Supplemental Online Appendix

Monopolistically Skewed Business Cycles: On the Origins of Distributional Fluctuations

A Data Preparation

This appendix describes the construction of the firm-quarter panels used in the empirical analysis.

A.1 Compustat Quarterly Sample

Sample restrictions. We start from quarterly Compustat North America and retain industrial-format, standard, domestic, consolidated records (`indfmt=INDL`, `datafmt=STD`, `popsrc=D`, and `consol=C`). Firms must be incorporated in the United States (`fic=USA`), observations must be reported in U.S. dollars (`curcdq=USD`), the current company-header SIC must be nonmissing and below 9000, and NAICS must be available. Calendar quarters come only from `datacqtr`. We require a valid `datadate` and retain the latest record within duplicate firm-quarter pairs. The baseline quarterly sample covers 1984Q1–2024Q4.

Real sales and growth rates. Nominal quarterly sales are deflated by the GDP price deflator (FRED series `GDPDEF`) to obtain real sales. We compute year-on-year log growth using the exact four-quarter lag. Both sales observations must be reported and positive. Note that we do not interpolate missing quarters.

Cleaning and outlier treatment. The cleaned growth sample requires positive current-quarter sales and assets, usable year-on-year log growth, and the four-quarter sales history needed to measure firm size. We trim the top and bottom 1 percent of log growth observations each quarter, then retain observations with a valid within-industry size classification in two-digit NAICS industries with at least ten firms that quarter. This cleaned sample is used throughout the analysis.

Firm size. We measure firm size in quarter t by average real sales over the preceding four quarters, requiring positive reported sales in each. Each quarter, we sort firms in the cleaned growth sample into size deciles within two-digit NAICS industries. For robustness, we also rank the same firms economy-wide. Subsample analyses retain these rankings unless stated otherwise.

Sample diagnostics. The four-quarter sales-history requirement excludes a further 1.8% of otherwise usable growth observations before trimming. The sample mainly covers firms with established reporting histories: the median observation is 11.0 years after a firm first appears in Compustat, and only 3.8% fall within its first two years.

Table A.1: Summary Statistics for Compustat Data

	Full Compustat	Cleaned Sample	QFR
Assets (mln. USD)	4,170	4,267	43.2
Sales (mln. USD)	414.2	423.3	10.8
Sales Growth (%)	7.3	7.2	0.63
25th percentile of Sales Growth (%)	-7.8	-7.1	-25.3
75th percentile of Sales Growth (%)	20.8	20.1	26.6
Net Leverage (%)	23.4	19.1	20.0
Short-term debt (%)	33.9	33.3	33.0
Obs./quarter	6,392	6,143	6,122
Unique firms	21,089	20,355	–
ρ(Sales Gr., GDP Gr.)	0.75	0.76	–
ρ(Sales Gr., Skew)	0.36	0.77	–
ρ(Skew, GDP Gr.)	0.18	0.51	–

Note: QFR: manufacturing, 1977Q3–2014Q1; statistics from Crouzet and Mehrotra (2020), Tables 1 and 3. Compustat: 1984Q1–2014Q1; sample cleaning follows Appendix A.1. Assets and sales are in 2009 dollars; Skew denotes cumulant skewness.

On average, firms appearing for the first or last time account for 1.8% and 2.0% of quarterly observations, respectively. These shares measure sample turnover rather than economic entry and exit, since firms may leave the sample when they no longer meet the inclusion criteria.

Sample comparison. Table A.1 compares the raw and cleaned Compustat growth samples with the Quarterly Financial Reports sample used by Crouzet and Mehrotra (2020).

A.2 Compustat Global Quarterly Sample

Universe and reporting conventions. We complement the North America sample with quarterly accounting data from Compustat Global over 2000Q1–2024Q4. We retain firms incorporated outside the United States, industrial-format consolidated statements reported under Compustat's historical standard, and observations with valid SIC codes below 9000 and two-digit NAICS codes. Calendar quarters come only from `datacqtr`. We require a valid `datadate` and retain the latest record within duplicate firm-quarter pairs.

Sales growth. We measure growth as the year-on-year log change in quarterly sales, deflated using country-specific annual GDP deflators. Sales must be positive in both quarters, with consistent country and currency reporting. We do not interpolate missing observations and trim the top and bottom 1 percent of growth observations within each country-quarter.

Accounting filters and size. As in the U.S. sample, we require positive sales and assets and measure firm size by average real sales over the preceding four quarters. Sales must be positive throughout this history, with the same country and currency as the current observation. Each quarter, we assign firms to size deciles within two-digit NAICS industries in each country, requiring at least ten firms per group. Unlike in the U.S. sample, these rankings are formed before restricting observations to usable growth rates and trimming growth tails. We retain only observations with a valid size classification.

B Additional Empirical Results

B.1 Aggregate Sales-Growth Turning Points

Figure B.1: Mean Growth Distribution and Turning Points

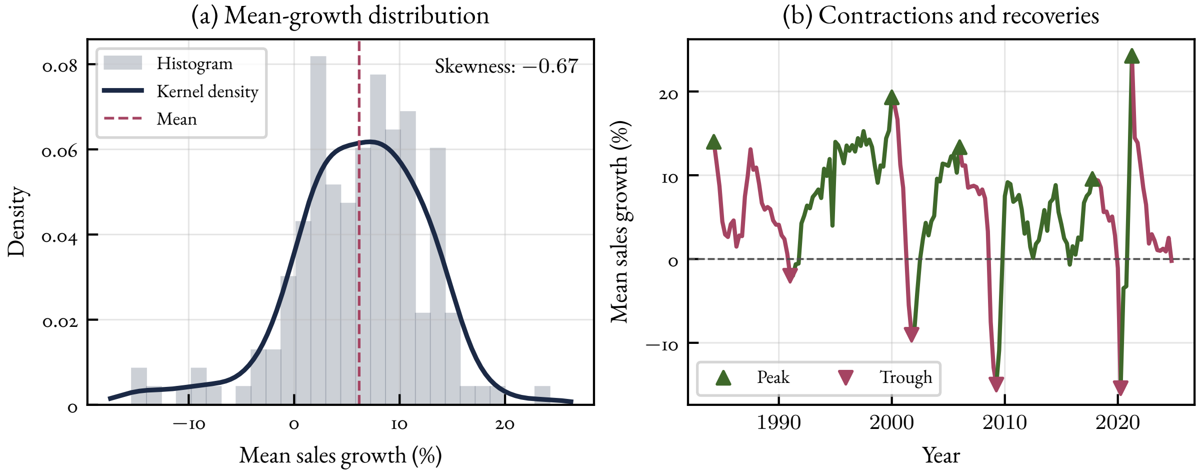

Note: Quarterly mean real-sales growth, 1984Q1–2024Q4. Panel (a) shows its distribution and kernel density; the dashed line marks its time average. Panel (b) shows declines in red and recoveries in green, with markers at confirmed peaks and troughs. Growth is in log-growth percentage points.

