Capacity Walls in Production Networks: Inflation after COVID*

Lukasz A. Drozd
Federal Reserve Bank of Philadelphia
lukasz.drozd@phil.frb.org

Marina M. Tavares
International Monetary Fund
mtavares@imf.org

July 2026

Abstract

We show that the post-pandemic inflation surge was driven in large part by the uneven destruction of sectoral workforces and the organizational capital embodied in them. Empirically, we establish a causal link between predetermined industry exposure to COVID-19 and both the 2020 supply-chain employment contraction and higher inflation at reopening. Conditional on exposure, industries whose supplier networks recovered faster experienced *less* inflation—the opposite of sectoral demand-pull. We quantify the mechanism in a nonlinear multisector New Keynesian model whose key new friction—the convex cost of rebuilding workforces—is disciplined by micro estimates of hiring costs measured outside the inflation episode. The calibrated model accounts for essentially all above-trend inflation in 2021 and roughly 70% at the 2022 peak. Crucially, the cross-sectoral unevenness of the contraction—not its aggregate size—drives the result: distributing the same total capacity destruction uniformly largely eliminates the inflationary momentum.

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, Makoto Nakajima, and Elisa Rubbo for helpful comments and discussions; Elisa Rubbo in particular for advice on modeling sticky prices in a network setup, and João Gomes for insightful suggestions and early feedback. We also thank Claire Jellison for research assistance at the early stages of the project. 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.

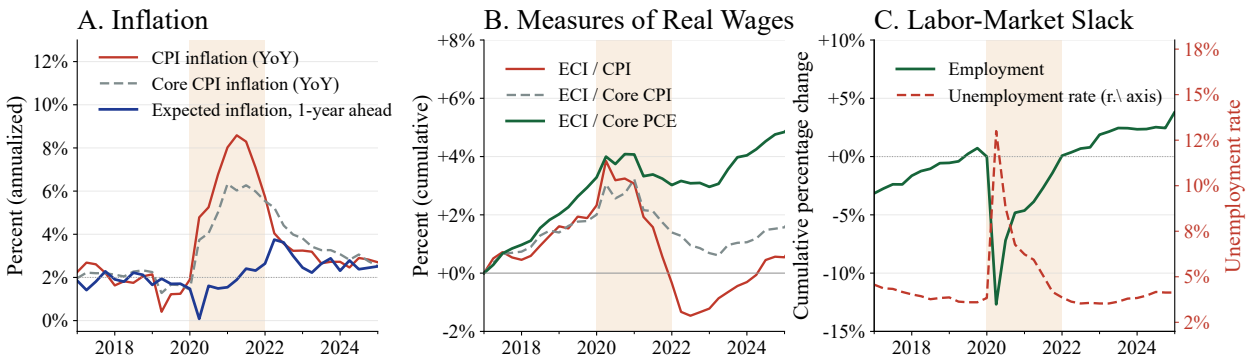

Figure 1: Inflation after COVID.

Notes: Panel A: one-year-ahead inflation expectations (FRED series `EXPINF1YR`), overlaid with year-over-year changes 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 relative to a 2017Q1 baseline. PCE deflators are `PCEPI` and `PCEPILFE`. Panel C: civilian employment level (`CE160V`, CPS household survey; left axis) expressed as the cumulative percentage change relative to its 2019 average; the dotted line marks the 2019 baseline. The right axis (dashed red) reports the civilian unemployment rate (`UNRATE`) on a common time grid. The shaded region marks 2020Q1–2021Q4. Quarterly averages of monthly series. Source: BLS and Federal Reserve Bank of Cleveland.

1 Introduction

Inflation accelerated through 2021 and peaked in mid-2022; yet, over most of this run-up, conventional measures of slack—unemployment, employment, the output gap—continued to show an economy operating below its pre-pandemic level. One-year-ahead inflation expectations rose far less than realized inflation, and composition-adjusted real wages declined (Figure 1). To see why this constellation of facts challenges conventional frameworks, recall that in the standard New Keynesian (NK) model, inflation is governed by the Phillips curve

$$\pi_t = \kappa\, x_t + \beta\, \mathbb{E}_t[\pi_{t+1}] + u_t,$$

where x_t is the output gap, $\mathbb{E}_t[\pi_{t+1}]$ is expected inflation, and u_t is a cost-push shock (β is the discount factor, $\beta \approx 1$). Three features of the episode discipline how this equation can be read: anchored expectations fix $\mathbb{E}_t[\pi_{t+1}]$; persistent measured slack rules out a large positive x_t; and declining real wages are difficult to reconcile with demand pulling prices up along a stable curve via the labor market. What remains is a supply-side shock—an upward shift of the curve (u_t), a locally steeper slope (κ), or both. We point to the supply origins of this shock: the pandemic fractured productive capacity unevenly across sectors, and as demand returned, sector-specific shortages propagated through the production network and contributed critically to inflation through its 2022 peak.

Our mechanism centers on the labor market. The pandemic did not merely con-

tract output; it destroyed sectoral workforces and the organizational capital embodied in them—the trained teams, matches, and routines built over years. This damage was highly uneven, driven by predetermined occupational exposure to the virus and therefore orthogonal to the subsequent recovery of demand. When demand rebounded, firms in the most exposed sectors faced a state variable that reopening could not instantly undo: workforces had to be rebuilt, and the cost of rebuilding rises sharply with its pace, as an extensive labor-economics literature documents (Linckh et al., 2025; Muehlemann and Strupler Leiser, 2018; Muehlemann and Pfeifer, 2016; Aepli et al., 2024). This convex rebuilding cost provides the microfoundation for u_t—and, in part, for the local steepening of κ. Because the damage was concentrated in a subset of sectors linked through input–output flows, downstream industries inherited their most-exposed suppliers' constraints, and the network geometry turned a sector-specific supply shock into an economy-wide inflation event.[1] We establish this mechanism causally in the cross-section of industries, and we quantify it in a nonlinear multisector New Keynesian model whose key friction is disciplined by micro evidence on hiring costs measured outside the inflation episode.

Our empirical analysis rests on a predetermined measure of the virus's impact on production. We construct an industry-level exposure index from the nature and intensity of workplace contacts in the occupations that make up each industry's pre-pandemic workforce—as scored by the COVID-19 Job Exposure Matrix (JEM) of Oude Hengel et al. (2022)—and propagate it upstream through input–output linkages, so that each industry inherits the exposure of its supply chain. The index is built strictly from pre-pandemic occupational characteristics and hence does not embed the endogenous employment contraction.

Two key results support a causal interpretation of the data. First, in an event-study difference-in-differences design, upstream exposure predicts the employment contraction of an industry's supply chain from early 2020 to early 2021.[2] Industries whose supply chains contracted most entered the reopening with capacity more fractured relative to the 2019 baseline, and their inflation should have been correspondingly higher. It was. The relationship operates through prices relative to costs: decomposing sectoral inflation

[1]Independent evidence from natural-disaster studies supports this propagation channel. Barrot and Sauvagnat (2016) document that a customer firm loses 2 to 3 percentage points of sales growth in the year after one of its suppliers is hit by a major U.S. natural disaster, even though the typical affected supplier accounts for only about 2.5% of the customer's input costs; each dollar of sales lost at the supplier induces roughly $2.40 of lost sales downstream, with propagation concentrated in specific, hard-to-replace inputs. See also Carvalho et al. (2021) and Miyauchi (2024).

[2]Measuring the contraction from early 2020 to early 2021 sidesteps the transient spring-2020 layoffs, which were largely reversed through recalls by the fall; annual averages would conflate this transient collapse with the persistent damage.

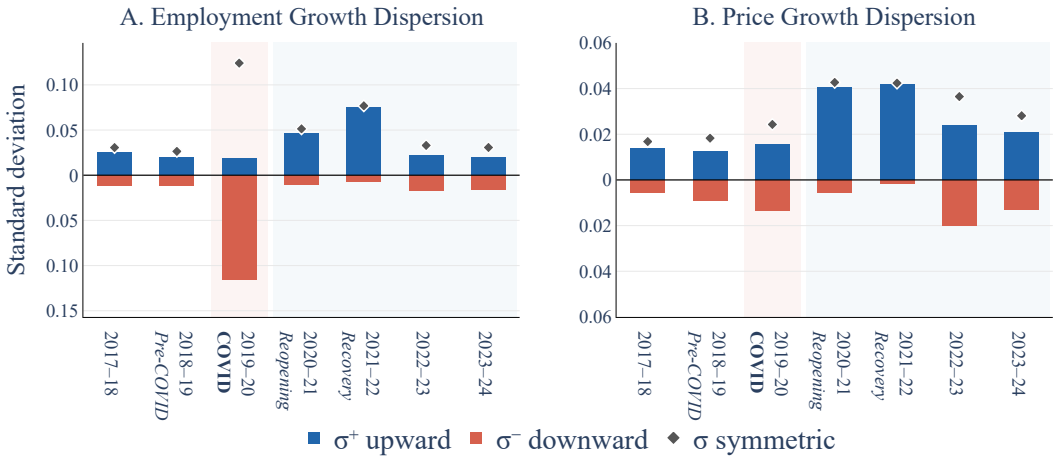

Figure 2: Within-sector dispersion of employment growth and price growth.

Notes: $\sigma(\Delta \log N)$ (standard deviation) is computed from QCEW NAICS-6 employment and aggregated across BEA sectors using 2017 employment weights. $\sigma(\Delta \log P)$ is computed from BEA-357 gross-output deflators and aggregated across sectors using Domar weights based on 2017 gross-output shares. Within each sector, industries are weighted by 2019 employment. Blue bars report $\sigma^+(x) := \sigma(\max\{0, x\})$; red bars report $\sigma^-(x) := \sigma(\min\{0, x\})$; diamonds report the symmetric σ. All year-pairs use raw QCEW annual averages on both sides, except the 2019–2020 slot: to capture the persistent 2020 damage after recalls from unemployment had stabilized, we replace that year-pair with the Jan–Feb 2020 to Jan–Feb 2021 log change per NAICS-6, i.e. $\Delta \log N_{19\to20} \equiv \log(N^{\text{Jan–Feb 2021}}/N^{\text{Jan–Feb 2020}})$. Values first winsorized at 3%.

into measured cost components and an unexplained margin residual, we find that the residual accounts for the exposure gradient. The same relationship appears in second moments—sectors in which exposure was more unevenly distributed experienced larger increases in within-sector price dispersion—and is visible in the raw data (Figure 2): the 2020 employment collapse was enormously dispersed across industries, the rebound was not, and the collapse was followed by a sharp and persistent rise in sectoral inflation dispersion. Our exposure measure connects these patterns—uneven destruction followed by uniform rebuilding—to the loss of organizational capital.

Second, we confront the leading alternative: a demand spike pushing activity against 2019 capacity. To separate the two stories, we examine the joint behavior of recovery speed and inflation across industries. Conditional on exposure, industries whose supplier networks recovered faster experienced *less* inflation. This sign pattern separates capacity destruction from a pure sectoral demand-pull story: had demand pushed activity against intact capacity, faster quantity recoveries would have coincided with *more* inflation, as demand traced out upward-sloping supply curves. The rebuilding mechanism predicts the opposite, because faster supplier recovery relaxes capacity constraints for a given initial disruption. Flat pre-trends, together with the subsequent fading of both relationships,

argue against fixed cross-industry confounds: a fixed confound would not switch on in 2020 and fade by 2024.

We interpret these facts through the lens of a nonlinear multisector New Keynesian model with sticky prices, sticky wages, and input–output linkages. The model formalizes the rebuilding friction: expanding employment above the workforce capacity inherited from the previous period incurs a convex cost, because new employees arrive without organizational capital and are therefore less productive. We refer to this friction as the *capacity wall* (CW). The central quantitative question is whether independently measured workforce-rebuilding costs are sufficiently convex, once propagated through the observed production network, to convert a dispersed recovery into aggregate inflation.

We discipline the network structure and input shares using the 2017 BEA input–output tables, which we organize into four vertical production chains, running from upstream suppliers through intermediate producers to downstream, consumer-facing industries. Within each node we embed asymmetric destructive shock that affects only a fraction of firms; this aggregation preserves tractability and capture both asymmetry and network propagation. The CW friction's level and curvature come from recent micro-level studies showing that hiring costs are both large and increasing in firms' own hiring intensity and local labor-market tightness (Linckh et al., 2025; Muehlemann and Strupler Leiser, 2018; Muehlemann and Pfeifer, 2016; Aepli et al., 2024). These costs are naturally interpreted as investments in organizational capital in the sense of Prescott and Visscher (1980) and Atkeson and Kehoe (2005): roughly 80% of measured hiring costs are nonwage expenditures on recruitment, post-match training, adaptation, informal instruction, and disruption to incumbent workers. We anchor the level using Linckh et al. (2025), who estimate that filling a skilled-worker vacancy costs approximately 22 weeks of wages,[3] and convert this per-hire expenditure into an annual user cost by applying the firm's real cost of capital and the annualized separation hazard implied by worker quit rates. For the curvature, Muehlemann and Pfeifer (2016) estimate that total hiring costs rise with hiring intensity with an elasticity of about 1.3, while Linckh et al. (2025) find that a one-standard-deviation increase in local labor-market tightness raises the largest post-match cost components—adaptation and disruption—by approximately 10–12%. The U.S. vacancy-to-unemployment ratio rose from less than 1.2 over 2017–19 to nearly 2.0 at its 2022 peak (Domash and Summers, 2022). Given the absence of evidence of broad mismatch between workers and jobs (Forsythe et al., 2022), we map this increase into a peak

[3]We define "skilled" workers broadly as those in occupations requiring occupation-specific or post-match training, rather than only college-educated workers. Applied to the U.S. wage bill, this group accounts for roughly four-fifths of compensation; we scale the calibrated cost accordingly. See Section 4. Muehlemann and Strupler Leiser (2018) report a somewhat lower share for Switzerland.

labor-market tightness approximately three standard deviations above its historical norm. Together, these inputs imply a peak rebuilding wedge of approximately 13%.

Our calibrated model assigns a substantial role to the CW. In cumulative terms, the mechanism accounts for essentially all above-trend inflation in 2021, roughly 70% at the 2022 peak, and about 35% by 2024. The central quantitative result, however, is that the geometry of the initial contraction matters. Approximately 60% of the cumulative inflation generated by the mechanism through 2024—and a large majority of the inflation it generated in 2021—arises from the uneven destruction of workforce-embodied organizational capital across sectors. A counterfactual that distributes the same aggregate capacity loss uniformly across sectors nearly eliminates the inflationary momentum in 2021. Although the direct flow effect of the friction wanes in 2023, its contribution to the aggregate price level persists well beyond the 2022 inflation peak. The model therefore rationalizes the apparent post-COVID steepening of the aggregate Phillips curve as a consequence of uneven rebuilding costs, rather than as a generic structural nonlinearity activated by any sufficiently large aggregate shock.

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.

Related literature. Supply disruptions are widely recognized as a central driver of post-COVID inflation, but the disruptions themselves are usually taken as given: with few exceptions, they have not been traced back to the COVID-19 shock as the source of the reopening surge. The most prominent exception is Guerrieri et al. (2022), who model the pandemic shutdown as a sectoral supply shock, though their mechanism—Keynesian demand shortages—operates during, rather than after, the contraction. Our contribution is to trace the reopening inflation to the pandemic shock itself: causally, in the cross-section of industries, and quantitatively, through a calibration disciplined by micro-level evidence on workforce-rebuilding costs. More broadly, we show that uneven capacity destruction can generate inflation during the recovery even when substantial aggregate slack remains, because inflation depends on both the aggregate capacity inherited from the shock and its distribution within the production network.

The most closely related work falls into three groups: models of pandemic-era capacity constraints and reallocation (Comin et al., 2023; Ferrante et al., 2023; Bohr, 2024; Fornaro, 2025; Rubbo, 2024);[4] theories of a nonlinear Phillips curve (Benigno and Eggerts-

[4]Bohr (2024) models firms' precautionary capacity buffers—ample before the pandemic and eroded by

son, 2023; Harding et al., 2023; Schmitt-Grohé and Uribe, 2022; Blanco et al., 2024); and quasi-experimental evidence on the slope of the Phillips curve (Cerrato and Gitti, 2025; Gitti, 2026).

In the first group, Comin et al. (2023) develop a multisector open-economy New Keynesian model with exogenous, occasionally binding capacity constraints and show that these constraints can explain about half of the goods-sector inflation surge in 2021–2022; our model generates similar constraints endogenously and provides a foundation for that core mechanism. Ferrante et al. (2023) study how the pandemic demand rotation from services to goods generates inflation in the presence of convex labor-reallocation costs; our model captures these forces through demand-driven destruction of organizational capital and offers a microfoundation for their convex adjustment costs. We also build on Rubbo (2024), who shows that shocks shifting expenditure toward inelastically supplied goods and factors act as endogenous cost-push shocks, and whose price-based decomposition attributes 2020 inflation to industry-specific shocks and the post-2021 period to aggregate factors. Relative to her work, we endogenize the supply-curve heterogeneity as pandemic-destroyed capacity, which yields a complementary reading of the post-2021 period: once the rebuilding friction propagates through the network and into wages, its relative-price fingerprint fades, so inflation that a price-based decomposition classifies as aggregate is, in our model, still the product of the COVID shock itself.[5]

In the second group, Benigno and Eggertsson (2023), Schmitt-Grohé and Uribe (2022), Harding et al. (2023), and Blanco et al. (2024) resolve the puzzle by positing a nonlinear Phillips curve, whose flat local pre-COVID slope is uninformative about the impact of large shocks. Benigno and Eggertsson (2023) locate the nonlinearity in the labor market: with search frictions and downward wage rigidity, the Phillips curve steepens sharply once the vacancy-to-unemployment ratio crosses into labor-shortage territory, and marginal (new-hire) labor costs decouple from more sluggish *average* wages. Schmitt-Grohé and Uribe (2022) also locate the nonlinearity on the wage side but through a different mechanism: heterogeneous downward nominal wage rigidity across workers renders the aggregate wage Phillips curve flat in slack territory—many workers pressed against the DNWR floor—and steep once inflation rises and the floor unbinds. Harding et al. (2023) locate it in

it—as the source of both the pre-COVID flatness and the post-COVID steepening of the Phillips curve; his constraint is capital-side and chosen ex ante, while ours is workforce-embodied and destroyed by the shock. Fornaro (2025) studies the complementary fiscal dimension of stimulus against binding supply constraints.

[5]Several other contributions are also indirectly related to our work. Moscarini and Postel-Vinay (2023) study the same episode from the worker side, documenting how job-ladder dynamics shaped pandemic-era labor-market reallocation. di Giovanni et al. (2022) extend the cross-sector reallocation logic to global value chains. Gagliardone et al. (2025) show that the reduced-form flatness of the Phillips curve reflects a low elasticity of marginal cost to the output gap in moderate-inflation regimes.

price setting: under a quasi-kinked demand schedule, firms raise prices more aggressively than they cut them, so cost-push shocks pass through more strongly once inflation is already elevated. Blanco et al. (2024) locate it in the timing of price adjustment: in a state-dependent (menu-cost) economy the frequency of repricing is endogenous, so a large cost shock pushes many firms across the menu-cost threshold at once and the aggregate response amplifies far beyond the Calvo benchmark—a mechanism they document in U.S. micro price data during the surge, with the fraction of prices changing per quarter rising from roughly 25% pre-COVID to 40–50% at the 2021–22 peak. Beaudry et al. (2025), in turn, argue that the empirical case for a generic nonlinear Phillips curve is fragile to how the post-COVID episode is treated; our mechanism supplies one economic reading of that fragility, in which the apparent steepening is episode-specific and reflects the geometry of the pandemic shock rather than a structural feature that would recur under any large aggregate shock. Our mechanism is related to all three theories, and closest in spirit to Benigno and Eggertsson (2023): the capacity wall is likewise a wedge between marginal and average labor costs that opens in a tight market. It differs, however, in measurement, in origin, and in the level at which it operates. In measurement: most of the wedge we measure is not paid out in wages—the bulk consists of nonwage investments in organizational capital—recruiting, training, adaptation, and disruption to incumbent workers—whose level and convexity we take from the micro labor literature, detached from the behavior of marginal wages. In origin: their wedge opens and closes with contemporaneous labor-market tightness, whereas ours is tied to a state variable—the depleted stock of organizational capital—created by the pandemic and unwound only as capacity is rebuilt, so its effects outlast the tightness that produced them. And in level: the wedge operates at the industry level and propagates through the production network, which is what allows us to bring causal, cross-sectional evidence to bear on its origin in the pandemic shock.

In the third group, Cerrato and Gitti (2025) provide quasi-experimental estimates of the U.S. Phillips-curve slope around the pandemic from cross-city variation. They document a stark steepening—from roughly 0.15 before COVID to above unity after it, the steepest since the mid-1970s—without which the unemployment gap would have contributed next to nothing to the surge, and they trace the steepening to a sharp fall in the labor-supply elasticity of the marginal worker together with a higher frequency of price changes. Our mechanism speaks directly to the first driver: the capacity wall makes the effective supply of labor to rebuilding firms steeply convex, so the collapsed elasticity they measure is the labor-market footprint our model predicts—and predicts to be temporary, receding as capacity is rebuilt. In a follow-up solo paper, Gitti (2026) shows that the regional

Phillips curve is kinked in labor-market tightness—flat in slack markets, steepening past the labor-shortage threshold, with the kink operating through new-hire wages and present already before COVID. Her evidence documents, with quasi-experimental variation, the convexity of marginal labor costs in tight markets that our calibration imposes from hiring-cost estimates.

Finally, our modeling and measurement build on the production-network literature (Acemoglu et al., 2012; Baqaee and Farhi, 2019; Rubbo, 2023), and in particular on Baqaee and Farhi (2022), who analyze sectoral supply and demand shocks in a disaggregated Keynesian economy applied to the COVID-19 crisis.[6]

2 Empirical Evidence

This section lays out the evidence linking predetermined industry exposure to COVID-19 to the subsequent employment contraction and inflation in the cross-section of U.S. industries. We first describe the data and measurement framework, and then present the reduced-form results.

2.1 Data and Measurement

Data sources and sample. We link eight data sources: (i) the Bureau of Economic Analysis (BEA) 2017 detail-level supply–use tables (SUTs) in producer prices; (ii) the corresponding annual BEA detail-level gross-output (GO) accounts; (iii) employment and payroll data from the Quarterly Census of Employment and Wages (QCEW); (iv) 2019 industry–occupation employment shares from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS); (v) import-price indices from the BLS International Price Program (IPP); (vi) import shares from the U.S. Census Bureau; (vii) the COVID-19 Job-Exposure Matrix (JEM) of Oude Hengel et al. (2022); and (viii) the BEA–BLS KLEMS data.

To measure capacity destruction and its propagation through production networks at a common level of industry detail, we harmonize all data sources to the BEA SUT classification. We merge the QCEW with the SUTs using BEA-published and NAICS-based crosswalks, and crosswalk the IPP and Census data at the finest level available. The JEM is originally classified by ISCO-08 occupation codes. We map it into SOC occupations,

[6]Our work is also related to aggregate decompositions of pandemic-era inflation in Bernanke and Blanchard (2025), Blanchard and Bernanke (2024), and Ball et al. (2022), and to related work by Arias et al. (2026) and Giannone and Primiceri (2024); qualitatively the shock we identify is consistent with their findings, but since our model is nonlinear we do not attempt a precise mapping to their linear-model-based VAR evidence.

aggregate the occupational scores to NAICS industries using 2019 OEWS employment shares, and then map the resulting industry measures into the SUT classification. The KLEMS data are used only to estimate the elasticity of substitution between labor and intermediate inputs, the one parameter required by our accounting framework beyond observed input–output and cost shares.[7]

We measure sectoral inflation using annual log changes in BEA detail-level industry gross-output deflators, which are internally consistent with the SUT methodology. For aggregate comparisons, we weight industry price changes using fixed 2017 PCE final-demand shares from the Use table in producer prices. Online Appendix A.1.7 compares the resulting measure with published PCE- and CPI-based inflation measures.

Our baseline analysis, both here and in the quantitative model, excludes four groups of sectors whose measured prices are heavily influenced by external forces, nonmarket pricing, or imputation: real estate, petroleum and mining, finance and insurance, and government. We refer to these groups collectively as the *exogenous*, or "EXO," sectors and drop them from the regressions. The EXO sectors nonetheless remain in the input–output matrix, so that transmission through the network is fully accounted for; in the structural model, their prices are treated as exogenous processes. We refer to PCE-weighted gross-output inflation as GO inflation and to the narrower index that zeroes out the EXO weights as N-GO inflation.

Predetermined industry exposure to the COVID shock. We construct industry exposure from the pre-pandemic occupational composition of each industry's workforce and from occupation-level characteristics that determine workplace exposure to SARS-CoV-2. The source is the Job-Exposure Matrix developed for Denmark, the Netherlands, and the United Kingdom at the ISCO-08 occupation level by Oude Hengel et al. (2022). We use the Danish tabulation: it is the most detailed in its occupational coverage and has been validated against population-wide epidemiological data, which show it to be highly predictive of infection outcomes (van der Feltz et al., 2023).[8]

For each occupation o, let e_o denote the average of the five JEM dimensions, each scored on a three-level infection-risk scale: the nature of workplace contacts, exposure

[7]The original 2017 detail-level SUT contains 402 industries. We aggregate these industries to 357 to bridge panel breaks introduced by the 2022 NAICS revision, which primarily affected retail trade. Online Appendix A.2 describes the crosswalks and aggregation procedure, and Online Appendix A.11 provides a complete tabulation of the data sources.

[8]Related occupation-based exposure measures used in the pandemic literature include the work-from-home feasibility index of Dingel and Neiman (2020) and the physical-proximity measure of Mongey et al. (2021). The JEM offers three advantages: it is multidimensional (six workplace-exposure dimensions rather than a single proximity score), it was designed by epidemiologists specifically to capture SARS-CoV-2 transmission risk, and its tabulation has been validated (van der Feltz et al., 2023).

to contaminated workspaces, work location, the feasibility of social distancing, and the feasibility of face covering. Let s_{io}^{2019} denote occupation o's share of employment in industry i. Industry i's own workforce exposure is $E_i \equiv \sum_o s_{io}^{2019} e_o$, computed after applying the occupation crosswalk and the 2019 OEWS shares at the NAICS level underlying the SUT classification; the vector $E = (E_1, \ldots, E_I)^\top$ collects these values across industries. Because both the occupational shares and the underlying workplace characteristics are determined independently of subsequent industry outcomes, E_i is predetermined with respect to the employment and price changes studied below.

An industry can be exposed directly through its own workforce or indirectly through disruptions among its suppliers. To capture the latter, let A denote the domestic intermediate-input matrix, whose construction we discuss below and detail in Online Appendix A.1. We define the supplier-network exposure index of industry i as the discounted sum of own exposures propagated upstream through the network:

$$\mathcal{E} \equiv \left(A + \delta A^2 + \delta^2 A^3 + \cdots\right) E = A L_\delta E, \qquad L_\delta \equiv (I - \delta A)^{-1}, \tag{1}$$

where $\delta \in [0, 1]$ discounts exposure transmitted through increasingly distant network layers. When $\delta = 0$, $\mathcal{E} = AE$ and the index includes only direct suppliers; when $\delta = 1$, it incorporates all upstream layers. Our baseline sets $\delta = 0.5$; Online Appendix A.10.1 reports robustness to alternative values.[9]

Inflation decomposition and cost tracing. To measure and decompose sectoral inflation, and to trace it through the U.S. production network, we separate observed price growth into measured labor costs, domestic and imported intermediate-input costs, and a residual, and we link industries together using the detail-level I–O tables. To that end, we first define industry i's measured margin M_i as the ratio of sales revenue to measured variable production cost:

$$M_i \equiv \frac{P_i Y_i}{C_i(Y_i)}, \qquad C_i(Y_i) \equiv \min_{N_i, m_i} \left\{W_i N_i + p_i m_i\right\} \quad \text{subject to producing } Y_i, \tag{2}$$

where N_i is employment, m_i is a materials composite, and p_i is its price index. This margin is the key residual of our analysis; it is identified by measuring the cost components on the right-hand side of equation (3). The cost function solves the cost-minimization problem

[9]The JEM does not cover 11 agricultural industries, which account for nearly all of the agricultural block. We exclude these industries, together with the EXO sectors, as regression observations, leaving $N = 330$ BEA detail-level industries. The uncovered agricultural industries are assigned zero exposure so that they can remain as nodes in the production network. All 357 industries therefore remain in the 357×357 matrix A; the sample restriction applies only to the regression observations.

above under a nested CES production structure that combines labor with intermediate inputs sourced from other industries. The intermediate composite, consistent with the use/make structure of the U.S. I–O accounts, combines domestic and imported commodities. Online Appendix A.1 states the exact aggregators and derives the decomposition in full.

Log-linearizing the resulting cost structure yields the headline identity, in log growth rates, that we use throughout:

$$\Delta \log P = \underbrace{\Delta \log M}_{\Delta\text{MARGIN}} + \underbrace{\underbrace{\operatorname{diag}(b)\,\Delta \log W}_{\Delta\text{WAGE}} + \underbrace{-\Delta \log(Y/N)}_{\Delta\text{PROD}}}_{\Delta\text{ULC}} + \underbrace{A\,\Delta \log P}_{\Delta\text{MAT}} + \underbrace{A^{\text{IMP}}\,\Delta \log P^{\text{IMP}}}_{\Delta\text{IMP}}. \tag{3}$$

All logarithms, ratios, and differences of vectors are understood element by element. On the left-hand side, $\Delta \log P$ is the vector of annual industry price changes measured from BEA gross-output deflators. The terms on the right-hand side are the measured margin residual defined above ΔMARGIN; unit labor costs ΔULC, comprising a wage and a labor-productivity term; and domestic ΔMAT and imported ΔIMP intermediate-input inflation.

Two primitives determine the loadings on labor and intermediate-input costs. The share s_i^N is labor's share of variable costs in industry i at the 2017 linearization point. Because capital is predetermined, capital income is treated as steady-state profits rather than as a variable cost. The parameter ρ is the elasticity of substitution between labor and the materials composite and is the only production elasticity that enters the first-order approximation. These primitives imply the CES-adjusted labor and materials coefficients $b_i = \rho + (1-\rho)s_i^N$ and $d_i = (1 - s_i^N)(1-\rho)$, which satisfy $b_i + d_i = 1$. We collect these coefficients in the vectors $b = (b_1, \ldots, b_I)^\top$ and $d = (d_1, \ldots, d_I)^\top$.

Combining the materials coefficients with the make and use structure of the SUT gives the domestic and imported input matrices:

$$A = \operatorname{diag}(d)(\zeta \odot \Omega)^\top \Theta^\top, \quad A^{\text{IMP}} = \operatorname{diag}(d)\left[(\mathbf{1} - \zeta) \odot \Omega\right]^\top. \tag{4}$$

Here, $\Omega^\top = (\omega_{ij})$ is the normalized BEA Use matrix, $\Theta = (\theta_{ij})$ is the BEA make-share matrix, $\zeta = (\zeta_{ij})$ contains domestic sourcing shares, and $\odot$ denotes the Hadamard product. The matrix $\Omega^\top$ maps commodity prices into the materials costs of the industries that use them, while $\Theta^\top$ maps industry output prices into domestic commodity prices. Imported prices are observed directly at the commodity level and therefore require no corresponding make-share mapping.

We estimate ρ using BEA–BLS KLEMS data and the cost-minimization condition linking the relative quantities of labor and intermediate inputs to their relative prices. Our headline estimate is $\widehat{\rho} = 0.39$, with a 95% confidence interval of [0.06, 0.72].[10]

We use equation (3) as an identity: all terms other than the margin are measured, and the margin residual is defined so that the equation holds for each industry at each point in time. Before proceeding, three remarks about the decomposition are in order.

Remark 1. *The materials term* ΔMAT *captures propagation through the production network. A unit increase in wages raises prices directly by* b*. To see this, consider a symmetric economy with common labor and materials coefficients* b *and* d*, and domestic sourcing share* ζ*. The induced increase in domestic materials prices then contributes* $b\zeta d$*, followed by* $b(\zeta d)^2$*, and so on. The total price response is therefore* $b + b\zeta d + b(\zeta d)^2 + \cdots = b/(1 - \zeta d) = b/(b + (1 - \zeta)d)$*, where the final equality uses* $b + d = 1$*. With fully domestic sourcing (*$\zeta = 1$*), the total response is exactly one-for-one. When some inputs are imported and their prices are held fixed (*$\zeta < 1$*), the response is smaller than one. Thus, the production network propagates heterogeneous cost shocks across industries and raises the response to a uniform wage shock. This conclusion holds for any* ρ *that implies nonnegative labor and materials coefficients, including the empirically relevant case* $0 \leq \rho \leq 1$*. Intuitively, this is because all variable costs are ultimately attributable to labor. The comparable total response to a uniform unit increase in margins is* $1/(1 - \zeta d)$*, and equal to* $1/b$ *under fully domestic sourcing.*

Remark 2. *The measured margin* M_i *in this identity is closer in definition to a gross profit margin than to net profit; put differently,* C_i *is closely related to the cost of goods sold in firm-level accounting. The reason is that the I–O accounts are establishment based rather than firm based. For multi-establishment firms, head-office activities are separated out and recorded under the Management of Companies sector (NAICS 55)—a sector whose revenue is largely imputed from its operating costs. Purchases of headquarters services by producing establishments capture only explicit transactions and are small. For the remaining sectors, measured labor costs therefore mostly reflect labor that firms include in cost of goods sold rather than in selling, general, and administrative expenses.*

Remark 3. *The attribution is not causal, because the productivity term is endogenous: any endogenous movement in* ΔPROD *feeds into measured unit labor costs. Suppose, for example, that demand rises and firms expand output while raising margins. Under diminishing marginal product of labor, output per worker* Y/N *falls, making* $-\Delta \log(Y/N)$ *positive and thereby increasing*

[10] This estimate is higher than estimates obtained at finer levels of disaggregation (e.g., Atalay, 2017; Boehm et al., 2019), making it a conservative baseline. Online Appendix A.1.6 provides estimation details.

measured unit labor costs. The decomposition then assigns part of a demand-driven expansion to ΔULC*, even though labor costs are not the primitive shock.*

2.2 Results

We organize the reduced-form evidence around two facts that we document below.

