Granular Responses to Aggregate Shocks: An Interpretation of Skewness over the Business Cycle

Niklas Schmitz*

February 27, 2024

Abstract

Skewness in the cross-sectional distribution of firms' sales growth rates is procyclical: In recessions, some firms experience particularly poor growth rate outcomes. This paper studies the role of these poor growth rate performers for aggregate fluctuations. There are two key findings. First, some *large* firms experience very poor sales growth rates in recessions. Because these firms are so large, the vast majority of the decline in sales levels in US recessions is driven by the firms with the worst growth rates. Second, several commonly studied aggregate shocks induce skewed responses across firms and can explain the close comovement of cross-sectional skewness with the business cycle. Importantly, the responses across the very largest firms in the US economy are also skewed, with some large firms showing strong responses to aggregate shocks. Because the firm size distribution is fat-tailed, these large firms account for a significant share of the aggregate fluctuations following aggregate shocks. This provides a new interpretation of how aggregate shocks induce business cycle fluctuations: aggregate shocks transmit through granular responses across firms, with much of the response accounted for by firms that are both large and responsive.

*Faculty of Economics, University of Cambridge. Address: Austin Robinson Building, Sidgwick Ave, CB3 9DD Cambridge, United Kingdom. E-mail address: ns751@cam.ac.uk. I am indebted to Chryssi Giannitsarou and Elisa Faraglia for their invaluable guidance. I thank Vasco Carvalho for extensive discussions and support. I thank Anmol Bhandari, Paul Bouscasse, Leonard Bocquet, Lukas Freund, Maren Froemel, Xavier Gabaix, Fatih Guvenen, Timo Haber, Sean Holly, Hanbaek Lee, Glenn Magerman, Shane Mahen, Mathis Momm, Alba Patozi, Charles Parry, Margit Reischer, and participants at the Cambridge Macro PhD Workshop for helpful comments. I thank Alba Patozi for sharing data with me.

1 Introduction

Skewness in the distribution of sales growth rates across firms (*micro skewness*) is strongly procyclical (Salgado et al. (2023)). Figure 1 plots the aggregate sales growth rate in the U.S. against the cross-sectional skewness across the sales growth rates of U.S. public firms. The correlation is 0.8. Skewness is low in recessions, suggesting some firms face particularly bad growth rates while there is limited potential for good growth rate outcomes.

This paper asks and answers two questions: 1) Does micro skewness matter for aggregate fluctuations? Yes, micro skewness matters for aggregate fluctuations because it captures the poor performance of some large firms in recessions. 2) What is the origin of procyclical micro skewness? I show that firms' heterogeneous responses to aggregate shocks induce procyclical skewness. Importantly, even the responses across the largest firms of the US economy are skewed such that some large firms respond strongly to aggregate shocks. Because the size distribution is fat-tailed, much of the aggregate effects of aggregate shocks can be traced to a handful of large and responsive firms. In this sense, the economy's granular response to aggregate shocks provides an interpretation for the procyclicality of micro skewness.

This paper starts by studying the implications of micro skewness for aggregate fluctuations. A

Figure 1: Micro skewness vs aggregate activity

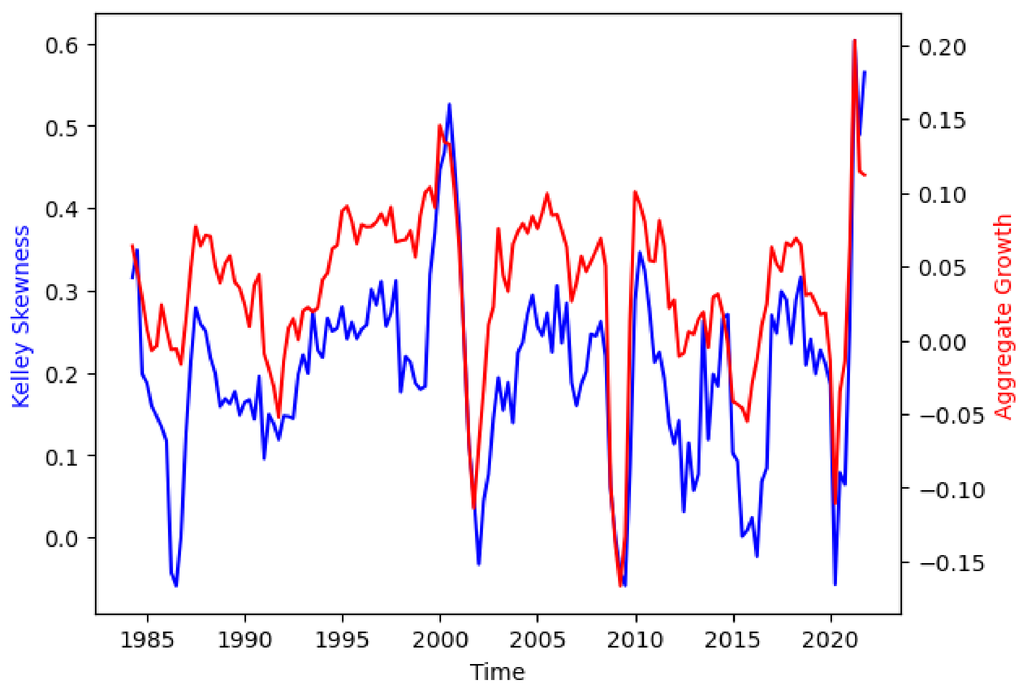

Note: The figure compares skewness across Compustat firms' year-over-year real sales growth to the aggregate sales growth rate. Skewness is estimated using 90% Kelley skewness. The sample period is 1984Q2-2021Q4.

stronger skew towards poor growth rate outcomes in recessions need not be relevant for aggregate sales growth if those firms with poor growth rate outcomes are small. However, there is evidence that at least some large firms experience very weak growth rates in downturns. Using Compustat data, I document the importance of firms that are both large and responsive in several ways. First, the vast majority of the decline in the level of sales in a recession is explained by the worst performers in terms of growth rates. Second, while the largest firms are less volatile than smaller firms, they face a probability of 18% to end up in the bottom 20% of the sales growth distribution in a recession. In contrast, they face almost no upside: The probability of experiencing a growth rate in the top 20% of growth rate outcomes is only 6%. This asymmetry between downside and upside risk exists in every recession in the sample but is only present for large firms. In this sense, large firms can be considered more procyclical, even though they are less volatile than small firms. Third, and in line with this interpretation, the time series of large firms' growth rates are on average less volatile than those of smaller firms but are not less skewed. Fourth, skewness across the largest firms is more closely associated with aggregate fluctuations than for other size bins.

What explains the procyclicality of micro skewness? I propose firms' heterogeneous responses to aggregate shocks as a mechanism. To motivate this hypothesis, I decompose the cross-section of sales growth rates into aggregate drivers and idiosyncratic components and measure the contribution of these components to skewness in growth rates. The aggregate component explains 80% of the business cycle pattern in micro skewness, even though it only accounts for a small share of overall variation in firm growth rates. Skewness in idiosyncratic shocks is therefore not the dominant driver of the skewness pattern. This result does not contrast with recent findings of skewness in the cross-section of TFP shocks (Salgado et al. (2023)). Instead, I argue that these shocks are not large enough to contribute to skewness in growth rates and are dominated by the skewness in aggregate drivers.

To move beyond correlations, I use local projections to study the comovement of the impulse responses of aggregate sales growth and micro skewness following six different types of aggregate shocks: monetary, oil supply, credit, uncertainty, sentiment, and TFP news shocks. These shocks reflect many of the most prominently studied aggregate shocks to explain business cycle fluctuations, are all different in nature, and are obtained using six different identification schemes. In response to any of the six shocks, aggregate growth and micro skewness decline in a closely correlated manner. I also compare the contributions of aggregate shocks to fluctuations in skewness against the contributions from a new series of idiosyncratic shocks. A variance decomposition shows that around 75% of the explained variance can be accounted for by aggregate shocks, compared to 25% for the idiosyncratic shocks. This supports the idea that idiosyncratic shocks can be important sources of business cycle fluctuations (Gabaix (2011), Carvalho & Grassi (2019)) while motivating this papers' focus on aggregate shocks.

Estimating the responses of micro skewness and aggregate sales growth bottom-up from firm-level IRFs confirms these results. I estimate firm-level impulse responses of their sales growth rates to aggregate shocks. Given the distribution of firm IRFs, the implied response of aggregate sales growth can be constructed as the size-weighted average of firm IRFs. Similarly, the response of micro skewness can be constructed from the cross section of firm IRFs. These bottom-up estimates show that both micro skewness and aggregate sales decline significantly following contractionary aggregate shocks. Therefore, heterogeneous responses to aggregate shocks provide an explanation for procyclical micro skewness.

Based on the distribution of firm IRFs, I show that the impulse responses of large firms are more skewed than the responses of small firms. The differences in the responses of skewness can be large. For both size groups, the size-weighted aggregate sales responses are closely correlated with the skewness response. However, because the size distribution is fat-tailed, the responses of large firms account for the vast majority of the aggregate sales growth response. While small firms may be more responsive to shocks on average (stronger unweighted mean response) and have more volatile responses, they have a minor contribution to economy-wide sales growth fluctuations following an aggregate shock.

The skewed responses of the largest firms provide a new narrative for how aggregate shocks induce aggregate fluctuations: A large share of the effect of an aggregate shock can be traced to a small number of very large firms with strong responses to the shock. Based on the firm-level impulse responses, I estimate that large firms (top 10% of the size distribution) with strong growth rate responses (bottom 20% of the IRF distribution) make up only 1.6% of all firms but account for one third of the sales response to an aggregate shock. This result is complementary to the granular hypothesis by Gabaix (2011): While he shows that idiosyncratic shocks to large firms can induce significant business cycle fluctuations, I argue that the effects of aggregate shocks are significantly driven by the responses of some large firms.

Why are some large firms so responsive to aggregate shocks? Identifying the sources of heterogeneity in impulse responses turns out to be a challenging task and this paper makes only a limited contribution in this regard. Most commonly studied firm characteristics do not seem to explain differences in IRF estimates well. Using both standard OLS and a random forest algorithm, profitability emerges as a potential exception: Firms with lower profitability experience significantly larger sales declines following contractionary aggregate shocks. However, the models struggle to explain most of the heterogeneity in the data and more work will be needed to pin down the sources of large firms' vulnerabilities.

The results of this paper are based on Compustat data. Compustat is the most widely used data set for US public firms and provides multiple benefits. It covers a long sample period at the quarterly frequency and contains rich information on firm characteristics. This is indispensible for studying the role of different firm characteristics in shaping the responsivess of firms to reliably identified aggregate shocks. The most obvious drawback of Compustat is the underpresentation of small firms. This bias does not affect the results of this paper since the focus is on the largest firms in the economy. However, in comparing the responses of the very largest firms to the rest of the sample, I will commonly refer to 'large' and 'small' firms throughout the text. This is for ease of reading only and these labels should be interpreted within the confines of the firm size distribution that Compustat provides. The 'small' firms are significantly larger than the average firm in the United States and 'large' firms represent the very top of the size distribution.

An extensive appendix confirms the procyclicality of micro skewness in the Compustat sample in detail. I pay particular attention to the robustness across different skewness measures and the role of outliers. Because the third moment is highly sensitive to outlier observations, I find that skewness measures based on quantiles provide a more reliable assessment of asymmetries in the distribution. Outliers are pervasive in the Compustat data and require careful cleaning procedures.

As I demonstrate, the wrong treatment of outliers can dramatically change the results for skewness measurement.

The appendix also provides additional motivation for studying micro skewness. The cross-sectional dispersion, a common object of study in the firm heterogeneity literature (for example, see Bloom et al. (2018)) is uncorrelated with aggregate fluctuations once controlling for micro skewness. The association between skewness and aggregate fluctuations increases with the level of granularity, and is significantly weaker when measuring skewness across 3-digit NAICS industries instead of across firms. The relation between micro skewness and aggregate growth is strongest contemporaneously, with weaker evidence for lead-lag relationships.

This paper connects to several strands of literature. Most directly, this paper can be placed within a small literature on the business cycle comovement of firm-level micro skewness. Higson et al. (2002) and Higson et al. (2004) provide evidence of countercyclical skewness for UK and German firms. In contrast, Salgado et al. (2023) offer a comprehensive analysis for over 40 countries and find strong evidence in favor of procyclical skewness. They offer skewed distributions of idiosyncratic shocks as a possible explanation for this fact. My own analysis confirms the procyclicality of skewness and suggests that differences in outlier treatment may explain the contrasting findings. Ilut et al. (2018) use data on US manufacturing establishments to provide evidence of concave firm employment responses to aggregate shocks. Similarly, my paper provides evidence on the role of heterogeneous sales responses to aggregate shocks in explaining procyclical skewness.

By focusing on the cross section of firm outcomes, my paper is also related to a large literature on uncertainty and dispersion. Bloom (2009) revived this literature, followed by important contributions such as Bachmann & Bayer (2014), Jurado et al. (2015), Bloom et al. (2018), Berger et al. (2020), and Ludvigson et al. (2021). Fernández-Villaverde & Guerrón-Quintana (2020) and Cascaldi-Garcia et al. (2023) provide comprehensive reviews. This literature often uses the countercyclical movements in cross-sectional dispersion as a proxy for uncertainty.[1] However, dispersion becomes acyclical once controlling for skewness, suggesting that changes to micro skewness are the more reliable business cycle fact. The same finding has already been made for the income distribution (Guvenen et al. (2014), Guvenen et al. (2022), Busch et al. (2022)).

This paper's focus on heterogeneous responses to aggregate shocks is directly motivated by a large literature on firms' heterogeneous responses to monetary shocks. Recent examples include Ottonello & Winberry (2020) and Cloyne et al. (2023), who pay great attention to identifying a particular firm characteristic that best explains differences across firm responses. While these authors focus on the details of the transmission mechanism of one particular shock, I take a broader perspective by focusing on a wide range of commonly studied aggregate shocks. This comes at the cost of less depth as I cannot identify one individual firm characteristic that explains firm's heterogeneous responses across six different shocks, with profitability being a potential exception.

Most importantly, this paper contributes to the literature on large firm dynamics. Gabaix (2011) proposes the granular hypothesis and Carvalho & Grassi (2019) demonstrate within a standard firm dynamics model that idiosyncratic shocks can account for a sizeable share of aggregate fluctuations.

[1] Kozeniauskas et al. (2018) demonstrate how shocks to cross-sectional dispersion relate to disagreement and uncertainty about macroeconomic aggregates.

Crouzet & Mehrotra (2020) use a representative sample of the US firm size distribution to show that the very largest firms of the US economy account for nearly all of the business cycle fluctuations in terms of sales. While large firms are less volatile than small firms, the authors argue that the difference in volatility is far outweighed by the larger size of large firms. Relatedly, I argue that small firms are more volatile but not necessarily more skewed: large firms' responses to contractionary shocks are predominantly negative, and among the largest firms some firms respond particularly strongly. Combined with a fat-tailed size distribution, this means that the response of aggregate sales to aggregate shocks is strongly driven by the response of (some) large firms. The strong correlation between cross-sectional skewness and aggregate fluctuations can therefore be interpreted as (at least partly) originating from the strong response of some large firms to aggregate shocks, which explains both the change in skewness in the cross section and the change in aggregate sales.

Lastly, the procyclicality of micro skewness may provide a micro-founded perspective on the 'plucking property' of the business cycle (McKay & Reis (2008), Dupraz et al. (2023)): recessions are deeper than expansions, resulting in negative skewness in the time series of aggregate growth. Similarly, Adrian et al. (2019) show that upside risks to aggregate growth are stable over time while downside risks fluctuate with the business cycle. To achieve negatively skewed aggregate growth with symmetric aggregate shocks, Ilut et al. (2018) suggest concave firm-level responses while Baqaee & Farhi (2019) and Dew-Becker et al. (2021) suggest aggregation in production networks if production inputs are complements. Alternatively, the findings from this paper can motivate further research into the properties of large and responsive firms and their role in shaping the asymmetric response of aggregate outcomes. Dew-Becker (2022) estimates asymmetries in conditional distributions of firm and aggregate outcomes using option data, finding that micro skewness is significantly related to macro volatility and concludes there should be a common shock driving both. I show that aggregate shocks resulting in micro skewness may provide an explanation for this fact when studying *realized* distributions, though it is unclear if the same holds for conditional distributions.

The paper proceeds as follows. Section 2 shortly describes the Compustat sample. Section 3 provides new stylized facts on the downside risks faced by large firms and their contribution to the fluctuations in aggregate sales. Section 4 studies the role of heterogeneous responses to aggregate shocks in explaining the comovement between aggregate sales and micro skewness. Section 5 concludes.

2 Data

The analysis uses data on US public firms from Compustat. Compustat is the benchmark firm-level data set for the United States, providing detailed balance sheet information at the quarterly frequency over a long sample period. Rich information on firm characteristics is necessary to comprehensively study the origins of heterogeneous firm responses to aggregate shocks. The quarterly frequency greatly improves the ability to identify relevant macroeconomic shocks relative to annual data. The long sample period enables me to cover multiple recessions and draw general conclusions about skewness facts in the US business cycle. Estimating impulse responses to aggregate shocks at the firm level also requires a sufficiently long time series for each firm. I am not aware of other data sets satisfying these criteria.[2]

Let $s_{i,t}$ be firm i's real sales in quarter t. Real sales will be the key measure of firm size in this paper. Year-over-year real sales growth is $g_{i,t} = \ln(s_{i,t}) - \ln(s_{i,t-4})$. The key business cycle indicator of this paper is aggregate real sales growth, constructed as the size-weighted average of existing firms' growth rates:

$$g_t = \frac{\sum_i g_{i,t} s_{i,t-4}}{\sum_i s_{i,t-4}}. \tag{1}$$

This definition of aggregate sales growth only considers firms that exist in both t and $t-4$ and therefore abstracts from entry and exit dynamics, which could affect the comovement of aggregate growth and micro skewness but are not the focus of this paper.

The main skewness measure is Kelley skewness:

$$ksk(G_t) = \frac{\left(Q_{0.9}^{G_t} - Q_{0.5}^{G_t}\right) - \left(Q_{0.5}^{G_t} - Q_{0.1}^{G_t}\right)}{Q_{0.9}^{G_t} - Q_{0.1}^{G_t}}, \tag{2}$$

where $G_t := \{g_{i,t}\}_{i=1,\dots,n_t}$ is the set of firm growth rates at time t. Kelley skewness compares the distance of the 90% quantile of the time-t distribution of firm growth rates ($Q_{0.9}^{G_t}$) from the median ($Q_{0.5}^{G_t}$) to the distance of the median from the 10% quantile, rescaled by the overall spread of the distribution ($Q_{0.9}^{G_t} - Q_{0.1}^{G_t}$). If the 90% quantile is further above the median than the 10% quantile is below the median, the distribution is right-skewed and Kelley skewness is positive. The measure ranges from -1 to 1. Kelley skewness allows for an easy decomposition of skewness movements into changes in upper and lower parts of the distribution and is significantly more robust to outliers than the third moment.

Compustat data features severe outliers for sales growth rates. I carefully check for these outliers and compare the implications of different cleaning methodologies and skewness definitions for the procyclicality of micro skewness. Appendix B contains the results. In summary, outliers in growth rates can be due to data mistakes, sales increases from low base levels, or M&A activity. The third moment is highly sensitive to the presence of these outliers. After removing the outliers and especially when using a quantile-based skewness measure such as Kelley skewness, micro skewness is highly procyclical across a variety of specifications. This supports the evidence from Salgado et al. (2023) and is in line with findings from the household income literature (Guvenen et al. (2014), Guvenen et al. (2022), Busch et al. (2022)).

Details on the sample construction are in Appendix A. Besides Compustat, I use data from CRSP for stock prices, I/B/E/S for sales forecasts, and Worldscope Fundamentals because of its good coverage of the date of incorporation. All variable definitions are listed in the appendix. The data cleaning filters out roughly half of the observations from the raw Compustat files. Since estimating firm-level impulse responses requires a sufficiently long time series for each firm, I focus on firms that have at least 40 consecutive observations for sales growth. This reduces the sample size further, see Figure 13. Despite the smaller sample size, the time series of micro skewness looks very similar, see Figure 14 in Appendix A.

Table 1 compares the full Compustat sample against the cleaned version of firm growth streaks.

[2] Ottonello & Winberry (2020) use Compustat data for the same reason.

For comparison, the table also reports summary statistics from the Quarterly Financial Reports (QFR), which have been used by Crouzet & Mehrotra (2020) to construct a representative sample of US firms in certain sectors. For example, the QFR can be used to construct a sample accurately reflecting the firm size distribution of US manufacturing firms, including private firms. Relative to this representative sample of manufacturing firms, the average firm in the Compustat data (which is not limited to manufacturing firms) is considerably larger, both in terms of assets (USD 6bn vs USD43mln) and sales (USD 740mln vs USD 11mln). The sales growth distribution in the QFR sample is more dispersed and more symmetric than in the Compustat sample with a mean growth rate closer to zero. Compared to the QFR, leverage and short-term debt are higher in raw Compustat data but lower in the cleaned data. The number of observations in the cleaned data is roughly half of the number of observations per quarter in the QFR. Importantly, although the data cleaning affects multiple firm characteristics on average, the correlation between aggregate sales growth and GDP growth is similar for both Compustat samples (0.68 vs 0.54). The correlation between micro skewness and aggregate sales growth, which is the key object of study in this paper, is virtually identical for both samples (0.84 vs 0.82).

While small firms are strongly underrepresented in Compustat, the results of this paper should

Table 1: Summary Statistics for Compustat Data

	Full Compustat	Cleaned Sample	QFR
Assets (mln. USD)	5,838	6,366	43.2
Sales (mln. USD)	467.4	741.9	10.8
Sales Growth (%)	7.3	6.2	0.63
$Q\,(\text{Sales Growth})_{0.25}$ (%)	-7.6	-5.6	-25.3
$Q\,(\text{Sales Growth})_{0.75}$ (%)	23.5	17.0	26.6
Net Leverage (%)	66	2	20
Short-term debt (%)	34	6	33
Obs./quarter	6,329	2,766	6,122
$\rho\,(\text{Sales Gr., GDP Gr.})$	0.68	0.54	–
$\rho\,(\text{Sales Gr., Skew})$	0.84	0.82	–

Statistics for QFR are for the manufacturing subset of Crouzet & Mehrotra (2020) from 1977Q3–2014Q1 and directly taken from their paper. Compustat statistics are for 1983Q3–2014Q1.

be unaffected by this bias. Small firms do exit more frequently, meaning the sample may understate the relative downside risk of large vs small firms. However, the focus of this paper is neither on entry/exit dynamics nor on small firms per se. Instead, I argue that the responses of large firms to aggregate shocks are also skewed and that this feature is crucial for the comovement of micro skewness and aggregate growth. The correlation between micro skewness and growth – which is the focus of this paper – is based on aggregate growth computed using only surviving firms, allowing to exclude entry/exit dynamics from both moments. Despite the underpresentation of small firms, sales concentration in the sample is still high. The largest 10% of firms account for 70% of sales on average, and the top 30% account for over 90% of sales. For comparison, the largest 1% of firms in the QFR sample of Crouzet & Mehrotra (2020) represent ca. 75% of total sales.

An important clarification for the remainder of this paper is in order. For ease of writing, the rest of this paper refers to ’small’ and ’large’ firms. These terms should be interpreted within the confines of the size distribution that Compustat allows to study, acknowledging that ’small’ firms in Compustat are significantly larger on average than small firms in a representative sample. I do not suggest that the data allows for an accurate comparison of truly large and truly small firms, but instead use the terms to refer to relative sizes within the data, arguing that the large firms within Compustat are still orders of magnitude larger than the other firms. Since the smallest firms in Compustat are often startups and may have differ from typical small firms across a range of features, I abstain from directly comparing the largest firms to the smallest firms in the data. Instead, I generally focus on comparing the top of the size distribution to the rest of the distribution.

Appendix B contains additional stylized facts on the procyclicality of skewness. I show that the relation between micro skewness and aggregate growth is strongest when micro skewness is measured across firms and becomes weaker at higher levels of aggregation. A univariate regression of aggregate growth on skewness measured across firms explains 70% of aggregate fluctuations. If skewness is measured across 2-digit NAICS industries instead, the R^2 declines to 20%. Micro skewness is highly correlated with contemporaneous aggregate growth and shows no significant lead-lag relationship beyond the four-quarter horizon. In addition, the procyclicality of skewness is responsible for previous findings of countercyclical dispersion. Cross-sectional dispersion has been a popular proxy in the large literature on uncertainty shocks (Bloom (2009), Bloom et al. (2018)). However, dispersion is acyclical conditional on controlling for skewness. This matches evidence from the household income literature (Guvenen et al. (2014)) and suggests that skewness instead of dispersion is the relevant moment of the cross-section to study. Taken together, these findings motivate the focus on the comovement of firm-level micro skewness with the business cycle.

3 The importance of micro skewness for macro fluctuations

Does micro skewness matter for the business cycle? One hypothesis is that micro skewness could matter for aggregate fluctuations if large firms experience bad growth rates. This section presents evidence supporting this idea. While large firms are commonly found to be less cyclical than small firms, this difference can be small, see Crouzet & Mehrotra (2020). Table 2 sorts firms by their level of real sales before a recession and reports the frequency with which these firms end up in the five quintiles of the cross-sectional growth rate distribution measured at the bottom of a recession.[3] Small firms in Compustat experience a U-shaped probability distribution: They face higher probabilities of ending up in either tail of the growth rate distribution in a recession than being in the middle of the distribution. Medium-sized firms in the sample (the middle 80% in terms of real sales) face a uniform distribution. Large firms’ probability distribution is unimodal: They are more likely to have a growth rate in the middle of the distribution than in the tails. However, the probability of the largest firms in Compustat to experience a poor growth rate is still sizeable: 17% of large firms have

[3] Recessions are defined as significant declines in aggregate sales growth and need not coincide with the NBER recession dates. The dates of the troughs are 1991Q4, 2001Q4, 2009Q2, 2015Q4, and 2020Q2. I cannot consider the sales decline in 1984 because there are not enough pre-recession observations to group firms into size bins.

a growth rate in the bottom 20% of the distribution, compared to 19% for medium-sized firms.[4]

Importantly, large firms are the only size bin which faces an asymmetry: While the downside

Table 2: Growth rate outcomes by size groups

	Size groups		
Growth Quintiles	Bottom 10%	Middle 80%	Top 10%
1	0.33	0.19	0.17
2	0.11	0.21	0.24
3	0.08	0.20	0.31
4	0.11	0.21	0.22
5	0.37	0.19	0.06
	Bottom 30%	Middle 40%	Top 30%
1	0.26	0.20	0.16
2	0.15	0.21	0.25
3	0.13	0.20	0.28
4	0.14	0.22	0.24
5	0.32	0.18	0.08

Size groups are defined based on average real sales over the three years preceding each recession. Recessions are defined as strong declines in aggregate sales growth, and the sales growth quintiles are measured at the trough of the sales growth decline.

Table 3: Growth rate outcomes for largest firms by recession

	Growth Quintiles for Top 10%				
Recession	1	2	3	4	5
1990	0.12	0.26	0.40	0.18	0.04
2000	0.12	0.20	0.35	0.25	0.07
2008	0.18	0.23	0.27	0.24	0.08
2014	0.24	0.27	0.27	0.20	0.02
2020	0.18	0.21	0.26	0.24	0.12
	Growth Quintiles for Top 30%				
	1	2	3	4	5
1990	0.13	0.26	0.32	0.22	0.07
2000	0.12	0.24	0.28	0.27	0.09
2008	0.17	0.26	0.25	0.23	0.09
2014	0.21	0.26	0.28	0.21	0.05
2020	0.16	0.21	0.27	0.25	0.11

Size groups are defined based on average real sales over the three years preceding each recession. The largest firms are defined as either the top 10% of firms (top panel) or the top 30% (bottom panel). Recessions are defined as strong declines in aggregate sales growth, and the sales growth quintiles are measured at the trough of the sales growth decline.

[4] Even though the data may understate the exit probability for small firms, it is still striking that small firms are less likely to remain in the middle of the growth rate distribution than at the top end or the bottom. A representative sample of firm exit would increase the presence of firms in the bottom quintile of the distribution and decrease transition probabilities to other bins equally, therefore not changing the relative probabilities of average versus good performance.

risk for large firms is comparable to that of smaller firms, they face virtually no upside. Only 6% of large firms have growth rates in the top 20% of the growth rate distribution. This pattern is robust to extending the definition of large firms to the largest 30% of firms, although the asymmetry weakens. Table 3 confirms that large firms face this asymmetry in every recession in the sample.

The asymmetry in growth rate outcomes for large firms does not occur outside of recessions.

Figure 2: Share of real sales by growth rate quintiles

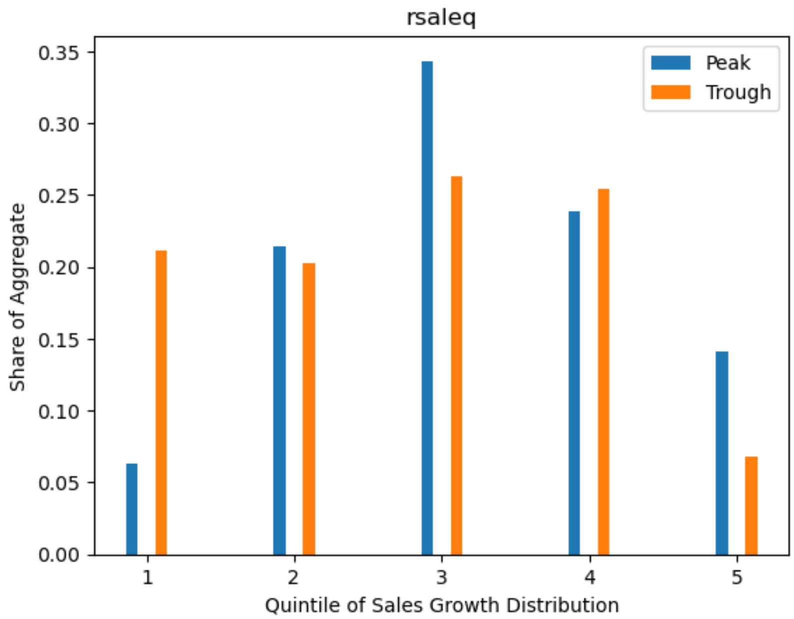

The share of real sales accounted for by the different quintiles of the cross-sectional growth rate distribution resembles a bell curve before recessions, as shown in Figure 2. At the trough of a recession, the share of real sales accounted for by firms in the bottom quintile of the growth rate distribution nearly quadruples relative to the peak, and the share of the top quintile halves. In other words, the representation of large firms increases in the bottom quintile and decreases in the top quintile when the economy is in a recession instead of in a boom.

