

Granularity, Networks, and Aggregate Price Volatility

Ellis Scharfenaker^{*†}
Ayixunbati Nuerlantai[‡]

September 28, 2026

Abstract

Idiosyncratic firm-level cost shocks affect aggregate price volatility when they survive aggregation within industries and propagate through the production network. We show that when firm shocks are demeaned within industries and aggregated with lagged sales weights, the relevant concentration term for the granular component includes a finite-firm correction to the Herfindahl. Under linear network propagation and standard independence assumptions, the variance of the granular-network component of inflation is proportional to a statistic combining within-industry concentration and network influence. Using Compustat firm data, BLS producer price indexes, and BEA input-output tables, we find that this component is positively and significantly related to realized producer-price inflation, that the correction removes about 30 percent of the per-industry Herfindahl term and 15 percent of the aggregate statistic, and that granular-network volatility concentrates in a handful of industries.

Keywords: Granularity, production networks, aggregate price volatility, Herfindahl index, input-output linkages, inflation

JEL Classification: E31, E32, L11, L16

1 Introduction

Recent work on aggregate fluctuations is organized around two mechanisms. Granularity holds that shocks to very large firms survive aggregation because the size distribution is skewed (Gabaix, 2011), with related evidence on firm-level volatility and its aggregate consequences in Herskovic et al. (2020). Network propagation holds that shocks to particular sectors have aggregate effects because production is organized through asymmetric input-output linkages (Acemoglu et al., 2012; Carvalho, 2014; Baqaee and Rubbo, 2023). The two are usually studied apart, and primarily for quantities rather than prices (Carvalho and Grassi, 2019; Grassi, 2018), though for prices the separation is artificial, since an idiosyncratic cost shock moves aggregate price volatility only when it survives aggregation within its industry and originates in a sector influential in the production network. The two mechanisms therefore interact for prices. Granularity determines which firm shocks survive within-industry aggregation, while the production network amplifies and redistributes their aggregate effects.

^{*}Department of Economics, University of Utah, 260 Central Campus Drive, Gardner Commons, RM 4100, Salt Lake City, UT, 84112, USA. Email: ellis.scharfenaker@economics.utah.edu

[†]London School of Economics and Political Science, Houghton Street, London, UK

[‡]Department of Economics, University of Utah. Email: ayixunbati.nuerlantai@utah.edu

This paper derives a statistic that joins them for the granular component of price fluctuations. When firm shocks are identified within industries, the variance of the industry-level granular shock is not governed by the Herfindahl alone. Following Gabaix (2011), we remove the equal-weight industry mean, separating the component that remains under symmetric firms from the granular shock generated by unequal firm weights. Under that convention the relevant concentration term becomes

$$g_{i,t-1} = H_{i,t-1} - \frac{1}{n_i},$$

with $H_{i,t-1}$ the within-industry sales Herfindahl and n_i the number of firms. Under linear network propagation, the model-implied variance of the granular-network component of inflation is proportional to

$$\Xi_t = \sum_{i=1}^N v_{i,t}^2 g_{i,t-1},$$

with $v_{i,t}$ the network influence of industry i . When idiosyncratic firm shocks share a common variance and are independent across firms, $\text{Var}(\pi_t^G) = \sigma_\varepsilon^2 \Xi_t$ holds under cross-industry independence, where π_t^G is the granular-network component of inflation, the part driven by the propagated firm-level shocks. Independence across industries discards the covariance between propagated granular shocks. That covariance is positive in the data, so the diagonal variance $\sigma_\varepsilon^2 \Xi_t$ is a lower bound on the variance of the propagated granular component.

The closest predecessor is Grassi (2018), who shows that aggregate output volatility from firm-level shocks depends on the interaction of sectoral Herfindahl indices and network position under oligopolistic competition. The present result moves from output to prices, derives the finite-firm correction governing the demeaned granular component, and tests whether that component appears in realized inflation. The statistic yields an industry ranking,

$$\text{GNC}_{i,t} = \frac{v_{i,t}^2 g_{i,t-1}}{\sum_k v_{k,t}^2 g_{k,t-1}},$$

the granular-network contribution, which combines within-industry concentration and network position and identifies the industries that dominate the diagonal variance of the granular-network component. The contribution concentrates in a few industries, with petroleum and coal products, nondurable wholesale, and transportation equipment at the head, and the finite-firm correction is quantitatively large, removing close to a third of the Herfindahl in the unweighted cross section and about a seventh once industries are weighted by influence.