Fact 1. COVID exposure predicts employment contraction and inflation relative to the 2019 baseline. Our headline regression is a pooled event-study specification over three windows, labelled as follows: pre-COVID, 2017–2019; Act I, 2019–2021; and Act II, 2021–2023. For each annualized outcome $y_{i,w}$, the regression equation is

$$y_{i,w} = \gamma_i + \theta_w + \beta_{\mathrm{I}} X_i D_w^{\mathrm{I}} + \beta_{\mathrm{II}} X_i D_w^{\mathrm{II}} + \varepsilon_{i,w}, \tag{5}$$

where γ_i and θ_w are industry and window fixed effects, and D_w^{I} and D_w^{II} are indicators for the Act I and Act II windows. The pre-COVID window is the omitted base: its interaction with the time-invariant regressor X_i is absorbed by the industry fixed effects. We estimate equation (5) separately for employment, prices, and the components of the price decomposition. The regressor X_i is always a measure of predetermined COVID exposure. For own-industry employment, we use the industry's direct exposure, $X_i = E_i$. For employment effects transmitted through the supplier network, industry prices, and their accounting components, we use network exposure, $X_i = \mathcal{E}_i$. We classify industries by their position in the production network. We define an industry's finalness as $\text{finalness}_i \equiv g_i/\kappa_i$, the ratio of its direct PCE weight to its Domar weight. Within each supply chain, industries with a finalness of at least 70 percent constitute the consumer-facing tier (h). We divide the remaining industries at the gross-output-weighted median of finalness into a middle tier (m) and an upstream tier (l). We refer to their union, (m+l), as the upstream and middle-tier sample. Section 2.1 and Online Appendix A.2 provide details on the construction. All regressions are weighted by 2017 industry value added, and standard errors are clustered by industry.

The outcome variables on the left-hand side include own ($\Delta \log N_i$) and network-propagated ($\Delta \log \mathcal{N}_i$) employment growth, gross-output inflation ($\Delta \log P_i$), and the domestic accounting components of the decomposition in (3). Table 1 reports estimates for the full sample and for the upstream and middle-tier ($m + l$) sample for all outcome variables.

Table 1 shows that more exposed supplier networks contracted sharply during Act I and subsequently recovered. The same exposure predicts higher inflation during Act I,

Table 1: Fact 1: industry inflation on supplier-network exposure $\mathcal{E}_i$ (extended).

Outcome	Act I $\hat{\beta}_{\text{Act I}}(t)$		Act II $\hat{\beta}_{\text{Act II}}(t)$	
	ALL (obs. = 330)	m+l-tier (obs. = 232)	ALL (obs. = 330)	m+l-tier (obs. = 232)
$\Delta \log N_i$ (on E_i)	−1.73	−3.37***	+0.43	+0.42
	(−1.60)	(−3.27)	(+0.37)	(+0.33)
$\Delta \log \mathcal{N}_i$	−2.53***	−2.62***	+1.74***	+1.16**
	(−5.47)	(−6.08)	(+3.41)	(+2.45)
$\Delta \log P_i$	+12.08**	+19.02***	+6.10	+1.94
	(+2.45)	(+3.55)	(+0.91)	(+0.19)
ΔMARGIN_i	+11.83	+24.81***	−16.84	−9.32
	(+1.33)	(+3.22)	(−1.35)	(−0.77)
ΔMARGIN_i ($\Delta \log(Y/N)$ controlled)	+12.51***	+20.89***	−2.49	−7.06
	(+2.58)	(+4.23)	(−0.37)	(−0.74)
ΔULC_i	−4.60	−11.93	+19.97*	+8.53
	(−0.65)	(−1.46)	(+1.87)	(+1.22)
ΔIMP_i	+0.25	+0.35*	+0.24***	+0.10
	(+1.58)	(+1.71)	(+3.38)	(+1.28)
ΔMAT_i	+4.60***	+5.80***	+2.73***	+2.63***
	(+5.13)	(+4.74)	(+5.79)	(+4.92)
$A \cdot \Delta\text{MARGIN}_i$	+3.09***	+4.12***	−0.93	−0.40
	(+4.18)	(+4.56)	(−1.05)	(−0.39)

Notes: Industry-level panel over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, Act II 2021–2023. WLS with 2017 industry value added; industry and window fixed effects; standard errors clustered by industry. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $\Delta \log \mathcal{N}_i$ replaces the own-industry LHS with the network-propagated (supplier-side) employment growth. The $\Delta \log(Y/N)$-controlled MARGIN row adds contemporaneous $\Delta \log(Y/N)_{i,w}$ as a regressor (fitted coefficient +0.867*** ALL, +0.807*** m+l-tier).

with a substantially larger effect among upstream and middle-tier industries. Given $\sigma_{\mathcal{E}} \simeq 0.105$, the point estimates imply that a one-standard-deviation increase in supplier exposure raised annual inflation by approximately 1.3 percentage points in the full sample and 2.0 percentage points in the $m + l$ sample.

After controlling for contemporaneous own-industry productivity growth, exposure predicts the measured margin residual even more strongly, particularly upstream—as expected given Remark 3.[11] The domestic-materials response is likewise positive, and the

[11] A natural question is whether the margin residual instead reflects unmeasured *marginal* labor costs—the wages of new hires, which can decouple from sluggish average wages in tight labor markets (Benigno and Eggertsson, 2023; Gitti, 2026). Two observations help relate that channel to our accounting and analysis. First, in these theories the marginal-wage wedge is a function of labor-market tightness, and measured sectoral tightness rose nearly uniformly—roughly uncorrelated with the size of each sector's collapse (see footnote 23)—so a tightness-driven wedge on its own would not readily generate the cross-sectional co-

supplier-margin pass-through term indicates that upstream margin residuals account for an important part of this channel. By Act II, most exposure coefficients have attenuated, consistent with a temporary disruption whose effects recede as capacity is rebuilt.

To display the timing more transparently, Figure 3 plots estimated coefficients β_t for each outcome variable from separate annual cross-sections without industry fixed effects, $y_{i,t} = \alpha_t + \beta_t X_i + \varepsilon_{i,t}$ for $t \in \{2018, \ldots, 2024\}$, where X_i is defined as above. The pre-COVID estimates serve as placebos, while the subsequent coefficients trace the employment, price, and cost responses year by year.

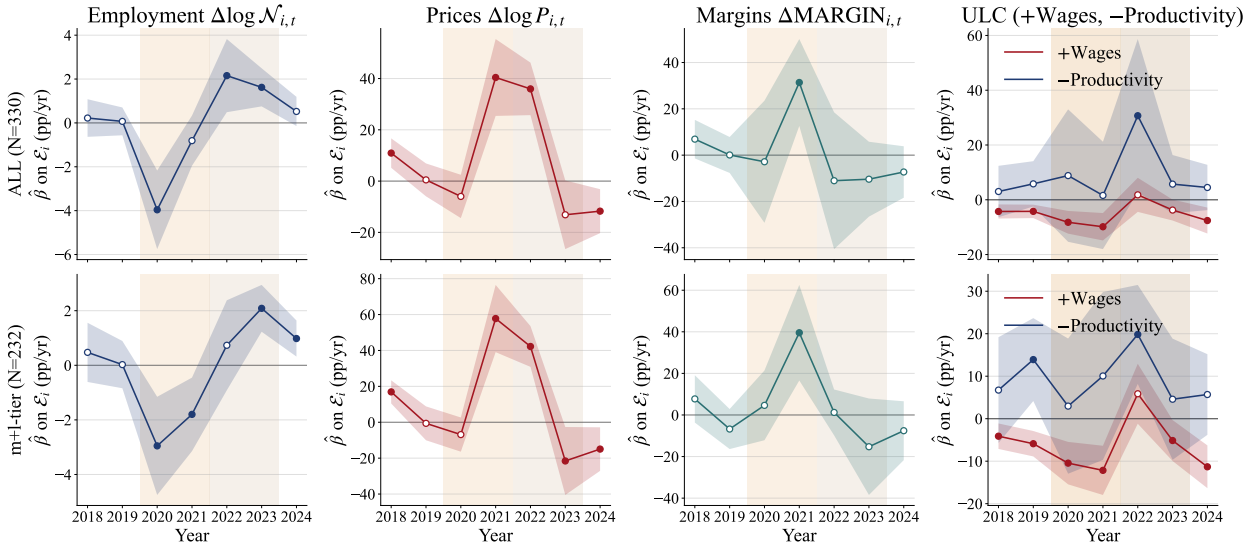

Figure 3: Dynamic effects of COVID-19 exposure.

Notes: The figure reports coefficients from annual cross-sectional regressions of the form $y_{i,t} = \alpha_t + \beta_t X_i + \varepsilon_{i,t}$, estimated separately by year. For own-industry employment, $X_i = E_i$; for network outcomes, $X_i = \mathcal{E}_i$. Regressions use 2017 value-added weights and HC1 standard errors. Bands report 95% confidence intervals.

The figure shows that, even without fixed effects, pre-trends are essentially absent: the relationship changes sharply in 2020 and subsequently attenuates. This reversal is difficult to explain with a time-invariant cross-sectional confound alone and supports an interpretation tied to the pandemic contraction and reopening.[12]

movement of the margin residual with predetermined exposure that we document. That said, at the current stage of measurement this question is difficult to answer. Measures of the wages of job switchers, defined as wage growth within the same individual, provide some evidence of an increased gap in 2022 (Federal Reserve Bank of Atlanta Wage Growth Tracker); compared to pre-pandemic fluctuations of this gap, however, the increase is not large. It is thus the extensive margin that would need to drive up labor costs—a larger number of job switchers—which our model captures through the higher observed quit rate.

[12]Adding interactions between own exposure E_i and the Act I and Act II indicators leaves the supplier-exposure coefficients largely unchanged. Own exposure is generally small and insignificant, with exceptions for prices in 2020 and wages during the recovery. Online Appendix Table 22 reports the results.

The exposure–inflation relationship also holds in second moments. The rebuilding mechanism predicts that sectors containing a more uneven distribution of exposed industries should experience a larger increase in price dispersion. For each BEA summary sector s, let $\sigma_s(\mathcal{E})$ denote the within-sector standard deviation of supplier exposure, and let $\Delta\sigma_s^+$ denote the change from the pre-COVID to the COVID period in the employment-weighted within-sector dispersion of cumulative upside inflation, $\max\{0, \Delta \log P_i\}$. We estimate the complementary cross-sectional regression $\Delta\sigma_s^+ = \alpha + \beta_\sigma\, \sigma_s(\mathcal{E}) + \varepsilon_s$ on the 38 BEA summary sectors that remain after excluding the EXO sectors and sectors not covered by the JEM.

Table 2: Fact 1: within-sector inflation dispersion on exposure dispersion.

Subsample	$\hat{\beta}\,(t)$	SE	p	R^2	Obs. (sectors)
ALL	$+87.63^{***}$ (+4.41)	19.86	0.000	0.475	38
m+l-tier	$+105.83^{***}$ (+2.98)	35.46	0.003	0.352	30

Notes: Cross-section of BEA-66 sectors with at least two endogenous industries per subsample. Dependent variable $\Delta\sigma_s^+$ and regressor $\sigma(\mathcal{E}_i)_s$ are defined in §2.2 (Fact 1). Subsamples defined in §2. WLS with sector-aggregated 2017 value-added weights; HC1 robust standard errors. Stars: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

As shown in Table 2, the estimated slope is positive and statistically significant at the 1% level. Exposure dispersion explains approximately 48% of the cross-sector variation in inflation dispersion in the full sample and 35% in the upstream and middle-tier sample. Thus, not only the average level of exposure but also its distribution across industries predicts subsequent price outcomes.

The Fact 1 decomposition attributes a substantial share of the exposure gradient in industry inflation to the *materials* line—the network-propagated cost of inputs from suppliers.[13] Because materials at industry i in window w decompose by identity as $(A\cdot g^P)_{i,w} = (A\cdot\text{Margin})_{i,w} + (A\cdot\text{ULC})_{i,w} + (A\cdot\text{Imports})_{i,w} + (A^2\cdot g^P)_{i,w}$, we can run separate panel-DiD regressions on each of these four supplier-side terms of decomposition—using the same specification as the main-text Fact 1—and by OLS linearity in the dependent variable their four coefficients partition the materials-total coefficient exactly.[14] Table 3 reports the result. In the initial disruption window (Act I, 2020–2021), the exposure gradient in materials is driven by *supplier margins* ($\hat{\beta} = +3.09$, $t = 4.18$); supplier ULC and supplier imports are effectively zero. In the recovery window (Act II, 2022–2023), the

[13] A Shapley R^2 decomposition across regressor blocks, reported in Online Appendix Table 13, confirms that predetermined exposure explains a large share of the materials component—the biggest among the four additive components of $\Delta \log P_i$.

[14] The partition is an algebraic property of OLS: because the four LHS variables sum to $(A \cdot g^P)_{i,w}$ by identity and the RHS design is common across regressions, running OLS of each term on the exposure interactions separately and summing the coefficients returns the same numbers as running OLS of $(A \cdot g^P)_{i,w}$ directly.

pattern reverses: supplier margins have unwound and are no longer significant, while *supplier ULC* carries the exposure ($\hat{\beta} = +2.05$, $t = 2.51$). Inflation propagated through the network as margins first and ULC later.

Table 3: Fact 1: network-margin decomposition.

	Act I ($\mathcal{E}_i \cdot D_w^{\text{shock}}$)		Act II ($\mathcal{E}_i \cdot D_w^{\text{rec}}$)	
Decomposition term	$\hat{\beta}$	(t)	$\hat{\beta}$	(t)
A·Margin (supplier markups)	$+3.09^{***}$	$(+4.18)$	-0.93	(-1.05)
A·ULC (supplier wages)	-0.60	(-0.93)	$+2.05^{**}$	$(+2.51)$
A·Imports (supplier imports)	$+0.09^{**}$	$(+2.17)$	-0.00	(-0.16)
$A^2 \cdot g^P$ (supplier of supplier)	$+1.26^{***}$	$(+6.33)$	$+0.57^{***}$	$(+5.37)$
MAT total ($A \cdot g^P$, identity sum)	$+3.84^{***}$	$(+4.92)$	$+1.68^{***}$	$(+4.13)$

Notes: Each row reports a separate WLS-VA panel DiD of the LHS variable (per-industry, per-window mean of the indicated single-pass component pre-multiplied by A) on two exposure interactions $\mathcal{E}_i \cdot D_w^{\text{shock}}$ and $\mathcal{E}_i \cdot D_w^{\text{rec}}$, with industry additive and window fixed effects; standard errors clustered on industry. The four rows partition the MAT-total coefficient exactly by OLS linearity in the dependent variable. Significance stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Exposure operator $\mathcal{E}_i = (W \cdot E)_i$ with $W = A(I - \delta A)^{-1}$ at $\delta = 0.50$. Sample: 330 endogenous BEA-357 industries.

Fact 2. Conditional on exposure, faster recovery predicts less inflation. Fact 1 could in principle reflect demand forces: demand in exposed industries may have risen above and beyond its 2019 level, in which case the price response would owe nothing to the destruction of organizational capital in the COVID contraction. This pure sectoral demand-pull account, however, predicts that industries with faster quantity recovery should experience *more* inflation, as demand moves them along upward-sloping supply or marginal-cost curves. In contrast, the rebuilding mechanism predicts the opposite: conditional on the initial disruption, industries whose supplier networks recover more quickly face smaller capacity shortfalls and should therefore show *less* inflation.

To test this, we extend the baseline regression with a supplier-network recovery velocity. We construct it by applying the same network operator used for exposure to average annual industry employment growth after reopening:

$$\mathcal{V}^{\text{wall}} \equiv AL_\delta \, \overline{g}^{N}_{2020:2022}, \qquad \mathcal{V}^{\text{post}} \equiv AL_\delta \, \overline{g}^{N}_{2022:2024}. \tag{6}$$

The first measure covers the period over which the rebuilding constraint is estimated to bind; the second is a post-recovery placebo, which should show a waning effect. The

augmented regression equation is

$$y_{i,w} = \gamma_i + \theta_w + \beta_{\mathrm{I}}^{\mathcal{E}} \mathcal{E}_i D_w^{\mathrm{I}} + \beta_{\mathrm{II}}^{\mathcal{E}} \mathcal{E}_i D_w^{\mathrm{II}} + \beta_{\mathrm{I}}^{\mathcal{V}} \mathcal{V}_i^{\mathrm{wall}} D_w^{\mathrm{I}} + \beta_{\mathrm{II}}^{\mathcal{V}} \mathcal{V}_i^{\mathrm{post}} D_w^{\mathrm{II}} + \varepsilon_{i,w}. \tag{7}$$

The discriminating coefficient is $\beta_{\mathrm{I}}^{\mathcal{V}}$: negative under the rebuilding mechanism, positive under demand pull.

Table 4: Fact 2: industry inflation on exposure $\mathcal{E}_i$ and recovery velocity $\mathcal{V}_i$.

	$\hat{\beta}^{\mathcal{E}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{E}}_{\text{Act II}}$	$\hat{\beta}^{\mathcal{V}^{\text{wall}}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}^{\text{post}}}_{\text{Act II}}$
$\Delta \log P_i$	$+12.54^{**}$	$+5.37$	-0.49	$+0.75$
	$(+2.48)$	$(+0.90)$	(-0.79)	$(+0.51)$
ΔMARGIN_i	$+12.60$	-21.17	-0.82	$+4.48$
	$(+1.43)$	(-1.56)	(-0.89)	$(+1.25)$
ΔULC_i	-5.42	$+23.58^{*}$	$+0.85$	-3.73
	(-0.74)	$(+1.89)$	$(+0.91)$	(-1.11)
ΔIMP_i	$+0.32^{*}$	$+0.20^{***}$	-0.07^{***}	$+0.04$
	$(+1.95)$	$(+2.68)$	(-3.48)	$(+1.61)$
ΔMAT_i	$+5.04^{***}$	$+2.76^{***}$	-0.46^{***}	-0.04
	$(+5.79)$	$(+5.72)$	(-3.80)	(-0.19)
$A \cdot \Delta\text{MARGIN}_i$	$+3.41^{***}$	-0.66	-0.34^{***}	-0.28
	$(+5.35)$	(-0.68)	(-4.91)	(-1.07)

Notes: Industry-level panel over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, Act II 2021–2023. WLS with 2017 industry value added; industry and window fixed effects; standard errors clustered on industry. $\mathcal{V}_i^{\text{wall}}$ and $\mathcal{V}_i^{\text{post}}$ are network-propagated employment growth over windows 2020–2022 and 2022–2024 respectively, defined in §2.2 eq. (6); each interacts only with its own-window dummy ($\mathcal{V}^{\text{wall}}$ with Act I, $\mathcal{V}^{\text{post}}$ with Act II). Stars: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. $N = 330$ industries.

Table 4 shows the results. The Act I recovery coefficient is negative and significant for prices, and similarly so across the price-side accounting components. It is particularly precise for domestic materials inflation and supplier-margin pass-through—of special interest because these capture the network transmission of prices. By contrast, the post-recovery coefficients are small and generally insignificant. These estimates do not rule out an important role for aggregate demand in the recovery, but they are inconsistent with a pure demand-pull explanation of the cross-section.

2.2.1 Unconditional evidence

We conclude this section by asking whether the industry-level forces documented above are large enough to account for a meaningful share of aggregate inflation. To answer this question, we aggregate the industry accounting identity and decompose observed

Table 5: Aggregate decomposition of U.S. inflation, 2017–2023.

	Placebo 2017–2019	Act I 2019–2021	Act II 2021–2023
GO inflation	+4.4	+7.1	+9.7
– Government	+0.2	+0.4	+0.4
– Oil/petroleum/metals	+0.1	+0.3	+0.2
– Real estate	+0.9	+0.8	+1.6
– Finance/insurance	+0.8	+0.9	+1.1
N-GO inflation	**+2.7**	**+5.2**	**+6.7**
Panel A: Static decomposition			
Margins	+0.4 (15%)	+2.4 (45%)	+0.2 (4%)
ULC	+1.4 (52%)	+1.4 (26%)	+4.2 (62%)
+ Wages	+3.0	+5.5	+3.2
– Productivity	−1.6	−4.2	+1.0
Imports	+0.0 (1%)	+0.1 (2%)	+0.1 (1%)
Materials (MAT)	+0.9 (33%)	+1.4 (26%)	+2.2 (32%)
Panel B: Network decomposition			
Margins	+0.2 (9%)	+3.3 (63%)	+0.2 (3%)
ULC	+2.3 (85%)	+1.6 (31%)	+6.1 (91%)
+ Wages	+4.3	+7.9	+4.5
– Productivity	−2.0	−6.3	+1.6
Imports	+0.1 (2%)	+0.2 (4%)	+0.2 (2%)

Notes: Entries in percentage points; bracketed values are share of N-GO inflation. Panel A (Static) treats domestic-materials inflation as a separate component. Panel B (Network) reallocates materials to their upstream margin, unit-labor-cost, and import-price origins across the production network via $L = (I - A)^{-1}$.

inflation into the contributions of unit labor costs, domestic and imported intermediate inputs, and the margin residual. This is an unconditional accounting exercise: it measures the aggregate contribution of each component without conditioning on industries' COVID exposure.

Let c denote the fixed 2017 PCE final-demand weight vector in the BEA Use table, expressed in producer prices. For our non-EXO gross-output measure, the weights corresponding to the four EXO blocks are set to zero, and the remaining weights are renormalized to sum to one. Mapping commodity weights into industry weights gives the industry PCE weight vector $g \equiv \Theta c$, so aggregate inflation is $\Pi = g^\top \Delta \log P$. Premultiplying equation (3) by $(\Theta c)^\top$ yields the *Static* decomposition

$$\Pi = \underbrace{(\Theta c)^\top \Delta \text{MARGIN}}_{\Pi_{\text{MARGIN}}} + \underbrace{(\Theta c)^\top \Delta \text{ULC}}_{\Pi_{\text{ULC}}} + \underbrace{(\Theta c)^\top \Delta \text{MAT}}_{\Pi_{\text{MAT}}} + \underbrace{(\Theta c)^\top \Delta \text{IMP}}_{\Pi_{\text{IMP}}}, \tag{8}$$

where unit-labor-cost decomposes to $\Pi_{\text{ULC}} = \Pi_{\text{WAGE}} + \Pi_{\text{PROD}}$, and where $\Pi_{\text{PROD}} \equiv (\Theta c)^{\top}[-\Delta \log(Y/N)]$. Unroll the domestic input–output recursion implicitly implied by the dependence of the cost of materials on prices of commodities, we obtain the *network-consistent* recursive decomposition (network decomposition, hereafter):

$$\Pi = (\Theta c)^{\top} L \left(\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}\right), \qquad L \equiv (I - A)^{-1}. \tag{9}$$

Here $(\Theta c)^{\top} L$ is the vector of industry Domar weights; equivalently, $\kappa \equiv L^{\top} g$, with κ_i denoting industry i's Domar weight. The network decomposition allocates the cost of materials to the components that are exogenous from the point of view of the I–O network.

Table 5 reports the above decompositions applied to raw sectoral inflation rates. It confirms the key finding implied by the regressions: during the first phase of inflation, labeled **Act I**—2019–2021, when N-GO inflation rose by 5.2 percentage points—the decomposition assigns a large role to the measured margin residual (Static: 2.4 pp, 45%; Network: 3.3 pp, 63%). Unit labor costs contributed as much as before the pandemic (Static: 1.4 pp; Network: 1.6 pp)—not because nominal wages were flat, but because of the sharp rise in measured labor productivity. This early increase in margins and the productivity surge are consistent with firm-level evidence (Glover et al., 2023). As noted, the productivity increase largely reflects composition effects from employment losses concentrated among younger and less experienced workers, as documented by the literature (Fernald and Li, 2022; Stewart, 2022; Crust et al., 2020). During **Act II**—2021–2023, when N-GO inflation rose by 6.7 percentage points—the accounting reverses: unit labor costs dominate (Static: 4.2 pp, 62%; Network: 6.1 pp, 91%). This large contribution, however, only partly owes to nominal wages. The key factor is stagnating productivity growth.

3 Model

This section develops a quantitative model to ask whether the destruction of workforce-embodied organizational capital during COVID was sufficiently large and uneven to account for the inflation observed during the recovery. The model's key state variable is each industry's workforce capacity at reopening. When employment expands beyond this capacity, firms incur convex rebuilding costs, creating a capacity wall that raises marginal costs. The input–output network propagates these localized cost pressures across industries, while nominal rigidities govern their transmission into prices and wages.

Time is discrete. The economy is initially in its pre-COVID steady state at $t = 0$, corresponding to 2019. The unanticipated pandemic contraction occurs at $t = 1$, corresponding

to 2020, and leaves industries with different workforce capacities at reopening. The recovery begins at $t = 2$, with $t = 2, 3, 4$ corresponding to 2021–2023. To accommodate movements in quantities and prices that are not generated by the cost-wall mechanism, the model includes four reduced-form wedges: $a^e_{i,t}$, $\xi^e_{i,t}$, τ^e_{ij}, and $\mu^e_{i,t}$. These wedges allow the model to account for other forces shaping the sectoral recovery in the data.

We first describe the production network and then the cost-wall mechanism. We then introduce households, wage and price setting, and the equilibrium conditions. Section 4 explains how the model's parameters, workforce capacities, and production network are disciplined by the data.

3.1 Environment and production network

The economy consists of twenty-four production nodes organized into twelve chain-by-tier composites, indexed by $i = 1, \ldots, 12$. Each composite belongs to one of four supply chains and occupies an upstream, middle, or consumer-facing position within that chain. We further divide each composite into two nodes, indexed by $j \in \{H, S\}$, according to the pandemic employment contractions experienced by their constituent industries. We refer to these as hit (H) and safe (S) nodes. This split preserves dispersion in workforce capacity among industries occupying similar positions in the production network.

Outside the endogenous network is an exogenous (EXO) block containing imported, energy, and housing-related inputs. EXO prices are taken as given, while firms choose the quantities of EXO inputs they purchase. Section 4 describes the empirical construction of the nodes, the EXO block, and the corresponding input–output coefficients.

Variables indexed by (i, j) refer to individual hit or safe nodes, variables indexed only by i refer to tier-level aggregates, and variables without either index refer to economywide aggregates. For parameters, omission of j indicates that the parameter is common to the hit and safe nodes within tier i.

Each node comprises a fixed measure of monopolistically competitive firms producing differentiated varieties and charging positive markups. Aggregate output of node (i, j) is produced according to

$$Y_{ij,t} = \exp(a^e_{i,t}) Z_{ij} V^{1-\alpha_i}_{ij,t}, \tag{10}$$

where Z_{ij} absorbs cross-sectional differences in productivity and capital, α_i is the capital share embedded in the reduced-form production technology, and $a^e_{i,t}$ is a time-varying effective productivity wedge. The variable-input composite combines labor $N_{ij,t}$ and

intermediate inputs $m_{ij,t}$:

$$V_{ij,t} = \left[\eta_{ij}^{1/\rho} N_{ij,t}^{(\rho-1)/\rho} + (1-\eta_{ij})^{1/\rho} m_{ij,t}^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)}, \tag{11}$$

where η_{ij} is the labor weight and ρ is the elasticity of substitution between labor and intermediate inputs. For tractability, the same elasticity governs the materials-input margin and substitution between hit and safe nodes within a tier. Within each tier, hit and safe output are aggregated according to

$$Y_{i,t} = \left[\sum_{j\in\{H,S\}} s_{ij}^{1/\rho} Y_{ij,t}^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)}. \tag{12}$$

Let $P_{i,t}$ denote the corresponding price index. The resulting tier composite can be purchased by final users or used as an intermediate input elsewhere in the network. The intermediate-input bundle used by node (i,j) combines inputs from the twelve endogenous tiers and the EXO block:

$$m_{ij,t} = \left[\sum_{k=1}^{12} \omega_{i,k}^{1/\rho} m_{ij,k,t}^{(\rho-1)/\rho} + (\omega_i^{EXO})^{1/\rho} (m_{ij,t}^{EXO})^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)} \tag{13}$$

where $m_{ij,k,t}$ denotes intermediates purchased by node (i,j) from endogenous supplier tier k, while $m_{ij,t}^{EXO}$ denotes its purchases from the EXO block. The corresponding unit price of the intermediate-input bundle is

$$p_{i,t}^{m} = \left[\sum_{k=1}^{12} \omega_{i,k} P_{k,t}^{1-\rho} + \omega_i^{EXO} \left(P_{i,t}^{EXO}\right)^{1-\rho}\right]^{1/(1-\rho)}, \tag{14}$$

where $P_{i,t}^{EXO}$ is the price index of the EXO inputs purchased by tier i, and $\sum_{k=1}^{12} \omega_{i,k} + \omega_i^{EXO} = 1$. Thus, $p_{i,t}^{m}$ summarizes the input-cost pressures faced by tier i. An increase in $P_{k,t}$ raises the production costs of its customers according to the input–output weight $\omega_{i,k}$, propagating the initial cost pressure through the network. Appendix A.4 reports the associated conditional input demands and price indices.

3.2 Capacity wall (CW)

Employment inherited from the previous period acts as workforce capacity. A firm can employ workers up to this capacity at the prevailing wage. Expanding beyond it re-

quires recruiting, screening, training, and reorganizing production. We capture these increasingly costly activities with a *quadratic* rebuilding-cost in employment above inherited capacity. We refer to this capacity threshold as the *capacity wall* (CW). For node (i, j), define employment above inherited workforce capacity, normalized by steady-state employment, as $dN^{+}_{ij,t} = \max\{(N_{ij,t} - N^{\text{CW}}_{ij,t})/N^{ss}_{ij}, 0\}$, and the associated rebuilding-cost premium as $\mathcal{W}_{ij,t} = \gamma_i \left(dN^{+}_{ij,t}\right)^2$, where γ_i governs the curvature and steepness of the capacity wall. Employment below $N^{\text{CW}}_{ij,t}$ generates no rebuilding premium. Once employment exceeds this threshold, the premium rises quadratically with the normalized amount of employment added beyond inherited capacity.

The initial reopening was unusual; however, many workers separated from firms in 2020 remained available for recall or rapid rehiring. We capture this *warm pool* by setting

$$N^{\text{CW}}_{ij,t} = \begin{cases} N_{ij,1} + \psi \max\left\{N^{ss}_{ij} - N_{ij,1},\, 0\right\}, & t = 2, \\ N_{ij,t-1}, & \text{otherwise}, \end{cases} \tag{15}$$

where $t = 1$ denotes 2020, $t = 2$ denotes 2021, N^{ss} denotes the employment level in 2019, and $\psi \in [0, 1]$ is the share of the 2020 employment shortfall that can be restored without incurring the additional rebuilding premium. When $\psi = 0$, all employment above its 2020 level is costly to rebuild; when $\psi = 1$, employment can return to its steady-state level before the wall binds.

The capacity wall raises the effective unit cost of labor:

$$w^{e}_{ij,t} = \tilde{w}_{ij} W_t \left(1 + \mathcal{W}_{ij,t}\right) \exp\left(1_{t=1} \tau^{e}_{ij}\right), \tag{16}$$

where W_t is the aggregate base wage, $\tilde{w}_{ij}$ captures cross-sectional wage differences, and τ^{e}_{ij} is a transitory COVID labor wedge that operates only in t=1 (2020) . Thus, two nodes paying the same base wage may face different effective labor costs during the recovery: the node that suffered the larger employment loss reaches its inherited capacity sooner and incurs a larger rebuilding premium.

Recruiting and training expenditures are investments in workforce-specific organizational capital whose benefits extend over the duration of the employment match. We therefore calibrate γ_i using rebuilding costs amortized over the expected match duration, as described in Section 4.

The capacity wall affects inflation through marginal cost. Combining the production

technology CES with cost minimization gives

$$MC_{ij,t} = \frac{\left[\eta_{ij}\left(w^{e}_{ij,t}\right)^{1-\rho} + (1-\eta_{ij})\left(p^{m}_{i,t}\right)^{1-\rho}\right]^{1/(1-\rho)}}{P_{ij,t}(1-\alpha_i)} \frac{V_{ij,t}}{Y_{ij,t}}. \tag{17}$$

Consequently, rapid employment rebuilding raises effective labor costs and marginal costs. The production network then transmits these pressures to downstream firms through higher intermediate-input prices.

3.3 Prices and expectations

Prices and base wages are sticky, and inflation expectations are strongly anchored, in the sense explained below.

Sticky prices. Each production node (i, j) comprises a sector of symmetric monopolistically competitive producers facing constant-elasticity demand with a within-node elasticity of $\varepsilon_i > 1$. Firms set prices subject to a Rotemberg-style quadratic adjustment cost relative to the nominal anchor $\bar{\Pi}$, plus an asymmetric additional term that activates only when inflation falls below that anchor. Firms form expectations of next-period inflation through the subjective operator defined below. The full derivation of the pricing equation from the firm's problem is provided in Online Appendix A.4.2.

For node (i, j), let $\Psi_i(\Pi) \equiv \Pi\left[\bar{\phi}_i(\Pi - \bar{\Pi}) + \lambda_i(\Pi - \bar{\Pi})^{-}\right]$ denote the asymmetric Rotemberg marginal adjustment-cost expression evaluated at inflation Π. The pricing first-order condition in per-unit-output form then reads

$$\Psi_i(\Pi_{ij,t}) = \varepsilon_i\left(MC_{ij,t} - \tfrac{\varepsilon_i - 1}{\varepsilon_i} + \mu^{e}_{i,t} + 1_{t=1}\bar{\mu}^{e}_{i}\right) + \beta\,\mathbb{E}_t\left[\Psi_i(\Pi_{ij,t+1})\,\tfrac{Y_{ij,t+1}}{Y_{ij,t}}\right], \tag{18}$$

where $\Pi_{ij,t} \equiv P_{ij,t}/P_{ij,t-1}$ is gross inflation and $(\cdot)^{-} \equiv \min\{0, \cdot\}$ captures asymmetric downward price rigidity. The parameter $\bar{\phi}_i$ governs the symmetric Rotemberg adjustment cost, λ_i governs downward rigidity, $\mu^{e}_{i,t}$ is the residual markup wedge, common to both nodes within tier i, and β is the annual discount factor, set to $\beta = 0.96$. The demand elasticity is set to a common value ε_i across tiers, implying a steady-state markup over the gross-output price equal to $\varepsilon_i/(\varepsilon_i - 1) - 1$. (In the calibration, $\bar{\phi}_i$ is chain-specific in the h tiers and equals $\bar{\phi}_{ml}$ in the m and l tiers; similarly, $\lambda_i = \lambda_h$ in the h tiers and $\lambda_i = \lambda_{ml}$ in the m and l tiers.)

