# Real-time forward-looking skewness over the business cycle

Ian Dew-Becker*

December 4, 2020

**Abstract**

This paper studies option-implied skewness at the firm and market level between 1980 and 2019. It reports the first real-time measure of conditional micro skewness available in the literature. Firm skewness is significantly procyclical, while market skewness has no clear connection with the business cycle in either direction. Credit spreads are also tightly linked to firm skewness. These facts are consistent with models driven by variation in firm-level tail risk, rather than aggregate disaster risk. However, in regressions and a VAR it is realized rather than implied skewness that drives output, consistent with models in which output is a concave function of micro shocks.

## 1 Introduction

### Background

A growing literature has developed evidence that the cross-sectional distribution of many micro variables, including earnings, employment, productivity, and output, changes significantly over time. The cross-sectional variance of outcomes is typically – though not always – countercyclical, while growing evidence implies that cross-sectional skewness is procyclical.

The evidence on these points – for skewness, in particular – has come from analyses of the distribution of *realized* outcomes. For example, research has measured the skewness of the distribution of annual employment growth at the firm level and annual income growth individual level. But in models in which skewness is a driving force (e.g. Salgado, Guvenen, and Bloom (2020)), it is agents' ex ante beliefs about *future* skewness that matter. That is, decisions today depend on beliefs about the skewness of micro shocks that will be realized

*Northwestern University and NBER.

tomorrow, rather than the skewness of the shocks that were already realized today. In addition to the fact that work so far has been based on realized rather than expected skewness, it has also been based on low-frequency data – almost exclusively annual observations of income, employment, or other outcomes.[1]

A large literature also studies time variation in the skewness of aggregate shocks, arguing that time-varying skewness can explain a number of puzzles in financial markets and potentially also drive business cycles.[2]

**Contribution**

This paper contributes to the literatures on both micro and aggregate skewness. It develops real-time forward-looking measures of the conditional skewness of both firm-level and aggregate shocks from option prices. The measures are available at up to the daily frequency.

I use data on firm-level option prices to measure investors' beliefs about the conditional skewness of both individual stock and market index returns on a daily basis since 1980 (then aggregated to the monthly frequency to compare to other economic time series). While the CBOE reports a market-level skewness index, it has received little attention, and this paper is novel for combining it with firm-level information.

In addition to providing the first high-frequency forward-looking measure of the skewness of micro outcomes, the length of the data sample is notable. Typical option-implied measures go back no earlier than 1990 or sometimes the mid-1980's. Having data since 1980 enables analysis of the behavior of micro and aggregate skewness across five business cycles, and gives a sample about the same length as the annual-frequency samples of realized skewness in, e.g. Salgado, Guvenen, and Bloom (2020) and Guvenen, Ozkan, and Song (2014).

To validate the skewness indexes, I show that they have significant predictive power for future realized skewness – i.e. the skewness of the returns that are actually realized. That holds even after controlling for lagged realized skewness. The option-implied skewness measures are also significantly positively correlated with annual measures of realized skewness from other sources, like employment and income growth, which shows that what the options measure is not simply something about frictions in financial markets, but that it in fact has a link to skewness in real outcomes.

With the skewness indexes in hand, the paper goes on to examine the relationship between conditional skewness and the business cycle.

[1]See Guvenen, Ozkan, and Song (2014), Harmenberg and Sieversten (2017), Busch et al. (2018), Ilut, Kehrig, and Schneider (2018), and Salgado, Guvenen, and Bloom (2020).

[2]E.g. Barro (2006), Gabaix (2012), Gourio (2012, 2013), and Wachter (2013).

**Results**

The paper's first key finding is that conditional skewness is significantly procyclical – near zero in booms and substantially negative in recessions – but only at the firm level. Skewness of the overall stock market is essentially uncorrelated with the business cycle, and actually slightly above average during NBER-dated recessions. At the firm level, though, and when isolating the contribution of just firm-specific shocks, the skewness is clearly countercyclical across a variety of measures.