The reallocation of firms across growth rate quintiles does not occur for many other firm characteristics. Figure 3 repeats the same exercise for firm age. Firm age in each growth rate bin does not change materially between a boom to a bust. In unreported robustness checks, I confirm that other characteristics including leverage, liquidity, Tobin's Q, dividend payer status, long-term debt, and book-to-market ratio also show no clear pattern.

The poor growth rate outcomes of large firms in recessions mean that those firms with the poorest growth rate realizations account for the majority of the decline in the level of sales in a downturn. Figure 4 demonstrates this result. This is not obvious: If large firms clustered in the middle of the sales growth distribution and the bottom quintile was only populated by the smallest firms, the sales decline could mostly be explained by the middle quintile of the growth rate distribution. The fact that sales levels decline most strongly due to the performance of the firms with the worst growth

Figure 3: Share of firm age by growth rate quintiles

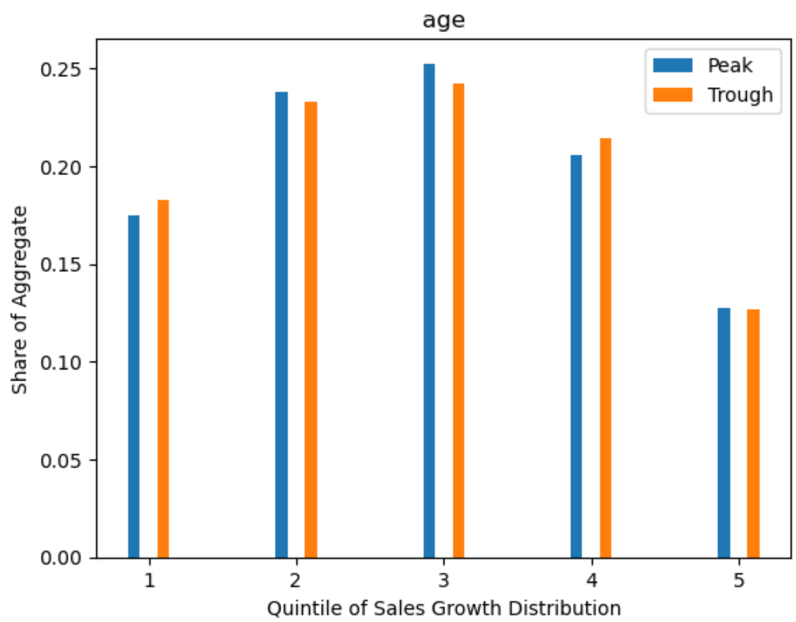

rates reflects that some large firms are subject to sizeable downside risks in recessions.

These facts may be surprising: Large firms are generally perceived as less volatile than small

Figure 4: Change in sales levels by growth quintiles

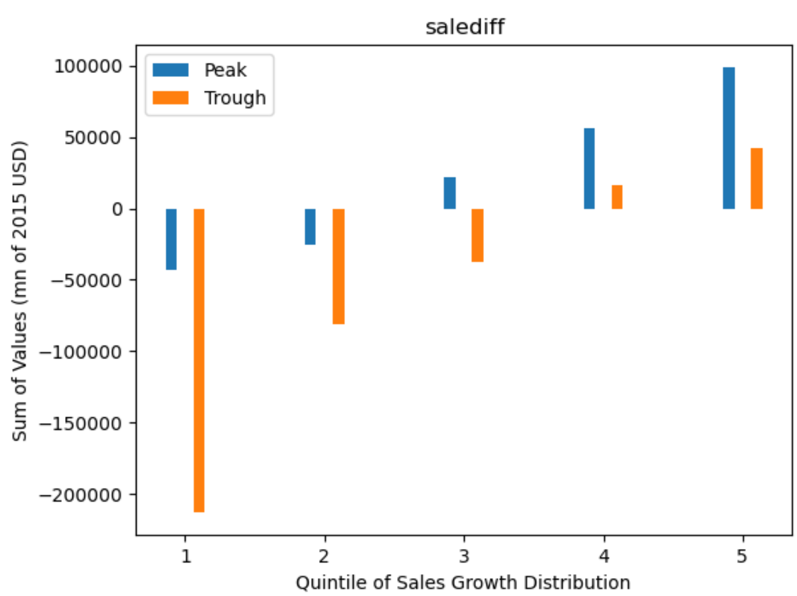

firms but appear to not face significantly lower downside risk. In this context, the distinction between volatility and skewness is simple yet crucial: Volatility is symmetric and can be large due to large positive or negative movements. The larger volatility of small firms is not explained solely by stronger downside risks for small firms but instead represents both higher upside *and* downside risks. To see this point more clearly, I compute the average skewness in the time series of firm growth rates within different size bins. Define this average as:

$$\overline{ksk}\,(J) = \frac{1}{n_J}\sum_{i \in J} ksk\,(G_i) \quad \text{with} \quad G_i := \{g_{i,t}\}_{t=1,\ldots T_i} \quad \text{for} \quad J = \{\text{top}, \text{bottom}, \text{middle}\}$$

Analogously, we can define the average dispersion, $\overline{disp}\,(J)$ using the dispersion measure $Q^{G_i}_{0.9} - Q^{G_i}_{0.1}$. Table 4 reports average skewness (in index points) and average dispersion (in percent) for different size bins. Since there is no obvious way to define the size groups for the time series of growth rates, I use two different approaches and show that the results are stable across them. The top panel groups firms into size bins based on the firms' real sales observed when entering the sample such that a firm is considered large if it was in the top x% of the real sales distribution in its first quarter of existence. This especially captures firms that are large throughout the sample but neglects those firms that became large over the sample period. The bottom panel groups firms into the respective size group if they are within that group for at least 80% of their quarters. Across both panels, small firms are significantly more volatile than the largest firms (ca. 60% vs 30% dispersion). However, average skewness in the time series of large firms is similar or even slightly more negative.

This pattern can inform the debate on the importance of large firms for aggregate fluctuations.

Table 4: Time series skewness and dispersion for large vs small firms

	Top 10%	Bottom 90%	Top 30%	Bottom 70%	Middle 80%
			Measured at start of sample		
Skew	-0.06	0.05	-0.02	0.06	0.04
Dispersion	33	60	36	64	54
			Within category 80% of time		
Skew	0.00	0.05	0.03	0.05	0.05
Dispersion	29	60	35	64	54

Skewness is in index points, dispersion is in percent.

Crouzet & Mehrotra (2020) find that large firms are less volatile than small firms but that the differences in volatility are small while differences in size are large. Aggregate fluctuations are therefore almost entirely explained by the largest firms. My finding is complementary: The time series of small firms is more volatile but not more skewed. The larger volatility of small firms is therefore not due to a stronger downward movement in recessions that is common across small firms. Instead, the higher volatility of small firms' growth rates may be largely driven by idiosyncratic components that is unrelated to the business cycle. These fluctuations average out across these firms, making the average growth rate across small firms less correlated with aggregate growth.

The previous findings raise the question if there can be a more effective business cycle statistic than skewness across the growth rates of all firms. Table 5 regresses aggregate sales growth on cross-sectional skewness computed only within certain size bins. Skewness across growth rate outcomes for the smallest 10 or even 30% of firms has no significant association with aggregate growth conditional on skewness in the other size bins. Skewness across growth rates of the largest 30% of firms is most strongly associated with the business cycle. Skewness for the largest 30% of firms explains 70% of the variation in aggregate sales growth, and skewness in the other size bins adds little explanatory power.

In summary, the largest firms in the Compustat data face similar downside risk to other firms

Table 5: Regressing aggregate activity on skewness by size groups

	10-80-10		30-40-30		Top 10%		Top 30%	
Size Group	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Low	0.05	0.05	0.11	0.07				
Middle	0.47	0.10	0.31	0.10				
High	0.43	0.09	0.51	0.09	0.80	0.10	0.84	0.10
R^2	0.73		0.74		0.64		0.70	

Size groups are defined based on average real sales over the previous three years. *10-80-10* separates the sample into the bottom 10%, middle 80%, and top 10% of firms by size, and computes skewness within each size group. *30-40-30* groups firms into bottom 30%, middle 40%, and top 30%. All variables are standardized. All regressions include a constant. Standard errors are Newey-West. R^2 is adjusted for the number of predictors.

during recessions but have significantly less upside potential. This asymmetry implies that sales concentrate more at the bottom end of the growth rate distribution in downturns. The decline in sales levels during downturns is therefore heavily driven by those firms with the worst growth rates. In general, the growth rates of large firms are less volatile than those of smaller firms but not less skewed. The performance of large firms may therefore be most informative about aggregate fluctuations. In fact, skewness across large firms is more closely associated with aggregate fluctuations than skewness among smaller firms. Skewness across firm outcomes therefore matters for aggregate fluctuations because some large firms are among the worst growth rate performers. The next section studies if those large firms may be performing poorly because they respond strongly to aggregate shocks.

4 The response of micro skewness to aggregate shocks

There are different reasons for why skewness in the cross section may arise. Firms could face skewed idiosyncratic shocks: $y_i = \beta\varepsilon + \varepsilon_i$, where ε is an aggregate component, and ε_i is an idiosyncratic component with a skewed distribution $\{\varepsilon_i\}_{i=1,\dots,N}$, as considered in Salgado et al. (2023). Alternatively, firms could have heterogeneous responses to an aggregate shock: $y_i = \beta_i * \varepsilon + \varepsilon_i$.[5] This second case is motivated by a growing literature documenting heterogeneous responses of firms to, for example, monetary policy shocks (see Ottonello & Winberry (2020) or Cloyne et al. (2023)).

The goal of this section is to quantify the importance of these two explanations for the procyclicality of micro skewness. The results can be summarized as follows. Variation in micro skewness can largely be accounted for by a common factor, even though this factor does not explain most of the variation in firm-level sales growth rates themselves. A wide range of aggregate shocks cause significant movements in micro skewness and aggregate sales growth and induce a close correlation between the two. Firms' heterogeneous responses to aggregate shocks appear to be an important

[5] A special case of this version is proposed by Ilut et al. (2018), who suggest firms face nonlinear but identical decision rules of the form $y_i = f(\beta\varepsilon + \varepsilon_i)$, where $f(\cdot)$ is concave.

driver of micro skewness, which I confirm with a variance decomposition of the skewness index that compares the contributions of aggregate versus idiosyncratic shocks, and with responses of skewness and aggregate sales growth constructed bottom-up from firm-level IRFs. The response of large firms to aggregate shocks is also skewed and large firms account for almost all of the sales growth response following aggregate shocks. This suggests the responsiveness of some large firms plays an important role in the effect of aggregate shocks on macroeconomic outcomes. It is less clear why some large firms are so responsive, but the role of profitability may be a good starting point for further exploration.

4.1 A skewness decomposition into common vs idiosyncratic factors

I start by studying the importance of skewness in idiosyncratic components. The evidence from the previous literature is mixed: Ilut et al. (2018) find no significant skewness in establishment-level TFP shocks, while Salgado et al. (2023) argue for strong procyclical skewness in TFP shocks computed using various methods. Both approaches suffer from a related criticism: Even if TFP shocks are not skewed, there may be other idiosyncratic shocks with a skewed distribution that drive skewness in sales growth rates; even if TFP shocks are skewed, their contribution to sales growth rates may be minute because the shocks are small.[6] Focusing on skewness in a particular idiosyncratic shock will therefore not provide conclusive evidence about whether skewed idiosyncratic shocks cause micro skewness unless the shock is found to be both skewed and explain a significant share of variation in sales growth rates.

Instead, my approach is more general: It starts from a generic decomposition of observed sales growth rates into aggregate and idiosyncratic components. I use principal component analysis to capture (possibly heterogeneous) responses of sales growth rates to aggregate factors and interpret the residual as capturing any idiosyncratic variation: $g_{i,t} = \gamma_i + a_{i,t} + \varepsilon_{i,t}$ with $a_{i,t} := \lambda_i F_t$.[7] The two components can then be used to study their impact on cross-sectional skewness. The results obtained this way are conservative in the sense that the idiosyncratic component may still contain aggregate fluctuations that firms could respond to in a nonlinear or heterogeneous fashion. However, the idiosyncratic component is certain to capture all firm-specific sources of variation.[8] If skewness in idiosyncratic shocks affects skewness in sales growth rates, the idiosyncratic component must explain a significant share of the skewness in growth rates. Any estimates from this approach for the importance of skewness in the idiosyncratic component in explaining skewness in growth rates therefore provide an upper bound.

Table 6 shows the results. Skewness in the aggregate components (across $a_{i,t}$) correlates closely with skewness in growth rates once sufficient aggregate factors are included in the decomposition, while the comovement between skewness in the idiosyncratic components and in the growth rates decreases strongly with the number of factors (rows 1 and 2).

[6] Panel regressions in Salgado et al. (2023) confirm this intuition. The skewness in TFP shocks explains virtually none of the variation in firm-level sales, employment, or investment growth as observed from the R^2 values of zero reported in Table 2 of their paper.

[7] The same approach is used in Herskovic et al. (2016) to extract the idiosyncratic component of sales growth rates.

[8] This is true except under a network perspective in which idiosyncratic shocks may cause comovement across firms that is perceived as aggregate fluctuations by the PCA algorithm. See Foerster et al. (2011) for a discussion of this point.

Table 6: Common vs Idiosyncratic Drivers of Skewness

No. Factors:	1	4	8
Correlations with skewness:			
$\rho\,(ksk_\varepsilon, ksk_g)$	0.66	0.36	0.33
$\rho\,(ksk_a, ksk_g)$	0.65	0.81	0.83
Decomposition of variation in skewness:			
R^2_ε	0.18	0.15	0.21
R^2_a	0.82	0.85	0.79
Fit of aggregate component:			
R^2_i q(0.25)	0.02	0.11	0.22
R^2_i q(0.5)	0.08	0.22	0.34
R^2_i q(0.75)	0.18	0.37	0.49
Observations		101,794	

Each column refers to a decomposition using a different number of principal components. The decomposition using the weighted PCA algorithm of Delchambre (2015) with zero weights for missing values and unit weights for all other observations. The first two rows measure the correlation of 90% Kelley skewness in sales growth rates with the skewness in the idiosyncratic components (ksk_ε) or the aggregate components (ksk_a). The following two rows decompose the variation in Kelley skewness into the contributions by skewness in the idiosyncratic part and skewness in the aggregate part. The last three rows show the 25, 50, and 75% quantile of the distribution across R^2 from firm-level time series regressions of the sales growth rate onto the aggregate component. The number of observations refers to the hypothetical balanced panel, of which 5.6% are missing.

Figure 5: Skewness in common vs idiosyncratic component

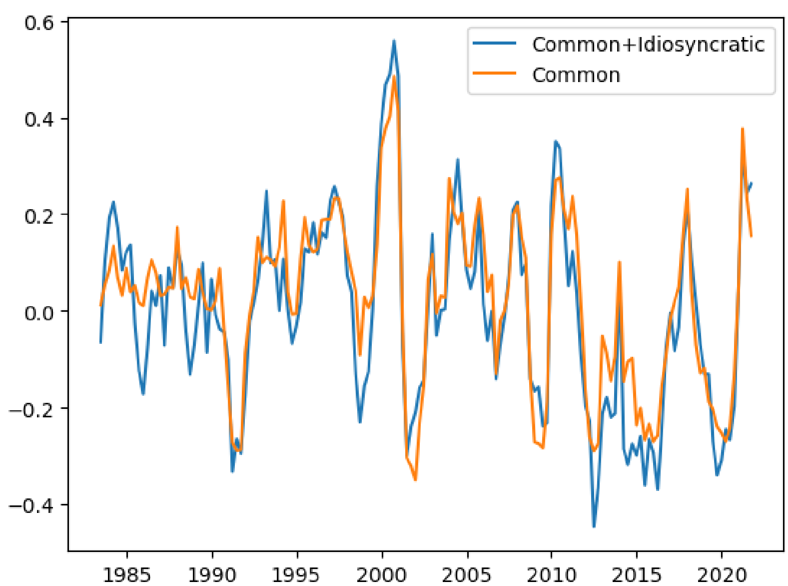

Correlations can be deceiving because comovement patterns may be strong while magnitudes of variation differ. To analyze which component explains most of the variation in Kelley skewness, I decompose the numerator of the skewness measure. The numerator is the component representing asymmetries in the distribution, while the denominator is solely a scaling factor ensuring Kelley skewness always lies between -1 and 1. Let the numerator of the Kelley skewness expression be $\eta(x) = Q^x_{0.9} - 2Q^x_{0.5} + Q^x_{0.1}$, where Q^x_τ indicates the τ-quantile of $x := \{x_i\}_{i=1,\ldots,N}$. The decomposition is then

$$\frac{\eta(g_t)}{Q^{g_t}_{0.9} - Q^{g_t}_{0.1}} = \frac{\eta(\gamma)}{Q^{g_t}_{0.9} - Q^{g_t}_{0.1}} + \frac{\eta(a_t)}{Q^{g_t}_{0.9} - Q^{g_t}_{0.1}} + \frac{\eta(\varepsilon_t)}{Q^{g_t}_{0.9} - Q^{g_t}_{0.1}} + \Delta_\gamma + \Delta_{a_t} + \Delta_{\varepsilon_t}, \tag{3}$$

where γ, a_t and ε_t refer to the distributions of the constant, the aggregate and the idiosyncratic component. Because the ordering of firms within these three distributions may change relative to the ordering of sales growth rates, the decomposition is not exact. The difference is captured by approximation errors

$$\Delta_a = \left(Q^{\tilde{a}_t}_{0.9} - Q^{a_t}_{0.9} + 2\left(Q^{\tilde{a}_t}_{0.5} - Q^{a_t}_{0.5}\right) + Q^{\tilde{a}_t}_{0.1} - Q^{\tilde{a}_t}_{0.1}\right) / (Q^{g_t}_{0.9} - Q^{g_t}_{0.1}) \tag{4}$$

with $Q^{\tilde{a}_t}_\tau$ denoting the aggregate component of the τ-quantile of the growth rate distribution g_t, and by Δ_γ and Δ_ε, which are defined analogously to Δ_a. Given these objects, we can compute partial contributions to explained variance in growth rate skewness. Of the skewness that is unexplained by the constant or the approximation error, the idiosyncratic component explains only 20%. The remaining 80% of unexplained variation are attributed to skewness in the common factors.[9] This decomposition result is stable across the number of aggregate factors used and even holds for the case of only one aggregate factor. To stress the importance of aggregate factors in driving micro skewness, Figure 5 shows that skewness in the idiosyncratic component adds little information beyond the procyclical pattern present in skewness of the common component.

The weak contribution of the idiosyncratic component is not due to a small size of that component. For most firms, the idiosyncratic component remains large after removing the aggregate factors. The last three rows of Table 6 show the 25,50, and 75% quantile of the distribution of R^2 values from firm-level time series regressions of the de-meaned sales growth rate onto the aggregate factors. Even when including eight factors, the aggregate component explains no more than 35% of time series variation for half the firms, and only explains more than 49% of variation for only 25% of firms. To emphasize: One aggregate factor explains 80% of the variation in micro skewness even though it only explains 10% of firm-level sales growth variation on average.

4.2 Growth-skewness comovement due to aggregate shocks

Can aggregate shocks explain the comovement of micro skewness and aggregate sales growth? This section estimates impulse responses of skewness and growth to monetary, oil, credit, uncertainty, sentiment, and TFP shocks. These shocks are different in nature and timing, constructed using varying identification schemes and sample periods. I find that all shocks induce a close comovement pattern between skewness and growth that is at least as strong as measured in the raw data.

[9] Because the skewness of the different components is not orthogonal, the explained variance attributed to each component depends on the ordering of the variables. The results presented here order the idiosyncratic component first to give conservative results for the aggregate component. Ordering the aggregate component before the idiosyncratic component yields 83% of explained variation for the aggregate component and 17% for the idiosyncratic component in the model with 8 factors.

I estimate the impulse responses of skewness and sales growth using local projections (Jordà (2005)):

$$y_{t+h} = \alpha_h + \beta_h \text{shock}_t + \sum_{\ell=1}^{L} \gamma'_{\ell,t} \text{controls}_{t-\ell} + e_{t+h} \tag{5}$$

for $h = 0, ...H$. The coefficients β_h are the impulse response of interest. The variable y is either micro skewness or aggregate sales growth. The shock series and controls are taken off-the-shelf from existing work. I now describe each shock in turn. Table 7 summarizes the regression specifications across the different shocks. Appendix C covers robustness checks and contains details on the variable definitions as well as data sources.

Monetary shock. I use the Bu et al. (2021) shocks, which are constructed to bridge periods

Table 7: Local projection specifications

Shock	Reference	Controls (lagged)	Sample period
Monetary	Bu et al. (2021)	Real GDP, GDP deflator,	
		Shadow Rate, EBP	1994Q1 – 2019Q4
Oil	Baumeister & Hamilton (2019)	Real GDP, GDP deflator, Oil price	1983Q3 – 2019Q4
Credit	Gilchrist & Zakrajšek (2012)	Real GDP, GDP deflator, EBP	1983Q3 – 2019Q4
Uncertainty	Ludvigson et al. (2021)	Real GDP, GDP deflator, VXO	1983Q1 – 2015Q4
Sentiment	Lagerborg et al. (2023)	ICE, real GDP, uncertainty,	
		Real stock prices	1983Q3 – 2019Q4
TFP	Ben Zeev & Khan (2015)	Real GDP per capita, real stock prices	
		per capita, labor productivity	1983Q3 – 2012Q1

All specifications include lags of the dependent variable and the shock series as controls and are estimated with two lags. 'ICE' is the University of Michigan Index of Consumer Expectations. Uncertainty is measured as the 12-month Jurado et al. (2015) uncertainty index.

of conventional and unconventional monetary policy. This is useful because the skewness series only starts in the mid-1980s while unconventional monetary policy became an important policy tool from 2008 onwards. Being restricted to a 1985-2008 sample period would make identification difficult, especially with quarterly data.[10] The shock is estimated with Fama-MacBeth regressions using changes in interest rates at different maturities around FOMC announcements such that the second-stage coefficient estimates are the monetary shock series. In my local projection specification, I include lags of real GDP and the GDP deflator (both as detrended log levels) as well as the Wu & Xia (2016) shadow rate and the excess bond premium as controls. The EBP captures financial conditions and is a useful control for the predictable component of the business cycle. I also include lags of the dependent variable and the shock as controls.

Oil supply shock. The oil supply shocks is identified following Baumeister & Hamilton (2019), who use carefully selected priors for demand and supply elasticities in the oil market (among priors for other coefficients) in a Bayesian VAR. Their identification scheme allows them to relax some identifying assumptions previously imposed in the literature, for example that the oil supply does

[10] Aggregating high-frequency shocks such as from Gertler & Karadi (2015) to the quarterly frequency has yielded insignificant effects of monetary policy on GDP. Using Romer & Romer (2004) shocks at the quarterly frequency also gives insignificant results for growth. This partly reflects that the strength of the effect of the Romer shocks depends on the inflation episode of the 1970s and early 1980s, see the discussion in Coibion (2012).

not respond on impact to shocks to the oil price. Under the new identification strategy, the authors find oil supply shocks to be a more important determinant of historical oil price movements than found in the previous literature. The shock series I use is the median of the posterior distribution. I add lags of the shock, GDP, the GDP deflator, the crude petroleum producer price index, and the dependent variable as controls.

Credit shock. The credit shock uses innovations in the excess bond premium following Gilchrist & Zakrajšek (2012). The excess bond premium is constructed from corporate bond spreads to proxy investor risk appetite and is orthogonal to the risk of corporate default. Gilchrist & Zakrajšek (2012) use a recursive identification strategy in a VAR to study the effect of EBP innovations on macroeconomic variables. They assume that indicators of economic activity do not respond to EBP shocks within the same quarter while financial variables can respond immediately. Following the equivalence result of Plagborg-Møller & Wolf (2021), I replicate this identification strategy within my local projections by controlling for contemporaneous and lagged values of real GDP and the GDP deflator while only controlling for lagged values of the 10-year US Treasury yield, the federal funds rate, and real stock prices.

Uncertainty shock. The identification of the uncertainty shock follows Ludvigson et al. (2021), who use restrictions on the time series of the structural shocks to jointly identify financial uncertainty, macroeconomic uncertainty, and output shocks. Given the VAR residuals, the authors randomly draw many candidates for the time series of the structural shocks and only retain those that satisfy restrictions motivated from economic theory and narratives of historical events. For example, financial uncertainty should be high in October 1987 ('Black Monday') and September 2008 (Lehman collapse).[11] The remaining shocks series can be used for set identification of the impulse responses. The authors find that financial uncertainty shocks are a source of business cycle fluctuations, while macroeconomic uncertainty is more likely to be an endogenous response to output shocks. To obtain a single shock series for the financial uncertainty shock, I use the 'maxG' solution, which jointly maximizes the inequalities associated with a subset of the constraints. The controls are lags of the shock, GDP, the GDP deflator, and the dependent variable.

Sentiment shock. While the previous shocks are related to economic fundamentals or financial conditions, business cycles may also be affected by fluctuations in consumer sentiment that are unrelated to economic conditions. Lagerborg et al. (2023) show that exogenous changes in consumer confidence can be recessionary. Their identification strategy relies on mass shootings in the United States, which are widely reported in the media and are shown to be predictors of downturns in sentiment. The authors show that the number of fatalities in mass shooting events can be viewed as exogenous to the state of the economy and used as a valid instrument to identify the effect of consumer confidence shocks on the business cycle. The authors estimate impulse responses in a proxy SVAR, and I extract the shock series from this system using the authors' replication codes. Similar to Lagerborg et al. (2023), I include lags of the shock, the University of Michigan Index of Consumer Expectations, real GDP, the Jurado et al. (2015) 12-month macroeconomic uncertainty index, real

[11] The idea behind the identification scheme is similar to the classic sign restrictions, except that the restrictions are directly imposed on the time series of the structural shocks as opposed to the shape or magnitude of the impulse response estimates.

stock prices, and the dependent variable in the local projections.

TFP news shock. News about future productivity can explain a significant share of business cycle variation, as shown in Beaudry & Portier (2006). I use shocks following the identification strategy of Ben Zeev & Khan (2015), who impose medium-run restrictions to identify news about investment-specific technology. Their shock is chosen to maximize the explained variance in (the inverse of) the relative price of investment in the medium term, while being orthogonal to both current TFP and the current relative price of investment. The authors find TFP news to account for a significant share of business cycle fluctuations. The impulse responses are estimated similar to the local projections of Ramey (2016), controlling for lags of the shock, real GDP per capita, real stock prices per capita, labor productivity, and the dependent variable.

Figure 6 shows the impulse response estimates for the six different shocks. All aggregate shocks

Figure 6: Comovement of growth and skew after aggregate shocks

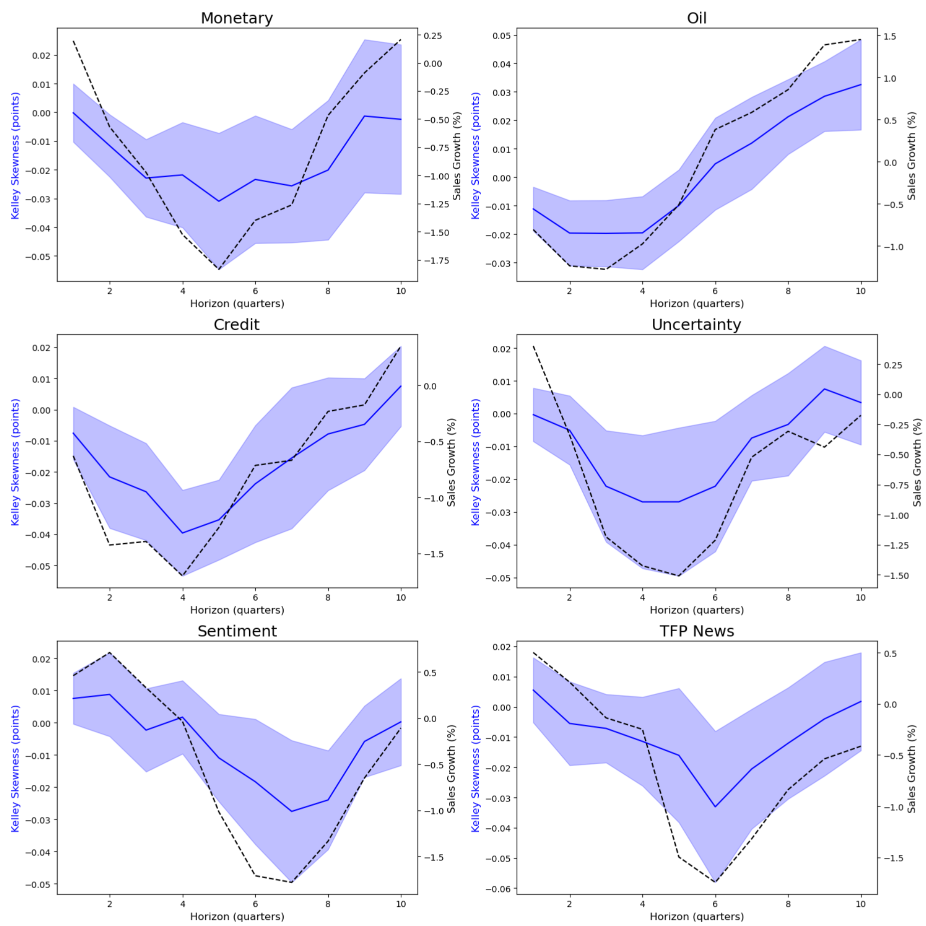

Note: The 90% confidence bands are based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and the TFP shock are reversed to be contractionary.

are associated with a subsequent decline in micro skewness (blue lines; left axis). Following a one standard deviation shock, the skewness index declines by between 0.02 and 0.06 points. The decline is strongest for the credit shock and weakest for the monetary shock. The peak effects occur 4 to 6 quarters after impact and is statistically significant across all shocks. The effects on skewness are not long-lived and die out after at most 10 quarters. The response of aggregate sales growth (black dashed lines; right axis) to the aggregate shocks looks very similar to the responses of skewness. The correlations of the impulse responses for a given shock range between 0.89 and 0.98. Aggregate shocks therefore appear capable of 1) inducing significant movements in micro skewness and 2) generating strong comovement between sales growth and skewness.

