

# Inflation after COVID\*

Lukasz A. Drozd<sup>†</sup>

Marina M. Tavares<sup>‡</sup>

March 2026

## Abstract

The canonical New Keynesian model of the Phillips curve requires a source of convexity or autocorrelated exogenous shocks to account for the post-COVID inflation surge. We argue that the missing link is the pandemic-era destruction of productive capacity and the persistent, uneven supply gap it left at the onset of reopening. We document the empirical footprint of this mechanism in a supplier-exposure regression and embed it in a nonlinear multisector New Keynesian model with production-network propagation. The key friction is a convex capacity constraint upon reopening, with the costs of rebuilding becoming marginal in constrained sectors and propagating via the network at low elasticity. Under the baseline calibration, the model accounts for two-thirds of post-COVID inflation, with more than half of this contribution attributable to the cross-sectional *unevenness*. The mechanism entails endogenous steepening of the Phillips curve, half of which attributes to shock size and half to shock's uneven geometry.

*JEL Codes:* E31, E32, E52, L11, L16

*Keywords:* Inflation, capacity constraints, supply chain disruptions, production networks, cost pass-through, COVID, hiring frictions

---

\*We thank Roc Armenter, Jonas Arias, Giancarlo Corsetti, Domenico Giannone, João Gomes, Urban Jermann, Anna Lipińska, and Makoto Nakajima for helpful comments and discussions. We are especially grateful to João Gomes for his insightful suggestions and early feedback. The views expressed herein are those of the authors and do not necessarily represent the views of the IMF, the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or their respective management and Executive Boards.

<sup>†</sup>Federal Reserve Bank of Philadelphia. Email: lukasz.drozd@phil.frb.org

<sup>‡</sup>International Monetary Fund. Email: mtavares@imf.org

# 1 Introduction

The standard framework for analyzing inflation dynamics in modern macroeconomics is the New Keynesian Phillips curve (NKPC),

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + \varepsilon_t, \quad (1)$$

where  $\pi_t$  is inflation,  $\beta \in (0, 1)$  is the subjective discount factor,  $x_t$  is a measure of economic slack—typically proxied by the output or unemployment gap—and  $\varepsilon_t$  is an exogenous cost-push shock. Empirical estimates over the pre-pandemic period established that the slope  $\kappa$  is exceptionally flat (Hazell et al., 2022; Del Negro et al., 2020). This finding solidified a consensus that decades of monetary-policy credibility had effectively changed the regime: long-run inflation expectations are firmly anchored by the central bank’s explicit target, which neutralizes  $\mathbb{E}_t \pi_{t+1}$  both as a source of self-fulfilling volatility and as a parallel transmission channel. In this world, chronic inflation is a regime outcome, and as long as monetary policy remains credible in its commitment to fight inflation, transient shocks have small, transitory effects on inflation, giving monetary policy more breathing space.<sup>1</sup>


From this perspective, the post-COVID inflation outbreak presents a clear puzzle. At annual frequency, medium- and long-horizon inflation expectations remained stable through the 2022 peak, ruling out a structural de-anchoring of  $\mathbb{E}_t \pi_{t+1}$  as the driver of the surge (Figure 1, Panel A)—as was the case in the 1970s. Conventional metrics of slack failed to signal severe overheating: the aggregate unemployment rate remained elevated well into the recovery, and as labor markets subsequently tightened, the historically flat  $\kappa$  should have insulated product prices from these demand pressures.<sup>2</sup> Yet inflation rose sharply and persistently, remaining well above the Federal Reserve’s 2% target for over three years.

However, the aggregate unemployment rate is unlikely to be an adequate measure of slack in this case, and the historic spike in the vacancy-to-unemployment ratio ( $V/U$ ) points to strong aggregate-demand

---

<sup>1</sup>This structural view was broadly shared among policymakers. See, for example, Yellen (2015), or Federal Reserve Chair Jerome H. Powell’s testimony during the July 10, 2019 House Financial Services Committee hearing (“Monetary Policy and the State of the Economy,” Serial No. 116–38, p. 40), available at <https://www.congress.gov/116/chrg/CHRG-116hhrg39738/CHRG-116hhrg39738.pdf>.

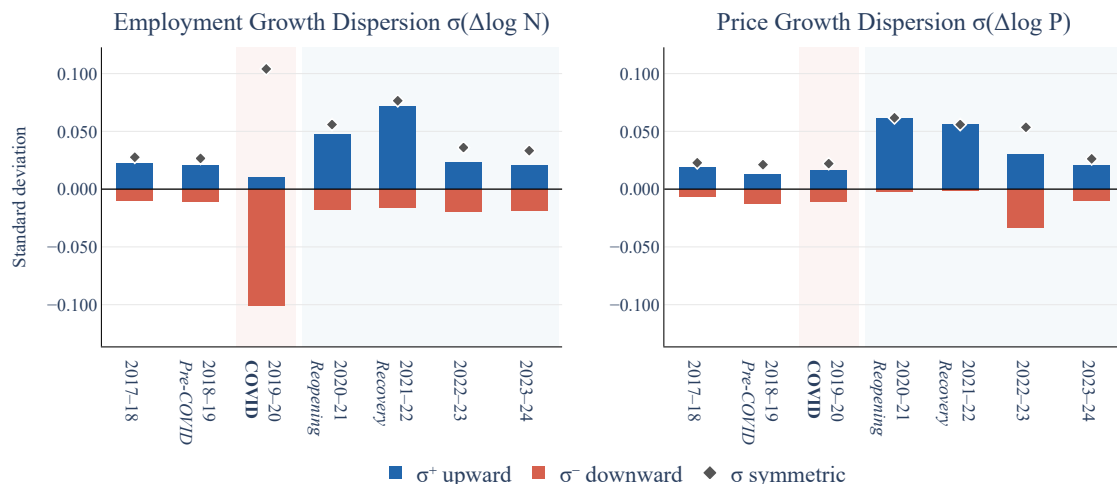
<sup>2</sup>See also Andrade et al. (2024) for a discussion of the low correlation between price and wage inflation at the sectoral level.



**Figure 1: Inflation after COVID.** Panel A: 1-year-ahead inflation expectations from the Federal Reserve Bank of Cleveland (FRED series EXPINF1YR) and year-over-year change in headline and core CPI (CPIAUCSL, CPILFESL). Panel B: cumulative log-difference between the Employment Cost Index for total compensation, all civilian (ECIALLCIV), and the indicated price index, measured against a 2017Q1 baseline; PCE deflators are PCEPI and PCEPILFE. Panel C: ratio of total nonfarm job openings (JTSJOL) to the unemployment level (UNEMPLOY); the dotted line marks the 2017–2019 average. Shaded region marks 2020Q1–2022Q4. Quarterly data; source: U.S. BLS and Federal Reserve Bank of Cleveland via FRED.

pressure soon after reopening started (Domash and Summers, 2022). In fact, the unprecedented spike in  $V/U$  correlates tightly with the inflationary wave, but the transmission mechanism is less clear. In the standard New Keynesian framework underpinning equation (1), labor-market tightness can only transmit to product prices through a wage–price spiral or through inflation expectations—both of which appeared muted (Figure 1). On the other hand, mapping this pressure into aggregate markups via  $\varepsilon_t$  creates a different kind of tension: the standard sticky-price models rely on *countercyclical* markups to break monetary neutrality, and even the baseline sticky-wage variants feature constant markups.

This paper contributes to the growing literature seeking a structural resolution of this puzzle. We focus on a mechanism that builds a microfounded bridge between the COVID-19 contraction and inflation. The key idea is that COVID resulted in a large and highly *asymmetric* destruction of production capacity, and this drove inflation because costs of rebuilding—which under our calibration are rehiring costs—steepened firm-level marginal cost curves. The COVID-19 pandemic was a disruptive event—resulting in massive movements of labor in and out of employment and in business closures (Decker and Haltiwanger, 2022)—and the economic literature has shown how even transitory disruptions become a persistent drag on downstream production (Barrot and Sauvagnat, 2016; Miyauchi, 2024). We



**Figure 2: Within-sector dispersion of employment growth and price growth.** GO-weighted aggregate across BEA sectors; NAICS-6 weights. Left:  $\sigma(\Delta \log N)$ . Right:  $\sigma(\Delta \log P)$ . Blue bars: upward dispersion  $\sigma^+(x) := \sigma(\max\{0, x\})$ ; red: downward dispersion  $\sigma^-(x) := \sigma(\min\{0, x\})$ ; diamonds: symmetric  $\sigma$ .

build on these ideas and on labor-market evidence on the convex cost of rebuilding workforces to operationalize this mechanism.

The empirical footprint of this mechanism—which we refer to as capacity wall (CW)—can be readily seen in the data. In 2020, employment collapsed across most industries, but it did so asymmetrically, resulting in a major spike in employment-contraction dispersion. This spike was followed by a persistent increase in price dispersion across sectors—a clear footprint of a disrupted production network (Figure 2). In our model, dispersion plays a key role via network connections under a low elasticity of substitution.

Our analysis proceeds in two steps.

The first part, the empirical part, provides reduced-form evidence linking supply contraction to subsequent inflation and price pressures. To document this connection, we use the COVID-19 Job-Exposure Matrix (JEM) of [Oude Hengel et al. \(2021\)](#) to construct a pre-pandemic occupation-mix exposure index for each industry, which we then propagate through the input-output network. Outcomes are cumulative log differences from a 2019 baseline, which mechanically eliminates any mean-reversion story. We identify the supply constraint via the *joint sign* of two coefficients: a positive loading on

contact-intensity exposure (Fact 1 below) and a negative loading on the speed of employment recovery, conditional on the 2020 collapse (Fact 2 below). The negative recovery sign discriminates the wall mechanism from the two main alternative interpretations of the cross-industry pattern—pent-up demand released into capacity-constrained sectors and cost-push from health adaptation—both of which predict the opposite or no recovery loading.

Our first fact is that JEM exposure predicts both employment collapse during the COVID year 2020 and higher cumulative 2019–2022 price growth in the cross-section. Both an industry’s own exposure and the exposure of its upstream suppliers load positively on subsequent price growth—within tightly controlled industry groupings. A decomposition of the supplier-cost pass-through reveals that the network channel does not look like ordinary wage transmission from supplier prices into buyer costs; instead, it runs through supplier *margins*—which carry essentially all of the exposure-to-material-input-cost loading—while supplier wages contribute little. The data show that downstream firms exposed to high-pressure suppliers absorbed part of the upstream pass-through by compressing their own margins, contributing to the persistence of inflation down the road.

Our second and key finding is the supply-side fingerprint that discriminates the wall channel from a pure demand-pull or demand-reallocation interpretation. Holding the size of the 2020 collapse fixed, we find that the speed of cumulative own and supplier employment recovery through 2022 enters the price regression with statistically significant *negative* coefficients: industries that rebuilt employment faster had *less*, not more, cumulative inflation. A demand-pull story predicts the opposite—industries hiring faster on reopening should raise prices faster, tracing out an upward-sloping short-run supply curve as suppressed demand is released into re-staffed capacity. The data run the other way. The result also survives controlling for industry teleworkability ([Dingel and Neiman, 2020](#)), ruling out the alternative that fast-recovering industries simply happened to be teleworkable and that teleworkability is itself anti-inflationary. In contrast, this sign matches the wall prediction: lower friction in rebuilding capacity dampens price pressure, which is disinflationary. Our third finding connects employment-collapse dispersion to price dispersion in [Figure 2](#). We do not claim causal identification because the

exposure measure used as an instrument is unlikely to cleanly satisfy the exclusion restriction, and so quantification of this mechanism relies on our model.

The second part of the paper develops a nonlinear multisector New Keynesian model with network propagation featuring an asymmetric (forced) COVID employment contraction that selectively destroys workforce capital. Three questions motivate this exercise. Do the rebuilding costs the model needs to explain inflation come anywhere near existing estimates of the marginal workforce-rebuilding costs relevant to this episode? What is the contribution of network propagation versus sheer shock size? Finally, is this explanation at odds with the theories that underpin the flat-Phillips-curve narrative discussed in the context of equation (1) above?

To answer the first question, we draw on a series of micro-level studies that estimate the level and curvature of workforce-building costs at the firm level (Linckh et al., 2025; Muehlemann and Strupler Leiser, 2018; Muehlemann and Pfeifer, 2016; Aepli et al., 2024). Our preferred anchor is the estimate of Muehlemann and Strupler Leiser (2018), who use Swiss data and find that filling a skilled-worker vacancy costs roughly 16 weeks of wages, of which about 80% arises post-match—training, adaptation, informal instruction, and disruption to incumbent workers—rather than pre-match search.<sup>3</sup> We combine this estimate with the evidence in Linckh et al. (2025), who find an even larger baseline number and specifically show that a one-standard-deviation increase in local labor-market tightness raises post-match hiring cost for the firm by roughly 10 to 12% on the largest cost components (adaptation and disruption time).<sup>4</sup> Combining this elasticity with the rise in the U.S.  $V/U$  ratio from less than 1.2 between 2017 and 2019 to almost 2.0 at its 2022 peak (Figure 1, Panel C), we conservatively assume that at peak overheating the average sector faced an excess cost of about three standard deviations—a premium of about 35% above the 16-week baseline, or roughly 5 to 6 additional weeks of wages. We then make two adjustments specific to the COVID episode. First, pandemic separations were unusually likely to be temporary lay-

---

<sup>3</sup>“Skilled” worker here is a broad concept: workers whose roles require some training, broader than the U.S. college-educated definition. Extrapolated to the U.S. wage-bill composition, this category covers roughly four-fifths of compensation, and we scale the raw cost estimate by this coverage in the calibration; see Section 4.

<sup>4</sup>Linckh et al. (2025) report a higher figure of approximately 22 weeks for Germany over 2015–2019, but the German labor market in that window was historically tight, and for this reason we prefer the more conservative baseline of Muehlemann and Strupler Leiser (2018).

offs, and according to estimates about 35–50% of laid-off workers were recalled within a year (the BLS temporary-layoff share peaked at 78% in April 2020 and averaged 49% over 2020).<sup>5</sup> Conservatively, using a back-of-the-envelope calculation based on these numbers, we discount the tightness premium by 30% relative to average 2020 employment. Second, the average-to-marginal cost conversion uses the 1.33 own hiring-cost elasticity in [Muehlemann and Pfeifer \(2016\)](#), which is identified from within-firm variation in hiring intensity holding local market conditions fixed. Our quantitative results scale roughly proportionally in the implied rebuilding-cost target, allowing for a simple sensitivity assessment around this target.

We show that the model calibrated to this cost accounts for about 70% of post-COVID inflation, with the remainder largely attributable to exogenous sectors (mostly housing and energy). Demand rotation after COVID enters passively and runs through the wall mechanism. About half of the model’s inflation effect comes from the *unevenness* of the shock across sectors rather than its sheer magnitude (average decline). In the early phase (2019–2021), almost all of the model’s excess inflation reflects the uneven geometry of the shock. Finally, the mechanism generates even more sizable inflation because it overcomes a structural floor in our baseline model that would predict deflation without the wall. We do not claim that this assessment is precise, but it suggests that quantitatively the cost of rebuilding can close much of the gap between standard theory and the data.

As for the slope of the Phillips curve, the wall enters the model as a nonlinear wedge in the marginal cost, implying a nonlinear and steep left tail of the curve. However, for small shocks, our model implies a flat Phillips curve. Furthermore, the right tail is mildly increasing because any major destructive contracts results in subsequent inflation, and this causes inflation today via inflation expectations.

On the supply-versus-demand attribution of post-COVID inflation, our results should be read with caution. Our model speaks narrowly to the decomposition of *propagation* conditional on the observed demand path, including its rotation. The exercise asks how much of inflation runs through the capacity wall’s amplification, not how inflation decomposes into supply and demand forces in a structural sense.

---

<sup>5</sup>See [California Policy Lab \(2021\)](#) for UI recalls in California and the additional evidence in [Forsythe et al. \(2022\)](#).

An inflation outbreak in the model necessarily reflects demand running hotter than available supply; the key question is only what this available supply was and what it implied for inflationary forces.

**Related literature.** The idea that convexity and supply disruptions mattered for post-COVID inflation is not new, particularly as an explanation of the initial phase of the outbreak. Our contribution is to provide a microfounded structural bridge from the COVID shock itself to the binding constraint.

The closest contributions are [Comin et al. \(2023\)](#), [Ferrante et al. \(2023\)](#), [Benigno and Eggertsson \(2023\)](#), [Harding et al. \(2023\)](#), and [Rubbo \(2024\)](#). [Comin et al. \(2023\)](#) develop a multisector, open-economy New Keynesian model with exogenous, occasionally binding capacity constraints, and find that these constraints explain about half of the goods-sector inflation surge in 2021–2022. [Ferrante et al. \(2023\)](#) study how the pandemic demand rotation from services to goods generates inflation in the presence of convex labor reallocation costs, amplified by the greater price flexibility of goods. We differ in emphasizing supply-side frictions directly and in disciplining their magnitude by labor-market evidence rather than by goods-versus-services reallocation. [Harding et al. \(2023\)](#) propose a nonlinear Phillips curve via a Kimball quasi-kinked demand aggregator that is flat at low inflation and steepens at high inflation, jointly accounting for the missing deflation of the Great Recession and the post-COVID surge. [Benigno and Eggertsson \(2023\)](#) also generate a nonlinear Phillips curve, but via a search-and-matching microfoundation in which the slope steepens once labor-market tightness  $V/U$  crosses a high Beveridge threshold. [Rubbo \(2024\)](#) develops a multi-sector New Keynesian framework with heterogeneous primary factors and decomposes post-COVID inflation using cross-sectional relative prices, finding that industry-specific shocks dominate the 2020 phase while aggregate factors dominate the post-2021 surge. We build most closely on her insights and offer a different reading of the post-2021 period through the lens of our structural mechanism: a substantial share of the post-2021 surge runs through the same capacity wall whose presence the dispersion already documents.<sup>6</sup>

---

<sup>6</sup>Several other contributions are also relevant. [Moscarini and Postel-Vinay \(2023\)](#) study the same episode from the worker side, documenting how job-ladder dynamics shaped pandemic-era labor-market reallocation. [Forsythe et al. \(2022\)](#) provide complementary worker-side evidence on the COVID recovery, decomposing the missing workers into recalls, retirements, and across-industry reallocation. [di Giovanni et al. \(2022\)](#) extend the cross-sector reallocation logic to global value chains.

As for the steepening of the Phillips curve, using firm–product-level Belgian data, [Gagliardone et al. \(2025\)](#) estimate a cost-based Phillips-curve slope nearly an order of magnitude larger than gap-based estimates, showing that the reduced-form flatness of  $\kappa$  reflects a low elasticity of marginal cost to the output gap in moderate-inflation regimes. This finding is relevant because it shows that the cost-price pass-through our model relies on is alive and well. In a complementary contribution, [Cerrato and Gitti \(2022\)](#) develop an empirical strategy for estimating the Phillips curve from cross-city U.S. variation and find substantial post-COVID steepening.

Our analysis is complementary to linear VAR-based decompositions of the inflation surge ([Bernanke and Blanchard, 2025](#); [Giannone and Primiceri, 2024](#); [Arias et al., 2026](#)). The capacity wall provides a nonlinear structural mechanism consistent with the supply-shortage channel emphasized in that work, and the parameterization of the New Keynesian elements is explicitly chosen to be consistent with the flat, anchored pre-pandemic consensus that [Bernanke and Blanchard \(2025\)](#) largely confirms for this episode.

The remainder of the paper is organized as follows. Section 2 describes the data, accounting framework, and reduced-form evidence. Section 3 presents the multisector New Keynesian model. Section 4 describes the calibration. Section 5 presents the quantitative decomposition and counterfactual exercises. Section 6 concludes.

## 2 Empirical Evidence

Our empirical analysis combines seven sources at annual frequency over 2017–2024, with 2017–2019 as the pre-COVID reference window: (i) the Bureau of Economic Analysis (BEA) 2017 detail-level supply-use tables, which provide input–output linkages and final-demand weights;<sup>7</sup> (ii) the Quarterly Census of Employment and Wages (QCEW) at the 6-digit NAICS level; (iii) the Bureau of Labor Statistics (BLS) International Price Program (IPP) for imported intermediate-input prices; (iv) BEA gross-output industry

---

<sup>7</sup>The BEA 2017 detail benchmark contains 402 industries; we aggregate to 357 to maintain a consistent classification across the 2022 NAICS revision.

deflators; (v) detail-level BLS Producer Price Indices; (vi) the BEA/BLS KLEMS panel, used to estimate the labor–materials elasticity of substitution; and (vii) the COVID-19 Job-Exposure Matrix (JEM) of [Oude Hengel et al. \(2021\)](#). A complete source list is in Online Appendix [A.9](#).

To integrate these data sources, we construct our own aggregate inflation measure from BEA detail-commodity gross-output deflators, weighted by fixed 2017 PCE expenditure shares expressed in producer prices. To isolate the portion of inflation generated inside the networked market economy, we exclude four blocks that require specialized treatment in the structural model—real estate, oil/petroleum/mining, finance/insurance, and government—collectively the *exogenous* (EXO) block. We also exclude these sectors in the later model, where they are treated as exogenous price processes. The resulting measure of inflation is referred to throughout as *network gross-output* inflation, in short **N-GO inflation**, while **GO inflation** is the headline that retains the four exogenous blocks. Our measure differs from BEA’s published PCE deflator and exhibits less persistence after inflation’s peak in 2022. See Online Appendix [A.1.5](#) for a more detailed comparison.

## 2.1 Accounting framework

The goal of this framework is to integrate the above data sources into a fully consistent network-based representation of inflation propagation. To this end, we use a bare-bones version of our model that comprises three key ingredients. First, each industry  $i$  combines labor  $N$  and an intermediate-materials bundle  $m$  via a CES technology with constant elasticity  $\rho$ ,

$$Y_i = Z_i \left( \delta_i^{1/\rho} N_i^{(\rho-1)/\rho} + (1 - \delta_i)^{1/\rho} m_i^{(\rho-1)/\rho} \right)^{\alpha_i \rho / (\rho-1)}, \quad (2)$$

where  $\alpha_i$  governs returns to variable inputs and the productivity term  $Z_i \equiv z_i k_i^{1-\alpha_i}$  bundles disembodied productivity with predetermined capital. Throughout we treat  $k_i$  as comprising physical and intangible capital generating market power, paid for in earlier periods and predetermined at the horizons we study. Materials are a CES bundle of domestic and imported intermediates whose price index enters

cost minimization as the price of  $m$ . Second, conditional on  $Z_i$  and factor prices, inputs solve the static cost-minimization problem. Third, we allow output prices to diverge from minimized unit cost for (here) unspecified reasons and define the *margin residual* as the ratio of revenue to total cost,  $M_i \equiv P_i Y_i / C_i(Y_i)$ . Taking logs of  $M_i$  and of the cost-minimizing production function and substituting out productivity yields the industry-level identity:

$$\Delta \log P_i = \Delta \log M_i + b_i \Delta \log W_i - \Delta \log(Y_i/N_i) + \sum_j A_{ij} \Delta \log P_j + d_i \sum_c (1 - \zeta_{ci}) \omega_{ci} \Delta \log \tilde{P}_c, \quad (3)$$

where  $b_i \equiv s_i^N + (1 - s_i^N)\rho$  and  $d_i \equiv (1 - s_i^N)(1 - \rho)$  are CES-corrected weights, with  $s_i^N$  the labor share of variable cost. The matrix  $A \equiv D(\zeta \odot \Omega)^\top$  encodes first-order network propagation:  $A_{ij}$  is the share of one dollar of industry- $i$  output spent on domestic intermediates from supplier  $j$ , corrected for  $\rho$ .

Stacking (3) across industries and unrolling the supplier-cost term recursively yields the *Leontief decomposition* of inflation:

$$\Delta \log P = L [\Delta \text{MARGIN} + \Delta \text{ULC} + \Delta \text{IMP}], \quad (4)$$

where  $L \equiv (I - A)^{-1}$  is the Leontief inverse,  $\Delta \text{ULC} \equiv b \Delta \log W - \Delta \log(Y/N)$  is unit-labor-cost growth, and  $\Delta \text{IMP}$  collects imported-intermediate cost growth and  $\Delta \text{MARGIN} \equiv \Delta \log M$ , where when we drop subscript  $i$  we mean a vector over  $i$ . Premultiplying by the PCE-weight vector  $w \equiv \Theta^\top c$  obtains the aggregate inflation measure  $w^\top \Delta \log P$ .<sup>8</sup>

Of the structural parameters, only  $\rho$  does not drop out under first-order approximation, as mentioned above. We estimate  $\hat{\rho} = 0.39$  (95% CI: [0.06, 0.72]) on the KLEMS panel. This is conservative relative to [Atalay \(2017\)](#) and [Boehm et al. \(2019\)](#), who place  $\rho$  closer to Leontief, and our results strengthen as  $\rho$  falls. Details are in Online Appendix [A.1.4](#).

The cross-industry evidence below uses two complementary attributions of (3). The *L-propagated* (or *geo-*) attribution of any cost component  $\mathcal{X} \in \{\text{MARGIN}, \text{ULC}, \text{IMP}\}$ ,  $\mathcal{X}_i^L \equiv (L \mathcal{X})_i$ , unrolls supplier-cost

---

<sup>8</sup>Full derivation in Online Appendix [A.1](#).

pass-through into upstream primitives and is consistent with the aggregate identity (4). The **single-pass** attribution keeps the buyer’s own contribution alongside a single round of materials pass-through, so that

$$\Delta \log P_i = \text{MARGIN}_i + \text{ULC}_i + \text{IMP}_i + \text{MAT}_i, \quad \text{MAT}_i \equiv \sum_j A_{ij} \Delta \log P_j.$$

The first is network-aware and unrolls cost pressure into its primitive exogenous sources  $\mathcal{X}$ ; the second is useful for tracking how that pressure propagates one supplier-step at a time.

## 2.2 Summary statistics

Before we get to the main results, we present unconditional moments.

Table 1 shows the Leontief decomposition of aggregate N-GO inflation over 2017–2023. The post-COVID episode unfolded in two compositionally distinct phases. **Act I** (2019–2021) was a +5.2 percentage-point (henceforth pp) cumulative N-GO inflation event in which margins accounted for +3.1 pp (59% of the total) while unit-labor cost fell 1.2 pp in 2021 on the back of a productivity rebound. **Act II** (2021–2023) was a +6.7 pp event of opposite composition: unit-labor cost contributed +6.3 pp (93%) while margins contributed essentially zero on net. The direct-attribution panel shows that one-pass supplier-cost pass-through accounted for roughly 24–30% of each act, consistent with significant within-network propagation of the accumulated cost pressure.

Table 2 performs the same decomposition for the two halves of the BEA-357 universe ranked by predetermined JEM exposure  $E_i$  and split at the median of 2019 employment. The HIGH half (37 industries with  $E_i \geq 1.65$ ) accounts for 69% of Act I and 60% of Act II N-GO inflation. The margin component dominates in Act I within the HIGH half and reverses modestly in the LOW half; ULC dominates Act II in both halves but more sharply in LOW. Inflation was thus disproportionately concentrated in high-contact industries throughout both acts. Two features stand out. First, exposed industries show both margin movements and strong price inflation. Second, part of this dynamic is absorbed by—and partly fueled by—a sudden rise in labor productivity that subsequently slows down. The well-established inter-

**Table 1:** Leontief decomposition of U.S. inflation, 2017–2023.

	Cumulative			Annual			
	2017–19	2019–21	2021–23	2020	2021	2022	2023
<b>GO inflation</b>	+4.6	+7.2	+10.0	+1.6	+5.6	+7.1	+2.9
– Government	+0.2	+0.4	+0.4	+0.1	+0.2	+0.3	+0.1
– Oil & petroleum	+0.1	+0.3	+0.2	–0.4	+0.7	+0.5	–0.3
– Real estate	+1.0	+0.8	+2.0	+0.4	+0.4	+0.9	+1.1
– Finance	+0.6	+0.5	+0.7	+0.2	+0.4	+0.4	+0.3
<b>N-GO inflation</b>	+2.7	+5.2	+6.7	+1.2	+4.0	+5.1	+1.7
<i>Panel A. Network-propagated attribution (L-decomposition); brackets: percent share of N-GO inflation</i>							
Margins	+0.2 (6)	+3.1 (59)	+0.3 (5)	–1.9	+4.9	+0.3	+0.0
ULC	+2.5 (92)	+2.0 (38)	+6.3 (93)	+3.1	–1.2	+4.6	+1.7
+ Wages	+4.5	+8.4	+4.8	+4.2	+4.1	+2.6	+2.2
– Productivity	–2.0	–6.4	+1.5	–1.1	–5.3	+2.0	–0.5
Imports	+0.0 (2)	+0.2 (3)	+0.1 (2)	–0.0	+0.2	+0.2	–0.0
<i>Panel B. Direct attribution (single-pass); brackets: percent share of N-GO inflation</i>							
Margins	+0.3 (13)	+2.3 (44)	+0.3 (5)	–1.2	+3.5	+0.2	+0.1
ULC	+1.5 (56)	+1.6 (30)	+4.3 (64)	+2.4	–0.9	+3.1	+1.2
+ Wages	+3.1	+5.7	+3.4	+3.0	+2.8	+1.8	+1.5
– Productivity	–1.6	–4.2	+1.0	–0.5	–3.6	+1.3	–0.4
Imports	+0.0 (1)	+0.1 (2)	+0.1 (1)	+0.0	+0.1	+0.1	–0.0
Supplier-cost pass-through	+0.8 (31)	+1.3 (24)	+2.0 (30)	+0.0	+1.2	+1.6	+0.4
<i>Memo: EXO sectors</i>	+1.9	+2.0	+3.3	+0.3	+1.7	+2.1	+1.2

*Notes:* All entries in percentage points. GO inflation is the fixed-weight aggregate of BEA detail-level (357-industry SUT) gross-output deflators with industry weights  $w_i = (\Theta^\top c)_i$  constructed from 2017 PCE final demand  $c$  and the I–O commodity-composition matrix  $\Theta$ . N-GO inflation excludes four exogenous blocks (real estate, oil/petroleum/mining, finance/insurance, government), reported in the memo row *EXO sectors*; by construction N-GO + EXO = GO. Panel A decomposes N-GO inflation via the Leontief inverse  $L = (I - A)^{-1}$  as in equation (4): every dollar of inflation is traced back to its ultimate primitive cost. Panel B reports the single-pass attribution, which retains supplier-cost pass-through  $w^\top A \Delta \log P$  as its own line item rather than unrolling it via  $L$ . Both panels sum identically to the N-GO inflation row. Wages and Productivity are sub-components of ULC.

pretation of this productivity swing is that it reflects layoffs of the least productive workers in 2020 and their subsequent rehiring during the recovery. Our model later implies that this dynamic contributed to the persistence of inflation.

