Retail inventories and inflation dynamics: The price margin channel*

Neil Mehrotra† Hyunseung Oh‡ Julio L. Ortiz§

September 15, 2026

Abstract

The initial increase in US inflation in 2021–22 was driven by large increases in retail price margins. We examine whether low inventories relative to sales contributed to this margin expansion. Using industry-level panel data and plausibly exogenous variation in supply conditions, we estimate the elasticity of retail margins with respect to finished-good inventories. We find a negative and statistically significant elasticity, indicating that lower inventories are associated with higher margins. We incorporate this inventory-margin channel into a New Keynesian Phillips curve (NKPC) framework. The augmented NKPC markedly improves empirical fit and helps account for two notable recent inflation episodes: the missing disinflation of 2009–11 and the COVID-era surge.

Keywords: Inflation; inventories; supply disruptions; Phillips curve.

JEL Classification Numbers: E31, E32, E22.

*We thank Chris Machol for excellent research assistance and George Alessandria, Olivier Coibion, Keshav Dogra, François de Soyres, Giuseppe Fiori, Aaron Flaaen, Joel Flynn, Nils Gornemann, Oleksiy Kryvtsov, Simon Mongey, Nitya Pandalai-Nayar, Mike Waugh, and Daniel Villar for helpful comments and suggestions. The views expressed are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of Minneapolis. All errors are our own. First draft: August 2025.

†Federal Reserve Bank of Minneapolis. E-mail: neil.mehrotra@mpls.frb.org

‡Federal Reserve Board. E-mail: hyunseung.oh@frb.gov

§Federal Reserve Board. E-mail: julio.l.ortiz@frb.gov

"Retail markups in a number of sectors have seen material increases in what could be described as a price-price spiral, whereby final prices have risen by more than the increases in input prices. The compression of these markups as supply constraints ease, inventories rise, and demand cools could contribute to disinflationary pressures."

—Lael Brainard (January 2023)[1]

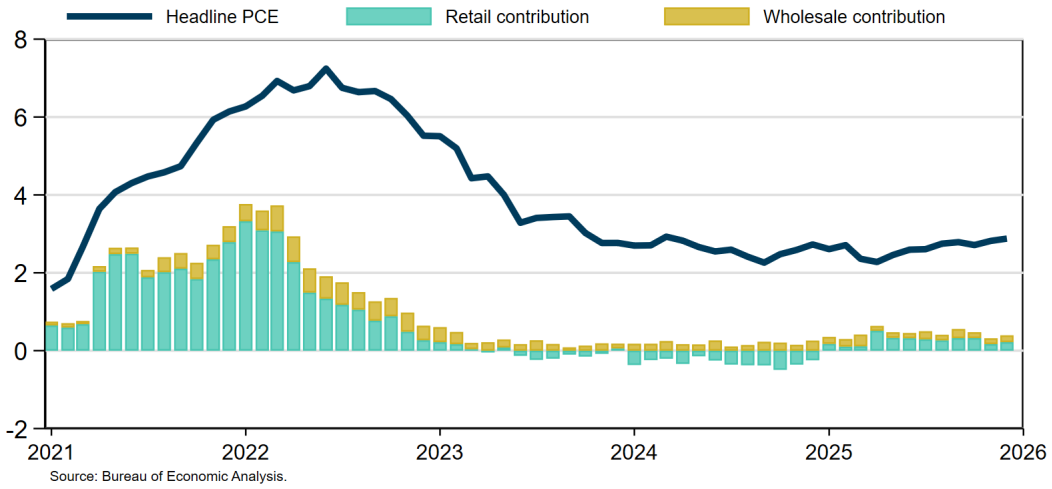
1 Introduction

The Phillips curve has long been a key pillar of the literature on inflation and the business cycle. Nevertheless, its empirical performance in accounting for inflation dynamics leaves considerable room for improvement. This limitation became especially apparent during the COVID-era inflation surge, which was initially concentrated in the goods sector and accompanied by sharp increases in retail and wholesale price margins, as measured by the corresponding PPI trade-margin indexes. Figure 1 shows the contribution of these margins—what Brainard referred to as the "price-price spiral"—to headline PCE inflation since 2021. During 2021–22, rising retail and wholesale price margins accounted for all of the increase in goods inflation and about half of the increase in overall inflation. This contribution is particularly striking because goods account for only about one-third of personal consumption expenditures.

Motivated by these facts, this paper studies how retailers' inventory positions shape price margins and what this relationship implies for inflation dynamics. The standard New Keynesian Phillips curve (NKPC) abstracts from inventory management and therefore cannot capture this inventory–margin channel. We provide industry-level evidence that lower inventory-to-sales ratios raise retail price margins, derive a microfounded NKPC with inventories, and show that the inventory-augmented NKPC markedly improves both its in-sample fit and its out-of-sample forecasting performance for U.S. inflation relative to standard specifications.

[1]Speech of Federal Reserve Vice Chair Brainard on January 19, 2023, available at https://www.federalreserve.gov/newsevents/speech/brainard20230119a.htm.

FIGURE 1: Trade contributions to PCE inflation



Notes: The figure shows the change in PPI retail and wholesale trade margins at the NAICS 3-digit level weighted by share of purchasers' value from the 2021 PCE bridge table

To establish the relevance of inventories for price dynamics, we must show that price margins rise when inventories fall relative to sales. We begin with a stylized model of a retail firm that jointly chooses prices and inventory stocks in the face of supply disruptions. We show that transitory shocks shift the firm's stocking policy without directly affecting its pricing decision, thereby enabling identification of the optimal pricing curve—how price margins vary with inventories—using instruments for inventory disruptions.

Empirically, we estimate the elasticity of retail price margins with respect to inventories using industry-level panel data spanning 2008q1 to 2026q1. Because pricing and stocking decisions are jointly determined, accurately estimating this elasticity requires an identification strategy. Motivated by our stylized model, we construct two shift-share variables that capture plausibly exogenous transitory supply shocks. We use these shift-share variables to instrument for changes in retailers' inventory-to-sales ratios. The first instrument reflects material supply shortages among upstream manufacturers from the Census Bureau's Quarterly Survey of Plant Capacity (QSPC). The second instrument reflects global shipping delays based on the Average Congestion Rate (ACR) of Bai et al. (2024). In both cases we regard supply shifts as occurring upstream, and map them to downstream retailers

using exposure shares based on a retail industry's supply use from a given manufacturing industry.

OLS estimates show that lower inventory-to-sales ratios are associated with higher retail price margins, even after controlling for cost variables. Instrumenting these ratios with upstream supply disruptions yields an even more negative elasticity, providing evidence that low inventory-to-sales ratios raise margins through state-dependent retail pricing.

To assess the macroeconomic implications, we embed retailer inventory management in a New Keynesian model and introduce deep habits at the level of individual retail goods, as in Ravn et al. (2006). Under deep habits, current sales build a customer base that supports future demand. When inventories are ample, retailers hold down margins to expand sales and build their customer base. When inventories are low relative to sales, they instead harvest their existing customer base by raising margins. The model yields two NKPC representations. The sales-based NKPC relates inflation to future marginal costs and the marginal benefit of sales but is difficult to estimate directly when markups vary. The stock-based NKPC instead relates inflation to observable inventory-to-sales ratios and expected marginal cost growth, providing a tractable specification for empirical analysis.

Estimating the stock-based NKPC using macroeconomic data, we find that incorporating the retail inventory-to-sales ratio markedly improves the model's empirical fit compared to standard NKPC specifications that only use current or expected marginal cost. The inventories specification comes closer to the model-implied specification for inflation, with the coefficient on lagged inflation becoming insignificant and the coefficient on expected inflation moving closer to one. We also establish superior forecasting performance of NKPC with inventories in forecasting one and two quarter-ahead inflation.

The stock-based NKPC helps account for two prominent recent inflation episodes—capturing both the missing disinflation of 2009–11 and the COVID-era inflation surge. Inventory-to-sales ratios were low in both periods, generating additional inflationary pressures that bring model-implied inflation closer to the observed data. These results highlight the role of inventory-induced movements in retail margins in shaping inflation dynamics.

Related literature. The COVID-era surge in inflation has sparked renewed interest in understanding inflation. Much of the literature seeks to only explain the COVID inflation

episode rather than offer a theory that can explain inflation dynamics over a longer time horizon. Theories of inflation due to goods-services reallocation, alternative measures of labor market tightness, convex cost curves, or energy prices have been prominent (for example, Ball et al. (2022), Bernanke and Blanchard (2025), Gagliardone and Gertler (2023), Benigno and Eggertsson (2023), Boehm and Pandalai-Nayar (2022)). A weakness of these labor market or cost-based based explanations is their inability to explain declines in the labor share and declines in real wages during the COVID inflation episode.

Given declines in the labor share and real wages, price markup shocks have been a candidate explanation for explaining inflation. Our paper offers a microfoundation for why price markups may be time-varying and an observable macroeconomic variable that captures time-varying markups. Smets and Wouters (2007) finds that most business cycle variation in inflation in their medium-scale estimated DSGE model is explained by price markup shocks. Mongey and Waugh (2025) provide a complementary explanation for price markup shocks during the COVID inflation episode from households trading up due to transfers. Flynn et al. (2026) shows how price markups rise in a model where firms commit to a certain quantity and face stockout risk.

A growing literature documents how production and transportation frictions contributed to rising costs across sectors (Alessandria et al., 2023; Boehm and Pandalai-Nayar, 2022; Comin et al., 2023; Ferrante et al., 2023; Ortiz, 2022). These studies emphasize upstream channels—such as global supply chain bottlenecks, convex production technologies, and sectoral reallocation—as key drivers of inflationary pressures. We build on this line of research by focusing on a downstream margin that has received comparatively less attention: the pricing behavior of retailers facing low inventories relative to sales. While prior work highlights how supply disruptions raise production costs, our contribution is to show that they can also amplify retail price margins when firms are constrained by limited inventory availability.

A paper closely related to our topic is Cavallo and Kryvtsov (2023), which uses online micro-level stockout data to study the inflationary effects of consumer product shortages during the pandemic. We complement their analysis by leveraging industry-level panel data, which offers two distinct advantages. First, our data encompasses a broader set of

retailers, including the motor vehicle dealer industry—a key contributor to goods inflation in the early stages of the pandemic. Second, we use retail gross margin data—the difference between the sales price and the acquisition cost—rather than retail prices, and we control for several industry-level cost measures, such as average hourly earnings. This allows us to analyze retailers' pricing behavior beyond what can be attributed to underlying cost pressures during this period.

The relationship between inventories and inflation has been explored in several studies, particularly in the context of business cycle models. Jung and Yun (2005), Lubik and Teo (2012), and Kryvtsov and Midrigan (2010, 2013) examine how inventory dynamics interact with pricing, often within frameworks that assume constant elasticity of substitution (CES) in demand. These papers typically find either a limited role for inventories in explaining inflation dynamics or a need for additional mechanisms—such as time-varying, counter-cyclical markups—to reconcile theory with the data, especially in response to monetary policy shocks. Our empirical analysis using industry-level data confirms that retail price margins respond in ways that are inconsistent with CES demand, supporting the need for richer pricing structures. Therefore, in contrast to this literature, our model explicitly incorporates standard features that generate time-varying price markups via deep-habit formation. We derive a representation of the NKPC in which the inventory-to-sales ratio serves as a sufficient statistic for retail markups. This approach allows us to empirically assess the role of time-varying markups in inflation dynamics without relying on direct and often noisy measures of price margins.

Our paper also contributes to the empirical Phillips curve literature, which has grappled with the declining explanatory power of traditional NKPC specifications in recent decades (Mavroeidis et al., 2014; Stock and Watson, 2020; Coibion and Gorodnichenko, 2025). A central debate in this literature concerns the inclusion of forward-looking variables that capture the information and beliefs shaping firms' pricing decisions. For example, Coibion and Gorodnichenko (2015) show that replacing professional inflation forecasts with household survey expectations can help resolve the 2009–11 missing disinflation puzzle, highlighting the value of expectation measures that better reflect firms' information set.[2] Reis (2023)

[2]Short-term inflation expectations jumped after Liberation Day but as of writing have not resulted in sharp rise in either headline or core inflation.

similarly argues for combining multiple indicators of inflation expectations to obtain early and more accurate signals of inflation dynamics. We find that the inventory-to-sales ratio is a superior forward-looking variable that encodes retailers' near-term beliefs and strategic pricing behavior. Incorporating this measure directly into the NKPC significantly improves its empirical fit, allowing it to account for key episodes such as the 2009–11 missing disinflation and the surge in goods inflation following the COVID-era supply disruptions.

The remainder of the paper is organized as follows. Section 2 presents a stylized retail firm model that motivates our identification strategy for estimating the inventory stock elasticity of the retail price margin. Section 3 implements this strategy using industry-level panel data. Section 4 extends the stylized model and derives a NKPC with inventories that is consistent with the empirically observed negative inventory elasticity of retail markups. Section 5 estimates the inventory-augmented Phillips curve to quantify the macroeconomic importance of retail pricing behavior, as reflected in the inventory-to-sales ratio. Section 6 concludes.

2 A stylized retail firm model

In this section, we present a stylized model of a retail firm to illustrate how temporary supply disruptions can be used to identify the firm's optimal pricing curve as a function of its inventory stock. The key insight is that while the firm's optimal *stocking* decision depends on both current and expected future costs of acquiring goods, its *pricing* decision depends only on the expected future cost. Therefore, changes in the current cost of acquiring goods—so long as they do not affect expectations of future costs—will shift the firm's stocking curve but leave its pricing curve unchanged. In particular, if supply disruptions are perceived to be transitory relative to other cost shocks, they can serve as valid instruments to trace out the firm's pricing curve. We first show this identification logic in a static one-period model, and then extend the insight to a dynamic two-period setting that makes explicit the role of intertemporal inventory decisions in separating pricing from stocking behavior.

2.1 One-period model

To build intuition about the retail firms' joint decisions over stocking and pricing, we consider a one-period model in which the firm purchases a stock of finished goods and sets a price at which to sell them. The firm simultaneously chooses the stock level and the selling price, taking as given the demand function and the unit cost of stocking. Demand for the good, denoted s, is given as $s = \boldsymbol{s}(a, p)$ where a is the stock and p is the price.[3] We assume the demand function satisfies $\partial \boldsymbol{s}/\partial a > 0$, $\partial^2 \boldsymbol{s}/\partial a^2 < 0$, and $\partial \boldsymbol{s}/\partial p < 0$. The first two conditions are standard in the inventory literature. For example, in stockout avoidance models, an additional unit of stock raises sales up to the point where stockouts are eliminated; in stock-elastic demand models, sales increase with an additional unit of stock but at a diminishing rate.

After sales, the firm may liquidate unsold inventory, $a - s$, incurring a cost. The total cost of stocking, net of liquidation, is given by:

$$(c + \xi)a - (1 - \gamma)c(a - s),$$

where c is the unit cost of stocking, ξ captures an additional shadow cost faced by the retailer (e.g., supply disruptions not embedded in c), and γ governs the liquidation cost. The parameter ξ is a stand-in for the difficulty in stocking by the retail firm not captured in the unit cost of stocking c.

The firm's profit is given by:

$$\phi = p\boldsymbol{s}(a, p) - (c + \xi)a + (1 - \gamma)c(a - \boldsymbol{s}(a, p)).$$

The firm chooses a and p to maximize this profit.