These findings confirm the priors one may have formed from considering Figure 1. Micro skewness moves closely with aggregate growth across many US recessions (including the Covid recession), suggesting the high correlation is a robust business cycle fact that does not only pertain to certain types of recessions. It is therefore encouraging to see that different types of shocks, all of which are considered potentially important drivers of the US business cycle, induce the procyclical skewness pattern.

4.3 Growth-skewness comovement due to idiosyncratic shocks

To compare the importance of aggregate shocks for procyclical skewness to the role of idiosyncratic shocks, I construct a new shock series reflecting the occurrence of size-weighted firm-level idiosyncratic shocks.[12] The shock series is constructed as follows:

1. Get the residuals $\varepsilon_{i,t}$ from the principal component regressions $g_{i,t} = \gamma_i + \lambda_i F_t + \varepsilon_{i,t}$ discussed above. The residuals are orthogonal to any aggregate factors.

2. Regress the residual on lagged year-over-year sales growth and quarter-over-quarter stock returns. This is to remove any easily predictable components from the residuals. Call the residuals from these regressions $\tilde{\varepsilon}_{i,t}$.

3. Run univariate regressions of the residuals $\tilde{\varepsilon}_{i,t}$ on I/B/E/S forecast errors $e_{i,t}$ and keep the fitted values from these regressions: $\hat{\varepsilon}_{i,t} = \hat{\alpha}_i + \hat{\beta}_i e_{i,t}$, where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are OLS coefficient estimates. The fitted values represent the part of the idiosyncratic residuals that was not forecasted by market analysts. I/B/E/S collects analyst forecasts for balance sheet items of publicly listed companies. For example, the data set may contain the forecast of a JP Morgan analyst for Apple's sales in the coming quarter.[13] I focus on quarterly sales forecasts made before the start of the quarter but at most 150 days in advance. The forecast error is the log difference between a given analyst forecast and the realized value. If there are multiple forecasts for a given company and quarter, I only keep the forecast with the smallest absolute error. This gives the most conservative estimate for the extent to which $\varepsilon_{i,t}$ was surprising to analysts. This step yields forecast errors and hence shocks $\hat{\varepsilon}_{i,t}$ for around 1000 firms per quarter. The skewness in the time series of forecast errors is closely correlated with the skewness in sales growth rates, as shown in Figure 24 in Appendix C.

[12] The idea that idiosyncratic shocks can have aggregate implications has been revived by Gabaix (2011). Gabaix & Koijen (Forthcoming) provide the methodology to aggregate idiosyncratic shocks to study their macro implications.

[13] Van Binsbergen et al. (2023) document that analyst forecasts contain important information that is not captured by publicly available information and useful for forecasting firm fundamentals.

4. Aggregate the individual shock series using real sales weights. Call the resulting shock series ξ_t.

5. Orthogonalize the shock series with respect to the macro shocks. To be conservative, I only consider the monetary shock, the oil supply shock, and the sentiment shock as aggregate shocks in this step. Given the nature and identification schemes of the shocks these three shocks can most clearly be considered pure aggregate shocks. In contrast, the credit shock may be driven by contributions of a small set of firms to the EBP, for example. The shock series ξ_t remains virtually unchanged in this step since it is barely correlated with any of the other shock series.

6. To ensure that the shock series cannot be forecasted using standard macroeconomic data, I regress ξ_t on eight factors from the FRED-QD data set by McCracken & Ng (2020). This data set contains over 200 economic time series providing a comprehensive picture of macroeconomic conditions. The residuals from this regression are the idiosyncratic shock series to be used in the local projections.

Figure 7: Results for idiosyncratic shock series

The left panel of Figure 7 shows the shock series. The shock series starts in 1994 due to the shorter sample for the I/B/E/S forecast errors. By construction, the shock series cannot be forecasted well by a broad set of macroeconomic variables. Table 22 in Appendix C reports the results from a regression of the shock series onto 11 macroeconomic variables used in the local projections and reports an adjusted R^2 of only 11%. The right panel of Figure 7 shows the impulse response of micro skewness (solid blue) and aggregate sales growth (dashed black) to a one standard deviation contractionary idiosyncratic shock. Skewness declines on impact in line with the decline in aggregate sales growth before recovering quickly after around six quarters. The order of magnitude of the impulse response estimates is in line with the responses to one standard deviation aggregate shocks.[14]

[14] An important step in constructing the shock series is removing the component that is predictable by FRED factors. If the predictable component is not removed, the shock series shows strong local minima around the burst of

The idiosyncratic shock series does not capture the effects of a particular source of idiosyncratic fluctuations but is instead a catch-all for any type of idiosyncratic shock such as product releases, strikes, natural disasters, or changes to the executive board. This allows to obtain a relevant shock series in the statistical sense. In contrast, many idiosyncratic shock series may be well identified but are too small to cause significant fluctuations at the macro level. I consider my approach conservative in that it starts the shock construction from firm-level sales growth rates and only removes those sources of variation that are either attributable to aggregate sources or forecastable. Given the new shock series, we can assess its contribution to the fluctuations in skewness relative to the contribution of aggregate shocks.

4.4 A variance decomposition of skewness fluctuations

What is the contribution of aggregate versus idiosyncratic shocks to the fluctuations in micro skewness? I estimate a variance decomposition for the skewness index using the approach of Gorodnichenko & Lee (2020). Their methodology proceeds in two steps. First, regress the target variable on all controls but not the current shock:

$$y_{t+h} = \alpha_h + \sum_{\ell=1}^{L} \gamma'_{\ell,t} \text{controls}_{t-\ell} + e_{t+h} \tag{6}$$

Then, for horizon h, the explained variance is the R^2 from the regression of the first-step residuals on the vector of shocks between t and $t+h$:

$$\hat{u}_{t+h} = \sum_{k=0}^{h} b'_{k,h} \text{shocks}_{t+k} + v_{t+h} \tag{7}$$

The result of this approach is the variance of the forecast error explained by the set of shocks considered. I compare the contributions of the monetary, oil supply, and sentiment shock against the contribution of the idiosyncratic shock series. I focus on these three aggregate shocks because they are most clearly interpretable and identified as aggregate shocks as opposed to possibly having granular origins. The first-stage regression controls for two lags of the shocks and of GDP growth, the EBP, and the skewness index. As recommended in Gorodnichenko & Lee (2020), I use a VAR-based bootstrap to address small-sample bias in the R^2 estimates.

Because the four shock series are not orthogonal, the variance explained by the shocks does not equal the sum of the variances explained by each individual shock. To achieve a decomposition, I approximate each shock z's partial contribution by subtracting the R^2 from equation 7 using the three other shocks from the explained variance using all four shocks: $\text{partial}_z = R^2_Z - R^2_{Z\setminus z}$, where $Z := \{\text{monetary}, \text{oil}, \text{sentiment}, \text{idiosyncratic}\}$ is the set of all shocks. Shock z's contribution to the explained variance is then $\text{contribution}_z = \omega_z R^2_Z$ with weights $\omega_z = \text{partial}_z / \left(\sum_{i\in Z} \text{partial}_i\right)$.[15]

Figure 8 plots the variance decomposition up to horizon ten of the skewness index. Within the

the dot-com bubble and during the Great Financial Crisis. Accordingly, the IRF estimates indicate a decline in the skewness index and aggregate sales growth that are twice as large as the baseline results. These results are available upon request.

[15] This is a one-step approximation to the type of relative importance calculation described in Grömping (2007).

Figure 8: Variance decomposition of micro skewness

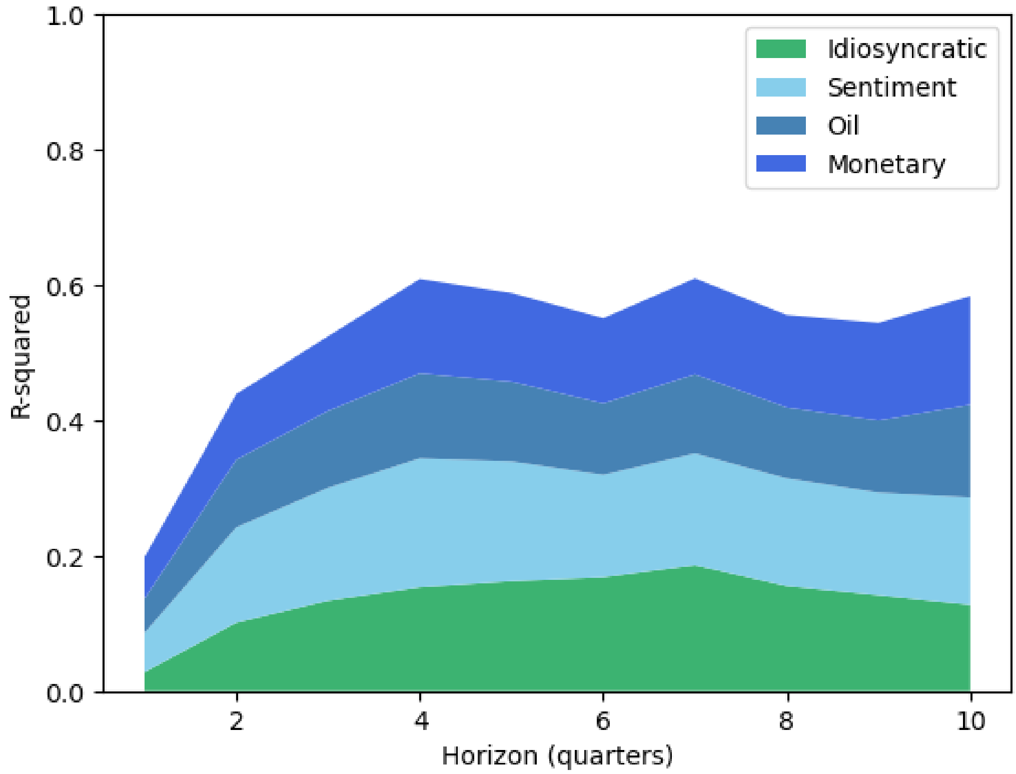

regression, up to 60% of the variation in forecast errors is explained by the four shocks. Explained variance is lower at short horizons but stable beyond the fourth quarter. Of the variance explained, the aggregate shocks account for roughly 75% compared to 25% for the idiosyncratic shock. The aggregate shocks themselves account for roughly 25% of variation each, with some variation over time. These results are in line with the PCA decomposition result from Section 4.1: Most of the variation in the skewness index is due to aggregate shocks, though a non-negligible share of the variation is accounted for by idiosyncratic shocks. These results confirm the importance of idiosyncratic shocks in aggregate fluctuations (Gabaix (2011), Carvalho & Grassi (2019)) while supporting this paper's focus on aggregate shocks.

4.5 Skewness across firm-level impulse responses

So far, this paper has used measures of aggregate growth and micro skewness as inputs to the local projections to study their impulse responses. Instead, this section estimates impulse responses of firm-level sales growth rates to aggregate shocks and then construct the response of micro skewness and aggregate sales growth bottom-up from the distribution of firm-level IRFs. All firm-level regressions control for lagged values of the shock and lagged GDP, as well as sales growth at the firm and the 2-digit NAICS level. In addition, I include shock-specific controls: shadow rate and leverage (monetary shock), GDP deflator (oil supply), excess bond premium (credit shock), Jurado et al. (2015) financial uncertainty (uncertainty), ICE consumer sentiment, macroeconomic uncertainty, and S&P500 stock prices (sentiment), and GDP per capita, labor productivity, and S&P500 stock prices per capita (TFP news). All controls are included with two lags. The only exception is a contemporaneous control for GDP growth in the credit shock regression, mirroring the specification in Gilchrist & Zakrajšek (2012).

The summary statistics for the distribution of firm-level impulse response estimates are reported

Table 8: Firm-level local projections - Summary statistics

	Monetary	Oil	Credit	Uncertainty	Sentiment	TFP News
# Streaks	4,120	5,332	2,813	5,115	5,296	4,829
# Firms	4,017	5,061	2,813	4,893	5,030	4,651
Avg. # Obs.	64	70	84	66	69	63
Avg. R^2	0.33	0.23	0.25	0.26	0.32	0.32
Sign. IRFs (%)	77	76	83	78	81	79
$Q_{0.1}^{\text{IRF}}$ (%)	-5.4	-13.9	-2.1	-5.0	-4.5	-5.2
$Q_{0.5}^{\text{IRF}}$ (%)	-0.16	-0.02	-0.33	-0.3	-0.29	-0.35
$Q_{0.9}^{\text{IRF}}$ (%)	5.0	13.1	1.2	4.0	4.1	4.7

The number of streaks can be larger than the number of unique firms. The average number of time series observations is measured for impact effect regressions and rounded to the nearest integer. The adjusted R-squared values are averaged across horizons and firms. The share of significant IRFs is the relative frequency of statistically significant IRFs for the peak of the impulse response estimates, measured using 90% confidence intervals based on Newey-West standard errors. Quantiles across firm-level IRFs are averaged across horizons. IRFs for the credit shock are only estimated for firms existing during the Great Financial Crisis.

in Table 8. Varying sample periods across the shocks and missing values for firm-specific controls (in particular leverage for the monetary shock) imply differences in sample sizes. The number of unique streaks is above 4,000 for all shocks except credit. The sample for the credit shock is smaller since I only consider streaks covering the Great Financial Crisis, which turns out to be crucial to identify the effects of credit shocks using the Gilchrist & Zakrajšek (2012) specification. The number of streaks can be larger than the number of unique firms in the sample since some firms can have multiple streaks in the data, although this does not happen frequently. The average time series is roughly 70 quarters long. The firm-level regression have average R^2 values of at least 23% and over 75% of impulse response estimates have statistically significant peak effects for each shock. The distribution of IRF estimates is widely dispersed with negative (unweighted) mean estimates, reflecting the contractionary nature of the shocks but large heterogeneity in terms of firm responses. These distributions look very similar when only considering IRFs with significant peak effects (results not shown).

To build some intuition for the impulse response estimates, Figures 25 - 30 in Appendix C show examples of firm-level IRFs for ExxonMobil, McDonald's, Marriott Hotels, Caterpillar, IBM, and Walt Disney. The impulse responses vary across firms in line with differences in sector exposure and cyclicality. For example, Caterpillar is a very cyclical company with strong impulse responses to aggregate shocks, while Walt Disney's sales are unresponsive to aggregate shocks. ExxonMobil responds strongly to oil supply shocks, while IBM is adversely affected by TFP shocks and Marriott by uncertainty shocks. As an additional sensitivity check, Figure 31 uses the oil supply shock to compare the impulse responses of four big oil corporations against non-oil firms. Oil corporations have considerably stronger responses to the oil shock. Taken together, the firm-level impulse responses appear to provide a reasonable picture about firm's responsiveness to aggregate shocks.

Based on the firm-level IRF estimates, I construct the response of skewness and aggregate sales

Figure 9: Comovement of bottom-up growth and skew after aggregate shocks

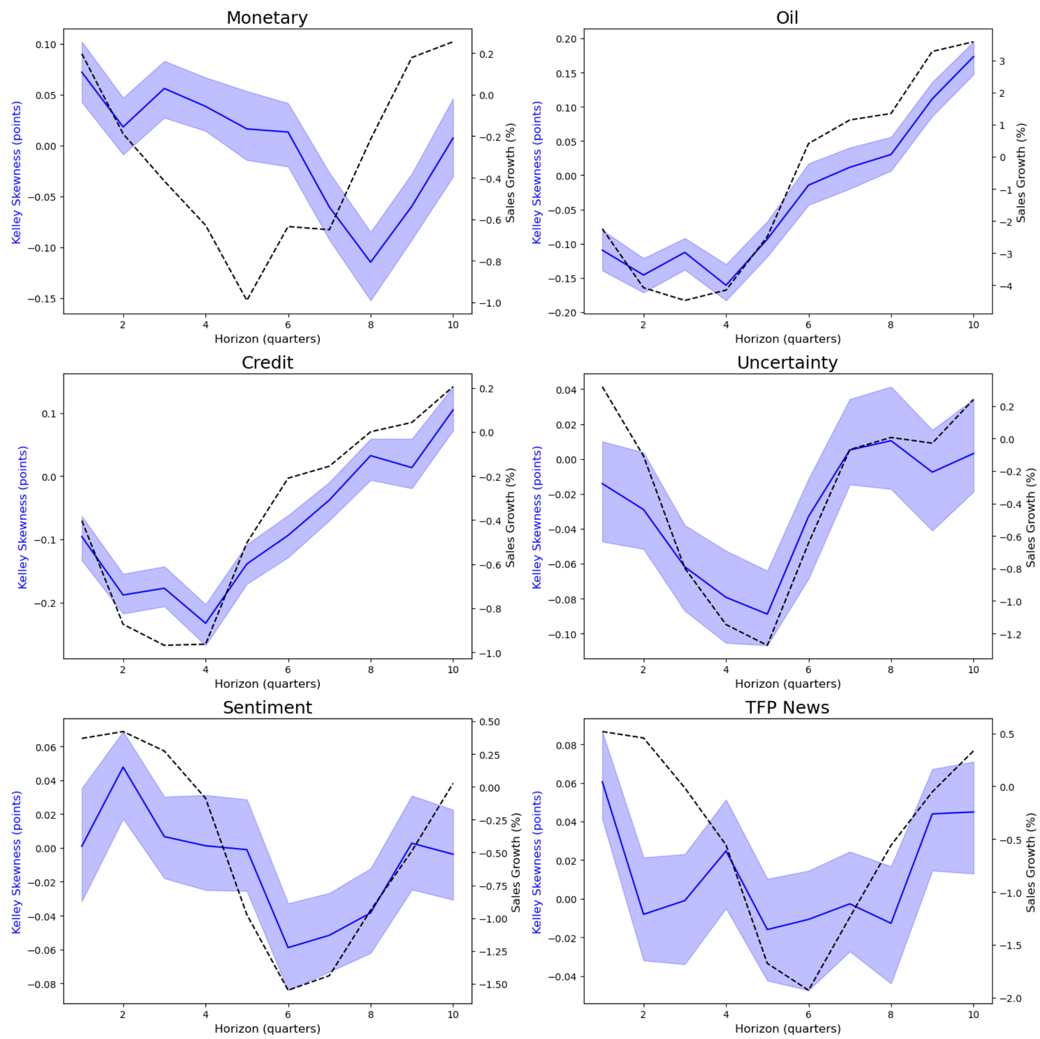

growth from the bottom up. The aggregate sales growth IRF is the size-weighted average of the firm IRFs, while the response of micro skewness is estimated from the cross section of firm IRFs. Testing for procyclical skewness in this exercise is significantly harder since individual firm IRFs are much more volatile than aggregate sales and the only source of procyclical skewness in response to a properly identified aggregate shock are heterogeneous responses across firms. The results are in Figure 9, where shaded areas are 90% confidence intervals based on a simple bootstrap with 2000 replications. Following a contractionary aggregate shock, micro skewness (solid blue) and aggregate sales growth (dashed black) show a closely correlated decline. This is especially true for the oil, credit and uncertainty shocks. The correlations of skewness and growth following a sentiment or TFP shock are also close but the evidence for a negative skewness response is less clear. The monetary shock leads to a severe contraction in skewness but only after eight quarters, with a positive resonse on impact. Sales growth declines earlier and is recovering while skewness bottoms out. Without putting too much weight on any individual impulse response estimate, the sum of findings across the six different shocks confirms that 1) micro skewness declines following contractionary aggregate shocks and 2) aggregate sales growth and micro skewness are strongly correlated following aggregate shocks. Figure 32 in Appendix C confirms that the results are robust to including more lags of the controls and adding lagged stock returns as controls. Using year-over-year stock returns instead of sales growth gives qualitatively similar results.

How does the response of large firms differ from the response of small firms? I split the sample

Figure 10: Large vs small firms: Bottom-up skewness and growth responses

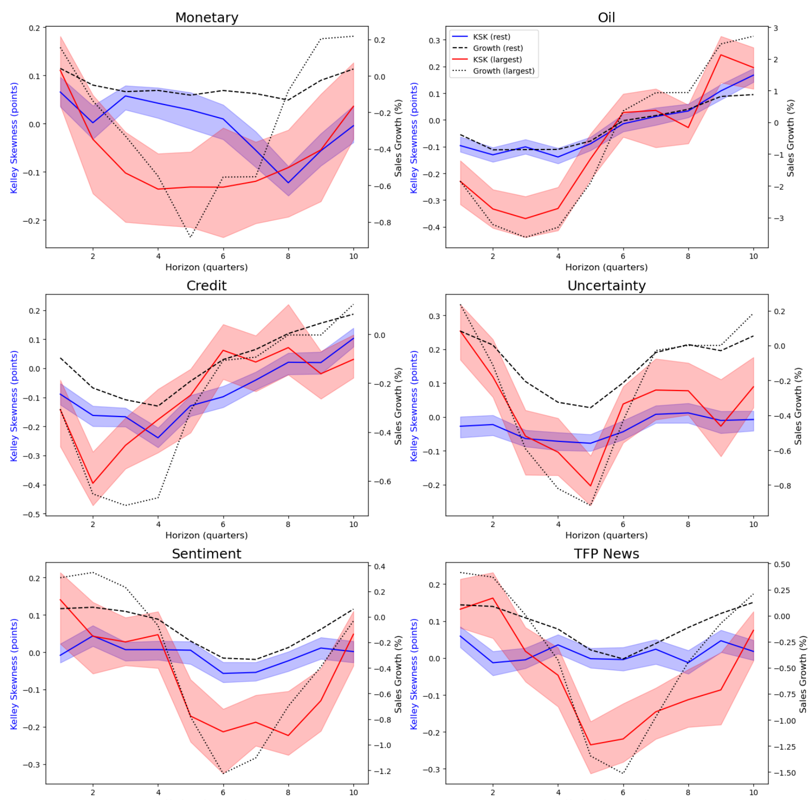

Note: The 90% confidence bands are based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and the TFP shock are reversed to be contractionary.

into two size groups (largest firms versus the rest) to study the impulse response of skewness across large vs small firms and compute their contribution to aggregate sales growth. Figure 10 shows the IRFs for the largest 10% of firms (defined by average real sales) and the IRFs for the bottom 90% of firms. By construction, the sum of the two lines equals the impulse response of aggregate sales growth shown in Figure 9. The black dotted (dashed) line shows large (small) firms' contribution to the impulse response of aggregate sales growth. The red (blue) line shows the impulse response of skewness across large (small) firms. The shaded areas are 90% confidence intervals.

The bottom-up skewness response of large firms is significantly negative across shocks and in line with the impulse responses for the skewness index (Figure 6). The response of the largest firms is more skewed than the response of the rest of the firms. The differences in skewness can be large. For example, the minimum of the skewness IRF in response to a one standard deviation sentiment shock is around -0.2 for the largest firms but only -0.04 for the smaller firms. In response to an oil shock, large firms' skewness declines by over 0.3 points, while smaller firms' skewness falls by 0.1 points at most. The differences are also large for the monetary and the TFP shock and less pronounced for the credit and the uncertainty shock. In any case, the response across large firms is *not less* skewed than the response of small firms.[16]

Since the firm size distribution is fat-tailed, the size-weighted sales response of large firms is significantly more contractionary than for small firms. The decline in aggregate sales growth following an aggregate shock is almost entirely due to the largest firms, confirming the findings of Crouzet & Mehrotra (2020) on the role of the largest firms for aggregate fluctuations.[17] Since the aggregate growth rate responses in this figure are weighted by firm size, my findings do not suggest that small firms are less responsive to shocks. Instead, their responses barely affect aggregate fluctuations since they receive small weights. To the extent that the response of small firms is less skewed than the response of large firms, small firms may also contribute less to aggregate fluctuations because there is a non-negligible share of small firms with positive impulse responses to negative shocks. The impulse responses of small firms may be more volatile than those of large firms, but the contribution of small firms to aggregate fluctuations can be weak if this volatility averages out across firms.

Figure 11 confirms this intuition by plotting the contribution of different size and growth rate bins to the aggregate growth response following an aggregate shock. For each shock, I consider the trough of the aggregate growth rate response and group the IRF estimates at this point into five quintile bins: from poorest performers (bin 1) to the best performers (bin 5). Weighting the impulse responses by the firms' sales weights allows to compute each bin's contribution to the year-over-year aggregate sales decline observed at the trough. I re-scale these contributions to sum to -100%. This yields a bar plot for each shock, which is shown in Figure 36 in Appendix C. Averaging the contribution of each bar across all six shocks yields Figure 11. Within each IRF bin, the figure also indicates the part of the contribution that is due to the largest firms (top 10% of the size distribution; blue bars) and the part due to the rest of firms (orange bars).

Following an aggregate shock, the aggregate sales growth decline is mostly driven by those firms with the poorest growth rates (lowest IRF values). The top 60% of the IRF distribution together

[16] One may be concerned about the reliability of IRF estimates for smaller firms since these firms are more volatile. Figure 35 in Appendix C shows a robustness check addressing these concerns. I separately construct a micro skewness index for the largest firms and the rest of firms and estimate the response of these two skewness indexes to aggregate shocks. The skewness index of the largest firms declines substantially, while the skewness index for the rest of firms shows a much more muted and often insignificant response.

[17] These findings are robust to comparing the top 10% of firms against the middle 80% or to comparing the top 30% of firms against the bottom 70%, although the differences across size groups are most pronounced in the baseline results (see Figures 33 and 34 in Appendix C).

Figure 11: Contributions of growth and size bins to aggregate growth decline

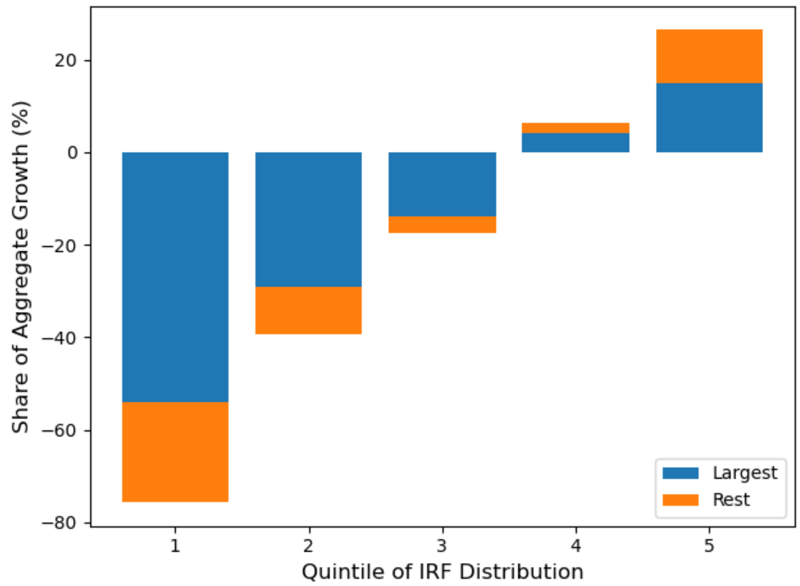

Note: Largest firms are the top 10% of the size distribution, which averages real sales over time for each firm. The contributions are re-scaled such that the bars add up to -100% and averaged across the six shocks.

have a weakly positive contribution to aggregate growth, while the bottom 20% account for roughly half of the aggregate growth response. Within each bin, the largest firms account for the majority of the response. This is true even among the worst performers: Over 70% of the contribution of the poor performers is due to firms that are both large and responsive. These large and responsive firms are important for the transmission of aggregate shocks to the economy even though they are rare. Only 1.6% of firms in the sample are considered to be both large and responsive, but they account for one third of the economy's response to aggregate shocks. Overall, the largest firms account for 78% of the aggregate sales decline, reflecting the fact that they are large and that their responses are more skewed than for the rest of firms.[18]

These findings have important implications for the interpretation of the skewness-growth comovement. If only skewness across small firms' IRFs showed a signficant response to aggregate shocks, procyclical skewness could be interpreted as a byproduct of the business cycle: Aggregate shocks drive aggregate fluctuations because they have a negative (but unskewed) impact on large firms on average. They happen to have a negatively skewed impact on small firms, which does not matter (much) for aggregate fluctuations. However, the results show that the response of large firms is skewed. This provides a new narrative for the transmission of aggregate shocks to the economy: A sizeable share of the effect of an aggregate shock can be traced to a small number of very large firms with strong responses to the shock. This is a complementary business cycle explanation to Gabaix (2011)'s granular hypothesis. The granular hypothesis argues that idiosyncratic shocks to large firms can explain a significant share of business cycle fluctuations. The results of this paper

[18]These results also hold for each shock individually, see Figure 36 in Appendix C.

suggest a *granular response to aggregate shocks*: Aggregate shocks explain a large part of aggregate fluctuations but a large part of their effect comes from the response of a handful of very large firms.