The total skewness faced by firms comes from a mixture of aggregate shocks – with acyclical skewness – and idiosyncratic shocks, which have procyclical skewness. The variation in firm skewness is quantitatively almost entirely driven by variation in the skewness of firm-specific rather than aggregate shocks, which explains why total firm skewness is also procyclical.

I next use cross-correlations to examine the lead-lag relationship between skewness and the business cycle. Skewness very clearly moves first, in the sense that it predicts measures of output better than output predicts skewness. A vector autoregression (VAR) confirms that result.

The natural next question is why skewness might drive the business cycle. A number of papers have proposed channels through which changes in risk might drive the business cycle (e.g. Christiano, Motto, and Rostagno (2014), though here the variation in risk is in the third instead of the second moment). In those models, changes in risk affect the business cycle through their effects on credit spreads. I show that there is a very tight relationship in the data between credit spreads and the conditional third moment faced by firms – they have a correlation of over 80 percent.

The final contribution of the paper is to study a VAR to understand in greater detail the dynamic effects of skewness shocks. In a VAR that includes just option-implied idiosyncratic skewness and macroeconomic aggregates, shocks to implied skewness can explain up to about one quarter of the variation in aggregate output, though the point estimates are closer to 10 percent. Next, we estimate a VAR including both implied and *realized* idiosyncratic skewness, where realized skewness is the skewness of the actual observed idiosyncratic shocks to firms, rather than the option-implied expectation. Realized skewness shocks can explain up to half the variation in aggregate employment and industrial production, with the point estimates around 35 percent. Furthermore, after controlling for realized skewness, implied skewness has only a very small impact on output.

What is particularly surprising about these results is that in that VAR, the option-implied skewness shocks still predict higher realized skewness in the future. That is, the VAR identifies a shock that predicts future realized skewness – so it is a shock to the conditional third moment (and in fact explains about 1/3 of the variation in the conditional moment) – but that has no impact on aggregate output. That result implies that it is not the conditional moment that really drives output but rather the skewness of realized shocks. Models in which aggregate output is a concave function of micro shocks, such as the complementary network models of Baqaee and Farhi (2020) and Dew-Becker, Tahbaz-Salehi, and Vedolin (2020) or the concave hiring model of Ilut, Kehrig, and Schneider (2018), are consistent with the VAR results.[3]

**Related literature**

As discussed above, this paper is related to work estimating time-variation in micro moments, including work on the cross-sectional distributions of income, employment growth, industrial production, sales, and stock returns. Much of that research studies not just the low-frequency properties of cross-sectional moments, but also how they vary over the business cycle and in relation to other variables. For cross-sectional variance, the results on cyclicality have been mixed, while the more recent work on skewness has so far found more consistent evidence that it is procyclical.

The evidence on the cyclicality of micro moments has been used to estimate and test structural models, such as those of Guvenen, Ozkan, and Song (2020) and Ferreira (2018). A key feature of such models is that they do not feature a representative agent or firm. Rather, micro heterogeneity must be modeled directly, and the heterogeneity has aggregate consequences. The moments developed here are therefore important for accurately calibrating the current generation of models. Since past work is typically measured at an annual frequency, and since it uses realized moments, which we show are noisy relative to true conditional moments, it has had significant shortcomings when it comes to estimation and testing. Our work is therefore important because it can be directly used to provide quarterly, monthly, or even daily information about dynamics of conditional moments faced by agents and firms in the economy, which was not previously possible.

