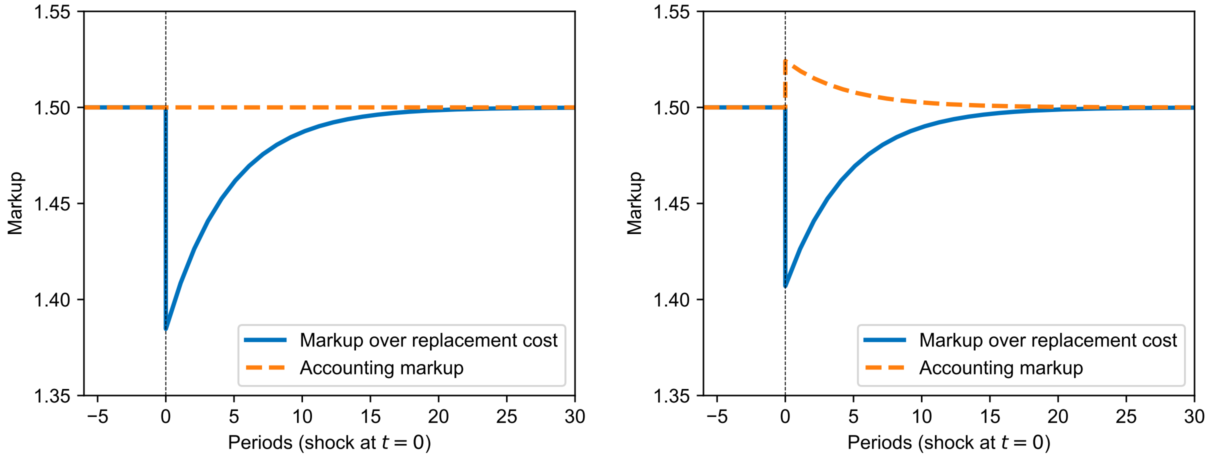

(a) Full accounting cost pricing ($\phi = 0$). **(b)** Partial accounting cost pricing ($\phi = 0.2$).

Note: The figure plots the accounting markup p_t/h_t and the markup over replacement costs p_t/c_t following a permanent increase in replacement costs by 10 percent. The markup is initially 1.5, and we assume accounting costs are given by average acquisition cost with a shipment–inventory ratio of $\delta = 0.2$. Prices follow $\hat{p}_t = \phi\hat{c}_t + (1 - \phi)\hat{h}_t$, with $\phi = 0$ in the left panel and $\phi = 0.2$ in the right panel.

ment costs falls in response to the cost shock. In other words, if costs increase during a boom, accounting markups may appear acyclical or procyclical even if true markups over replacement costs are countercyclical.

6 Macroeconomic Rigidity

In this section, we evaluate the extent to which accounting cost pricing can generate delayed responses to shocks in an input–output model calibrated to the U.S. economy. We first extend our results on price dynamics from the previous sections to a model with general input–output linkages and calibrate the model using detailed information on the inventory turnover rates of disaggregated industries. Then, we simulate delays in the pass-through of industry cost shocks to consumer prices. Finally, we simulate the 2025 tariff shocks in our model and compare the model-implied price dynamics to the data.

6.1 Setup

Following Baqaee and Farhi (2020) and Rubbo (2023), we consider an economy consisting of N industries and a single factor (value added). A continuum of identical firms in each industry produce varieties of an industry output by combining value added with intermediate inputs from other industries in a constant returns, neoclassical production function. Each industry's output is a CES aggregate over individual firm varieties. Household consumption is a CES bundle of industry goods. We treat the price of value added as exogenous and log-linearize the model around a zero inflation, perfect foresight steady state. We denote the steady-state input–output matrix by Ω, where Ω_{ij} is industry i's expenditures on inputs from industry j as a share of total expenditures.

To characterize how cost shocks propagate through this economy, we define a pricing rule, $\mathcal{P}_i : \{\hat{mc}_{it+\tau}\}_{\tau=-\infty}^{\infty} \mapsto \hat{p}_{it}$, as a map from the path of an industry's marginal costs to its prices (both log-linearized around their steady-state levels). For example, with flexible and competitive pricing, the pricing rule is simply $\hat{p}_{it} = \hat{mc}_{it}$. Each industry's marginal costs in turn depend on the prices of their suppliers. So, we can define an equilibrium as a path of prices that solves the fixed point

$$\hat{p}_{it} = \mathcal{P}_i\left(\left\{\hat{mc}_i\left(\hat{\boldsymbol{p}}_{t+\tau}\right)\right\}_{\tau=-\infty}^{\infty}\right), \qquad \text{for all } i \text{ and } t,$$

where $\hat{\boldsymbol{p}}_t = (\hat{p}_{1t}, ..., \hat{p}_{Nt})$ is the vector of all industry prices at time t and $\hat{mc}_i(\cdot)$ is industry i's marginal cost as a function of contemporaneous industry prices.

Proposition 6 provides equilibrium conditions in the network economy under each pricing rule. These equilibrium conditions also provide a natural way to solve for the path of prices in a calibrated model.

Proposition 6 (Network price dynamics). *Let $\{\hat{\boldsymbol{c}}_t\}_{t=0}^{\infty}$ be a path of exogenous cost shocks, where $\hat{\boldsymbol{c}}_t = (\hat{c}_{1t}, \ldots, \hat{c}_{Nt})$ is the vector of cost shocks to each industry at time t. The following expressions characterize the dynamics of industry prices $\hat{\boldsymbol{p}}_t = (\hat{p}_{1t}, \ldots, \hat{p}_{Nt})$ under different pricing rules.*

1. ***Average acquisition cost (AAC) pricing:*** *Industry prices follow*

$$\hat{\boldsymbol{p}}_t = (I - \Delta\Omega)^{-1}\left[(I-\Delta)\hat{\boldsymbol{p}}_{t-1} + \Delta\hat{\boldsymbol{c}}_t\right],$$

where $\Delta = \operatorname{diag}(\delta_1, \ldots, \delta_N)$ is a diagonal matrix of industry shipment–inventory ratios.

2. ***FIFO pricing:*** *The price of each industry i follows*

$$\hat{p}_{it} = \sum_j \Omega_{ij}\hat{p}_{j,t-d_i} + \hat{c}_{i,t-d_i},$$

where $d_i = \lceil \delta_i^{-1} \rceil$ is the inventory turnover time for industry i.

3. ***Calvo pricing:*** *Industry prices follow the second-order difference equation*

$$\hat{\boldsymbol{p}}_t = \left(I - K\Omega + \beta\,(I-\Theta)^2\right)^{-1}\left[(I-\Theta)\,\hat{\boldsymbol{p}}_{t-1} + \beta\,(I-\Theta)\,\hat{\boldsymbol{p}}_{t+1} + K\hat{\boldsymbol{c}}_t\right],$$

where β is the discount factor, $\Theta = \mathrm{diag}\,(\theta_1, \ldots, \theta_N)$ is a diagonal matrix of industry Calvo adjustment probabilities, and $K = \Theta(I - \beta(I-\Theta))$.

Given the same path of cost shocks $\{\hat{\boldsymbol{c}}_t\}_{t=0}^{\infty}$, the paths of prices for each industry under AAC pricing and Calvo frictions coincide if $\delta_i = \theta_i$ for all i and $\beta = 0$.

Under AAC pricing, the matrix $(I - \Delta\Omega)^{-1}$ acts as a modified Leontief multiplier: it captures how cost increases propagate through the input–output network, dampened by the fact that downstream firms gradually fold new input prices into their average costs. Under FIFO pricing, a cost shock to an upstream sector cascades downstream at a speed determined by the cumulative inventory holding periods along the supply chain. The longer the chain of production, the longer the aggregate price level takes to absorb the shock. Assuming that firms set prices at the beginning of the period and keep them fixed within the period, the FIFO inventory delays are given by $d_i = \lceil \delta_i^{-1} \rceil$. Finally, under Calvo frictions, because reset prices today depend on expectations of prices tomorrow, the equilibrium condition is on the entire path of prices.

From Proposition 6, it is easy to see that when $\beta = 0$ and $\Delta = \Theta$, the AAC and Calvo price paths coincide. When $\beta = 0$, the relative speeds of pass-through under AAC and Calvo pricing therefore depend on a comparison of Δ and Θ. If inventory delays are smaller than the average length of price spells, then prices adjust faster under AAC pricing, and vice versa. Relative to AAC pricing, FIFO pricing with delays $d_i = \lceil \delta_i^{-1} \rceil$ tends to result in slower price adjustment (though not always). Under AAC, prices react immediately but incompletely to cost shocks, since cost shocks affect the blended average cost of inventories; under FIFO, prices only adjust once the full inventory delay period has passed.

6.2 Calibrating the Model with Accounting Cost Pricing

Table 5 describes the data sources we use to calibrate the model. We create a trade-adjusted input–output table, building on the 405-industry input–output table from the Bureau of Economic Analysis (BEA), use shipment–inventory ratios for each industry from the Economic Census, and separate industries engaging in accounting cost pricing and replacement cost pricing using the Rauch (1999) product categorization.

Table 5: Data sources for calibration.

Parameter	Description	Source
Ω	Input–output matrix	BEA 2017 405-industry input–output table
δ_i	Shipment–inventory ratios	2017 Economic Census inventories, COGS
d_i	Inventory turnover time	$\lceil \delta_i^{-1} \rceil$
ϕ_i	Weight on replacement costs	Rauch (1999): $\phi_i = 1$ if in $\{w, r\}$, $\phi_i = 0$ if n
β	Discount factor	$(0.96)^{1/12}$
θ_i	Calvo frictions	Pasten et al. (2020) industry-level estimates

Building a trade-adjusted input–output table. In our framework, the fact that industry goods spend time sitting in the inventories of wholesale distributors or retailers before being delivered to customers can contribute to delays in pass-through. A challenge with input–output tables produced by the BEA is that they treat wholesale and retail trade as margin industries. This means that the BEA input–output table reflects only inputs used to provide value-added services by the wholesale and retail trade sectors and does not record merchandise sold through wholesale and retail trade as inputs to those industries (sales through the wholesale and retail trade sector are logged directly as sales of the producing industry to the purchasing industry).