The first order condition with respect to a (stocking) is:

$$p\frac{\partial \boldsymbol{s}}{\partial a} - (c + \xi) + (1 - \gamma)c\left(1 - \frac{\partial \boldsymbol{s}}{\partial a}\right) = 0,$$

[3] For simplicity, we normalize aggregate demand to 1.

which simplifies to:

$$(p - (1-\gamma)c)\frac{\partial \boldsymbol{s}}{\partial a} = \gamma c + \xi. \tag{1}$$

Absent both a shadow cost of stocking ($\xi = 0$) and a liquidation cost ($\gamma = 0$), the right-hand side of (1) becomes zero. In that case, as long as $p > (1-\gamma)c$, the firm would optimally stock an infinite quantity ($a \to \infty$), driving $\partial \boldsymbol{s}/\partial a \to 0$. Therefore, at least one of ξ or γ must be strictly positive to yield a finite interior solution. The resulting stocking curve determines the optimal level of the stock as a function of the price, conditional on the cost parameters c, ξ, and γ.

The first order condition with respect to price is:

$$s = -(p - (1-\gamma)c)\frac{\partial \boldsymbol{s}}{\partial p}. \tag{2}$$

This pricing equation pins down the optimal price as a function of the level of the stock, taking c and γ as given. For example, suppose the demand function has constant price elasticity:

$$\boldsymbol{s}(a, p) = u(a) \times p^{-\eta},$$

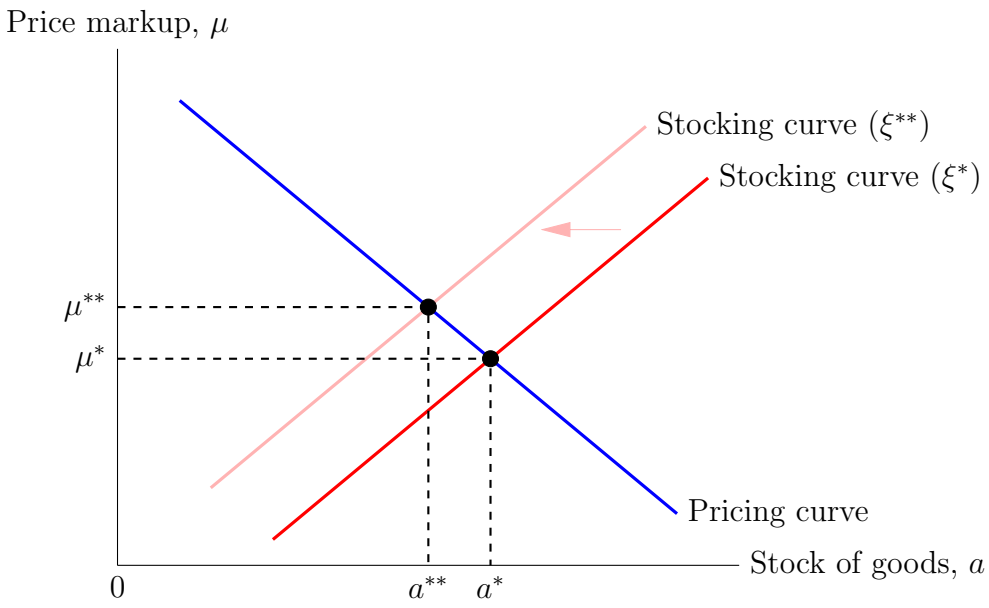
where $u(a)$ is increasing in a and $\eta > 1$. Then the price markup is constant and given by:

$$\mu \equiv \frac{p}{(1-\gamma)c} = \frac{\eta}{\eta - 1},$$

which is the standard result under CES demand. In this case, the pricing curve is flat, and changes in the stock have no effect on the markup.

Taken together, the optimal stocking and pricing curves jointly determine the firm's equilibrium stock and price (or markup). A key feature of this simple model is that the shadow cost ξ appears only in the stocking condition—not in the pricing condition. As a result, shifts in ξ move the stocking curve without affecting the pricing curve, allowing us to trace out the firm's pricing behavior. Figure 2 illustrates this identification strategy. When $\xi = \xi^*$, the firm's equilibrium stock and markup are (a^*, μ^*). When ξ increases to

FIGURE 2: Illustration of a Retailer's Equilibrium



Notes: This plot is illustrative. A finished-goods inventory model with a negatively sloped stocking curve would yield qualitatively similar implications.

$\xi^{**} > \xi^{*}$, the equilibrium shifts to (a^{**}, μ^{**}). Supply disruptions that affect ξ not captured by c therefore serve as valid instruments to recover the pricing curve. If the pricing curve is flat—as in the case of constant price elasticity—then changes in ξ would shift the stock without altering the markup. In contrast, if the pricing curve is downward sloping, as shown in the figure, then a higher ξ leads to lower stock and a higher markup. Verifying this inverse relationship between inventory and price margins is the focus of the empirical analysis that follows.

2.2 Two-period model

We now extend the logic of the one-period model to a dynamic, two-period setting in which inventories can be carried across periods. This extension verifies that the identification strategy remains valid when the firm makes intertemporal stocking decisions.

Assume the retail firm lives for two periods, facing the same demand function in each period as in the one-period model, $s = \boldsymbol{s}(a, p)$. In the first period, the firm purchases a

stock of goods a_1, sets a price p_1, and sells $s_1 = \boldsymbol{s}(a_1, p_1)$. The remaining inventory $a_1 - s_1$ is carried over to the second period, subject to a depreciation rate δ. In the second period, the firm purchases n additional units, bringing the total available stock to:

$$a_2 = (1 - \delta)(a_1 - s_1) + n.$$

It then sets price p_2 and sells $s_2 = \boldsymbol{s}(a_2, p_2)$ units. As in the one-period model, leftover inventory at the end of the second period is liquidated at cost, governed by the parameter γ.

The firm maximizes the sum of discounted profits:

$$\pi = p_1 s_1 - (c_1 + \xi_1)a_1 + \beta \mathbb{E}_1 \left[p_2 s_2 - c_2(a_2 - (1 - \delta)(a_1 - s_1)) + (1 - \gamma)c_2(a_2 - s_2) \right],$$

where β is the discount factor, and $\mathbb{E}_1$ denotes expectations conditional on information available in period 1. We assume that the second-period supply disruption ξ_2 is unanticipated, so $\mathbb{E}_1 \xi_2 = 0$.

The first-order condition with respect to a_1 (initial stocking) is:

$$(p_1 - \beta(1 - \delta)\mathbb{E}_1 c_2) \frac{\partial \boldsymbol{s}(a_1, p_1)}{\partial a_1} = c_1 + \xi_1 - \beta(1 - \delta)\mathbb{E}_1 c_2,$$

which mirrors the one-period stocking condition in equation (1).

Similarly, the first-order condition with respect to p_1 is:

$$s_1 = -(p_1 - \beta(1 - \delta)\mathbb{E}_1 c_2) \frac{\partial \boldsymbol{s}(a_1, p_1)}{\partial p_1},$$

which corresponds to the pricing condition in equation (2).

As in the one-period case, the shadow cost ξ_1 enters only the stocking equation and not the pricing equation. This separation confirms that temporary supply disruptions—captured by changes in ξ_1—remain valid instruments for identifying the slope of the pricing curve in a dynamic environment.

2.3 Empirical proxy for the stock of goods: inventory-to-sales ratio

In the next section, we turn to the data and implement this identification strategy using industry-level panel data to estimate the elasticity of retail price margins with respect to the stock of goods. In empirical work, we proxy the stock using the end-of-period inventories. Indeed, our stylized model implies a tight relationship between the stock of goods, a, and end-of-period inventories, inv, as given by:

$$inv = a - s.$$

Letting $is = inv/s$ denote the inventory-to-sales ratio and $\overline{IS}$ its long-run average, we obtain the following log-linear approximation:

$$\hat{a} = \hat{s} + \frac{\overline{IS}}{\overline{IS} + 1}\widehat{is},$$

where hatted variables denote log-deviations from their respective averages. This expression implies that changes in the stock-sales ratio and the inventory-to-sales ratio are proportionally related. As such, conditional on sales, fluctuations in the stock of goods are directly reflected in movements of the inventory-to-sales ratio. Accordingly, we use the inventory-to-sales ratio as an empirical proxy for the stock of goods in our model—controlling for sales—as it conveys equivalent information about stock dynamics and is both directly observable and widely used in the empirical literature.

3 Inventories and prices at the industry level

In this section, we use industry-level panel data to estimate how the retail and wholesale margin price reacts to changes in the inventory stock.[4]

[4]Since both retailers and wholesalers primarily manage finished-good inventories and typically do not transform the products as manufacturers do, we also include data from the wholesale trade industry for completeness of our analysis.

3.1 Industry-level data

We assemble quarterly data for eight industries across the retail and wholesale trade sectors, which we collectively refer to as the retail sector for simplicity, spanning 2008q1 to 2026q1. From the Bureau of Labor Statistics (BLS), we collect industry-specific Producer Price Index (PPI) data, average hourly earnings, and total hours worked. Note that in the trade sectors, the PPI reflects gross margins—that is, the difference between the sales price and the acquisition cost of goods sold by retailers. From the Bureau of Economic Analysis (BEA), we incorporate information on real inventories and real sales, which are sourced from the National Income and Product Accounts tables. We merge the BLS and BEA data at the three-digit North American Industry Classification System (NAICS) level, which represents the most granular classification available across these data sets. Further details on data construction are provided in the Appendix.

3.2 Industry-level regression framework

To formally examine the relationship between inventory holdings and changes in the industry gross margin, we estimate the following regression,

$$\log(PPI_{i,t}) = \beta \log(IS_{i,t}) + X_{i,t}\gamma + \varepsilon_{i,t}. \tag{3}$$

The outcome variable is the gross margin, measured as the log of the PPI. The coefficient of interest, β, captures the elasticity of the gross price margin with respect to inventory-to-sales. All control variables are contained in $X_{i,t}$.

Because stocking and pricing decisions are jointly determined, estimating equation (3) via ordinary least squares (OLS) would yield a biased estimated of β. In particular, as shown in the stylized model, unobserved persistent or anticipated shocks can shift both the pricing and stocking curves. Shifts in the pricing curve induce a positive correlation between inventory stocks and price margins. As a result, we expect the OLS estimate of β to be biased upward.

To estimate equation (3), we therefore adopt an identification strategy which exploits plausibly exogenous and transitory supply-side variation which we discuss in turn.

3.3 Identification with shift-share instruments

We seek to establish causality by instrumenting for the retailer's inventory-to-sales ratio. Motivated by our stylized model, we construct two instruments that exploit plausibly exogenous variation in upstream suppliers' ability to deliver finished-goods inventories to downstream retailers. Delivery disruptions can arise from (i) production problems—including suppliers' difficulties sourcing raw materials—and (ii) delays in shipping finished goods. Our instruments capture both channels by measuring retailers' exposure to shocks that affect upstream manufacturers' capacity to produce and deliver goods.

Each instrument takes the form of a shift-share variable (Bartik, 1991). Identification comes from plausibly exogenous shifts at the manufacturing level that reflect supply-side disruptions. Conditional on the specified controls, we assume (i) these shifts are orthogonal to unobserved shocks in downstream retail outcomes, and (ii) the mapped shifts affect retail price margins only through inventories. Fixing the exposure shares *ex ante* makes them predetermined but does not ensure exogeneity, since shares can co-move with persistent demand or pricing dynamics. Following Borusyak et al. (2022), we therefore place identification on the exogeneity of the shifts: under their quasi-experimental shift-share framework, a share-weighted average of shifts is as good as random when the shifts are exogenous, so instrument validity hinges on the shifts rather than on the exogeneity of the exposure shares. We next define the two shift-share instruments used in our analysis.

3.3.1 Reported material shortages from the QSPC

Production delays in manufacturing typically show up in capacity utilization data. Indeed, capacity utilization have taken center stage in recent work on inflationary pressures arising from production bottlenecks (e.g., Boehm and Pandalai-Nayar, 2022; Comin et al., 2023). Yet utilization is endogenous, reflecting both demand and supply forces. Unlike Boehm and Pandalai-Nayar (2022), who emphasize demand-driven movements in utilization and their interaction with supply conditions, our goal is to isolate upstream supply disturbances that shift utilization orthogonally to contemporaneous retail demand (conditional on controls).

Our first instrument exploits the establishment-level reason codes in the QSPC—the survey underlying official capacity utilization statistics. A key advantage of the QSPC is that

it records *why* establishments operate below capacity, allowing us to isolate supply-driven variation in utilization. For each three-digit NAICS manufacturing industry, we collect the share of underutilizing establishments that cite "insufficient supply of materials."[5]

Firm-reported material shortages have been shown to track supply-driven disturbances in production (Balleer and Noeller, 2024; Braun et al., 2024). Building on this evidence, we take the raw shortage shares and, industry-by-industry, residualize them on their own lag, measures of contemporaneous aggregate demand and industry-specific demand, and quarter fixed effects for seasonality. The industry-specific demand proxy is the common component of real value added and output price deflators (extracted separately by industry) that loads positively on both. This procedure removes predictable demand movements and serial correlation, leaving an innovation that reflects unexpected, supply-driven disruptions.

Formally, our QSPC shift-share instrument is defined as:

$$Z_{i,t}^{QSPC} = \sum_{j} s_{ij,t_0} \times g_{j,t}^{QSPC} \tag{4}$$

The shift variable, g_{jt}^{QSPC}, is defined as the residualized share of plants in upstream industry j reporting material shortages (as described above). Meanwhile, the share variable, $s_{ij,t_0} \equiv \omega_{ij,t_0}/\omega_{i,t_0}$, is defined as the retailer i's exposure to a given manufacturing industry j in period t_0. The numerator, ω_{ij,t_0}, is the value of inputs purchased by retail industry i from manufacturing industry j in the base year t_0.[6] The denominator, $\omega_{i,t_0} = \sum_j \omega_{ij,t_0}$, represents the total value of inputs (manufacturing and non-manufacturing) used by retailer i. The term s_{ij,t_0} thus reflects the input share sourced from manufacturing industry j within retailer i's total input bundle.

3.3.2 Bai et al. (2024)'s supply chain disruption index

We next consider an instrument for global supply disruptions based on the ACR index of Bai et al. (2024), available from January 2016 to January 2026. The ACR index measures

[5] See https://www.census.gov/programs-surveys/QSPC.html for details. The relevant ratio appears in Table 3 of the published survey results.

[6] We calibrate ω_{ij,t_0} using data from the BEA Industry Input-Output Use Table, with $t_0 = 2007$, which is before the start of our sample period.

global port congestion and is regarded as an exogenous indicator of shipping disruptions. Our ACR shift-share instrument is defined as:

$$Z_{i,t}^{ACR} = \sum_j \left(s_{ij,t_0} \times \frac{m_{j,t_0}}{m_{t_0}} \right) \times g_{t-1}^{ACR}, \tag{5}$$

The shift variable, g_{t-1}^{ACR}, is defined as the residual from a regression of the ACR index on its own lag.[7] Because the ACR shift variable is a time series, we define our share variable as a composite measure of exposure where s_{ij,t_0} denotes the retailers' exposure from the BEA 2007 Use Table, as before, and $m_{j,t_0}/m_{t_0}$ denotes the manufacturers' import share in 2007. Thus, the share variable is a constant exposure weight mapping exogenous global supply disruptions to manufacturers according to import exposure, and to retailers based on the use shares.[8]

3.4 Regression results

We estimate equation (3) via OLS as well as with our instrumental variables, using the aforementioned industry-level data across retailers. The control variables specified in $X_{i,t}$ consist of the log of the following industry-level variables: real sales, average hourly earnings, total hours worked, and an exposure-weighted sum of upstream manufacturing suppliers' PPI. The upstream PPI for each retail industry is defined as:

$$PPI_{i,t}^{\text{upstream}} = \sum_j s_{ij,t_0} \times PPI_{j,t}, \tag{6}$$

where $PPI_{j,t}$ denotes the PPI of manufacturing industry j in period t. We also specify a range of fixed effects in $X_{i,t}$. First, we specify industry fixed effects to control for time-invariant unobserved retailer characteristics as well as the time-invariant component of the shifts. In addition, we specify industry-specific linear time trends to absorb gradual, deterministic retail industry-specific trends over time. Finally, we specify time fixed effects.

