SKEWED BUSINESS CYCLES

Sergio Salgado
Fatih Guvenen
Nicholas Bloom

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

ABSTRACT

Using firm-level panel data from the US Census Bureau and almost fifty other countries, we show that the skewness of the growth rates of employment, sales, and productivity is procyclical. In particular, during recessions, they display a large left tail of negative growth rates (and during booms, a large right tail of positive growth rates). We find similar results at the industry level: industries with falling growth rates see more left-skewed growth rates of firm sales, employment, and productivity. We then build a heterogeneous-agent model in which entrepreneurs face shocks with time-varying skewness that matches the firm-level distributions we document for the United States. Our quantitative results show that a negative shock to the skewness of firms' productivity growth (keeping the mean and variance constant) generates a persistent drop in output, investment, hiring, and consumption. This suggests the rising risk of large negative firm-level shocks could be an important factor driving recessions.

Sergio Salgado
The Wharton School
University of Pennsylvania
2325 Steinberg Hall - Dietrich Hall
Philadelphia, PA 19104-6367
ssalgado@wharton.upenn.edu

Fatih Guvenen
Department of Economics
University of Minnesota
4-101 Hanson Hall
1925 Fourth Street South
Minneapolis, MN, 55455
and NBER
fatihguvenen@gmail.com

Nicholas Bloom
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
nbloom@stanford.edu

1 Introduction

This paper studies the cyclicality of the distribution of the growth rate of firm-level outcomes. In the previous literature, recessions have been characterized as a combination of a negative first-moment (mean) shock and a positive second-moment (uncertainty) shock (Bloom, 2014). In this paper, we argue that recessions are also accompanied by negative third-moment (skewness) shocks implying that, during economic downturns, a subset of firms does extremely badly, leading to a left tail of large negative outcomes. This is often accompanied by a deceleration of growth for a subset of 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 shrinking and the right tail expanding. Consequently, the skewness of the growth rates is procyclical.

Using firm-level panel data from the US Census and Compustat and panel data on firms from thirty-nine other countries, we show that the cross-sectional skewness of the distribution of several firm-level outcomes, such as sales growth, employment growth, productivity growth, and stock returns, is strongly procyclical. As an illustration of our main empirical result, the top panel of Figure 1 displays the distribution of firms' employment growth from the US Census Bureau's Longitudinal Business Database (LBD). The solid line shows the empirical density of employment growth pooling observations from the most recent two recession years, 2001–02 and 2008–09. The dashed line instead shows the density for the subsequent expansion years, in this case, years 2003 to 2006 and 2010 to 2014. One can clearly see that, relative to expansion periods, the distribution of employment growth during recessions has a thicker left tail, whereas the right tail exhibits little change, indicating an increase in dispersion that is mostly due to a widening left tail.

This asymmetric change in the distribution of employment growth from expansion to recession years can be quantified using the Kelley skewness (Kelley, 1947), a measure that is robust to the presence of outliers. This measure is defined as the difference between the 90th-to-50th log percentiles differential (a measure of dispersion in the right tail) and the 50th-to-10th log percentiles differential (a measure of dispersion in the left tail) divided by the 90th-to-10th log percentiles differential (a measure of the total dispersion of the distribution). For a distribution with a compressed upper half and a dispersed lower half (i.e., a left skew), the Kelley skewness is negative. In the case of the top panel of Figure 1, we find a decline in the dispersion of employment growth above the median from 0.22 to 0.20 from expansion to recession years, whereas the dispersion below the median increases from 0.17 to 0.25. This asymmetric change in the tails generates a decline in the Kelley skewness from 0.10 during expansions to –0.12 during recessions.[1] Similarly, the bottom panel of Figure 1 displays the distribution of sales

---

[1]Put differently, a Kelley skewness of 0.10 indicates that during expansion, 45% of all the dispersion is

growth for Compustat firms for recessions years (2001-02 and 2008-09) and for expansion years (2003 to 2006, and 2010 to 2014). As in the case of employment growth, here we also find that recessions are characterized by a widening left tail, which gives rise to both an increase in dispersion and a decline in the skewness of the sales growth distribution. Furthermore, using firm and establishment-level panel data, we show that the skewness of stock returns, and firms' productivity shocks, are also left skewed during recession. This suggests that the underlying demand and productivity shocks experienced by firms has a cyclical skew.

We find the same empirical pattern at the two-digit NAICS industry level: the *within-industry* skewness of firm-level employment growth, sales growth, stock returns, and productivity growth is positively correlated with the industry growth rate. Moreover, the same pattern is also seen globally. Using firm-level data for over thirty-nine countries that are both geographically and economically diverse, we show that the skewness of the same firm-level variables within each country is robustly procyclical with respect to that country's business cycle.

Given these empirical patterns, we then evaluate to what extent these cyclical fluctuations in the skewness of firm-level shocks can account for business cycle variations in investment, hiring, and output. To investigate this we study a set of vector regression (VAR) estimations which show that skewness shocks to the distribution of firms' stock returns predict large drops in aggregate economic activity. In particular, a drop in firm-level skewness precedes large declines in industrial production, investment and employment. Of course VAR estimates do not identify causal impacts, but they do highlight that increasing left-tailed skewness of firm-level shocks forecasts significantly lower growth rates.

In the second part of the paper we take a different modeling approach to evaluate the potential macro impact of fluctuations in micro-skewness. We build a heterogeneous-agents model where the key feature is the presence of a large number of entrepreneurs that face shocks with time-varying variance and time-varying skewness. To capture the potentially non-linear response of firms to shocks, we assume that entrepreneurs are risk-averse, face a combination of convex and non-convex adjustment costs to capital, and can invest in capital and a risk-free asset. We numerically solve the model and choose the parameters of the firm's productivity process so that our modeled economy matches the average skewness of the sales growth distribution we observe among US firms during expansionary periods and the large decline in skewness observed during a typical recession.

In our main quantitative exercise we study the aggregate effects of a pure skewness shock—that is, a decline in the skewness of firms' productivity shocks–while keeping the mean and variance constant. Our model predicts that a change in the skewness of the distribution of firm-level

accounted for by firms with employment growth below the median, whereas during recessions, this share increases to 56%.

shocks alone, matching the decline in the skewness of sales growth we observe among US firms, would reduce GDP by 1.7%. Consistent with our VAR evidence, the decline in aggregate economic activity is quite persistent as GDP stays below its pre-shock level several quarters after the shock. This is in contrast to the standard uncertainty shock analyzed in the literature that typically generates a sharp drop and rapid rebound of GDP. The significant and persistent drop in output is driven by a decline in capital investment, which is the result of three forces. First, the presence of a fixed cost to capital adjustment creates a real options effect that reduces the incentives of firms to invest when skewness declines. This is a reflection of the Bernanke (1983) *"Bad News Principle"*, that only outcomes about the bad state of the world matter for option value to delay investment. Second, the drop in skewness makes capital riskier, inducing an increase in investment in the risk-free asset. Third, relative to the standard uncertainty shock (a symmetric increase in dispersion), in our model a decline in skewness results in a widening left tail of the firm productivity distribution without a corresponding widening of the right tail (an asymmetric increase in dispersion). This ameliorates the impact of the *Oi-Hartman-Abel* effect from uncertainty shocks, suppressing the overshoot from pure uncertainty shocks.[2]

In summary, our results indicate that a negative shock to the skewness of firms' productivity (that keeps the mean and variance constant) can generate a moderate recession by itself. Of course, recessions are likely driven by a combination of shocks to multiple moments - for example, a negative, positive and negative shock to the first, second and third moments respectively - so we see our paper as highlighting the additional contribution of left-tail micro-skewness in driving recessions.

This paper is related to several strands of literature. First and foremost, our paper relates to the study of the effects of uncertainty on firms' decisions. Several papers have shown that an increase in uncertainty can have important macroeconomic implications in the presence of adjustment costs, risk aversion, or financial frictions.[3] Our results are complementary to this literature as we show that dispersion of firms' outcomes–a standard measure of uncertainty–increases as a results of a widening left-tail.

Second, several authors have suggested that rare disasters—presumably arising from an asymmetric distribution of shocks—can generate large fluctuations in economic activity, such as the Great Recession. Reviving the ideas introduced first by Rietz (1988), Barro (2006) uses a panel of countries to estimate the probability of large disasters and argue that these low-probability events can have substantial implications for aggregate economic activity and

---

[2]See a discussion about the Oi-Hartman-Abel in the survey article of Bloom (2014).

[3]See, for example, Arellano *et al.* (2018), Fernandez-Villaverde *et al.* (2011), Schaal (2017), Bachmann and Bayer (2013), Bachmann and Bayer (2014), Gilchrist *et al.* (2014), Jurado *et al.* (2015), Leduc and Liu (2016), Basu and Bundick (2017), Berger *et al.* (2017), Alfaro *et al.* (2018), Kozeniauskas *et al.* (2018), and Bloom *et al.* (2018).

asset pricing. Several papers have confirmed the importance of fluctuation in disaster risk for aggregate economic activity.[4] The results of our paper can be seen as evidence that rare disasters also occur at the microeconomic level, and because firms are not typically perfectly insured against micro-economic risk these have real effects.

Finally, our paper contributes to a growing literature that focuses on the cyclical patterns of the skewness (and other higher order moments) of worker and firm outcomes, such as labor earnings (e.g. Guvenen *et al.* (2014), Busch *et al.* (2017), and Harmenberg and Sievertsen (2017)), firm productivity (Kehrig (2011)), employment growth (e.g. Ilut *et al.* (2018) and Decker *et al.* (2015)), and stock returns (e.g. Harvey and Siddique (2000), Oh and Wachter (2018), and Ferreira (2018)).

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 the model and Section 5 presents our quantitative results. Section 6 concludes.

2 Data and Measurement

2.1 Data and Sample Selection

Our analysis is based on five large dataset that encompassing firm- and establishment-level information for the United States and for almost forty other countries.[5] The breadth of our dataset allows us to provide a detailed description of the cyclical patterns of the distribution of firm-level outcomes and productivity growth.[6]

First, we extract panel data on employment at the firm and establishment level from the US Census Bureau's Longitudinal Business Database (LBD). The LBD provides high-quality measures of employment, wage bill, industry, and age for the entire US nonfarm private sector linked over time at the establishment level from 1976 to 2015. From the LBD, we construct employment at the firm and establishment levels and use it to calculate cross-sectional moments of the distribution of employment growth at narrow firm population groups. The LBD contains over 6 million firms per year, which for measuring higher-order moments like skewness is a major advantage.

Second, we obtain output, input, and productivity data for a panel of manufacturing establishment combining information from the US Census of Manufacturing and the Annual Survey

---

[4]See for instance Gabaix (2008, 2012), Gourio (2008, 2012, 2013), Wachter (2013), Kilic and Wachter (2015), Kozlowski *et al.* (2018, 2016), Venkateswaran *et al.* (2015), among others.

[5]Table A.6 in Appendix B.3 shows the list of countries in our dataset.

[6]Replication files for the empirical section can be downloaded from here.

of Manufacturing (ASM) covering years from 1976 to 2015. From the merged dataset we select establishment with at least ten years of valid observations on employment and sales, which by the ASM methodology oversamples larger establishments (and thus implicitly larger firms). These datasets also include a measure of total factor productivity which is calculated by the US Census Bureau.

Third, we draw panel data of publicly traded firms from Compustat. Although this dataset contains mostly large established firms, it provides several additional variables which are important in our analysis. In particular, we use data on quarterly and annual sales, annual employment, and daily stock prices from 1970 to 2017, 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.* (2006).

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 cross-country publicly traded firm-level panel data containing sales and employment information between 1986 and 2016 from the Osiris dataset collected by the Bureau van Dijk (BvD). In order to maintain a homogeneous sampling criteria, we only consider firms with ten or more years of data. Additionally, we restrict our sample to country-year bins with more than one hundred firms, countries with at least ten years of data, and years with five countries or more. Our main results are based on an unbalanced panel of firms spanning thirty nine countries from 1991 to 2015. 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 twenty-nine countries from 1985 to 2017.

Finally, we obtain additional firm-level data from the Amadeus dataset also collected by the BvD. This dataset comprises a smaller sample of countries, for a shorter timespan, but with rich firm-level information for small and large firms, both publicly traded and privately held. In particular, Amadeus provides information on sales, employment, value added, capital, and labor input cost so that we can estimate TFP at the firm level. Our sample contains information for twenty one European countries starting in the mid 2000s.

Table I summarizes the data sources and the main sample characteristics.[7] Additional details on data construction, selection criteria, and moment calculation for each dataset used in our analysis can be found in Appendix B.[8]

---

[7]Appendix Table A.6 shows a list of the countries we consider in our analysis and the data available for each of them.

[8]The online appendix available from here and at the author's websites provides further details about the underlying data and provides the material to replicate the results presented in the empirical section of the paper.

2.2 Measuring Dispersion and Skewness

For most of our results, we measure the growth rate of a firm-level outcome as the log-difference between period t and $t + k$ where t is a quarter for stock returns and a year in the case of employment, sales, and productivity. For both dispersion and skewness, we use quantile-based measures that are robust to outliers, which are common in micro datasets. As we shall see, they also have magnitudes that are easy to interpret. Our measure of dispersion is the differential between the 90th and 10th percentiles, denoted by $P9010_t$, where t is a quarter or a year depending on the dataset. Additionally, we use the differentials between the 90th and 50th percentiles, $P9050_t$, and between the 50th and 10th percentiles, $P5010_t$, as measures of dispersion in the right and left tails, respectively. Finally, our preferred measure of skewness is the Kelley skewness (Kelley, 1947), which is defined as

$$KSK_t = \underbrace{\frac{P90_t - P50_t}{P90_t - P10_t}}_{\text{Right Tail Share}} - \underbrace{\frac{P50_t - P10_t}{P90_t - P10_t}}_{\text{Left Tail Share}} \qquad \in [-1, 1]. \tag{1}$$

As seen here, the Kelley measure 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 value of Kelley skewness indicates that the left tail accounts for more than one-half of the total dispersion and the distribution is negatively skewed, and vice versa for a positive value.[9]

3 Skewness over the Business Cycle

In this section, we show that the skewness of the distribution of firm-level growth is positive during expansions but becomes negative during recessions in both the United States (Section 3.1) and across countries (Section 3.2); we then confirm that our results hold within industries (Section 3.3) and for firms' productivity shocks (Section 3.4).

3.1 US Evidence

The first contribution of our paper is to show that the skewness of the growth rates of firm–level outcomes varies over time and is strongly procyclical. We start by considering the evolution of the Kelley skewness of the distribution of firms' employment growth from the US Census Bureau's LBD, which is displayed in the top panel of Figure 2. To calculate Kelley

---

[9]An important drawback of this measure of skewness is that it 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 98th and 2nd. We have explored some of these alternative choices and did not find them to matter for our results (see Appendix Figure A.1). Additional measures of skewness can be found in Kim and White (2004).

skewness, we weight observations by firm employment so that our measure reflects the underlying firm-size distribution.[10] Figure 2 shows, first, that the skewness of employment growth, on average, is positive and around 0.10 for most of the sample period. Second, the skewness of the distribution employment growth is strongly procyclical, declining from an average of 0.11 at the peak of the typical recession to around –0.10 at the trough. Noticeable, the Great Recession represents the largest declined in the skewness of employment growth over the entire sample period, with a Kelley Skewness declining to a value of -0.20. This indicates that during the recession, 60% of the total mass of the distribution is accounted for by firms below the median.[11] Similarly, the bottom of Figure 2 shows the cross-sectional skewness of annual sales growth for Compustat firms. Relative to the LBD, this is a more selective set of mostly large firms. Nevertheless, we find that the skewness of the distribution of sales growth is positive on average and declines around 0.20 points during a recession.

The procyclicality in the skewness of firm growth is driven by a rapid change in the relative weight of the tails of the distribution. This can be observed in the top panel of Figure 3, where we plot the time series of the $P5010_t$ (black line with squares) and the $P9050_t$ (blue line with circles) of the employment growth distribution using data from the LBD. The bottom panel of Figure 3 shows the same statistics for the sales growth distribution from Compustat. Two important aspects are worth noticing. First, during expansionary periods, the right tail outweighs the left tail (most of the time, $P9050_t$ is above $P5010_t$), generating a distribution of firms growth that is positively skewed during expansions. Second, for both employment and sales growth, recessions are episodes in which the $P5010_t$ increases, indicating a left tail that stretches out, whereas the $P9050_t$ decreases, indicating a right tail that contracts. This asymmetric change in the tails drives the drop in the skewness of firms' employment and sales growth.[12] Appendix Figure A.1 confirms the robustness of these results to a range of different approaches and sample selection. In particular, we show that the skewness of employment growth in the LBD is procyclical if we divide firms in groups of different size or age, if we look at establishment employment growth, if we explicitly consider the entry and exit of firms, and if we consider different measures of skewness.

To have a better sense of the magnitude of the change in skewness and its relation with the

---

[10]In particular, we weight the employment growth of firm i in period t by the average employment in periods t and $t+1$, that is, $\overline{E}_{i,t} = 0.5 \times (E_{i,t} + E_{i,t+1})$. The results for publicly traded firms are unweighted since most of the firms are large.

[11]The procyclicality of the skewness of employment growth in the LBD has been discussed in different forms in previous papers (e.g. Davis and Haltiwanger (1992) and Ilut *et al.* (2018)). We complement these studies by looking at the cyclicality of the skewness of employment growth within industries, age, and size categories.

