NBER WORKING PAPER SERIES

SKEWED BUSINESS CYCLES

Sergio Salgado
Fatih Guvenen
Nicholas Bloom

Working Paper 26565
http://www.nber.org/papers/w26565

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019, Revised July 2026

This is a major revision of NBER WP #26565. For helpful comments and suggestions we thank seminar participants at the Austin, Bank of Mexico, Central Bank of Chile, CESifo, Chicago, Columbia, Cornell, Dallas Fed, ESWC, HK City, Insper, Michigan, Montreal, New York Fed, Ohio State, Philadelphia Fed, Queen's, Richmond Fed, SED, SITE, Toronto, UCL, UCLA, and Wharton. Part of this research was performed at a Federal Statistical Research Data Center under FSRDC Projects 1694 and 2019. All results based on US Census data have been reviewed to ensure that no confidential information is disclosed. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the US Census or the National Bureau of Economic Research. All results have been reviewed to ensure that no confidential information is disclosed. The replication packet for the empirical results of the paper is available from: https://sergiosalgado.net/home/research/

Skewed Business Cycles
Sergio Salgado, Fatih Guvenen, and Nicholas Bloom
NBER Working Paper No. 26565
December 2019, Revised July 2026
JEL No. E3

ABSTRACT

Using Census and Compustat firm panel data, we show that the skewness of four measures — the growth rate of sales, employment, productivity, and firm value — is robustly procyclical. All four display a fat left tail of large negative outcomes during recessions and a compressed right tail. This holds across US industries and over 60 countries. Using a firm dynamics model with adjustment costs, risk aversion, and financial frictions solved in general equilibrium, we analyze the impact of first-, second-, and third-moment shocks. A second-moment shock generates a sharp downturn followed by a rapid recovery, whereas concentrating the same total variance in the left tail—a negative third-moment shock—produces a persistent recession matching the depth and slow recovery observed in US data. Modeling recessions as a combination of negative first-moment, positive second-moment, and negative third-moment shocks is critical to match both the cross-sectional moments of firm growth and the aggregate response to recessions.

Sergio Salgado
University of Pennsylvania
The Wharton School
ssalgado@wharton.upenn.edu

Fatih Guvenen
University of Toronto
Department of Economics
and Minnesota Economics
Big Data Institute
and also NBER
fatihguvenen@gmail.com

Nicholas Bloom
Stanford University
Department of Economics
and NBER
nbloom@stanford.edu

1 Introduction

This paper studies how the distribution of firm-level shocks changes over the business cycle. Recessions have classically been characterized as a combination of a negative first-moment (mean) shock and/or a positive second-moment (uncertainty) shock. We argue that recessions are a combination of both of these plus a large negative third-moment (skewness) shock.

Empirically, we find that during downturns a subset of firms does extremely poorly, leading to a left tail of large negative outcomes. At the same time, another subset of firms experiences muted growth rates, compressing the right tail of positive outcomes. In this sense, negative skewness captures a combination of elevated "downside risk" and muted "upside surprises." For example, although major disruptions—such as the Great Recession or the COVID-19 pandemic—impact all firms, a subset in certain industries (e.g., airlines, automotive, or movie theaters) fared much worse. Hence, recessions can be viewed as periods of heightened occurrence of *firm-level disasters*. This is often accompanied by a deceleration of growth for firms at the top end, leading to a compression of the right tail of positive outcomes. The opposite patterns happen during expansions, with the left tail of firm growth shrinking and the right tail expanding, indicating an increase in rapid growth potential. Consequently, the skewness of firms' growth rates is procyclical.

Using firm-level panel data from the US Census Bureau, Compustat, Orbis, and Global Compustat data on firms from the US and 65 other countries, we show that the cross-sectional skewness of four firm-level growth measures—growth rates of sales, employment, productivity, and stock value—is robustly procyclical. As an illustration of our main empirical result, Figure 1 displays the distribution of firm sales growth from the US Census' Longitudinal Business Database (LBD). The solid line shows the empirical density of firm output growth from the most recent three recessions: the 2001 slowdown, the 2008 Financial Crisis, and the 2020 COVID recession. The dashed line shows the density for the expansion years around these recessions, the periods 2002–07 and 2010–19. These have both been normalized to a median growth rate of zero. Figure 1 shows that the distribution of firm growth during recessions has a thicker left tail than during expansions, indicating an increase in dispersion that is mostly coming from the left tail.

This asymmetric change in the distribution of firm growth from expansion to recession years can be quantified using the Kelley skewness, a measure of skewness that is robust to outliers.[1] In this case, the Kelley skewness of sales growth is 0.03 in expansion years (a mildly positively skewed distribution) and –0.16 in recession years (a strongly negatively skewed distribution). This value of 0.03 for skewness during expansions indicates that 49% of the overall dispersion is accounted for by firms with output growth below the median, whereas during recessions, a Kelley value of –0.16 indicates that this share increases to about 58%. Similarly, we show that employment growth, productivity growth, and stock value growth also all display procyclical skewness across US firms. This is robust to splits of firm size, age, industry, and region.

Moving to the US industry data level, we find a very similar pattern: the *within-industry* skewness of firm-level sales, employment, productivity, and stock value growth is strongly correlated with industry growth rates. Finally, we examine data across 52 other countries, showing again that firm-level sales, employment, productivity, and stock-value growth skewness is strongly correlated with GDP growth.

Collectively, these results show that the skewness of firm growth is procyclical across all four measures: skewness rises toward positive values in expansions and falls sharply toward more negative values in recessions. Hence, a natural question is whether this asymmetric change in the distribution of idiosyncratic risk has quantitatively important consequences for aggregate dynamics. In particular, does the shape of the cross-sectional distribution, beyond its first two moments, affect the depth and persistence of recessions? Standard uncertainty models generate sharp downturns that reverse quickly once dispersion returns to normal; if the additional risk is instead concentrated in the left tail, the aggregate response may be qualitatively different.

To address this question, in the second half of the paper we build a firm dynamics model in which risk-averse entrepreneurs hire labor subject to convex and non-convex labor adjustment costs and face financial frictions in the form of costly external equity issuance. Since our empirical analysis covers the universe of US establishments — the vast majority of which are privately held — we follow the tradition

[1] The Kelley skewness is the difference between the 90th-to-50th percentile differential (a measure of dispersion in the right tail) and the 50th-to-10th percentile differential (a measure of dispersion in the left tail) divided by the difference between the 90th and 10th percentiles (a measure of the total dispersion of the distribution). For a distribution with a compressed upper half and a dispersed lower half, the Kelley skewness is negative (i.e., a left-skewed distribution).

of heterogeneous-entrepreneur models (Quadrini (2000); Angeletos (2007); Cagetti and De Nardi (2006)) in which owners bear undiversifiable business risk that shapes their production decisions.

The key assumption is that idiosyncratic firm productivity follows a first-order autoregressive process with disaster-mixture innovations: with regime-dependent probability p_s, a firm draws from a left-shifted component that generates large negative productivity realizations—a firm-level disaster. During recessions the disaster probability increases, fattening the left tail and compressing the right tail of the cross-sectional productivity distribution while simultaneously raising its variance. We solve the model in general equilibrium using the Krusell and Smith (1998) algorithm with endogenous wages and estimate aggregate impulse responses to recession onset using Jordà (2005) local projections on simulated time series.

To isolate the role of distribution shape, we construct three recession specifications that progressively add changes to the cross-sectional moments of the idiosyncratic shock distribution. In the Mean-only case, average TFP declines but the dispersion and skewness of idiosyncratic shocks remain at their expansion parameters; in this case, GDP falls by 2.5% one quarter after the onset and recovers gradually. The Mean+Uncertainty case adds a symmetric increase in idiosyncratic variance—matching the total dispersion of sales growth observed during a typical recession—which deepens the initial downturn to nearly 5% through precautionary hiring freezes and widened (S,s) inaction bands. However, the symmetric spread also generates large positive productivity draws that, through the convexity of revenue in productivity—the Oi-Hartman-Abel effect—produce a rapid recovery that returns GDP above its pre-recession level, a pattern inconsistent with the data.

Finally, our baseline case—Mean+Skewness—replaces the symmetric increase in dispersion with the calibrated disaster-mixture process, concentrating the same total variance in the left tail via a higher disaster probability. GDP declines by approximately 5% two quarters after the onset and remains roughly 1.5% below its pre-recession level three years later, closely tracking the empirical impulse response estimated from US data since 1970. Two mechanisms account for this persistence. First, heightened left-tail risk induces risk-averse entrepreneurs to hoard cash at the expense of hiring—the precautionary channel. Second, the compressed right tail eliminates the large positive productivity realizations that, under symmetric uncertainty, generate the Oi-Hartman-Abel effects that drive rapid recoveries. This Mean+Skewness specification mimics the aggregate dynamics and also the micro data. It yields a recession-year increase in cross-sectional dispersion and the sharp negative shift in Kelley skewness of sales, employment, productivity, and equity return growth.

Contribution to the literature. Our paper extends the rare disasters literature to the firm level. A large body of work shows that low-probability, large negative shocks to aggregate consumption or output can account for key asset pricing puzzles and shape business cycle dynamics Rietz (1988); Barro (2006)). Kozlowski *et al.* (2020) show that observing a single tail event permanently shifts agents' beliefs about tail risk, generating persistent stagnation—a mechanism that complements our focus on the direct effects of time-varying disaster probability at the firm level.[2] We document that these low-probability, large negative shocks also occur at the micro level—in individual firms' sales, employment, productivity, and market value growth—and that such left-tail events amplify macroeconomic fluctuations through the interaction of risk aversion, adjustment costs, and financial frictions.

Second, we contribute to a growing literature on procyclical skewness in firm-level and individual-level data. Davis and Haltiwanger (1992) show that job destruction is countercyclical and more volatile than job creation, implying a negatively skewed employment growth distribution in downturns. Kehrig (2011) documents that revenue productivity dispersion rises in recessions, and Guvenen *et al.* (2014) show that recessions affect the skewness—not merely the variance—of individual earnings changes. More recently, Dew-Becker (2024) shows that option-implied firm-level skewness is significantly procyclical, and Decker *et al.* (2016) document a secular decline in high-growth young firms that historically contributed to positive skewness. Jaimovich *et al.* (2023) show that the standard Gaussian specification is at odds with the fat-tailed, non-Gaussian nature of observed firm dynamics, and that higher moments of the shock process have sizable aggregate consequences. On the theoretical side, asymmetric firm responses to good and bad news have a long tradition—from the bad-news principle of Bernanke (1983) to the concave hiring rules generated by ambiguity aversion in Ilut *et al.* (2018).

Finally, our work relates to the literature on asymmetric business cycles. McKay and Reis (2008) document that contractions in employment are briefer and more violent than expansions, and Jensen *et al.* (2020) show that business cycle asymmetry

[2]See also Gourio (2008, 2012), Gabaix (2012), and Kozeniauskas *et al.* (2018).

has deepened over three decades, driven by rising leverage. Using 150 years of data across advanced economies, Jordà *et al.* (2024) find that all business cycles resemble mini-disasters: growth is pervasively fat-tailed with negative skewness. A related literature shows that input-output linkages can generate negatively skewed aggregate fluctuations even from symmetric firm-level shocks, including Acemoglu *et al.* (2012) modeling production networks, Dew-Becker (2023) examining complementarity across inputs, and Miranda-Pinto *et al.* (2023) building off firm-to-firm networks. Our paper provides a distinct and complementary channel: we document that firm-level shocks themselves become more left-skewed in recessions, contributing an additional source of aggregate skewness beyond network amplification.

The rest of the paper is organized as follows. Section 2 describes the data we use and the statistics discussed in the empirical section. Section 3 shows the main empirical results of our paper, that is, that the skewness of several firm-level outcomes and productivity shocks is procyclical. Section 4 describes a simple model to account for the empirical evidence. Section 5 summarizes our results.

2 Data and Measurement

2.1 Data and Sample Selection

Our analysis is based on five large datasets that encompass firm- and plant-level information for the United States and for a combined sample of 52 other countries. The breadth of our datasets allows us to provide a detailed description of the cyclical patterns of the distribution of the growth rate of firm-level outcomes and firm productivity.

First, we use panel data on firm employment and sales from the US Census Bureau's LBD. The LBD provides high-quality measures of employment, industry, and age for the entire US non-farm private sector, linked over time at the establishment level from 1976 to 2022. Precise estimation of the tail percentiles underlying the Kelley skewness measure requires large samples, so having several million observations per year for the aggregate sample, and for sub-sample cuts by age, size, industry, and other observables is critical. Starting in 1998, the LBD also contains information on firm revenues which we also use in our analysis (Haltiwanger *et al.*, 2019).[3] From the

[3]The revenue data at Census is collected at the firm-level and is available from 1998 to 2022. Approximately 30% of firms that have employment information do not have revenue information.

LBD, we calculate cross-sectional moments of the distribution of employment and sales growth within narrow firm population groups. The LBD contains an average of 4.6 million firms per year which, for measuring higher-order moments like skewness, is a major advantage. Our primary unit of analysis is the firm; therefore, we aggregate employment across all the establishments of a particular firm in a given year.

Second, we obtain data on employment, revenues, capacity utilization, and productivity for a panel of manufacturing establishments by combining information from three Census data sets, the Census of Manufacturing (CM), the Annual Survey of Manufactures (ASM), and the Quarterly Survey of Plant Capacity Utilization (QPC), covering years from 1976 to 2021. From the merged dataset, we select establishments with at least ten years of non-missing, non-negative observations on employment and output which, by the ASM methodology, oversamples larger establishments (and thus, implicitly, larger firms). We use this dataset to measure establishment-level (revenue) productivity and corresponding cross-sectional moments of the distribution of productivity shocks.

Third, we draw panel data of publicly traded firms from Compustat. Although this dataset contains mostly large firms, it provides several additional variables that are helpful in our analysis. In particular, we use annual data on sales and employment, and daily stock prices from 1970 to 2024, and we restrict attention to a sample of firms with more than ten years of data to reduce the types of compositional issues identified in Davis *et al.* (2007). Because Compustat reports firms' entire global operations rather than only their US component, it also lets us confirm that firm-growth skewness holds in global sales, employment, and productivity.[4]

Fourth, we study whether the patterns we document for the United States are also observed in other countries, both developed and developing. To that end, we use within-country firm-level panel data containing sales and employment information between the mid 1980s and 2019 from the Osiris dataset collected by Moody's. In order to maintain homogeneous sampling criteria, we consider firms with ten or more years of data and restrict our sample to country-year bins with more than one hundred

A significant fraction of these firms (about 25%) are in the NAICS 81 code (Other Services). The cross-sectional moments for the employment growth distribution do not change if we focus on the sample of firms with employment and revenue data. See Haltiwanger *et al.* (2016) for additional details.

[4]This is important because many large US firms have substantial activity outside the US, engaging in transfer pricing to shift activity (and profits) across borders, which could potentially be cyclical.

firms. Our main results are based on an unbalanced panel of firms spanning 50 countries from 1984 to 2019. We complement this dataset with information on firm-level stock prices obtained from the Global Compustat dataset. Applying similar selection criteria, we obtain a sample of daily stock price information for firms in 51 countries from 1986 to 2024.

Finally, we use firm-level panel data from Amadeus—which is also part of Orbis—containing a smaller number of countries, for a shorter timespan, but with rich firm-level information for small and large firms, both publicly traded and privately held. Crucially, this dataset provides information on sales, employment, intermediate inputs, capital stock, and labor input cost so that we can estimate firm-level (revenue) TFP for a sample of 10 European countries starting in the mid-1990s.[5] All combined, our dataset spans 65 countries and the US with different degrees of coverage and data availability. Additional details on data construction, list of countries, sample selection, and moment calculation for each dataset used in our analysis can be found in Appendix A.

Measuring Dispersion and Skewness. To minimize the impact of outliers, we measure the growth rate of firm-level outcomes, such as employment, output, and revenue per worker, by the arc-percent change between period t and $t+1$, defined as $2 \times (x_{t+1} - x_t) / (x_{t+1} + x_t)$. This measure is bounded between –2, in the case a firm exits the market, and 2, in the case a firm enters the market. Returns are computed as the percentage change of prices within a given period.

To further control for potential outliers, we measure dispersion and skewness using quantile-based measures. As we shall see, they also have magnitudes that are easy to interpret. Our measure of dispersion is the 90th to 10th percentiles differential, denoted by $P9010_t$, where t is a quarter or a year depending on the dataset. Our preferred measure of skewness is the Kelley skewness, which is defined as

$$KSK_t = \underbrace{\frac{P9050_t}{P9010_t}}_{\text{Right Tail Share}} - \underbrace{\frac{P5010_t}{P9010_t}}_{\text{Left Tail Share}} \qquad \in [-1, 1].$$

[5] We use the 2021 vintage of Orbis, accessed through Wharton Research Data Services. In practice our data cover from the early 1990s to 2019 with substantially better coverage starting in 2005. For detailed information on constructing a consistent dataset, see Kalemli-Özcan *et al.* (2024). Table A1 in the Appendix summarizes the data sources and provides basic sample statistics for each of the datasets we use in our analysis. Appendix Table A.1 shows a list of the countries we consider in our analysis and the data available for each of them.

This measure is useful as it provides a simple decomposition of the share of total dispersion that is accounted for by the left and the right tails of a distribution. A negative Kelley skewness indicates that the left tail accounts for more than half of the total dispersion and the distribution is negatively skewed, and vice versa for a positive value.[6]

3 Skewness over the Business Cycle

In this section, we show that the skewness of key measures of firm growth is procyclical—rising in expansions and falling sharply in recessions—in the United States (Section3.1), across firm characteristics (Section 3.2), within industries (Section 3.3), and across several countries (Section 3.4).

3.1 US Macro Evidence

The first contribution of our paper is to document that the skewness of the growth rates of firm-level outcomes varies over the business cycle. Figure 2 has four panels showing (clockwise) the distributions of firm growth of sales, employment, sales-per-employee, and stock market value. To better appreciate the change in the densities of each variable, we plot each density between -1.5 and 1.5 and adjust it to have a median of zero, so the recessions and expansions overlap.[7] It is clear that relative to expansions, recessions display a negative skewness, with a left tail that is thicker. In all cases, skewness is negative in recessions and for sales, employment, and sales-per-employee it is positive in expansions. To get a sense of the significance of the changes in the sales growth distribution, consider a 50% decline in sales. The results in the top-left plot of Figure 2 indicate that such a large decline is 2.6 times more likely during recessions than in expansions, without a significant change in the likelihood of a positive increase of the same magnitude, as one would predict after a *symmetric* increase in the dispersion of sales growth. Simple time-series analysis confirms this

[6] Notice that the Kelley skewness is invariant to 20 percent of the observations in the sample (the top and bottom 10 percent of the distribution are not considered). In principle, the Kelley skewness can be computed using any two symmetric percentiles, such as the 95th and 5th or 97.5th and 2.5th percentiles. Results are similar using these choices and to using the third standardized moment.

[7] Recall that the arc-percent change assigns a value of -2 for firms exiting the economy and a 2 for firms entering. These observations are not shown in Figure 1 and 2 but have a significant effect on the Kelley skewness. The values reported in each plot, however, include these observations. To better visualize the change in the tails of the distribution, Figure A1 in the appendix shows the log-density plot for each variable for the full range.

cyclical skewness is statistically significant. Table 1 reports regressions where the left-hand side is the skewness of the growth rate of our four measures against the growth rate of GDP. As expected, the skewness of all variables strongly comoves with the business cycle.

The first three columns of Table 1 run yearly regressions on Census LBD data. For example, column 1 regresses the cross-sectional skewness of firm sales growth against annual GDP growth. There are an average of 4.6 million firm-level observations per year used to measure skewness, and the regression uses 24 years of data from 1998 to 2021. The coefficient of 0.068 indicates that for a 1 standard deviation decrease in GDP growth, the Kelley skewness in the cross-sectional distribution of firm sales growth declines by 0.068. As a benchmark, the difference in Kelley skewness between the last three recessions (2001, 2008/9 and 2020) and expansions in Figure 1 is 0.19. Columns 2 and 3 run similar regressions for employment and sales-per-employee growth, reporting similar results. Finally, column 4 regresses the cross-sectional skewness of quarterly stock-returns (value growth) from CRSP data against quarterly GDP growth, spanning on average 4,000 firms, again finding a highly significant positive coefficient.

3.2 Firm Size, Age, Industry, and Region

It is well known that small and young firms account for an oversized fraction of employment growth in the economy (Haltiwanger *et al.*, 2016) and are the most cyclical (Fort *et al.*, 2013). If that is the case, one would expect these firms to show a stronger cyclicality pattern in skewness relative to large, well-established firms. Similarly, some sectors might be more responsive than others to economic fluctuations (e.g., construction versus healthcare), perhaps driving the aggregate changes in the skewness of firms' growth rates that we find. To answer these questions, we start by exploiting the large LBD sample size to divide firms into narrow population groups. Perhaps surprisingly, we find that the distribution of firm growth shows similarly cyclical patterns, irrespective of the size, age, industry, or other characteristics of the firm.

To see this, the top-left panel of Figure 3 shows that the skewness of employment growth is procyclical within firm-size groups.[8] Quantitatively, the skewness of

[8]Notice that since firms' growth is calculated between periods t and $t+1$ we avoid contemporaneous correlation between firm growth and size, by classifying firms by lagged average employment.

employment growth declines almost the same during recessions, whether we focus on small and medium-sized firms of fewer than 50 or if we look at large firms with 1,000 workers or more. We also find that skewness is similarly procyclical between young and old firms, as shown in the top right panel. In the bottom two panels we split by major industry and by region, again finding strongly procyclical skewness of firm-level employment growth. Appendix Figure A2 shows similar patterns for sales growth although for a shorter sample period.

3.3 US Industry Evidence

We now examine within-industry business cycles and skewness. This analysis provides more variation in our core US data by exploiting differences in growth and skewness across industries. To do this, we take our US Census data and divide it into about 280, 4-digit NAICS sectors, and within each sector generate the annual Kelley skewness of firm growth rates and the average firm-level sales growth rates. This analysis is possible because we have on average 4.6 million firm observations per year in the LBD, yielding an average of 16,000 firms per industry-year.

The results are plotted in the top row of Figure 5 that shows the relation between the skewness of firm growth and the average growth within the industry, controlling for industry and year effects.[9] We see a strong upward slope: industries with higher-than-average mean output growth rates experience significantly more positively skewed sales and employment growth rates. In terms of magnitudes, the top left panel of Figure 4 shows that when the average industry sales growth is –0.08, the Kelley skewness is around –0.1, indicating that 55% of the total dispersion in sales growth within an industry is accounted for by the left tail of the distribution. In contrast, when the average sales growth is 0.08, Kelley skewness is 0.10, so now it is the right tail that accounts for about 55% of the total dispersion. The effects for employment growth are even stronger, as shown in the right panel.

Sales and employment growth, however, depend on the shocks that firms experience and their responses to these shocks. Industry-level information allows us to look more closely at variations in productivity. In particular, we use data for a sample of manufacturing establishments in the United States that combines records from the

[9] Although Figure 5 and Table 2 absorb year fixed effects, they combine between- and within-industry variation. Figures A3 and A4 provide evidence that the growth-skewness relationship holds even within industries after purging all common time-series variation, confirming that the pattern is not driven entirely by macro-factors like monetary or fiscal conditions.

CM and the ASM. The Census Bureau also provides a readily available measure of establishment-level (revenue) productivity, which we use in our analysis (Foster *et al.*, 2016). As the bottom-left panel of Figure 4 shows, the skewness of firms' productivity growth is negative in industries experiencing average sales declines.[10]

Finally, we examine data on firm value growth. The bottom right panel plots the skewness of firm-level stock returns by NAICS 2-digit industry-year for a sample of publicly traded firms from Compustat, against their average industry growth rates.[11] We see that industries with positive industry growth rates tend to have a positively skewed value growth distribution, whereas the opposite happens when industry growth is negative.

These industry-level results are confirmed in Table 2, which displays a series of industry panel regressions in which the dependent variable is the Kelley skewness of the growth of sales, employment, revenue-per-worker, and stock returns within an industry-year cell, and the explanatory variable is industry-level output growth rates.[12] The results in this case are based on a sample of publicly traded firms from Compustat, which allows us to cover a longer time span. We also include TFP growth measures from the Census ASM. In all cases, we find that periods of high output growth are periods in which the skewness of firm growth is significantly positive. The bottom rows of Table 2 report the average skewness of the dependent variables in industry expansions and contractions, showing this is negative in contractions for all variables, and positive in expansions for sales, employment, sales/employee and total factor productivity growth.

3.4 Cross-Country Evidence

We then move to investigate firm skewness for a combined sample of 65 countries that are geographically and economically diverse, spreading over five continents including

[10]Standard measures of productivity might be affected by changes in capacity utilization (Basu, 1996). To control for possible changes we regress firm productivity on the number of hours worked in the firm and electricity use—common capacity measures—and use the residual as our firm-level productivity cleansed from capacity effects. These results are shown in Appendix A.4.

[11]The returns panel uses 2-digit rather than 4-digit NAICS cells because the publicly traded sample in Compustat is too sparse at the 4-digit level to yield reliable within-cell skewness estimates.

[12]Figure 5 and Table 2 absorb year fixed effects, thereby purging the common aggregate cycle, while combining the remaining between- and within-industry variation in skewness. Because the four panels in Figure 5 rest on different industry partitions---4-digit NAICS for sales and employment, manufacturing establishments for productivity, and 2-digit NAICS for returns---the slope magnitudes are not directly comparable across panels. Figures A3 and A4 show that the *within industry* skewness of firm growth is also procyclical.

high-income countries (such as Germany, Japan, or France) and middle-income countries (such as Peru, Egypt, or Thailand).

Figure 5 shows binscatters for four different country-by-year firm-level skewness measures against country GDP growth. Data are plotted after removing country and year fixed effects, which purges out any systematic differences in country growth rates or common global shocks like the 2008 Financial Crisis. The top panels show the Kelley skewness of the firm sales and employment growth distribution against country GDP growth rates, revealing a strongly positive relationship similar to the one we find across industries. In the bottom left panel we examine the skewness of TFP shocks, which again is procyclical across different countries. In this case, we use firm-level data for a sample of European countries for which we have rich enough data on output, employment, capital and intermediate inputs data to measure firm-level (revenue) TFP.[13]

Finally, in the bottom-right panel we complete our analysis by looking at the variation of the skewness of annual returns across different countries using a sample of publicly traded firms from Global Compustat. Although the coverage of the stock market varies significantly across countries, it is clear that in periods where the economy is growing fast, the distribution of stock returns is more positively skewed, whereas periods of economic downturns are also periods in which returns skew negative.

Table 3 repeats the cyclicality regression discussed above for the United States but this time exploiting the panel dimension of the cross-country dataset to assess the cyclicality of skewness in international data. The dependent variable is the skewness of employment growth, sales growth, productivity growth, or value growth (stock returns) within a given country-year cell. The business cycle is captured by the log GDP per capita growth in the respective country, which we have rescaled to have a unit variance to facilitate the comparison with the rest of the results. The regressions also include a full set of time and country fixed effects to control for common economic conditions that might affect all countries simultaneously or for fixed differences across countries. The results are consistent with our time-series

[13]Our firm-level data from Amadeus include information on small and large firms, both publicly traded and privately held, from seventeen European countries. For a smaller sample of these countries, namely, Germany, Spain, Finland, France, Italy, Norway, Poland, Portugal, Sweden, and Ukraine, we have enough information to estimate firm-level TFP. Appendix A.4 describes in full detail the sample selection and estimation procedure.

findings of procyclical skewness for all four measures with similar levels of statistical significance.[14]

4 The Quantitative Impact of Skewness Shocks

In this section, we develop a heterogeneous firms general equilibrium model to quantify the importance of skewness shocks for business cycle fluctuations and compare them with standard mean and uncertainty shocks. Despite its parsimony, the model generates a meaningful contribution of skewness shocks to aggregate fluctuations, which is distinct from that of a mean or an uncertainty shock.

4.1 The Model

In our model economy there is a unit mass of risk-averse entrepreneurs, each operating a decreasing-returns-to-scale technology, and a representative household that consumes and supplies elastic labor. We follow a tradition of heterogeneous-entrepreneur models in which undiversifiable idiosyncratic risk directly affects production decisions (Cagetti and De Nardi (2006); Angeletos (2007); Boar *et al.* (2025); Quadrini (2000)). As Greenwald and Stiglitz (1993) establish, firms facing costly external finance behave as if risk-averse in their input choices even absent an entrepreneur structure—a property our model inherits through the combination of equity issuance costs and a cash-carrying wedge.

Time is discrete at quarterly frequency. The aggregate state, described below, follows an exogenous two-state Markov chain $s_t \in \{E, R\}$ (expansion, recession) that governs the distribution of firm-level productivity shocks. Output of firm i in state s is

$$y_i = A_s \, \theta_i \, z_i \, l_i^{\alpha}, \tag{1}$$

where A_s is aggregate productivity ($A_E \geq A_R$), θ_i is a permanent revenue scaling factor, z_i is firm-specific productivity, and l_i are workers.[15] The firm workforce attrition

[14]In a robustness exercise, we drop 2008 and 2009 from the sample to assess whether the results are driven by the Great Recession. The coefficient on GDP growth for sales skewness changes from 0.034 to 0.036, and for employment skewness from 0.057 to 0.058, both remaining significant at the 1% level.

[15]We keep the model tractable by abstracting from capital, although this is unlikely to change our main conclusions for two reasons. First, Bloom (2009) shows that labor and capital adjustment costs generate qualitatively similar dynamics in response to uncertainty shocks. Second, including capital adjustment costs would likely amplify the skewness channel, as investment provides an additional margin to hiring through which skewness shocks affect firm decisions.

rate is δ_L and new hires produce next period. Hence, the law of motion of labor is

$$l_i' = (1-\delta_L)l_i + h_i, \tag{2}$$

where l_i' is the labor input in the next period and $h_i = l_i' - (1-\delta_L)\, l_i$ is net hiring. Each firm draws a permanent type $\theta_i \sim \log\mathcal{N}(0, \sigma_\theta^2)$ at entry that generates persistent size differences: high-θ firms operate at larger scale at every productivity level. Because θ_i is constant over time, it cancels in growth rates and affects only the cross-sectional firm size distribution.

Firms are owned by risk-averse entrepreneurs with CRRA preferences $u(c) = (c^{1-\gamma}-1)/(1-\gamma)$ and discount the future at rate β. The entrepreneur model is relevant for our data: the LBD covers the universe of US private and public establishments, for which two-thirds of employment is in privately held firms (Davis *et al.*, 2007), and Moskowitz and Vissing-Jørgensen (2002) document that approximately 75% of this private equity is held by households for whom it constitutes at least half of total net worth.[16] Adjusting employment incurs three types of costs:

$$\text{AC}(h_i) = \psi_H\, w\, (h_i)^+ + \psi_F\, w\, (-h_i)^+ + \psi_S\, y_i\, \mathbb{I}_{h_i \neq 0}, \tag{3}$$

where $(x)^+ = \max(x, 0)$. The first two terms are linear per-worker hiring and firing costs proportional to the wage, and the third is a fixed disruption cost proportional to revenue that generates S-s inaction bands.

Financial frictions. Entrepreneurs can hold cash $n \geq 0$ that earns a below-market return $R_n = \kappa/\beta$ with $\kappa < 1$, creating a carrying cost for liquidity. When the firm's net cash flow is negative, the entrepreneur can raise external equity at cost given by

$$\Phi(d_i) = \eta_1\, |d_i| + \eta_2 \left(\frac{d_i}{y_i}\right)^2 \cdot y_i, \qquad d_i < 0, \tag{4}$$

following Hennessy and Whited (2007), combining a proportional component with a scale-invariant quadratic that penalizes large deficit-to-sales ratios disproportionately.

[16] For the other one third of employment in public firms Panousi and Papanikolaou (2012) show that even in these publicly traded firms, investment declines with idiosyncratic risk, and more so when managers hold a larger equity stake—indicating that managerial risk aversion affects firm-level decisions well beyond the private-firm sector.