The paper is also related to studies using stock options to estimate investor views about conditional distributions. The most widely used option-implied moment is the VIX index, which measures the conditional standard deviation of returns on the S&P 500. The CBOE also produces a SKEW index, replicated here, that measures the conditional skewness of the S&P 500 (see also Kozhan, Neuberger, and Schneider (2013)). Those are both aggregate moments – determined by the distribution of aggregate outcomes – but the theoretical literature just discussed is entirely about the cross-section. Dew-Becker and Giglio (2020) provide

[3] Gourieroux et al. (2020) suggest another possibility, that skewed shocks at the firm level could be self-exciting and have aggregate effects. They provide supporting evidence from derivatives prices.

a long time-series of firm-level implied volatility – a cross-sectional analog to the VIX – studying its cyclicality and using it to calibrate and test structural models. This paper takes that work a step further by measuring cross-sectional skewness, which is significantly more difficult technically. There is also a broader finance literature studying firm-level skewness, typically in the context of option pricing or return forecasting.[4]

Kozlowski, Veldkamp, and Venkateswaran (2020) examine market-level option-implied skewness (from the CBOE SKEW index) in a model of learning and show how it can affect the business cycle. I find, though, that it is primarily firm-level skewness that is cyclical, which suggests one direction for potentially extending their ideas.

The remainder of the paper is organized as follows. Section 2 discusses the data and methods used to estimate conditional skewness and section 3 validates the skewness measure. Sections 4 and 5 study the time-series of skewness and its cyclicality, and section 6 estimates a VAR. Finally, section 7 compares and contrasts the behavior of skewness and volatility, and section 8 concludes.

# 2 Data and methods

This section describes the theoretical foundation of the option-implied skewness measure, then discusses how it is feasibly implemented empirically. We next turn to the measurement of realized skewness, and then discuss data sources for the subsequent empirical analysis.

## 2.1 Theoretical structure

The fundamental goal is to construct an estimate of the conditional skewness of economic outcomes at the aggregate and individual firm levels. I focus on stock prices as the outcome of interest because they have liquid options associated with them over a long period, which are available at high frequency and help reveal investor perceptions of skewness. In addition, stock prices are useful because they summarize the information available to investors about the future profitability of US corporations. It would of course be preferable to have measure of skewness of some more fundamental concept, such as productivity, but unfortunately productivity is not directly measurable, nor does it have an associated options market to help reveal agents' probability distributions.

[4] Ferreira (2018), Jondeau, Zhang, and Zhu (2019), and Oh and Wachter (2019), among others, study realized skewness in the cross-section. Conrad, Dittmar, and Ghysels (2013) study firm-level option-implied skewness in the context of return forecasting. See also Bakshi, Kapadia, and Madan (2003) for foundational work on firm-level skewness and option prices.

By skewness, I specifically mean the scaled third moment. For a generic random variable $x_{t+1}$, define

$$skew_t\left(x_{t+1}\right) \equiv \frac{E_t\left[\left(x_{t+1}-E_t x_{t+1}\right)^3\right]}{E_t\left[\left(x_{t+1}-E_t x_{t+1}\right)^2\right]^{3/2}} \tag{1}$$

where $E_t$ denotes the statistical expectation conditional on information available at time $t$. There are other measures of skewness, such as Kelley skewness, but I focus on the scaled third moment because it involves expectations of smooth functions of the random variable $x$, which will be relatively easy to calculate from asset prices.

### 2.1.1 Skewness from option prices

Except under extreme and unrealistic assumptions, asset prices do not reveal statistical expectations. Rather, they can be used to reveal so-called "risk-neutral" expectations, which correspond to expectations weighted by state prices (Arrow–Debreu prices). Intuitively, they contain an adjustment for risk premia. The well-known VIX index, for example, is, under weak assumptions, a risk-neutral conditional standard deviation of stock returns over the next 30 days.

In what follows, I will refer to "option-implied expectations", and use the expectation operator, but I will always mean risk-neutral expectations. If risk premia are constant, they will only induce a level shift in the skewness. If they vary over time, though, they may cause the option-implied measure to fail to perform as a good statistical predictor of realized skewness. It will therefore be central in what follows to ensure that option-implied skewness actually forecasts realized skewness. Analogously, Dew-Becker and Giglio (2020) show that option-implied *volatility*, at both the firm and market levels, forecasts realized volatility, even after controlling for other predictors.