[12]This asymmetric change is also observed at higher frequencies. For instance, Figure A.2 shows a similar set of results for the annual change of quarterly sales.

cycle, columns (1) to (3) of Table II show a set of time series regressions of the form

$$KSK_t = \alpha + \beta \Delta GDP_t + \delta t + \epsilon_t, \tag{2}$$

where the dependent variable is the Kelley skewness of the cross-sectional distribution of different firm-level outcomes. In all regressions, the independent variable is the log change of real GDP per capita—which we have normalized to have unitary variance so the coefficients are comparable across columns—and t is a linear trend. The estimated coefficients are positive and large for all three variables—employment growth, sales growth, and stock returns—and also economically and statistically significant (at the 1% level for the first two and the 5% level for the third). For example, the estimated coefficient of 4.6 in column (1) implies that a two standard deviation—or about a 4%—drop in GDP per capita growth is associated with a 0.09 fall in the Kelley skewness of the firm employment growth distribution. Column (2) shows a similar result for sales growth with a larger coefficient (5.4). Column (3) shows a smaller coefficient for stock returns (2.1) that is still highly significant.[13] The change in skewness of stock-returns over the cycle also suggests the decrease in skewness of sales and employment growth is driven, at least in part, by a rise in negatively skewed external shocks (e.g. productivity or demand shocks) rather than skewed firm control variable (like investment and employment). In order to shed additional light on the cyclicality of the skewness of firms' shocks, in Section 3.4 we directly test whether firm-level productivity shocks are also left skewed during recessions.

3.2 Cross-Country Evidence

Is the procyclical skewness we have documented so far a pattern specific to the United States, or is it also observed in other countries? The second contribution of our paper is to answer this question using firm-level panel data covering almost fifty countries that are both geographically and economically diverse, spreading over five continents including developing countries (such as the United States, Germany, Japan, and others) and developed countries (such as Peru, Egypt, Thailand, and others).

The top panel of Figure 4 displays the empirical density of the distribution of log sales growth (in US dollars as of 2005) for a unbalanced panel of firms pooling all the countries in our sample from 1991 to 2015. The solid red line is the density of the growth rate of sales during recession periods, where a recession is defined as a year in which the growth rate of

---

[13]As we show in Appendix A, these results are remarkably robust to different specifications and across several firm-level outcomes. In particular, Table A.2 shows that the skewness remains strongly procyclical if we residualize the firms' outcomes by the firms' observable characteristics and fixed unobserved heterogeneity and if we consider the growth of log sales-per-worker (more closely related to firm productivity). Furthermore, we confirm that the dispersion in firms' growth is countercyclical (Appendix Table A.1) but we do not find significant business cycle variation in the kurtosis (right panel of Table A.2).

GDP is in the first decile of the country-specific GDP growth distribution. The dashed blue line is the density of sales growth during expansion periods defined as years in which GDP growth is above the first decile of the country-specific distribution of GDP growth. Similar to the results presented in Figure 1, the dispersion in sales growth increases little during recession years, with $P9010_t$ rising slightly from 0.82 to 0.85. However, this modest increase masks larger changes in each tail: the left tail stretches out, with $P5010_t$ rising from 0.36 to 0.43, and the right tail shrinks, although by a smaller amount, with $P9050_t$ falling from 0.46 to 0.43. The opposite moves of each tail dispersion partially cancel out each other, leading to the smaller rise in $P9010_t$ just mentioned. In contrast, for skewness, the contraction of the upper tail and the expansion of the lower tail dispersion reinforce each other to generate a larger decline in Kelley skewness, which falls from 0.12 to 0.0.

To have a clearer picture of the cyclical changes of the skewness over the business cycle, the bottom left panel of Figure 4 shows a scatter plot in which the x-axis is the average of firm employment growth within a country-year bin, and the y-axis is the Kelley skewness of the same firm-level outcome. The data points align nicely along a straight line over a wide range of average employment growth rates (ranging from –0.15 to 0.20), confirming the strong positive relationship between skewness and the within-country business cycle. Our results indicate that, when the average firm employment growth is –0.15 (typically during a big recession) the Kelley skewness is –0.30, implying that two-thirds of the mass of the distribution of employment growth is accounted for by the left tail. In contrast, when the average employment growth is 0.10, the skewness is 0.30, indicating the opposite split, with two-thirds of the total dispersion now being accounted for by the right tail. The bottom right panel of Figure 4 shows a similar result for sales growth. Importantly, to construct these figures we have controlled for country- and time-fixed effects, so these results are not driven by fixed characteristics of the countries considered in the sample or by global shocks—such as the Great Recession—that can affect all countries at the same time.[14]

In columns (4) to (6) in the center panel of Table II we repeat 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. This time, the dependent variable is the skewness of employment growth, sales growth, or 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 unitary variance to facilitate the comparison

---

[14]One important concern is that our cross-country results are based exclusively on publicly traded firms. Interestingly, we also find remarkably similar results using an unbalanced panel of firms, private and publicly traded, drawn from the BvD Amadeus dataset, as Figure A.3 in the Appendix shows. Relative to our baseline sample, the BvD Amadeus dataset covers a much larger sample of firms, but over a shorter period of time (2000 to 2015 for most countries) over a smaller sample of European countries.

with the rest of the results. The regressions also include a full set of time and country fixed effects to control for aggregate economic conditions that might affect all countries simultaneously or fixed differences across countries. The results confirm our previous findings of procyclical skewness for all three variables with similar levels of statistical significance. Compared with the United States, the estimated coefficient is slightly higher for employment (5.9 across countries versus 4.6 for the United States), somewhat lower for sales (3.1 versus 5.4), and nearly identical for stock returns. These results further confirm the procyclical nature of skewness in firm-level outcomes.

3.3 Industry-Level Evidence

We now turn to industry-level data from the United States and investigate whether the skewness of firms' outcomes is procyclical within different industries. To this end, using LBD data, the top panel of Figure 5 shows a scatter plot of the skewness of employment growth against the average employment growth within an industry-year cell, where an industry is defined at the two-digit NAICS level. In this case, a positive correlation indicates that periods of low economic activity at the *industry* level are associated with a negative shift in skewness within that industry, and vice versa for periods of high economic activity. As in the country-level results, we include a full set of time and industry fixed effects, so that the results are driven by within-industry changes rather than aggregate changes in growth rates. In terms of magnitudes, the top panel of Figure 5 shows that when the average industry employment growth is –0.08, the Kelley skewness is around -0.20, indicating that 60% of the total dispersion in employment growth within an industry is accounted for by the left tail of the distribution. When the average employment growth is 0.08 instead, the Kelley is skewness is 0.20, indicating that the right tail accounts for 60% of the total dispersion. Similarly, the bottom panel of Figure 5 shows that the within-industry skewness of sales growth is higher when the average growth in that industry is higher.[15] Hence, industries that grow faster are also sectors in which the skewness of firm-level outcomes is positive.

To further examine the relation of the industry cycle and the skewness of sales growth, employment growth, and quarterly returns, columns (8) to (10) of Table II display a series of industry panel regressions in which the dependent variable is the Kelley skewness of the log

---

[15]Similarly to the aggregate results discussed in Section 3.1 and Section 3.2, the change in the skewness of the within-industry distribution of firms' growth is driven by an asymmetric response of the right and left tails to the industry business conditions. This is clearly seen in Figure A.4 in Appendix A that shows that $P9050_t$ of log sales growth is positively correlated with the within-industry cycle, increasing during periods of high economic activity within the industry. In contrast, the $P5010_t$ is negatively correlated with the industry cycle, increasing during periods of low economic activity within industry level. Interestingly, the dispersion in both tails of the distribution shows a hockey-stick pattern rising sharply as the average sales growth moves away from zero. This uneven within-industry pattern drives the positive correlation between skewness and the economic conditions within an industry depicted in Figure 5.

growth of different firm-level outcomes within an industry-year cell. In this case, we capture the within-industry business conditions by the average log sales growth in an industry-year cell. To facilitate the interpretation of magnitudes, we have rescaled the sales growth within each industry to have a variance of one so that the regression coefficient can be interpreted as the effect of a change in the within-industry sales growth of one standard deviation and can be easily compared to the coefficients of columns (1) to (3) in Table II. These results again show a strongly procyclical skewness.[16]

3.4 Firm-Level Productivity Evidence

The evidence we have provided so far indicates that the skewness of the distribution of firm-level outcomes is procyclical, declining during periods of low economic activity. Such pattern, however, can be the result of shocks with time-varying higher order moments (i.e. time-varying skewness) and/or the endogenous responses of firms to symmetric shocks that generate asymmetric changes in firm's outcomes. In order to separate the importance of shocks from endogenous decisions, in this section we study the cyclical properties of the distribution of firms' productivity shocks, and we show and that the skewness of this distribution is procyclical, across countries and for manufacturing establishments in the United States.

We first use firm-level data for a subsample of European countries for which we have rich enough data to measure revenue TFP obtained from the Amadeus dataset collected by the BvD.[17] In particular, for each country, we estimate firm-level log productivity, z_i, as

$$z_{i,t} = \log Y_{i,t} - \alpha_K \log K_{i,t} - \alpha_L \log L_{i,t},$$

where $Y_{i,t}$ is the deflated value added of firm i in year t, $K_{i,t}$ is a measure of the capital stock, and $L_{i,t}$ is a measure of labor input. As is standard in the productivity literature (e.g. Syverson (2011)) we assume constant returns to scale at the firm-level (so $\alpha_K = 1 - \alpha_L$) and measure α_L as the industry-country labor share (the ratio of the total wage bill to total value added).

Once we have calculated $z_{i,t}$, we obtain a measure of firms' productivity shocks, denoted by $\epsilon_{i,t}$, from the residual of the following firm-level panel regression,

$$z_{i,t} = \beta_0 + \beta_1 z_{i,t-1} + \mu_i + \delta_t + \epsilon_{i,t}, \tag{3}$$

---

[16]We also find a similar positive and statistically significant relationship between industry cycles and skewness when we consider each industry separately. Appendix Figure A.5 shows the slope coefficient of a set of within-industry time series regressions of the Kelley skewness of firms' growth on the within-industry average firm growth. Notice that, although there is substantial heterogeneity across industries, for all of them the coefficient on the average firm growth is positive and economically and statistically significant.

[17]Our firm-level data comprises information of small and large firms, both publicly traded and privately held from seventeen European countries, namely, Germany, Denmark, Spain, Finland, France, United Kingdom, Greece, Hungary, Ireland, Island, Italy, Netherlands, Norway, Poland, Portugal, Sweden, and Ukraine. Appendix B.4 describes in full detail the sample selection and estimation procedure.

where μ_i is a firm fixed effect and δ_t is a year fixed effect. We then calculate different moments of the distribution of $\epsilon_{i,t}$ within a country-industry-year cell.

In order to have a simple comparison to our previous results, the top panel of Figure 6 shows a scatter plot in which each observation is a year-industry-country cell. In the x-axis we plot the average productivity shock and in the y-axis we plot the Kelley skewness. In this figure, we have controlled for country, industry, and year fixed effects and therefore, our results are neither driven by fixed differences across countries and industries, nor by aggregate economic fluctuations. As before, we find that skewness and average shock are positively correlated. In terms of magnitudes, an average decline of productivity of 0.05 is associated to a decline of 0.05 in the skewness of the distribution. As we show in Appendix B, the procyclicality of the skewness of firm's shocks is robust to changes in the estimation method we use to calculate productivity, holds for each individual country in our sample, and it is robust to changes in the measure of within-industry cycle. In particular, as we show in column (7) of Table II, the skewness of firms' productivity shocks is positively correlated with the average sales growth within a country-industry cell.[18]

We complement our cross-country results using data for a sample of manufacturing establishments in the United States that combines records from the Census of Manufacturing and the Annual Survey of Manufacturing spanning the years 1976 to 2015. We take the measures of log-productivity as reported by the Census and we obtain an estimate from firms' productivity shocks from the residuals of a panel regression as in equation 3. Because Census data only contains information about manufacturing firms, here we divide our establishment sample in 3-digit NAICS cells within a year and we calculate the average and the Kelley skewness of the productivity shock within each industry-year cell. As the bottom panel Figure 6 shows, the skewness of firms' shocks is negative in industries experiencing average productivity declines. Furthermore, regression results shown in column (11) of Table II indicate a positive and statistically significant relation between the industry cycle (measured by within-industry sales growth) and the skewness of productivity shocks.

In conclusion, these results–along with the procyclical skewness of firms stock-returns reported in Section 3.1–indicate that the shocks driving firm growth also have procyclical skewness. This procyclical skewness could be driven, for instance, by rising bankruptcy during recessions, which would generate left-skewed demand shocks (e.g. if a major customer goes

---

[18]For further robustness, in Appendix B.4 we use three additional measures of productivity. In the first, we reestimate the productivity residuals, $z_{i,t}$, by running an OLS panel regression using firm-level data at the country level; Second, we estimate $z_{i,t}$ using the method developed by Olley and Pakes (1996); third, we estimate labor productivity by regressing firm value added on employment and a firm industry fixed effect. All these methods deliver similar results, qualitatively and quantitatively: the skewness of firm-level productivity shocks is strongly procyclical.

bankrupt this will generate a large left-tail shock). Another could be the underlying driving process itself heterogeneously impacts firms–that is, a few firms lose badly in recessions and a few firm gain heavily in booms–which is similar in spirit to the granularity work in Gabaix (2011).

3.5 VAR Estimates of the Impact of a Skewness Shock

To have a preliminary assessment of the impact of changes in the cross-sectional skewness on aggregate economic activity, we estimate a range of monthly VARs using data for the United States from January of 1964 to December of 2015. To make our results comparable to the existing literature, we consider a standard set of variables including–in the following order– the S&P500 stock market index, a measure of stock-market volatility, a measure of firm-level stock market skewness, the Federal Funds Rate, the average of hourly earnings, the consumer price index, the level of hours, the level of employment, and an index of industrial production. We focus on the impact of a change in the skewness of stock market returns on industrial production and employment. The skewness of stock returns–measured by the cross-sectional Kelley skewness of daily returns within a month–is included third in VAR so as to ensure that the impact of first- and second-order moments–proxied by the S&P500 and the 90th-to-10th percentiles differential of stock returns within a month respectively–are pealed out before looking at the impact of a skewness shock.[19]

The upper panel of Figure 7 shows a persistent and economically significant decline of industrial production after a skewness shock: industrial production drops 0.5% four months after the shock, and reaches a declines of 0.8% after a year. Consistent with previous studies, a volatility shock generates 0.5% decline in industrial production after six months. Relative to a skewness shock, however, the economy recovers rapidly seven months after the jump in volatility. The bottom plot of Figure 7 displays similar patterns for aggregate employment. As we show in Appendix C, the significant and persistent decline of economic activity after a skewness shock is remarkably robust to several alternative specifications, variable ordering, measures of skewness, detrending (Figure A.8), or if we estimate the effect of a skewness shock using local projections as in Jordà (2005) (Figure A.9).

In summary, our empirical analysis shows that the skewness of firm-level outcomes declines sharply during recessions, both at the country and at the industry level. Motivated by this robust evidence, in the next section we study a heterogeneous agents model that we use to

---

[19]All variables with the exception of the Federal Funds Rate, the measure of volatility, and the measure of skewness are in logs. All variables, with the exception of the measures of volatility and skewness, are detrended using the Hodrick-Prescott filter with smoothing parameter equal to 129,600. We do not detrend volatility or skewness to facilitate the comparison to the rest of the empirical analysis. As we show in Appendix C, however, considering detrended measures of volatility and skewness does not change our results.

evaluate the macroeconomic importance of the large swings in the skewness we observe in the data.

4   Model

Given the evidence discussed in the previous section, a natural question is to ask whether the observed fluctuations in the skewness of firms' outcomes and productivity have any implication for aggregate economic activity. To answer this question, we build a heterogeneous-agent model populated by a large number of infinitely lived households/entrepreneurs who combine capital and labor using a technology which is subject to aggregate and idiosyncratic productivity shocks to produce a homogeneous good. We make two modeling choices that are important in generating large impacts of skewness shocks, but which we also think are empirically reasonable.

First, entrepreneurs are not able to insure against idiosyncratic shocks, so they are exposed to idiosyncratic risk. This seems plausible–very few businesses are able to insure fully (or even partially) against the risks they face, and since the managers of most firms have significant equity stakes in their businesses, they are exposed to business risk.[20] Second, entrepreneurs are able to save in the form of both the capital good and a one-period bond with a risk-free return. The risk-free return can be thought of as a government or foreign sector, which at least from the perspective of the entrepreneur, provides a return which is independent of idiosyncratic risk. We now describe each component of the model in more detail.