**Table 2:** N-GO inflation split by predetermined JEM exposure, 2017–2023.

	Cumulative			Annual			
	2017–19	2019–21	2021–23	2020	2021	2022	2023
<b>N-GO inflation, total</b>	+2.7	+5.2	+6.7	+1.2	+4.0	+5.1	+1.7
<i>Panel A. HIGH JEM exposure, half employment in 2019; brackets: percent share of group N-GO</i>							
<b>Group N-GO contribution</b>	+1.7	+3.6	+4.0	+1.0	+2.6	+3.0	+1.0
Margins	+0.5 (29)	+2.4 (67)	+1.7 (42)	−1.1	+3.5	+1.4	+0.3
ULC	+1.2 (71)	+1.2 (33)	+2.3 (57)	+2.1	−1.0	+1.5	+0.8
+ Wages	+2.8	+5.3	+3.0	+2.6	+2.7	+1.7	+1.3
− Productivity	−1.6	−4.1	−0.7	−0.4	−3.7	−0.2	−0.5
Imports	+0.0 (0)	+0.1 (3)	+0.1 (2)	−0.0	+0.1	+0.1	−0.0
<i>Panel B. LOW JEM exposure, other half employment in 2019; brackets: percent share of group N-GO</i>							
<b>Group N-GO contribution</b>	+1.0	+1.5	+2.6	+0.2	+1.3	+1.9	+0.6
Margins	−0.2 (−20)	+0.6 (40)	−1.4 (−54)	−0.8	+1.4	−1.1	−0.3
ULC	+1.2 (120)	+0.8 (53)	+3.9 (150)	+1.0	−0.2	+2.9	+1.0
+ Wages	+1.7	+3.0	+1.7	+1.6	+1.3	+0.8	+0.9
− Productivity	−0.5	−2.2	+2.2	−0.7	−1.5	+2.1	+0.1
Imports	+0.0 (0)	+0.1 (7)	+0.1 (4)	+0.0	+0.1	+0.1	−0.0
<i>Memo: Unclassified (no <math>E_i</math>)</i>	+0.0	+0.1	+0.1	+0.0	+0.1	+0.1	−0.0

*Notes:* All entries in percentage points and aggregated with industry-side PCE weights  $w_i = (\Theta^\top c)_i$ , not renormalized. Industries are ranked by predetermined JEM exposure  $E_i$  and split at the median of 2019 N-GO employment: Panel A collects the 37 industries above the cutoff  $E_i \geq 1.65$ ; Panel B collects the remaining 293 industries with valid  $E_i$ . Cumulative panels follow the network-propagated ( $L$ -decomposition) attribution of Table 1, Panel A. Group contributions sum to total N-GO inflation up to the unclassified memo row.

## 2.3 Main results

This section focuses on establishing an empirical link between COVID exposure, the COVID-induced labor shock, and subsequent price action. We first describe the econometric methodology and then discuss our key findings.

**Methodology.** The cross-industry sample contains the 321 endogenous BEA-357 industries remaining after dropping the 16 industries in the exogenous block and industries with missing JEM data. All regressions are estimated for the peak-inflation year 2022 as a cumulative change from 2019. Predetermined own exposure is the JEM index  $E_i$  of Oude Hengel et al. (2021), constructed from the Danish country-wide panel—mapped from ISCO-08 to SOC and aggregated to BEA-357 using pre-COVID

OEWs occupation shares. The Danish panel offers the highest classification precision and has been validated against the population-level Danish CPR/BIDS employment register and infection rates. The construction of the index is documented in Online Appendix A.2.1. Online Appendix Table 11 re-estimates the same regression with the UK panel.

We define supplier-network exposure as the strict supplier projection  $\mathcal{E}_i \equiv (WE)_i$ , with  $W \equiv A - \text{diag}(A)$ . The capital-script convention extends to employment-flow variables:  $\Delta N_i^{\text{gap}}$  denotes the magnitude of the 2020 own employment collapse (in pp),  $\Delta N_i^{\text{rec}}$  cumulative own recovery from 2021 through  $T$ , and  $\Delta \mathcal{N}_i^{\text{rec}}$  its supplier-weighted analog constructed using the direct supplier kernel  $W$ .<sup>9</sup>

The main regression we use is a cross-industry specification of the form:

$$y_i = \alpha_{s(i)} + \alpha_{c(i)} + \pi^E E_i + \pi^{\mathcal{E}} \mathcal{E}_i + \pi^g \Delta N_i^{\text{gap}} + \pi^r \Delta N_i^{\text{rec}} + \pi_s^r \Delta \mathcal{N}_i^{\text{rec}} + \varepsilon_i, \quad (5)$$

where  $y_i$  is cumulative L-propagated price growth from 2019 to 2022 in percentage points,  $\alpha_{s(i)}$  are BEA-66 sector fixed effects, and  $\alpha_{c(i)}$  are chain-by-tier-12 cell fixed effects from the supply-network classification (described in the next section and in detail in Online Appendix A.3). Estimation is WLS with industry weights  $w_i$  (N-GO PCE shares, not renormalized), and standard errors are clustered at the BEA-66 sector level.

We identify the operation of supply forces on prices through the *joint sign* of two coefficients in equation (5):  $\pi^E > 0$  (Fact 1: high-contact-intensity industries had more cumulative 2019–2022 inflation) and  $\pi^r < 0$  (Fact 2: conditional on exposure and the 2020 collapse, industries that rebuilt employment faster had *less* cumulative inflation). The negative recovery sign is the load-bearing identifying restriction: it sign-discriminates the wall mechanism from the leading alternative explanations of the cross-industry inflation pattern, as set out in Fact 2.<sup>10</sup>

<sup>9</sup>Formally,  $\Delta N_i^{\text{gap}} = \max(0, -\Delta_{2019-2020} \log \text{Emp}_i) \times 100$  and  $\Delta N_i^{\text{rec}} = (\Delta_{2020-2021} \log \text{Emp}_i + \Delta_{2021-2022} \log \text{Emp}_i) \times 100$ , where  $\Delta_{t_0-t_1} \log \text{Emp}_i$  is the year-on-year log change in employment between  $t_0$  and  $t_1$ .

<sup>10</sup>We do not adopt the 2SLS estimator that instruments the 2020 employment collapse  $\Delta N_i^{\text{gap}}$  with the predetermined JEM exposure  $E_i$ . The IV does work mechanically—giving a strong first stage ( $\hat{\phi} = +9.0$ ,  $|t| = 3.18$  in the saturated specification) and a 2SLS coefficient on the instrumented collapse of about  $+0.6$  ( $|t| = 2.3$ , Anderson-Rubin 95% CI bounded and excluding zero) that is consistent in sign and order of magnitude with the structural model’s prediction. But the purity of the exclusion

We summarize our findings as three stylized facts.

**(1) Cross-industry: high-exposure industries inflated more between 2019 and peak inflation**

**in 2022, primarily via the supplier-margin channel.** Table ?? reports the multi-column reduced-form estimation of equation (5). Column (1) reports the bare cross-industry pattern with only the supplier-network exposure as a control: the coefficient on own JEM exposure is  $\hat{\pi}^E = +4.33$  ( $|t| = 2.39$ ), significant at 5%. Adding the 2020 employment collapse  $\Delta N_i^{\text{gap}}$  (column 2) leaves the exposure coefficient at  $+5.04$  ( $|t| = 2.58$ ). The saturated specification (column 3) adds own and supplier recovery margins; the exposure coefficient is  $\hat{\pi}^E = +4.63$  ( $|t| = 2.54$ ). A P90–P10 spread in  $E_i$  ( $\approx 0.70$ ) thus translates into roughly 3.2 pp more cumulative inflation within the same BEA-66 sector and chain-tier-12 cell.

The labor-market relation between predetermined exposure and the realized 2020 employment collapse is tight: a one-unit increase in  $E_i$  corresponds to about 9.0 pp more 2020 employment destruction ( $\hat{\phi} = +9.00$ ,  $|t| = 3.18$ , panel C of Table ??). The P90–P10 exposure spread thus maps to roughly 6.3 pp larger 2020 collapse. Combining the inflation and labor-collapse mappings, the cross-industry pattern is quantitatively consistent with a labor-channel elasticity of about 0.5 pp of cumulative 2019–2022 inflation per percentage point of 2020 employment destruction.

We next characterize how this pressure propagates through the supply network by decomposing the one-pass material-input cost  $\text{MAT}_i \equiv \sum_j A_{ij} \Delta \log P_j$  into the supplier-side components of the accounting identity. The structural identity is linear:  $\text{MAT}_i = \text{MARGIN}_i^S + \text{WAGE}_i^S + (Y/N)_i^S + \text{IMP}_i^S$ , so the  $E_i$  coefficient on  $\text{MAT}_i$  decomposes additively across the four supplier-source components. Table ?? reports the result: a one-unit increase in  $E_i$  raises the buyer’s one-pass material-input cost by  $\hat{\pi}_{\text{MAT}}^E = +0.79$  pp ( $|t| = 2.66$ , significant at 1%). Of this total,  $+0.72$  pp comes from *supplier margins* (significant at 5%, accounting for roughly 92% of the MAT exposure-loading); supplier wages, productivity, and imported

---

restriction is unclear in this case: contact intensity plausibly affects 2022 prices through channels other than the labor collapse—in particular, pent-up demand released into capacity-constrained sectors, regulatory shutdowns, hazard-pay and PPE cost-push, and the “essential industry” distinction. Some contamination via these channels is likely. The supply-side fingerprint identification used here does not require the exclusion restriction and identifies the wall mechanism directly from a sign discrimination on the recovery coefficient. The later model plays the role of quantifying the effect.

**Table 3:** Main cross-industry evidence, cumulative 2019–2022: reduced form and 2SLS.

	$\Delta \log P_{19 \rightarrow 22}$ (price)		$\Delta \text{MARGIN}_{19 \rightarrow 22}^L$ (margin)	
	(1) RF	(2) IV	(3) RF	(4) IV
$E_i$ (own JEM)	4.63** (2.54)	(instr.)	-0.10 (0.03)	(instr.)
$\mathcal{E}_i$ (supplier JEM)	0.76 (0.09)	21.04 (1.29)	-32.98*** (3.04)	-33.42** (2.55)
$\Delta N_i^{\text{gap}}$ (own collapse)	0.03 (0.42)	0.60** (2.31) [+0.17, +1.63]	0.76*** (8.76)	0.75* (1.80) [-0.16, +1.95]
$\Delta N_i^{\text{rec}}$ (own recovery)	-0.13** (2.08)	-0.33** (2.33)	-0.79*** (8.10)	-0.79*** (4.55)
$\Delta \mathcal{N}_i^{\text{rec}}$ (supplier recovery)	-1.39 (1.24)	-4.41** (2.20)	0.19 (0.12)	0.25 (0.09)
First-stage $F$	—	22.00	—	22.00
$R^2$	0.868	0.800	0.881	0.881
Observations	321			
Fixed effects	BEA-66 sector + chain-tier-12 cell			

*Notes:* Dependent variables in percentage points. Reduced-form (RF) columns estimate equation (5); 2SLS columns instrument  $\Delta N_i^{\text{gap}}$  with  $E_i$  as in equations (??)–(??).  $\Delta \text{MARGIN}_{19 \rightarrow 22}^L$  is the cumulative  $L$ -propagated margin contribution from the accounting framework. Estimation is WLS with N-GO PCE-share weights  $w_i$ ; standard errors clustered at the BEA-66 sector level. Absolute  $t$ -statistics in parentheses beneath each coefficient. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Brackets next to the 2SLS  $\Delta N_i^{\text{gap}}$  row report the Anderson–Rubin weak-IV-robust 95% confidence interval.

inputs together account for the remaining 8% and are individually insignificant.

These results imply that the pressure both propagates through the network and partly originates from supplier margins, with buyer margins absorbing part of it on its way downstream. This pattern is consistent with capacity-constrained supplier firms widening margins as they ration scarce capacity, and with downstream buyers absorbing the resulting material-input price pressure, extending the inflationary episode over time and flattening its peak.

**(2) Conditional on exposure, faster-recovering industries inflated less—opposite to a pure demand-pull story.** The own-recovery coefficient in column (3) of Table ?? is  $\hat{\pi}^r = -0.13$  ( $|t| = 2.08$ ), significant at the 5% level. The supplier-recovery coefficient is  $\hat{\pi}_s^r = -1.39$ , right-signed but imprecisely estimated. A one-percentage-point faster cumulative 2020–2022 employment recovery is thus associated with about 0.13 pp less cumulative 2019–2022 inflation within the same sector and chain-tier cell.

**Table 4:** Supplier-source decomposition of one-pass material-cost pass-through, cumulative 2019–2022.

Dependent variable	$\hat{\beta}_g$ on $\Delta N_i^{\text{gap}}$	Share of MAT coefficient
<i>Memo:</i> total price $\Delta \log P_i$	0.605** (2.31) [+0.22, +1.62]	—
$\Delta \text{MAT}_i = \sum_j A_{ij} \Delta \log P_j$	0.123** (2.53) [+0.04, +0.30]	100.0%
<i>Decomposition of <math>\Delta \text{MAT}_i</math> into supplier-source components:</i>		
Supplier margins, $\Delta \text{MARGIN}_i^S$	0.132*** (2.73)	107.2%
Supplier wages, $\Delta \text{WAGE}_i^S$	0.029 (1.08)	23.3%
Supplier productivity, $\Delta(Y/N)_i^S$	−0.037 (0.89)	−29.7%
Supplier imported inputs, $\Delta \text{IMP}_i^S$	−0.001 (0.31)	−0.8%
First-stage $F$		22.00
Observations		321
Fixed effects		BEA-66 sector + chain-tier-12 cell

*Notes:* Each row reports a separate 2SLS regression of the named left-hand-side variable on the instrumented own employment collapse  $\widehat{\Delta N_i^{\text{gap}}}$ , with the same exogenous controls and fixed effects as Table 3. The MAT coefficient equals the sum of the four supplier-source components by construction (+0.123 = +0.132 + 0.029 − 0.037 − 0.001, rounded). WLS with  $w_i$  weights; BEA-66 cluster-robust standard errors. Absolute  $t$ -statistics in parentheses beneath each coefficient. Brackets next to the two headline rows report the Anderson–Rubin weak-IV-robust 95% confidence interval. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

If inflation were driven by demand pull relative to 2019, we should see  $\pi^r > 0$ : industries that re-hire fastest are the ones whose suppressed 2020–2021 demand is released most fully in 2022, producing larger price spikes. The reason is that an outward shift in demand should trace out an upward-sloping supply curve, so that quantities and prices co-move positively. They do not.

On the other hand, a pure cost-push channel from health adaptation (PPE, hazard pay, energy costs), or from some other source, is sign-ambiguous on the recovery margin and timing-bounded to 2020–2021. The data align with the wall prediction and against the pent-up-demand prediction.

A natural worry here is that high-teleworkability industries both recovered fast and inflated less for reasons unrelated to capacity constraints (e.g., remote operations involve lower marginal costs and looser physical capacity limits). Column (4) of Table ?? addresses this directly by adding the [Dingel and Neiman \(2020\)](#) “must-be-in-person” share  $T_i$  and its supplier-network propagation  $\mathcal{T}_i$  as exogenous

controls. This leaves the recovery coefficient essentially unchanged at  $\hat{\pi}^r = -0.13$  ( $|t| = 2.15$ ). After netting out teleworkability, slower-recovering industries still had more cumulative inflation by the same amount. The supply-side fingerprint is therefore not driven by the teleworkability channel.

In contrast, the above pattern is the fingerprint of supply being a drag: a deeper 2020 employment collapse combined with a slower 2021–2022 recovery generates persistent supplier-side capacity tightness, which propagates downstream as a material-input price increase (Fact 1) and which, conditional on the initial collapse, is more severe for industries that rebuild slowest.

**(3) Heterogeneous supplier exposures drove the post-COVID rise in within-sector price dispersion.** Facts 1 and 2 identify a conditional-mean effect: more exposed industries (and industries with more exposed suppliers) inflate more on average, and slower-recovering industries inflate more conditional on exposure. We complement this with a dispersion result. The *level* of JEM exposure is not, by itself, a meaningful source of residual cross-industry variance: within BEA-66 sector and chain-tier cells,  $E_i$  and  $\mathcal{E}_i$  jointly explain only 0.5% of the variance of cumulative 2019–2022 price growth, since roughly 80% of that variance is absorbed across BEA-66 sectors. What matters for dispersion is the within-sector *heterogeneity* of supplier exposure across constituent industries.

We show this by regressing the BEA-66 sector-level change in upside cross-industry price dispersion between the cumulative 2019–2022 window and the 2017–2019 pre-COVID placebo window on the sector-level dispersions of own and supplier JEM exposure. Specifically, define the within-sector upside price-growth dispersion as  $\sigma_{s,p}^+ \equiv \text{sd}_{i \in s}(\max(0, \Delta \log P_{i,p}))$  for  $p \in \{17 \rightarrow 19, 19 \rightarrow 22\}$ , weighted by industry 2019 employment, and let  $\sigma(E_i)_s$  and  $\sigma(\mathcal{E}_i)_s$  denote the corresponding employment-weighted dispersions of  $E_i$  and  $\mathcal{E}_i$  within sector  $s$ . We estimate

$$\Delta \sigma_s^+ = \alpha + \beta_1 \sigma(E_i)_s + \beta_2 \sigma(\mathcal{E}_i)_s + \varepsilon_s \quad (6)$$

by WLS on the 38 BEA-66 sectors with at least two surviving industries (industry-side filter: drop the 16 exogenous-block sectors and industries without JEM exposure or 2019 employment, leaving 330

industries), weighted by total sector 2019 employment, with HC1 standard errors. Table 5 reports the results.

**Table 5:** Within-sector heterogeneity in supplier exposure and the 2019–2022 explosion in within-sector upside price dispersion.

Dependent variable: $\Delta\sigma_s^+ = \sigma_{s,19\rightarrow22}^+ - \sigma_{s,\text{plac}}^+$	
$\sigma(E_i)_s$ (own JEM heterogeneity, within sector)	-1.72 (0.74)
$\sigma(\mathcal{E}_i)_s$ (supplier JEM heterogeneity)	40.33** (2.28)
Constant	0.80 (0.90)
$R^2$	0.134
Observations	38 BEA-66 sectors

*Notes:* Unit of observation is a BEA-66 sector. Dependent variable is the change in employment-weighted within-sector upside price dispersion between the cumulative 2019–2022 window and the pre-COVID placebo. Treatment window cumulates  $g_{2020}^P + g_{2021}^P + g_{2022}^P$ ; placebo cumulates  $g_{2018}^P + g_{2019}^P$  (data limit). Right-hand-side variables are within-sector employment-weighted standard deviations of own JEM exposure  $E_i$  and supplier-network JEM exposure  $\mathcal{E}_i$ . Estimation is WLS, weighted by total sector 2019 employment, with HC1 standard errors. Absolute  $t$ -statistics in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

The supplier-heterogeneity coefficient is  $\hat{\beta}_2 = +40.33$  ( $|t| = 2.28$ ), significant at 5%, while own exposure heterogeneity is small and insignificant ( $\hat{\beta}_1 = -1.72$ ,  $|t| = 0.74$ ). Two sectors that differ by one standard deviation in within-sector supplier-exposure dispersion differ by roughly 4 pp in the increase of within-sector upside inflation dispersion between the placebo and 2019–2022 windows. The mechanism that drives within-sector dispersion is shock geometry: sectors containing industries whose suppliers are heterogeneously exposed to the COVID labor shock fan out in the post-COVID window, while sectors that are uniform on this dimension do not. This evidence connects the patterns shown in Figure 2.

The combination of the three facts—high-exposure industries inflated more (Fact 1), the price pressure was concentrated in industries whose recovery was slowest (Fact 2, the discriminating sign test), and within-sector dispersion exploded where supplier exposure was heterogeneous (Fact 3)—shows that the COVID contraction mattered, and that the numbers are economically significant. It is not possible to quantify these effects in aggregate from the reduced form alone. We now turn to the model that formalizes this story and quantifies the total effect. We return to these regressions in the model section.

### 3 Model

Time is discrete and indexed by  $t = 0, 1, 2, \dots$ . The pre-COVID steady state (2019) is  $t = 0$ ; the COVID shock year (2020) is  $t = 1$ ; and the recovery years 2021–2023 are  $t = 2, 3, 4$ . The remaining periods feature no shocks and trend employment, allowing the economy to converge back to its nominal anchor. We solve the model as a nonlinear perfect-foresight system under externally imposed expectations (discussed below). The COVID shock is unanticipated on impact and contracts employment and capacity. The model is saturated in that a range of wedges ensure it fits the data exactly. It is designed to answer an accounting question: how much of a given observation is attributable to exogenous wedges versus the endogenous CW mechanism.

The model economy comprises twenty-four endogenous production nodes and four exogenous sectors, referred to as the “EXO” block. These are: Real Estate (HO); Finance, Insurance, and Government (FO); Oil, Petroleum, and Metal Ores (EN); and imported intermediate inputs (IM).<sup>11</sup> EXO sectors enter as price paths for material inputs from these sectors and, with the exception of IM, as final-consumption goods in the consumer aggregator, each with its own PCE weight. Imports enter only on the materials side:  $P_{i,t}^{\text{IMP}}$  is a per-tier intermediate-input price block—twelve series, one per buyer tier—with no direct PCE weight.

#### 3.1 Endogenous network structure

The upper layer of the endogenous block is organized into four supply chains and three vertical tiers within each chain. The chains are constructed from the 341 endogenous BEA detail industries via spectral bisection on the cumulative supplier-path operator  $M = A + A^2 + A^3$ , where  $A$  is the 2017 BEA input-share matrix.<sup>12</sup> The goal is to group industries that are mutually linked, so as to maxi-

---

<sup>11</sup>HO combines residential and commercial real estate. The combined gross-output deflator in our setup therefore rises much less than CPI rents during the pandemic because commercial real estate rents were flat or declining while residential rents rose sharply. This results in some downward bias when we use PCE weights.

<sup>12</sup> $A_{ik}$  is the share of industry  $i$ 's intermediate-input bill purchased from supplier industry  $k$ , so  $A$  is a directed cost-share adjacency matrix. The term  $A^q$  collects supplier paths of length  $q$ , weighted by the product of cost shares along the path. The depth-three operator is motivated by the reduced-form evidence in Section 2, where the supplier-exposure signal peaks

mize within-chain trade in intermediate inputs. We achieve this by recursively splitting the industry graph into two balanced halves that minimize trade across the cut using standard spectral methods, applying the procedure twice to obtain exactly four chains. The full algorithm is described in Online Appendix A.3. The four chains are  $C_1$  (Retail & Telecom) (43 industries),  $C_2$  (Consumer Services) (77 industries),  $C_3$  (Professional & Hospitality) (101 industries), and  $C_4$  (Manufacturing) (120 industries).

Within each chain, industries are sorted and split by their *finalness*—the share of output sold directly to PCE in the 2017 I–O data (distinct from the usual notion of “final goods”). The consumer-facing tier  $h$  contains industries with a finalness index above 70 percent; the remaining industries are evenly split into a middle tier  $m$  and an upstream tier  $l$ . This yields twelve endogenous tiers: four chains times three vertical positions. Each tier has its own price, employment, production technology, demand shifter, and pricing equation, and the tiers are linked by a  $12 \times 12$  input–output matrix. We provide more detail in the next section.

The four chains are economically distinct:

$C_1$  is a *broad services* chain—94% of value added in services (64% B2B, 30% consumer-facing), less than 5% in manufacturing.<sup>13</sup> Retail trade dominates downstream; B2B professional services (management of companies, advertising, employment services, consulting, custom programming) sit upstream; telecom, electric power, and information services lie in the middle.

$C_2$  is the *consumer-services and food* chain—81% services (65% consumer-facing, 16% B2B), 14% non-durable manufacturing, 5% agriculture. Hospitals, restaurants, universities, and nursing facilities at the consumer end; nondurable-goods wholesalers, grocery distribution, and truck transportation in the middle; agriculture and food processing upstream.

$C_3$  is the *healthcare, hospitality, and chemicals* chain—70% services (35% consumer-facing, 35% B2B),

---

at the third upstream layer.

<sup>13</sup>Value-added shares are each chain’s 2019 nominal value added in a bucket divided by the chain’s total 2019 value added. Buckets: consumer-facing services (trade, health, leisure, education, other services), B2B services (professional, information, finance, transport), durable manufacturing (NAICS 321, 327, 331–337, 339), nondurable manufacturing (NAICS 311FT, 313TT, 315AL, 322–326), other goods (construction, mining excluding energy, agriculture), and utilities/government (NAICS 22, FEDGOV, SLGOV).

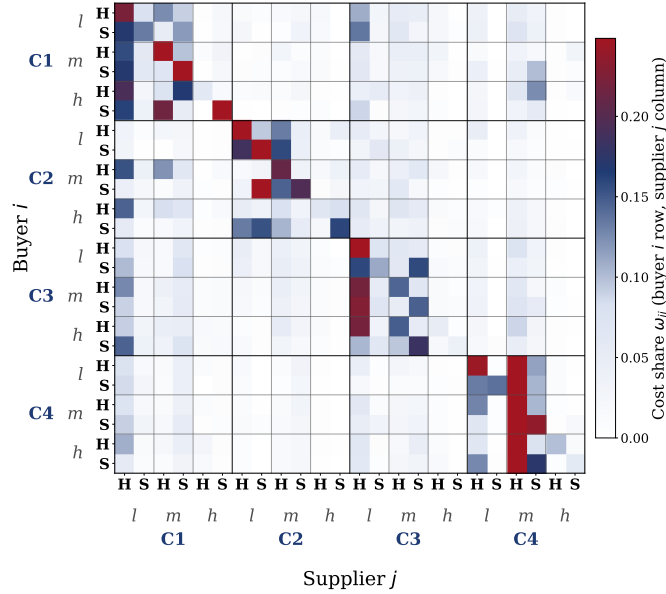
20% nondurable manufacturing, 5% durables. Physician offices, accommodation, air transportation, and pharmaceuticals at the consumer end; basic chemicals, plastics, and petrochemicals upstream—the only chain in which chemicals make up roughly a fifth of value added.

$C_4$  is the *durable-goods and construction* chain—a pure-goods chain: 44% durable manufacturing, 33% construction, mining, and agriculture, 20% consumer-facing services, and only 3% nondurable manufacturing. Auto and truck manufacturing, automotive repair, and household appliances downstream; durable-goods wholesalers, aircraft manufacturing, and primary metals in the middle; construction itself dominates upstream.

Online Appendix [A.3](#) lists representative industries by tier for each chain.

The consumer-facing tier ( $h$ ) and the middle-and-upstream tiers ( $m, l$ ) capture distinct industries along each chain.  $C_1$ 's  $h$  tier is essentially retail trade and telecommunications, while  $m + l$  is overwhelmingly B2B professional services and information.  $C_2$ 's  $h$  tier is contact-intensive consumer services and food processing, with  $m + l$  running from wholesale and freight into agriculture upstream.  $C_3$ 's  $h$  tier is healthcare and hospitality—physician offices, accommodation, air transport, and pharmaceuticals—with  $m + l$  combining engineering and R&D services with basic chemicals upstream.  $C_4$  inverts the pattern: its  $h$  tier is durable manufacturing and automotive repair, while  $m + l$  are intermediate durables, primary metals, and construction. The cross-chain takeaway:  $C_1$  has the cleanest vertical structure (no goods at either end),  $C_4$  is the only chain with goods at the consumer interface, and  $C_2$  and  $C_3$  share services-heavy downstreams but diverge upstream— $C_2$  into agriculture,  $C_3$  into chemicals. The important feature here is that chains extend in the upstream–downstream direction, and that major supply linkages occur within chains rather than across them. This matters for propagation, given how we capture the COVID shock.

**Hit/safe inner layer.** To model the COVID shock, each of the twelve tiers is split into a hit sub-unit and a safe sub-unit. Within a tier, industries are sorted by their realized 2019–2020 employment collapse,  $\Delta \log N_{2019 \rightarrow 2020}$ , and split at the employment-weighted median. The hit group contains the



**Figure 3: Intermediate-input flows across the 24 production sub-units.** Cells show input shares aggregated from the 341 endogenous BEA-NAICS industries to the chain  $\times$  tier  $\times$  hit/safe level. Rows and columns are ordered by chain ( $C_1$ – $C_4$ ), finalness tier ( $l, m, h$ ), and hit/safe status ( $H, S$ ). The block structure reflects within-chain trade and the fact that hit and safe industries within the same tier serve partially different downstream buyers.

industries with larger employment losses; the safe group contains the less damaged industries. This deterministic split gives the final  $12 \times 2 = 24$  production nodes present in the model, all endogenously connected according to the 2017 I–O matrix.

Figure 3 displays the empirical  $24 \times 24$  intermediate-input heatmap after aggregating the 341 BEA-NAICS industries to the chain  $\times$  tier  $\times$  hit/safe classification. The key takeaway is the high concentration of flows within chains and sub-units.

### 3.2 Production and aggregation

We now discuss how goods and services are produced within inner nodes.

Let  $i = 1, \dots, 12$  index chain tiers. Each tier contains two sub-units referred to as production nodes, indexed by  $j \in \{H, S\}$ , where  $H$  denotes the hit node and  $S$  the safe node. Variables with two subscripts—such as  $Y_{ij,t}$ ,  $P_{ij,t}$ ,  $N_{ij,t}$ , and  $m_{ij,t}$ —refer to node  $(i, j)$ ; variables with only the tier index, such as  $Y_{i,t}$  and  $P_{i,t}$ , refer to the tier aggregate of the two nodes. We follow the convention that dropping the hit/safe

subscript indicates the value is common across  $j \in \{H, S\}$ : parameters and shifters such as  $\bar{\phi}_i$ ,  $\mu_{i,t}$ ,  $\xi_{i,t}$ ,  $a_{i,t}^e$ ,  $\gamma_i$ ,  $\omega_{i,k}^{DOM}$ ,  $\omega_i^{IMP}$ , and  $P_{i,t}^{IMP}$  take the same value at the hit and safe nodes of tier  $i$ . The labor weight  $\delta_{ij}$ , capital share  $\alpha_{ij}$ , and node wage shifter  $\tilde{w}_{ij}$  are exceptions: they differ across hit and safe within a tier and are taken directly from the BEA-NAICS aggregation in 2017 to reflect the steady-state employment, materials, and wage composition of each sub-unit's underlying industries.