The expectations operator in (18) is adaptive with respect to aggregate inflation, which

keeps model-implied inflation expectations consistent with the data (Figure 1, Panel A). Specifically, we parameterize the expectations operator as follows, assuming that firms retain perfect foresight over firm-level quantities:

$$\mathbb{E}_t\left[\Psi_i(\Pi_{ij,t+1})\,\frac{Y_{ij,t+1}}{Y_{ij,t}}\right] \equiv \Psi_i(\Pi^e_{t+1|t})\,\frac{Y_{ij,t+1}}{Y_{ij,t}}. \tag{19}$$

Where $\Pi^e_{t+1|t}$ is the convex-combination forecast defined below, while $Y_{ij,t+1}/Y_{ij,t}$ is extracted from the model's perfect-foresight transition path. This separability matters little quantitatively; it is equivalent to assuming a log-linear expectations operator.

The pricing equation carries two additive markup wedges, strictly separating the 2020 collapse from the 2021–2024 recovery. The 2020 wedge $\bar{\mu}^e_i$ is active in the collapse year only ($1_{t=1}\bar{\mu}^e_i$). The recovery wedge $\mu^e_{i,t}$ closes the price residual that remains after the structural components have acted. This distinction is central to the later accounting: to isolate the model's endogenous contribution, we zero out the recovery wedge ($\mu^e_{i,t} = 0$) and re-solve the model while keeping the structural 2020 wedge $\bar{\mu}^e_i$ in place. Retaining $\bar{\mu}^e_i$ keeps the model consistent with the relatively mild deflation of the COVID year 2020, which it would otherwise fail to match.[15]

Base wage. The base wage is sticky and follows a reduced-form wage Phillips curve:

$$\frac{W_t}{W_{t-1}} = \bar{\Pi}\left(\frac{N_t}{N^{ss}}\right)^{\phi_w}\left(\frac{LS^{ss}}{LS_{t-1}}\right)^{\iota_a}\left(\prod_{k=1}^{n}\frac{\Pi_{t-k}}{\bar{\Pi}}\right)^{\iota_w}, \tag{20}$$

where ϕ_w governs the sensitivity of base-wage growth to aggregate employment, and ι_w governs backward indexation to cumulative inflation deviations. Taking logs, the indexation term is $\iota_w \sum_{k=1}^{n}(\log \Pi_{t-k} - \log \bar{\Pi})$, equivalent to the multiplicative product above. We set $n = 4$, so current wage growth responds to inflation over the preceding four years; this multi-year window allows the real wage to track cumulative inflation while keeping per-period pass-through small. Our baseline model implies a fairly flat wage Phillips curve in response to moderate shocks that are "even" across sectors; indeed, our results would be almost unchanged under ϕ_w close to zero.

The middle term in equation (20) serves as a technical long-run anchor on the aggregate labor share $LS_t \equiv W_t N_t/(P_t \mathcal{Y}_t)$, utilizing a negligible weight $\iota_a = 0.05$ that pulls LS_t

[15]The attribution of the CW to inflation relative to the model's floor is almost identical without this wedge, but the model's implied floor inflation path shifts down due to severe deflation in 2020. The deflation stems from a massive productivity surge in 2020, which the a^e wedge accounts for by construction. Under this assumption, the productivity surge kicks in in 2021, when the CW begins to bind, and offsets its inflationary effect from 2021 onwards.

toward its steady-state value LS^{ss} — small enough that the anchor has negligible effect over 2019–2024 while ensuring convergence to the terminal steady state. By the firm's first-order condition at steady state, the real wage equals the marginal product of labor, so $LS^{ss} = \text{MPL}^{ss}\, N^{ss}/\mathcal{Y}^{ss}$. This technical anchor pulls W/P toward MPL without requiring an explicit analytical expression for MPL within our nested-CES production network, and it stabilizes the model's convergence to the terminal steady state.[16]

Aggregate employment is $N_t = \sum_{i=1}^{12} \sum_{j\in\{H,S\}} N_{ij,t}$, with $N^{ss} = \sum_{i=1}^{12} \sum_{j\in\{H,S\}} N_{ij}^{ss}$. Aggregate PCE-weighted gross inflation is $\Pi_t = P_t/P_{t-1}$, where we report $\log \Pi_t \times 100\%$ as annual inflation in percentage points, and $\sum_{s=0}^{t} \log \Pi_s \times 100\%$ as cumulative inflation by year t.

Paid wages. The labor wedge disconnects prices from average costs, generating a surplus for the firm. The split of this surplus between profits and wages is undetermined within the model. We parameterize paid wages through a wedge x^e that divides the surplus as follows:

$$w_{ij,t} = \tilde{w}_{ij}\, W_t \left(1 + x_{ij}^{e} \mathcal{W}_{ij,t}\right), \tag{21}$$

where $x_{ij}^{e} \in [0, 1]$ is the wall-to-wage leakage parameter. We take the conservative route of not attributing the exogenous markup wedge $\mu_{i,t}^{e}$ to wages, even though model–data comparisons suggest that doing so would help explain the greater persistence of wage inflation in the data after 2023. Our focus is on inflation between 2019 and 2023, and we prefer to keep the ingredients relevant to the core mechanism explicit.

Inflation expectations. We parameterize expectations as

$$\Pi_{t+1|t}^{e} = \nu_{\text{forward}}\, \Pi_{t+1}^{\text{model}} + \nu_{\text{past}}\, \Pi_t + \nu_{\text{target}}\, \bar{\Pi}, \qquad \nu_{\text{forward}} + \nu_{\text{past}} + \nu_{\text{target}} = 1, \tag{22}$$

where the weights ν_* sum to unity. The calibrated anchored-expectations regime places substantial weight on the long-run target ($\bar{\Pi}$). This three-component structure—forward-looking, backward-looking, and anchored—is the reduced-form counterpart of the theory of anchoring under central-bank communication developed in Drozd (2026), where the coefficient ν_{target} on the long-run target arises as the weight private agents place on the stated inflation objective relative to purely backward or forward updating. We also consider the perfect-foresight case ($\nu_{\text{forward}} = 1$) and report the results in Online Appendix A.9.2; that variant is at odds with the muted expectations response in Figure 1, Panel A.

[16] Because all wedges are held constant after 2024, labor productivity at the end of the shock period may differ from its initial steady-state level.

Market clearing, final demand, and MPC. For each tier $k = 1, \ldots, 12$, output equals final demand plus intermediate demand from all nodes:

$$Y_{k,t} = D_{k,t}(\mathcal{Y}_t) + \sum_{i=1}^{12} \sum_{j \in \{H,S\}} \omega_{i,k} \left(\frac{P_{k,t}}{p_{i,t}^m} \right)^{-\rho} m_{ij,t}. \tag{23}$$

In consumer-facing tiers, final demand primarily represents personal consumption expenditures; in middle and upstream tiers, it may also include investment, government consumption, exports, and inventory accumulation. We do not distinguish among final-demand categories or model a household sector explicitly. Instead, we parameterize final demand using a time-varying wedge and a marginal-propensity-to-consume (MPC) parameter that generates an endogenous feedback loop from aggregate income back to sectoral demand—the standard Keynesian multiplier channel:

$$D_{k,t}(\mathcal{Y}_t) = D_k^{ss} \exp\left(\xi_{k,t}^e + \text{MPC} \cdot \left(\log \mathcal{Y}_t - \log \mathcal{Y}_t^{\text{data}} \right) \right), \tag{24}$$

where $\xi_{k,t}^e$ is a saturating final-demand shifter (Section 4), $\mathcal{Y}_t^{\text{data}}$ is the observed real aggregate output path, and D_k^{ss} is the tier's 2017 final-demand level from the SUT—PCE plus government consumption, investment, exports, and inventory change, net of imports—held fixed throughout the transition. Real aggregate output is the gross-output sum weighted by each tier's steady-state final-demand share:

$$\mathcal{Y}_t = \sum_{k=1}^{12} \frac{D_k^{ss}}{Y_k^{ss}} Y_{k,t}, \tag{25}$$

where Y_k^{ss} is the steady-state level of tier-k gross output. By the SUT identity, D_k^{ss}/Y_k^{ss} is the fraction of tier k's 2017 gross output flowing to final users net of imports, so $\mathcal{Y}_t$ is equivalent to real GDP at fixed 2017 weights. The fully saturated model matches the data by construction, so $\mathcal{Y}_t = \mathcal{Y}_t^{\text{data}}$ and final demand reduces to $D_k^{ss} \exp(\xi_{k,t}^e)$. The MPC term activates only in counterfactuals that drop the saturating wedges, allowing the model to generate endogenous Keynesian demand propagation when we identify wedge attributions.

The EXO block. The EXO block contains imported, energy, and housing-related categories that lie outside the endogenous production network. Their prices are exogenous, but endogenous producers choose their purchases from the EXO block as part of the materials bundle, with tier-specific weights. Selected EXO categories also enter the aggregate

price index directly through their PCE shares; imported inputs are excluded from this direct component and affect the aggregate price index only through their influence on endogenous producers' costs. The fixed-weight aggregate price index combines the twelve endogenous tier composites with the direct PCE contribution of the EXO block:

$$P_t = \sum_{k=1}^{12} c_k^{\text{PCE}} P_{k,t} + c^{\text{PCE-EXO}} P_t^{\text{PCE-EXO}}, \qquad \sum_{k=1}^{12} c_k^{\text{PCE}} + c^{\text{PCE-EXO}} = 1. \tag{26}$$

In counterfactuals that remove EXO price shocks, both the tier-specific EXO input-price indices and the direct PCE-EXO price index are held at their steady-state values, producing the N-GO price index defined in Section 2.

Online Appendix A.4 presents the remaining price indices and conditional input demands used above.

3.4 Equilibrium

Given the initial steady state, the paths of exogenous EXO prices, and the reduced-form wedges, an equilibrium is a sequence of node-level prices, employment, production, and intermediate-input choices, together with the aggregate wage and inflation rate, such that: firms minimize production costs; workforce capacity and rebuilding costs evolve according to the capacity-wall equations; prices satisfy the pricing condition; the base wage and inflation expectations follow their respective laws of motion; and output markets clear throughout the production network.

4 Calibration

The model is annual, and its steady state is calibrated to 2019 values. All targets are expressed relative to trend at the chain level.[17] The model is solved as a nonlinear perfect-foresight transition, with the COVID shock unanticipated (an MIT shock) on impact. The perfect-foresight horizon runs through 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 take the model to the data, we follow the business-cycle-accounting methodology of Chari et al. (2007). We introduce exogenous wedges $a_{i,t}^e$, $\xi_{i,t}^e$, τ_{ij}^e, and $\mu_{i,t}^e$ that align

[17] To obtain a consistent comparison with the data, all targets are detrended at the *chain level* against a 2017–2024 trend, separately for output (real value added), employment, gross-output deflators, and wages. Online Appendix A.6 provides the details.

the equilibrium conditions with the observed sectoral paths of gross output, gross-output deflators, and employment. The resulting exact fit is an accounting inversion, not the basis of our quantitative attribution. To isolate the inflation generated by the endogenous rebuilding-cost mechanism, we set the residual markup wedge $\mu^e_{i,t}$ to zero and re-solve the model while retaining the other empirically recovered wedges. The resulting counterfactual price path measures the inflation generated by the observed recovery through the cost wall and the production network, absent residual markup movements.

The New Keynesian components 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 $\phi_w = 0.15$. This sits at the upper end of the range of identified post-2000 flat-Phillips-curve estimates surveyed in Furlanetto and Lepetit (2024) and discussed in Section 5.2, with Hazell et al. (2022) (quarterly slope $\kappa \approx 0.006$, ≈ 0.025 annualized) and Del Negro et al. (2020) (post-1990 slopes statistically close to zero) on the lower side of that range. Backward wage indexation is modest at $\iota_w = 0.20$ and applies to the four-year cumulative (sum-in-logs) inflation deviation $\sum_{k=1}^{n}(\log \Pi_{t-k} - \log \bar{\Pi})$ rather than to a single lag, so that wages substantially catch up with cumulative inflation over time. The price block has no backward indexation ($\delta_{\text{index}} = 0$).

Table 6 summarizes the key parameter values. Panel A lists externally imposed parameters; Panel B lists calibrated wedges and their target moments. The remainder of the section describes the details.

Network structure. The four chains and three tiers are functional aggregates, designed to maximize within-chain upstream-to-downstream resolution in a small system. To obtain the chains, we recursively split the BEA detail-level I–O industry graph into balanced halves that minimize trade across the cut, applying the spectral bisection twice to obtain four chains. Empirically, each chain sources 42–51% of its intermediates from within its own boundaries (C_1 47%, C_2 42%, C_3 41%, C_4 51%); for comparison, a random four-way partition would yield $\approx 25\%$.[18]

Within each chain, industries are then sorted by their *finalness* and split into three tiers. Finalness, defined for each BEA sector in the SUT, is the share of the sector's direct impact on PCE relative to its total Domar impact: $\text{finalness}_i \equiv g_i/\kappa_i$, where g_i is the sector's direct PCE weight (from Section 2) and κ_i its Domar weight. The consumer-facing tier h contains industries with finalness above 70 percent within each chain. The remaining industries are split at the gross-output median into a middle tier m (higher finalness) and

[18]The within-chain intermediate share for chain C is $\sum_{i\in C}\sum_{j\in C}\omega_{ij}\,\text{VA}_i \,/\, \sum_{i\in C}\text{VA}_i$, where ω_{ij} is industry j's share of industry i's intermediate input bill (rows of Ω sum to one).

Table 6: Calibration summary

Panel A: Imposed parameters			
ρ	CES substitution elasticity	0.39	KLEMS
γ_i	CW intensity	per tier	Text
c^{wall}	Average wall premium, 2021–22	0.13	Text
q_0, q_1, T	Quit rates (normal, tight); horizon	0.0225, 0.030, 36 mo	JOLTS
ψ	Warm-pool share	0.20	Text
$\bar{\phi}_{c,h}$	Rotemberg price-adjustment cost, h tier	0.6–4.7	PSW
$\bar{\phi}_{ml}$	Rotemberg price-adjustment cost, m, l tiers	2.86	PSW
$\lambda_{c,h}$, λ_{ml}	Downward-rigidity coefficient	1.9–14.1, 8.6	Text
ϕ_w	Wage Phillips slope	0.15	Text
ι_w	Wage indexation, per lag	0.20	Text
n	Wage indexation horizon (years)	4	Text
ι_a	Labor-share long-run anchor weight	0.05	Text
δ_{index}	Price indexation	0	Text
ε	Within-node demand elasticity	5	Text
β	Discount factor	0.96	Annual
α_i	Tier capital share	per tier	BEA 2017 SUT
η_{ij}	Node labor share	per node	BEA 2017 SUT
$\tilde{w}_{ij}$	Steady-state wage premium	per node	BEA 2017 SUT
$\omega_{i,k}$	Domestic I–O coefficients	12×12	BEA 2017 SUT
ω_i^{EXO}	EXO-input share, per tier	per tier	BEA 2017 SUT
D_k^{ss}/Y_k^{ss}	Tier final-demand share, net of imports	per tier	BEA 2017 SUT
Panel B: Calibrated objects			
$a_{i,t}^e$ (12×5)	Productivity wedge	Matches tier-level N path	
$\xi_{i,t}^e$ (12×5)	Final-demand wedge	Matches tier-level Y path	
$\bar{\tau}_i^e$ (12)	Tier-mean labor wedge, 2020	Matches tier-level N in 2020	
$\Delta\tau_c^e$ (4)	Hit/safe labor-wedge spread, 2020	Matches chain dispersion $\sigma^-(\Delta \log N)$, QCEW	
$\bar{\mu}_i^e$ (12)	Structural 2020 markup wedge	Matches tier-level P_{GO} in 2020	
x_{ij}^e (24)	Wage-margin attribution share	Matches node wage path, 2021–24	
$\mu_{i,t}^e$ (12×5)	Residual markup wedge	Matches tier-level P_{GO} path, 2021–24	

Notes: The table lists the model's primitive parameters and calibrated wedges; the c^{wall}-derivation intermediates (r, Q_0, Q_1, η, κ_{skill}) and the expectations weights (ν_{forward}, ν_{past}, ν_{target}) are discussed in the text. i indexes the 12 chain-by-tier composites and $j \in \{H, S\}$ indexes the hit/safe node. The CES elasticity ρ governs all three CES margins (labor–materials, intermediate-input bundle, and hit/safe). The within-node demand elasticity $\varepsilon = 5$ implies a steady-state markup $\varepsilon/(\varepsilon - 1) = 1.25$. The wall premium $\mathcal{W}_{ij,t} = \gamma_i (dN_{ij,t}^+)^2$ depends only on the current period. Targets for the calibrated wedges in Panel B are matched to the detrended series of Online Appendix A.6. Sensitivity to parameter perturbations ($\rho = 0.5$, $\sigma_w = 5$, $\rho_m = 0.5$, λ halved, c^{wall} reduced by 22%, $\psi = 0.30$, $\iota_w = 0$) is reported in Online Appendix A.9.

an upstream tier l (lower finalness).

Figure 4 illustrates the flow of value through the network. The diagram is *directional*: each path runs from a supplier node on the left to a buyer node on the right, with thickness proportional to the dollar value of intermediate trade between them (minor linkages are omitted for readability). As the figure shows, the four chains are highly self-contained.

In terms of industry content, the four chains organize into a clean typology, each with

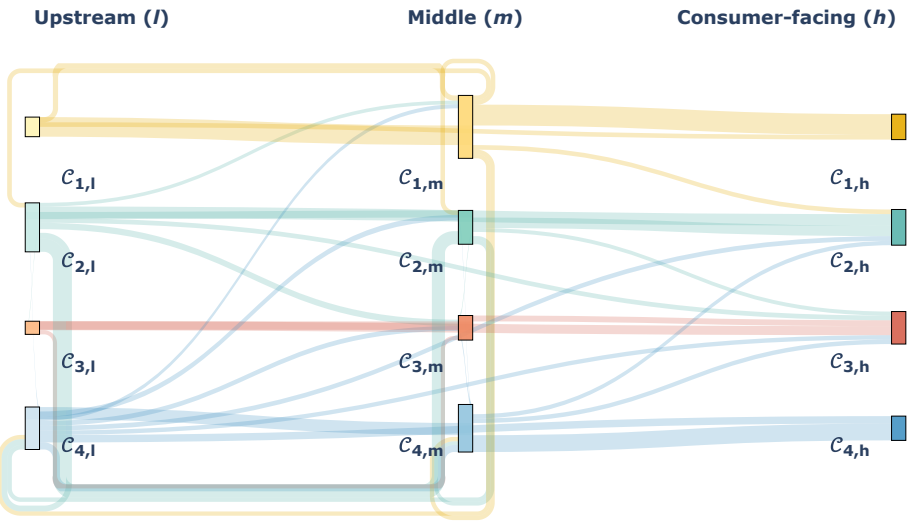

Figure 4: Intermediate-input flows across the 12 chain × tier nodes.

Notes: Directional Sankey diagram of intermediate-input trade across the 12 production cells (4 chains × 3 finalness tiers). Each ribbon runs from a supplier node on the left to a buyer node on the right, with thickness proportional to the dollar value of intermediate trade between them; minor linkages are omitted for readability. Ribbon color marks the chain of origin: C_1 blue, C_2 red, C_3 teal, C_4 gold. The four within-chain color blocks identify the self-contained chain structure; cross-color ribbons are cross-chain spillovers. Source: 2017 BEA Use and Make tables aggregated to the chain × tier classification of Appendix A.2.

an identifiable downstream anchor and a distinct class of upstream inputs. C_1 is the *professional / information-services chain*; C_2 is the *personal-services and food chain*; C_3 is the *health-and-hospitality chain, chemistry-fed*; and C_4 is the *manufacturing-and-construction chain*. We take these one at a time.

C_1 (35% of endogenous value added; 94% services, 2% durable manufacturing) is the professional / information-services chain. It links retail trade and telecom downstream industries to professional and business-to-business services upstream (custom programming, management consulting, legal services, advertising). Its B2B share is very large, with around 59% of its gross output sold as intermediates to other chains. Its h-tier (4 industries) carries 14.1% of total direct PCE share—81% of the chain's total direct PCE share.[19]

C_2 (26% of non-EXO value added (VA); 73% services, 8% nondurable food manufacturing, 5% agriculture) is the personal-services and food chain. It connects hospitals, restaurants, and universities downstream to grocery wholesale, food processors, truck freight, and primary agriculture upstream. Its defining feature is a long, physical supply

[19]Within-chain VA shares (h/m/l): C_1 31%/41%/28%; C_2 65%/18%/17% (top-loaded); C_3 36%/37%/27%; C_4 12%/41%/47% (upstream-heavy).

chain feeding labor-intensive consumer services. The h-tier (49 industries) covers 27.5% of PCE—the heaviest consumer-facing block, with the largest contributors being hospitals (6.2% PCE), restaurants (4.7% PCE combined), nursing and residential care (1.4%), education (1.1%), dentists (1.0%), and family services (0.9%).

C_3 (21% of VA; 63% services, 14% nondurable chemicals/pharma manufacturing, 9% utilities) is the health-and-hospitality chain, fed by chemistry. It connects physician offices, outpatient care, accommodation, and air transportation downstream to drug wholesale, basic chemicals, plastics, and engineering services upstream. Pharmaceutical and chemical manufacturing alone accounts for roughly one-sixth of the chain's value added. Around 52% of its gross output is sold as intermediates—meaningfully B2B, though less so than C_1 (59%). The h-tier (26 industries) carries 12.0% of direct PCE, anchored by physician offices (3.7% PCE), pharmaceutical preparations (1.5% PCE), accommodation (1.2% PCE), outpatient care (1.2% PCE), and air transportation (1.0% PCE).

Finally, C_4 (18% of VA; 31% durable manufacturing, 39% construction—dominated by the single CNSTRCT aggregate—18% wholesale, 2% nondurable) is the manufacturing-and-construction chain. It links auto manufacturing and communication-equipment finals downstream to construction, machinery wholesale, and primary metals upstream. Its h-tier (23 industries) carries only 3.8% of direct PCE share—reflecting the small share of durable goods in PCE. Its final sales are nonetheless larger overall, since much of them are investment goods. Structurally the chain is upstream-heavy: 47% of its value added sits in the l-tier (construction, primary metals, heavy machinery), with only 12% in h (auto repair, AutoMan, household appliances, computer and electronics finals).[20] This completes the functional block of the network.

Production. The same elasticity ρ governs the labor–materials margin, the materials-input margin, and the substitution between hit and safe nodes within a tier. Following the estimate in Section 2, we set $\rho = 0.39$. As shown there, to first order this simplification is innocuous, and the approximation error is small for plausible values of these elasticities (below unity).

Warm pool. An unusually large share of pandemic layoffs, and even of business closures, were temporary in 2020. We account for this in two ways. First, we measure the employment decline between January–February 2020 and January–February 2021, which nets out temporary layoffs already recalled and enforces consistent employment measurement

[20] See Online Appendix A.2 for more details.

in subsequent years. (To construct labor productivity Y/N, we continue to use QCEW average annual employment, since gross output is measured annually.)

Second, we account for the recalls still outstanding in early 2021. According to BLS data, in January 2021 there were 10.1 million unemployed, versus 5.8 million in January 2020, and 2.7 million workers on temporary layoff, versus 0.65 million in January 2020. Taking January 2020 to represent a slightly overheated labor market, the excess unemployment pool was $4.3 = 10.1 - 5.8$ million; of this, excess temporary layoffs ($2.1 = 2.7 - 0.65$ million) constituted about 50 percent. The recall rate from temporary layoffs is estimated in the search literature at 30–50 percent for a typical recession. Because many recalls had already occurred, we assume a mid-range recall rate of 40 percent, which gives a warm-pool share of $\psi = 20\%$.[21]

CW slope. Calibrating the CW slope is the key step. To avoid overstating the results, whenever the alternatives are ambiguous we opt for the more conservative one. As shown in the Online Appendix, the sensitivity of the results to this parameter is roughly proportional to its value, so extrapolation to alternative values is straightforward.

Let c^{wall} denote the average wall wedge over the peak rebuilding window 2021–2022, expressed as a fraction of the base wage W in the model. We choose γ_i so that, at the sectoral level, this average wedge equals the target c^{wall} in every tier:

$$\gamma_i \;=\; \frac{c^{\text{wall}}}{\overline{(dN^+)^2}_i}, \qquad \overline{(dN^+)^2}_i \;=\; \frac{\sum_{t\in\mathcal{T}} \left(dN^+_{i,t}\right)^2 \mathbf{1}\{dN^+_{i,t} > 0\}}{\sum_{t\in\mathcal{T}} \mathbf{1}\{dN^+_{i,t} > 0\}}, \tag{27}$$

where i denotes the tier and $\mathcal{T} = \{2021, 2022\}$ collects the peak rebuilding years. This implies that, over the peak rebuilding window, a firm receiving a marginal order prices it to cover the labor-rebuilding cost required to produce it, c^{wall}.

To find c^{wall}, we use evidence from the labor literature on the convexity of hiring premia in an extraordinarily tight labor market.[22] The key pieces of evidence come from Linckh et al. (2025), Muehlemann and Strupler Leiser (2018), and Muehlemann and Pfeifer (2016). In particular, Linckh et al. (2025) estimate that a one-standard-deviation increase

[21] The standard assumed recall rate is between 30% and 50% in the labor literature. Many recalls took place in 2020; our baseline recall rate is 40 percent. This calculation uses BLS data available at https://www.bls.gov/opub/ted/2021/temporary-layoffs-remain-high-following-unprecedented-surge-in-early-2020.htm (BLS, *The Economics Daily*: "Temporary layoffs remain high following unprecedented surge in early 2020," February 10, 2021). These data are from the CPS and are seasonally adjusted. Additional evidence on recalls using UI data (for the state of California) can be found in California Policy Lab (2021).

[22] 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 γ_i therefore requires no markup, labor-share, or MPL adjustment in our framework.

in vacancy-to-unemployment tightness in the local labor market raises post-match hiring costs by roughly 10 to 12% for the largest cost components (adaptation and disruption time). We take this number as the baseline and observe that the U.S. V/U ratio rose from about 1.2 in late 2019 to about 2.0 at its 2022 peak, an increase of roughly three standard deviations relative to its historical distribution (Domash and Summers, 2022).[23]

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. We adopt the Linckh et al. (2025) 22-week baseline as our normal-times per-hire cost and set $Q_0 = 22/52 \approx 0.423$ for a normalized unit wage; the shorter Muehlemann and Strupler Leiser (2018) baseline delivers a lower c^{wall}, and we report it as a robustness variant in Online Appendix A.9. Applying the per-SD tightness premium of Linckh et al. (2025) to the three-standard-deviation U.S. tightening from 2019 to the 2022 peak gives a tight-market premium of $1 + 3 \times 0.12 = 1.36$ over Q_0, or $Q_1 = 30/52 \approx 0.577$—about eight additional weeks of wages on top of the baseline. Regional evidence in Gitti (2026), showing that new-hire wages respond to tightness only above the labor-shortage threshold, is consistent with this convexity.

A hire is capital to the firm: it pays Q once, then draws the worker's services until the match dissolves at hazard $\delta(q)$—the classic quasi-fixed-factor view of labor (Oi, 1962), treated as an investment priced into firm value as in Merz and Yashiv (2007). We assume that firms price in the user cost of one unit of "hire capital," which is $(r + \delta(q))Q$ per year, where $\delta(q) = 12[-\log(1 - q)]$ is the annualized separation hazard implied by the monthly quit rate q and r is the firm's real cost of capital. We use $r = 0.06$, consistent with a standard real-equity-premium calibration; we work with a before-financing measure, so this parameter is not tied to the household discount rate β elsewhere in the model. We evaluate at $q_0 = 0.0225$ (2017–2019 JOLTS average) and $q_1 = 0.030$ (recovery quit-rate peak), which imply $\delta(q_0) = 0.273$ and $\delta(q_1) = 0.366$.

A second, distinct source of convexity comes from firm-side hiring technology: with

[23] We calibrate the wall intensity tier by tier (equation (27))—implying that every tier faces approximately the same wedge at the peak—rather than imposing a single economy-wide γ. Cross-sector evidence supports this assumption. JOLTS openings over CPS unemployment by prior supersector level, sit +2.5 to +4.9 SD above their 2001–2019 baseline across all eleven supersectors, averaging about +3.5 SD—slightly *above* the 3 SD used here, so the aggregate anchor is conservative. The shift is uncorrelated with the size of each sector's downturn: regressed on the 2020 employment collapse, the slope is statistically zero ($\beta = -0.012$ SD per pp of collapse, $R^2 < 0.01$, $N = 11$, WLS by 2019 employment). Leisure and hospitality, which lost 23 pp, tightened no more in V/U than transportation, which lost less than 1 pp. The reason is mechanical: a deep collapse raises vacancies and the prior-industry unemployment pool in step, moving the numerator and denominator of V/U together. This is consistent with Forsythe et al. (2022), who find similar patterns using more granular job-postings data.

$C(H) = H^{\eta}$ and $\eta = 1.3$ (Muehlemann and Pfeifer, 2016), the marginal cost of a hire is $C'(H) = \eta H^{\eta-1} = \eta\, C(H)/H$, so at the tight-market hire rate q_1 the average-to-marginal wedge on top of the normal-times C' scales by $(q_1/q_0)^{\eta-1}$. With $q_1/q_0 = 1.33$ and $\eta - 1 = 0.3$, the within-firm convexity factor is $(q_1/q_0)^{\eta-1} \approx 1.09$.[24]

The estimate of c^{wall} is further scaled down by the share of the wage bill covered by skilled workers in the relevant sense: by our estimate, roughly four-fifths of the wage bill falls under the Linckh et al. (2025) definition of workers requiring some training. We assume that hiring workers outside this definition exhibits no convexity and set $\kappa_{\text{skill}} = 0.8$.

Combining these elements yields

$$c^{wall} = \kappa_{\text{skill}}\, \eta \left[(q_1/q_0)^{\eta-1} \big(r + \delta(q_1)\big) Q_1 \;-\; \big(r + \delta(q_0)\big) Q_0 \right], \tag{28}$$

where each bracketed term is the per-year user cost of one unit of hire capital under the relevant regime, and the difference isolates the *additional* rebuilding cost that the tight-market episode imposes relative to steady state. Evaluating at $Q_0 = 22/52$, $Q_1 = 30/52$, $q_0 = 0.0225$, $q_1 = 0.030$, and $r = 0.06$, the normal-times user cost is $(r + \delta_0)Q_0 \approx 0.141$; the tight-market user cost, including the within-firm convexity factor, is $(q_1/q_0)^{\eta-1}(r+\delta_1)Q_1 \approx 0.268$; so $c^{wall} = 0.8 \times 1.3 \times (0.268 - 0.141) \approx 0.13$.

Price stickiness. We calibrate the Rotemberg adjustment costs to the monthly Calvo price-change frequencies of Pasten et al. (2020) (hereafter PSW). The primitive is each industry's monthly non-adjustment probability $1 - f_m^i$, with the natural reading that stickier industries have flatter sectoral Phillips curves and pass less of a cost shock through to their customers. To make this monthly evidence comparable to our annual model, we match the twelve-month cumulative pass-through of a permanent cost shock in the monthly Calvo model to that of the annual Rotemberg model. This yields an annual non-adjustment probability $\lambda_a^i = (1 - f_m^i)^{12}$, which the standard log-linearized Calvo–

[24] The hiring-cost convexity estimated by Muehlemann and Pfeifer (2016)—$\eta = 1.3$ on skilled-worker hires—provides a within-firm intensity-convexity margin and is distinct from the across-market tightness convexity captured by the Linckh et al. (2025) elasticity above. We apply η here as a within-firm average-to-marginal conversion, where the variation identifying it is firms' own scale of hiring. The rich set of industry fixed effects and controls in Muehlemann and Pfeifer (2016) absorbs local market conditions. Blatter et al. (2012) report broadly consistent convexity using a Swiss firm-level sample — marginal hiring costs rise by 45% between the 1st and 10th hire and reach 24 weeks of wages for the 30th hire against an average of 10–17 weeks (a specification with significant linear, quadratic, and cubic components). Using US aggregate corporate-sector data and Q-theory GMM, Merz and Yashiv (2007) find $\hat{\eta}_2 = 3.4$ for hiring (approximately cubic) and explicitly reject the pure-quadratic benchmark $\eta = 2$ from below. Both estimates lie well above the $\eta = 1.3$ we adopt from Muehlemann and Pfeifer (2016), so our headline c^{wall} calibration is conservative — best read as a lower bound on the rebuilding-cost mechanism.

Rotemberg equivalence maps into a Rotemberg adjustment cost $\bar{\phi}_i$.[25] PSW frequencies match 218 of our 357 BEA industries directly; we recover the remaining gaps with a NAICS/BEA-66 stem fallback and hand-coded NAICS-2 analogs for twenty uncovered service industries, using 2019 QCEW wage bills as a NAICS-6 proxy for value-added shares throughout (Online Appendix A.7 provides the details).

Downward price rigidity. The price block features a strong asymmetry: prices adjust upward more readily than they fall. This feature limits deflation in the pandemic year, when energy prices collapsed, and otherwise plays a minor role. Guided by the structural-asymmetry literature (Eichenbaum et al., 2011; Karadi and Reiff, 2019; Cavallo et al., 2024), we target a 3–5× marginal-cost asymmetry between price decreases and increases. 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, $(\lambda_{C1,h}, \lambda_{C2,h}, \lambda_{C3,h}, \lambda_{C4,h}) = (1.88, 9.95, 14.07, 6.26)$ and $\lambda_{ml} = 8.58$. Online Appendix A.9 reports sensitivity to this parameter.

Inflation expectations. We require the model to fit the lagging behavior of inflation expectations shown in Figure 1, Panel A. To calibrate this, we estimate the three-weight projection in equation (22) on the Cleveland Fed market-based series (Haubrich et al., 2012), with the University of Michigan household survey serving as a cross-check. The estimates support the headline *adaptive-anchored* regime $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (0,\ 0.25,\ 0.75)$. Online Appendix A.8 details the estimation and reports the perfect-foresight case as robustness.

MPC. We set MPC = 0.4, in line with the average annual MPC of 0.37 in the heterogeneous-agent, incomplete-markets model of Dupor et al. (2023), which is disciplined by cross-regional estimates of consumer-spending responses to ARRA government spending.