4.1   Entrepreneurs

4.1.1   Production Technology

The production function of entrepreneur j is given by

$$y_{j,t} = A_t e_{j,t} k_{j,t}^{\alpha} n_{j,t}^{\nu}, \text{ with } \alpha + \nu < 1,$$

so it displays decreasing returns to scale. The aggregate productivity shock, A_t, follows the first-order autoregressive process

$$\log A_t = \rho_A \log A_{t-1} + \sigma_\eta \eta_t,$$

where η_t is a Gaussian innovation with zero mean and unitary variance. The idiosyncratic productivity process $e_{j,t}$ is given by

$$e_{j,t} = \rho e_{j,t-1} + \epsilon_{j,t}, \tag{4}$$

---

[20]Even in very large publicly traded firms in the United States, top executives own substantial equity stakes; Furthermore, most private firms are owned by the manager or their family.

where $\epsilon_{j,t}$ has zero mean, time-varying variance, denoted by $\sigma_{\epsilon,t-1}$, and time-varying skewness, denoted by $\gamma_{\epsilon,t-1}$. Notice that we have assumed that the distribution of innovations in period t depends on the values of the variance and skewness observed in period $t-1$. This timing captures the "news shock" aspect of firm-level risks in the model: an increase in dispersion or a decline in the skewness of firms' shocks represents news about the characteristics of the distribution of innovations in the future but not a change in the distribution from which the current realizations of $\epsilon_{j,t}$ are drawn.

4.1.2 Capital Adjustment Costs

We consider a flexible combination of convex and non-convex adjustment costs to capital. To this end, let $i_{j,t}$ denote *net* investment in capital:

$$i_{j,t} = k_{j,t+1} - (1 - \delta) k_{j,t}, \tag{5}$$

where δ is the depreciation rate of capital. Capital adjustment costs are given by the sum of a fixed *disruption cost*, ϕ_1, paid by the entrepreneur for any net investment or disinvestment, a quadratic adjustment cost, ϕ_2, and a *resale cost* for net disinvestment (partial irreversibility), ϕ_3. Therefore, the total adjustment cost function for capital input is

$$\phi\left(k_{j,t+1}, k_{j,t}\right) = \phi_1 \mathbb{I}_{|i_{j,t}|>0} y_{j,t} + \frac{\phi_2}{2} \left(\frac{i_{j,t}}{k_{j,t}}\right)^2 + (1 - \phi_3) |i_{j,t}| \, \mathbb{I}_{i_{j,t}<0}, \tag{6}$$

where $\mathbb{I}$ is an indicator function.

4.1.3 The Problem of the Entrepreneur

Entrepreneurs do not value leisure and value consumption streams according to the standard CRRA utility function specified below. Entrepreneurs supply labor to their own firm (they cannot work for someone else's firm). They can save in capital and in a risk-free asset that pays an interest rate r_t. Denote the entrepreneur's value function by $V\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right)$ where $k_{j,t}$ is the entrepreneur's stock of physical capital, $a_{j,t}$ is the beginning-of-the-period holdings in the risk-free asset, and $e_{j,t}$ is her idiosyncratic productivity.

For notational simplicity, define the vector of aggregates states as $\Omega_t \equiv (A_t, \sigma_{\epsilon,t-1}, \gamma_{\epsilon,t-1}, \mu_t)$ where A_t is the aggregate productivity level, $\sigma_{\epsilon,t-1}$ and $\gamma_{\epsilon,t-1}$ are the variance and the skewness of the distribution of idiosyncratic shock, respectively, and μ_t is the distribution of entrepreneurs

15

over idiosyncratic states. Then, we can write the dynamic problem of the entrepreneur as

$$V\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right) = \max_{\substack{\{c_{j,t}, k_{j,t+1}, \\ a_{j,t+1}, n_{j,t}\}}} \left\{\frac{c_{j,t}^{1-\xi}}{1-\xi} + \beta\mathbb{E}V\left(k_{j,t+1}, a_{j,t+1}, e_{j,t+1}; \Omega_{t+1}\right)\right\}, \quad (7)$$

$$\text{s.t. } c_{j,t} + i_{j,t} + a_{j,t+1} \leq y_{j,t} - w_t\left(\Omega_t\right)n_{j,t} - \phi\left(k_{j,t+1}, k_{j,t}\right) + \left(1 + r_t\left(\Omega_t\right)\right)a_{i,t},$$

$$i_{j,t} = k_{j,t+1} - (1-\delta)k_{j,t},$$

$$\mu_{t+1}\left(k_{j,t+1}, a_{j,t+1}, e_{j,t+1}\right) = \Gamma\left(\Omega_t\right),$$

$$k_{j,t} > 0, a_{j,t} \geq 0, n_{j,t} > 0,$$

given the laws of motion for A_t, $\sigma_{\epsilon,t}$, and $\gamma_{\epsilon,t}$. The term $w_t \equiv w\left(\Omega_t\right)$ denotes the wage rate in the economy. In what follows, we assume the interest rate on the risk-free asset is fixed, that is $r\left(\Omega_t\right) = r$.[21] Let $C^e\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right)$, $K^e\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right)$, $N^e\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right)$, and $A^e\left(k_{j,t}, a_{j,t}, e_{j,t}; \Omega_t\right)$ denote the policy rules for consumption, next period capital, current period labor, and the risk-free asset for the entrepreneurs.

4.2 Non-Entrepreneurial Households

The economy is populated by a large number of identical hand-to-mouth households that consume C_t units of the homogeneous good and supply labor elastically which we denote by N_t. More concretely, the non-entrepreneurial households solve the static problem

$$U\left(C_t, N_t\right) = \max_{C_t, N_t} \left\{\frac{C_t^{1-\sigma}}{1-\sigma} - \psi\frac{N_t^{1-\gamma}}{1-\gamma}\right\}, \quad (8)$$

$$C_t \leq w_t N_t,$$

given the law of motion of the aggregate state, Ω_t. Denote by $C\left(\Omega_t\right)$ and $N\left(\Omega_t\right)$ the optimal choices of consumption and labor for the non-entrepreneurial household.

4.3 Recursive Competitive Equilibrium

Given the exogenous process for aggregate productivity, A, the exogenous process of the variance and skewness of e_j, the interest rate of the risk-free asset, r, and the evolution of the idiosyncratic productivity processes for the entrepreneurs, $\{e_j\}_{j\in J}$, a recursive competitive equilibrium for this economy is a set of policy functions $\left\{\left\{C_j^e, K_j^e, N_j^e, A_{j,}^e\right\}_{j\in J}, C, N\right\}_{t=0}^{\infty}$, a wage function $\{w\}$, and value functions $\{V, U\}$ such that

---

[21]This implies that we will not solve the interest rate in equilibrium. The wage rate, however, is such that the labor market clears in each period.

i) the policy and value functions solve (7) and (8), respectively; ii) the labor market clears, that is,

$$\int N^e(k_j, a_j, e_j; \Omega) \, d\mu(k_j, a_j, e_j) = N(\Omega);$$

and iii) the mapping $\Gamma(\omega)$ that determines the evolution of the joint distribution of e_j, k_j, and a_j is consistent with the policy functions, the evolution of the aggregate productivity process, and the evolution of the process of σ_ϵ and γ_ϵ.

4.4 Parameters and Estimation

In this section, we describe the quantitative specification of our modeled economy. To solve the entrepreneurs' problem, we employ non-linear methods similar to Krusell and Smith (1998). Most of our parameters are standard in the macro literature, and we take them from the existing estimates when possible. However, the parameters governing the stochastic process of firms' productivity are novel to our analysis, and we use a simulated method of moments approach to estimate them.

Frequency, Preferences, and Aggregate Productivity

We set the time period to be a quarter. For the entrepreneurs, we set ξ equal to 6.0. The household's discount rate, β, is set to $0.95^{0.25}$, whereas the interest rate on the risk-free asset is set to match an annual return of 2%. For the non-entrepreneurial sector, we set σ to 2. For the labor supply of the non-entrepreneurial households, we fix a value of γ to 1.5, and we choose ψ so that they spend an average of 33% of their time working.

The exponents of the capital and labor inputs in the entrepreneur's technology are set to $\alpha = 0.25$ and $\nu = 0.5$. The capital depreciation rate, δ, is set to match an annual depreciation of 14%. As for the adjustment cost parameters, we set the fixed adjustment cost of capital, ϕ_1, equal to 1.5%, a quadratic adjustment cost, ϕ_2, equal to 7.0, and a resale cost, ϕ_3, equal to 34.0% (Bloom, 2009).

We assume that aggregate productivity follows a standard first-order autoregressive process with an autocorrelation of 0.95 and normally distributed innovations with mean 0 and a standard deviation of 0.75%, similar to the quarterly values used in other papers in the literature (Khan and Thomas, 2008). The top panels of Table III summarize the set of calibrated parameters.

Idiosyncratic Productivity

To capture time-varying risk, we assume that the economy transitions between two states. The first is a low-risk state (denoted by L), which corresponds to periods in which the variance of the innovations of the idiosyncratic shocks is low and the skewness is positive, as we observe in

non-recession periods. The second state is a high-risk state (denoted by H), which corresponds to periods in which the variance of the innovations of the idiosyncratic shocks is high and the skewness is negative, as we observe during a typical recession. Low- and high-risk states alternate following a first-order Markov process.

Since high and low risk periods differ in the skewness of measured productivity, we need to depart from the standard assumption of Gaussian shocks. Although there are several alternatives to model idiosyncratic shocks with time-varying higher order moments, here we take a simple approach and we assume that, conditional on the risk state of the economy, the innovations of the firms' idiosyncratic productivity process, $\epsilon_{j,t}$, are drawn from a mixture of two normally distributed random variables, that is,

$$
\epsilon_{j,t} \sim \begin{cases} N\left(\mu^s, \sigma_1^s\right) & \text{with prob } p^s \\ N\left(-\frac{p^s}{1-p^s}\mu^s, \sigma_2^s\right) & \text{with prob } 1-p^s, \end{cases} \tag{9}
$$

where s denotes the risk state of the economy, $s \in \{H, L\}$. Hence, to fully characterize the stochastic process faced by firms, we need to find ten parameters, namely, $\{\mu^s, \sigma_1^s, \sigma_2^s, p^s\}$ with $s \in \{H, S\}$, and the parameters governing the transition probabilities between low- and high-risk periods, denoted by π_L and π_H, respectively.[22]

Since we do not directly observe the productivity process faced by the firms, we choose the parameters of the stochastic process of firms' productivity to match the features of the US data described in the empirical section of the paper. In particular, we take the data of quarterly sales growth from Compustat, and we search for parameters of the stochastic process so that the cross-sectional distribution of sales growth derived from the model reproduces the observed average values of the 90th-to-50th log percentiles differential, the 50th-to-10th log percentiles differential, the Kelley skewness, and the 90th-to-10th log percentiles differential of the quarterly sales growth distribution during expansion periods and the same set of moments for recession periods for a total of eight moments.[23] The probability of being in the high-risk state in the next period conditional on being in the high-risk state in this period, π_H, is set to be equal to the fraction of recession quarters that are followed from another recession quarter in the data, $\pi_H = 0.84$, whereas the transition probability of the low risk state, π_L, is set so that the share of expansion quarters following another expansion quarter is 0.95. Recession and expansion

periods in the data correspond to the recession quarters defined by the NBER from 1970 to 2014.

Based on our estimations, we find that in periods of low risk, the variance of the idiosyncratic productivity shocks, η, is equal to 4.85%, whereas the coefficient of skewness is equal to 0.85. In contrast, in periods of high-risk, the variance of the productivity shocks is equal to 6.85%, and the coefficient of skewness is negative and equal to -1.14. The bottom panels of Table III displays our estimates for the different parameters of the idiosyncratic productivity process, whereas Table IV shows the targeted and model-simulated moments.[24] Furthermore, our model is consistent with the standard business cycle statistics in terms of the cyclicality and volatility of aggregate output, consumption, investment, and employment (see Table A.4 in the Appendix).

5   Quantitative Results

In this section we evaluate the macroeconomic impact of a change in the skewness of the distribution of firms shocks. We start by showing that our model is able to separate a skewness shock from a mean and variance shock, and we then evaluate the impact of the change in skewness in terms of output and other macroeconomic aggregates.

5.1   Idiosyncratic Shocks and Model Fit

To evaluate the effects of a decrease in the skewness of firm-level shocks, we independently simulate 1,000 economies, each of 300 quarters' length. For the first 150 periods, the economy remains in the low-risk state, and then all economies are hit by a change in the level of risk (i.e. a decrease in the skewness of firm-level shocks, an increase in the dispersion of firm-level shocks, or both at the same time). From that period on, all economies evolve normally. We then average different macroeconomic outcomes across all simulated economies and calculate the impact of the change in risk as the log percentage deviation of a given macro variable relative to its value in the period previous to the shock.

We begin by analyzing the response of the distribution of firm productivity growth after a change in aggregate risk. The left panels of Figure 8 display moments of the distribution of firms' idiosyncratic productivity growth, $\Delta e_{j,t} = e_{j,t} - e_{j,t-4}$, for three cases. In the first case, the economy moves from the low-risk state to the high-risk state, leading to an increase in the variance and a decrease in the skewness of idiosyncratic shocks (blue line with circles),

---

[24]The variance of a random variable η, which is distributed as a mixture of two normally distributed random variables, is given by $Var(\eta) = \mathbb{E}(\eta^2) - \mathbb{E}(\eta)^2$, whereas the skewness is given by $Skew(\eta) = \left(\mathbb{E}(\eta^3) - 3\mathbb{E}(\eta)Var(\eta) - \mathbb{E}(\eta)^3\right)/Var(\eta)^{\frac{3}{2}}$. Here $\mathbb{E}(\eta)$ is the first moment of the η given by $\mathbb{E}(\eta) = p_1\mu_1 + p_2\mu_2$. Similarly, $\mathbb{E}(\eta)^2 = p_1\left(\mu_1^2 + \sigma_1^2\right) + p_2\left(\mu_2^2 + \sigma_2^2\right)$ and $\mathbb{E}(\eta^3) = p_1\left(\mu_1^3 + 3\mu_1\sigma_1^2\right) + p_2\left(\mu_2^3 + 3\mu_2\sigma_2^2\right)$ are the second and third moments.

which corresponds to what it is observed during a typical recession. In the second case, the increase in risk leads only to a decrease in the skewness of idiosyncratic shocks (black line with diamonds), and finally, in the third case, the increase in risk leads to an increase in the variance of idiosyncratic shocks only, which is the typical uncertainty shock studied in the literature (red line with triangles).[25] The top left panel of Figure 8 shows that the average firm in our model does not experience a change in productivity when risk changes. This ensures that our results are not driven by a change in average productivity and are driven solely by changes in the shape of the distribution of productivity shocks. Then, comparing the black line in the middle and bottom left panels, one can see that our model is able to generate a pure change in the skewness, that is, a change in the productivity distribution that reflects only a decrease in the skewness but a muted change in the mean and the variance of the firm-level productivity distribution.[26] Similarly, our model can generate a pure uncertainty shock (the red line with triangles in the middle panels of Figure 8).

We now analyze the response of the sales growth distribution–our empirical target–to a change the dispersion and skewness of firms' shocks. The right panels of Figure 8 show the average, the dispersion, and the skewness of the annual change in quarterly sales implied by the model calculated as $\Delta y_{j,t} = \log y_{j,t} - \log y_{j,t-4}$. It is not surprising that a change in risk that combines a simultaneous increase in the variance and a decrease in the skewness of firm-level productivity shocks generates an increase in the cross-sectional dispersion of sales growth and a large decrease in skewness (blue line with circles in the middle and bottom right panels). Comparing the case in which only dispersion changes—which is the typical uncertainty shock—with the case in which only the skewness changes—the baseline case we discuss in the following section—one can see that by considering a shock with time-varying skewness, the model is able to capture the asymmetric response of the tails of the sales growth distribution (compare the red line with triangles to the blue line with circles in the bottom right panel).[27]

5.2   The Macroeconomic Effect of a Skewness Shock

We now study the macroeconomic effect of a decrease in the skewness of firm-level productivity. For doing that, we shock the economy with a change in the skewness of the innovations of $e_{j,t}$ and calculate the response of different macroeconomic aggregates as the percentage change

---

[25]To make this comparison, we reestimate the parameters of the stochastic process in (9) to separate the changes in dispersion (a symmetric increase in risk) from changes in dispersion and skewness (an asymmetric increase in risk). Appendix Table A.5 shows the estimation targets for each case.

[26]The median firm, however, experiences an increase in productivity after a decline in the skewness. Disentangling the mean and the median of the distribution of firms' shocks allow us to keep the mean and variance constant after a change in skewness.

[27]As we show in Appendix Figure A.7b, the Kelley skewness of the employment growth distribution also declines after the drop in the skewness of firms' shocks. Importantly, as we show in Figure A.7a in the Appendix, the dispersion and the skewness of sales growth do not change after a decline in aggregate productivity, A_t.

relative to their value prior to the shock. In our exercise, when the economy receives a skewness shock that drives the skewness of idiosyncratic shocks from positive to negative, we keep the mean and variance of the idiosyncratic productivity process constant at their low-risk level, so our results reflect a pure change in the skewness of the distribution of firms' shocks.

Figure 10a shows that output declines by 1.4% four quarters after a skewness shock and 1.7% after eight quarters. This is a significant decline in aggregate economic activity considering that only the shape of the distribution of firm-level shocks has changed. Moreover, the decline in output is quite persistent, staying below its pre-shock level even after twelve periods after the shock. This is in contrast with the typical uncertainty shock that generates a decrease in output and a rapid rebound a few quarters after the shock. In our model, the drop in output is generated by the rapid and persistent decline in capital investment after a change in skewness. The top right panel of Figure 10b shows that capital investment drops around 15% during the first quarter after the shock and stays below its pre-shock level for several quarters. Labor does not drop in the first period after the shock because labor input is fully flexible and news about the future conditions of risk does not change firms' hiring decisions.[28] In contrast, consumption declines rapidly in response to the decrease in the skewness of firm-level shocks, dropping by around 1% relative to its pre-shock level, whereas the accumulation of the risk-free asset increases because productive capital is now riskier.