The granular industry shocks $z_{i,t}$ built from the within-industry residuals, propagated through the Leontief inverse and weighted by influence, form a constructed granular-network component

$$\hat{\pi}_t^G = \sum_i v_{i,t} z_{i,t}.$$

This component tracks realized producer-price inflation positively and significantly. The slope is not on a directly comparable scale, since the constructed shocks are in cost-residual rather than price units and observed inflation also contains the propagated equal-weight industry means and other determinants. We therefore read the slope as an association rather than as a pass-through coefficient. Whether the finite-firm statistic forecasts realized volatility better than the Herfindahl is a separate question, and the raw correlations do not differ statistically while the two statistics correlate above 0.98, so neither has separately identified incremental predictive content once both enter a common regression. We provide a structural account of how granular cost shocks propagate

toward aggregate prices, with the ranking and the finite-firm correction emerging as descriptive outputs.

The result bears on current accounts of inflation. The seller's inflation literature places price adjustment in dominant firms and upstream sectors, where cost shocks, market power, and strategic pricing interact (Weber et al., 2024), and a related literature documents heterogeneous pass-through across firms and industries (Bils et al., 2012; Gopinath and Itskhoki, 2010; Amiti et al., 2019). Pastén et al. (2024) study how price rigidity interacts with network structure across sectors; the present analysis works within sectors and supplies the finite-firm shocks the cross-sector framework treats as exogenous. Alvarez-Blaser et al. (2025) find that firm-level granular residuals account for much of inflation variance in advanced economies, without a structural sufficient statistic or network channel. The statistic derived here supplies the missing structural counterpart.

2 A model-implied statistic for the granular-network component of price volatility

Consider an economy with industries $i = 1, \dots, N$. Industry i contains n_i firms indexed by $j = 1, \dots, n_i$. Let the cost shock to firm j in industry i be

$$\eta_{ij,t} = u_{i,t} + \varepsilon_{ij,t}, \quad \text{Var}(u_{i,t}) = \sigma_u^2, \quad \mathbb{E}[\varepsilon_{ij,t}] = 0, \quad \text{Var}(\varepsilon_{ij,t}) = \sigma_\varepsilon^2,$$

where $u_{i,t}$ is an industry-wide cost innovation and $\varepsilon_{ij,t}$ is an idiosyncratic firm-level innovation. The idiosyncratic innovations are independent across firms within an industry and independent of the vector of industry-wide innovations $\mathbf{u}_t$. We place no restriction on the mean of $u_{i,t}$ or on its covariance across industries. Let $w_{ij,t-1}$ denote firm j 's lagged sales share within industry i , so that

$$\sum_{j=1}^{n_i} w_{ij,t-1} = 1.$$

To isolate firm-specific cost pressure within an industry, first remove the equal-weight industry mean,

$$\tilde{\eta}_{ij,t} = \eta_{ij,t} - \bar{\eta}_{i,t}^{\text{EW}}, \quad \bar{\eta}_{i,t}^{\text{EW}} = \frac{1}{n_i} \sum_{j=1}^{n_i} \eta_{ij,t}.$$

The industry-level granular shock is then constructed by aggregating these demeaned firm shocks with lagged sales weights

$$z_{i,t} = \sum_{j=1}^{n_i} w_{ij,t-1} \tilde{\eta}_{ij,t} = \sum_{j=1}^{n_i} \left(w_{ij,t-1} - \frac{1}{n_i} \right) \eta_{ij,t}. \quad (1)$$

The industry-level granular shock $z_{i,t}$ is the within-industry analogue of the granular residual in Gabaix (2011), applied to cost shocks rather than productivity growth; see also Bräuning et al. (2022) for related evidence on concentrated cost-price linkages.

Since $\sum_j (w_{ij,t-1} - 1/n_i) = 0$, the industry-wide innovation cancels from the second equality and

$$z_{i,t} = \sum_{j=1}^{n_i} \left(w_{ij,t-1} - \frac{1}{n_i} \right) \varepsilon_{ij,t},$$

so the granular shock is driven by the idiosyncratic innovations alone. The second equality also makes clear that the shock survives aggregation only when firm sales shares differ from the equal-weight benchmark. If all firms are equal-sized, then $w_{ij,t-1} = 1/n_i$ for all j and $z_{i,t} = 0$.