We thus reroute the input–output network to treat industries in wholesale and retail trade like other productive industries. In our trade-adjusted input–output table, trade sectors purchase goods from producing industries, add a wholesale/retail margin using inputs, and resell to consumers at a new price. We model specialized wholesale and retail trade sectors for each producing industry with nonzero trade margins, and so our final table contains 843 industries inclusive of these trade industries.

Appendix B provides a detailed account of the steps we take to construct the trade-adjusted input–output table. Broadly, we take three steps, reversing the steps that the BEA uses to construct trade margins in the input–output accounts (Horowitz and Planting 2009). First, for each industry with non-zero retail and wholesale selling margins, we hand-construct a mapping from the industry to wholesale and retail product lines in the 2017 Economic Census. We then use the gross margin of these product lines from the Economic Census to estimate the share of purchases from an industry that are made through wholesale and retail channels versus directly from industry suppliers. Finally, we add these wholesale and retail sellers to the input–output table and reroute the flow of goods through them.

Shipment–inventory ratios. We calculate shipment–inventory ratios as the ratio of cost of goods sold to end-of-year inventories for each NAICS-6 industry from the 2017 Economic Census. The Economic Census reports inventories for industries in mining (NAICS sector 21), manufacturing (sector 31–33), wholesale trade (sector 42), and retail trade (sector 44–45).

Appendix Figure A5 shows a histogram of inventory–shipment ratios across NAICS-6 industries. The average NAICS-6 industry holds just under two months of shipments in inventories; the average manufacturing industry holds 2.5 months of inventory. There is a tail of industries with substantially higher inventory stocks, such as wineries and distilleries (which hold 21 and 10 months of shipments in inventory), biological product manufacturing (which holds 12 months of shipments in inventory), and aircraft manufacturing and aerospace products and parts manufacturing (which hold 11 and 9 months of shipments in inventory).[20]

Industries with replacement cost pricing. Section 3 finds that firms in industries with organized exchanges or reference prices tend to engage in replacement cost pricing, while firms with administered prices tend to engage in accounting cost pricing. We reflect this in the calibration by modeling the prices of industries whose primary products are categorized as w (organized exchange) or r (reference prices) by Rauch (1999) as having flexible prices. Only industries whose primary products are categorized as n, i.e., having administered prices, are subject to delays in price adjustment from inventory turnover.

Other parameters. We calibrate a Calvo model with sticky prices using industry-level frequencies of price adjustment measured in PPI microdata by Pasten et al. (2020). As in standard input–output calibrations, we assume that price stickiness for the wholesale and retail trade sectors only applies to trade margins. For the Calvo model, we assume that firms use a discount factor of $\beta = (0.96)^{1/12}$, corresponding to an annual interest rate of 4 percent.

6.3 Transmission of Industry Cost Shocks to Consumer Prices

We first explore how long the model predicts industry cost shocks take to transmit to consumer prices. For an industry i, we consider a one-time, permanent shock to i's

[20] Appendix Figure A6 shows that our measured inventory–shipment ratios have a correlation of $\rho = 0.68$ with those measured by Antràs and Tubdenov (2025), who calculate these ratios using Compustat data. The benefit of the Economic Census data is that they also reflect inventory holdings of private firms and capture a broader set of industries (501 NAICS-6 industries, compared to 269 in Antràs and Tubdenov 2025).

Figure 12: Delayed effects of industry cost shocks on consumer prices in AAC calibration.

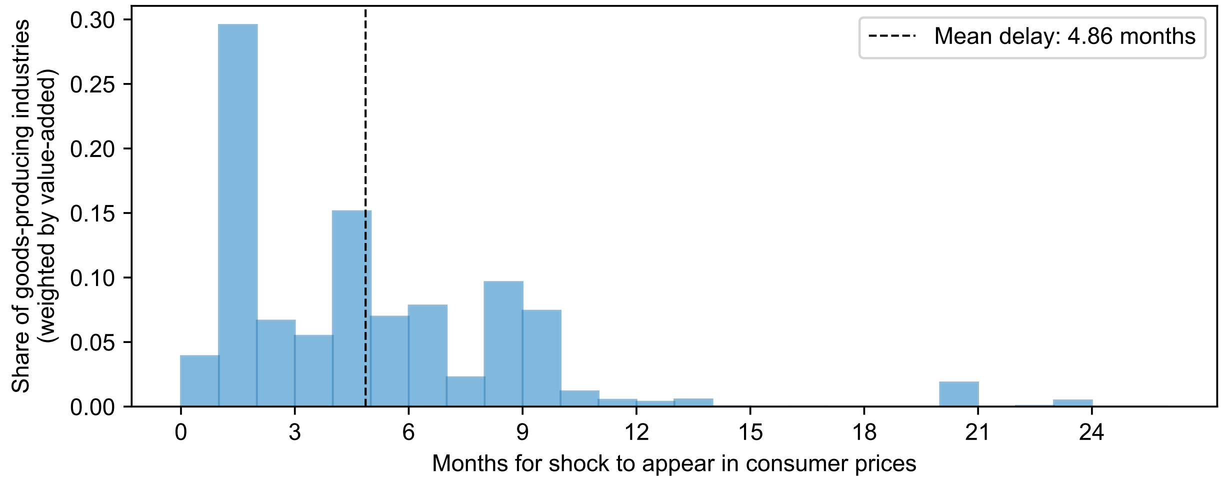

Note: For each goods-producing industry, we simulate a one-time, permanent shock to production costs and calculate the cumulative pass-through delay, $\sum_{t=0}^{\infty}(1 - \hat{p}_{\text{PCE},t}/\hat{p}_{\text{PCE},\infty})$. The figure shows a histogram of these delays across industries weighted by each industry's value added in the calibration with average acquisition cost (AAC) accounting.

production costs. We then calculate the delay in the pass-through of the cost shock to consumer prices as $\sum_{t=0}^{\infty}(1 - \hat{p}_{\text{PCE},t}/\hat{p}_{\text{PCE},\infty})$, where $\hat{p}_{\text{PCE},t}$ is the price of the personal consumption expenditures basket at time t, measured as a log-deviation from steady state. This statistic captures the average lag between the shock and the response of consumer prices, weighting each period by its contribution to eventual pass-through. For example, if pass-through is zero before period d and complete from period d onward, this statistic is equal to d.

Figure 12 shows that, in the model with AAC pricing, cost shocks to goods-producing industries take about 5 months on average to be passed through to consumer prices. The calibration with FIFO pricing yields longer lags in shock transmission: Appendix Figure A7 shows that cost shocks to goods-producing industries take about 6 months on average to be passed through to consumer prices.

These lags are longer than industries' individual inventory–shipment ratios because, in addition to taking time to be incorporated into an industry's output, these lags include the time that it takes for an industry's output to flow through the production of downstream industries and through the inventories of trade sectors before being purchased by consumers. This simulation captures an idea frequently invoked in descriptions of pricing in the popular press that production lags and inventories delay the pass-through of cost changes to downstream prices. See, for example, coverage of coffee bean prices in 2025:

> [I]ndustry experts estimate it takes at least nine months for raw bean prices to filter through to coffee drinkers [....] because U.S. roasters typically hold about two to three months' worth of bean stocks on average and need another two to three months to roast and package their products.[21]

Or coverage of how consumer chocolate prices reflect the prices manufacturers paid for cocoa:

> Cocoa—the key ingredient in chocolate—has been steadily falling in price from its all-time high in 2024. [...] But that drop hasn't trickled down to what consumers are paying. [...] The reason: The Easter chocolate on shelves today was produced using cocoa purchased at its extreme high.[22]

In our model, accounting cost pricing means that these production lags and inventory turnover rates indeed affect the speed of pass-through as the news coverage suggests. The result is substantial delays between industry cost shocks and the adjustment of consumer prices, even without sticky prices per se.

6.4 Pass-Through of 2025–2026 Tariffs

We use the calibrated model to simulate the effect of the 2025–2026 tariffs on consumer prices. We compare the pass-through of tariffs to consumer prices predicted by the model to a model with Calvo price adjustment frictions and to the data.

Additional data and assumptions. Our simulation requires an industry-by-month series of tariff shocks. We use effective tariff rates by imported commodity from the US International Trade Commission (USITC). Since these customs data are reported by commodity, rather than by purchasing industry, we construct the industry-level tariff shocks as the weighted average of the effective tariff rate on each imported input, with weights given by the import expenditure shares of each commodity for each industry from the 2017 BEA Imports Matrix.

In addition to simulating the log-linearized model—where only industries' steady-state inventories matter—we also consider how the transmission of these shocks to consumer prices would change if we let industry inventories fluctuate from their steady-state levels due to forward buying. In particular, for each imported commodity c, we use monthly imports deflated by the BLS MXP Import Price Deflator series to read import

[21] *Reuters*, "US coffee drinkers face higher prices even after Trump's tariff reset."
[22] *CNN Business*, "Cocoa prices are falling—but not in time for Easter."

quantities off the data Y_{ct}. We seasonally adjust these imports using the 2021–2024 series and take the log-deviation from October 2024 as the difference from steady state, $d \log Y_{ct}$. We assume excess inventory is stockpiled by the importing industry.[23]

For the alternative model with Calvo frictions, we also need to take a stand on firms' expectations of future tariff shocks. We consider two versions. The first version assumes that changes in tariffs each period are unanticipated, though the effects of already-present tariffs on industries' future price adjustments are fully anticipated. The second version assumes that the tariff path is expected by firms starting in November 2024.

Results. Figure 13 plots the evolution of the price index for personal consumption expenditures under AAC and FIFO accounting cost pricing. We compare these price paths to a model with entirely flexible prices, as well as to the price paths that obtain in the Calvo calibrations where firms either are surprised by each successive tariff announcement or fully anticipate the tariff path starting in November 2024.