Importantly, in the first quarter after the shock, the response of investment and consumption is not driven by a change in the skewness of the realizations of $e_{j,t}$ received by the firms—recall our timing assumption in equation (4)—but by a change in the perception about the risk in the economy: at the moment of the shock, entrepreneurs receive news that in the future the distribution of $e_{j,t}$ will be left-skewed, and their endogenous responses drive a decline in investment and consumption. A decrease in skewness triggers a precautionary increase in entrepreneurs' savings, but since capital is riskier, investment in the risk-free asset surges, as shown in the bottom right panel of Figure 10a. We conclude that a decline in the skewness of the distribution of idiosyncratic shocks can by itself generate a persistent drop in aggregate economic activity.

5.3 Skewness and Uncertainty Shocks

The results shown in the previous section trace the macroeconomic impact of a change in the skewness of the distribution of firm's shocks while keeping the variance of these shocks constant. As a consequence, the dispersion of the sales growth distribution generated by our model remains more or less invariant after a skewness shock (see center panel of Figure 8). Our empirical evidence, however, showed that recessions are characterized by an increased in

---

[28]Adding labor adjustment costs will trigger an automatic response of labor to changes in risk, increasing the aggregate impact of a change in variance and skewness.

dispersion paired with a sharp decline in the skewness of the distribution of firms outcomes and productivity. Hence, in this section, we analyze the joint impact of an increase in dispersion and a decline in the skewness as we see a typical recession.

As Figure 10 indicates (blue line with squares), the joint impact of a variance shock and a skewness shock magnifies and accelerates the impact of output relative to the impact of a pure skewness shock: output, in this case, declines up to 2% four quarters after the shock. This additional decline in output is explained by a larger decline in investment and consumption, and a surge in the investment in the risk-free asset. The combined effect of variance and skewness accelerates the recovery after the shock as output starts to recover rapidly six quarters after the shock. Hence, our results suggest that a joint change of the dispersion and skewness of firm's productivity shock–which is consistent with the observed changes in dispersion and skewness of firm-level outcomes–can generate aggregate dynamics that are similar to what is observed in a typical recession.

5.4 Understanding the Macroeconomic Impact of Skewness Shocks

How do the model ingredients interact with a change in the skewness of firms' productivity shocks? To answer this, in this section we perform a series of experiments, changing different parameters or assumptions in the model–while keeping all other parameters at their baseline level–in order to isolate their contribution to the results discussed in the previous section.

News Shock

In our baseline results, in the period in which a change in risk occurs, firms do not experience a change in the actual realizations of shocks but only receive news that in the next period, the skewness of productivity shocks will be lower. After that, firms' productivity distribution changes as the shocks are drawn from a left-skewed distribution. We compare this baseline case to one in which we keep the underlying distribution of firms shocks fixed so that we can evaluate the pure effect of a change in news about the future risk conditions.[29] This exercise is similar to an increase in the probability of a *disaster* (Gourio, 2008; Barro and Ursua, 2011), although in our case it represents an increase in *disasters* at the microeconomic level. The red line with triangles in Figure 10 shows that a shock that only represents news about the future skewness of the distribution of firms' shocks generates a decline in output of about 0.5%, which is around one-third of the overall decline in our baseline results. The first-period impact on

---

[29]In particular, we simulate our model using the same realizations of the aggregate risk process used in our baseline analysis. In period T all economies receive a skewness shock, however, in this case, we keep the parameters determining the actual underlying idiosyncratic productivity process fixed at their pre-shock low-risk level. While this violates rationality–firms expect more skewed shocks but this never arises–it is a useful device for distinguishing the expectation and the realization impacts of a skewness shock.

investment in capital and in the risk-free asset is the same as in the baseline results as these are forward-looking variables that rapidly react to future risk conditions.

Adjustment Costs and Risk Aversion

In our model, capital adjustment costs and risk aversion play an important role in the propagation of aggregate and idiosyncratic shocks. On the one hand, fixed adjustment costs to capital create inaction regions that expand during periods of high uncertainty making firms more cautions and freezing investment (Bloom, 2009). On the other hand, risk-averse entrepreneurs might prefer to reduce the size of their firms as physical capital becomes riskier when skewness drops, and therefore invest a larger fraction of their wealth in the risk-free asset, further reducing capital investment. In order to quantify the relative importance of these two channels, in this section we compare the impact of an increase in risk in two economies, one in which we maintain the level of risk aversion of entrepreneurs as in the baseline case but we allow flexible capital investment by shutting down all adjustment costs; and a second economy in which adjustment costs are as in the baseline case but entrepreneurs value consumption using a risk-neutral linear utility function.

Figure 10 shows the response of different macro aggregates after a decrease in skewness when we shut down the adjustment costs. In this case, the output response is stronger and steeper relative to the baseline case (compare the line with + symbols to baseline in the top left panel of Figure 10). Hence, capital adjustment costs dampen the impact of a skewness shock. The decline in output is driven by a decline in capital investment, which drops as entrepreneurs scale back their firms in response to a decline in the skewness of shocks, moving their wealth into the risk-free asset as the bottom right panel of Figure 10 shows.

When we consider entrepreneurs with linear utility–basically making them risk-neutral–the impact of a decrease in the skewness of firms' shocks generates a much larger decrease in output (compare the line with x symbols to the baseline results in Figure 10). Consumption, however, increases after the shock as entrepreneurs reduce their capital investment due to the increase in the inaction regions generated by the adjustment costs and the increase in risk. Investment in the risk-free asset remains almost unaltered since risk-neutral individuals invest most of their wealth in their firm. The increase in aggregate consumption is counterfactual as a typical recession is characterized by a concurrent decrease in skewness and aggregate consumption. Hence, we conclude that to obtain plausible business cycle fluctuations stemming from a decrease in the skewness of firms' productivity shocks, a combination of risk-averse entrepreneurs and adjustment costs is appropriate.

Changes in the Returns of the Risk-Free Asset

Throughout our analysis, we have assumed that the interest rate of the risk-free asset is fixed and does not respond to aggregate economic conditions. As we showed in the previous section, entrepreneurs respond to a skewness shock by reducing their position in the risky asset–physical capital–and increasing their investment in the risk-free asset. This means that our quantitative results could change if the interest rate of the risk-free asset drops enough to counteract the aggregate impact of a skewness shock.

One way to evaluate whether changes in the interest rate have a large quantitative impact on our results is to fully endogenize the interest rate of the risk-free asset. This comes at the additional cost of having to solve a larger general equilibrium loop. Instead, we take a simpler route. In particular, we consider a pseudo-GE case in which the annual return of the risk-free asset declines by 110 basis points when the economy is at the high-risk state (high variance and negative skewness). Importantly, this is fully incorporated in the solution of the entrepreneurs' problem as they correctly predict that during periods of high aggregate risk, the interest rate of the risky-free asset is lower.

The line with circles in Figure 10 displays the evolution of different macroeconomic aggregates after a risk shock that decreases the skewness of productivity shocks paired with a decline in the interest rate of the risk-free asset. Interestingly, the concurrent decrease in the interest rate has little impact on the overall drop in aggregate economic activity generated by a skewness shock, although the results move in the expected direction with capital investment declining less and investment in the risk-free asset increasing less relative to the baseline case. There are two factors that explain the rather small relative impact of the decline in the interest rate on our baseline results. The first is the relatively high risk aversion of entrepreneurs, which combined with the large swings of the distribution of the productivity shocks, generates strong incentives for the entrepreneurs to move their wealth into the risk-free asset, despite its lower returns. The second relates to the fixed adjustment costs. It is well established that fixed adjustment costs generate regions of inaction that increase after an uncertainty shock. Therefore, after a decrease in skewness, entrepreneurs become less responsive to changes in the interest rate. Hence, we conclude that changes in the interest rate do not have a large impact on our quantitative results.

6   Conclusions

This paper studies how the distribution of the growth rate of firm-level variables changes over the business cycle. Using firm-level panel data for the United States from Census and non-Census datasets and firm-level panel data for almost forty other countries, we reach three

main conclusions. First, recessions are characterized by a large drop in the skewness of firm-level employment growth, sales growth, productivity growth, and stock returns. Second, the decline in the skewness of firms' outcomes is a phenomenon observed not only in the United States but also in other countries, both developed and developing. Finally, we also find strong procyclicality of the skewness at the industry level, in that industries with low and high growth rates experience a negative and positive skewness of shocks, respectively.

In the second part of our paper, we analyze the impact of a change in the skewness of firms' idiosyncratic productivity in the context of a heterogeneous-agent model. We assume that the exogenous idiosyncratic productivity process faced by entrepreneurs is subject to time-varying skewness, and we choose the parameters of this model to match the evolution of the dispersion and skewness of the sales growth distribution in the United States. Our results suggest that a change in the skewness of the firm-level productivity distribution can by itself generate a significant decline in aggregate economic activity (even though the mean and variance of firms' shocks are held constant). In fact, in our modeled economy, a decline in the skewness of firms' shocks of the magnitude observed in a typical US recession generates a drop in GDP of 1.7%. The combined impact of a variance and a skewness shock generates an even large decline in output (-2.0%), consumption (-2.0%), and investment (-40.0%), suggesting higher moment micro-shocks can play a major role in explaining business cycle dynamics.

References

Alfaro, I., Bloom, N. and Lin, X. (2018). *The Finance Uncertainty Multiplier.* Working Paper 24571, National Bureau of Economic Research.

Arellano, C., Bai, Y. and Kehoe, P. (2018). Financial markets and fluctuations in uncertainty. *Federal Reserve Bank of Minneapolis Staff Report.*

Bachmann, R. and Bayer, C. (2013). 'wait-and-see' business cycles? *Journal of Monetary Economics*, **60** (6), 704–719.

— and — (2014). Investment Dispersion and the Business Cycle. *American Economic Review*, **104** (4), 1392–1416.

Barro, R. and Ursua, J. (2011). *Rare Macroeconomic Disasters.* NBER Working Papers 17328, National Bureau of Economic Research, Inc.

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

Basu, S. and Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, **85** (3), 937–958.

Berger, D., Dew-Becker, I. and Giglio, S. (2017). *Uncertainty shocks as second-moment news shocks.* Working Paper 23796, National Bureau of Economic Research.

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.

— (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, **28** (2), 153–76.

—, Floetotto, M., Jaimovich, N., Saporta-Eksten, I. and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, **86** (3), 1031–1065.

Busch, C., Domeij, D., Guvenen, F. and Madeira, R. (2017). *Asymmetric Business Cycle Risk and Government Policy.* Working paper, University of Minnesota.

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., Miranda, J., Foote, C. and Nagypal, E. (2006). Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. *NBER Macroeconomics Annual*, **21**, 107–179.

Decker, R., Haltiwanger, J., Jarmin, R. and Miranda, J. (2015). *Where has all the skewness gone? The decline in highgrowth (young) firms in the US.* Working Paper 21776, National Bureau of Economic Research.

Fernandez-Villaverde, J., Guerron-Quintana, P., Rubio-Ramirez, J. F. and Uribe, M. (2011). Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review*, **101** (6), 2530–2561.

Ferreira, T. (2018). Stock market cross-sectional skewness and business cycle fluctuations.

Gabaix, X. (2008). Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macrofinance. *American Economic Review*, **98** (2), 64–67.

— (2011). The granular origins of macroeconomic fluctuations. *Econometrica*, **79** (3), 733–772.

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

GILCHRIST, S., SIM, J. W. and ZAKRAJŠEK, E. (2014). *Uncertainty, financial frictions, and investment dynamics*. Tech. Rep. 20038, National Bureau of Economic Research.

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.

— (2013). Credit risk and disaster risk. *American Economic Journal: Macroeconomics*, **5** (3), 1–34.

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

HARMENBERG, K. and SIEVERTSEN, H. H. (2017). *The Labor-Market Origins of Cyclical Income Risk*. Working paper, Institute for International Economic Studies.

HARVEY, C. R. and SIDDIQUE, A. (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, **55** (3), 1263–1295.

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.

JORDÀ, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, **95** (1), 161–182.

JURADO, K., LUDVIGSON, S. C. and NG, S. (2015). Measuring uncertainty. *American Economic Review*, **105** (3), 1177–1216.

KALEMLI-OZCAN, S., SORENSEN, B., VILLEGAS-SANCHEZ, C., VOLOSOVYCH, V. and YESILTAS, S. (2015). *How to construct nationally representative firm level data from the ORBIS global database*. Tech. rep., National Bureau of Economic Research.

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

KELLEY, T. L. (1947). *Fundamentals of Statistics*. Harvard University Press.

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.

KILIC, M. and WACHTER, J. A. (2015). *Risk, Unemployment, and the Stock Market: A Rare-Event-Based Explanation of Labor Market Volatility*. NBER Working Paper 21575.

KIM, T.-H. and WHITE, H. (2004). On more robust estimation of skewness and kurtosis. *Finance Research Letters*, **1** (1), 56–73.

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. (2016). The tail that wags the economy: Belief-driven business cycles and persistent stagnation.

—, — and — (2018). *The tail that keeps the riskless rate low*. Tech. rep., National Bureau of Economic Research.

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

LEDUC, S. and LIU, Z. (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, **82**, 20–35.

MCKAY, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics*, **88**, 1–14.

OH, S. and WACHTER, J. A. (2018). *Cross-sectional Skewness*. Working paper, Wharton School at University of Pennsylvania.

OLLEY, G. S. and PAKES, A. (1996). The dynamics of productivity in the telecommunications equipment. *Econometrica*, **64** (6), 1263–1297.

ORLIK, A. and VELDKAMP, L. (2014). *Understanding uncertainty shocks and the role of black swans*. Tech. Rep. 20445, National Bureau of Economic Research.

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

SCHAAL, E. (2017). Uncertainty and unemployment. *Econometrica*, **85** (6), 1675–1721.

SYVERSON, C. (2011). What determines productivity? *Journal of Economic literature*, **49** (2), 326–65.

VENKATESWARAN, V., VELDKAMP, L. and KOZLOWSKI, J. (2015). *The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation*. 2015 Meeting Papers 800, Society for Economic Dynamics.

WACHTER, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, **68** (3), 987–1035.

YASAR, M., RACIBORSKI, R. and POI, B. (2008). Production function estimation in stata using the olley and pakes method. *The Stata Journal*, **8** (2), 221–231.

FIGURE 1 – THE SKEWNESS OF FIRM OUTCOMES IS LOWER DURING RECESSIONS

(A) Census LBD: Log Employment Growth
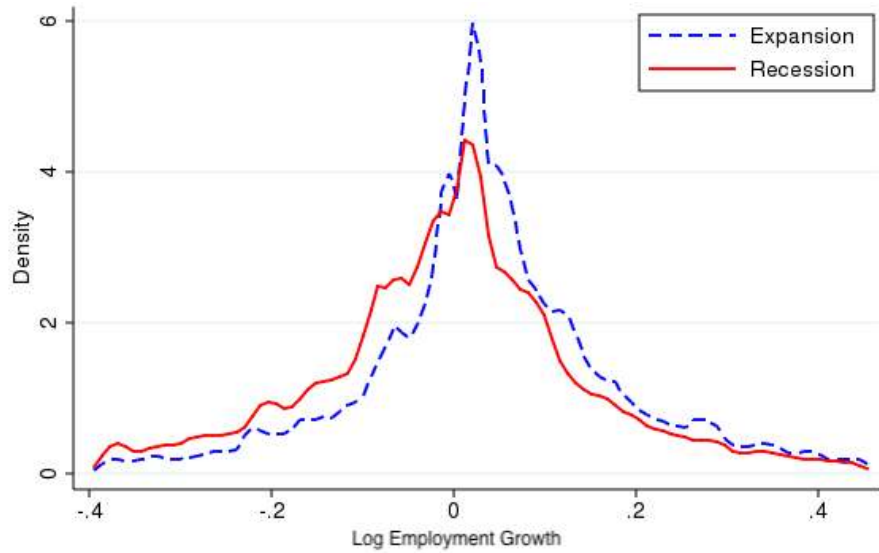
(B) Compustat: Log Sales Growth

Note: The top panel of Figure 1 shows the employment-weighted empirical density of the distribution of firms' log employment growth between years t and $t + 1$ constructed from the LBD. The lower panel shows the empirical density of the distribution of firms' log sales growth between years t and $t + 1$ constructed from Compustat. Each density has been rescaled to have a median of zero and unitary variance. The blue-dashed line shows the density of a pooled sample of Expansion years (2003 to 2006 and 2010 to 2014) whereas the red-solid line shows the density of a pooled sample of Recession years (2001 and 2008). In the top panel, the unscaled 10th percentile of the log employment growth distribution during expansion (recession) periods is -0.17 (-0.27), the 50th is 0.01 (-0.02), and the 90th is 0.23 (0.18). In the bottom panel, the corresponding moments are -0.22 (-0.47), 0.05 (-0.03), and 0.45 (0.33). See Appendix B for additional details on the sample construction and moment calculations in the LBD and Compustat.

FIGURE 2 – THE SKEWNESS OF FIRM-LEVEL OUTCOMES IS PROCYCLICAL

(A) Census LBD: Skewness of Log Employment Growth

(B) Compustat: Skewness of Log Sales Growth

Note: The top panel of Figure 2 shows the time series of the employment-weighted cross-sectional Kelley skewness of the distribution of firms' log employment growth between years t and $t+1$ constructed from the LBD. The bottom panel shows the time series of the cross-sectional Kelley skewness of the distribution of firms' log sales growth between years t and $t+1$ constructed from Compustat. Shaded areas represent the share of the year (in quarters) declared as a recession by the NBER. See Appendix B for details on the sample construction and moment calculations in the LBD and Compustat.