Let $\mathbf{A}$ denote the $N \times N$ direct requirements matrix and let

$$\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$$

be the Leontief inverse. Define industry i 's sales-weighted cost shock as

$$c_{i,t} = \sum_{j=1}^{n_i} w_{ij,t-1} \eta_{ij,t}.$$

Since the sales weights sum to one, the demeaning in (1) gives

$$c_{i,t} = z_{i,t} + \bar{\eta}_{i,t}^{\text{EW}}.$$

The sales-weighted industry cost shock therefore consists of the granular component $z_{i,t}$ and the equal-weight industry mean. Let $\mathbf{c}_t$, $\mathbf{z}_t$, and $\bar{\boldsymbol{\eta}}_t^{\text{EW}}$ denote the corresponding vectors across industries. Treating $\mathbf{c}_t$ as the vector of industry unit-cost innovations, under unit cost pass-through and constant value-added prices, log-linearized cost minimization implies

$$\Delta \mathbf{p}_t = \mathbf{L}' \mathbf{c}_t = \mathbf{L}' \mathbf{z}_t + \mathbf{L}' \bar{\boldsymbol{\eta}}_t^{\text{EW}}.$$

Economically, $\mathbf{L}'$ cumulates the direct and indirect cost requirements embodied in downstream prices, so a unit cost shock in industry i raises the prices of industries that use i directly or indirectly as an input.

Aggregate inflation is the sales-share-weighted average of industry price changes,

$$\pi_t = \mathbf{s}_{t-1}' \Delta \mathbf{p}_t,$$

where $\mathbf{s}_{t-1}$ is the vector of lagged industry sales shares in aggregate output. Define the influence vector

$$\mathbf{v}_t = \mathbf{L} \mathbf{s}_{t-1},$$

where $v_{i,t}$ is the marginal effect of a unit cost shock in industry i on aggregate inflation. Aggregate inflation can then be written as

$$\pi_t = \mathbf{v}_t' \mathbf{c}_t = \mathbf{v}_t' \mathbf{z}_t + \mathbf{v}_t' \bar{\boldsymbol{\eta}}_t^{\text{EW}}. \quad (2)$$

Writing

$$\pi_t^G = \mathbf{v}_t' \mathbf{z}_t = \sum_{i=1}^N v_{i,t} z_{i,t}, \quad (3)$$

we define π_t^G as the granular-network component of aggregate inflation. Section 3 constructs its empirical counterpart $\hat{\pi}_t^G = \mathbf{v}_t' \mathbf{z}_t$ and tests whether it is associated with realized producer-price inflation.

The variance of the granular shock follows from (1). Expanding the squared effective weights and using $\sum_j w_{ij,t-1} = 1$ gives

$$\text{Var}(z_{i,t}) = \sigma_\varepsilon^2 g_{i,t-1}, \quad g_{i,t-1} = H_{i,t-1} - \frac{1}{n_i} \quad (4)$$

where $H_{i,t-1} \equiv \sum_j w_{ij,t-1}^2$ is the within-industry sales Herfindahl and n_i is the number of firms in industry i (see Appendix A for details). The term $g_{i,t-1}$ is nonnegative, equals zero when firms are equal-sized, increases with within-industry concentration for fixed n_i , and reduces to the Herfindahl when n_i is large. It coincides with the numerator of the normalized Herfindahl index used in industrial organization, but here it arises as an algebraic consequence of demeaning rather than a definitional choice. The factorization relies on a common idiosyncratic variance σ_ε^2 across firms and on independence of the idiosyncratic innovations across firms within an industry. It places no restriction on the industry-wide innovation $u_{i,t}$, which the demeaning removes.

The complete sales-weighted cost shock $c_{i,t}$ has variance $\sigma_u^2 + \sigma_\varepsilon^2 H_{i,t-1}$, so the Herfindahl governs the idiosyncratic contribution to the variance of the complete shock, while the finite-firm correction $g_{i,t-1}$ governs the variance of its granular component; the two concentration terms coincide as n_i grows.

If the granular industry shocks are independent across industries, the cross terms in the variance of π_t^G vanish and

$$\text{Var}(\pi_t^G) = \sigma_\varepsilon^2 \sum_{i=1}^N v_{i,t}^2 g_{i,t-1}. \quad (5)$$

Independence across industries discards the off-diagonal covariance $\sum_{i \neq j} v_{i,t} v_{j,t} \text{Cov}(z_{i,t}, z_{j,t})$. When this term is positive the retained diagonal term is a lower bound on the variance of the propagated granular component. Under the maintained assumptions, the scale-free statistic

$$\Xi_t = \sum_{i=1}^N v_{i,t}^2 g_{i,t-1} \quad (6)$$

summarizes the model-implied variance of the granular-network component. Influence enters squared because $v_{i,t}$ scales industry i 's contribution to π_t^G , so under cross-industry independence it scales the corresponding contribution to variance by its square. This is the convexity that yields the sales Herfindahl in Gabaix (2011) and the squared influence norm in Acemoglu et al. (2012).