[25] $\bar{\phi}_i \approx \dfrac{(\varepsilon - 1)\,\lambda_a^i}{(1 - \lambda_a^i)(1 - \beta\lambda_a^i)}$, with $\beta = 0.96$ and $\varepsilon = 5$, used only for the $1 - f_m^i \to \bar{\phi}_i$ conversion—the model takes $\bar{\phi}$ directly as a parameter. The same primitive maps to the per-industry annual pass-through $\delta_i = 1 - (1 - f_m^i)^{12}$ used in the empirical network operator L of Section 2. The full derivation, including the construction of the twelve-month pass-through invariant, is in Online Appendix A.7.

Saturating wedges. Five saturating wedges ensure that the model fits the sectoral data: the final-demand shifter $\xi^e_{i,t}$, the supply-efficiency wedge $a^e_{i,t}$, the labor wedge $\tau^e_{ij,t}$ (decomposed into a tier-mean part and a hit–safe contrast $\Delta\tau^e_i$), the structural 2020 markup $\bar{\mu}^e_i$, and the backfilling residual markups $\mu^e_{i,t}$. They are recovered recursively and play distinct roles. The conceptual division of labor is approximately as follows: $\xi^e_{i,t}$ moves the final-demand curve for the tier's (detrended) composite output, ensuring that the model matches the sectoral gross-output path $Y_{i,t}$; $a^e_{i,t}$ then moves supply efficiency, matching the sectoral employment path $N_{i,t}$ given $Y_{i,t}$ (so together, ξ^e pins down how much is produced in a period and a^e pins down how many workers are needed to deliver that Y at the matched labor productivity Y/N). As noted, we match detrended series because the model is stationary. The tier mean of τ^e enters labor demand log-linearly alongside a^e and is therefore redundant; its level is fully absorbed into $a^e_{i,2020}$. What matters is $\Delta\tau^e_i \equiv \tau^{e,\text{safe}}_i - \tau^{e,\text{hit}}_i$, which targets the excess downward dispersion of 2019–2020 employment growth relative to a pre-COVID baseline. The data target is the standard deviation of the collapse, calculated at the NAICS 6-digit level for each model chain. The 2020 price $P_{i,2020}$ is closed by the structural markup wedge $\bar{\mu}^e_i$, active only in the collapse year. The residual markup wedge $\mu^e_{i,t}$ closes the remaining 2021–2024 price gap and is iterated until convergence. $\mu^e_{i,t}$ and $\bar{\mu}^e_i$ serve distinct roles: $\bar{\mu}^e_i$ is kept in the counterfactuals of Section 5, whereas $\mu^e_{i,t}$ is residual and is zeroed out to identify the contribution of the wall to post-pandemic inflation. Online Appendix A.5 shows analytically that the wedges are locally identified.

A final parameter, the wage-margin attribution share x^e_{ij} (24 nodes), governs the fraction of the wall premium $\mathcal{W}_{ij,t}$ that flows into measured wages via equation (21); because it enters no other equilibrium condition, it is calibrated post-solution without affecting the allocation.[26]

Technical parameters. The model includes two additional parameters that play no substantive role but are technically necessary to close it: ι_a and ε. The first is set at a very low level, so as not to interfere with the solution over the period under consideration (2019–2024); it ensures that real wages catch up with labor productivity at the terminal point (so that $W/P = MPL$). The second determines the steady-state markup of firms, which, expressed as a share of gross output, is lower than typical estimates: the I–O-based ratio of surplus to gross output across endogenous industries (value-added weighted) is

[26] Per tier, x^e_{ij} is chosen by grid search over $[0, 1]$ as the largest value at which model-implied measured wages do not overshoot the observed tier wage path over 2021–2024 (tier wages from ECI/QCEW), i.e., closest-without-overshoot. The $[0, 1]$ restriction is imposed by the search domain, and hit and safe nodes within a tier share the tier's x^e_{ij}.

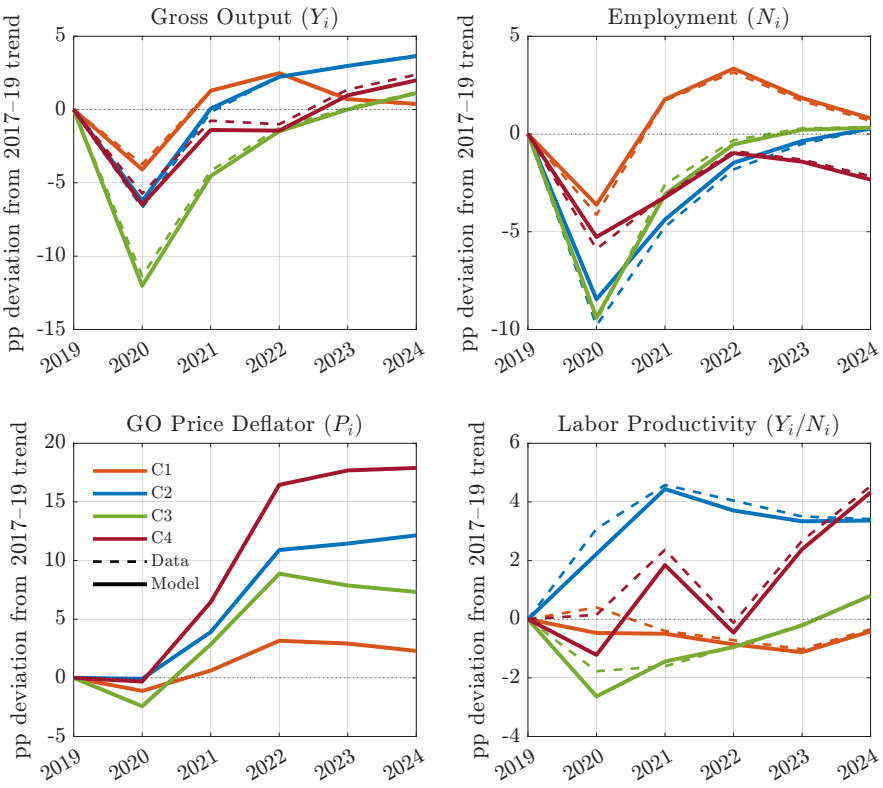

Figure 5: Chain-level fit under the saturated calibration.

Notes: Gross output, employment, labor productivity, and gross-output deflators by chain. Dashed: data. Solid: model with μ active. Cumulative log points relative to 2019, detrended.

24 percent, consistent with the assumed elasticity $\varepsilon = 5$.[27]

By construction, the saturated calibration ensures that the model matches the data; Figure 5 displays the resulting fit.

5 Quantitative Results

This section presents the model's quantitative results: the structural decomposition of post-COVID inflation, the implied steepening of the Phillips curve, and model–data comparisons for the reduced-form facts of Section 2.

[27] There are several ways of calculating it, yielding results between 20% and 30%. We use the residual formula $(GO_i - CE_i - MAT_i)/GO_i$ applied at the industry level to the 2017 detail use I–O table in producer prices, where GO_i is industry gross output, CE_i is compensation of employees, and MAT_i are intermediate inputs. This leaves taxes on production in the numerator and denominator. Industry-level ratios are winsorized at the 3rd and 97th percentiles and aggregated across the 341 endogenous BEA industries with value-added weights, yielding 23.7% (24.2% without winsorization).

5.1 Main result

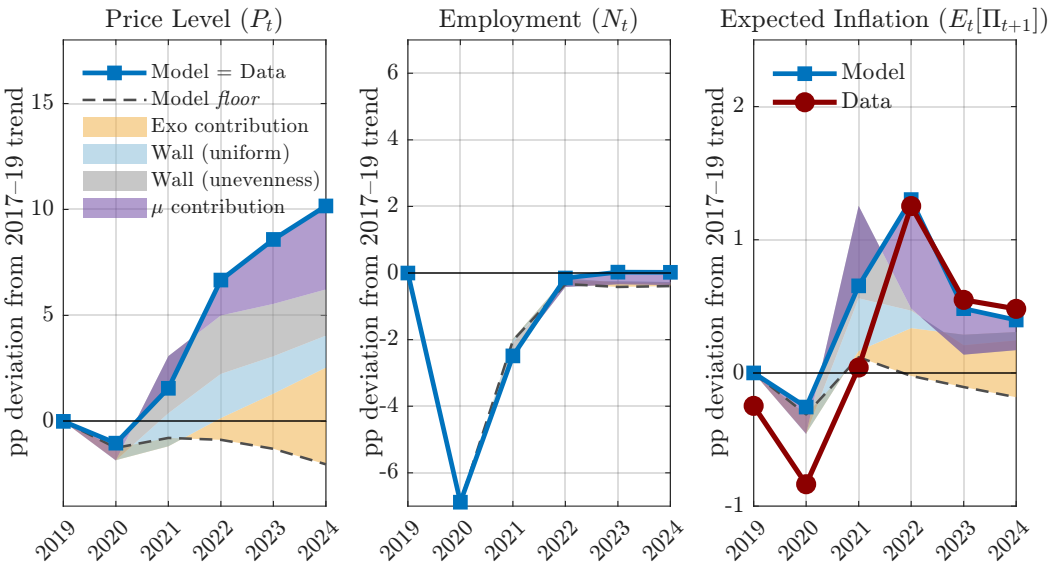

Figure 6: Baseline Model Decomposition.

Notes: Each panel plots pp deviation from the 2017–19 trend over 2019–2024 and decomposes the model path relative to the structural floor ($\gamma = 0$, no exo, no μ, dashed gray) into four stacked contributions: EXO shocks (amber), uniform-wall (light blue), wall-unevenness (gray), and the residual μ wedge (purple). The top blue line with square markers is the saturated Model, which equals Data by construction. The dark-red line in the expected-inflation panel is the Cleveland Fed 1-year expectations series.

Figure 6 decomposes the model paths of PCE-weighted gross-output inflation, employment, and expected inflation. We construct the counterfactuals sequentially. First, we remove the recovery markup wedge by setting $\mu^e = 0$. We then remove the EXO-sector shocks. Finally, we set $\gamma = 0$, eliminating workforce rebuilding costs. The resulting path is the model's *structural floor*.

To separate the effect of aggregate workforce depletion from that of its uneven distribution across nodes, the figure also reports an *even-wall* counterfactual. In this variant, every node enters 2021 with the same proportional reduction in workforce capacity, $N_u^{\text{CW}} = N_u^{\text{ss}}\,(N_{2020}/N^{\text{ss}})$. Thus, the aggregate loss of workforce capacity is unchanged, but its distribution across nodes is equalized. The gap between the structural floor and the even-wall path measures the effect of the aggregate capacity loss—the *uniform-wall* component. The additional gap between the even-wall path and the model with the actual, uneven wall measures the contribution of *wall unevenness*. Their sum is the total contribution of the cost wall.

As the figure shows, the CW is the leading structural contributor to cumulative above-trend price-level growth over the recovery, and it is sharply *front-loaded*: it accounts for

about +4.3 pp of cumulative inflation in 2021, when firms are rebuilding hardest, and settles at +3.7 pp by 2024. CW unevenness carries most of the 2021 contribution and about 60% of the 2024 total. The residual markup wedge μ^e, by construction zero in 2020–2021, builds slowly and by 2024 slightly exceeds the wall itself (+3.9 pp), reflecting a persistent gap that the CW alone cannot close as firms complete rebuilding. The model tracks inflation expectations closely, though it still overpredicts them in 2020 and 2021.

The structural floor is mildly deflationary from 2021 onward for two reasons. First, employment remains below its 2019 baseline, putting downward pressure on the base wage for several years, as the standard Phillips curve implies. Second, labor-productivity growth runs above trend over 2020–2024, which lowers inflation in the long run. The terminal price level is roughly 2% below its initial steady-state level in 2050, and the model implies mildly below-trend inflation between 2025 and 2029, catching up only in the long run through $\iota_a > 0$. The perfect-foresight model yields similar results but less persistent inflation after 2022; results are reported in Online Appendix A.9.2.

The effect on the price level extends well beyond 2024, for several reasons. First, prices are sticky and downwardly rigid, so inflation does not immediately reflect the initial spike in the CW-implied labor wedge; prices are also slow to incorporate the disinflationary forces that emerge later, mostly beginning in 2025, as reported in the appendix. Second, adaptive expectations shut down the front-loading of inflationary forces through expectations of the future, which changes the timing of inflation propagation. Third, after a prolonged period of elevated prices, part of the price increase filters directly into wages through backward-looking catch-up wage growth, which is modeled with additional lags. Most importantly, these standard features interact with the path of measured labor productivity, which, as noted in the discussion of Table 5, accelerated rapidly following the pandemic. This acceleration insulated prices from the CW-implied rise in margins early on, shifting the inflationary impulse forward in time as productivity growth slowed in 2022 and 2023. Under this dynamic, price increases had more time to become embedded in wages and therefore became more persistent. Although the productivity boom–bust cycle is exogenous in the model, it is closely linked to the CW mechanism: faced with an unprecedented contraction and the need to downsize, firms laid off their least productive workers (Crust et al., 2020; Fernald and Li, 2022; Stewart, 2022).

One interpretation of the post-pandemic inflation outbreak is that it reveals a nonlinear Phillips-curve response to large shocks. Our model only partly supports this narrative: much of the model-implied steepening of the Phillips curve comes from the uneven geometry of the shock, which may be specific to this crisis.

Figure 7 plots the Phillips curve implied by the model when it is hit by a one-time,

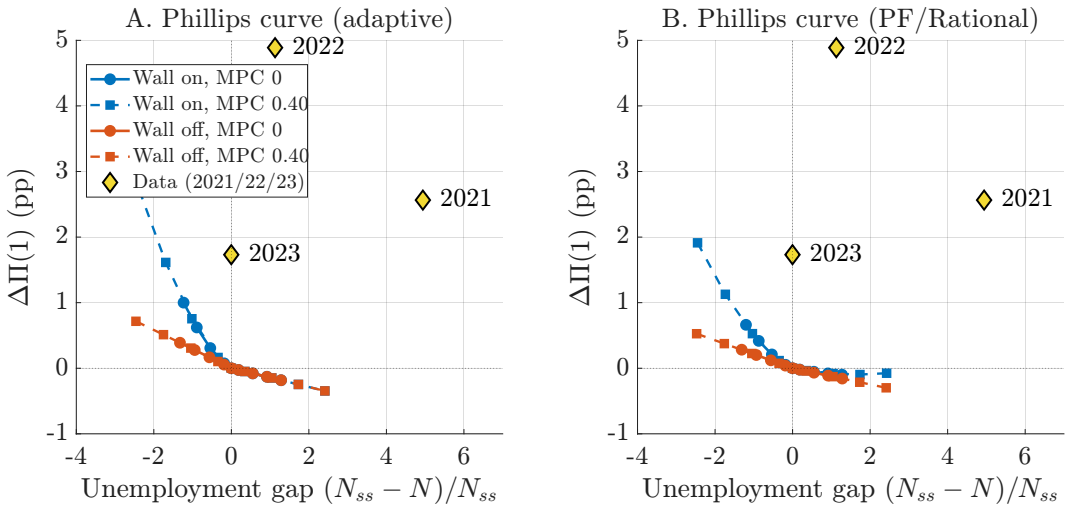

Figure 7: Model Phillips curve under uniform iid ξ^e demand shocks.

Notes: Impact-period ($t = 1$) inflation $\Delta\Pi(1)$ is plotted against the unemployment-gap proxy $-\Delta N(1)$, in percentage points. Blue: calibrated wall; orange: structural counterfactual $\gamma_i = 0$. Panels (a) and (b) report adaptive expectations $(0, 0.25, 0.75)$ and PF/Rational expectations $(1, 0, 0)$, respectively.

symmetric aggregate demand shock (ξ) applied at the steady state.[28] The Phillips curve generated by the model is steeper than leading near-zero post-2000 empirical estimates suggest.[29] At a 0.7% expansionary pulse, the wall-on slope reaches 0.82 under adaptive expectations, approximately 2.8× the wall-off baseline of 0.30. Quasi-experimental evidence supports both the direction and the rough magnitude of this shift: Cerrato and Gitti (2025) estimate that the slope of the U.S. Phillips curve rose from roughly 0.15 before COVID to above unity afterward. Our wall-on slope—a moment we do not target—accounts for a sizable share of that measured steepening, though the comparison is necessarily approximate given the different slope concepts involved. Consistent with the mechanism, Cerrato and Gitti (2025) trace the steepening to a collapse in the labor-supply elasticity of the marginal worker—precisely the footprint of a binding rebuilding constraint—while Gitti (2026) finds that the regional Phillips curve steepens past the labor-shortage threshold through new-hire wages. The key point, however, is that the actual data for 2021, 2022,

[28] To generate this plot, we sweep the impulse over $\xi^e \in \{\pm 0.1\%, \pm 0.3\%, \pm 0.5\%, \pm 0.7\%\}$ symmetrically for all tiers and plot the impact-period tuple $(-\Delta N(1), \Delta\Pi(1))$, starting from the steady state. Online Appendix Figures 13–16 in Section A.9 show the underlying impulse-response paths at the ±0.7% shock-size benchmarks: Figures 13–14 for the expansionary pulse and Figures 15–16 for the contractionary pulse.

[29] For instance, Hazell et al. (2022) estimate a quarterly slope of $\kappa \approx 0.006$ (≈ 0.025 annualized) using state-level instrumental variables, and Del Negro et al. (2020) find a posterior narrowly centered on zero across most post-1990 specifications. Furlanetto and Lepetit (2024) survey this literature, noting that while the values calibrated in standard DSGE models tend to be slightly larger than these empirical estimates, they remain small. Notably, the Federal Reserve's workhorse FRB/US model, widely used for Tealbook forecasting scenarios, assumes an aggregate slope well below even the Hazell et al. (2022) estimate.

and 2023 lie far away from this locus. We therefore caution against drawing conclusions from this episode about the general response of inflation to large shocks.

Under perfect foresight, the model implies an upward-sloping Phillips curve in the right tail, as shown in the right panel. Intuitively, a negative demand shock today contracts employment, triggering wall-implied inflation in the future; if anticipated, this future inflation raises current inflation. The result is stagflation.

Model mechanism. The mechanism behind these results is straightforward. The key object is the calibration target c^{wall} in the calibration of the γ_i's. Recall that we set this target so that, on average, all tiers face this marginal cost of rebuilding their workforce between 2021 and 2022. Because the model's saturating wedges are disciplined to match the sectoral paths of employment, the CW-generated labor wedge necessarily *equals* c^{wall} across all tiers in 2021 and 2022. How this calibration-implied cost impulse translates into inflation is then governed by the New Keynesian elements of the model: it is only partly passed through on impact, with the network further attenuating pass-through. The μ^e wedge does not affect employment much, because it shifts the distribution of income rather than its level; as a result, the assumed value of the MPC matters little for the results (not shown).

5.2 Model validation and extensions

The remainder of the section discusses broader implications of the mechanism and explores additional properties of the model.

ULC–price pass-through steepening. Our setup assumes that the CW-implied labor wedge "leaks" into paid wages through the mechanical attribution implied by equation (21). The same labor-market evidence suggests that this leakage should amount to about 20% of the wall-implied wedge on average. Our calibration specifies this leakage-parameter path and targets sectoral wages in the data. Although the attribution is mechanical, it is supported by evidence and has a natural interpretation: as firms face the CW, profits rise, hiring becomes intensive, and workers' bargaining power rises as well, increasing wage pressure.

This feature allows us to relate the model to the elevated ULC–price pass-through observed in the data during the outbreak. Online Appendix A.10.3 documents the empirical fact: estimating the cross-industry OLS regression $\Delta \log P_{i,w} = \alpha_{s(i)} + \beta_w \Delta \log \mathrm{ULC}_{i,w} + \varepsilon_{i,w}$ by window with BEA-66 sector fixed effects, the pass-through slope nearly doubles from

pre-COVID to Act I ($\hat{\beta}$ rises from $+0.14^{***}$ to $+0.26^{**}$ in the full sample, and from $+0.15^{***}$ to $+0.31^{*}$ in the $m + l$-tier subsample), and remains elevated in Act II among upstream and middle-tier industries ($\hat{\beta}_{m+l} = +0.29^{***}$; see Table 23). We repeat this specification on the model side, running the OLS regression $\Delta \log P_u = \alpha + \beta\, \Delta \mathrm{ULC}_u + \varepsilon_u$ across the model's $N = 24$ hit/safe nodes, evaluated at the structural baseline ($\mu^e = 0$). Under adaptive expectations, the estimated slope is $\hat{\beta} = +0.53$ ($t = 3.10$) over Act I and $+0.37$ ($t = 3.21$) over Act II; under perfect foresight, it is $\hat{\beta} = +0.48$ ($t = 3.26$) and $+0.32$ ($t = 3.20$). The model thus reproduces the qualitative Act I–Act II profile of the empirical pass-through—an elevated slope during the outbreak, with the upstream/middle-tier subsample carrying the persistence into Act II.

Inflation, employment collapse, and recovery velocity. We next examine the model's predictions for the link between upstream supply-chain recovery and inflation (Fact 2), as well as the link between employment collapse and inflation (Fact 1).

We pool data over the $2019 \to 2022$ window and regress cumulative price growth on three supplier-network terms: the own 2020 employment collapse $-\Delta \log N_i^{2020}$ (positive for a collapse), its supplier-network propagation $-\Delta \log \mathcal{N}_i^{2020} = \Omega \cdot (-\Delta \log N^{2020})$, and the supplier-network recovery velocity $\mathcal{V}_i = \Omega \cdot \Delta \log N_{\text{per yr}}^{2020 \to 2022}$. The data counterpart is estimated at BEA-357 industry resolution with BEA-Summary sector fixed effects, both by WLS-OLS and by 2SLS that instruments the own collapse with the predetermined JEM exposure E_i. The model side estimates the identical three-regressor OLS specification at $N = 12$ chain-by-tier units with chain fixed effects. The data–model comparison is therefore approximate and qualitative in nature.

A note of caution is in order regarding the IV setup. The exclusion restriction cannot be asserted with certainty. For example, direct supply restrictions, including mandated closures and capacity limits, might have affected prices independently of measured employment changes. Under such a leaky instrument, the true model would be $\Delta \log P_i = \gamma \cdot \Delta \log N_i^{2020} + \delta \cdot E_i + \mathbf{x}_i^{\top} \mathbf{\Phi} + \varepsilon_i$, where $\delta \neq 0$ captures the direct non-labor effect of E_i on prices. Standard IV algebra implies $\hat{\beta}^{\mathrm{IV}} = \gamma + \delta / \hat{\pi}_E$, where $\hat{\pi}_E = -6.39$ is the first-stage slope from Table 7. What makes the IV informative in this case is that the implied bias, $-\delta/6.39$, is small relative to $\hat{\beta}^{\mathrm{IV}} = -0.60$. Flipping the sign would require $|\delta| \approx 3.8$ pp per unit of E_i, a within-sector contamination that seems implausibly large.[30]

The discriminating coefficient on the supplier-recovery velocity $\mathcal{V}$ is *negative and statistically significant in every column*: data OLS -2.67^{**}, data IV -4.21^{***}, and model OLS -9.09^{**}

[30]Weak-instrument-robust inference (inverted Anderson–Rubin confidence set) is reported in Online Appendix A.10.4; all tests survive at the 5% level despite a first-stage F-statistic (8.68) below conventional cutoffs.

Table 7: Fact 2 recovery regression: data versus model.

	Data (BEA-357, $N = 321$)		Model (12 chain×tier units)	
Coefficient	OLS	IV (2SLS)	Adaptive	PF/Rational
$-\Delta \log N_i^{2020}$	0.04	-0.61^{**}	-0.02	-0.02
	(+0.71)	(-2.24)	(-0.15)	(-0.15)
$-\Delta \log \mathcal{N}_i^{2020}$	-1.67^{*}	-0.91	-0.54	-0.54
	(-1.89)	(-0.99)	(-0.49)	(-0.49)
$\mathcal{V}_i$	-2.67^{**}	-4.30^{***}	-9.09^{**}	-9.09^{**}
	(-2.04)	(-3.12)	(-2.69)	(-2.69)
Fixed effects	BEA-Summary	BEA-Summary	chain (C1–C4)	chain (C1–C4)
Estimator	WLS, clustered SE	2SLS, clustered SE	OLS	OLS
N	321	321	12	12

Notes: Dependent variable: cumulative 2019–2022 industry inflation $\Delta \log P$ (pp). Endogenous regressor: $\Delta \log N_i^{2020}$, IV-instrumented by predetermined E_i. Network controls $\mathcal{E}_i$, $\Delta \log \mathcal{N}_i^{2020}$, $\mathcal{V}_i$ per §2.2 at $\delta = 0.0$. Data-side FE: BEA-Summary sector; model-side FE: chain (C1–C4) at $N = 12$ chain×tier units, hit/safe collapsed ($\mathcal{E}_i$ has no model counterpart, marked '—'). First-stage Wald $F = 7.70$; weak-IV diagnostics in Online Appendix A.10.4. Cluster-on-BEA-Summary SE. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

under both expectations specifications. Faster supplier recovery in the model thus also predicts *less* cumulative inflation. The supplier-network collapse term enters negatively on both sides as well, and the own-labor pass-through is sharply identified at -0.60^{**} in the data IV but is absorbed into the supplier-network terms at the model's $N = 12$ aggregation, where cross-tier inflation variation is carried by the supplier-recovery channel.

Price dispersion. Figure 8 compares the model-implied price-growth dispersion $\Delta\sigma^{+}(\Delta \log P)$ to the data counterpart of the dispersion evidence in Fact 1. Because the model features zero dispersion in the steady state, the relevant comparison is *excess* dispersion: we subtract the 2017–2019 baseline dispersion from the data values and denote the resulting object by $\Delta\sigma$.

The model matches the sectoral dispersion reasonably well on average, but it does not reproduce every sectoral pattern. The largest miss is C4, which shows large dispersion in the data but not in the model. Although C4 collapsed highly asymmetrically across hit and safe nodes in 2020, with hit-minus-safe gaps of 2.8–5.5 pp in $\log N$, its recovery at our level of aggregation was fairly symmetric between hit and safe sectors, so the model shows little price dispersion there. The visible micro-level dispersion in the data likely reflects other disruptive forces within this chain: raw-material shortages, semiconductor bottlenecks, transportation backlogs, and pent-up selective demand are well-documented candidates in manufacturing. This miss suggests that shock asymmetry may have contributed more

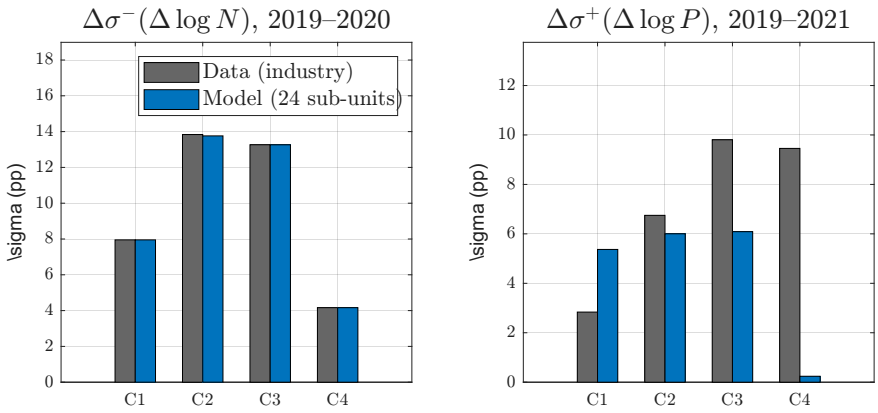

Figure 8: Within-chain dispersion: data vs. model.

Notes: The figure compares data and model moments by chain. Left panel: $\Delta\sigma^{-}(\Delta \log N)$, the 2019–2020 downward employment-growth dispersion net of its pre-COVID average. Right panel: $\Delta\sigma^{+}(\Delta \log P)$, the 2019–2021 upward price-growth dispersion net of its pre-COVID average. Employment moments use NAICS-6 industries within chain, weighted by 2019 employment. Price moments use the 341 endogenous BEA industries within chain, Domar-weighted using 2017 shares. NAICS-6 employment log changes are winsorized at 3%. Gray bars report data; blue bars report the model aggregation at the 24 chain-node level.

to inflation through this sector than our model predicts.

Inflation decomposition by component. We conclude by comparing the inflation decomposition of Section 2.2.1 in the model and the data at the common 24-node aggregation. For comparability, the model series add back the trends removed during calibration. We report the Network (L) decomposition, which attributes every dollar of inflation to its ultimate primitive cost via the Domar weights $\kappa^{\top} = (\Theta c)^{\top} L$.

As Table 8 shows, the model reproduces the key Act I pattern: above-trend inflation rises primarily through profit-margin residuals, while productivity growth acts as a major drag on inflation, as in the data.[31] Act II is informative for a different reason: the model underpredicts wage inflation and correspondingly overshoots the margin contribution. This gap indicates that our wage-leakage attribution is conservative: as the wall recedes, the model falls short of the catch-up wage inflation observed in the data. The missing wage inflation toward the end of the sample is central to understanding why inflation is more persistent in the data than in the model.

[31]This margin-then-ULC sequence corresponds broadly to the shortages-then-labor pattern documented in Bernanke and Blanchard (2025) for the U.S. and Blanchard and Bernanke (2024) across eleven economies; the mapping is qualitative, not an identity.

Table 8: 24-sub-unit Network decomposition of GO inflation: Data, Model, and Model $\mu = 0$.

	Act I (2019–2021)			Act II (2021–2023)		
	Data	Model	Model μ=0	Data	Model	Model μ=0
GO inflation	**+5.20**	**+6.06**	**+8.03**	**+6.72**	**+9.48**	**+3.08**
Margins	+3.28	+6.56	+8.25	+0.19	+13.11	+8.59
ULC	+1.61	−0.90	−0.61	+6.13	−4.32	−6.21
+ Wages	+7.93	+9.60	+9.60	+4.55	+1.61	+1.61
− Productivity	−6.32	−10.49	−10.21	+1.58	−5.93	−7.83
Imports	+0.20	+0.39	+0.39	+0.15	+0.70	+0.70

Notes: Network decomposition at the 24-sub-unit level (12 tier composites × hit/safe) for Model and Model μ^e=0; fixed 2017 N-PCE weights, 341 endogenous BEA-detail industries. Every dollar of inflation is attributed to its ultimate primitive cost via the Domar weights $\kappa^\top = (\Theta c)^\top L_{24\times24}$, with $L_{24\times24} = (I - A_{24\times24})^{-1}$ built from `omega_24x24_with_imp_mon.csv`; there is no separate Materials (MAT) row because upstream input inflation is unrolled back into upstream Margins, ULC, and Imports. Windows: Act I = $g_{2020} + g_{2021}$, Act II = $g_{2022} + g_{2023}$. Wages use the no-overshoot Pipeline 1 wmeas. Entries in pp. Identity sums hold to machine precision on the model side per column.

6 Conclusion

We have shown that, as the economy reopened, firms in the most disrupted sectors faced convex costs of rebuilding depleted workforces. These rebuilding costs propagated through input–output linkages, converting a sectoral shock into economy-wide inflation. Quantitatively, the mechanism accounts for essentially all above-trend inflation in 2021 and roughly 70% at the 2022 peak, and most of its force derives from the unevenness of the destruction rather than its aggregate size. Its main shortfall is persistence: inflation remained elevated beyond the period over which the rebuilding constraints bind. The same friction rationalizes the apparent post-COVID steepening of the aggregate Phillips curve.

Our results do not imply that demand was unimportant. Monetary and fiscal support undoubtedly accelerated the surge; fractured supply would not have ignited inflation on its own. But the two sides do not enter additively: demand and supply rebuilding interact, and the relevant policy question is not whether demand fueled inflation but what path of demand would have achieved recovery at the lowest social cost—counting both the length of the recovery and the burden of inflation. Answering that question requires counterfactuals from a structural model rather than accounting decompositions, and the unevenness of the contraction is central to it: policy that leans against rebuilding-driven inflation trades a cooler price level against a slower reconstruction of capacity. Here we have taken the first step, providing a consistent account of the inflation as the

product of COVID's destructive force; we analyze the optimal monetary policy response in companion work (Drozd and Tavares, 2026).

A cross-country comparison is beyond the scope of the current paper, but the transatlantic policy divergence hints that the preservation of organizational capital might have mattered. The United States and Europe made opposite choices in 2020: the United States insured *workers*, through expanded unemployment benefits and transfers, while Europe insured *jobs*, through short-time-work schemes that at their peak covered on the order of a fifth of dependent employment (Giupponi et al., 2022).[32] Quasi-experimental evidence shows that short-time-work schemes preserve worker–firm matches, and that their value is greatest precisely when shocks are transitory and matches embody specific capital (Giupponi and Landais, 2023)—the organizational capital at the center of our mechanism. Europe therefore entered the reopening with its productive capacity largely intact. Consistent with the rebuilding mechanism, European inflation remained well below U.S. inflation through 2021; Europe's own 2022 surge was driven mostly by energy and food prices following the invasion of Ukraine—a different shock (Blanchard and Bernanke, 2024). The comparison is only suggestive, because it is confounded by the substantially larger U.S. fiscal transfers, and part of Europe's muted response may reflect structurally flatter wage Phillips curves.

Notably, measured labor-market tightness rose in much of Europe as it did in the United States (Blanchard and Bernanke, 2024), yet the inflationary consequences differed sharply. This contrast underscores our paper's broader lesson: what matters is the granular state of the supply side into which demand returns. Inflationary pressure during recoveries from large disruptions depends not only on how much capacity is lost, but on where it is lost.

References

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Aepli, 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.

[32]The closest U.S. analog, the Paycheck Protection Program, preserved few employment relationships at very high cost (Autor et al., 2022).

Arias, Jonas E., Juan F. Rubio-Ramírez, and Minchul Shin, "Are Fiscal Transfers Inflationary?," Working Paper 26-23, Federal Reserve Bank of Philadelphia 2026.

Atalay, Enghin, "How Important Are Sectoral Shocks?," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 254–280.

Atkeson, Andrew and Patrick J. Kehoe, "Modeling and Measuring Organization Capital," *Journal of Political Economy*, 2005, *113* (5), 1026–1053.

Autor, David, David Cho, Leland D. Crane, Mita Goldar, Byron Lutz, Joshua K. Montes, William B. Peterman, David D. Ratner, Daniel Villar Vallenas, and Ahu Yildirmaz, "The $800 Billion Paycheck Protection Program: Where Did the Money Go and Why Did it Go There?," NBER Working Paper 29669, National Bureau of Economic Research January 2022.