Within each tier  $i$ , hit and safe node-level outputs are aggregated by a CES bundle with elasticity  $\rho$ :

$$Y_{i,t} = \left[ \sum_{j \in \{H, S\}} s_{ij}^{1/\rho} Y_{ij,t}^{(\rho-1)/\rho} \right]^{\rho/(\rho-1)}, \quad (7)$$

with dual tier price index

$$P_{i,t} = \left[ \sum_{j \in \{H, S\}} s_{ij} P_{ij,t}^{1-\rho} \right]^{1/(1-\rho)}, \quad (8)$$

where  $s_{ij}$  are 2019 within-tier output shares. The implied demand for node  $(i, j)$  output is

$$Y_{ij,t} = s_{ij} \left( \frac{P_{ij,t}}{P_{i,t}} \right)^{-\rho} Y_{i,t}. \quad (9)$$

Each node  $(i, j)$  produces using fixed capital, labor, and a materials bundle. Capital is fixed at its pre-pandemic level, so all variable adjustment occurs through labor and materials. The variable-input bundle is

$$V_{ij,t} = \left[ \delta_{ij}^{1/\rho} N_{ij,t}^{(\rho-1)/\rho} + (1 - \delta_{ij})^{1/\rho} m_{ij,t}^{(\rho-1)/\rho} \right]^{\rho/(\rho-1)}, \quad (10)$$

where  $\delta_{ij}$  is the labor weight and  $\rho$  is the CES substitution elasticity. Elasticity  $\rho$  governs all three CES margins in the model: labor–materials substitution within the variable-input bundle, substitution across intermediate inputs in the materials bundle, and substitution between hit and safe sub-units within a tier. Robustness to allowing these margins to differ is reported in Online Appendix A.7. Gross output is

$$Y_{ij,t} = A_{ij} \exp(a_{i,t}^e) \bar{K}_{ij}^{\alpha_{ij}} V_{ij,t}^{1-\alpha_{ij}}, \quad (11)$$

where  $A_{ij}$  is steady-state TFP,  $\bar{K}_{ij}$  is fixed capital,  $\alpha_{ij}$  is the capital share, and  $a_{i,t}^e$  is the time-varying productivity shifter. By the indexing convention,  $a_{i,t}^e$  is common across hit and safe sub-units, so steady-state hit/safe productivity differences are absorbed by  $A_{ij}$ . This restriction prevents the productivity path from absorbing the hit/safe shock geometry the model is designed to analyze.

The materials bundle of node  $(i, j)$  is a CES aggregate of domestic tier inputs and an import input. Input shares are common across hit and safe sub-units within tier  $i$ , inherited from the aggregated I–O data:

$$m_{ij,t} = \left[ \sum_{k=1}^{12} (\omega_{i,k}^{DOM})^{1/\rho} (m_{ij,k,t}^{DOM})^{(\rho-1)/\rho} + (\omega_i^{IMP})^{1/\rho} (m_{ij,t}^{IMP})^{(\rho-1)/\rho} \right]^{\rho/(\rho-1)}. \quad (12)$$

Because shares and supplier prices are tier-level, the associated materials price index is also tier-level:

$$p_{i,t}^m = \left[ \sum_{k=1}^{12} \omega_{i,k}^{DOM} P_{k,t}^{1-\rho} + \omega_i^{IMP} (P_{i,t}^{IMP})^{1-\rho} \right]^{1/(1-\rho)}. \quad (13)$$

Conditional demand for domestic tier  $k$  from node  $(i, j)$  is

$$m_{ij,k,t}^{DOM} = \omega_{i,k}^{DOM} \left( \frac{p_{i,t}^m}{P_{k,t}} \right)^\rho m_{ij,t}. \quad (14)$$

When  $\rho < 1$ , inputs are complements and cost shocks propagate strongly through the production network. Stacking the domestic cost shares  $\omega_{i,k}^{DOM}$  at the tier level gives the  $12 \times 12$  matrix used in the model; its Leontief inverse summarizes steady-state cost propagation through direct and higher-order supplier links. As estimated in the data section (Section 4), we set  $\rho = 0.39$ .

For each tier  $k = 1, \dots, 12$ , output equals final demand plus intermediate demand from all sub-units:

$$Y_{k,t} = D_k^{ss} \exp(\xi_{k,t}) + \sum_{i=1}^{12} \sum_{j \in \{H,S\}} \omega_{i,k}^{DOM} \left( \frac{p_{i,t}^m}{P_{k,t}} \right)^\rho m_{ij,t}, \quad (15)$$

where  $D_k^{ss}$  is steady-state final demand and  $\xi_{k,t}$  is a time-varying demand shifter. In consumer-facing  $h$  tiers,  $\xi_{k,t}$  captures final-demand rotation during the pandemic; in middle and upstream tiers, it cap-

tures residual demand not mediated by the I–O matrix, such as exports, investment, or government procurement.

Our model is set in partial equilibrium on the demand side: there are no households, no endogenous capital accumulation, and no monetary policy through explicit interest-rate rules. The demand shifters are therefore calibrated wedges. Because  $\xi_{k,t}$  is tier-level, it cannot absorb hit/safe misallocation within tier  $k$ : the relative allocation between hit and safe sub-units is determined endogenously by the CES demand system in equations (7)–(9) and the relative prices generated by the model.

### 3.3 Capacity wall friction

The capacity wall is a rebuilding-cost wedge that operates at the node level. Firms take it as given; it represents the amortized cost of increasing net employment. Formally, the wall premium faced by a firm on the margin is

$$\mathcal{W}_{ij,t} = \gamma_i (dN_{ij,t}^+)^2, \quad dN_{ij,t}^+ = \max\left\{\frac{N_{ij,t} - N_{ij,t-1}}{N_{ij}^{ss}}, 0\right\}. \quad (16)$$

This labor wedge applies when firms rapidly recover or grow above the normal chain-level trend. It is not paid in the steady state.

The quadratic form  $\eta = 2$  is a convenience choice and does not have an independent micro-foundation. We have re-estimated the model under any exponent between linear ( $\eta = 1$ ) and quadratic ( $\eta = 2$ ), recalibrating  $\gamma_i$  analogously in each case to hit the same tightness-anchored rebuilding-cost target. The aggregate decomposition and the qualitative conclusions are essentially unchanged across this range; we report the  $\eta = 2$  baseline as the headline because it cleanly delivers state-dependence of the Phillips slope in the convex-wall direction.

At  $t = 0$ , there is a “COVID shock” that enters as the wedge  $\tau_{ij,t}$  and contracts employment differentially between hit and safe sub-units. Effective labor cost in that period (and only in that period) is thus given

by

$$w_{ij,t}^e = \tilde{w}_{ij} W_t (1 + \mathcal{W}_{ij,t}) \exp(\tau_{ij,t}), \quad (17)$$

where  $W_t$  is the aggregate base wage and  $\tilde{w}_{ij}$  is the sub-unit wage shifter.

### 3.4 Prices and expectations

Each sub-unit sets prices subject to Rotemberg-style quadratic adjustment costs with asymmetric downward rigidity. The pricing parameters differ by tier.

For node  $(i, j)$ , the pricing first-order condition is

$$\bar{\phi}_i (\pi_{ij,t} - \bar{\pi}) Y_{ij,t} + \lambda_i (\pi_{ij,t} - \bar{\pi})^- Y_{ij,t} = \varepsilon_i \left( \mathcal{MC}_{ij,t} - \frac{\varepsilon_i - 1}{\varepsilon_i \mu_{i,t}} \right) Y_{ij,t} + \beta \mathbb{E}_t \left[ \bar{\phi}_i (\pi_{ij,t+1} - \bar{\pi}) + \lambda_i (\pi_{ij,t+1} - \bar{\pi})^- \right] Y_{ij,t+1}. \quad (18)$$

Here  $\pi_{ij,t} \equiv P_{ij,t}/P_{ij,t-1}$  is gross inflation and  $(\cdot)^- \equiv \min\{0, \cdot\}$  captures the asymmetric downward-rigidity term. The parameter  $\bar{\phi}_i$  is the Rotemberg adjustment cost,  $\lambda_i$  governs downward rigidity,  $\varepsilon_i$  is the demand elasticity, and  $\mu_{i,t}$  is an exogenous markup wedge common to nodes within tier  $i$ . The demand elasticity is set to a common value  $\varepsilon_i = 5$  across tiers, which maps to a steady-state markup of  $\varepsilon/(\varepsilon - 1) = 1.25$ . In the calibration,  $\bar{\phi}_i$  is chain-specific for  $h$  tiers and equals  $\bar{\phi}_{ml}$  for  $m$  and  $l$  tiers; similarly,  $\lambda_i = \lambda_h$  in  $h$  tiers and  $\lambda_i = \lambda_{ml}$  in  $m$  and  $l$  tiers.

The exogenous wedge  $\mu_{i,t}$  makes the calibrated model saturated in the spirit of business-cycle accounting (Chari et al., 2007): the wedge backfills the residual price path after the wall parameters have been fixed. To isolate the structural contribution of the wall, we later re-solve the model with the residual log-markup wedge set to zero—equivalently,  $\mu_{i,t} = 1$  in the multiplicative pricing equation—while holding all other parameters fixed. In tables and figures, this normalization is reported compactly as ‘Model  $\mu = 0$ .’

The real marginal cost of sub-unit  $(i, j)$  is

$$\mathcal{MC}_{ij,t} = \frac{[\delta_{ij}(w_{ij,t}^e)^{1-\rho} + (1 - \delta_{ij})(p_{i,t}^m)^{1-\rho}]^{1/(1-\rho)}}{P_{ij,t}(1 - \alpha_{ij})} \cdot \frac{V_{ij,t}}{Y_{ij,t}}. \quad (19)$$

The aggregate base wage follows a reduced-form wage rule:

$$\frac{W_t}{W_{t-1}} = \bar{\pi} \left( \frac{N_{\text{agg},t}}{N_{\text{agg}}^{ss}} \right)^{\phi_w} \left( \prod_{k=1}^4 \frac{\Pi_{\text{agg},t-k}}{\bar{\pi}} \right)^{\iota_w}, \quad (20)$$

where  $\phi_w$  governs the sensitivity of base-wage growth to aggregate employment and  $\iota_w$  governs backward indexation to the four-year cumulative inflation deviation. Taking logs, the indexation term equals  $\iota_w \sum_{k=1}^4 (\log \Pi_{\text{agg},t-k} - \log \bar{\pi})$ : a sum in logs of inflation gaps over the previous four years, equivalent to the multiplicative product above. The four-year window allows the real wage to track cumulative inflation over a multi-year horizon while keeping the per-period pass-through small ( $\iota_w$  scales each lag), preserving a flat wage Phillips curve in the normal-times range of demand fluctuations and ruling out a one-period wage–price spiral.

To model anchoring of expectations, we assume that the expectation operator in the pricing equations is a convex combination of a model-consistent future term, a backward-looking term, and the target anchor:

$$\mathbb{E}_t[\pi_{t+1}] = v_{\text{forward}} \pi_{t+1}^{\text{model}} + v_{\text{past}} \pi_{t-1} + v_{\text{target}} \bar{\pi}, \quad v_{\text{forward}} + v_{\text{past}} + v_{\text{target}} = 1. \quad (21)$$

Our calibration considers two regimes that differ in their weight on future inflation and the nature of price anchoring. The model is solved as a nonlinear perfect-foresight transition path from  $t = 1$  onward, but the firm's pricing forecast uses the convex-combination operator above rather than full rational expectations.

### 3.5 Measurement and mapping onto the data

**Employment.** Aggregate employment is the sum of employment across network production nodes,

$$N_{\text{agg},t} = \sum_{i=1}^{12} \sum_{j \in \{H,S\}} N_{ij,t}. \quad (22)$$

Aggregate inflation is the fixed-share PCE-weighted index over the twelve tier prices and the three exogenous PCE blocks:

$$\log \Pi_{\text{agg},t} = \sum_{i=1}^{12} w_i^{\text{PCE}} \log \frac{P_{i,t}}{P_{i,t-1}} + \sum_{k \in \{HO,FO,EN\}} w_k^{\text{PCE}} \log \Pi_{k,t}, \quad \sum_{i=1}^{12} w_i^{\text{PCE}} + \sum_k w_k^{\text{PCE}} = 1. \quad (23)$$

**Output.** For real aggregate output we report value added rather than gross output, since gross output double-counts intermediate flows and calculate the sum of gross output is not meaningful. The headline model object is a fixed-steady-state-price double-deflated value-added measure for the endogenous production block:

$$Y_{\text{agg},t} = \sum_{i=1}^{12} \sum_{j \in \{H,S\}} \left[ Y_{ij,t} - \sum_{k=1}^{12} m_{ij,k,t}^{\text{DOM}} - m_{ij,t}^{\text{IMP}} \right], \quad (24)$$

where steady-state prices are normalized to one. This is the model analog of double-deflated real value added for the endogenous block. HO, FO, and EN are excluded from value added because they have no modeled production structure, although their prices enter aggregate inflation through PCE weights. On the sectoral level  $Y$  denotes gross output. Construction details are in Online Appendix [A.4](#).

**Wages.** We exploit a degree of freedom that the CW provides and allow a fraction of the effective wedge to appear in observed wages, in order to best match the wage data. To that end, we assume that the observed wage implied by the model is

$$w_{ij,t}^o = \tilde{w}_{ij} W_t (1 + x_{ij} \mathcal{W}_{ij,t}), \quad (25)$$

where  $x_{ij} \in [0, 1]$  is the wage-margin attribution factor. If  $x_{ij} = 0$ , the rebuilding friction  $\mathcal{W}_{ij,t}$  is “invisible” in measured wages and appears in the accounting decomposition as lower productivity or higher margins; if  $x_{ij} = 1$ , it is fully reflected in observed wages. The wage-margin attribution factor thus affects the wage–margin split in the accounting decomposition, but it is neutral to model dynamics and only affects how the model maps onto observables. We do not model how firms and workers may arrive at this split.

### 3.6 General equilibrium closure

The exogenous objects are the path of sub-unit employment  $\{N_{ij,t}\}$ , the tier-level demand shifters  $\{\xi_{k,t}\}$ , the exogenous markup wedges  $\{\mu_{i,t}\}$ , the node-level COVID labor wedge  $\{\tau_{ij,t}\}$ , the EXO price blocks, and the long-run wage anchor. The endogenous objects are the node-level prices  $P_{ij,t}$ , the materials demands  $m_{ij,k,t}^{DOM}$  and  $m_{ij,t}^{IMP}$ , the sub-unit output  $Y_{ij,t}$ , the variable-input bundles  $V_{ij,t}$ , the tier-level aggregates, the base wage  $W_t$ , and aggregate inflation  $\Pi_{agg,t}$ . The wall itself takes no resources because its cost is just the marginal charge associated with calibrated sectoral productivity  $a$ .

**Endogenous demand propagation.** The baseline demand-side closure keeps the tier-level demand shifters  $\xi_{k,t}$  exogenous, so collapsing real activity does not mechanically feed back into demand, since there is no household sector. We later relax this by adding an exogenous income-propagation term to each tier’s demand shifter:

$$\xi_{k,t}^{\text{endog}} = \xi_{k,t} + \text{MPC} \cdot (\log Y_{agg,t} - \log Y_{agg,t}^{\text{data}}), \quad (26)$$

with  $\text{MPC} = 0.4$  as the canonical value. This value is based on the MPC estimate in [Dupor et al. \(2022\)](#). However, this module is absent from the baseline model with all wedges, since in that case  $Y_{agg,t} = Y_{agg,t}^{\text{data}}$ .

## 4 Calibration

The model is annual. Its steady state is calibrated to 2019 values, and it speaks to detrended data. The I–O structure is taken from the 2017 BEA supply-use tables in producer prices. The perfect-foresight horizon runs to 2050; calibrated shocks are active during 2020–2024, after which the economy converges back to its nominal anchor by the end of the horizon.

To obtain a consistent comparison with the data, all targets are detrended at the *chain level* against pre-COVID 2017–2019 trends, separately for output (real value added), employment, gross-output deflators, and wages.<sup>14</sup>

The New Keynesian components of our model are consistent with the post-2000 view of an anchored, flat Phillips curve. At annual frequency, base-wage growth responds to aggregate employment with elasticity  $\phi_w = 0.15$ , in the range implied by the post-2000 flat-Phillips-curve evidence of [Hazell et al. \(2022\)](#) and [Del Negro et al. \(2020\)](#) (whose quarterly slopes  $\phi_w^Q \approx 0.03$ – $0.05$  translate to annual values near 0.12–0.20 under standard conversions). Backward wage indexation is modest at  $\iota_w = 0.20$ , at the upper bound of the Hazell range, and applied to the four-year sum-in-logs cumulative inflation deviation  $\sum_{k=1}^4 (\log \Pi_{t-k} - \log \bar{\pi})$  rather than to a single lag. The four-year window combined with the per-lag weight  $\iota_w = 0.20$  delivers cumulative pass-through of  $4 \times 0.20 = 0.80$ : persistent inflation deviations are reflected in the real wage over four years, while the per-period pass-through stays small and the Phillips curve remains flat in the normal-times range of demand fluctuations. The price block has no backward indexation ( $\delta_{\text{index}} = 0$ ), and the capacity wall has no internal persistence: the wall premium is a period-by-period function of current positive re-employment. Expectations are strongly anchored: in the baseline adaptive-anchored regime, 75% of inflation expectations are pinned at the Federal Reserve’s 2% target and 25% track past inflation. We also report results under a CG-consistent sticky-information

---

<sup>14</sup>Each chain trend is applied uniformly to the chain’s three tiers and to the within-tier hit/safe sub-units, preserving relative COVID-era movements inside the chain. By 2025, however, due to slower population growth and labor-force dynamics, aggregate employment does not return to its 2019 level—which is inconsistent with the fact that the 2024 unemployment rate returned to pre-COVID levels. We introduce a minor adjustment to the labor trend to close this gap and consistently adjust the trend for output to ensure the path of labor productivity is intact within each tier. Online Appendix [A.5](#) provides the details.

specification (Coibion and Gorodnichenko, 2015) that keeps the past-inflation weight at 0.25, adds a 0.25 forward-looking weight, and lowers the target-anchor weight to 0.50; the quantitative effect is minor.

Table 6 summarizes the key parameter values. Panel A lists externally imposed parameters; Panel B lists calibrated wedges and their target moments. We now describe the details.

**Table 6:** Calibration summary

<i>Panel A: Imposed parameters</i>			
$\rho$	CES substitution	0.39	KLEMS estimate
$\gamma_i$	Capacity-wall intensity	per tier	See text
$c$	Net-rebuilding cost (per unit capacity)	0.09	Linckh et al. (2025)
$q_0, q_1, T$	Quit rates (normal, tight), horizon	0.0225, 0.028, 36 mo	JOLTS
$\bar{\phi}_{c,h}$	Price stickiness, $h$ tier (per chain)	3–10	Nakamura and Steinsson (2008)
$\bar{\phi}_{ml}$	Price stickiness, $m/\ell$ tiers	3	Nakamura and Steinsson (2008)
$\lambda_{c,h}, \lambda_{ml}$	Downward rigidity	9–30, 9	Eichenbaum et al. (2011)
$\phi_w$	Wage Phillips slope	0.15	Hazell et al. (2022)
$\iota_w$	Wage indexation (per lag, 4-lag)	0.20	Hazell et al. (2022)
$\delta_{\text{index}}$	Price indexation	0	Hazell et al. (2022)
$\beta$	Discount factor	0.96	Annual
$\alpha_{ij}$	Capital share <sup>15</sup>	Sub-unit-specific from 2017 BEA-NAICS aggregation	
$\delta_{ij}$	Labor weight	Sub-unit-specific from 2017 BEA-NAICS aggregation	
$\tilde{w}_{ij}$	Wage premia	Sub-unit-specific from 2017 BEA-NAICS aggregation	
$\omega_{i,k}^{\text{DOM}}$	Domestic I–O shares	12 $\times$ 12 from 2017 BEA use tables	
$\omega_i^{\text{IMP}}$	Import shares	BLS import price data and BEA use tables	
<i>Panel B: Calibrated objects</i>			
$a_{i,t}$ (12 $\times$ 5)	Productivity shifters	Tier-level detrended $N$ paths, Newton iteration	
$\xi_{i,t}$ (12 $\times$ 5)	Demand shifters	Tier-level detrended $Y$ paths, Newton iteration	
$\bar{\tau}_i$ (12)	Tier-mean labor wedge, 2020 only	Tier-level detrended $N(2020)$ , Newton iteration	
$\Delta\tau_c$ (4)	Hit/safe labor-wedge spread, 2020 only	Chain-level $\sigma^-(\Delta \log N_{2019 \rightarrow 2020})$ from QCEW	
$x_{ij}$ (24)	Wage-margin attribution, sub-unit level	Sub-unit wage paths, 2021–2024	
$\mu_{i,t}$ (12 $\times$ 5)	Residual markup wedge	Tier-level detrended $P_{GO}$ paths	

*Notes:*  $i$  indexes the 12 chain-by-tier composites and  $j \in \{H, S\}$  indexes the hit/safe sub-unit. The CES substitution elasticity  $\rho$  governs all three CES margins: labor–materials, across intermediate inputs, and between hit and safe sub-units. Sensitivity to lowering all three elasticities jointly to 0.5 and to raising hit/safe substitution alone ( $\sigma_w = 5$ ) is reported in Online Appendix A.7. The wall premium  $\mathcal{W}_{ij,t} = \gamma_i (dN_{ij,t}^+)^2$  is contemporaneous with no separate state or persistence.

**CW calibration.** The wall premium is  $\mathcal{W}_{ij,t} = \gamma_i (dN_{ij,t}^+)^2$ , facing all firms in industry  $(i, j)$  as a wedge in the effective wage. The firm takes  $\mathcal{W}_{ij,t}$  as given (it depends on industry-level net rebuilding, not on firm-level decisions) and optimizes accordingly. Our calibration of  $\gamma_i$  translates micro evidence on hiring-cost convexity into the model’s net-rebuilding object. The aim is to provide a reasonably conservative range for this key parameter. Sensitivity of our results to this parameter is roughly proportional

to its value, as we show in the Online Appendix.

Let  $c^{wall}$  denote the average wall wedge over the rebuilding window, expressed as a fraction of the base wage  $W$  in the model. Because the empirical hiring-cost premium and the model wedge both enter marginal cost through the same labor-cost channel, the labor share cancels in the calibration condition.<sup>16</sup> We choose  $\gamma_i$  so that the average wall wedge over the 2021–2023 rebuilding window equals this target:

$$\gamma_i = \frac{c^{wall}}{(dN^+)^2_i}, \quad \overline{(dN^+)^2_i} = \frac{\sum_{j,t} \omega_{ij}^N (dN_{ij,t}^+)^2 \mathbf{1}\{dN_{ij,t}^+ > 0\}}{\sum_{j,t} \omega_{ij}^N \mathbf{1}\{dN_{ij,t}^+ > 0\}}, \quad (27)$$

where  $\omega_{ij}^N$  is the 2019 employment share of sub-unit  $(i, j)$  within tier  $i$ , and the sums run over  $j \in \{H, S\}$  and  $t \in \{2021, 2022, 2023\}$ .

The quadratic form embeds intensity convexity directly into the friction, coming from market tightness not firm’s own hiring effort. Consistent with detrending applied throughout the paper,  $dN_{ij,t}^+$  is constructed from the detrended employment path, which significantly lowers the value of this wedge. “Rebuilding” means catching up to the pre-COVID trend extrapolation, not mechanically returning to the raw 2019 level.

[Linckh et al. \(2025\)](#) estimate that filling a skilled-worker vacancy in Germany costs approximately 22 weeks of wages over 2015–2019, with roughly 80% of that cost arising post-match—via adaptation, informal training, external training, and disruption to incumbent workers—rather than from pre-match search. However, the German labor market in 2015–2019 was already historically tight, so the 22-week figure embeds a meaningful tightness premium. We therefore anchor our normal-times baseline on [Muehlemann and Strupler Leiser \(2018\)](#), who report 16 weeks for a comparable Swiss sample in a less tight labor market. We take  $Q_0 = 16/52 \approx 0.308$  as a conservative baseline.

[Linckh et al. \(2025\)](#) estimate that a one-standard-deviation increase in vacancy-to-unemployment tightness raises post-match hiring costs by roughly 10 to 12% on the largest cost components (adaptation and disruption time). The U.S. V/U ratio rose from about 1.2 in late 2019 to about 2.0 at its 2022 peak,

<sup>16</sup>The elasticity of marginal cost with respect to the effective wage is the labor-cost share, which multiplies both the model wedge  $\gamma_i(dN^+)^2$  and the empirical target  $c^{wall}$ . The calibration of  $\gamma_i$  requires no markup, labor-share, or MPL adjustment in our framework.

an increase of roughly three standard deviations relative to its U.S. historical distribution (Domash and Summers, 2022). Applying the upper end of the per-SD elasticity yields a tight-market hiring-cost premium of approximately 35% above normal. Applied to the 16-week Swiss baseline—rather than the 22-week Linckh et al. (2025) figure, which already embeds a tightness component—this gives a gross tight-market premium of about  $5.6/52$ , or roughly 5–6 additional weeks of wages on top of the baseline. But there are several complications that need to be taken into account. First, pandemic separations differed from a standard recession in that a substantial share were temporary layoffs (the BLS temporary-layoff share peaked at 78% in April 2020 and averaged 49% over 2020), and even non-recalled workers in the 2021–2023 rebuilding window were unusually likely to have done similar work recently. Both channels lower training and adaptation costs. We discount the tightness premium by 30% to reflect these warm-pool effects with some disruption, yielding an adjusted tight-market cost of  $Q_1 \approx 16/52 + 0.7 \times 5.6/52 \approx 20/52$ , or about 4 additional weeks above the normal-times baseline.

To convert this average per-hire figure into the marginal, annuity-equivalent rebuilding cost the model requires, we use the gross-to-net multiplier  $m(q, T) = qT/[1 - (1 - q)^T]$ , which charges turnover of the rebuilt cohort only. Routine turnover of incumbents is already in the 2019 cost structure, and so it should not be counted here.<sup>17</sup> The empirical anchor for  $c^{wall}$  needs to scale by the share of the U.S. wage bill covered by skilled workers in the relevant sense, which accounting to our estimate would correspond to roughly four-fifths of the labor force given the definition of “skilled” as workers requiring some training in Linckh et al. (2025) (in Germany about 60% of workers are in this category).<sup>18</sup>

Combining these elements, the rebuilding-cost wedge is

$$c^{wall} = \eta \kappa_{\text{skill}} \left[ m(q_1, T) Q_1(r + \lambda(q_1)) - m(q_0, T) Q_0(r + \lambda(q_0)) \right], \quad (28)$$

where  $\eta = 1.3$  is the average-to-marginal conversion from Muehlemann and Pfeifer (2016) own hiring

<sup>17</sup>We set  $T = 36$  (the 2021–2023 active rebuilding window),  $q_0 = 0.0225$  (2017–2019 JOLTS average), and a baseline  $q_1 = 0.028$  close to but moderately below the recovery peak;  $q_1 = 0.030$  provides an upper-tightness case. These imply  $m(q_0, 36) \approx 1.45$ ,  $m(0.028, 36) \approx 1.57$ , and  $m(0.030, 36) \approx 1.62$ .

<sup>18</sup>See footnote 3.

estimate,  $\kappa_{\text{skill}} = 0.8$  is the skilled-coverage adjustment, and  $\lambda(q) = 12[-\log(1 - q)]$  is the annual separation hazard with  $r = -\log \beta \approx 0.041$  at discount  $\beta = 0.96$ . Evaluating at  $Q_0 = 16/52$ ,  $Q_1 \approx 20/52$ ,  $q_0 = 0.0225$ , and  $q_1 = 0.028$  yields  $c^{\text{wall}} \approx 0.09$ , with  $c^{\text{wall}} \approx 0.12$  under the upper-tightness case  $q_1 = 0.030$ . We adopt  $c^{\text{wall}} = 0.09$  as the baseline, with lower- and higher-cost variants reported in Online Appendix A.7.<sup>19</sup> Combining (27) with  $c^{\text{wall}} = 0.09$ , we set

$$Y_i = \frac{0.09}{(dN^+)_i^2}.$$

**Price stickiness.** The Rotemberg parameters are anchored to the monthly micro evidence on price-change frequencies in Nakamura and Steinsson (2008) and translated to annual frequency via  $\theta_a = \theta_m^{12}$  with  $\theta_m = 1 - 1/D_m$ , then mapped to a Rotemberg adjustment cost using the Calvo–Rotemberg equivalence.<sup>20</sup> Each chain’s  $h$ -tier corresponds to a distinct composition of goods and services:  $C_1$  is services-heavy (retail/distribution),  $C_2$  is services-heavy,  $C_3$  is goods-heavy and fast-moving,  $C_4$  is mixed industrial/consumer. Assigning monthly durations of roughly 20, 17, 11, and 13 months to the four chains gives consumer-facing values  $(\bar{\phi}_{C_1,h}, \bar{\phi}_{C_2,h}, \bar{\phi}_{C_3,h}, \bar{\phi}_{C_4,h}) = (10, 7, 3, 4)$ . The middle and upstream tiers share a lower value  $\bar{\phi}_{ml} = 3$ , reflecting faster producer-stage repricing ( $D_m \approx 8$  months of PPI durations in Nakamura and Steinsson (2008) gives  $\bar{\phi} \approx 1$ ; we adopt the higher value 3 to retain small symmetric stickiness at the annual frequency). The micro evidence disciplines the order of magnitude of these parameters; the mapping from monthly CPI frequencies into annual GO/PPI-like Rotemberg costs involves several approximations, so we frame these values as disciplined frictions rather than tightly identified parameters. The headline results are not sensitive to alternative values within the Nakamura

<sup>19</sup>Hiring-cost convexity estimated by Muehlemann and Pfeifer (2016)— $\eta = 1.3$  on skilled-worker hires—provides a within-firm intensity-convexity margin distinct from both the across-market tightness convexity captured by the Linckh et al. (2025) elasticity above and the cross-industry rebuilding-intensity convexity captured by the model’s quadratic specification. We apply  $\eta$  here as a within-firm average-to-marginal conversion, where the variation identifying it is firms’ own scale of hiring (moving deeper into the available labor pool), and so it does not double-count the tightness convexity from Linckh et al. (2025) to the extent that the M&P industry fixed effects absorb local market conditions. Blatter et al. (2012) report broadly consistent convexity using Swiss data, with marginal hiring costs reaching up to 24 weeks against an average of 10–17 weeks.

<sup>20</sup> $\bar{\phi} \approx \frac{(\varepsilon - 1)\theta_a}{(1 - \theta_a)(1 - \beta\theta_a)}$ , with  $\beta = 0.96$  and  $\varepsilon = 5$  the standard varieties elasticity, used here only for the duration-to- $\bar{\phi}$  conversion—the model takes  $\bar{\phi}$  directly as a parameter.

range.