[7] We lag the shift variable by one quarter, specifying its $t-1$ value, to account for possible delays in the transmission of global supply shocks to retailers.

[8] Relative to our QSPC instrument, the ACR instrument features only one shift in each period and leverages heterogeneous exposure to the shift such that $Z_{i,t}^{ACR}$ varies across retailers and time. In this case, the exogeneity of g_t^{ACR} is sufficient for consistent estimation.

TABLE 1: Industry-level regression results

Dependent variable: Gross price margin						
Estimation method	OLS	2SLS: QSPC			2SLS: ACR	2SLS: QSPC, ACR
		Baseline	Pre-COVID	Transitory		
IS ratio	-0.058	-1.821***	-0.760**	-1.066***	-1.196***	-1.064***
	(0.092)	(0.558)	(0.356)	(0.384)	(0.259)	(0.203)
First-stage F-stat		16.35	3.051	14.85	6.39	17.29
AR Wald F (p-value)		3.1E-5	0.038	7.1E-5	0.015	1.7E-9
J-test p-value						0.525
Observations	582	582	384	582	318	318
Sample	2008q1-2026q1	2008q1–2026q1	2008q1–2019q4	2008q1–2026q1	2016q2-2026q1	2016q2-2026q1

Notes: Driscoll–Kraay standard errors are reported in parentheses. Instruments for the IV columns are: reported material supply shortages from the Quarterly Survey of Plant Capacity (QSPC) and Average Congestion Rate (ACR). * p<0.10, ** p<0.05, *** p<0.01.

For the QSPC instrumental variable regression, however, we specify an interaction of the sum of exposure shares with time fixed effects.[9]

3.4.1 OLS estimation

We first estimate equation (3) using OLS. The results are reported in the OLS column of Table 1. The OLS estimate in suggests that changes in the inventory-to-sales ratio are negatively related to changes in the retailer's gross price margin, however, the point estimate is not statistically significant.

3.4.2 Two-stage least squares (2SLS) estimation

We next implement instrumental variables (IV) estimation using 2SLS, introducing the instruments first individually and then jointly. The corresponding results are reported in Table 1. When instrumenting for retail inventories using the QSPC IV, we estimate an elasticity of -1.82.

[9]In our context, the exposure shares do not sum to one as we measure a retail industry's exposure based on its supply use from only manufacturers (ω_{ij,t_0}) even though its total supply use (ω_{i,t_0}) includes non-manufacturers. Without controlling for sum of share-by-time fixed effects, our QSPC instrument would be correlated with with sum of share which would render our QSPC instrument invalid.

One concern with our QSPC IV might be that it could reflect persistent rather than transitory upstream supply shifts, particularly around the COVID episode. If so, then the instrument may violate the exclusion restriction since persistent upstream supply shifts can plausibly shift expected future marginal costs. We address this concern in a few of ways. First, we note that the inclusion of time fixed effects in our regression allows us to control for persistent aggregate shocks. Second, with regard to COVID, we note that pandemic-related supply disruptions were widely perceived to be transitory in nature, at least initially. To verify that our results are not driven by the COVID episode, the third column of Table 1 reports the estimated elasticity using data through 2019q4. Over this pre-COVID sample, we estimate a smaller though statistically significant elasticity of -0.68.

We next address this concern a different way by making use of the full sample and decomposing the QSPC variable at the manufacturing-level, into permanent and transitory components (Beveridge and Nelson, 1981). After doing so, we define the QSPC shift variable, $g_{j,t}^{QSPC}$, as the transitory component. We provide further details in Appendix B.3. The estimated elasticity when using the transitory-QSPC IV is -1.07.

Next, we estimate the regression by specifying the ACR-based IV. The results are reported in the sixth column of Table 1. Under the ACR-based IV, the estimated elasticity is approximately -1.20. Finally, we use both IVs to jointly instrument for retail inventory-to-sales.[10] The results are reported in the final column. When specifying both IVs, we estimate an elasticity of -0.98. A test of overidentifying restrictions yields a p-value of 0.31, failing to reject the null that our model is misspecified and providing additional support to the validity of our estimates.

Overall, the estimated elasticities are statistically significant and range from -0.76 to -1.82. Furthermore, the first-stage F-statistics suggest that the instruments are generally strong and in instances in which our instruments appear to be weak, inference remains supported by the Anderson-Rubin Wald test, which is robust to weak instruments.

Relative to OLS, the IV estimates are more negative, indicating that the OLS estimates are biased upward. This pattern is consistent with the intuition described above and with the mechanism illustrated in the stylized model. In particular, an upward shift in the pricing

[10] We specify the version of the QSPC IV that uses the transitory component of the QSPC shift variable.

TABLE 2: Industry-level regression results including other input cost controls

Dependent variable: Gross price margin					
Estimation method	OLS	2SLS: QSPC		2SLS: ACR	2SLS: QSPC, ACR
		Baseline	Transitory		
IS ratio	-0.728***	-3.547***	-3.113**	1.373	-0.870***
	(0.246)	(1.298)	(1.458)	(2.518)	(0.241)
First-stage F-stat		10.23	8.604	0.468	59.85
AR Wald F (p-value)		0.003	0.007	0.446	6.9E-5
J-test p-value					0.158
Observations	106	106	106	106	106
Sample	2018q4-2026q1	2018q4–2026q1	2018q4–2019q4	2018q4–2026q1	2018q4–2026q1

Notes: Driscoll–Kraay standard errors are reported in parentheses. Instruments for the IV columns are: reported material supply shortages from the Quarterly Survey of Plant Capacity (QSPC) and Average Congestion Rate (ACR). * p<0.10, ** p<0.05, *** p<0.01.

curve results in higher margins and greater inventory stocks. Because OLS does not account for these unobserved shifts in the pricing curve, it induces a positive correlation between the inventory-to-sales ratio and gross margins, thereby biasing the OLS estimates upward.

3.4.3 Robustness

Retailers may face additional costs beyond the labor and acquisition costs that we explicitly control for. If these costs are omitted, we cannot rule out that the estimated increase in gross margins reflects higher unobserved input costs. To address this concern, we estimate a version of our regressions that includes controls for other non-labor and non-capital costs, drawn from the BLS satellite series of *net inputs to industry price indexes*. These data are available only from December 2018 for NAICS 441, 445, and 452, and from December 2020 for the remaining retail and wholesale industries.

Because these series are available only over a shorter horizon, the sample size is reduced and the estimates are correspondingly noisier. Nonetheless, as reported in Table 2, the results remain broadly consistent with our baseline findings, indicating that our main conclusions are robust to the inclusion of these additional cost controls.[11]

[11]The estimates based on the ACR IV are no longer significant. This reflects low power due to the shorter

We report further robustness checks in the Appendix. These include first-stage regression results and a set of balance tests for the QSPC instrument at both at the manufacturing and retail levels. We also redefine the QSPC instrument to address potential measurement concerns.

3.5 Interpretation of the pricing response

Our findings suggest that retail price margins tend to rise in response to exogenous supply disruptions, even after accounting for a broad set of controls, including measures of input costs. In the context of the NKPC literature, this evidence aligns with a strand of research emphasizing the role of time-varying retail price markups in shaping inflation dynamics. The next section builds on this perspective by further exploring its implications.

That said, we acknowledge that alternative interpretations of our results are possible and warrant further investigation. First, given our reliance on industry-level data, we are unable to observe the composition of retail goods. Retailers may respond to supply disruptions by shifting toward higher-margin products, which would mechanically raise measured price margins without requiring active price increases. Second, we cannot rule out unobserved cost pressures. In particular, absent data on retailers' expectations of future marginal costs, it remains possible that firms raise prices to maintain rather than expand markups. Third, we remain agnostic on manufacturer markups and do not study their interaction with retailer markups—a topic examined in detail by Alvarez-Blaser et al. (2025) using data from a large global manufacturer. These alternative mechanisms highlight the need for richer data to fully disentangle pricing responses, which we leave for future work.

4 The NKPC with inventories

Our empirical investigation reveals a significantly positive response in the retail price margin to supply disruptions that adversely affect the stocking behavior of the retail industry. Building on this relationship, we examine whether the systematic link between the dynamics of the retail inventory-to-sales ratio and the retail price markup can help improve the

sample and is not due to the inclusion of the non labor-non capital input cost control variable.

empirical performance of the NKPC.

To address this question, we extend the stylized model introduced in Section 2. Specifically, we incorporate inventory management into the retail sector and generate state-dependent retail markups using the deep-habits framework of Ravn et al. (2006). These mechanisms generate time-varying price margins that align with our empirical findings and are widely utilized in the literature on inflation and business cycles.

Subsequently, we derive two representations of the linearized NKPC with inventories. We demonstrate that incorporating the retail inventory-to-sales ratio as an observable variable in the NKPC effectively captures time-varying price markups, which are otherwise difficult to measure directly.

4.1 Demand function

We first derive the retailers' demand function. To allow desired retail markups to vary over time even under flexible prices, we introduce deep habits at the level of individual retail varieties, following Ravn et al. (2006). Aggregate habit-adjusted consumption is given by

$$C_t = \left[\int_0^1 v_{j,t}^{1/\eta} \left(s_{j,t} - \theta s_{j,t-1} \right)^{\frac{\eta-1}{\eta}} dj \right]^{\frac{\eta}{\eta-1}},$$

where $s_{j,t}$ denotes real sales of variety $j \in [0,1]$, and $v_{j,t} > 0$ is a variety-specific demand shifter. The parameter $\theta \in [0,1]$ governs the strength of deep habits, while $\eta > 1$ is the elasticity of substitution across habit-adjusted varieties.

Expenditure minimization yields demand for variety j:

$$s_{j,t} = v_{j,t} C_t \left(\frac{p_{j,t}}{P_t} \right)^{-\eta} + \theta s_{j,t-1}, \tag{7}$$

where the price index is

$$P_t = \left[\int_0^1 v_{j,t} p_{j,t}^{1-\eta} dj \right]^{\frac{1}{1-\eta}}.$$

The price elasticity of demand faced by retailer j is therefore

$$-\frac{\partial s_{j,t}}{\partial p_{j,t}}\frac{p_{j,t}}{s_{j,t}} = \eta\frac{s_{j,t} - \theta s_{j,t-1}}{s_{j,t}}.$$

When $\theta = 0$, demand reduces to the standard CES case with constant elasticity η. With deep habits, however, the elasticity depends on current sales relative to the inherited habit stock. Moreover, because current sales become part of next period's habit stock, a retailer's current price affects its future demand. Deep habits therefore generate state-dependent desired markups even when prices are fully flexible.

4.2 The retail firm problem

The monopolistically competitive retail firm j in period t purchases $y_{j,t}$ quantity of manufacturing goods at a unit price Q_t. It differentiates the good which is then stocked and sold to households at a price $p_{j,t}$. The stock of retail goods, $a_{j,t}$, is the sum of the undepreciated inventories from the previous period and the goods purchased in the current period:

$$a_{j,t} = (1 - \delta)inv_{j,t-1} + y_{j,t}, \tag{8}$$

where $inv_{j,t-1}$ is the end-of-period inventories in period $t-1$ and δ is the inventory depreciation rate. In turn, the end-of-period inventories evolve as follows:

$$inv_{j,t} = a_{j,t} - s_{j,t}. \tag{9}$$

Dividing each side by $s_{j,t}$, we obtain a relationship between the inventory-to-sales ratio and the stock-sales ratio. Since they differ only by a constant, the two ratios move together. The inventory-to-sales ratio is defined as:

$$is_{j,t} = \frac{inv_{j,t}}{s_{j,t}}.$$

Demand for the retail good is given in (7). Following Bils and Kahn (2000) and Jung

and Yun (2005), we assume that demand is also elastic to the stock of goods for sale:

$$v_{j,t} = \left(\frac{a_{j,t}}{A_t}\right)^{\zeta}, \tag{10}$$

where ζ is the elasticity of habit-adjusted sales, $s_{j,t} - \theta s_{j,t-1}$, with respect to relative stock.[12] The aggregate stock is defined as

$$A_t = \left(\int_0^1 a_{j,t}^{\zeta} dj\right)^{\frac{1}{\zeta}}.$$

Combining (7) with (10) clarifies the role of deep habits. A decline in relative stock reduces the habit-adjusted component of demand, $s_{j,t} - \theta s_{j,t-1}$, while leaving demand inherited from past sales unchanged. It therefore reduces the price-sensitive share of current sales, lowers the price elasticity of demand, and raises the retailer's desired margin. When $\theta = 0$, inventories only shift the level of demand and the elasticity remains constant at η, so the pricing curve is flat.

Nominal price rigidity is introduced *a la* Rotemberg (1982). In detail, the retail firm faces a quadratic price adjustment cost in the following form:

$$\frac{\nu_p}{2}\left(\frac{p_{j,t}}{p_{j,t-1}} - 1\right)^2 P_t S_t, \tag{11}$$

where ν_p represents the degree of nominal price rigidity and $S_t = \int_0^1 s_{j,t} dj$ is aggregate sales.

The retail firm maximizes the expected discounted sum of profits:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_{0,t} \Phi_{j,t},$$

[12] We do not employ a stockout avoidance model, as the presence of deep habits introduces rich dynamics that make it difficult to track the relevant idiosyncratic state variables. As noted by Crouzet and Oh (2016) in the context of news shocks and by Kryvtsov and Midrigan (2013) for monetary policy shocks, the stock-elastic demand specification and the stockout avoidance model exhibit similar equilibrium implications for inventory dynamics, particularly with respect to the endogeneity of inventory-to-sales ratios.

where the period-t profit of the retailer is

$$\Phi_{j,t} = p_{j,t}s_{j,t} - Q_t y_{j,t} - \frac{\nu_p}{2}\left(\frac{p_{j,t}}{p_{j,t-1}} - 1\right)^2 P_t S_t.$$

The retailer is subject to the above conditions (7)-(10). The full optimality conditions are presented in the Appendix.

4.3 Deriving the NKPC with inventories

In standard New Keynesian models without inventories, every unit sold must be produced or acquired contemporaneously. The firm's sales condition therefore links the marginal value of sales directly to current real marginal cost. Combining this condition with the firm's pricing decision yields the standard NKPC.

With inventories, this direct link breaks down. Retailers can satisfy current demand by drawing down existing stocks rather than by purchasing additional goods contemporaneously. A current sale therefore affects the inventory stock carried forward and is linked to the expected cost of replenishing inventories in the future. This additional intertemporal margin gives rise to alternative representations of the NKPC.

4.3.1 Equilibrium pricing condition

The log-linearized optimal pricing condition in symmetric equilibrium is

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} - \frac{1}{\nu_p}(\hat{\kappa}_t + \hat{c}_t - \hat{s}_t), \tag{12}$$

where hatted values denote log-deviations from the noninflationary steady state, and κ_t is the real marginal benefit of an additional unit of sales in period t—the Lagrange multiplier on the demand constraint (7). Current inflation therefore depends on expected future inflation, the real marginal benefit of sales, and the gap between habit-adjusted consumption and sales. This condition also applies in the model without inventories. The key difference between models with and without inventories lies in how demand is met, as summarized by κ_t. In the model with inventories, κ_t is not tied one-for-one to contemporaneous real

marginal cost because retailers can satisfy demand by drawing down existing inventories rather than immediately purchasing additional goods to restock.