Ball, Laurence, Daniel Leigh, and Prachi Mishra, "Understanding U.S. Inflation During the COVID Era," *Brookings Papers on Economic Activity*, 2022, *Fall 2022*, 1–54.

Baqaee, David and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

_ and _, "Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis," *American Economic Review*, 2022, *112* (5), 1397–1436.

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.

Beaudry, Paul, Chenyu Hou, and Franck Portier, "On the Fragility of the Nonlinear Phillips Curve View of Recent Inflation," NBER Working Paper 33522, National Bureau of Economic Research 2025.

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.

Bernanke, Ben S. and Olivier J. Blanchard, "What Caused the US Pandemic-Era Inflation?," *American Economic Journal: Macroeconomics*, 2025, *17* (3), 1–35.

Blanchard, Olivier J. and Ben S. Bernanke, "An Analysis of Pandemic-Era Inflation in 11 Economies," Working Paper 32532, National Bureau of Economic Research May 2024.

Blanco, Andrés, Corina Boar, Callum J. Jones, and Virgiliu Midrigan, "Non-Linear Inflation Dynamics in Menu Cost Economies," NBER Working Paper 32094, National Bureau of Economic Research 2024.

Blatter, Marc, Samuel Muehlemann, and Samuel Schenker, "The costs of hiring skilled workers," *European Economic Review*, January 2012, *56* (1), 20–35.

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.

Bohr, Clement E., "Capacity Buffers: Explaining the Retreat and Return of the Phillips Curve," 2024. ECB Forum on Central Banking, Sintra.

California Policy Lab, "Analysis of Unemployment Insurance Claims in California," Technical Report, California Policy Lab, UCLA and UC Berkeley December 2021.

Carvalho, Vasco M., Makoto Nirei, Yukiko U. Saito, and Alireza Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara, "Large Shocks Travel Fast," *American Economic Review: Insights*, 2024, *6* (4), 558–574.

Cerrato, Andrea and Giulia Gitti, "The Return of the Phillips Curve: Evidence from US Cities," June 2025. Working paper, Federal Reserve Bank of Boston and Collegio Carlo Alberto.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "Business Cycle Accounting," *Econometrica*, 2007, *75* (3), 781–836.

Coibion, Olivier and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 2012, *120* (1), 116–159.

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

Crust, Erin E., Mary C. Daly, and Bart Hobijn, "The Illusion of Wage Growth," *FRBSF Economic Letter*, 2020, *2020* (26), 01–05.

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.

Dingel, Jonathan I. and Brent Neiman, "How Many Jobs Can Be Done at Home?," *Journal of Public Economics*, 2020, *189*, 104235.

Domash, Alex and Lawrence H. Summers, "How Tight Are U.S. Labor Markets?," Working Paper 29739, NBER 2022.

Drozd, Lukasz A., "Built by Actions, Lost in Words: Central Bank Communication under Inflation Targeting," July 2026. Working paper, Federal Reserve Bank of Philadelphia.

_ **and Marina M. Tavares**, "Spending the Anchor: How Monetary Policy Should Respond to Destructive Recessions," Working Paper, Federal Reserve Bank of Philadelphia July 2026. Forthcoming.

Dupor, Bill, Marios Karabarbounis, Marianna Kudlyak, and M. Saif Mehkari, "Regional Consumption Responses and the Aggregate Fiscal Multiplier," *The Review of Economic Studies*, November 2023, *90* (6), 2982–3021.

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

Fernald, John G. and Huiyu Li, "The Impact of COVID on Productivity and Potential Output," *Federal Reserve Bank of San Francisco Working Paper*, 2022, (2022-19).

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

Fornaro, Luca, "Fiscal Stimulus with Supply Constraints," *IMF Economic Review*, 2025, *73* (3), 708–732.

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.

Furlanetto, Francesco and Antoine Lepetit, "The Slope of the Phillips Curve," Finance and Economics Discussion Series 2024-043, Board of Governors of the Federal Reserve System 2024. Forthcoming as a chapter in the *Handbook of Inflation*, eds. G. Ascari and R. Trezzi.

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.

Giannone, Domenico and Giorgio E. Primiceri, "The Drivers of Post-Pandemic Inflation," Working Paper 32859, National Bureau of Economic Research 2024.

Gitti, Giulia, "Nonlinearities in the Regional Phillips Curve with Labor Market Tightness," July 2026. Working paper, Collegio Carlo Alberto and University of Turin.

Giupponi, Giulia and Camille Landais, "Subsidizing Labour Hoarding in Recessions: The Employment and Welfare Effects of Short-Time Work," *Review of Economic Studies*, 2023, *90* (4), 1963–1994.

_ , _ , and Alice Lapeyre, "Should We Insure Workers or Jobs During Recessions?," *Journal of Economic Perspectives*, 2022, *36* (2), 29–54.

Glover, Andrew, José Mustre del Río, and Alice von Ende-Becker, "How Much Have Record Corporate Profits Contributed to Recent Inflation?," *Federal Reserve Bank of Kansas City Economic Review*, 2023, *108* (1), 1–13.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," *American Economic Review*, 2022, *112* (5), 1437–1474.

Harding, Martín, Jesper Lindé, and Mathias Trabandt, "Understanding Post-COVID Inflation Dynamics," *Journal of Monetary Economics*, 2023, *140*, S101–S118.

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.

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.

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*, 2022, *48* (1), 61–70.

Karadi, Peter and Adam Reiff, "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 111–146.

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" 2025.

Merz, Monika and Eran Yashiv, "Labor and the Market Value of the Firm," *American Economic Review*, 2007, *97* (4), 1419–1431.

Miyauchi, Yuhei, "Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade," *Econometrica*, 2024, *92* (6), 1869–1905.

Mongey, Simon, Laura Pilossoph, and Alexander Weinberg, "Which Workers Bear the Burden of Social Distancing?," *Journal of Economic Inequality*, 2021, *19* (3), 509–526.

Moscarini, Giuseppe and Fabien Postel-Vinay, "The Job Ladder: Inflation vs. Reallocation," Working Paper 31466, National Bureau of Economic Research 2023.

Muehlemann, Samuel and Harald Pfeifer, "The Structure of Hiring Costs in Germany: Evidence from Firm-Level Data," *Industrial Relations: A Journal of Economy and Society*, 2016, *55* (2), 193–218.

_ **and Mirjam Strupler Leiser**, "Hiring Costs and Labor Market Tightness," *Labour Economics*, 2018, *52*, 122–131.

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.

Oi, Walter Y., "Labor as a Quasi-Fixed Factor," *Journal of Political Economy*, 1962, *70* (6), 538–555.

Pasten, Ernesto, Raphael Schoenle, and Michael Weber, "The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy," *Journal of Monetary Economics*, 2020, *116*, 1–22.

Prescott, Edward C. and Michael Visscher, "Organization Capital," *Journal of Political Economy*, 1980, *88* (3), 446–461.

Rubbo, Elisa, "Networks, Phillips Curves, and Monetary Policy," *Econometrica*, 2023, *91* (4), 1417–1455.

_ , "What Drives Inflation? Lessons from Disaggregated Price Data," Working Paper 32194, National Bureau of Economic Research 2024.

Schmitt-Grohé, Stephanie and Martín Uribe, "Heterogeneous Downward Nominal Wage Rigidity: Foundations of a Nonlinear Phillips Curve," Working Paper 30774, National Bureau of Economic Research 2022.

Stewart, Jay, "Why Was Labor Productivity Growth So High During the COVID-19 Pandemic? The Role of Labor Composition," *International Productivity Monitor*, Spring 2022, (42), 87–103.

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.

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.

Online Appendix

Not for Publication

Capacity Walls in Production Networks: Inflation after COVID

Lukasz A. Drozd Marina M. Tavares

A.1 Inflation Accounting Framework: Derivation

This appendix derives the inflation-accounting identity used in Section 2.1. We begin with the nested production structure, derive industry unit costs, map commodity prices into industry materials costs using the BEA supply–use tables, and finally stack the resulting industry equations into the network system used in the empirical analysis.

A.1.1 Technology and variable cost

There are I industries, indexed by i, producing J commodities, indexed by j. Industry i combines labor N_i and a composite of material inputs m_i through the constant-returns CES aggregator

$$V_i = \left[\eta_i^{1/\rho} N_i^{(\rho-1)/\rho} + (1-\eta_i)^{1/\rho} m_i^{(\rho-1)/\rho}\right]^{\rho/(\rho-1)}, \tag{A.1}$$

where ρ is the elasticity of substitution between labor and materials and η_i is the labor-distribution parameter. Industry output is

$$Y_i = Z_i V_i^{\alpha_i}, \qquad Z_i \equiv z_i K_i^{1-\alpha_i}, \tag{A.2}$$

where α_i is the variable-input share, z_i is total factor productivity, and K_i is predetermined capital.

The materials composite combines domestically produced and imported commodities:

$$\begin{aligned} m_i = \Bigg[& \sum_{j=1}^{J} (\zeta_{ij}\omega_{ij})^{1/\rho_m} \widetilde{m}_{ij}^{(\rho_m-1)/\rho_m} \\ & + \sum_{j=1}^{J} \left((1-\zeta_{ij})\omega_{ij}\right)^{1/\rho_m} \left(\widetilde{m}_{ij}^{\mathrm{IMP}}\right)^{(\rho_m-1)/\rho_m} \Bigg]^{\rho_m/(\rho_m-1)}. \end{aligned} \tag{A.3}$$

Here $\widetilde{m}_{ij}$ and $\widetilde{m}_{ij}^{\mathrm{IMP}}$ denote, respectively, domestic and imported purchases of commodity j by industry i; ω_{ij} is commodity j's share in industry i's material bundle; ζ_{ij} is the domestic sourcing share; and ρ_m is the elasticity of substitution within the materials bundle.

A domestic commodity can in turn be produced by several industries. Commodity j is assembled from industry outputs according to

$$\widetilde{Y}_j = \left[\sum_{i=1}^{I} \theta_{ij}^{1/\rho_Y} Y_i^{(\rho_Y-1)/\rho_Y}\right]^{\rho_Y/(\rho_Y-1)}, \tag{A.4}$$

where θ_{ij} is industry i's share in the domestic production of commodity j, and ρ_Y is the elasticity of substitution across industries producing the same commodity.

Let p_i denote the dual price index associated with the materials bundle (A.3). Conditional on producing Y_i, industry i's measured variable cost is

$$C_i(Y_i) \equiv \min_{N_i, m_i} \left\{ W_i N_i + p_i m_i \right\} \quad \text{subject to} \quad Z_i V_i^{\alpha_i} \geq Y_i, \tag{A.5}$$

where W_i is the industry wage. Because the suballocation among domestic and imported commodities has already been summarized by the dual price p_i, equation (A.5) is the outer cost-minimization problem. We define the measured margin as revenue relative to measured variable cost:

$$M_i \equiv \frac{P_i Y_i}{C_i(Y_i)}, \tag{A.6}$$

so that

$$\Delta \log P_i = \Delta \log M_i + \Delta \log \left(\frac{C_i(Y_i)}{Y_i} \right). \tag{A.7}$$

A.1.2 Labor, materials, and industry unit cost

The unit price of the variable-input composite (A.1) is

$$P_i^V = \left[\eta_i W_i^{1-\rho} + (1-\eta_i) p_i^{1-\rho} \right]^{1/(1-\rho)}, \tag{A.8}$$

with the usual Cobb–Douglas limit when $\rho = 1$. Since V_i is constant returns in labor and materials,

$$C_i(Y_i) = P_i^V V_i, \qquad V_i = \left(\frac{Y_i}{Z_i} \right)^{1/\alpha_i}. \tag{A.9}$$

The first-order condition for the allocation between labor and materials is

$$\frac{m_i}{N_i} = \frac{1-\eta_i}{\eta_i} \left(\frac{W_i}{p_i} \right)^{\rho}, \tag{A.10}$$

or, equivalently,

$$\frac{p_i m_i}{W_i N_i} = \frac{1-\eta_i}{\eta_i} \left(\frac{p_i}{W_i} \right)^{1-\rho}. \tag{A.11}$$

Labor's share of variable cost is

$$s_i^N \equiv \frac{W_i N_i}{C_i(Y_i)} = \frac{\eta_i W_i^{1-\rho}}{\eta_i W_i^{1-\rho} + (1-\eta_i) p_i^{1-\rho}}. \tag{A.12}$$

Under the symmetric quantity normalization $N_i = m_i$, equation (A.10) implies

$$s_i^N = \frac{\eta_i^{1/\rho}}{\eta_i^{1/\rho} + (1-\eta_i)^{1/\rho}}, \qquad \eta_i = \left[1 + \left(\frac{1 - s_i^N}{s_i^N}\right)^{\rho}\right]^{-1}. \tag{A.13}$$

This normalization is needed only to recover the structural parameter η_i; the accounting identity itself depends on the observed cost share s_i^N, not separately on η_i.

Empirically, we measure s_i^N from the BEA Use Table as

$$s_i^N = \frac{CE_i}{\text{IntermediateUse}_i + CE_i + \text{Taxes}_i},$$

where CE_i is compensation of employees and Taxes_i denotes taxes on production and imports *net of subsidies* (BEA series V00200). This convention defines labor's share as the share of compensation of employees in the non-surplus production cost.[33]

For any variable X_i, define the annual log change

$$\Delta \log X_i \equiv \log X_{i,t} - \log X_{i,t-1}.$$

The decomposition is first order: all cost and revenue shares are evaluated at the 2017 expansion point and held fixed over time. The Divisia log changes in the quantity and price of the variable composite are therefore

$$\Delta \log V_i = s_i^N \Delta \log N_i + (1 - s_i^N) \Delta \log m_i, \tag{A.14}$$

$$\Delta \log P_i^V = s_i^N \Delta \log W_i + (1 - s_i^N) \Delta \log p_i. \tag{A.15}$$

Taking log changes in equation (A.10) gives

$$\Delta \log m_i = \Delta \log N_i - \rho \left(\Delta \log p_i - \Delta \log W_i\right). \tag{A.16}$$

Substituting into equation (A.14) yields

$$\Delta \log V_i = \Delta \log N_i - (1 - s_i^N) \rho \left(\Delta \log p_i - \Delta \log W_i\right). \tag{A.17}$$

[33] During 2020–2021 the BEA classifies Paycheck Protection Program loan forgiveness as employer subsidies and records it within Taxes_i: compensation of employees registers the payroll that PPP helped firms maintain, while the offsetting fiscal support reduces Taxes_i on the same value-added line. Because our shares are calibrated to the 2017 BEA Detail Supply–Use tables, PPP accounting does not affect the calibration.

Since $C_i = P_i^V V_i$,

$$\begin{aligned}\Delta \log\left(\frac{C_i}{Y_i}\right) &= \Delta \log P_i^V + \Delta \log V_i - \Delta \log Y_i \\ &= \left[s_i^N + (1 - s_i^N)\rho\right] \Delta \log W_i \\ &\quad + (1 - s_i^N)(1 - \rho)\Delta \log p_i - \Delta \log(Y_i/N_i).\end{aligned} \tag{A.18}$$

Define

$$b_i \equiv \rho + (1 - \rho)s_i^N, \qquad d_i \equiv (1 - s_i^N)(1 - \rho). \tag{A.19}$$

Note that $b_i + d_i = 1$ for all parameter values. Industry unit-cost growth can then be written compactly as

$$\Delta \log\left(\frac{C_i}{Y_i}\right) = b_i \Delta \log W_i + d_i \Delta \log p_i - \Delta \log(Y_i/N_i). \tag{A.20}$$

Neither α_i nor Z_i appears separately in equation (A.20). Their effects on the quantity of variable inputs required per unit of output are summarized by observed labor productivity Y_i/N_i, together with the relative-input condition (A.16).

A.1.3 Commodity prices and the production network

The dual price index associated with the materials aggregator (A.3) is

$$\begin{aligned}p_i = \Bigg[&\sum_{j=1}^{J} \zeta_{ij} \omega_{ij} \widetilde{P}_j^{1-\rho_m} \\ &+ \sum_{j=1}^{J} (1 - \zeta_{ij}) \omega_{ij} \left(P_j^{\text{IMP}}\right)^{1-\rho_m} \Bigg]^{1/(1-\rho_m)},\end{aligned} \tag{A.21}$$

where $\widetilde{P}_j$ is the domestic price of commodity j and P_j^{IMP} its import price. Similarly, the dual price index associated with the commodity aggregator (A.4) is

$$\widetilde{P}_j = \left[\sum_{i=1}^{I} \theta_{ij} P_i^{1-\rho_Y}\right]^{1/(1-\rho_Y)}. \tag{A.22}$$

The parameters ω_{ij}, ζ_{ij}, and θ_{ij} are calibrated so that they coincide with the corresponding expenditure and make shares at the 2017 expansion point. First-order log changes in

equations (A.21) and (A.22) are therefore

$$\Delta \log p_i = \sum_{j=1}^{J} \omega_{ij} \zeta_{ij} \Delta \log \widetilde{P}_j + \sum_{j=1}^{J} \omega_{ij} (1 - \zeta_{ij}) \Delta \log P_j^{\text{IMP}}, \tag{A.23}$$

$$\Delta \log \widetilde{P}_j = \sum_{i=1}^{I} \theta_{ij} \Delta \log P_i. \tag{A.24}$$

Thus ρ_m and ρ_Y do not enter the first-order price-index equations. This is the envelope property of CES price indices: at the expansion point, their elasticities with respect to component prices equal the corresponding expenditure or revenue shares. The elasticities ρ_m and ρ_Y continue to govern quantity substitution and second-order terms; they do not disappear from our later model, as it is solved non-linearly. This approximation, however, turns our fairly accurate, as long elasticity ρ is preserved—as is the case here.

In matrix notation, let $\Omega, \zeta \in \mathbb{R}^{J \times I}$, with $\Omega_{ji} = \omega_{ij}$ and $\zeta_{ji} = \zeta_{ij}$, and let $\Theta \in \mathbb{R}^{I \times J}$, with $\Theta_{ij} = \theta_{ij}$. Then

$$\begin{aligned} \Delta \log p = (\zeta \odot \Omega)^\top \Theta^\top \Delta \log P \\ + \left[(\mathbf{1} - \zeta) \odot \Omega \right]^\top \Delta \log P^{\text{IMP}}, \end{aligned} \tag{A.25}$$

where $\odot$ denotes the Hadamard product and $\mathbf{1}$ is a conformable matrix of ones. Under this orientation, $(\zeta \odot \Omega)^\top$ maps domestic commodity prices into industry materials costs, while $\Theta^\top$ maps industry output prices into domestic commodity prices. Imported prices enter directly and do not pass through the domestic make matrix.

A.1.4 The stacked inflation identity

Substituting equation (A.20) into equation (A.7) gives the scalar industry equation

$$\Delta \log P_i = \Delta \log M_i + b_i \Delta \log W_i - \Delta \log(Y_i / N_i) + d_i \Delta \log p_i. \tag{A.26}$$

Define the domestic and imported intermediate-input matrices

$$A \equiv \text{diag}(d)(\zeta \odot \Omega)^\top \Theta^\top, \qquad A^{\text{IMP}} \equiv \text{diag}(d) \left[(\mathbf{1} - \zeta) \odot \Omega \right]^\top. \tag{A.27}$$

Component-wise,

$$A_{ii'} = d_i \sum_{j=1}^{J} \zeta_{ji} \omega_{ji} \theta_{i'j},$$

the CES-corrected share of industry i's variable cost accounted for by domestic intermedi-

ate inputs ultimately produced by supplier industry i'. The imported matrix A^{IMP} takes no make fold: it retains a commodity column index and multiplies the commodity-level import-price vector $\Delta \log P^{\text{IMP}}$ directly.

Stacking equation (A.26) across industries and using equation (A.25) yields the vector pricing identity of equation (3):

$$\begin{aligned}\Delta \log P = \underbrace{\Delta \log M}_{\Delta\text{MARGIN}} + \underbrace{\underbrace{\operatorname{diag}(b)\Delta \log W}_{\Delta\text{WAGE}} + \underbrace{-\Delta \log(Y/N)}_{\Delta\text{PROD}}}_{\Delta\text{ULC}} \\ + \underbrace{A\Delta \log P}_{\Delta\text{MAT}} + \underbrace{A^{\text{IMP}}\Delta \log P^{\text{IMP}}}_{\Delta\text{IMP}}.\end{aligned} \tag{A.28}$$

The structural derivation is a first-order approximation evaluated at fixed 2017 shares. In the empirical implementation, we define $\Delta \log M$ as the residual required to close equation (A.28), so the decomposition holds exactly by construction. The measured margin residual can absorb nonwage costs omitted from measured payroll and materials expenditures, changes in unmeasured tax or cost wedges, timing differences between input and output prices, and measurement error; it should not be interpreted mechanically as net profit or market power.

A.1.5 Fixed point and aggregate decompositions

Moving the domestic-materials term to the left-hand side gives

$$(I - A)\Delta \log P = \Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}. \tag{A.29}$$

Provided the domestic input–output system is productive, the price-equation Leontief inverse exists:

$$L \equiv (I - A)^{-1} = I + A + A^2 + \cdots. \tag{A.30}$$

Industry inflation can then be written as

$$\Delta \log P = L\left(\Delta\text{MARGIN} + \Delta\text{ULC} + \Delta\text{IMP}\right). \tag{A.31}$$

Let c denote the fixed 2017 PCE final-demand vector expressed in commodity space. The corresponding industry weight vector is

$$g \equiv \Theta c, \tag{A.32}$$

and aggregate inflation is

$$\Pi = g^\top \Delta \log P = (\Theta c)^\top \Delta \log P. \tag{A.33}$$

Premultiplying equation (A.28) by $g^\top$ produces the Static decomposition:

$$\Pi = \underbrace{g^\top \Delta \text{MARGIN}}_{\Pi_{\text{MARGIN}}} + \underbrace{g^\top \Delta \text{ULC}}_{\Pi_{\text{ULC}}} + \underbrace{g^\top \Delta \text{MAT}}_{\Pi_{\text{MAT}}} + \underbrace{g^\top \Delta \text{IMP}}_{\Pi_{\text{IMP}}}. \tag{A.34}$$

The unit-labor-cost component is

$$\Pi_{\text{ULC}} = \underbrace{g^\top \operatorname{diag}(b) \Delta \log W}_{\Pi_{\text{WAGE}}} + \underbrace{g^\top \left[-\Delta \log(Y/N) \right]}_{\Pi_{\text{PROD}}}. \tag{A.35}$$

Under this notation a rise in labor productivity makes $\Pi_{\text{PROD}} < 0$, while a decline in labor productivity makes $\Pi_{\text{PROD}} > 0$.

Premultiplying the unrolled identity (A.31) by $g^\top$ gives the Network decomposition:

$$\Pi = \underbrace{g^\top L \Delta \text{MARGIN}}_{\Pi^L_{\text{MARGIN}}} + \underbrace{g^\top L \Delta \text{ULC}}_{\Pi^L_{\text{ULC}}} + \underbrace{g^\top L \Delta \text{IMP}}_{\Pi^L_{\text{IMP}}}. \tag{A.36}$$

The two decompositions answer different accounting questions. The Static decomposition retains observed domestic materials inflation as a separate proximate component and therefore describes the costs directly faced by each industry. The Network decomposition reallocates domestic materials inflation to its upstream margin, unit-labor-cost, and import-price origins across all rounds of the production network. The row vector $g^\top L$ gives the resulting Domar-type weights.

A.1.6 Substitution elasticity ρ

The only substitution elasticity that enters the first-order inflation identity is ρ, the elasticity between labor and the materials composite. We estimate it using BEA–BLS KLEMS data and the log form of the relative-demand condition (A.11),

$$\log\left(\frac{p_i m_i}{W_i N_i}\right) = FE_i + FE_t + (1-\rho)\log\left(\frac{p_i}{W_i}\right) + u_{it}, \tag{A.37}$$

where the industry intercept $\log((1-\eta_i)/\eta_i)$ is absorbed by FE_i and FE_t captures aggregate common factors. The slope on $\log(p_i/W_i)$ identifies $1-\rho$. Our headline estimate is

$$\widehat{\rho} = 0.39, \qquad 95\%\ \text{CI} = [0.06,\ 0.72]. \tag{A.38}$$

This estimate is higher than estimates typically obtained at finer levels of disaggregation (e.g., Atalay, 2017; Boehm et al., 2019). It is therefore a conservative baseline for our purposes: the network mechanism strengthens as ρ falls. The code and data are part of the replication package, with details described in the attached Jupyter notebook.

A.1.7 Comparison to official measures of inflation

Table 9 compares our measure of inflation to CPI, headline PCE, and Core PCE.

Table 9: Cumulative inflation from 2019 (log-%, pp): paper's measure vs. BEA and BLS published indices.

	Headline measures			Ex-volatile measures	
Year	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 Θ. **BEA PCE** is the BEA-published Personal Consumption Expenditures price index (FRED `PCEPI`). **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`).

A.2 Supply-chain construction

This appendix specifies the spectral-bisection algorithm that partitions the 341 endogenous BEA detail industries into four supply chains and twelve chain × finalness tiers, and the within-tier hit/safe refinement that produces the 24 production sub-units used in the main text. A companion spreadsheet (`chain_classification.xlsx`) lists every BEA detail industry together with its assigned chain, tier, and hit/safe fraction.

A.2.1 Spectral-bisection algorithm

Let A be the 341×341 structural intermediate-input operator constructed from the 2017 BEA Use and Make tables as defined in Section 2.1,

$$A \;=\; \operatorname{diag}(d)\,(\zeta \odot \Omega)^{\top}\,\Theta^{\top},$$

where entry $A_{ii'}$ is the share of one dollar of industry-i's output paid to domestic intermediates ultimately sourced from supplier industry i'. We use A itself as the supplier-reach kernel for chain construction: this is the same network operator that drives the price decomposition in Section 2.1, so chain construction and price propagation share a single object.

We symmetrize A,

$$\tilde{A} \;=\; \tfrac{1}{2}\big(A + A^{\top}\big),$$

zero its diagonal, and form the normalized graph Laplacian $\mathcal{L} = I - D^{-1/2}\,\tilde{A}\,D^{-1/2}$, where D is the diagonal matrix of $\tilde{A}$ row sums. The Fiedler vector f—the eigenvector corresponding to the second-smallest eigenvalue of $\mathcal{L}$—encodes the bisection that minimizes weighted trade across the cut. Industries are sorted by f_i, and the cut is placed at the gross-output-balanced midpoint: the threshold $f^{\star}$ at which $\sum_{i:\, f_i \le f^{\star}} \mathrm{GO}_i$ first crosses $\frac{1}{2}\sum_i \mathrm{GO}_i$. The two halves therefore have approximately equal economic weight rather than equal industry count. Recursive application yields the four chains $C_1, \ldots, C_4$, sorted by descending gross output and containing 44, 80, 98, and 119 industries respectively.

Within each chain, industries are tiered by *finalness*—as defined in text (direct PCE weight to Domar weight). The top tier (h) contains industries with finalness at or above 70%; the remaining industries are split at the gross-output-weighted median of finalness into a middle tier (m, higher finalness) and an upstream tier (l, lower finalness). The combination yields the twelve chain $\times$ tier production cells of the model.

A.2.2 Hit/safe refinement within each cell

Each tier is further split into a *hit* node and a *safe* node on the basis of the 2020 employment collapse. Within each tier, industries are sorted by $\Delta \log N_{i,\,2019\to2020}$ in ascending order. The industry whose cumulative 2019 employment first crosses 50% fixes the threshold; industries at or below it are classified *hit*, the remainder *safe*.

In seven of the twelve tiers ($C_{1,h}$, $C_{1,l}$, $C_{3,h}$, $C_{3,l}$, $C_{4,h}$, $C_{4,m}$, $C_{4,l}$), a single BEA-detail line dominates the tier by employment, so the median rule yields a hit-employment share outside assumed tolerance band [40%, 60%]. For these triggered tiers we apply the same

sort-and-cut rule at NAICS-6 granularity using QCEW BEA-357 × NAICS-6 employment data, then aggregate back to BEA codes via a convex combination weighted by 2019 wagebill of each constituent NAICS-6 industry. The procedure is as follows:

1. Pool the NAICS-6 sub-industries of every BEA in the cell.
2. Sort the pool by $\Delta \log N^{\text{NAICS-6}}_{j,\,2019\to2020}$ ascending; cumulative 2019 NAICS-6 employment $\leq 50\%$ of the tier total is classified as *hit*, the remainder as *safe*.
3. For each BEA i in the tier, we compute the within-BEA hit fraction as

$$\text{hitfrac}_i \;=\; \frac{\sum_{j\in\text{NAICS-6}(i)\,\cap\,\text{hit}} N_{j,\,2019}}{\sum_{j\in\text{NAICS-6}(i)} N_{j,\,2019}} \;\in\; [0,1].$$

4. BEA i contributes $\text{hitfrac}_i \cdot \{Y, N, \text{wagebill}, \omega\text{-row}, \omega\text{-col}\}_i$ to the hit node and $(1 - \text{hitfrac}_i)$ of the same to the safe node.

A.2.3 Input-flow heatmap

Figure 9 shows the obtained 24×24 intermediate-input heatmap, where node (i, j) is the share of destination sub-unit j's intermediate-input bill sourced from i; darker shades indicate stronger linkages. It complements the chain-level Sankey diagram in the main text by resolving the within-chain × tier × hit/safe structure of trade. The four denser blocks along the diagonal are the visual signature of the chain construction: trade clusters within chains, exactly as the within-chain shares quantify. The lighter but non-empty off-diagonal entries are cross-chain spillovers, which the model carries through at noticeably lower intensity than the within-chain flows.

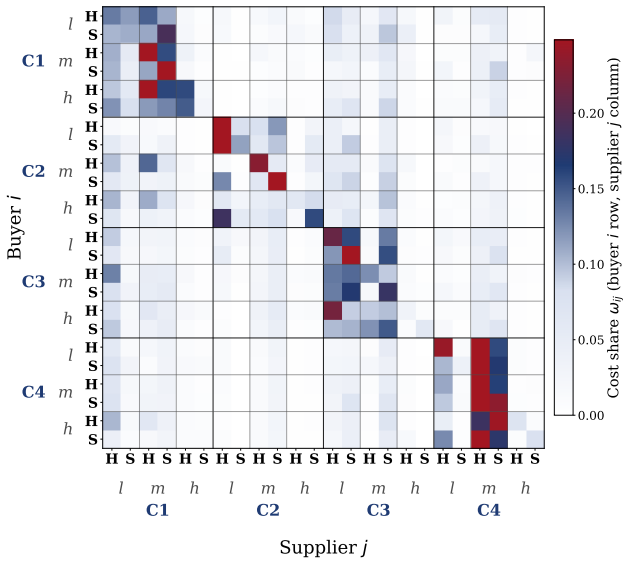

Figure 9: Intermediate-input flows across the 24 production nodes.

Notes: Cells show input shares aggregated from the 341 endogenous BEA-NAICS industries to the chain × tier × 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.

A.3 JEM exposure of BEA detail-level supply chains

This appendix details the construction of the industry-level COVID exposure index E_i used in the cross-industry reduced-form analysis of Section 2 (equation (1)). The construction proceeds in three stages. We also discuss the population-level validation of the Danish country panel of the underlying Job-Exposure Matrix and report the weak-instrument diagnostic battery for the data–model IV moment that appears in Section 5.

A.3.1 Construction

The COVID-19 Job-Exposure Matrix (Oude Hengel et al., 2022) 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, each rated on a 0–3 scale of infection risk, grouped into three blocks: *transmission risk* (Number of contacts, Nature of contacts, Contaminated workspaces, Location), *mitigation measures* (Social distance, Face Covering), and *precarious work* (Income insecurity, Migrants). We retain the first two blocks — the six workplace-exposure dimensions — and drop the two precarious-work dimensions, which measure worker financial vulnerability rather than physical contact risk and would confound the "physical exposure" interpre-

tation of E_i. Because the underlying epidemiological evidence and working-environment conditions differ modestly across the three country panels, we take the unweighted cross-country mean of scores per ISCO code as the headline; the Danish panel alone has been validated against a population-scale infection register (see *Validation against the Danish population register* below), and this validation supplies the epistemological support for the JEM signal even when the cross-country mean is used.

The index E_i is built bottom-up in three stages, all from pre-pandemic inputs—no realized 2020 or post-2020 outcome enters. The construction of the industry-level index has three steps:

Step 1: Occupation-level exposure via first principal component. *Source:* `Supplementary file COVID-19 JEM.xlsx` released with Oude Hengel et al. (2022), Denmark, Netherlands, and UK country sheets, one row per ISCO-08 4-digit occupation with eight numerical columns. We retain the first six columns — the four transmission-risk dimensions (Number of contacts, Nature of contacts, Contaminated workspaces, Location) and the two mitigation-measure dimensions (Social distance, Face Covering) — and drop the last two precarious-work columns (Income insecurity, Migrants), which do not describe physical contact risk. Within each country panel we average duplicate ISCO codes; we then take the unweighted cross-country mean of scores per ISCO code to obtain a 436×6 occupation-by-dimension matrix.

We aggregate the six dimensions into a single occupation-level exposure score using the *first principal component* (PC1) of the six-dimension matrix rather than an equal-weighted mean, so that the weighting is determined by the covariance structure of the JEM data rather than by an ad-hoc choice. We standardize each of the six columns (mean 0, unit variance across ISCOs), form the 6×6 correlation matrix, and take the eigenvector associated with the largest eigenvalue. The sign convention is fixed so that all loadings are positive (higher score = higher exposure). PC1 explains 73.5% of the variance across the six dimensions; loadings are ≈ 0.44 for the five transmission-risk and social-distance dimensions and 0.29 for face covering — reflecting that face covering has the lowest average correlation with the other exposure signals (0.4–0.6 versus 0.7–0.85 for the other pairs). For interpretability, we apply an affine rescaling so the resulting PC1 score has the same mean and standard deviation as the corresponding equal-weighted mean of the six dimensions; the rescaling is a monotone transform of PC1 and leaves the occupation-level exposure on the original $[0, 3]$ ordinal scale of the underlying JEM scores. The PC1-based score has a Pearson correlation of 0.995 with the equal-weighted mean across ISCOs, so downstream regression coefficients shift by less than 1%; the composite is defended by a

data-driven weighting instead of an arbitrary uniform one. Formally:

$$\text{exposure}_o \;=\; \alpha \;+\; \beta \cdot \left(v^\top \tilde{\text{score}}_o\right), \qquad v \;\in\; \arg\max_{\|v\|=1} v^\top \Sigma v, \qquad \text{exposure}_o \in [0,3],$$

where $\tilde{\text{score}}_o \in \mathbb{R}^6$ is the standardized six-dimension score vector for ISCO occupation o, Σ is the 6×6 correlation matrix across the ISCO cross-section, v is the PC1 eigenvector, and (α, β) is the affine rescaling to the original $[0, 3]$ exposure scale.