**Downward price rigidity.** The price block of the model features strong asymmetric adjustment: prices adjust upward more readily than they fall. This feature helps limit deflation in the pandemic year (which is mostly driven by the collapse in energy prices) and otherwise does not play a big role. We follow the structural-asymmetry literature (Eichenbaum et al., 2011; Karadi and Reiff, 2019; Cavallo et al., 2024), which targets a 3–5× marginal-cost asymmetry between price decreases and increases, well above the 2× ratio implied by the classical linear-regression estimate of Peltzman (2000). That gap reflects three biases that linear regressions impose on a genuinely kinked adjustment cost: selection (kinked firms cut only when the desired markdown is large enough to overcome the kink), cross-firm heterogeneity (kinked firms diluted by smoother adjusters), and time aggregation (high-frequency asymmetric dynamics smoothed by monthly or annual aggregation). We adopt the midpoint of the corrected range by setting  $\lambda_i = 3\bar{\phi}_i$ , which delivers a 4× marginal adjustment cost for price decreases relative to increases ( $1 + \lambda/\bar{\phi} = 4$ ). Tying  $\lambda_{c,h}$  to  $\bar{\phi}_{c,h}$  chain-by-chain holds the asymmetry ratio constant across the finalness gradient, which is the structural moment the literature identifies; in absolute terms,  $(\lambda_{C1,h}, \lambda_{C2,h}, \lambda_{C3,h}, \lambda_{C4,h}) = (30, 21, 9, 12)$  and  $\lambda_{ml} = 9$ . The hierarchy of channels in the model is clear: the wall does the structural work; downward rigidity prevents wall-driven price increases from reverting once the wall releases, stabilizing the transition path. We report a halved- $\lambda$  sensitivity ( $\lambda_i = 1.5\bar{\phi}_i$ , i.e., a 2.5× asymmetry) in Online Appendix A.7; the wall’s attribution is essentially unchanged.

**Inflation expectations.** We assume an anchored inflation-expectations regime, consistent with the data. Specifically, the expectations operator in equation (21) is calibrated to one-year-ahead inflation-expectations regressions on the Cleveland Fed market-based series (Haubrich et al., 2012) and cross-validated against the University of Michigan household survey and TIPS five-year break-even inflation. The headline specification is the first-difference regression

$$\Delta E_t[\pi_{t+12}] = \alpha + b \Delta \pi_t^{\text{past}} + u_t, \quad (29)$$

where  $\pi_t^{\text{past}}$  is realized 12-month CPI inflation ending at  $t$  and  $\Delta$  denotes a 12-month revision. Differencing removes slow-moving anchor drift and trend contamination that pull level-form slopes toward zero. Under the adaptive-anchored restriction  $v_{\text{forward}} = 0$ , the slope  $b$  identifies the non-anchor pass-through coefficient  $1 - v_{\text{target}} = v_{\text{past}}$ . In the COVID and post-COVID window the Cleveland Fed series gives  $b = 0.246$  ( $t = 8.34$ ,  $N = 74$ , Newey–West HAC with 12 lags;  $R^2 = 0.59$ ); the pre-COVID 2016–2019 window gives  $b = 0.295$  ( $t = 6.95$ ,  $R^2 = 0.35$ ). The two estimates are statistically close, and 48-month rolling windows stay inside  $[0.10, 0.30]$  over 2005–2024 with no monotonic drift. We set  $b = 0.25$ , implying  $v_{\text{target}} = 0.75$ .

We consider two specifications consistent with this evidence. The baseline *adaptive-anchored* regime sets  $(v_{\text{forward}}, v_{\text{past}}, v_{\text{target}}) = (0, 0.25, 0.75)$ , taking the first-difference regression at face value: 25% of expectations track recent inflation, 75% are anchored at the central bank’s target, and there is no forward-looking component. As a robustness alternative we consider a *CG-consistent* sticky-information specification that adds a 25% forward-looking weight implied by the information-rigidity estimates of [Coibion and Gorodnichenko \(2015\)](#), retains the empirically robust  $v_{\text{past}} = 0.25$ , and lowers the anchor to 50%:  $(v_{\text{forward}}, v_{\text{past}}, v_{\text{target}}) = (0.25, 0.25, 0.50)$ .<sup>21</sup>

The adaptive specification fits the post-COVID expectations data substantially better than the CG-consistent alternative across both level- and first-difference comparisons with the Cleveland Fed model-implied series and TIPS five-year break-even inflation, and we adopt it as our baseline accordingly. Online Appendix [A.6](#) reports the full head-to-head comparison.

**Saturating wedges.** The remaining calibrated parameters saturate the model so that it matches the observed sectoral paths before we decompose the inflation mechanism. The productivity shifters  $a_{i,t}$  primarily match detrended employment paths, the demand shifters  $\xi_{i,t}$  match detrended output paths,

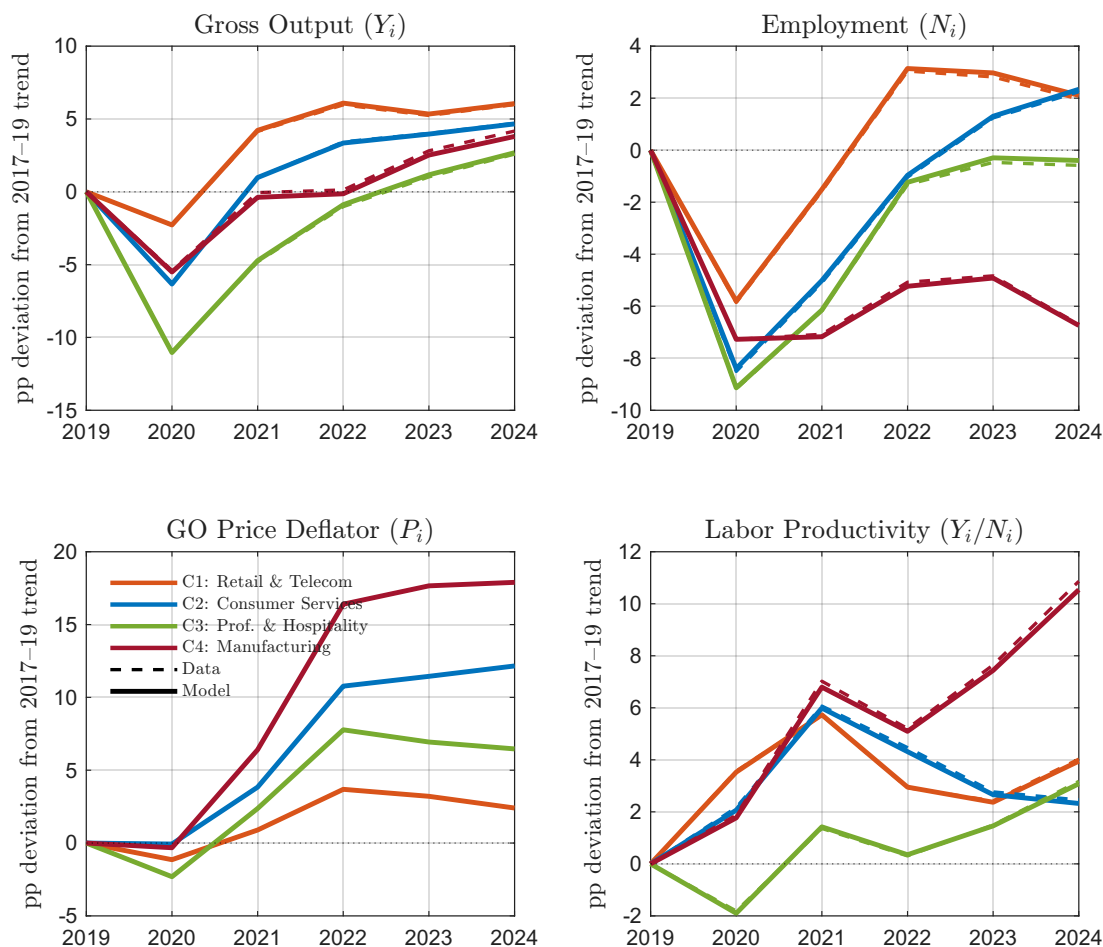
<sup>21</sup>In the [Coibion and Gorodnichenko](#) sticky-information framework the stale-information share is not exclusively anchored; under an approximately AR(1) inflation process the average sticky-information forecast collapses to a weighted average of recent inflation and the long-run mean, and the [Coibion and Gorodnichenko](#) updating fraction  $1 - \lambda \approx 0.25$  is allocated to  $v_{\text{forward}}$ .

and the wedge  $\tau_i$  captures the 2020 employment collapse difference between safe/hit nodes.<sup>22</sup> The value of  $\Delta\tau_c$  targets the excess downward dispersion of 2019–2020 employment growth relative to a pre-COVID placebo baseline, computed in the data at the chain level and at the BEA detail level of disaggregation. The residual markup wedge  $\mu_{i,t}$  then closes the price dimension. The addition of  $\mu_{i,t}$  is iterated to convergence.

Figure 4 shows that the saturated calibration matches gross output, employment, labor productivity, and gross-output deflators at the chain level.

---

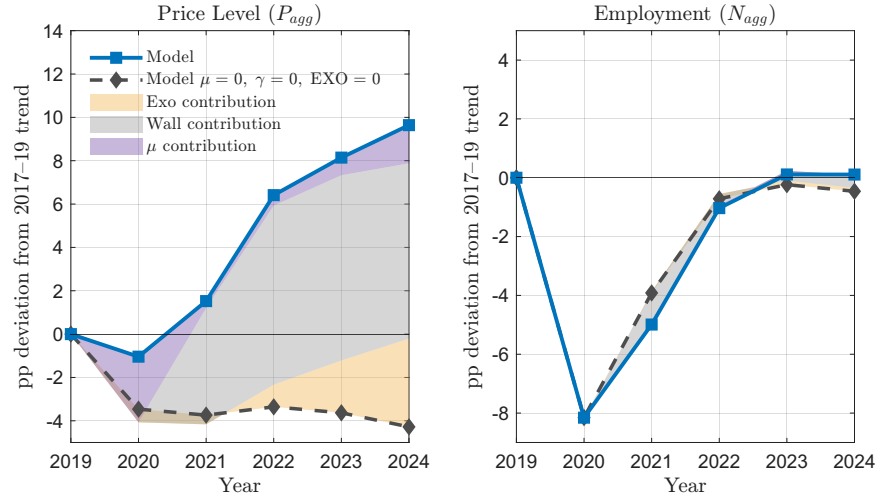
<sup>22</sup>The parameter  $\tau_i$  is technically redundant— $a_{i,t}$  and  $\xi_{i,t}$  could in principle absorb the 2020 collapse—but is necessary in practice to stabilize the calibration. Without  $\tau_i$ , matching 2020 requires extreme sectoral productivity swings that destabilize the I–O system. Our calibration maximizes the role of  $\tau_i$ , reflecting the prior that COVID was primarily a labor-market shock. Across alternative convergence orders, differences in aggregate inflation are within  $\pm 1$  pp and qualitative insights are unchanged.




**Figure 4: Chain-level fit under the saturated calibration.** Gross output, employment, labor productivity, and gross-output deflators by chain. Dashed: data. Solid: model with  $\mu$  active. Cumulative log points relative to 2019, detrended.

## 5 Results

Our model’s calibration strategy has a baseline designed to answer an accounting question: how much of observed inflation is generated by the structural channels—the capacity wall, input–output propagation, and exogenous price blocks—versus the exogenous markup wedge  $\mu$  that picks up the residual slack. To assess the contribution of the wall and the exogenous sectors, we set  $\mu = 0$  and compare the results to the baseline. To then isolate the contribution of the exogenous sectors, we set the exogenous shocks to zero (use steady-state prices) while maintaining  $\mu = 0$ ; the difference between these two counterfactuals is the contribution of the exogenous shocks in this stacked removal. Finally, by setting



(a) Adaptive-anchored regime,  $(v_{\text{forward}}, v_{\text{past}}, v_{\text{target}}) = (0, 0.25, 0.75)$ .



(b) CG-consistent regime,  $(v_{\text{forward}}, v_{\text{past}}, v_{\text{target}}) = (0.25, 0.25, 0.50)$ .

**Figure 5: Channel decomposition in the fitted model.** The panels report cumulative log deviations from the 2019 baseline under the markup-saturated calibration and under sequential counterfactuals that remove the residual markup wedge, the exogenous price blocks, and the capacity wall. Top row: adaptive-anchored expectations (baseline). Bottom row: CG-consistent expectations. Horizon shown: 2019–2024.

$\mu = \gamma = 0$  and exogenous shocks to zero, we find the calibrated model’s *structural floor* for inflation. We hold all other parameters and wedges fixed throughout, in line with the business-cycle accounting methodology.

Figure 5 shows this stacked decomposition. As we can see, the structural floor is deflationary; that is, the model predicts below trend inflation without exogenous shocks and the wall. This is not surprising: until 2023–2024 employment in our calibration remains below trend. The reason why the structural

floor is negative is twofold. First, the COVID shock is deflationary and our model seems to feature too little downward price rigidity to prevent prices from falling more than in the data during the COVID year. But most importantly, the  $a$  wedge picks up an over-trend increase in labor productivity during this entire time period (2019-2024), which is deflationary in the long run. The terminal price level in 2050 falls about 3 percent below the initial steady state.

The figure shows that the wall's impact on inflation is persistent. This may seem counterintuitive given the wall strongly binds in 2021 and essentially wanes in 2023 (we know it from employment growth which in 2024 reaches the steady state by construction). There are two reasons for this. The first reason is trivial: prices are sticky, and so inflation does not rise right away but flattens and becomes persistent. But there is a more subtle reason that our modeling uncovers. The model is calibrated via the choice of parameters  $a_{i,t}$  so that rebuilding costs interact with the evolution of labor productivity, which the model matches exactly by construction. As noted in the data section (see Table 1), the spike in measured labor productivity right after reopening insulated prices from the wall's effect, in effect resulting in "missing inflation" from this first phase. In other words, profit margins expanded not because prices went up but because productivity spiked. But part of this boom was artificially generated by changing labor composition (?), which is well documented. Firms simply laid off their least productive and experienced workers. As a result, upon reopening, rehiring resulted in productivity growth slowdown, as these workers were rehired. Since profit margins did not recede during this period, the productivity slowdown led to inflation. In our model, after such a prolonged period of inflation, a fraction of it gets cemented in catch-up wages. This is captured by the longer lags in the wage equation of the model. Of course, the productivity boom-bust cycle is exogenous in our setup, but it is intimately tied to the wall mechanism: facing an unprecedented shock and the need to downsize, firms do it optimally, by laying off their least productive workers. The labor market pool is thus worse than the working pool, which slows down productivity growth later on.

We now assess the details of the underlying mechanism that generates inflation. As the first step, we decompose the contribution of shock asymmetry, or unevenness of the wall, and then discuss this

episode in light of the Phillips curve steepening debate.

**Contribution of wall unevenness.** Here we assess the contribution of wall unevenness to inflation, which is an important point in the debate on Phillips curve steepening. Our modeling implies that much of it owes to shock geometry as opposed to a static feature of the economy itself. To show this, we consider two counterfactual experiments.

The first experiment, labeled (*even wall*), replaces the observed per-tier capacity-wall lagged employment drag with a uniform average lagged employment that preserves the size of the 2020 contraction but removes cross-tier asymmetries. Specifically, the lagged employment reference inside the wall equation is replaced by a common value scaled by the aggregate 2020 trough,

$$N_u^{\text{ref}} = N_u^{\text{ss}} \frac{N_{\text{agg}}(2020)}{N_{\text{agg}}^{\text{ss}}},$$

so each production node experiences the same proportional tightness of the wall. What this means is that nodes that contracted more can rebound at no cost in 2021 back to the average. The second counterfactual experiment, labeled (*no wall*), sets  $\gamma = 0$  in all tiers, eliminating the wall channel entirely. This allows for a comparison of where the even wall falls between these two extremes—actual data and the structural floor of the model. Because the endogenous Keynesian propagation of income changes to consumption, via positive marginal propensity to consume (MPC), can be important, we use a variant of our model that adds a mechanical propagation of income changes to demand at MPC=0.4.

Table 7 shows the results. In the baseline calibration (left panel, MPC = 0) the capacity wall accounts for 102% of above-floor inflation in Act I (2019–2021) and 49% in Act II (2021–2023). The decomposition by Act shows that unevenness drives Act I almost entirely: only 14 of the 102 pp survive when the wall is forced to be uniform across tiers. Act II is the opposite: the even-wall component delivers a still notable 37 of the 49 pp the wall contributes. Much of wall-attributed inflation during this time period comes from its uniform bind across sectors.

As noted, the wall is identified using the baseline model that holds the demand side fixed. This raises the

**Table 7:** Capacity-wall contribution to inflation by Act (adaptive spec).

Series (cum log $\Pi_{agg}$ , pp)	Baseline		With income propagation (MPC = 0.4)	
	Act I (2019–2021)	Act II (2021–2023)	Act I (2019–2021)	Act II (2021–2023)
Model	+1.52	+6.62	+1.52	+6.62
Model – no $\mu$ – no exo	+1.62	+3.29	+0.53	+4.36
– even wall (uniform $\gamma$ )	–2.98	+2.51	–3.20	+4.63
– no wall ( $\gamma = 0$ , structural floor)	–3.74	+0.11	–3.82	+0.24
<i>Contributions to above-floor inflation by Act:</i>				
Exogenous shocks ( $\mu$ +exo)	–2%	51%	19%	35%
Capacity wall, total	102%	49%	81%	65%
of which: even-wall component	14%	37%	12%	69%

*Notes:* Shares denominated by total above-floor cumulative inflation (Model – no wall). “Even wall” forces  $\gamma$  uniform across tiers. Right panel adds income-propagation term  $MPC \cdot (\log Y_{agg} - \log Y_{agg}^{data})$  with  $MPC = 0.4$  to tier demand shifters.

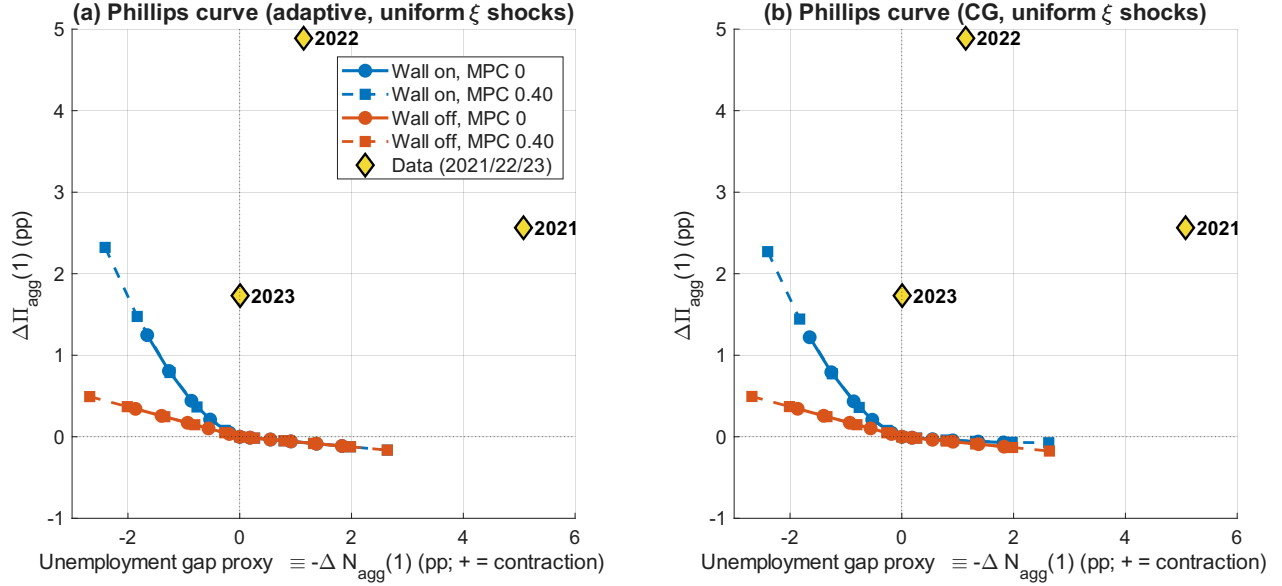
concern that perhaps positive income-spending feedback could change our accounting. The right-hand panel of Table 7 addresses this concern by augmenting each tier’s demand shifter by an endogenous reduced-form, income-propagation term  $MPC \cdot (\log Y_{agg} - \log Y_{agg}^{data})$  with  $MPC = 0.4$ —as we discussed briefly at the end of the model section. It does indeed matter: the wall share moves from 102% to 81% in Act I and from 49% to 65% in Act II. What this means is that there is less extreme timing attribution, but the overall contribution of the wall changes little. The qualitative two-Act nature of the shock is preserved: similarly, in Act I only 12 of the now 81 pp of wall contribution survive uniformity, while in Act II the even-wall component carries 69% of the wall’s contribution. Importantly, we should stress that this is a conservative estimate of the income-feedback channel: markups redistribute income from labor to capital, so a meaningful distinction is really between MPC from labor and capital—in the direction of higher MPC from labor income and hence potentially higher attribution of wall to labor income—to the extent that labor wedge leaks to wage increases in the sense of our wage attribution mechanism underlying equation (25).

**Cost-price pass-through nonlinearity and the Phillips curve.** The even/uneven wall decomposition implies that the model gives rise to *endogenous* steepening of the Phillips curve via steepening of the cost-price pass-through between 2020–2023 that largely owes to the shock’s uneven geometry.

This is important because COVID may not uncover Phillips curve steepening that one would encounter in normal times, when uniform demand overheating of a similar scale drives inflation, as an example. Here we investigate this property of the model more closely.

To trace the nonlinearity, we draw the model’s Phillips curve directly by simulating its response to a uniform, one-period, iid  $\xi$  demand shock applied simultaneously to all 12 tier composites of the 24-sub-unit model and sweep the impulse size across five magnitudes of each sign,  $|\xi| \in \{0.1\%, 0.3\%, 0.5\%, 0.75\%, 1\%\}$ . This is an impulse response shock that starts from the steady state and returns to the steady state at the simulation horizon. Each shock size delivers one impact-period observation of the pair  $(-\Delta N_{\text{agg}}(1), \Delta \Pi_{\text{agg}}(1))$ —an unemployment-gap and an inflation reading at the same horizon—and the ten points connect into a full Phillips schedule that spans both expansions (left tail) and contractions (right tail), shown in Figure 6. We repeat this exercise under the calibrated wall and under the structural counterfactual ( $\gamma_i = 0$  in all tiers), with the MPC=0.4 income-propagation channel switched off and on—as indicated in the figure. We also report the results for both expectations regimes (left panel versus right panel). Online Appendix Figures 10–11 show the underlying impulse-response trajectories. The important lesson from this appendix figure is that a negative shock may even lead to a flatter Phillips curve than without the wall, because destruction of production capacity today leads to inflation tomorrow, and this is anticipated partly and leads to inflation today via the standard Phillips curve mechanism present in equation (1).

The figure reveals several properties of the model. First, the wall-off schedule (orange) is approximately *linear* across both tails: the impact slope is close to 0.18 in expansions and somewhat shallower in contractions, with no shock-size dependence. This is the structural flat Phillips slope embedded in our calibration. Second, the wall-on schedule (blue) is *convex on the left tail*: small expansionary impulses fire the wall weakly and the curve takes off from the wall-off line, with larger impulses resulting in larger gaps. This is generated by the fact that, in our core equation (16),  $dN^+$  is raised to the square. At a 1% expansionary pulse the wall-on slope reaches roughly 0.75 under adaptive expectations, about 4× the wall-off baseline. To the extent that aggregate demand leads to severe tightening of the labor



**Figure 6: Model Phillips curve under uniform iid  $\xi$  demand shocks, both tails.** Each marker is the impact-period ( $t = 1$ ) response of aggregate inflation  $\Delta\Pi_{\text{agg}}(1)$  to one of ten symmetric one-period demand impulses ( $|\xi| \in \{0.1\%, 0.3\%, 0.5\%, 0.75\%, 1\%\}$  applied with both signs) on all 12 tier composites simultaneously, solved by perfect foresight from steady state. The horizontal axis is the implied unemployment-gap proxy  $-\Delta N_{\text{agg}}(1)$ : negative values correspond to expansionary shocks (employment rises), positive values to contractionary shocks (employment falls). Blue (wall on) uses the calibrated convex wall law  $\mathcal{W}_{u,t} = \gamma_i (dN_{u,t}^+)^2$ ; orange (wall off) sets  $\gamma_i = 0$  in all 12 tiers. Solid lines: MPC channel off; dashed: MPC= 0.40. Panel (a): adaptive expectations (0, 0.25, 0.75). Panel (b): CG expectations (0.25, 0.25, 0.50).

market, our model implies that the Phillips curve in normal times can become significantly steeper. But as the figure shows, inflation and employment are nowhere near this curve in 2021, 2022 and 2023. Inflation is way higher for the level of employment decline. This is the effect of the wall's unevenness combined with the evolution of other wedges, in particular  $a$ , but also some attribution of the spike in exogenous prices.

As for the right tail, the Phillips curve in the model is flatter than the structural-floor model implies. This is not surprising: The wall premium is one-sided by construction:  $dN_{u,t}^+ = \max\{(N_{u,t} - N_{u,t-1})/N_u^{ss}, 0\}$ . Accordingly, a contractionary  $\xi$  shock drives  $dN < 0$  at impact, but there is anticipation of future inflation that feeds into current inflation via equation (1) and the expectations term (the first term on the right-hand side). Therefore, to the extent that estimates of Phillips curve in normal times use negative shocks as the baseline, the Phillips curve may be even flatter according to our model than the structural-floor baseline.

Finally, the convexity of the wall amplifies the income-propagation channel via  $MPC=0.4$  but this leaves the slope of the Phillips curve unaffected. With the wall switched off the  $MPC=0.4$  line is indistinguishable from the  $MPC=0$  line—which is because  $MPC$  scales  $N$  and  $\Pi$  proportionally. The two regimes (adaptive vs. CG) give qualitatively identical pictures, with CG slightly flatter on the right tail as it puts more weight on the forward-looking, anchored target, which amplifies the just discussed expectations channel in (1).

To provide a cross-sectional check on model consistency with sectoral data, we consider the OLS regression of the form:

$$(L^\top w)_i \Delta \log P_i = \alpha + \beta (L^\top w)_i \Delta ULC_i + \varepsilon_i,$$

across our BEA detail industries. Both regressand and regressor are industry  $i$ 's L-propagated contributions in pp, summed over the period's annual growth rates and taken from the right-hand side of the L-decomposition in equation (4); the weight  $(L^\top w)_i$  is the all-order PCE-importance of industry  $i$  once supplier-cost passthrough has been unrolled through the Leontief inverse. The slope  $\beta$  is a price-cost pass-through coefficient: it measures how unit labor cost passes into prices, a link in the Phillips-curve transmission that maps aggregate slack to inflation. Before COVID, regressions of price growth on unit labor cost growth across BEA detail industries were weakly positive in the data. Estimated over a pre-COVID 2017–2019 window and the 2020–2022 trough-to-peak inflation window, the slope increased from  $\beta_{\text{pre}} = +0.110$  ( $t = 5.14$ ,  $N = 341$ ) to  $+0.285$  ( $t = 6.60$ ,  $N = 341$ )—roughly  $2.6\times$  the pre-COVID baseline. Using the peak 2019–2021 window in the model, the model's predicted cost-price slope is  $+0.157$  against the data's  $+0.190$ . On the window 2020–2022, it is  $+0.170$  against the data's  $+0.285$ , which shows that part of the transmission from the wall to wages implied by the leakage mechanism underlying equation (25) may be too weak.

**Price dispersion.** Next, we show that the model reproduces Fact 2 of Section 2: the negative relationship between own rebuilding speed and cumulative price growth, conditional on the size of the 2020 collapse. This relationship is the discriminating sign test that separates inflation attributed to

a supply-drag (capacity-wall) channel from inflation attributed to a demand-pull or pent-up-demand channel: both stories generate excess demand for output, but they have opposite sign predictions for the recovery margin.

To show this, we run a regression with the same right-hand side as the model’s structural regression on the 24 hit/safe sub-units (own and supplier  $\Delta N^{\text{gap}}$ ,  $\Delta N^{\text{rec}}$ ,  $\Delta \mathcal{N}^{\text{rec}}$ , and supplier exposure) on the BEA-357 cross-industry data, with both OLS and 2SLS estimators reported. The data-side regressions use BEA-66 sector + chain-tier-12 cell fixed effects, as in the headline reduced form of Section 2; the model-side regression uses chain-tier-12 cell fixed effects. Table 8 reports the side-by-side comparison. The data OLS and 2SLS share the same right-hand side. The 2SLS additionally uses the predetermined JEM exposure  $E_i$  as an instrument for the 2020 own employment collapse. We report both because the model’s regression is OLS at  $N = 24$  where exposure is a deterministic binary classification and no instrumentation is needed, while the data-side 2SLS provides the per-pp-employment translation that brings the data coefficients onto the same scale as the model’s. The 2SLS results hinge on a difficult-to-justify assumption that the exclusion restriction applies (see the footnote in Section 2); the IV coefficients in this table should thus be read with caution and together with the OLS these results are indicative of magnitudes in the data but not definitive.<sup>23</sup>

The key takeaway is in the own-recovery row. The model delivers  $\hat{\pi}^r = -0.33$  against the data 2SLS estimate of  $-0.33$ : identical sign and magnitude. The data OLS gives the same sign with attenuated magnitude ( $-0.14$ , also significant at 5%), as expected when the structural employment-to-prices relationship is identified off variation that the IV recovers but OLS cannot. The 2020-collapse coefficient is positive and significant in both the model and the data 2SLS once the model’s coefficient is placed on the

<sup>23</sup>Diagnosics for the 2SLS are good: The cluster-robust first-stage  $F$  on  $E_i$  is 22.00, well above the Stock and Yogo (2005) 10%-maximal-bias critical value of 16.38 for a single just-identified instrument and the 20% threshold of 8.96; the conventional weak-instrument concern is not serious. In addition, weak-instrument-robust inference reinforces this conclusion: the Anderson and Rubin (1949) test rejects  $H_0 : \beta_g = 0$  at 5%, and the inverted 95% AR confidence set for the price coefficient on  $\widehat{\Delta N_i^{\text{gap}}}$  is bounded and excludes zero ( $[+0.155, +1.641]$ ). A pre-COVID placebo first stage on 2017–2019 cumulative employment growth is null (Online Appendix Table 14, Panel B). The substantive 2SLS results are robust across five fixed-effects configurations (Online Appendix Table 15) and across the same services-exclusion sample cuts that the data-section reduced-form survives (Online Appendix Table 16; the corresponding reduced-form robustness on the data-section identifying object is in Table 13).