4.6 The origins of heterogeneous responses

Why are some large firms so vulnerable in response to aggregate shocks? More generally, what explains the heterogeneous responses of firms to aggregate shocks? This section presents some tentative results on the most promising firm characteristics that explain my findings. Explaining the heterogeneity across impulse responses turns out to be challenging and there does not appear to be a single firm characteristic accounting for most of the variation across firms. However, I present some evidence that those large firms with weak ex-ante profitability may be more likely to suffer when an aggregate shock hits.

To make progress, I identify the firm characteristics that best predict heterogeneity in the

Table 9: Predictors of heterogeneous responses among the largest firms

	Monetary	Oil	Credit	Uncertainty	Sentiment	TFP News
Age	0.01	0.62	-0.11	-0.04	-0.10	0.03
	(0.22)	(0.62)	(0.10)	(0.16)	(0.15)	(0.18)
Size	-0.69	0.15	0.41	0.44	0.48	0.18
	(0.54)	(1.35)	(0.25)	(0.42)	(0.38)	(0.41)
Leverage	4.69	1.11	0.66	2.55	2.65	-8.03
	(6.12)	(13.97)	(1.81)	(6.81)	(4.36)	(8.74)
Liquidity	1.03	0.97	-0.14	0.22	0.04	1.15
	(0.48)	(1.23)	(0.20)	(0.46)	(0.53)	(0.45)
Dividend payer	0.03	0.98	0.05	0.09	0.15	0.16
	(0.17)	(0.45)	(0.05)	(0.13)	(0.11)	(0.14)
Fixed assets	-0.37	-1.09	-0.39	-0.34	-0.38	0.03
	(0.24)	(0.61)	(0.11)	(0.22)	(0.19)	(0.28)
Short-term debt	-1.94	1.15	-0.30	-0.81	-1.03	3.55
	(2.37)	(6.62)	(0.74)	(2.83)	(1.94)	(3.45)
Long-term debt	-3.62	-1.97	-0.44	-2.26	-4.60	8.02
	(5.35)	(12.14)	(1.62)	(8.08)	(9.00)	(8.52)
Sales / Assets	0.19	1.49	0.17	0.47	0.17	0.20
	(0.31)	(0.76)	(0.17)	(0.24)	(0.31)	(0.26)
Profitability	2.75	4.46	0.66	1.41	1.78	1.94
	(0.99)	(2.78)	(0.46)	(1.06)	(0.71)	(0.76)
R&D	-0.26	3.20	0.00	0.37	0.21	-0.10
	(0.83)	(1.46)	(0.40)	(0.69)	(0.64)	(0.62)
Inventory	0.24	-0.18	-0.18	0.32	0.11	0.25
	(0.38)	(0.90)	(0.20)	(0.31)	(0.38)	(0.37)
Observations	412	534	282	512	513	483
adj. R^2	0.42	0.32	0.11	0.11	0.14	0.17

Dependent variable: Minimum of IRF across horizons, in percent. Firm characteristics are averaged over time for each firm, then standardized across firms. Dividend payer is a dummy. All regressions include a constant and controls for the number of recession quarters, aggregate growth, industry growth (2-digit NAICS), pre-recession firm growth, and average firm growth. Standard errors are heteroskedasticity robust.

minimum of the impulse response (across IRF horizons). The predictors are a set of common firm characteristics, where the time series of each characteristic is averaged for a given firm and then standardized across firms.[19] For each shock, this yields a cross section of IRF minima and firm characteristics. Table 9 shows the results. Within the set of large firms, characteristics such as age and firm size are not significant predictors of cross-firm differences in IRFs. The only successful predictor turns out to be profitability (measured as return on assets), which is significant for most shocks. For example, a one standard deviation increase in ROA raises the minimum of the impulse response to a monetary shock by 2.75 percentage points. More profitable firms therefore respond less negatively to contractionary shocks.[20]

In general, the broad set of firm characteristics explains only a small share of the variation in impulse response estimates. To test if the poor fit is due to nonlinearities, I estimate a random forest (Breiman (2001)) on the data. The random forest is a supervised machine learning algorithm that naturally handles state dependence in the form of threshold effects. It repeatedly searches for a value of a characteristic that allows to split the data in a way to make optimal constant predictions within the created partitions ('leafs'). For example, the algorithm may decide that the best way to separate firms' IRFs into two groups is to separate large from small firms. The random forest therefore allows for state dependence in a naturally interpretable way that closely maps into the common understanding of state dependence in macroeconomics. In contrast to other nonlinear machine learning models such as deep neural networks, the random forest has been shown to perform well for tabular data with a moderate number of observations and is less sensitive to hyperparameter tuning (Goulet Coulombe et al. (2022)).

The random forest consists of a collection of regression trees. Let $x_{i,j}$ be a characteristic j of firm i and denote its IRF value as y_i. The number of *regions* is N_m. Each tree splits the sample of firm characteristics x into non-overlapping regions R_m such that all firms within the same region obtain the same prediction:

$$\hat{y}_i = f(x_i) = \sum_m c_m \mathbb{I}_{\{x_i \in R_m\}}, \tag{8}$$

with constants

$$c_m = \frac{1}{N_m} \sum_{y_i : x_i \in R_m} y_i \tag{9}$$

and regions defined as rectangular hyperregions in the predictor space

$$R_m = \{x_i : k^m_{j,l} < x_{i,j} \leq k^m_{j,h} \forall x_{i,j} \in x_i\}. \tag{10}$$

The goal of the algorithm is to find the characteristic-specific borders $k^m_{j,l}$ and $k^m_{j,h}$ that yield the best prediction.

Without further restrictions, this tree is able to achieve zero forecast error by splitting the sample into as many leafs as observations. This would clearly overfit the data. The random forest handles

[19] The construction of the firm characteristics is described in Appendix C.

[20] Alternative ways of cutting the data confirm these findings. Among large firms, the probability of firms with high ROA (top 10% of ROA distribution) to end up in the bottom quintile of the growth rate distribution at the trough of a recession is 12% compared to 24% for the firms with the lowest ROA (bottom 10%). Similarly, (small and large) firms with strong responses to aggregate shocks (bottom half of IRF distribution) have significantly lower ROA than those firms with IRFs in the top half of the distribution. This findings is robust across shocks but weakens when only considering the very largest firms.

this problem in two ways. First, certain hyperparameters, such as the number of maximum number of leafs in a tree or the maximum depth of a tree, allow to avoid overfitting of a given tree and improve out-of-sample fit. Second, the random forest predictions are averages over an ensemble of regression trees. Regressions tree predictions differ from each other because each tree is trained on a slightly different (bootstrapped) sample and because the firm characteristics used to split the sample at a given node are a random subsample of the full set of characteristics. The number of trees and the share of characteristics to be considered for each split are additional hyperparameters to be chosen.

The model is trained using a grid search over different hyperparameter choices. The number of trees (`n_estimators`) is chosen from $\{50, 100, 150, 200, 300\}$, the share of characteristics to consider for a given split (`max_features`) is $\{0.25, 0.3, 0.4, 0.5\}$, the maximum depth of a single tree (`max_depth`) can be $\{5, 10, 20, 30\}$, and the maximum number of leaf nodes (`max_leaf_nodes`) is $\{3, 6, 9, 12\}$. These hyperparameter ranges are chosen to allow for a good bias-variance tradeoff. To train the model, I randomly sample 75% of the full sample as a training sample. The remaining 25% of the data are used as a test sample to evaluate the fit of the model on unseen data. To increase the number of observations, I focus on all impulse responses as opposed to only those of the largest firms. After training, the R^2 values on the test sample are ranging from 14% to 26%, roughly matching the *in-sample* performance of OLS regressions.

In contrast to linear regression settings, random forests are nonparametric and do not provide easily interpretable regression coefficients that can be used to assess variable contributions to the prediction. Shapley values (Shapley (1953)) allow to distribute the prediction success to the different predictors.[21] For a given observation y_i, the Shapley value measures the average marginal contribution of a predictor $x_{i,j}$ across all possible coalitions S of predictors:

$$\phi_j^{(i)}(v) = \sum_{S \subseteq \mathcal{X}_i \setminus \{x_{i,j}\}} \frac{|S|!\,(|\mathcal{X}_i| - |S| - 1)\,!}{|\mathcal{X}_i|!} \left(v\left(S \cup \{x_{i,j}\}\right) - v(S) \right), \tag{11}$$

where $\mathcal{X}_i = \{x_{i,1}, ..., x_{i,p}\}$ is the set of firm characteristics of firm i, p indexes the last predictor, and the value function

$$v(S) = \int f\left(x_{i,1}, ..., x_{i,p}\right) d\mathbb{P}_{\mathcal{X}_i \notin S} - \mathbb{E}_X\left(f(X)\right) \tag{12}$$

computes the difference between the prediction of coalition S and the average predicted value. The average absolute contribution of a firm characteristic j to explaining heterogeneity in IRFs is then

$$\phi_j = \frac{1}{N} \sum_{i=1}^{N} |\phi_j^{(i)}|. \tag{13}$$

Larger values indicate a larger contribution to the prediction. Unlike regression coefficients, the Shapley values do not indicate the sign of the contribution. An exact evaluation of the Shapley value is computationally intense since the number of coalitions grows exponentially in the number of predictors. Lundberg & Lee (2017) provide an efficient implementation that reduces the computational burden significantly. I follow their approach.

Figure 12 shows the characteristics with the largest Shapley values for the different shocks. Size

[21] Molnar (2020) provides an introduction to the concept for the interpretation of machine learning models.

Figure 12: Shapley values of different firm characteristics

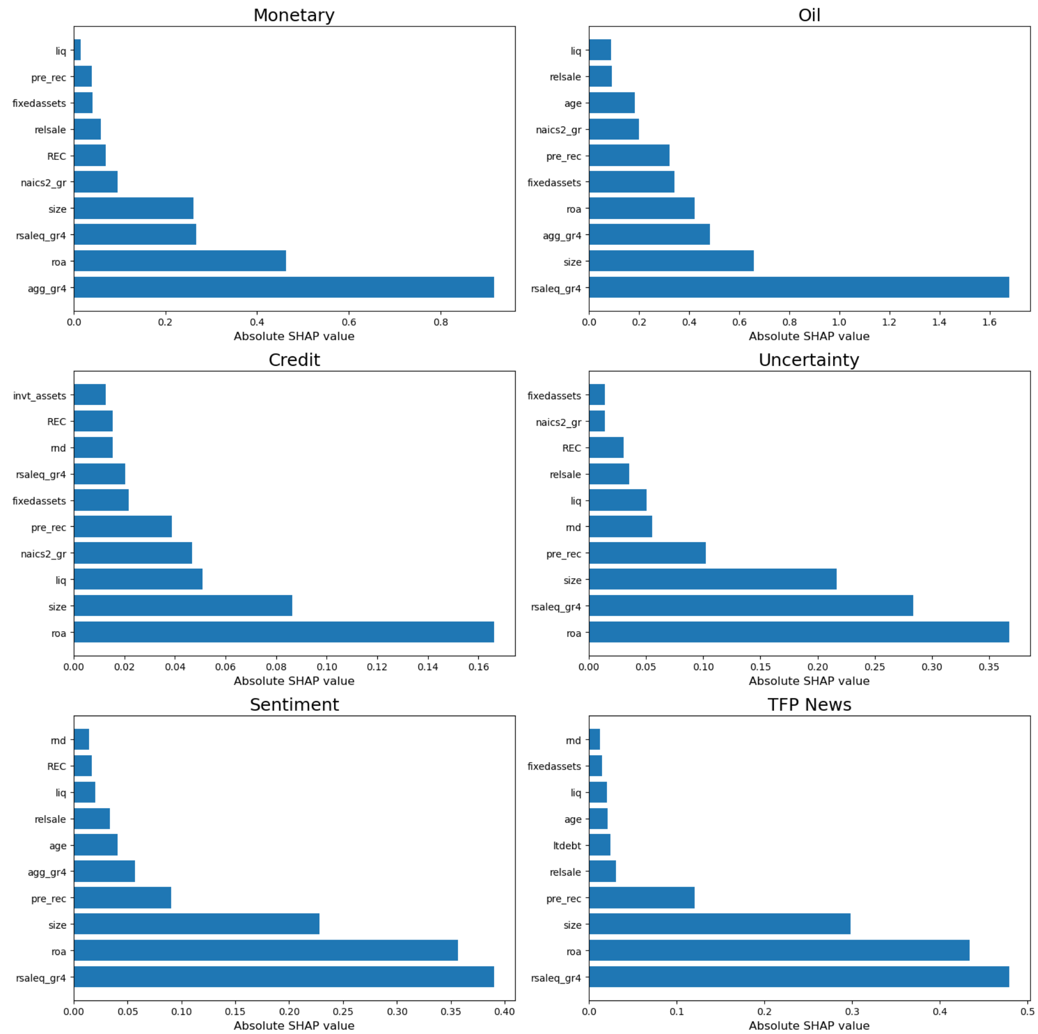

Note: The variable names are as follows. liq: liquidity. pre_rec: pre-recession sales growth. fixedassets: fixed assets. relsale: sales over assets. REC: recession dummy. naics2_gr: industry sales growth. size: real sales. rsaleq_gr4: firm sales growth. roa: return on assets (profitability). agg_gr4: Aggregate sales growth. age: (log) firm age. invt_assets: inventories over assets. rnd: R&D expenses over assets. ltdebt: long-term debt over assets. See Appendix C for a full list of variable definitions.

is generally among the most important predictors, reflecting that large firms have weaker responses on average. Apart from average sales growth rates, which are mostly included as a control, the only other predictor that repeatedly features among the top predictors is the return on assets, a common measure of firm profitability. This confirms the findings from Table 9 in a state-of-the-art machine learning regression allowing for flexible forms of nonlinearity.

How should we interpret the role of profitability in moderating the responsiveness of large firms to aggregate shocks? Higher profitability could reflect structural factors such as higher pricing power. Firms with high markups may have more space to pass on increases in input costs or could engage in predatory pricing behaviour to gain market share during downturns. High markups may also be a reflection of operating in a market niche with a stable customer base due to specific needs or preferences. This may reduce the volatility of such firms in general. In contrast, firms with low or negative profitability may be vulnerable to aggregate shocks because they have limited balance sheet capacity to deal with financial blows, or low pricing power may mean that these firms are forced to follow market trends of declining prices in an environment of declining demand. Future work could explore the relation between profitability (or markups specifically) and firms' responsiveness to shocks in more detail, and may also want to identify further characteristics that can account for a larger share of the heterogeneity in responsiveness across firms.

5 Conclusion

This paper studies the origins and implications of procyclical micro skewness. There are two key findings. First, micro skewness matters for aggregate fluctuations because some large firms experience very poor growth rates in recessions. This means that most of the decline in sales levels in a recession is accounted for by those firms with the poorest growth rate outcomes. Second, aggregate shocks induce a close comovement between aggregate sales growth and micro skewness because firms respond heterogeneously to these shocks. Importantly, the response of the very largest public firms in the US economy is also skewed such that some large firms respond strongly to aggregate shocks. Since the firm size distribution is fat-tailed, the aggregate effects of aggregate shocks are largely explained by a small number of very large firms with strong responses. I confirm this finding for six different types of aggregate shocks, identified with six different identification schemes.

The findings of this paper suggest that the sources of vulnerability of large firms provide a fruitful area for future work. A natural next step is to identify the firm characteristics that best predict the responsiveness of large firms. This paper has taken a small step in that direction and proposed profitability as a potential candidate for future study. In ongoing work, I am taking a more detailed look at the mechanisms through which profitability may affect firm vulnerabilities.

References

Adrian, T., Boyarchenko, N. & Giannone, D. (2019), 'Vulnerable growth', *American Economic Review* **109**(4), 1263–89.

Axtell, R. L. (2001), 'Zipf distribution of us firm sizes', *Science* **293**(5536), 1818–1820.

Bachmann, R. & Bayer, C. (2014), 'Investment dispersion and the business cycle', *American Economic Review* **104**(4), 1392–1416.

Baqaee, D. R. & Farhi, E. (2019), 'The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem', *Econometrica* **87**(4), 1155–1203.

Baumeister, C. & Hamilton, J. D. (2019), 'Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks', *American Economic Review* **109**(5), 1873–1910.

Beaudry, P. & Portier, F. (2006), 'Stock prices, news, and economic fluctuations', *American economic review* **96**(4), 1293–1307.

Ben Zeev, N. & Khan, H. (2015), 'Investment-specific news shocks and us business cycles', *Journal of Money, Credit and Banking* **47**(7), 1443–1464.

Berger, D., Dew-Becker, I. & Giglio, S. (2020), 'Uncertainty shocks as second-moment news shocks', *The Review of Economic Studies* **87**(1), 40–76.

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.

Bloom, N. (2014), 'Fluctuations in uncertainty', *Journal of Economic Perspectives* **28**(2), 153–176.

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

Breiman, L. (2001), 'Random forests', *Machine learning* **45**, 5–32.

Bu, C., Rogers, J. & Wu, W. (2021), 'A unified measure of fed monetary policy shocks', *Journal of Monetary Economics* **118**, 331–349.

Busch, C., Domeij, D., Guvenen, F. & Madera, R. (2022), 'Skewed idiosyncratic income risk over the business cycle: Sources and insurance', *American Economic Journal: Macroeconomics* **14**(2), 207–42.

Carvalho, V. M. & Grassi, B. (2019), 'Large firm dynamics and the business cycle', *American Economic Review* **109**(4), 1375–1425.

Cascaldi-Garcia, D., Sarisoy, C., Londono, J. M., Sun, B., Datta, D. D., Ferreira, T., Grishchenko, O., Jahan-Parvar, M. R., Loria, F., Ma, S. et al. (2023), 'What is certain about uncertainty?', *Journal of Economic Literature* **61**(2), 624–654.

Catherine, S., Sodini, P. & Zhang, Y. (2022), 'Countercyclical income risk and portfolio choices: Evidence from sweden', *Swedish House of Finance Research Paper No. 20-20* .

Cloyne, J., Ferreira, C., Froemel, M. & Surico, P. (2023), 'Monetary policy, corporate finance and investment', *Journal of the European Economic Association* .

Coibion, O. (2012), 'Are the effects of monetary policy shocks big or small?', *American Economic Journal: Macroeconomics* **4**(2), 1–32.

Crouzet, N. & Mehrotra, N. R. (2020), 'Small and large firms over the business cycle', *American Economic Review* **110**(11), 3549–3601.

Decker, R. A., Haltiwanger, J., Jarmin, R. S. & Miranda, J. (2016), 'Where has all the skewness gone? the decline in high-growth (young) firms in the us', *European Economic Review* **86**, 4–23.

Delchambre, L. (2015), 'Weighted principal component analysis: a weighted covariance eigendecomposition approach', *Monthly Notices of the Royal Astronomical Society* **446**(4), 3545–3555.

Dew-Becker, I. (2022), 'Real-time forward-looking skewness over the business cycle', *NBER Working Paper No. 30478* .

Dew-Becker, I., Tahbaz-Salehi, A. & Vedolin, A. (2021), 'Skewness and time-varying second moments in a nonlinear production network: theory and evidence', *NBER Working Paper No. 29499* .

Dupraz, S., Nakamura, E. & Steinsson, J. (2023), A plucking model of business cycles, Technical report.

Favara, G., Gilchrist, S., Lewis, K. F. & Zakrajšek, E. (2016), 'Updating the recession risk and the excess bond premium', *FEDS Notes* .

Fernández-Villaverde, J. & Guerrón-Quintana, P. A. (2020), 'Uncertainty shocks and business cycle research', *Review of economic dynamics* **37**, S118–S146.

Ferreira, T. R. (2018), 'Stock market cross-sectional skewness and business cycle fluctuations', *International Finance Discussion Papers 1223* .

Foerster, A. T., Sarte, P.-D. G. & Watson, M. W. (2011), 'Sectoral versus aggregate shocks: A structural factor analysis of industrial production', *Journal of Political Economy* **119**(1), 1–38.

Friedrich, B., Laun, L. & Meghir, C. (2021), 'Income dynamics of immigrants and natives in sweden 1985-2016', *NBER Working Paper No. 28527* .

Gabaix, X. (2011), 'The granular origins of aggregate fluctuations', *Econometrica* **79**(3), 733–772.

Gabaix, X. & Koijen, R. S. (Forthcoming), 'Granular instrumental variables', *Journal of Political Economy* .

Gertler, M. & Gilchrist, S. (1994), 'Monetary policy, business cycles, and the behavior of small manufacturing firms', *The Quarterly Journal of Economics* **109**(2), 309–340.

Gertler, M. & Karadi, P. (2015), 'Monetary policy surprises, credit costs, and economic activity', *American Economic Journal: Macroeconomics* **7**(1), 44–76.

Gilchrist, S. & Zakrajšek, E. (2012), 'Credit spreads and business cycle fluctuations', *American economic review* **102**(4), 1692–1720.

Gorodnichenko, Y. & Lee, B. (2020), 'Forecast error variance decompositions with local projections', *Journal of Business & Economic Statistics* **38**(4), 921–933.

Goulet Coulombe, P., Leroux, M., Stevanovic, D. & Surprenant, S. (2022), 'How is machine learning useful for macroeconomic forecasting?', *Journal of Applied Econometrics* **37**(5), 920–964.

Groeneveld, R. A. (1996), The relative efficiency of several statistics measuring skewness, *in* H. N. Nagaraja, P. K. Sen & D. F. Morrison, eds, 'Statistical Theory and Applications: Papers in Honor of Herbert A. David', Springer, pp. 43–53.

Groeneveld, R. A. & Meeden, G. (1984), 'Measuring skewness and kurtosis', *Journal of the Royal Statistical Society: Series D (The Statistician)* **33**(4), 391–399.

Grömping, U. (2007), 'Estimators of relative importance in linear regression based on variance decomposition', *The American Statistician* **61**(2), 139–147.

Guvenen, F., Ozkan, S. & Song, J. (2014), 'The nature of countercyclical income risk', *Journal of Political Economy* **122**(3), 621–660.

Guvenen, F., Pistaferri, L. & Violante, G. L. (2022), 'Global trends in income inequality and income dynamics: New insights from grid', *Quantitative Economics* **13**(4), 1321–1360.

Herskovic, B., Kelly, B., Lustig, H. & Van Nieuwerburgh, S. (2016), 'The common factor in idiosyncratic volatility: Quantitative asset pricing implications', *Journal of Financial Economics* **119**(2), 249–283.

Higson, C., Holly, S. & Kattuman, P. (2002), 'The cross-sectional dynamics of the us business cycle: 1950–1999', *Journal of Economic Dynamics and Control* **26**(9-10), 1539–1555.

Higson, C., Holly, S., Kattuman, P. & Platis, S. (2004), 'The business cycle, macroeconomic shocks and the cross-section: the growth of uk quoted companies', *Economica* **71**(282), 299–318.

Hinkley, D. V. (1975), 'On power transformations to symmetry', *Biometrika* **62**(1), 101–111.

Ilut, C., Kehrig, M. & 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. & Ng, S. (2015), 'Measuring uncertainty', *American Economic Review* **105**(3), 1177–1216.

Kim, T.-H. & White, H. (2004), 'On more robust estimation of skewness and kurtosis', *Finance Research Letters* **1**(1), 56–73.

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

Lagerborg, A., Pappa, E. & Ravn, M. O. (2023), 'Sentimental business cycles', *The Review of Economic Studies* **90**(3), 1358–1393.

Liu, G. (2020), 'Data quality problems troubling business and financial researchers: A literature review and synthetic analysis', *Journal of Business & Finance Librarianship* **25**(3-4), 315–371.

Ludvigson, S. C., Ma, S. & Ng, S. (2021), 'Uncertainty and business cycles: exogenous impulse or endogenous response?', *American Economic Journal: Macroeconomics* **13**(4), 369–410.

Lundberg, S. M. & Lee, S.-I. (2017), 'A unified approach to interpreting model predictions', *Advances in neural information processing systems* **30**.

McCracken, M. & Ng, S. (2020), 'Fred-qd: A quarterly database for macroeconomic research', *NBER Working Paper No. 26872* .

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

Molnar, C. (2020), *Interpretable machine learning*, Lulu.com.

Oh, S. S. & Wachter, J. A. (2022), 'Cross-sectional skewness', *The Review of Asset Pricing Studies* **12**(1), 155–198.

Ottonello, P. & Winberry, T. (2020), 'Financial heterogeneity and the investment channel of monetary policy', *Econometrica* **88**(6), 2473–2502.

Plagborg-Møller, M. & Wolf, C. K. (2021), 'Local projections and vars estimate the same impulse responses', *Econometrica* **89**(2), 955–980.

Pruitt, S. & Turner, N. (2020), 'Earnings risk in the household: Evidence from millions of us tax returns', *American Economic Review: Insights* **2**(2), 237–254.

Ramey, V. A. (2016), 'Macroeconomic shocks and their propagation', *Handbook of Macroeconomics* **2**, 71–162.

Romer, C. D. & Romer, D. H. (2004), 'A new measure of monetary shocks: Derivation and implications', *American economic review* **94**(4), 1055–1084.

Salgado, S., Guvenen, F. & Bloom, N. (2023), 'Skewed business cycles', *Working paper* .

Shapley, L. S. (1953), 'A value for n-person games', *Contributions to the Theory of Games (AM-28), Volume II* .

Van Binsbergen, J. H., Han, X. & Lopez-Lira, A. (2023), 'Man versus machine learning: The term structure of earnings expectations and conditional biases', *The Review of Financial Studies* **36**(6), 2361–2396.

Wu, J. C. & Xia, F. D. (2016), 'Measuring the macroeconomic impact of monetary policy at the zero lower bound', *Journal of Money, Credit and Banking* **48**(2-3), 253–291.

Appendix A Data Preparation

I start from the entire Compustat database at the quarterly frequency. After the download, the data has 1,928,055 quarter-firm observations and covers the period 1961Q1 - 2022Q3. The date is defined using the item `datacqtr`, not the fiscal quarter. The unique firm identifier is `gvkey`. I drop firms that are not incorporated (variable `fic`) or headquartered (`loc`) in the United States. I remove any companies with an SIC code above 9000, which includes non-operating establishments. I drop any observations with negative nominal sales (`saleq`) and remove all duplicates of the firm-quarter identifier (`gvkey` and `datacqtr`).

Nominal sales are deflated with the GDP price deflator (USAGDPDEFQISMEI on FRED) to obtain real sales $s_{i,t}$ of firm i in quarter t. If a firm shows a missing value of real sales in a period that is surrounded by non-missing sales observations, I fill the missing value via linear imputation. If two missing values are adjacent, no imputation is performed. Real sales growth is the year-over-year growth rate of quarterly real sales: $g_{i,t} = \frac{s_{i,t}-s_{i,t-4}}{s_{i,t-4}}$. Analogously, real sales growth rates computed as log differences are $\tilde{g}_{i,t} = \ln(s_{i,t}) - \ln(s_{i,t-4})$. Aggregate real sales growth is

$$g_t = \frac{\sum_i s_{i,t-4} g_{i,t}}{\sum_i s_{i,t-4}} \tag{14}$$

This way of computing aggregate sales ensures that only growth rates of firms are considered that experience non-missing sales in both quarters. It is not biased by the entry of new firms or exit of dying firms. Similarly, I compute growth rates for different industries at the 2,3,4, and 5 digit NAICS level. An alternative measure of aggregate fluctuations used in this paper is real GDP growth, which is defined as the year-over-year growth rate of real GDP (item GDPC1 on FRED).

I construct several variables for firm characteristics, following Ottonello & Winberry (2020). Leverage is the ratio of total debt (sum of items `dlcq` and `dlttq` to total assets (`atq`). Net leverage is the ratio of total debt minus net current assets (`actq`) to total assets. Liquid assets ("liquidity ") is the ratio of cash and short-term investments (`cheq`) to total assets.

This yields the *full sample*. The full sample of non-missing sales growth rate observations has 1,145,568 firm-quarter observations and covers the period 1962Q1 – 2022Q3.

Additional steps yield the *cleaned sample*, which aims to remove sales growth rate outliers:

1. Remove any firm-quarter observations with negative current assets (`actq`), total assets (`atq`), or liquid assets.
2. To control for acquisitions, remove a firm-quarter observation if acquisitions account for more than 5 percent of total assets in the current or any of the three preceding quarters. This ensures that the year-over-year growth rate is not contaminated by previous acquisitions.
3. Remove the observation if net current assets relative to total assets falls outside of $[-10, 10]$.
4. Remove observations with leverage above 10 or below zero.
5. Remove any observations with percentage sales growth rate g below -1 (I do not apply this filter to log growth rates $\tilde{g}$).

6. Remove any observations where the ratio of sales to total assets is in the top 0.1% of observations. This is to clean any sales growth observations that may be due to mistakes in the data.

7. To further account for growth rate outliers, I remove the top and bottom 1% of growth rate observations in each quarter.

8. Since the data on acquisitions is only available from 1983Q3 onwards, I remove all earlier observations.

The resulting sample covers the period from 1983Q3 until 2021Q4 and has 499,249 firm-quarter observations. I merge this sample with information on stock prices (variable `PCLOSE`) and the first date of trading (`BEGDAT`) from CRSP using the `PERMCO-GVKEY` linking table. I also merge the sample with information on dates of incorporation from Worldscope Fundamentals (variable `INCORPDAT`) using the CUSIP identifier. This allows me to construct firm age as the minimum across 1) the date of the first observation in Compustat, 2) the first date of trading from CRSP, and 3) the date of incorporation as indicated in Worldscope Fundamentals. This approach makes use of the well-populated and accurate information in Worldscope while avoiding negative firm ages, as discussed in Cloyne et al. (2023).