The work by van der Feltz et al. (2023) validate the Danish panel of the JEM against population-scale individual-level SARS-CoV-2 infection records. As noted, four features of the Danish data make this panel more suitable than the other two (UK and especially NL):

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. *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. 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. *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.

Step 2: ISCO-08 → SOC-2018 mapping. The OEWS data used in Stage 3 are 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 exposure_o across all ISCO-08 codes that map to s through the chained crosswalk. The

resulting score covers 865 SOC-2018 codes.

Step 3: Aggregation to BEA-357 industries. We use the OEWS occupation-employment matrix at the BEA-357 level (a 345×789 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,

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

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 $E_i = 1.32$, standard deviation 0.33, and range [0.34, 2.31]. Top-exposed industries are grocery and food retail, 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. Agriculture sectors are not included, and we drop them from the analysis.

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 → SOC-2010 and SOC-2010 → SOC-2018) predate 2020. The OEWS occupation matrix uses pre-COVID employment shares (vintage ≤ 2019). The exposure index E_i is fully determined by industry i's pre-COVID occupation composition crossed with intrinsic occupation-level transmission risk.

A.4 Additional Model Details

This appendix provides the price indices and conditional input demands associated with the production structure in Section 3.1, and derives the nonlinear Rotemberg pricing equation of Section 3.3. Market-clearing conditions, aggregate-output construction, MPC-driven final demand, and the aggregate price index appear in Section 3.4.

A.4.1 Price indices and conditional demands

Within each tier i, hit and safe output are combined according to equation (12). The corresponding unit-expenditure function gives the tier-level price index

$$P_{i,t} = \left[\sum_{j \in \{H,S\}} s_{ij} \, P_{ij,t}^{1-\rho} \right]^{1/(1-\rho)}, \tag{A.39}$$

where $\sum_{j \in \{H,S\}} s_{ij} = 1$. Conditional demand for the output of node (i, j) is

$$Y_{ij,t} = s_{ij} \left(\frac{P_{ij,t}}{P_{i,t}} \right)^{-\rho} Y_{i,t}. \tag{A.40}$$

Thus, relative price movements reallocate demand between the hit and safe nodes within a tier, with elasticity ρ.

Similarly, cost minimization over the materials bundle in equation (13) gives the intermediate-input price index faced by nodes in tier i:

$$p_{i,t}^{m} = \left[\sum_{k=1}^{12} \omega_{i,k} \, P_{k,t}^{1-\rho} + \omega_{i}^{\mathrm{EXO}} \left(P_{i,t}^{\mathrm{EXO}} \right)^{1-\rho} \right]^{1/(1-\rho)}, \tag{A.41}$$

where

$$\sum_{k=1}^{12} \omega_{i,k} + \omega_{i}^{\mathrm{EXO}} = 1.$$

The resulting conditional demands for inputs sourced from endogenous tier k and from the EXO block are

$$\begin{aligned} m_{ij,k,t} &= \omega_{i,k} \left(\frac{P_{k,t}}{p_{i,t}^{m}} \right)^{-\rho} m_{ij,t}, \\ m_{ij,t}^{\mathrm{EXO}} &= \omega_{i}^{\mathrm{EXO}} \left(\frac{P_{i,t}^{\mathrm{EXO}}}{p_{i,t}^{m}} \right)^{-\rho} m_{ij,t}. \end{aligned} \tag{A.42}$$

Because the input–output weights vary across purchasing tiers, $p_{i,t}^{m}$ summarizes the supplier-price pressures relevant for tier i. Consequently, a price increase originating in one part of the network affects downstream marginal costs according to both the direct input share and the possibilities for substitution across suppliers.

A.4.2 Rotemberg pricing equation

This subsection derives the nonlinear Rotemberg pricing equation underlying the price block of Section 3.3 (equation (18)) from a fully specified monopolistically-competitive firm problem. The derivation makes explicit the exact nonlinear term $\Pi_{ij,t}\left(\Pi_{ij,t} - \bar{\Pi}\right)$ that arises because the firm chooses a price *level* and inflation enters the adjustment cost. We also derive its asymmetric downward-rigidity counterpart and show that a first-order approximation around the zero-net-inflation steady state recovers the familiar local New Keynesian Phillips curve.

A.4.2.1 Variety aggregator and demand

Each production node (i, j) is a composite of a continuum of monopolistically-competitive varieties indexed by $\ell \in [0, 1]$. The node composite good is

$$Y_{ij,t} \;=\; \left[\int_0^1 y_{ij,t}(\ell)^{(\varepsilon_i-1)/\varepsilon_i}\, d\ell\right]^{\varepsilon_i/(\varepsilon_i-1)}, \tag{A.43}$$

with associated price index

$$P_{ij,t} \;=\; \left[\int_0^1 p_{ij,t}(\ell)^{1-\varepsilon_i}\, d\ell\right]^{1/(1-\varepsilon_i)}. \tag{A.44}$$

Cost minimization by downstream users delivers the standard constant-elasticity demand for variety ℓ,

$$y_{ij,t}(\ell) \;=\; \left(\frac{p_{ij,t}(\ell)}{P_{ij,t}}\right)^{-\varepsilon_i} Y_{ij,t}. \tag{A.45}$$

$Y_{ij,t}$ is the object that enters the tier level aggregator. Thus ε_i governs substitution across varieties *within* node (i, j), while ρ governs substitution across hit and safe nodes and across inputs in the rest of the model.

A.4.2.2 Rotemberg adjustment cost

Firms face a quadratic cost of changing prices relative to the nominal anchor $\bar{\Pi}$, with an additional asymmetric piece active only when inflation falls *below* that anchor (downward rigidity):

$$\Gamma_i(\Pi) \;=\; \tfrac{\bar{\phi}_i}{2}\left(\Pi - \bar{\Pi}\right)^2 \;+\; \tfrac{\lambda_i}{2}\left[\left(\Pi - \bar{\Pi}\right)^-\right]^2, \qquad x^- \equiv \min\{x, 0\}, \tag{A.46}$$

with the derivative given by

$$\Gamma_i'(\Pi) \;=\; \bar{\phi}_i\big(\Pi - \bar{\Pi}\big) \;+\; \lambda_i\big(\Pi - \bar{\Pi}\big)^{-}. \tag{A.47}$$

A.4.2.3 Firm problem

A firm from hit/safe sector (i, j) chooses price $p_{ij,t}(\ell)$ taking $P_{ij,t}$, $Y_{ij,t}$, the real marginal cost $MC_{ij,t}$, and the residual and structural-2020 pricing wedges $\big(\mu^e_{i,t},\ \bar{\mu}^e_i\big)$ as given. Let firm-level gross inflation be

$$\Pi_{ij,t}(\ell) \;\equiv\; \frac{p_{ij,t}(\ell)}{p_{ij,t-1}(\ell)}. \tag{A.48}$$

The firm's recursive problem is given by

$$\begin{aligned} V_{ij,t}\big(p_{ij,t-1}(\ell)\big) = \max_{p_{ij,t}(\ell)} \Bigg\{ & \left[\frac{p_{ij,t}(\ell)}{P_{ij,t}} - MC_{ij,t} - \mu^e_{i,t} - \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\right] y_{ij,t}(\ell) \\ & - \Gamma_i\big(\Pi_{ij,t}(\ell)\big)\, Y_{ij,t} \;+\; \beta\, \mathbb{E}_t\, V_{ij,t+1}\big(p_{ij,t}(\ell)\big) \Bigg\}, \end{aligned} \tag{A.49}$$

subject to the demand curve (A.45). The adjustment cost is measured in units of the node composite and scaled by node demand $Y_{ij,t}$; this is the standard Rotemberg convention that delivers the nonlinear Phillips curve below. The residual wedge $\mu^e_{i,t}$ and the structural-2020 wedge $\bar{\mu}^e_i$ enter additively: a positive realization raises the wedge between price and marginal cost the firm targets, and therefore raises inflationary pressure.

A.4.2.4 First-order conditions

Differentiating (A.49) and using the demand derivative

$$\frac{\partial y_{ij,t}(\ell)}{\partial p_{ij,t}(\ell)} \;=\; -\varepsilon_i\, \frac{y_{ij,t}(\ell)}{p_{ij,t}(\ell)}, \tag{A.50}$$

the derivative of current inflation with respect to the current price,

$$\frac{\partial \Pi_{ij,t}(\ell)}{\partial p_{ij,t}(\ell)} \;=\; \frac{1}{p_{ij,t-1}(\ell)} \;=\; \frac{\Pi_{ij,t}(\ell)}{p_{ij,t}(\ell)},$$

and the envelope effect of today's price on next period's gross inflation,

$$\frac{\partial \Pi_{ij,t+1}(\ell)}{\partial p_{ij,t}(\ell)} \;=\; -\frac{p_{ij,t+1}(\ell)}{p_{ij,t}(\ell)^2} \;=\; -\frac{\Pi_{ij,t+1}(\ell)}{p_{ij,t}(\ell)},$$

the first-order condition evaluated at a symmetric equilibrium, $p_{ij,t}(\ell) = P_{ij,t}$, $y_{ij,t}(\ell) = Y_{ij,t}$, and $\Pi_{ij,t}(\ell) = \Pi_{ij,t} \equiv P_{ij,t}/P_{ij,t-1}$, gives

$$\begin{aligned} 0 \;=\; & \left[1 - \varepsilon_i + \varepsilon_i\Big(MC_{ij,t} + \mu^e_{i,t} + \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\Big)\right] Y_{ij,t} \\ & - \Gamma_i'\big(\Pi_{ij,t}\big)\,\Pi_{ij,t}\,Y_{ij,t} \;+\; \beta\,\mathbb{E}_t\left[\Gamma_i'\big(\Pi_{ij,t+1}\big)\,\Pi_{ij,t+1}\,Y_{ij,t+1}\right]. \end{aligned} \tag{A.51}$$

Rearranging terms delivers the nonlinear Rotemberg Cost-Price Equation:

$$\begin{aligned} & \left[\bar{\phi}_i\big(\Pi_{ij,t} - \bar{\Pi}\big) + \lambda_i\big(\Pi_{ij,t} - \bar{\Pi}\big)^-\right]\Pi_{ij,t}\,Y_{ij,t} \\ & \quad = \; \varepsilon_i\Big(MC_{ij,t} - \tfrac{\varepsilon_i - 1}{\varepsilon_i} + \mu^e_{i,t} + \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\Big) Y_{ij,t} \\ & \qquad + \; \beta\,\mathbb{E}_t\left\{\left[\bar{\phi}_i\big(\Pi_{ij,t+1} - \bar{\Pi}\big) + \lambda_i\big(\Pi_{ij,t+1} - \bar{\Pi}\big)^-\right]\Pi_{ij,t+1}\,Y_{ij,t+1}\right\}. \end{aligned} \tag{A.52}$$

Defining the auxiliary function

$$\Psi_i(\Pi) \;\equiv\; \left[\bar{\phi}_i\big(\Pi - \bar{\Pi}\big) + \lambda_i\big(\Pi - \bar{\Pi}\big)^-\right]\Pi, \tag{A.53}$$

and note that equation (A.52) can be written as

$$\Psi_i(\Pi_{ij,t})\,Y_{ij,t} \;=\; \varepsilon_i\Big(MC_{ij,t} - \tfrac{\varepsilon_i - 1}{\varepsilon_i} + \mu^e_{i,t} + \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\Big) Y_{ij,t} \;+\; \beta\,\mathbb{E}_t\left[\Psi_i(\Pi_{ij,t+1})\,Y_{ij,t+1}\right]. \tag{A.54}$$

Dividing by $Y_{ij,t}$ delivers the per-unit form, in which the future Rotemberg term carries the output-scaling ratio $Y_{ij,t+1}/Y_{ij,t}$:

$$\Psi_i(\Pi_{ij,t}) \;=\; \varepsilon_i\Big(MC_{ij,t} - \tfrac{\varepsilon_i - 1}{\varepsilon_i} + \mu^e_{i,t} + \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\Big) \;+\; \beta\,\mathbb{E}_t\left[\Psi_i(\Pi_{ij,t+1})\,\frac{Y_{ij,t+1}}{Y_{ij,t}}\right]. \tag{A.55}$$

The key nonlinear correction relative to a standard log-linearized Rotemberg condition is $\Pi_{ij,t}$ term multiplying inflation gap $\Pi_{ij,t} - 1$. The intuition is this: The future term enters because a higher price today mechanically lowers next period's gross inflation $\Pi_{ij,t+1}$, making future price adjustment cheaper; expected future adjustment costs therefore feed into the current pricing decision. In the special case $\bar{\Pi} = 1$, the current-period adjustment-cost term reduces to $\bar{\phi}_i\,\Pi_{ij,t}(\Pi_{ij,t} - 1) + \lambda_i\,\Pi_{ij,t}(\Pi_{ij,t} - 1)^-$, and the same object appears inside the conditional expectation at $t + 1$.

The marginal-cost term is the firm's desired inverse-markup gap. Absent adjustment costs, downward rigidity, and wedges, the firm would choose the static monopolistic-competition markup: $MC_{ij,t} = (\varepsilon_i - 1)/\varepsilon_i$. When marginal cost rises above this desired inverse markup, firms want to raise prices. With Rotemberg they do so gradually.

Because the model is solved as a deterministic perfect-foresight transition, the future quantity $Y_{ij,t+1}$ that enters equation (A.54) is evaluated at the model-implied path. Inflation expectations, however, need not coincide with model-realized future inflation; per the calibrated expectations operator of Section 3.3 (equation (22)), the price-setter forms a convex combination of a model-consistent, a backward-looking, and a target-anchored component. The subjective pricing-expectations operator is thus equivalently given by (separable):

$$\mathbb{E}_t^s \left[\Psi_i(\Pi_{ij,t+1})\, Y_{ij,t+1} \right] \;\equiv\; \Psi_i\big(\Pi^e_{t+1|t}\big)\, Y_{ij,t+1}, \tag{A.56}$$

where

$$\Pi^e_{t+1|t} \;=\; \nu_{\text{forward}}\, \Pi^{\text{model}}_{t+1} \;+\; \nu_{\text{past}}\, \Pi_t \;+\; \nu_{\text{target}}\, \bar{\Pi}, \qquad \nu_{\text{forward}} + \nu_{\text{past}} + \nu_{\text{target}} = 1, \tag{A.57}$$

and Π^{model}_{t+1} denotes the model-consistent value of next-period aggregate inflation along the perfect-foresight path. When $\nu_{\text{forward}} = 1$ and $\nu_{\text{past}} = \nu_{\text{target}} = 0$, the operator collapses to the standard nonlinear perfect-foresight Rotemberg term,

$$\mathbb{E}_t^s \left[\Psi_i(\Pi_{ij,t+1})\, Y_{ij,t+1} \right] \;=\; \Psi_i(\Pi_{ij,t+1})\, Y_{ij,t+1}.$$

Substituting (A.56) into (A.54) gives

$$\Psi_i(\Pi_{ij,t})\, Y_{ij,t} \;=\; \varepsilon_i\Big(MC_{ij,t} - \tfrac{\varepsilon_i - 1}{\varepsilon_i} + \mu^e_{i,t} + \bar{\mu}^e_i \cdot \mathbf{1}\{t = 2020\}\Big) Y_{ij,t} \;+\; \beta\, \Psi_i\big(\Pi^e_{t+1|t}\big)\, Y_{ij,t+1}, \tag{A.58}$$

with $\Psi_i(\cdot)$ as defined in equation (A.53).

A.5 Calibration of structural wedges

Let $g_{X,i,t} \equiv \Delta \log X_{i,t} = \log X_{i,t} - \log X_{i,t-1}$, so that $g_{X,i,2020}$ pertains to the 2019→2020 log change. The baseline is *saturated* in the sense of Chari et al. (2007): the exogenous wedges $(\xi^e_{i,t},\ a^e_{i,t},\ \tau^e_{ij,t},\ \bar{\mu}^e_i,\ \mu^e_{i,t})$ are recovered so that the model reproduces, tier by tier, the observed paths of output, employment, the within-chain dispersion of the 2020 collapse, and prices. The 2020 collapse year is treated separately from the 2021–2024 recovery via a dedicated structural markup wedge, $\bar{\mu}^e_i$. We show that all wedges can be locally point-

identified. This result implies that the Newton minimization of the data-model distance is well-defined and does not contain flat directions around the calibrated solution.

We consider first-order approximation: log-linearized structural equations around the calibrated steady state. Because the full model is nonlinear, the result establishes only *local* identification. Global identification is a numerical matter and is supported by convergence of the calibration algorithm.

Step 1: Demand wedge → output. Log-linearizing market clearing (equation (23)) around the steady state gives

$$g_{Y,i,t} = \theta_i^D \, g_{\xi^e,i,t} + \Xi_i\big(\mathbf{g}_{P,t}\big),$$

where $\theta_i^D \equiv D_i^{ss}/Y_i^{ss}$ is tier i's steady-state final-demand share, $\mathbf{g}_{P,t} \equiv (g_{P,k,t})_{k=1}^{12}$ is the vector of tier-level price growth rates (using the bold-vector convention introduced in Section 2), and $\Xi_i(\cdot)$ collects the input–output and substitution terms that this price vector contributes to tier i's output. Holding prices fixed, the demand wedge $g_{\xi^e,i,t}$ maps one-to-one into $g_{Y,i,t}$, so matching the observed output path identifies $\xi^e_{i,t}$. The income-propagation term in equation (24) vanishes at the calibrated solution by construction, since the model's aggregate output matches the data and the bracket $(\log \mathcal{Y}_t - \log \mathcal{Y}_t^{\text{data}})$ is zero. The MPC channel is irrelevant for identifying ξ^e; it matters only in counterfactuals, where the decomposition distinguishes wedge attribution from structural forces.

Step 2: Productivity vs. labor-wedge indeterminacy. With $g_{Y,i,t}$ pinned down and capital predetermined (so $g_{\bar{K}} = 0$ in equation (10), with capital investment treated as a sunk cost incurred in previous periods), the production identity reads

$$g_{Y,i,t} = a^e_{i,t} + (1-\alpha_i)\, g_{V,i,t}, \qquad g_{V,i,t} = ls_i \, g_{N,i,t} + (1-ls_i)\, g_{m,i,t},$$

where we suppress the within-tier j index and use the tier-average capital share α_i and labor share ls_i. Under the linearization this aggregation is exact to first order, with α_i common within tier and η_{ij} differing across hit and safe nodes. Here V is the labor–materials bundle of equation (11) and $ls_i = W_i N_i / C_i$ is labor's steady-state share of variable cost (the bundle's labor share). The labor–materials first-order condition (equation (A.16)) pins down the input ratio:

$$g_{m,i,t} - g_{N,i,t} = \rho\big(g_{w^e,i,t} - g_{p^m,i,t}\big) = \rho\, \bar{\tau}^e_{i,t} + (\text{terms not involving } \tau^e),$$

where ρ is the common CES elasticity and the wedge enters through the effective wage, $\log w^e_{ij,t} = \log \tilde{w}_{ij} + \log W_t + \log(1 + \mathcal{W}_{ij,t}) + \tau^e_{ij,t}$ (equation (16)). In $t = 2020$ rebuilding is absent ($\mathcal{W} \approx 0$), so the node-specific term in g_{w^e} is $\tau^e_{ij,t}$ itself. We write $\bar{\tau}^e_{i,t} \equiv s_{iH}\, \tau^e_{iH,t} + s_{iS}\, \tau^e_{iS,t}$ for the output-share-weighted tier mean, with s_{ij} the within-tier output shares of equation (12). Substituting back, employment is linear in both wedges, implying

$$g_{N,i,t} = -\left(\tfrac{1}{1-\alpha_i}\, a^e_{i,t} + \kappa_\tau\, \bar{\tau}^e_{i,t}\right) + \text{const}, \qquad \kappa_\tau \equiv (1 - ls_i)\,\rho,$$

where "const" collects terms independent of (a^e, τ^e) and the negative sign reflects that τ^e enters w^e as a labor tax (a higher wedge substitutes activity away from labor, lowering N). Because a^e and $\bar{\tau}^e$ enter additively and linearly, for a given output path the 2020 employment target $g_{N,i,2020}$ pins down only the combination $\frac{1}{1-\alpha_i}\, a^e_{i,2020} + \kappa_\tau\, \bar{\tau}^e_{i,2020}$, not its two parts. The tier-mean wedge is therefore redundant: its level is absorbed into $a^e_{i,2020}$, which the employment target pins on its own. We therefore set $\bar{\tau}^e_{i,2020} \equiv 0$ without loss of generality.

Step 3: Hit/safe contrast pins down the collapse geometry. What survives is the within-tier differential wedge,

$$\Delta\tau^e_i \equiv \tau^e_{iS,2020} - \tau^e_{iH,2020}.$$

By construction $\Delta\tau^e_i$ is orthogonal to the (a^e, ξ^e) block: it redistributes g_N between the hit and safe nodes of tier i while leaving the tier mean—and hence a^e_i and $g_{Y,i}$—unchanged. Aggregated across the three tiers of a chain, these contrasts generate the within-chain dispersion of the 2020 collapse, which we measure at NAICS-6 detail (aggregated to split the tier 50–50) by

$$\sigma^-_c \equiv \text{std}\left[\min\{0,\ g_{N,u,2020}\}\right]_{u \in c}.$$

This is the downward employment-collapse dispersion across the disaggregated units u of chain c (the right-panel object in Figure 2). In the calibration, it enters as excess downward dispersion relative to normal times (2017–2019), denoted $\Delta\sigma^-_c$. We tie the per-tier contrasts to a single chain-level amplifier χ_c through a data-fixed profile,

$$\Delta\tau^e_i = \chi_{c(i)}\ g_i^{\text{QCEW}},$$

where g_i^{QCEW} is the per-tier hit–safe collapse-excess shape taken directly from the industry employment data and held fixed in the calibration. With one scalar χ_c per chain and one moment σ^-_c per chain, the system $(\chi_c)_{c=1}^4$ is exactly identified by $\{\Delta\sigma^-_c\}_{c=1}^4$; we solve it by Newton iteration.

Step 4: Structural 2020 markup wedge. The 2020 collapse year receives separate treatment from the 2021–2024 recovery: the supply-chain rebuilding channel is dormant ($\mathcal{W} \approx 0$), the wall mechanism has not yet activated, and cost-side pass-through to prices reflects the one-time response to the collapse rather than the persistent dynamics that follow. We therefore introduce a per-tier *structural* markup wedge $\bar{\mu}_i^e$, active only at $t = 2020$, entering the pricing block (equation (18)) additively alongside the residual markup $\mu_{i,t}^e$:

$$\text{NKPC RHS}_{i,t} = \varepsilon_{c(i)}\left(\text{MC}_{i,t} - \frac{\varepsilon_{c(i)} - 1}{\varepsilon_{c(i)}} + \mu_{i,t}^e + \bar{\mu}_i^e \cdot \mathbf{1}\{t = 2020\}\right) Y_{i,t}.$$

With output, employment, and the collapse geometry already pinned down by Steps 1–3, the structural wedge $\bar{\mu}_i^e$ is the only remaining degree of freedom at 2020, and is recovered by closing the 2020 tier-composite price target $g_{P,i,2020}^{\text{data}}$. We recover it via a damped fixed-point map driven by the NKPC sensitivity $\varepsilon_{c(i)}/\bar{\phi}_{c(i)}$ and bound it symmetrically, $\bar{\mu}_i^e \in [-1, 1]$. The bounds rarely bind in the calibrated solution; convergence is monotone within about 10 iterations.

Step 5: Residual markup path, 2021–2024. With $(\xi^e, a^e, \chi, \bar{\mu}_i^e)$ matched to output, employment, the collapse dispersion, and the 2020 price target, the 2021–2024 price paths $g_{P,i,t}$ are determined by the pricing block. Any remaining discrepancy relative to observed prices is closed by the per-tier residual markup $\mu_{i,t}^e$, which is recovered in a second pass and iterated to convergence. In the saturated model $\mu_{i,t}^e$ therefore measures the component of the 2021–2024 price path the structural block does not explain—the interpretation used in the main text and in footnote 23 of Section 4. By construction $\mu_{i,2020}^e$ is not pinned by data (the 2020 target is absorbed by $\bar{\mu}_i^e$ in Step 4) and converges to a value near zero in the calibrated solution, confirming the clean separation of the two markup wedges.

A.6 Detrending of data inputs

All calibration targets are constructed as excess of a single pre-COVID 2017–2019 log-linear trend, applied uniformly at the aggregate level to every tier and hit/safe node. We chose the pre-COVID window so that none of the post-2020 damage is absorbed into the trend. An alternative tier-own detrending is emitted alongside as robustness (in `tier_detrended_targets_own.csv`) and yields quantitatively similar results.

Tier-level measurement. Real gross output Y_i is BEA GO in 2017\$; N_i is QCEW average annual employment. We construct N for 2020 from the Jan–Feb 2020 to Jan–Feb 2021

growth rate rather than the 2019–2020 annual averages, so the transient March–April 2020 layoffs that were largely reversed by year-end do not contaminate the persistent damage measure. Tiers are indexed by i and chains by c throughout.

Aggregate trend construction. For each variable $X \in \{Y, N, P_{GO}, Y/N, W, WN\}$ we compute a single economy-wide annual log growth rate over 2017–2019,

$$g_X^{\text{trend}} \;=\; \tfrac{1}{2}\log\!\left(X_{2019}^{\text{agg}} \,/\, X_{2017}^{\text{agg}}\right),$$

where the aggregate is the sum across all 24 endogenous tiers: for Y and N the raw sum; for P_{GO} the implicit deflator $\sum_i P_i^{GO}(t)\, Y_i(t) / \sum_i Y_i(t)$; for W the ratio of the economy-wide wage bill to total employment; for WN the sum of tier wage bills; and Y/N is computed as $g_Y^{\text{trend}} - g_N^{\text{trend}}$. Each tier's calibration target for year t is the tier's own annual log growth minus this aggregate trend. Multi-year excess uses two- or three-year cumulative log changes minus the corresponding multiple of the annual trend.

Applied aggregate trend rates. Computed from the current BEA vintage on the 2017–2019 window:

	Y (GO)	N	P_{GO}	Y/N	W	WN
Annual trend (%/yr)	+2.55	+1.66	+1.65	+0.89	+3.44	+5.10

Exogenous price paths. The exogenous-block paths are detrended on their own pre-COVID 2017–2019 rate, the same convention applied to endogenous tier prices:

$$\varepsilon_e(t) = \log P_e(t) - \log P_e(t-1) - g_e^{\text{pre}}, \qquad g_e^{\text{pre}} = \tfrac{1}{2}\log\!\big(P_e(2019)/P_e(2017)\big),$$

with $g_{RE}^{\text{pre}} = +2.83\%$, $g_{FG}^{\text{pre}} = +2.81\%$, and $g_{EN}^{\text{pre}} = +1.51\%$ per year, corresponding to the real-estate, finance-and-government, and energy blocks. The per-tier import-price shocks $\varepsilon_{i,t}^{IMP}$ are built the same way from BLS import-price indices aggregated to the tier basis. The shocks are zeroed for $t > 2024$, holding the exogenous price level flat over the model's 45 post-window periods; because these prices follow a random walk, that level shift is permanent.

A.7 Rotemberg price stickiness parameter calibration

We translate the monthly BLS price-change frequencies of Pasten et al. (2020) into the annual Rotemberg adjustment cost $\bar{\phi}$ by matching one-year pass-through of a permanent cost shock across the two protocols.

A.7.1 Mapping problem

Under monthly Calvo with reset probability f_m, a permanent unit cost shock passes through as $p_t = (1-f_m)\,p_{t-1} + f_m$, so cumulative twelve-month pass-through is $1-(1-f_m)^{12}$. This presumes a permanent, fully-realized shock (constant frictionless target) and no strategic complementarity, so the reset price equals the target each period and the Calvo root is exactly $1 - f_m$.

The annual Rotemberg block adjusts as $p_y = \lambda_a\, p_{y-1} + (1 - \lambda_a)$, with first-year pass-through $1 - \lambda_a$. Matching the twelve-month response gives the annual non-reset probability

$$\lambda_a = (1 - f_m)^{12}.$$

The first-order Calvo–Rotemberg slope equivalence (zero trend inflation, log-linearized) then maps λ_a to $\bar{\phi}$ via $\kappa = (\varepsilon - 1)/\bar{\phi}$:

$$\bar{\phi} = (\varepsilon - 1)\,\frac{\lambda_a}{(1 - \lambda_a)\,(1 - \beta\,\lambda_a)}, \qquad \beta = 0.96,\ \varepsilon = 5. \tag{A.59}$$

Sticky industries (high λ_a) receive a large $\bar{\phi}$ and a flat NKPC; flexible industries a small $\bar{\phi}$ and a steep one.

We aggregate to the model's chain × tier-12 cells at the annual horizon, where pass-through is linear in the per-industry annual non-reset probability. For cell c,

$$\lambda_a^c = \sum_{i \in c} s_i\,(1 - f_m^i)^{12}, \tag{A.60}$$

each industry frequency compounded to the annual horizon *first*, then averaged with 2019 QCEW wage-bill weights s_i (as proxy for value added shares). Note that averaging $1 - f_m^i$ before compounding would understate stickiness through the convexity of $(\cdot)^{12}$. Cell $\bar{\phi}^c$ follows from (A.59) evaluated at λ_a^c. The wage-bill weights s_i proxy the expenditure shares the price index strictly requires; the two coincide when within-cell labor shares are uniform. Table 10 reports per-cell values, with $1 - f_m^c \equiv (\lambda_a^c)^{1/12}$ the effective non-reset rate.

Table 10: Per-cell Rotemberg adjustment cost $\bar{\phi}$ from Pasten et al. (2020) monthly Calvo frequencies via (A.60)–(A.59). The ψ_a column is the slope object $\lambda_a^c/[(1-\lambda_a^c)(1-\beta\lambda_a^c)]$ before the $(\varepsilon-1)$ scaling.

Cell	$1-f_m^c$	monthly reset freq.	ψ_a	$\bar{\phi}=4\,\psi_a$	n PSW-exact / total
$C1_h$	0.839	16.1%	0.157	0.63	0/4
$C1_m$	0.906	9.4%	0.629	2.52	10/20
$C1_l$	0.930	7.0%	1.202	4.81	10/19
$C2_h$	0.917	8.3%	0.829	3.32	19/47
$C2_m$	0.877	12.3%	0.324	1.30	5/10
$C2_l$	0.802	19.8%	0.082	0.33	16/20
$C3_h$	0.929	7.1%	1.172	4.69	17/29
$C3_m$	0.906	9.4%	0.629	2.51	16/31
$C3_l$	0.924	7.6%	1.011	4.04	30/41
$C4_h$	0.899	10.2%	0.522	2.09	14/22
$C4_m$	0.880	12.0%	0.348	1.39	52/67
$C4_l$	0.903	9.7%	0.583	2.33	22/31

Notes: Model parameters use a five-key collapse: per-chain h-tier values ($\bar{\phi}_{C1}=0.63$, $\bar{\phi}_{C2}=3.32$, $\bar{\phi}_{C3}=4.69$, $\bar{\phi}_{C4}=2.09$) and one wage-bill-weighted average across the eight $m+l$ cells, $\bar{\phi}_{ml}=2.86$.

A.8 Inflation-expectations estimation

This appendix calibrates the three-weight expectations operator used in the model.

A.8.1 Specification and identification

The parameterized expectations operator from equation (22) of the main text is

$$E_t[\Pi_{t+12}] \;=\; \nu_{\text{forward}}\,\Pi_{t+12} \;+\; \nu_{\text{past}}\,\Pi_t^{\text{past}} \;+\; \nu_{\text{target}}\,\bar{\Pi}, \qquad \nu_{\text{forward}}+\nu_{\text{past}}+\nu_{\text{target}}=1. \tag{A.61}$$

$E_t[\Pi_{t+12}]$ is the observed one-year-ahead expectation (Cleveland Fed model-implied or Michigan survey median), Π_t^{past} is realized 12-month CPI inflation ending at t, Π_{t+12} is realized 12-month inflation over the following year (a rational-expectations proxy for the forward-looking component under perfect foresight), and $\bar{\Pi}=2\%$ is the central bank's long-run target.[34] The economic content of equation (A.61) is a convex-combination reduced form: agents' contemporaneous forecast projects onto three components—the component of realized future inflation they would internalize under perfect foresight, the past inflation they observe, and the credible anchor at which they expect the central bank to return prices. The weights are what we want to identify.

[34]We use headline 12-month CPI for Π_t^{past} because TIPS-based market expectations are CPI based.

Constrained OLS. The adding-up constraint $\nu_{\text{forward}} + \nu_{\text{past}} + \nu_{\text{target}} = 1$ turns (A.61) into a two-parameter regression. Subtracting $\bar{\Pi}$ from both sides and substituting $\nu_{\text{target}} = 1 - \nu_{\text{forward}} - \nu_{\text{past}}$ delivers

$$E_t[\Pi_{t+12}] - \bar{\Pi} \;=\; \nu_{\text{forward}} \left(\Pi_{t+12} - \bar{\Pi}\right) \;+\; \nu_{\text{past}} \left(\Pi_t^{\text{past}} - \bar{\Pi}\right) \;+\; u_t, \tag{A.62}$$

an OLS regression with no intercept on the deviations from the anchor. We recover $\hat{\nu}_{\text{target}} = 1 - \hat{\nu}_{\text{forward}} - \hat{\nu}_{\text{past}}$ residually; its standard error follows by the delta method from the joint covariance of the two estimated coefficients. Standard errors are Newey–West HAC at 12 lags to match the overlap horizon of Π_{t+12}.

Equation (A.62) is a linear *projection*, not a structural identification of what agents actually do: it gives the least-squares decomposition of observed expectations onto three observable components. Because the forward regressor Π_{t+12} is realized future inflation—an ex-post noisy realization of the information agents would ideally condition on—the projection weight $\hat{\nu}_{\text{forward}}$ is subject to classical attenuation bias whenever true forward-looking behavior is present. So a small $\hat{\nu}_{\text{forward}}$ is consistent with agents being genuinely forward-looking but conditioning on a signal noisier than Π_{t+12}; it is not proof of a non-forward-looking regime. The regression is therefore best read as a description of how well each observable component tracks the observed expectations series, which is what the model's calibration requires.