Figure B.1 documents the asymmetry of aggregate mean sales growth discussed in Section 2. Panel (a) shows its negatively skewed distribution, with cumulant skewness of -0.67. Panel (b) traces the sharp contractions and subsequent recoveries, including the large positive growth spikes following the 2009 and 2020 troughs. We identify turning points using an adaptation of the algorithm in Dupraz et al. (2025). The algorithm tracks a running maximum until growth falls by at least fourteen percentage points, confirming a peak, then tracks a running minimum until growth rises by the same amount, confirming a trough. This procedure isolates substantial contractions and recoveries while ignoring smaller reversals. These episodes illustrate the impact-and-recovery interpretation of the distributional dynamics and are consistent with the plucking evidence in Dupraz et al. (2025).

B.2 Economy-Wide Size Gradients

Figure B.2 repeats Figure 3 using economy-wide size deciles. This ranks firms by their position in the full Compustat cross section rather than within their two-digit NAICS industry. The sample and size measure follow Appendix A.1 and all other moment and inference choices are identical to Figure 3.

Figure B.2: Economy-Wide Opposing Size Gradients in Growth Moments

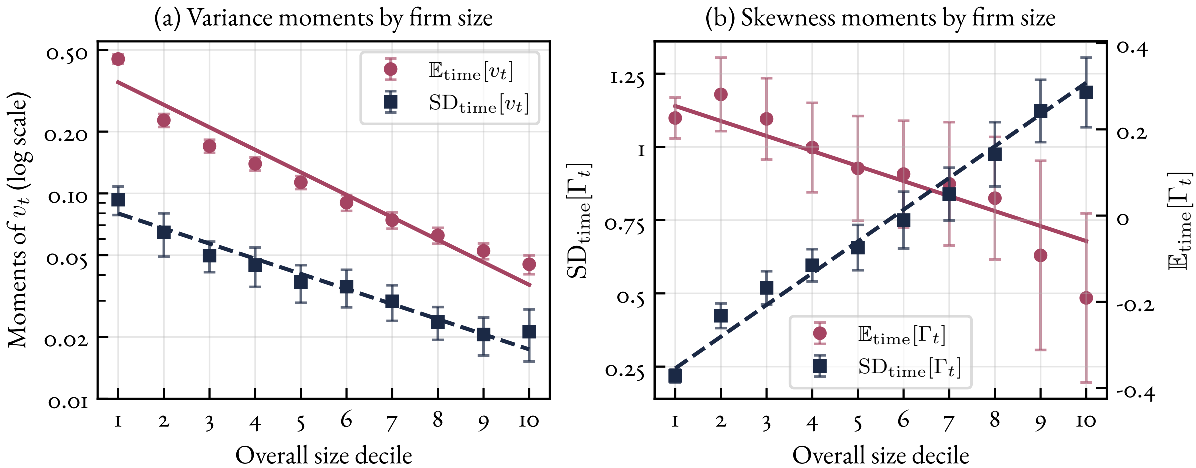

Note: Economy-wide size deciles, 1984Q1–2024Q4. Circles and squares show time-series means and standard deviations of annual log-growth variance in Panel (a) and cumulant skewness in Panel (b). Bars are 90% delta-method Newey–West intervals with four lags. Precision-weighted fits are log-linear in Panel (a) and linear in Panel (b).

Figure B.2 shows essentially the same opposing gradients. Both the time-series mean and standard deviation of growth variance fall with economy-wide size, whereas the time-series standard deviation of cumulant skewness rises sharply. The result therefore does not depend on whether size is measured relative to firms in the same broad industry or relative to the full Compustat cross section.

B.3 Compustat Global Results

We assess the external validity of the main stylized facts using the non-U.S. Compustat Global sample described in Appendix A.2, covering 2000Q1–2024Q4. Time-series aggregates and size-decile moments weight countries equally after calculating moments within countries. We require at least 50 firms per country-quarter and at least 20 qualifying quarters per country, yielding 40 countries and 2,854 country-quarter observations.

Figure B.3 repeats the analysis in Figure 2, expressing each country's mean growth, variance, and skewness as deviations from its own time-series average. Panel (a) averages country-centered growth and country-adjusted cumulant skewness equally across countries in each quarter. Their raw time-series correlation is 0.56. Panel (b) groups the country-quarter observations into 50 equal-count bins and plots the mean growth and variance in each bin. As in Figure 2, we impose a symmetric V-shaped fit by regressing country-adjusted variance on an intercept and absolute country-centered growth. The regression uses the underlying country-quarter observations and yields $\widehat{\beta} = 0.368$. Confidence bands use standard errors clustered by country.

Figure B.3: Growth, Skewness, and Dispersion in Compustat Global

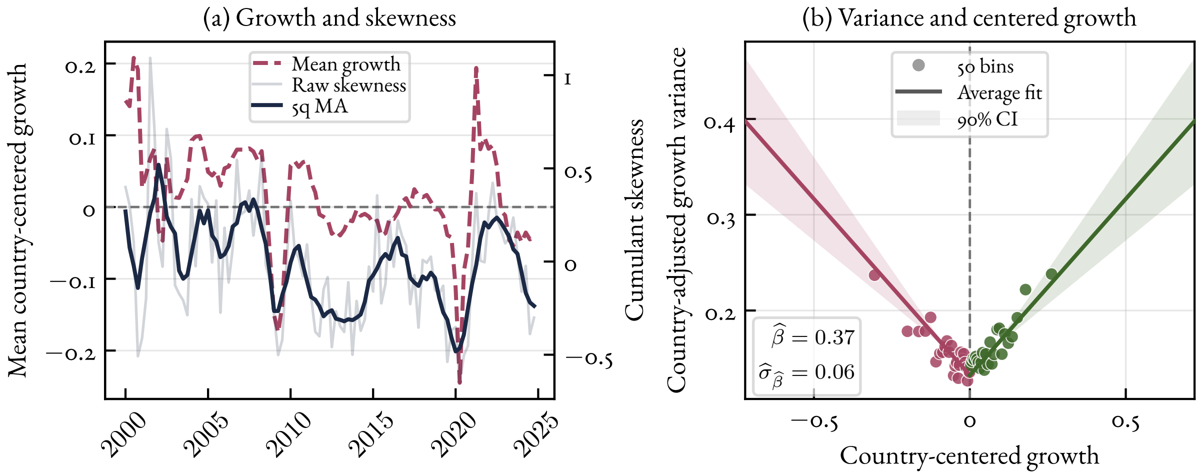

Note: Compustat Global, 2000Q1–2024Q4. Panel (a) equally weights country-adjusted moments: dashed red is mean growth; solid blue is a centered five-quarter average of skewness. Panel (b) shows 50 equal-count bins and a symmetric V-shaped fit with 90% country-clustered confidence bands. Growth is in log units; variance is in squared log units.

Figure B.4 applies the same country-balanced logic to the size gradients in Figure 3. Size deciles follow Appendix A.2. We use unstandardized real-sales growth, calculate moments within each country after pooling firms with the same size decile, and then average the country-specific moments equally. Moments require at least 30 firms per country-quarter-decile, and quarterly averages require at least five countries. These requirements yield 21 contributing countries over 2005Q1–2024Q4.

Panel (a) shows that both the time-series mean and time-series standard deviation of growth variance broadly decline with firm size. Panel (b) shows a noisy decline in mean skewness and a broad increase in its time-series standard deviation across size deciles. Thus the Global sample exhibits the same broad size dependence without assigning disproportionate weight to countries with larger listed-firm populations.