FIGURE 3 – THE DISPERSION OF LEFT TAIL OF FIRM-LEVEL OUTCOMES IS COUNTERCYCLICAL

(A) Census LBD: Dispersion of Log Employment Growth

(B) Compustat: Dispersion of Log Sales Growth

Note: The top panel of Figure 3 shows the time series of the cross-sectional dispersion of the distribution of firms' log employment growth between years t and $t+1$ constructed from the LBD. The bottom panel shows the time series of the cross-sectional dispersion of the distribution of firms' annual log sales growth between years t and $t+1$ constructed from Compustat. Shaded areas represent the share of the year (in quarters) declared as a recession by the NBER. See Appendix B for details on the sample construction and moment calculations in the LBD and Compustat.

31

FIGURE 4 – THE SKEWNESS OF FIRM-LEVEL OUTCOMES IS PROCYCLICAL WITHIN COUNTRIES

(A) BvD Osiris: Cross-Country Log Sales Growth

(B) BvD Osiris: Log Employment Growth and Log Sales Growth by Country-Year


Note: The top panel of Figure 4 shows the empirical density of firms' log sales growth in US dollars between years t and $t+1$ constructed from the BvD Osiris dataset. Each density has been rescaled to have a median of zero and unitary variance. The red-solid line is the empirical density over all the observations of firms during recession years, defined as years in which the country is in the first decile of the country-specific distribution of the growth rate of GDP per capita (74K firm-year observations). The blue-dashed line is the empirical density over all the observations of firms during expansion periods (523K firm-year observations) which are years not classified as recessions. The unscaled 10th percentile of the sales growth distribution during expansion (recession) periods is -0.31 (-0.42), the 50th percentile is 0.06 (0.00), and the 90th percentile is 0.51 (0.44). The bottom left (right) panel displays a scatter plot showing the relation between the within-country average firm log employment growth between years t and $t+1$ (log sales growth) and the within-country Kelley skewness of firms' log employment growth between years t and $t+1$ (log sales growth) constructed from the BvD Osiris dataset. The regression slope is equal to 1.59 (0.50) which is significant to the 1%. Scatter plots controlling for time and country fixed effects. See Appendix B for details on the sample construction and moment calculations.

32

FIGURE 5 – THE SKEWNESS OF FIRM-LEVEL OUTCOMES IS PROCYCLICAL WITHIN INDUSTRIES

(A) Census LBD: Log Employment Growth by Industry-Year


(B) Compustat: Log Sales Growth by Industry-Year

Note: The top panel of Figure 5 displays a scatter plot showing the relation between the within-industry business cycle, measured by the average growth rate of employment, and the within-industry skewness, measured by the Kelley skewness of firms' employment growth constructed from the LBD. Each dot is a quantile of the industry-year distribution of the average log employment growth, where an industry is defined by a two-digit NAICS group. scatter plots controlling for industry and time fixed effects. The slope coefficient is equal to 1.99 which is statistically significant at the 1%. The bottom panel shows the same statistics for sales growth distribution from Compustat. The slope coefficient is 1.33 which is statistically significant at the 1%. See Appendix B for details on the sample construction and moment calculations in the LBD and Compustat.

FIGURE 6 – THE SKEWNESS OF PRODUCTIVITY SHOCKS IS PROCYCLICAL WITHIN INDUSTRIES

(A) BvD Amadeus: Firm-Level Productivity Shocks by Industry-Year

(B) Census ASM: Establishment-Level Productivity Shocks by Industry-Year
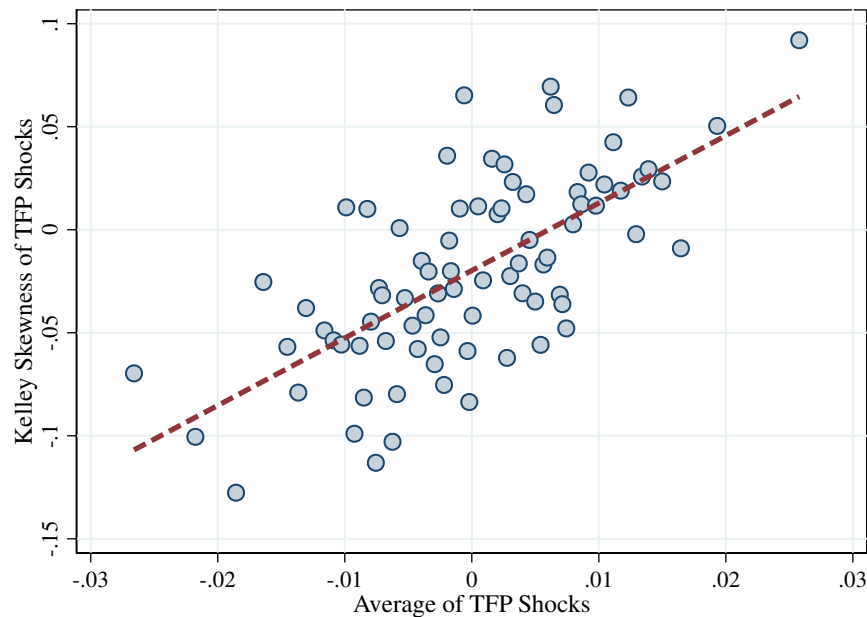
Note: The top panel of Figure 6 displays a scatter plot showing the relation between the skewness of firms' productivity shocks within a country-industry-year cell, and the average productivity shock within the same cell constructed from the BvD Amadeus dataset. Each dot is a quantile of the industry-year distribution of average TFP shocks, where an industry is defined by a two-digit NAICS group. To reduce the impact of outliers on our estimates, we winsorize the distribution of Kelley skewness and average growth at the top and bottom 0.05. The slope coefficient is 1.43 which is statistically significant at the 1%. The bottom panel displays similar statistics calculated from the US Census data for a sample of manufacturing establishments. The slope coefficient is 3.27 which is statistically significant at the 1%. Binscatter plots controlling for industry, country, and time fixed effects. See appendix B.4 for additional details on the sample selection and the TFP estimation procedure.

FIGURE 7 – MACROECONOMIC IMPACT OF A SKEWNESS SHOCK

(A) Industrial Production

(B) Employment

Note: The top panel of Figure 7 shows the impact of a shock to the skewness of daily stock returns of two standard deviations (line with squares) and the impact of a shock to the volatility of daily stock returns (line with circles). Dashed lines show the corresponding 95% confidence intervals. The skewness (volatility) is measured as the Kelley skewness (90th-to-10th log percentiles differential). The standard deviation of the time-series of the Kelley skewness (90th-to-10th log percentiles differential) is 9.5% (1.95%). The bottom panel of Figure 7 shows the impact on aggregate employment. See appendix C for details on the data and robustness.

Table I – Data and Sample Characteristics

Source	Country	Sample Period	Frequency	Comments
Census Longitudinal Business Database	United States	1976-2015	Annual	Employment for entire nonfarm private sector
Census and Annual Survey of Manufacturing Firms	United States	1976-2015	Annual	Employment, sales, and TFP for manufacturing establishments
Compustat	United States	1970-2017	Quarterly	Employment, sales, and stock prices for publicly traded firms
BvD Osiris	Several Countries	1991-2015	Annual	Employment and sales for publicly traded firms in 39 countries
BvD Amadeus	Several Countries	1996-2018	Annual	Sales, employment, and TFP for firms in 21 countries
Global Compustat	Several Countries	1985-2017	Daily	Stock Prices for firms in 29 countries

Note: Table I shows the list of datasets and time-frame used in the analysis. Data availability varies country by country. See Table A.6 in the Appendix for a complete list of countries and available data.

TABLE II – THE SKEWNESS OF FIRMS OUTCOMES IS LOWER DURING RECESSIONS AND RISES IN EXPANSIONS

Dependent Variable:				Kelley Skewness of Log Growth of Firms Outcomes							
Sample:	United States			Cross-Country				Cross-Industry			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Outcome:	Emp.	Sales	Returns	Emp.	Sales	Returns	TFP	Emp.	Sales	Returns	TFP
$\Delta GDP_{i,t}$	4.64***	5.40***	2.09**	5.94***	3.10***	2.16**					
	(1.45)	(1.06)	(1.03)	(1.36)	(1.07)	(0.87)					
$\Delta S_{j,t}$							1.14***	6.43***	13.44***	1.36**	2.20***
							(0.35)	(1.16)	(1.41)	(0.51)	(1.04)
R^2	0.32	0.23	0.07	0.27	0.38	0.41	0.14	0.27	0.49	0.24	0.26
N	39	47	184	701	720	2,428	3,503	1,045	1,046	4,133	457
Period	76-14	70-17	70-16	91-15	91-15	70-17	99-18	70-17	70-17	70-16	76-15
Freq.	Yr	Yr	Qtr	Yr	Yr	Qtr	Yr	Yr	Yr	Qtr	Yr
F.E.	-	-	-	Yr/Ctry	Yr/Ctry	Qtr/Ctry	Yr/Ind/Ctry	Yr/Ind	Yr/Ind	Qtr/Ind	Yr/Ind
Source	LBD	CSTAT	CSTAT	BvD	BvD	GCSTAT	Amadeus	CSTAT	CSTAT	CSTAT	ASM
Sample	-	231K	650K	357K	633K	5,800K	357K	231K	231K	733K	-

Note: The left panel of Table II shows a set of time series regressions for the United States in which the dependent variable is the Kelley skewness of the distribution of one-year firm log employment growth from the LBD (column 1), one-year log sales growth (column 2) from Compustat (CSTAT), and one-year stock returns (column 3) from CSTAT. In each regression, the independent variable is the one-year log GDP per capita growth. LBD moments are weighted by firm size measured by the average employment of the firm between years t and $t+1$. All regressions include a linear trend. Newey-West standard errors in parentheses below the point estimates. Columns (4) to (7) show a set of country-panel regressions in which the dependent variable is the Kelley skewness of different firm-level outcomes. Employment and sales data come from BvD Osiris dataset, stock returns are from BvD Osiris dataset, stock returns are from Global Compustat (GCSTAT), and TFP shocks are estimated using data from the BvD Amadeus dataset. In each regression, the independent variable is the one-year of log GDP per capita growth. Standard errors are clustered at the country level. Columns (8) to (11) show a series of industry-panel regressions in which the dependent variable is the Kelley skewness of the within-industry distribution of firms-level outcomes. In columns (8) to (10) an industry is defined as a 2-digit NAICS cell; in column (11) an industry is a 3-digits NAICS cell within manufacturing (NACIS 31-33). Results shown in column (11) come from the US Census of Manufacturing (CM) and the Annual Survey of Manufacturing Firms (ASM). In each regression, the independent variable is the average log sales growth within an industry-year cell. Standard errors are clustered at the industry level. The row labeled Sample shows the underlying sample of firms/periods used to calculate the cross-sectional moments. Underlying sample sizes in LBD and the Census' CMF/ASM are not disclosed. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

37

TABLE III – PARAMETERIZATION

Preferences and Technology		
γ	0.45	Frisch elasticity of labor supply
ψ	2.5	Leisure preference, non-entrepreneurs spend 1/3 time working
σ	2.0	Risk aversion, non-entrepreneurial sector
ξ	6.0	Risk aversion, entrepreneurs
β	$0.95^{0.25}$	Annual discount factor of 95%
r	0.005	Annual return of risk-free asset of 2%
α	0.25	CRS production, markup of 33%
ν	0.50	CRS labor share of 2/3, capital share of 1/3
δ	3.8%	Annual depreciation of capital stock of 14.4%
ρ_a	0.95	Quarterly persistence of aggregate productivity
σ_a	0.75%	Standard deviation of innovation of aggregate productivity
ρ	0.95	Quarterly persistence of idiosyncratic productivity
Adjustment costs		
ϕ_1	1.50%	Fixed cost of changing capital stock
ϕ_2	7.0	Quadratic cost of changing capital stock
ϕ_3	34.0%	Resale loss of capital
Estimated Parameter of Idiosyncratic Stochastic Process		
σ_1^L	1.45	Standard deviation of first mixture in low-risk periods (%)
σ_2^L	7.55	Standard deviation of second mixture in low -risk periods (%)
μ^L	-0.92	Mean of first mixture in low-risk periods (%)
p^L	63.67	Probability of first mixture in low-risk periods (%)
σ_1^H	4.37	Standard deviation of first mixture in high-risk periods (%)
σ_2^H	9.06	Standard deviation of second mixture in high-risk periods (%)
μ^H	1.98	Mean of first mixture in high-risk periods (%)
p^H	78.28	Probability of first mixture in high-risk periods (%)
Transition Probabilities of Risk States		
π_L	0.97	Quarterly probability of remaining in low risk state
π_H	0.84	Quarterly probability of remaining in high-risk state

Note: The top two panels of Table III shows the calibrated parameters referring to preferences, technology, and adjustment costs. The two bottom panels of Table III shows the parameters of the stochastic process of firm-level productivity. We target moments of the annual change of quarterly sales in Compustat. The parameters for low-risk periods (denoted by an upper script L) are obtained by targeting the $P9010_t$, $P9050_t$, and the $P5010_t$ percentiles differential, and Kelley skewness of the log sales growth distribution for all the expansion years between 2000 and 2014. The parameters for high-risk periods (denoted by an upper script H) are obtained by targeting the same set of moments for years 2001 and 2008 (full recession years). See Table IV for comparison between the targeted and model generated moments. The transition probability π_L is calculated as the share of expansion quarters that were followed by another expansion quarter, whereas π_H is calculated as the share of recession quarters that were followed by another recession quarter using data from 1970 to 2014.

TABLE IV – RISK PROCESS MOMENTS

|	$P90 - P10$	$P90 - P50$	$P50 - P10$	KSK	Yrs
Data					
Low-Risk	0.54	0.30	0.24	0.10	03-06;10-14
High-Risk	0.70	0.31	0.39	-0.11	01,08
$\Delta (H - L)$	0.16	0.01	0.15	-0.20	-
Model					
Low-Risk	0.48	0.27	0.20	0.15	-
High-Risk	0.58	0.26	0.32	-0.10	-
$\Delta (H - L)$	0.10	-0.01	0.12	-0.25	-

Note: The top panel of Table IV shows cross-sectional moments log quarterly sales growth between quarters t and $t+4$ from Compustat for low-risk periods–quarters in the years 2003 to 2006 and quarters in the years 2010 to 2014–and high-risk periods–quarters in years 2001 and 2008. Quarters in years 2002 and 2009 are discarded for not representing full recession years. The model moments, shown in the lower panel of Table IV, are calculated from a 5,000-quarters simulation with the first 500 quarters discarded.