Results. Table 11 reports the constrained-OLS estimates on monthly Cleveland Fed market-based expectations (Haubrich et al., 2012) and on the Michigan household survey, across post-COVID, pre-COVID, and full-sample windows.

Table 11: Constrained-OLS estimates of the three-weight expectations operator.

Series	Sample	N	$\hat{\nu}_{\mathrm{f}}$	$\hat{\nu}_{\mathrm{p}}$	$\hat{\nu}_{\mathrm{t}}$	R^2	RMSE
Cleveland Fed	**COVID + after (2020–2024)**	59	-0.13	0.28	0.85	0.69	0.47
			(0.03)	(0.02)	(0.02)		
Cleveland Fed	Pre-COVID (2016–2019)	48	-0.04	0.43	0.61	0.12	0.34
			(0.12)	(0.13)	(0.18)		
Cleveland Fed	Full (2003–2024)	263	-0.05	0.25	0.80	0.44	0.50
			(0.04)	(0.03)	(0.04)		
Michigan	COVID + after (2020–2024)	59	0.22	0.45	0.33	0.66	0.56
			(0.04)	(0.04)	(0.04)		
Michigan	Pre-COVID (2016–2019)	48	-0.25	-0.16	1.41	-12.22	0.63
			(0.23)	(0.37)	(0.43)		
Michigan	Full (2003–2024)	263	0.10	0.44	0.46	-0.98	1.04
			(0.09)	(0.10)	(0.10)		
Stability across the surge (splits of the COVID+ window):							
Cleveland Fed	2020-01–2021-12	24	-0.13	0.30	0.83		
			(0.04)	(0.08)	(0.05)		
	2022-01–2024-12	35	-0.22	0.31	0.91		
			(0.14)	(0.04)	(0.10)		
	Δ (late − early)		-0.09	0.01	0.08		
	t-stat		[-0.61]	[0.10]	[0.70]		
Michigan	2020-01–2021-12	24	0.25	0.30	0.45		
			(0.03)	(0.06)	(0.04)		
	2022-01–2024-12	35	0.27	0.47	0.27		
			(0.20)	(0.04)	(0.16)		
	Δ (late − early)		0.01	0.17	-0.18		
	t-stat		[0.07]	[2.31]	[-1.11]		

Notes: Constrained OLS of $E_t[\Pi_{t+12}] - \bar{\Pi} = \nu_{\mathrm{f}}(\Pi_{t+12} - \bar{\Pi}) + \nu_{\mathrm{p}}(\Pi_t^{\mathrm{past}} - \bar{\Pi}) + u_t$ with no intercept and $\bar{\Pi} = 2\%$; recovering $\hat{\nu}_{\mathrm{t}} = 1 - \hat{\nu}_{\mathrm{f}} - \hat{\nu}_{\mathrm{p}}$ residually. Newey–West HAC standard errors (in parentheses) at 12 lags for the main-window rows, and at $\min(12, N-2)$ for the split-sample rows. Π_t^{past} is the trailing 12-month log CPI inflation; Π_{t+12} is the realized 12-month log CPI inflation over the following year. The stability sub-block splits the COVID+ window at 2022-01; the t-statistic on the difference in each coefficient across sub-windows uses the pooled standard error $\sqrt{\mathrm{SE}_{\mathrm{late}}^2 + \mathrm{SE}_{\mathrm{early}}^2}$. Rolling 24-month estimates for the Cleveland Fed series over 2019–2024 are emitted as `expectations_constrained_ols_rolling.csv` (not shown). Cleveland Fed: model-implied one-year-ahead expectation (Haubrich et al., 2012). Michigan: University of Michigan Survey of Consumers one-year median. CPI: FRED `CPIAUCSL`.

Two features of the projection on Cleveland Fed post-COVID data stand out. First, the projection loads little on realized future inflation: $\hat{\nu}_{\mathrm{forward}} = -0.13$ (SE 0.03). As noted above, attenuation from the noisy forward regressor makes this a lower bound on the true forward channel rather than proof of its absence—but for calibration purposes, what the projection tells us is that the observed Cleveland Fed series does not track realized future inflation closely enough for a forward-loading operator to improve fit materially. Second,

the anchor projection dominates: $\hat{\nu}_{\text{target}} = 0.85$ against a past weight of $\hat{\nu}_{\text{past}} = 0.28$. The Cleveland Fed series is heavily tethered to the 2% anchor, with only about a quarter of the projection on rolling past inflation—a purely adaptive specification would put much more weight on recent inflation. The full-sample estimates deliver the same picture: $\hat{\nu}_{\text{forward}} = -0.05$, $\hat{\nu}_{\text{past}} = 0.25$, $\hat{\nu}_{\text{target}} = 0.80$.

We take from this the calibration $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (0,\ 0.25,\ 0.75)$ for the adaptive-anchored regime. Setting $\nu_{\text{forward}} = 0$ is a modeling choice consistent with the projection: the small point estimate leaves little fit to give up, and the attenuation caveat cuts the other way as well—forcing zero loses nothing the noisy projection would have identified. Given that choice, imposing $\nu_{\text{past}} + \nu_{\text{target}} = 1$ returns weights close to the post-COVID and full-sample point estimates. Robustness to alternative anchors $\bar{\Pi} \in \{1.50,\ 1.75,\ 2.00,\ 2.25\}\%$ leaves $\hat{\nu}_{\text{forward}}$ near zero and $\hat{\nu}_{\text{past}}$ in a tight band $[0.26,\ 0.33]$; the target weight compensates for the anchor level through the adding-up constraint, but the two data-identified projections are essentially invariant to the anchor choice (sensitivity table emitted alongside as `expectations_constrained_ols_sensitivity.csv`).

Households versus markets. The Michigan survey tells a different story than the market series. In the post-COVID window, household expectations load more heavily on both forward-looking and past inflation than market expectations do, and correspondingly less on the anchor: $\hat{\nu}_{\text{forward}} = 0.22$, $\hat{\nu}_{\text{past}} = 0.45$, $\hat{\nu}_{\text{target}} = 0.33$. This is consistent with the well-documented fact that households anchor less strongly to the central-bank target than professional forecasters or market participants do (Coibion and Gorodnichenko, 2012). The model's baseline calibration uses market-based expectations because equilibrium prices are set by firms, whose expectations—like those of the professional forecasters and market participants that populate the Cleveland Fed's data inputs—are closer to the market series than to household surveys.

Stability of the projection through the surge. The lower sub-block of Table 11 splits the COVID+ window at 2022-01 and re-estimates the projection separately on the two halves. The Cleveland Fed projection is stable across the split: the past weight moves by 0.01 ($t = 0.10$) and the target weight by 0.08 ($t = 0.70$), both consistent with sampling variation. Rolling 24-month estimates tell the same story. The Michigan series does move: the past weight shifts by 0.17 ($t = 2.31$)—statistically significant at conventional levels—and the anchor weight correspondingly falls by −0.18. Households became meaningfully more backward-looking in 2022 than they had been in 2020–2021, loading more on realized past inflation and less on the 2% target—the empirical footprint of the household de-

anchoring worry that concerned policymakers in mid-2022, and one that does not appear in the market series.

A.9 Model robustness and sensitivity

This appendix reports robustness and sensitivity results for the model of Section 5, organized in four progressing subsections: (i) parameter sensitivity of the headline wall contribution; (ii) baseline under perfect-foresight expectations; (iii) warm-pool sensitivity ($\psi = 0.30$), also reported under both regimes; and (iv) Phillips-curve impulse responses.

A.9.1 Parameter sensitivity of the wall contribution

Experiment	Wall Act I (pp)	Wall Act II (pp)
Baseline	+4.26	−0.02
Higher warm-pool floor ($\psi = 0.30$)	+3.32	+0.14
Low cost anchor ($c = 0.10$, $\gamma \times 0.78$)	+3.76	−0.26
All elasticities at 0.5 ($\rho_b = \rho_m = \sigma_w = 0.5$)	+3.80	+0.12
High hit/safe substitution ($\sigma_w = 5$)	+2.16	+0.53
Lower downward rigidity ($\lambda \times 0.5$ per chain)	+3.90	−0.23
High labor–materials substitution ($\rho_m = 0.5$)	+4.17	+0.00
Wage backward indexation off ($\iota_w = 0$)	+4.35	−0.52

Notes: The two reported numbers are the capacity-wall channel's contribution to cumulative PCE inflation Π, in pp, computed as a counterfactual difference: cumulative inflation under ($\mu^e = 0$, EXO off, γ active) minus cumulative inflation under ($\mu^e = 0$, EXO off, $\gamma = 0$). Equivalently, the "Wall" column of the stacked structural decomposition in Figure 6, baseline column. For each row, the model is fully re-solved end-to-end with one parameter perturbed (column header) and all others held at the baseline. **Wall Act I** is the wall-channel contribution to cumulative Π over 2019–2021 (the 2021 entry of the "Wall" column). **Wall Act II** is the Act-II increment: the wall-channel contribution from 2021 to 2023 (the 2023 entry minus the 2021 entry). Both are pp-level contributions, NOT shares; positive values mean the wall mechanism added cumulative inflation in that window. Baseline calibration: $c_{\text{wall}} = 0.13$, wall premium $\mathcal{W}_{ij,t} = \gamma_i (dN^{+}_{ij,t})^2$, warm-pool floor $\psi = 0.20$, $\rho_m = \rho_b = \sigma_w = 0.39$, per-chain $\lambda = 3\bar{\phi}$, $\iota_w = 0.20$.

Table 12: Wall-channel robustness (adaptive expectations spec).

Table 12 reports the cumulative CW contribution by Act under the adaptive-anchored expectations regime across a range of reparameterizations: warm-pool sensitivity (ψ raised from the baseline 0.20 to 0.30), three CES-elasticity perturbations (joint $\rho = 0.5$; high hit/safe substitution $\sigma_w = 5$; high labor–materials substitution $\rho_m = 0.5$), the structural-asymmetry parameter λ halved, the wall-premium anchor c reduced by about 22%, and wage backward indexation turned off ($\iota_w = 0$). Each exercise is also reported under the perfect-foresight expectations specification where applicable.

Findings (parameter sweep). Reducing the per-net-hire rebuilding cost c scales down both Act I and Act II wall contributions roughly proportionally. Halving the downward-rigidity parameter λ leaves the results essentially unchanged. Raising the hit/safe substitution σ_w to 5 nearly halves the 2019–2021 reopening-window wall contribution, confirming that network rigidity is essential for the results. Raising all three CES elasticities jointly to 0.5 moves the contributions in both windows but the results are qualitatively similar: the wall contribution falls from 4.1 to 3.7 in Act I, with Act II largely unchanged. Raising labor–materials substitution alone to $\rho_m = 0.5$ has little impact. Turning off wage backward indexation ($\iota_w = 0$) shrinks the Act II contribution because the wage block carries less of the earlier impulse forward. Sticky-price parameters have virtually no bearing on the results—a flat Phillips curve would deliver similar numbers. Network rigidity and low network elasticity, by contrast, are essential: the wall contribution weakens substantially once these elasticities are raised much above 0.5.

A.9.2 Baseline under perfect-foresight expectations

The main model uses the adaptive-anchored expectations regime $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (0, 0.25, 0.75)$. This subsection reports the results from the perfect-foresight (PF) case $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (1, 0, 0)$, which leads to counterfactual expected inflation dynamics but is model-consistent. All other parameters are held at their headline values; the structural markup wedges $\bar{\mu}_i^e$ and the residual markup wedges $\mu_{i,t}^e$ are re-calibrated using the same targets. Our implementation assumes that perfect foresight applies from 2021 onwards, and that 2020 expectations are adaptive to avoid aggressive front-loading of inflation in the pandemic year. (This would otherwise be absorbed by $\bar{\mu}^e$, and the model would yield artificially low inflation.)

Inflation-by-channel decomposition (baseline calibration). Figure 10 is the perfect-foresight counterpart of the adaptive-anchored main-text Figure 6, at the baseline warm-pool floor $\psi = 0.20$. Two differences stand out beyond 2021. First, perfect foresight attenuates the wall-unevenness contribution beyond 2022, because forward-looking pricing partly smooths the sector-specific cost shocks over time; the effect is much smaller in 2021. Second, the structural floor is materially lower in the perfect-foresight model: the data imply that productivity has increased by about 2% above the assumed trend over 2020–2023, which is only slowly corrected by ι_a; perfect foresight anticipates this source of deflation and front-loads it, dragging the overall contribution down relative to the data (though not as much relative to the model's own floor). We thus consider these results

qualitatively similar to the baseline model, with the caveat that there is somewhat less persistence of inflation and a somewhat smaller unevenness contribution beyond 2022 under perfect foresight.

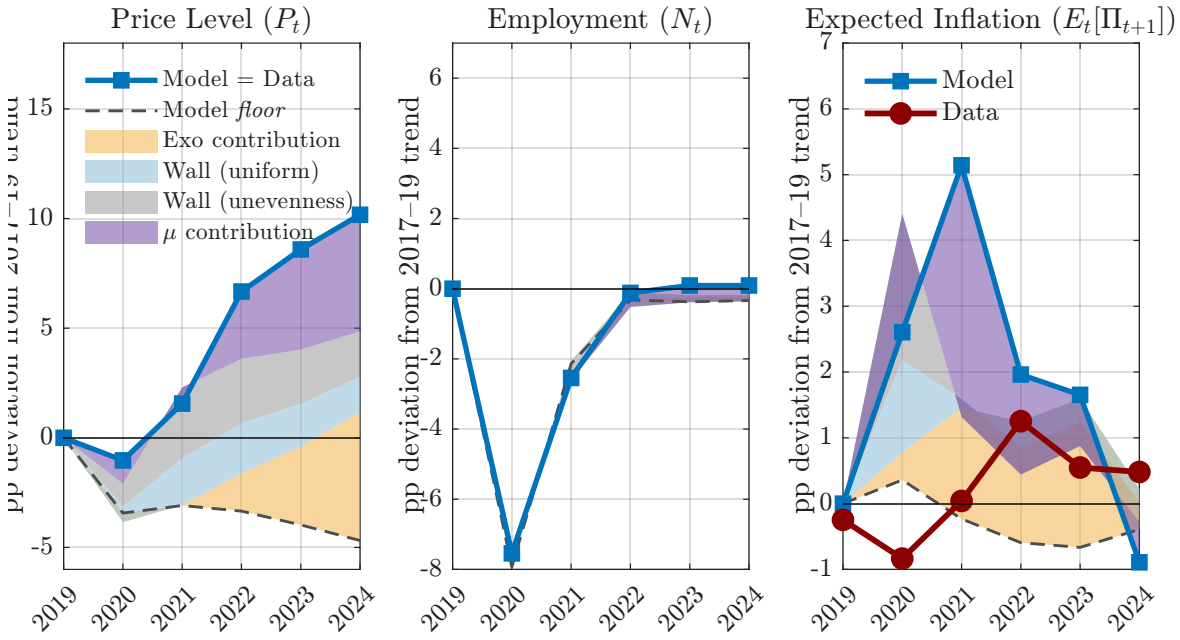

Figure 10: Inflation-by-channel decomposition under perfect foresight (baseline $\psi = 0.20$).

Notes: Perfect-foresight twin of Figure 6, re-run with $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (1, 0, 0)$; same construction, same warm-pool floor $\psi = 0.20$, same wall calibration. The stacked decomposition reports cumulative log deviations from the 2019 price level under sequential removal of the markup wedge μ^e, the EXO shocks, and the CW, and splits the CW channel into a uniform-wall component and a wall-unevenness component. Comparing with the warm-pool cut in Figure 12: raising ψ from 0.20 to 0.30 leaves the *total* wall contribution essentially unchanged, but reduces the wall-unevenness contribution substantially beyond 2022 (with a much smaller effect in 2021).

Extended horizon: real-side aggregates. Figure 11 extends the horizon to 2029 and reports the underlying real-side response of the model under both expectations regimes. Each panel plots cumulative log deviations from the 2019 baseline, in percentage points, PCE-weighted across the 12 tier composites. The top-left panel shows the PCE price level P, the top-right panel aggregate output Y, and the bottom-left panel aggregate employment N. The bottom-right panel decomposes wages into three series: the base wage W paid in the labor market (blue circles), the effective wage W^{eff} that firms face on the cost side including the wall premium (orange squares), and the measured wage W^{meas} recorded in the data, which captures the fraction of the wall premium that leaks to workers via the per-tier parameter x_u (green diamonds). The yellow band over 2020–2024 marks the calibration window in which μ^e, the EXO shocks, the wall, and the markup drift are

anchored to the data; the light-gray band over 2025–2029 marks the "hold" window in which all wedges and exogenous shocks are held at their 2024 terminal values and the model propagates forward under no further forcing. Two features of the long-horizon response stand out. Aggregate employment N does not return to its 2019 level within the plotted window, so the wall mechanism continues to operate. The three-way wage gap in the bottom-right panel closes only gradually in the hold band, as the wall premium decays and the leakage parameter x_u closes the gap between the base and measured wages. The two regimes yield similar long-run paths despite very different expectations behavior, consistent with the comparison of the price-level decompositions in Figure 12.

A.9.3 Warm-pool sensitivity ($\psi = 0.30$)

The warm-pool floor ψ governs the fraction of the 2019→2020 employment dip that firms can recall costlessly before the capacity wall fires. Raising ψ from the baseline 0.20 to 0.30 assumes a larger fraction of pre-pandemic worker–firm matches is preserved through the contraction. Figure 12 reports the full structural decomposition under $\psi = 0.30$ for both expectations regimes. The comparison with the baseline (main-text Figure 6 for the adaptive regime; Figure 10 for perfect foresight) yields two conclusions. First, the *total* wall contribution is essentially unchanged: the capacity wall remains the leading structural contributor to above-trend inflation, and its Act I magnitude is barely dented. Second, the *unevenness* component of the wall—the additional inflation attributable to the cross-sectional concentration of workforce destruction—is substantially reduced beyond 2022, with a much smaller effect in 2021. Both features hold under adaptive and perfect-foresight expectations. Intuitively, a larger warm pool preserves more matches in the deeply-hit sectors and narrows the cross-sectional dispersion of workforce destruction; the aggregate rebuilding burden is little changed, but its uneven distribution across nodes—which is what amplifies inflation through the network—is compressed as the recovery matures.

A.9.4 Phillips-curve impulse responses

Expansionary +1% demand shock. Figures 13 and 14 show perfect-foresight paths that feed the 1% expansionary point on the joint Phillips schedule in main-text Figure 7. Each figure shows a five-panel summary of the model's response to a uniform, one-period, +1% iid ξ demand shock applied to all 12 tier composites simultaneously, plotted over a 16-period horizon and starting from the model's steady state. Every panel contains four cases: wall on vs. wall off (calibrated γ_i vs. $\gamma_i = 0$) crossed with the MPC income-propagation

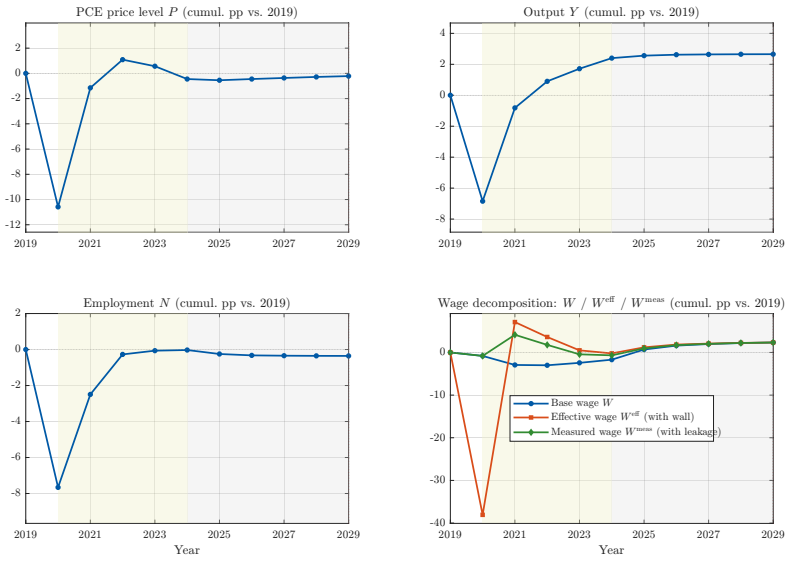

(a) Adaptive expectations (headline), $\nu = (0,\ 0.25,\ 0.75)$.

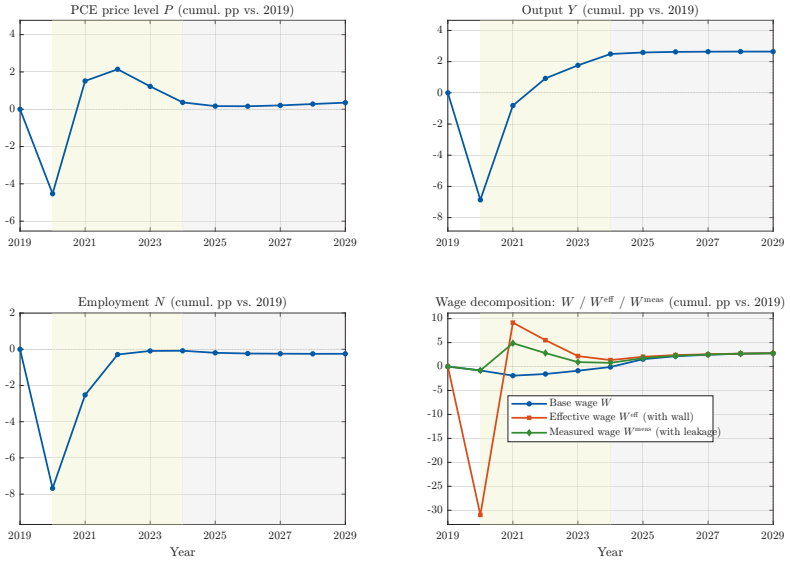

(b) Perfect foresight (robustness), $\nu = (1,\ 0,\ 0)$.

Figure 11: Extended-horizon real-side response of the model under both expectations regimes.

Notes: Cumulative log deviations from the 2019 baseline, in percentage points, PCE-weighted across the 12 tier composites. Panels: PCE price level P (top-left), output Y (top-right), employment N (bottom-left), and the three-way wage decomposition W / W^{eff} / W^{meas} (bottom-right). Yellow band (2020–2024) marks the calibration window; light-gray band (2025–2029) marks the "hold" window in which all wedges are held at their 2024 terminal values. Same shocks, same calibration window, and same wall mechanism across the two panels; only the expectations regime ν differs.

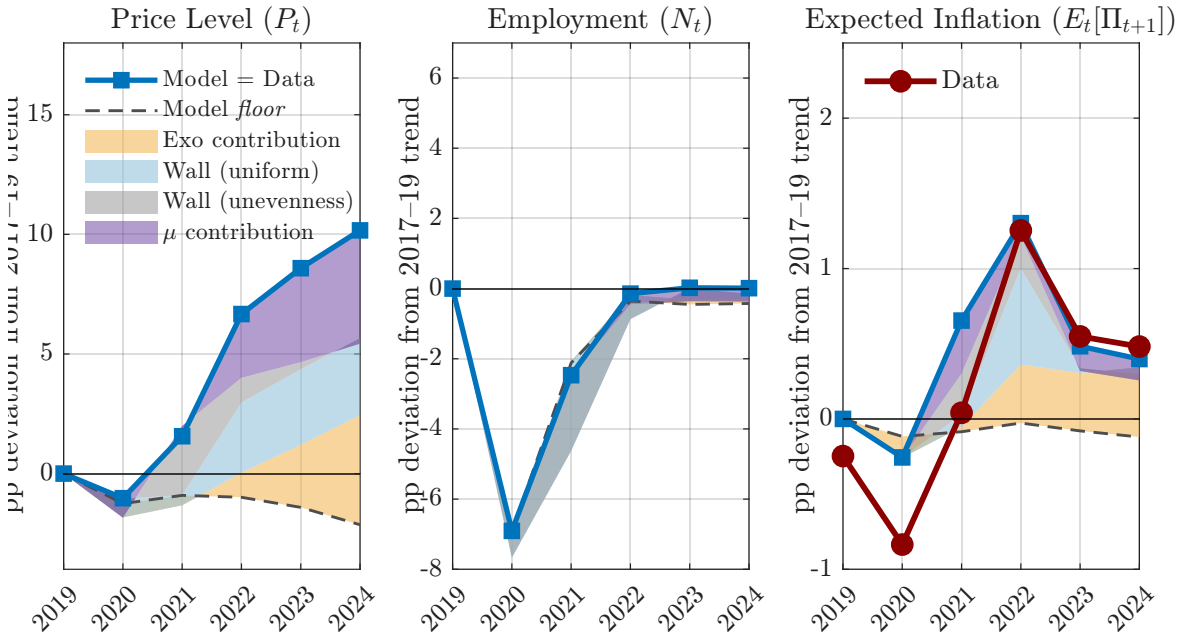

(a) Adaptive expectations (headline), $\nu = (0, 0.25, 0.75)$.

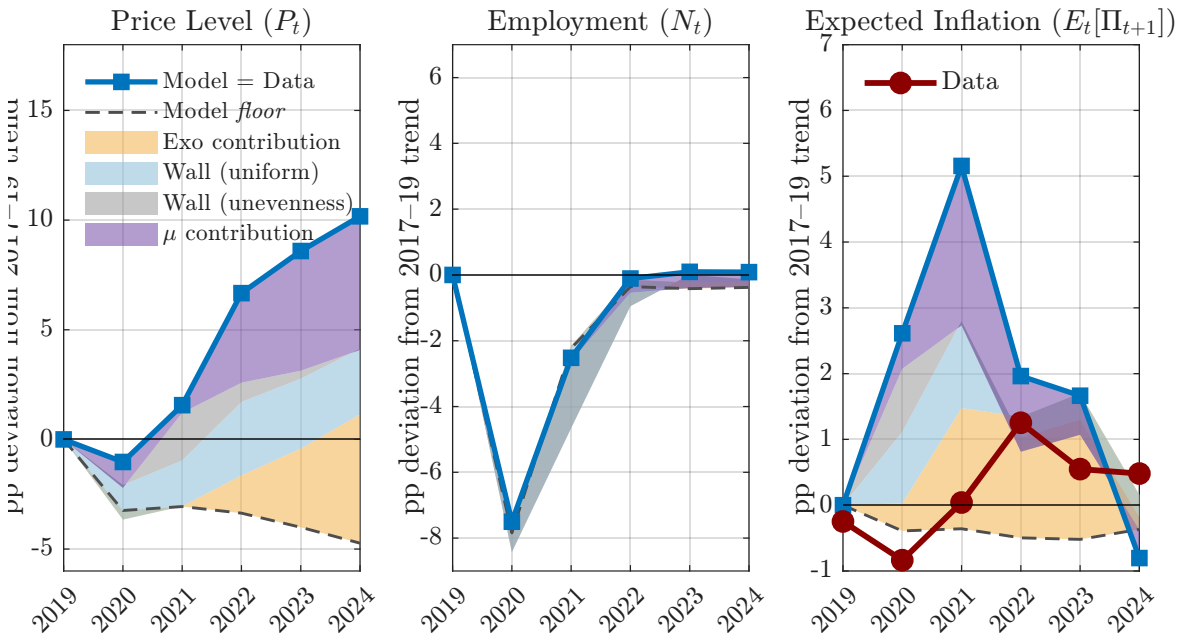

(b) Perfect foresight, $\nu = (1, 0, 0)$.

Figure 12: Inflation-by-channel decomposition under the $\psi = 0.30$ warm-pool floor.

Notes: Warm-pool sensitivity twin of Figure 6, re-run with the warm-pool floor raised to $\psi = 0.30$ (baseline $\psi = 0.20$). Same construction: cumulative log deviations from the 2019 price level, decomposed under sequential removal of the markup wedge μ^e, the EXO shocks, and the CW ($\gamma = 0$). Panel (a): adaptive-anchored expectations (headline regime). Panel (b): perfect-foresight expectations.

channel off (MPC= 0) and on (MPC= 0.40). Figure 14 reports the perfect-foresight case (1, 0, 0).

The main-text Phillips schedule (Figure 7) corresponds to impact-period ($t = 1$) values from panels (a) and (c) of these IRFs, and it plots $\Delta\Pi(1)$ against the unemployment-gap proxy $-\Delta N(1)$. Repeating this for eight shock sizes ($|\xi| \in \{0.1\%, 0.3\%, 0.5\%, 0.7\%\}$ applied with both signs) and each of the four arms (wall on/off $\times$ MPC $\in \{0, 0.4\}$) gives the curve in that figure.

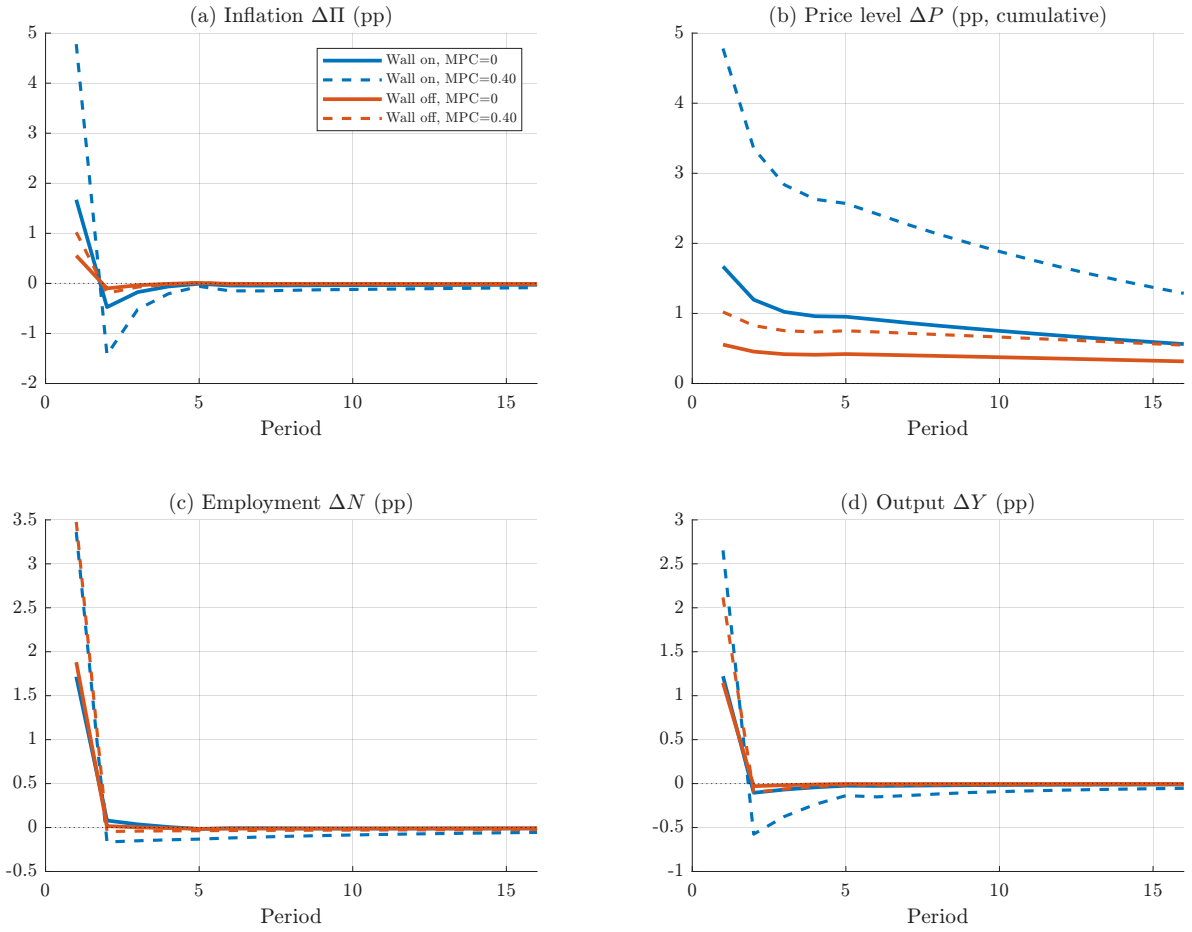

Figure 13: Impulse responses to a +1% iid ξ demand shock, adaptive spec.

Notes: Five panels report inflation, the cumulative price level, employment, output, and the mean CW premium 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 γ); orange: wall off ($\gamma = 0$). Solid: MPC= 0; dashed: MPC= 0.40.

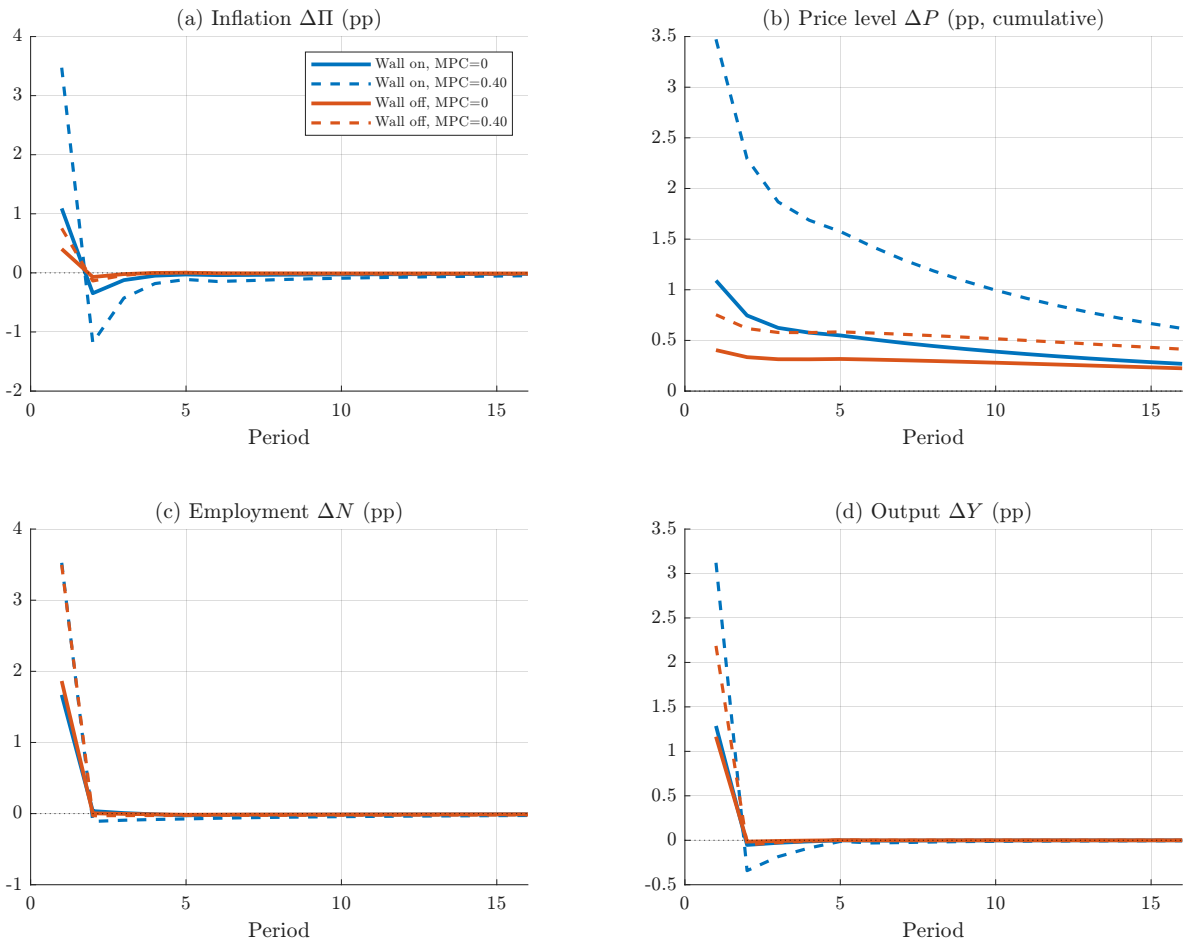

Figure 14: Impulse responses to a +1% iid ξ demand shock, perfect-foresight spec.