To be able to work with within-firm time series variation in some parts of my analysis, I perform a final step of cleaning to yield the *streak sample*. As in Ottonello & Winberry (2020), I only keep growth rate streaks of 40 consecutive quarters, and remove all other observations. This yields a sample of 2,520 unique growth streaks for 2,376 unique companies. 144 companies have two sales growth rate streaks in the data. The sample period is 1983Q3 until 2021Q4 and there are 151,701 firm-quarter observations. To approximate a balanced panel, the *long* sample only consider firms within the clean sample that have observations before 1985Q1 and after 2021Q1. This leaves 661 unique firms. The firms may have missing values within their time series of sales growth rates, but the share of missing values in the sales data is only 5.6%. In contrast, the streak sample has almost 50% missing values relative to a balanced panel.

Figure 13 shows the number of firm-level observations per quarter for the different samples. The full sample contains less than 2000 per quarter firms in the 1960s. The number of firms jumps strongly in the 1980s to around 6000 observations per quarter and keeps rising until the burst of the dot-com bubble, when the number of observations per quarter peaks at over 8000. Since then, the number of firms has declined and reaches around 4000 in 2020. In the cleaned sample, the number of firms is half as high as in the full sample but follows a similar pattern over time. The streak sample drops more than half of the observations of the cleaned sample and has around 1000 observations per quarter for most quarters. The number of firms is lower toward the start and the end of the sample. By construction, the number of observations per quarter is very constant over time for the long sample.

Despite differences in in the number of observations, the skewness pattern across samples is very similar. Figure 14 computes Kelley skewness in sales growth rates for each of the four samples.

Figure 13: Number of Sales Growth Observations per Quarter

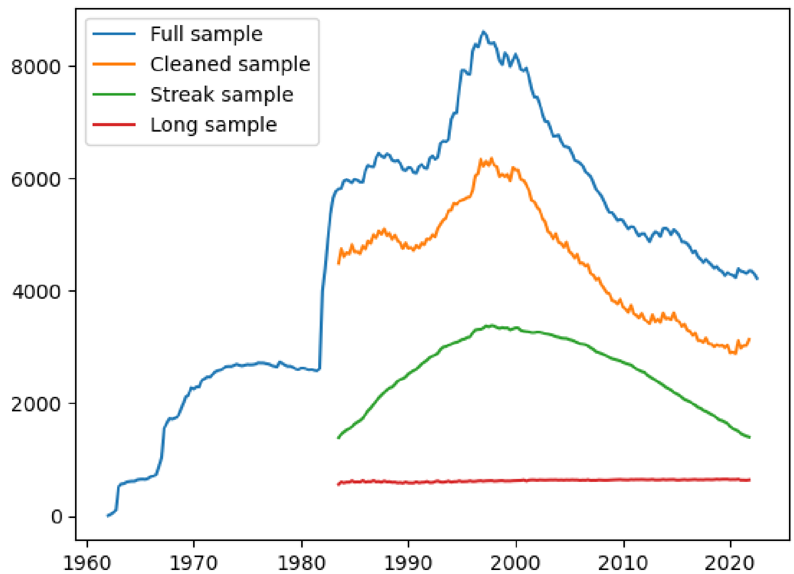

Note: The full sample of growth rates covers 1962Q1-2022Q3, and the cleaned sample covers 1983Q3-2021Q4.

Figure 14: Kelley skewness for different samples

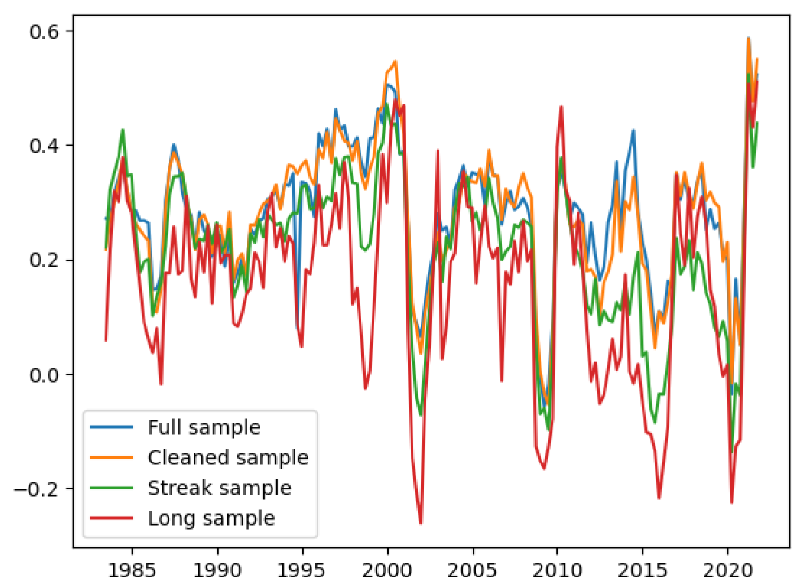

Note: Skewness is computed using 90% Kelley skewness. The sample period is 1983Q3-2021Q4.

Appendix B Checking the procyclicality of micro skewness

This section describes results on skewness measurement and the relation of cross-sectional skewness across firms with the business cycle. A priori, this relationship is ambiguous in terms of sign and strength. Existing theories of business cycle fluctuations offer no clear prior information to guide the analysis. For example, production network models as in Baqaee & Farhi (2019) show that an increase in dispersion over firm outcomes reduces output, but this is true irrespective of any asymmetry in the distribution. Such theories are also consistent with the idea of symmetric cross-sectional distributions and no correlation between skewness and the business cycle. In contrast, procyclical skewness may tell a story about firm-specific disasters in recessions and/or idiosyncratic outperformance of some firms in booms. Countercyclical skewness could occur if the majority of firms perform poorly in recessions while some are able to achieve high growth rates, yielding a stronger right tail.

The goal of this section is to cleanly establish the correlation between skewness and the business cycle. This is non-trivial because skewness can be measured in different ways, and some skewness measures are highly sensitive to outliers. Therefore, I revisit earlier results on skewness in the cross-section of firm sales growth and add to the existing evidence by paying particular attention to different forms of skewness measures and the role of outliers. The main result is that skewness across firms has a positive correlation with the business cycle.

My main data source is Compustat. Compustat contains rich data on firm characteristics, including balance sheet information, that allow to study potential drivers of cross-firm skewness in detail. Because the data is of quarterly frequency, I can study business cycle dynamics with greater detail than with annual data. Especially when identifying shocks that may drive the business cycle, the quarterly frequency enables cleaner identification and more power. In addition, Compustat data is available from the early 1960s. It covers most post-WW2 US recessions. In addition, some firms have long individual time series such that I can exploit time series variation within firms. Compustat also covers many industries, which is an advantage over existing work that focuses on specific sectors such as manufacturing. Since different industries may exhibit heterogeneous co-variation with the business cycle, focusing on a subset of industries is not desirable a priori. I am not aware of any other US firm-level dataset that satisfies these requirements. Because Compustat is the most commonly used firm-level dataset for the US, I can directly put my findings in the context of the broader literature.

A potential drawback is that Compustat only covers publicly listed firms, which are larger on average, may face different financial constraints than private/small firms, and hence can have different cyclical behaviour (Axtell (2001), Gertler & Gilchrist (1994)). While smaller firms may be more cyclical than large firms, Crouzet & Mehrotra (2020) show that this difference is dominated by the dispersion of firm size and therefore too small to meaningfully affect macroeconomic aggregates. This suggests that focusing on Compustat firms can be sufficient to understand aggregate fluctuations. Since Compustat does not provide employment data at the quarterly frequency, I focus on sales only. Liu (2020) points to general concerns with the data quality in Compustat. Addressing these issues is beyond the scope of this paper.

I describe the sample construction and data cleaning procedures in Appendix A. The *full* sample applies minimal cleaning procedures and yields a sample of 1,145,568 firm-quarter observations covering the period 1962Q1 - 2022Q3. The *cleaned* sample removes observations with high leverage, acquisitions, negative assets, and potential data mistakes. It also removes the top and bottom 1% of growth rate observations in each quarter to account for potential outliers. The cleaned sample has 499,249 firm-quarter observations and covers the period 1984Q2 - 2021Q4. I start from the full sample to give the most comprehensive picture of skewness dynamics, before later confirming the results for the cleaned sample.[22]

Before detailing different skewness measures, I simply plot the density of sales growth rate outcomes in NBER recessions versus non-recession periods. Figure 15 shows the result using the cleaned sample.[23] The orange line is the density outside of recessions and the blue line shows the distribution for recession periods. In recessions, sales growth rates are clearly lower on average (equal-weighted mean of 5.9% versus 13.6%). Although it is harder to see from the figure, the standard deviations are similar across distributions and actually higher outside of recessions (58.7% in recessions versus 62.6% outside of recessions). To focus the comparison on the skewness, the green line adjusts the recession density such that it has the same mean and variance as the non-recession distribution. The asymmetry of the distribution clearly changes. The rest of this section studies this change in more detail.

Figure 15: Estimated Densities for Recession vs Non-Recession Quarters

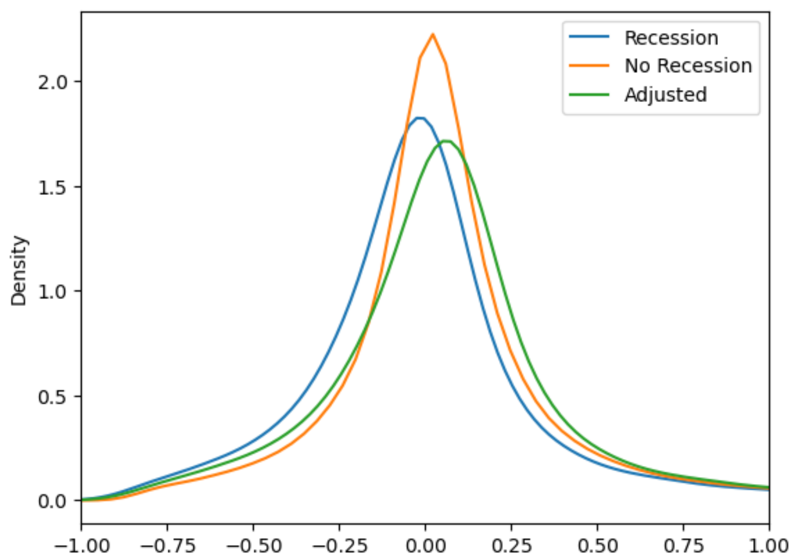

Note: Densities are estimated from the cleaned sample, which covers the period 1984Q2 – 2021Q4. Recessions are defined as by the NBER. The adjusted density has the same mean and variance as the non-recession distribution.

[22]In addition, Appendix A shows how to construct a *streak* sample that only contains observations that are part of a firm's sales growth streak of at least 40 consecutive quarters without a missing observation. This is useful for later parts of the paper that aim to exploit within-firm time series variation. The same appendix also constructs a *long* sample that approximates a balanced panel with only few missing values. This section does not make use of either of these samples.

[23]As I detail below, the full sample has many outliers, which make it difficult to reliably compare the distributions. The cleaned sample yields reliable results without requiring further outlier adjustments.

B.1 Correlation of skewness with the business cycle

To measure skewness more formally, I consider different variants of skewness coefficients. All variants measure the asymmetry of a distribution. Given observations $\{x_i\}_{i=1,\dots,N}$, the third moment is defined as

$$SK = \mathbb{E}\left(\frac{x_i - \mu}{\sigma}\right)^3, \tag{15}$$

where $\mu = \mathbb{E}(x_i)$ and $\sigma^2 = \mathbb{E}\left(x_i - \mu\right)^2$. The third moment can be arbitrarily large and need not be finite for fat-tailed distributions.[24] Since observations are taken to the third power, the measure is sensitive to outliers, and its estimation with small samples can be noisy.

Alternatively, the quantiles of the distribution can be used to measure its asymmetry. Hinkley (1975) proposes

$$SK_2 = \frac{(Q_{1-\alpha} - Q_{0.5}) - (Q_{0.5} - Q_\alpha)}{Q_{1-\alpha} - Q_\alpha}, \tag{16}$$

where Q_α is the $\alpha \times 100\%$ quantile of the distribution of $\{x_i\}$. The measure is positive if the deviation of the $1-\alpha$ quantile from the median is larger than the deviation of the α quantile from the median. The skewness measure does not require any assumptions on the existence of moments. Since the quantiles are always finite, they yield a finite measure of skewness even in the presence of fat tails. The rescaling by overall dispersion $(Q_{1-\alpha} - Q_\alpha)$ ensures that skewness lies within $[-1, 1]$. A common choice is $\alpha = 0.1$, which is often referred to as Kelley skewness. While its estimation is more robust to outliers, the main drawback is that it does not consider information about the tails of the distribution, which may be of interest in certain cases. Additionally, the choice of α is arbitrary.

Groeneveld & Meeden (1984) suggest a generalization that is independent of the choice of α:

$$SK_3 = \frac{\int_0^{0.5}\{Q_{1-\alpha} + Q_\alpha - 2Q_{0.5}\}d\alpha}{\int_0^{0.5}\{Q_{1-\alpha} - Q_\alpha\}d\alpha} = \frac{\mu - Q_{0.5}}{\mathbb{E}|x_i - Q_{0.5}|} \tag{17}$$

For simplicity, I sometimes refer to this measure as "GM skewness". Groeneveld (1996) shows that the estimator of SK_3 can be vastly more efficient than the estimator of SK if the data is at least moderately fat-tailed (Student t-distribution with less than 10 degrees of freedom). Kim & White (2004) compare the skewness measures introduced above and caution against the use of the third moment due its sensitivity to outliers.

For these reasons, much of the modern literature on skewness in macroeconomics and finance makes use of robust skewness measures, most commonly Kelley skewness (SK_2). Guvenen et al. (2014), Pruitt & Turner (2020), Friedrich et al. (2021), Busch et al. (2022), and Guvenen et al. (2022) use Kelley skewness to study asymmetries in the income distribution. Decker et al. (2016), Ilut et al. (2018), and Salgado et al. (2023) use Kelley skewness to study asymmetries in distributions of firm outcomes like sales or employment growth. Ferreira (2018) measures skewness in the distribution of stock returns using Kelley skewness. Some recent papers in the finance literature use the third moment instead to measure skewness in the cross section, for example Catherine et al. (2022) and Oh & Wachter (2022).

[24] For example, a Pareto distribution with density function $f(x) = \frac{\alpha x_m^\alpha}{x^{\alpha+1}}$ for $x \geq x_m$ has finite expected value iff $\alpha > 1$ and finite third moment iff $\alpha > 3$.

Following the recent literature, the main skewness measure of this paper is Kelley skewness (SK_2). For robustness, I also consider the third moment (SK) and the Groeneveld & Meeden (1984) measure (SK_3). This choice is motivated by two features of the data. First, the distribution of real sales growth is fat-tailed in Compustat, as I discuss in Section B.2. Second, the number of firm-level sales growth observations is limited. It never exceeds 10,000 per quarter in the full sample, and lies between 1,500 and 2,000 in the cleaned sample (see Figure 13). In addition, Kelley skewness can easily be interpreted and decomposed into downside ($Q(1-\alpha)-Q(0.5)$) and upside ($Q(0.5)-Q(\alpha)$) drivers of skewness.

To study the correlation of cross-sectional skewness with the business cycle, I use two variables of aggregate fluctuation. The first is the aggregate sales growth rate. This measures the real, year-over-year growth rate in aggregate sales measured across Compustat firms. The second variable is real year-over-year GDP growth. Both growth rates are in real terms, at the quarterly frequency, and measured relative to the previous year's quarter. See Appendix A for details. I focus on these variables as they can most naturally be compared to firm-level sales. While aggregate sales can be constructed bottom-up from firm level sales, GDP can be expected to be less directly related to firm-level sales since GDP is constructed from domestic output instead of worldwide sales of US firms.

Table 10 shows the correlation of the within-quarter skewness across firms' sales growth rates with the business cycle. The top panel uses aggregate sales growth as the measure of aggregate activity, while the bottom panel uses GDP growth. The different columns refer to different measures of skewness: Kelley skewness computed with α taking values of 0.9, 0.95 and 0.99, the third moment, and the GM skewness.

The correlation is positive for the 90% and 95% Kelley skewness measures, with values ranging from 0.32 to 0.64. Going further into the tail reduces the correlation between skewness and the business cycle and yields results close to zero for the 99% Kelley skewness, the third moment, and the GM skewness. The results are similar for both business cycle measures, but the correlations are consistently weaker when using GDP growth. This is to be expected since firm-level sales are more closely related to aggregate sales than to GDP.

The table also splits the sample into two subsamples: the period before 1984 and the period from 1984 onwards. The second subsample coincides with the cleaned sample, which will be used in later analyses. This allows to understand to which extent the results may be driven by the choice of the sample period. The correlation of cross-firm skewness with the business cycle is weaker in the early subsample than in the post-1984 period. This holds across all skewness measures. For the post-1984 sample, the correlations are all positive and reach values up to 0.86 for the 90% Kelley skewness. The correlations are again lower when using the third moment and the GM skewness, and generally indicate no strong association with aggregate activity. The 99% Kelley skewness shows a positive correlation for this subsample, with a value of 0.42 and 0.45 for sales growth and GDP growth, respectively.

To understand the differences in the correlations across subsamples better, Figure 16 plots rolling window correlations between skewness and aggregate sales growth over time. The window width is fixed at 40 quarters, yielding a time series starting in the early 1970s. The blue line shows the rolling correlation between the 90% Kelley skewness and aggregate sales growth. The orange line uses 95% Kelley skewness, and the green line uses the third moment. Each point on the line shows the correlation between skewness and aggregate sales over the previous 40 quarters.

The blue and orange line move closely together. The differences in correlation values across subsamples are due to the sharp drop in correlation in the early 1980s: the correlation is relatively stable over most of the sample with values ranging between 0.5 and 0.9. When computing the correlation with data between the early 1970s and the early 1980s, however, the correlation decreases sharply to a low of around 0 (90% Kelley) or -0.3 (95% Kelley). The subsequent increase in the correlation is equally sharp, and the correlation for the period 1980-1990 is again between 0.6 and 0.8. The association between firm-level skewness and the business cycle has therefore been stable over the past decades except for correlations computed on 1972-1987 data.

The correlation of the third moment with the business cycle behaves similarly to the Kelley measure until the year 2000. While the correlation values are generally weaker, they show the same pattern of decline and recovery in the 1980s. The main difference to the Kelley measure occurs for the post-2000 period, in which the third moment has no significant association with aggregate sales growth. These differences are likely due to the role of outliers, to which I turn next.

Table 10: Correlations of skewness with the business cycle

	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
With sales growth					
Full sample	0.64	0.41	-0.07	-0.04	-0.00
Pre-1984	0.49	0.27	-0.16	-0.05	0.1
Post-1984	0.86	0.78	0.42	0.08	0.22
With GDP growth					
Full sample	0.52	0.32	-0.11	-0.14	-0.09
Pre-1984	0.50	0.30	-0.09	-0.21	0.03
Post-1984	0.69	0.68	0.45	0.06	0.11

The first three columns compute skewness using the Kelley measure (SK_2), with different choices for α. The last column computes the Groeneveld & Meeden (1984) skewness coefficient (SK_3). The pre-1984 sample is 1962Q1-1983Q4. The post-1984 sample is 1984Q1-2022Q3.

B.2 Outliers in Compustat data

I perform several robustness checks with different data cleaning procedures. The Compustat data contains many extremely large sales growth rate observations. In the full sample, the largest growth rate is $24,353,692\%$ for the firm Wynn Resorts in 2005Q3. Wynn Resorts develops and operates high-end casinos and hotels. It has been founded in 2002 and traded on the Nasdaq since then, but only opened its first resort in 2005Q2. Sales reported in Compustat are virtually zero before the opening date, and jump to a reported value of $243,536,925,850$ US dollars once the resort opens.[25]

This points toward two problems in the Compustat data. First, reported sales sometimes appear unrealistically large, which is likely due to data errors. Second, firms with previously low sales volume that suddenly take up operations experience gigantic sales growth rates. The implications for the

[25] For comparison, Apple reported global revenue of 117bn US dollars for 2023Q1, see Apple website.

Figure 16: Rolling Window Correlations of Skewness with Aggregate Activity

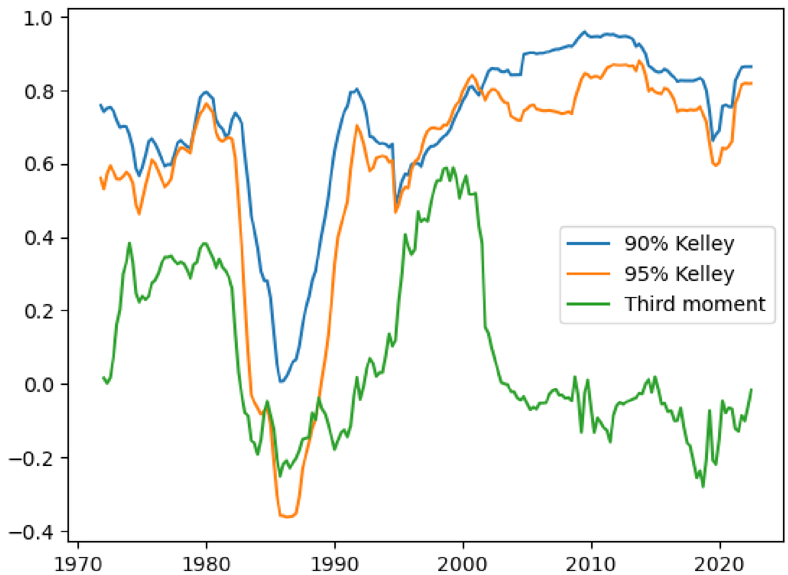

Note: Each line represents the rolling-window correlation between aggregate sales growth and the respective measure of firm-level skewness. The window has a fixed width of 40 quarters. Each point shows the correlation over the past 40 quarters.

estimation of skewness can be large: Removing the sales growth rate of Wynn Resorts changes the estimated third moment in 2005Q3 from 79.5 to 63.5, despite there being 6328 observations in the quarter.

Table 11 compares characteristics of the largest sales growth rate observations to the rest of the full sample.[26] For the largest 50, 100, 500, 1000, and 5000 outliers in terms of the sales growth rate, the table shows the average (across firms and quarters) of the sales growth rate, sales over assets, lagged real sales, total real assets, acquisitions over assets, and the share of observations occuring from the year 2000 onwards. By construction, mean sales growth rates are vastly higher in the outlier sample. The 50 largest outliers show an average growth rate of over one million percent. Even the largest 5000 growth rate observations have a mean growth rate of $25,891\%$. In the rest of the sample, the (equal-weighted) mean sales growth rate falls as more outliers are removed and reaches 21% when 5000 outliers are cleaned.

Mistakes in the reported value of sales can manifest as disproportionately large sales relative to total assets. The ratio of sales over total assets is 9.4 for the top 50 outliers, whereas it is 0.3 for the rest of the sample. The value falls as larger samples of outliers are considered but stays constant for the rest of the sample, suggesting that the high sales-to-assets ratios among outliers are driven by the very largest sales growth observations.

Large growth rate observations can also occur when firms have previously experienced virtually no sales but were already listed on a US stock exchange. For example, this can be the case for developers

[26] Additional characteristics that remain unreported show no clear patterns between outliers and the rest of the sample. For example, the outliers considered here are not concentrated in a certain industry, do not only occur for firms with high leverage, and do not cluster in the Covid period.

such as in the case of Wynn Resorts explained above. Another example in the data are biotech companies that sell shares to finance the development of a new product: As soon as the product hits the market, sales soar from previously miniscule levels. The data confirms that four-quarter lagged real sales (which enter the growth rate calculation) are considerably smaller among the outlier observations than in the rest of the sample. Again, this is especially true for the top 50 outliers. The outlier firms are also considerably smaller in terms of total assets. The top 1000 outlier observations correspond to firms with real assets of about 200mn US dollars, whereas the average balance sheet size in the rest of the sample is over 4bn USD.

Firms may also experience a sudden surge in sales if they merge with or acquire another firm. To test if mergers and acquisitions explain some of the outliers, I compute the share of acquisition expenditure relative to total assets. While the top 50 outlier firms do not have a higher expenditure on acquisitions than the rest of the sample, the top 1000 outliers spend 3.5% of total assets on acquisitions compared to 1.2% in the rest of the sample. This suggests that a significant share of outliers may be explained by $M\&A$ activity.

Lastly, to study why the outlier-sensitive third moment shows lower correlations of skewness with the business cycle for the post-2000 period (as shown in Figure 16), the bottom panel of Table 11 contains the share of observations that occur in the year 2000 or after. While only 37% of observations in the full sample occur in the post-2000 period, among outliers these years account for 52 to 78 percent of observations. Especially the very largest outliers are clustered in the 2000s and 2010s, which may explain the weak correlation between aggregate activity and skewness as measured by the third moment.

In summary, sales growth outliers in the Compustat data have vastly larger values than the rest of the sample and may bias the skewness estimation significantly. Extreme outliers can be due to mistakes in the data (high sales relative to total assets), a low base level of sales, or due to acquisitions. The next section aims to control for these factors and understand their implications on the correlation between skewness and the business cycle.

B.3 Outlier correction and other robustness checks

To get as clear a picture on the role of outliers as possible, I use different methods to control for outliers and compare their effect on the correlation between cross-sectional skewness and the business cycle. The first set of controls removes the top and bottom $x\%$ from the sample. This addresses outliers across the entire sample, but can potentially remove outliers asymmetrically within a given period. Since skewness is a measure of asymmetry, this method may fail to accurately control for outliers if there remain some observations that bias the skewness estimates in certain quarters. A second set of controls removes the top and bottom $x\%$ of observations for each quarter. This necessarily removes as many observations at the top and the bottom of the distribution in any given quarter. The drawback is that it might fail to remove some outliers in certain quarters while removing observations in other quarters that wouldn't necessarily be categorized as outliers when considering the entire sample.

Table 12 contains the results. Panel a uses aggregate sales growth as the business cycle indicator,

Table 11: Characteristics of largest sales growth rate outliers

	Number of outliers				
	50	100	500	1000	5000
Mean sales growth rate					
Outliers	1,265,899%	738,290%	198,418%	110,001%	25,891%
Rest of sample	78%	69%	47%	38%	21%
Sales over Assets					
Outliers	9.4	9.7	2.8	2.2	1.2
Rest of sample	0.3	0.3	0.3	0.3	0.3
Lagged real sales					
Outliers	562	766	6,089	8,102	91,340
Rest of sample	479,371	479,391	479,546	479,741	481,294
Total real assets					
Outliers	424,798	260,620	256,979	202,233	224,414
Rest of sample	4,041,914	4,042,091	4,043,334	4,044,791	4,057,315
Acquisitions over Assets					
Outliers	0.5%	2.3%	3.5%	3.1%	2.5%
Rest of sample	1.2%	1.2%	1.2%	1.2%	1.2%
Share of post-2000 observations					
Outliers	78%	73%	63%	60%	52%
Rest of sample	37%	37%	37%	37%	37%

All values are averages over all firm-quarter observations in the respective samples. Rest of sample refers to the full sample without the growth rate outliers. Some statistics are subject to necessary data cleaning. For computing sales over assets, I remove all observations with negative or zero assets. For acquisitions over assets, I remove all observations with negative or zero assets and all observations with negative acquisitions. Acquisitions over assets are computed for the current and the three preceding quarters. Real sales and total real assets are in thousands of 2015 US dollars, using the GDP deflator.

panel b uses GDP growth. The values for x are 0.1, 1, and 2. I compute the same skewness measures as above. The results are similar to those from Table 10. The correlation estimates are generally strongest for the 90% Kelley skewness measure, for the post-1984 period, and for aggregate sales growth as opposed to GDP growth as a measure of aggregate activity. Even after removing the top and bottom 2% of outliers, the correlation between the third moment and sales growth is -0.37 for the post-1984 period. When removing the top and bottom 2% of outliers in every quarter, the correlation changes to 0.65 for the same period and skewness measure. This suggests that the sales growth rates in Compustat are frequently large enough to make the third moment a rather unreliable measure of the asymmetry in the cross section.

In contrast, the skewness measures based on quantiles paint a more consistent picture. The correlations are generally positive. The 99% Kelley skewness sometimes suggests correlations around zero and potentially slightly negative, but this behaviour is driven by the pre-1984 period. In the sample starting in 1984, the correlations are positive across all variants for outlier control. For the post-1984 sample, the correlations range from 0.33 to 0.86 and are mostly above 0.6. This holds for both aggregate sales growth and GDP growth as measures of the business cycle. The quantile-based skewness measures therefore indicate a strong positive correlation between skewness and the business cycle. Recessions are periods in which the left tail of the cross-firm sales growth distribution widens, such that some firms experience particularly bad outcomes. This supports the results of Salgado et al. (2023).

Earlier findings in Higson et al. (2002) and Higson et al. (2004) suggest that skewness across firms is *countercyclical* instead of procyclical. Higson et al. (2002) use annual Compustat data from 1950-1999 and compute skewness via the third moment. To control for outliers, they remove all observations with growth rates larger than 25% in absolute value (robustness checks include cutoffs of 50% and 100%). Their main finding is a negative association between GDP growth and their skewness estimate. Higson et al. (2004) perform a similar analysis for a sample of UK quoted companies from 1967 to 1999, again applying the symmetric cutoff rule to remove outliers. They confirm their earlier evidence of countercyclical skewness.