4.3.2 Sales-based NKPC

The log-linearized optimal sales condition is

$$\hat{\kappa}_t = \theta\beta\mathbb{E}_t(\hat{r}_{t,t+1} + \hat{\kappa}_{t+1}) - [\eta(1-\theta) - 1 + \theta\beta]\,\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}), \tag{13}$$

where $r_{t,t+1}$ is the real stochastic discount factor between period t and $t+1$ and mc_t the real marginal cost of replenishing a retail good in period t, equivalently the Lagrange multiplier on the law of motion for the inventory stock (8).

To interpret this condition, it is helpful to begin with the benchmark New Keynesian model without inventories and with CES preferences across varieties ($\theta = 0$). The sales condition then reduces to

$$\hat{\kappa}_t = -(\eta - 1)\widehat{mc}_t.$$

Because the markup without inventories is the inverse of real marginal cost, this relation shows that the real marginal benefit of sales is proportional to the markup. Substituting this expression into (12) yields the familiar Rotemberg-type NKPC:[13]

$$\hat{\pi}_t = \beta\mathbb{E}_t\hat{\pi}_{t+1} + \left(\frac{\eta - 1}{\nu_p}\right)\widehat{mc}_t.$$

In this benchmark, the sales and stocking conditions coincide, since all sales must be met by contemporaneous purchases of goods. This equivalence implies a one-for-one mapping between marginal cost and the marginal benefit of sales, which is why the standard NKPC can be written directly in terms of real marginal cost.

With inventories, however, this one-for-one mapping breaks down. Additional sales need not be met by contemporaneous purchases. Instead, the retailer recognizes that sales can be satisfied by drawing down inventories, so the optimal sales decision depends on the

[13] Without habits ($\theta = 0$), consumption equals sales so that $\hat{c}_t = \hat{s}_t$.

expected discounted future marginal cost rather than on the current marginal cost. When $\theta > 0$, habit formation adds an additional intertemporal channel: the expected discounted marginal benefit of future sales also matters. Combining (12) and (13) yields the following sales-based NKPC.

Proposition 1 (Sales-based NKPC)**.** *The sales-based NKPC is derived by combining* (12) *and* (13)*:*

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + B_1 \mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}) - B_2 \mathbb{E}_t(\hat{r}_{t,t+1} + \hat{\kappa}_{t+1}) + B_3(\hat{s}_t - \hat{c}_t),$$

where

$$B_1 = \frac{1}{\nu_p}\left[\eta(1-\theta) - 1 + \theta\beta\right], \quad B_2 = \frac{\theta\beta}{\nu_p}, \quad B_3 = \frac{1}{\nu_p}.$$

In the CES case ($\theta = 0$), the sales-based NKPC simplifies to

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \left(\frac{\eta - 1}{\nu_p}\right) \mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}).$$

There are two main takeaways. First, without habits, expected future inflation and expected discounted real marginal cost are sufficient statistics for current inflation. Second, with habits, the sales-based NKPC is generally difficult to estimate empirically because it requires separate measures of the real marginal benefit of sales and the marginal cost.

4.3.3 Stock-based NKPC

Inventories also give rise to the equilibrium stocking condition, which characterizes how retailers choose inventory holdings. The log-linearized condition is

$$\hat{a}_t = \hat{c}_t + \frac{1}{1 - \beta(1-\delta)}[\beta(1-\delta)\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}) - \widehat{mc}_t] + \hat{\kappa}_t. \tag{14}$$

This condition highlights three key determinants of the retailer's stocking decision. First, *aggregate demand* ($\hat{c}_t$) matters, as retailers target a desired inventory-to-sales ratio; higher expected sales translate into higher desired stocks. Second, the *intertemporal substitution*

channel in stocking captures the difference between the expected discounted future marginal cost and the current marginal cost: when costs are expected to rise, firms have an incentive to buy more today and carry inventories forward. This channel parallels the mechanism emphasized by Crouzet and Oh (2016). Third, the term $\hat{\kappa}_t$ reflects the *real marginal benefit of sales*: holding more stock increases sales capacity and boosts expected profits, reinforcing the incentive to accumulate inventories.

Substituting out $\hat{\kappa}_t$ and denoting $\overline{IS}$ as the steady-state inventory-to-sales ratio, we obtain the following stock-based NKPC.

Proposition 2 (Stock-based NKPC)**.** *The stock-based NKPC is derived by combining* (12) *and* (14)*:*

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \tilde{B}_1[\beta(1-\delta)\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}) - \widehat{mc}_t] - \tilde{B}_2 \widehat{is}_t,$$

where

$$\tilde{B}_1 = [\nu_p(1-\beta(1-\delta))]^{-1}, \quad \tilde{B}_2 = \overline{IS}(\nu_p(\overline{IS}+1))^{-1}.$$

Two insights stand out. First, inflation is driven by three observable statistics: expected future inflation, the expected change in discounted real marginal cost, and the inventory-to-sales ratio. Second, this representation is agnostic to the source of endogenous markup variation since these demand effects are absorbed by the inventory-to-sales ratio. The dynamics of the ratio thus serve as a sufficient statistic for markup behavior in a wide class of models, including CES and deep habits.

4.4 Discussion

We find that the stock-based NKPC offers greater flexibility by remaining agnostic about the sources of time-varying markups. Its key observable is the inventory-to-sales ratio, which—after purging the intertemporal substitution channel—primarily captures the inflationary role of the marginal benefit of sales, which is closely related to desired markups. This provides an empirically tractable proxy for time-varying markups in inflation dynamics, which we analyze in the next section.

It is also worth noting that, while our industry-level IV results exploit exogenous variation from short-lived supply disruptions, the stock-based NKPC does not take a stand on the origin of shocks. Its relative improvement in fit should therefore not be interpreted as evidence of the importance of supply disruptions *per se*, since any inflationary disturbance can propagate through inventories. Put differently, while our IV strategy isolates shifts in the stocking curve, the fluctuations in the inventory-to-sales ratio that enter the stock-based NKPC reflect both shifts of the curve and endogenous movements along it.

5 Empirical Phillips curve with inventories

In this section, we estimate the stock-based NKPC using macroeconomic data and show that it outperforms both the canonical specification and the sales-based NKPC under CES demand. The improvement comes from including the inventory-to-sales ratio as an observable, which captures the time-varying markup channel missing from standard NKPCs. We further show that this specification accounts for two prominent episodes—the 2009–11 missing disinflation and the COVID-era surge—underscoring the role of inventories as an indicator of inflationary pressures. Note that this section provides a reduced-form empirical assessment of aggregate inflation dynamics; causal effects of the inventory channel are established earlier in the industry-level regressions.

5.1 Empirical specification and data

We estimate three reduced-form (semi-structural) Phillips curve specifications with theory-motivated exclusion and equality restrictions:

$$\hat{\pi}_t = \alpha_0 + \alpha_1 \hat{\pi}_{t-1} + \alpha_2 \mathbb{E}_t \hat{\pi}_{t+1} + \alpha_3 \widehat{mc}_t + \alpha_4 (\mathbb{E}_t \hat{r}_{t,t+1} + \mathbb{E}_t \widehat{mc}_{t+1}) + \alpha_5 \widehat{is}_t + \varepsilon_t, \tag{15}$$

where ε_t denotes the regression error. The first sets $\alpha_4 = \alpha_5 = 0$, yielding the canonical NKPC without inventories, as in Galí and Gertler (1999), with both backward- and forward-looking inflation expectations. The second sets $\alpha_3 = \alpha_5 = 0$, corresponding to the sales-based NKPC under CES demand (Proposition 1). The third imposes the equality restriction

$\alpha_3 = -\alpha_4$, delivering the stock-based NKPC (Proposition 2) with expected marginal cost growth expressed under zero inventory depreciation and no average time discounting.[14]

To estimate these regressions, we compile macroeconomic variables at quarterly frequency. Since retail inventory management is primarily relevant within the goods sector, our analysis focuses on accounting for goods inflation, though we also report results using headline inflation for comparison. Accordingly, our baseline regression defines $\hat{\pi}_t$ as the annualized quarterly inflation of the goods component of Personal Consumption Expenditures (PCE). For expected inflation $\mathbb{E}_t\hat{\pi}_{t+1}$, we use both the mean professional inflation forecast from the Survey of Professional Forecasters (SPF) and the mean household inflation forecast from the Michigan Survey, each reflecting year-ahead expectations. To our knowledge, no survey provides separate forecasts specifically for goods inflation. By estimating an unrestricted coefficient on inflation expectations, we allow for proportionality between expectations of headline and goods inflation.

To proxy for the real marginal cost, $\widehat{mc}_t$, we use the unemployment gap (scaled by -1), a commonly used measure of economic slack (e.g., Stock and Watson, 2020). Specifically, we define the unemployment gap as the difference between the quarterly unemployment rate and the quarterly estimate of the Non-Accelerating Inflation Rate of Unemployment (NAIRU) from the Congressional Budget Office (CBO). Because our Phillips curve specification includes expected future real marginal cost, $\mathbb{E}_t\widehat{mc}_{t+1}$, using the unemployment gap is particularly useful, as we can construct a comparable expectation-based proxy. We define the expected future unemployment gap as the difference between the SPF's year-ahead unemployment forecast and the year-ahead NAIRU estimate from the CBO.

Alternative proxies for the real marginal cost used in the literature include the output gap, the labor share, the vacancy-to-unemployment ratio, and unit labor cost. Because the NKPC model with inventories requires a measure of expected future real marginal cost, however, constructing such a measure using these alternatives is more challenging.[15]

We proxy the expected real stochastic discount factor, $\mathbb{E}_t\hat{r}_{t,t+1}$, by the SPF's year-ahead

[14] Although results are robust to relaxing the restriction, estimating α_3 and α_4 separately induces multicollinearity due to the high correlation between current and expected marginal cost.

[15] Nevertheless, in the Appendix we show that our main results are robust to the choice of the marginal cost proxy in a canonical NKPC specification augmented with the inventory-to-sales ratio, suggesting that our findings are unlikely to be driven by the specific proxy used.

inflation forecast minus the effective Federal Funds rate.

Finally, for the inventory-to-sales ratio, $\widehat{is}_t$, we aggregate inventory and sales data from the retail and wholesale sectors. We detrend this variable using the Hamilton (2018) filter.

The sample spans 1983q3–2026q1. The inventory-to-sales ratio is constructed from 1980q4 onward—when wholesale inventories become available—and this earlier start is used to implement the Hamilton filter, which regresses the ratio on its 8th–11th lags.

5.2 NKPC regression results

We now present regression estimates for the three Phillips curve specifications.

5.2.1 Evidence for the stock-based NKPC

Our derivation of multiple NKPC representations is useful because, under CES demand, the sales-based and stock-based specifications should be empirically equivalent. If their empirical fit differs in favor of the stock-based specification, this provides evidence against CES and, through the lens of our model, implies that aggregate inflation dynamics require channels of time-varying markups beyond those captured by nominal rigidity.

Table 3 reports OLS estimates of the three specifications in equation (15), using goods and headline PCE inflation as dependent variables and SPF inflation expectations. For goods inflation, the stock-based NKPC in column (3) attains a noticeably higher adjusted R-squared than both the canonical specification (column 1) and the sales-based NKPC under CES demand (column 2). The same pattern holds for headline inflation: column (6) dominates columns (4) and (5). The better fit of the stock-based NKPC is consistent with time-varying markup forces absent from the CES sales-based specification. Deep habits provide one microfoundation for these forces. Taken together with the industry-level evidence, these results provide plausible evidence that time-varying markups matter for inflation dynamics.[16]

[16]In the Appendix, we show that other metrics such as the Akaike information criterion or the one-quarter-ahead root mean squared errors, which also generally favor the stock-based NKPC.

TABLE 3: NKPC Estimates for Goods and Headline Inflation

	Goods PCE Inflation			Headline PCE Inflation		
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged inflation	0.359***	0.361***	0.052	0.467***	0.475***	0.207*
	(0.128)	(0.128)	(0.099)	(0.130)	(0.136)	(0.121)
Exp. inflation	0.317	0.290	1.137**	0.297*	0.225*	0.633***
	(0.276)	(0.250)	(0.455)	(0.154)	(0.124)	(0.227)
Marginal cost	0.057			0.086		
	(0.145)			(0.064)		
Exp. marginal cost		0.020			0.046	
		(0.093)			(0.048)	
Exp. marginal cost growth			0.047			0.091
			(0.162)			(0.094)
IS ratio			-0.435***			-0.188***
			(0.087)			(0.046)
Constant	-0.077	-0.047	-2.131*	0.532*	0.647**	0.071
	(0.843)	(0.761)	(1.146)	(0.305)	(0.301)	(0.533)
Observations	172	172	172	172	172	172
Adjusted R-squared	0.136	0.135	0.389	0.317	0.315	0.503

Notes: Variable definitions are provided in the main text. Expected inflation is from the SPF; results using the Michigan Survey are reported in the Appendix. Newey–West standard errors are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

5.2.2 Why the stock-based NKPC fits better: The role of inventories

Why does the stock-based NKPC fit better? Inspecting columns (3) and (6), the inventory-to-sales ratio enters with a negative and statistically significant coefficient, as predicted by equation (14), while variables related to marginal cost dynamics—including both current and expected marginal cost—are not significant. The magnitudes are economically meaningful: a one-standard-deviation decline in the inventory-to-sales ratio is associated with roughly a 2 percentage-point increase in goods inflation and a 0.9 percentage-point increase in headline inflation. By comparison, a one-standard-deviation rise in SPF expected infla-

tion is associated with increases of about 1.1 and 0.6 percentage points, respectively. These results indicate that the inventory-to-sales ratio plays an important role in accounting for inflation dynamics, likely operating through goods prices.

In column (3), once the inventory-to-sales ratio is included, the coefficient on SPF expectations strengthens substantially, while the lagged goods inflation term becomes negligible. This suggests that lagged goods inflation had been proxying for omitted inventory dynamics, attenuating the role of expectations. With inventories included, the SPF term becomes more informative, consistent with goods-price dynamics being shaped jointly by professional forecasts and retailers' inventory management decisions.

Using the Michigan Survey for inflation expectations yields a complementary message, as shown in the Appendix. The coefficients of the goods NKPC are stable, and lagged goods inflation is insignificant even in specifications without inventories, consistent with evidence that household expectations—rather than backward-looking indexation—play a central role in the Phillips curve (Coibion et al., 2018). Adding the inventory-to-sales ratio in this setting leaves coefficients largely unchanged but increases explanatory power, indicating that inventories provide incremental information largely orthogonal to the channel captured by Michigan expectations. In this sense, inventories complement household expectations by improving fit without undermining the structural stability of the NKPC.

Our findings are robust across a wide range of specifications and data choices.[17] Re-estimating the model with alternative marginal cost proxies commonly used in the literature confirms that the inventory-to-sales ratio remains a strong and significant predictor of inflation.[18] We further estimate the NKPC using core goods inflation—adjusting the inventory-to-sales ratio accordingly—and again obtain results consistent with our main conclusion. Finally, testing a nonlinear specification of the inventory term shows that the linear mechanism continues to dominate.

TABLE 4: Semipartial R-squared (Unique Contribution of Regressors)

	Goods PCE Inflation	Headline PCE Inflation
Lagged inflation	0.18	2.61
Exp. inflation	3.91	4.56
Exp. marginal cost growth	0.04	0.58
IS ratio	25.3	19.09
Total R-squared	40.34	51.47

Notes: Semipartial R-squared values report each regressor's marginal contribution conditional on all others. All values are percentages.