At their peak in October 2025, tariffs in the flexible-price economy would have increased consumer prices by 0.8 percentage points. Under AAC pricing, the pass-through of tariffs to consumer prices is delayed, though it rises to about 0.6 percentage points by October 2025. Interestingly, forward buying has relatively little effect on further delaying pass-through. Instead, the bulk of the delays in the model are caused by firms' existing input and inventory stocks. Similar results obtain, with slightly longer delays in pass-through, under FIFO pricing. On the other hand, the model with Calvo frictions generates a faster pass-through of the tariffs to consumer prices, and particularly so if we make the (strong) assumption that firms anticipated the path of tariffs starting in November 2024.

We compare predicted and observed price changes for 212 subcomponents of the personal consumption expenditures price index. Minton et al. (2026) use the following specification to measure the pass-through of tariff shocks to PCE category price indices,

$$\Delta \hat{p}_{jt}^{\text{data}} = \sum_{h=0}^{H} \beta_h \Delta \hat{p}_{jt-h}^{\text{flex}} + \alpha_j + \lambda_t + \varepsilon_{jt}, \tag{14}$$

where $\Delta \hat{p}_{jt}^{\text{data}}$ is the change in the log price index for PCE category j in month t, $\Delta \hat{p}_{jt-h}^{\text{flex}}$ is

[23]Under average acquisition cost accounting, given $d \log Y_{tc}$, the book value of industry i's inventory of input c follows the law of motion,

$$\hat{h}_{tic} = \hat{h}_{t-1ic} \frac{\kappa_i - 1 + \sum_{k=0}^{t-1} d \log Y_{kc}}{\kappa_i + \sum_{k=0}^{t} d \log Y_{kc}} + \frac{1 + d \log Y_{tc}}{\kappa_i + \sum_{k=0}^{t} d \log Y_{kc}} \hat{c}_{tc}, \qquad \text{where} \quad \kappa_i = \begin{cases} \delta_i^{-1} & \text{if } I_i^{ss} > 0 \\ 0 & \text{if } I_i^{ss} = 0 \end{cases}.$$

An analogous expression can be derived for the law of motion of accounting costs under FIFO.

Figure 13: Response of the consumer price index to the 2025 tariff shock.

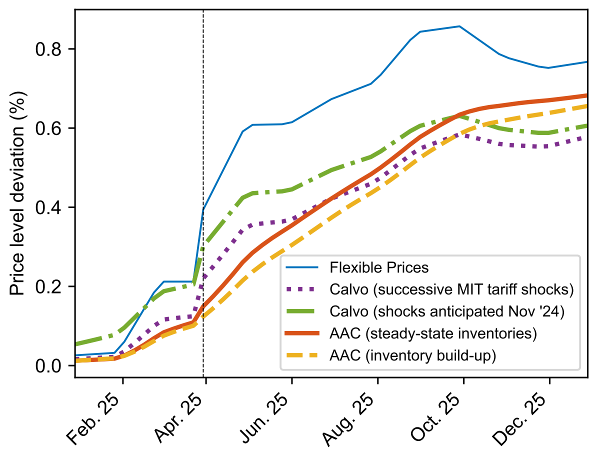

(a) AAC. **(b)** FIFO.

Note: The left panel plots the response of the PCE price index, in percentage points, to the effective tariff rate series derived from USITC data under AAC pricing. The right panel plots the response under FIFO pricing. The dashed vertical line marks April 2$^{\text{nd}}$, 2025 ("liberation day"). The numbers in parentheses report the cumulative gap between flexible-price and accounting-cost price indices, $\sum_t (\hat{p}^{\text{flex}}_{\text{PCE},t} - \hat{p}_{\text{PCE},t})$.

Figure 14: Pass-through to consumer prices of the 2025 tariff shocks.

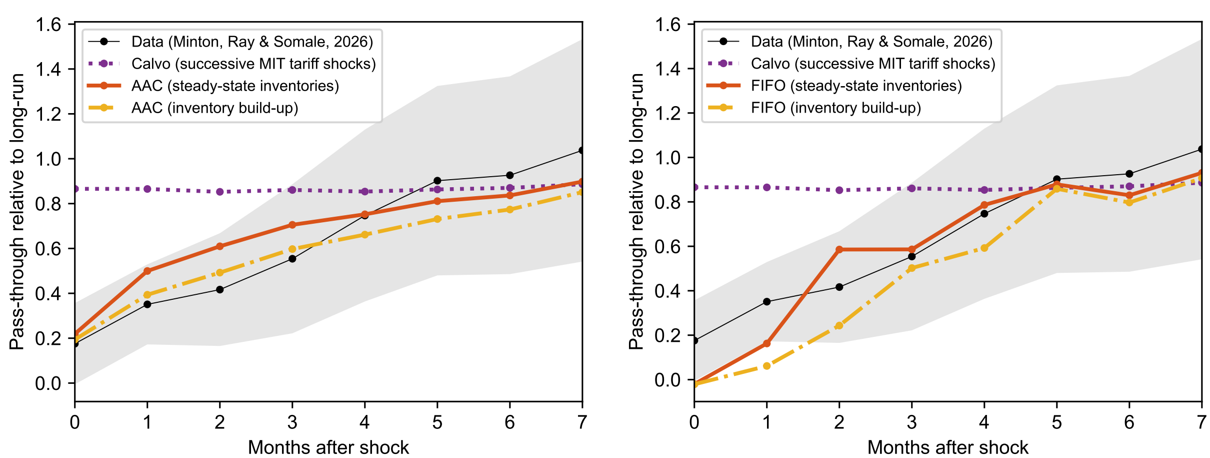

(a) AAC. **(b)** FIFO.

Note: The left panel plots the pass-through estimates from Equation (14) for AAC pricing. The right panel plots the same estimates under FIFO pricing. In both panels, the plotted lines are the cumulative sum of the distributed lag coefficients, $\sum_{k=0}^{h} \beta_k$. The empirical estimate is Minton et al. 2026, Figure 3. The numbers in parentheses denote the cumulative pass-through delay calculated from the regression estimates over the lag horizon H, $\sum_{h=0}^{H} \left(1 - \sum_{k=0}^{h} \beta_k\right)$.

Table 6: Response of PCE category inflation to 2025 tariffs: Models *vs.* data.

	PCE category monthly inflation: Data				
	(1)	(2)	(3)	(4)	(5)
Inflation in AAC Model	0.782**			1.338**	
	(0.337)			(0.415)	
Inflation in FIFO Model		0.649**			0.666**
		(0.194)			(0.197)
Inflation in Calvo Model			−0.085	−0.574**	−0.155
			(0.124)	(0.204)	(0.110)
Date FEs	Yes	Yes	Yes	Yes	Yes
PCE Category FEs	Yes	Yes	Yes	Yes	Yes
N	6572	6572	6572	6572	6572
R^2	0.04	0.04	0.04	0.04	0.04

Note: The AAC and FIFO models are the calibrated models allowing for inventory build-up. The Calvo model assumes that changes in tariffs each month are unanticipated by firms. Standard errors are two-way clustered by PCE category and month. * indicates significance at 10%, ** at 5%.

the change in the log price index for PCE subcomponent j in month $t-h$ in a model with flexible prices, and α_j and λ_t are PCE category and time fixed effects. This specification uses both cross-sectional and time series variation in the hypothetical long-run pass-through of tariffs to each PCE category to estimate the average speed of pass-through from tariffs to consumer prices. The estimates obtained by Minton et al. (2026) are plotted in both panels of Figure 14. They find that pass-through is sluggish at first, but reaches complete pass-through at a horizon of seven months.

We run the same specification across PCE categories in our model, using the model-generated price changes $\Delta\hat{p}_{jt}^{\text{model}}$ to replace $\Delta\hat{p}_{jt}^{\text{data}}$ on the left-hand side of (14). Figure 14 overlays our estimates on the empirical results from Minton et al. (2026). Both the AAC and FIFO models predict delays in pass-through roughly in line with the empirical estimates. The calibrations with forward buying of imports generate slightly more sluggish pass-through. The figure also shows the pass-through estimated on the Calvo model where firms are surprised by new tariff shocks in each period. That model predicts a high degree of pass-through on impact, though the convergence to complete pass-through is very slow. The high pass-through on impact is caused by the fact that many of the tariffs apply to final goods used for personal consumption expenditures; in the accounting cost models, these tariffs still take time to affect consumer prices due to the inventories of wholesalers and retailers importing these goods.

While specification (14) estimates an aggregate speed of adjustment, it does not tell us whether the model successfully predicts cross-sectional differences in the speed of pass-through. To check whether this is the case, we regress monthly inflation in the data for each PCE category on monthly inflation for that category predicted by the model, with PCE category and time fixed effects:

$$\Delta\hat{p}_{jt}^{\text{data}} = \beta\Delta\hat{p}_{jt}^{\text{model}} + \alpha_j + \lambda_t + \varepsilon_{jt}. \tag{15}$$

Table 6 reports the results. Both the AAC and FIFO predictions are positively and significantly associated with PCE-category inflation, with coefficients of $\beta \approx 0.78$ and 0.65, respectively. The coefficient on the Calvo prediction is small, negative, and not statistically distinguishable from zero. The accounting cost models continue to be positively associated with inflation when included alongside the Calvo prediction. These results indicate that accounting cost pricing captures variation in category-level inflation beyond category and month fixed effects, although there is still a large share of variation in month-to-month inflation that neither the accounting cost pricing nor Calvo pricing models can explain.

7 Conclusion

This paper proposes a theory of *accounting cost pricing*. In a simple principal–agent model, we show that these behaviors can arise because of a need to monitor managerial decisions. While the incentives used by the principal can limit value-eroding managerial discretion, they also distort managers' pricing decisions to incorporate information about historical input costs and fixed costs that are irrelevant to the optimal price and ignore valuable information about current replacement costs and demand conditions that are germane.

Accounting cost pricing can predict similar price dynamics to sticky-price models—though the parameters that discipline delayed adjustment under accounting cost pricing depend on the lag between when inputs or inventories are purchased and when they are shipped out to customers, rather than the length of price spells. But the dependence on how historical purchase costs work through accounting costs of goods sold also predicts some key differences in how firms' prices respond to slow-moving persistent and transitory shocks, and how much supply chains can exacerbate delays in price adjustment. In the context of the 2025–2026 tariffs, accounting cost pricing, combined with forward buying, can explain why pass-through to inflation was gradual, even when firms anticipated tariffs well in advance.