Figure B.4: Growth Variance and Skewness by Firm Size in Compustat Global

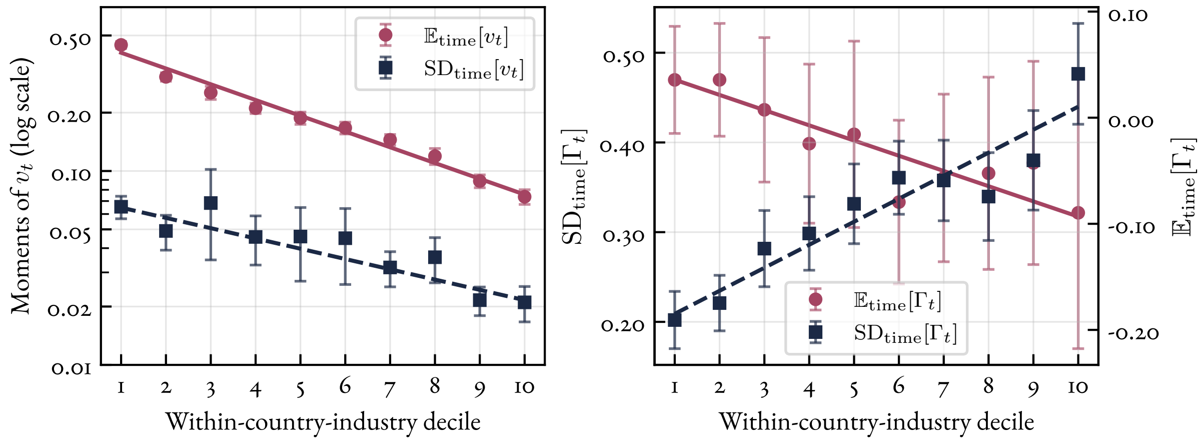

Note: Countries receive equal weight within each quarter-decile. Circles and squares show time-series means and standard deviations of annual log-growth variance in Panel (a) and cumulant skewness in Panel (b). Bars are 90% delta-method Newey–West intervals with four calendar-quarter lags. Precision-weighted fits are log-linear for variance and linear for skewness.

B.4 Inflation Exposure and Growth-Distribution Changes

The inflation exposures of Davis et al. (2025) allow us to repeat the exposure-group comparison for another episode. Figure B.5 applies the procedure in Section 5.4, comparing outer exposure quintiles (1 and 5) with middle quintiles (2–4). The balanced sample spans 2021Q2–2023Q2 and contains 1,086 firms across nine industries. We measure size and moment changes relative to 2022Q1, holding exposure groups, industry–size cells, and common composition weights fixed.

Figure B.5: Inflation Exposure and Growth-Distribution Changes

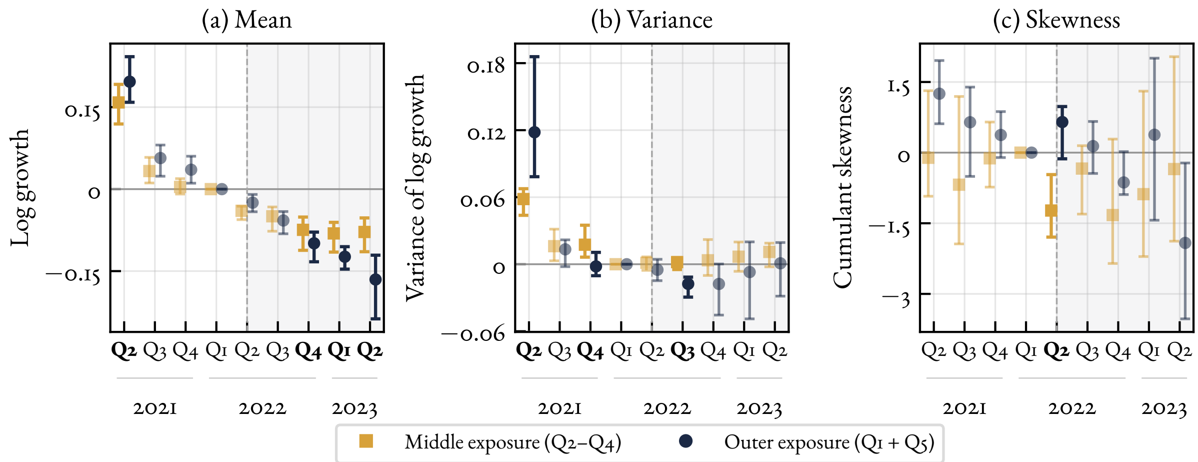

Note: Changes from 2022Q1 at common industry–size shares. Yellow squares: exposure quintiles 2–4; blue circles: 1 and 5. Mean and variance use log and squared log units. Bars: 90% pointwise percentile intervals from 5,000 industry-pairs bootstrap draws. Bold labels and darker markers and whiskers denote significant outer-minus-middle changes. Shading begins at 2022Q2.

The figure shows a different pattern from COVID. In 2022Q2, variance changes little in either group, while skewness rises among outer-exposure firms and falls among middle-exposure firms. The outer-minus-middle variance contrast is -0.006, with a 90 percent pointwise interval of $[-0.021, 0.009]$. The skewness contrast is positive at 1.88, with an interval of $[0.67, 2.52]$, reversing the COVID ordering.

Two features limit what we can learn from this reversal. First, the variance paths already differ before 2022Q2. Pre-event skewness contrasts, by comparison, do not individually exclude zero at the 10 percent level. Second, inflation rose over several quarters, overlapping the pandemic recovery and multiple news events. The exposure measure itself combines equity-market reactions in May and September 2022, including news about earnings and expected interest rates as well as inflation. The comparison therefore captures changes over an inflation episode whose onset and information content are less clearly defined than those of COVID.

B.5 Risk-Exposure Moments by Firm Size

This subsection checks whether size-dependent exposure heterogeneity can account for the empirical size gradients in growth variance and skewness. We use the DHS pandemic and inflation exposures of Davis et al. (2025), centered within each shock and size group, as proxies for the model's exposure loadings. Their return-based sign convention need not coincide with that of the model's cost loadings.

The diagnostic follows directly from Proposition 1. Equation (9) shows that growth variance depends on the slope of the real-sales policy, $\mathcal{R}'(0)$, and on the cross-sectional dispersion of exposure, $\mathbb{V}[\lambda]$. Hence a systematic size gradient in $\mathbb{V}[\lambda]$ could mechanically contribute to a size gradient in growth variance. Equation (10) shows that skewness is scaled by real-sales curvature, $\rho^\star_{\mathcal{R}}$, but also contains the exposure term $\mathbb{V}[\lambda^2]$, normalized by the $3/2$ power of overall shock variance. When idiosyncratic risk is negligible, the corresponding exposure diagnostic is $\mathbb{V}[\lambda^2]/\mathbb{V}[\lambda]^{3/2}$, equal to $(2/3)C_\lambda$ in Section 3.3. If the distribution of λ_i's were sharply different between small and large firms, these exposure moments could partly substitute for differences in $\mathcal{R}'(0)$ and $\rho^\star_{\mathcal{R}}$. We therefore sort firms into size deciles and ask whether the relevant exposure moments vary systematically with size.

A second diagnostic concerns the maintained symmetry approximation in Proposition 1. An asymmetric exposure distribution could generate growth skewness even before local curvature of $\mathcal{R}$ becomes important; its sign depends on how exposures translate into growth responses. We therefore also examine the skewness of the exposure proxies across size groups.