Each industry's granular-network contribution is

$$\text{GNC}_{i,t} = \frac{v_{i,t}^2 g_{i,t-1}}{\sum_{k=1}^N v_{k,t}^2 g_{k,t-1}}, \quad (7)$$

its share of Ξ_t rather than of total realized price volatility, since the omitted covariance is attributable to industry pairs and observed inflation also contains non-granular components. Splitting influence into its direct expenditure share $s_{i,t-1}$ and the excess influence $v_{i,t} - s_{i,t-1}$ generated by the network decomposes Ξ_t into

$$\Xi_t = \sum_i s_{i,t-1}^2 g_{i,t-1} + 2 \sum_i s_{i,t-1} (v_{i,t} - s_{i,t-1}) g_{i,t-1} + \sum_i (v_{i,t} - s_{i,t-1})^2 g_{i,t-1}, \quad (8)$$

a pure granularity, an interaction, and pure network amplification term (see Appendix A). The interaction term is large only for industries that combine high direct weight with strong network amplification and surviving granular variance, so it isolates the complementarity that an industry which is concentrated but peripheral, or central but internally symmetric, does not generate.

Granularity determines whether firm shocks survive within-industry aggregation, while network structure determines their indirect amplification and the distribution of their aggregate effects across industries.

3 Empirical illustration

We compute $g_{i,t-1}$ from Compustat firm-quarter sales shares from 2004-Q4 to 2024-Q4, estimate firm-level cost residuals $\hat{\eta}_{ij,t}$ from a fixed-effects autoregression in operating expenses, and construct $v_{i,t}$ from the 2017 BEA direct requirements matrix.¹ The dataset consists of 76 quarters of 83 industries. The granular industry shocks $z_{i,t}$ follow from (1), and their propagated aggregate component is $\hat{\pi}_t^G$. The realized-inflation comparison uses the NAICS3 industries with usable PPI observations, an average of 33 industries per quarter, and mean lagged-sales-share coverage of 0.725 (minimum 0.672). Figure 1 plots realized inflation and the fitted values from its regression on $\hat{\pi}_t^G$.




Figure 1: Realized producer-price inflation and fitted values from the granular-network regression. The fitted series is $b_0 + b_1 \hat{\pi}_t^G$ from the regression of realized inflation π_t on the constructed granular-network component $\hat{\pi}_t^G = \sum_i v_{i,t} z_{i,t}$. The shaded region is a 95 percent band based on Newey–West standard errors.

Equation (2) decomposes model-implied inflation into the granular-network component π_t^G and the propagated equal-weight industry means. We map the first component to the data by forming $\hat{\pi}_t^G = \sum_i v_{i,t} z_{i,t}$ and regressing realized producer-price inflation on it.² Over 76 quarters, the constructed component tracks realized inflation with a correlation of 0.464, accounting for about 21.5 percent of the variance of quarterly inflation. Newey–West standard errors reject a zero slope in favor of a positive association ($b_1 = 0.486$, $t = 2.6$, $p = 0.011$). The slope absorbs the scale of the operating-expense residuals together with any incomplete pass-through and cannot be interpreted as a structural pass-through coefficient. The estimated intercept is -0.003 with a standard error of 0.006. Because the dependent variable is total observed inflation while $\hat{\pi}_t^G$ captures only its granular-network component, a zero intercept is not imposed by (2); the intercept and regression residual also absorb the propagated equal-weight industry means and other determinants of inflation. The

¹Input-output coefficients are stable across benchmark years, so the choice of benchmark has little effect on the network influence vector.

²The constructed series includes upstream industries that have a cost shock but no producer-price series; dropping them leaves the result essentially unchanged.

fitted series has lower amplitude than realized inflation, as expected when the regressor is in cost-residual rather than price units (Gopinath and Itskhoki, 2010; Amiti et al., 2019). The shocks enter from the cost side, the network is held fixed, and $\hat{\pi}_t^G$ is constructed without reference to inflation, so the co-movement is evidence that the granular-network component appears in realized inflation rather than an artifact of estimation, while the fitted line plotted in Figure 1 is estimated from the regression. Replacing the linear trend in the first-stage cost regression with calendar-time dummies leaves the correlation at 0.454, so the result does not depend on the trend specification.

The statistic Ξ_t retains only the diagonal of the variance of the granular-network component π_t^G . Splitting the sample variance of $\hat{\pi}_t^G$ into the sum of industry-component variances and the residual cross-industry covariance, the off-diagonal carries 39.8 percent of the total and is positive, so the full sample variance is about 1.7 times the diagonal. The largest off-diagonal contributions run between energy and materials sectors, led by oil and gas extraction, petroleum and coal products, and chemicals, the same upstream cluster that heads the ranking. The cross-industry independence maintained in Section 2 is therefore conservative for the variance of the granular-network component, and the discarded covariance is a measured quantity rather than an unexamined residual.