References

Afrouzi, H., J. P. Flynn, and C. Yang (2024). What can measured beliefs tell us about monetary non-neutrality? Technical Report 32541, National Bureau of Economic Research.

Anderson, E., S. Rebelo, and A. Wong (2025). Markups across space and time. *Review of Economic Studies* (rdaf098).

Antràs, P. and V. Tubdenov (2025). Measuring the average period of production. *AEA Papers and Proceedings 115*, 624–630.

Arcidiacono, P., P. B. Ellickson, C. F. Mela, and J. D. Singleton (2020). The competitive effects of entry: Evidence from supercenter expansion. *American Economic Journal: Applied Economics 12*(3), 175–206.

Baqaee, D. R. and E. Farhi (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics 135*(1), 105–163.

Baumol, W. J. (1959). *Business Behavior, Value, and Growth*. New York: MacMillan.

Bertrand, M. and S. Mullainathan (2003). Enjoying the quiet life? Corporate governance and managerial preferences. *Journal of Political Economy 111*(5), 1043–1075.

Bewley, T. F. (2025). *Price Setting*. John Wiley & Sons.

Bils, M. and Y. Chang (2000). Understanding how price responds to costs and production. *Carnegie-Rochester Conference Series on Public Policy 52*, 33–77.

Blinder, A., E. R. Canetti, D. E. Lebow, and J. B. Rudd (1998). *Asking about prices: A new approach to understanding price stickiness*. Russell Sage Foundation.

Burstein, A., V. M. Carvalho, and B. Grassi (2025). Bottom-up markup fluctuations. *The Quarterly Journal of Economics 140*(4), 2619–2684.

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

Cagan, P. (1979). *Persistent Inflation: Historical and policy essays*. Columbia University Press.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Carreras-Valle, M.-J. and S. M. Lee (2026). Tariff front-running. Working paper.

Eyster, E., K. Madarász, and P. Michaillat (2021). Pricing under fairness concerns. *Journal of the European Economic Association 19*(3), 1853–1898.

Fabiani, S., M. Druant, I. Hernando, C. Kwapil, B. Landau, C. Loupias, F. Martins, T. Matha, R. Sabbatini, H. Stahl, and A. Stokman (2006). What firms' surveys tell us about price-setting behavior in the euro area. Working Paper 808, MPRA.

Gagnon, E. and D. Lopez-Salido (2020). Small price responses to large demand shocks. *Journal of the European Economic Association 18*(2), 792–828.

Giroud, X. and H. M. Mueller (2010). Does corporate governance matter in competitive industries? *Journal of Financial Economics 95*(3), 312–331.

Giroud, X. and H. M. Mueller (2011). Corporate governance, product market competition, and equity prices. *The Journal of Finance 66*(2), 563–600.

Goldberg, P. and R. Hellerstein (2008). A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market. *American Economic Review 98*(2), 423–429.

Gordon, R. J. (1981). Output fluctuations and gradual price adjustment. *Journal of Economic Literature 19*(2), 493–530.

Gordon, R. J. (1990). What is New-Keynesian economics? *Journal of Economic Literature 28*(3), 1115–1171.

Graham, J. R., C. R. Harvey, and S. Rajgopal (2005). The economic implications of corporate financial reporting. *Journal of Accounting and Economics 40*(1-3), 3–73.

Hall, R. L. and C. J. Hitch (1939). Price theory and business behavior. *Oxford Economic Papers* (2), 12–45.

Hart, O. D. (1983). The market mechanism as an incentive scheme. *The Bell Journal of Economics 14*(2), 366–382.

Horowitz, K. J. and M. A. Planting (2009). Concepts and methods of the US input–output accounts. Technical report, Bureau of Economic Analysis.

Hortaçsu, A., O. R. Natan, H. Parsley, T. Schwieg, and K. R. Williams (2024). Organizational structure and pricing: Evidence from a large US airline. *Quarterly Journal of Economics 139*(2), 1149–1199.

Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review 76*(4), 728–741.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Kohler, T., J.-P. L'Huillier, G. Phelan, and M. Weiss (2026). Why do supply disruptions lead to inflation (while demand booms do not)? Survey evidence from the COVID pandemic. Working paper.

Ling, Q. and D. P. Tinkelman (2024). Costs and benefits of the LIFO-FIFO choice. *Journal of Corporate Accounting & Finance 35*(3), 11–20.

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics 117*(4), 1295–1328.

Minton, R., M. Ray, and M. Somale (2026). Detecting tariff effects on consumer prices in real time - part ii. Feds notes, Federal Reserve Board of Governors.

Minton, R. and B. Wheaton (2022). Hidden inflation in supply chains: Theory and evidence. Working paper.

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

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

Okun, A. M. (1981). *Prices and Quantities: A Macroeconomic Analysis*. The Brookings Institution.

Pasten, E., R. Schoenle, and M. Weber (2020). The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics 116*, 1–22.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics 48*(1), 7–35.

Rogerson, W. P. (2008). Intertemporal cost allocation and investment decisions. *Journal of Political Economy 116*(5), 931–950.

Rogerson, W. P. (2011). On the relationship between historic cost, forward looking cost, and long run marginal cost. *Review of Network Economics 10*(2).

Rotemberg, J. J. (2005). Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. *Journal of Monetary Economics 52*(4), 829–852.

Rotemberg, J. J. (2011). Fair pricing. *Journal of the European Economic Association 9*(5), 952–981.

Rubbo, E. (2023). Networks, Phillips curves and monetary policy. *Econometrica 91*(4), 1417–1455.

Sangani, K. (2025). Complete pass-through in levels. Working paper.

Schoenle, R. (2026). Discussion of ”complete pass-through in levels”. NBER Economic Fluctuations and Growth Winter Meeting.

Shi, L., Q. Liu, and M. Kim (2024). Does sunk-cost affect prices? Evidence from the US airline industry.

Journal of Economics & Management Strategy 33(4), 801–824.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy 88*(1), 1–23.

Woodford, M. (2003). Imperfect common knowledge and the effects of monetary policy. *Knowledge, information, and expectations in modern macroeconomics: In honor of Edmund S. Phelps 25*(1), 4.

Online Appendix to *Pricing Off Accounting Costs*

(Not for publication)

A Additional Tables and Figures

Table A1: Inventories and prices: Examples from earnings calls.

Company	Date	Quote
Wayfair	Oct 31, 2019	As suppliers communicate pricing decisions to us, our pricing algorithms seek to maintain a balance between factoring in our higher wholesale costs and keeping us market priced. In practice, this [...] means that for a short period of time, inventory-carrying retailers may have a slight cost advantage. This should dissipate over the next couple of quarters as the market rebalances, pre-tariff inventory is depleted and all retailers incur higher wholesale costs reflective of tariffs.
Boeing	Apr 23, 2025	So we have the opportunity for those duties that we pay, those tariffs that we pay, to recover that when we deliver the aircraft. Now, you know, it creates a little bit of a cash flow timing issue that we'll have to work through.
Global Industrial	July 29, 2025	We anticipate implementing additional pricing actions as inventory affected by tariffs moves through our cost of sales. [....] Management of our margin profile remains a key area of focus as we move through the current cycle we have started to see the timing benefit of pre tariff inventory decline. [...O]ur goal is to manage price cost neutral.
Floor and Decor	July 31, 2025	And looking across spectrum of publicly traded, you know, flooring companies and people that sell our category, we feel like we're doing better and we feel like we are taking share. I do think that the independent channels have taken price earlier, quicker because they had to and we're able because the way we turned this is when the benefit of slow turning inventory is beneficial. We were able to hold off and take less price increase.
Oil States International	Oct 31, 2025	Now it is also our belief that everybody has the same supply sources which are generally foreign. We're all experiencing the same thing. But there's also been a build up of inventory as activity is slowed. So I think the industry has to work through the pre-tariff inventory. But then it is my view if the tariffs hold they're going to have to be passed on to customers. It's one of timing.

Newell Brands	Aug 1, 2025	We have identified 10 other categories where we have a relative sourcing advantage versus competition based on country of origin and or where we have existing tariff-free inventory available for incremental promotions. Collectively, we have successfully secured tariff-related relative sourcing advantage or tariff-free inventory wins with over 30 customers across nearly every domestic channel where we go to market. [....] Broadly what we're seeing is that in most categories where everybody is affected by tariffs, pricing is moving up, but it's not moving up at the same on the same timing. We're not seeing generally speaking more aggressive promoted prices. But what we are seeing is in some cases competitors are delaying the date of the price increase because they still have inventory onshore that hasn't been subject to the tariffs. We expect this situation on pricing from a competitive standpoint to really get a lot clearer over the next three to six months as the pre-tariff inventory sort of runs out and the price increases become more visible at retail.
Whirlpool	Oct 28, 2025	So we know there's quite a bit of an inventory overhang. Inventory which was at pre tariff cost. That's an important one. Of course by definition as you go through Q3 and Q4 with the anticipation that import volumes come down, that excess inventory at one point will quote unquote flush through the system, we would expect that to be happening kind of towards the end of Q4. [...] I would expect in 26 to see industry behavior which is more reflective of normal shipment patterns and particularly more important the underlying cost base.
Light & Wonder	Nov 5, 2025	The reality is we've worked through the majority of the pre-tariff inventory at this point and so we will now start to see impacts to that mid to low, sorry, mid to high single digit millions as we go forward. That's something that we and the team are going to continue to look at through the margin enhancement initiatives to see what we can mitigate.
Academy Sports and Outdoors	Dec 9, 2025	The goal was when we first learned these accelerated tariffs is we were looking at it saying, okay, there's a lot of inventory on this side of the water at those pre-accelerated tariff prices. If we can pull those into our warehouses and DCs, that should allow us to be priced at last year's level on a lot of these items going into holiday.