5.2.3 Orthogonal contribution of inventories

To isolate the contribution of the inventory–sales ratio in the stock-based NKPC and compare it with other regressors, we report the semipartial R-squared. Following the Frisch-Waugh-Lovell logic, we orthogonalize each regressor with respect to the other controls and then computing its explanatory power. This statistic captures the regressor's unique (orthogonal) contribution and, in general, does not sum to total R-squared across regressors. Recall that, in our framework, the component of the inventory-to-sales ratio orthogonal to expected marginal cost growth primarily reflects movements in time-varying markups; accordingly, its semipartial R-squared provides a conservative measure of the markup channel.

Table 4 reports results for both goods and headline inflation using expectations from either the SPF or the Michigan Survey. For SPF expectations (Panel A), the inventory-to-sales ratio accounts for a substantial 24.4 percent, far exceeding the marginal cost and inflation expectations channels in terms of orthogonal contributions. This indicates that variation in the inventory-to-sales ratio orthogonal to marginal cost and inflation expectations is quantitatively important for goods inflation. A similar, though smaller, role appears for headline inflation, where the inventory-to-sales ratio's semipartial R-squared is about

[17] See Appendix for details.

[18] As discussed earlier, we estimated the canonical NKPC augmented with the inventory-to-sales ratio, rather than the stock-based NKPC, because expected marginal cost proxies are not available for this robustness check.

18.9 percent.

5.3 The stock-based NKPC and recent inflation episodes

Having established the superior fit of the stock-based NKPC in the full sample, we next examine whether the framework can also account for two specific inflation episodes that have attracted considerable attention in the literature.

5.3.1 The 2009–11 missing disinflation puzzle

The 2009–11 missing disinflation puzzle has been widely discussed in the literature, with a prominent resolution proposed by Coibion and Gorodnichenko (2015). Using the SPF, they show that a version of the canonical NKPC substantially underpredicts inflation over this period. However, when substituting inflation expectations from the Michigan Survey—where household expectations rose more sharply—the puzzle largely disappears. This implies that the puzzle is conditional on the choice of the inflation expectation measure, particularly the SPF.

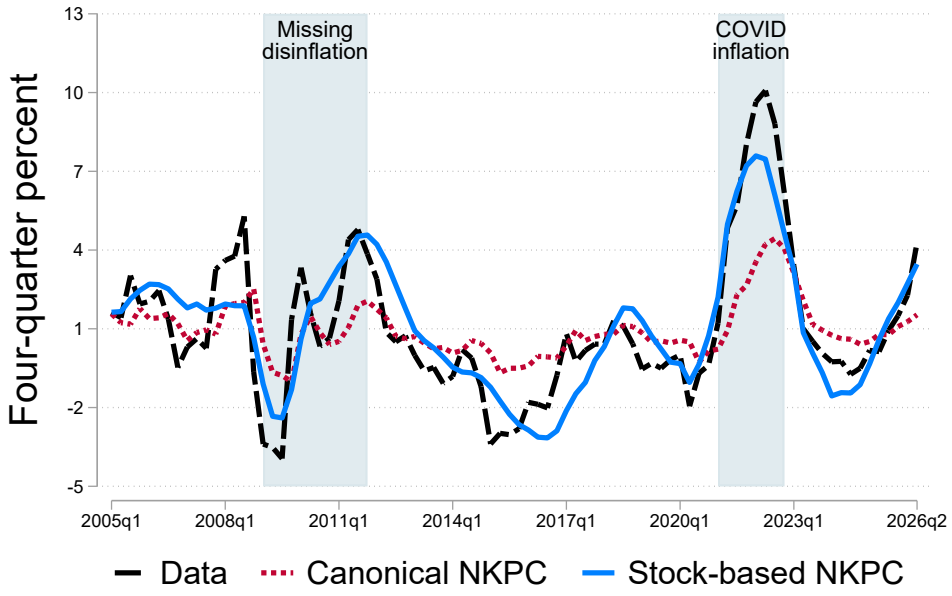
We offer an alternative explanation using the stock-based NKPC while still conditioning on SPF expectations. Panel B of Figure 3 plots actual headline PCE inflation alongside inflation predicted by both the canonical and stock-based NKPCs. The canonical specification predicts four-quarter inflation to remain below 2 percent between 2009 and 2011 (the first shaded region), while observed inflation reached 3 percent. In contrast, the stock-based NKPC successfully tracks the actual inflation path. This improvement coincides with a period of below-trend inventory-to-sales ratios, consistent with widespread inventory liquidation by financially constrained firms following the 2008 credit shock, as analyzed by Kim (2021).

5.3.2 The COVID-era inflation surge

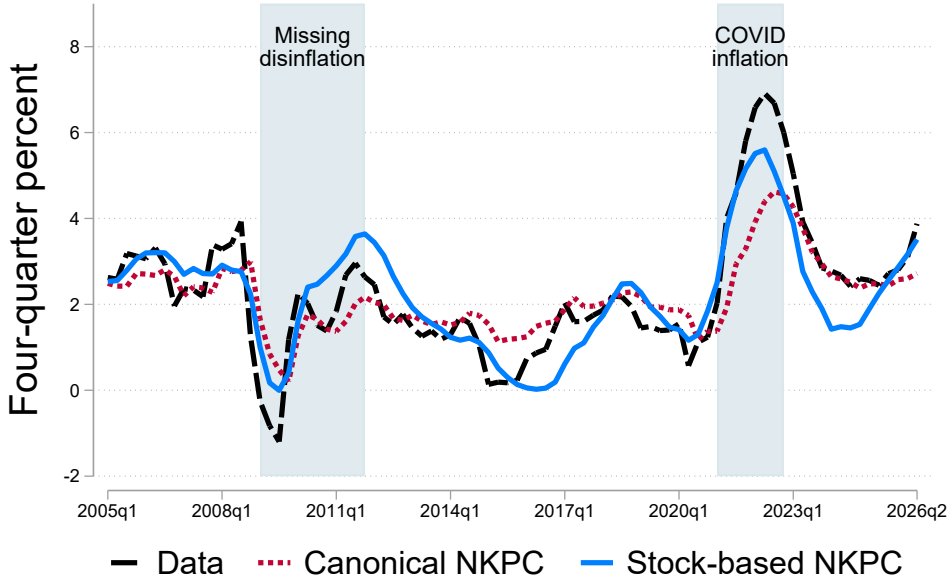
The stock-based NKPC also captures the surge in inflation—particularly for goods—during the COVID period. As shown in Panels A and B of Figure 3, goods inflation rose to 10 percent between 2021 and 2022 (the second shaded region), a surge that the stock-based NKPC better tracks in both panels. In contrast, the canonical NKPC explains only a

FIGURE 3: Model Fit

(A) Goods PCE Inflation



(B) Headline PCE Inflation



Notes: Panel A shows the model fit to goods PCE inflation. Panel B shows the model fit to headline PCE inflation. In both panels, the first shaded region corresponds to 2009q1–2011q4 and the second to 2021q1–2022q4. Data extend through 2026q2.

smaller portion of the rise, predicting a peak inflation of roughly 4 percent when using SPF expectations (Panel A) and 6 percent with Michigan expectations (Panel B). Moreover, the stock-based NKPC better captures the sharp reversal in goods inflation following the peak, underscoring the complementary role of the inventory-to-sales ratio in accounting for goods inflation dynamics, especially during periods of swift adjustment.

Turning to headline inflation, the tight inventory-to-sales ratio, which played a central role in the surge in goods inflation, also contributed significantly to the rise in headline inflation. The stock-based NKPC not only better matches the magnitude of the headline surge than the canonical specification, but also more accurately predicts its earlier timing—showing a peak that arrives sooner and aligns more closely with the data. This improvement is pronounced when using SPF expectations (Panel C), and while the magnitude remains similar with Michigan expectations (Panel D), the earlier predicted peak remains consistent with the observed timing.[19]

6 Conclusion

This paper investigates whether the retail price margin channel—whereby low inventories relative to sales amplifies retail price margins—can help explain observed inflation dynamics. By modeling the retailer's joint decision over pricing and inventory management, and constructing industry-level instruments from supply chain disruptions that plausibly affect only the retailer's stocking behavior, we identify the elasticity of price margins with respect to inventory stocks. We estimate this elasticity to be negative, implying that tighter retail inventories are systematically associated with higher price margins. This evidence supports the view that retail prices are shaped not only by cost-push factors but also by the strategic pricing behavior of firms facing inventory constraints.

Integrating this mechanism into a NKPC framework, we show that the inventory-to-sales ratio serves as a powerful and observable proxy for retail markups. The resulting stock-based NKPC delivers a substantially improved empirical fit, successfully capturing key inflationary episodes such as the 2009–11 missing disinflation and the COVID-era surge

[19] In the Appendix, COVID-era pseudo-out-of-sample inflation forecasts show a similar pattern, though all models exhibit a delayed pickup in inflation relative to the data.

that originated from the goods sector. These findings suggest that incorporating retail inventory dynamics enhances our understanding of inflation—particularly during periods when supply chains are constrained—by capturing a critical source of time-varying markups often overlooked in standard NKPC models.

Looking ahead, several avenues for further research remain. One is to examine how inventory dynamics interact with nonlinear NKPC specifications, for example by allowing the slope of the Phillips curve to vary across regimes of inventory tightness or demand pressure. Another is to investigate whether the inventory-based mechanism can account for the time variation in the slope of the NKPC documented in the empirical literature. More broadly, extending the framework to incorporate sectoral heterogeneity or to model the interaction between retail inventories and upstream supply chain frictions could yield new insights into how microeconomic bottlenecks shape aggregate inflation dynamics.

References

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl, "The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022," *Journal of International Economics*, 2023, *146*, 103788.

Alvarez-Blaser, Santiago, Alberto Cavallo, Alexander MacKay, and Paolo Mengano, "Markups and Cost Pass-through Along the Supply Chain," Working Paper 34110, National Bureau of Economic Research August 2025.

Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti, "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory," Working Paper 32098, National Bureau of Economic Research February 2024.

Ball, Laurence, Daniel Leigh, and Prachi Mishra, "Understanding US inflation during the COVID-19 era," *Brookings Papers on Economic Activity*, 2022, *2022* (2), 1–80.

Balleer, Almut and Marvin Noeller, "Monetary Policy in the Presence of Supply Constraints: Evidence from German Firm-level Data," Discussion Paper 18732, CEPR 2024.

Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," W. E. Upjohn Institute for Employment Research 1991.

Benigno, Pierpaolo and Gauti B Eggertsson, "It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve," Technical Report, National Bureau of Economic Research 2023.

Bernanke, Ben and Olivier Blanchard, "What caused the US pandemic-era inflation?," *American Economic Journal: Macroeconomics*, 2025, *17* (3), 1–35.

Beveridge, Stephen and Charles R. Nelson, "A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle," *Journal of Monetary Economics*, 1981, *7*, 151–174.

Bils, Mark and James A. Kahn, "What Inventory Behavior Tells Us about Business Cycles," *American Economic Review*, June 2000, *90* (3), 458–481.

Boehm, Christoph E. and Nitya Pandalai-Nayar, "Convex Supply Curves," *American Economic Review*, December 2022, *112* (12), 3941–69.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Braun, Robin, Aaron Flaaen, and Sinem Hacioglu-Hoke, "Supply vs Demand Factors Influencing Prices of Manufactured Goods," FEDS Notes, Board of Governors of The Federal Reserve System February 2024.

Cavallo, Alberto and Oleksiy Kryvtsov, "What Can Stockouts Tell Us about Inflation? Evidence from Online Micro Data," *Journal of International Economics*, 2023, *146*, 103769.

Coibion, Olivier and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 197–232.

_ and _ , "Inflation, Expectations and Monetary Policy: What Have We Learned and to What End?," Working Paper 33858, National Bureau of Economic Research May 2025.

_ , _ , and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, 2018, *56* (4), 1447–1491.

Comin, Diego A, Robert C Johnson, and Callum J Jones, "Supply Chain Constraints and Inflation," Working Paper 31179, National Bureau of Economic Research April 2023.

Crouzet, Nicolas and Hyunseung Oh, "What Do Inventories Tell Us about News-driven Business Cycles?," *Journal of Monetary Economics*, May 2016, *79*, 49–66.

Ferrante, Francesco, Sebastian Graves, and Matteo Iacoviello, "The Inflationary Effects of Sectoral Reallocation," *Journal of Monetary Economics*, 2023, *140*, S64–S81.

Flynn, Joel P, George Nikolakoudis, and Karthik Sastry, "Pricing and production without the invisible hand," Technical Report, National Bureau of Economic Research 2026.

Gagliardone, Luca and Mark Gertler, "Oil prices, monetary policy and inflation surges," 2023.

Galí, Jordi and Mark Gertler, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 1999, *44* (2), 195–222.

Hamilton, James D., "Why You Should Never Use the Hodrick-Prescott Filter," *The Review of Economics and Statistics*, 12 2018, *100* (5), 831–843.

Jung, Yongseung and Tack Yun, "Monetary Policy Shocks, Inventory Dynamics, and Price-setting Behavior," Working Paper Series 2006-02, Federal Reserve Bank of San Francisco 2005.

Kim, Ryan, "The Effect of the Credit Crunch on Output Price Dynamics: The Corporate Inventory and Liquidity Management Channel," *The Quarterly Journal of Economics*, February 2021, *136* (1), 563–619.

Kryvtsov, Oleksiy and Virgiliu Midrigan, "Inventories and Real Rigidities in New Keynesian Business Cycle Models," *Journal of the Japanese and International Economies*, June 2010, *24* (2), 259–281.

_ and _ , "Inventories, Markups, and Real Rigidities in Menu Cost Models," *The Review of Economic Studies*, January 2013, *80* (1), 249–276.

Lubik, Thomas A. and Wing Leong Teo, "Inventories, Inflation Dynamics and the New Keynesian Phillips Curve," *European Economic Review*, April 2012, *56* (3), 327–346.

Mavroeidis, Socrates, Mikkel Plagborg-Møller, and James H. Stock, "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 2014, *52* (1), 124–188.

Mongey, Simon and Michael E. Waugh, "Pricing inequality," Technical Report, National Bureau of Economic Research 2025.

Ortiz, Julio L., "Spread Too Thin: The Impact of Lean Inventories," International Finance Discussion Papers 1342, Federal Reserve Board 2022.

Ravn, Morten, Stephanie Schmitt-Grohé, and Martin Uribe, "Deep Habits," *The Review of Economic Studies*, 2006, *73*, 195–218.

Reis, Ricardo, "Four Mistakes in the Use of Measures of Expected Inflation," *AEA Papers and Proceedings*, May 2023, *113*, 47–51.

Rotemberg, Julio J., "Sticky Prices in the United States," *Journal of Political Economy*, December 1982, *90* (6), 1187–1211.

Smets, Frank and Rafael Wouters, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 2007, *97* (3), 586–606.

Stock, James H. and Mark W. Watson, "Slack and Cyclically Sensitive Inflation," *Journal of Money, Credit and Banking*, 2020, *52* (S2), 393–428.