The symmetric cutoff rule may explain the results of countercyclical skewness. Since growth rates are positive on average, it is not clear that applying a rule that removes any growth rates symmetrically around zero is desirable. Instead, one could apply these cutoffs symmetrically around the mean growth rate of the sample. More importantly, business cycle fluctuations imply that observations above 25% year-over-year growth should be expected more frequently in boom times relative to recessions, and growth rates below -25% should be expected to occur more regularly in recessions than in booms. The fixed cutoff values imply more observations in the left tail are removed during recessions (leading to higher skewness), while more observations in the right tail are removed during booms (leading to lower skewness). This can bias the correlation between skewness and the business cycle downward.

I confirm this intuition in Table 13. Using the sample of quarterly growth rates from Compustat, I find that applying a symmetric cutoff rule like in Higson et al. (2002) biases the skewness estimation and leads to significantly lower correlation estimates. The correlation becomes more negative for skewness measures that rely more strongly on the tail, and as the cutoff rule becomes more stringent. The 90% Kelley measure gives a correlation with sales growth of 0.64 for the 100% cutoff rule, but a correlation of -0.27 for the 25% cutoff rule. For the given cutoff rules, most of the correlations are negative and can reach values up to -0.81. However, it appears that the strength of these results is driven by the bias introduced through the cutoff rules themselves.

The cutoff rules strongly change certain properties of the sample. Table 14 reports several descriptive statistics for the samples associated with the different methods for outlier removal. The average growth rates when applying the cutoff rules are considerably lower since a larger share of highly positive growth rates is removed. The cutoff rule also enforces symmetry around zero in terms of the maximum and minimum values, which is not the case for the other methods. Importantly, the cutoff rules remove big parts of the original sample. For example, the 25% cutoff rule drops over 370,000 observations, or more than 30% of the sample. Enforcing the 25% cutoff is equivalent to removing the top 22% of growth rate observations and dropping the lowest 10% of observations. I conclude that skewness estimation can be highly sensitive to the choice of data cleaning procedure. The evidence presented so far speaks strongly in favour of procyclical skewness across a variety of estimates and samples, whereas results of countercyclical skewness appear to be driven by biases introduced through certain methods for outlier removal.

Since percentage growth rates are bounded below by -100% but are not bounded above, skewness estimates may be biased disproportionately by positive outliers. To test if this biases the correlation estimates, Table 15 contains results using growth rates computed as log differences, which are not bounded below. The correlations are positive across skewness measures. The values decrease for measures that rely more on the tail and are generally higher in the post-1984 period. The results do not depend strongly on the business cycle measure and are very similar for aggregate sales versus GDP growth. Removing outliers by dropping the 2% of largest and smallest observations in each quarter increases the correlations for especially for the third moment and the GM skewness. In the post-1984 period, the correlations range from 0.72 to 0.84 for aggregate sales and 0.56 to 0.72 for GDP growth. Overall, the results based on log differences support those obtained using percentage growth rates.

Table 15 also shows correlation estimates using weighted growth rates. Since many outliers are due to firms with small base levels of sales, I weight the sales growth rates with lagged sales. This is equivalent to simply using four-quarter differences of real sales ($s_t - s_{t-4}$) instead of growth rates. Using first differences puts higher weight on large firms, which are likely to have a larger impact on aggregate fluctuations. Just like for sales growth rates, it is a priori not clear that the skewness in the cross section should be correlated with the business cycle. Recessions could materialize as shocks that reduce firm sales by the same rate across all firms, in which case the skew of the distribution of sales changes would remain unaffected. However, the data suggest a clear association between changes in the asymmetry of the sales change distribution and the business cycles. The correlations are positive across measures and get stronger when outliers are removed.

So far, the analysis has focused on the full sample to provide the most comprehensive perspective on the Compustat data. I confirm that the results also hold for the *cleaned* and the *streak* sample, which will be used in later parts of the analysis. Table 16 shows the results. Panel a computes correlations between skewness and aggregate sales growth for the three different samples using different skewness measures. The correlations are higher for the cleaned sample and the streak sample, especially for skewness measures that are more sensitive to the tails, like the 99% Kelley skewness and the third moment. Correlation estimates range from 0.13 to 0.82, and from 0.73 to 0.82 when excluding the 99% Kelley skewness and the third moment. Panel b repeats the analysis using GDP growth. The results are very similar and confirm the procyclicality of cross-firm skewness.

Panel c of Table 16 shows the mean, minimum, and maximum sales growth rate observations in the sample. While many growth rate observations have been dropped compared to the full sample, as is evident from the number of observations shown in panel c, the range of growth rates is still large in the cleaned data. The minimum growth rate is -95%, while the maximum growth rates are 1,782% and 1,283% for the cleaned and the streak sample, respectively. The (equal-weighted) mean growth rates are 13% and 7%, which contrasts with the average growth rate of 134% in the full sample including outliers.

Panel d reports several sample characteristics that have been shown to differ systematically between outliers and the rest of the sample, see Table 11. All samples have similar sales over assets, but lagged real sales are highest for the streak sample. Total real assets fall from the full to the cleaned sample, but are higher in the streak sample than in the cleaned sample. These differences arise because the streak sample only considers firms that exist for at least 40 quarters, which are likely to be larger than younger firms. Average acquisitions over assets decrease from 1.2% in the full sample to 0.2% in the clean sample and to 0.1% in the streak sample. This is in line with the observation that many outliers had been associated with acquisitions.

To conclude this subsection, Figure 17 plots the time series of skewness against the business cycle. The left panel plots skewness against aggregate sales growth and the right panel against GDP growth. Skewness is estimated from the cleaned sample using the Kelley measure with $\alpha = 0.9$. As the correlation estimates suggest, it becomes visible that skewness is strongly procyclical. Every downturn in aggregate sales is associated with a closely following decline in the cross-sectional skewness: Micro and macro skewness are closely associated. The pattern is similar but somewhat weaker for GDP growth. In particular, since the decline in aggregate sales around 2015 is not associated with a decline in GDP growth rates, the skewness measure indicates a downturn during that period which is not reflected by GDP growth as a business cycle indicator. The correlation between GDP growth and micro skewness is still high. The weaker association is rather a reflection of the link between sales and GDP than on the association between micro and macro skewness, which is very strong when focusing on sales growth (left panel).

A detailed interpretation of the business cycle co-movement of the sales distribution requires to consider not just the changes in skewness over time, but also the level of skewness. Figure 17 indicates cross-sectional skewness is positive in general and the distribution becomes more symmetric in recessions. However, this behaviour may be driven by the lower bound on growth rates when computing percentage growth. To interpret the level of skewness, Figure 18 compares two different skewness estimates against aggregate sales growth. The left panel computes skewness based on year-over-year log differences of real sales and the right panel computes skewness based on year-over-year differences in real sales levels. Neither measure imposes a lower or upper bound on the underlying data. The results indicate that cross-sectional skewness is positive in expansions but turns negative during recessions. This supports the interpretation of recessions as being associated with firm-level disasters (left skew), while booms coincide with strong growth for some firms (right skew).[27]

The following subsections will study additional aspects of the relation between micro and macro skewness. To simplify the analysis, I proceed by working with the cleaned sample and using the 90% Kelley measure. This measure is widely used in the recent literature, and the analysis presented in this subsection shows that it accurately captures the positive correlation between skewness and the business cycle. While it presents the strongest correlation results compared to the other measures, even measures that go far into the tails of the distribution yield qualitatively (and often quantitatively) very similar results.

B.4 Lead-lag relationship

So far, the analysis has focused on the contemporaneous correlation of skewness with the business cycle. This section provides additional evidence on potential lead-lag relationships. While there is mild evidence of short-term lead and lag correlations, the main result is that the contemporaneous correlation is the strongest and is therefore the main focus of the analysis.

[27] As shown in Table 15, using percentage growth rates, log differences, or differences in levels of real sales leads to identical interpretations for the procyclicality of cross-sectional skewness. All measures are highly correlated and indicate very similar changes in skewness, but they give different interpretations for the level of skewness.

Figure 17: Cross-sectional skewness vs aggregate activity

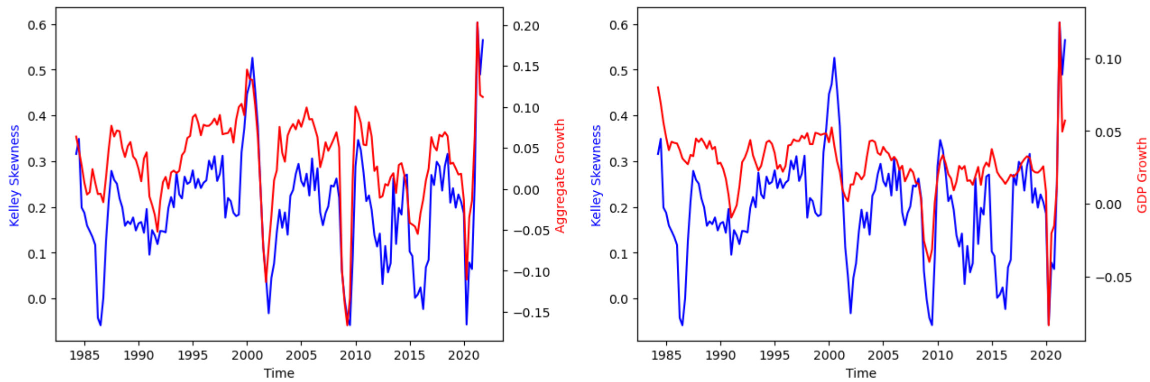

Note: The left figure compares skewness to year-over-year aggregate sales growth. The right panel plots the same skewness estimate against year-over-year GDP growth. Skewness is estimated using 90% Kelley skewness. The sample period is 1984Q2-2021Q4.

Figure 18: Cross-sectional skewness vs aggregate activity - Alternative skewness series

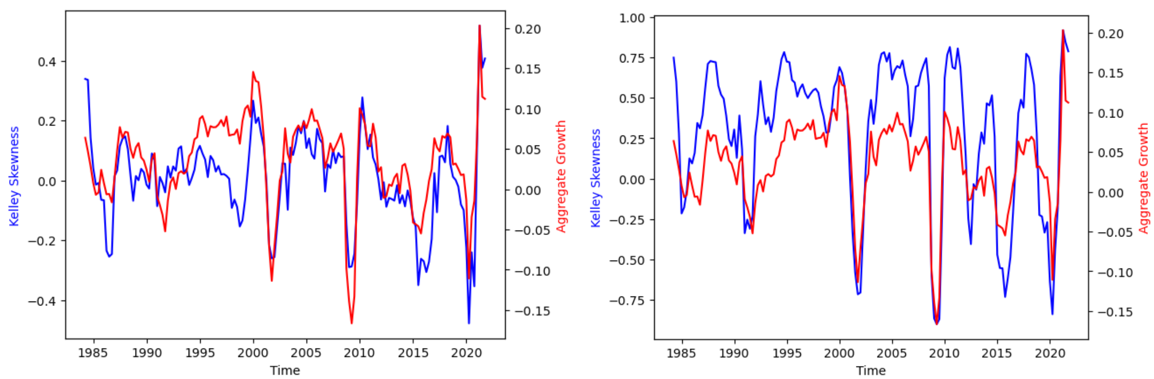

Note: Aggregate activity is measured by year-over-year aggregate sales growth. The left panel computes skewness based on log differences of real sales. The right panel plots skewness estimates for four-quarter differences of real sales $(s_t - s_{t-4})$. Skewness is estimated using 90% Kelley skewness. The sample period is 1984Q2-2021Q4.

Figure 19 plots the correlations $\rho_j = corr\left(\text{skew}_{t+j}, \text{growth}_t\right)$ for $j = -24, ..., 0, ..., 24$. Skewness is the 90% Kelley measure, and aggregate growth refers to sales growth. The contemporaneous correlation is highest at 0.82, as estimated before. Aggregate growth and skewness also have a high correlation when skewness is lagged by up to two quarters or leads by up to three quarters. Lagged correlations are lower than lead correlations. For example, lagging skewness by one period yields a correlation of 0.59, while leading it by one quarter yields a correlation of 0.77.

To control for lags more formally, Table 17 shows estimation results from regressions of aggregate activity on skewness. All variables are standardized to make the coefficient sizes comparable across variables. The left panel uses aggregate sales growth as the dependent variable, while the three columns on the right use GDP growth. A large share of the variation in aggregate sales growth can be explained by skewness alone: Regressing sales growth on skewness yields an adjusted R^2 of 0.68. Including four lags of the skewness measure does not improve the linear fit and does not decrease the coefficient on contemporaneous skewness. The lagged coefficients are generally insignificant, except for a significant yet small coefficient estimate on two-quarter lagged skewness. Controlling for lagged sales growth (column 3) improves the fit slightly and yields a significant coefficient on the first two lags. The coefficient on skewness still remains significant, and the magnitude remains large: All else equal, a one standard deviation increase in skewness is associated with a 0.46 standard deviation increase in aggregate sales growth. The results for GDP growth as the dependent variable are similar. Due to the weaker link between sales and GDP, the R^2 is generally lower and the coefficient on skewness is lower. Nevertheless, it remains significant at the 5% level across regressions. The only exception is the regression controlling for lagged GDP growth (column 6), where the coefficient on skewness is only significant at the 10% level and drops to 0.22.

Figure 19: Lead-lag correlations of Skewness vs Aggregate Growth

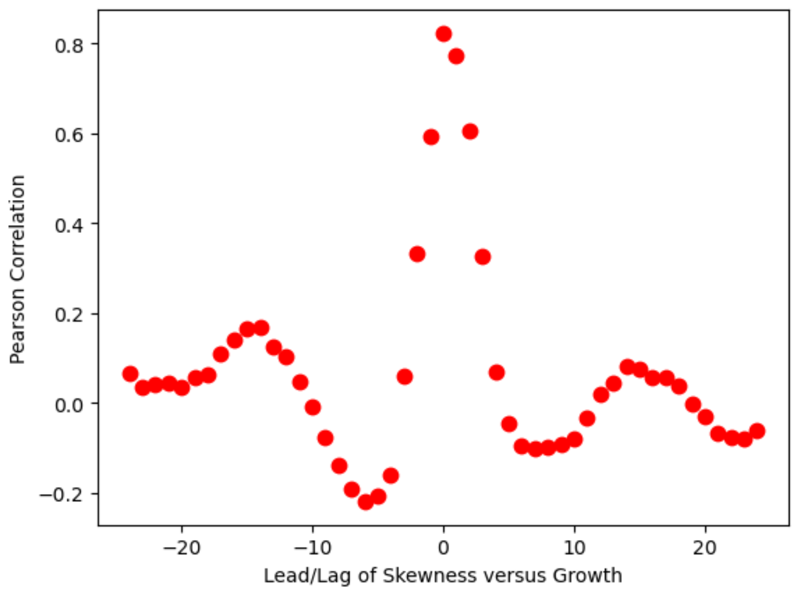

Note: The estimates are based on the cleaned sample with a sample period of 1984Q2 – 2021Q4.

B.5 Level of aggregation

This section provides the motivation for focusing on skewness across firm outcomes. If industry shocks were the main driver of business cycle fluctuations, it could be sufficient to focus on skewness across industries. For example, the analysis of Dew-Becker et al. (2021) focuses on industries. However, this can miss important variation if firm-level shocks are the origins of aggregate fluctuations or, more generally, if there is additional information at the firm level that is lost through aggregation to the industry level. The following results indicate that the association between skewness and aggregate fluctuations is strongest at the firm level.

Figure 20 shows the results of univariate regressions of aggregate sales growth on skewness estimated at different levels of aggregation. I compute skewness in the cross-section at the firm level and across 2-, 3-, 4-, and 5-digit NAICS industries. The left panel shows the coefficient estimates from regressions with standardized data. The black error bars indicate 95% confidence intervals. The coefficient rises from around 0.4 at the 2-digit NAICS level to around 0.8 at the firm level. The linear relation between cross-sectional skewness and aggregate activity is decreasing with the level of aggregation at which skewness is estimated. The right panel shows the corresponding R^2 values for the regressions. The linear fit is again decreasing with the level of aggregation. At the 2-digit NAICS level, the R^2 is below 0.2. At the firm level, the R^2 is almost 0.7.

To further show that skewness across firms is associated with aggregate fluctuations beyond industry effects, I compute skewness across demeaned firm growth rates and re-estimate the univariate regressions of aggregate growth on skewness. The demeaned growth rate of firm i in industry j at time t is $g^m_{i,j,t} = g_{i,j,t} - \bar{g}_{j,t}$, where $\bar{g}_{j,t}$ is the sales growth rate in industry j. I compute the demeaned growth rates for NAICS industries at the 2-, 3-, 4-, and 5-digit level. The regressions results are shown in Figure 21. The left panel shows the coefficient on skewness and the right panel show the R^2. Demeaning at the 2-digit level leaves the coefficient on skewness above 0.5 and the R^2 is above 0.25. Clearly, industry-level variation accounts for some of the association of skewness with the business cycle, but a significant share of variation remains to be explained by firm-level skewness.

With lower levels of aggregation, the coefficient falls and the linear fit decreases. At the 5-digit level, the coefficient is 0.2 and the R^2 is 0.04. However, industries at the 5-digit level are very narrow and there are only seven firms per quarter in the average 5-digit NAICS industry. The median number of firms per quarter and 5-digit NAICS industry is three. In contrast, 2-digit NAICS industries have 144 firms on average. The weaker results for increasingly narrow industry definitions need to be interpreted with this in mind.

To further support the idea that firm-level skewness is a relevant object of interest, I also estimate multivariate regressions with the different skewness measures from Figure 20 (not based on demeaned growth rates) simultaneously included as explanatory variables. Table 18 shows that including skewness across 2-digit NAICS industries adds no explanatory power beyond firm-level skewness and leaves the coefficient on firm-level skewness essentially unchanged (column 1). When including skewness estimates for all levels of aggregation, the R^2 increases only slightly and the coefficient on (standardized) firm-level skewness is much larger than on the other (standardized) skewness measures. Regressing aggregate sales growth on the difference between firm-level skewness and skewness across 2-digit NAICS industries results in a smaller but strongly significant coefficient estimate. The linear fit is considerably lower. The results are very similar for GDP growth as the dependent variable, except that the R^2 values are generally lower: Firm-level skewness is most closely associated with aggregate activity, even after controlling for skewness across industries at different levels of aggregation.

Overall, the results of this subsection suggest measuring skewness at the firm level yields the most informative measure for aggregate activity. A part of the skewness captured at the firm level is driven by skewness across industries. Controlling for industry-level shocks will therefore be important. Nevertheless, focusing only on the industry level may miss important variation, and in fact only explains a small share of the variability in aggregate activity.

Figure 20: Growth vs Skew by Level of Aggregation

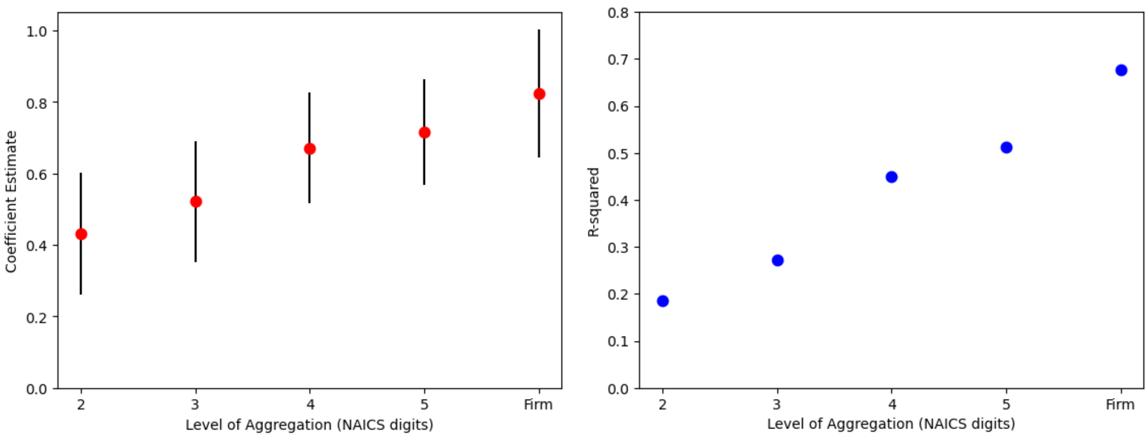

Note: The left figure plots coefficients from univariate regressions of aggregate sales growth on skewness, where skewness is computed for different levels of aggregation. The black bars show 95% confidence intervals. The right panel shows the corresponding R-squared value for each regression. The data is from the cleaned sample, covering the period 1984Q2 – 2021Q4. All regressions include a constant.

Figure 21: Growth vs Skew across Firms - Controlling for Industry Growth

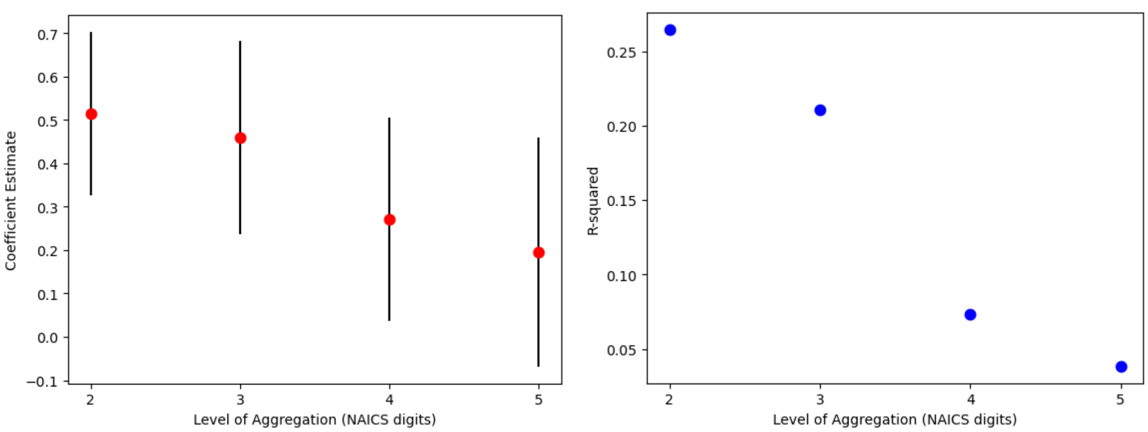

Note: The left figure plots coefficients from univariate regressions of aggregate sales growth on skewness across firm-level sales growth rates, controlling for industry growth at different levels of aggregation. The black bars show 95% confidence intervals. The right panel shows the corresponding R-squared value for each regression. The data is from the cleaned sample, covering the period 1984Q2 – 2021Q4. All regressions include a constant.

B.6 Skewness vs Dispersion

A large literature uses cross-sectional dispersion as a proxy for uncertainty in the economy and studies its business cycle implications (see Bloom (2014) for a summary and Bloom et al. (2018) for a recent application). Dispersion is commonly found to be countercyclical: The spread across firm outcomes increases in recessions and decreases during booms. Are the results on the procyclicality of skewness robust to controlling for time-varying dispersion?

I estimate regressions of aggregate activity on dispersion and skewness to answer this question. To be consistent with the skewness measurement, I compute dispersion as the difference between the

90th and the 10th percentile of the cross-sectional distribution: $D_t = Q_{0.9,t} - Q_{0.1,t}$. Skewness is estimated as the 90% Kelley skewness, and aggregate activity is measured by aggregate sales growth or GDP growth. Table 19 shows the results. All variables are standardized to make the interpretation of coefficients easier.

Dispersion itself explains 17% of the variation in aggregate sales and is significantly correlated with the business cycle (see column 1). The coefficient and the R^2 are lower than in an equivalent univariate regression of aggregate activity on skewness, see the leftmost column of Table 17. Including both dispersion and skewness simultaneously shows that dispersion is not significantly associated with business cycle fluctuations conditional on controlling for skewness (column 2). The coefficient estimate for dispersion is -0.04 compared to 0.84 for skewness. Compared to the univariate regression of aggregate activity on skewness (Table 17), including dispersion in the regression adds no explanatory power in terms of adjusted R^2 and does not affect the coefficient estimate on skewness. This result is robust to removing the Covid period (column 3) and is qualitatively identical when using GDP growth as the business cycle indicator (columns 4-6).

In contrast to a large literature that focuses on the cross-sectional dispersion of firm outcomes, these results suggest that cross-sectional skewness may be the more important object to understand business cycle fluctuations. These findings echo the work of Guvenen et al. (2014), who find that the association of cross-sectional dispersion in the income distribution with the business cycle is driven by skewness. The results are also in line with Figure 15, which shows that dispersion (measured as the standard deviation) is similar across recession and non-recession periods, whereas the mean and skewness are changing significantly.

Focusing on skewness instead of dispersion can have strong implications for theoretical work. In the uncertainty literature, an increase in cross-sectional dispersion has theoretically ambiguous effects on output (see Bloom (2014) for a discussion).[28] Firms may be risk-loving if they are able to contract in recessions and expand in booms to benefit from growth opportunities (Oi-Hartmann-Abel effect). Alternatively, the real options theory posits that firms adopt a wait-and-see approach if the future becomes more uncertain and they face adjustment costs, which leads to a contraction. In contrast to changes in uncertainty, skewness captures an increase in upside or downside risks, which can be expected to have unambiguous effects on firm decisions (all else equal). Supporting this intuition, Salgado et al. (2023) find that shocks that increase cross-sectional skewness have stronger contractionary effects than dispersion shocks.

In the context of real options theory, Bernanke (1983) establishes the 'bad news principle': The investment decision of a firm is based on the expected value of future bad outcomes, irrespective of the potential for good outcomes. In this sense, downside risk is the key driver of a decline in investment. Motivated by this view, Table 20 studies the relation between aggregate activity and different parts of the cross-sectional distribution. The business cycle indicator is aggregate sales growth. The left columns use independent variables computed from the distribution of percent sales growth rates, while

[28]The object of interest in the uncertainty literature is dispersion in the space of future possible outcomes. Since the distribution of expected future outcomes is difficult to measure, many papers use proxies such as cross-sectional dispersion in firm-level realizations. In a model with many firms, an increase in uncertainty for each individual firm manifests as a larger dispersion of shocks hitting the firms, and hence a larger dispersion across firm outcomes. In this sense, cross-sectional dispersion of firm outcomes is related to uncertainty about the future.

the right columns use variables based on the distribution of log differences in sales. The difference between the 90th percentile and the median (Q(0.9)–Q(0.5)) is positively associated with aggregate activity, while the difference between median and 10th percentile is negatively associated with sales growth. In contrast to the bad news principle, the upper half of the growth rate distribution is related to aggregate activity. The findings are similar across both types of dependent variables, and remain similar when considering individual quantiles instead of differences across quantiles. The median is generally most strongly associated with aggregate activity. The differences between coefficients on the 90th and the 10th percentile are small. The R^2 values are high across regressions and vary between 0.59 and 0.85. In summary, all parts of the cross-sectional distribution are associated with aggregate activity.

While all parts of the sales growth distribution are associated with aggregate activity, the different parts can exhibit different time series behaviour. Figure 22 shows the 90th percentile of (log) sales growth minus the median, and the median minus the 10th percentile.[29] Dispersion increases in recessions, but mostly because bad growth rate outcomes become more frequent. Most of the volatility of changes in the growth rate distribution is concentrated below the median, while the difference between the 90th and the 50th percentile stays relatively constant. The two notable exceptions are the dot-com boom in the early 2000s and the post-Covid recovery, during which some firms experienced very strong growth rate realizations. Overall, the co-movement of dispersion with aggregate activity is driven by skewness. While upside potential changes little during recessions, downturns are periods with larger potential for bad growth rate realizations.

Appendix C Details: Response of micro skewness to aggregate shocks

Local projection specifications and robustness. Table 21 describes the construction of all data entering the local projections, including the shocks. I use existing data or the authors' replication codes for all shock series. For the baseline specifications, micro skewness and aggregate sales growth are computed using the streak sample as described in Appendix A. Figure 23 reports the results from several robustness checks on the impulse responses of micro skewness to aggregate shocks. All alternative specifications yield very similar results, sometimes so close that the different impulse responses cannot be clearly distinguished from the plot because they coincide for the first three decimals. The black dashed line reports the baseline results from Figure 6. The blue lines are impulse responses from regressions using four lags instead of two. Regressions underlying the orange lines include lagged values of aggregate sales growth as controls. Impulse responses for skewness computed using the cleaned sample instead of the streak sample are shown in green.

Idiosyncratic shock series. Figure 24 plots the skewness in year-over-year sales growth rates (blue; left axis) against skewness in I/B/E/S sales growth rate forecast errors (red; right axis). Clearly,

[29] I use growth rates computed from first differences of log (real) sales for this figure since using percent growth rates yields a figure that is dominated by volatility in the upper tail. This is because percent growth rates are bounded below by -1 while they face no upper bound. Log growth rates make the distribution of possible growth rate outcomes more symmetric.

Figure 22: Upside vs Downside Risk

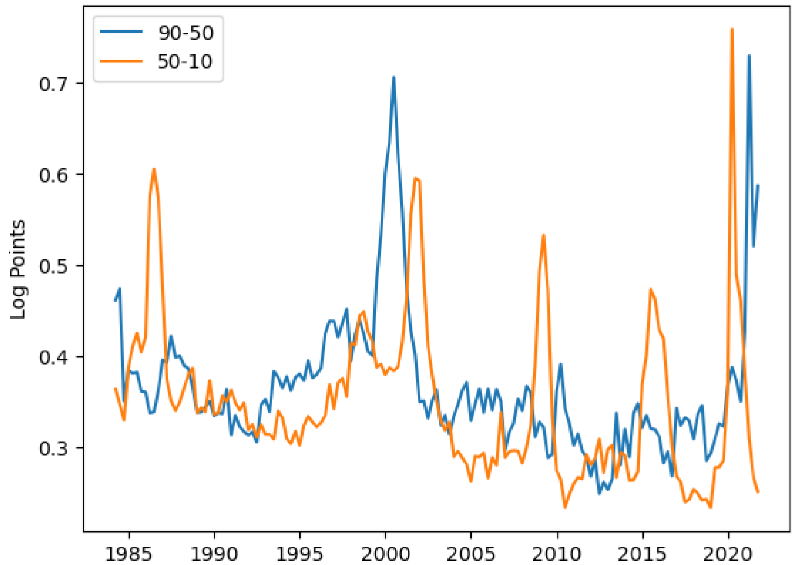

Note: Upside risk is the difference between the 90th and the 50th percentile. Downside risk is the difference between 50th and the 10th percentile. The quantiles are computed for the distribution of log growth rates. The data is from the cleaned sample, covering the period 1984Q2 – 2021Q4.

the two series are closely correlated throughout multiple downturns. It appears that skewness in forecast errors may be leading skewness in sales growth rates.