Idiosyncratic productivity. Recent work documents that the standard Gaussian specification for idiosyncratic firm-level productivity is at odds with the data (Jaimovich *et al.* (2023) and Melcangi and Sarpietro (2024)). At the macro level, Gourio (2012) models time-varying aggregate disaster risk and shows that fluctuations in the probability of rare disasters can generate realistic business cycle and asset pricing dynamics. We follow this approach and assume that the firm-level log productivity follows an AR(1) with disaster-shock innovations:

$$\log z_i^{'} = \rho_z \log z_i + \eta_i^{'}, \qquad \eta \sim (1 - p_s) \cdot \mathcal{N}(\mu_{\text{base},s}, \sigma_s^2) + p_s \cdot \mathcal{N}(\mu_{\text{base},s} + \mu_D, \sigma_s^2), \quad (5)$$

where p_s is the disaster probability (regime-dependent), $\mu_D < 0$ is the disaster mean shift (regime-invariant), and σ_s is the common-component standard deviation which is regime-dependent: σ_E in expansions, σ_R in recessions. The base mean $\mu_{\text{base},s}$ is also regime specific and chosen so that $\mathbb{E}[\exp(\eta) \mid s] = 1$ in each regime, ensuring that the innovation process does not shift mean productivity in levels. The mixture structure separates two sources of idiosyncratic risk: the baseline component $\mathcal{N}(\mu_{\text{base},s}, \sigma_s^2)$, which governs the bulk of the distribution, and the disaster component, which with probability p_s shifts the innovation mean by $\mu_D < 0$, generating left-tail events. Since p_s and σ_s are regime-dependent, the shape of the cross-sectional productivity distribution varies over the business cycle.

The Entrepreneurs' Problem. The state of firm i is (s, z, l, n): aggregate regime, idiosyncratic productivity, employment, and cash. The firm chooses next-period employment l' and cash holdings n' to maximize

$$V(s, z, l, n) = \max_{l', n' \geq 0} \left\{ u(c) + \beta \, \mathbb{E}\left[V(s', z', l', n') \,\middle|\, s, z \right] \right\}, \quad (6)$$

subject to $c > 0$. The firm's payout is

$$d_i = y_i - w\, l_i - \text{AC}(h_i) + R_n\, n_i - n_i', \quad (7)$$

where w is the equilibrium wage, taken as given by the firm. Entrepreneur consumption equals the firm payout $c = d_i$ subject to $c > 0$. Since the entrepreneur has CRRA utility, the firm endogenously maintains strictly positive payouts in equilibrium. In practice, the equity-issuance cost function $\Phi(d_i)$ in Equation 4 applies in the deficit region $d_i < 0$: the value function in that region reflects flow utility $u(d_i - \Phi(d_i))$, which is strictly worse than $u(d_i)$ and therefore further steepens the precautionary incentive near $d_i = 0$. In equilibrium, firms adjust employment and cash holdings to avoid deficits entirely, creating a strong precautionary motive for holding liquid assets.

Households and labor market. The supply side of the labor market is populated by a representative household with GHH preferences (Greenwood *et al.*, 1988):

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(C_t - \frac{\theta}{1 + 1/\varepsilon} N_t^{1+1/\varepsilon} \right), \tag{8}$$

where C_t is aggregate consumption, N_t is aggregate hours, and $\varepsilon > 0$ is the Frisch elasticity of labor supply. The intratemporal optimality condition yields an inverse labor supply schedule $w_t = \theta \, N_t^{1/\varepsilon}$, which, after normalizing around the steady state ($w_{ss} = 1$, $N_{ss} = L_{ss}$), gives the labor market clearing condition

$$w_t = \left(\frac{L_t}{L_{ss}} \right)^{1/\varepsilon}, \tag{9}$$

where $L_t = \sum_i l_{i,t}$ is predetermined aggregate employment. Under GHH preferences, labor supply depends only on the real wage and not on consumption, eliminating wealth effects. This ensures that wages respond to aggregate employment—a slow-moving, predetermined variable—rather than to contemporaneous output, which includes the immediate aggregate productivity jump A_s on recession entry.

In general equilibrium, the wage w_t enters the firm's problem as an additional aggregate state variable. Firms form rational expectations about future wages using a regime-dependent forecasting rule that depends on the current wage and aggregate output. Similarly to Khan and Thomas (2008) and Bloom *et al.* (2018), we solve for the equilibrium using the algorithm of Krusell and Smith (1998). Additional details are provided in Appendix B.[17]

[17] The general equilibrium closes the labor market but treats the interest rate as exogenous: $R_n = \kappa/\beta$ is a fixed parameter calibrated to match the median cash-to-sales ratio. The parameter κ captures structural frictions in the firm's financial environment—agency costs and tax disadvantages of holding liquid assets—rather than a market-clearing rate. Because all three recession specifications face the same R_n, any endogenous adjustment to the safe rate would operate symmetrically across specifications and would not affect the relative comparison between the uncertainty and skewness channels. We also verify that our results are stable when varying κ around its calibrated value

4.2 Calibration

Table 4 lists all model parameters and the calibration targets. Production and preference parameters follow standard values in the literature. We set risk aversion to $\gamma = 2$ and an annual discount factor of 0.95. We set $\varepsilon = 5$, which roughly implies that a 1% change in aggregate employment moves wages by 0.2%. This ensures that the bulk of cyclical adjustment occurs on the employment margin, consistent with the well-documented sluggishness of real wages over the business cycle.

The aggregate regime transition is governed by p_{ee} and p_{rr}. We set $p_{ee} = 0.97$ and $p_{rr} = 0.95$, which captures the fact that after the onset of a recession, GDP, employment, and the skewness of sales growth remain below pre-recession levels for several quarters, as we show in Figures 8 and 9 below.

The six shock and financial parameters—σ_E, σ_R, μ_D, p_E, p_R, and κ—are calibrated jointly via simulated method of moments. At each candidate parameter vector, the procedure solves the value function, simulates a long panel of firms, classifies each calendar year as expansion or recession using the approximate NBER convention (≥ 2 recession quarters), and computes cross-sectional moments of annual firm-level sales growth separately by regime.[18] The four spread moments (P90–P50 and P50–P10 in expansion and recession) are targeted from the LBD, and the median cash-to-sales ratio is targeted from Compustat. The permanent heterogeneity parameter σ_θ is calibrated separately to match the top-10% employment share in the LBD. Finally, simulated sales and employment observables are perturbed by multiplicative measurement error $\exp(\varepsilon_{me})$ with $\varepsilon_{me} \sim \mathcal{N}(0, \sigma_{me}^2)$ and $\sigma_{me} = 0.05$ before computing cross-sectional moments. We calibrate this noise level by combining information from the Census ASM and LBD.[19] The measurement error does not affect the simulation dynamics — only the moment computation used for calibration.

(Appendix Figure A6).

[18]This classification follows the NBER practice of identifying recessions by sustained contraction across multiple indicators. In the model, a year with two or more quarters in the recession regime ($s = R$) corresponds to a year in which the economy spends at least half its time in the contractionary state. A similar convention is used when computing cross-sectional moments from the data (Section 3).

[19]We compute the cross-dataset correlation for log revenue and for log employment by matching a sample of single-establishment manufacturing firms in the LBD to their corresponding records in the ASM. We restrict to single-establishment firms because the LBD only reports revenue at the firm level. The correlation across the two datasets is 0.95 for both variables. Given this, we set a log standard deviation of $\sigma_{me} = 0.05$.

Identification. The regime-dependent standard deviation σ_s pins down the upper spread of sales growth (P90–P50), which reflects the right tail of the innovation distribution: σ_E targets P90–P50 in expansions and σ_R targets P90–P50 in recessions. The disaster parameters (μ_D, p_E, p_R) jointly control the lower spread (P50–P10): higher p_s widens the left tail without affecting the right tail. Three parameters target two moments (P50–P10 in expansion and recession). The parameter μ_D is restricted to be regime-invariant—the severity of firm-level disaster events is a structural feature of the economy, while only the frequency p_s varies across regimes. This cross-equation restriction disciplines the estimation: the *ratio* of P50–P10 across regimes is informative about p_R/p_E, while the *levels* pin down the product $p_s \cdot |\mu_D|$, and hence μ_D itself. The cash wedge κ is separately identified by the median cash-to-sales ratio through the carrying cost $1 - \kappa/\beta$. The remaining parameter, σ_θ, is calibrated to match the share of employment at the top 10% of firms.

Model Fit. Figure 6 compares the model's unconditional cross-sectional moments of annual sales growth with the LBD targets in expansions and recessions. We consider three specifications: a Mean-Only case with a 1.5% aggregate TFP decline, a Mean+Uncertainty case that adds a symmetric variance increase, and our baseline Mean+Skewness case that instead loads the additional dispersion onto the left tail. The baseline closely matches all four targeted spreads: P90–P50 is 0.21 in both expansions and recessions, while P50–P10 rises from 0.22 to 0.37, exactly as in the data. These spreads imply a model Kelley skewness of +0.03 in expansions and −0.21 in recessions. The expansion value matches the +0.03 reported in Figure 1 from the LBD, confirming that the baseline specification reproduces the positive skewness observed during normal times. By contrast, the Mean-Only specification leaves the distribution essentially unchanged, and the Mean+Uncertainty specification raises total dispersion but does so too symmetrically: it widens the right tail excessively and the left tail too little.

Figure 7 shows the dynamics of cross-sectional moments around recession episodes, plotting deviations from the pre-recession average. Four patterns are worth noting. First, the response of employment growth mirrors that of sales growth across all specifications, consistent with the tight link between output and labor demand under decreasing returns. Second, equity returns exhibit a sharp drop at recession onset followed by a rapid recovery, reflecting their forward-looking nature — the value function capitalizes the expected duration of the recession state, so the decline is concentrated in the transition period. Third, Kelley skewness of sales and employment growth—the third column—turns sharply negative at recession onset under the Mean+Skewness, while it remains near zero under other specifications, including the Mean+Uncertainty case. The model-implied Kelley skewness of employment growth displays the same sign pattern: $+0.07$ in expansions and -0.13 in recessions, compared with $+0.07$ and -0.22 in the LBD (Figure A1).

Finally, the second and third moments of productivity and stock value growth closely mimic the nature of the underlying driving process. This reflects the "smooth pasting" optimality conditions, which imply that even in models with strong non-linearities (like in models with fixed costs of adjustment), value growth is continuous in the underlying driving process (e.g., Dixit (1993)). Similarly, labor productivity growth has a close correspondence to the underlying driving process. Hence, using the variance and skewness of firm equity returns and productivity is helpful for identifying the higher-order moments of the underlying driving process over the cycle. Indeed, while different combinations of linear, fixed or quadratic costs may impact the employment and sales cross-sectional second and third moments, they have no material impact on the productivity or value-growth moments (Appendix Figure A5). Collectively, the baseline Mean+Skewness specification reproduces the central empirical regularity documented in Section 3: a procyclical Kelley skewness that is mildly positive in expansions and sharply negative in recessions, for both sales and employment growth.

4.3 Results

In this section, we evaluate the aggregate consequences of left-tail idiosyncratic risk by progressively isolating the role of a skewness shock. Our main finding is that the skewness channel accounts for a substantial and persistent fraction of the decline in GDP during recession. We quantify the aggregate impulse response to recession onset using Jordà (2005) Local Projections (LP) on simulated aggregate time series and compare the model responses to empirical LP estimates from US FRED real GDP data for the 1970–present sample. We solve the model in general equilibrium with endogenous wages using the Krusell and Smith (1998) algorithm described in Appendix B. Firms form rational expectations about future wages via a regime-dependent log-linear forecasting rule. The algorithm iterates between solving the firm's value function at $M = 13$ wage grid points, simulating the economy with market-clearing wages from (9), and updating the forecasting coefficients until convergence. Under the GHH preferences with $\varepsilon = 5$, the converged forecasting rule achieves an R^2 of 0.95, with equilibrium wages deviating by less than $\pm 5\%$ from their steady-state level and zero boundary hits on the wage grid.[20] Further computational details are provided in Appendix B.5.

Decomposition Design. Following our estimation procedure, we consider three specifications that decompose the full recession into its components. The Mean-Only specification reduces aggregate productivity to A_R while keeping the idiosyncratic shock distribution at its expansion parameters ($p = p_E$, $\sigma = \sigma_E$). The Mean+Uncertainty specification adds a mean-preserving increase in innovation variance: the disaster probability stays at p_E while σ increases to $\sigma_{\text{adj}} = \sqrt{\text{Var}_{\text{rec}} - p_E(1 - p_E)\mu_D^2}$, where $\text{Var}_{\text{rec}} = \sigma_R^2 + p_R(1 - p_R)\mu_D^2$, so that total innovation variance matches the full recession. The Mean+Skewness specification—our baseline—applies the full calibrated recession shock ($p = p_R$, $\sigma = \sigma_R$), with the same total variance as the Mean+Uncertainty case but concentrated in the left tail via a higher disaster probability.

4.4 The Impact of a Skewness Shock

Before showing our key results, Figure 8 plots the impact of pure first, second, and third moment shocks. For comparison we also include the empirical path of GDP during a recession.[21]

The first moment (Mean-Only) specification—where TFP is dropping by 1.5%, which is typical for a US recession (Fernald, 2014)—generates a GDP decline of 2.5% one quarter after the shock with a gradual recovery. In comparison, a second moment shock (Uncertainty Only) generates a sharp decline of 3%, arising from two channels. The first is precautionary behavior: as uncertainty rises, risk-averse entrepreneurs reduce hiring and hoard cash to avoid raising costly external finance in the event of large negative shocks. The second operates through the (S,s) inaction bands created by fixed adjustment costs: higher uncertainty raises the option value of waiting, widening the range of productivity realizations over which firms choose

[20] In our baseline results we consider a Frisch elasticity of $\varepsilon = 5$ which is consistent with the well-documented acyclicality of US real wages. Lower values (e.g., $\varepsilon \in \{1, 3\}$) would generate counterfactually large wage swings over the business cycle.

[21] Further details on the empirical local projections are provided in Appendix C.

not to adjust employment. This initial drop is followed by a rapid overshoot which reflects the Oi–Hartman–Abel mechanism: revenue is convex in productivity (see Oi (1961), Hartman (1972), and Abel (1983)), so that cross-firm dispersion in productivity increases aggregate output. The third-moment (Skewness Only) shock generates a similar drop in activity, again due precautionary behavior and real options effects. But crucially, because the right tail is compressed, the drop is not followed by a large overshoot because of the lack of an Oi-Hartman-Abel effect. As a result, a Skewness Only shock generates a 2% decline in GDP with almost no overshoot. None of these three shocks in isolation, however, can generate the large decline in GDP observed during a typical recession shown by the black dotted line in Figure 8: a 5% decline in GDP one quarter after the onset of the recession.

Figure 9 shows that a Mean+Skewness recession—one that matches the cross-sectional firm growth moments we document—generates a decline of comparable magnitude to actual recessions. To build intuition, consider first a pure decline in aggregate TFP of 1.5% (blue solid line in Figure 8). This does not generate a drop in GDP large enough to match the empirical data in our model. Adding a skewness shock increases the size of the drop as firms pause hiring because of real options, financial frictions, and risk aversion effects. In contrast, the Mean+Uncertainty shock generates a similar drop on impact but recovers too rapidly toward its pre-recession level as the Oi-Hartman-Abel overshoot takes hold. Hence, the recession is too short-lived and generates a counterfactual overshoot. In contrast, a Mean+Skewness shock both fits the microdata and also the aggregate path of output.

To quantify the contribution of each amplification channel, we re-solve the general equilibrium Mean+Skewness economy shutting off one mechanism at a time and re-estimate the local projections on the simulated data (Appendix Figure A7). Removing risk aversion produces the largest reduction in the GDP response: the trough shrinks by roughly half, and the economy recovers within eight quarters. Eliminating adjustment costs reduces the GDP trough by about one-third, consistent with the role of (S,s) inaction bands in delaying employment adjustment. Finally, shutting off financial frictions has the smallest effect, reducing the trough by roughly one-quarter.

Finally, we examine the response of aggregate cash holdings, which shed light on the precautionary mechanism (Appendix Figure A8). Under the Mean-Only specification, firms draw down reserves to smooth through the downturn: aggregate cash falls by roughly 6% four quarters after the onset. When we add uncertainty or skewness shocks this pattern reverses, as firms increase cash holdings for precautionary reasons, matching the countercyclical cash holdings seen in US firms (Alfaro *et al.*, 2024).

5 Conclusions

This paper documents that the cross-sectional distribution of firm-level growth rates—in sales, employment, productivity, and firm value—becomes systematically more negatively skewed during recessions. This pattern is pervasive: it appears across firms of different sizes and ages, across US industries, and across 65 countries. We show that this fact has quantitatively important implications for business cycles. In a heterogeneous-firm model, concentrating idiosyncratic risk in the left tail, rather than increasing dispersion symmetrically, turns a transitory downturn into a persistent recession with no recovery overshoot. Risk aversion and adjustment costs are the key amplification mechanisms. Our results add empirical alignment to business cycle models by identifying a powerful additional recessionary force. A long literature has sought mechanisms that can drive, amplify, and propagate recessions, in part because observed first- and second-moment shocks often appear too small to explain the depth and persistence of downturns. A recessionary decline in the skewness of idiosyncratic firm shocks provides such a force while also matching firm micro-data.

References

ABEL, A. B. (1983). Optimal investment under uncertainty. *The American Economic Review*, **73** (1), 228–233.

ACEMOGLU, D., CARVALHO, V. M., OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** (5), 1977–2016.

ALFARO, I., BLOOM, N. and LIN, X. (2024). The finance uncertainty multiplier. *Journal of Political Economy*, **132** (2), 385–439.

ANGELETOS, G.-M. (2007). Uninsured idiosyncratic investment risk and aggregate saving. *Review of Economic Dynamics*, **10** (1), 1–30.

BARRO, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, **121** (3), 823–866.

BASU, S. (1996). Procyclical productivity: increasing returns or cyclical utilization? *The Quarterly Journal of Economics*, **111** (3), 719–751.

BERNANKE, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, **98** (1), 85–106.

BLOOM, N. (2009). The impact of uncertainty shocks. *Econometrica*, **77** (3), 623–685.

—, FLOETOTTO, M., JAIMOVICH, N., SAPORTA-EKSTEN, I. and TERRY, S. J. (2018). Really uncertain business cycles. *Econometrica*, **86** (3), 1031–1065.

BOAR, C., GOREA, D. and MIDRIGAN, V. (2025). The aggregate costs of uninsurable business risk. *American Economic Review*, forthcoming.

CAGETTI, M. and DE NARDI, M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy*, **114** (5), 835–870.

DAVIS, S. J. and HALTIWANGER, J. (1992). Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics*, **107** (3), 819–863.

—, —, JARMIN, R. and MIRANDA, J. (2007). Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. In D. Acemoglu, K. Rogoff and M. Woodford (eds.), *NBER Macroeconomics Annual 2006*, vol. 21, MIT Press, pp. 107–180.

DECKER, R., HALTIWANGER, J., JARMIN, R. and MIRANDA, J. (2016). Where has all the skewness gone? the decline in high-growth (young) firms in the u.s. *European Economic Review*, **86**, 4–23.

DEW-BECKER, I. (2023). Tail risk in production networks. *Econometrica*, **91** (4), 1315–1360.

— (2024). Real-time forward-looking skewness over the business cycle. *Review of Economic Dynamics*, **54**.

DIXIT, A. (1993). Choosing among alternative discrete investment projects under uncertainty. *Economics letters*, **41** (3), 265–268.

Fernald, J. G. (2014). *A Quarterly, Utilization-Adjusted Series on Total Factor Productivity*. Working Paper 2012-19, Federal Reserve Bank of San Francisco.

Fort, T. C., Haltiwanger, J., Jarmin, R. S. and Miranda, J. (2013). How firms respond to business cycles: The role of firm age and firm size. *IMF Economic Review*, **61** (3), 520–559.

Foster, L., Grim, C., Haltiwanger, J. and Wolf, Z. (2016). Firm-level dispersion in productivity: is the devil in the details? *American Economic Review*, **106** (5), 95–98.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, **127** (2), 645–700.

Gourio, F. (2008). Disasters and recoveries. *American Economic Review*, **98** (2), 68–73.

— (2012). Disaster risk and business cycles. *American Economic Review*, **102** (6), 2734–2766.

Greenwald, B. and Stiglitz, J. E. (1993). Financial market imperfections and business cycles. *Quarterly Journal of Economics*, **108** (1), 77–114.

Greenwood, J., Hercowitz, Z. and Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. *American Economic Review*, **78** (3), 402–417.

Guvenen, F., Ozkan, S. and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, **122** (2), 612–660.

Haltiwanger, J., Jarmin, R., Kulick, R., Miranda, J., Penciakova, V., Tello-Trillo, C. *et al.* (2019). *Firm-level Revenue Dataset*. Tech. rep., Center for Economic Studies, US Census Bureau.

—, Jarmin, R. S., Kulick, R. and Miranda, J. (2016). High growth young firms: contribution to job, output, and productivity growth. In *Measuring entrepreneurial businesses: current knowledge and challenges*, University of Chicago Press, pp. 11–62.

Hartman, R. (1972). The effects of price and cost uncertainty on investment. *Journal of Economic Theory*, **5** (2), 258–266.

Hennessy, C. A. and Whited, T. M. (2007). How costly is external financing? evidence from a structural estimation. *The Journal of Finance*, **62** (4), 1705–1745.

Ilut, C., Kehrig, M. and Schneider, M. (2018). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, **126** (5), 2011–2071.

Jaimovich, N., Terry, S. J. and Vincent, N. (2023). *The Empirical Distribution of Firm Dynamics and Its Macro Implications*. NBER Working Paper 31337.

Jensen, H., Petrella, I., Ravn, S. H. and Santoro, E. (2020). Leverage and deepening business-cycle skewness. *American Economic Journal: Macroeconomics*, **12** (1), 245–281.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, **95** (1), 161–182.

Jordà, O., Schularick, M. and Taylor, A. M. (2024). Disasters everywhere: The costs of business cycles reconsidered. *IMF Economic Review*, **72** (1), 116–151.

Kalemli-Özcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., Volosovych, V. and Yeşiltaş, S. (2024). How to construct nationally representative firm-level data from the orbis global database. *American Economic Journal: Macroeconomics*, **16** (2), 353–374.

Kehrig, M. (2011). *The Cyclicality of Productivity Dispersion*. Working Paper 11-15, Center for Economic Studies, U.S. Census Bureau.

Khan, A. and Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, **76** (2), 395–436.

Kozeniauskas, N., Orlik, A. and Veldkamp, L. (2018). What are uncertainty shocks? *Journal of Monetary Economics*, **100**, 1–15.

Kozlowski, J., Veldkamp, L. and Venkateswaran, V. (2020). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, **128** (8), 2839–2879.

Krusell, P. and Smith, A. A., Jr. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, **106** (5), 867–896.

McKay, A. and Reis, R. (2008). The brevity and violence of contractions and expansions. *Journal of Monetary Economics*, **55** (4), 738–751.

Melcangi, D. and Sarpietro, S. (2024). *Nonlinear Firm Dynamics*. Staff Report 1088, Federal Reserve Bank of New York.

Miranda-Pinto, J., Silva, A. and Young, E. R. (2023). Business cycle asymmetry and input-output structure: The role of firm-to-firm networks. *Journal of Monetary Economics*, **137**, 1–20.

Moskowitz, T. J. and Vissing-Jørgensen, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review*, **92** (4), 745–778.

Oi, W. Y. (1961). The desirability of price instability under perfect competition. *Econometrica*, **29** (1), 58–64.

Panousi, V. and Papanikolaou, D. (2012). Investment, idiosyncratic risk, and ownership. *Journal of Finance*, **67** (3), 1113–1148.

Quadrini, V. (2000). Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics*, **3** (1), 1–40.

Rietz, T. A. (1988). The equity risk premium: a solution. *Journal of Monetary Economics*, **22** (1), 117–131.

Figure 1: Firm Sales Growth has Procyclical Skewness

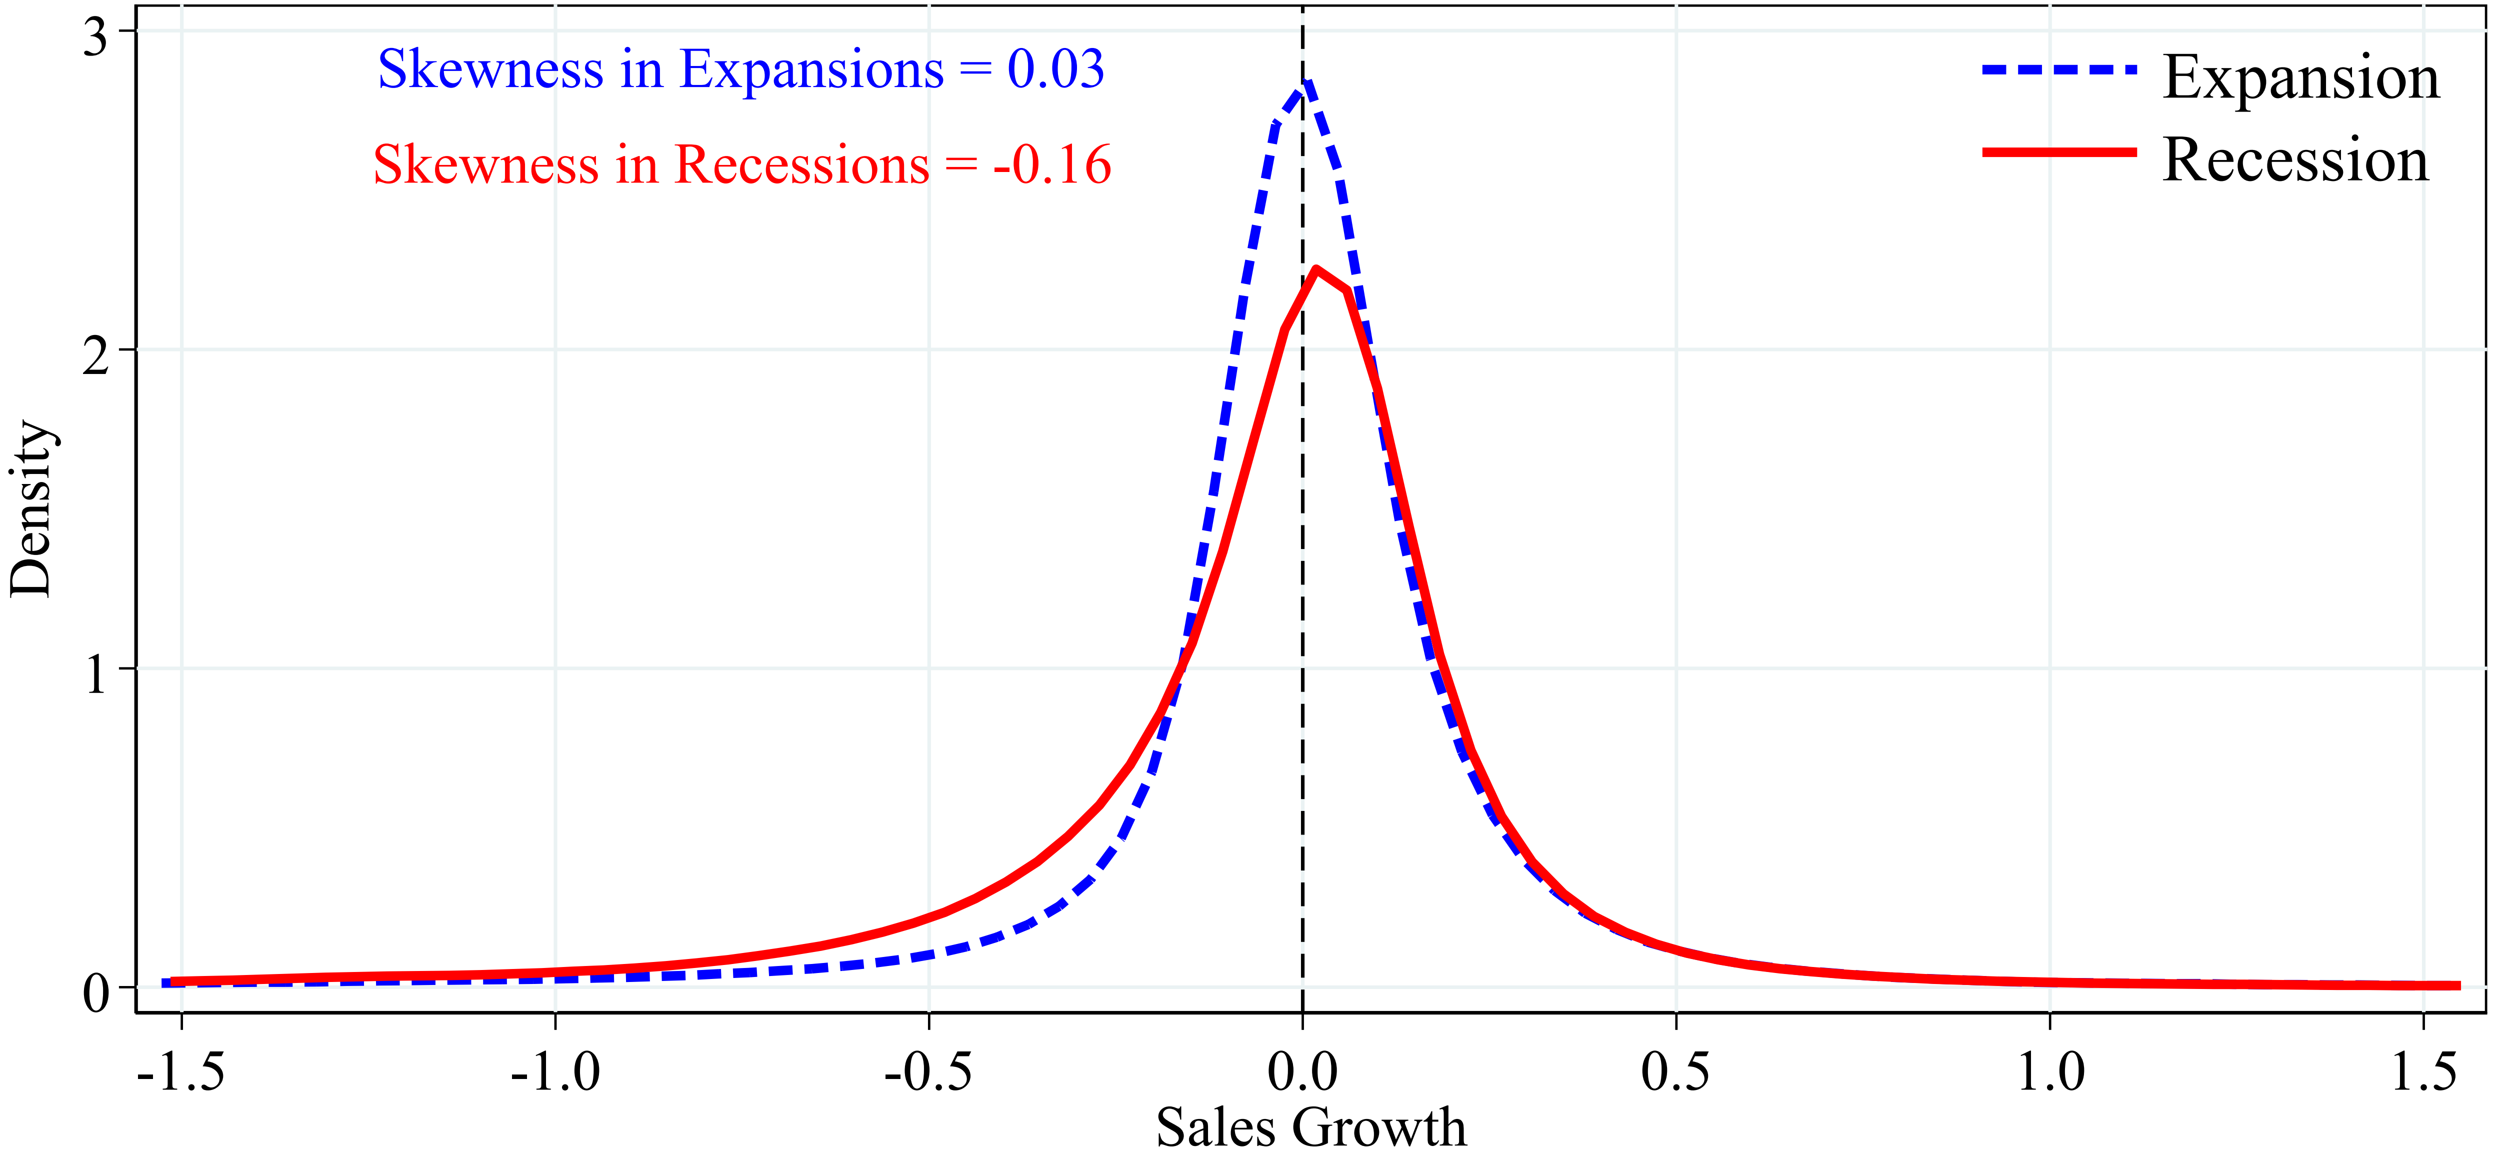

Notes: Distribution of firm sales growth from the US Census' Longitudinal Business Database for firms with 20 employees or more. Establishments combined into firms using firmid. Each density has been adjusted to have a median of zero. Blue-dashed line for expansion years (2003-06 and 2010-19, approximately 590,000 firm-year observations), red-solid line recession years (2001, 2008, and 2020, approximately 570,000 firm-year observations). Sales growth is $\Delta x_t = \frac{x_{t+1} - x_t}{0.5(x_{t+1} + x_t)}$. Skewness is measured using Kelley Skewness. Cross-sectional moments weighed by average sales. Firm sales calculated by summing across all establishments within the firm.