More than one method of constructing option-implied skewness has been proposed in the literature. For the sake of consistency with the measure of realized skewness analyzed below, I use the formula of Kozhan, Neuberger, and Schneider (2013),

$$SKEW_t \equiv 3\frac{\int_0^\infty \frac{K-F_t}{K^2 F_t} O_t\left(K\right)dK}{\left[\int_0^\infty \frac{1}{K^2} O_t\left(K\right)dK\right]^{3/2}} \tag{2}$$

where $O_t\left(K\right)$ is the price of an out-of-the-money option at strike $K$ on date $t$. Kozhan, Neuberger, and Schneider (2013) show that the leading term of the Taylor expansion of $SKEW_t$ is equal to the conditional skewness of the distribution of stock returns over the life of the options. The CBOE uses a similar formula for its S&P 500 skew index. The CBOE formula has a tighter link to the conditional third moment but a weaker link to feasible measures of realized skewness. Figure A.1 shows that the Kozhan, Neuberger, and Schneider (2013) and CBOE measures are highly similar, though not identical.[5]

### 2.1.2 Idiosyncratic skewness

In studying cross-sectional volatility, Campbell et al. (2001) model returns on individual firms with a simple single-factor specification,

$$r_{i,t} = r_{m,t} + \varepsilon_{i,t} \tag{3}$$

Treating each firm's loading on the market as $\approx 1$ makes the calculations simpler and avoids having to calculate a market loading for each stock (in the context of volatility, Dew-Becker and Giglio (2020) show that adjusting for stock-specific market loadings makes little difference). It is then straightforward to show that

$$skew\left(r_{i,t}\right) = \left(\frac{\sigma_{m,t}^2}{\sigma_{i,t}^2}\right)^{3/2} skew\left(r_{m,t}\right) + \left(1 - \frac{\sigma_{m,t}^2}{\sigma_{i,t}^2}\right)^{3/2} \underbrace{\left(skew\left(\varepsilon_{i,t}\right) + 3\frac{E\left[r_{m,t}\varepsilon_{i,t}^2\right]}{\sigma_{\varepsilon,t}^3}\right)}_{\text{Idiosyncratic skewness}} \tag{4}$$

where $\sigma_{m,t}^2$ is the conditional variance of the market return and $\sigma_{i,t}^2$ is the conditional variance of firm $i$'s return. $skew\left(r_{i,t}\right)$ and $skew\left(r_{m,t}\right)$ are calculated using the formula (2) above (and described further in section 2.2.2). Similarly, it is possible to estimate the conditional variances from option price data. The last term, $skew\left(\varepsilon_{i,t}\right) + 3\frac{E\left[r_{m,t}\varepsilon_{i,t}^2\right]}{\sigma_{\varepsilon,t}^3}$, then appears as a residual and represents the contribution of the firm-specific shocks, $\varepsilon_{i,t}$, to firm-level skewness. I refer to that term as idiosyncratic skewness. Idiosyncratic skewness has two sources: skewness in the firm-specific shocks themselves, and covariance between the magnitude of the shocks, $\varepsilon_{i,t}^2$, and the return on the overall market. When cross-sectional realized variance, as measured by $\varepsilon_{i,t}^2$, is higher in periods when the common component of returns is negative, then individual stock returns will have relatively long left tails and hence left skewness.

[5]For the purpose of estimating risk premia, the details of the construction become highly important. When examining correlations with, e.g., the business cycle, though, a shift in the mean becomes irrelevant, as do very high frequency differences.

## 2.2 Empirical implementation

### 2.2.1 Data

Skewness, from the scaled third moment, can in principle be calculated based on options on any underlying. To measure the skewness of aggregate outcomes, I calculate a skew index for the S&P 500 using data on CME S&P 500 futures options for the period 1983–1995 and data on CBOE S&P 500 options from Optionmetrics subsequently. When the two data sources overlap (1996–2015) they give nearly identical results (see figure A.1).

This paper's main innovation is to calculate a skewness index for *firm-level* outcomes using options on individual stocks. For the period 1/1996–12/2019, the data is from Optionmetrics, and for 1/1980–6/1995 from the Berkeley Options Database, which records all trades and quotes for all options traded on the CBOE over that period. See Dew-Becker and Giglio (2020) for a more extensive description of the Berkeley data.