To test the predictability of the idiosyncratic shock series ξ_t, Table 22 reports the results from a regression of the shock series on eleven different macroeconomic time series (all lagged) that are used in the local projections. The adjusted R^2 is only 11%, suggesting that the shock series is widely unpredictable. This is true even though the set of predictors includes forward-looking variables such as Treasury yields, stock returns, and the excess bond premium.

Firm-level impulse responses. To build intuition for the firm-level impulse responses, Figures 25 - 30 show the firm-level impulse responses for ExxonMobil, McDonald's, Marriott Hotels, Caterpillar, IBM, and Walt Disney. These six companies represent different sectors of the US economy and different degrees of cyclicality. This is reflected in their impulse responses. ExxonMobil generally shows strong responses to aggregate shocks and is particularly exposed to the oil supply shock. In contrast, the sales growth of McDonald's is much less volatile and shows an insignificant response to oil supply shocks. Marriott's sales respond weakly to monetary shocks but significantly to uncertainty shocks. Caterpillar is a very cyclical company and its large impulse responses reflect this. IBM responds more to monetary and TFP shocks than to oil or uncertainty shocks. Lastly, Disney's sales show weak responses to all aggregate shocks. As an additional validation exercise, I compare the impulse responses of different companies to a one-standard deviation oil supply shock, which is the only shock clearly identifiable as originating from a particular sector of the economy. Figure 31 shows the impulse responses of four big oil corporations in blue and compares their impulse responses to those of IBM, McDonald's, and Disney. Energy corporations show significantly larger

Figure 23: Robustness: Comovement of growth and skew after aggregate shocks

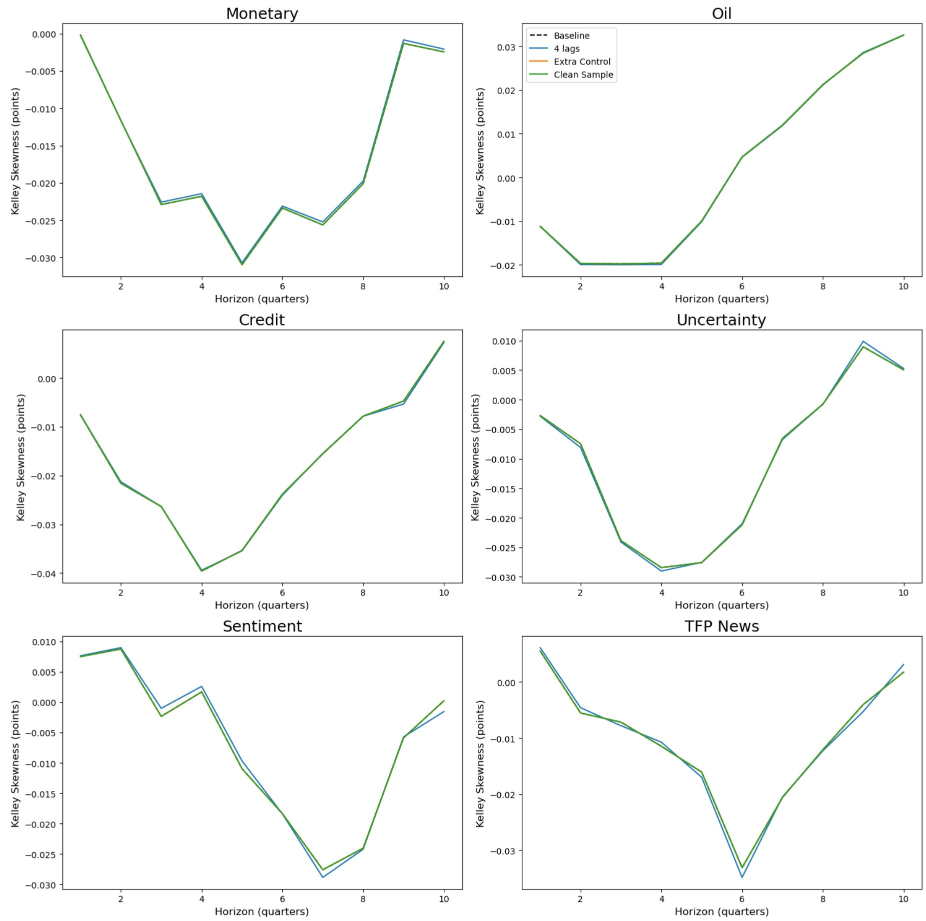

Note: Red shaded areas are 90% confidence bands for the results using the Covid sample, based on Newey-West standard errors. Robustness check for the Covid sample has not been conducted for the TFP shocks due to data availability. Shock magnitudes are normalized to be one standard deviation. The signs of the TFP shocks are reversed to be contractionary.

sales responses to oil supply shocks. Taken together, the firm level IRFs appear reasonable in terms of their qualitative features.

Robustness of firm-level IRFs. Figure 32 shows that the bottom-up impulse responses of micro skewness are robust to different specifications. Including four lags of all controls (blue lines) or adding lagged QoQ stock returns (orange lines) yields very similar results. The green lines show the impulse responses of skewness when computing firm-level IRFs for year-over-year stock returns as the dependent variable. The response of skewness in stock returns is similar to the response of skewness in sales growth rates, though the magnitudes can be larger. The response of return skewness to the TFP news shock poses an exception from this rule.

Robustness of large vs small firm IRFs. Figure 33 repeats the analysis of Figure 10 but excludes the bottom 10% of the firm size distribution from the analysis. Similarly, Figure 34 compares the responses of the top 30% against the bottom 70%. The results confirm the main points of the

Figure 24: Skewness in sales growth rates and I/B/E/S forecast errors

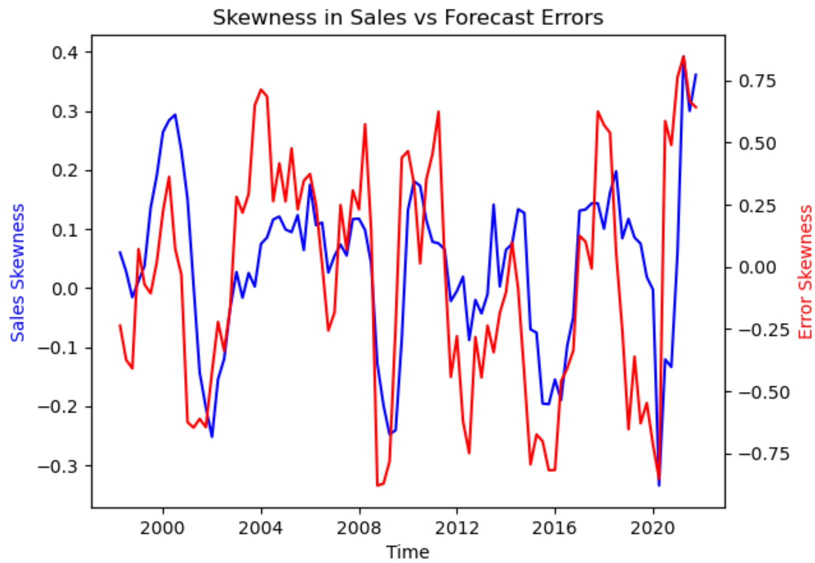

baseline results: The impulses responses of the largest firms in the Compustat sample are not less skewed than the responses of the smaller firms. Because the largest firms account for the vast majority of aggregate sales, they account for most of the decline in aggregate sales following an aggregate shock.

Because small firms feature more volatile sales growth time series, their impulse response estimates may be more volatile as well and the bottom-up skewness response for these firms may be less reliable. I argue that this concern does not affect the main results in this paper. First, the key insight of the analysis is that the response of the largest firms is also skewed, which holds independently of the result for the smaller firms. Second, many of the smaller firms are still very large in the context of the overall US firm size distribution, making the estimation of firm-level IRFs an arguably reasonable approach, especially if one is willing to accept the firm-level IRFs for the very largest firms. Third, I conduct a robustness check confirming my results without relying on the estimation of firm-level IRFs. I separately construct a micro skewness index for the largest firms (top 10%) and the rest of the firms. Firms are grouped into the two size bins each quarter and firm size is measured as the average real sales over the previous four quarters. Given the two indexes, I estimate their response to an aggregate shock using the local projection specification from equation 5. Figure 35 shows the results. Skewness for the largest firms declines significantly following a contractionary aggregate shock, while skewness for the rest of firms does not show a clear decline for any of the shocks. Some of the estimates for smaller firms' skewness even appear to indicate an increase in the skewness index as opposed to a decrease, though the estimates are often insignificant and generally smaller in magnitude than for skewness of the largest firms.

Figure 36 shows the contributions of different growth rate (IRF) and size bins to the decline in aggregate sales following an aggregate shock. The methodology is outlined in Section 4.5. Across the six different shocks, the following results hold true: 1) The majority of the sales decline is driven by firms with the poorest growth rate outcomes. 2) Within each bin, the majority of the response is explained by the very largest firms in the sample (top 10%). This is true even among the worst performers. 3) These large and responsive firms account for roughly one third of the overall economy's response to a shock, even though they only make up about 1.5% of all firms in the sample. The aggregate growth response is least concentrated for the uncertainty shock, where large and responsive firms account for 23% of the aggregate response, and is most concentrated for the oil shock (38%).

Construction of firm characteristics to explain IRF heterogeneity. The quarterly firm characteristics are constructed as follows. Real variables are obtained by deflating with the GDP deflator.

- Age: Log years since inception, defined as the minimum of the start of trading data in CRSP (`begdat`), the first data point in Compustat, and the date of incorporation recorded in Worldscope Fundamentals
- Size: Log real sales (`saleq`)
- Leverage: Current liabilities (`dlcq`) plus long-term liabilities (`dlttq`), divided by total assets (`atq`)
- Liquidity: Cash and short-term investments (`cheq`), divided by total assets (`atq`)
- Dividend payer: Equals one if dividends (`dvpq`) are larger than zero, equals zero otherwise.
- Fixed assets: Cost of fixed property (`ppentq`), divided by total assets (`atq`)
- Short-term debt: Short-term liabilities (`dlcq`), divided by one-quarter lagged total assets (`atq`)
- Long-term debt: Long-term liabilities (`dlttq`), divided by one-quarter lagged total assets (`atq`)
- Sales / Assets: Nominal sales (`saleq`), divided by total assets (`atq`)
- Profitability (ROA): Income before extraordinary expenses (`ibq`), divided by total assets (`atq`)
- R&D: Research and development expenses (`xrdq`), divided by lagged total assets (`atq`)
- Inventory: Total inventories (`invtq`), divided by lagged total assets (`atq`)
- Recession indicator: Equals one if the quarter in the firm's sample is an NBER recession quarter
- Aggregate growth: Year-over-year growth in aggregate real sales (see equation 14)
- Industry growth: Year-over-year growth in industry-level real sales, with industries defined at 2-digit NAICS level
- Firm growth: Year-over-year growth in firm-level real sales
- Pre-recession firm growth: Rank in the distribution of year-over-year firm-level sales growth rates if the quarter is at most three quarters before the start of a NBER recession, missing value otherwise

All firm characteristics are averaged over time to obtain a single observation per characteristic and firm. Since the distribution of shocks may be unequal within each firm's sample, each quarter is weighted by a share indicating the size of that quarter's shock: $\omega_t = |\text{shock}_t|/(\sum_\tau |\text{shock}_\tau|)$. This is to reflect that the values of firm characteristics prevailing when large shocks hit matter more for explaining differences across impulse responses than firm characteristics prevailing when shocks are small. The results are similar when taking simple averages and available upon request.

Figure 25: ExxonMobil: Sales Growth Responses to Aggregate Shocks

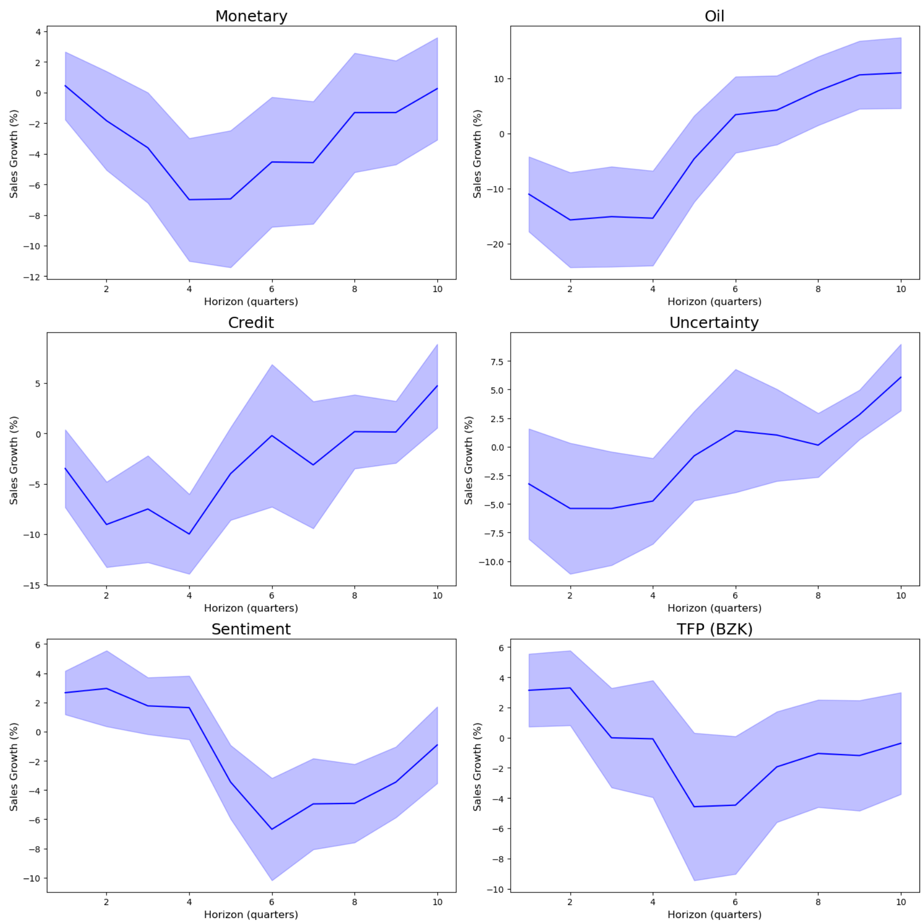

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Figure 26: McDonald's: Sales Growth Responses to Aggregate Shocks

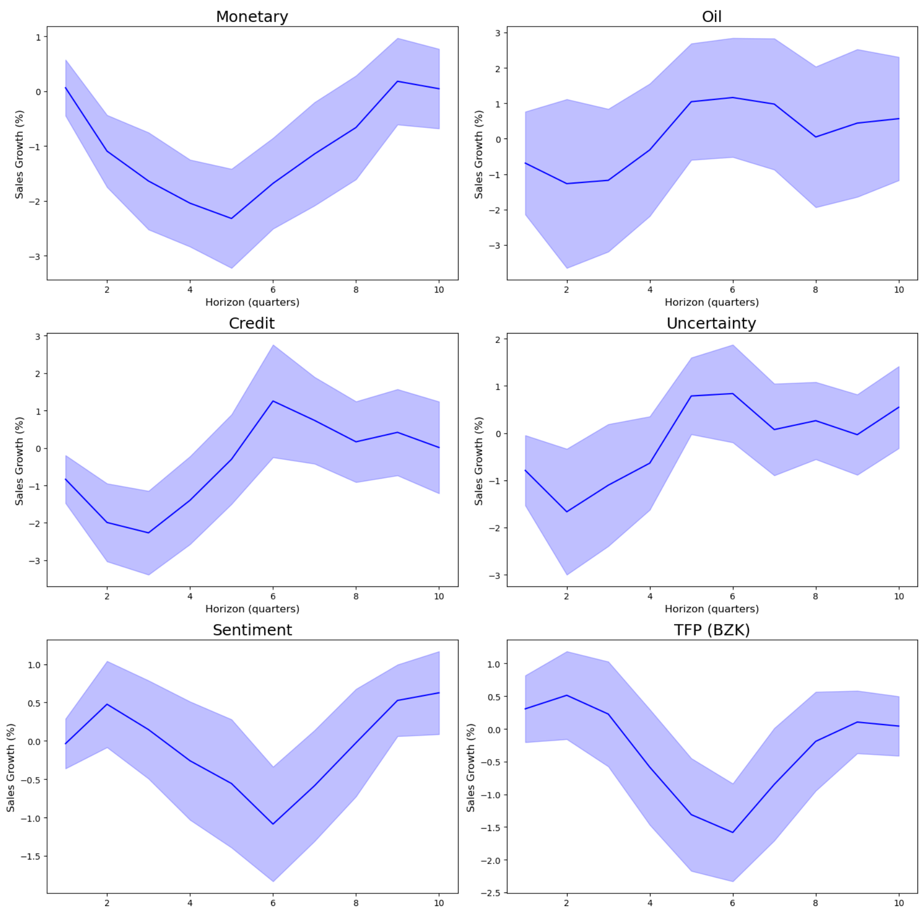

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Table 12: Correlations of skewness with the business cycle - Outlier cleaning

	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
	Panel a: Correlations with aggregate sales growth				
Top/bottom 0.1%					
Full sample	0.64	0.41	-0.06	-0.14	0.05
Pre-1984	0.49	0.27	-0.16	-0.05	0.03
Post-1984	0.86	0.78	0.42	0.08	0.62
Top/bottom 1%					
Full sample	0.63	0.42	-0.04	-0.10	0.39
Pre-1984	0.50	0.30	-0.11	0.00	0.27
Post-1984	0.85	0.79	0.47	-0.51	0.82
Top/bottom 2%					
Full sample	0.63	0.43	-0.05	-0.11	0.48
Pre-1984	0.51	0.31	-0.08	0.06	0.37
Post-1984	0.86	0.80	0.33	-0.37	0.83
Top/bottom 0.1% each quarter					
Full sample	0.64	0.41	-0.07	-0.18	0.05
Pre-1984	0.49	0.28	-0.17	-0.38	0.02
Post-1984	0.86	0.78	0.46	-0.02	0.53
Top/bottom 1% each quarter					
Full sample	0.65	0.47	0.05	-0.08	0.39
Pre-1984	0.49	0.33	-0.08	-0.22	0.29
Post-1984	0.86	0.80	0.64	0.23	0.82
Top/bottom 2% each quarter					
Full sample	0.66	0.54	0.20	0.12	0.52
Pre-1984	0.52	0.39	0.06	-0.06	0.39
Post-1984	0.86	0.82	0.72	0.65	0.84
	Panel b: Correlations with GDP growth				
Top/bottom 0.1%					
Full sample	0.52	0.32	-0.10	-0.24	0.02
Pre-1984	0.50	0.30	-0.09	-0.19	0.15
Post-1984	0.69	0.68	0.51	-0.36	0.56
Top/bottom 1%					
Full sample	0.51	0.32	-0.10	-0.09	0.32
Pre-1984	0.51	0.31	-0.10	-0.01	0.36
Post-1984	0.67	0.66	0.47	-0.56	0.68
Top/bottom 2%					
Full sample	0.51	0.32	-0.11	-0.08	0.40
Pre-1984	0.51	0.31	-0.10	0.09	0.42
Post-1984	0.66	0.66	0.34	-0.41	0.69
Top/bottom 0.1% each quarter					
Full sample	0.54	0.34	-0.06	-0.20	0.03
Pre-1984	0.54	0.34	0.01	-0.16	0.20
Post-1984	0.69	0.68	0.47	-0.12	0.41
Top/bottom 1% each quarter					
Full sample	0.55	0.39	0.03	-0.07	0.33
Pre-1984	0.55	0.39	0.05	0.02	0.39
Post-1984	0.69	0.68	0.61	0.22	0.71
Top/bottom 2% each quarter					
Full sample	0.55	0.44	0.15	0.08	0.43
Pre-1984	0.55	0.44	0.15	0.10	0.45
Post-1984	0.69	0.68	0.65	0.56	0.72

The first three columns compute skewness using the Kelley measure (SK_2), with different choices for α. The last column computes the Groeneveld & Meeden (1984) skewness coefficient (SK_3). The pre-1984 sample is 1962Q1-1983Q4. The post-1984 sample is 1984Q1-2022Q3.

Table 13: Correlations of skewness with the business cycle - Cutoff rules

	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
	Panel a: Correlations with aggregate sales growth				
100% cutoff					
Full sample	0.64	0.48	0.00	0.05	0.58
Pre-1984	0.60	0.51	0.22	0.18	0.56
Post-1984	0.74	0.57	-0.63	-0.35	0.72
50% cutoff					
Full sample	0.37	0.09	-0.63	-0.50	0.28
Pre-1984	0.45	0.20	-0.46	-0.37	0.34
Post-1984	0.34	-0.13	-0.79	-0.67	0.24
25% cutoff					
Full sample	-0.27	-0.55	-0.76	-0.72	-0.35
Pre-1984	-0.07	-0.38	-0.71	-0.69	-0.17
Post-1984	-0.58	-0.75	-0.81	-0.78	-0.64
	Panel b: Correlations with GDP growth				
100% cutoff					
Full sample	0.53	0.41	0.08	0.11	0.51
Pre-1984	0.56	0.44	0.17	0.16	0.53
Post-1984	0.47	0.35	-0.51	-0.31	0.50
50% cutoff					
Full sample	0.36	0.16	-0.53	-0.52	0.30
Pre-1984	0.50	0.33	-0.37	-0.41	0.43
Post-1984	0.09	-0.30	-0.72	-0.66	0.05
25% cutoff					
Full sample	-0.11	-0.37	-0.62	-0.65	-0.19
Pre-1984	0.14	-0.13	-0.51	-0.61	0.05
Post-1984	-0.59	-0.68	-0.71	-0.67	-0.59

The first three columns compute skewness using the Kelley measure (SK_2), with different choices for α. The last column computes the Groeneveld & Meeden (1984) skewness coefficient (SK_3). The pre-1984 sample is 1962Q1-1983Q4. The post-1984 sample is 1984Q1-2022Q3.

Table 14: Descriptive statistics - Samples for correlation estimation

	Mean Growth	Minimum Growth	Maximum Growth	Observations
Full sample				
Full sample	134%	-100%	24,353,693%	1,145,568
Pre-1984	34%	-100%	633,842%	199,064
Post-1984	155%	-100%	24,353,693%	946,504
Top/bottom 0.1%				
Full sample	36%	-99%	12,182%	1,143,276
Pre-1984	14%	-99%	12,160%	198,978
Post-1984	41%	-99%	12,182%	944,298
Top/bottom 1%				
Full sample	16%	-82%	680%	1,122,656
Pre-1984	9%	-82%	678%	197,857
Post-1984	17%	-82%	680%	924,799
Top/bottom 2%				
Full sample	12%	-69%	303%	1,099,744
Pre-1984	8%	-69%	303%	196,078
Post-1984	13%	-69%	303%	903,666
Top/bottom 0.1% each quarter				
Full sample	39%	-100%	49,960%	1,143,031
Pre-1984	13%	-100%	21,436%	198,594
Post-1984	44%	-100%	49,960%	944,437
Top/bottom 1% each quarter				
Full sample	17%	-94%	1,639%	1,122,411
Pre-1984	8%	-86%	898%	194,995
Post-1984	18%	-94%	1,639%	927,416
Top/bottom 2% each quarter				
Full sample	13%	-86%	803%	1,099,490
Pre-1984	7%	-76%	370%	191,009
Post-1984	14%	-86%	803%	908,481
100% cutoff				
Full sample	4%	-100%	100%	1,078,962
Pre-1984	5%	-100%	100%	194,198
Post-1984	4%	-100%	100%	884,764
50% cutoff				
Full sample	3%	-50%	50%	963,423
Pre-1984	4%	-50%	50%	182,132
Post-1984	3%	-50%	50%	781,291
25% cutoff				
Full sample	2%	-25%	25%	774,776
Pre-1984	3%	-25%	25%	153,089
Post-1984	2%	-25%	25%	621,687

Table 15: Correlations of skewness with the business cycle - Additional considerations

	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
	Panel a: Correlations with aggregate sales growth				
Log differences					
Full sample	0.71	0.63	0.29	0.13	-0.00
Pre-1984	0.56	0.46	0.20	0.13	0.10
Post-1984	0.83	0.78	0.53	0.20	0.22
Log differences + Top/bottom 2% each quarter					
Full sample	0.72	0.69	0.54	0.56	0.52
Pre-1984	0.57	0.53	0.37	0.38	0.39
Post-1984	0.84	0.80	0.74	0.72	0.84
Weighted rates					
Full sample	0.90	0.93	0.95	0.56	-0.00
Pre-1984	0.89	0.92	0.93	0.57	0.10
Post-1984	0.91	0.93	0.95	0.56	0.22
Weighted rates + Top/bottom 2% each quarter					
Full sample	0.89	0.92	0.94	0.90	0.01
Pre-1984	0.88	0.91	0.92	0.90	0.12
Post-1984	0.91	0.93	0.95	0.93	0.24
	Panel b: Correlations with GDP growth				
Log differences					
Full sample	0.60	0.54	0.29	0.16	-0.09
Pre-1984	0.59	0.48	0.28	0.18	0.03
Post-1984	0.63	0.61	0.50	0.18	0.11
Log differences + Top/bottom 2% each quarter					
Full sample	0.61	0.58	0.46	0.43	0.43
Pre-1984	0.61	0.56	0.40	0.35	0.45
Post-1984	0.64	0.61	0.59	0.56	0.72
Weighted rates					
Full sample	0.66	0.64	0.53	0.28	-0.09
Pre-1984	0.66	0.61	0.49	0.33	0.03
Post-1984	0.69	0.66	0.57	0.26	0.11
Weighted rates + Top/bottom 2% each quarter					
Full sample	0.68	0.67	0.62	0.55	-0.06
Pre-1984	0.70	0.68	0.59	0.55	0.09
Post-1984	0.69	0.68	0.64	0.60	0.13

The first three columns compute skewness using the Kelley measure (SK_2), with different choices for α. The last column computes the Groeneveld & Meeden (1984) skewness coefficient (SK_3). The pre-1984 sample is 1962Q1-1983Q4. The post-1984 sample is 1984Q1-2022Q3.

Table 16: Correlations of skewness with the business cycle - Full vs cleaned samples

	Panel a: Correlations with aggregate sales growth				
	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
Full sample	0.64	0.41	-0.07	-0.04	-0.00
Cleaned sample	0.82	0.77	0.46	0.13	0.74
Streak sample	0.76	0.73	0.53	0.41	0.77
Long sample					
	Panel b: Correlations with GDP growth				
	90% Kelley	95% Kelley	99% Kelley	Third moment	GM Skewness
Full sample	0.64	0.32	-0.11	-0.14	-0.09
Cleaned sample	0.82	0.57	0.32	0.01	0.51
Streak sample	0.76	0.62	0.54	0.29	0.64
Long sample					
	Panel c: Sample characteristics				
	Mean Growth	Min. Growth	Max. Growth	Observations	
Full sample	134%	-100%	24,353,693%	1,145,568	
Cleaned sample	13%	-95%	1,782%	499,249	
Streak sample	7%	-95%	1,283%	151,701	
Long sample					
	Panel d: Sample characteristics				
	Sales over Assets	Lagged real sales	Total real assets	Acquisitions over assets	
Full sample	0.33	479,352	4,041,760	1.2%	
Cleaned sample	0.31	505,409	2,033,001	0.2%	
Streak sample	0.32	688,856	2,891,703	0.1%	
Long sample					

In panels a and b, the first three columns compute skewness using the Kelley measure (SK_2), with different choices for α. The last column computes the Groeneveld & Meeden (1984) skewness coefficient (SK_3). In panels c and d, all values are averages over all firm-quarter observations in the respective samples. Some statistics in panel d are subject to necessary data cleaning. For computing sales over assets, I remove all observations with negative or zero assets. For acquisitions over assets, I remove all observations with negative or zero assets and all observations with negative acquisitions. Acquisitions over assets are computed for the current and the three preceding quarters. Real sales and total real assets are in thousands of 2015 US dollars, using the GDP deflator. The full sample covers 1962Q1 – 2022Q3. The cleaned, streak, and long sample each cover 1983Q3 – 2021Q4.

Table 17: Regressions controlling for lagged growth

	Sales Growth			*GDP Growth*		
$Skew_t$	**0.82**	**0.90**	**0.46**	**0.47**	**0.55**	0.22
	(0.09)	(0.10)	(0.10)	(0.12)	(0.18)	(0.12)
$Skew_{t-1}$		0.04			-0.02	
		(0.07)			(0.10)	
$Skew_{t-2}$		**-0.19**			-0.18	
		(0.08)			(0.10)	
$Skew_{t-3}$		0.00			0.07	
		(0.08)			(0.10)	
$Skew_{t-4}$		-0.02			-0.00	
		(0.08)			(0.10)	
$Growth_{t-1}$			**0.70**			**0.50**
			(0.14)			(0.19)
$Growth_{t-2}$			**-0.27**			0.05
			(0.08)			(0.07)
$Growth_{t-3}$			-0.01			0.22
			(0.09)			(0.17)
$Growth_{t-4}$			-0.04			-0.30
			(0.06)			(0.16)
$Constant$	**-0.78**	**-0.63**	**-0.39**	0.18	0.25	0.14
	(0.20)	(0.21)	(0.14)	(0.47)	(0.29)	(0.17)
R^2	0.68	0.69	0.79	0.35	0.37	0.59

Skewness is estimated using the 90% Kelley measure and the cleaned sample. All data is standardized. The sample period is 1984Q2 – 2021Q4. Standard errors are Newey-West. Bold font indicates statistical significance at the 5% level. R^2 is adjusted for the number of regressors.