John Deere	Feb 19, 2026	You know, we've seen, you know, competition, you know, as of late, last quarter or two, you know, take some price increase, signal some price increase. The timing of that, we're keeping an eye on, you know, our, many of our competitors still have a lot of inventory in the, in the field. So as those price increases for 2026 start to manifest themselves in transaction price, you know, there's a bit of a lag there that, that, that creates some pressure.

Table A2: Effect of airlines' accounting costs of fuel on ticket prices: Robustness and extensions.

	Log Fare$_{irt}$				
	Strategic complem.		Asymmetry	Sticky fares	
	(1)	(2)	(3)	(4)	(5)
Log Spot Fuel Price$_t$	0.348**	0.350**	0.354**	0.356**	0.355**
	(0.028)	(0.028)	(0.030)	(0.030)	(0.030)
Fuel Cost Diff$_{it}$	0.373**	0.436**		0.408**	0.409**
	(0.136)	(0.107)		(0.115)	(0.115)
Passenger Share on Route$_{irt}$	−0.100**				
	(0.017)				
Fuel Cost Diff$_{it}$ × Passenger Share$_{irt}$	0.230**				
	(0.108)				
Fuel Cost Diff for All Airlines on Route$_{rt}$		0.104**			
		(0.038)			
Fuel Cost Diff$_{it}$ (Positive)			0.653**		
			(0.168)		
Fuel Cost Diff$_{it}$ (Negative)			0.360**		
			(0.174)		
$\Delta_{t,t-4}$ Log Spot Price × Share Tickets Bought ≥3 Weeks in Advance$_i$				−0.153*	
				(0.081)	
$\Delta_{t,t-4}$ Log Spot Price × Share Tickets Bought ≥3 Months in Advance$_i$					−0.801*
					(0.434)
Carrier × Route FEs	Yes	Yes	Yes	Yes	Yes
N	1765168	1765168	1765168	1765168	1765168
R^2	0.58	0.58	0.58	0.58	0.58

Note: The outcome variable, Log Fare$_{irt}$, is the log of the average fare for carrier i on route r in quarter t in the DB1B sample. Spot Fuel Price$_t$ is the spot price of jet fuel from the EIA, and Fuel Cost Diff$_{it}$ = Log Fuel Cost$_{it}$ − Log Spot Fuel Price$_t$ is the log difference between carrier i's accounting cost of jet fuel in quarter t and the spot price. Passenger Share$_{irt}$ is the share of passenger seats flown by carrier i on route r in quarter t, as a share of total route capacity from the BTS T100 data. $\Delta_{t,t-4}$Log Spot Price is the change in the log spot price from quarter $t-4$ to quarter t. The shares of tickets bought more than 21 days and more than 90 days in advance are calculated for each carrier using the DB1C data. Standard errors are two-way clustered by route and quarter. * indicates significance at the 10% level, ** at 5%.

Table A3: Volatility of airlines' operating margins, before and after adjustment for gap between accounting cost of fuel and spot price.

Carrier	Pass-through at 1 year	Volatility of operating margins: Reported	Volatility of operating margins: Adj. for spot price	Pct. change
Southwest	0.70	0.064	0.084	**+31.2**
Delta	0.78	0.079	0.087	**+10.4**
American	0.95	0.067	0.071	**+5.4**
Hawaiian	0.97	0.108	0.110	**+2.0**
Frontier	0.99	0.171	0.170	−0.5
Allegiant	1.01	0.081	0.079	−2.4
JetBlue	1.03	0.078	0.085	**+10.2**
United	1.03	0.072	0.072	−0.1
Alaska	1.04	0.098	0.101	**+3.3**
Spirit	1.10	0.136	0.135	−0.8

Note: The column "Pass-through at 1 year" is the cumulative pass-through at four quarters ($\sum_{k=0}^{3} \beta_{ik}$) for each carrier estimated using specification (1). Operating margins are calculated as operating profits (sales less COGS less SG&A) as a fraction of sales, and the volatility of operating margins is calculated over 1990–2026 excluding the COVID-19 pandemic period from 2020–2021. The adjusted column uses adjusted profits = operating profits $+((\text{Fuel Cost}_{it} - \text{Spot Fuel Price}_t) \times \text{Fuel Consumption}_{it})$. Here fuel costs and spot prices are measured in dollars per gallon, and fuel consumption is measured in gallons. The final column reports the percentage change in the volatility of operating margins before and after the adjustment, with bold entries indicating positive values.

Table A4: Descriptive statistics about firm respondents in the survey conducted by Schoenle (2026).

Q: Which best describes your role?		
Answer	Count	Share
Owner / co-owner	106	0.26
General manager	83	0.21
Supervisor	69	0.17
Store manager	46	0.12
Pricing manager / revenue manager	36	0.09
Other	60	0.15
Q: Are you personally involved in deciding the posted price(s) for your main product(s) or service(s)?		
Yes, I decide the prices	260	0.65
Yes, I implement prices decided by others – but have some leeway such as timing or size of changes	140	0.35
Q: What is your main product or service category?		
Professional services (e.g., repair, cleaning)	104	0.26
Grocery / food retail	67	0.17
Personal services (e.g., barbershop, salon, spa)	46	0.12
Building materials / hardware	41	0.10
Restaurants / prepared food	35	0.09
Retail gasoline / fuel	20	0.05
Other	87	0.22
Q: How many locations does your business operate?		
1	170	0.42
2–5	127	0.32
6–20	65	0.16
21+	38	0.10

Table A5: Compustat sample filters.

Filter	Description	Observations	Unique firms
Starting sample	–	2 106 695	45 401
Industry	Excludes firms with missing NAICS codes or in NAICS sector 51 (information), 52–53 (finance, insurance, and real estate)	1 272 910	23 607
Key fields	Following fields are non-empty and strictly positive: Inventory stock (`INVTQ`) in current quarter and same quarter of prior year, sales (`SALEQ`), and cost of goods sold (`COGSQ`).	801 669	18 187
Comparable observations	`COMPSTQ` is empty, and year-over-year change in total assets is below 30%.	580 386	17 771
Inventory valuation method	Primary method is FIFO or AAC	423 153	14 686

Table A6: Estimating the extent of replacement cost pricing: Robustness.

	Log(Sales / COGS)					
	Levels			First Differences		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Purchases / COGS)	0.145**	0.135**		0.116**	0.122**	
	(0.013)	(0.019)		(0.012)	(0.016)	
Log(Purchases / COGS) for LIFO firms in industry		0.004			0.003	
		(0.007)			(0.010)	
Log(Purchases / COGS) × AAC			0.177**			0.128**
			(0.016)			(0.019)
Log(Purchases / COGS) × FIFO			0.124**			0.124**
			(0.020)			(0.016)
Firm × Quarter-of-Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Calendar Quarter FEs	Yes			Yes		
N	382 211	257 763	382 211	345 468	228 312	345 468
R^2	0.87	0.88	0.86	0.10	0.12	0.09

Note: Columns 1–3 report results from (6), and columns 4–6 report results from (7). We exclude observations in the bottom and top 1 percent of values for the dependent and independent variables. The control for Log(Purchases / COGS) for LIFO firms in firm *i*'s industry uses the median Log(Purchases / COGS) for all LIFO firms in the same NAICS-4 industry as *i* in the same calendar quarter. Regressions are weighted by sales. Driscoll-Kraay standard errors with eight lags in parentheses. * indicates significance at 10%, ** at 5%.

Figure A1: Excerpt from Delta Air Lines 2009 Q1 earnings call.

	1Q	2Q	3Q	4Q	2009
Call options	7%	12%	29%	19%	17%
Collars	33%	26%	4%	0%	16%
Swaps	37%	37%	23%	17%	28%
Total	77%	75%	56%	36%	61%
Downside participation		37%	73%	83%	
All-in projected fuel price / gallon*	$2.26	$2.08	$1.91	$1.73	$1.99
Hedge loss/gallon included in projected price	$0.71	$0.45	$0.18	$0.03	$0.34

* Includes tax and transportation costs of approximately $0.17/gallon, and call option premiums

Note: "Delta hedged 77% of its fuel consumption for the March 2009 quarter, which drove $684 million in realized fuel hedge losses for the period. As a result, Delta's average fuel price for the March 2009 quarter was $2.26 per gallon, including hedge losses, taxes and transportation." Reported on pg. 6 of Delta Air Lines earnings call on April 21, 2009. `https://www.sec.gov/Archives/edgar/data/27904/000101968709001435/delta_8k-ex9901.htm`.

Figure A2: Carrier hedging is imperfectly correlated with share of early ticket purchases.

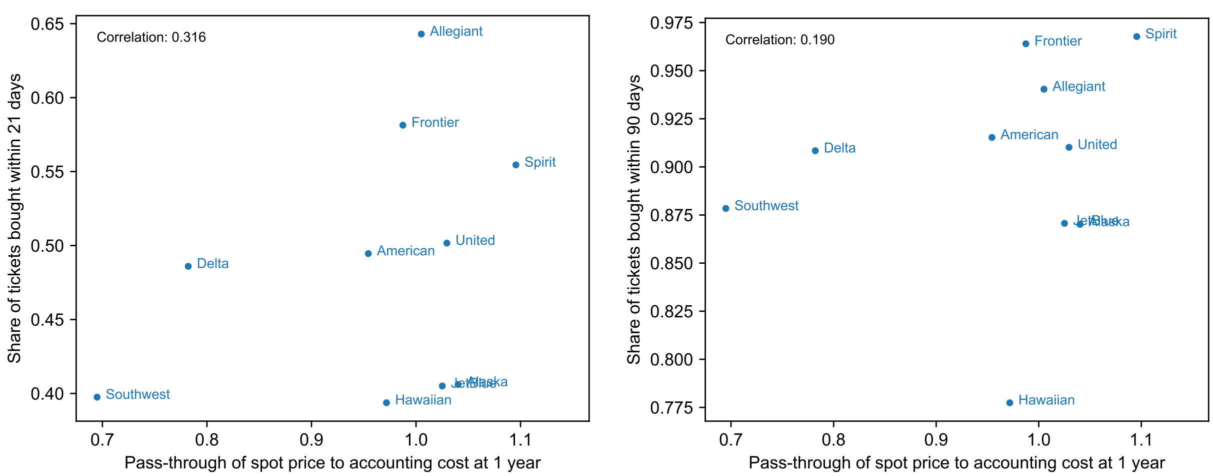

(a) Share bought within 21 days of departure. **(b)** Share bought within 90 days of departure.