Whether $z_{i,t}$ isolates cost shocks rather than demand bears on the reading of the structural result. The correlation between the firm residual $\hat{\eta}_{ij,t}$ and contemporaneous firm revenue growth is 0.34, and it remains 0.32 after the within-industry demeaning used in the shock construction. These correlations show that the residual co-moves with revenue growth but do not identify how much of that co-movement is demand-driven, since revenue growth can itself respond to cost-induced price changes. We therefore read $z_{i,t}$ as a cost-pressure residual that may contain a nonzero demand component, rather than as a pure cost shock.

The per-industry-quarter mean of $g_{i,t-1}/H_{i,t-1}$ is 0.703, so the finite-firm correction removes about 30 percent of the Herfindahl term, and once industries are weighted by influence the removed share falls to about 15 percent, since high-influence industries carry more firms and are less granular than the average industry. The two statistics nevertheless co-move closely, correlating at 0.988.

The three-way decomposition of Ξ_t shows how the mechanisms combine. Pure granularity accounts for 58.5 percent, the interaction between concentration and network position for 32.2 percent, and pure network amplification for 9.3 percent, so granularity sets the level of the diagonal statistic while the network reshapes which industries generate it. An interaction share above thirty percent is too large to treat as a correction to a pure-granularity benchmark. Pure granularity is led by petroleum, transportation equipment, and computers and electronics, the interaction term by petroleum and nondurable wholesale trade, and pure network amplification by oil and gas extraction, nondurable wholesale trade, and transportation equipment; see Appendix Table D.1.

The finite-firm statistic tracks realized producer-price inflation volatility over the sample. Measuring realized volatility on a rolling twelve-quarter window, the contemporaneous correlation is 0.621, and the Herfindahl approximation reaches 0.596 at the same window. A test for the difference between the two correlations does not reject equality ($Z = 1.57$, $p = 0.117$). In an encompassing regression of realized volatility on both statistics, neither is significant, with $t = 0.8$ for the finite-firm statistic and $t = -0.4$ for the Herfindahl, so the two statistics cannot be separated as predictors at this sample size. The comovement survives deseasonalization and comparison against a broader price index; Appendix Appendix C reports the supporting results.

Figure 2 plots the standardized statistic against standardized realized inflation volatility at the twelve-quarter, contemporaneous specification. The two series track each other over the cycle.

NAICS3	Industry	GNC	Cum. GNC	H_i	g_i	v_i
324	Petroleum and coal products	21.14	21.14	0.225	0.177	0.106
424	Wholesale, nondurable goods	14.02	35.16	0.123	0.105	0.111
336	Transportation equipment	13.13	48.29	0.094	0.084	0.122
334	Computers and electronics	9.86	58.15	0.065	0.063	0.118
517	Telecommunications	6.15	64.29	0.216	0.200	0.054
325	Chemicals	5.66	69.96	0.039	0.037	0.122
211	Oil and gas extraction	5.41	75.37	0.099	0.091	0.069
621	Ambulatory health care	3.95	79.32	0.395	0.378	0.028
518	Data processing and hosting	3.52	82.83	0.284	0.273	0.034
423	Wholesale, durable goods	1.81	84.64	0.051	0.038	0.068
311	Food manufacturing	1.78	86.43	0.087	0.072	0.049
541	Professional and technical services	1.32	87.75	0.033	0.027	0.068
561	Administrative and support services	1.07	88.82	0.057	0.044	0.049
331	Primary metals	0.86	89.68	0.101	0.072	0.034
333	Machinery	0.85	90.53	0.057	0.049	0.041
<i>Time-averaged decomposition of Ξ_t</i>						
Pure granularity						58.5
Interaction						32.2
Pure network amplification						9.3

Table 1: Top panel reports the top industries by average granular network contribution. GNC is the time-averaged granular network contribution in percent, and Cum. GNC is the cumulative share after ranking industries by GNC. Bottom panel reports the time-averaged decomposition of Ξ_t into pure granularity, interaction, and pure network amplification.




Figure 2: Finite-firm granular-network statistic and realized inflation volatility over time. Both series are standardized. Realized volatility is measured on a 12-quarter rolling window, the specification that maximizes the correlation between the two ($r = 0.621$, $n = 66$).

4 Conclusion

Within-industry demeaning of firm-level cost shocks introduces a finite-firm correction, rather than the Herfindahl alone, the concentration term governing the granular component of price volatility. The model-implied variance $\sigma_\varepsilon^2 \Xi_t$ combines this correction with network influence in a single statistic. Granularity determines whether an idiosyncratic shock survives within-industry aggregation, while network structure determines its indirect amplification and the distribution of its aggregate effects.

The central evidence is that the constructed granular-network component, propagated through the Leontief inverse and weighted by influence, is positively and significantly associated with realized producer-price inflation.