Online Appendix to *Retail inventories and inflation dynamics: The price margin channel*

Neil Mehrotra* Hyunseung Oh† Julio Ortiz‡

Table of Contents

A Appendix to Section 1 **3**

B Appendix to Section 3 **4**

B.1 Data sources and construction of variables 4

B.1.1 Aggregate variables 4

B.1.2 Manufacturing-level variables 4

B.1.3 Retail-level variables 6

B.2 Summary statistics 8

B.3 Extracting transitory component of QSPC variable 9

B.4 First-stage regression results 10

B.5 Balance tests 11

B.6 Adjusted QSPC instrument 13

C Appendix to Section 4 **15**

C.1 Deriving the demand function 15

C.2 Full optimality conditions 16

C.2.1 Baseline with inventories 16

C.2.2 Without inventories 18

D Appendix to Section 5 **21**

D.1 Specification with Michigan Survey inflation expectations 21

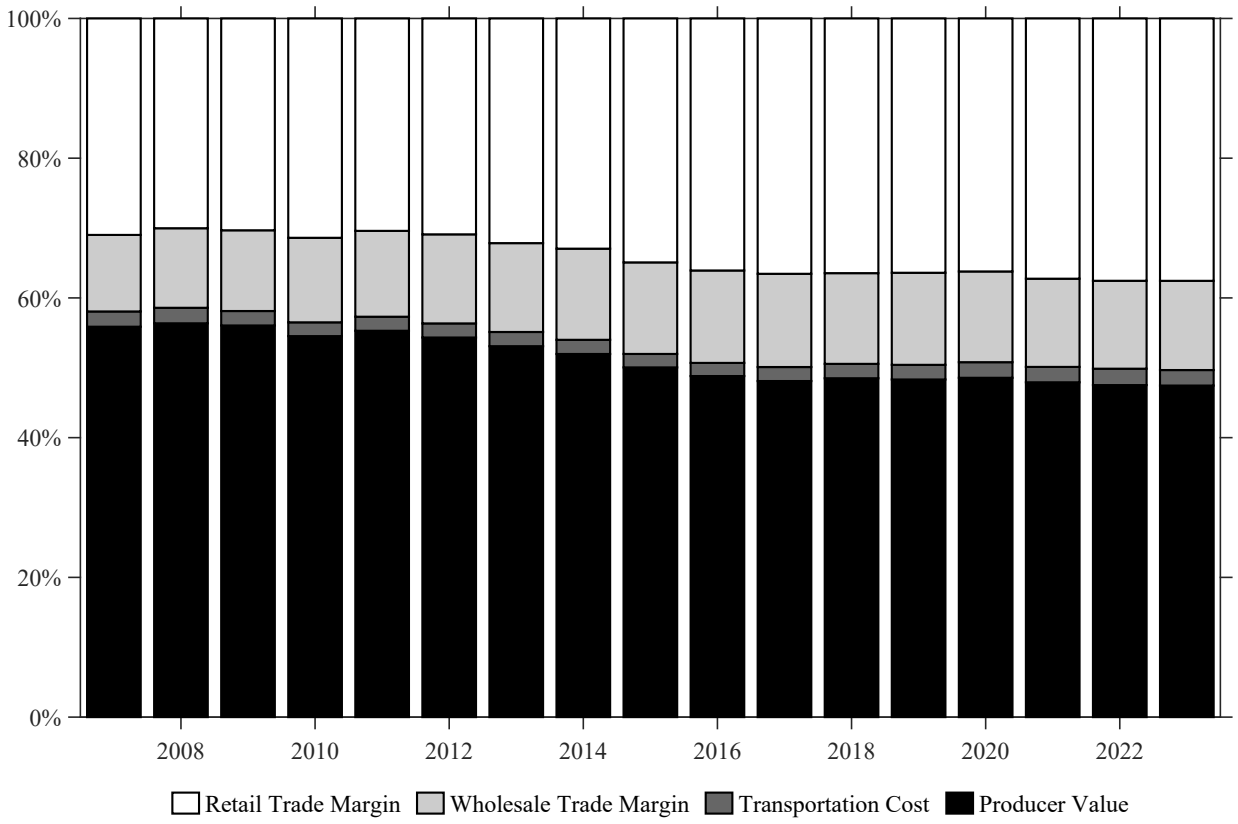
*Federal Reserve Bank of Minneapolis. E-mail: neil.mehrotra@mpls.frb.org.
†Federal Reserve Board. E-mail: hyunseung.oh@frb.gov.
‡Federal Reserve Board. E-mail: julio.l.ortiz@frb.gov.

D.2 NKPC model fit and forecast performance 24
D.3 Alternative marginal cost proxies 27
D.4 Core goods PCE inflation 30
D.5 Introducing a nonlinear inventory specification 32

A Appendix to Section 1

To calculate the share of retail and wholesale trade margins in the purchasers' value of PCE goods, we use the PCE bridge summary table published by the BEA. Figure A.1 shows the annual composition between 2007 and 2023. Producers' value accounts for roughly 50 percent of purchasers' value, a share that has gradually declined over time. Retail margins represent between 30 and 40 percent, with their share rising in recent years. Wholesale margins remain around 10 percent, while transportation costs account for only a small share.

FIGURE A.1: Breakdown of Purchasers' Value for PCE Goods



Notes: Data come from the PCE bridge table published by the BEA. Bars show producers' value, transportation cost, wholesale trade margin, and retail trade margin as shares of purchasers' value for PCE goods.

B Appendix to Section 3

B.1 Data sources and construction of variables

In this section, we detail the data used in our empirical analysis, the sources through which we collect this information, and steps taken to clean or construct certain variables.

B.1.1 Aggregate variables

The two aggregate time series used in our analysis are the ACR and demand-driven inflation.

Bai et al. (2024) ACR measure. The ACR is a weighted average of the congestion rates for the top 50 container ports worldwide. The congestion rate for a given port is defined as the number of delayed ship visits relative to the total number of ship visits.

B.1.2 Manufacturing-level variables

BLS. We collect PPI data from the BLS.

BEA. We construct our measure of import exposure for manufacturing industries by using the BEA's import matrix, a supplement to the I-O tables. The import matrix provides us with the dollar value imports of intermediate inputs from U.S. manufacturing industries. Using this information, we construct import shares in the year 2007.

For the construction of our industry-specific demand proxy, detailed below, we also collect industry-level quarterly gross output quantities and prices from the BEA.

Census Bureau. Plants responding to the QSPC report the reasons why they produce below capacity. We focus on responses citing "insufficient material supply" as the constraint. These data are publicly available at a quarterly frequency beginning in 2013q1, and available upon request back to 2008q1. Among the 21 three-digit NAICS manufacturing industries, 14 have a complete time series for the share of plants operating below capacity and reporting insufficient material supply, while the remaining 7 industries have

breaks in coverage. To ensure that we can accurately estimate and residualize the demand-driven and serially correlated components of the data, we restrict our analysis to the 14 industries with complete time series. These industries are:

1. NAICS 311 – Food manufacturing
2. NAICS 321 – Wood product manufacturing
3. NAICS 324 – Petroleum & coal products manufacturing
4. NAICS 325 – Chemical manufacturing
5. NAICS 326 – Plastics & rubber product manufacturing
6. NAICS 327 – Nonmetallic mineral product manufacturing
7. NAICS 331 – Primary metal manufacturing
8. NAICS 332 – Fabricated metal product manufacturing
9. NAICS 333 – Machinery manufacturing
10. NAICS 334 – Computers & electronic product manufacturing
11. NAICS 335 – Electrical equipment, appliances & component manufacturing
12. NAICS 336 – Transportation equipment manufacturing
13. NAICS 337 – Furniture & related product manufacturing
14. NAICS 339 – Miscellaneous manufacturing

Constructing industry-specific demand. Our QSPC shift variable is constructed by residualizing the raw QSPC data on an industry-specific demand proxy. This demand proxy is constructed using principal component analysis. For each manufacturing industry, we collect quarterly gross output quantities and prices from the industry accounts available through the BEA. We calculate the quarterly percent change in gross output quantities and prices, and extract a component that loads positively onto these two series. Intuitively, a demand shock would lead prices and quantities to move in the same direction.

B.1.3 Retail-level variables

BLS. We collect average hourly earnings, hours worked, and PPI data from the BLS. Note that the PPI data for retail and wholesale industries is the gross margin (i.e., the difference between the current selling price and the current aquisiton price for a good).

The PPI series for NAICS retailer 449 (Furniture, home furnishings, electronics, and appliance retailers) is only available starting in December 2022 (based on NAICS 2022 definitions). To obtain a full history of observations for this 3-digit industry, we combine the PPI of NAICS 442 (furniture stores, from the 2017 definitions) and NAICS 443 (electronic stores, from the 2017 definitions). The combined series is a weighted average of the PPI of NAICS 442 and 443 where we use a constant weight to construct the average. Specifically, we take the following steps:

1. Download the discontinued PPI series for retailers 442 and 443 from the BLS website.
2. Download the current PPI series for retailer 449 from the BLS website.
3. Download gross retail margin data from the Annual Retail Trade Survey: 2022 (re-stated). This "restated" release uses NAICS 2017.
4. Compute weights for retailers 442 and 442, defined as the average margin for each of these two industries from 2019-2022. The weights are about 0.7 and 0.3, respectively.
5. Normalize PPI for industries 442 and 443 such that they equal 100 in December 2022. Then construct a weighted average of the two series, where the weights are taken from the previous step.
6. Splice the PPI series for industry 449.

BEA. We collect real sales as well as real inventory-to-sales ratios from the National Income and Product Accounts Underlying Tables (Table 2BU and 3BU), available through the BEA.

In addition, we collect our measure of retail exposure from the 2007 I-O use table. From the I-O use table, we obtain the value of inputs purchased by a given retail industry from a given manufacturing industry in 2007. We construct the retail exposure share by

dividing by the total value of inputs (manufacturing and non-manufacturing) used by a given retailer.

B.2 Summary statistics

Table B.1 presents the summary statistics for each variable and lists the wholesale and retail industries included in our analysis.

TABLE B.1: Summary Statistics and Industry Names

Panel A: Summary Statistics	Mean	Median	Std.
Inventory-to-sales ratio	-66.81	-62.15	31.47
PPI (gross margin)	494.29	492.08	27.94
Average hourly earnings	305.98	307.65	28.69
Real sales	1211.46	1210.49	93.70
Hours worked	1215.55	1234.16	49.59
Exposure-weighted mfg. PPI	51.95	48.83	14.80
Panel B: Industry Names			
Industry	NAICS code		
Merchant wholesalers, durable goods	423		
Merchant wholesalers, nondurable goods	424		
Motor vehicle & parts dealers	441		
Building materials & garden equipment & supply dealers	444		
Food and beverage stores	445		
Furniture, home furnishings, electronics, & appliance retailers	449		
General merchandise stores	455		
Clothing & clothing accessories retailers	458		

Notes: Summary statistics are based on data for the sample period spanning 2008q1 to 2026q1. Industry classifications follow NAICS 2022 definitions.

B.3 Extracting transitory component of QSPC variable

To extract the transitory component of the QSPC variable, we estimate an unobserved components model for each manufacturing industry for which we have QSPC data. The model is specified as

$$
\begin{aligned}
QSPC_{jt} &= \tau_{jt} + c_{jt} + X_{jt}\gamma_j \\
\tau_{jt} &= \tau_{jt-1} + \eta_{jt}, \qquad \eta_{jt} \sim N(0, \sigma^2_{\eta,j}) \\
c_{jt} &= \rho_j c_{jt-1} + u_{jt}, \qquad u_{jt} \sim N(0, \sigma^2_{u,j})
\end{aligned}
$$

The set of controls specified in X_{jt} include the industry-specific demand proxies and a COVID dummy. The transitory component is c_{jt}. We estimate the model for each manufacturing industry, j, and specify the QSPC shift variable as $g^{QSPC}_{jt} = \widehat{c}_{jt}$.

B.4 First-stage regression results

In the first-stage regressions corresponding to Table 1 (reported in Table B.2), both QSPC material shortages and shipping congestion (ACR)—instruments designed to isolate supply-side disruptions—significantly predict declines in the retail inventory-to-sales ratio.

TABLE B.2: First-stage Regression Results

Dependent variable: Inventory-to-sales ratio					
Instrument	QSPC			ACR	QSPC, ACR
	Baseline	Pre-COVID	Transitory		
QSPC IV	-5.981***	-2.555*	-8.374***		-6.157***
	(1.333)	(1.317)	(1.997)		(0.959)
ACR IV				-57.94**	6.916
				(20.74)	(15.32)
Observations	584	384	584	320	320
Adj. R-squared	0.977	0.994	0.978	0.984	0.989
Sample	2008q1–2026q1	2008q1-2019q4	2008q1-2026q1	2016q2–2026q1	2016q2–2026q1

Notes: Driscoll–Kraay standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.5 Balance tests

To further support the validity of our QSPC-based shift-share instrument, Table B.3 reports a set of falsification tests at both the manufacturer and retailer levels.

One concern is that the QSPC-based shift variable may not adequately control for demand conditions at the manufacturer level. To assess this, we regress two standard demand proxies—new orders-to-shipments and shipments—on the QSPC shift variable. Panel A reports the results at one- and two-quarter horizons, and in all cases we find no statistically significant relationship.

Another concern is that the instrument may not fully account for potential confounders at the retailer level. To address this, we regress three retail outcomes—real sales, non-capital and non-labor input costs, and labor productivity—on the QSPC shift-share variable.[4] The first three rows of Panel B show that none of these variables are related to the instrument.

Finally, we conduct a pre-trend test by regressing lagged outcomes on the QSPC instrument. The fourth row of Panel B reports these results, which indicate that the QSPC-based shift-share instrument does not predict outcomes prior to the shifts.

[4]Non-capital and non-labor input costs are from the BLS satellite series of *net inputs to industry price indexes*, available from December 2018 for NAICS 441, 445, and 452, and from December 2020 for the remaining retail and wholesale industries (see https://www.bls.gov/ppi/input-indexes).

TABLE B.3: Balance Tests

	QSPC IV			ACR IV		
	Coefficient	Std. error	Obs.	Coefficient	Std. error	Obs.
Panel A: Manufacturer-level						
New orders	-0.035	0.129	508	-2.918	4.113	280
Shipments	-0.023	0.084	1009	-1.234	1.021	840
Lagged PPI	-0.002	0.001	1009	-0.019	0.014	839
Panel B: Retailer-level						
Real sales	-2.418	1.434	584	3.411	15.293	320
Non-labor, non-capital costs	0.420	0.297	106	-2.440	2.894	106
Labor productivity	1E-5	8.3E-6	584	0.0003	0.0002	320

Notes: Each row reports regressions in which the row variable is the dependent variable. Panel A reports regressions of the row variable against the QSPC shift variable, controlling for manufacturing industry fixed effects and time fixed effects. Panel B reports regressions of the row variable against the QSPC shift-share variable, controlling for the same controls as in regression (3) of the main text. Driscoll-Kraay standard errors are reported in the std. error column. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

B.6 Adjusted QSPC instrument

Ideally, we would observe the extent to which capacity utilization declines among plants operating below capacity. Instead, our QSPC variable of interest is constructed as the share of plants reporting material shortages relative to all plants operating below capacity. As a result, it is possible, at least in principle, for our QSPC measure to rise even as overall capacity utilization increases. To address this potential measurement concern, we normalize the QSPC variable by capacity utilization in each manufacturing industry, yielding an "adjusted" QSPC variable. We then residualize this adjusted variable as described above. This adjustment mildly strengthens the instrument, as reported in Table B.4, suggesting that our baseline QSPC instrument is relatively robust to potential mismeasurement.