Note: The x-axis shows the pass-through of spot fuel prices to accounting costs of fuel per gallon at four quarters, estimated using (1). The y-axes show the share of passenger tickets purchased within 21 days and within 90 days of departure for each airline across all routes in the available DB1C sample. The figure reports the correlation of the two statistics across the ten airlines.

Figure A3: Responses to "Suppose your input cost rises today, but you still have inventory bought at lower cost. When do you raise your sales price?"

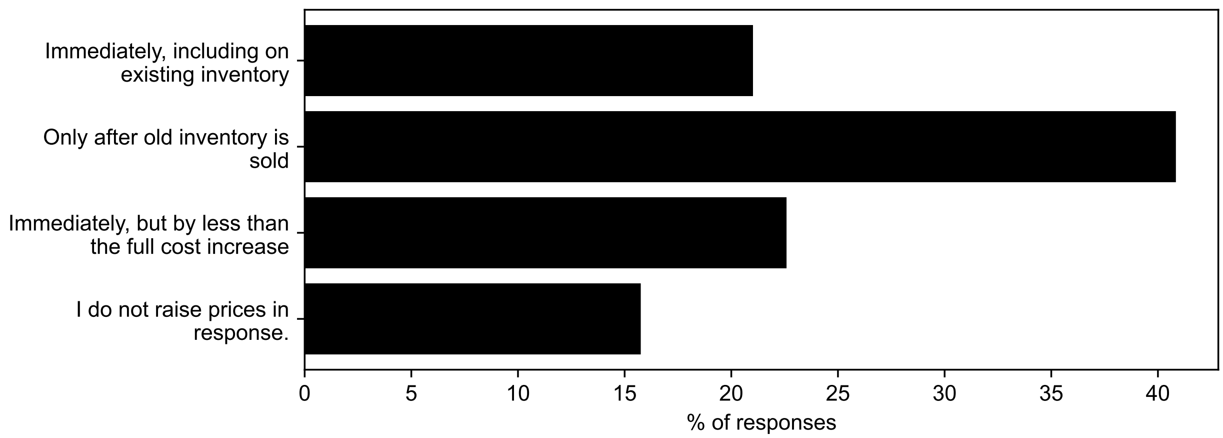

(a) All respondents.

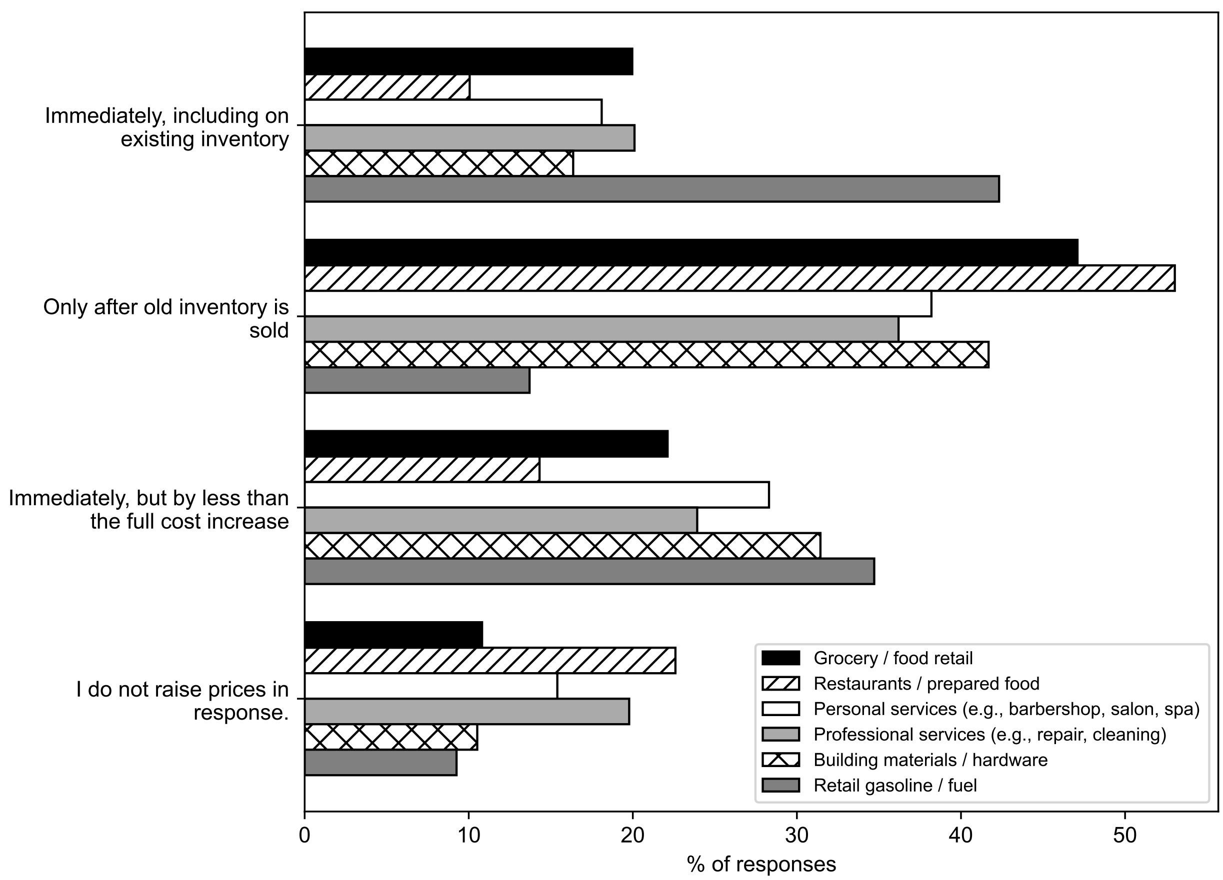

(b) Respondents, split by industry of main product or service.

Note: Responses from a survey of 400 firm price-setters conducted by Schoenle (2026).

Figure A4: Responses to "Suppose your overhead costs go up, for example on rent, insurance, salaried staff or equipment leases. How do increases in such overhead costs affect your pricing—compared to increases in per-unit input costs on materials, wholesale goods, shipping per item, or hourly labor per job?"

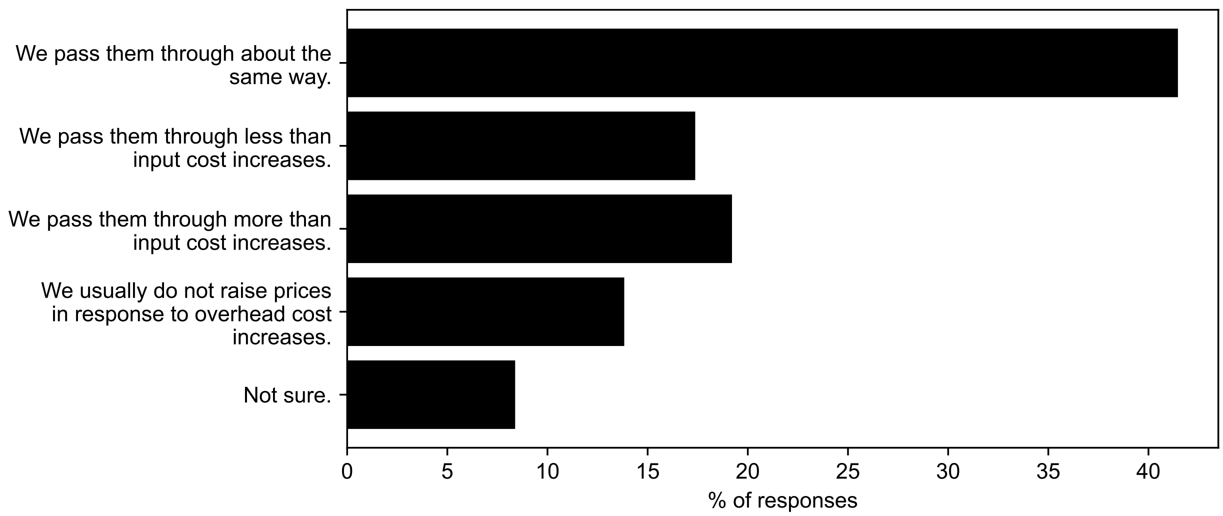

Note: Responses from a survey of 400 firm price-setters conducted by Schoenle (2026).

Figure A5: Histogram of inventory–shipment ratios across NAICS-6 industries.

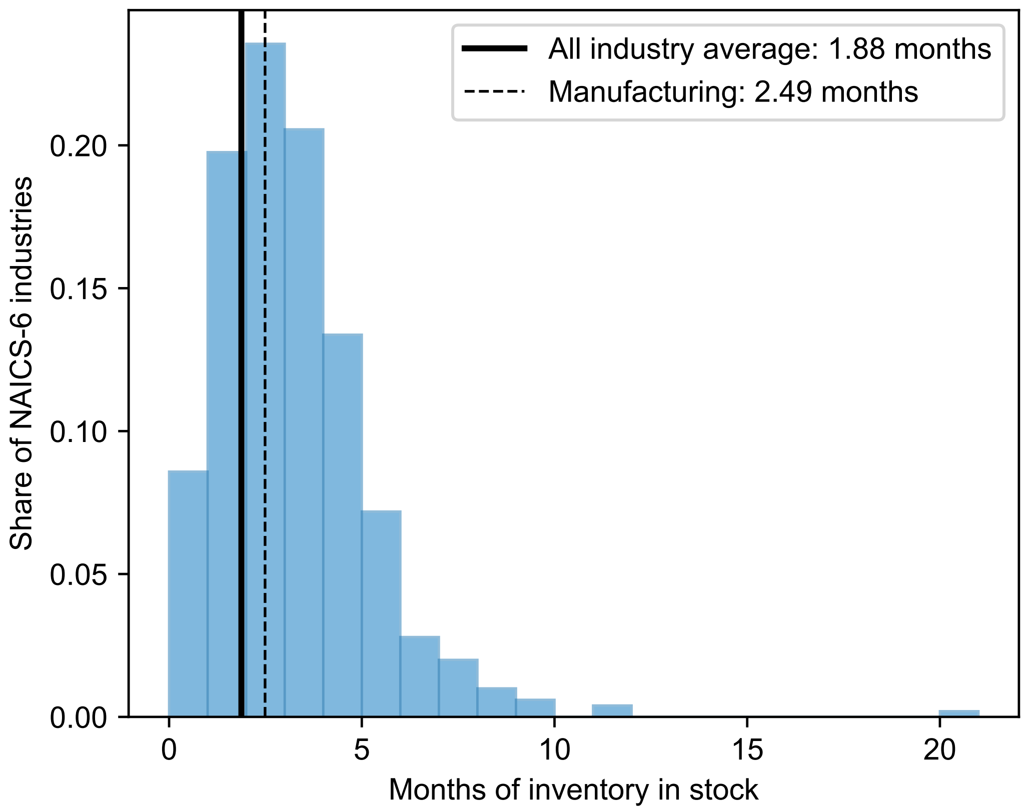

Note: The histogram includes 501 NAICS-6 industries in the following sectors: mining (NAICS 21), manufacturing (NAICS 31–33), wholesale trade (NAICS 42), and retail trade (NAICS 44–45).