We match the untrimmed DHS exposures to the Compustat sample and use the within-industry size percentiles defined in Appendix A.1. The reference quarters are 2020Q1 for COVID and 2022Q2 for inflation, so the underlying sales history precedes the respective events. Separately for each measure, we sort matched firms by their size percentile into ten equal-count groups. The samples contain 1,432 COVID and 1,447 inflation firms.

Figure B.6 reports the exposure moments for the DHS pandemic and inflation measures. Let x_{ic} denote the raw DHS proxy and $\bar{x}_{cj}$ its mean within shock c and size decile j. For firms in that group, define the centered proxy $\lambda_{ic} = x_{ic} - \bar{x}_{cj}$. Panel (a) reports the raw-proxy mean $\bar{x}_{cj}$; the remaining

Figure B.6: DHS exposure moments by firm size

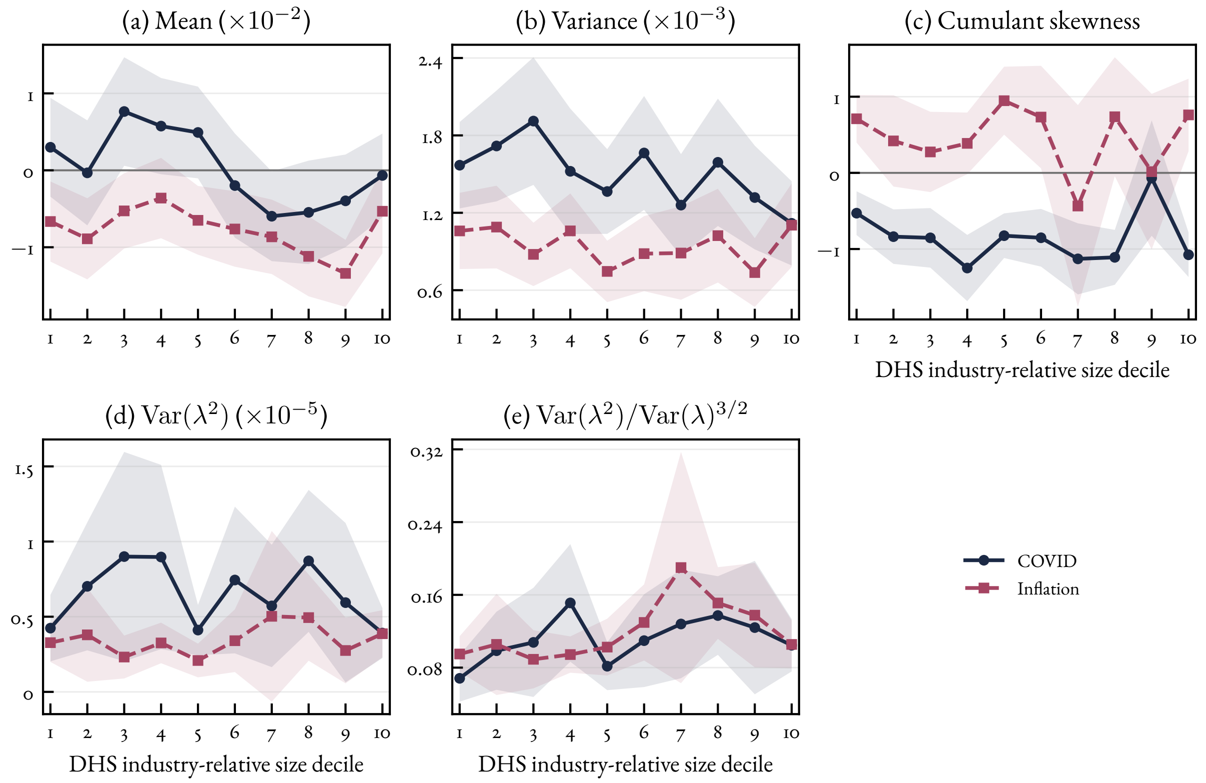

Note: Ten equal-count groups per exposure measure, ordered by within-industry size percentile. Panel (a) reports raw-proxy means; the remaining panels use exposures centered within each shock and size group. Shaded regions are pointwise 95% normal intervals using standard errors from 1,999 firm-bootstrap draws within fixed groups.

panels report variance, cumulant skewness, $\mathbb{V}[\lambda^2]$, and $\mathbb{V}[\lambda^2]/\mathbb{V}[\lambda]^{3/2}$. Bootstrap inference resamples firms within fixed shock–size groups and recenters exposures in each draw.

The exposure profiles are uneven across size groups for both shocks. Inflation exposure variance has no clear size ordering, and its scaled $\mathbb{V}[\lambda^2]$ statistic peaks in the seventh group. For COVID, exposure variance tends to decline with size. Its scaled $\mathbb{V}[\lambda^2]$ statistic varies unevenly: it is higher on average across the upper five deciles but peaks in the fourth decile. For inflation, the upper five deciles also have a higher average scaled statistic and slightly lower average exposure variance. Some comparisons for both shocks therefore point in the direction required by an exposure-based explanation. However, the uneven profiles, wide confidence intervals for the higher-order moments, and differences across shocks leave the evidence inconclusive about a common exposure explanation of the opposing growth-moment gradients in Equations (9) and (10). The skewness panel qualifies the maintained symmetry approximation. COVID exposures are left-skewed in every bin, whereas inflation exposures are right-skewed in nine of the ten groups. The COVID pattern means that asymmetric exposure incidence may contribute to pandemic growth skewness. Its skewness is not monotone by size; this evidence alone does not establish a separate explanation of the small-large gradient.

C Firm-Level Responses to Aggregate Shocks: Estimation and Inference

This appendix documents the shocks, firm projections and distributional comparisons in Section 5.2. It derives the adjustment for uncertainty in estimated firm responses, explains how coefficient error affects skewness, and reports sensitivity to variance adjustment.

C.1 Shock Sources and Controls

Selected innovations. The main business-cycle shock is the original replication-package construction of Angeletos et al. (2020), ending in 2017Q4. Its sign is positive when unemployment rises. Credit innovations follow the excess-bond-premium construction of Gilchrist and Zakrajšek (2012): a recursive quarterly VAR with two lags, ordered as consumption growth, nonresidential investment growth, GDP growth, GDP-deflator inflation, the excess bond premium, stock-price growth, the ten-year Treasury yield and the federal funds rate. We refit the VAR through 2019Q4 and use its orthogonalized excess-bond-premium innovation.

Oil supply news uses the pre-Covid monthly series of Känzig (2021), summed over three complete months. Monetary policy uses the U.S. monetary component of Jarociński and Karadi (2020), which separates policy tightening from central-bank information using interest-rate and stock-price surprises. We use the updated author decomposition restricted to 1994–2019. Investment news follows Ben Zeev and Khan (2015): we refit the nine-state VAR through 2019Q4, retaining its four lags and 60-quarter news-identification objective. Financial uncertainty uses the max-G financial shock of Ludvigson et al. (2021), refitted through December 2019 with CPI-deflated gold returns in the identifying input. The refits retain the authors' identifying restrictions and use revised historical inputs.