Figure 2: Firm Sales, Employment, Productivity and Stock Market Value Growth have Procyclical Skewness

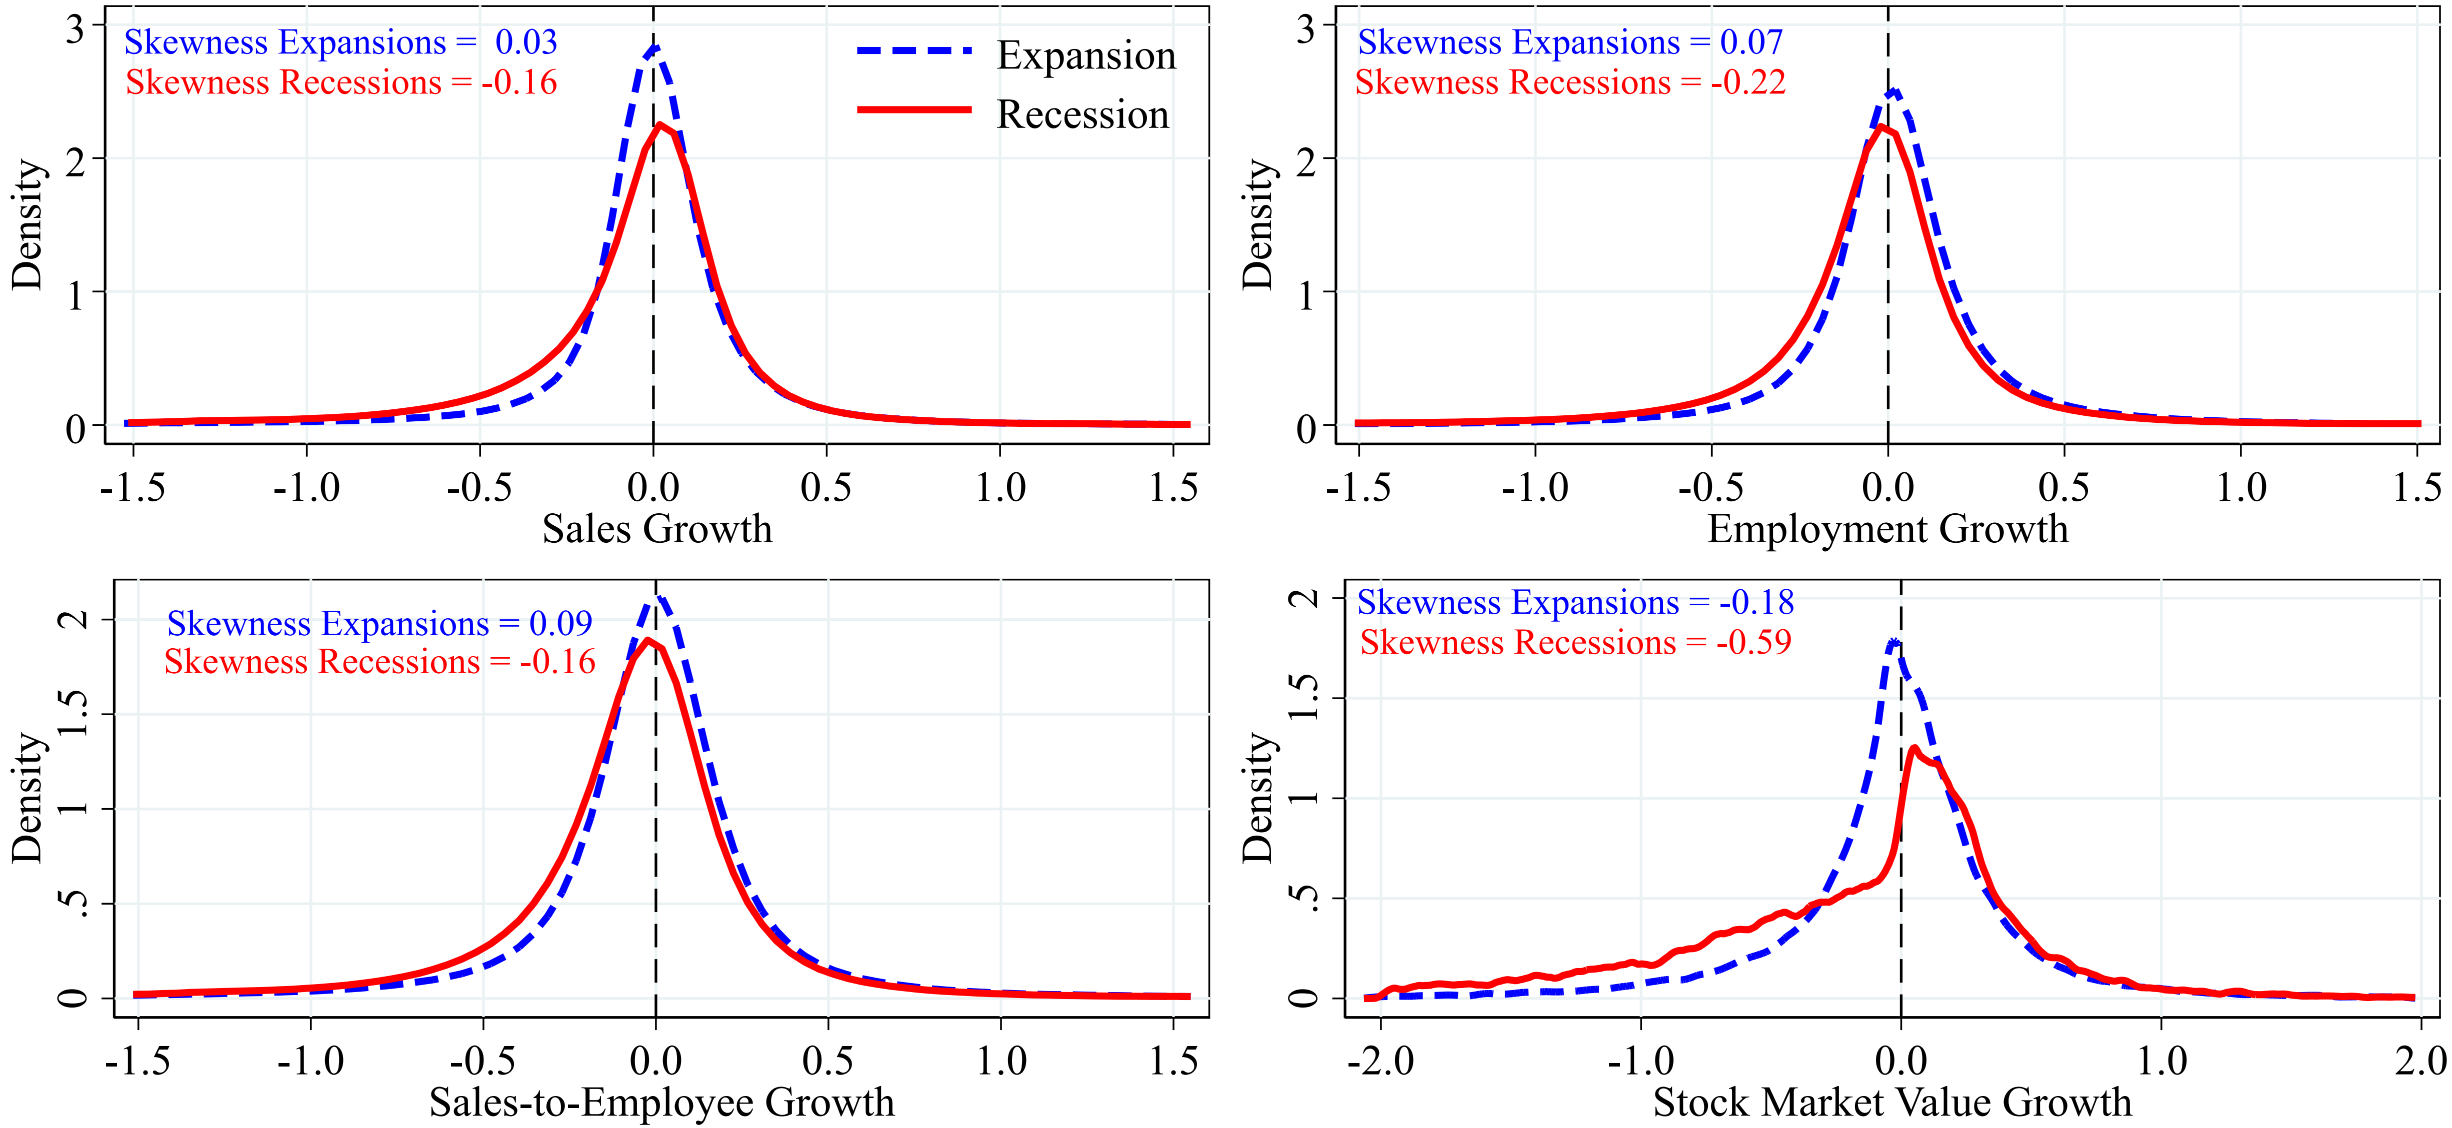

Notes: Distribution of firm growth constructed from the US Census' LBD for firms with 20 employees or more for sales, employment, and sales-per-employee growth. Firm employment sums across all establishments. Stock Market Value Growth for publicly traded firms from CRSP. Each density has adjusted to a median of zero. Blue-dashed line sample of expansion years (2002-07 and 2010-19, approximately 590,000 firm-year observations), red-solid line recession years (2001, 2008, and 2020, approximately 570,000 firm-year observations). Skewness is Kelley Skewness. Cross-sectional moments of sales and sales-to-employee growth weighted by average sales. Employment growths weighted by average employment.

Figure 3: Firm Employment Growth Skewness is Cyclical across Firm Size, Age, Industry and Region splits

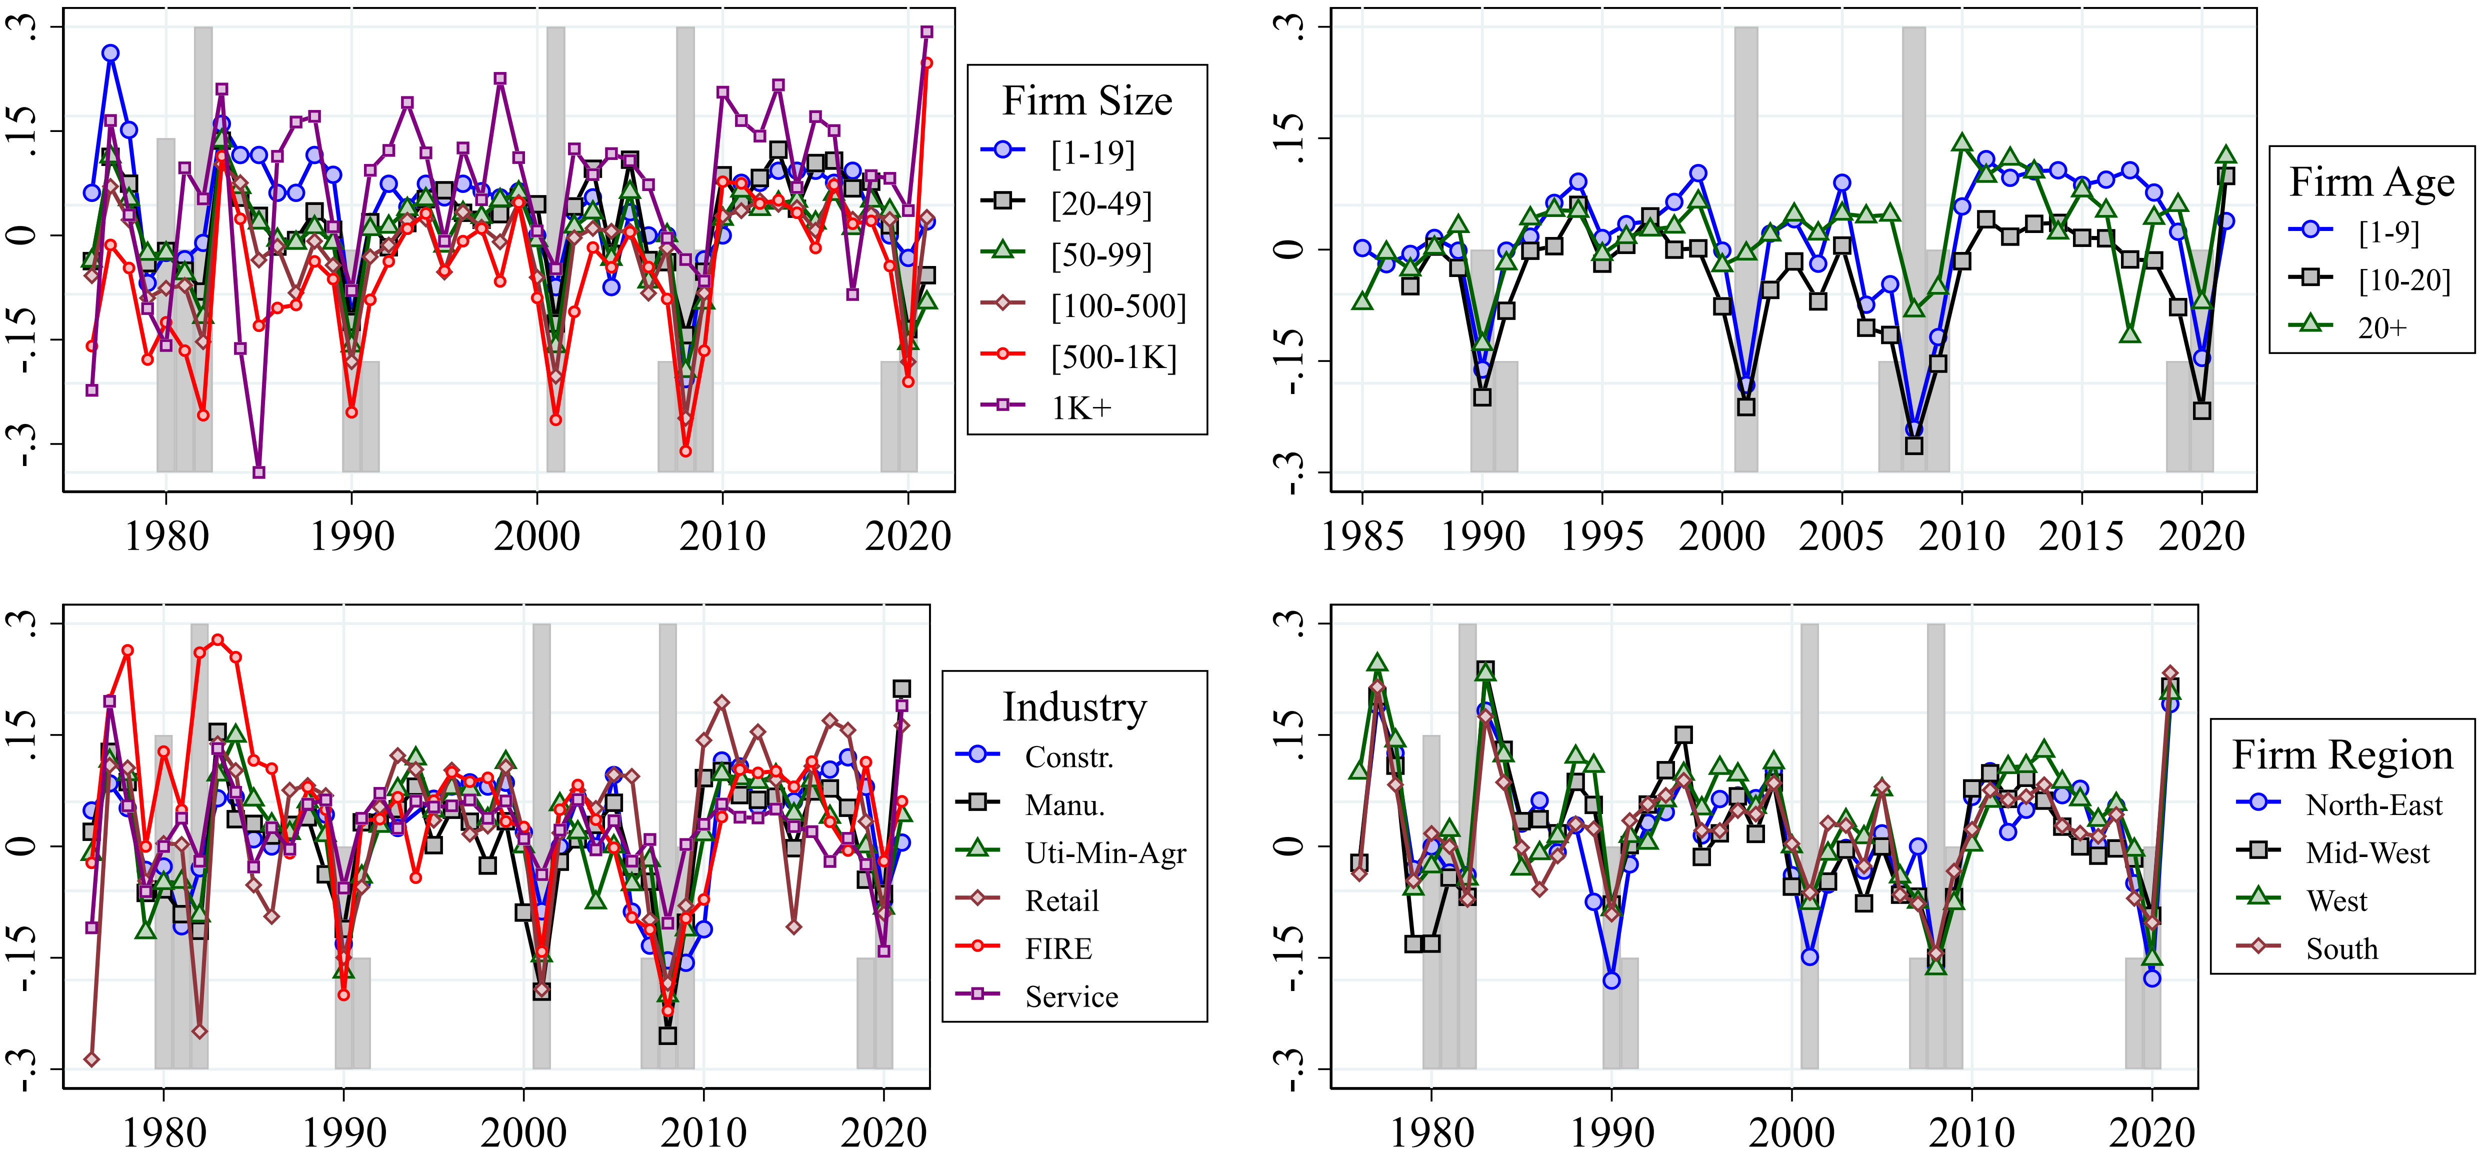

Notes: This figure shows the time series of the cross-sectional Kelley skewness of the distribution of firm employment growth for different firm groups using data from the US Census' Longitudinal Business Database (LBD), containing an average of about 4.6 million firm observations per year. Shaded areas are the share NBER recession quarters within a year. Regions are the four Census Divisions. Employment growth is shown as it is the only indicator available back to 1976.

Figure 4: Firm Sales, Employment, Productivity and Stock Value Growth have Industry Level Procyclical Skewness

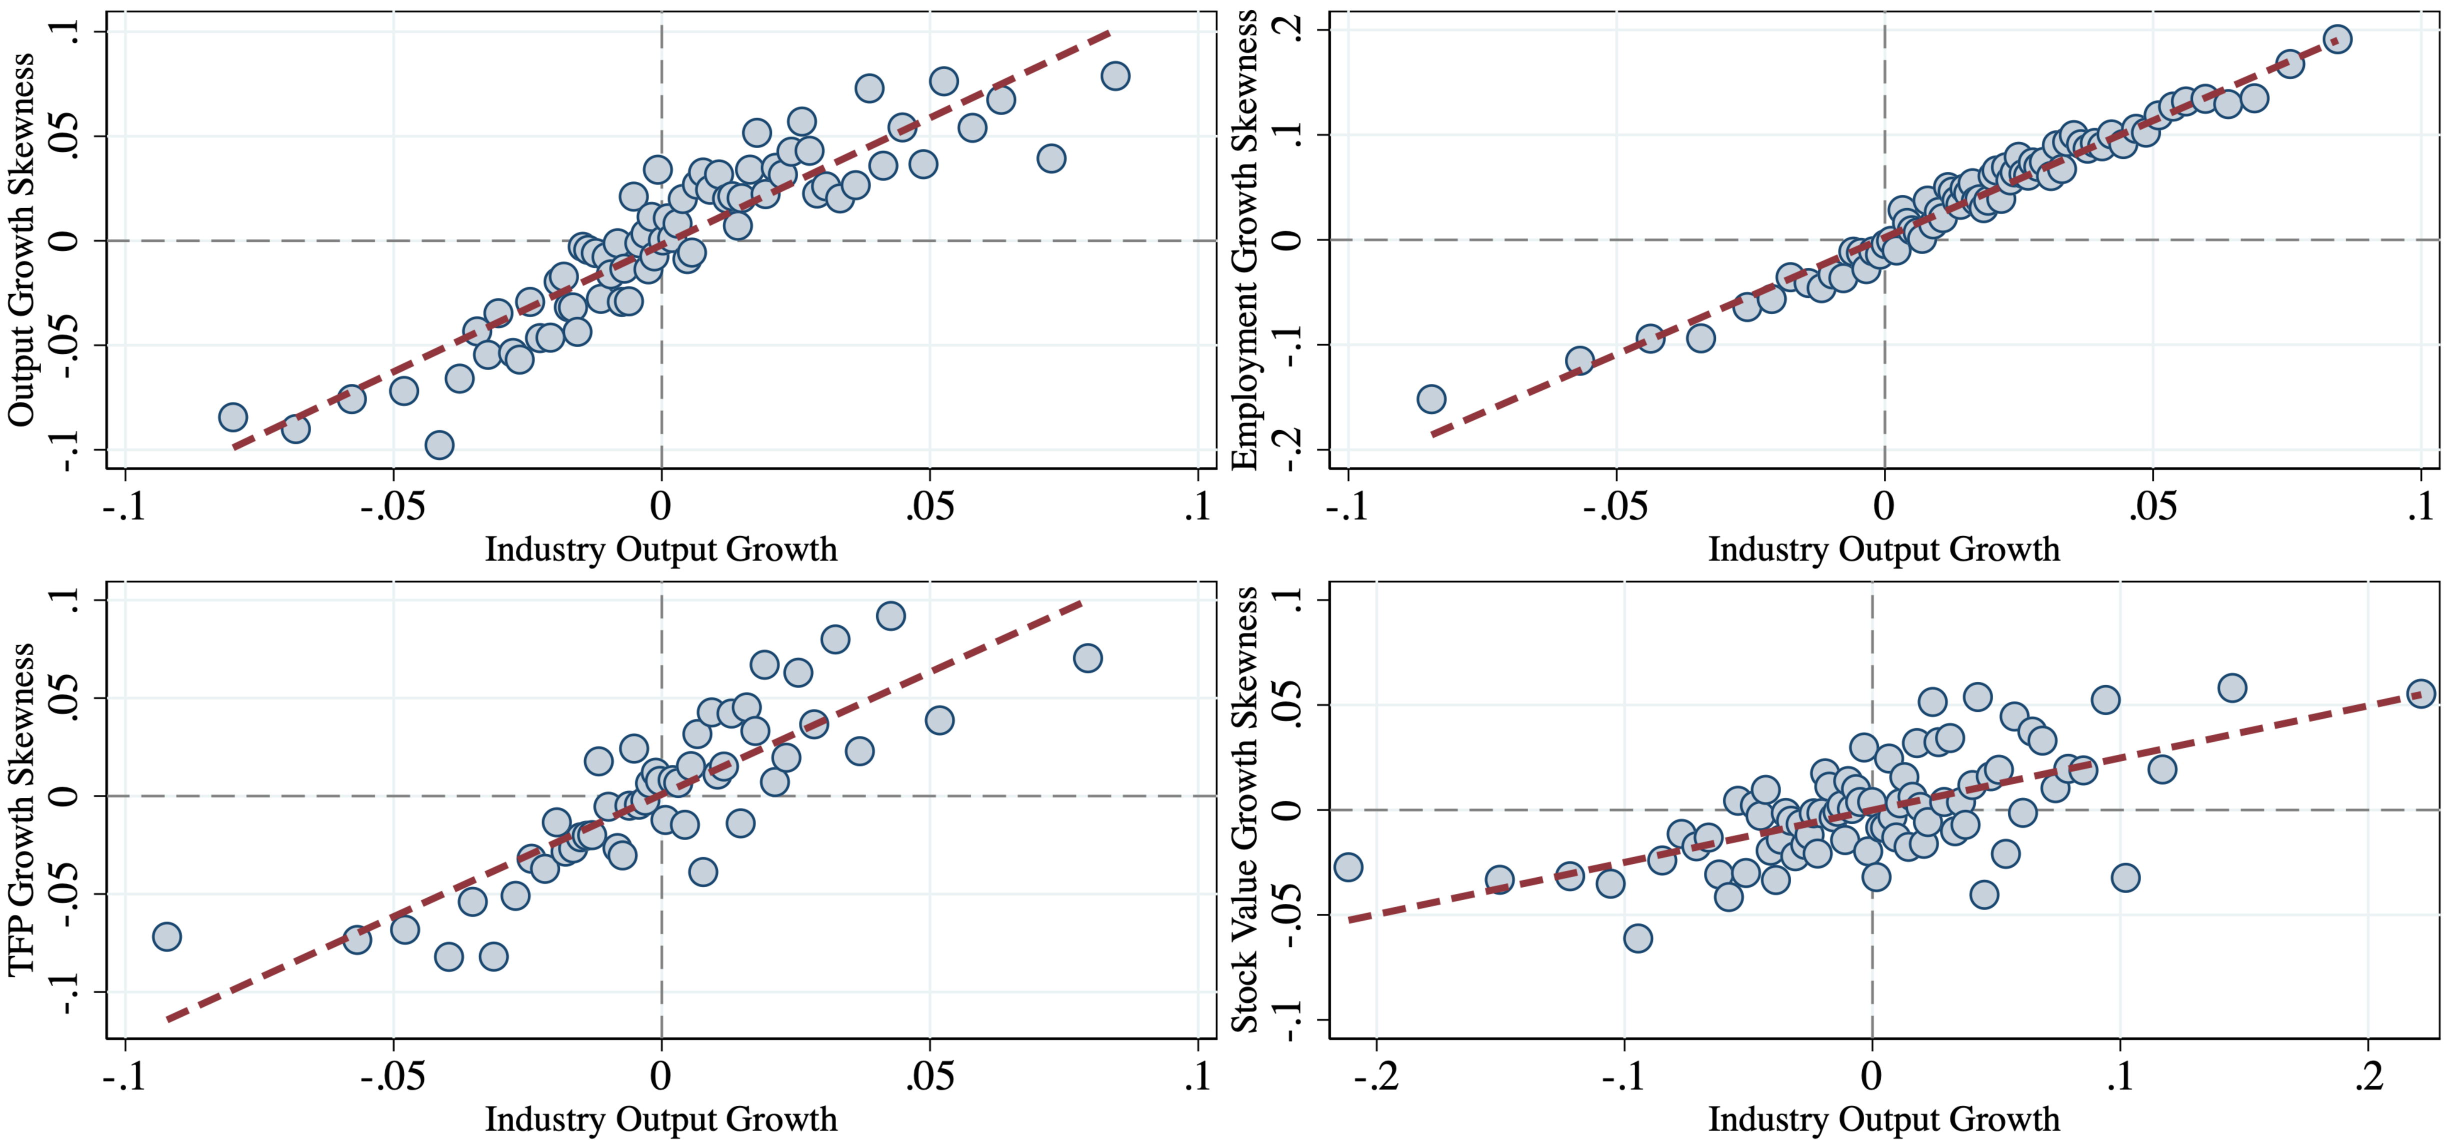

Notes: This figure shows bin-scatters of the Kelley skewness of firm sales growth (top left panel), firm employment growth (top right panel), establishment TFP growth (bottom left panel), and firm stock returns (bottom right panel) against the average output growth within an industry-year. Census data uses 4-digit NAICS cells. Returns data uses 2-digit NAICS. Top panels use data from the US Census' Longitudinal Business Database (LBD) aggregated across establishments to the firm level. Bottom panels use Census ASM and CRSP data respectively. "Industry Output Growth" is average sales growth. Binscatters control for industry and year effects.

Figure 5: Firm Sales, Employment, Productivity and Stock Value Growth have Country Level Procyclical Skewness

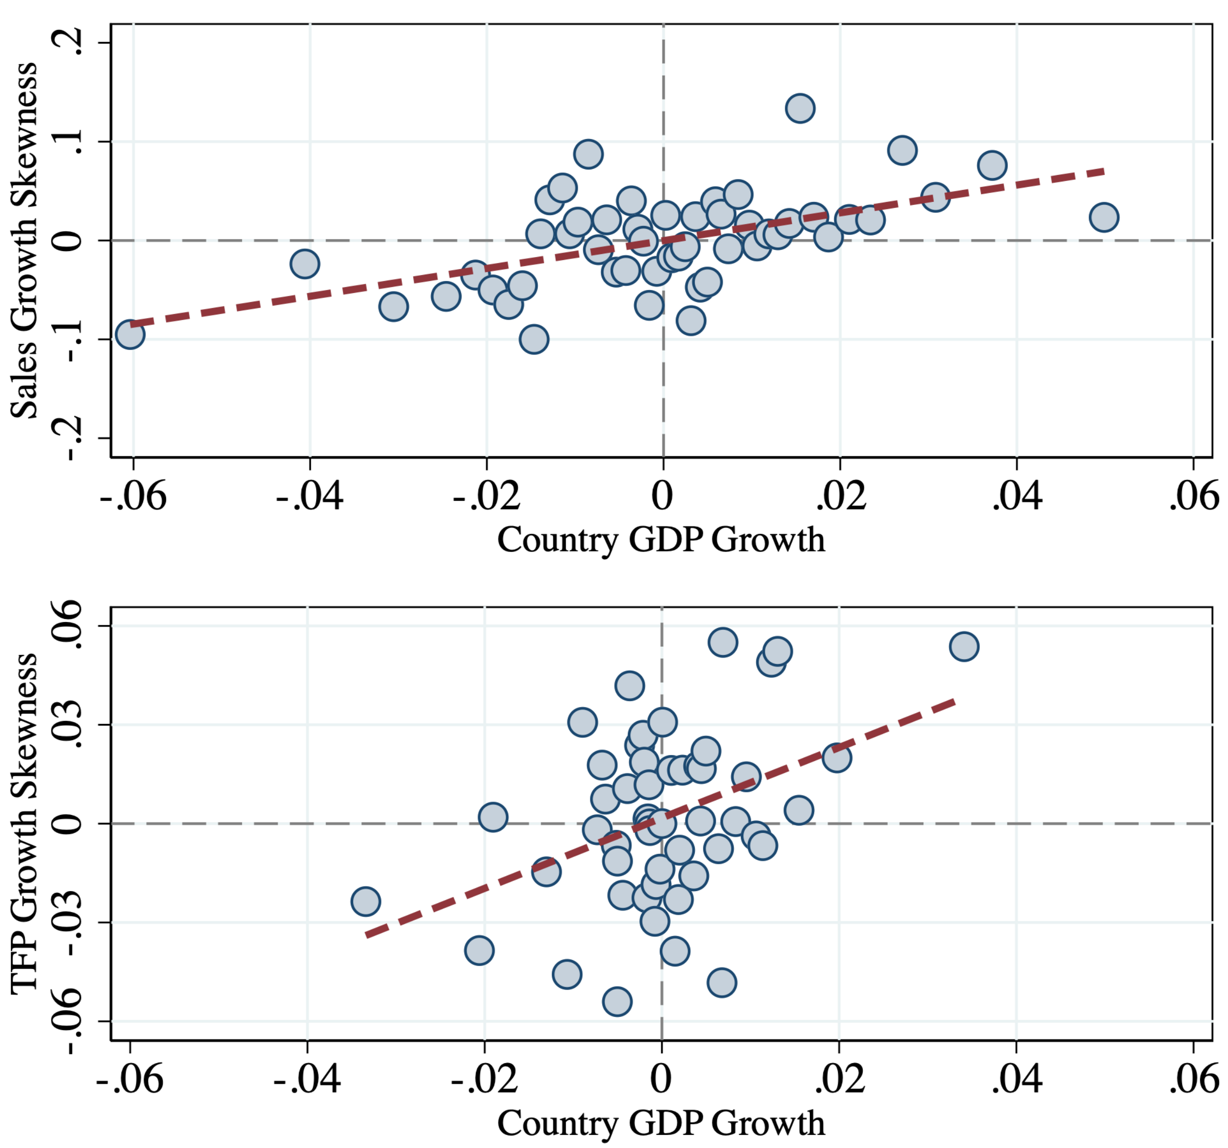

Notes: This figure shows bin-scatters of the Kelley skewness of firm growth against GDP growth. Top panels use data from Osiris for 42 countries covering years 1984 to 2021, for an average sample of 11 thousand firms per year. Bottom left panel uses data from Amadeus with 2-digit industries in 11 countries covering years 1999 to 2020, for a sample of about 320 thousand firms per year. Bottom right panel uses data from Global Compustat for stock returns for 37 countries covering the years 1985 to 2020 for a sample of about 38 thousand firms per year. Binscatters control for country and year fixed effects.