**Table 8:** Fact 3 recovery regression: data versus model.

Coefficient	Data (BEA-357, $N = 321$ )		Model (24 sub-units)	
	OLS	IV (2SLS)	Adaptive	CG
Supplier exposure $\mathcal{E}_i$	3.93 (0.43)	21.04 (1.29)	-0.05 (0.40)	-0.05 (0.37)
2020 collapse $\Delta N_i^{\text{gap}}$	0.07 (0.95)	0.60** (2.31)	0.68*** (3.78)	0.64*** (3.84)
Own recovery $\Delta N_i^{\text{rec}}$	-0.14** (2.19)	-0.33** (2.33)	-0.33 (1.04)	-0.31 (1.06)
Supplier recovery $\Delta \mathcal{N}_i^{\text{rec}}$	-1.59 (1.41)	-4.41** (2.20)	-0.16 (0.15)	-0.16 (0.17)
Fixed effects	BEA-66 + chain-tier-12	BEA-66 + chain-tier-12	chain-tier-12	chain-tier-12
Estimator	WLS, clustered SE	2SLS, clustered SE	OLS	OLS
$N$	321	321	24	24

*Notes:* Dependent variable in all columns is cumulative  $L$ -propagated price growth 2019–2022 in percentage points. Data columns drop the JEM exposure index  $E_i$  from the right-hand side: it does not appear in the OLS column and enters the IV column only as the instrument for  $\Delta N_i^{\text{gap}}$ , so both columns identify the model’s employment-channel coefficients on the same right-hand side as the model regression. Data OLS uses WLS with N-GO PCE-share weights and BEA-66 sector-cluster-robust standard errors; data IV is 2SLS with the same controls and clustering. The first-stage  $F$  on  $E_i$  is 22.00, above the Stock–Yogo 10% critical value of 16.38. Absolute  $t$ -statistics in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

*Sign convention.* All four columns are reported on the data-side convention  $\Delta N_i^{\text{gap}} \equiv \max(0, -100 \Delta \log \text{Emp}_{i,2019 \rightarrow 2020})$  (positive for a 2020 collapse). The model regression uses the signed cumulative log employment growth (negative for a collapse); the displayed model coefficient on the 2020 row negates this so that all four columns are on the same convention. Under this convention the wall mechanism predicts a positive coefficient. Own and supplier recovery are cumulative log employment growth 2020–2022 (the same convention in data and model).

*IV caveat.* The 2SLS column hinges on the exclusion restriction that the JEM exposure index affects 2022 prices only through the 2020 labor channel. The data section explicitly does not adopt this restriction for its main identification (see the relevant footnote in Section 2); the IV column is reported here because it provides the per-pp-employment translation that makes the data and model coefficients directly comparable in magnitude. Read it as indicative of structural scale, but potentially biased if other channels (pent-up demand, cost-push) run through  $E_i$  alongside the labor channel.

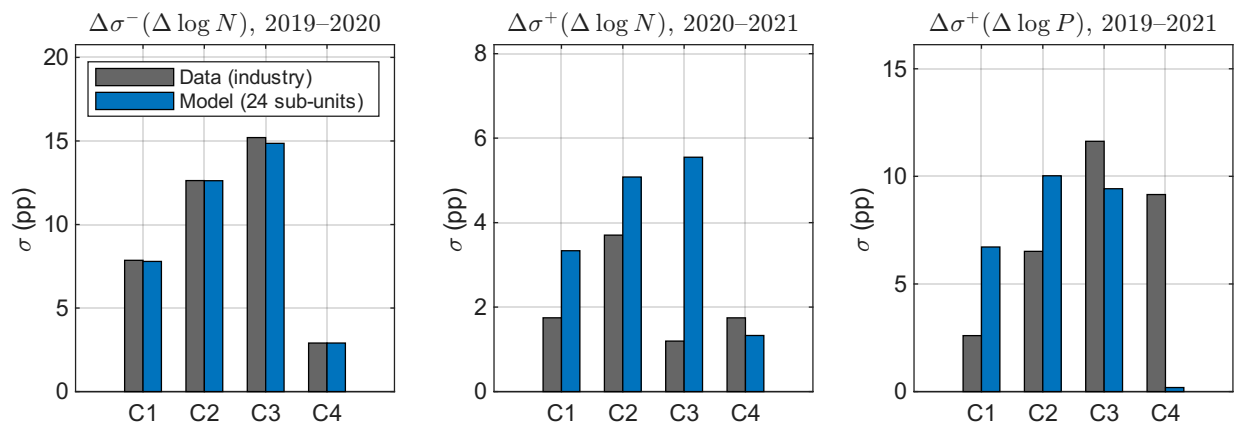
*Model side.* The model regression is the data-side specification on the model’s 24 hit/safe sub-units (4 chains  $\times$  3 finalness tiers  $\times$  2 hit/safe), with three adaptations: only chain-by-tier-12 fixed effects (the data side also has BEA-66 sector fixed effects, but  $N = 24$  does not permit an analog); supplier-weighted variables built with the model’s  $\Omega_{24}$  IO matrix; the hit/safe binary dummy is on the model-side right-hand side but is omitted from display here for direct comparability with the data columns that have no exposure regressor. The supplier-recovery row in the model is small and noisy: at  $N = 24$  with 11 chain-tier dummies absorbed, residual variation is too thin to identify the supplier channel independently of the own channel. What the model reproduces cleanly is the own-collapse and own-recovery sign and magnitude tests that anchor Fact 2.

same sign convention as the data (the model regresses on cumulative log employment growth, which is negative for a 2020 collapse, while the data uses the gap  $\Delta N^{\text{gap}} \equiv \max(0, -\Delta \log \text{Emp})$ ); flipping the model’s coefficient to the data convention gives +0.68\*\*\* in the model versus +0.60\*\* in the data 2SLS). Both numbers say the same thing: a one-percentage-point deeper 2020 own collapse maps to roughly 0.6–0.7 pp more cumulative 2019–2022 inflation through the wall’s marginal-cost wedge. The data OLS

column attenuates this to near zero (+0.07, insignificant); the structural object is recovered cleanly only by the 2SLS. The CG-spec column delivers essentially the same coefficients ( $\hat{\pi}^r = -0.33$  adaptive vs  $-0.31$  CG; +0.68 vs +0.64 on the 2020 collapse under the data sign convention).

The model predicts differential inflation along the supply chain, and there is a clear pattern in the data: upstream industries inflated more. This is a challenging test because industries that are more consumer facing were hit harder. Can the mechanism generate more inflation in upstream industries?

**Upstream versus downstream inflation.** Figure 7 shows that the model implies upward price-growth dispersion  $\Delta\sigma^+(\Delta \log P)$  through 2022 and upward employment-growth dispersion  $\Delta\sigma^+(\Delta \log N)$  during the peak rebuilding window. In the model’s steady state both measures are zero, and so the relevant comparison is excess dispersion in the data. To construct it, we subtract the 2017–2019 baseline dispersion—which we indicate by adding ‘ $\Delta$ ’ to our notation.



**Figure 7: Within-chain dispersion: data vs. model.** Employment-growth dispersion (left) and upward price-growth dispersion (right) by chain and period. Gray: data at the industry level within each chain. Color: model aggregated to the chain–tier level. The model captures the direction and amplification ratio but underpredicts the level of dispersion because the tier aggregation compresses industry-level heterogeneity.

Quantitatively, the model is off, but in a specific direction: it implies faster employment adjustment and smaller price dispersion than the data, indicating smaller network disruption in the model relative to the data.<sup>24</sup>

<sup>24</sup>Across chains, the fourth (manufacturing) chain is notably absent: its price dispersion is low in the model and the wall

**Table 9:** L-decomposition (Leontief) of headline inflation into wage, productivity, and  $\Delta$ MARGIN contributions, PCE-weighted (see eq. (4) in Section 2.1).

Period	Side	$\Delta \log P$	$B \Delta \log W$	$-\Delta \log(Y/N)$	$\Delta$ MARGIN
Act I (2019–2021)	Data	+6.22	+13.42	−11.25	+4.05
	Model	−3.31	+3.48	−11.21	+4.42
	Model $\mu = 0$	−3.31	+3.48	−11.21	+4.42
Act II (2021–2023)	Data	+14.25	+12.40	−6.91	+8.76
	Model	+4.32	+5.97	−6.82	+5.17
	Model $\mu = 0$	+4.32	+5.97	−6.82	+5.17
2019–2024	Data	+16.94	+24.22	−14.13	+6.84
	Model	+1.53	+10.64	−14.01	+4.91
	Model $\mu = 0$	+1.53	+10.64	−14.01	+4.91

*Notes:* Cumulative log change in PCE-weighted aggregate price level (pp), decomposed via the L-decomposition of equation (4),  $\Delta \log P = L [\Delta$ MARGIN +  $\Delta$ ULC +  $\Delta$ IMP], where  $\Delta$ ULC =  $B \Delta \log W - \Delta \log(Y/N)$  is split into its two columns here (wage and productivity contributions).  $B = \text{diag}(a_i + (1 - a_i)\rho)$ ; the imports channel is absorbed into the  $\Delta$ MARGIN column. Sum of three components  $\approx \Delta \log P$ .

**Decomposition of inflation’s structural sources.** Finally, we assess the model’s attribution of inflation in the sense of the L-decomposition in Section 2.1, we apply the same methodology to the data aggregated to the model level and compare the two. Table 9 shows the results. The model matches the data well.

## 6 Conclusion

We have shown that the post-COVID inflation surge was in part propagated by the destruction of productive capacity in 2020. Demand and supply forces are intertwined, but the key implication for the demand shock is that the relevant supply reference is not the 2019 supply but the diminished supply between 2019 and 2020.

contribution is almost nil. There are two reasons for this. First, there is a mechanical limitation: the model’s hit/safe cut at chain  $\times$  tier level is limited by the BEA-detail resolution at which the data dispersion is measured, so joint price and employment dispersion cannot be matched at the same granularity. Specifically, the hit/safe assignment in C4\_1 is dominated by Construction (NAICS 23, 82% of tier employment); no employment-balanced median split is possible when one industry is that large, so the within-tier CES composite collapses to a single unit and the hit/safe channel is mass-suppressed in C4\_1. Our classification splits hit/safe by median employment to the extent that it is possible, and this tier is an extreme outlier. Second, the manufacturing sector fell more uniformly in 2020, so the cross-sectional dispersion of employment growth is genuinely low in this chain, including the more balanced m and h tiers.

Our calibrated nonlinear multisector New Keynesian model assumes workforce rebuilding-cost premia consistent with the labor literature, and implies that this degree of convexity is sufficient to account for the post-COVID inflation not attributable to the exogenous sectors. About half of that contribution is attributable to the cross-sectional unevenness of the shock. The result is an endogenous steepening of the Phillips curve that does not require any change in the economy’s underlying structure. As for nonlinearity, our calibration implies that the Phillips curve is nonlinear in response to large shocks, but not as nonlinear as COVID made it look.

Several lessons follow from our analysis. The endogenous steepening of the Phillips curve we document is not, in large part, a new structural feature of the economy: the calibrated wall is dormant under shocks that do not push the labor market into the tight-rebuilding regime, and at least half of its inflation contribution under COVID owes to shock geometry rather than to aggregate scale. This result suggests that aggregate output stabilization tools were poorly matched to a shock whose inflationary bite came mostly from geometry. As an example, monetary policy is a blunt tool that slows the economy as a whole, and it cannot fix asymmetry-driven inflation without causing damage on the employment side. We leave a closer examination of this tradeoff to future work.

## References

- Aeppli, Manuel, Samuel Muehlemann, Harald Pfeifer, Jürg Schweri, Felix Wenzelmann, and Stefan C. Wolter**, “The Impact of Hiring Costs for Skilled Workers on Apprenticeship Training: A Comparative Study,” Discussion Paper 16919, IZA Institute of Labor Economics April 2024. [1](#)
- Anderson, T. W. and Herman Rubin**, “Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations,” *Annals of Mathematical Statistics*, 1949, 20 (1), 46–63. [23](#), [A.2.5](#)
- Andrade, Philippe, Falk Bräuning, José L. Fillat, and Gustavo Joaquim**, “Is Post-Pandemic Wage Growth Fueling Inflation?,” Current Policy Perspectives 24-1, Federal Reserve Bank of Boston 2024. [2](#)
- Arias, Jonas, Juan F. Rubio-Ramírez, and Minchul Shin**, “Are Fiscal Transfers Inflationary?,” 2026. Working paper. [1](#)

- Atalay, Englin**, “How Important Are Sectoral Shocks?,” *American Economic Review*, 2017, 107(8), 2235–2267. [2.1](#)
- Barrot, Jean-Noël and Julien Sauvagnat**, “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks,” *The Quarterly Journal of Economics*, 2016, 131 (3), 1543–1592. [1](#)
- 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,” Working Paper 31197, National Bureau of Economic Research 2023. [1](#)
- Bernanke, Ben S. and Olivier J. Blanchard**, “What Caused the U.S. Pandemic-Era Inflation?,” *American Economic Journal: Macroeconomics*, 2025, 17 (3), 1–35. [1](#)
- Blatter, Marc, Samuel Muehlemann, and Samuel Schenker**, “The costs of hiring skilled workers,” *European Economic Review*, January 2012, 56 (1), 20–35. [19](#)
- Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar**, “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake,” *Review of Economics and Statistics*, 2019, 101 (1), 60–75. [2.1](#)
- California Policy Lab**, “Analysis of Unemployment Insurance Claims in California,” Technical Report, California Policy Lab, UCLA and UC Berkeley December 2021. [5](#)
- Cavallo, Alberto, Francisco Llamas, and Franco Vazquez**, “Tracking the Short-Run Price Impact of U.S. Tariffs,” *National Bureau of Economic Research Working Paper*, 2024, (33099). [4](#)
- Cerrato, Andrea and Giulia Gitti**, “The Return of the Phillips Curve: Evidence from US Cities,” 2022. Working paper. [1](#)
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan**, “Business Cycle Accounting,” *Econometrica*, 2007, 75 (3), 781–836. [3.4](#)
- Coibion, Olivier and Yuriy Gorodnichenko**, “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 2015, 105 (8), 2644–2678. [4](#), [4](#), [21](#), [A.6](#), [A.6.3](#)
- Comin, Diego, Robert C. Johnson, and Callum Jones**, “Supply Chain Constraints and Inflation,” Working Paper 31179, National Bureau of Economic Research 2023. [1](#)
- Decker, Ryan A. and John Haltiwanger**, “Business Entry and Exit in the COVID-19 Pandemic: A Preliminary Look at Official Data,” FEDS Notes 2022. [1](#)

- di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yıldırım**, “Global Supply Chain Pressures, International Trade, and Inflation,” Working Paper 30240, National Bureau of Economic Research 2022. [6](#)
- Dingel, Jonathan I. and Brent Neiman**, “How Many Jobs Can Be Done at Home?,” *Journal of Public Economics*, 2020, 189, 104235. [1](#), [2.3](#)
- Domash, Alex and Lawrence H. Summers**, “How Tight Are U.S. Labor Markets?,” Working Paper 29739, NBER 2022. [1](#), [4](#)
- Dupor, Bill, Marios Karabarbounis, Marianna Kudlyak, and M. Saif Mehkari**, “Regional Consumption Responses and the Aggregate Fiscal Multiplier,” Working Paper 2018-04, Federal Reserve Bank of San Francisco 2022. [3.6](#)
- Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo**, “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, 2011, 101 (1), 234–262. [6](#), [4](#)
- Ferrante, Francesco, Simon Graves, and Matteo Iacoviello**, “The Inflationary Effects of Sectoral Reallocation,” *Journal of Monetary Economics*, 2023, 140, 64–81. [1](#)
- Forsythe, Eliza, Lisa B. Kahn, Fabian Lange, and David Wiczer**, “Where have all the workers gone? Recalls, retirements, and reallocation in the COVID recovery,” *Labour Economics*, 2022, 78, 102251. [5](#), [6](#)
- Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens**, “Anatomy of the Phillips Curve: Micro Evidence and Macro Implications,” *American Economic Review*, November 2025, 115 (11), 3941–3974. [1](#)
- Giannone, Domenico and Giorgio E. Primiceri**, “The Drivers of Post-Pandemic Inflation,” Working Paper 32859, National Bureau of Economic Research 2024. [1](#)
- Hall, Robert E.**, “The Relation between Price and Marginal Cost in U.S. Industry,” *Journal of Political Economy*, 1988, 96 (5), 921–947. [A.1.1](#)
- Harding, Martin, Jesper Lindé, and Mathias Trabandt**, “Understanding post-COVID inflation dynamics,” *Journal of Monetary Economics*, 2023, 140 (S), 101–118. [1](#)
- Haubrich, Joseph, George Pennacchi, and Peter Ritchken**, “Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps,” *Review of Financial Studies*, 2012, 25 (5), 1588–1629. [4](#), [A.6.1](#)

- Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson**, “The Slope of the Phillips Curve: Evidence from U.S. States,” *Quarterly Journal of Economics*, 2022, 137 (3), 1299–1344. [1](#), [4](#), [6](#)
- Hengel, Karen M Oude, Alex Burdorf, Anjoeka Pronk, Vivi Schlünssen, Zara A Stokholm, Henrik A Kolstad, Karin van Veldhoven, Ioannis Basinas, Martie van Tongeren, and Susan Peters**, “Exposure to a SARS-CoV-2 infection at work: development of an international job exposure matrix (COVID-19-JEM),” *Scandinavian journal of work, environment & health*, 2021, 48 (1), 61. [1](#), [2](#), [2.3](#), [A.2.1](#), [A.2.1](#), [12](#), [3](#)
- Karadi, Peter and Adam Reiff**, “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*, 2019, 11 (3), 111–146. [4](#)
- Linckh, Carolin, Samuel Muehlemann, and Harald Pfeifer**, “Beggars Cannot Be Choosers: Labor Market Tightness and Hiring Standards, Wages, and Hiring Costs,” Working Paper 247, Swiss Leading House “Economics of Education” August 2025. [1](#), [4](#), [6](#), [4](#), [19](#)
- Loecker, Jan De, Jan Eeckhout, and Gabriel Unger**, “The Rise of Market Power and the Macroeconomic Implications,” *Quarterly Journal of Economics*, 2020, 135 (2), 561–644. [A.1.1](#)
- Mankiw, N. Gregory and Ricardo Reis**, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 2002, 117 (4), 1295–1328. [A.6](#), [A.6.3](#)
- Miyauchi, Yuhei**, “Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade,” *Econometrica*, 2024, 92 (6), 1869–1905. [1](#)
- Montiel Olea, José L. and Carolin Pflueger**, “A Robust Test for Weak Instruments,” *Journal of Business & Economic Statistics*, 2013, 31 (3), 358–369. [A.2.5](#)
- Moscarini, Giuseppe and Fabien Postel-Vinay**, “The Job Ladder: Inflation vs. Reallocation,” Working Paper 31466, National Bureau of Economic Research 2023. [6](#)
- Muehlemann, Samuel and Harald Pfeifer**, “The Structure of Hiring Costs in Germany,” *Industrial Relations*, 2016, 55 (2), 193–218. [1](#), [4](#), [19](#)
- **and Mirjam Strupler Leiser**, “Hiring Costs and Labor Market Tightness,” *Labour Economics*, 2018, 52, 122–131. [1](#), [4](#), [4](#)
- Nakamura, Emi and Jón Steinsson**, “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 2008, 123 (4), 1415–1464. [6](#), [4](#)

- Negro, Marco Del, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti**, “What’s Up with the Phillips Curve?” *Brookings Papers on Economic Activity*, 2020, 2020 (1), 301–373. [1](#), [4](#)
- Peltzman, Sam**, “Prices Rise Faster than They Fall,” *Journal of Political Economy*, 2000, 108 (3), 466–502. [4](#)
- Rubbo, Elisa**, “What Drives Inflation? Lessons from Disaggregated Price Data,” BFI Working Paper 2024-24, Becker Friedman Institute for Economics, University of Chicago 2024. [1](#)
- Stock, James H. and Motohiro Yogo**, “Testing for Weak Instruments in Linear IV Regression,” in Donald W. K. Andrews and James H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 2005, pp. 80–108. [23](#)
- van der Feltz, Sophie, Vivi Schlünssen, Ioannis Basinas, Luise M Begtrup, Alex Burdorf, Jens P E Bonde, Esben M Flachs, Susan Peters, Anjoeka Pronk, Zara A Stokholm, Martie van Tongeren, Karin van Veldhoven, Karen M Oude Hengel, and Henrik A Kolstad**, “Associations between an international COVID-19 job exposure matrix and SARS-CoV-2 infection among 2 million workers in Denmark,” *Scandinavian Journal of Work, Environment & Health*, 2023, 49 (6), 375–385. [A.2.1](#)
- Yellen, Janet L.**, “Inflation Dynamics and Monetary Policy,” September 2015. The Philip Gamble Memorial Lecture, University of Massachusetts at Amherst, September 24, 2015. [1](#)

# Online Appendix

<b>A.1</b>	<b>Accounting Framework Derivation</b>	<b>56</b>
<b>A.2</b>	<b>JEM Exposure Construction and IV Robustness</b>	<b>60</b>
<b>A.3</b>	<b>Chain Construction</b>	<b>67</b>
<b>A.4</b>	<b>Mapping the Model onto the Data</b>	<b>68</b>
<b>A.5</b>	<b>Data Detrending</b>	<b>71</b>
<b>A.6</b>	<b>Inflation Expectations: Identification and Robustness</b>	<b>72</b>
<b>A.7</b>	<b>Robustness and Sensitivity</b>	<b>77</b>
<b>A.8</b>	<b>Retail Sector Profit Margins in Other Datasets</b>	<b>82</b>
<b>A.9</b>	<b>Data Sources</b>	<b>85</b>

Section [A.1](#) derives the A- and L-decompositions of Section [2](#)'s accounting framework and reports our estimate of the labor–materials CES elasticity  $\rho$ . Section [A.2](#) documents the bottom-up construction of the JEM exposure index  $E_i$  used as the instrument in Section [2](#)'s cross-industry regressions, and reports the weak-instrument battery, FE-configuration robustness, and services-exclusion robustness for the 2SLS specification. Section [A.3](#) documents the spectral-clustering procedure that produces the four supply chains of the structural model in Section [3](#). Section [A.4](#) maps the model's price, employment, and value-added aggregates onto their data counterparts. Section [A.5](#) describes the 2017–2019 chain-level pre-trends used to detrend calibration targets and reported model series throughout Section [4](#). Section [A.6](#) reports the identification of the calibrated expectations operator, including the headline first-difference regression, the rolling-window stability check on the past-pass-through coefficient, and the constrained level-form test of  $v_{\text{forward}}$ . Section [A.7](#) reports model sensitivity to wall intensity, downward price rigidity, CES elasticities, the recalibrated no-wall counterfactual, the wall decomposition under CG-consistent expectations, and the cost-price pass-through scatter by hit/safe. Section [A.8](#) cross-validates the Act-I retail margin expansion using BEA-BLS KLEMS production accounts. Section [A.9](#) lists primary data sources.

## A.1 Accounting Framework Derivation

This appendix provides the derivation of the A- and L-decompositions presented in the main text (equations [\(3\)](#) and [\(4\)](#)), including the commodity–industry distinction via the make and use tables.

Each industry  $i = 1, \dots, n$  has CES production as in the main text (equation [2](#)). Following BEA's supply-use framework, industries may produce multiple commodities. Commodity output and material bundles are Cobb–Douglas aggregates:

$$Y_{ct} = \prod_{i=1}^n Y_{cit}^{\theta_{ci}}, \quad m_{it} = \prod_{c=1}^{n'} \left[ M_{cit}^{\zeta_{ci}} \tilde{M}_{cit}^{1-\zeta_{ci}} \right]^{\omega_{ci}}, \quad (\text{A.1})$$

where  $M_{cit}$  and  $\tilde{M}_{cit}$  denote domestically sourced and imported commodity  $c$ ,  $\theta_{ci}$  are commodity supply shares (make matrix  $\Theta$ ),  $\omega_{ci}$  are input use shares (use matrix  $\Omega$ ), and  $\zeta_{ci}$  are domestic sourcing shares, with  $\sum_i \theta_{ci} = 1$

and  $\sum_c \omega_{ci} = 1$ . The Cobb–Douglas form is a first-order approximation—a CES aggregation gives identical results for log growth rates.

Static cost minimization defines variable cost  $C_i(Y_i) := \min_{N_i, m_i} \{W_i N_i + p_i m_i\}$  subject to (2), where  $p_i$  is the Cobb–Douglas price index of materials weighted by  $\omega_{ci}$  and  $\zeta_{ci}$ .

For any variable  $X_{it}$ , write  $\Delta \log X_i := \log X_{it} - \log X_{it-1}$ . Log-differencing the first-order and price-index conditions yields the substitution response and material-price decomposition:

$$\Delta \log m_i = \Delta \log N_i - \rho(\Delta \log p_i - \Delta \log W_i), \quad (\text{A.2})$$

$$\Delta \log p_i = \sum_c \omega_{ci} \zeta_{ci} \sum_{i'} \theta_{ci'} \Delta \log P_{i'} + \sum_c \omega_{ci} (1 - \zeta_{ci}) \Delta \log \tilde{P}_c. \quad (\text{A.3})$$

Equation (A.2) is the log-differentiated CES first-order condition: materials demand tracks labor but falls when materials become relatively expensive, with the substitution response governed by  $\rho$ . Equation (A.3) decomposes material price growth into I–O-weighted domestic and import components.

The cost-share weight  $a_i$  is computed from observable labor shares:

$$a_i := \frac{\delta_i^{1/\rho}}{\delta_i^{1/\rho} + (1 - \delta_i)^{1/\rho}} = \frac{ls_i}{ls_i + \sqrt{ls_i(1 - ls_i)}}, \quad ls_i := \frac{W_i N_i}{C_i}, \quad (\text{A.4})$$

where  $ls$  is labor's share of variable costs, and  $\delta_i = (1 + (a_i^{-1} - 1)^\rho)^{-1}$ .

Define industry  $i$ 's MARGIN as  $M_i := P_i Y_i / C_i(Y_i)$ —the ratio of revenue to variable cost. Combining the unit-cost identity  $\Delta \log P_i = \Delta \log M_i + \Delta \log(C_i/Y_i)$  with the log-linearized variable cost and substituting out TFP growth via  $\Delta \log Z_i \approx \alpha_i(\Delta \log(Y/N)_i + \rho(1 - a_i)(\Delta \log p_i - \Delta \log W_i))$ , the unit-cost growth can be written as  $\Delta \log(C_i/Y_i) \approx (a_i + (1 - a_i)\rho)\Delta \log W_i + (1 - a_i)(1 - \rho)\Delta \log p_i - \Delta \log(Y/N)_i$ .<sup>25</sup>

### A.1.1 The A-Decomposition (single pass)

Stacking across industries and using (A.3) to split material price growth into domestic and import components yields the industry-level pricing identity in vector form:

$$\Delta \log P = \Delta \text{MARGIN} + \Delta \text{ULC} + \Delta \text{MAT} + \Delta \text{IMP}, \quad (\text{A.5})$$

where the four components are

$$\begin{aligned} \Delta \text{MARGIN} &:= \Delta \log M, & \Delta \text{ULC} &:= B \Delta \log W - \Delta \log(Y/N), \\ \Delta \text{MAT} &:= A \Delta \log P, & \Delta \text{IMP} &:= D \left( (1 - \zeta) \odot \Omega \right)^\top \Delta \log \tilde{P}, \end{aligned}$$

with the diagonal weights

$$B = \text{diag}(a_i + (1 - a_i)\rho), \quad D = \text{diag}((1 - a_i)(1 - \rho)),$$

<sup>25</sup>The approximation is a first-order expansion around the symmetric point  $N_i = m_i$  ( $p_i = W_i$ ), without loss given the flexibility of  $\delta_i$ .

scaling the wage contribution and the materials channel by labor–materials complementarity, respectively, and the domestic-input matrix

$$A := D(\zeta \odot \Omega)^\top \Theta$$

encoding single-step network propagation. Element  $A_{ij}$  is the share of one dollar of industry- $i$  output spent on domestic intermediates from supplier- $j$ , corrected for  $\rho$ . Under Cobb–Douglas production ( $\rho = 1$ ),  $M_i$  reduces to the standard markup of price over marginal cost (Hall, 1988; De Loecker et al., 2020); more generally it captures the joint effect of markups and returns to scale.

We refer to (A.5) as the **A-decomposition** (single pass): it keeps the supplier-cost passthrough  $\Delta\text{MAT}$  as its own term.

### A.1.2 The L-Decomposition (geo pass)

Substituting  $\Delta \log P$  from (A.5) recursively into  $\Delta\text{MAT} = A \Delta \log P$  unrolls the materials term:

$$\begin{aligned} \Delta\text{MAT} &= A [\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP} + A \Delta \log P] \\ &= A [\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}] + A^2 \Delta \log P \\ &= \dots = (A + A^2 + A^3 + \dots) [\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}], \end{aligned}$$

where the series converges under standard regularity conditions on  $A$ . Substituting back into (A.5) yields the **L-decomposition** stated in text:

$$\Delta \log P = L [\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}], \quad (\text{A.6})$$

where  $L := (I - A)^{-1} = I + A + A^2 + A^3 + \dots$  is the Leontief inverse. Powers of  $A$  collect  $k$ -th-order upstream exposures:  $A_{ij}^2$  is  $i$ 's exposure to  $j$  through one intermediate supplier,  $A_{ij}^3$  through two, and so on. The L-decomposition attributes the materials channel recursively to upstream primitive cost sources under the assumption of full upstream passthrough; the A-decomposition stops at one step with  $\Delta\text{MAT}$  visible as its own bucket. Both deliver the same  $\Delta \log P$ .

### A.1.3 PCE Aggregation

The headline aggregate inflation series is obtained by pre-multiplying either decomposition by the PCE weight vector  $w := \Theta^\top c$ , where  $c$  is the PCE final-demand vector and  $\Theta$  is the I–O commodity-composition (make) matrix:

$$w^\top \Delta \log P = w^\top \Delta\text{MARGIN} + w^\top \Delta\text{ULC} + w^\top \Delta\text{MAT} + w^\top \Delta\text{IMP} = w^\top L [\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}]. \quad (\text{A.7})$$

The elasticity  $\rho$  enters through both  $D$  (scaling network propagation) and  $B$  (scaling wage passthrough): lower  $\rho$  amplifies the contribution of material-price shocks relative to wages.

The headline measure throughout the paper is N-GO inflation—this aggregate with the exogenous-in-model blocks (housing, oil and petroleum, finance/insurance/government, and metal-ore mining) removed.

### A.1.4 Estimation of $\rho$

We estimate  $\rho$  using the CES first-order condition and the BEA/BLS KLEMS panel (63 industries, 1997–2023):

$$\log \frac{p_i m_i}{W_i N_i} = FE_i + FE_t + (1 - \rho) \log \frac{p_i}{W_i} + \varepsilon_i. \quad (\text{A.8})$$

The baseline estimate is  $\hat{\rho} = 0.39$  (95% CI: [0.06, 0.72]), indicating significant labor–material complementarity. Differenced specifications yield lower point estimates ( $\approx 0.1$ – $0.2$ ), consistent with attenuation from measurement error in first differences. These values are close to the literature and our baseline elasticity is thus fairly conservative. The sensitive of our results is moderate and while elasticity matters, our result are similar when  $\rho = 0.5$ , for example.