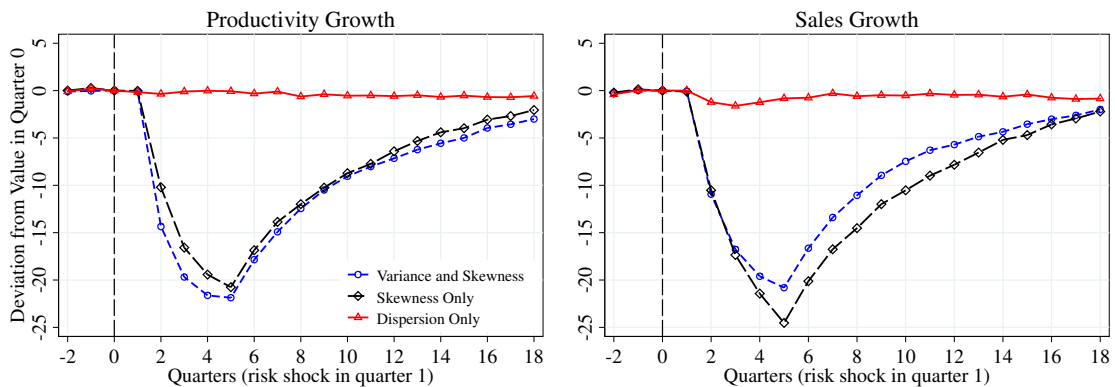FIGURE 8 – PRODUCTIVITY AND SALES GROWTH AFTER AN INCREASE IN RISK

(A) Average
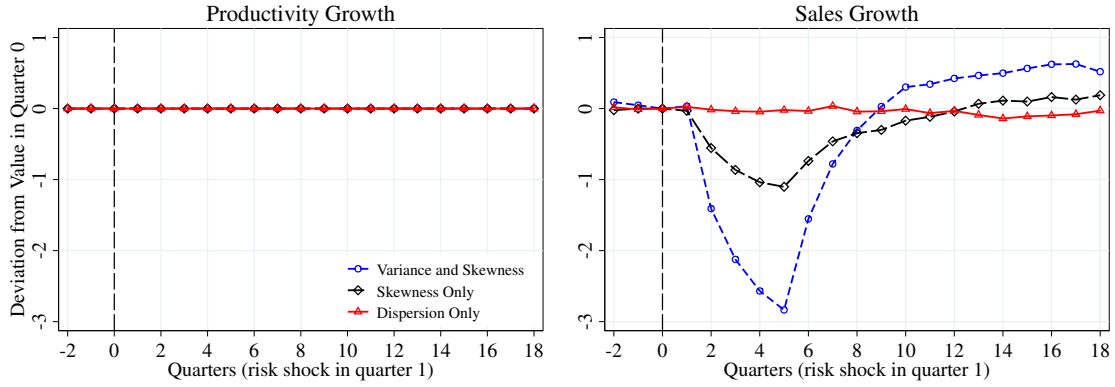


(B) P90-P10
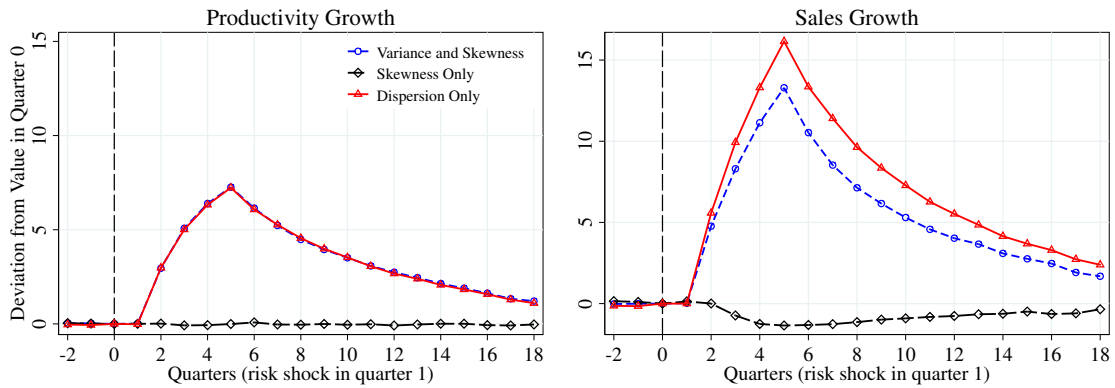
(C) Kelley skewness

Note: The top left panel of Figure 8 shows the model-generated average of the one-year productivity growth distribution ($\Delta e_{j,t} = e_{j,t+4} - e_{j,t}$), whereas the top right shows the average of the log sales growth distribution ($\Delta y_{j,t+4} = \log y_{j,t+4} - \log y_{j,t}$) for different risk shocks. The middle and bottom panels show the dispersion and skewness. Each plot is based on independent simulations of 1,000 economies of 300-quarter length. In each simulation, we assume that the economy is in the low-risk state for 150 periods. We then impose a risk shock in quarter 151, allowing normal evolution of the economy afterwards. We plot the deviation relative to the moment value in quarter 0. The red line with triangles traces the impact of an increase in the variance of firms' shocks; the black line with diamonds traces the impact of a drop in the skewness of firms' shocks; the blue line with circles traces the joint impact of an increase in variance and a decrease in skewness.

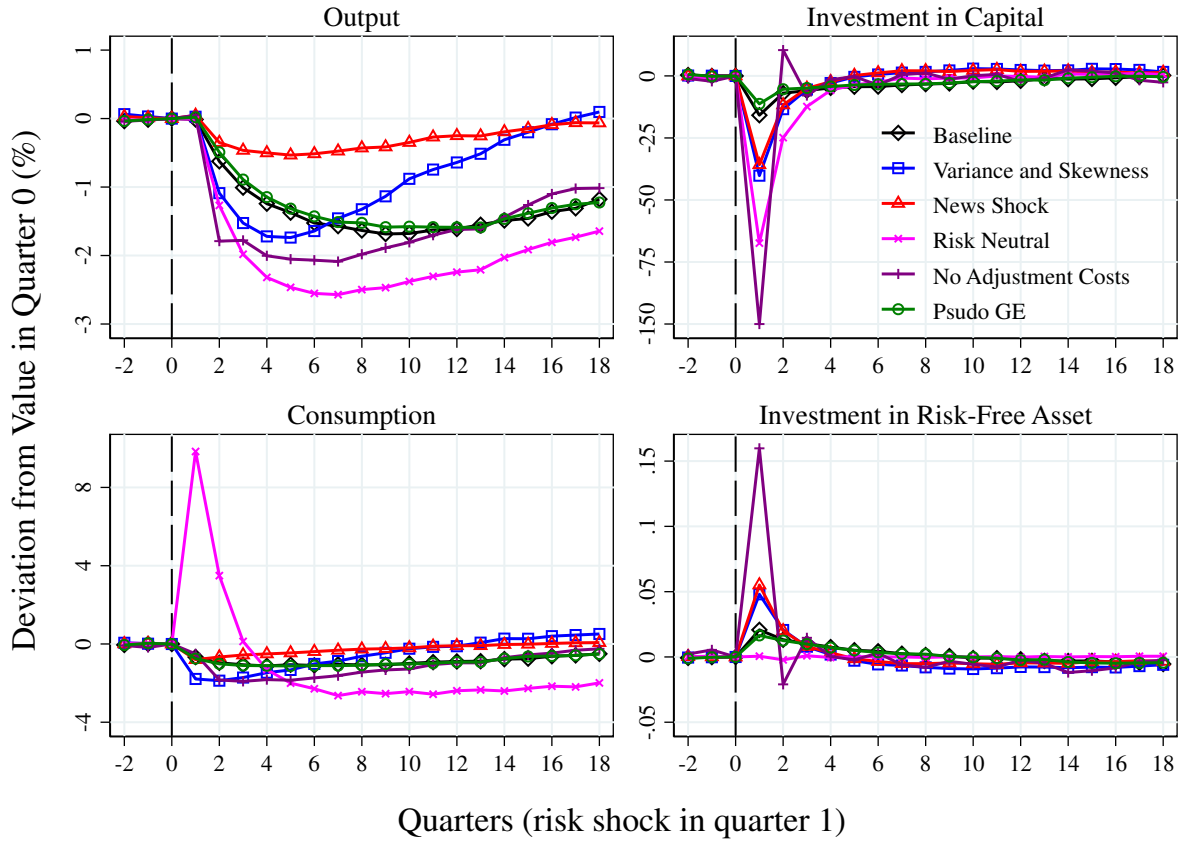FIGURE 9 – EFFECT OF SKEWNESS SHOCK ON OUTPUT

(A) Output

(B) Macroeconomic Aggregates
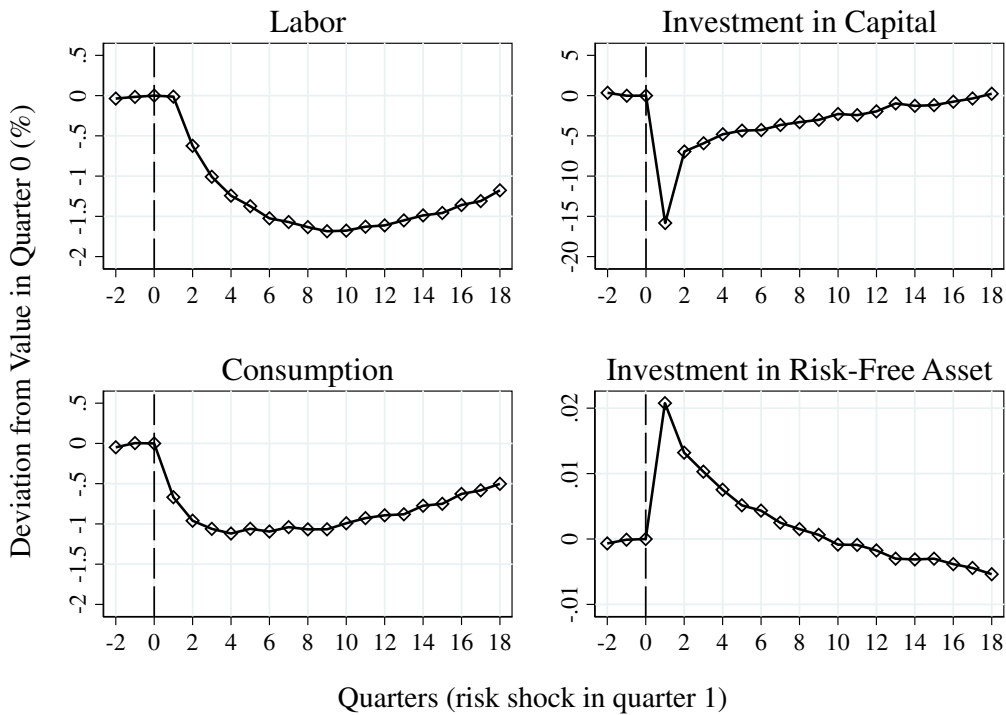
Note: Figure 9 shows the effect of a decline in the skewness of firm idiosyncratic productivity. The plot is based on independent simulations of 1,000 economies of 300-quarter length. In each simulation, we assume that the economy is in the low-risk state for 150 periods. We then impose a drop in the skewness of firms' shocks in quarter 151, allowing normal evolution of the economy afterwards. We plot the log percentage deviation of each macroeconomic aggregate from its value in quarter 0. Top panel shows the effects of output, whereas the bottom panel shows the impact on labor, investment in capital, consumption, and investment in the risk-free asset.

FIGURE 10 – EFFECT OF SKEWNESS AND VARIANCE SHOCKS ON MACRO AGGREGATES

Quarters (risk shock in quarter 1)

Note: Figure 10 shows the effect of a decline in the skewness of firms' idiosyncratic productivity for several different parameterizations of the model. Each plot is based on independent simulations of 1,000 economies of 300-quarter length. In each simulation, we assume that the economy is in the low-risk state for 150 periods. We then impose a risk shock in quarter 151, allowing normal evolution of the economy afterwards. We plot the log percentage deviation of each macroeconomic aggregate from its value in quarter 0. The Baseline (diamonds) is the estimated effect under the baseline parameterization; Variance and Skewness (squares) traces the impact of a skewness shock pared with a variance shock; News shock (triangles) traced the impact after a change in the skewness of shocks that does not change the realizations of shocks faced by firms; Risk neutral (x symbols) traces the impact of a skewness shock in the case that entrepreneurs have linear utility functions; No Adjustment costs (+ symbols) traces the impact of a skewness shock in the case all adjustment costs have been set to 0; Psudo GE (circles) traces the impact of a skewness shock paired with a decline in the returns of the risk-free asset. Labor is omitted since it follows the same pattern of output.

A   Appendix: Robustness Results

TABLE A.1 – DISPERSION OF FIRMS OUTCOMES IS HIGHER DURING RECESSIONS

90th-to-10th Log Percentiles Differential of the Log Growth of Firms' Outcomes

	United States						Cross-Country		
	Firm Sales		Stock Returns		Firm Emp.		Firm Sales	Firm Stock	Firm Emp.
	(1)	(2)	(3)	(4)	(5)		(7)	(8)	(9)
	One Year	Three Year	One Year	Three Year	One Year		Growth	Returns	Growth
$\Delta GDP_{i,t}$	-3.91***	2.55**	-3.93**	-4.78***	0.93*		-0.79	-1.84	-0.76
	(1.14)	(0.99)	(1.62)	(1.78)	(0.50)		(0.59)	(1.79)	(0.74)
N	184	182	180	180	39		838	4,306	824
Freq.	Qtr	Qtr	Qtr	Qtr	Yr		Yr	Qtr	Yr
F.E.	N	N	N	N	N		Yr/Ctry	Qtr/Ctry	Yr/Ctry
Source	CSTAT	CSTAT	CSTAT	CSTAT	LBD		BvD	GCSTAT	BvB

Note: The left panel of Table A.1 shows a series of time series regressions in which the dependent variables are the 90th-to-10th log percentiles differential of the one-year and three-year growth rate of sales (columns 1 and 2), stock returns (columns 3 and 4), and employment growth (columns 5) for a sample of firms from Compustat (columns 1 to 4) and the LBD (columns 5). Compustat data cover the period 1970 to 2017 whereas LBD data covers the period 1976 to 2015. In each regression, the independent variable is the annual growth rate of quarterly GDP per capita. All regressions include a linear trend. Newey-West standard errors in parentheses below the point estimates. The right panel of Table A.1 shows a series of country-panel regressions where the dependent variable is the within-country P90-P10 log percentiles differential of firm-level sales growth, stock returns, or employment growth. The independent variable is the growth rate of GDP per capita at the country level. Sales and employment data are obtained from the BvD Osiris database, whereas stocks returns are obtained from Global Compustat. All regressions consider a full set of time and country fixed effects. The row labeled Sample shows the underlying sample of firms used to calculate the cross-sectional moments. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A.2 – HIGHER ORDER MOMENTS OF FIRM-LEVEL OUTCOMES

	Kelley Skewness					Crow-Siddiqui Kurtosis			
	Residual Sales		Sales per Employee		Sales Deviation	Firm Sales		Stock Returns	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	One Year	Three Years	One Year	Three Years		One Year	Three Years	One Year	Three Years
$\Delta GDP_{i,t}$	3.80***	1.48	3.01**	4.17***	1.46**	0.36***	−0.19	0.16***	−0.21*
	(1.42)	(0.91)	(1.23)	(0.97)	(0.58)	(0.08)	(0.12)	(0.06)	(0.13)
N	178	178	174	166	47	184	182	180	180
Freq.	Qtr	Qtr	Qtr	Qtr	Yr	Qtr	Qtr	Qtr	Qtr
Sample	500K	500K	500K	500K	113K	640K	640K	650K	650K
Source	CSTAT	CSTAT	CSTAT	CSTAT	CSTAT	CSTAT	CSTAT	CSTAT	CSTAT

Note: The left panel of Table A.2 shows a series of time series regressions for the United States in which the dependent variable is the Kelley Skewness of the one-year and three-year growth rate of residualized sales growth (columns 1 and 2) and the growth rate of sales per employee (columns 3 and 4) for a sample of firms from Compustat. In columns (1) and (2), we have orthogonalized the growth rates of sales from time fixed-effects, firm-fixed effect, size, and other firm-level observable characteristics. Column (5) shows the correlation of GDP growth and the cross-sectional skewness of the deviation of annual firms' sales from an HP trend. Compustat data cover the period 1970 to 2017. The dependent variable in columns (6) to (9) is the Crow-Siddiqui measure of kurtosis defined as $CKU_t = \frac{P97.5_t - P02.5_t}{P75_t - P25_t}$. In each regression, the independent variable is the annual growth rate of quarterly GDP per capita. All firm-level moments were calculated weighting the growth rate observations by firm size measured by the average sales of the firm between periods t and $t+k$. All regressions include a linear trend. Newey-West standard errors in parentheses below the point estimates. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A.3 – CENSUS LBD: CROSS SECTIONAL MOMENTS OF LOG EMPLOYMENT GROWTH

Moment:	(1) Kelley Skewness	(3) Coefficient of Skewness	(3) $P9010_{jt}$	(4) $P9050_{jt}$	(5) $P5010_t$
$Ave\Delta E_{j,t}$	0.92***	8.15**	-1.38***	-2.47***	1.09***
	(0.16)	(4.05)	(0.34)	(0.42)	(0.08)
R^2	0.56	0.31	0.58	0.66	0.52
N	900	900	900	900	900
F.E.	Yr/Ind	Yr/Ind	Yr/Ind	Yr/Ind	Yr/Ind
Period	76-14	76-14	76-14	76-14	76-14

Note: The results in Table A.3 are based in firm-level data from the Census' LBD dataset. Each column shows the results of an industry panel OLS regressions in which the dependent variables is a different moment of the of employment growth distribution and the independent variable is the average employment growth within an industry-year cell. An industry is a 2-digit NAICS industry group. Standard errors clustered at the industry level, below the point estimates in parenthesis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A.4 – BUSINESS CYCLE STATISTICS