Figure 6: Model Matches Micro Moments of Sales Growth Over the Business Cycle

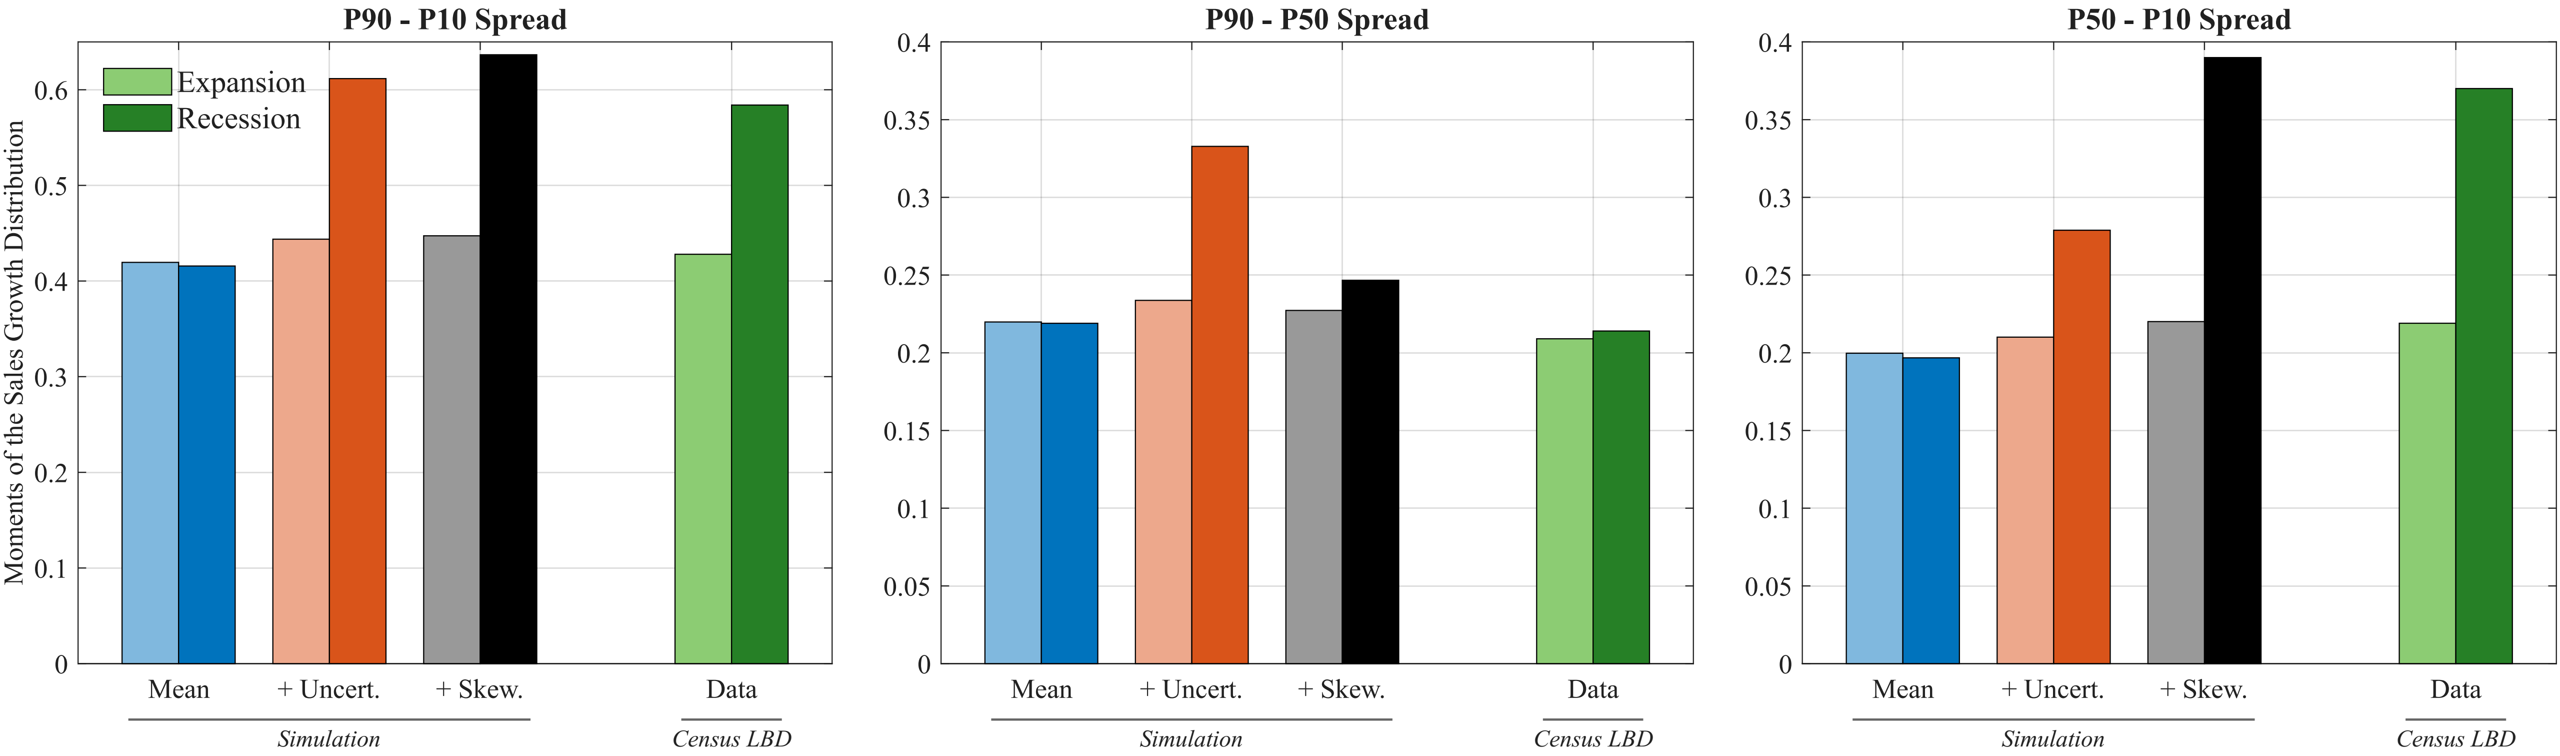

Notes: P90 is the 90th percentile, P50 is the 50th percentile and P10 is the 10th percentile of the annual sales growth distribution. Simulation data based on a simulation of 6,000 firms over 40,000 quarters (500-quarter burn-in discarded). Calendar years are classified as expansion or recession using the norm of 2 or more quarters of negative growth. Each panel shows a different spread moment of annual sales growth: P90-P50 (upper spread), P50-P10 (lower spread), and P90-P10 (total spread). Light and dark bars distinguish expansion and recession years, respectively. The first three bar groups in each panel correspond to the three model specifications (Mean equals first-moment TFP shock only, +Uncertainty equals a first- and second-moment TFP shock, and +Skewness equal a first-, second-, and third-moment TFP shock). The rightmost bar group shows Census LBD firm data targets. Simulated observables include 5% multiplicative measurement error before computing moments.

Figure 7: Growth Moments of Firm Sales, Employment, TFP, and Returns after a Skewness Shock

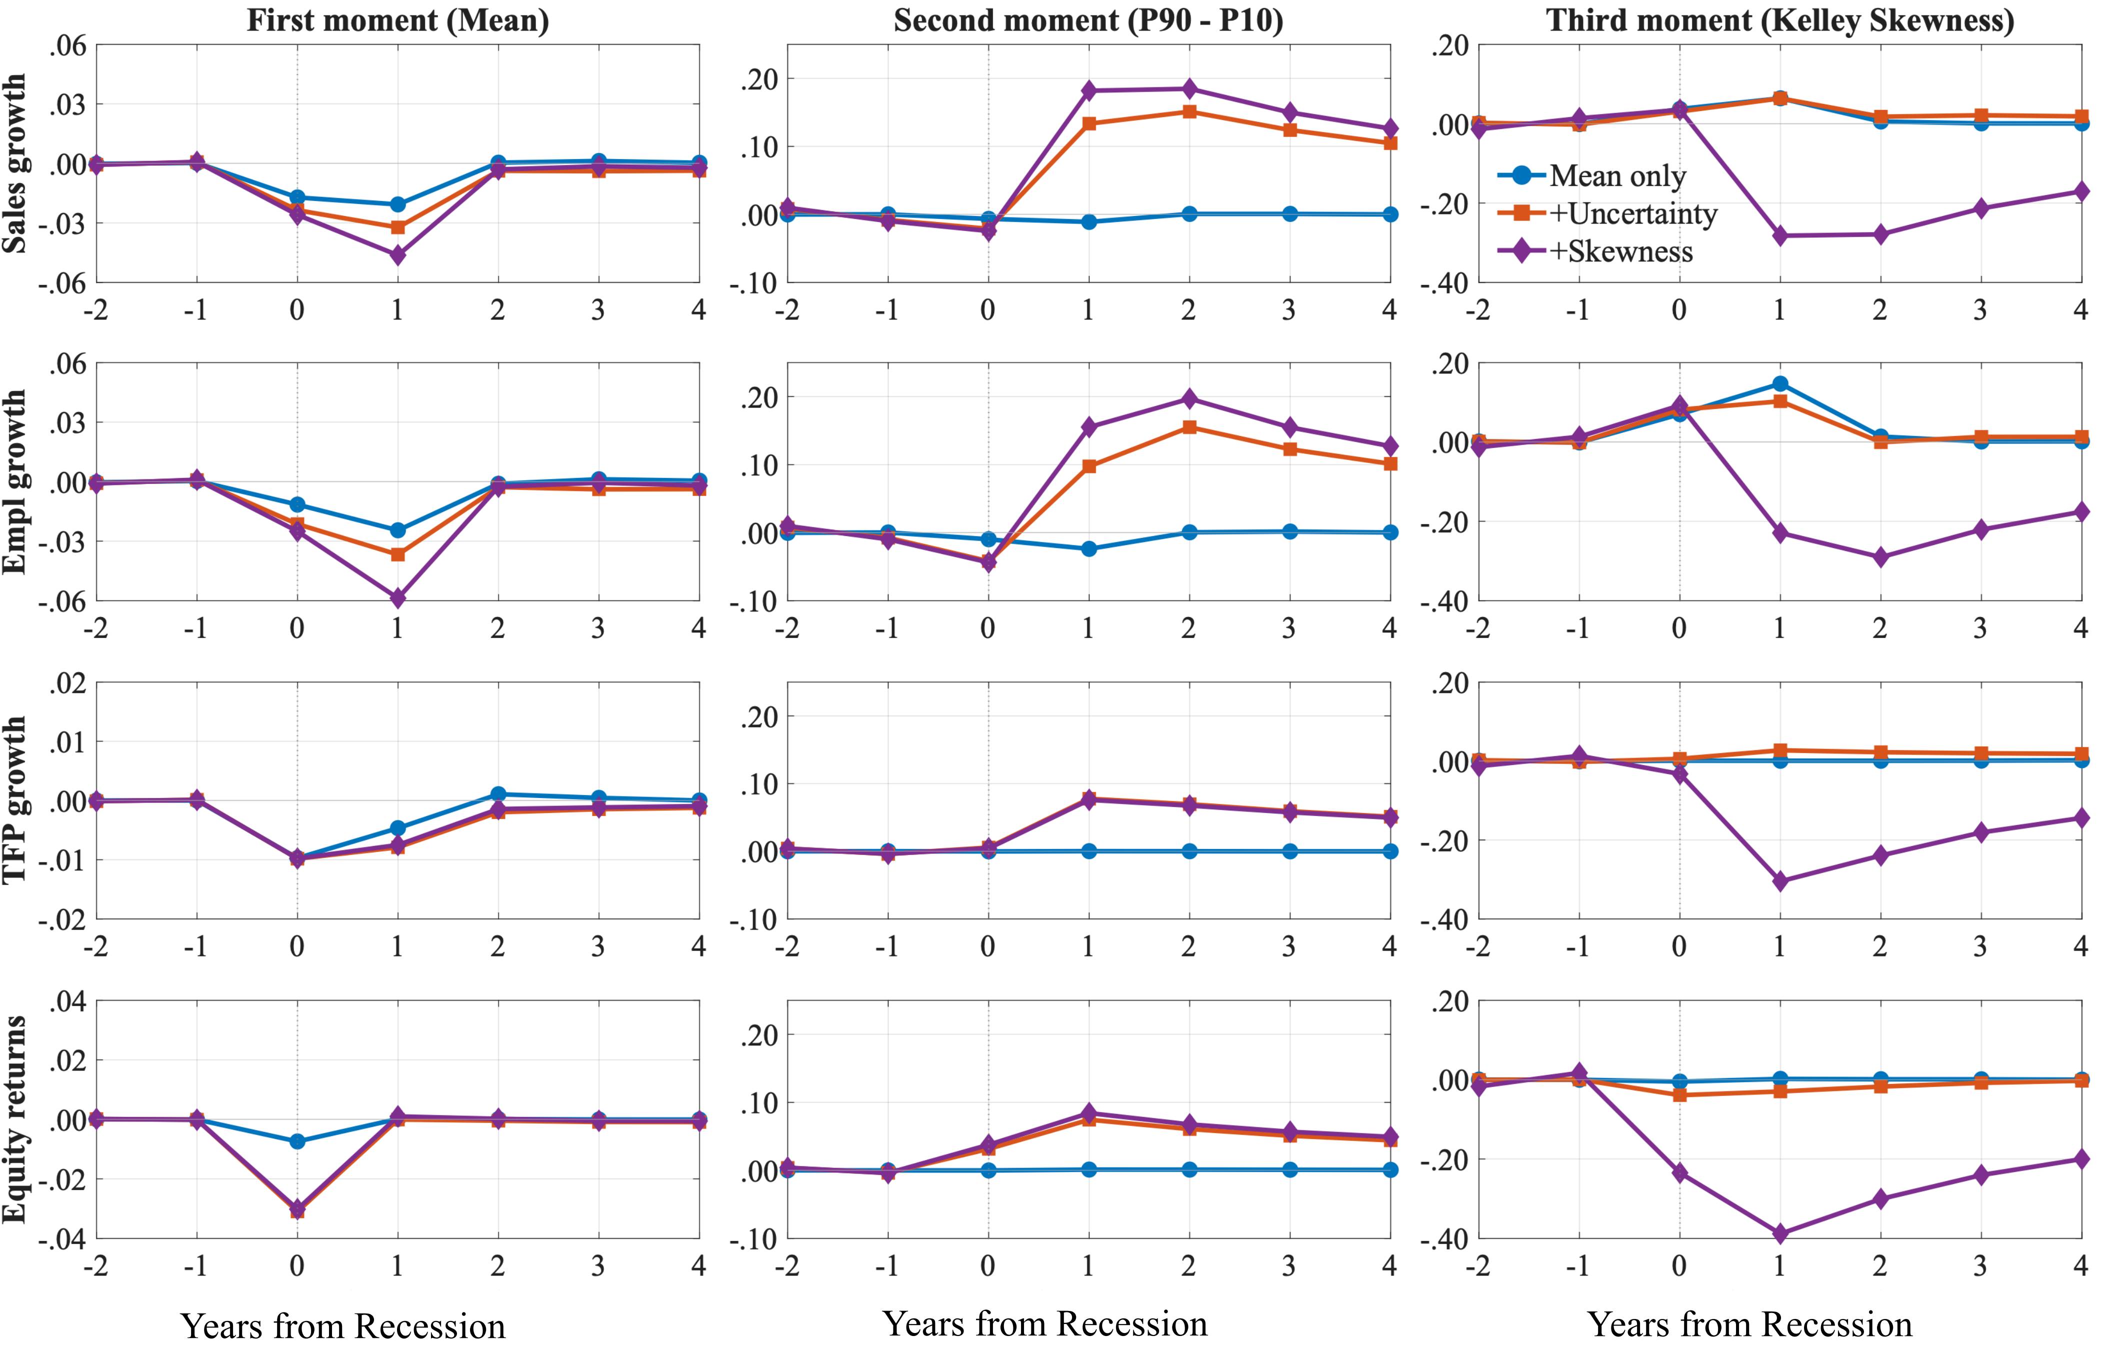

Notes: Based on simulation of 6,000 firms over 40,000 quarters with four variance-matched specifications. Recession episodes are identified by the first recession year of each spell; the window spans years -2 to +5 relative to onset. Each panel plots deviations from the pre-recession average (years -2 and -1). Rows correspond to four variables: annual sales growth, employment growth, TFP growth (log(Y/L)), and equity returns. Columns show the first moment (mean), second moment (P90-P10), and third moment (Kelley skewness). Four lines per panel correspond to: mean + uncertainty, mean + skewness, only uncertainty, and only skewness.

Figure 8: The GDP Response to First, Second, and Third Moment Shocks

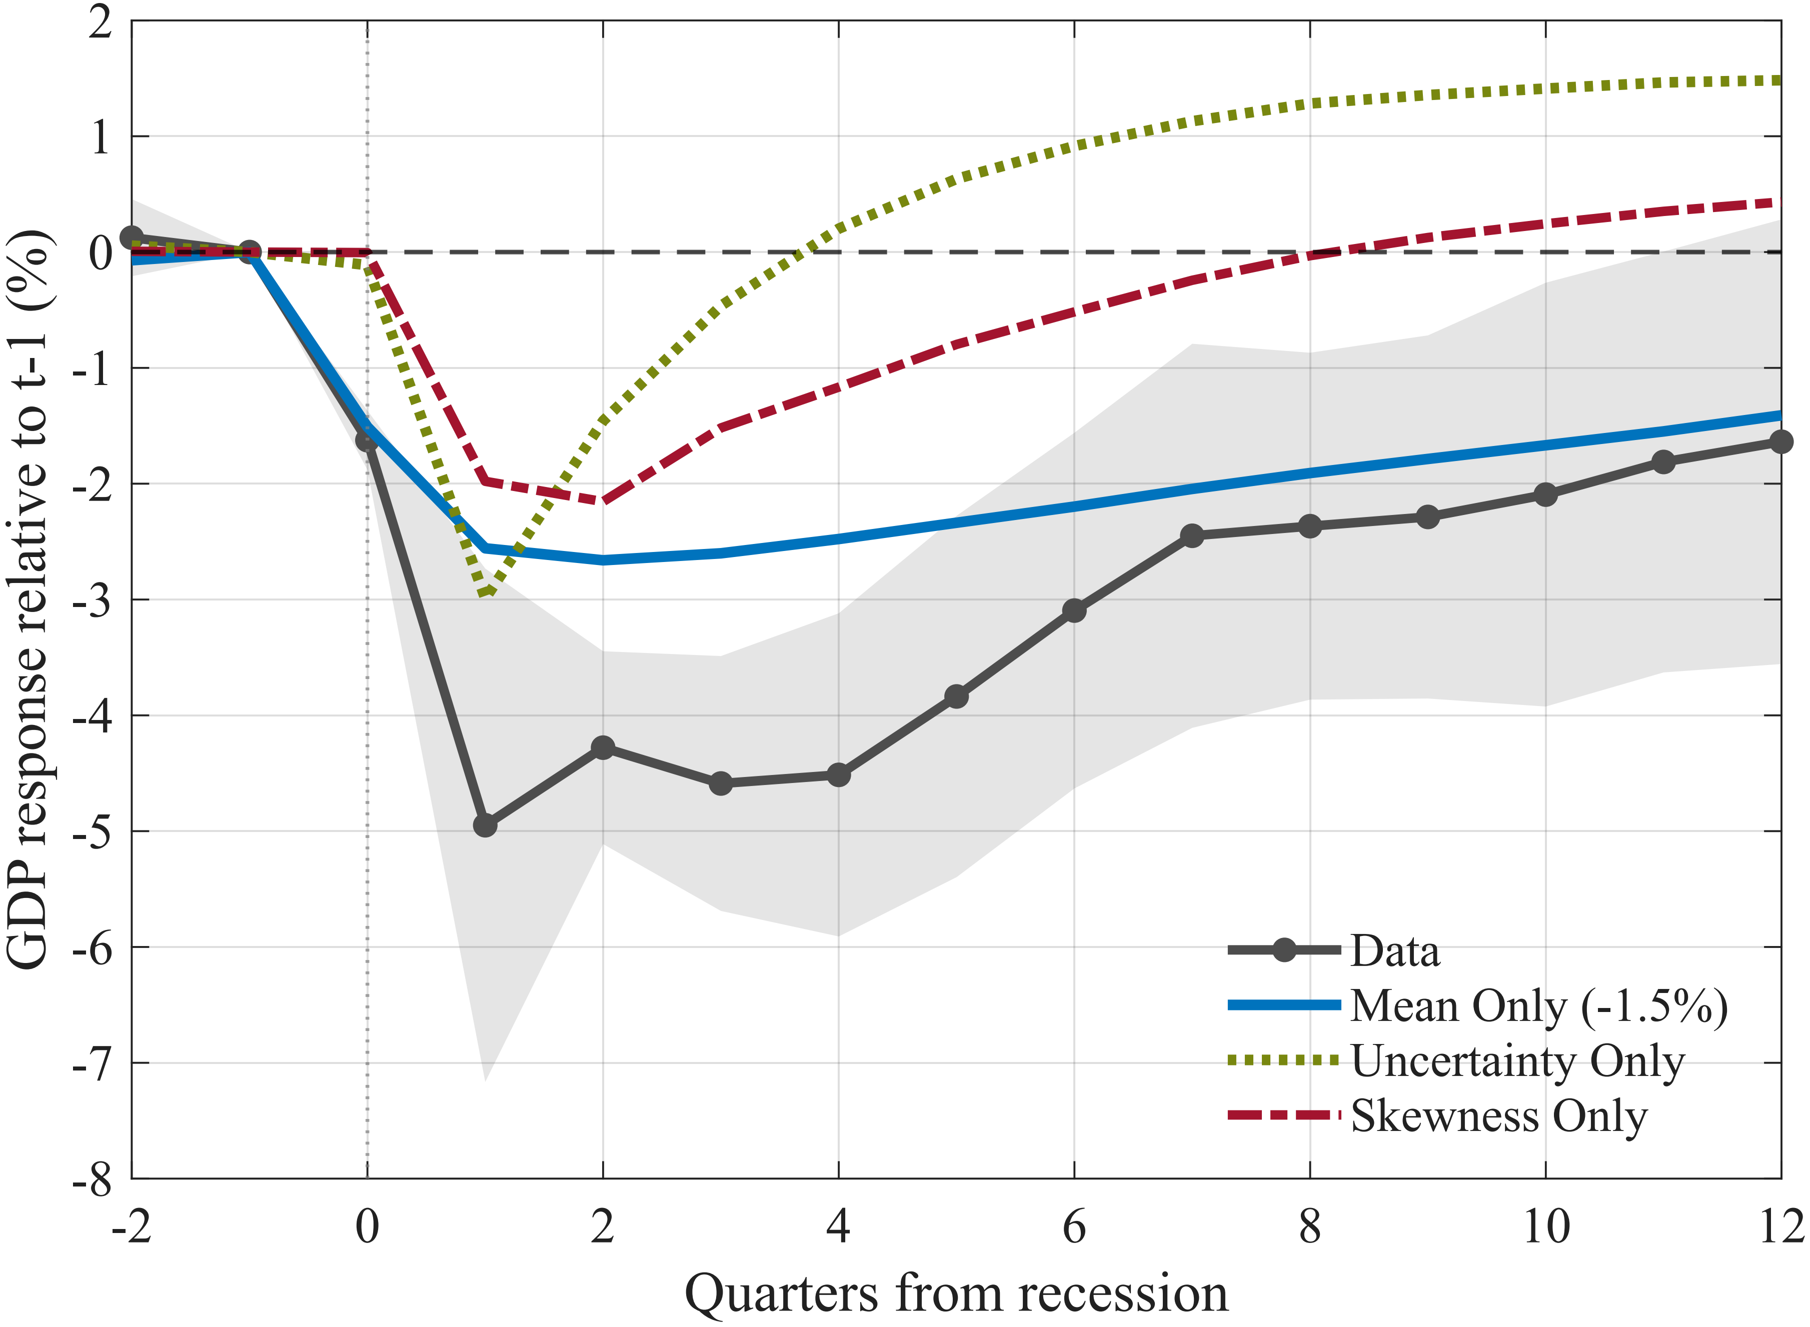

Notes: Jorda (2005) local projections estimated on simulated aggregate time series from a panel of 6,000 firms over 40,000 quarters. The dependent variable is the cumulative log-change in aggregate GDP from quarter t-1 to t+h. The regressor of interest is a recession-onset indicator (transition from expansion to recession regime). Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey-West HAC with bandwidth max(|h|+1, 34). Gray circles show the empirical LP estimated on U.S. FRED data (real GDP), 1970-present sample.

Figure 9: The GDP Response to First, Second and, Third Moment Shocks

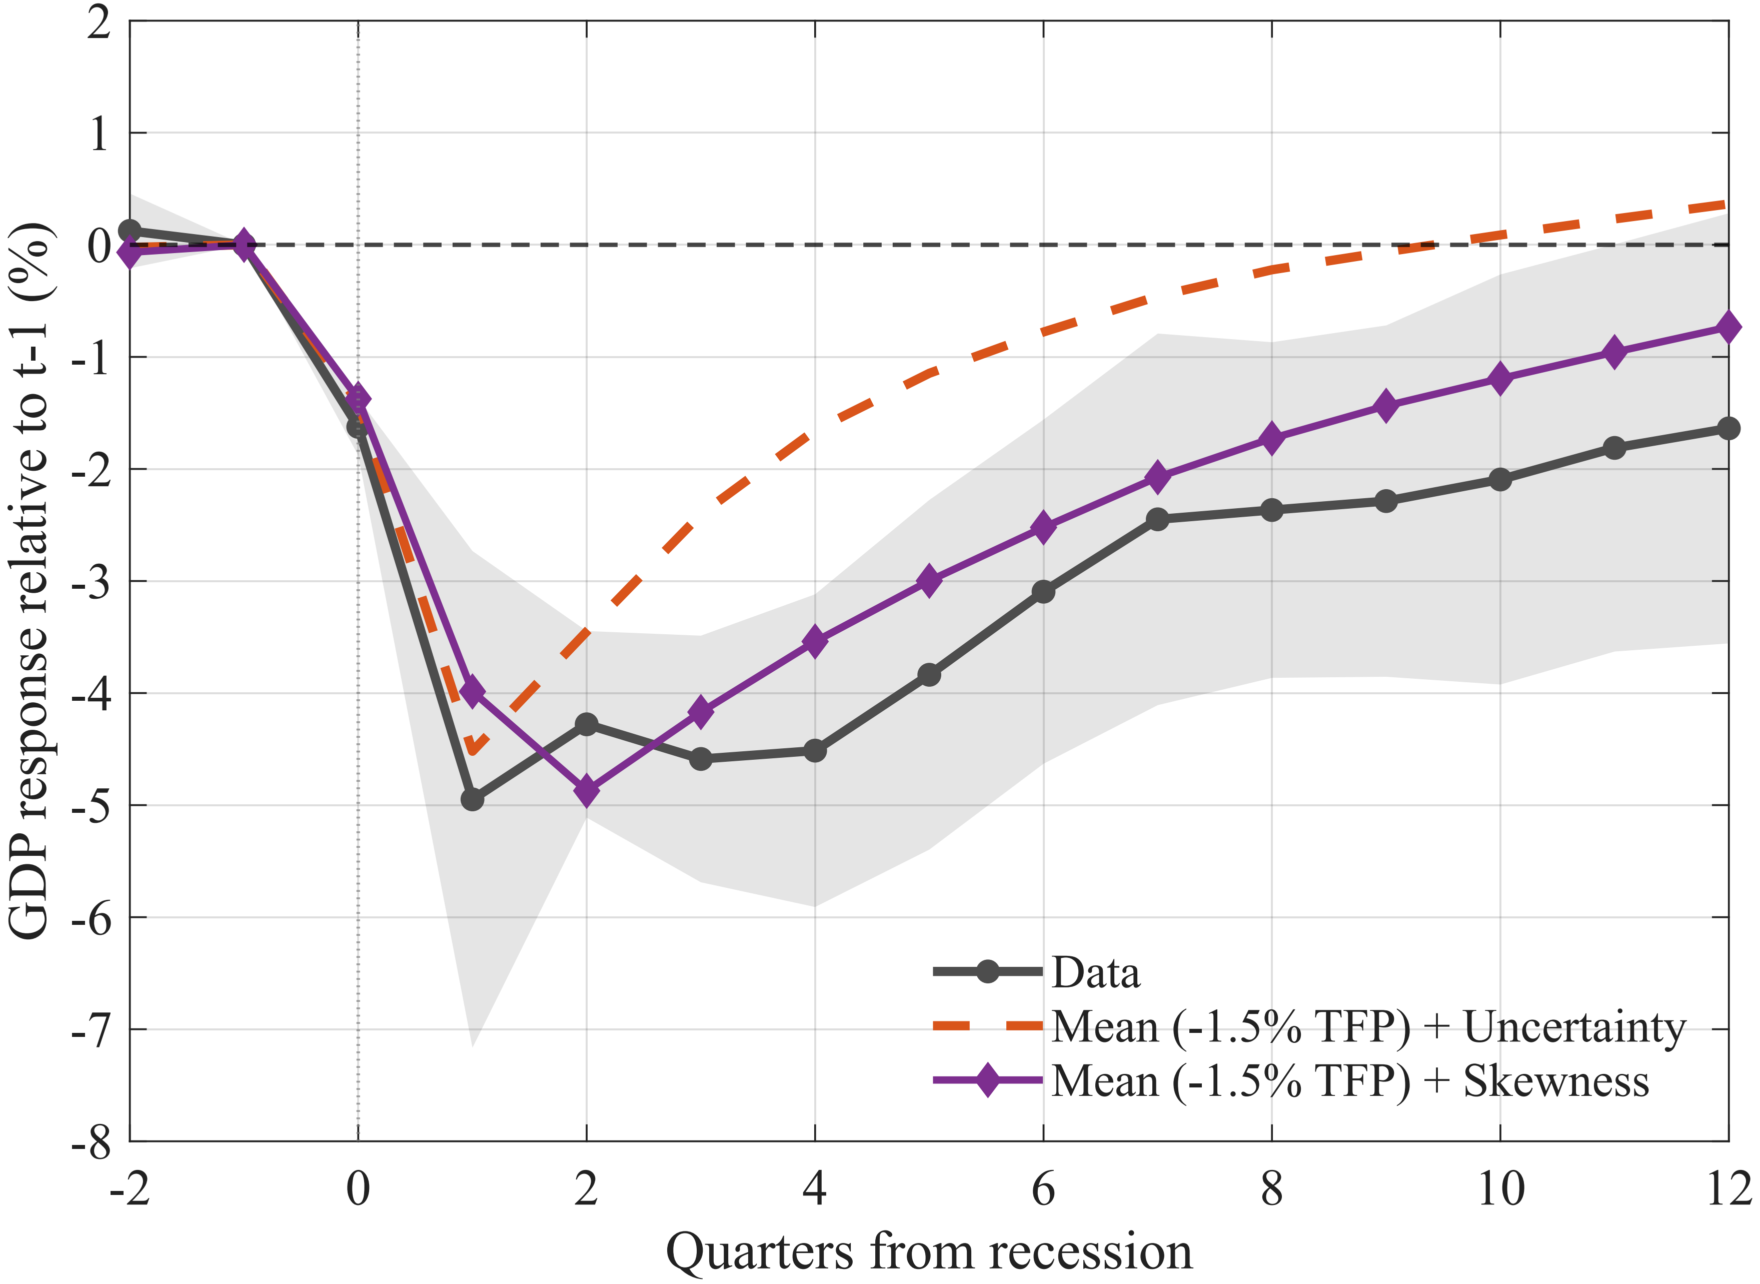

Notes: Jorda (2005) local projections estimated on simulated aggregate time series from a panel of 6,000 firms over 40,000 quarters. The dependent variable is the cumulative log-change in aggregate GDP from quarter t-1 to t+h. The regressor of interest is a recession-onset indicator (transition from expansion to recession regime). Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey-West HAC with bandwidth max(|h|+1, 34). Gray circles show the empirical LP estimated on U.S. FRED data (real GDP), 1970-present sample. The two lines correspond to a Mean + Uncertainty shock and a Mean + Skewness shock.