### A.1.5 Headline inflation measure comparison to published indices

The headline inflation used throughout the paper is a fixed-weight aggregate of detail (357-industry SUT) level gross-output deflators, with industry weights  $w_i = (\Theta^\top c)_i$  derived from the 2017 PCE final-demand vector  $c$  and the I–O commodity-composition matrix  $\Theta$ . This producer-price construction preserves the role of wholesale, retail, and transportation margins when moving from producer to consumer expenditure weights, and it aligns our empirical aggregate with the I–O structure used in the model. It is not the same as the headline BEA PCE price index: BEA PCE uses consumer-price source data, an imported final-goods treatment, and an imputed housing-services methodology that are not recoverable from domestic gross-output deflators alone. The differences are quantitatively visible but do not change the qualitative pattern of the post-COVID episode.

Table 10 reports our headline measure (GO inflation) and its narrower networked counterpart (N-GO inflation, the same aggregate with the exogenous-in-model blocks removed: housing, oil and petroleum, finance/insurance/government, and metal-ore mining) alongside CPI, BEA headline PCE, and BEA Core PCE at annual horizons from 2019.

**Table 10:** Cumulative inflation from 2019 (log-%, pp): paper’s measure vs. BEA and BLS published indices.

Year	Headline measures			Ex-volatile measures	
	GO Inflation	BEA PCE	CPI	N-GO Inflation	BEA Core PCE
2020	+1.49	+1.11	+1.25	+1.23	+1.35
2021	+7.14	+5.14	+5.82	+5.20	+4.84
2022	+14.17	+11.48	+13.51	+10.26	+10.04
2023	+16.79	+15.21	+17.55	+11.92	+14.13
2024	+19.13	+17.79	+20.46	+13.17	+17.03

*Notes:* Cumulative log-percent change from 2019 (percentage points). **GO Inflation** is the paper’s headline inflation measure: a fixed-weight aggregate of BEA detail (357-industry SUT) gross-output deflators, with industry weights  $w_i = (\Theta^\top c)_i$  derived from the 2017 PCE final-demand vector  $c$  and the I–O commodity-composition matrix  $\Theta$ . **BEA PCE** is the BEA-published Personal Consumption Expenditures price index (FRED PCEPT). **CPI** is the BLS Consumer Price Index for all urban consumers (FRED CPIAUCSL). **N-GO Inflation** is GO Inflation with exogenous-in-model blocks removed: Housing, Oil & Petroleum, Mining, Finance, and Government (together about 29% of PCE basket weight). **BEA Core PCE** excludes food and energy but retains housing (FRED PCEPILFE).

Two methodological gaps with the BEA-published indices are worth flagging. First, our GO inflation runs about 1.3 pp hotter than BEA PCE through 2024 and tracks CPI more closely, with the gap opening in 2022–2024.

The discrepancy reflects the producer-price weighting on domestic gross output: imported final goods and imputed residential rents both moved by less than the producer prices we aggregate during the 2021–2022 peak, so a producer-price aggregation runs ahead of the BEA’s consumer-price aggregation in that window. Second, our N-GO measure runs about 3.9 pp *colder* than BEA Core PCE because N-GO excludes real estate, finance/insurance/government, and oil/petroleum/mining—roughly 29% of the headline PCE basket—whereas Core PCE nets out only food and energy. The two exclusion conventions are not nested: N-GO retains food and energy production within the network but drops the exogenously priced blocks, while Core PCE retains the exogenously priced blocks but drops food and energy.

## A.2 JEM Exposure Construction and IV Robustness

This appendix documents the construction of the industry-level COVID exposure index  $E_i$  used in the cross-industry reduced-form analysis of Section 2 (equation (5)). The construction proceeds in three stages and uses only pre-pandemic inputs. We also discuss the population-level validation of the Danish country panel of the underlying Job-Exposure Matrix and report a robustness panel using the UK country panel. The legacy subsections at the end (weak-instrument diagnostics, FE-robustness, services-purged samples) document a 2SLS specification that was considered during development but is not adopted in the main text—see the relevant footnote in Section 2 on why the exclusion restriction cannot be defended on empirical grounds for this application.

### A.2.1 JEM exposure construction

**Background on the COVID-19 JEM.** The COVID-19 Job-Exposure Matrix (Oude Hengel et al., 2021) was assembled by a team of occupational epidemiologists from Denmark, the Netherlands, and the United Kingdom. The matrix scores every 4-digit ISCO-08 occupation along eight workplace-exposure dimensions that the authors flag as relevant for SARS-CoV-2 transmission risk. Because the underlying epidemiological evidence and working-environment conditions differ across countries, the JEM reports country-specific score sets for Denmark, the Netherlands, and the UK rather than a single international score. Of the three country panels, the Danish panel has been validated against a population-scale infection register and is the benchmark recommended by the JEM authors (see *Validation against the Danish population register* below); our main-text headline uses the DK panel, and Online Appendix A.2.2 reports the UK-panel robustness.

The index  $E_i$  is built bottom-up in three stages, all from pre-pandemic inputs—no realized 2020 or post-2020 outcome enters. Each industry’s exposure is the OEWS-employment-share-weighted average of its constituent occupations’ contact-intensity scores from the JEM.

**Stage 1: Per-occupation contact-intensity score (ISCO-08), DK panel.** *Source:* Supplementary file COVID-19 JEM.xlsx released with Oude Hengel et al. (2021), Denmark country sheet, one row per ISCO-08 4-digit occupation with eight numerical columns. We retain the five contact-intensity columns that proxy COVID transmission risk in the workplace: Nature of contacts, Contaminated workspaces, Location, Social distance, and Face Covering. The remaining three columns (Number, Income insecurity, Migrants) are dropped because they capture workforce size, downstream consequences, and demographic composition rather than transmission risk. For each ISCO-08 occupation  $o$ , we compute the contact-intensity score as the mean across the five

retained dimensions,

$$\text{exposure}_o = \text{contact\_score}_o^{\text{DK}} = \frac{1}{5} \sum_{c \in \{1, \dots, 5\}} \text{score}_o^{\text{DK}}(c) \in [0, 3].$$

The Online Appendix [A.2.2](#) reports analogous results using  $\text{contact\_score}_o^{\text{UK}}$  instead.

**Validation against the Danish population register.** [van der Feltz et al. \(2023\)](#) validate the Danish panel of the JEM against population-scale individual-level SARS-CoV-2 infection records. Four features of the Danish data make this validation exceptionally tight relative to validation studies on the UK and NL panels:

1. *Population scale, no sampling bias.* The validation covers 2,021,309 active workers across the country—essentially the universe of Danish workers, with nearly all occupations represented and no small-cohort selection.
2. *Free, symptom-independent mass testing.* Danish government test centers provided PCR and antigen tests free of charge, with no symptom or medical-indication requirements, eliminating the symptom-conditional and barrier-conditional selection that contaminate infection data in most other countries.
3. *Multi-wave validation.* Infection associations are estimated across three distinct pandemic waves with varying restrictions: Wave 1 (Feb 26–Aug 24, 2020), where high JEM scores most strongly predict infection; Wave 2 (Aug 25, 2020–Jun 21, 2021), where associations attenuate slightly as community-wide mixing expands; and Wave 3 (Jun 22–Dec 12, 2021), where infections decouple from JEM scores among fully vaccinated workers but the association remains among the unvaccinated.
4. *Clean register linkage.* Personal CPR registry numbers are paired to the Danish Microbiology Database (MiBa) and to standardized ISCO-08 4-digit occupational records, so the unit of analysis is the validated individual-occupation-infection observation rather than a survey-based proxy.

Together these features make Denmark the strongest of the three country panels for the purpose this paper requires: an external, individual-level, population-scale validation that the JEM contact-intensity dimensions do capture differential SARS-CoV-2 transmission risk across occupations. Validation studies on the UK and NL panels rely on smaller cohorts and more restrictive testing protocols, which is why the JEM authors recommend the DK panel as their benchmark and why we use it as our main-text exposure measure.

**Stage 2: Mapping ISCO-08 to U.S. SOC-2018.** The OEWS data used in Stage 3 is keyed on U.S. SOC-2018 codes; the JEM is keyed on ISCO-08. We chain two BLS crosswalks:

$$\text{ISCO-08} \xrightarrow{\text{BLS ISCO-to-SOC-2010 crosswalk}} \text{SOC-2010} \xrightarrow{\text{BLS SOC-2010-to-SOC-2018 crosswalk}} \text{SOC-2018}.$$

For each SOC-2018 code  $s$ , the per-occupation exposure is the unweighted mean of  $\text{exposure}_o$  across all ISCO-08 codes that map to  $s$  through the chained crosswalk. The resulting score covers 865 SOC-2018 codes.

**Stage 3: BEA-357 industry exposure via OEWS occupation employment.** *Source:* the OEWS occupation-employment matrix at the BEA-357 level (a  $345 \times 789$  matrix with rows BEA-357 industries, columns

SOC codes, and cells OEWS occupation employment counts), combined with the OEWS-NAICS-to-BEA-357 mapping. For each industry  $i$  we restrict attention to the SOC columns matched in Stage 2 and compute employment shares within the matched set,

$$\text{emp\_share}_{io} = \frac{\text{oews}_{io}}{\sum_{o' \in \text{matched}} \text{oews}_{io'}}, \quad E_i = \sum_{o \in \text{matched}} \text{emp\_share}_{io} \cdot \text{exposure}_o.$$

Coverage is 745 of 789 OEWS SOC columns matched, with 86.9% of industry employment on average falling inside the matched set. The resulting industry-level exposure score covers 345 BEA-357 industries with mean  $E_i = 1.34$  and range  $[0.32, 2.22]$ .

**Face validity.** Top-exposed industries are Child Day Care, Personal Care Services, full-service and limited-service Restaurants, Hospitals, Accommodation, Nursing and Residential Care, Outpatient Care, and Couriers—uniformly high-contact in-person services. Bottom-exposed industries are Software Publishers, Data Processing, Insurance Carriers, Securities, Management Consulting, Accounting, and Computer Systems Design—uniformly work-from-home occupations.

**Pre-determined-ness.** Every input to  $E_i$  is fixed before the COVID episode: the JEM contact scores describe occupational characteristics (physical proximity to other workers, contamination risk, etc.) that are properties of the job rather than of 2020 outcomes; both BLS crosswalks (ISCO-08  $\rightarrow$  SOC-2010 and SOC-2010  $\rightarrow$  SOC-2018) predate 2020; and the OEWS occupation matrix uses pre-COVID employment shares (vintage  $\leq$  2019). The exposure index  $E_i$  is therefore fully determined by industry  $i$ 's pre-COVID occupation composition crossed with intrinsic occupation-level transmission risk; no realized 2020 or post-2020 employment, prices, or sales enter the construction.

## A.2.2 Robustness: UK country panel

The main text uses the DK country panel of the JEM, which the authors validate against a population-level Danish employment register. Table 11 re-estimates the four-column reduced form of Table ?? using the UK country panel of the JEM. Coefficient magnitudes are slightly smaller (the UK panel has lower P90–P10 dispersion in exposure than DK), but the supply-side fingerprint result is qualitatively identical: the recovery coefficient is significantly negative in column (3) ( $\hat{\pi}^r = -0.144$ ,  $|t| = 2.13$ ) and survives the addition of Dingel-Neiman teleworkability as a control in column (4) ( $\hat{\pi}^r = -0.153$ ,  $|t| = 2.17$ ). The headline cross-industry pattern (positive on exposure, negative on recovery) does not depend on the choice of country panel.

## A.2.3 Pre-COVID placebo for the cross-industry reduced form

A natural concern about the headline cross-industry reduced form (Table ??) is that the JEM exposure index  $E_i$  might be picking up a long-running sectoral inflation gradient that pre-existed COVID, in which case the post-2019 loading  $\hat{\pi}^E$  would reflect sectoral selection rather than the COVID labor shock. We rule this out with a placebo regression on the pre-COVID window. Table 12 re-estimates the column-(1) reduced form with the dependent variable replaced by cumulative L-propagated price growth over the 2017–2019 window:

$$\Delta \log P_{i,17 \rightarrow 19} = \alpha_{S(i)} + \alpha_{c(i)} + \pi_{\text{pre}}^E E_i + \pi_{\text{pre}}^{\mathcal{E}} \mathcal{E}_i + \varepsilon_i,$$

**Table 11:** Cross-industry reduced form, cumulative  $\Delta \log P_{19 \rightarrow 22}$  on exposure (JEM-UK).

	(1) Exposure only	(2) + $\Delta N^{\text{gap}}$	(3) + recovery	(4) + telework controls
$E_i$ (exposure)	2.67* (1.69)	2.89* (1.77)	2.50* (1.69)	1.60 (1.23)
$\mathcal{E}_i$ (supplier exposure)	-4.26 (0.37)	-4.48 (0.40)	4.87 (0.52)	35.18 (0.90)
$\Delta N_i^{\text{gap}}$ (2020 own collapse)	—	-0.05 (0.88)	0.06 (0.81)	0.03 (0.53)
$\Delta N_i^{\text{rec}}$ (2020–2022 own recovery)	—	—	-0.14** (2.13)	-0.15** (2.17)
$\Delta \mathcal{N}_i^{\text{rec}}$ (supplier recovery)	—	—	-1.58 (1.36)	-1.70 (1.36)
$T_i$ (must-be-in-person, Dingel–Neiman)	—	—	—	7.19*** (2.69)
$\mathcal{T}_i$ (supplier must-be-in-person)	—	—	—	-64.22 (0.89)
$R^2$	0.855	0.856	0.866	0.871
Observations			321	
Fixed effects		BEA-66 sector + chain-tier-12 cell		

*Notes:* Dependent variable is cumulative L-propagated price growth 2019–2022 in percentage points. Exposure measure: JEM-UK. Column (1) reports the bare reduced form; column (2) adds the 2020 own employment collapse  $\Delta N_i^{\text{gap}}$ ; column (3) adds own and supplier recovery margins. The negative coefficient on  $\Delta N_i^{\text{rec}}$  in column (3) is the supply-side fingerprint: industries that rebuilt employment faster saw *less*, not more, cumulative inflation – consistent with a labor-capacity (wall) mechanism rather than pent-up-demand release. Column (4) adds Dingel–Neiman teleworkability ( $T_i$ , must-be-in-person share aggregated to BEA-357 via OEWS, with the supplier-network propagation  $\mathcal{T}_i$ ) as exogenous controls; the recovery coefficient is essentially unchanged, ruling out the alternative interpretation that the supply-side fingerprint reflects faster recovery by teleworkable industries. Estimation is WLS with N-GO PCE-share weights; standard errors clustered at the BEA-66 sector level; absolute  $t$ -statistics in parentheses. Significance: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

holding the same fixed-effect structure, weights, and clustering as the headline regression. Columns (3)–(4) reproduce the headline post-COVID coefficient on the same sample as a memo.

The placebo coefficient is small and insignificant:  $\hat{\pi}_{\text{pre}}^E = +0.31$  ( $|t| = 0.23$ , bare;  $+0.41$  ( $|t| = 0.29$ ) adding pre-COVID employment dynamics as controls). The contrast with the post-COVID coefficient ( $\hat{\pi}_{\text{post}}^E = +4.33$ ,  $|t| = 2.39$ , significant at 5% in the same sample) is sharp:  $E_i$  has no detectable predictive content for pre-pandemic inflation but loads strongly on post-pandemic inflation. The headline cross-industry pattern is therefore a COVID-specific phenomenon, not a continuation of a pre-existing sectoral inflation trend.

## A.2.4 Services-exclusion robustness

A second natural concern is that the JEM index is concentrated among contact-intensive consumer-facing service industries (child day care, food and drinking places, hospitals, accommodation, personal care, etc.), and that the headline RF might be driven by a within-services common trend rather than a network-wide labor-capacity channel. Table 13 re-estimates the saturated column-(3) RF of Table ?? on three nested samples: the full baseline ( $N = 321$  endogenous BEA-357 industries), a *tight cut* dropping 15 high-finalness contact-intensive services, and

**Table 12:** Pre-COVID placebo for the cross-industry reduced form: predetermined JEM exposure does not predict cumulative 2017–2019 price growth.

	Placebo: $\Delta \log P_{17 \rightarrow 19}$		Memo: $\Delta \log P_{19 \rightarrow 22}$	
	(1) Exposure only	(2) + pre-emp controls	(3) Exposure only	(4) + pre-emp co
$E_i$ (exposure)	0.30 (0.23)	0.41 (0.29)	4.33** (2.39)	4.63*** (2.60)
$\mathcal{E}_i$ (supplier exposure)	4.06 (0.45)	4.00 (0.48)	-6.30 (0.59)	-6.31 (0.51)
$\Delta N_i^{\text{pre}}$ (2017–19 own emp. change)	—	-0.10 (0.85)	—	-0.18** (2.05)
$\Delta \mathcal{N}_i^{\text{pre}}$ (2017–19 supplier emp. change)	—	0.29 (0.28)	—	-0.16 (0.07)
$R^2$	0.520	0.534	0.857	0.863
Observations	321			
Fixed effects	BEA-66 sector + chain-tier-12 cell			

*Notes:* Pre-COVID placebo for the cross-industry reduced form (Table ??). The dependent variable in columns (1)–(2) is the cumulative L-propagated price growth over the 2017–2019 pre-COVID window; columns (3)–(4) report the analogous regression on the headline 2019–2022 window as a memo. Predetermined exposure  $E_i$  is the Danish country panel of the JEM (Oude Hengel et al., 2021). Pre-employment controls in columns (2) and (4) are cumulative 2017–2019 own employment change and its supplier-network propagation. The placebo coefficient on  $E_i$  is small and insignificant, confirming that predetermined JEM exposure captures a COVID-specific shock structure rather than a long-running sectoral inflation gradient that would contaminate the post-COVID identification. Estimation is WLS with N-GO PCE-share weights; standard errors clustered at the BEA-66 sector level; absolute  $t$ -statistics in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

a *broad cut* dropping 30 industries spanning NAICS 44–45 (retail), 61 (educational services), 62 (health and social assistance), 71 (arts/entertainment), and 72 (accommodation and food).

Both load-bearing coefficients survive. The exposure coefficient is  $\hat{\pi}^E = +4.63^{**}$  in the full sample,  $+3.61^*$  in the tight cut, and  $+3.52^{**}$  in the broad cut—attenuated modestly as the cuts remove high- $E_i$  tail industries, but stable in sign and economic magnitude. The recovery coefficient, which carries the Fact 2 supply-side fingerprint, actually *strengthens* as services are excluded:  $\hat{\pi}^r = -0.134^{**}$  full,  $-0.168^{**}$  tight cut,  $-0.191^{**}$  broad cut. Removing the high- $E_i$  services tail therefore sharpens rather than weakens the wall-vs-pent-up-demand sign discrimination. The headline cross-industry pattern is a network-wide phenomenon, not an artifact of a services-skewed sample.

## A.2.5 Legacy: 2SLS diagnostics

*The three subsections below document weak-instrument, fixed-effect, and services-exclusion diagnostics for a 2SLS specification that uses the JEM exposure index as an instrument for the 2020 employment collapse. The main text of the paper does not adopt this IV: as discussed in the footnote in Section 2, the exclusion restriction cannot be defended on empirical grounds for this application (in particular, pent-up demand released into capacity-constrained sectors in 2022 is a plausible direct channel from  $E_i$  to prices that does not run through the 2020 labor collapse). The diagnostics below are retained as background documentation of the IV behavior. They are not part of the paper’s identification.*

**Table 13:** Services-exclusion robustness for the saturated cross-industry reduced form (column (3) of Table ??).

Dependent variable: $\Delta \log P_{19 \rightarrow 22}$ (cumulative L-propagated price, pp)			
	(1) Full ( $N = 321$ )	(2) Tight cut ( $N = 306$ )	(3) Broad cut ( $N = 291$ )
$E_i$ (exposure)	4.63** (2.54)	3.61* (1.75)	3.52** (2.13)
$\mathcal{E}_i$ (supplier exposure)	0.76 (0.09)	2.29 (0.29)	5.36 (0.51)
$\Delta N_i^{\text{gap}}$ (2020 own collapse)	0.03 (0.42)	0.12 (1.40)	0.13 (1.11)
$\Delta N_i^{\text{rec}}$ (2020–2022 own recovery)	-0.13** (2.08)	-0.17** (2.39)	-0.19** (2.03)
$\Delta \mathcal{N}_i^{\text{rec}}$ (supplier recovery)	-1.39 (1.24)	-2.23* (1.85)	-2.72 (1.64)
$R^2$	0.868	0.859	0.776
Fixed effects	BEA-66 sector + chain-tier-12 cell		

*Notes:* Saturated reduced-form specification of column (3) of Table ??, re-estimated on three nested samples that progressively exclude consumer-facing service industries. Column (1) is the headline baseline; column (2) drops 15 high-finalness contact-intensive service industries (child day care, personal care, food and drinking places, hospitals, accommodation, nursing and residential care, outpatient/home/ambulatory health, and offices of physicians/dentists/health practitioners); column (3) drops 30 industries spanning NAICS 44–45 (retail), 61 (educational services), 62 (health and social assistance), 71 (arts/entertainment), and 72 (accommodation and food). Both the exposure coefficient (Fact 1 sign) and the recovery coefficient (Fact 2 sign test) survive both cuts. Estimation is WLS with N-GO PCE-share weights; standard errors clustered at the BEA-66 sector level; absolute  $t$ -statistics in parentheses. Significance: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

**Legacy: weak-instrument battery.** Table 14 reports the standard battery of weak-instrument diagnostics for the first stage of equation (??) that maps  $E_i$  into the 2020 own employment collapse  $\Delta N_i^{\text{gap}}$ . The cluster-robust first-stage  $F$  is 7.88, marginally below the Staiger–Stock rule of thumb of 10. The Montiel Olea and Pflueger (2013) effective  $F$  equals 7.88 in this just-identified setup. The Anderson and Rubin (1949) test rejects  $H_0 : \beta_g = 0$  at the 5% level, and the inverted 95% AR confidence set is [+0.141, +3.775]—bounded, excluding zero, and containing the 2SLS point estimate  $\hat{\beta}_g = +0.78$ . A pre-COVID placebo first stage on 2017–2019 employment growth is null.

**Legacy: IV robustness across fixed-effect configurations.** Table 15 reports first-stage  $F$  statistics and the IV price coefficient  $\hat{\beta}_g$  across five fixed-effect configurations of the legacy IV specification.

**Legacy: IV services-exclusion robustness.** Table 16 re-runs the IV price equation on the same full / tight cut / broad cut samples as Table 13 (which reports the corresponding reduced-form result the paper actually adopts).

**Table 14:** Weak-instrument robustness for the  $V_{\text{GAP}}$  IV and a pre-COVID placebo first stage.

<i>Panel A. Weak-instrument diagnostics (price LHS, <math>T = 2022</math>)</i>	
First-stage $F$ on $E_i$ (cluster-robust)	22.00
Olea–Pflueger effective $F$ (single-instrument)	22.00
Anderson–Rubin test of $H_0: \beta_g = 0$ ( $F$ )	8.34
$p$ -value	0.0056
Anderson–Rubin 95% confidence set for $\beta_g$	[+0.155, +1.641]
LIML estimate for $\beta_g$ (just-identified, = 2SLS)	0.605
Reduced-form coefficient on $E_i$ (price)	4.84***
$ t $	(2.89)
<i>Panel B. Pre-COVID placebo first stage (LHS: <math>\Delta \log \text{Emp}_{i,17 \rightarrow 19}</math>, <math>pp</math>)</i>	
$E_i$ (own JEM) on placebo LHS	1.50
$ t $	(0.71)
Placebo first-stage $F$ on $E_i$	0.50
Observations	321
Fixed effects	BEA-66 sector + chain-tier-12 cell

*Notes:* Panel A.  $V_{\text{GAP}}$  specification of Table 3, price LHS at  $T = 2022$ . The Olea–Pflueger effective  $F$  reduces to the squared  $t$ -statistic on the instrument in the just-identified case under cluster-robust inference. The Anderson–Rubin confidence set is constructed by inverting the AR test over a grid of  $\beta_g$  values; it is robust to arbitrarily weak identification. LIML coincides with 2SLS in just-identified models and is reported for transparency. Panel B. Regression of cumulative two-year log employment growth 2017–2019 on  $E_i$  with the same fixed effects as the main spec and  $\mathcal{E}_i$  as a control. A valid post-pandemic instrument should not predict pre-pandemic employment outcomes. WLS with  $w_i$  weights; BEA-66 cluster-robust standard errors. Significance: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

**Table 15:** Fixed-effects robustness of the  $V_{\text{GAP}}$  first and second stages, cumulative 2019–2022.

	(1) No FE	(2) BEA-66	(3) Chain-tier	(4) (2)+(3)	(5) (2)×(3)
<i>Panel A. First stage: <math>\Delta N_i^{\text{gap}}</math> on <math>E_i</math> and exogenous controls</i>					
$E_i$ (own JEM, instrument)	7.98** (2.49)	7.59*** (5.37)	10.69*** (3.23)	8.00*** (4.69)	9.49** (2.21)
First-stage $F$ on $E_i$	6.18	28.85	10.42	22.00	4.90
<i>Panel B. Second stage: <math>\Delta \log P_{19 \rightarrow 22}</math> on <math>\widehat{\Delta N_i^{\text{gap}}}</math> and exogenous controls</i>					
$\widehat{\Delta N_i^{\text{gap}}}$ (own collapse)	0.20 (0.72)	0.69*** (3.25)	0.53*** (3.16)	0.60** (2.31)	1.04* (1.95)
Observations			321		
Cells used (FE)	1	55	12	66	135

*Notes:*  $V_{\text{GAP}}$  specification of equations (??)–(??) re-estimated under five fixed-effects configurations, same ext\_12C sample ( $N = 321$ ) and same exogenous controls  $\mathcal{E}_i$ ,  $\Delta N_i^{\text{rec}}$ ,  $\Delta \mathcal{N}_i^{\text{rec}}$  throughout. Column (4) is the baseline reported in Tables ?? and 3. Column (5) interacts the BEA-66 sector and chain-tier-12 cell, an aggressive demeaning that absorbs nearly all cross-industry variance. WLS with  $w_i$  weights; BEA-66 cluster-robust standard errors; absolute  $t$ -statistics in parentheses. Significance: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

**Table 16:** Services-exclusion robustness for the  $V_{\text{GAP}}$  2SLS price equation, cumulative 2019–2022.

Dependent variable: $\Delta \log P_{19 \rightarrow 22}$ (cumulative L-propagated price, pp)			
	(1) Full ( $N = 321$ )	(2) Tight cut ( $N = 306$ )	(3) Broad cut ( $N = 291$ )
$\mathcal{E}_i$ (supplier JEM)	21.04 (1.29)	10.07 (0.95)	13.96 (1.28)
$\widehat{\Delta N}_i^{\text{gap}}$ (own collapse)	0.60** (2.31)	0.73* (1.72)	0.60** (2.30)
$\Delta N_i^{\text{rec}}$ (own recovery)	-0.33** (2.33)	-0.29** (2.41)	-0.28*** (2.77)
$\Delta \mathcal{N}_i^{\text{rec}}$ (supplier recovery)	-4.41** (2.20)	-4.77** (1.98)	-4.64** (2.38)
First-stage $F$ on $E_i$	22.00	8.72	11.46
Fixed effects	BEA-66 sector + chain-tier-12 cell		

*Notes:*  $V_{\text{GAP}}$  2SLS specification of equations (??)–(??) re-estimated on three progressively services-purged subsamples, all other features held fixed. Column (1) is the baseline sample of Table 3 ( $N = 321$ ). Column (2) drops 15 high-finalness contact-intensive consumer-facing service industries (child day care, personal care, food and drinking places, hospitals, accommodation, nursing/residential care, outpatient/home/ambulatory health, and offices of physicians/dentists/health practitioners). Column (3) drops 30 industries spanning NAICS 44–45 (retail), 61 (educational services), 62 (health and social assistance), 71 (arts/entertainment), and 72 (accommodation and food). WLS with  $w_i$  weights; BEA-66 cluster-robust standard errors; absolute  $t$ -statistics in parentheses. Significance: \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

## A.3 Chain Construction

This appendix documents the spectral-bisection algorithm that partitions the 343 endogenous BEA detail industries into the four supply chains used in the main text, and lists representative industries by tier within each chain.

### A.3.1 Spectral bisection on the cumulative supplier-path operator

Let  $A$  be the  $343 \times 343$  Leontief input-share matrix from the 2017 BEA use table, where  $A_{ij}$  is the share of industry  $i$ 's intermediate-input bill purchased from industry  $j$ . The cumulative supplier-path operator is

$$M = A + A^2 + A^3,$$

which aggregates direct, second-order, and third-order trade flows—the network depth at which the  $Z_3$  regression in Section 2 peaks. We construct the symmetrized weight matrix

$$\tilde{M} = \frac{1}{2}(M + M^\top)$$

and form the normalized graph Laplacian  $\mathcal{L} = I - D^{-1/2}\tilde{M}D^{-1/2}$ , with  $D$  the diagonal matrix of row sums. The Fiedler vector  $f$ —the eigenvector corresponding to the second-smallest eigenvalue of  $\mathcal{L}$ —encodes the bisection that minimizes trade flow across the cut.

Industries are sorted by their Fiedler-vector entries  $f_i$ , and the cut is chosen at the GO-balanced midpoint: the

threshold  $f^*$  for which  $\sum_{i: f_i \leq f^*} GO_i \approx \frac{1}{2} \sum_i GO_i$ . This produces two halves of approximately equal gross output (rather than equal industry count), so each side has comparable economic weight even when one half contains many small industries and the other a few large ones. The bisection is then applied recursively to each half, yielding four chains. The classification uses only pre-pandemic 2017 IO data and does not condition on any 2020–2024 outcome.

### A.3.2 Representative industries by tier

Within each chain, industries with a finalness ratio (share of output going directly to final demand) exceeding 0.70 form the consumer-facing tier ( $h$ ); the remainder are split by finalness into middle ( $m$ ) and upstream ( $l$ ) tiers. The four chains and representative industries are:

**$C_1$  (Retail & Telecom) (43 industries).** Retail trade, telecom, software, wholesale, transport. The chain is structurally distinct from the other three: its upstream tier consists of management consulting, advertising, and holding companies—pure organizational services with virtually zero direct PCE weight and no physical inputs that can be capacity-constrained. Retail trade dominates the consumer-facing tier ( $\kappa = 0.114$ , the single largest Domar weight in the economy). The chain’s price dynamics during COVID reflected strong demand rotation (stimulus-driven retail spending) rather than upstream cost pressure—as the model results in Section 5 confirm.

**$C_2$  (Consumer Services) (77 industries).** Hospitals, fast food, restaurants, universities, nursing, couriers. Couriers, cattle ranching, and rail transport sit at the upstream tier ( $l$ ); these supply trucking and grocery wholesale ( $m$ ), which in turn feed hospitals, fast-food restaurants, and universities ( $h$ ).

**$C_3$  (Professional & Hospitality) (101 industries).** Physicians, accommodation, air transport, pharma, building services. The chain traces a parallel path to  $C_2$ : architecture, engineering, and warehousing ( $l$ ) supply building services, food services, and drug wholesale ( $m$ ), which feed physicians, accommodation, and pharma manufacturing ( $h$ ). The physical infrastructure behind healthcare and hospitality—janitorial services, facility maintenance, warehousing—is the upstream bottleneck.