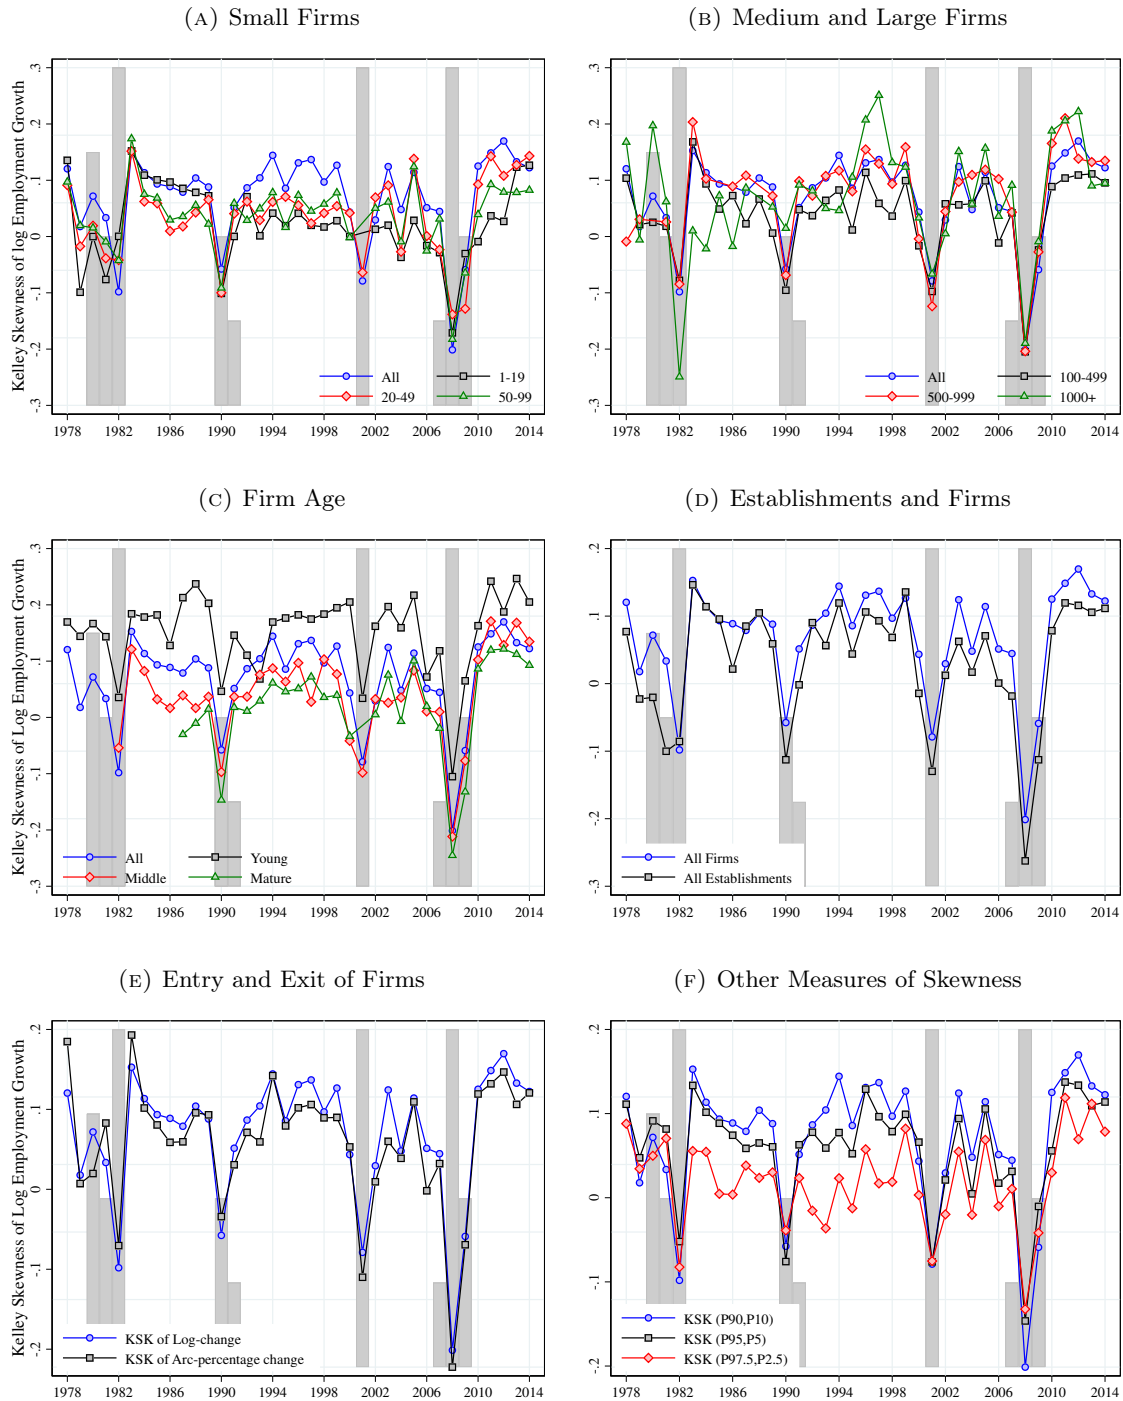|	Data			Model		
	$\sigma(x)$	$\sigma(y)/\sigma(x)$	$\rho(x,y)$	$\sigma(x)$	$\sigma(y)/\sigma(x)$	$\rho(x,y)$
Output	1.47	1.00	1.00	2.00	1.00	1.00
Capital Investment	6.86	4.64	0.91	9.38	4.69	0.30
Consumption	1.21	0.82	0.87	1.81	0.91	0.65
Hours	1.89	1.28	0.87	2.00	1.00	1.00

Note: The left panel of Table A.4 displays business cycle statistics for quarterly US data covering 1970Q1 to 2017Q4. The column $\sigma(x)$ is the standard deviation of the log variable in the first column. The column $\sigma(y)/\sigma(x)$ is the standard deviation of the variable relative to the standard deviation of log output. All business cycle data are current as of February 3, 2019. Output is real gross domestic product (FRED GDPC1), investment is real gross private domestic investment (FRED GPDIC1), consumption is real personal consumption expenditures (FRED PCECC96), and hours is total non-farm business sector hours (FRED HOANBS). The second panel contains business cycle statistics computed from a simulation of the model of 5,000 quarters with the first 500 periods discarded. All series are HP-filtered with smoothing parameter 1,600, in logs expressed as percentages.

TABLE A.5 – TARGETED MOMENTS FOR NUMERICAL COMPARISON

	$P9010$	$P9050$	$P5010$	KSK
Low-Risk	0.54	0.30	0.24	0.10
High-Risk	0.70	0.31	0.39	-0.11
Only Skewness	0.54	0.243	0.297	-0.11
Only Variance	0.70	0.39	0.31	0.10

Note: Table A.5 shows the target used in the estimation of the firm-level productivity process. Rows labeled "Low-Risk" and "High-Risk" are used in the baseline estimation. The values for "Only Skewness" are used to estimate the parameters when the economy is shocked with a change in the skewness only. Similarly, the values for "Only Variance" are used to estimate the parameters when the economy is assumed to be shocked only by a change in the variance of firms' shocks while keeping the skewness constant.

FIGURE A.1 – Census LBD: Skewness of Log Employment Growth is Robustly Procyclical

(A) Small Firms

(B) Medium and Large Firms

(C) Firm Age

(D) Establishments and Firms

(E) Entry and Exit of Firms

(F) Other Measures of Skewness

Note: Figure A.1 is based on firm- and establishment-level from the Census' LBD dataset. The top panels show the Kelley skewness of the distribution of firms' log employment growth within different firm size groups. The center-left panel shows the skewness of the distribution of firms' log employment growth within different firm age groups. Young firms are those less than five years old, middle-aged firms are those between six and ten years old, and mature firms are those of more than ten years old. Firms already in the sample in 1976 are not considered in any of these groups. Shaded areas represent the share of the year (in quarters) declared as recession by the NBER. All moments weighted by average employment at the firm or establishment level. See Appendix B for details on the sample construction and moment calculations in the LBD.

46

(A) Compustat: Skewness of Log Sales Growth Distribution

(B) Compustat: Upper and Lower Tail Dispersion of Log Sales Growth

Note: The top panel of Figure A.2 shows the time series of the cross-sectional Kelley skewness of the distribution of the annual growth rate of quarterly sales for a sample of firms from Compustat. The bottom panel of Figure A.2 shows the time series 90th-to-50th log percentiles differential and the 50th-to-10th log percentiles differential of the annual log quarterly sales growth for a sample of firms from Compustat. The shaded areas represent NBER recession quarters. See Appendix B.2 for additional details on the sample construction and moment calculations.

FIGURE A.3 – SKEWNESS OF FIRM-LEVEL OUTCOMES INCLUDING PRIVATE FIRMS IS PROCYCLICAL

(A) BvD Amadeus: Log Employment Growth

(B) BvD Amadeus: Log Sales Growth

Note: Figure A.3 shows scatter plots of the Kelley skewness and average of the employment growth and sales growth distribution within a country-year cell. The figure is based on an unbalanced panel of firms from the BvD Amadeus database in the following European countries: AUT, BEL, BLR, CHE, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, IRL, ISL, ITA, NLD, NOR, POL, PRT, SWE, UKR. The data cover years 2000 to 2015. BvD Amadeus contains private and publicly traded firms. See Appendix B.4 for additional details the data construct and moment calculation.

(A) Compustat: Right Tail Dispersion of Log Sales Growth

(B) Compustat: Left Tail Dispersion of Log Sales Growth


Note: The Figure A.4 displays a scatter plot showing the relation between the within-industry business cycle, measured by the average growth rate of sales growth, and the within industry dispersion of sales growth constructed from Compustat data. The top panel shows the 90th-to-50th log percentiles differential whereas the bottom panel shows the 50th-to-10th log percentiles differential.

(A) Compustat: Firm-Level Employment Growth

$$Kelley_{j,t} = \alpha + \beta Me_{j,t} + \varepsilon_{j,t}$$

(B) Compustat: Firm-Level Sales Growth

$$Kelley_{j,t} = \alpha + \beta Me_{j,t} + \varepsilon_{j,t}$$

Note: Figure A.5 shows the coefficients and confidence intervals for within-industry regression of the cross-sectional Kelley skewness on the average growth of employment (top panel) and sales (bottom panel) for a sample of publicly traded firms from Compustat. Each industry regression includes a linear trend. Confidence intervals are calculated at 95% of significance. Industries are defined as two-digit NAICS. In each plot, the dashed line is the coefficient of a panel regression of the within industry skewness and average firm growth controlling for time and fixed effect. See Appendix B.1 for additional details on the sample construction and moment calculations in Compustat.

(A) Aggregate Productivity Shock does not Affect Dispersion or Skewness of Sales Growth
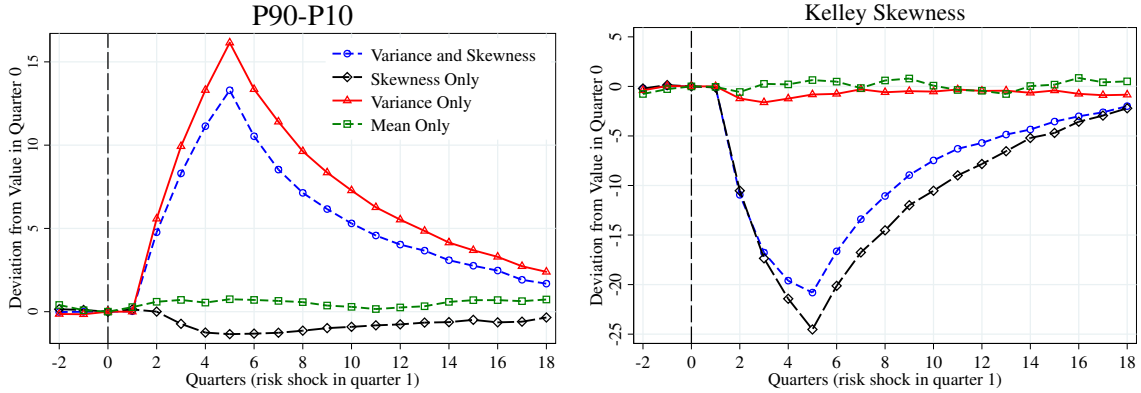
(B) Skewness of Employment and Sales Growth

(C) Right and Left Tail Dispersion of Sales Growth
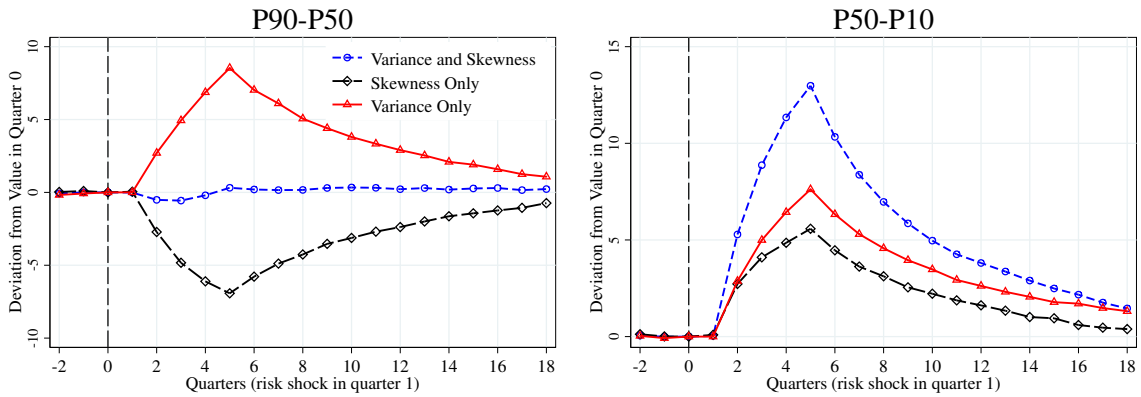
Note: Figure A.6 shows different model-generated moments of the sales growth and employment growth distribution. Each plot is based on independent simulations of 1,000 economies of 300-quarter length. In each simulation, we assume that the economy is in the low-risk state for 150 periods. We then impose a drop in the skewness of firms' shocks in quarter 151, allowing normal evolution of the economy afterwards. We plot the deviation of each macroeconomic aggregate from its value in quarter 0. The red line with triangles trances the impact of an increase in the variance of firms' shocks; the black line with diamonds trances the impact of a drop in the skewness of firms' shocks; the blue line with circles trances the joint impact of an increase in variance and a decrease in skewness.

B Appendix: Data

This appendix describes the data sources, the sample selection, and the calculations of the moments we use for our empirical analysis. In Section B.1 we describe the firm- and establishment-level data for the United States obtained from the Census Bureau's Longitudinal Business Database (LBD). Section B.2 describes our sample of Compustat firms. For our cross-country comparisons, we use firm-level data available in the Bureau van Dijk's Osiris database and Global Compustat which we describe in Section B.3. Finally, Section B.4 describes our sample and TFP estimation for our sample of firms from the BvD Amadeus dataset and the establishment-level data for the US Census of Manufacturing and the Annual Survey of Manufacturing. The online appendix and replication packet—available in this link https://bit.ly/2ZfaMVR and on the authors' websites—contains further details and corresponding do-files for most of our calculations.

B.1 United States: Longitudinal Business Database

We construct measures of employment growth at the firm- and establishment-level using the Census Bureau's LBD. The LBD covers the universe of establishment in the nonfarm private sector in the United States from 1976 to 2015. It provides detailed establishment and firm-level information on employment, payroll, 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 nonnegative, 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. We measure the growth rate of employment of firm j in period t as the log-difference between periods t and $t + k$, $g_{j,t}^e = \log E_{j,t+k} - \log E_{j,t}$ where $k \in \{1, 5\}$. In order to capture the entry and exit of firms, we replace by a 0 the level of employment in the period before the establishment is observed for the first time and the period after the establishment 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 $t + k$ which is given by $g_{j,t}^{arc} = \frac{E_{j,t+k} - E_{j,t}}{0.5 \times (E_{j,t+k} + E_{j,t})}$.

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 a particular firm, which violates the disclosure criteria imposed by the 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 (around 6 million observations per year) 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 are weighted by the average employment of the firm (or establishment) between periods t and $t + k$, that is, $\overline{E}_{j,t} = 0.5 \times (E_{j,t+k} + E_{i,t})$.

We also use the LBD to compare the empirical distribution of employment growth between recession and expansion years using a kernel density estimation. The sample selection is the same used in the rest of our results; however, the Census Bureau requires to drop the bottom and top 5% of the distribution when estimating empirical densities. The kernel densities presented in Figure 3 were calculated using the remaining sample.

B.2 United States: Compustat

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 accessed through the Wharton Research Data Services (WRDS).

For our results at the annual frequency, we obtain firm-level data from 1970 to 2017. The raw annual dataset contains 500,004 year/firm observations. We drop all observations with negative sales (Compustat variable sale), duplicated entries, and firms incorporated outside the United States (Compustat variable fic equal to "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 $k \in \{1, 3, 5\}$. This leaves us with 266,192 firm/year observation (sales growth) between 1970 and 2016, with an average of 5,663 firms per year. 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 -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).

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 2017q4 . 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 equal to "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), which further reduces the sample to 819,977 observations between 1970q4 and 2017q2, with an average of 5,359 firms per quarter. Finally, in each quarter we calculate different cross-sectional moments discussed in the main body of this document. Our main sample considers firms with at least 10 years of data (40 quarters), although our results remain robust if we drop this restriction or if we consider firms with 25 years of data (more than 100 quarters). 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 entering 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).

B.3 Cross-Country: BvD Osiris and Global Compustat

Cross-country firm-level panel data on sales and employment come from the Bureau van Dijk's Osiris database.[30] 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 2018. 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

---

[30]See Kalemli-Ozcan *et al.* (2015) for additional details on the Orbis dataset.

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 45 countries between 1989 and 2015. 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 2018 for 48 countries. We drop all duplicated observations and drop all firms with less 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 44 countries from 1985 to 2017 from which we drop all country quarters with less than 100 firms. The final data contains a total of 29 countries from 1985 to 2017. 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.6 shows the list of countries available in our dataset and the data available for each country.

TABLE A.6 – Data Availability by Country

Source	BvD Osiris Sales	BvD Osiris Emp	Global Compustat Returns	BvD Amadeus Sales	BvD Amadeus Emp	BvD Amadeus TFP
ARG	×	×				
AUS	×	×	×			
AUT			×	×	×	
BEL	×	×	×	×	×	
BLR	×			×	×	
BMU	×	×				
BRA	×	×	×			
CAN	×	×				
CHE	×	×	×	×	×	
CHL	×	×	×			
CHN	×	×				
DEU	×	×	×	×	×	×
DNK	×	×	×	×	×	×
EGY	×	×				
ESP	×	×	×	×	×	×
FIN	×	×	×	×	×	×
FRA	×	×	×	×	×	×
GBR	×	×	×	×	×	×
GRC	×	×	×	×	×	×
HKG	×	×				
HUN	×			×	×	×
IDN	×	×	×			
IND	×	×	×			
IRL	×		×	×	×	×
IRN	×	×				
ISL		×			×	×
ISR	×	×	×			
ITA	×	×	×	×	×	×
JPN	×	×	×			
KOR	×	×	×			
MEX	×	×				
MYS	×	×				
NLD	×	×	×	×	×	×
NOR	×	×	×	×	×	×
NZL	×	×	×			
PAK	×	×				
PER	×	×				
PHL	×	×				
POL	×	×	×	×	×	×
PRT	×	×		×	×	×
RUS	×	×	×			
SGP	×	×				
SWE	×	×	×	×	×	×
THA	×	×				
TUR	×	×	×			
UKR		×				
ZAF	×	×	×	×	×	×

Note: Table A.6 shows data available for each country (identified by its isocode). US data sources are omitted.