Granular-network volatility is concentrated in a small group of industries led by petroleum, nondurable wholesale trade, and transportation equipment, where firm-level shocks both survive aggregation and occupy influential positions in the production network. Industry size and network position therefore reinforce one another in determining aggregate price volatility.

References

- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.
- Alvarez-Blaser, S., Auer, R., Lein, S. M., and Levchenko, A. A. (2025). The granular origins of inflation. NBER Working Paper 33404, National Bureau of Economic Research.
- Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402.
- Baqae, D. and Rubbo, E. (2023). Micro propagation and macro aggregation. *Annual Review of Economics*, 15:91–123.
- Bils, M., Klenow, P. J., and Malin, B. A. (2012). Reset price inflation and the impact of monetary policy shocks. *American Economic Review*, 102(6):2798–2825.
- Bräuning, F., Fillat, J. L., and Joaquim, G. (2022). Cost-price relationships in a concentrated economy. Current policy perspectives, Federal Reserve Bank of Boston.
- Carvalho, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48.
- Carvalho, V. M. and Grassi, B. (2019). Large firm dynamics and the business cycle. *American Economic Review*, 109(4):1375–1425.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.
- Gopinath, G. and Itskhoki, O. (2010). Frequency of price adjustment and pass-through. *The Quarterly Journal of Economics*, 125(2):675–727.
- Grassi, B. (2018). IO in I-O: Size, industrial organization, and the input-output network make a firm structurally important. Working Paper 619, IGIER, Bocconi University.
- Herskovic, B., Kelly, B., Lustig, H., and Van Nieuwerburgh, S. (2020). Firm volatility in granular networks. *Journal of Political Economy*, 128(11):4097–4162.
- Pastén, E., Schoenle, R., and Weber, M. (2024). Sectoral heterogeneity in nominal price rigidity and the origin of aggregate fluctuations. *American Economic Journal: Macroeconomics*, 16(2):318–352.
- Weber, I. M., Lara Jauregui, J., Teixeira, L., and Nassif Pires, L. (2024). Inflation in times of overlapping emergencies: Systemically significant prices from an input–output perspective. *Industrial and Corporate Change*, 33(2):297–341.

Appendix A Derivations

A.1 Finite-firm correction

The granular shock (1) is a weighted sum of independent idiosyncratic innovations with effective weights $w_{ij,t-1} - 1/n_i$, so

$$\text{Var}(z_{i,t}) = \sigma_\varepsilon^2 \sum_{j=1}^{n_i} \left(w_{ij,t-1} - \frac{1}{n_i} \right)^2.$$

Expanding the square and using $\sum_j w_{ij,t-1} = 1$,

$$\sum_{j=1}^{n_i} \left(w_{ij,t-1} - \frac{1}{n_i} \right)^2 = \sum_{j=1}^{n_i} w_{ij,t-1}^2 - \frac{2}{n_i} \sum_{j=1}^{n_i} w_{ij,t-1} + \frac{n_i}{n_i^2} = H_{i,t-1} - \frac{1}{n_i},$$

which gives (4), with n_i replaced by the survivor count $m_{i,t}$ in the empirical implementation.

A.2 Complete versus granular cost shock

Since $\bar{\eta}_{i,t}^{\text{EW}} = u_{i,t} + n_i^{-1} \sum_j \varepsilon_{ij,t}$, the complete sales-weighted cost shock is $c_{i,t} = u_{i,t} + \sum_j w_{ij,t-1} \varepsilon_{ij,t}$, and under the maintained assumptions

$$\text{Var}(c_{i,t}) = \sigma_u^2 + \sigma_\varepsilon^2 H_{i,t-1}, \quad \text{Var}(z_{i,t}) = \sigma_\varepsilon^2 g_{i,t-1}.$$

The industry-wide innovation contributes to the variance of the complete shock but cancels from its granular component. The Herfindahl governs the idiosyncratic part of the former, the finite-firm correction the latter, and the Herfindahl approximation is obtained by dropping $1/m_{i,t}$ from the granular-component variance.

A.3 Decomposition

Writing $v_{i,t} = s_{i,t-1} + (v_{i,t} - s_{i,t-1})$, where $s_{i,t-1}$ is the direct expenditure share and $v_{i,t} - s_{i,t-1}$ is the excess influence generated by the network,

$$\Xi_t = \sum_i s_{i,t-1}^2 g_{i,t-1} + 2 \sum_i s_{i,t-1} (v_{i,t} - s_{i,t-1}) g_{i,t-1} + \sum_i (v_{i,t} - s_{i,t-1})^2 g_{i,t-1}, \quad (9)$$

the pure-granularity, interaction, and network-amplification terms reported in Section 3.