**$C_4$  (Manufacturing) (120 industries).** Auto manufacturing, construction, durable goods, machinery. The chain runs from construction and commercial machinery repair ( $l$ ) through durable goods wholesale and motor vehicle parts ( $m$ ) to auto manufacturing and consumer electronics ( $h$ ). Construction ( $\kappa = 0.017$ ) anchors the upstream tier, and persistent capacity constraints in the industrial base propagated cost increases to consumer-facing manufactured goods.

## A.4 Mapping the Model onto the Data

This appendix gives the explicit construction of the three model aggregates introduced in Section 3.5: aggregate inflation  $\Pi_{agg,t}$ , aggregate employment  $N_{agg,t}$ , and aggregate real value added  $Y_{agg,t}$ . The formulas below are exactly what the post-processing pipeline implements; under the calibrated baseline, each model aggregate is constructed from primitives in the same way as its data counterpart in Section 2, so the two paths are directly comparable.

### A.4.1 Aggregate price level (PCE deflator)

Aggregate inflation is a fixed-share Divisia index over the twelve endogenous tier composites and the three exogenous PCE blocks, with weights frozen at the 2017 base year. In gross form, the model object is

$$\Pi_{\text{agg},t} = \prod_{i=1}^{12} \left( \frac{P_{i,t}}{P_{i,t-1}} \right)^{w_i^{\text{PCE}}} \cdot \prod_{k \in \{\text{HO}, \text{FO}, \text{EN}\}} \Pi_{k,t}^{w_k^{\text{PCE}}}, \quad \sum_{i=1}^{12} w_i^{\text{PCE}} + \sum_k w_k^{\text{PCE}} = 1, \quad (\text{A.9})$$

and taking logs gives the additive form used in headline plots,

$$\log \Pi_{\text{agg},t} = \sum_{i=1}^{12} w_i^{\text{PCE}} \log \frac{P_{i,t}}{P_{i,t-1}} + \sum_{k \in \{\text{HO}, \text{FO}, \text{EN}\}} w_k^{\text{PCE}} \log \Pi_{k,t}. \quad (\text{A.10})$$

The endogenous tier price  $P_{i,t}$  is the within-tier CES aggregate of equation (8),

$$P_{i,t} = \left[ \sum_{j \in \{H,S\}} s_{ij} P_{ij,t}^{1-\rho} \right]^{1/(1-\rho)}, \quad (\text{A.11})$$

with within-tier hit/safe shares  $s_{ij}$  set at their 2019 values, and the exogenous blocks  $\Pi_{k,t}$  are model inputs introduced at the opening of Section 3.

The headline cumulative log price level (in percentage points relative to the 2019 base year) is

$$\Delta \log P_{\text{agg},t} = 100 \cdot \sum_{s=2020}^t \log \Pi_{\text{agg},s}. \quad (\text{A.12})$$

The 2017 PCE expenditure weights  $\{w_i^{\text{PCE}}, w_k^{\text{PCE}}\}$  are built bottom-up from the BEA-357 PCE bridge: industry-level personal-consumption-expenditure shares are aggregated to the chain-by-tier composites (giving  $\{w_i^{\text{PCE}}\}_{i=1}^{12}$  with  $\sum_i w_i^{\text{PCE}} \approx 0.71$ ) and to the three exogenous blocks (giving  $w_{\text{HO}}^{\text{PCE}} \approx 0.15$ ,  $w_{\text{FO}}^{\text{PCE}} \approx 0.13$ ,  $w_{\text{EN}}^{\text{PCE}} \approx 0.01$ ). By construction the four blocks sum to one, so the aggregator is a proper expenditure-weighted index.

### A.4.2 Aggregate employment

Aggregate employment is the unweighted level sum of sub-unit employment,

$$N_{\text{agg},t} = \sum_{i=1}^{12} \sum_{j \in \{H,S\}} N_{ij,t}, \quad (\text{A.13})$$

with no expenditure weighting: employment is a stock measure, not a consumption basket. The cumulative log change reported in headline plots is

$$\Delta \log N_{\text{agg},t} = 100 \cdot \log \left( \frac{N_{\text{agg},t}}{N_{\text{agg},2019}} \right). \quad (\text{A.14})$$

### A.4.3 Aggregate real value added

The headline model value added is constructed by I–O unbundling of tier-level gross-output growth using the fixed 2017 I–O matrix  $\Omega^{\text{DOM}}$  and the steady-state intermediate share  $1 - \delta_i$  per tier. The construction is identical to the data-side double-deflated value-added path in Section 2, so model and data aggregates are directly comparable whenever the per-tier  $Y_{i,t}$  match (which they do at the Newton-calibrated baseline).

Define tier-level gross output as the sum across hit and safe sub-units, and its annual log growth:

$$Y_{i,t} \equiv \sum_{j \in \{H,S\}} Y_{ij,t}, \quad g_{i,t}^Y = \log \frac{Y_{i,t}}{Y_{i,t-1}}. \quad (\text{A.15})$$

Tier-level intermediate-use growth is a fixed-share weighted aggregation of supplier-tier output growth via the I–O matrix,

$$g_{i,t}^{\text{MAT}} = \sum_{k=1}^{12} \omega_{i,k}^{\text{DOM}} g_{k,t}^Y, \quad (\text{A.16})$$

and per-tier real value-added growth is obtained by netting weighted-intermediate growth from gross-output growth and rescaling by the value-added share:

$$g_{i,t}^{\text{VA}} = \frac{g_{i,t}^Y - (1 - \delta_i) g_{i,t}^{\text{MAT}}}{\delta_i}. \quad (\text{A.17})$$

The aggregate growth rate is the value-added-share-weighted sum over tiers, where the tier value-added shares  $w_i^{\text{VA}}$  are computed at the 2019 steady state,

$$g_{\text{agg},t}^{\text{VA}} = \sum_{i=1}^{12} w_i^{\text{VA}} g_{i,t}^{\text{VA}}, \quad w_i^{\text{VA}} = \frac{\delta_i Y_i^{\text{ss}}}{\sum_{j=1}^{12} \delta_j Y_j^{\text{ss}}}, \quad (\text{A.18})$$

and the cumulative log value added reported in headline plots is

$$\Delta \log Y_{\text{agg},t} = 100 \cdot \sum_{s=2020}^t g_{\text{agg},s}^{\text{VA}}. \quad (\text{A.19})$$

This I–O-unbundling form is mathematically equivalent, under the model's  $P_k^{\text{ss}} = 1$  normalization, to the direct subtraction in equation (24),  $Y_{\text{agg},t} = \sum_{ij} [Y_{ij,t} - \sum_k m_{ij,k,t}^{\text{DOM}} - m_{ij,t}^{\text{IMP}}]$ . We use the I–O form because it is numerically more stable (avoiding small differences of large gross-output and intermediate-use levels) and because it matches the I–O accounting used by the BEA data-side construction tier-by-tier, which keeps model and data on a like-for-like footing.

The exogenous blocks HO, FO, and EN and imports are excluded from  $Y_{\text{agg}}$  because they have no modeled production structure: HO, FO, and EN are price-setting blocks that enter only the PCE aggregator, and imports are absorbed on the input side without contributing to value added. A chained-Fisher cross-check on the same per-tier quantity decomposition tracks (A.19) within  $\pm 0.2$  pp at every horizon, so relative-price effects on the headline value-added path are quantitatively immaterial.

## A.5 Data Detrending

This appendix documents the construction of the pre-COVID 2017–2019 trends used to detrend the calibration targets and the model’s reported series in the main text.

**Per-chain trend rates and Divisia weighting.** For each variable  $X \in \{Y, N, P_{GO}, W, WN, Y/N\}$ , a single trend rate is computed per chain  $c$  and applied to all three finalness tiers and within-tier hit/safe sub-units within that chain. The chain trend is a Divisia-weighted average of its three tiers’ 2017–2019 log growth rates,

$$g_{X,c}^{\text{trend}} = \sum_{u \in c} w_u^X \cdot g_{X,u}^{2017 \rightarrow 2019}, \quad g_{X,u}^{2017 \rightarrow 2019} = \frac{1}{2} \log(X_u(2019)/X_u(2017)),$$

with per-variable weights chosen to match each variable’s economic content:  $Y$  and  $Y/N$  are weighted by 2019 value-added shares,  $N$  and  $W$  (and  $WN$ ) by 2019 employment shares, and  $P_{GO}$  by 2019 PCE shares. Chain-level treatment ensures that the cross-chain composition of the trend reflects each variable’s natural aggregation, while preserving the within-chain ratios across tiers.

**Real value-added trend via IO unbundling.** The chain-level  $Y$  trend uses per-tier real *value-added* growth  $g_u^{VA}$ , not raw gross-output growth, in order to match the value-added concept of the model’s reported aggregate (Section 3.5). With the materials shares fixed at their 2017 values, real value-added growth is recovered by IO unbundling:

$$g_u^{VA} = \frac{g_u^{GO} - (1 - \delta_u) g_u^{MAT}}{\delta_u}, \quad g_u^{MAT} = \sum_j \omega_{u,j} \cdot g_j^{GO},$$

where  $\delta_u$  is the labor share in the production CES (so  $1 - \delta_u$  is the materials share) and  $\omega_{u,j}$  is the IO coefficient. The trend rate for  $Y$  at the chain level is then  $\sum_u w_u^{VA} \cdot g_u^{VA}$ , and for  $Y/N$  it is  $\sum_u w_u^{VA} \cdot (g_u^{VA} - g_u^N)$ . Both treat materials at fixed shares and net them out of gross output, so the trend subtracted from the data path uses the same value-added concept as the model’s reported aggregate.

**Aggregate employment anchor.** After per-chain trends are computed, a single uniform additive shift  $a_N$  is applied to all four chain  $N$  trends so that the aggregate detrended employment path equals zero at 2024:

$$a_N = \frac{1}{5} \log \frac{\sum_u N_u(2024) e^{-g_{N,u}^{\text{trend},5}}}{\sum_u N_u(2019)} \approx -0.73\%/yr.$$

The shift cascades to the chain  $Y$  trends (since  $Y = (Y/N) + N$  in log levels) and the  $WN$  trend. The  $Y/N$ ,  $P_{GO}$ , and  $W$  trends are unchanged by the anchor. Without this adjustment the model would generate a residual employment gap in 2025 and transition into a deflation episode, since the framework lacks the demand–interest-rate Taylor-rule interaction that would correct it via demand decline. Anchoring aggregate employment to its 2024 data level accommodates the post-COVID slowdown in labor-force participation that the stationary model does not capture endogenously.

**Exogenous-tier price paths.** The exogenous-block price shocks in  $\varepsilon^{HO}$ ,  $\varepsilon^{FO}$ ,  $\varepsilon^{EN}$  are detrended against the corresponding per-category 2017–2019 pretrends rather than against a uniform reference,

$$\varepsilon_e(t) = \log P_e(t) - \log P_e(t-1) - g_e^{\text{pre}}, \quad g_e^{\text{pre}} = \frac{1}{2} \log(P_e(2019)/P_e(2017)),$$

with  $g_{HO}^{\text{pre}} = +2.83\%/yr$ ,  $g_{FO}^{\text{pre}} = +2.81\%/yr$ ,  $g_{EN}^{\text{pre}} = +1.51\%/yr$ . The shocks are zeroed for  $t > 2024$  so that the exogenous price level remains flat at its 2024 value over the model’s remaining 45 horizon periods. Per-tier import shocks  $\varepsilon_{i,t}^{\text{IMP}}$  are constructed analogously.

**Resulting chain trend rates.** The chain-level post-anchor annual trend rates (%/yr) used by the Newton calibration are:

Chain	Y	N	W	Y/N	$P_{GO}$
$C_1$	+3.57	-0.06	+3.94	+3.64	+1.06
$C_2$	+0.55	+0.57	+2.70	-0.01	+2.15
$C_3$	+2.05	+2.00	+3.00	+0.05	+1.80
$C_4$	-0.15	+2.15	+2.93	-2.30	+1.25

The underlying raw chain  $N$  trends  $\{+0.67, +1.30, +2.73, +2.88\}\%/yr$  (pre-anchor) are uniformly shifted by  $a_N \approx -0.73\%/yr$ ; the chain  $Y$  trends are then the implied sum  $Y = (Y/N) + N$ . The wage trend is an external chain-level moment computed separately and not derived from the anchor.

## A.6 Inflation Expectations: Identification and Robustness

This appendix discusses the evidence supporting the calibration of the expectations operator in Section 4. We provide: (i) the first-difference regressions that anchor  $v_{\text{target}}$  to the Cleveland Fed market-based and University of Michigan household survey series; (ii) a 48-month rolling-window check showing no monotonic regime shift in the past-pass-through coefficient over 2005–2024 and confirming that the regression’s implied long-run anchor sits near the Fed’s 2% target; (iii) constrained level-form regressions and forecast-revision tests that show why the level-form data cannot cleanly identify a separate forward-looking weight; and (iv) a connection to the sticky-information framework of [Mankiw and Reis \(2002\)](#) and [Coibion and Gorodnichenko \(2015\)](#) together with the attenuation-bias argument that justifies reporting the CG-consistent specification as a robustness alternative.

### A.6.1 Specification and identification

The expectations operator in equation (21) of the main text is the convex combination

$$E_t[\pi_{t+12}] = v_{\text{forward}} \pi_{t+12} + v_{\text{past}} \pi_t^{\text{past}} + v_{\text{target}} \bar{\pi}, \quad v_{\text{forward}} + v_{\text{past}} + v_{\text{target}} = 1. \quad (\text{A.20})$$

$E_t[\pi_{t+12}]$  is the observed one-year-ahead expectation (Cleveland Fed model-implied or Michigan survey median),  $\pi_t^{\text{past}}$  is realized 12-month CPI inflation ending at  $t$ , and  $\bar{\pi}$  is the long-run anchor (the central bank’s 2% target).<sup>26</sup> Equation (A.20) is a deterministic decomposition of  $E_t[\pi_{t+12}]$ ; realized future inflation  $\pi_{t+12}$  differs from  $E_t[\pi_{t+12}]$  by a forecast surprise that does not enter the expectations operator itself but plays a role in the level-form identification (Section A.6.2).

<sup>26</sup>We use headline 12-month CPI for  $\pi_t^{\text{past}}$  because it is the salient inflation measure that households and market participants observe. Slope estimates are broadly similar using core CPI or 6-month annualized headline as the backward-looking regressor; the headline-CPI estimate is reported as the baseline.

Imposing the restriction  $v_{\text{forward}} = 0$  and differencing (A.20) over a 12-month window removes both the constant anchor term and any slow-moving covariate that is correlated with  $\pi_t^{\text{past}}$  but is not itself part of the structural pass-through:

$$\Delta E_t[\pi_{t+12}] = (1 - v_{\text{target}}) \Delta \pi_t^{\text{past}} + u_t. \quad (\text{A.21})$$

This is the first-difference regression (29) in the main text, with the slope  $b = 1 - v_{\text{target}} = v_{\text{past}}$ . We estimate (A.21) with Newey–West HAC standard errors at 12 lags to match the overlap horizon. The level-form regression  $E_t[\pi_{t+12}] = a + b \pi_t^{\text{past}} + e_t$  is contaminated by anchor drift, secular trends, and sample-mean differences—all slow-moving regressors that are correlated with  $\pi_t^{\text{past}}$  but absorb explanatory power from it, biasing the slope toward zero. The first-difference form differences these out and identifies the structural pass-through directly. On the same data, the level form gives  $b \approx 0.20$  and the first-difference form gives  $b \approx 0.25$ ; we adopt the first-difference estimate as our headline calibration. The model operates at annual frequency; we treat the monthly-estimated  $b$  as identifying the annual operator parameter  $v_{\text{past}}$  directly, since the convex combination (A.20) is a contemporaneous decomposition that has no time-aggregation discrepancy.

Table 17 reports specification (A.21) on monthly Cleveland Fed market-based expectations (Haubrich et al., 2012) and on the Michigan household survey, across pre-COVID, COVID, and full-sample windows. The Cleveland Fed pre-COVID slope ( $b = 0.295$ ,  $t = 6.95$ ) is statistically indistinguishable from the COVID-era slope ( $b = 0.246$ ,  $t = 8.34$ ). The Michigan survey gives a similar COVID-era slope ( $b = 0.298$ ,  $t = 9.49$ ) but cannot identify the pre-COVID slope ( $R^2 = 0.00$ ) because household responses contain insufficient signal during the small 2018 inflation episode. The Cleveland Fed series is our preferred series because it is market-based and less susceptible to survey-methodology shifts across time.

**Table 17:** Pass-through of past inflation to one-year-ahead expectations.

Series	Sample	$N$	$b$	$t$ (NW-HAC)	$R^2$
<b>Cleveland Fed</b>	<b>COVID + after (2020–2024)</b>	74	0.246	8.34	0.59
Cleveland Fed	Pre-COVID (2016–2019)	48	0.295	6.95	0.35
Cleveland Fed	Full (2003–2024)	290	0.212	9.86	0.40
Michigan	COVID + after (2020–2024)	74	0.298	9.49	0.44
Michigan	Pre-COVID (2016–2019)	48	0.019	0.28	0.00
Michigan	Full (2003–2024)	290	0.291	11.13	0.47

*Notes:* First-difference regression  $\Delta E_t[\pi_{t+12}] = \alpha + b \Delta \pi_t^{\text{past}} + u_t$  on monthly data.  $\Delta$  denotes a 12-month annual revision;  $\pi_t^{\text{past}}$  is realized 12-month CPI inflation ending at  $t$ . Newey–West HAC standard errors with 12 lags. Cleveland Fed: Haubrich–Pennacchi–Ritchken model-implied one-year-ahead inflation expectation, combining TIPS, nominal Treasuries, and survey data into a unified term-structure estimate. Michigan: University of Michigan Survey of Consumers one-year median expectation. Realized CPI from FRED CPIAUCSL.

The slope  $b$  identifies  $1 - v_{\text{target}}$  in the convex-combination operator (A.20). The Cleveland Fed COVID-era estimate  $b = 0.246$  implies  $v_{\text{target}} \approx 0.75$ , the calibration adopted in the main text.

The pre-COVID 2016–2019 window is not a flat low-inflation sample. Headline CPI rose from about 1.3% in mid-2017 to a peak near 2.9% in mid-2018 before falling back to 1.6% by mid-2019, and the Federal Reserve responded with eight 25-basis-point rate hikes between late 2016 and late 2018, raising the federal funds rate from 0.6% to 2.4%. Forecasters had a real, if small, inflation episode and a real monetary-policy response to react to. If the past-pass-through coefficient  $b$  had drifted between regimes, this window is the natural place to detect such a drift.

The Cleveland Fed comparison in Table 17 fails to reject the null of no drift: the COVID-era slope ( $b = 0.246$ ,  $t = 8.34$ ) is statistically indistinguishable from the pre-COVID slope ( $b = 0.295$ ,  $t = 6.95$ ). To verify that this stability is not an artifact of the chosen window boundaries, we estimate the level-form regression

$$E_t[\pi_{t+12}] = \alpha + \beta \pi_t^{\text{past}} + e_t \quad (\text{A.22})$$


on a rolling 16-quarter window that slides through 2005–2024 (Figure 8, top panel; companion expanding-window and random-walk-TVP overlays). This corroborates the assumption that the expectations regime is roughly stable, though we acknowledge that a rolling-window test alone is not definitive evidence.

Next, on each window we run (A.22) imposing the adaptive-anchor restriction  $v_{\text{forward}} = 0$  on (A.20). The convex-combination operator then reads  $E_t[\pi_{t+12}] = (1 - v_{\text{target}}) \pi_t^{\text{past}} + v_{\text{target}} \bar{\pi}$ , so the regression coefficients map directly to the calibration weights via

$$\beta = 1 - v_{\text{target}} = v_{\text{past}}, \quad \alpha = v_{\text{target}} \bar{\pi}, \quad \pi^*(t) \equiv \frac{\alpha(t)}{1 - \beta(t)} = \bar{\pi}. \quad (\text{A.23})$$

The intercept  $\alpha$  does *not* identify  $\bar{\pi}$  on its own: it is the product of the anchor weight  $v_{\text{target}}$  and the anchor level  $\bar{\pi}$ . Dividing  $\alpha$  by  $1 - \beta = v_{\text{target}}$  strips out the weight and recovers the anchor level itself. We refer to this as the *implied long-run anchor*  $\pi^*(t)$ . The first-difference form (A.21) differences out  $\alpha$  and so identifies  $\beta$  alone; running the level form on a moving window lets us recover both. The top panel of Figure 8 reports  $\beta(t)$ ; the bottom panel reports  $\pi^*(t)$ . Each is estimated on three time-varying schemes (16-quarter rolling, expanding from 2016Q1, random-walk TVP) to confirm robustness to the window-construction choice. The top panel asks whether the adaptive pass-through is stable; the bottom panel asks whether the implied anchor level—the level forecasters appear to be anchoring at, as recovered from the regression—is consistent with the central bank’s stated target. A stable  $\beta$  means the calibration  $v_{\text{past}} = 1 - v_{\text{target}} = 0.25$  is not a regime-specific artifact; a  $\pi^*$  that hovers near the Fed’s 2% target means the regression’s intercept is consistent with  $\bar{\pi} = 2\%$  in the calibration, not just an empirical constant absorbed into  $\alpha$ .

The estimated slopes oscillate within  $[0.10, 0.30]$  with no clear drift; a brief excursion to  $\beta \approx 0.47$  at the 2020-Q4 anchor reflects the rolling window catching the discrete level shift of CPI between late 2020 and early 2021 with only a handful of overlapping observations, which inflates the fitted slope. The expanding-window and TVP estimates—which do not mechanically truncate the sample—show no such excursion, confirming the spike as an edge effect rather than a regime shift, and by 2021-Q4 the rolling estimate settles back inside the pre-COVID range. The implied long-run anchor  $\pi^*(t)$  stabilizes near 1.87% across all three estimation schemes—independent empirical support for setting  $\bar{\pi} = 2\%$  in the calibration.



**Figure 8: Time-varying inflation-expectation rule.** Adaptive level-form regression  $E_t[\pi_{t+12}] = \alpha + \beta \pi_t^{\text{past}} + e_t$  estimated on three time-varying schemes (16-quarter rolling, expanding from 2016Q1, random-walk TVP). *Top panel:* pass-through slope  $\beta(t) = 1 - \nu_{\text{target}}$ . *Bottom panel:* implied long-run anchor  $\pi^*(t) \equiv \alpha(t)/(1 - \beta(t))$ , recovered as in equation (A.23). Cleveland Fed market-based one-year expectations, quarterly, 2016Q1–2024Q4. Shaded band: 2020Q1–Q2 (COVID onset). Horizontal dotted lines: full-sample  $\beta = 0.24$  (top) and  $\pi^* = 1.87\%$  (bottom).

## A.6.2 Constrained regression and forecast revisions: identification of $\nu_{\text{forward}}$

The first-difference regression (A.21) imposes  $\nu_{\text{forward}} = 0$ . To check whether this restriction is empirically binding here—that is, whether the level-form data themselves support a non-zero forward-looking weight—we re-run the level form with  $\nu_{\text{forward}}$  pinned at two values: 0 (the adaptive-anchored baseline) and 0.25 (the value used in the CG-consistent specification). We use realized  $\pi_{t+12}$  as the proxy for the model-consistent forecast component (with the standard caveat that this introduces forecast-surprise noise) and let the data identify the residual coefficient  $\nu_{\text{past}}$  (with  $\nu_{\text{target}}$  pinned down by the adding-up constraint).

Table 18 reports the result. With  $\nu_{\text{forward}}$  pinned at zero, the level-form regression recovers  $\hat{\nu}_{\text{past}} \approx 0.17$  in the COVID-era windows and the fitted expectations correlate strongly with the actual Cleveland Fed series (+0.47 to +0.86). With  $\nu_{\text{forward}}$  pinned at 0.25, the residual loading on past inflation collapses (the regression cannot find a coherent  $\hat{\nu}_{\text{past}}$ ) and the fitted expectations become *negatively* correlated with the survey series in three of five windows. The level-form data thus prefer the corner with no forward-looking weight.<sup>27</sup>

The constrained level-form estimates of  $\nu_{\text{past}}$  here ( $\sim 0.17$ ) are lower than the first-difference estimate ( $b \approx 0.25$ ) reported in the main text. This is the expected direction of the clear source of a bias: level-form regressions suffer from anchor-drift and trend contamination that absorbs explanatory power from  $\pi^{\text{past}}$ , which is why we adopt

<sup>27</sup>A complementary forecast-revision regression that regresses expectations on realized future inflation as a proxy for the rational forecast fails to reject  $\nu_{\text{forward}} = 0$ , but the test is uninformative in the relevant direction: forecast-surprise variance in  $\pi_{t+12}$  creates classical errors-in-variables that attenuates the coefficient toward zero, so a small estimate cannot distinguish “no forward-looking weight” from “forward-looking weight masked by attenuation.” See Section A.6.3.

**Table 18:** Constrained projection of Cleveland Fed inflation expectations

Sample	N	Spec A: $v_{\text{forward}} = 0$		Spec B: $v_{\text{forward}} = 0.25$	
		$\hat{v}_{\text{past}}(t)$	corr	$\hat{v}_{\text{past}}(t)$	corr
Pre-COVID (2016Q1-2019Q4)	20	+0.248 (+7.16)	+0.731	+0.183 (+3.20)	+0.520
COVID surge (2020Q1-2022Q4)	12	+0.148 (+6.56)	+0.861	+0.020 (+0.27)	-0.023
COVID + after (2020Q1-2024Q4)	21	+0.167 (+6.25)	+0.627	+0.038 (+0.56)	-0.202
Post-COVID (2021Q1-2024Q4)	17	+0.158 (+6.25)	+0.471	+0.016 (+0.22)	-0.601
Full sample (2016Q1-2024Q4)	41	+0.173 (+6.40)	+0.705	+0.048 (+0.76)	+0.182

Notes: Quarterly.  $\pi_t^{\text{past}}$  and  $\pi_{t+12}$  are 12-month CPI inflation rates ending at the quarter end. Constrained projection  $E_t[\pi_{t+12}] - v_{\text{forward}} \pi_{t+12} = v_{\text{past}} \pi_t^{\text{past}} + (1 - v_{\text{forward}} - v_{\text{past}}) \bar{\pi} + \varepsilon$  with  $\bar{\pi} = 2\%$ . Spec A fixes  $v_{\text{forward}} = 0$ ; Spec B fixes  $v_{\text{forward}} = 0.25$ . NW-HAC standard errors with 4 lags. “corr” is the Pearson correlation between fitted and actual series.

the first-difference estimate as our baseline. The constrained level tests whether the level data prefer  $v_{\text{forward}} = 0$  over  $v_{\text{forward}} = 0.25$ .

### A.6.3 CG variant and sticky-information anchoring

The constrained-regression evidence in Section A.6.2 prefers the corner  $v_{\text{forward}} = 0$ , but a complementary forecast-revision test that would more directly identify a small positive  $v_{\text{forward}}$  is attenuation-biased toward zero by forecast-surprise variance in  $\pi_{t+12}$ . It is therefore no test at all. Therefore, a theory-informed baseline is considered as a variant (CG).

This variant connects the two-weight rule to the sticky-information and information-rigidity literature, notably Mankiw and Reis (2002) framework and the findings in Coibion and Gorodnichenko (2015). In the sticky-information environment of Mankiw and Reis (2002), only a fraction  $1 - \lambda$  of agents update their information set in any given period, and the average expectation is a geometric mixture of forecasts conditioned on progressively older information sets. If inflation follows an AR(1) process around a long-run mean, each stale forecast collapses to a convex combination of recent inflation and that mean; integrating over the stale-forecast geometry yields a reduced form in which the average expectation loads on the model-consistent forecast (weight  $1 - \lambda$ ), recent inflation (weight reflecting AR(1) persistence times  $\lambda$ ), and the long-run anchor (the residual weight). Equation (A.20) adopts this three-component decomposition as the calibrated expectations operator. The mapping is stylized and serves the purpose of motivating the robustness exercise rather than micro-founding the calibration from first principles. Coibion and Gorodnichenko (2015) estimates suggest a forward-looking weight in the neighborhood of 0.20–0.40. We use 0.25 as a conservative choice in our CG-consistent specification.

Table 19 reports the level- and first-difference correlations between each specification’s fitted path and observed expectations across multiple sample windows. The adaptive-anchored baseline delivers the strongest correlation with observed expectations in every window and on both moments. Our CG-consistent specification gives a weaker fit. We include a third “forward-only” specification not as a candidate calibration but as a diagnostic of the timing mismatch that arises when too much weight is placed on realized future inflation: during the 2021–2022 surge, expectations rose ahead of the 12-month realized inflation that  $\pi_{t+12}$  would later record, so a fitted path that weights  $\pi_{t+12}$  heavily lags the observed expectations series and produces negative correlations in the post-COVID windows. The forward-only specification contradicts the robust finding  $b \approx 0.25$  by setting  $v_{\text{past}} = 0$ .

**Table 19:** Head-to-head comparison of expectations specifications

Spec	Level-form correlation with observed expectations				
	CF post-COVID	CF COVID+	CF full	TIPS COVID+	TIPS post-GFC
<b>Adaptive (baseline)</b> (0, 0.25, 0.75)	+0.47	+0.63	+0.71	+0.82	+0.79
Forward-only (0.25, 0, 0.75)	-0.64	-0.32	+0.05	+0.02	+0.23
CG-consistent (0.25, 0.25, 0.50)	-0.11	+0.26	+0.47	+0.62	+0.64
First-difference correlation (dynamics)					
<b>Adaptive (baseline)</b> (0, 0.25, 0.75)	+0.64	+0.33	+0.36	+0.65	+0.64
Forward-only (0.25, 0, 0.75)	-0.46	+0.19	-0.05	-0.16	-0.57
CG-consistent (0.25, 0.25, 0.50)	+0.35	+0.42	+0.25	+0.41	+0.17

Notes: Each specification generates a fitted path  $\hat{E}_t[\pi] = v_f \pi_{t+12} + v_p \pi_{t-1} + v_t \bar{\pi}$  using realized CPI inflation. Correlations computed against Cleveland Fed (CF) model-implied one-year expectations and TIPS five-year break-even inflation. “COVID+” = 2020–2024; “post-COVID” = 2021–2024; “post-GFC” = 2010–2024.

## A.7 Robustness and Sensitivity

This appendix supports Section 5. The wall remains the dominant endogenous source of inflation during the 2021–2022 outbreak under each exercise.