Table 18: Growth vs Skewness at different levels of aggregation

	Sales Growth			*GDP Growth*		
	(1)	(2)	(3)	(4)	(5)	(6)
Firm	**0.82**	**0.72**		**0.66**	**0.65**	
	(0.10)	(0.14)		(0.15)	(0.20)	
NAICS 2	0.01	**0.21**		-0.12	0.32	
	(0.06)	(0.10)		(0.11)	(0.18)	
NAICS 3		0.08			-0.18	
		(0.11)			(0.16)	
NAICS 4		**-0.23**			-0.21	
		(0.09)			(0.12)	
NAICS 5		0.09			-0.03	
		(0.05)			(0.12)	
Firm - NAICS 2			**0.40**			**0.39**
			(0.11)			(0.10)
Constant	**-0.77**	**-0.68**	-0.01	0.13	0.14	**0.62**
	(0.21)	(0.24)	(0.24)	(0.32)	(0.34)	(0.22)
R^2	0.67	0.71	0.15	0.35	0.38	0.14

Skewness is estimated using the 90% Kelley measure and the cleaned sample. All data is standardized. The sample period is 1984Q2 – 2021Q4. Standard errors are Newey-West. Bold font indicates statistical significance at the 5% level. R^2 is adjusted for the number of regressors.

Table 19: Skewness vs Dispersion

	Sales Growth			*GDP Growth*		
	(1)	(2)	(3)	(4)	(5)	(6)
Dispersion	**0.42**	-0.04	-0.03	**0.44**	0.17	0.22
	(0.09)	(0.09)	(0.11)	(0.13)	(0.16)	(0.14)
Skewness		**0.84**	**0.86**		**0.50**	**0.34**
		(0.12)	(0.14)		(0.19)	(0.15)
Constant	**-1.31**	**-0.63**	-0.69	**-0.81**	-0.41	-0.31
	(0.45)	(0.29)	(0.38)	(0.60)	(0.57)	(0.61)
R^2	0.17	0.68	0.64	0.19	0.37	0.30

Dispersion is measured as the difference between the 90th and the 10th percentile of sales growth rate outcomes in a given quarter. Skewness is measured as the 90% Kelley skewness. All variables are standardized. All regressions use the cleaned sample. The sample period is 1984Q2–2021Q4, except for columns (3) and (6), which remove the Covid period. The sample period is then 1984Q2–2019Q4. R-squared values are adjusted for the number of regressors. All standard errors are Newey-West. Bold font indicates statistical significance at the 5% level.

Table 20: Aggregate Activity vs Cross-Sectional Quantiles

	Growth Rates			*Log Differences*		
	(1)	(2)	(3)	(4)	(5)	(6)
$Q_{0.9}$–$Q_{0.5}$	**0.28**			**0.45**		
	(0.04)			(0.05)		
$Q_{0.5}$–$Q_{0.1}$	**-0.55**			**-0.42**		
	(0.12)			(0.07)		
$Q_{0.9}$		**0.11**	**0.18**		**0.20**	**0.30**
		(0.03)	(0.03)		(0.04)	(0.03)
$Q_{0.5}$		**0.52**			**0.44**	
		(0.16)			(0.14)	
$Q_{0.1}$		**0.22**	**0.40**		**0.16**	**0.26**
		(0.06)	(0.05)		(0.03)	(0.03)
Constant	**0.07**	0.03	**0.06**	0.02	-0.00	-0.00
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
R^2	0.59	0.83	0.80	0.67	0.85	0.83

The left columns use year-over-year real sales growth as the dependent variable, while the right columns approximate growth rates using year-over-year log differences of real sales. All regressions use the cleaned sample. The sample period is 1984Q2–2021Q4. R-squared values are adjusted for the number of regressors. All standard errors are Newey-West. Bold font indicates statistical significance at the 5% level.

Table 21: Data for local projections

Variable	Transformation from raw data	Data source
Real GDP	Log level	FRED (`GDPC1`)
GDP Deflator	Log level	FRED (`GDPDEF`)
Real oil price	Quarterly average of monthly data, deflated	FRED (`WPU0561` & `GDPDEF`)
GDP per capita	Real GDP ('rgdp') per population ('civpop')	Ramey (2016) TFP data
Labor productivity	Real GDP ('rgdp') per hours worked ('tothours')	Ramey (2016) TFP data
Shadow rate	Quarterly average of monthly data	Atlanta Fed*
Stock prices	Shiller stock prices divided by GDP deflator	Ramey (2016) & FRED
Stock prices per capita	Stock prices per population ('civpop')	Ramey (2016)
VXO	Quarterly average of daily data	FRED (`VXOCLS`)
Uncertainty Index	Log level of Jurado et al. (2015) index	Lagerborg et al. (2023)
Consumer Expectations	Log level	Lagerborg et al. (2023)
Monetary shock	Quarterly sum of monthly data	Fed Board**
Oil shock	Quarterly average of monthly data	Baumeister***
Credit shock	Quarterly average of monthly data	Favara et al. (2016)†
Uncertainty shock	Quarterly average of monthly 'maxG' shock	Ludvigson et al. (2021)
Sentiment shock	Quarterly average of monthly data	Lagerborg et al. (2023)‡
TFP News Shock	Level of Ben Zeev & Khan (2015) shock	Ramey (2016) TFP data
Micro skewness	Own construction based on streak sample	-
Sales growth	Own construction based on streak sample	-

(*) The shadow rate data is available at: `https://www.atlantafed.org/cqer/research/wu-xia-shadow-federal-funds-rate`. (**) The Bu et al. (2021) monetary shocks are from `https://www.federalreserve.gov/econres/feds/a-unified-measure-of-fed-monetary-policy-shocks.htm`. (***) The updated shock series of Baumeister & Hamilton (2019) is available at `https://sites.google.com/site/cjsbaumeister/datasets?authuser=0`. (†) The credit shock is constructed by replicating the eight-variable Gilchrist & Zakrajšek (2012) recursive VAR (Section IV.B) for the sample 1973Q1 – 2019Q4. (‡) The sentiment shock is extracted from the author's proxy SVAR estimated from 1965:1 until 2018:11. The instrument is the number of fatalities (larger than or equal to 7) excluding the 2017 Las Vegas shooting, which Lagerborg et al. (2023) show to be stronger for this sample period than their baseline instrument.

Table 22: Regression of idiosyncratic shocks on macro variables

Variable	Coef.	SE	Variable	Coef.	SE
Sales	0.05**	0.02	FFR	-0.06	0.04
Skew	-0.95	1.13	10yr TY	0.10**	0.05
GDP	-0.06	0.05	EBP	-0.11	0.08
Inflation	0.14	0.09	Stocks	-0.00	0.00
Productivity	0.04	0.03	ICE	0.00	0.00
Oil	-0.00	0.00	Constant	-0.50**	0.21

Standard errors are Newey-West. Number of observations is 100 quarters. Adjusted R^2: 11.3%. (**) indicates statistical significance at the 5% level. All predictors are lagged by one quarter. All variables in yoy growth rates except 10yr TY, FFR, and EBP, which are in levels.

Figure 27: Marriott International: Sales Growth Responses to Aggregate Shocks

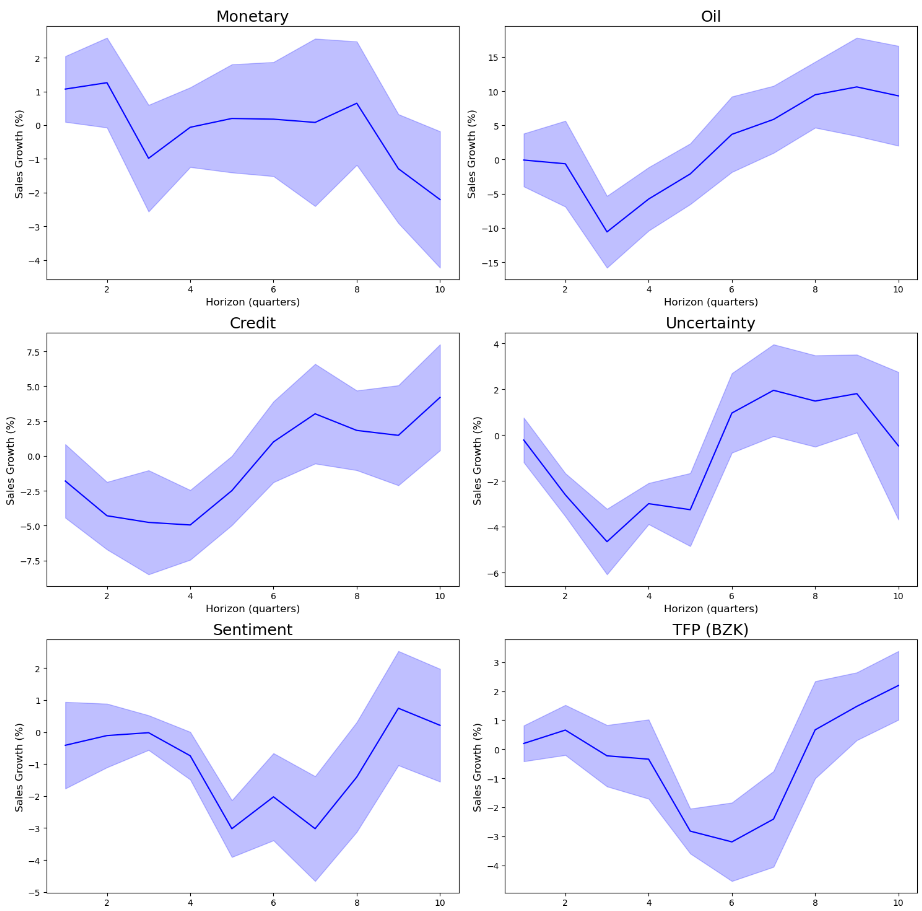

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Figure 28: Caterpillar: Sales Growth Responses to Aggregate Shocks

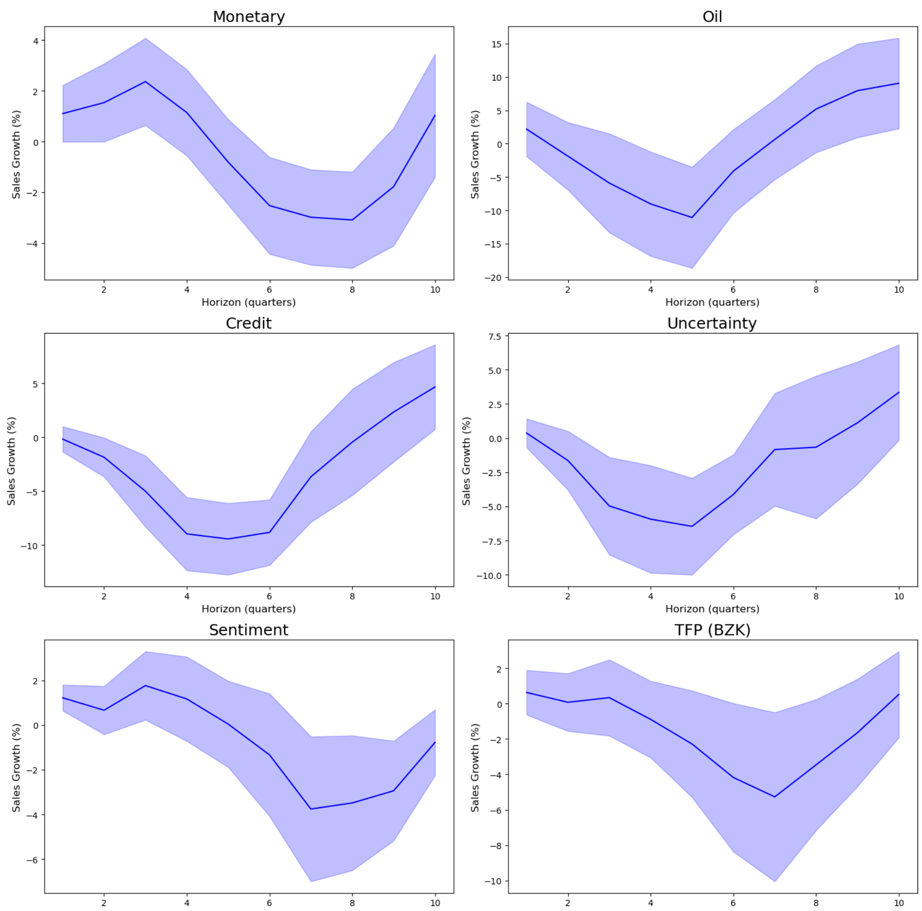

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Figure 29: IBM: Sales Growth Responses to Aggregate Shocks

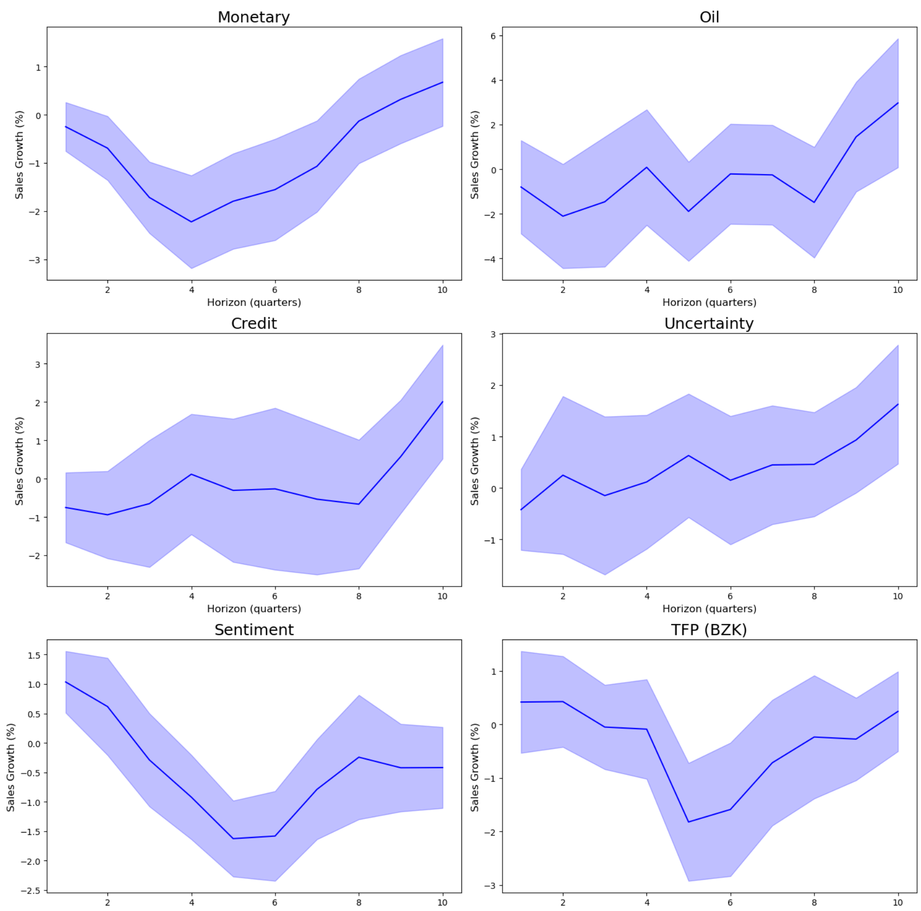

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Figure 30: Walt Disney: Sales Growth Responses to Aggregate Shocks

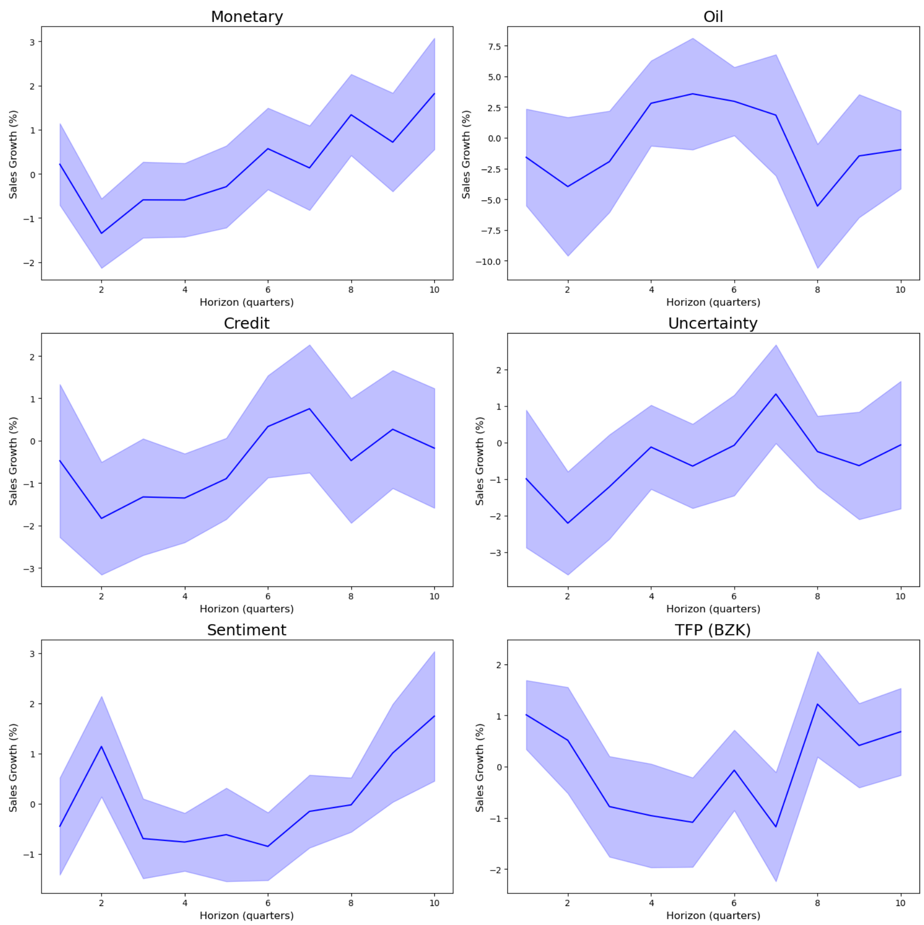

Note: Blue shaded areas are 90% confidence bands, based on Newey-West standard errors. Shock magnitudes are normalized to be one standard deviation. The signs of the sentiment and TFP shocks are reversed to be contractionary.

Figure 31: Sales Growth Responses to Oil Supply Shock

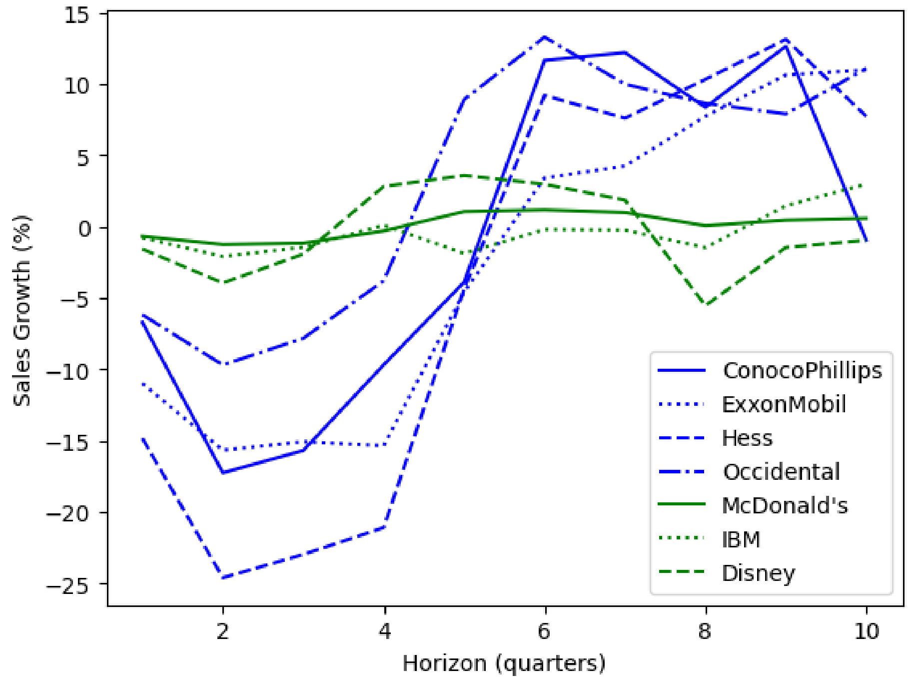

Note: Blue impulse responses refer to oil companies, while green are non-oil companies. Shock magnitudes are normalized to be one standard deviation.

Figure 32: Robustness: Comovement of bottom-up skewness and growth after aggregate shocks

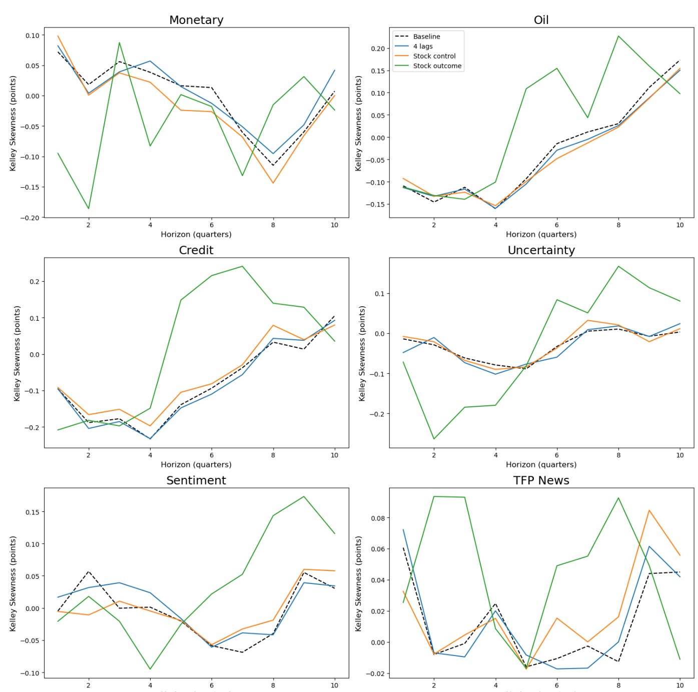

Figure 33: Top 10% vs middle 80% firms: Bottom-up skewness and growth responses

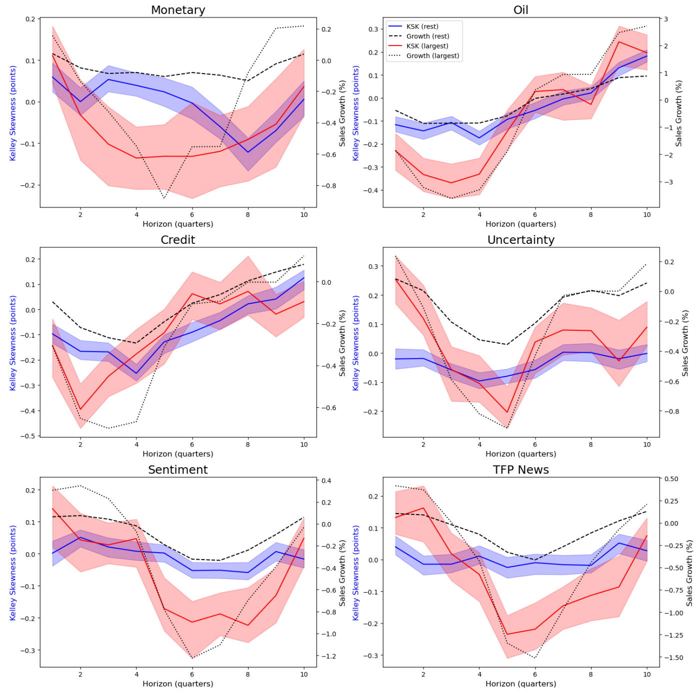

Figure 34: Top 30% vs bottom 70% firms: Bottom-up skewness and growth responses

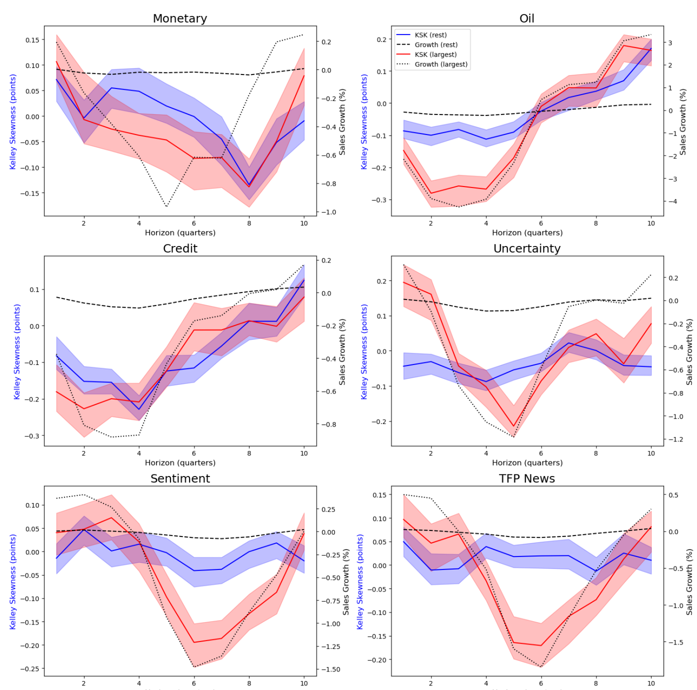

Figure 35: Top 10% vs bottom 90% of firms: Skewness index responses

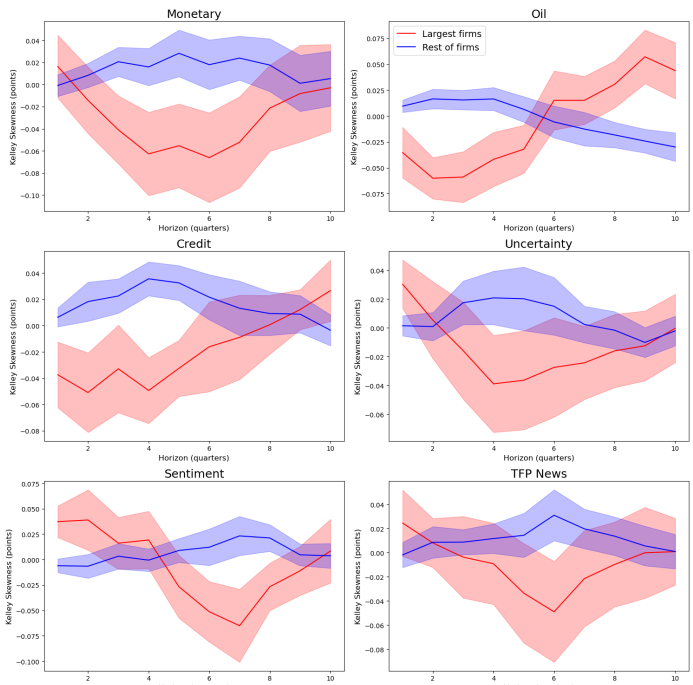

Note: The red lines show the impulse response of the skewness index for the top 10% of firms by size (real sales) and the blue lines show the impulse responses of the skewness index for the rest of firms. The 90% confidence intervals are computed using Newey-West standard errors.

Figure 36: Contributions of growth and size bins to aggregate growth decline - All shocks

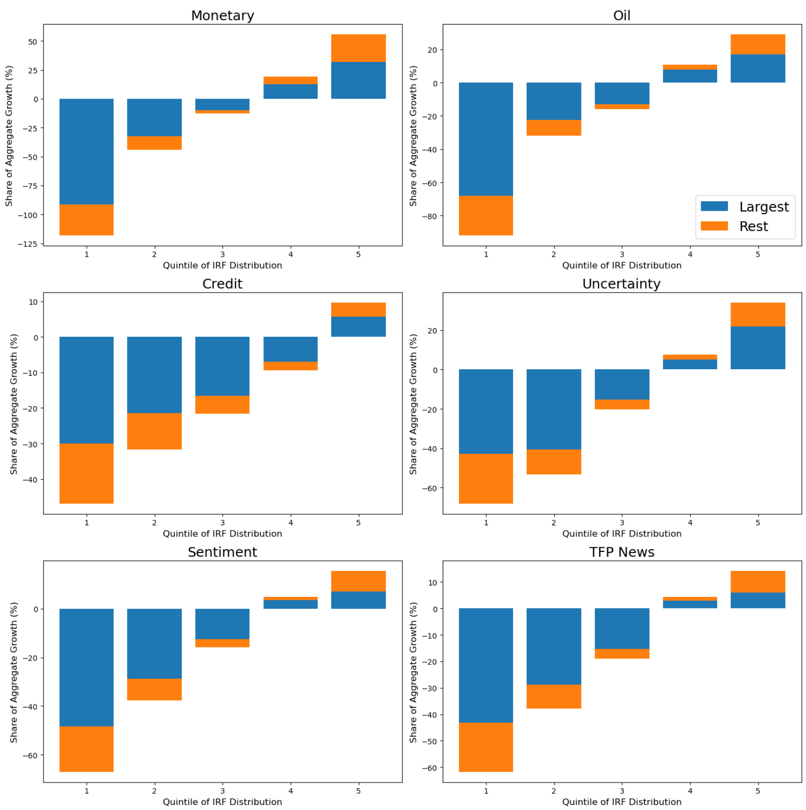

Note: Largest firms are the top 10% of the size distribution, which averages real sales over time for each firm. The contributions are re-scaled such that the bars add up to -100%.