Figure A6: Correlation between inventory–shipment ratios from 2017 Economic Census and measures from public firms by Antràs and Tubdenov (2025).

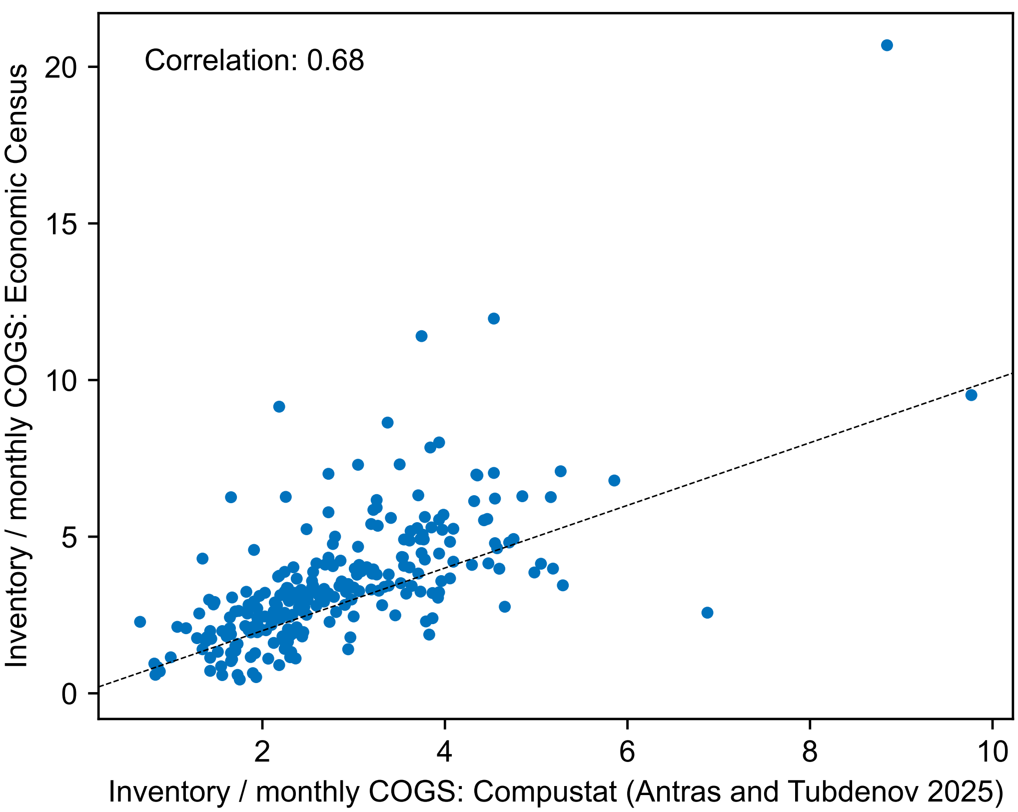

Note: Each point represents a NAICS-6 industry. The dotted line is the 45-degree line. The plot contains 246 NAICS-6 industries where we can construct measures of the inventory–COGS ratio from the 2017 Economic Census and where Antràs and Tubdenov (2025) report the average inventory–COGS ratio for public firms (which they refer to as the average period of production).

Figure A7: Delayed effects of industry cost shocks on consumer prices in FIFO calibration.

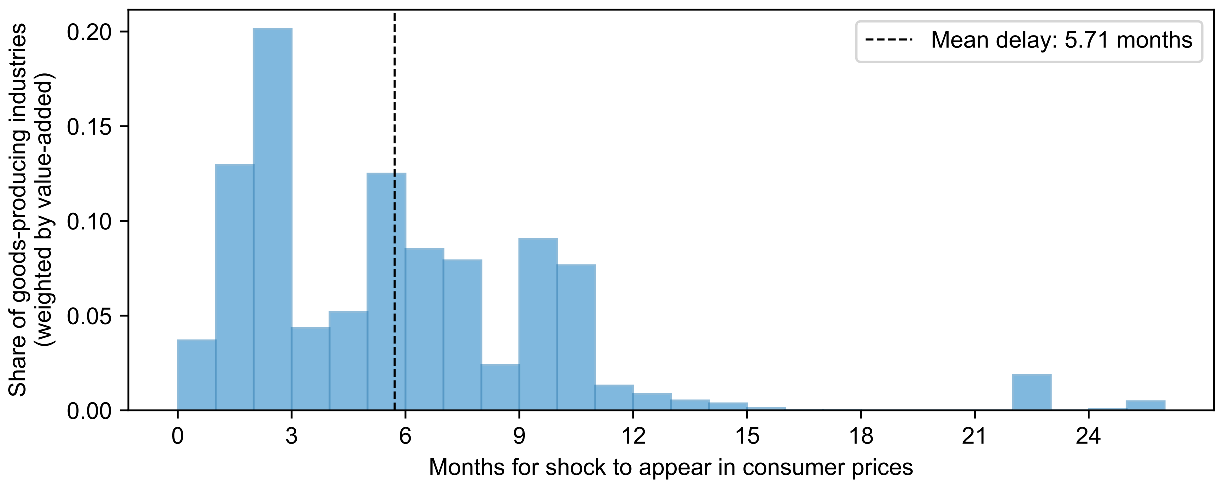

Note: For each goods-producing industry, we simulate a one-time, permanent shock to production costs and calculate the cumulative pass-through delay, $\sum_{t=0}^{\infty}(1 - \hat{p}_{\text{PCE},t}/\hat{p}_{\text{PCE},\infty})$. The figure shows a histogram of these delays across industries weighted by each industry's value added in the calibration with first-in-first-out (FIFO) accounting.

B Constructing the Trade-Adjusted Input–Output Matrix

This appendix describes how we construct the trade-adjusted input–output matrix used in the quantitative analysis. The adjustment makes wholesale and retail intermediation explicit: a trade industry purchases merchandise, combines it with the inputs required to distribute it, and resells it to the next buyer. This representation allows inventories held by wholesalers and retailers to delay the transmission of upstream cost changes. We implement the construction using the 2017 detailed BEA supply and use tables and the corresponding *Margins Before Redefinitions* table, together with product-line sales and gross margins from the 2017 Economic Census.

From trade margins to merchandise flows. In the BEA accounts, trade output is measured by the margin earned on merchandise sales; merchandise purchased for resale is not an intermediate input of the trade industry. Chapter 8 of Horowitz and Planting (2009) describes how the BEA assigns trade margins to commodities and uses. A central ingredient is the margin rate of the most appropriate kind of business: the industry that primarily sells a product line supplies a proxy for that product line's margin rate. The BEA combines these rates with product-line sales and reconciles the resulting margins to aggregate controls. Our construction works backward from the published margins, using matched trade-industry margin rates to infer merchandise sales. Because the published data do not identify every distribution route, this is a reconstruction under the allocation assumptions below, rather than an exact inversion of all BEA estimation and balancing steps.

Let c index commodities and b index buyers in the margins table, including both industries and final uses. For each commodity–buyer pair, denote purchasers' expenditure by E_{cb}, producer value by P_{cb}, transportation costs by T_{cb}, and wholesale and retail margins by W_{cb} and R_{cb}. The valuation identity is

$$E_{cb} = P_{cb} + T_{cb} + W_{cb} + R_{cb}. \tag{16}$$

The margins are dollar amounts attributable to a transaction, not the gross sales of the intermediaries involved. In particular, a positive wholesale margin can reflect a wholesaler selling to the buyer directly or selling to a retailer who subsequently serves that buyer.

Matching commodities to trade margin rates. We hand-construct a concordance from BEA commodities to wholesale and retail product lines in the North American Product Classification System (NAPCS). Separately for each channel, we identify the principal selling industry for the matched product line using the Economic Census: we select the NAICS industry with the largest reported contribution to the product line's sales, breaking ties by its dollar sales of that line. We then assign the commodity the gross margin rate of this industry. For channel $k \in \{W, R\}$, the assigned rate is

$$g_c^k = \frac{\text{gross margin of the matched trade industry}}{\text{sales of the matched trade industry}}. \tag{17}$$

For wholesalers, the denominator uses sales on own account where available; otherwise it uses reported sales. Thus, g_c^k is a margin as a fraction of sales, not a markup over acquisition cost. Missing detailed margin rates are filled from the nearest available parent NAICS industry, searching from six to three digits. Wholesale trade agents and brokers (NAICS 425120) are excluded from the product-line selection because they lack the required margin data. The concordance also associates each new trade industry with a BEA trade-industry code for matching industry characteristics.

This procedure allows margin rates to differ across commodities but holds a commodity's assigned rate fixed across buyers. Buyer-specific variation in observed margins relative to expenditure therefore translates into variation in the inferred distribution channels. The principal-industry match is used to assign margin rates and industry characteristics; the operating-input mix of the new trade industries is constructed separately below.