Table 1: The Skewness of Firm Growth is Procyclical, US Time Series

	(1) $Sales_t$	(2) $Emp._t$	(3) Rev/Emp_t	(4) $Returns_t$
ΔGDP_t	0.068***	0.065***	0.042*	0.040***
	(0.017)	(0.012)	(0.022)	(0.010)
R^2	0.46	0.52	0.34	0.13
Obs.	24	46	24	241
Period	97-20	77-21	97-20	61-21
Freq.	Yr	Yr	Yr	Qtr
Source	LBD	LBD	LBD	CRSP
Expansion	0.03	0.07	0.09	-0.18
Recession	-0.16	-0.22	-0.16	-0.59
Yr. Obs.	4.6M	4.6M	4.5M	4.0K

Notes: Time-series regressions for the United States where the dependent variable is the Kelley skewness of the distribution of one-year firm employment growth, sales growth and sales/employment growth from the US Census' Longitudinal Business Database (LBD) (columns 1 to 3) and one-year stock returns from CRSP (column 4). Cross-sectional moments on sales and revenue per worker weighted by average sales. Employment growth weighted by average employment. GDP growth standardized to have mean zero and standard deviation of 1. Newey-West standard errors below point estimates. Regressions include a linear trend. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: The Skewness of Firm Growth is Procyclical Across US Industries

	(1)	(2)	(3)	(4)	(5)
	$Sales_{j,t}$	$Emp._{j,t}$	$Rev/Emp._{j,t}$	$TFP_{j,t}$	$Returns_{j,t}$
$\Delta Sales_{j,t}$	0.610***	0.577***	0.113*	0.022***	0.020***
	(0.161)	(0.188)	(0.056)	0.010	(0.003)
R^2	0.62	0.66	0.14	0.26	0.41
Obs.	931	938	918	457	3,214
Period	70/24	70/24	70/24	76-19	70/24
Freq.	Yr	Yr	Yr	Yr	Qtr
F.E.	Yr/Ind	Yr/Ind	Yr/Ind	Yr/Ind	Qtr/Ind
Source	CSTAT	CSTAT	CSTAT	ASM	CSTAT
Expansion	0.13	0.11	0.11	0.03	-0.07
Recession	-0.04	-0.13	-0.13	-0.03	-0.15
Yr. Obs.	4.1K	4.1K	4.1K	30.0K	4.1K

Notes: Industry-panel regressions. Dependent variable is the Kelley skewness of the within-industry distribution of firms-level outcomes from Compustat (CSTAT) and the US Annual Survey of Manufacturing (ASM). Industry j at 2-digit NAICS level for all but column (4). ASM industry at 4-digit NAICS. Independent variable is average sales growth within an industry-year. Standard errors clustered at the industry level. Compustat sample consider firms with 10+ years of data. Regressions include a linear trend. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: The Skewness of Firm Growth is Procyclical Across Countries

	(1) Sales$_{j,t}$	(2) Emp.$_{j,t}$	(3) TFP$_{j,t}$	(4) Returns$_{j,t}$
$\Delta GDP_{j,t}$	0.034***	0.068***	0.018***	0.017***
	(0.009)	(0.012)	(0.004)	(0.002)
R^2	0.378	0.275	0.521	0.310
Obs.	989	957	2648	5418
Period	84-18	84-18	99-18	85-20
Freq.	Yr	Yr	Yr	Qtr
F.E.	Yr/Ctry	Yr/Ctry	Yr/Ctry	Qtr/Ctry
Source	BvD	BvD	Amadeus	GCSTAT
Expansion	0.14	0.07	0.03	-0.06
Recession	-0.02	-0.04	0.01	-0.14
Sample	12.3K	21.0K	411K	38K

Notes: Country-level panel regressions in which the dependent variable is the Kelley skewness of different firm-level outcomes in country k in period t . Employment and sales data come from Osiris dataset (42 countries), TFP shocks are estimated using data from the Amadeus dataset (11 countries), and stock returns are from Global Compustat (GCSTAT, 37 countries). In each regression, the independent variable is the one-year GDP per capita growth. Standard errors are clustered at the country level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Model Parameters

Parameter	Value	Source	Parameter	Value	Source
Production & labor			*Productivity process*		
α	0.65	Aggr. labor share	ρ_z	0.95	Khan & Thomas (2008)
A_E	1.0	Normalisation	σ_E	0.032	P90–10 sales, Exp
A_R	0.985	1.5% TFP drop	σ_R	0.025	P90–10 sales, Rec
δ_L	0.044	BLS/JOLTS	μ_D	-0.28	P50–10 sales
			p_E	0.012	P90–10 sales, Exp
			p_R	0.055	P50–10 sales, Rec
Adjustment costs Bloom (2009)			*Financial frictions*		
ψ_H	0.018	Hiring cost	κ	0.9686	Med. cash/sales
ψ_F	0.018	Firing cost	η_1	0.06	Hennessy & Whited (2007)
ψ_S	0.021	Disruption cost	η_2	0.50	Hennessy & Whited (2007)
Aggregate regime			*Preferences*		
p_{ee}	0.97	Avg. NBER expansion	γ	2.0	Risk aversion
p_{rr}	0.95	Khan & Thomas (2008)	β	$0.95^{1/4}$	Annual disc. 5%
			ε	5.0	Frisch elast.
Heterogeneity & measurement					
σ_θ	0.498	Top 10% empl. share	σ_{me}	0.05	Census LBD/ASM

Notes: Table shows calibrated parameters of the model. The "Source" column indicates the targeted moments or external estimate used to discipline each parameter. Sales growth moments from LBD. Cash/Sales ratio from Compustat. Cross-sectional moments computed from simulation after applying 5% multiplicative measurement error to firm-level sales and employment, following Census microdata.

Supplementary

Online Appendix

Figure A1: Firm Sales, Employment, Productivity and Stock Market Value Growth have Procyclical Skewness

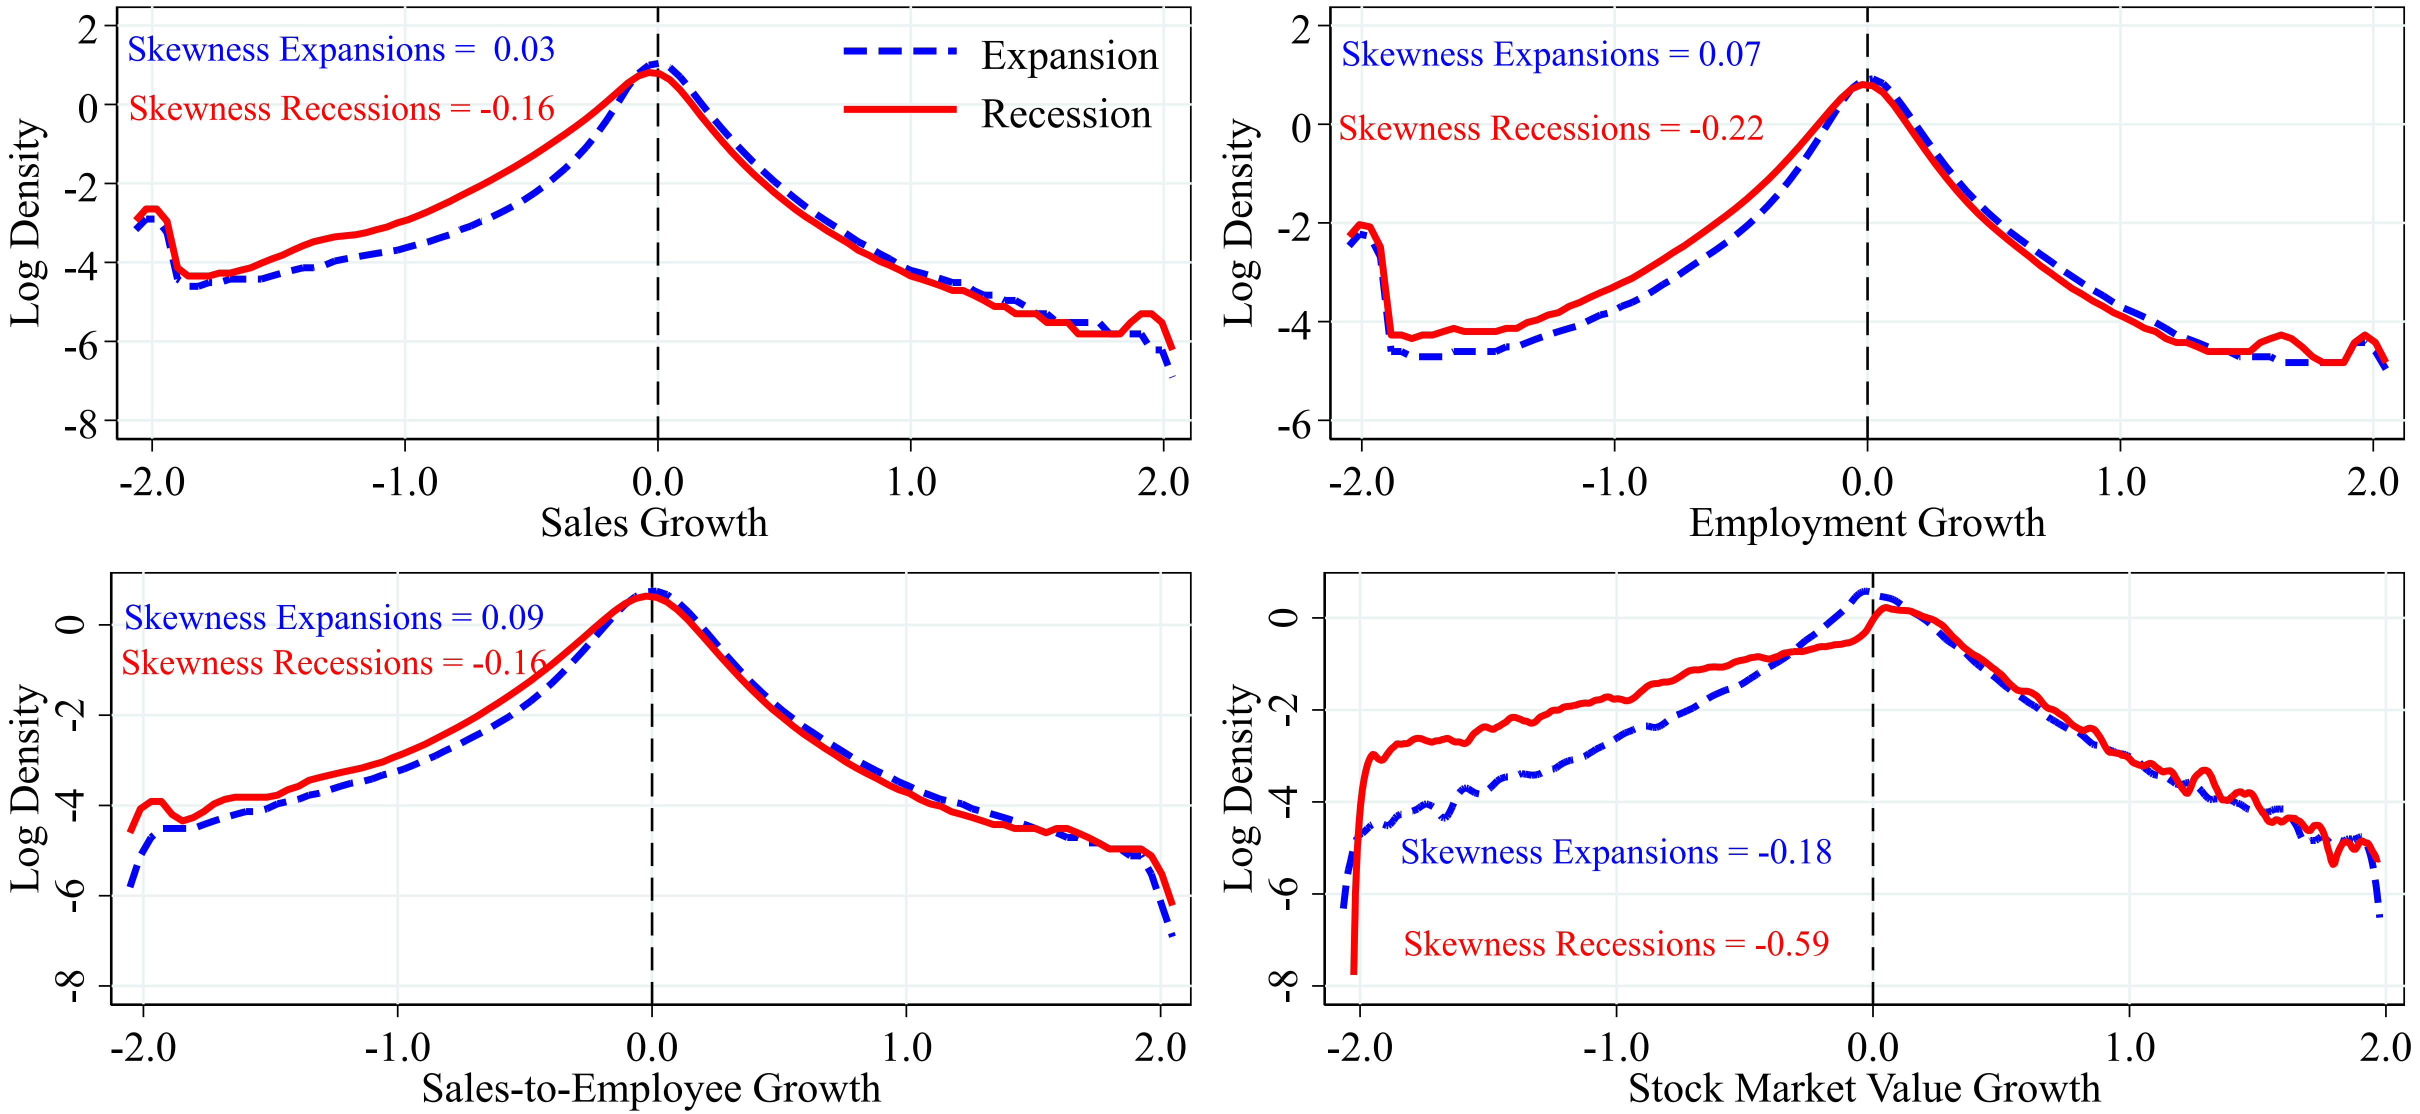

Notes: Distribution of firm growth constructed from the US Census' LBD for firms with 20 employees or more for sales, employment, and sales-per-employee growth. Firm employment sums across all establishments. Stock Market Value Growth for publicly traded firms from CRSP. Each density has adjusted to a median of zero. Blue-dashed line sample of expansion years (2003 to 2006 and 2010 to 2019) red-solid line sample of recession years (2001, 2008, and 2020). We exclude 2002, 2007, and 2009 from the sample since are not full recession years. Skewness is Kelley Skewness. Cross-sectional moments of sales and sales-to-employee growth weighted by average sales. Employment growths weighted by average employment.

Figure A2: Firm Sales Growth Skewness is Cyclical for all Firm Size, Age, Industry and Region Splits

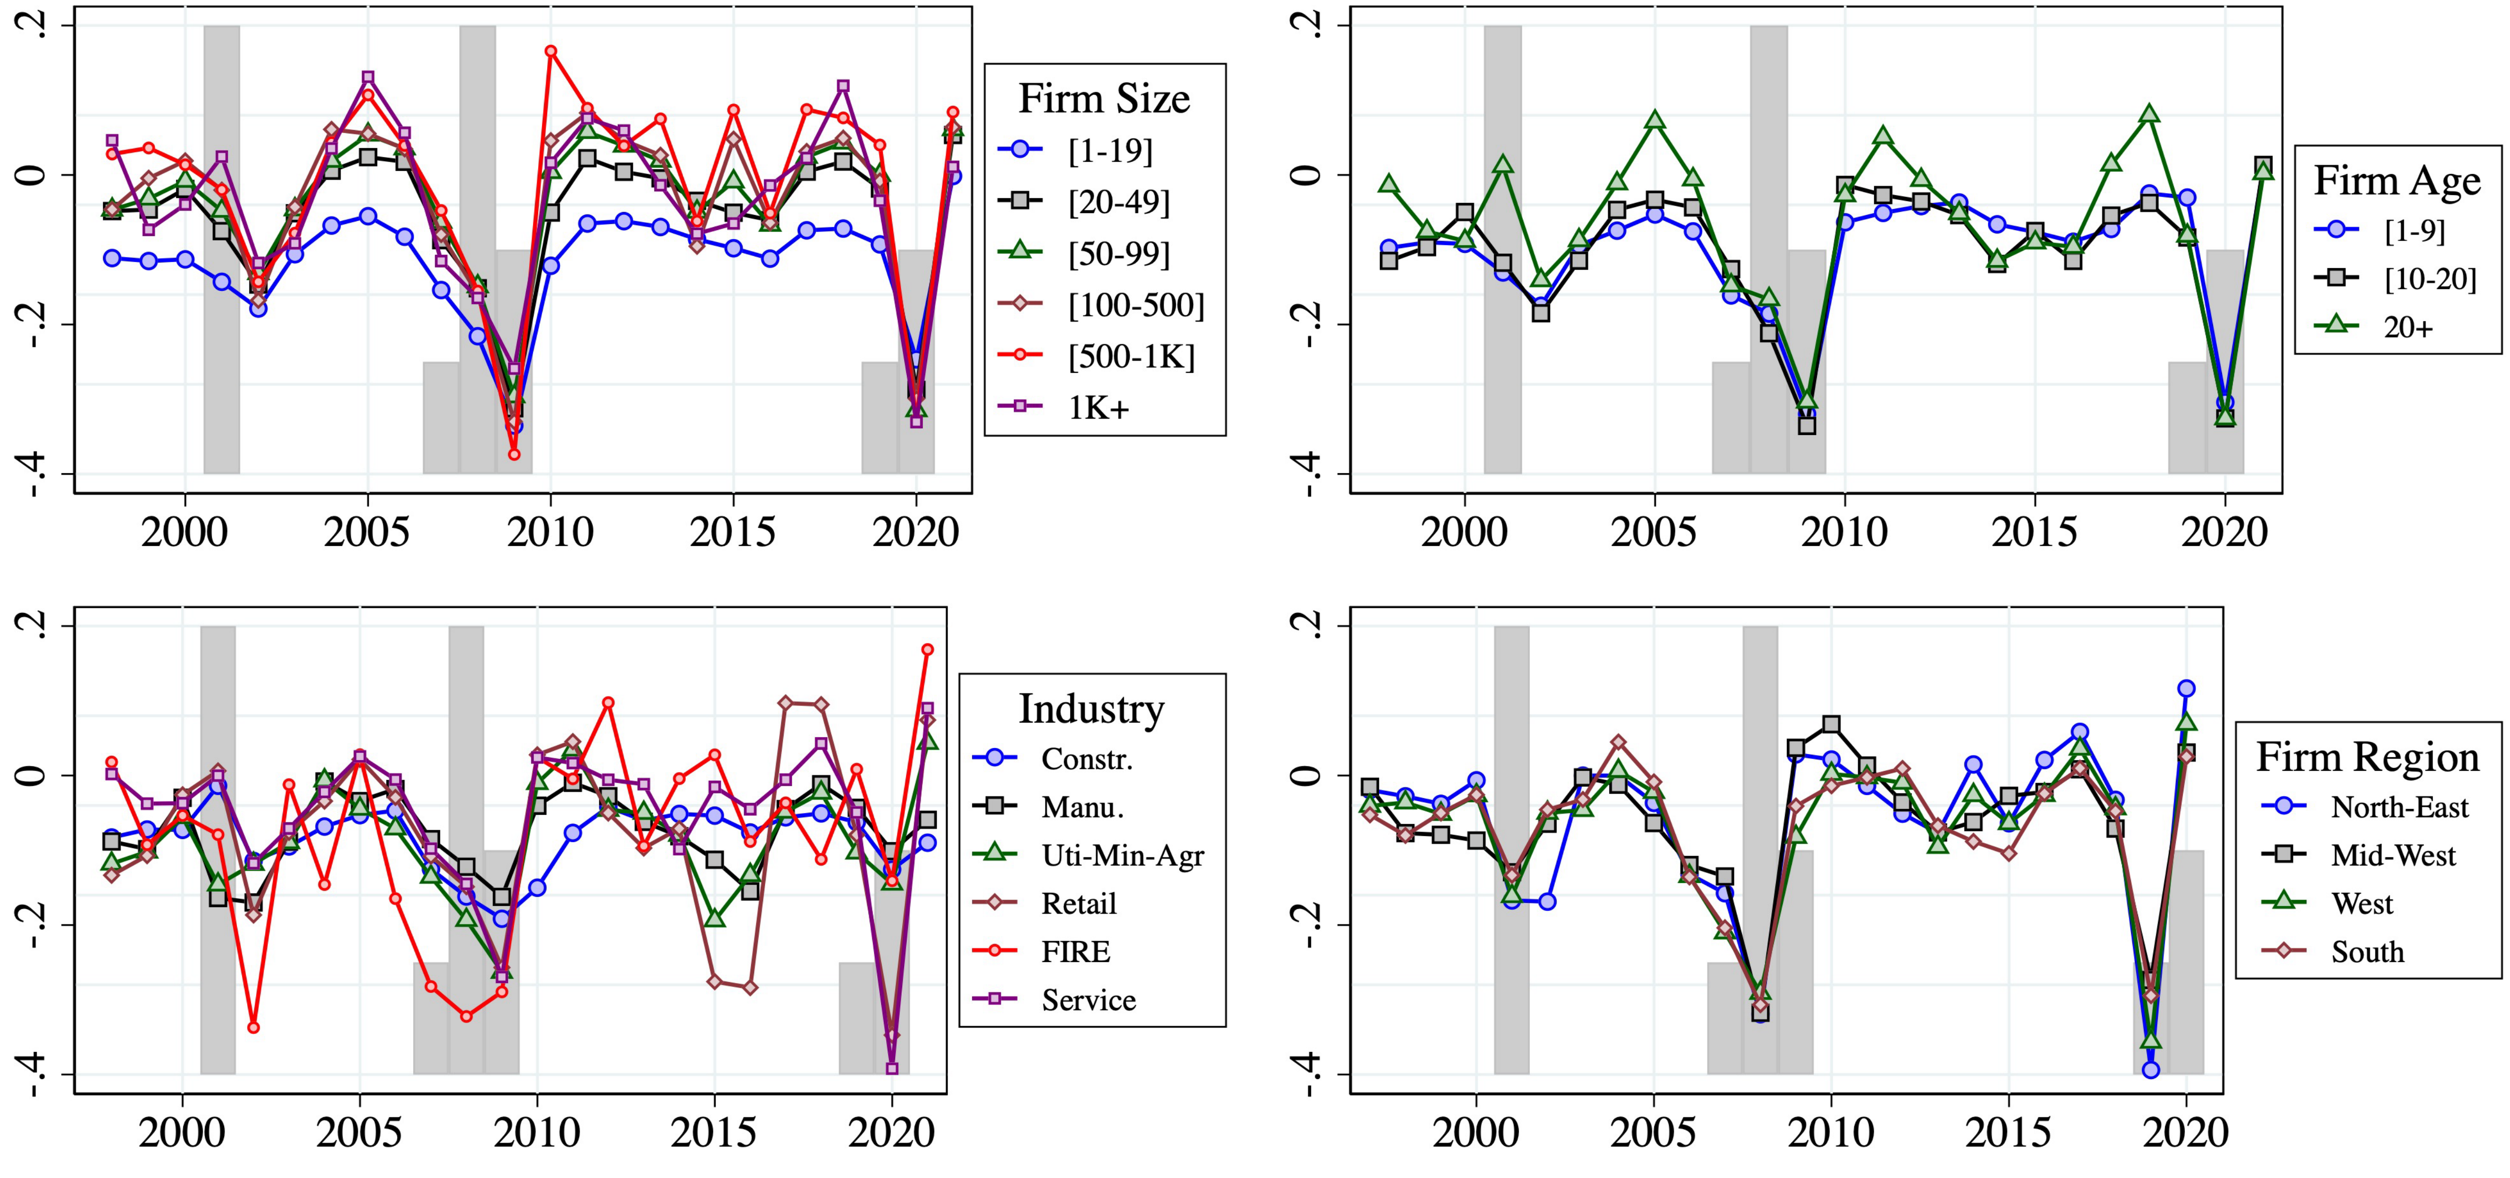

Notes: This figure shows the time series of the cross-sectional Kelley skewness of the distribution of firm sales growth for different firm groups using data from the US Census' Longitudinal Business Database (LBD), containing an average of about 4.6 million firm observations per year. Shaded areas are the share NBER recession quarters within a year. Regions are the four Census Regions.

Appendix Figure A3: The Skewness of Output Growth is Procyclical Within US Industries

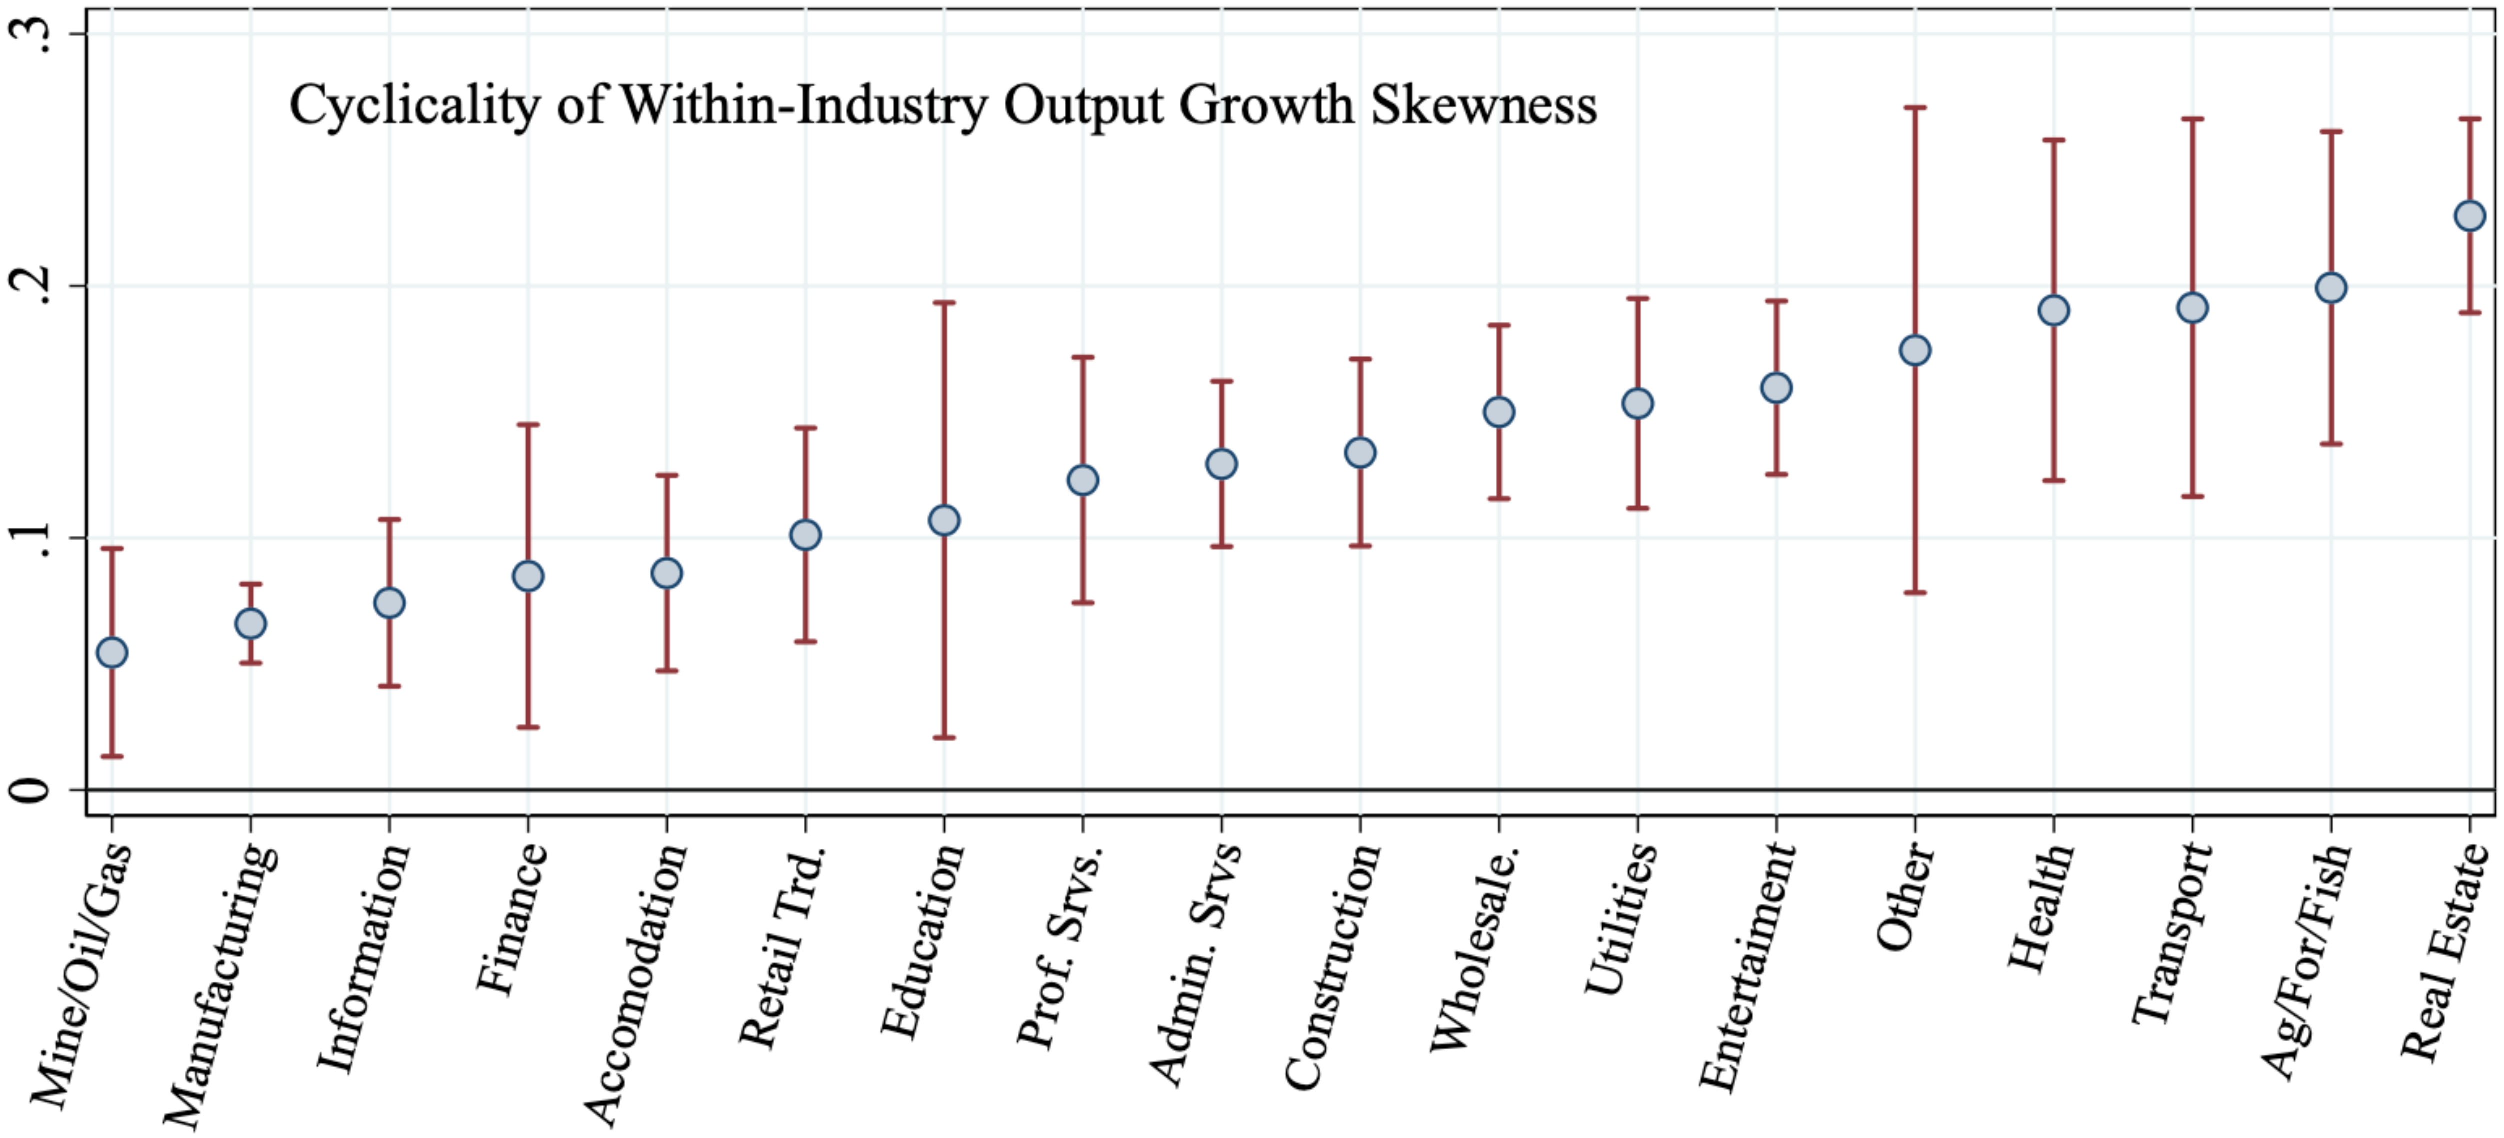

Notes: This figure shows the coefficients and 95% confidence intervals for within-industry regressions of the cross-sectional Kelley skewness on the average growth of sales for a sample of publicly traded firms from Compustat. Each industry regression includes a linear trend.