B.4 TFP Estimation

B.4.1 Cross-Country: BvD Amadeus

In this appendix, we describe in detail the construction of our measure of firm-level TFP using data from Amadeus. We consider a set of countries, namely, Germany, Denmark, Spain, Finland, France, United Kingdom, Greece, Hungary, Ireland, Island, Italy, Netherlands, 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 trough WRDS. Our data contains 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 parenthesis):

- Sales (TURN),
- Operating revenues (OPRE),
- Employment (EMPL),
- Cost of Employees (STAF),
- Cost of Material (MATE),
- Total Fixed Assets (FIAS),
- Industry (NAICS and SIC codes),
- Exchange rate from local currency to Euros (EXCHANGE2).

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).[31] We also drop all observations with missing, zero, or negative values in either of the following variables: OPRE, MATE, FIAS, and STAF. We also drop all observation 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). Firms in Sweden report information in their local currency, which we transform to Euros using the exchange rate also reported by Amadeus.

B.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 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,

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

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

---

[31]The first two characters in the firm identifier in Amadeus refer to the country.

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-sector-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 $\alpha_{K,j,k,t} = 1 - \alpha_{L,j,k,t}$.[32] 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 $((1 - \alpha_{L,j,k,t}) \log K_{i,j,k,t} + \alpha_{L,j,k,t} \log E_{i,j,k,t})$.

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 BvD Amadeus was 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 *et al.* (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 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, follow through is we relax these conditions.

B.4.3 Additional Evidence on the Skewness of Productivity Shocks

As we discussed in Section 3.4, the skewness of productivity shocks is robustly negative during periods of low economic activity within a country or an industry. Here we show some additional robustness results. Figure A.7 shows that the positive relation between the skewness of the productivity shocks and the business conditions is robustly positive, independently of the estimation method one uses to calculate firm-level productivity. For comparison, the top left panel repeats our main results shown in Figure 6.

Comparing the slopes in the plots in Figure A.7 is difficult to appreciate whether some measures of productivity lead to more cyclical measures of skewness of TFP shocks since each plot has a different

---

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

TABLE A.7 – POSITIVE CORRELATION OF SALES GROWTH AND SKEWNESS OF FIRM'S SHOCKS

	(1)	(2)	(3)	(4)
Estimation Method:	Factor Shares	Panel Regression	Olley and Pakes	Labor Productivity
Ave. Sales Growth	1.21***	1.10***	1.18***	1.256***
	(0.40)	(0.31)	(0.40)	(0.37)
R^2	0.18	0.14	0.31	0.17
N	3,873	3,873	3,873	3,873

Note: Table A.7 shows a set of country-industry panel regressions in which the dependent variable is the Kelley skewness of firm productivity shocks calculated using the four different methods described in Section B.4.2. In all regressions, the explanatory variable is the average sales growth within the same bin. All regressions control for country, industry, and year fixed effects. Standard errors (below the point estimates) are clustered at the country level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

x-axis. In order to have a more direct comparison between the skewness across different estimation methods, Table A.7 shows a series of panel regressions in which the independent variable is the skewness of TFP shocks for each of the four methods described in Section B.4.2 and the main regressor is the average of the real sales growth (log change of operative revenues) within a country-industry-year cell. The coefficient associated to the average sales growth is positive and statistically significant at the 1% in all cases and of the same order of magnitude. This indicates that in periods in which industries experience a decline in sales, the skewness of the productivity shocks affecting the firms in that industry is negative as well.
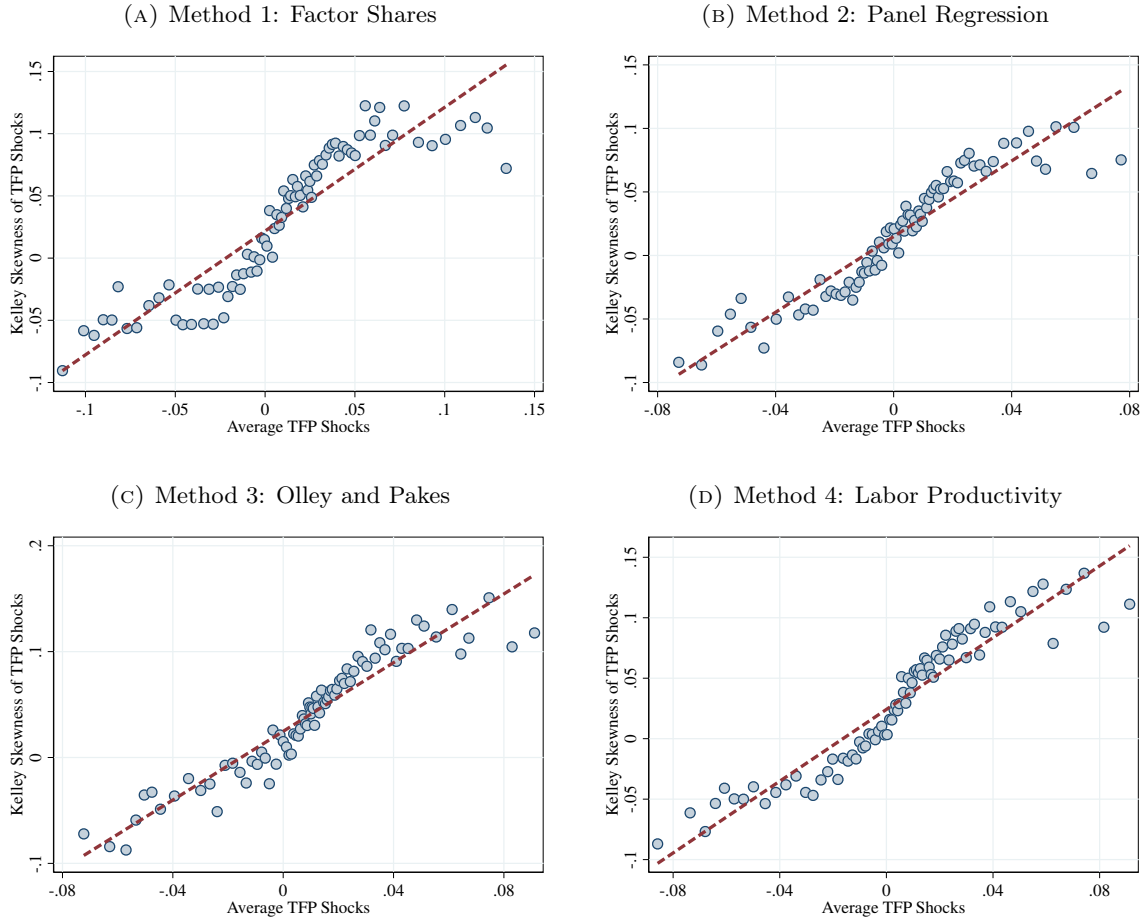
Finally, Table A.8 shows that the skewness of firm's shocks is procyclical at the country level. In particular, we show the results of an industry-panel regression for each country in our sample using the average TFP shock as our measure of business condition. The results are indicative that the procyclicality of the skewness of firm shocks is not driven by any particular country in our sample and it is stronger in countries such as Germany and France.

B.4.4 United States: US Census and Survey of Manufacturing

Here we describe the sample selection and moment construction using data from the US Census of Manufacturing (CM) and the Survey of Manufacturing firms (ASM). The CM, which is part of the Economic Census, is conducted every five years, in every year ended in 2 or 5 and was first implemented in 1963. It covers all establishments with at least one paid employee in the manufacturing sector (NAICS 31-33) for a total sample between 300,000 and 400,000 establishments per Census. Information is delivered by firms at the establishment level and Census provides a unique identifier (lbdnum) which we use to follow establishments over time. The Census Bureau complements the CM data with the ASM every year the Economic Census is not conducted since 1973. Relative to the CM, the ASM is skewed towards large firms as it covers all establishments of firms covered by the CM above a certain threshold and a smaller sample of small and medium size firms. The average sample of firms in our sample is around 30,000 establishments per year for a total of around 1.1 million establishment/year observations. The merged CM/ASM contains consistent data in industry, sales, employment, capital expenditures, materials, and others. Importantly, from 1976 to 2015, the data contains measures of log-productivity prepared by the Bureau of Labor Statistics which we directly use in our analysis.

To keep a consistent sample selection across datasets, we consider establishments for ten or more years of data. Since the ASM sample is refreshed every Census year, this sample selection criteria naturally select large and stable firms. Our results, however, are robust to the changes in the 10 years threshold.

FIGURE A.7 – THE SKEWNESS OF FIRMS' SHOCKS IS ROBUSTLY PROCYCLICAL

(A) Method 1: Factor Shares

(B) Method 2: Panel Regression

(C) Method 3: Olley and Pakes

(D) Method 4: Labor Productivity

Note: Figure A.1 is based on a sample of firms from BvD Amadeus. Each plot shows a scatter plot polling information across all countries, industries, and years in the sample. In each plot, the y-axis is the Kelley skewness of the within country-industry-year distribution of firm productivity shocks whereas the x-axis is the average productivity shock with the same cell. Productivity shocks are calculated using four different methods: Factor shares (top left plot), panel regression (top right plot), Olley and Pakes (bottom left), and labor productivity (bottom right). The slopes (standard errors) sorted from top left to bottom right are the following (standard errors in parenthesis): 0.69 (0.04), 1.79 (0.07), 1.97 (0.11), and 1.41 (0.07). To create this figure, we winsorize the distribution of Kelley skewness and average growth at the top and bottom 1%. Scatter plots controlling for country, industry, and time fixed effects.

We construct measures of employment growth, sales growth, and productivity growth as the log-change between years t and $t-1$. Productivity shocks instead are calculated as the residual of the following regression,

$$z_{i,t} = \beta_0 + \beta_1 z_{i,t-1} + \mu_i + \delta_t + \epsilon_{i,t},$$

where μ_i is a establishment fixed effects and δ_t is a year-fixed effect. Similarly to what we do in the LBD sample, we construct percentiles of the distribution of firm's growth and productivity residual as the average level of the corresponding variable over a 2 percentage point band around the desired percentile. For instance, to calculate the 10th percentile of the distribution of TFP residuals within a industry-year cell, we select a sample of establishment between located between the 8 and the 12 percentiles of the distribution. Then, our measure of the 10th percentile is the average over all the establishment in this sample. We proceed in a similar fashion for the 50th and the 90th percentiles.

ISO	(1) DEU	(2) DNK	(3) ESP	(4) FIN	(5) FRA	(6) GBR	(7) GRC	(8) HUN	(9) IRL
Sales Growth	5.27***	4.47***	1.69***	2.27***	3.24***	2.25***	1.27***	0.87***	1.70***
	(0.70)	(1.02)	(0.32)	(0.14)	(0.45)	(0.14)	(0.24)	(0.24)	(0.19)
R^2	0.66	0.63	0.60	0.64	0.56	0.74	0.47	0.76	0.66
N	208	73	392	245	334	275	179	271	186

Continuation

ISO	(10) ISL	(11) ITA	(12) NLD	(13) NOR	(14) POL	(15) PRT	(16) SWE	(17) UKR
Sales Growth	2.56***	1.29***	2.12***	2.04***	1.87***	1.74***	2.86***	1.56***
	(0.31)	(0.19)	(0.21)	(0.35)	(0.36)	(0.17)	(0.228)	(0.209)
R^2	0.84	0.55	0.72	0.51	0.52	0.61	0.65	0.55
N	102	306	152	208	264	234	228	229

Note: Table A.8 shows a set of industry panel regressions in which the dependent variable is the Kelley skewness of firm productivity shocks. Firm-level productivity was calculated as the residuals of a firm-level panel regression (the second method described in section B.4.2). In each column, the independent variable is the average TFP shock within an industry. Each regression includes a set of industry and time fixed effects. Standard errors (below the point estimates) are clustered at the industry level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Then, we use these percentiles to calculate dispersion and skewness. This ensure that our results do not disclose any sensitive information. Furthermore, in order to further avoid any disclosure concerns, none of our results are based in industry-year cells that do not pass the basic disclosure criteria required by the Census Bureau. In practice, since industries are defined at the 3-digit NAICS levels, we do not use any industry-year cell that has 200 establishment observations or less.

C   VAR data and Robustness

In this section, we describe in additional detail the data and methods used to estimate the VAR and impulse responses discussed in Section 7 of the main text and we provide some additional robustness results. The variables we consider in the analysis are the log of the S&P500 stock market index (closing value of last trading date of the month), a measure of stock-market volatility and a measure of stock market skewness (both explained below), the Federal Funds Rate (FRED variable FEDFUNDS), the log-level of the average of hourly earnings (FRED variable AHETPI), the log-consumer price index (FRED variable CPIAUCSL), the log-level of hours (FRED variable AWHMAN), the log-level of employment (FRED variable PAYEMS), and the log of an index of industrial production (FRED variable INDPRO).

We construct measures of volatility and skewness using daily returns from publicly traded firms obtained from CRSP dataset acceded through WRDS. In particular, for each firm i we calculate the day-to-day log-change of the stock price within a month m, $d_{i,m}$ and then we calculate the difference between the 90th-to-10th percentile differential and the Kelley skewness using all observation of daily returns over all the firms within a month m. As an alternative measure, we consider four-weeks log-change of daily prices of firm i within a month m, and then we calculate the 90th-to-10th percentile

differential and the Kelley skewness within a month. All variables in our baseline results are HP detrended using a smoothing parameter of $\lambda = 129,600$, with the exception of the measures of volatility and skewness

Figure A.8 displays a series of robustness results for our VAR analysis. In particular, the top panel A.8 repeats the impulse response in our baseline results (square symbols); the resulting impulse response calculated after dropping all data during and after the Great Recessions (diamond symbols); a case in which we drop the measure of volatility and consider only skewness (x symbols); a case in which we reverse the order of the VAR considering first the measure of skewness, then volatility, and then the S&P500 (triangle symbols); a case in which we consider monthly returns rather than daily returns to construct our measures of volatility and skewness (hollow square symbols); and a case in which we have HP filtered the measures of volatility and skewness (v symbols). In all cases, we find that a decline in the skewness of firms' returns–that can be interpreted as an exogenous shock to firms– generates a persistent decline in industrial activity that overshoots after 24 months. Figure A.8 shows two additional cases, a four variables VAR in which we consider the S&P500, volatility, skewness, and the industrial production indicator or employment (circle symbols), and a five variables VAR in which we consider the S&P500, volatility, skewness, the industrial production indicator, and employment (+ symbols). The effect of a skewness shock in these cases is even stronger relative to the baseline results. The bottom panel of Figure A.8 shows the response of log-employment. Hence the response of industrial production and employment to a skewness shock seems to be robust to different sample selection or variable ordering.

For additional robustness, Figure A.9 shows the response of industrial production and employment to a skewness shock estimated using the Local Projection Method proposed by Jordà (2005). Here, we consider the same variables as in our baseline VAR but we run a set of OLS time series regressions of the form,

$$y_{t+h} \quad = \quad \beta_0 + \beta_{1,h} vol_t + \beta_{2,h} skew_t + \Gamma X_t + \epsilon_t$$

for $h = 0, .., 16$ using monthly data. We then plot the value of $\beta_{1,h}$ and $\beta_{2,h}$–and corresponding confidence intervals–for industrial production and employment. As before, the measure of market volatility is the 90th-to-10th log percentiles differential of the within-month daily returns and the measure of skewness is the Kelley skewness form the same distribution. The results shown in Figure A.9 are similar to those obtained using standard VAR methods, both qualitatively and quantitatively. Hence, we conclude that the response of industrial production and employment to a skewness shock is robust to different estimation methods.

FIGURE A.8 – ROBUSTNESS: MACROECONOMIC IMPACT OF A SKEWNESS SHOCK

(A) Industrial Production

(B) Employment

Note: Figure A.8 shows the impact on industrial production (top panel) and employment (bottom panel) of a two-standard deviation shock to the skewness of the stock returns under different specifications. Baseline (filled squares) considers the same specification as in the main body of the text; Pre 2008 (diamonds) consider the same specification but estimates the VAR using data pre 2008 only; Only skewness (x) drops the measure of volatility from the Baseline; Reverse (triangles) estimates a VAR in which we order skewness first, volatility second, and then the S&P500; Four variables (circles) keeps the S&P500, volatility, skewness, and industrial production (or employment); Five variables (+) adds back employment in both estimations. Monthly (hollow squares) considers a case in which we estimate skewness and volatility using monthly returns; HP filtered (v) consider a case in which the baseline measures of volatility and skewness are HP-filtered using a smoothing parameter of $\lambda = 129,600$.

62

FIGURE A.9 – LOCAL PROJECTIONS: MACROECONOMIC IMPACT OF A SKEWNESS SHOCK

(A) Industrial Production

(B) Employment

Note: Figure A.9 shows the impact of a two-standard deviation shock to the skewness of the stock returns estimated using Local Projections (Jordà, 2005). The data period is 1964 to 2008.