TABLE B.4: Industry Regression Results: Adjusted QSPC

	Dependent variable: Gross price margin		
Estimation method	OLS	IV	
Instrument		Adj. QSPC	Adj. QSPC, ACR
Panel A: One-quarter Difference ($k = 1$)			
IS ratio	-0.143***	-0.518***	-0.616***
	(0.038)	(0.205)	(0.186)
First-stage F-stat		6.77	13.24
AR Wald F (p-value)		0.070	<0.001
Hansen J (p-value)			0.910
Observations	536	536	232
Sample	2008q1–2025q1	2008q1–2025q1	2017q2–2024q2
Panel B: Two-quarter Difference ($k = 2$)			
IS ratio	-0.121**	-0.725***	-0.595***
	(0.048)	(0.198)	(0.068)
First-stage F-stat		16.46	14.35
AR Wald F (p-value)		0.005	<0.001
Hansen J (p-value)			0.200
Observations	536	528	224
Sample	2008q1–2025q1	2008q2-2025q1	2017q3–2024q2

Notes: Driscoll–Kraay standard errors are reported in parentheses. Instruments for the IV columns are: reported material supply shortages from the Quarterly Survey of Plant Capacity (QSPC) and Average Congestion Rate (ACR). * p<0.10, ** p<0.05, *** p<0.01.

C Appendix to Section 4

In this section, we provide details on how we derived the non-CES demand function and several necessary equations based on that. We also provide the retailers' full optimality conditions.

C.1 Deriving the demand function

Households minimize the cost in choosing each variety to satisfy the habit-adjusted consumption variety:

$$c_{j,t} \equiv s_{j,t} - \theta s_{j,t-1},$$

taking as given the variety price as well as habits formed in the previous period:

$$\min_{c_{j,t}} \int_0^1 p_{j,t} c_{j,t} dj \quad \text{subject to} \quad \left[\int_0^1 v_{j,t}^{1/\eta} \left(c_{j,t}^{\frac{\eta-1}{\eta}} \right) dj \right]^{\frac{\eta}{\eta-1}} = C_t.$$

The first order condition for each variety can be written as

$$p_{j,t} = \lambda_t^p \left(\frac{v_{j,t} C_t}{c_{j,t}} \right)^{1/\eta},$$

where λ_t^p is the lagrange multiplier on the aggregation constraint. Defining the price index $P_t \equiv \lambda_t^P$ as

$$P_t = \left(\int_0^1 v_{j,t} p_{j,t}^{1-\eta} dj \right)^{\frac{1}{1-\eta}},$$

then at the optimum, we have

$$P_t C_t = \int_0^1 p_{j,t} c_{j,t} dj.$$

Therefore, the demand function could be expressed as

$$c_{j,t} = v_{j,t}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta} C_t,$$

In terms of real sales, the demand function is

$$s_{j,t} = v_{j,t}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta} C_t + \theta s_{j,t-1}. \tag{C.1}$$

C.2 Full optimality conditions

C.2.1 Baseline with inventories

The lagrangian of the retailer is written as follows:

$$\begin{aligned}\mathcal{L} = \sum_{t=0}^{\infty} \mathbb{E}_0 \Lambda_{0,t} \Bigg[& p_{j,t}s_{j,t} - Q_t y_{j,t} + P_t mc_{j,t}\left\{y_{j,t} + (1-\delta)(a_{j,t-1} - s_{j,t-1}) - a_{j,t}\right\} \\ & + P_t \kappa_{j,t}\left\{\left(\frac{a_{j,t}}{A_t}\right)^{\zeta}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta} C_t + \theta s_{j,t-1} - s_{j,t}\right\} - \frac{\nu_p}{2}\left(\frac{p_{j,t}}{p_{j,t-1}} - 1\right)^2 P_t S_t \Bigg].\end{aligned}$$

Taking the first order conditions, we get

$$\begin{aligned}
[p_{j,t}] &: s_{j,t} = \eta\kappa_{j,t}\left(\frac{a_{j,t}}{A_t}\right)^{\zeta}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta-1} C_t + \nu_p \frac{P_t S_t}{p_{j,t-1}}\left(\frac{p_{j,t}}{p_{j,t-1}} - 1\right) \\
&\quad - \nu_p \mathbb{E}_t \frac{\Lambda_{t,t+1} p_{j,t+1} P_{t+1} S_{t+1}}{p_{j,t}^2}\left(\frac{p_{j,t+1}}{p_{j,t}} - 1\right), \\
[a_{j,t}] &: mc_{j,t} = \zeta\kappa_{j,t}\left(\frac{a_{j,t}}{A_t}\right)^{\zeta-1}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta}\frac{C_t}{A_t} + (1-\delta)\mathbb{E}_t \frac{\Lambda_{t,t+1}P_{t+1}}{P_t} mc_{j,t+1}, \\
[s_{j,t}] &: \frac{p_{j,t}}{P_t} = \kappa_{j,t} + \mathbb{E}_t \frac{\Lambda_{t,t+1}P_{t+1}}{P_t}((1-\delta)mc_{j,t+1} - \theta\kappa_{j,t+1}), \\
[y_{j,t}] &: mc_{j,t} = \frac{Q_t}{P_t}, \\
[mc_{j,t}] &: a_{j,t} = y_{j,t} + (1-\delta)(a_{j,t-1} - s_{j,t-1}), \\
[\kappa_{j,t}] &: s_{j,t} = \left(\frac{a_{j,t}}{A_t}\right)^{\zeta}\left(\frac{p_{j,t}}{P_t}\right)^{-\eta} C_t + \theta s_{j,t-1}.
\end{aligned}$$

The nominal stochastic discount factor between period t and $t+1$ could be written as the household's subjective discount factor β, the effective real stochastic discount factor $r_{t,t+1}$, and the price indices:

$$\Lambda_{t,t+1} = \beta r_{t,t+1} \frac{P_t}{P_{t+1}}.$$

In a symmetric equilibrium, and replacing the stochastic discount factor with the complete market equation, the first-order conditions could be rewritten as

$$\begin{aligned}
[p_{j,t}] &: S_t = \eta\kappa_t C_t + \nu_p S_t \pi_t(\pi_t - 1) - \nu_p \beta \mathbb{E}_t r_{t,t+1} S_{t+1} \pi_{t+1}(\pi_{t+1} - 1), \\
[a_{j,t}] &: mc_t = \zeta\kappa_t \frac{C_t}{A_t} + \beta(1-\delta)\mathbb{E}_t r_{t,t+1} mc_{t+1}, \\
[s_{j,t}] &: 1 = \kappa_t + \beta \mathbb{E}_t r_{t,t+1}((1-\delta)mc_{t+1} - \theta\kappa_{t+1}), \\
[y_{j,t}] &: mc_t = \frac{Q_t}{P_t}, \\
[mc_{j,t}] &: A_t = Y_t + (1-\delta)(A_{t-1} - S_{t-1}), \\
[\kappa_{j,t}] &: S_t = C_t + \theta S_{t-1}.
\end{aligned}$$

The key steady-state conditions are

$$\begin{aligned}
S &= \eta\kappa C, \\
(1 - \beta(1-\delta))mc &= \zeta\kappa \frac{C}{A}, \\
1 &= (1 - \beta\theta)\kappa + \beta(1-\delta)mc, \\
C &= (1-\theta)S.
\end{aligned}$$

This implies that

$$\begin{aligned}
\kappa &= \frac{1}{\eta(1-\theta)}, \\
mc &= \frac{1 - (1-\beta\theta)\kappa}{\beta(1-\delta)}, \\
\zeta &= \eta(1 - \beta(1-\delta))mc \frac{A}{S}.
\end{aligned}$$

Therefore, the optimal pricing condition could be log-linearized as

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} - \frac{1}{\nu_p}(\hat{\kappa}_t + \hat{c}_t - \hat{s}_t).$$

The optimal stocking condition could be log-linearized as

$$\widehat{mc}_t = (1 - \beta(1-\delta))(\hat{\kappa}_t + \hat{c}_t - \hat{a}_t) + \beta(1-\delta)\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}).$$

Substituting out $\hat{\kappa}_t$, we get

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \frac{1}{\nu_p(1-\beta(1-\delta))}[\beta(1-\delta)\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}) - \widehat{mc}_t] - \frac{1}{\nu_p}(\hat{a}_t - \hat{s}_t).$$

Note that the optimal stocking condition could also be written as

$$\hat{a}_t - \hat{c}_t = \hat{\kappa}_t + \frac{1}{1-\beta(1-\delta))}[\beta(1-\delta)\mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}) - \widehat{mc}_t],$$

meaning that the stock to habit-adjusted sales ratio depends on both the expected marginal cost changes and the shadow value of the sales equation, which is linked to the price markup. In turn, the optimal sales condition could be expressed as

$$\hat{\kappa}_t = \theta\beta \mathbb{E}_t(\hat{r}_{t,t+1} + \hat{\kappa}_{t+1}) - [\eta(1-\theta) - 1 + \theta\beta]\, \mathbb{E}_t(\hat{r}_{t,t+1} + \widehat{mc}_{t+1}).$$

C.2.2 Without inventories

The lagrangian of the retailer without inventories is written as follows:

$$\mathcal{L} = \sum_{t=0}^{\infty} \mathbb{E}_0 \Lambda_{0,t} \Bigg[p_{j,t} s_{j,t} - Q_t y_{j,t} + P_t mc_{j,t} \{y_{j,t} - s_{j,t}\}$$
$$+ P_t \kappa_{j,t} \left\{ \left(\frac{p_{j,t}}{P_t} \right)^{-\eta} C_t + \theta s_{j,t-1} - s_{j,t} \right\} - \frac{\nu_p}{2} \left(\frac{p_{j,t}}{p_{j,t-1}} - 1 \right)^2 P_t S_t \Bigg].$$

Taking the first order conditions, we get

$$\begin{aligned}
[p_{j,t}] : s_{j,t} &= \eta\kappa_{j,t}C_t\left(\frac{p_{j,t}}{P_t}\right)^{-\eta-1} + \nu_p\frac{P_tS_t}{p_{j,t-1}}\left(\frac{p_{j,t}}{p_{j,t-1}}-1\right) \\
&\quad - \nu_p\mathbb{E}_t\frac{\Lambda_{t,t+1}p_{j,t+1}P_{t+1}S_{t+1}}{p_{j,t}^2}\left(\frac{p_{j,t+1}}{p_{j,t}}-1\right), \\
[s_{j,t}] : \kappa_{j,t} &= \frac{p_{j,t}}{P_t} - mc_{j,t} + \mathbb{E}_t\frac{\Lambda_{t,t+1}P_{t+1}}{P_t}\theta\kappa_{j,t+1}, \\
[y_{j,t}] : mc_{j,t} &= \frac{Q_t}{P_t}, \\
[mc_{j,t}] : y_{j,t} &= s_{j,t}, \\
[\kappa_{j,t}] : s_{j,t} &= \left(\frac{p_{j,t}}{P_t}\right)^{-\eta}C_t + \theta s_{j,t-1}.
\end{aligned}$$

In a symmetric equilibrium, the first-order conditions could be rewritten as

$$\begin{aligned}
[p_{j,t}] : S_t &= \kappa_t\eta C_t + \nu_p S_t\pi_t(\pi_t-1) - \nu_p\beta\mathbb{E}_t r_{t,t+1}S_{t+1}\pi_{t+1}(\pi_{t+1}-1), \\
[s_{j,t}] : \kappa_t &= 1 - mc_t + \theta\beta\mathbb{E}_t r_{t,t+1}\kappa_{t+1}, \\
[y_{j,t}] : mc_t &= \frac{Q_t}{P_t}, \\
[mc_{j,t}] : Y_t &= S_t, \\
[\kappa_{j,t}] : S_t &= C_t + \theta S_{t-1}.
\end{aligned}$$

The key steady-state conditions are

$$\begin{aligned}
S &= \kappa\eta C, \\
1 &= (1-\beta\theta)\kappa + mc, \\
C &= (1-\theta)S.
\end{aligned}$$

This implies that

$$\begin{aligned}
\kappa &= \frac{1}{\eta(1-\theta)}, \\
mc &= 1 - (1-\beta\theta)\kappa.
\end{aligned}$$

Therefore, the optimal pricing condition could be log-linearized as

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} - \frac{1}{\nu_p}(\hat{\kappa}_t + \hat{c}_t - \hat{s}_t).$$

The optimal sales condition could be expressed as

$$\hat{\kappa}_t = \theta\beta \mathbb{E}_t(\hat{r}_{t,t+1} + \hat{\kappa}_{t+1}) - [\eta(1-\theta) - 1 + \theta\beta]\, \widehat{mc}_t.$$

D Appendix to Section 5

In this section, we assess the robustness of our results along four dimensions: (i) replacing SPF-based inflation expectations with those from the Michigan Survey; (ii) employing alternative proxies for real marginal cost; (iii) estimating the NKPC with core goods PCE inflation as the dependent variable; and (iv) introducing a nonlinear inventory specification. Finally, we also examine pseudo-out-of-sample performance.

D.1 Specification with Michigan Survey inflation expectations

Our main results are robust to using inflation expectations from the Michigan Survey instead of the SPF. Table D.1 is analogous to Table 3 in the main text but replaces SPF expectations with those from the Michigan Survey. For headline PCE inflation, the Michigan Survey improves the fit of all three NKPC specifications in columns (4)-(6) relative to their SPF-based counterparts, as measured by the adjusted R-squared. Nevertheless, our main results hold: the stock-based NKPC in column (6) delivers the best fit and retains a significantly negative coefficient on the inventory-to-sales ratio. For goods PCE inflation, the improvement in fit from adding the inventory-to-sales ratio—comparing column (3) to column (1)—is even larger than in the headline case, as reflected in the adjusted R-squared.

Figure D.1 is analogous to Figure 3.

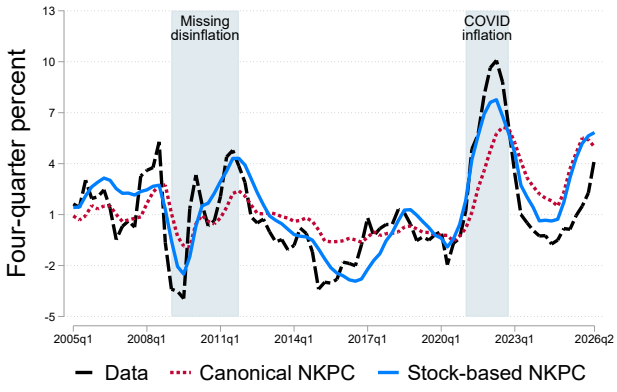
TABLE D.1: NKPC Estimates for Goods and Headline Inflation (Michigan Expectations)

	Goods PCE Inflation			Headline PCE Inflation		
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged inflation	0.213*	0.213*	0.052	0.299**	0.281**	0.170
	(0.115)	(0.115)	(0.110)	(0.122)	(0.132)	(0.131)
Exp. inflation (Michigan)	1.069**	1.069**	0.738*	0.571**	0.564**	0.440**
	(0.427)	(0.430)	(0.377)	(0.243)	(0.241)	(0.220)
Marginal cost	-0.024			0.077		
	(0.150)			(0.062)		
Exp. marginal cost		-0.010			0.066	
		(0.091)			(0.048)	
Exp. marginal cost growth			0.260*			0.216***
			(0.145)			(0.082)
IS ratio			-0.341***			-0.137***
			(0.078)			(0.039)
Constant	-3.469**	-3.454**	-2.128	-0.575	-0.569	0.037
	(1.595)	(1.605)	(1.410)	(0.730)	(0.746)	(0.689)
Observations	172	172	172	172	172	172
Adjusted R-squared	0.242	0.242	0.399	0.405	0.412	0.518

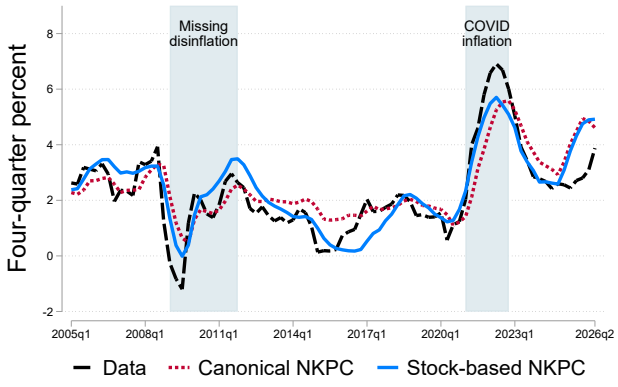
Notes: Variable definitions are provided in the main text. Expected inflation is from the Michigan Survey. Newey–West standard errors are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

FIGURE D.1: Model Fit

(A) Goods PCE Inflation



(B) Headline PCE Inflation



Notes: Panel A shows the model fit to goods PCE inflation. Panel B shows the model fit to headline PCE inflation. In both panels, the first shaded region corresponds to 2009q1–2011q4 and the second to 2021q1-2022q4. Data extend through 2026q2.