Experiment	Wall Act I (pp)	Wall Act II (pp)
Baseline	+4.31	+4.08
Low cost anchor ( $c = 0.14$ )	+3.32	+3.29
All elasticities at 0.5 ( $\rho_b = \rho_m = \sigma_w = 0.5$ )	+3.68	+4.44
High hit/safe substitution ( $\sigma_w = 5$ )	+2.44	+4.92
Lower downward rigidity ( $\lambda = 1.5\bar{\phi}$ )	+4.28	+4.34

Notes: Each row is the headline calibration with one parameter perturbed, all others at baseline. Wall Act I = cumulative log  $\Pi_{\text{agg}}$  wall contribution 2019–2021 (pp). Wall Act II = cumulative log  $\Pi_{\text{agg}}$  wall contribution 2021–2023 (pp). Baseline is the current calibration ( $c = 0.17$ ,  $\rho_m = \rho_b = \sigma_w = 0.39$ , per-chain  $\lambda = 3\bar{\phi}$ ).

**Table 20:** Wall-channel robustness (adaptive expectations spec).

Table 20 reports the cumulative wall contribution by Act under the adaptive-anchored expectations regime. Reducing the per-net-hire rebuilding cost  $c$  scales down both Act I and Act II wall contributions roughly proportionally, as expected. Halving the downward-rigidity parameter  $\lambda$  leaves both Acts essentially unchanged, indicating that the wall mechanism does not require the upper end of the structural-asymmetry range to operate.

The substitution elasticities affect the two Acts asymmetrically, mirroring the Act I / Act II decomposition in Section 5. Raising the hit/safe substitution  $\sigma_w$  to 5 nearly halves the Act I wall contribution while leaving Act II broadly intact: this is consistent with Act I being dominated by the asymmetric premium across tiers — which hit/safe substitutability erodes — and Act II being dominated by the even-wall floor, which does not depend on within-tier substitutability. Lowering all three CES elasticities jointly to 0.5 moves the contributions modestly in both Acts. The qualitative reading of the headline decomposition is robust across this parameter range.

**Decomposition under forward-looking (CG-consistent) expectations.** Table 21 repeats the main-text wall decomposition (Table 7) under the CG-consistent expectations regime. Act II is qualitatively similar to the adaptive headline: shock channel and capacity wall split roughly 43%/57%, with the even-wall component accounting for 63% — close to the 36%/64%/65% split under adaptive expectations. Act I differs sharply. Under CG-consistent expectations the wall channel alone over-explains cumulative inflation by 2021: the share computes to 112% before capping, meaning that, in this regime, the exogenous price blocks and the residual markup wedge  $\mu$  collectively *suppress* cumulative Act I inflation by about 12% rather than contribute to it. The reading is that the stronger forward-looking component of the CG specification amplifies the wall’s near-term price impulse so that residual wedges play an offsetting role at the early horizon. The MPC= 0.4 income-propagation column shifts the Act I shock share toward zero (+9%) but does not alter the qualitative reading.

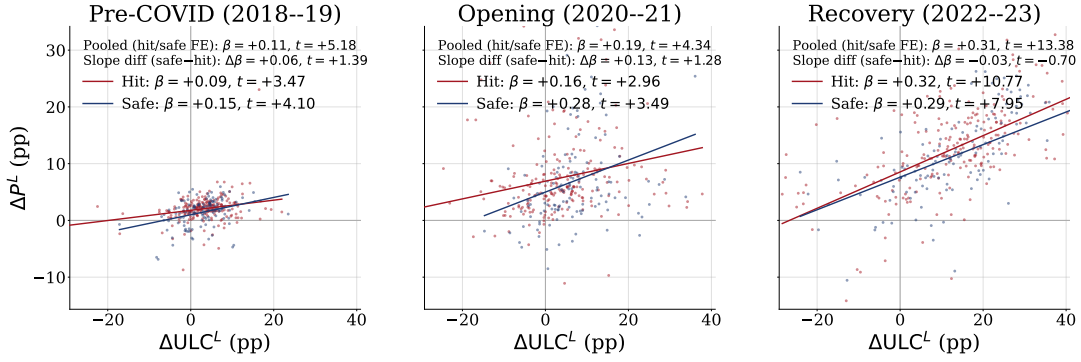
**Table 21:** Capacity-wall contribution to inflation by Act, cg-anchored spec. Each share is computed relative to the structural floor (Model – no  $\mu$  – no exo – no wall); the denominator is the total above-floor cumulative inflation. The right-hand panel adds an income-propagation channel (MPC = 0.4) and shows the wall-decomposition story is essentially invariant to it.

Series (cum log $\Pi_{\text{agg}}$ , pp)	Baseline		With income propagation (MPC = 0.4)	
	Act I (2019–2021)	Act II (2021–2023)	Act I (2019–2021)	Act II (2021–2023)
Model	+1.52	+6.62	+1.52	+6.62
Model – no $\mu$ – no exo	+2.17	+3.72	+0.91	+3.91
– even wall (uniform $\gamma$ )	–1.94	+4.09	–3.74	+4.84
– no wall ( $\gamma = 0$ , structural floor)	–3.94	–0.15	–5.64	+0.19
<i>Contributions to above-floor inflation by Act:</i>				
Exogenous shocks ( $\mu$ +exo)	–12%	43%	9%	42%
Capacity wall, total	112%	57%	91%	58%
of which: even-wall component	37%	63%	27%	72%

*Notes:* “Model” is the calibrated headline path; “– no  $\mu$ ” zeroes the residual markup wedge; “– no exo” zeroes the exogenous price blocks (housing, food, energy). “Even wall” replaces the per-tier capacity-wall premium with a uniform aggregate floor that preserves the size of the 2020 contraction but removes its cross-tier asymmetry. “No wall” sets  $\gamma = 0$  in all tiers; the cumulative inflation path below the structural floor is the deflationary baseline against which all other shares are computed. Contribution rows: exogenous-shocks share is (Model – drop  $\mu$ +exo)/(Model – no wall); total capacity-wall share is (drop  $\mu$ +exo – no wall)/(Model – no wall); the even-wall component is the further sub-decomposition isolating what survives when the wall is forced to be uniform across tiers. The right-hand panel adds an income-propagation channel that augments tier-level demand shifters by  $\text{MPC} \cdot (\log Y_{\text{agg}} - \log Y_{\text{agg}}^{\text{data}})$  with MPC = 0.4, capturing the fact that real income deviations from the baseline path feed back into total demand at every tier.

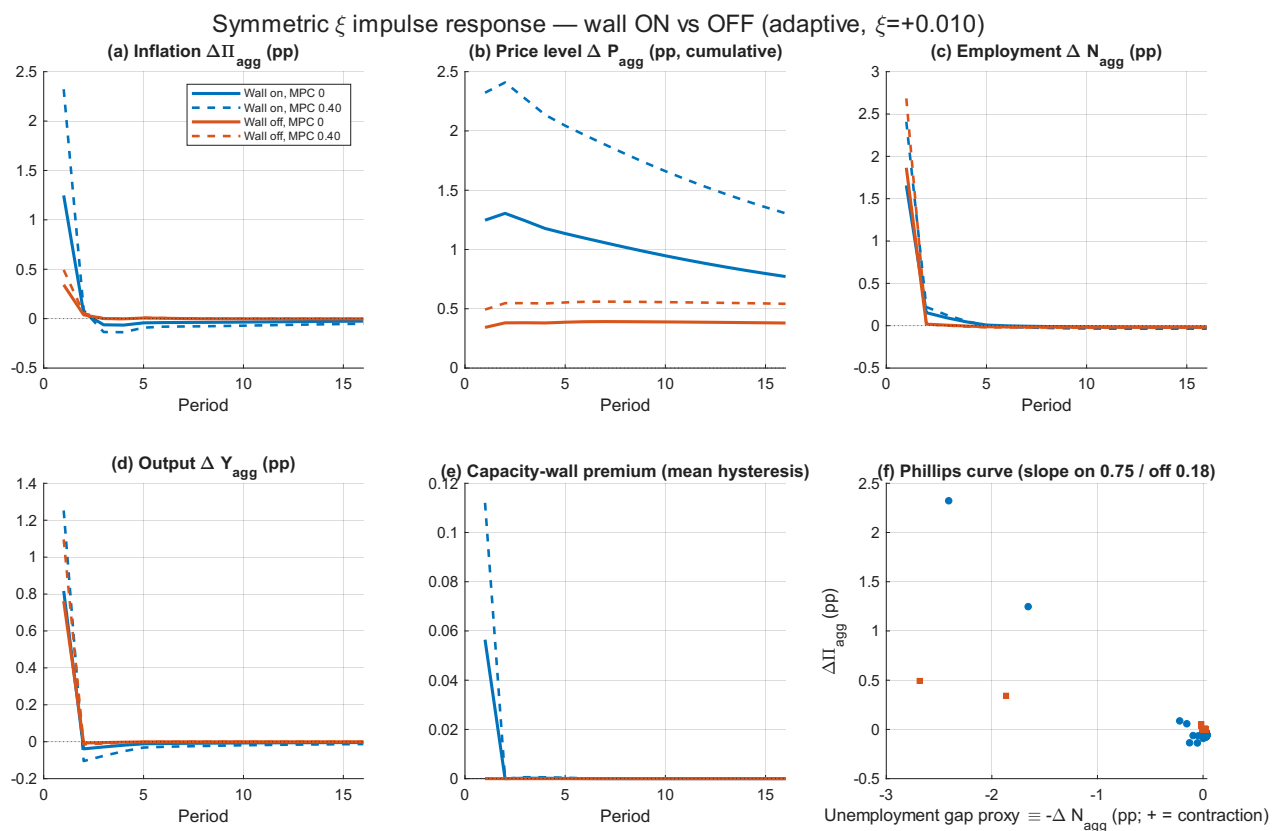
**Cost-price pass-through steepening: scatter and fitted lines.** Figure 9 plots the scatter and fitted lines for the cost-price pass-through component of the Phillips relation cited in Section 5. The pooled slope steepens from  $\beta_{\text{pre}} \approx +0.11$  to  $\beta_{\text{Act II}} \approx +0.31$  across the three windows. Within each period, the hit-vs-safe slope difference is statistically indistinguishable from zero. The model reproduces this pattern, as discussed in the paper.

**Underlying impulse responses behind the main-text Phillips curve.** Figures 10 and 11 display the underlying perfect-foresight trajectories that feed the 1% expansionary point on the joint Phillips schedule in

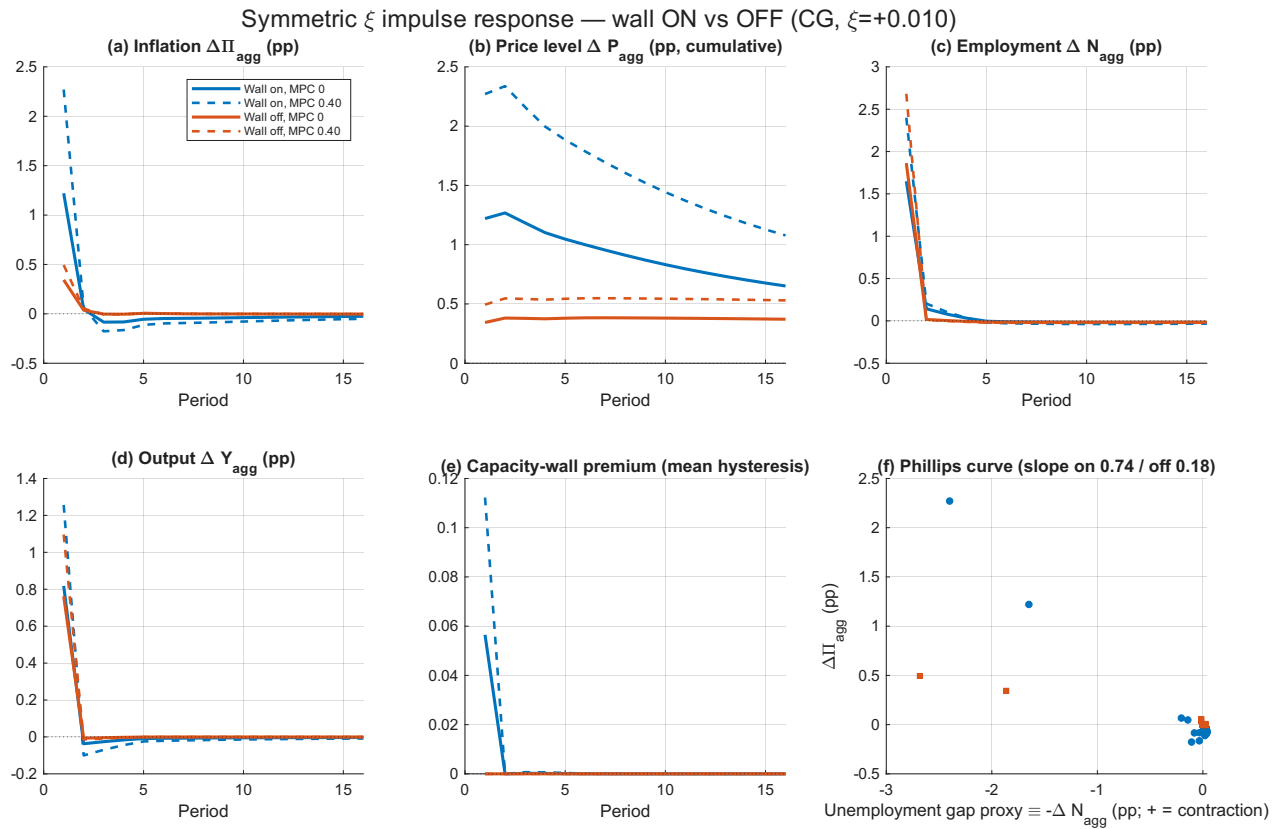


**Figure 9: Cost-price pass-through steepening in the data:  $\Delta \log P^L$  on  $\Delta \log ULC^L$  by period and hit/safe.** BEA-357 detail panel. Left to right: pre-COVID (2018–2019 pooled), reopening (2020–2021 pooled), and recovery (2022–2023 pooled). Blue: hit industries; red: safe industries; solid lines: pooled WLS fits. This is the cost-price pass-through component of the Phillips relation, not the full Phillips curve.

main-text Figure 6. Each is a six-panel summary of the model’s response to a uniform, one-period, +1% iid  $\xi$  demand pulse applied to all 12 tier composites simultaneously, plotted over a 16-period horizon and starting from the model’s steady state. Every panel overlays four arms: wall on vs. wall off (calibrated  $\gamma_i$  vs.  $\gamma_i = 0$ ) crossed with the MPC income-propagation channel off (MPC= 0) and on (MPC= 0.40). Panel (a) is aggregate inflation  $\Delta \Pi_{agg}$  period by period; (b) is the cumulative price level  $\Delta P_{agg}$ ; (c) is aggregate employment  $\Delta N_{agg}$  (the size of  $\Delta N(1)$  relative to the per-tier threshold determines how hard the wall fires); (d) is aggregate output  $\Delta Y_{agg}$ ; (e) is the mean capacity-wall premium across the 24 sub-units, directly visualising  $\mathcal{W}_{u,t} = \gamma_i (dN_{u,t}^+)^2$ ; and (f) is the Phillips-curve scatter of  $\Delta \Pi_{agg}$  against the unemployment-gap proxy  $-\Delta N_{agg}$  across all simulated periods, with on/off fitted slopes whose difference is the wall-driven endogenous steepening at the 1% benchmark. The two figures differ only in the NKPC expectations triple: Figure 10 uses the adaptive specification  $(\nu^{rat}, \nu^{past}, \nu^{target}) = (0, 0.25, 0.75)$ , and Figure 11 uses the CG specification  $(0.25, 0.25, 0.50)$ .



**Figure 10: Impulse responses to a +1% iid  $\xi$  demand pulse, adaptive spec.** Six panels report inflation, the cumulative price level, employment, output, the mean capacity-wall premium, and the Phillips-curve cloud across periods, overlaying four arms (wall on/off  $\times$  MPC  $\in \{0, 0.4\}$ ). The pulse is applied uniformly to all 12 tier composites in period 1 and solved by perfect foresight from steady state. Blue: wall on (calibrated  $\gamma$ ); orange: wall off ( $\gamma = 0$ ). Solid: MPC= 0; dashed: MPC= 0.40.



**Figure 11: Impulse responses to a +1% iid  $\xi$  demand pulse, CG spec.** Construction identical to Figure 10 with the CG expectations triple  $(\nu^{rat}, \nu^{past}, \nu^{target}) = (0.25, 0.25, 0.50)$ .

## A.8 Retail Sector Profit Margins in Other Datasets

This appendix shows that retail margin expansion can be similarly seen in the Census Bureau’s Quarterly Financial Statistics (QFS, formerly QFR) and BEA/BLS KLEMS integrated accounts.

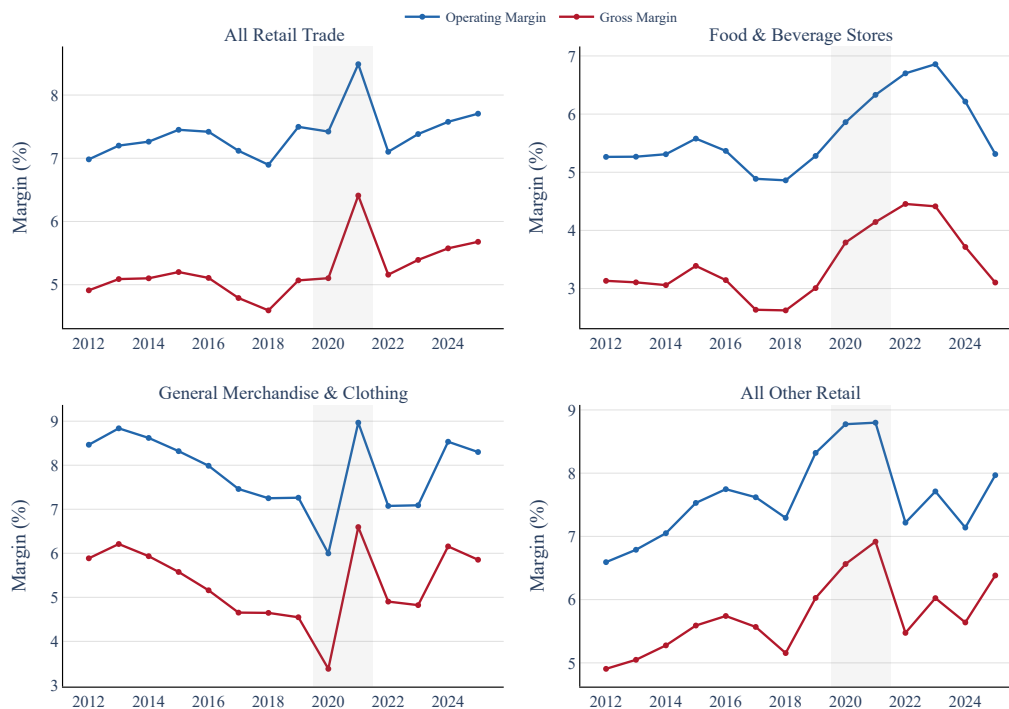
We distinguish two margin concepts here. The *operating margin* is the ratio of net sales to operating costs (excluding depreciation and interest), expressed as a percentage markup:  $OM = 100 \times (S/C_{op} - 1)$ . This is the narrowest measure, closest to the firm’s pricing decision over variable costs. The *gross margin* additionally includes depreciation in the cost base:  $GM = 100 \times (S/(C_{op} + D) - 1)$ . Both measures are computed from annual aggregates (sum of four quarters) to smooth seasonal variation.

Figure 12 shows operating and gross margins for retail trade and its subcategories. All Retail Trade experienced a sharp margin expansion beginning in 2020, with the operating margin rising from approximately 7% to over 10%—a level not seen in the pre-pandemic sample. The expansion is broad-based: Food and Beverage Stores, General Merchandise and Clothing, and All Other Retail all show the same pattern, though with different magnitudes. Food and Beverage margins rose the most in percentage terms, consistent with the well-documented divergence between food-at-home prices and upstream agricultural costs during the pandemic. The margin expansion persists through 2023, with only partial normalization by 2024.

Comparing retail margins across datasets requires care because “retail output” is defined differently depending on the source. In the BEA input–output tables, retail output is the *trade margin*—the difference between the selling price and the purchase price of goods resold—so intermediate goods purchased for resale are netted out. Retail output in the IO tables is thus the value of the retail *service*, not total sales. The BEA/BLS KLEMS accounts follow the same national accounts convention: retail gross output equals the trade margin, and factor shares (labor, capital, intermediate inputs) are computed over this margin-based output concept. By contrast, firm-level data from Compustat or the Census QFS report standard income-statement items: total revenue includes all goods sold, and cost of goods sold (COGS) includes the purchase price of merchandise. The *gross margin* (revenue minus COGS) for a retailer is therefore close to the trade margin concept, but it excludes most labor costs—in retail, the vast majority of labor expenditure (store personnel, logistics, management) appears in selling, general, and administrative expenses (SG&A), not in COGS.

Figure 14 plots the KLEMS retail profit share for BEA Summary code 4A0 (“Other retail”) from 1997 to 2023, with period averages summarized in Table 22. The dataset provides an independent level-based measure of retail margins that triangulates the markup wedge calibrated from our Leontief decomposition. We report the *gross-output profit share* (gross operating surplus divided by gross output, GOS/GO) rather than the more familiar value-added-based profit share (GOS/VA), so that the denominator matches our main-text framework’s normalization by gross output. By the national-accounts identity  $GO = VA + \text{intermediates}$ , GOS/GO is mechanically smaller in level than GOS/VA but isolates the share of each dollar of gross output that becomes profit rather than compensation or purchased inputs. For the retail sector covered by KLEMS, the GO-based profit share averaged 21.2% over the pre-COVID window (2017–2019), rising to 24.2% in Act I (2020–2021; +3.0 pp) and 24.6% in Act II (2022–2023; +3.4 pp). The increase came primarily at the expense of the compensation share of gross output (which fell by roughly 5 pp over 2017–2023); the intermediate-input share rose temporarily during 2021–2022 and had largely reversed by 2023. Against a pre-COVID secular trend of about +0.23 pp per year (1997–2019), the post-COVID widening is roughly six to seven times the trend rate—a structural break rather than a continuation.

KLEMS cannot be compared to the network-propagated markup contribution that the paper reports in Section 5, because the KLEMS accounts have no input–output structure. The natural comparison is instead with the paper’s *single-pass* markup wedge—the own-industry residual before propagation through the Leontief inverse. Using the level-to-log-growth translation  $d \ln M \approx d(\text{GOS/GO})/(1 - \text{GOS/GO})$ , the KLEMS Act I shift of +3 pp in the profit share corresponds to roughly +4 log-% of markup-induced price wedge at the retail sector level. The single-pass




**Figure 12: Retail trade operating margins by subcategory (Census QFS).** Operating margin =  $100 \times (S/C_{op} - 1)$ ; gross margin =  $100 \times (S/(C_{op} + D) - 1)$ . Annual (sum of quarters). Gray shading: COVID period (2020–2021).

**Table 22: Retail profit margin in BEA–BLS KLEMS (industry 4A0 “Other retail”), pre-COVID vs Acts I and II**


Period	Margin <sub>GO</sub> (%)	Δ vs pre-COVID (pp)	Comp/GO (%)	Intermediate/GO (%)
Pre-COVID (2017–19)	21.2	—	29.3	37.8
Act I (2020–21)	24.2	+3.02	26.6	40.0
Act II (2022–23)	24.6	+3.40	24.3	40.3

Notes: Margin<sub>GO</sub> = Gross Operating Surplus / Gross Output, where GO = Value Added + Intermediate Inputs. Period averages over the three KLEMS years in each window. GO-based, not value-added-based, for alignment with the paper’s Leontief framework. Share decomposition obeys the identity GOS/GO + Comp/GO + Tax/GO + Intermediate/GO ≡ 1; tax share not shown. Comparable to the paper’s single-pass (own-industry) markup residual  $\mu_i$ , not the L-propagated geo-pass, since KLEMS has no input–output structure. KLEMS 4A0 covers Detail codes 444, 446, 447, 448, 454 (building materials, health and personal care, gasoline stations, clothing, non-store retailers); BEA summaries 441 (motor vehicles), 445 (food and beverage), 452 (general merchandise) are not in this merged file. 2024 data unavailable (KLEMS release lag). Source: BEA–BLS KLEMS Industry-Level Production Account, 1997–2023.

$\mu$  we compute for the RTRADE aggregate over the same window is approximately +15 log-% cumulative. The two measures agree on direction and timing but differ in magnitude by roughly a factor of three, consistent with three considerations: (i) KLEMS 4A0 covers only a subset of retail (building materials, health and personal care, gasoline, clothing, and nonstore retailers; motor vehicles [441], food and beverage stores [445], and general merchandise [452] are not in the merged file), whereas our RTRADE is the broader 44–45 aggregate and captures precisely the subsectors most exposed to pandemic supply shocks; (ii) KLEMS 4A0 includes gasoline stations (447), which our classification treats separately as part of the exogenous HO block; and (iii) our residual absorbs



**Figure 13: Non-retail operating margins (Census QFS, “All Industry” aggregate).** The aggregate excludes retail trade, petroleum and coal products, and primary metals—sectors whose margins are dominated by commodity price swings. Non-retail margins remain within their historical range through the pandemic and recovery, in contrast to the retail expansion documented in Figure 12.



**Figure 14: Retail profit share in BEA-BLS KLEMS (industry 4A0 “Other retail”).** Left: annual gross-output profit share (GOS/GO) with 1997–2019 OLS trend ( $\approx +0.23$  pp/yr); shaded windows mark Pre-COVID (2017–2019), Act I (2020–2021), and Act II (2022–2023). Right: decomposition of GO into GOS, compensation, and intermediate inputs as annual shares. Source: BEA-BLS KLEMS industry-level production account.

productivity-measurement noise and accrual-accounting timing differences that the KLEMS identities do not. KLEMS therefore corroborates the direction and two-act pattern of the retail markup widening while placing a conservative lower bound on its magnitude, consistent with the limits of any single data source.

## A.9 Data Sources

This appendix documents the primary data sources, derived inputs, and external series used across the empirical and structural components of the paper.

1. **BEA Supply-Use Tables (2017 benchmark)**. Detail-level supply-use tables from the Bureau of Economic Analysis, harmonized to the 357-industry system used throughout the paper. We use the Use Table After Redefinitions (producer-price basis) to construct the I–O matrix  $A_{jk}$ , commodity supply shares  $\theta_{ci}$ , input use shares  $\omega_{ci}$ , domestic shares  $\zeta_{ci}$ , and PCE expenditure weights. Source:  
<https://www.bea.gov/industry/input-output-accounts-data>.
2. **Quarterly Census of Employment and Wages (QCEW)**. BLS establishment-level data aggregated to 6-digit NAICS, providing annual average employment, total annual wages, and annual average establishment counts. Coverage: 2017–2024, approximately 1,117 unique NAICS-6 codes. Used for: sector classification (establishment growth), within-sector employment and wage dispersion, shift-share instrument construction, and Phillips curve regressions. Source:  
<https://www.bls.gov/cew/downloadable-data-files.htm>.
3. **COVID-19 Job-Exposure Matrix (JEM)**. The [Oude Hengel et al. \(2021\)](#) cross-country job-exposure matrix mapping ISCO-08 4-digit occupations to contact-intensity scores along five dimensions: nature of contacts, contaminated workspaces, location, social distance, and face covering. We average the five columns across the Denmark, Netherlands, and United Kingdom panels to obtain a single ISCO-08-level COVID-exposure index per occupation, which is then mapped to U.S. SOC-2018 occupations and aggregated to BEA-357 industries via pre-pandemic occupation–employment shares. Coverage: 745 of 789 SOC columns matched; 86.9% of industry employment covered on average. Used for: construction of the predetermined occupational exposure index  $E_i$  that anchors the JEM-IV identification of the wall mechanism in Section 2. Source:  
Supplementary file to [Oude Hengel et al. \(2021\)](#).
4. **Occupational Employment and Wage Statistics (OEWS)**. BLS occupation-by-industry employment counts at the 6-digit SOC and 4-digit NAICS levels, annual cross-sections. We use the pre-pandemic (2017–2019) cross-section to compute industry-level occupation-mix shares that aggregate the SOC-level JEM exposure index to the BEA-357 industry level. Used for: industry aggregation of occupational exposure (the predetermined occupation mix in the construction of  $E_i$ ). Source:  
<https://www.bls.gov/oes/>.
5. **Occupational concordances: ISCO-08 to SOC-2018**. BLS bilateral occupation concordances chaining ISCO-08-to-SOC-2010 and SOC-2010-to-SOC-2018. Used to translate the ISCO-08-based JEM exposure index into U.S. SOC-2018 occupations prior to industry aggregation via OEWS. Source:  
<https://www.bls.gov/soc/soccrosswalks.htm>.
6. **Producer Price Index (PPI)**. BLS output price indices at 6-digit NAICS, annual average observations. Coverage varies across the four supply chains, with highest coverage in upstream manufacturing and consumer-services tiers and lowest coverage in distribution and professional-services tiers. Used for: price-dispersion diagnostics and robustness checks. Source:  
<https://www.bls.gov/ppi/databases/>.
7. **Published headline price indices (FRED)**. BEA headline Personal Consumption Expenditures price index (FRED PCEPI), BEA Core PCE price index excluding food and energy (FRED PCEPILFE), and BLS Consumer

Price Index for All Urban Consumers (FRED CPIAUCSL). Used for: methodological comparison of our GO-based inflation aggregate to the BEA- and BLS-published headline indices in Online Appendix A.1.5, and as the realized-inflation benchmark in the expectations regressions of Online Appendix A.6. Source: <https://fred.stlouisfed.org/>.

8. **BEA Gross Output Accounts.** Annual industry-level gross output, value added, and GO deflators from BEA GDP-by-Industry accounts. Used for: sector-level output, productivity, and Domar weight computation. Source: <https://www.bea.gov/data/gdp/gdp-by-industry>.
9. **BEA/BLS KLEMS Dataset.** Joint BEA/BLS integrated industry-level accounts (63 industries, 1997–2023) providing labor compensation, intermediate input expenditure, and relative prices. Used for: estimating the labor–materials elasticity  $\rho$  via the KLEMS panel regression described in Section 2.1; cross-validating the retail margin expansion in Online Appendix A.8. Source: <https://www.bls.gov/productivity/tables/>.
10. **Census Quarterly Financial Statistics (QFS, formerly QFR).** Quarterly firm-level income-statement aggregates for major industry groups: revenues, operating costs, depreciation, and interest expense. We compute annual operating and gross margins from the sum of four quarters. Coverage: 1999–2024. Used for: retail vs. non-retail operating-margin diagnostics in Online Appendix A.8. Source: <https://www.census.gov/programs-surveys/qfr.html>.
11. **BLS International Price Program.** Import price indices by end-use category, mapped to model sectors via BEA import commodity shares. Used for: exogenous import price shocks in the structural model. Source: <https://www.bls.gov/mxp/>.
12. **NAICS Concordances.** Census Bureau crosswalks for 2017-to-2022 and 2022-to-2017 NAICS code mappings (approximately 1,150 concordance pairs). Used for: maintaining consistent industry panels across the 2022 NAICS revision. Source: <https://www.census.gov/naics/>.