Allocating each buyer's purchases across channels. For a positive expenditure cell, define $[x]_0^1 \equiv \min\{1, \max\{0, x\}\}$. We first infer the share of expenditure on purchases from retailers:

$$s_{cb}^R = \left[\frac{R_{cb}}{g_c^R E_{cb}}\right]_0^1, \qquad X_{cb}^R = s_{cb}^R E_{cb}. \tag{18}$$

Without a binding cap, dividing the observed retail margin by the retail margin rate gives retailers' gross sales to the buyer. We next allocate expenditure not assigned to retail between wholesale and direct purchases:

$$s_{cb}^W = (1 - s_{cb}^R)\left[\frac{W_{cb}}{g_c^W (E_{cb} - X_{cb}^R)}\right]_0^1, \qquad X_{cb}^W = s_{cb}^W E_{cb}, \tag{19}$$

$$s_{cb}^D = 1 - s_{cb}^R - s_{cb}^W, \qquad X_{cb}^D = s_{cb}^D E_{cb}. \tag{20}$$

Here X_{cb}^W denotes wholesale sales directly to buyer b, excluding sales to retailers. The three expenditures sum to the original purchasers' expenditure. For cells with positive expenditure and valid positive margin rates, the shares lie between zero and one. Missing or indeterminate implied shares are set to zero in the implementation; a missing channel match therefore does not generate purchases through that channel. The original accounting adjustments in the source tables, including negative entries, are retained rather than globally replaced by zero.

The ordering in (18)–(20) is an identifying assumption. Retail sales are inferred first, and wholesale margins are then attributed as far as possible to wholesale sales directly to the buyer. Only wholesale margins left over after this allocation are used to infer wholesale sales to retailers. Consequently, the construction does not assume that all retail merchandise passes through a wholesaler.

Reconstructing purchases for resale. Retailers' merchandise acquisition cost associated with cell (c, b) is their inferred sales less the observed retail margin:

$$C^R_{cb} = X^R_{cb} - R_{cb}. \tag{21}$$

Let $\widetilde{W}_{cb}$ denote the wholesale margin remaining after accounting for wholesale sales directly to b. The implementation calculates

$$\widetilde{W}_{cb} = W_{cb} - \left\lfloor g^W_c X^W_{cb} \right\rfloor, \qquad q_{cb} = \min\left\{1, \frac{\widetilde{W}_{cb}}{g^W_c C^R_{cb}}\right\}. \tag{22}$$

The floor operation reproduces the rounding convention in the code, in the monetary units of the margins table. An undefined or infinite residual share is set to zero. The residual share is capped above at one; the implementation does not impose an additional lower bound at this step. For regular positive cells, q_{cb} is the fraction of retailers' merchandise costs sourced from wholesalers. Retailers' purchases from wholesalers and directly from producers are, respectively,

$$H_{cb} = q_{cb} C^R_{cb}, \qquad D^R_{cb} = (1 - q_{cb}) C^R_{cb}. \tag{23}$$

Total wholesale sales associated with the transaction are $X^W_{cb} + H_{cb}$. Subtracting the observed wholesale margin gives wholesalers' purchases from the upstream commodity supplier:

$$C^W_{cb} = X^W_{cb} + H_{cb} - W_{cb}. \tag{24}$$

The procedure retains the reported dollar margins even when a channel share is capped. In such cells, the realized margin-to-sales ratio need not equal the Census proxy g^k_c.

These flows satisfy a useful accounting identity:

$$X_{cb}^{D} + D_{cb}^{R} + C_{cb}^{W} = E_{cb} - R_{cb} - W_{cb} = P_{cb} + T_{cb}. \tag{25}$$

Thus, inserting the trade stages preserves the buyer's expenditure and separates out the two observed trade margins. The adjustment does not separately reroute transportation margins: transportation costs remain embedded in the upstream commodity value in (25), while transportation industries remain in the network as ordinary industries. Accordingly, the reconstructed upstream merchandise flow is valued at producer value plus transportation costs, rather than at producer value alone.

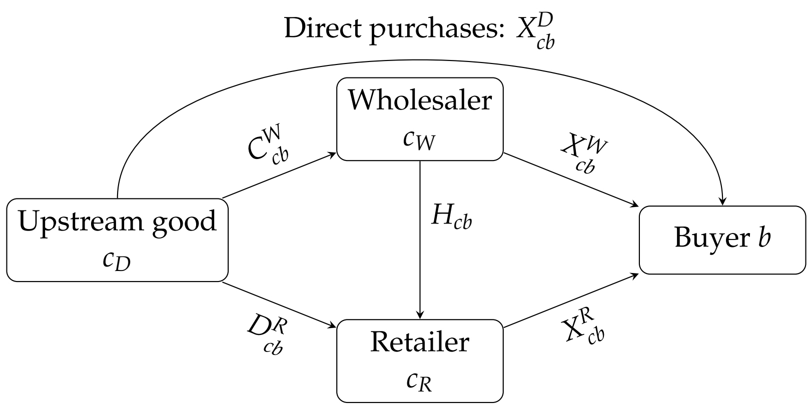

Figure B8: Reconstructed distribution channels for a commodity–buyer pair

Notes: Arrows indicate the direction of merchandise sales, with labels giving transaction values. The buyer can purchase directly, from a wholesaler, or from a retailer; the retailer can source directly or through a wholesaler. Wholesale sales less merchandise costs equal W_{cb}, and retail sales less merchandise costs equal R_{cb}. Operating inputs and value added supporting these margins are allocated separately and are omitted from the diagram.

Building the augmented use table. We retain a direct version of every original industry and commodity, and add commodity-specific wholesale and retail versions wherever the concordance identifies the corresponding channel. In the code these have suffixes `DI`, `WH`, and `RT`. The new wholesale and retail commodities represent the good resold at that stage, including its acquisition cost and the intermediary's margin. They are distinct from the original BEA commodities representing trade services alone.

For each original buyer, the augmented use table records X_{cb}^{D}, X_{cb}^{W}, and X_{cb}^{R} as purchases of the direct, wholesale, and retail versions of c. We then sum the reconstructed merchandise purchases over buyers: the retailer for c purchases $\sum_b D_{cb}^{R}$ of c_D and $\sum_b H_{cb}$ of c_W, and the wholesaler for c purchases $\sum_b C_{cb}^{W}$ of c_D. These merchandise requirements and the margin totals below are aggregated over the transactions in the margins file before selecting the model's industry and commodity lists. The model retains the original industries and personal consumption expenditures (PCE) as buyers; other final-use categories are not separate nodes in the final matrix.

Allocating trade operating inputs and value added. Merchandise purchases are only part of trade industries' costs. We also redistribute the inputs used to provide distribution services, including value added, from the original trade industries to the new commodity-specific intermediaries. This redistribution uses a common input mix within wholesale trade and a separate common input mix within retail trade.

To describe it, let B_a^k be total purchases of augmented commodity a by all original industries in trade channel k, after splitting their own purchases across distribution channels. Let $B^k = \sum_a B_a^k$, and let $M_c^W = \sum_b W_{cb}$ and $M_c^R = \sum_b R_{cb}$ be the margins assigned to the new wholesaler or retailer for commodity c. We allocate operating inputs according to

$$O_{a,c_k} = \frac{B_a^k}{B^k} M_c^k, \qquad k \in \{W, R\}. \tag{26}$$

Each new intermediary consequently receives operating inputs and value added totaling its assigned margin, in addition to its merchandise acquisition costs. For example, its total costs are $\sum_b C_{cb}^R + M_c^R = \sum_b X_{cb}^R$ for retail, and $\sum_b C_{cb}^W + M_c^W = \sum_b (X_{cb}^W + H_{cb})$ for wholesale. These identities apply to the constructed flows before restricting to retained commodity rows.

Not all activity of the original trade industries is captured by merchandise margins. Define $\alpha_k = \sum_{c \in C_k} M_c^k / B^k$, where C_k is the set of commodities with an added industry in channel k. We retain the fraction $1 - \alpha_k$ of each original trade industry's input purchases in that industry's direct version. The remaining fraction is redistributed through (26). This preserves aggregate operating-input purchases within each trade channel while retaining trade activity outside the allocated margins. In particular, the construction does not assign each commodity-specific intermediary the detailed input mix of its matched NAICS industry.

From augmented make and use tables to expenditure shares. The augmented make table preserves the original supply relationships among direct industries and direct commodities, including secondary production. Each added wholesale or retail commodity is supplied by its corresponding commodity-specific trade industry. Where an added commodity has no corresponding industry, the code assigns it to other durable goods merchant wholesalers (BEA 423A00); this is the fallback for scrap wholesale. Only commodity supply proportions enter the final calculation, so the level assigned to an added make-table entry is immaterial when it has a single supplier.

Let $\widetilde{U}_{ai}$ denote purchases of augmented commodity a by industry i, and $\widetilde{V}_{aj}$ its supply by industry j. Define

$$B_{ia} = \frac{\widetilde{U}_{ai}}{\sum_{a'} \widetilde{U}_{a'i}}, \qquad D_{aj} = \frac{\widetilde{V}_{aj}}{\sum_{j'} \widetilde{V}_{aj'}}. \tag{27}$$

We set $B_{ia} = 0$ for a buyer with zero total costs. The trade-adjusted input–output matrix is

$$\Omega_{ij} = \sum_a B_{ia} D_{aj}, \qquad \Omega = BD. \tag{28}$$

Thus, as in the main text, rows index purchasing industries and columns index supplying industries. The denominator is total recorded input expenditure, including the retained value-added components, rather than merchandise costs alone. Employee compensation and gross operating surplus are supplied by a common value-added node, which purchases no inputs. A PCE node purchases the household consumption bundle, so its row records consumption expenditure shares over supplying industries. Rows with nonzero total expenditure sum to one when all purchased commodities have an assigned supplier.

The resulting matrix makes the additional stages of distribution explicit in the cost network. An upstream cost change can reach the buyer directly, through wholesale trade, through retail trade, or through both stages in sequence. The pricing and inventory assumptions of the model then determine the timing of transmission along each of these reconstructed paths.