Signs, scales and dates. For innovation ϵ_t^s, set

$$x_t^s = \omega_s \epsilon_t^s / \widehat{\mathrm{SD}}(\epsilon^s), \qquad \omega_s \in \{-1, 1\}, \tag{C.1}$$

with ω_s chosen so positive values denote adverse innovations. The scale is the population standard deviation over 1984Q1–2019Q4, except 1984Q1–2017Q4 for the main business cycle and 1994Q1–2019Q4 for monetary policy. It is fixed across firms, horizons and bootstrap draws. The adverse component is $\max\{x_t^s, 0\}$; the favorable component is $\min\{x_t^s, 0\}$, so the response to a negative one-unit innovation is minus its coefficient. Both components enter jointly.

Macro controls. Every firm LP contains two lags of each state and two lags of own growth, with spell-specific slopes. All shocks use log real GDP, log GDP deflator and the excess bond premium in levels; the firm projections omit spell trends. Spell-specific slopes allow heterogeneous responses to these prior macroeconomic conditions alongside the common time fixed effects.

C.2 Firm Spells and Relative Responses

Estimation sample. A *spell* is an uninterrupted sequence of quarterly observations for one firm. We apply the accounting restrictions and quarterly growth trim in Appendix A. Growth is $g_{jt} = \log S_{jt} - \log S_{j,t-4}$, where S is reported sales deflated by the GDP deflator. Growth, sales-size histories, industry codes, shocks and controls must be observed at the dates needed for the regression. A missing or growth-trimmed observation ends a spell; subsequent usable observations form a new spell. A change in the firm's size group does not end a spell.

Within each shock's sample window, we retain every such sequence containing at least 40 quarters of usable growth. A firm can contribute several nonoverlapping spells, each of which is estimated separately using the full uninterrupted sequence. Within-industry size ranks need not be available in every quarter for a spell to enter response estimation. A length-L spell supplies $L - 2 - h$ observations at horizon h: 38 at $h = 0$ and 32 at $h = 6$ when $L = 40$. Lags and forward outcomes are constructed within each spell. The firm-specific design must have full column rank and positive residual degrees of freedom at $h = 6$; earlier-horizon designs contain those rows. The roster uses the common support of the light and paper-control specifications considered during sample construction. Every length-eligible candidate passes the rank and degrees-of-freedom checks. The same spells enter every horizon, while observation counts decline with h.

Size groups for the moment comparison. We use quarterly within-two-digit-NAICS ranks of average real sales over $t - 4, \ldots, t - 1$, computed in the broad accounting sample. To enter the small-firm (large-firm) group, the firm must have an observed rank in every quarter of the spell and remain in the lower (upper) half throughout, including quarters used for lags or forward outcomes. Spells in which a firm crosses the median or lacks ranks in some quarters still enter response estimation and the average response used to select the horizon of the largest contraction. They are excluded from the comparison of response moments across size groups. A firm can contribute an earlier spell to the small-firm group and a later spell to the large-firm group. Industry membership follows the firm's reported code at each date.

Table C.1: Firm-Spells and Stable Size Groups

	Firms	Spells	Stable small	Stable large	Reporting spells
Main business cycle	6,720	6,936	2,011	2,505	4,470
Credit spread	6,860	7,102	2,065	2,541	4,555
Oil supply news	6,860	7,102	2,065	2,541	4,534
Monetary policy	5,125	5,181	1,557	1,963	3,488
Investment news	6,860	7,102	2,065	2,541	4,535
Financial uncertainty	6,860	7,102	2,065	2,541	4,535

Note: Firms and spells count all estimated responses. Stable-size columns count spells before cell eligibility. Reporting spells count distinct spells in the selected-horizon reporting cells. No coefficient trimming is applied. A firm may contribute more than one spell.

Time fixed effects and normalization. Equation (28) is joint OLS with spell-specific intercepts, slopes, and common time fixed effects. Because the shocks vary only over time, a common shift in the signed shock slopes can be absorbed by the time fixed effects. We normalize the equal-spell mean of each signed slope to zero. This identifies relative responses and leaves cell-centered variance and skewness unchanged. Spell-specific macro coefficients capture differences in firms' responses to prior aggregate conditions; their common component is absorbed by the time fixed effects.

Selecting the horizon of the largest average contraction. We select the horizon at which the estimated average response is most negative. Time fixed effects absorb the common response to an aggregate shock, so the main regression identifies only relative responses, whose mean is normalized to zero at every horizon. To estimate the average response needed for horizon selection, we therefore use an auxiliary regression that omits time fixed effects but otherwise uses the same observations and regressors. The response moments are computed from the main regression with time fixed effects. With J_s estimated spells, choose

$$h_s^\star = \arg\min_{h\in\{0,\dots,6\}} \frac{1}{J_s}\sum_{j=1}^{J_s} \widehat{b}_{jh}^{s,+,\text{aux}}, \tag{C.2}$$

breaking ties at the earliest horizon. Every estimated spell has equal weight, including movers and spells with incomplete size labels. A firm with two spells contributes twice. We select the horizon using untrimmed mean responses before forming the size-group moment contrasts. Selected horizons are 1, 3, 5, 4, 5 and 3 for the main business-cycle, credit, oil, monetary-policy, investment-news and uncertainty shocks, respectively.

C.3 Cell Moments and Aggregation

We retain all estimated coefficients without coefficient trimming; the quarterly growth trim used to construct the regression sample is unchanged.

Reporting support. Form industry–date cells at positive-shock origins usable at $h_s^\star$. Require ten stable-small and ten stable-large firms per cell and at least five qualifying industries per date. Keep these cells and their aggregation weights fixed across bootstrap draws. All point and bootstrap cells contain at least ten coefficients and have positive raw variance. A spell coefficient can appear at several dates, so its total weight reflects its presence in the eligible reporting cells.

Variance adjustment. Estimated shock responses contain sampling error. Even if true responses were identical within a cell, their estimates would be dispersed. This source of variation matters for size comparisons when response estimates differ in precision across small and large firms.

Suppress shock and horizon indices. For a fixed cell I of n firms, write $\widehat{\boldsymbol{b}}_I = \boldsymbol{b}_I + \boldsymbol{e}_I$, where $\boldsymbol{b}_I$ contains true adverse-shock responses and $\boldsymbol{e}_I$ their estimation errors. Let $M_I = I_n - \mathbf{1}_n\mathbf{1}_n'/n$. If coefficient errors are centered conditional on the true responses and have covariance $\Omega_{e,I}$, then

$$\begin{aligned}\mathbb{E}\left[\frac{1}{n}\widehat{\boldsymbol{b}}_I' M_I \widehat{\boldsymbol{b}}_I \,\middle|\, \boldsymbol{b}_I\right] &= \frac{1}{n}\boldsymbol{b}_I' M_I \boldsymbol{b}_I + \frac{1}{n}\mathbb{E}[\boldsymbol{e}_I' M_I \boldsymbol{e}_I \mid \boldsymbol{b}_I] \\ &= \frac{1}{n}\boldsymbol{b}_I' M_I \boldsymbol{b}_I + \frac{1}{n}\operatorname{tr}(M_I \Omega_{e,I}).\end{aligned} \tag{C.3}$$

The cross term vanishes because coefficient errors are centered. The trace term is the contribution of estimation uncertainty to centered coefficient variance, including covariance between firms' errors. Subtracting its estimate gives Equation (29).