Appendix Figure A4: The Skewness of Employment Growth is Procyclical Within US Industries

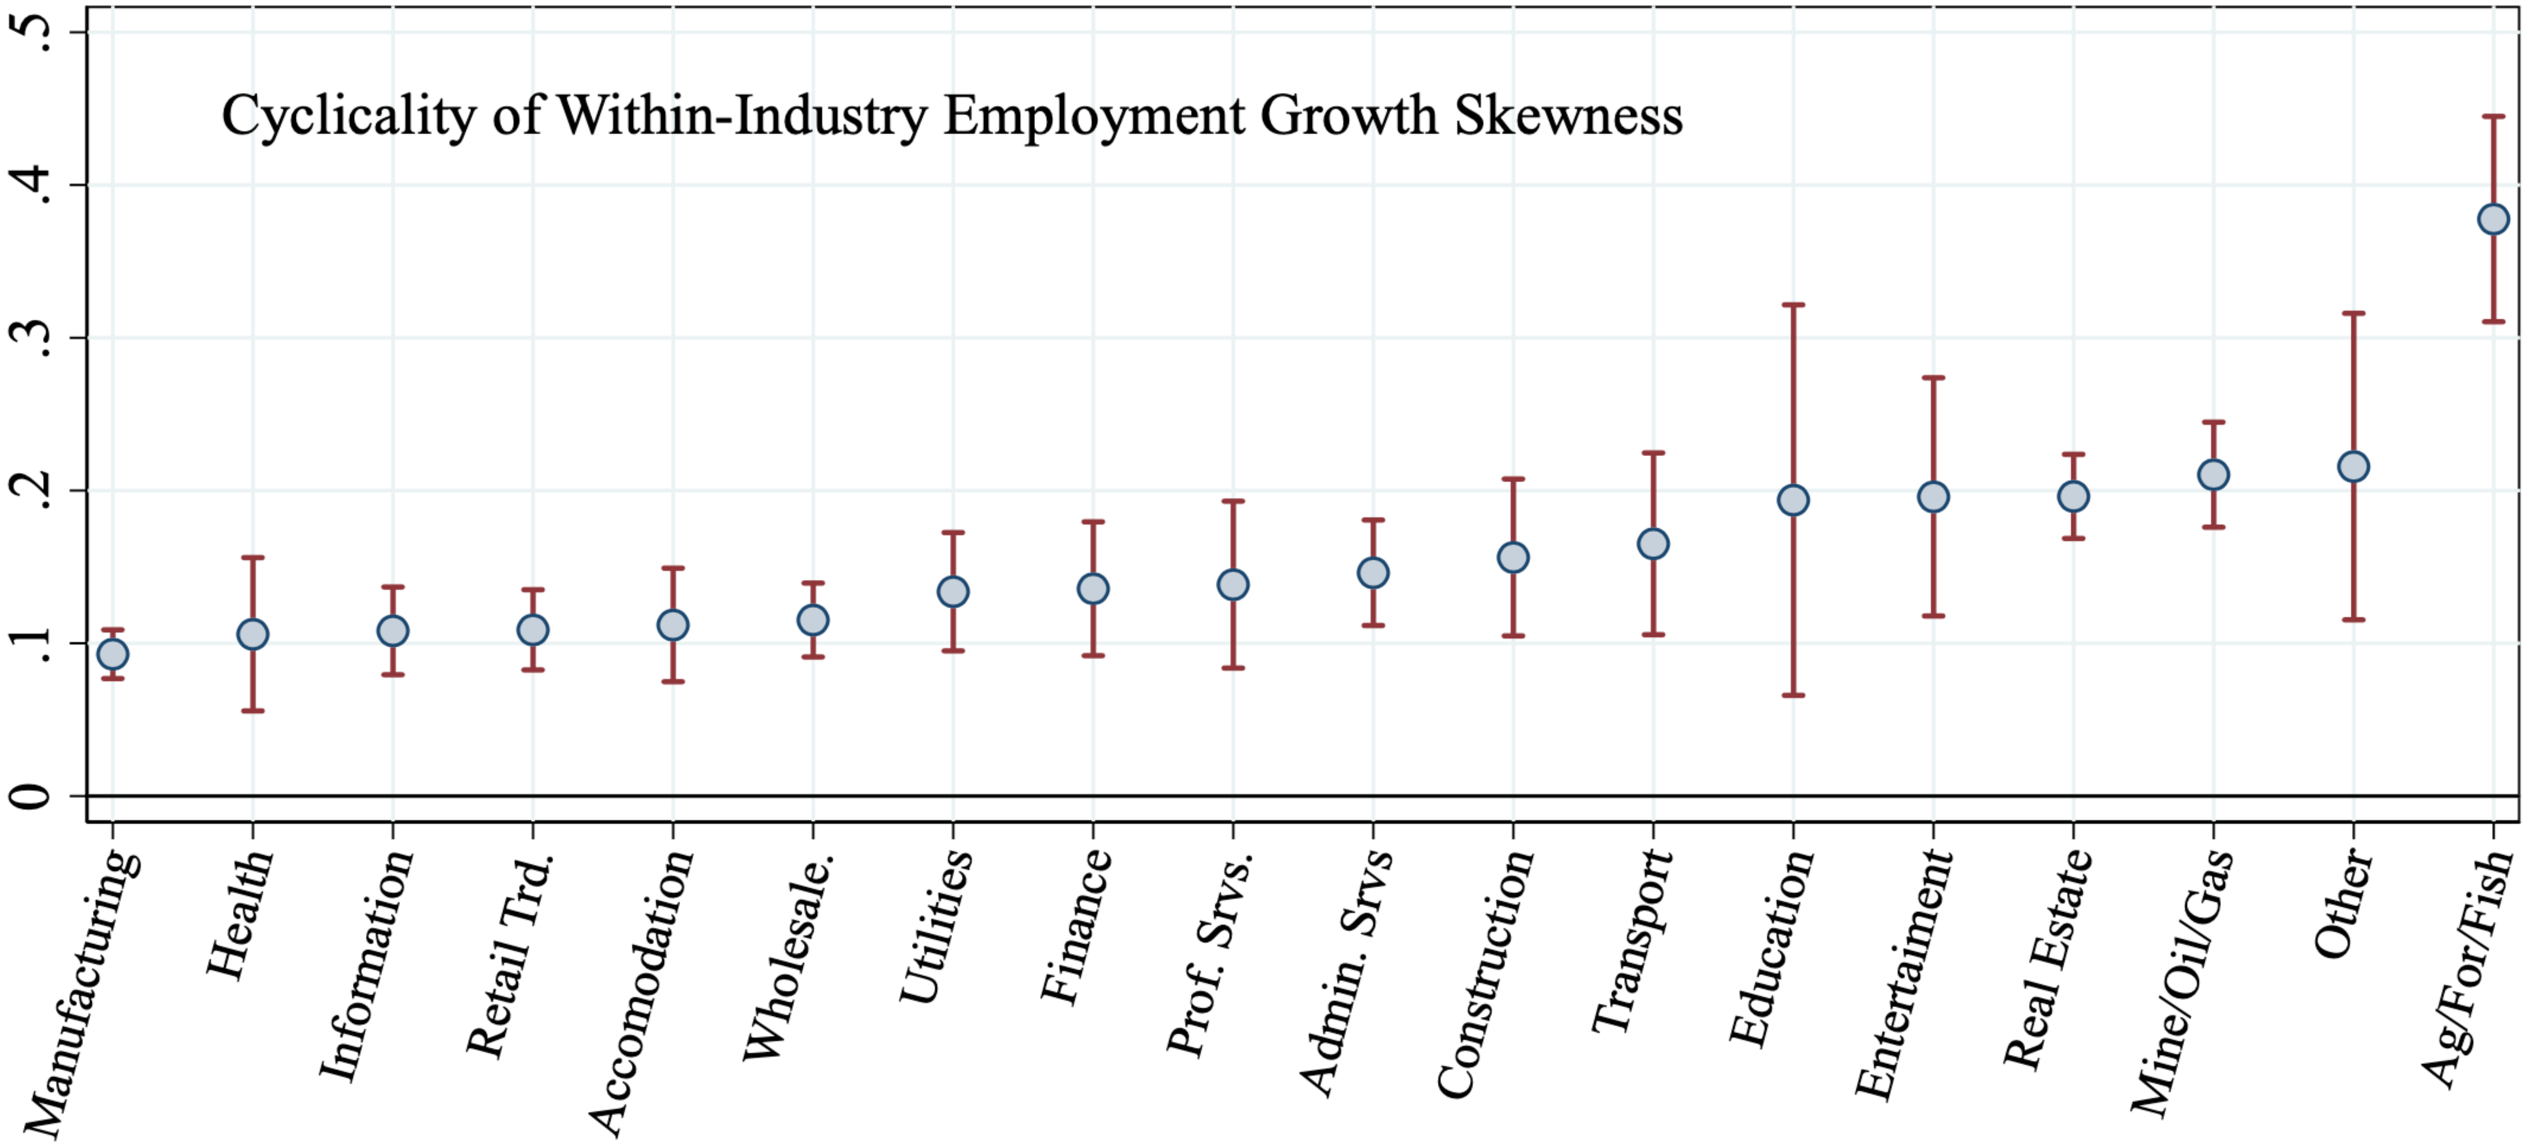

Notes: This figure shows the coefficients and 95% confidence intervals for within-industry regressions of the cross-sectional Kelley skewness on the average growth of employment for a sample of publicly traded firms from Compustat. Each industry regression includes a linear trend.

Figure A5: Cross-sectional Growth after Mean + Uncertainty shock under Different Adjustment Costs

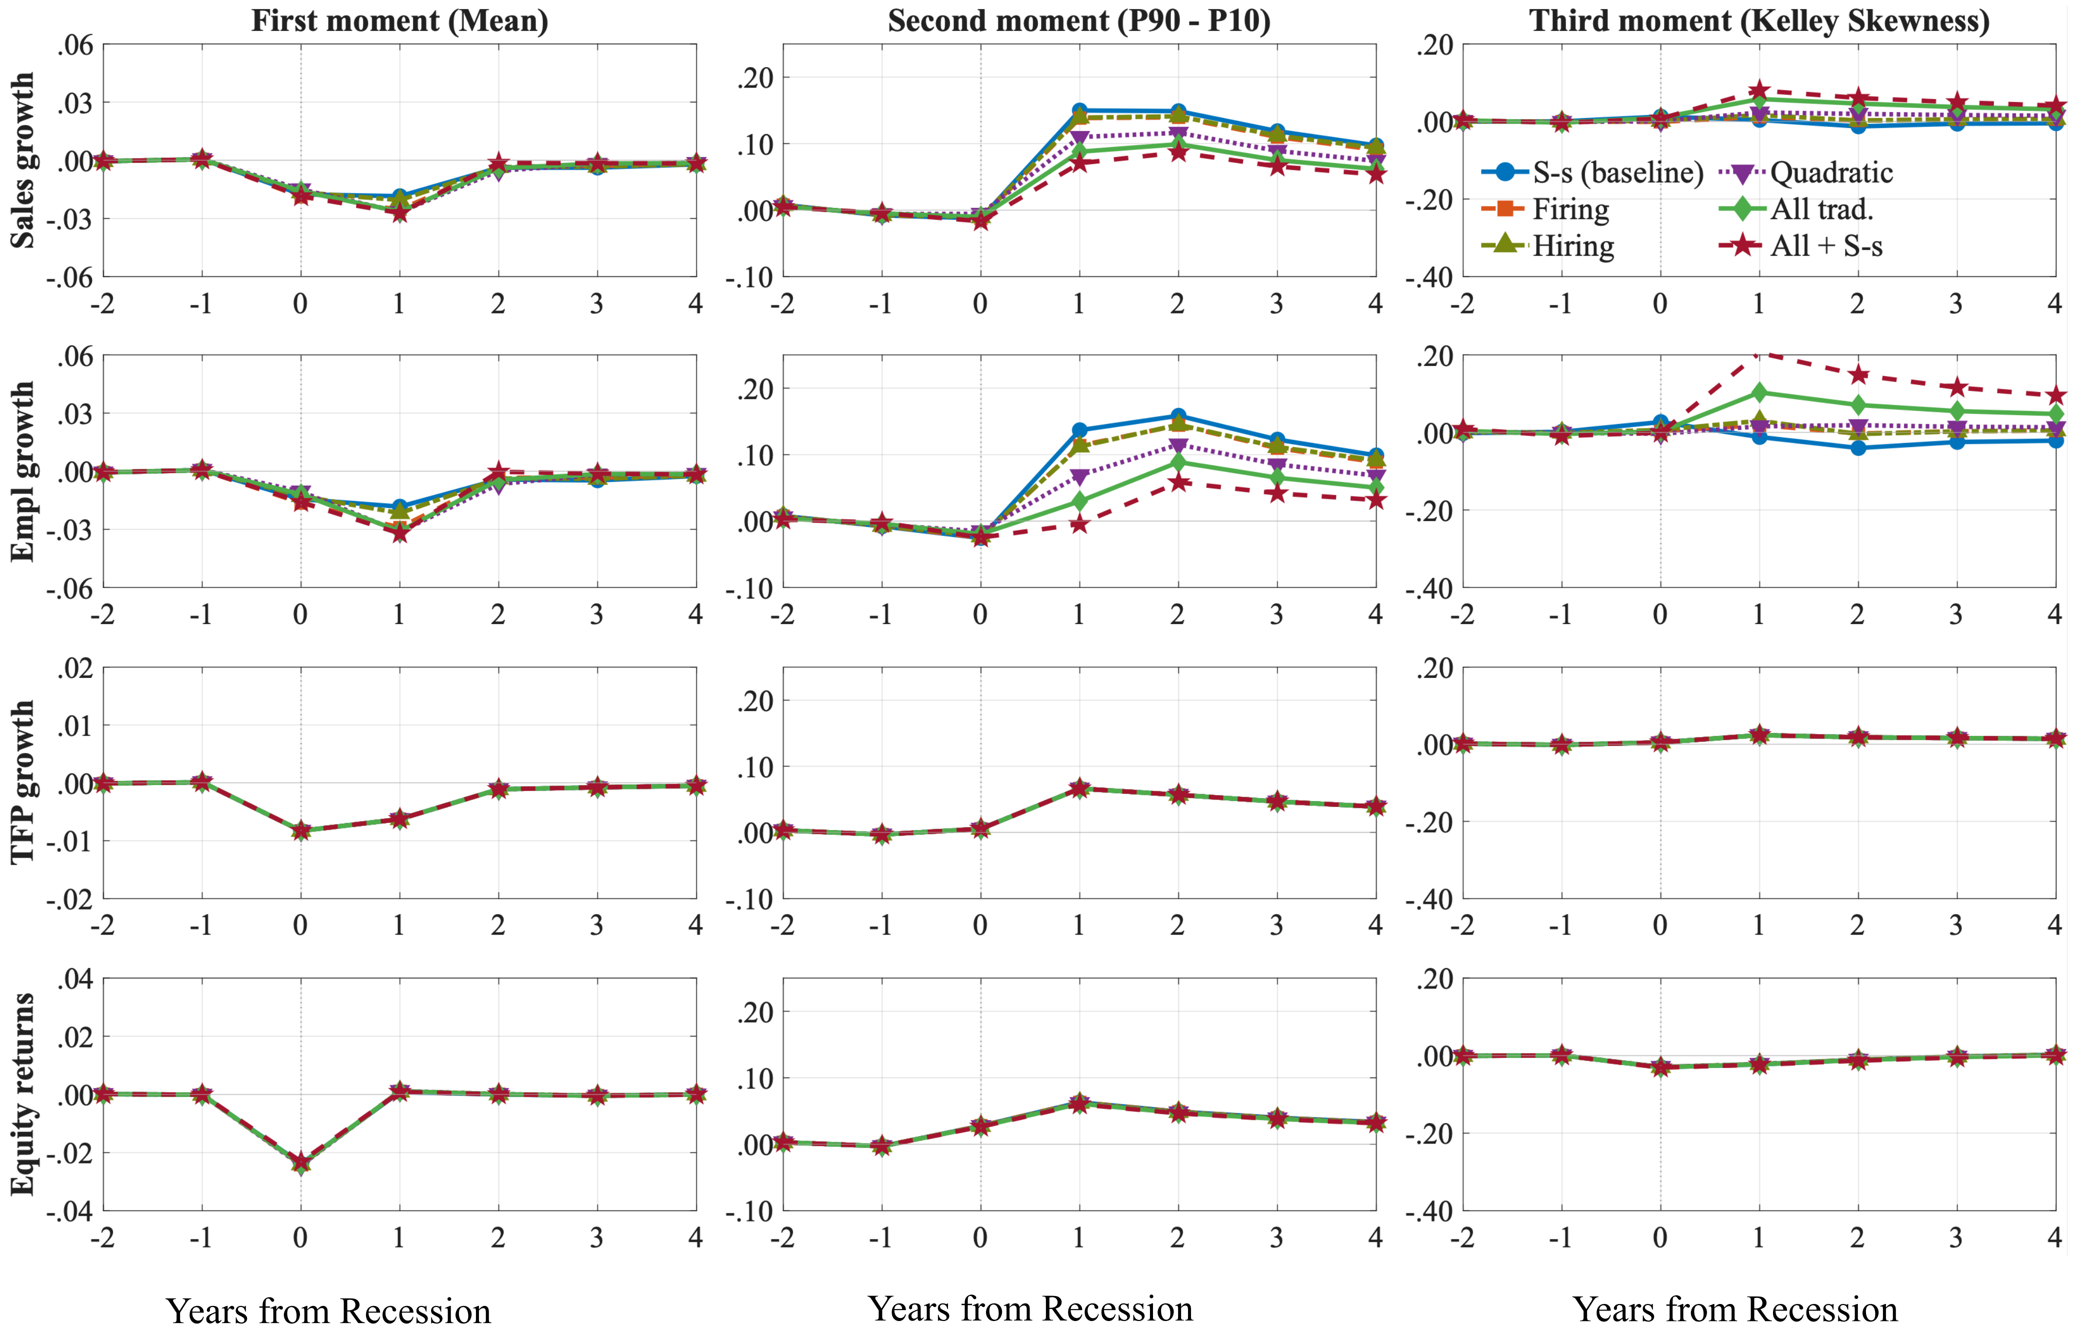

Notes: Cross-sectional moments after a mean + uncertainty shock based on a General Equilibrium simulation of 6,000 firms over 40,000 quarters. Model is solved separately under six adjustment cost configurations. Each panel plots deviations from the pre-recession average (years -2 and -1). Rows correspond to four variables: annual sales growth, employment growth, TFP growth (log(Y/L)), and equity returns. Columns show the first moment (mean), second moment (P90-P10), and third moment (Kelley skewness). Six lines per panel correspond to: S-s disruption only (baseline), firing cost only, hiring cost only, quadratic cost only, these three costs combined, and a final version that includes all costs (All +S-s).

Figure A6: The GDP Response to a Third Moment Shock for Different Values of Rn

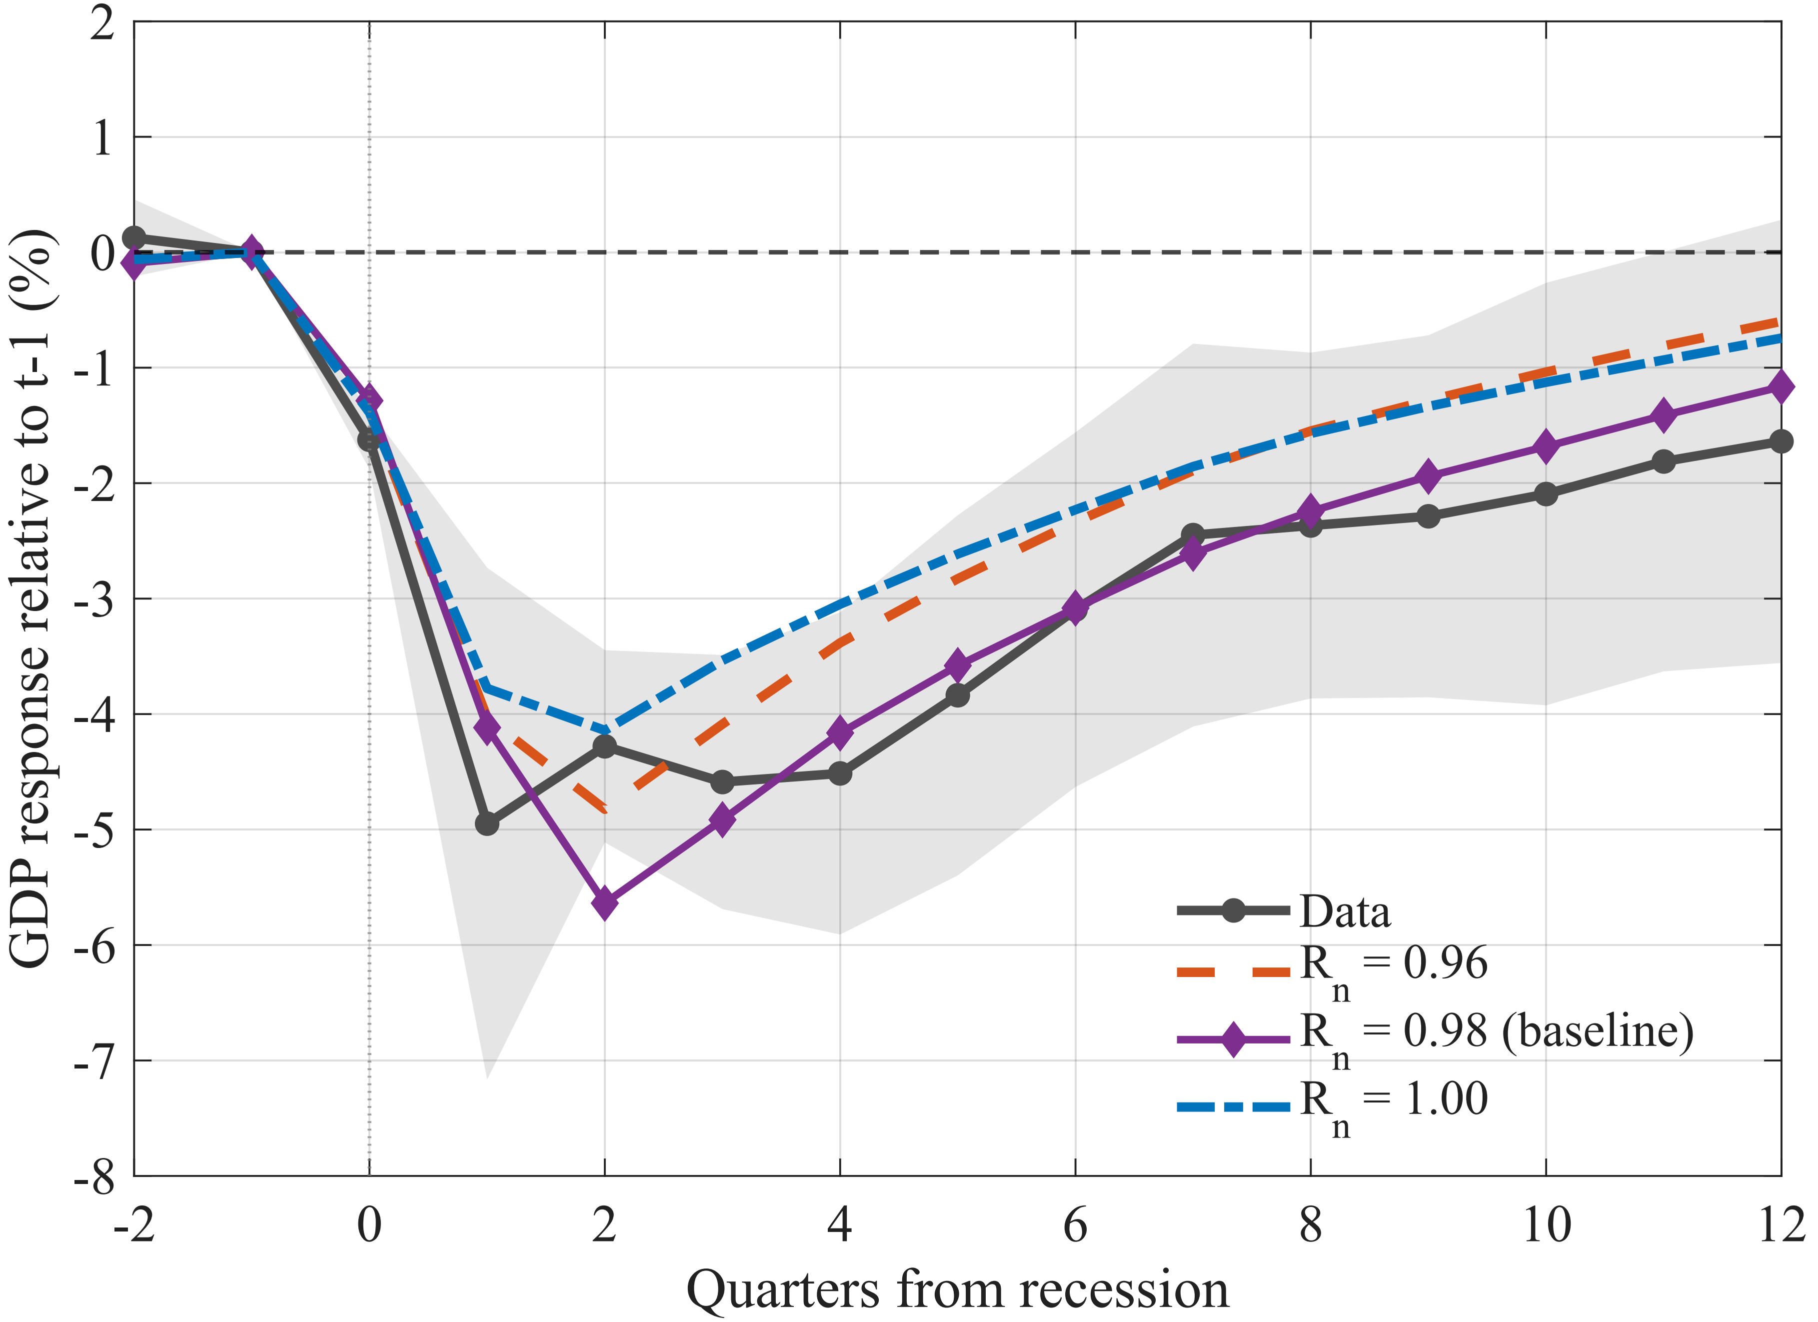

Notes: Jorda (2005) local projections estimated on simulated aggregate time series from a panel of 6,000 firms over 40,000 quarters. The dependent variable is the cumulative log-change in aggregate GDP from quarter t-1 to t+h. The regressor of interest is a recession-onset indicator (transition from expansion to recession regime). Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey-West HAC with bandwidth max(|h|+1, 34). Gray circles with shaded bands show the empirical LP estimated on U.S. FRED data (real GDP), 1970-present sample. The three lines correspond to the Mean+Skewness specification under different values of the cash carrying-cost parameter: κ=0.95 (Rn=0.96), κ=0.97 (Rn=0.98, baseline), and κ=0.99 (Rn=1.00). All other parameters are held at their baseline values.