D.2 NKPC model fit and forecast performance

Table D.2 reports three model comparison measures: the Akaike (AIC) and Bayesian (BIC) information criteria, which gauge in-sample fit with a penalty for model complexity, and the one-quarter-ahead root mean squared error (RMSE), a pseudo-out-of-sample accuracy measure computed from a recursive regression starting with 1981q3-1995q1 and expanding through 2026q2. For all three metrics, lower values indicate better performance. Across goods and headline inflation, and using either SPF or Michigan inflation expectations, all measures favor the stock-based NKPC.

TABLE D.2: In-sample Fit and One-quarter-ahead Forecast Performance

	Canonical NKPC	CES-demand NKPC	Stock-based NKPC
Panel A: Goods PCE Inflation			
AIC	891.25	891.34	832.58
BIC	903.84	903.93	848.32
RMSE	3.72	3.72	3.40
Panel B: Headline PCE Inflation			
AIC	611.74	612.09	557.87
BIC	624.33	624.68	573.61
RMSE	1.71	1.71	1.57

Notes: The CES-demand NKPC corresponds to the sales-based specification under CES demand. AIC denotes the Akaike Information Criterion and BIC the Bayesian Information Criterion. RMSE is the one-quarter-ahead root mean squared error, computed from a recursive regression beginning with the sample period 1981q3–1995q1. Lower AIC and BIC values indicate better in-sample fit, while lower RMSE indicates superior out-of-sample performance.

For a formal pseudo-out-of-sample predictive-accuracy comparison between the stock-based and CES-demand NKPCs, we report the Diebold-Mariano (DM) statistics. Again, we re-estimate each specification recursively, starting with 1981q3-1995q and expanding through 2025q1, and generate one-quarter-ahead forecasts. DM statistics are computed from squared forecast errors. Table D.3 reports results for our four specifications, where positive values favor the stock-based NKPC over the CES-demand NKPC.

In the full sample, DM statistics generally favor the stock-based specification, although

power is limited given the reduced pseudo-out-of-sample size. By subsample, the stock-based NKPC delivers significantly higher predictive accuracy for goods inflation in the latter period (2010q1–2026q1). For headline inflation, gains are significant in the earlier period (1995q1-2009q4) and less clear thereafter. Overall, the DM tests tend to favor the stock-based NKPC, with weaker power in shorter windows.

TABLE D.3: Diebold-Mariano Tests: Recursive Pseudo-out-of-sample (One-quarter-ahead)

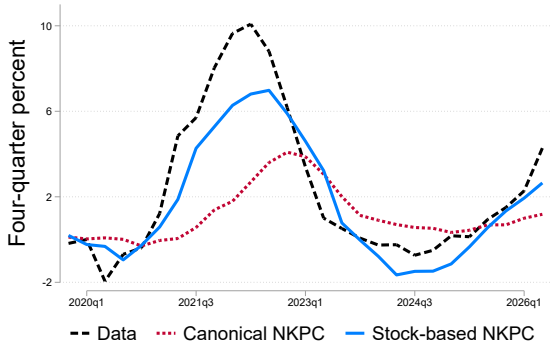
	1995q1-2026q1	1995q1-2009q4	2010q1-2026q1
Goods PCE	1.415	-0.263	1.571
Headline PCE	1.239	1.692*	0.940

Notes: The table reports Diebold-Mariano statistics. Positive values indicate lower squared forecast errors for the stock-based NKPC relative to the CES-demand NKPC. One-quarter-ahead forecasts are generated from recursive estimations starting with 1981q3–1995q1 and expanding through 2026q2. Diebold-Mariano statistics are based on quadratic loss differentials. * p<0.10, ** p<0.05, *** p<0.01.

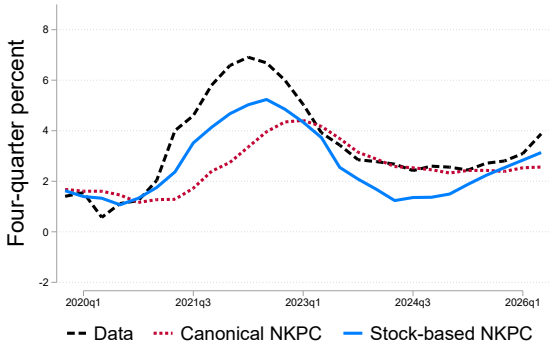
COVID-era pseudo-out-of-sample forecasts. Given the limited power of DM tests, the COVID-era inflation surge provides a useful episode for pseudo-out-of-sample comparisons, since the models are estimated without exposure to this period. Figure D.2 plots one-quarter-ahead pseudo-out-of-sample forecasts against realized goods and headline inflation from 2019q4 onward, using both SPF and Michigan inflation expectations. Quarterly forecasts are transformed into four-quarter inflation for the plots. For comparability with the main text, we show the stock-based and canonical NKPCs (the CES-demand NKPC behaves similarly to the canonical NKPC). The patterns mirror our in-sample results: the stock-based NKPC tracks the data more closely, especially during the COVID-era run-up in goods inflation. In particular, for headline inflation with Michigan expectations, the stock-based NKPC also anticipates an earlier pickup and earlier peak—albeit a lower one than the canonical NKPC—followed by a faster reversal.

FIGURE D.2: Pseudo-out-of-sample Model Fit (One-quarter-ahead)

(A) Goods PCE Inflation



(B) Headline PCE Inflation



Notes: Panel A plots one-quarter-ahead pseudo–out-of-sample forecasts for goods PCE inflation. Panel B show the corresponding forecasts for headline PCE inflation.

D.3 Alternative marginal cost proxies

We estimate the canonical NKPC regression with and without the inventory-to-sales ratio across several commonly used proxies for real marginal cost. Table D.4 reports results using goods PCE inflation, and Table D.5 reports results using headline PCE inflation. In both tables, columns (1) and (2) use the inverse of the unemployment rate gap, constructed as described in the main text; columns (3) and (4) use the labor share gap; and columns (5) and (6) use the unit labor cost gap, with both gap measures constructed using the Hamilton filter. Finally, columns (7) and (8) use the CBO's output gap. Across all four proxies, including the inventory-to-sales ratio as an observable improves the fit of the NKPC: the adjusted R-squared value in columns (2), (4), (6), and (8), are consistently higher than those in columns (1), (3), (5), and (7), respectively.

TABLE D.4: NKPC Estimates for Goods PCE Inflation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged inflation	0.359***	0.0275	0.351***	0.0252	0.363***	-0.00145	0.338***	0.0299
	(0.128)	(0.0991)	(0.126)	(0.103)	(0.124)	(0.0931)	(0.128)	(0.100)
Exp. inflation	0.317	1.201***	0.411	1.392***	0.344	1.344***	0.266	1.151***
	(0.276)	(0.412)	(0.328)	(0.433)	(0.326)	(0.433)	(0.266)	(0.395)
Unemployment rate gap (inverse)	0.0571	0.270**						
	(0.145)	(0.137)						
Labor share gap			-0.157	0.115				
			(0.117)	(0.107)				
Unit labor cost gap					-0.00695	0.185*		
					(0.131)	(0.0968)		
Output gap (CBO)							0.294*	0.279**
							(0.170)	(0.133)
IS ratio gap		-0.454***		-0.473***		-0.471***		-0.433***
		(0.0815)		(0.0829)		(0.0813)		(0.0779)
Constant	-0.0767	-2.057*	-0.371	-2.738**	-0.196	-2.570**	0.170	-1.992*
	(0.843)	(1.077)	(0.952)	(1.120)	(0.970)	(1.127)	(0.840)	(1.029)
Observations	172	172	169	169	169	169	172	172
R-squared	0.136	0.403	0.144	0.406	0.135	0.412	0.157	0.408

Notes: Newey-West standard errors are reported in parentheses. * denotes 10% significance, ** denotes 5% significance, *** denotes 1% significance.

Table D.5: NKPC Estimates for Headline PCE Inflation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged inflation	0.467***	0.133	0.467***	0.205*	0.476***	0.0988	0.430***	0.143
	(0.130)	(0.117)	(0.124)	(0.120)	(0.121)	(0.106)	(0.134)	(0.121)
Exp. inflation	0.297*	0.814***	0.350*	0.832***	0.281*	0.845***	0.294*	0.774***
	(0.154)	(0.198)	(0.181)	(0.235)	(0.165)	(0.206)	(0.151)	(0.187)
Unemployment rate gap (inverse)	0.0862	0.225***						
	(0.0638)	(0.0680)						
Labor share gap			-0.0981*	-0.00240				
			(0.0551)	(0.0604)				
Unit labor cost gap					0.0402	0.175***		
					(0.0519)	(0.0479)		
Output gap (CBO)							0.182**	0.205***
							(0.0802)	(0.0664)
IS ratio gap		-0.206***		-0.190***		-0.213***		-0.189***
		(0.0416)		(0.0472)		(0.0425)		(0.0391)
Constant	0.532*	-0.0277	0.315	-0.400	0.489	-0.177	0.644**	-0.0108
	(0.305)	(0.419)	(0.363)	(0.484)	(0.355)	(0.456)	(0.321)	(0.397)
Observations	172	172	169	169	169	169	172	172
R-squared	0.317	0.535	0.325	0.507	0.313	0.539	0.342	0.537

Notes: Newey-West standard errors are reported in parentheses. * denotes 10% significance, ** denotes 5% significance, *** denotes 1% significance.

D.4 Core goods PCE inflation

Table D.6 reports NKPC regression results when using core goods inflation as the dependent variable rather than overall goods inflation. To ensure consistency, we remove food and energy components from both wholesale and retail inventories and sales. Specifically, we exclude groceries, alcoholic beverages, and petroleum from wholesale measures, and drop food and beverage stores from retail inventories and sales as well as gasoline stations from retail sales. Constructing these core-consistent measures requires more detailed category-level data, which are only available beginning in 1992, reducing the sample to 122 observations. Within this sample, our results continue to hold for core goods PCE: the coefficient on the inventory-to-sales ratio remains negative and statistically significant, and the adjusted R-squared improves markedly for the stock-based NKPC relative to the canonical NKPC.

TABLE D.6: NKPC Estimates for Core Goods PCE Inflation

	(1)	(2)	(3)
Lagged core inflation	0.622***	0.614***	0.327***
	(0.122)	(0.125)	(0.119)
Exp. inflation (SPF)	-0.366	-0.178	0.783
	(0.308)	(0.302)	(0.484)
Marginal cost	0.0327		
	(0.0755)		
Exp. marginal cost		-0.0182	
		(0.0386)	
Exp. marginal cost growth			-0.105
			(0.0892)
IS ratio gap			-0.186***
			(0.0629)
Constant	0.894	0.420	-1.829
	(0.834)	(0.779)	(1.157)
Observations	127	127	127
R-squared	0.354	0.354	0.513

Notes: All the variables are described in the main text. For example, the marginal cost is proxied by the inverse of the unemployment gap and the expected marginal cost is defined as the sum of the expected real interest rate and the expected inverse of the unemployment gap. The expected marginal cost growth is the difference between the expected marginal cost and the marginal cost. Newey–West standard errors are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

D.5 Introducing a nonlinear inventory specification

In this section we consider one specific form of nonlinearity in the inventory channel: an interaction between the inventory–sales ratio gap and an indicator for whether the gap is negative or positive. This specification is motivated by the idea that inventory shortages may affect markups differently from inventory surpluses, and thus captures a potentially important asymmetry of interest to the literature. We emphasize, however, that this is only one way of modeling nonlinearities in the NKPC. A fuller exploration of alternative nonlinear specifications is an important question for future research and lies beyond the scope of this paper. Our main contribution is to document robust evidence of a strong linear effect of inventories on inflation dynamics, a relationship that has not been firmly established in the prior literature.

Columns (2) and (5) of Tables D.7 report the stock-based NKPC estimates, fitted to goods and headline PCE, respectively, when the regression is augmented with a nonlinear inventory term that interacts the inventory–sales ratio gap with an indicator for whether the gap is negative. The coefficient on the linear inventory-to-sales ratio gap remains negative in both regressions, consistent with the baseline linear results. The interaction term also enters with a negative sign, suggesting possible nonlinear effects when inventories fall short of trend.

The evidence for such nonlinearity, however, is weak. The nonlinear term is estimated imprecisely, with absolute t-statistics close to or below one, and its inclusion reduces the precision of the linear inventory coefficient, which loses conventional statistical significance. Nevertheless, the linear inventory-to-sales ratio gap continues to emerge as the primary inventory-related driver of inflation dynamics in these regressions.

To assess whether the weaker significance reflects collinearity with the nonlinear interaction rather than the disappearance of the linear effect, we re-estimate the specification using an interaction term for positive inventory-to-sales ratio gaps. In columns (3) and (6) of the tables, the linear inventory coefficient remains negative and statistically significant, while the nonlinear interaction is small and insignificant. Unlike the negative-gap specification, the positive-gap interaction does not appear strongly collinear with the linear inventory term in our data, which is why the linear coefficient remains precisely estimated.

These results suggest that the weaker significance in the negative-gap specification reflects multicollinearity, not the absence of the linear inventory channel. Overall, the evidence points to a robust linear effect of inventories on inflation, with only limited and imprecise signs of potential nonlinearities.

TABLE D.7: Inventory Nonlinearity in Stock-based NKPC

	Goods PCE inflation			Headline PCE inflation		
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged goods PCE inflation	0.0518	0.0440	0.0440			
	(0.0994)	(0.0892)	(0.0892)			
Lagged PCE inflation				0.207*	0.199*	0.199*
				(0.121)	(0.110)	(0.110)
Exp. inflation	1.137**	1.086**	1.086**	0.633***	0.619***	0.619***
	(0.455)	(0.446)	(0.446)	(0.227)	(0.224)	(0.224)
Exp. marginal cost growth	0.0467	0.111	0.111	0.0911	0.113	0.113
	(0.162)	(0.138)	(0.138)	(0.0941)	(0.0833)	(0.0833)
IS ratio gap	-0.435***	-0.299*	-0.527***	-0.188***	-0.142	-0.219***
	(0.0869)	(0.176)	(0.0725)	(0.0457)	(0.0863)	(0.0367)
IS ratio gap < 0		-0.228			-0.0768	
		(0.187)			(0.0865)	
IS ratio gap > 0			0.228			0.0768
			(0.187)			(0.0865)
Constant	-2.131*	-2.444**	-2.444**	0.0709	-0.0282	-0.0282
	(1.146)	(1.146)	(1.146)	(0.533)	(0.514)	(0.514)
Observations	172	172	172	172	172	172
R-squared	0.389	0.394	0.394	0.503	0.504	0.504

Notes: Newey-West standard errors are reported in parentheses. * denotes 10% significance, ** denotes 5% significance, *** denotes 1% significance.