We implement the correction using the estimated covariance matrix of coefficient errors, $\widehat{\Omega}_{e,I}$. This matrix captures both uncertainty in individual response estimates and covariance between firms' estimation errors. For the coefficients in each reporting cell, adjusted variance is

$$\widehat{V}_I^{\text{adj}} = \frac{1}{n}\widehat{\boldsymbol{b}}_I' M_I \widehat{\boldsymbol{b}}_I - \frac{1}{n}\operatorname{tr}(M_I \widehat{\Omega}_{e,I}). \tag{C.4}$$

The first term, denoted $V_I(\widehat{\boldsymbol{b}})$, is the observed dispersion of estimated responses; the second is its estimated component attributable to estimation error. Appendix C.4 describes how the calendar-block procedure supplies the covariance estimate and lower confidence bounds. All moments use denominator n.

The correction estimates an economic variance by subtracting estimated noise from observed dispersion. If the noise estimate exceeds observed dispersion, the corrected estimate is negative. We retain such estimates to show the behavior of the subtraction; their interpretation is that the estimated noise exhausts the observed dispersion. An excessively large noise correction lowers the point estimate, especially in groups with less precise coefficients, and can therefore alter the estimated size gradient as well as its uncertainty.

Size-group contrasts. For moment $M \in \{V^{\text{adj}}, \Gamma\}$, let T_s be the retained adverse dates and K_{st} their eligible industries. Report

$$\Delta M_s^\star = \frac{1}{|T_s|}\sum_{t\in T_s}\frac{1}{|K_{st}|}\sum_{k\in K_{st}}\left(\widehat{M}_{skt,\text{small}} - \widehat{M}_{skt,\text{large}}\right). \tag{C.5}$$

The model predicts positive variance and skewness contrasts. The latter means that small-firm responses are less negatively skewed. Skewness is calculated within each cell before averaging the small-minus-large differences.

C.4 Bootstrap Inference and Coefficient Reliability

Calendar-block draws. We use a blockwise wild construction (Yeh, 1998; Shao, 2010) with OLS residuals. Divide the master calendar into nonoverlapping six-quarter blocks beginning in 1984Q1. For each of 1,000 draws, assign independent Rademacher signs to the 24 blocks through 2019Q4. Signs are shared across firms, spells, horizons and shocks, preserving common movements across firms and dependence within calendar blocks.

For each draw, multiply the OLS residuals by the assigned block signs and add them to the fitted outcomes. Refit Equation (28) on these outcomes, re-estimating the shock coefficients, nuisance coefficients and time fixed effects with the original regressors and normalizations. Let $\widehat{\boldsymbol{b}}^{*(r)}$ contain the resulting adverse-shock coefficients and define their deviations from the original estimates as $\boldsymbol{e}^{(r)} = \widehat{\boldsymbol{b}}^{*(r)} - \widehat{\boldsymbol{b}}$. The covariance of these deviations implied by the resampling scheme supplies $\widehat{\Omega}_{e,I}$ for each retained cell I.

Form moments in the original reporting cells in every draw. For the fixed set I of coefficients, with $n = |I|$, the variance and skewness draws are

$$\begin{aligned} V_I^{*(r)} &= V_I(\widehat{\boldsymbol{b}} + \boldsymbol{e}^{(r)}) - V_I(\boldsymbol{e}^{(r)}) - \frac{1}{n}\operatorname{tr}(M_I \widehat{\Omega}_{e,I}), \\ \Gamma_I^{*(r)} &= \Gamma_I(\widehat{\boldsymbol{b}} + \boldsymbol{e}^{(r)}). \end{aligned} \tag{C.6}$$

Perturbing the coefficients adds another layer of estimation noise. The variance draw subtracts the realized variance of that perturbation and the original plug-in noise correction, computing both on the fixed cell. We aggregate the small-minus-large contrasts with the original cell weights. Standard errors are the sample standard deviations of the 1,000 contrast draws. For $H_0 : \Delta \leq 0$ against $H_1 : \Delta > 0$, the one-sided confidence-c lower bound is the point estimate minus z_c standard errors, where z_c is the standard-normal quantile and $c \in \{0.90, 0.95, 0.99\}$. These pointwise bounds condition on the source shocks, regressors, firm-spell sample, size classifications and selected horizon.

Skewness and coefficient reliability. We report the unadjusted skewness of the retained coefficients,

$$\widehat{\Gamma}_I = \frac{n^{-1}\sum_{j\in I}(\widehat{b}_j - \bar{\widehat{b}}_I)^3}{V_I(\widehat{\boldsymbol{b}})^{3/2}}, \tag{C.7}$$

where $\bar{\widehat{b}}_I$ is the cell mean. For a randomly drawn firm j within a size–industry–date group, write $\widehat{b}_j = b_j + e_j$, where b_j is the true response and e_j its estimation error. Adding independent, centered, unskewed estimation error leaves the third central moment of the response unchanged but increases its variance from $V(b_j)$ to $V(b_j) + V(e_j)$. Thus

$$\Gamma(\widehat{b}_j) = r^{3/2}\Gamma(b_j), \qquad r := \frac{V(b_j)}{V(b_j) + V(e_j)} = \frac{\mathrm{SNR}}{1 + \mathrm{SNR}}, \qquad \mathrm{SNR} := \frac{V(b_j)}{V(e_j)}. \tag{C.8}$$

Coefficient noise attenuates standardized skewness toward zero in this benchmark. Equal reliability across the two size groups within a cell multiplies their skewness difference by the same positive factor, preserving its sign. Differences in reliability imply different attenuation factors.

For a fixed cell with positive estimated signal, estimated reliability is $\widehat{r}_I = \widehat{V}_I^{\mathrm{adj}} / V_I(\widehat{\boldsymbol{b}})$: the share of observed coefficient variance attributed to response heterogeneity. Table C.2 reports noise shares and attenuation factors formed from the weighted group variance totals.

Estimated noise shares range from 44 to 78 percent and attenuation factors from 0.10 to 0.42.

Averaged equally across the six shocks, attenuation factors are nearly identical for small and large firms, at 0.227 and 0.234, respectively. Small-firm attenuation factors are lower for credit, monetary policy and investment news, and higher for the other three shocks.

Table C.2: Coefficient Noise Shares and Skewness Attenuation

	Noise share (percent)		Attenuation factor	
Shock	Small	Large	Small	Large
Main business cycle	68.6	76.0	0.176	0.118
Credit spread	60.8	53.4	0.245	0.318
Oil supply news	64.3	69.0	0.213	0.173
Monetary policy	78.2	72.8	0.102	0.142
Investment news	64.3	59.0	0.213	0.262
Financial uncertainty	44.3	46.7	0.415	0.389
Average			0.227	0.234

Note: Diagnostics use the baseline block(6) noise correction and weighted group variance totals. Noise share is estimated noise divided by observed variance; attenuation is the corrected-to-observed variance ratio to the power $3/2$. The latter is a benchmark under independent, centered, unskewed coefficient error, not an exact correction to average cell-level skewness. The final row averages attenuation factors equally across the six shocks.

C.5 Robustness Comparisons

We vary coefficient-noise subtraction while retaining the six shocks, controls, firm projections and selected horizons of Figure 5. All comparisons use the same 1,000 calendar-block wild draws with OLS residuals, fixed reporting cells and weights, and no coefficient trimming.