Figure A7: The GDP Response to a Third Moment Shock for Different Assumptions

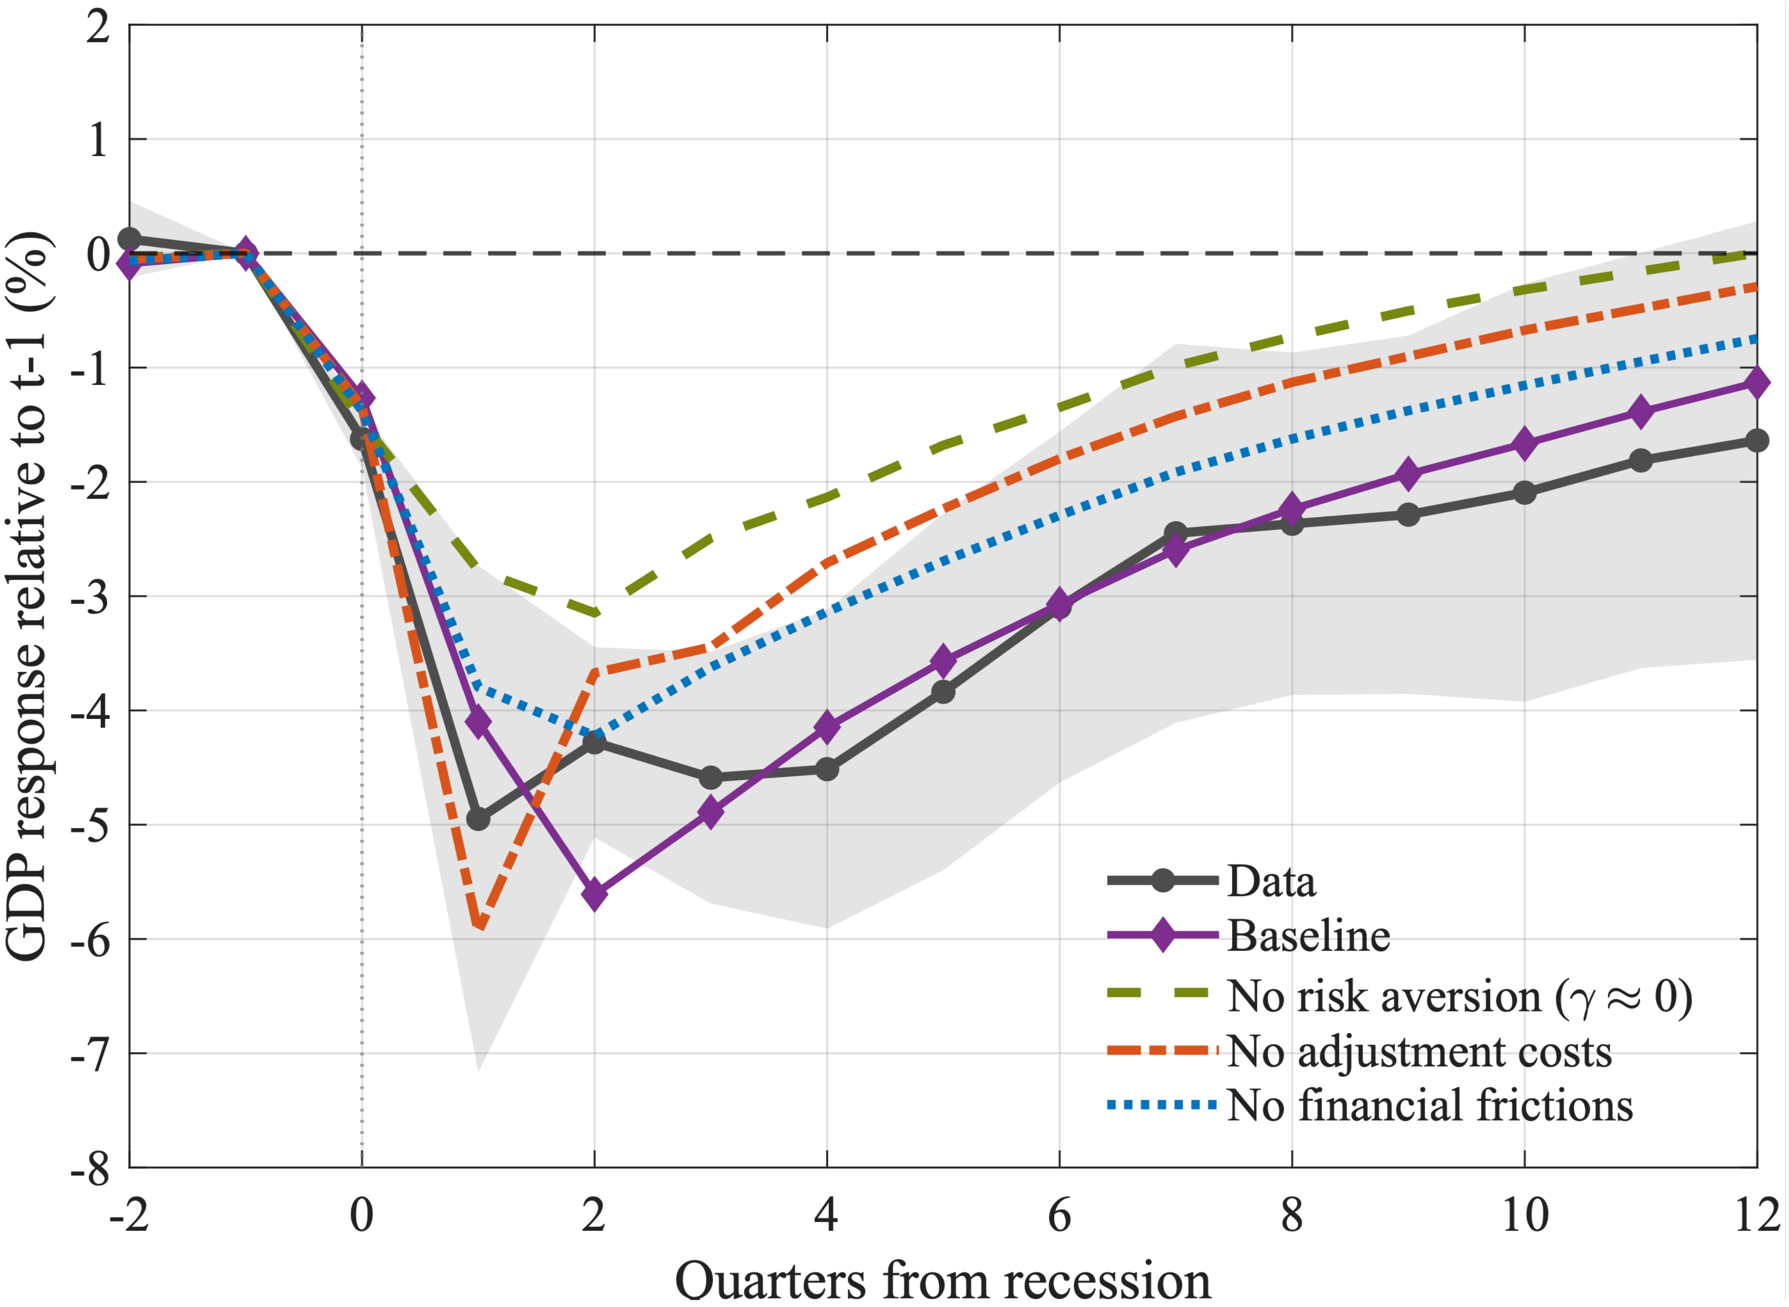

Notes: Jorda (2005) local projections estimated on simulated aggregate time series from a panel of 6,000 firms over 40,000 quarters. The dependent variable is the cumulative log-change in aggregate GDP from quarter t-1 to t+h. The regressor of interest is a recession-onset indicator (transition from expansion to recession regime). Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey-West HAC with bandwidth max(|h|+1, 34). Gray circles with shaded bands show the empirical LP estimated on U.S. FRED data (real GDP), 1970-present sample.

Figure A8: Response of Cash to First, Second, and Third Moment Shocks

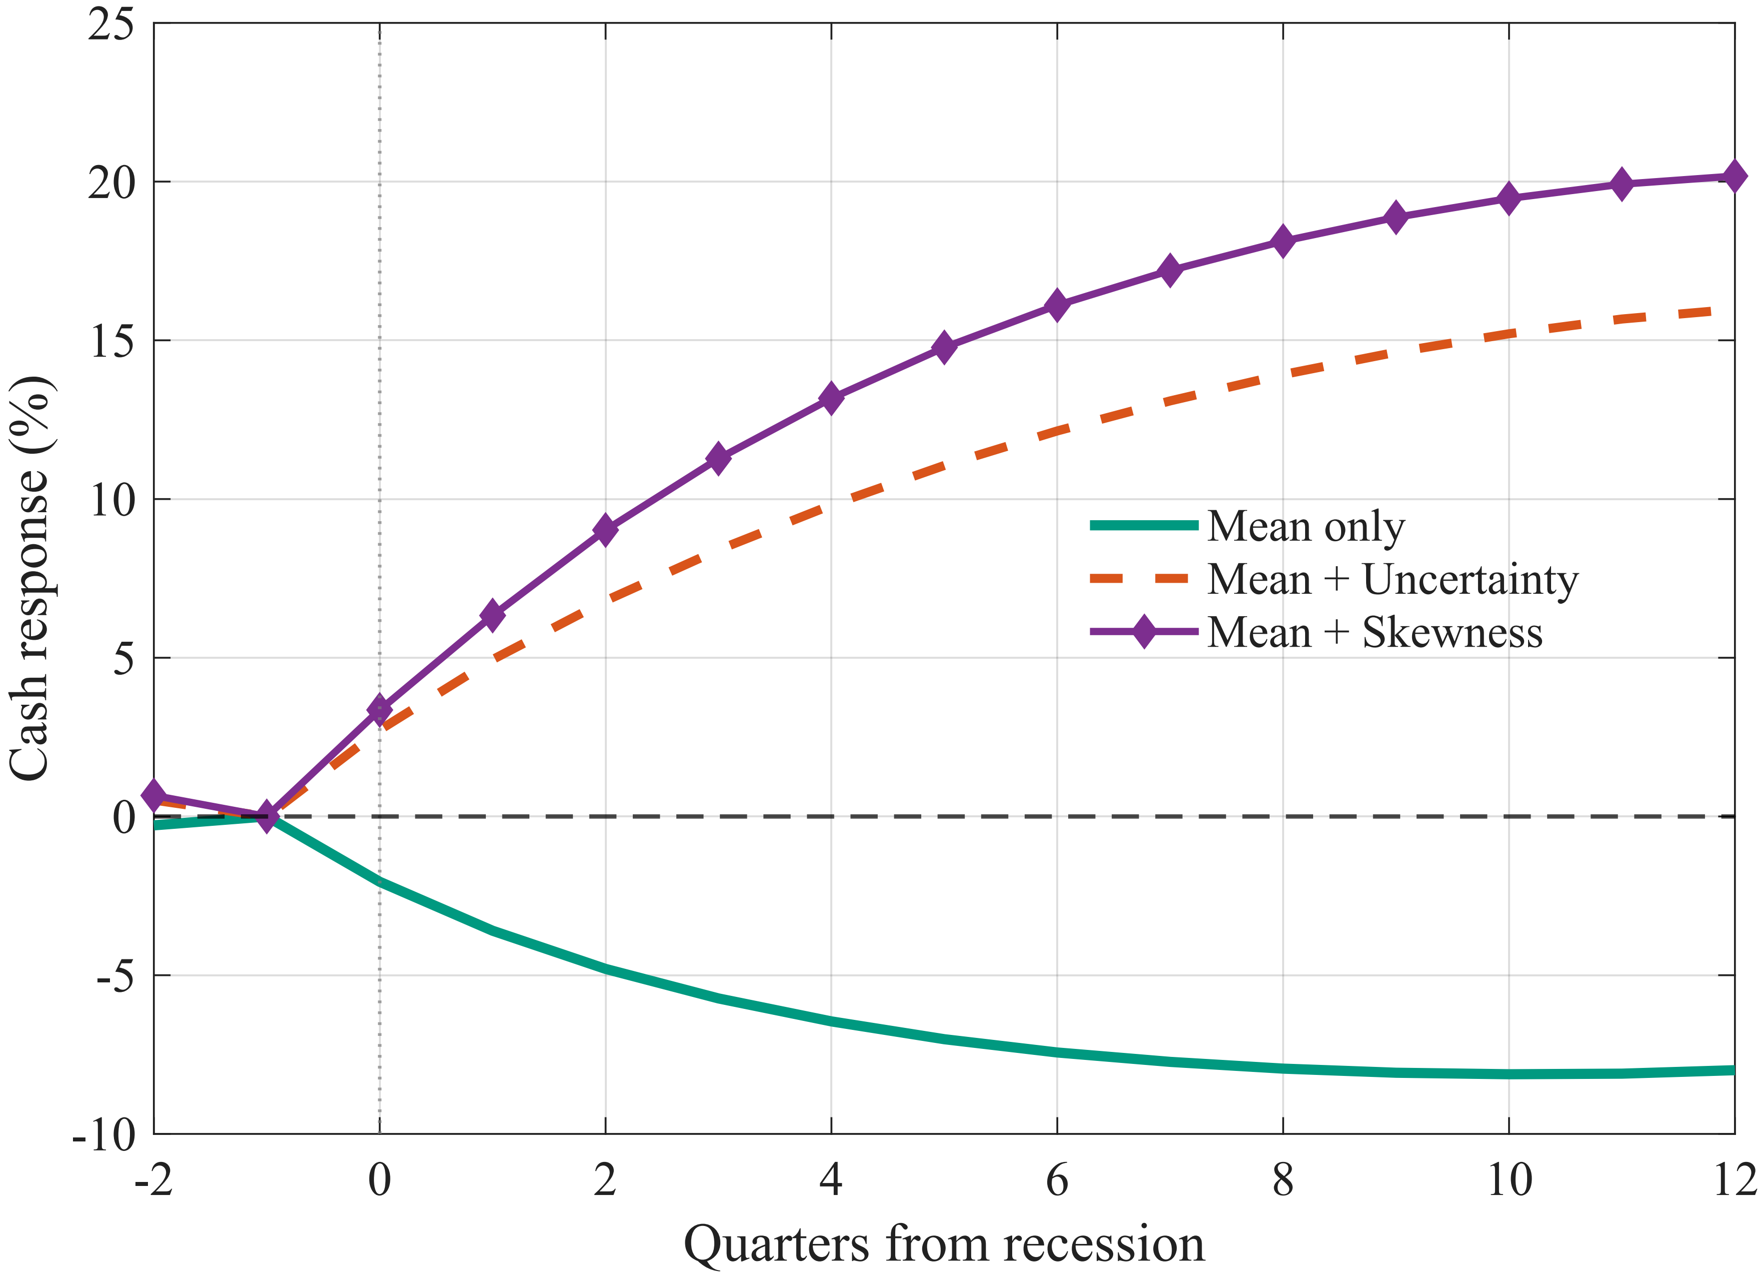

Notes: Jorda (2005) local projections estimated on simulated aggregate time series from a panel of 6,000 firms over 40,000 quarters. The dependent variable is the cumulative log-change in entrepreneur's cash holdings from quarter t-1 to t+h. The regressor of interest is a recession-onset indicator (transition from expansion to recession regime). Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey-West HAC with bandwidth max(|h|+1, 34). The three lines correspond to a Mean + Uncertainty shock, a Mean + Skewness shock, and a Mean-only shock with a decline of 1.5% in TFP.

Table A1: Data Sources and Sample Characteristics

Variable		Compustat	LBD	Census LBD-Rev	ASM/CM	BvD Osiris	BvD Amadeus
Employment	Mean	9,364	23	-	89	6,257	25
	P10	36	1	-	12	35	1
	P50	1,078	4	-	70	631	4
	P90	17,900	24	-	334	10,650	29
	P99	145,000	177	-	-	108,256	302
Sales (000s)	Mean	2,912	-	5,059	1,100	1.08	9.14
	P10	8	-	75	0.0	0.14	0.08
	P50	325	-	426	6.05	0.72	0.44
	P90	5,384	-	3,425	229	1.40	6.86
	P99	49,686	-	36,940	-	16.90	9.89
Frequency		Annual	Annual	Annual	Annual	Annual	Annual
Period		1970-2024	1978-2022	1998-2022	1976-2021	1991-2019	1996-2019
Obs. (M)		0.23	4.6	4.5	0.25	0.60	39.7
Unit. of Obs.		Firm	Firm	Firm	Estab.	Firm	Firm
Firm Type		Pub.	Pub./Priv.	Pub./Priv.	Pub./Priv.	Pub.	Pub./Priv.
Countries		US	US	US	US	Multiple	Multiple
Sectors		All	Non Farm	Non Farm	Manuf.	All	All

Notes: Table shows the list of datasets and time-frames used in the analysis. Sample statistics correspond to 2010 for comparability. All monetary values are expressed in US dollars of 2010. We omit data from Global Compustat since it does not contain information on employment or sales. LBD sample statistics are aggregated at the firm-level. Total observations correspond to all firm-year observations across all years in the sample with valid observations of the variable of interest (sales for Orbis and Compustat, employment for the LBD, and revenue for LBD-Rev). ASM results are calculated using sample weights. The 99th percentile of establishments sales and employment are not reported to avoid disclosure of sensitive information.

A Data Description

This appendix describes the data sources, the sample selection, and the calculations of the moments we use for our empirical analysis. In Section A.1 we describe the firm- and establishment-level data for the United States obtained from the Census Bureau's Longitudinal Business Database (LBD). Section A.2 describes our sample of Compustat firms. For our cross-country comparisons, we use firm-level data available in the Moody's Osiris database and Global Compustat which we describe in Section A.3. Finally, Section A.4 describes our sample and TFP estimation for our sample of firms from the Amadeus dataset and the establishment-level data for the US Census of Manufacturing and the Annual Survey of Manufactures.

A.1 United States: Longitudinal Business Database

We construct measures of employment growth at the firm-level using the Census Bureau's Longitudinal Business Database (LBD). The LBD covers the universe of establishments in the non-farm private sector in the United States from 1976 to 2021. It provides detailed establishment and firm-level information on employment, payroll, establishment location, firm age, industry, legal form of organization, and others. Crucially, firm and establishment identifiers in the LBD allow us to construct measures of employment growth at different time horizons. From the LBD, we select a sample of establishments that, in a given year, have non-negative, non-missing employment and payroll and have valid industry data. We then sum up the employment within the same firm across all establishments to construct an annual measure of employment. Starting in 1998, the LBD also contains information on firm-level revenue (LBDR). The coverage from the LBD to the LBDR is not complete. There are at least 30 percent of firms that have employment information (in LBD) but do not have revenue information (LBDR). A significant fraction of these firms (about 25%) are in the Other Services (except Public Administration) sector (NAICS 81). Nevertheless, the cross-sectional moments for the employment growth distribution do not change when we focus on the sample of firms that have employment and revenue data (results available in Census). See Haltiwanger *et al.* (2016) for additional details on Revenue LBD data. We deflate nominal revenue measures to 2012 prices using the CPI (FRED series CPIAUCSL).

We measure the growth rate of employment of firm j in period t as the log-difference between periods t and $t+1$, $g^e_{j,t} = \log E_{j,t+1} - \log E_{j,t}$. In order to capture the entry and exit of firms, we replace by a 0 the level of employment in the period before the first establishment is observed for the first time and the period after the last establishment of the firm is observed for the last time in the sample. We then calculate the growth rate of employment using the arc-percentage change between periods t and $t+1$ which is given by $g^{arc}_{j,t} = \frac{E_{j,t+1}-E_{j,t}}{0.5\times(E_{j,t+1}+E_{j,t})}$. We take the arc-percent change as our baseline measure. We proceed in the same way for revenue growth.

To calculate the Kelley skewness we require the computation of specific percentiles of the distribution of employment growth. Notice that a percentile provides information about *one* particular firm, which violates the disclosure criteria imposed by the US Census Bureau. Hence, to avoid the disclosure of any sensitive information, we calculate the *pth* percentile of the employment growth distribution as the employment-weighted average on a band of $+1$ and -1 percent around percentile *pth*. For instance, the 90th percentile of the distribution is the weighted average of the employment growth across all observations between the 89th and 91st percentiles of the distribution, both ends included. We proceed in the same way to construct the 10th and 50th percentiles of the distribution and use these values to calculate the Kelley skewness, the 90th-to-10th log percentiles differential and the rest of the measures of dispersion. The massive sample size of the LBD (more than 4.6 million observations per year on average) ensures that the sample used to calculate each of the percentiles is large enough to have an accurate approximation to the actual quantiles of the distribution. All moments of the employment growth distribution are weighted by the average employment of the firm (or establishment in the case of establishment level results) between periods t and $t+k$, that is, $\overline{E}_{j,t} = 0.5 \times (E_{j,t+k} + E_{j,t})$. Similarly, we weight revenue growth measures by the average real revenue between periods t and $t+k$.

We also use the LBD to compare the empirical distribution of employment and revenue growth between recession and expansion years using a kernel density estimation. The sample selection is the same as that used in the rest of our results with a few exceptions. First, the densities are calculated using a pooled sample of firms for recession years (the 2001 recession, the Great Recession, and the COVID pandemic) or expansion years (2003 to 2006 and 2010 to 2019). Second, we select a sample of firms that have at least 20 workers or more in year t.

A.2 United States: Compustat and CRSP

We construct time series of the cross-sectional dispersion and skewness of the sales growth distribution, the employment growth distribution, the stock returns distribution, and others using data of publicly traded firms from Compustat and CRSP accessed through the Wharton Research Data Services (WRDS).

For our results at the annual frequency, we obtain firm-level data from 1970 to 2023. We drop all observations with negative sales (Compustat variable sale), duplicated entries, and firms incorporated outside the United States (Compustat variable fic different from "USA"). We also drop all observations that do not have a SIC classification or with a classification above 90. We deflate nominal variables using CPI (FRED series CPIAUCSL) and we calculate the growth rate of sales and employment (Compustat variable emp) as the log change between year t and $t+k$ with $k \in \{1, 3, 5\}$. Our main sample considers firms with at least 10 years of data (not necessarily continuous), but our results remain robust if we drop this restriction or if we consider firms with at least 25 years of data. When accounting for entry and exit of firms using the arc-percentage change, for each firm we add an observation upon entry (equal to 2) and one additional observation upon exit (equal to -2) under the assumption that before and after exit, the firm has a value of sales or employment equal to 0. We consider entry firms as newly listed firms, while exiting firms are those delisted in a particular period, independent of the reason (M&A, bankruptcy, or any other).

To measure the degree of financial constraints, we use two measures: i) lagged average leverage and ii) the lagged average interest coverage ratio. The first is calculated as the ratio between short- and long-term debt (sum of DLTT and DLC) and total value of assets. We then calculate the lagged level of leverage as the average value of the leverage between periods $t-1$ and $t-2$. Finally, we winsorise this measure at the value of 2. To calculate the interest coverage ratio we take EBIT divided by total interest expenses (TIE). If TIE is equal to 0, then the coverage ratio is set to a large number (coverage ratio of infinity). We then calculate the lagged level coverage ratio as the average value between periods $t-1$ and $t-2$.

For our results based on quarterly data, we begin by retrieving firm-level data of net sales and stock prices at the annual and quarterly frequency, and employment at the annual frequency, from 1964q1 to 2023q4. The raw dataset of sales (Compustat variable saleq) and stock prices (Compustat variable prccq) contains more than 1.7 million quarter-firm observations with an average of approximately 4,660 firms per quarter. We drop all observations with negative sales, duplicated observations, and firms incorporated outside the United States (Compustat variable fic different from "USA"). We also drop all observations that do not have a SIC classification or with a classification above 90. Then, we deflate nominal sales by the CPI (FRED series CPIAUCSL), and we calculate the growth rate of sales as the log-difference and the arc percentage change between quarter t and $t+k$ with $k \in \{4, 12, 20\}$. This leaves us with around 1 million sales growth (log-difference) observations. For our main results, we consider firms with at least 10 years of data on quarterly sales (40 quarters, not necessarily continuous).

A.3 Cross-Country: Osiris-Amadeus and Global Compustat

Cross-country firm-level panel data on sales and employment come from the Bureau van Dijk's Osiris database.[22] Osiris is a database of listed public companies, commodity-producing firms, banks, and insurance companies from over 190 countries. The combined industrial company dataset which we use in our analysis contains financial information for up to 20 years and 80,000 companies.

The raw dataset contains 977,412 country/firm/year observations from 1982 to 2022. We drop all observations with missing or negative sales, all duplicated entries, and all firms with missing NAICS classification. We transform all observations into US dollars using the exchange rate reported in the same database. Then, we deflate nominal sales using US annual CPI and calculate the growth rate of real sales as the log change and arc percentage change between years t and $t+k$ with $k \in \{1, 3\}$. This leaves us with 748,574 observations (log change of sales). We further restrict the sample to firms with more than 10 years of data; country/year cells with more than 100 observations; countries with more than 10 years of data; and years with more than 5 countries. This sample selection reduces the dataset to an unbalanced panel of 678,563 observations in 50 countries between 1989 and 2018. We complement this data with real GDP in US dollars from the World Bank's World Development Indicators database.

The cross-country data on daily stock prices come from the Global Compustat database (GCSTAT), which provides standardized information on publicly traded firms for several countries at annual, quarterly, and daily frequencies. The raw data contain firm-level observations of daily stock prices between 1985 and 2024. We drop all duplicated observations and drop all firms with fewer than 2000 observations (firms with approximately 10 years of data). Then we calculate daily price returns as the log-difference of the stock price between two consecutive trading days. We apply a similar sample selection, keeping firms with at least 10 years of daily price data. The total sample contains an unbalanced panel of countries from 1985 to 2024 from which we drop all country quarters with fewer than 100 firms.

[22] See Kalemli-Özcan *et al.* (2024) for additional details on the Orbis dataset.

The final data contain a total of 51 countries from 1985 to 2024. Then, within each quarter, we calculate the cross-sectional moments of the daily stock price distribution. We complement this dataset with per capita GDP growth from the World Bank's World Development Indicators and quarterly GDP growth from the OECD Stats. Table A.1 shows the list of countries available in our dataset and the data available for each country.

A.4 TFP Estimation

A.4.1 Cross-Country: Amadeus

In this appendix, we describe in detail the construction of our measure of firm-level TFP using data from BvD Amadeus. See Kalemli-Özcan *et al.* (2024) for details about constructing a representative dataset using BvD Amadeus data. We consider a set of countries, namely, Germany, Spain, Finland, France, Italy, Norway, Poland, Portugal, Sweden, and Ukraine, for which firm-level information is available for enough industries and sectors. For each country in the sample, we retrieve firm-level panel data from Amadeus through WRDS. Our data contain a large range of firms, from small to very large firms (V+L+M+S: plus Small Companies dataset), both publicly traded and privately held. The main variables we use in our analysis are the following (Amadeus names of variables in parentheses):

- Sales (TURN),
- Operating revenues (OPRE),
- Employment (EMPL),
- Cost of Employees (STAF),
- Cost of Material (MATE),
- Total Fixed Assets (FIAS),
- Industry (NAICS and SIC codes),

In order to estimate firm-level productivity for a large number of firms within each country, we perform a simple sample selection. For each country, we drop duplicates, observations without information on industry (NAICS), and firms with discrepancies between the country identifier and the firm identifier (INDR; note that the first two characters in the firm identifier in Amadeus refer to the country.) We also drop all observations with missing, zero, or negative values in any of the following variables: OPRE, MATE, FIAS, and STAF. We also drop all observations with zero or negative value of VA = OPRE − MATE which is our measure of value added. We deflate all monetary values by the country-specific CPI (obtained from the World Bank).

A.4.2 Estimating TFP

The literature has considered several different methods to measure TFP at the firm-level (Syverson, 2011), and in this section we consider a few standard methods. If we assume that the firm's production function is Cobb-Douglas, we can estimate the firm-level productivity, $z_{i,j,t}$, as the residual of the following equation,

$$\log y_{i,j,k,t} = \alpha_K \log K_{i,j,k,t} + \alpha_L \log E_{i,j,k,t} + z_{i,j,k,t}, \tag{10}$$

Table A.1 – Data Availability by Country

	Osiris		GCSTAT	Amadeus				Osiris		GCSTAT	Amadeus		
	Sales	Emp	Returns	Sales	Emp	TFP		Sales	Emp	Returns	Sales	Emp	TFP
ARG	x		x				ITA	x	x	x	x	x	x
AUS	x	x	x				JOR	x	x				
AUT			x				JPN	x	x	x			
BEL	x	x	x				KOR	x	x	x			
BGD	x						KWT	x					
BGR			x				LKA	x	x				
BMU	x	x					LTU			x			
BRA	x	x	x				LUX			x			
CAN	x	x	x				LVA			x			
CHE	x	x	x				MEX	x	x	x			
CHL	x	x	x				MYS	x	x				
CHN	x	x	x				NGA	x					
COL	x		x				NLD	x	x	x			
CYM	x	x					NOR	x	x	x	x	x	x
CZE			x				NZL	x		x			
DEU	x	x	x	x	x	x	PAK	x	x				
DNK	x	x	x				PER	x					
EGY	x						PHL	x	x	x			
ESP	x	x	x	x	x	x	POL	x	x	x	x	x	x
EST			x				PRT			x	x	x	x
FIN	x	x	x	x	x	x	ROU	x	x	x			
FRA	x	x	x	x	x	x	RUS	x	x	x			
GBR	x	x	x				SAU	x		x			
GRC	x	x	x				SGP	x	x	x			
HKG	x	x	x				SVK			x			
HRV			x				SVN			x			
HUN			x				SWE	x	x	x	x	x	x
IDN	x	x	x				THA	x	x	x			
IND	x	x	x				TUR	x	x	x			
IRL			x				UKR				x	x	x
IRN	x	x					VNM	x	x				
ISL			x				ZAF	x	x	x			
ISR	x	x	x										

Notes: The table shows data availability for each country (identified by its ISO3 code). An "x" indicates that the country appears in the final estimation sample for that source and variable after all sample restrictions (at least 10 years of data, at least 100 firms per country-year or country-quarter cell). Osiris provides annual sales and employment for 50 and 41 countries, respectively. Global Compustat (GCSTAT) provides daily stock prices used to compute quarterly return moments for 51 countries. Amadeus provides sales, employment, and revenue TFP for 10 European countries. US data sources (Census LBD, Compustat, CRSP, ASM) are omitted.

where $y_{i,j,k,t}$ is the value added of firm i, in industry j, in country k, in year t; $K_{i,j,k,t}$ is the deflated measure of fixed assets and $E_{i,j,k,t}$ is a measure of labor input (employees or wage bill).

We use four different methods to estimate $z_{i,j,t}$. The first method—which we use in our main empirical results—uses country-industry factor shares to estimate α_L and α_K. In particular, we calculate the total wage bill and total value added at the country-industry-year level. Industries are defined by two-digit NAICS. To ensure our measure of factor shares is calculated with enough firms, we restrict our estimates to years in which the country-industry cell contains more than one hundred observations and periods with more than five sectors within a country-year. We then obtain the labor share as

$$\alpha_{L,j,k,t} = \frac{\sum_{i \in I_{j,k,t}} w_{i,j,k,t}}{\sum_{i \in I_{j,k,t}} y_{i,j,k,t}},$$

where $I_{j,k,t}$ is the set of firms in the industry-country-year cell and $w_{i,j,k,t}$ is the cost of employees at the firm-level (STAF). Then, we calculate the capital share as $\alpha_{K,j,k,t} = 1 - \alpha_{L,j,k,t}$.[23] We then apply these factor shares in equation (10) to obtain our first measure of productivity as the difference between $\log y_{i,j,k,t}$ and

$$\left((1 - \alpha_{L,j,k,t}) \log K_{i,j,k,t} + \alpha_{L,j,k,t} \log E_{i,j,k,t}\right).$$

In the second method, we obtain $z_{i,j,t}$ as the residuals of a firm-level OLS panel regression. In order to control for differences in labor quality across firms, we use the wage bill (STAF) at the firm level as a measure of labor input. We then run an OLS panel regression to obtain $\hat{z}_{i,j,k,t}$ for each firm.

The third approach uses the methodology developed by Olley and Pakes (1996) to estimate $\hat{z}_{i,j,k,t}$. This method has stricter data requirements, and therefore we further restrict our within-country sample to firms with information about investment expenditure (change in the value of total fixed assets, FIAS), and firms with at least 5 years of data. Furthermore, because the data available in Amadeus were increasingly populated until 2005, we consider information only after that year. To obtain the Olley and Pakes (1996) estimates we use the Stata command OPREG as implemented by Yasar, Raciborski, and Poi (2008).

The fourth method abstracts from capital differences across firms and proxies a measure of labor productivity. In particular, we obtain labor productivity as the residual of the following equation estimated using OLS within each country

$$\log y_{i,j,k,t} = \tilde{\alpha}_L \log E_{i,j,k,t} + \mu_i + \tilde{z}_{i,j,k,t}. \tag{11}$$

Then, for each productivity measure, we estimate firm-level productivity shocks as the residual of the following OLS panel regression within each country

$$\hat{z}_{i,j,k,t} = \beta_{0,k} + \beta_{1,k} \hat{z}_{i,j,k,t-1} + \mu_i + \delta_t + \epsilon_{i,j,k,t}, \tag{12}$$

where μ_i and δ_t are firm and year fixed effects respectively. In order to reduce the impact of outliers that normally appear in micro data, for each country, we winsorize each measure

[23] In this calculation, we use the nominal values of value added and cost of employees.

of productivity shock at the top and bottom 1%. Additionally, we use the average of the real sales growth within a bin (defined by country, industry, or year) as a measure of business conditions. Then, for each measure of productivity shock, we calculate the average shock within a country-industry-year bin and different percentiles of the distribution. To further ensure our results are not driven by outliers at the country-industry level, after we have obtained these percentiles, we trim the measures of Kelley skewness and the average productivity shocks at the top and bottom 1% and we restrict our sample to country-industry-year bins with more than 100 firms. Our results, however, hold if we relax these conditions. The results across all these methods are remarkably consistent as shown in Figure A.1.