Notes: Construction identical to Figure 13 with the perfect-foresight expectations triple $(\nu_{\text{forward}}, \nu_{\text{past}}, \nu_{\text{target}}) = (1, 0, 0)$.

Contractionary −1% demand shock. Figures 15 and 16 mirror the previous exercise for a contractionary shock.

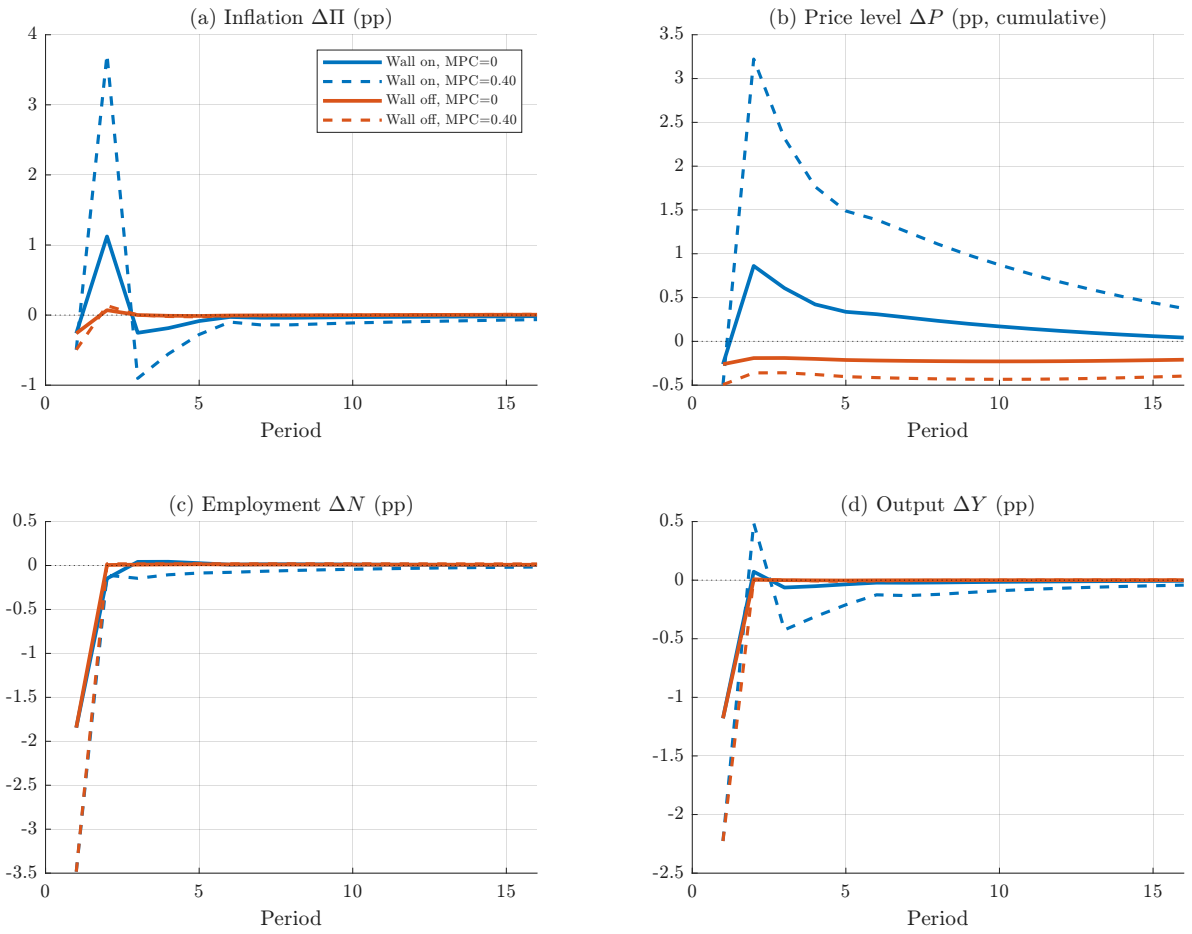

Figure 15: Impulse responses to a -1% iid ξ demand shock, adaptive spec.

Notes: Construction identical to Figure 13, with the sign of the impulse flipped. Wall premium (panel e) is zero at impact for all arms because the wall is one-sided in $dN^+ = \max\{dN, 0\}$; any positive premium in later periods reflects the recovery overshoot firing the wall in $t \geq 2$.

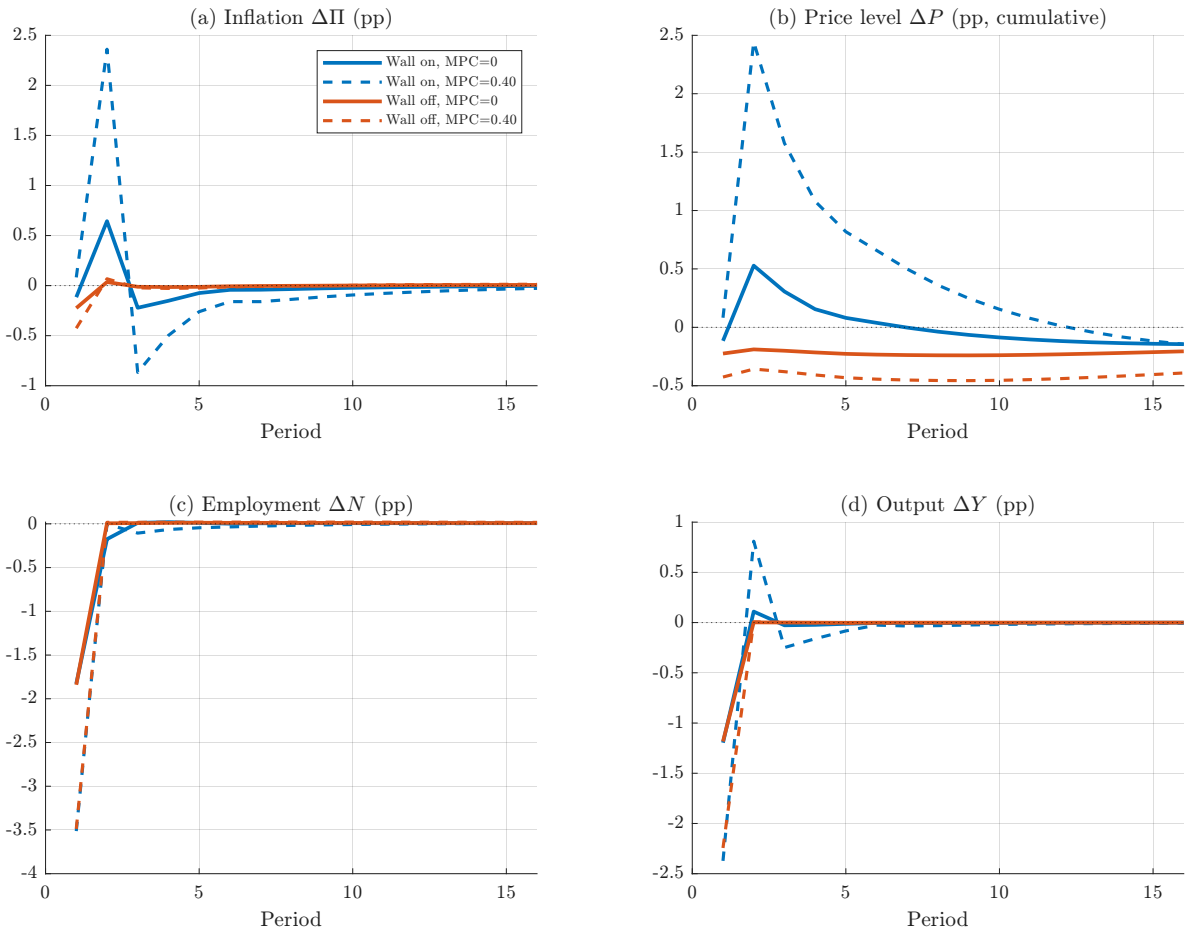

Figure 16: Impulse responses to a −1% iid ξ demand pulse, perfect-foresight spec.

Notes: Construction identical to Figure 15 with the perfect-foresight expectations triple (1, 0, 0).

A.10 Regression robustness

Shapley R^2 decomposition for Fact 1. Table 13 reports the total R^2 of the Fact 1 specification partitions across regressor blocks.

Table 13: Fact 1, Shapley R^2 decomposition.

Outcome	Industry FE	Window FE	**Exposure DiD**	Total R^2
$\Delta \log P_i$	0.372	0.071	**0.097**	0.541
ΔMARGIN_i	0.345	0.024	**0.033**	0.402
ΔULC_i	0.399	0.031	**0.043**	0.473
ΔIMP_i	0.481	0.018	**0.029**	0.529
ΔMAT_i	0.308	0.182	**0.265**	0.755

Notes: Sample of $N = 330$ endogenous BEA-357 industries pooled across three windows (pre 2017–2019, Act I 2019–2021, Act II 2021–2023); panel observations = 990. Outcomes are per-year-averaged in pp/year. The five outcomes form the additive direct identity $\Delta \log P_i = \Delta\text{MARGIN}_i + \Delta\text{ULC}_i + \Delta\text{IMP}_i + \Delta\text{MAT}_i$. Shapley values average over all 3! = 6 orderings in which the blocks could be added to the model and sum to the total R^2 by construction. Stars and significance not applicable to R^2 shares.

A.10.1 Robustness to δ

The regressions of §2.2 fix the per-step network discount at $\delta = 0.5$. The operator $A\,L_\delta$ that constructs $\mathcal{E}$ and $\mathcal{V}$ is monotone in δ: $\delta = 0$ retains only the direct supplier link ($\mathcal{E} = A\,E$), and $\delta = 1$ propagates fully through the Leontief inverse ($\mathcal{E} = A(I - A)^{-1}E$). The baseline value $\delta = 0.5$ interpolates between these endpoints. Tables 14–19 report the Facts 1 and 2 estimates at the two endpoints, side-by-side with the $\delta = 0.5$ main-text headline.

Signs and statistical significance are preserved across $\delta \in \{0, 0.5, 1\}$ for every headline cell. Magnitudes shift by 15–20% monotonically. The Fact 1 inflation slope on $\mathcal{E}$ moves from +12.5 ($\delta = 0$) to +10.5 ($\delta = 1$) in the full sample, and from +20.8 to +15.5 in the $m + l$-tier subsample; both endpoints remain significant at the 5% level or better. Fact 2's wall-discriminator $\hat{\beta}^{\mathcal{V}^{\text{wall}}}_{\text{Act I}}$ on $\Delta \log P$ moves from −1.29 to −1.00, remaining negative and significant at the 10% level throughout. The post-wall placebo $\hat{\beta}^{\mathcal{V}^{\text{post}}}_{\text{Act II}}$ remains close to zero and insignificant across all δ, as it should. Fact 1's dispersion test is significant at the 1% level at every δ, with the regression R^2 rising mildly from 0.450 at $\delta = 0$ to 0.495 at $\delta = 1$.

A.10.2 Supplier-network vs. own exposure

Table 22 reports the double-exposure robustness referenced in Fact 1: adding own-exposure interactions leaves the supplier-network coefficients substantively unchanged.

Table 14: Fact 1: industry inflation on supplier-network exposure $\mathcal{E}_i$.

	pre-COVID (omitted)		Act I $\hat{\beta}_{\text{Act I}}$ (t)		Act II $\hat{\beta}_{\text{Act II}}$ (t)	
Outcome	ALL ($N=330$)	m+l-tier ($N=232$)	ALL ($N=330$)	m+l-tier ($N=232$)	ALL ($N=330$)	m+l-tier ($N=232$)
$\Delta \log P_i$	0.00	0.00	$+12.53^{**}$	$+20.77^{***}$	+7.22	+2.65
	(0.00)	(0.00)	(+2.25)	(+3.54)	(+0.94)	(+0.23)
ΔMARGIN_i	0.00	0.00	+11.44	$+26.49^{***}$	−16.68	−8.60
	(0.00)	(0.00)	(+1.15)	(+3.30)	(−1.24)	(−0.62)
ΔULC_i	0.00	0.00	−3.61	−11.76	$+21.94^{*}$	+9.41
	(0.00)	(0.00)	(−0.47)	(−1.38)	(+1.95)	(+1.15)
ΔIMP_i	0.00	0.00	+0.24	+0.33	$+0.23^{***}$	+0.11
	(0.00)	(0.00)	(+1.48)	(+1.63)	(+3.34)	(+1.36)
ΔMAT_i	0.00	0.00	$+4.46^{***}$	$+5.72^{***}$	$+1.73^{***}$	$+1.73^{***}$
	(0.00)	(0.00)	(+5.18)	(+5.14)	(+3.37)	(+3.17)
$A\cdot\Delta\text{MARGIN}_i$	0.00	0.00	$+3.80^{***}$	$+5.07^{***}$	−1.16	−0.58
	(0.00)	(0.00)	(+4.47)	(+5.10)	(−1.16)	(−0.50)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, and Act II 2021–2023. Weighted least squares with 2017 industry value added as weights. Industry and window fixed effects throughout (so $\alpha_{\text{pre-COVID}} = \beta_{\text{pre-COVID}} = 0$ by construction). Standard errors clustered on industry. The identification targets are $\beta_{\text{Act I}}$ and $\beta_{\text{Act II}}$. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.10.3 Cost–price pass-through in the pandemic windows

The main-text ULC discussion in Section 5.2 refers to an elevated cost–price pass-through observed in the data during the outbreak. Table 23 documents this fact directly. For each window $w \in \{\text{pre-COVID, Act I, Act II}\}$ we estimate the cross-industry OLS regression

$$\Delta \log P_{i,w} = \alpha_{s(i)} + \beta_w \, \Delta \log \text{ULC}_{i,w} + \varepsilon_{i,w},$$

with BEA-66 sector fixed effects $\alpha_{s(i)}$, weighted least squares using 2017 industry value added, and standard errors clustered on BEA-66 sector. Cumulative log changes are computed within window ($\Delta \log X_{i,w} = \sum_{t \in w} \Delta \log X_{i,t}$). The full sample covers 330 BEA detailed industries (dropping 16 EXO and 11 with missing JEM exposure); the $m + l$-tier subsample restricts to non-consumer-facing industries within the identified supply chains. The estimated pass-through slope $\hat{\beta}_w$ nearly doubles from pre-COVID to Act I, and remains elevated in Act II among upstream and middle-tier industries. It is against these estimates that the model's Act I–Act II pass-through comparison is benchmarked.

Table 15: Fact 1: industry inflation on supplier-network exposure $\mathcal{E}_i$.

	pre-COVID (omitted)		Act I $\hat{\beta}_{\text{Act I}}$ (t)		Act II $\hat{\beta}_{\text{Act II}}$ (t)	
Outcome	ALL ($N = 330$)	m+l-tier ($N = 232$)	ALL ($N = 330$)	m+l-tier ($N = 232$)	ALL ($N = 330$)	m+l-tier ($N = 232$)
$\Delta \log P_i$	0.00	0.00	$+10.48^{***}$	$+15.47^{***}$	$+4.72$	$+1.53$
	(0.00)	(0.00)	$(+2.74)$	$(+3.58)$	$(+0.93)$	$(+0.20)$
ΔMARGIN_i	0.00	0.00	$+10.16$	$+19.91^{***}$	-12.87	-6.02
	(0.00)	(0.00)	$(+1.54)$	$(+3.24)$	(-1.27)	(-0.64)
ΔULC_i	0.00	0.00	-2.96	-8.64	$+15.96^{*}$	$+6.21$
	(0.00)	(0.00)	(-0.55)	(-1.32)	$(+1.78)$	$(+1.16)$
ΔIMP_i	0.00	0.00	$+0.15$	$+0.20$	$+0.15^{***}$	$+0.06$
	(0.00)	(0.00)	$(+1.36)$	$(+1.44)$	$(+3.07)$	$(+1.02)$
ΔMAT_i	0.00	0.00	$+3.14^{***}$	$+4.00^{***}$	$+1.48^{***}$	$+1.28^{***}$
	(0.00)	(0.00)	$(+4.86)$	$(+4.51)$	$(+5.21)$	$(+3.69)$
$A \cdot \Delta \text{MARGIN}_i$	0.00	0.00	$+2.66^{***}$	$+3.61^{***}$	-0.52	-0.16
	(0.00)	(0.00)	$(+4.38)$	$(+4.97)$	(-0.73)	(-0.20)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, and Act II 2021–2023. Weighted least squares with 2017 industry value added as weights. Industry and window fixed effects throughout (so $\alpha_{\text{pre-COVID}} = \beta_{\text{pre-COVID}} = 0$ by construction). Standard errors clustered on industry. The identification targets are $\beta_{\text{Act I}}$ and $\beta_{\text{Act II}}$. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.10.4 Weak-instrument tests for Section 5

This subsection provides weak IV tests for Section 5.

On the data 2SLS moment of Table 7, the cluster-robust first-stage Kleibergen–Paap rk-Wald F on the excluded instrument E_i is 5.04 — below the conventional $F > 10$ threshold and below the Stock and Yogo (2005) critical values for 10% (16.38) and 30% (9.08) maximum 2SLS bias. Conventional Wald-based inference on $\beta_{\Delta \log N}$ is therefore unreliable, and we report weak-instrument-robust Anderson–Rubin inference. The AR test rejects $\mathcal{H}_0 : \beta_{\Delta \log N} = 0$ at the 5% level (AR = 4.18, $p = 0.046$), and the inverted AR 95% confidence set is $(-\infty, -0.02]$ — with a strictly negative upper bound that excludes zero. The identifying claim of Table 7 survives weak-IV-robust inference despite the first-stage F being below conventional cutoffs; the AR confidence set, not the naive Wald CI, is the correct inferential statement. With one instrument and one endogenous regressor, LIML coincides with 2SLS, so no additional information is recovered from a k-class alternative.

Table 16: Fact 2: industry inflation jointly on supplier-network exposure and supplier-recovery velocity.

	$\hat{\beta}^{\mathcal{E}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act II}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act II}}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act II}}}_{\text{Act II}}$
$\Delta \log P_i$	$+14.28^{***}$	-1.29^{*}	-1.29^{**}	-1.07	-0.47
	$(+2.84)$	(-1.95)	(-2.51)	(-0.70)	(-0.33)
ΔMARGIN_i	$+16.77^{*}$	-2.71^{**}	-1.94	-3.81	$+1.73$
	$(+1.95)$	(-2.08)	(-0.84)	(-1.13)	$(+0.66)$
ΔULC_i	-7.16	$+1.92$	$+0.80$	$+2.48$	-2.21
	(-0.90)	$(+1.61)$	$(+0.35)$	$(+0.97)$	(-0.84)
ΔIMP_i	$+0.28^{*}$	-0.08^{***}	-0.03^{***}	-0.01	$+0.01$
	$(+1.71)$	(-3.69)	(-3.00)	(-0.17)	$(+0.90)$
ΔMAT_i	$+4.38^{***}$	-0.42^{***}	-0.13^{**}	$+0.26$	$+0.00$
	$(+4.98)$	(-3.48)	(-2.25)	$(+1.36)$	$(+0.00)$
$A \cdot \Delta\text{MARGIN}_i$	$+3.52^{***}$	-0.39^{***}	-0.19	$+0.43^{**}$	-0.28
	$(+4.90)$	(-3.52)	(-1.06)	$(+2.46)$	(-1.12)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, Act II 2021–2023. WLS with 2017 industry value added as weights; industry and window fixed effects; standard errors clustered on industry. $\boldsymbol{\mathcal{V}}_{\text{Act I}}$ is built over the wall-measurement window 2020–2022; $\boldsymbol{\mathcal{V}}_{\text{Act II}}$ over the post-wall placebo window 2022–2024 (definition in §2.2 eq. (6)). Sign $\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act I}} < 0$ is the discriminating test for the wall mechanism. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $N = 330$ industries.

Table 17: Fact 2: industry inflation jointly on supplier-network exposure and supplier-recovery velocity.

	$\hat{\beta}^{\mathcal{E}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act II}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act II}}}_{\text{Act I}}$	$\hat{\beta}^{\mathcal{V}_{\text{Act II}}}_{\text{Act II}}$
$\Delta \log P_i$	$+12.29^{***}$	-1.00^{*}	-1.14^{**}	-0.79	-0.54
	$(+3.31)$	(-1.91)	(-2.45)	(-0.58)	(-0.42)
ΔMARGIN_i	$+15.14^{**}$	-1.95^{*}	-1.51	-3.17	$+1.59$
	$(+2.34)$	(-1.89)	(-0.83)	(-1.05)	$(+0.67)$
ΔULC_i	-6.40	$+1.38$	$+0.49$	$+2.14$	-2.17
	(-1.02)	$(+1.47)$	$(+0.27)$	$(+0.95)$	(-0.92)
ΔIMP_i	$+0.23^{**}$	-0.07^{***}	-0.02^{***}	-0.01	$+0.01$
	$(+2.09)$	(-3.72)	(-3.22)	(-0.37)	$(+0.62)$
ΔMAT_i	$+3.32^{***}$	-0.36^{***}	-0.09^{**}	$+0.24$	$+0.02$
	$(+5.10)$	(-3.65)	(-2.13)	$(+1.40)$	$(+0.20)$
$A \cdot \Delta\text{MARGIN}_i$	$+2.71^{***}$	-0.33^{***}	-0.13	$+0.36^{**}$	-0.23
	$(+5.27)$	(-3.63)	(-0.94)	$(+2.21)$	(-0.97)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, Act II 2021–2023. WLS with 2017 industry value added as weights; industry and window fixed effects; standard errors clustered on industry. $\boldsymbol{\mathcal{V}}_{\text{Act I}}$ is built over the wall-measurement window 2020–2022; $\boldsymbol{\mathcal{V}}_{\text{Act II}}$ over the post-wall placebo window 2022–2024 (definition in §2.2 eq. (6)). Sign $\hat{\beta}^{\mathcal{V}_{\text{Act I}}}_{\text{Act I}} < 0$ is the discriminating test for the wall mechanism. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $N = 330$ industries.

Table 18: Fact 1: within-sector inflation dispersion on within-sector exposure dispersion.

Subsample	$\hat{\beta}(t)$	SE	p	R^2	Obs. (sectors)
ALL	$+92.18^{***}$ (+3.68)	25.07	0.000	0.450	38
m+l-tier	$+124.47^{***}$ (+3.10)	40.16	0.002	0.357	30

Notes: Cross-section of BEA-66 sectors with at least two endogenous industries in each subsample. Dependent variable $\Delta\sigma_s^+$ and regressor $\sigma(\mathcal{E}_i)_s$ are defined in §2.2 (Fact 1). Subsamples: ALL pools all kept industries; m+l-tier restricts to industries within the identified supply chains of §2 that do *not* sell 70 percent or more of output directly to PCE (i.e., excludes the consumer-facing top-tier within each chain). WLS with sector-aggregated 2017 value-added weights; HC1 robust standard errors. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 19: Fact 1: within-sector inflation dispersion on within-sector exposure dispersion.

Subsample	$\hat{\beta}(t)$	SE	p	R^2	Obs. (sectors)
ALL	$+70.76^{***}$ (+5.41)	13.08	0.000	0.495	38
m+l-tier	$+79.56^{***}$ (+2.94)	27.02	0.003	0.349	30

Notes: Cross-section of BEA-66 sectors with at least two endogenous industries in each subsample. Dependent variable $\Delta\sigma_s^+$ and regressor $\sigma(\mathcal{E}_i)_s$ are defined in §2.2 (Fact 1). Subsamples: ALL pools all kept industries; m+l-tier restricts to industries within the identified supply chains of §2 that do *not* sell 70 percent or more of output directly to PCE (i.e., excludes the consumer-facing top-tier within each chain). WLS with sector-aggregated 2017 value-added weights; HC1 robust standard errors. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 20: Fact 1: industry inflation on supplier-network exposure $\mathcal{E}_i$.

	pre-COVID (omitted)		Act I $\hat{\beta}_{\text{Act I}}$ (t)		Act II $\hat{\beta}_{\text{Act II}}$ (t)	
Outcome	ALL ($N = 330$)	m+l-tier ($N = 232$)	ALL ($N = 330$)	m+l-tier ($N = 232$)	ALL ($N = 330$)	m+l-tier ($N = 232$)
$\Delta \log P_i$	0.00	0.00	$+12.53^{**}$	$+20.77^{***}$	$+7.22$	$+2.65$
	(0.00)	(0.00)	$(+2.25)$	$(+3.54)$	$(+0.94)$	$(+0.23)$
ΔMARGIN_i	0.00	0.00	$+11.44$	$+26.49^{***}$	-16.68	-8.60
	(0.00)	(0.00)	$(+1.15)$	$(+3.30)$	(-1.24)	(-0.62)
ΔULC_i	0.00	0.00	-3.61	-11.76	$+21.94^{*}$	$+9.41$
	(0.00)	(0.00)	(-0.47)	(-1.38)	$(+1.95)$	$(+1.15)$
ΔIMP_i	0.00	0.00	$+0.24$	$+0.33$	$+0.23^{***}$	$+0.11$
	(0.00)	(0.00)	$(+1.48)$	$(+1.63)$	$(+3.34)$	$(+1.36)$
ΔMAT_i	0.00	0.00	$+4.46^{***}$	$+5.72^{***}$	$+1.73^{***}$	$+1.73^{***}$
	(0.00)	(0.00)	$(+5.18)$	$(+5.14)$	$(+3.37)$	$(+3.17)$
$A \cdot \Delta\text{MARGIN}_i$	0.00	0.00	$+3.80^{***}$	$+5.07^{***}$	-1.16	-0.58
	(0.00)	(0.00)	$(+4.47)$	$(+5.10)$	(-1.16)	(-0.50)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, and Act II 2021–2023. Weighted least squares with 2017 industry value added as weights. Industry and window fixed effects throughout (so $\alpha_{\text{pre-COVID}} = \beta_{\text{pre-COVID}} = 0$ by construction). Standard errors clustered on industry. The identification targets are $\beta_{\text{Act I}}$ and $\beta_{\text{Act II}}$. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 21: Fact 1: industry inflation on supplier-network exposure $\mathcal{E}_i$.

	pre-COVID (omitted)		Act I $\hat{\beta}_{\text{Act I}}$ (t)		Act II $\hat{\beta}_{\text{Act II}}$ (t)	
Outcome	ALL ($N=330$)	m+l-tier ($N=232$)	ALL ($N=330$)	m+l-tier ($N=232$)	ALL ($N=330$)	m+l-tier ($N=232$)
$\Delta \log P_i$	0.00	0.00	$+10.48^{***}$	$+15.47^{***}$	$+4.72$	$+1.53$
	(0.00)	(0.00)	(+2.74)	(+3.58)	(+0.93)	(+0.20)
ΔMARGIN_i	0.00	0.00	$+10.16$	$+19.91^{***}$	-12.87	-6.02
	(0.00)	(0.00)	(+1.54)	(+3.24)	(−1.27)	(−0.64)
ΔULC_i	0.00	0.00	-2.96	-8.64	$+15.96^{*}$	$+6.21$
	(0.00)	(0.00)	(−0.55)	(−1.32)	(+1.78)	(+1.16)
ΔIMP_i	0.00	0.00	$+0.15$	$+0.20$	$+0.15^{***}$	$+0.06$
	(0.00)	(0.00)	(+1.36)	(+1.44)	(+3.07)	(+1.02)
ΔMAT_i	0.00	0.00	$+3.14^{***}$	$+4.00^{***}$	$+1.48^{***}$	$+1.28^{***}$
	(0.00)	(0.00)	(+4.86)	(+4.51)	(+5.21)	(+3.69)
$A\cdot\Delta\text{MARGIN}_i$	0.00	0.00	$+2.66^{***}$	$+3.61^{***}$	-0.52	-0.16
	(0.00)	(0.00)	(+4.38)	(+4.97)	(−0.73)	(−0.20)

Notes: Industry-level panel pooled over three windows: pre-COVID 2017–2019 (omitted reference), Act I 2019–2021, and Act II 2021–2023. Weighted least squares with 2017 industry value added as weights. Industry and window fixed effects throughout (so $\alpha_{\text{pre-COVID}} = \beta_{\text{pre-COVID}} = 0$ by construction). Standard errors clustered on industry. The identification targets are $\beta_{\text{Act I}}$ and $\beta_{\text{Act II}}$. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 22: Supplier-network vs. own exposure: double-exposure robustness.

	$\mathcal{E}_i$ only		Joint: $\hat{\beta}^{\mathcal{E}}$		Joint: $\hat{\beta}^{E}$	
Outcome	Act I	Act II	Act I	Act II	Act I	Act II
$\Delta \log P_i$	$+12.08^{**}$	$+6.10$	$+10.83^{**}$	$+4.59$	$+0.61$	$+0.74$
	(+2.45)	(+0.91)	(+2.19)	(+0.61)	(+0.77)	(+1.18)
ΔMARGIN_i	$+11.83$	-16.84	$+12.97$	-21.12	-0.56	$+2.09$
	(+1.33)	(−1.35)	(+1.40)	(−1.47)	(−0.42)	(+1.17)
ΔULC_i	-4.60	$+19.97^{*}$	-6.98	$+22.69^{*}$	$+1.16$	-1.32
	(−0.65)	(+1.87)	(−0.90)	(+1.77)	(+1.16)	(−0.71)
ΔMAT_i	$+4.60^{***}$	$+2.73^{***}$	$+4.59^{***}$	$+2.81^{***}$	$+0.00$	-0.04
	(+5.13)	(+5.79)	(+4.65)	(+5.59)	(+0.03)	(−0.80)
$A\cdot\Delta\text{MARGIN}_i$	$+3.09^{***}$	-0.93	$+2.91^{***}$	-0.73	$+0.09$	-0.10
	(+4.18)	(−1.05)	(+3.71)	(−0.73)	(+0.99)	(−0.88)

Notes: Pooled three-window panel DiD of §2.2 (pre-COVID 2017–2019 omitted base) with own-exposure interactions $E_i \cdot D^{\text{I}}_w$, $E_i \cdot D^{\text{II}}_w$ added alongside the supplier-network interactions. WLS with 2017 value-added weights; industry and window fixed effects; standard errors clustered on industry. VA-weighted correlation of E_i and $\mathcal{E}_i$: 0.29. $N = 330$ industries. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 23: Cost–price pass-through elevates in the pandemic windows.

Subsample	Pre-COVID 2017–2019	Act I 2019–2021	Act II 2021–2023
ALL	+0.145***	+0.262**	+0.134
	(+3.83)	(+2.32)	(+1.53)
m+l-tier	+0.150***	+0.313*	+0.287***
	(+4.04)	(+1.89)	(+5.31)
Memo, ALL sample (cross-industry):			
sd($\Delta \log$ ULC), pp	11.37	10.00	15.06
sd($\Delta \log P$), pp	4.70	8.15	7.60
corr($\Delta \log P, \Delta \log$ ULC)	+0.29	+0.24	+0.39
N industries	ALL = 330, h-tier = 98, m+l-tier = 232 (across ~57 BEA-66 sectors)		

Notes: Cross-section of BEA Detailed industries (drop 16 EXO and 11 missing JEM exposure). For each window, $\Delta \log P_{i,w}$ regressed on $\Delta \log \text{ULC}_{i,w}$ with BEA-66 sector fixed effects $\gamma_{s(i)}$ absorbing between-sector mean shocks. Weighted least squares with 2017 industry value added as weights; standard errors clustered on BEA-66 sector. Cumulative-log values within each window ($\Delta \log X_{i,w} = \sum_{t \in w} \Delta \log X_{i,t}$). Subsamples: ALL covers all kept industries; m+l-tier restricts to industries within the identified supply chains of §2 that do *not* sell 70 percent or more of output directly to PCE (i.e., excludes the consumer-facing top-tier within each chain). Stars: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

A.11 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 θ_{ci}, input use shares ω_{ci}, domestic shares ζ_{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. (2022) 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. (2022).

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.1.7, and as the realized-inflation benchmark in the expectations regressions of Online Appendix A.8. 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 ρ via the KLEMS panel regression described in Section 2.1. Source:
https://www.bls.gov/productivity/tables/.

10. **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/.

11. **Pasten–Schoenle–Weber (2020) BLS-wide monthly price-change frequencies.** Monthly Calvo reset frequencies θ_m^j at the BEA detail-industry level from Pasten et al. (2020), covering 341 industries. Used for: calibrating both the per-cell Rotemberg adjustment cost $\bar{\phi}$ in the model NKPC (Section 3.3) and the per-industry annual pass-through δ_i in the heterogeneous network operator L used in the cross-industry reduced-form regressions of Section 2. Pass-through-preserving Calvo-to-annual-Rotemberg mapping and BEA-357 crosswalk methodology are documented in Online Appendix A.7. Source: Pasten et al. (2020).

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/.