The sample overall covers the period 1980–2019 (with a six-month gap for firm options in 1995), including five recessions. It is comparable in length to the samples used in Guvenen, Ozkan, and Song (2014), Busch et al. (2017), and Salgado, Guvenen, and Bloom (2020). In addition, it is longer than the sample that the CBOE uses for its skew index, which extends only to 1990, and for the VXO (S&P 100 volatility), which extends to 1986. I will show that the extra data is important for understanding trends over time, the effects of major events in 1987 and 1990, and for giving an appreciably greater number of recessions to analyze. Appendix A.1 discusses the underlying data in more detail.

### 2.2.2 Calculating integrals

For any given underlying, options are not traded at a continuum of strikes. We therefore must interpolate and extrapolate the option prices in order to calculate integrals, which we do through a Gaussian process regression described in Appendix A.1. The Gaussian process regression treats option prices within and across firms as being driven by a joint Gaussian process, which yields estimates of (unobserved) option prices at arbitrary strikes for any given firm based on that own firm's options and also information from option prices for other firms. Using the fitted prices, we can calculate the integrals in (2) numerically to arbitrary accuracy. For the S&P 500 we also use a Gaussian process, but do not incorporate any information from firm-level options, so the Gaussian process in that case just a method for interpolating between available strikes.

We estimate skewness for each stock on every day as long as the stock has at least one available option traded. Again, the Gaussian process makes this possible. In the case where a stock only has a single option traded, the estimation will use the information from that single option essentially to pin down the level of its option prices, while the relationship between option prices and strikes will simply mimic the average observed across other stocks on the same date. When more options for a given firm are available on a particular date, those prices eventually fully determine the pricing function for that firm.

For the S&P 500, the options on the index itself are the only source of information. We therefore require at least six options on each date, with at least two on each side of the underlying price. Because the BODB sample covers only the largest firms, I restrict attention to only the 200 largest firms by market capitalization in the Optionmetrics sample. That also helps ensure that the options used are liquid and reduces selection bias among smaller stocks.[6] Figure A.1 shows that the results are nearly identical using just the top 100 firms instead.

Formally, I calculate the skew index, $SKEW_{i,t}$ for each firm $i$ on every date $t$, and then calculate value-weighted average skewness as

$$SKEW_{firm,t} = \frac{\sum_i SKEW_{i,t} mkt_{i,t}}{\sum_i mkt_{i,t}} \tag{5}$$

where $mkt_{i,t}$ is the market capitalization of firm $i$ on date $t$. Weighting by firm employment instead of market capitalization has minimal effects (see also Dew-Becker and Giglio (2020)). $SKEW_{S\&P,t}$ is then the S&P 500 skewness index.

### 2.2.3 Idiosyncratic skewness

Using equation (4), we define

$$SKEW_{idio,t} = \left(1 - \frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{-3/2} \left(SKEW_{firm,t} - \left(\frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{3/2} SKEW_{mkt,t}\right) \tag{6}$$

where $IV_{mkt,t}$ and $IV_{firm,t}$ are the market- and firm-level option-implied volatilities (standard deviations) of future returns. There are numerous ways to calculate option-implied volatility. I use the Black–Scholes formula because it requires only a single option price for a given stock or index.[7]

[6] Among smaller firms, those with traded options may be systematically different from those without options: more volatile, and potentially with more highly non-normal returns that give investors reason to want to trade options.

[7] The model-free implied volatility for individual firms is somewhat noisier than the Black–Scholes implied volatility, which creates substantial volatility in $SKEW_{idio,t}$ in the early part of the sample.