Appendix B Data sources and variable constructions

The analysis uses four datasets.

B.1 Data sources

1. *Compustat North America (Fundamentals Quarterly)*. Firm-level data come from S&P Global Compustat, covering U.S.-incorporated firms ($fic = USA$) from 2004Q4 through 2024Q4. Firms are identified by permanent *gvkey*, and NAICS3 industry is taken from the first three digits of the Compustat NAICS code. We require $naics \geq 100$ and drop NAICS2 sectors 22, 44–45, 52, and 91–92, which are regulated or not primarily market-determined. Industries and industry-quarters with fewer than two firms are excluded. The estimation sample begins in 2005Q1, since four lags of cost are needed for the shock regression. We use operating expenses (*xoprq*) and net revenues (*saleq*), require both to be strictly positive, and on duplicate *gvkey*-quarter records retain the most recent *datadate*. The resulting panel covers 67 NAICS3 industries.

2. *BLS Producer Price Index*. NAICS3 PPI series are converted to quarterly frequency and merged to the Compustat shock panel by industry and calendar quarter. Coverage statistics for the resulting aggregate inflation series are reported in Section 3.

3. *BEA Input-Output Table*. Production-network linkages come from the 2017 BEA Benchmark direct requirements matrix aggregated to NAICS3, an 83×83 matrix **A**. We treat the 2017 matrix

as time-invariant and compute $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$. The IO system has $N_{\text{IO}} = 83$ nodes and the Compustat panel has $N_{\text{CS}} = 67$; after intersecting with the PPI sample, 60 industries remain. IO nodes outside the matched Compustat sample are assigned zeros when constructing full-network objects.

4. *FRED All-Commodities PPI*. For the supplemental comparison in Section 3, we use the monthly PPI for All Commodities (*PPIACO*, NSA, 1982 = 100), accessed via the *fredr* R package, averaged to quarterly frequency and converted to log changes.

B.2 Variable constructions

Within-industry sales weights $w_{ij,t}$ are each firm's share of its NAICS3 industry-quarter total revenue, computed from Compustat quarterly net sales (*saleq*). Let $\mathcal{F}_{i,t}$ denote the firms present in both $t-1$ and t with an estimated cost residual, let $m_{i,t} = |\mathcal{F}_{i,t}|$ be their number, and let the lagged weights be renormalized over $\mathcal{F}_{i,t}$ so that they sum to one over the firms entering the shock. The within-industry Herfindahl and finite-firm correction apply (4) to that set,

$$H_{i,t-1} = \sum_{j \in \mathcal{F}_{i,t}} w_{ij,t-1}^2, \quad g_{i,t-1} = H_{i,t-1} - \frac{1}{m_{i,t}},$$

so the weights, the equal-weight mean, the firm count, and the Herfindahl are all defined over one firm set and (4) holds exactly.

Firm-level cost shocks are residuals $\hat{\eta}_{ij,t}$ from

$$\log x_{ij,t} = \sum_{k=1}^4 \beta_k \log x_{ij,t-k} + \alpha_j + \gamma t + \eta_{ij,t},$$

where $x_{ij,t}$ is operating expenses (*xoprq*), α_j a firm fixed effect, and t a common linear trend. Industry-quarter fixed effects are deliberately omitted so that the industry-quarter common component remains in the residual and is removed afterward by equal-weight demeaning, matching Eq. (1).

The industry-level granular shock is

$$z_{i,t} = \sum_{j \in \mathcal{F}_{i,t}} w_{ij,t-1} (\hat{\eta}_{ij,t} - \bar{\eta}_{i,t}^{\text{EW}}), \quad \bar{\eta}_{i,t}^{\text{EW}} = \frac{1}{m_{i,t}} \sum_{j \in \mathcal{F}_{i,t}} \hat{\eta}_{ij,t},$$

which equals $\sum_{j \in \mathcal{F}_{i,t}} (w_{ij,t-1} - 1/m_{i,t}) \hat{\eta}_{ij,t}$ and has variance $\sigma_\varepsilon^2 g_{i,t-1}$. The corresponding undemeaned sales-weighted industry cost shock is

$$c_{i,t} = \sum_{j \in \mathcal{F}_{i,t}} w_{ij,t-1} \hat{\eta}_{ij,t} = z_{i,t} + \bar{\eta}_{i,t}^{\text{EW}}.$$

The empirical analysis isolates $z_{i,t}$ and constructs the granular-network component $\hat{\pi}_t^G = \sum_i v_{i,t} z_{i,t}$. Recomputing the correction term over the full lagged industry panel leaves Ξ_t correlating at 0.96 with the baseline, with a mean ratio of 1.02.