A.4.3 Controlling for Capacity Utilization

It is well known that failing to control for changes in capacity utilization can lead to wrong estimates on the cyclicality of firm and aggregate productivity (Basu, 1996). This might be even more important using firm-level data and when trying to measure higher-order moments. To control for changes in firm-level utilization that might bias our results, we residualize the log of measured productivity for different measures of utilization available in the ASM, namely, average number of hours for productive workers and energy utilization. In particular, we have used firm-level capacity utilization administrative data from the Census' Quarterly Survey of Plant Capacity Utilization (QPC) and the Annual Survey of Manufactures (ASM). We proceed in three steps which we describe below.

First, we use two questions from QPC to calculate what fraction of the total capacity is used by a plant in any given time. This measure of capacity utilization is only available for those plants in the QPC. Therefore, we merge the QPC data with the ASM which contains information on output and inputs, including energy usage and hours worked. With these variables, we construct three measures of plant capacity utilization: i) average number of hours per worker, ii) the ratio between real energy expenditure and real value of machinery and equipment, and iii) the ratio between total hours worked and the real value of machinery and equipment. Figure A.2 shows that all these measures are highly correlated with the QPC measure of capacity utilization, which we take as our benchmark. Hence, in our analysis, we use the ASM-based measures of capacity utilization given that they appear to be tightly correlated with QPC data (which is not available for our fuller dataset). This delivers three credible measures of capacity utilization for the entire ASM data sample.

Second, we check if the ASM-based measures of capacity utilization are procyclical, which is an important validation step. For this, we use plant-level data and correlate the growth rate of the three ASM-based measures of capacity utilization we described above with the growth rate of the real value of shipments (revenue growth). As expected, we find a positive and strong correlation between revenue growth and capacity utilization growth. These results are displayed in Figure A.3. Each figure shows mean sales growth at the industry level and the mean growth of capacity utilization for different measures controlling for year and industry fixed effects. The tight positive slope indicates that in periods in which the sales growth is above the industry average, utilization is also above its industry average.

Third, using our capacity utilization measures, we regress the log-TFP measure available in Census on these measures of capacity utilization using a rich set of polynomials to

Figure A.1 – Baseline: Skewness of TFP shocks is Procyclical

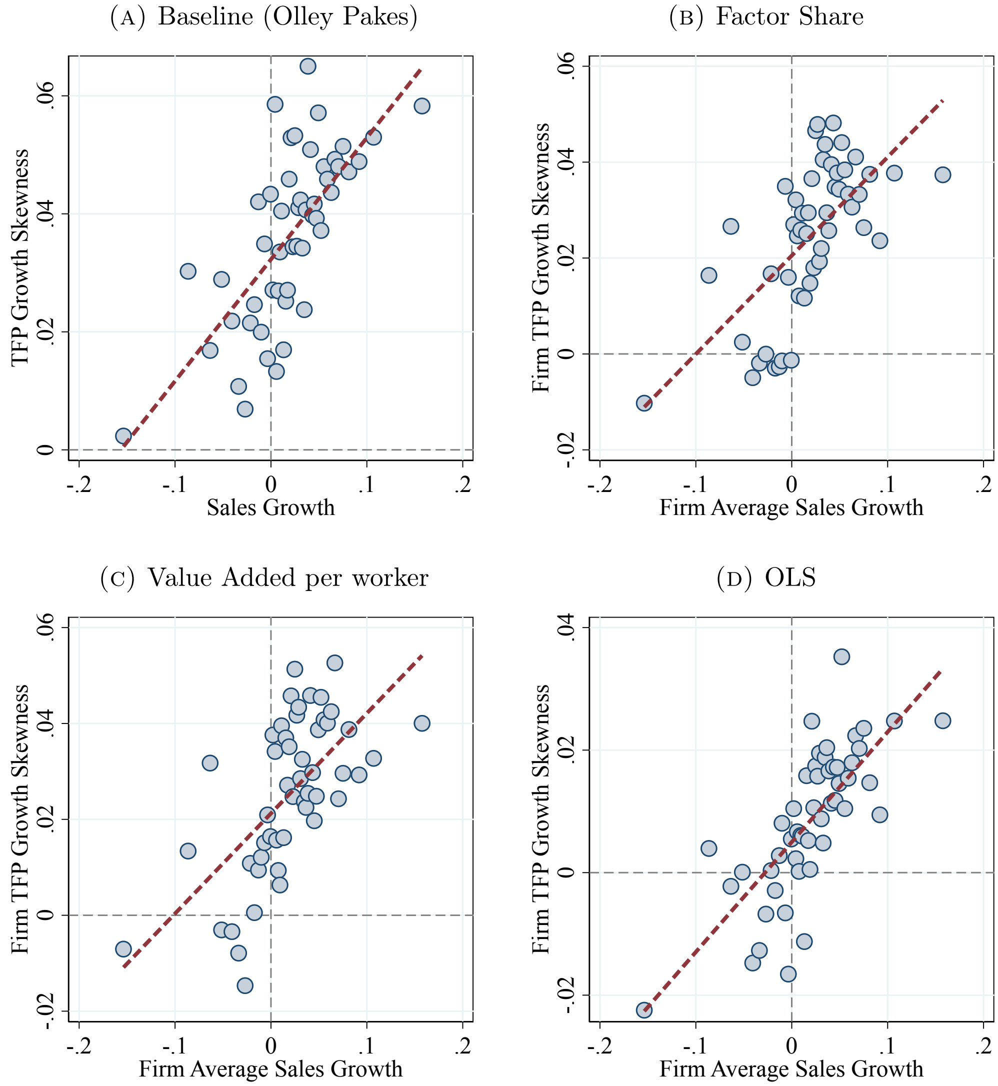

Notes: This figure shows the relation between the average firm sales growth within a year-country-industry and the skewness of firm productivity shocks within the same bins calculated using different methods. Binscatters control for year, country, and industry fixed effects. Industries defined as two-digit NAICS.

account for potential non-linearities in the relation between plant capacity utilization and productivity. This should cleanse any impact of capacity utilization from our measurement of TFP given this flexible functional form. In particular, we use multiple specifications which span a linear term for capacity utilization to a third-degree polynomial. We use the residuals of these regressions as our utilization-adjusted measure of TFP and we analyze whether the skewness of this measure is procyclical. To do so, we calculate the mean and Kelley skewness of TFP growth within a year/3-digit NAICS code and we regress these two measures.

Figure A.2 – Correlation of Capacity Utilization in Census' QPC and ASM

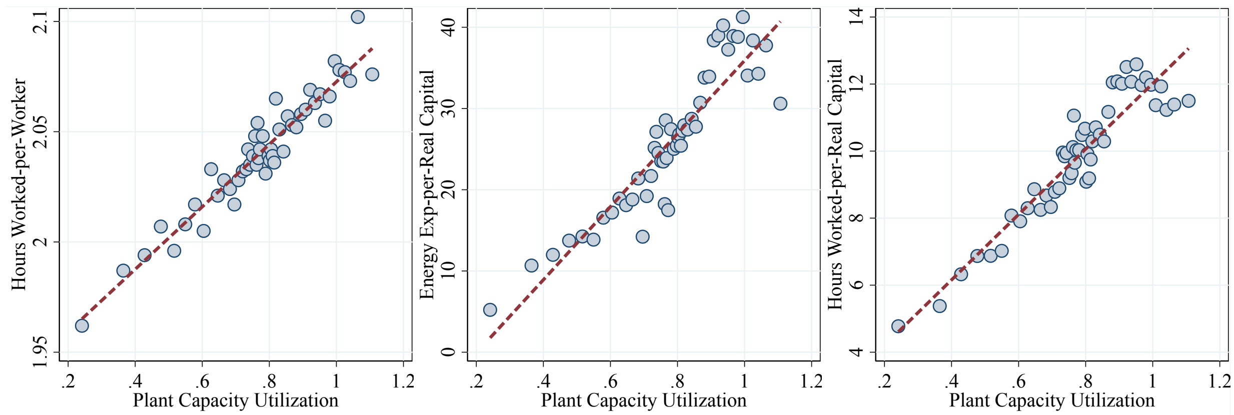

Notes: FigureA.2 shows industry-year level binscatters of the measure of plant capacity utilization from QPC (x-axis) and a proxy for capacity utilization from ASM (y-axis). Binscatters control for year and industry fixed effects. Industries are at the 3-digit NAICS level.

Figure A.3 – Capacity Utilization is Procyclical in ASM

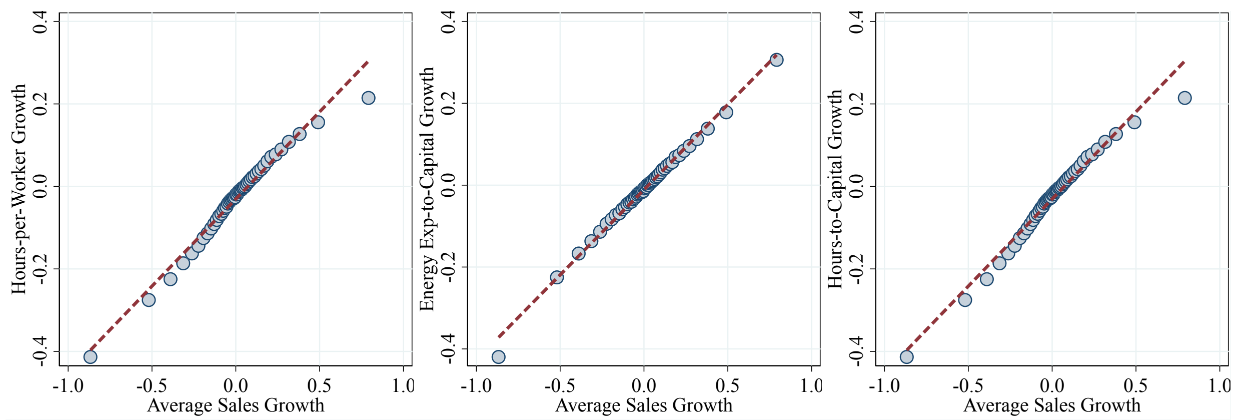

Notes: FigureA.3 shows industry-year level binscatters of the measure of plant capacity utilization growth (y-axis) and plant sales growth (log-change of total value of shipments) from ASM. Binscatters control for year and industry fixed effects. Industries are at the 3-digit NAICS level.

Table A.2 – Skewness of Firm TFP Growth is Procyclical

(a) TFP Growth

Specification	β	SE	*t-stat*	R^2
Baseline	1.158	0.061	18.75	0.131
Residual 1	1.164	0.062	18.75	0.129
Residual 2	1.164	0.062	18.74	0.129
Residual 3	1.167	0.062	18.78	0.129
Residual 4	1.168	0.062	18.8	0.129
Residual 5	1.164	0.062	18.91	0.133
Residual 6	1.151	0.062	18.68	0.130
Residual 7	1.168	0.062	18.8	0.129

(b) Sales Growth

Specification	β	SE	*t-stat*	R^2
Baseline	0.564	0.050	11.52	0.080
Residual 1	0.587	0.050	11.95	0.080
Residual 2	0.587	0.050	11.95	0.081
Residual 3	0.589	0.050	11.99	0.081
Residual 4	0.590	0.050	12.01	0.081
Residual 5	0.576	0.050	11.75	0.083
Residual 6	0.569	0.050	11.66	0.081
Residual 7	0.590	0.050	12.01	0.081

Notes: TableA.2 shows industry panel regressions (3-digit NAICS) of the Kelley skewness of TFP growth on the average TFP growth and average sales growth within an industry using data from the Census Annual Survey of Manufactures. All regressions are weighted by sample weights and include year and industry fixed effects. We use three measures of capacity utilization: Hours per worker, Energy per Capital, and Hours per Capital. The baseline is the correlation between TFP shock skewness and mean TFP shock or mean sales growth without any adjustment for capacity utilization. Residual 1 refers to log-TFP residualized using a model in which the three measures of capacity utilization enter linearly; Residual 2 adds a quadratic term on each measure, Residual 3 adds a cubic term, and Residual 4 adds a quartic term. Residual 5 uses a quartic polynomial only on Hours per worker, Residual 6 a quartic on Energy per Capital, and Residual 7 uses Hours per Capital.

Table A.2a shows the results of these regressions controlling for year and industry fixed effects. The first row (baseline) calculates the relation between TFP shock skewness and mean TFP shock or mean sales growth without any controls for capacity utilization. Residual 1 refers to log-TFP residualized using a model in which the three measures of capacity utilization enter linearly; Residual 2 adds a quadratic term on each measure, Residual 3 adds a cubic term, and Residual 4 adds a quartic term. Residual 5 uses a quartic polynomial only on Hours per worker, Residual 6 a quartic on Energy per Capital, and Residual 7 uses Hours per Capital.

We repeat these regressions using the average sales growth within an industry instead of the average TFP growth. The results, shown in Table A.2b, indicate a positive relation between the skewness of TFP growth and the average sales growth at the industry level. We conclude that, although capacity utilization is an important contributor to the observed changes in measured productivity, controlling for its fluctuations does not materially affect the procyclical nature of the skewness of the distribution of firm shocks.

B Computational Appendix

B.1 Value Function Iteration

The firm's value function is solved numerically by value function iteration (VFI) on a discretized state space $(s, z, l, n) \in \{E, R\} \times \mathcal{Z} \times \mathcal{L} \times \mathcal{N}$ taking the wage w as given. In general equilibrium, the wage becomes an additional aggregate state and the VFI is solved jointly at a grid of wage values (Section B.5). In both cases (the partial and general equilibrium), we have that the idiosyncratic productivity grid $\mathcal{Z}$ contains $N_z = 51$ equally spaced points in logs, spanning ± 3 unconditional standard deviations of $\log z$:

$$\log z \in \left[-3 \frac{\sigma_\eta}{\sqrt{1-\rho_z^2}}, \; +3 \frac{\sigma_\eta}{\sqrt{1-\rho_z^2}} \right],$$

where $\sigma_\eta = \sqrt{\sigma_s^2 + p_s(1-p_s)\mu_D^2}$ uses the regime with the widest innovation variance. The employment grid $\mathcal{L}$ contains $N_l = 100$ points, log-spaced between the frictionless optimal scale at the lowest and highest z grid points (with a small buffer). The cash grid $\mathcal{N}$ uses $N_n = 21$ coarse points (log-spaced in the positive region, plus a zero point).

Choice of next-period cash. Because the cash choice n' is continuous, we evaluate the Bellman equation on a finer grid of $N_n^{\text{fine}} = 101$ candidate n' values. For each candidate, the continuation value is obtained by linear interpolation along the cash dimension using precomputed bracket indices and weights on the coarse grid. This avoids the discretization error of restricting n' to the coarse grid while keeping storage at $N_n = 21$.

Choice of next-period employment. Employment l' is chosen from the grid $\mathcal{L}$. At each state (s, z, l, n), the algorithm evaluates two branches: (i) *adjust*, which optimizes jointly over all $l' \in \mathcal{L}$ and $n' \in \mathcal{N}^{\text{fine}}$, paying the adjustment costs in (3); and (ii) *no-adjust*, which sets $l' = (1-\delta_L)\, l$ (interpolated between adjacent grid points) and optimizes over n' only. The firm adjusts whenever the adjust branch yields higher value.

Convergence acceleration. Each Bellman iteration is followed by $N_H = 50$ Howard (policy iteration) steps that update the value function at the current policy without reoptimizing, which substantially accelerates convergence. The value function update is dampened with mixing parameter $\lambda = 0.5$:

$$V^{(k+1)} = \lambda\, T(V^{(k)}) + (1 - \lambda)\, V^{(k)},$$

where T denotes the Bellman operator. When stalling or oscillation is detected (distance metric unchanged for 10 consecutive iterations, or period-two cycling), the algorithm reduces λ and disables Howard steps to break the cycle. Convergence is declared when the sup-norm distance between successive iterates falls below 10^{-3}.

B.2 Expectation Step

The conditional expectation $E[V(s', z', l', n') \mid s, z]$ integrates over two sources of uncertainty: the next-period aggregate state s' (discrete Markov transition) and the idiosyncratic innovation η (disaster-mixture distribution). The innovation η is drawn from the mixture distribution parameterized by the current regime s; the next-period regime s' enters only through the Markov-weighted continuation value. The transition over s' is handled by direct summation over the two states weighted by the Markov transition probabilities $\Pr(s' \mid s)$.

The integration over the continuous innovation η uses Gauss-Hermite quadrature with $N_{GH} = 7$ nodes. Because the innovation distribution is a two-component mixture (Equation 5), we evaluate the quadrature separately for each component — baseline $\mathcal{N}(\mu_{\text{base},s}, \sigma_s^2)$ and disaster $\mathcal{N}(\mu_{\text{base},s} + \mu_D, \sigma_s^2)$— and combine them with weights $(1 - p_s)$ and p_s. This yields $2 \times N_{GH}$ quadrature nodes per regime, each mapping the current $\log z$ to a candidate $\log z'$:

$$\log z' j = \rho_z \log z + \mu_{c,s} + \sigma_s \sqrt{2}, x_j, \qquad j = 1, \ldots, N_{GH},$$

where x_j are the standard Gauss-Hermite nodes, $\mu_{c,s} \in \{\mu_{\text{base},s}, ; \mu_{\text{base},s} + \mu_D\}$ is the component mean, and the factor $\sqrt{2}$ converts from the physicist's to the probabilist's convention. At each node, the value function is evaluated by linear interpolation along the $\log z$ dimension. The normalized Gauss-Hermite weights $w_j/\sqrt{\pi}$ (summing to one) are premultiplied by the component probability $(1 - p_s)$ or p_s, so the full set of $2N_{GH}$ weights sums to one. Finally, the Jensen correction constant $\mu_{\text{base},s}$ in Equation 5 ensures $E[\exp(\eta)] = 1$ and is given by

$$\mu_{\text{base},s} = -\tfrac{1}{2}\sigma_s^2 - \log\left[(1 - p_s) + p_s \exp(\mu_D)\right].$$

B.3 Permanent Heterogeneity

Permanent revenue types θ_i are discretized using Gauss-Hermite quadrature with $N_\theta = 5$ points on the distribution $\log \theta \sim \mathcal{N}(0, \sigma_\theta^2)$. The grid is shifted so that the minimum node maps to $\theta = 1$, ensuring all firms have positive profits. Because θ enters production multiplicatively (Equation 1), the optimal employment of a type-θ firm scales as $\theta^{1/(1-\alpha)}$ relative to the $\theta = 1$ baseline. Accordingly, the employment and cash grids are scaled proportionally for each θ point, which preserves the bracket indices used for interpolation across types. The value function is solved independently for each θ point.

B.4 Simulation

We simulate a panel of $N = 6{,}000$ firms over $T = 40{,}000$ quarters, discarding the first 500 quarters as burn-in. At each quarter, the aggregate state evolves according to the Markov transition matrix, and each firm draws an idiosyncratic innovation from the disaster-mixture distribution in Equation (5). Given the realized state (s_t, z_{it}) and the beginning-of-period employment l_{it} and cash n_{it}, the firm's policy functions—obtained by interpolating the converged VFI solution—determine next-period employment l_{it+1} and cash n_{it+1}.

Annual aggregation. Quarterly firm-level variables are aggregated to annual frequency following Census and BLS conventions. Annual sales equal the sum of four quarterly sales within the calendar year. Annual employment equals the mean of four quarterly employment levels, which avoids a timing lag from the time-to-build friction (Q1 employment reflects the previous year's hiring decisions). Annual growth rates are computed as arc-percent changes of annual aggregates.

Recession classification. Calendar years are classified as recession or expansion using an NBER-style convention: a year is labeled a recession if two or more of its four quarters fall in the recession regime ($s = R$). This definition aligns with the empirical practice of classifying recession years by the presence of sustained contraction, rather than a single negative quarter.

Cross-sectional moments. In each annual period, we compute the P90–P50, P50–P10, and P90–P10 interquantile spreads and the Kelley skewness of firm-level sales and employment growth rates. These moments are then averaged separately across expansion and recession years. Simulated sales and employment observables are perturbed by multiplicative measurement error $\exp(\varepsilon_{me})$ with $\varepsilon_{me} \sim \mathcal{N}(0, \sigma_{me}^2)$ and $\sigma_{me} = 0.05$ before computing growth rates, matching the noise present in administrative microdata.

Local projections. Aggregate impulse responses are estimated using Jordà (2005) local projections on the simulated aggregate time series. The dependent variable at horizon h is the cumulative log-change in aggregate GDP (or employment) from quarter $t-1$ to $t+h$. The regressor of interest is a recession-onset indicator. Controls include four lags of levels and growth rates, and a linear trend. Standard errors use Newey–West HAC with bandwidth $\max(h+1, 34)$.

B.5 Krusell–Smith General Equilibrium

In general equilibrium, the wage w_t is an additional aggregate state variable. We solve for the equilibrium using the algorithm of Krusell and Smith (1998), which approximates the infinite-dimensional distribution of firms by a low-dimensional sufficient statistic—in our case, the current wage w_t and aggregate output Y_t—and iterates until the forecasting rule is consistent with the equilibrium wage path.

Forecasting rule. Firms forecast next-period wages using a regime-dependent rule that conditions on the current wage and aggregate output:

$$\log w_{t+1} = b_0(s, s') + b_1(s, s') \log w_t + b_3(s, s') \log Y_t, \tag{13}$$

where (s, s') indexes the current-to-next regime transition. This yields twelve coefficients—b_0, b_1, and b_3 for each of the four $(s, s') \in \{E, R\} \times \{E, R\}$ cells—that are updated at each KS iteration.

Because $\log Y_t$ is not directly available in the VFI—only the wage grid is a state variable—we use an auxiliary linear regression $\log Y_t = \alpha_0(s) + \alpha_1(s) \log w_t$ estimated from the simulated data. This auxiliary prediction is substituted into (13) during the bracket precomputation step of the VFI, so that the forecasted wage depends only on the wage grid point and the current regime. Notice that this auxiliary regression is only used in the VFI step: the coefficients in step 3 below are updated using the actual simulated Y_t.

VFI at the wage grid. The firm's value function is solved jointly at $M = 13$ wage grid points spanning $[0.95, 1.05]$. The continuation value at each grid point interpolates the value function at the *forecasted* wage w_{t+1} (bracket interpolation across the wage grid), so the forecasting rule (13) enters the Bellman equation directly.

KS iteration. The algorithm iterates between three steps:

1. *Value function solution*: solve the firm's Bellman equation at all wage grid points, taking the forecasting coefficients as given.

2. *Simulation*: simulate the economy with $N = 6{,}000$ firms over $T = 40{,}000$ quarters. At each period, aggregate employment $L_t = \sum_i l_{i,t}$ is predetermined by firm decisions in $t - 1$. The market-clearing wage follows the GHH inverse labor supply: $w_t = (L_t/L_{ss})^{1/\varepsilon}$, where L_{ss} is mean expansion employment from the PE simulation and $\varepsilon = 5$.

3. *Coefficient update*: regress $\log w_{t+1}$ on $[1,\ \log w_t,\ \log Y_t]$ by OLS, separately for each (s_t, s_{t+1}) cell, where $\log Y_t$ is the realized aggregate output from the simulation. The auxiliary regression of $\log Y_t$ on $[1,\ \log w_t]$ is also re-estimated by regime. The updated coefficients are dampened toward the previous iteration with mixing parameter 0.95.

The iteration continues until the maximum absolute change in forecasting coefficients falls below 10^{-4}.

Convergence diagnostics. Table B.3 reports the converged coefficients and the R^2 of the forecasting rule for each regime transition. The R^2 ranges from 0.947 ($E \to E$) to 0.964 ($R \to R$), indicating that the current wage and aggregate output jointly provide a tight summary of the information in the firm distribution that is relevant for forecasting future wages.

After convergence, we evaluate the forecasting rule along the simulated wage path. The mean absolute forecast error is 0.031% of w, with a maximum of 0.35%. We also compute the Den Haan (2010) metric---the mean absolute percentage deviation between the simulated wage path and the path implied by iterating the forecasting rule forward from initial conditions. The Den Haan statistic is 0.0035, well below the 1% threshold that is standard in the literature. Equilibrium wages range from 0.980 to 1.007 (mean 0.992, std 0.005), with zero boundary hits on the wage grid. These diagnostics confirm that the forecasting rule provides an accurate approximation of the rational expectations equilibrium for this economy.

Table B.3 – Forecasting Rule: Converged Coefficients and Accuracy

	$E \to E$	$E \to R$	$R \to E$	$R \to R$
b_0	-5.21	-5.23	-2.49	-2.52
b_1	-2.39	-2.40	-0.71	-0.73
b_3	0.66	0.66	0.31	0.32
R^2	0.95	0.95	0.96	0.96

Notes: The table reports the converged coefficients and R^2 from OLS regressions of $\log w_{t+1}$ on $[1,\ \log w_t,\ \log Y_t]$, estimated on 40,000 simulated quarters and conditioned on the (s_t, s_{t+1}) regime transition.

B.6 Estimation

The model is calibrated in partial equilibrium with fixed wages ($w = 1$). The cross-sectional moments targeted in calibration (P90–P50 and P50–P10 of firm sales growth) are driven by idiosyncratic shocks and are approximately invariant to the small equilibrium wage movements that arise in general equilibrium. After calibration, we solve the general equilibrium using the Krusell–Smith algorithm described in Section B.5.

The five shock parameters $(\sigma_E, \sigma_R, \mu_D, p_E, p_R)$ and the cash wedge κ are estimated by minimizing the weighted distance between simulated and empirical moments. The estimation proceeds in two stages.

Stage 1: Shock parameters. The five-dimensional vector $(\sigma_E, \sigma_R, \mu_D, p_E, p_R)$ is estimated by minimizing

$$Q(\boldsymbol{x}) = \left[\hat{\boldsymbol{m}}(\boldsymbol{x}) - \boldsymbol{m}^{\text{data}}\right]' W \left[\hat{\boldsymbol{m}}(\boldsymbol{x}) - \boldsymbol{m}^{\text{data}}\right],$$

where $\hat{\boldsymbol{m}}(\boldsymbol{x})$ is the four-dimensional vector of simulated spread moments (P90–P50 and P50–P10 in expansion and recession), $\boldsymbol{m}^{\text{data}}$ contains the corresponding LBD targets, and W is a diagonal weighting matrix with entries $1/m_j^2$ (inverse-squared magnitude, floored at 0.02). Each function evaluation requires solving the VFI and simulating the full panel. The minimization uses MATLAB's `patternsearch`, a derivative-free algorithm suited to objectives that are noisy due to simulation randomness. The search uses reduced grids ($N_z = 41$, $N_l = 80$, $N_n^{\text{fine}} = 31$) for speed; final verification uses the full-resolution grids.

Stage 2: Cash wedge. Conditional on the shock parameters from Stage 1, κ is estimated by one-dimensional pattern search targeting the median cash-to-sales ratio (Compustat target: 0.08). The permanent heterogeneity parameter σ_θ is calibrated semi-analytically: for each candidate σ_θ, the employment distribution implied by $l \propto \theta^{1/(1-\alpha)}$ is computed via Monte Carlo, and the value yielding the target top-10% employment share (LBD target: 0.65) is found by interpolation.

C Empirical Local Projections

In this section, we describe the empirical specification of the Local Projections in our analysis. We apply the same regression for the macro time series—which we describe

TABLE C.4 – Recession Onsets in the Sample

Onset ($D_s = 1$)	Common label
1974Q3	First oil shock / Great Inflation
1980Q2	Volcker recession (double-dip, leg 1)
1981Q4	Volcker disinflation (double-dip, leg 2)
1990Q4	1990–91 recession
2008Q3	Global Financial Crisis
2020Q1	COVID-19

below—and the model-generated time series. For each outcome y (in logs) and horizon $h \in \{-8, \dots, +16\}$, we estimate a separate local-projection regression (Jordà, 2005). For response horizons $h \geq 0$,

$$y_{s+h} - y_{s-1} \;=\; \beta_h \, D_s \;+\; \sum_{j=1}^{4} \Gamma_{h,j} \, X_{s-j} \;+\; \alpha_h \;+\; \varepsilon_{s,h}, \tag{14}$$

where D_s is the recession-onset dummy of Section C.1 and X_{s-j} includes four lags of y in levels and growth rates and a linear trend. We normalize to the quarter before onset by setting $\beta_{-1} = 0$, and for the pre-event horizons $h < -1$ we drop the lags so that β_h for $h < -1$ is the difference between the mean of $y_{s+h} - y_{s-1}$ and its mean in non-onset quarters. Coefficients are reported in percentage points (100× the log deviation from the $s-1$ value). We use Newey--West standard errors (Newey and West, 1987) with the Montiel Olea and Plagborg-Møller (2021) bandwidth $m(T, h) = \max\big(|h| + 1,\ \lfloor 1.3\, T^{1/3} \rfloor\big)$, and plot symmetric 90% and 95% bands.

C.1 Recession-onset indicator

Let $g_s = \Delta \log \mathrm{GDPC1}_s$ be quarterly real GDP growth. We date a recession onset at the first quarter of a confirmed contraction:

$$D_s \;=\; I\,[\, g_{s-1} \geq 0 \;\wedge\; g_s < 0 \;\wedge\; g_{s+1} < 0 \,]\,. \tag{15}$$

So $D_s = 1$ when a non-negative quarter is followed by two negative ones: $s - 1$ is the last expansion quarter, s the onset, and $s+1$ the confirmation. The rule uses only ex-post GDP, never fires mid-downturn, and ignores isolated single-quarter declines. Over the central sample, the indicator selects the six onsets in Table C.4.

C.2 Sample and data

The sample runs from 1970Q1 onward ($T \approx 225$ quarters); we begin in 1970 because earlier real-output and TFP vintages are less comparable across data releases. With six onsets, the panel yields $n = 220$ regressions at $h = 0$, declining by one per horizon. The three outcomes are real GDP per capita, nonfarm-business hours per capita, and total factor productivity (TFP) — respectively the headline output, labor, and productivity series in the

Table C.5 – Data Time Series

Symbol	Series (source ID)	Provider	Freq.	Transformation
y	Real GDP,	BEA (FRED)	Q, SA	$\log(y/N)$
h	Hours	BLS (FRED)	Q, SA	$\log(h/N)$
z	TFP	Fernald (FRBSF)	Q	cumulate to level, log
N	Population	BLS/CPS (FRED)	M	3-month mean

Notes: The table provides data series, providers, and transformations. FRED series are retrieved through the FRED API; the Fernald TFP series is downloaded from the FRBSF. Real GDP is chained 2017$ (GDPC1), the hours series is Nonfarm Business Hours (HOANBS). Population (CNP16OV) is used for per-capita scaling.

model-vs-data comparison. Real GDP (GDPC1, originating with the Bureau of Economic Analysis) and nonfarm-business hours (HOANBS, Bureau of Labor Statistics) are quarterly and seasonally adjusted; we scale each by the working-age population (CNP16OV, from the BLS Current Population Survey, aggregated from monthly to quarterly as a three-month average) and take logs. All FRED series are retrieved programmatically through the FRED API.

Total factor productivity is taken from Fernald's quarterly TFP dataset for the US business sector (Fernald, 2014), distributed by the Federal Reserve Bank of San Francisco. We use the TFP growth series (dTFP, reported as annualized log changes in percentage points), convert it to quarterly log changes, and cumulate to a level index normalized to unity at the start of the sample. Table C.5 summarizes the series.

References.

Cerra, V. and Saxena, S. C. (2008). Growth dynamics: The myth of economic recovery. *American Economic Review*, 98(1):439--457.

Den Haan, W. J. (2010). Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):4--27.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789--1823.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703--708.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263--1297.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, 49(2):326--365.

Yasar, M., Raciborski, R., and Poi, B. (2008). Production function estimation in Stata using the Olley and Pakes method. *Stata Journal*, 8(2):221--231.