Coefficient noise. Figure C.1 reports unadjusted variance on the left and HAC(8)-adjusted variance on the right, leaving the coefficients, labels and reporting weights unchanged. The unadjusted bootstrap variance is $V_I(\widehat{\boldsymbol{b}} + \boldsymbol{e}^{(r)})$. HAC(8) replaces the plug-in noise covariance with the Bartlett-weighted covariance of calendar-quarter coefficient scores through eight lags; the raw six-quarter bootstrap is unchanged. All six variance contrasts have positive 95 percent lower bounds under both alternatives. The dispersion ordering is therefore already present before the noise correction. Comparing the raw and adjusted moments shows how much of measured dispersion is attributed to uncertainty in the estimated responses. Skewness is unchanged from Figure 5.

Figure C.1: Unadjusted and HAC(8)-Adjusted Coefficient Variance

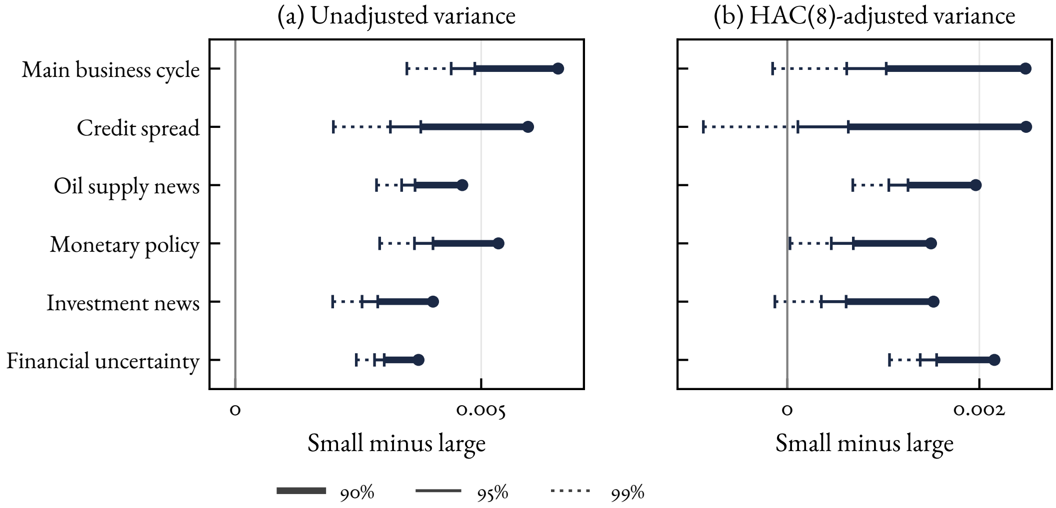

Note: Small-minus-large variance contrasts using the Figure 5 coefficients, stable within-industry size labels and industry/date weights, without coefficient trimming. The left panel omits noise subtraction; the right uses HAC(8) for the plug-in correction, with the same raw six-quarter bootstrap. Bars show one-sided 90, 95 and 99 percent pointwise lower bounds from 1,000 draws. Variance is in squared log-growth units; panel axes use different scales.

D Relative Policy Estimates

D.1 Estimation

We apply Proposition 8 to quarterly growth moments in the baseline within-industry size deciles, requiring at least 30 firms per cell.[25] Mean growth m_t is demeaned over the full sample before selecting estimation dates. As in Section 5.3, we retain the 90% of balanced quarters with the smallest m_t^2, leaving 147 quarters.

To limit noise in the curvature adjustment, we fit a smooth positive curve to the estimated variance ratios:

$$\left(\frac{\widehat{\alpha}_j}{2\widehat{\beta}_j}\right)^F = \exp(a + bj + cj^2 + dj^3). \tag{D.1}$$

The fit uses nonlinear least squares in levels. Fitted values enter only the adjustment factor A_{jt}; relative idiosyncratic volatility uses the raw ratios. Curvature comparisons require $A_{jt} \geq 0.05$ in both the relevant decile and Decile 5, with at least eight qualifying quarters. Each baseline curvature comparison retains 139 of the 147 quarters.

The upward curvature gradient persists when we vary this cutoff between 0.01 and 0.075, although individual estimates change. Varying exposure shape over $\mathcal{K}_\lambda \in \{1, 2, 4\}$ has little effect on the estimates. The cutoff comparisons reselect dates; the exposure-shape comparisons keep the baseline dates.

[25] We use sample variance and the standardized third central moment.

Table D.1: Rank Statistics for Relative Policy Estimates

	Spearman			Kendall		
Outcome	ρ_s	90% range	Opp. sign	τ_b	90% range	Opp. sign
Policy slope	-0.988	$[-1.000, -0.952]$	< 0.001	-0.956	$[-1.000, -0.867]$	< 0.001
Idiosyncratic volatility	-0.952	$[-0.976, -0.758]$	< 0.001	-0.867	$[-0.911, -0.556]$	< 0.001
Curvature magnitude	0.939	$[0.709, 0.976]$	< 0.001	0.822	$[0.556, 0.911]$	< 0.001

Note: Central 90% ranges from 50,000 joint log-Gaussian draws using Newey–West covariance estimates with four calendar-quarter lags. Opp. sign denotes the share of draws with the opposite sign or zero. Estimates are normalized at Decile 5.

D.2 Confidence Intervals

We use the delta method with Newey and West (1987) covariance estimates and four calendar-quarter lags. We calculate covariances from the moments' influence functions, using numerical derivatives for the fitted curve. This accounts for dependence between each estimate and its Decile-5 reference and, for curvature, uncertainty in the fitted variance ratios. The 90% normal intervals condition on the selected dates and full-sample centering of mean growth. Decile 5 is fixed at one by normalization.

D.3 Rank Statistics

Table D.1 summarizes uncertainty about the ordering across size deciles. We recompute Spearman's correlation and Kendall's τ_b in 50,000 draws from a joint log-Gaussian approximation using the estimated covariance matrix of the log estimates (Krinsky and Robb, 1986). Drawing in logs preserves positivity, and Decile 5 remains fixed at one. The table reports the central 90% of the simulated rank statistics and the shares with the opposite sign to the estimate. These describe the stability of the rankings under the fitted distribution.

D.4 Industry Composition

The three size gradients also appear when we rank firms across the whole economy (Figure D.1).

Larger industries can still dominate the pooled moments. We therefore estimate the system separately in broad industry groups, preserving the within-NAICS2 size ranks, and give each group's estimates equal weight. Each group is normalized at its own Decile 5 before averaging. Six groups have enough observations for this comparison.[26]

Figure D.2 shows that all three gradients also appear in the equal-weight averages. The same directions persist with a 30-firm cell minimum or estimation by individual two-digit industry.

[26] We require 20 firms per cell, 40 balanced quarters below the common 90th-percentile cutoff of economy-wide m_t^2, and finite estimates for all ten deciles. The retained groups cover natural resources and mining; manufacturing; trade, transportation, and utilities; information; finance and real estate; and professional and business services.