Lagged industry sales shares are $s_{i,t-1} = S_{i,t-1} / \sum_k S_{k,t-1}$ across matched industries, with $S_{i,t-1}$ lagged industry revenue. The influence vector is $v_{i,t} = [\mathbf{L} \mathbf{s}_{t-1}]_i$, where $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$ is the Leontief inverse of the 2017 BEA direct requirements matrix, with zeros assigned to IO nodes outside the Compustat sample. Ξ_t and $\text{GNC}_{i,t}$ combine $g_{i,t-1}$ and $v_{i,t}$ at the industry-quarter level.

Quarterly PPI inflation is $\pi_{i,t} = \Delta \log P_{i,t}$ on the quarterly-average price index, defined for consecutive quarters, and aggregate inflation is $\pi_t = \sum_i s_{i,t-1} \pi_{i,t}$ across matched industries. Realized volatility is the rolling standard deviation of π_t at windows of 4, 6, 8, 10, 12 quarters. The FRED comparison uses *PPIACO* aggregated to quarterly frequency by the same procedure.

Appendix C Lead-window correlations

Table C.1 reports the correlation between each statistic and realized inflation volatility across rolling-window lengths W and leads k . The finite-firm statistic peaks at $W = 12$, $k = 0$, and the Herfindahl approximation peaks at the same cell.

Window W	Lead k						
	0	1	2	3	4	5	6
<i>Finite-firm statistic Ξ_t</i>							
4	0.164	0.253	0.334	0.387	0.347	0.231	0.188
6	0.245	0.289	0.312	0.333	0.364	0.386	0.343
8	0.316	0.355	0.386	0.381	0.363	0.359	0.372
10	0.488	0.453	0.436	0.440	0.444	0.411	0.366
12	0.621	0.619	0.603	0.527	0.494	0.483	0.456
<i>Herfindahl approximation $\sum_i v_{i,t}^2 H_{i,t-1}$</i>							
4	0.140	0.250	0.336	0.387	0.335	0.202	0.148
6	0.208	0.269	0.302	0.330	0.358	0.374	0.318
8	0.268	0.323	0.363	0.362	0.347	0.345	0.355
10	0.450	0.411	0.400	0.410	0.415	0.385	0.341
12	0.596	0.590	0.569	0.483	0.450	0.444	0.418

Table C.1: Correlation between each statistic and realized inflation volatility across rolling-window lengths W and leads k . Positive k means the statistic leads realized volatility by k quarters. The peak cell of each panel is in bold. The two statistics correlate at 0.988, and their peak correlations, 0.621 and 0.596, cannot be separated as predictors by the encompassing regression of Section 3.

C.1 Robustness of the volatility comparison

Ξ_t is built from unadjusted quarterly Compustat shares and is not seasonally adjusted. A test for a quarterly seasonal in Ξ_t does not reject its absence ($F = 1.13$, $p = 0.34$), and removing calendar-quarter means from both series raises the correlation from 0.621 to 0.639, so the sawtooth in the statistic is not driving the result. Against the all-commodities Producer Price Index (PPIACO) from FRED, which covers the full economy rather than the Compustat–PPI subsample, the contemporaneous correlation rises from 0.132 at the four-quarter window to 0.48 at the twelve-quarter window.

In HAC distributed-lag regressions of twelve-quarter realized volatility on the current and four lagged values of each statistic, the effect concentrates at one- and two-quarter lags, peaking at one. The two statistics fit comparably, with $R^2 = 0.467$ for the finite-firm statistic and 0.431 for the Herfindahl, a gap consistent with their high correlation and not separable in the encompassing regression.

Appendix D Decomposition

Table D.1 reports the leading industries within each component of the decomposition in (8). Pure granularity is dominated by sectors with highly unequal within-industry firm-size distributions; the interaction term highlights sectors where concentration and network position reinforce one another; and pure network amplification brings in sectors whose aggregate relevance comes from network position rather than extreme within-industry concentration.

Component	Industry	Share (%)
<i>Pure granularity</i>		
	Petroleum and coal products	26.93
	Transportation equipment	13.65
	Computers and electronics	12.78
	Wholesale, nondurable goods	11.99
	Telecommunications	6.61
<i>Interaction</i>		
	Petroleum and coal products	17.45
	Wholesale, nondurable goods	16.89
	Transportation equipment	14.29
	Oil and gas extraction	8.92
	Chemicals	7.58
<i>Pure network amplification</i>		
	Oil and gas extraction	12.99
	Wholesale, nondurable goods	11.97
	Transportation equipment	7.16
	Administrative and support services	6.26
	Chemicals	6.14

Table D.1: Top industries within each component of the decomposition of Ξ_t . Shares report each industry’s average contribution to a given component as a percentage of that component’s average total over time.