Figure D.1: Economy-Wide and Within-Industry Relative Policy Estimates

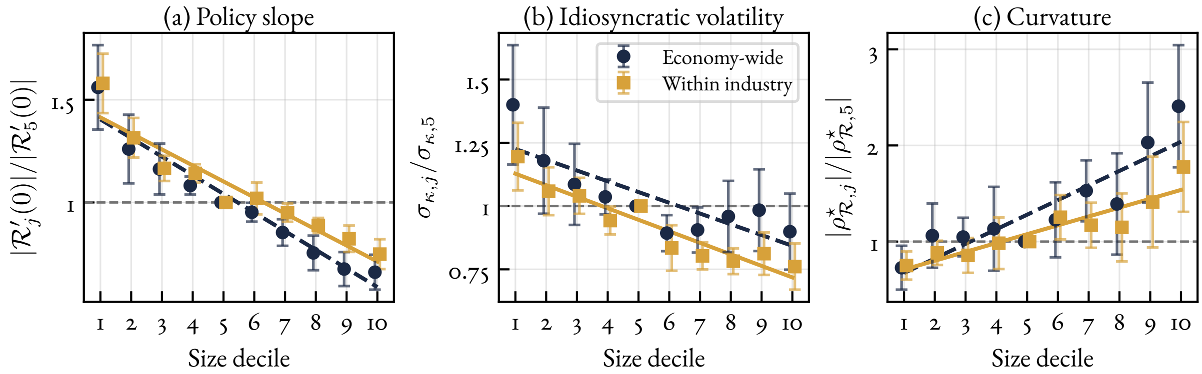

Note: Blue circles use economy-wide size ranks; yellow squares use ranks within two-digit NAICS industries. Estimates are normalized at Decile 5. Bars show 90% delta-method intervals with four calendar-quarter Newey–West lags; lines are descriptive OLS fits.

Figure D.2: Relative Policy Estimates by Industry

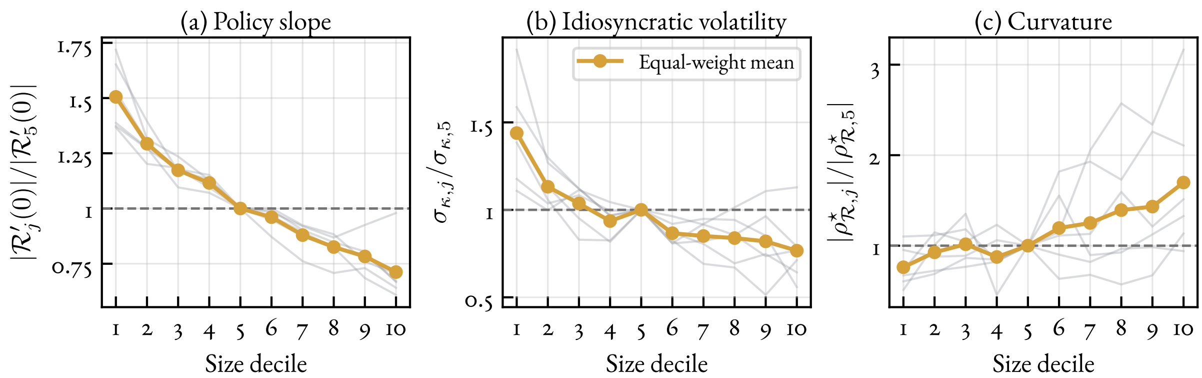

Note: Gray lines show estimates for six broad industry groups; the yellow line is their equal-weight mean. Size ranks are within two-digit NAICS industries, with each group normalized at Decile 5.

E Theory Extensions

E.1 Breaking Symmetry: Finite Public Types

Propositions 2 and 3 extend to finitely many public types. Partition the unit mass of firms into positive-mass sets I_τ, and let τ collect a type's baseline cost and residual-demand environment. Assume demand is type-anonymous—invariant to relabeling firms within I_τ—and impose the symmetry, independence, bounded-support, and moment conditions of Section 3.1 conditional on type. In particular, shocks are centered within type. The homogeneity argument in the proof of Proposition 2 is unchanged, so every type has its own policy but the same exact aggregate scale:

$$q_{i\tau}(u, \epsilon_{i\tau}) = e^{-\bar{\lambda}u/\eta} x_\tau(\epsilon_{i\tau}), \qquad \log r_{i\tau,t} = -\frac{\bar{\lambda}}{\eta} u_t + \mathcal{R}_\tau(\epsilon_{i\tau,t}) + \zeta_{i\tau,t}. \tag{E.1}$$

For the price residual, impose uniformly over the finite type set the type-anonymous analogue of the second-order functional expansion used above: for a receiving type τ, rivals' quantities enter to second order through the first two log-deviation moments of every source type, with a uniform cubic remainder. Writing $\delta_{i\tau,t}$ for own log-quantity deviation and $M_{\sigma t}, S_{\sigma t}$ for those two source moments, use the coefficient convention $\zeta_{i\tau,t} = \sum_\sigma \{a_{\tau\sigma}(\delta_{i\tau,t}) M_{\sigma t} + \frac{1}{2} b_{\tau\sigma}(\delta_{i\tau,t}) S_{\sigma t}\} + O(s_t^3)$. Let $\sigma^2_{\epsilon,\sigma,t} = \sigma^2_{\kappa,\sigma} + u_t^2 \sigma^2_{\lambda,\sigma}$ and $s_t := \max_\sigma \sigma_{\epsilon,\sigma,t}$. Conditional centering and a Taylor expansion of $\mathcal{Q}_\sigma$ imply that the type-σ first and second raw log-quantity moments are, respectively, $\frac{1}{2}\mathcal{Q}''_\sigma(0)\sigma^2_{\epsilon,\sigma,t} + O(s_t^3)$ and $\mathcal{Q}'_\sigma(0)^2 \sigma^2_{\epsilon,\sigma,t} + O(s_t^3)$. Substitution in the functional expansion gives

$$\begin{aligned} \zeta_{i\tau,t} &= Z_{\tau t} + \nu_{i\tau,t}, & \|\nu_{\tau t}\|_{L^3(I_\tau)} &= O(s_t^3), \\ Z_{\tau t} &= \sum_\sigma \chi_{\tau\sigma} \sigma^2_{\epsilon,\sigma,t}, & \chi_{\tau\sigma} &:= \frac{1}{2}\left[a_{\tau\sigma}\mathcal{Q}''_\sigma(0) + b_{\tau\sigma}\mathcal{Q}'_\sigma(0)^2\right], \end{aligned} \tag{E.2}$$

where $a_{\tau\sigma}$ and $b_{\tau\sigma}$ are the corresponding type-moment derivatives at the benchmark. Thus $Z_{\tau t}$ may differ across receiving types, but is common within each type. After within-type centering it affects mean growth only. With $s := \max\{s_t, s_{t-1}\}$,

$$\mathrm{Var}_\tau(g_t) = \mathrm{Var}_\tau(\Delta \mathcal{R}_\tau(\epsilon_t)) + O(s^4), \qquad \mu_{3,\tau}(g_t) = \mu_{3,\tau}(\Delta \mathcal{R}_\tau(\epsilon_t)) + O(s^5).$$

Consequently, the paper's key cross-sectional variance, curvature, and skewness results continue to hold type by type. If types share the same inverse-demand curve and shock process but differ in baseline costs, the local operating-quantity comparison in Propositions 6 and 7 also ranks their policy slopes and curvature near the choke quantity. Pooling types introduces the standard between-type mean terms, so an arbitrary empirical bin requires the group-level residual condition in Equation (3) rather than automatic exact cancellation.