There are two points to note about the definition (6). First, unlike the market- and firm-level indexes, $SKEW_{idio,t}$ is not an estimate of a third moment, but rather the *contribution* of idiosyncratic shocks, $\varepsilon_{i,t}$ to firm-level skewness (unlike the second moment, it is not possible to decompose the firm-level third moment into third moments of the common and idiosyncratic shocks). It is a mix of a third moment – $E\left[\varepsilon_{i,t}^3\right]$ – and a cross-moment, $E\left[r_{m,t} r_{i,t}^2\right]$.

Second, unlike $SKEW_{firm,t}$, $SKEW_{idio,t}$ is not an average across firms. Rather, it is calculated from the average moments across firms. That helps reduce the impact of measurement error. $SKEW_{idio,t}$ should thus, formally, be thought of as the contribution of idiosyncratic shocks to skewness for a hypothetical firm that has the average level of skewness and volatility observed across firms on date $t$.

## 2.3 Realized skewness

The skewness measures are conditional expectations, so to validate them I will compare them to the realized sample moments. As the indexes are for monthly skewness, they are naturally compared to realized cubed monthly returns. That single cubed return is very noisy, though, making it a weak test of the accuracy of the index as a true expectation.[8] Neuberger (2012) shows that it is possible to construct a realized third moment from daily data using implied volatility. Using the case where returns aggregate linearly for simplicity (i.e. $r_{monthly} = \sum_{days} r_{daily}$; Neuberger (2012) discusses the extension to multiplicative aggregation)

$$E\left[r_{monthly}^3\right] = E\left[\sum_{days} r_{daily}^3 + 3 r_{daily} \Delta\sigma_{daily}^2\right] \tag{8}$$

[8]In the case of volatility, and assuming a Normal distribution, the squared monthly return is equal to the true expectation multiplied by a $\chi_1^2$ random variable, which has a standard deviation of $\sqrt{2}$ – so a single observation is an extremely noisy measure of the true expectation.

For volatility, it is possible to take advantage of the fact that, if returns are unpredictable, $E\left[r_{monthly}^2\right] = E\left[\sum_{days} r_{daily}^2\right]$, which follows simply from the additivity of the variances of uncorrelated random variables. That equation simply says that the expectations of monthly squared returns and the sum of daily squared returns are the same. Under a Normal distribution, instead of having a $\chi_1^2$ error, we have $\chi_{20}^2/20$ (for 20 trading days in a month), which has a standard deviation smaller by a factor of $\sqrt{20}$. In the case of cubed returns, the monthly realized third moment is no longer equal in expectation to the sum of the daily realized third moments,

$$E\left[r_{monthly}^3\right] \neq E\left[\sum_{days} r_{daily}^3\right] \tag{7}$$

which is why the Neuber (2012) formula is different.

where $\Delta\sigma^2_{daily}$ is the daily change in the conditional variance of the return, which as above, comes from option-implied volatility. Equation (8) reveals two sources for skewness at monthly frequencies: high-frequency skewness, through $r^3_{daily}$, and the realized covariance between realized returns and changes in implied volatility (because unlike with variances, $E\left[r^3_{monthly}\right] \neq E\left[\sum_{days} r^3_{daily}\right]$). If volatility rises when returns are negative, then there is a larger chance for cumulative returns to become extremely negative, generating negative skewness. On the other hand, if volatility rises when daily returns are positive, then monthly returns will be more positively skewed.

Based on (8), average realized skewness at the firm and market level in month $t$ are

$$RSKEW_{firm,t} = \sum_i mkt_{i,t} \frac{\sum_{days \in t} r^3_{i,daily} + 3r_{i,daily}\Delta\sigma^2_{i,daily}}{\left(\sum_{days \in t} r^2_{i,daily}\right)^{2/3}} \tag{9}$$

$$RSKEW_{S\&P,t} = \frac{\sum_{days \in t} r^3_{S\&P,daily} + 3r_{S\&P,daily}\Delta\sigma^2_{S\&P,daily}}{\left(\sum_{days \in t} r^2_{S\&P,daily}\right)^{2/3}} \tag{10}$$

Similarly to how $SKEW_{idio,t}$ is constructed, idiosyncratic realized skewness is then also a residual, with

$$RSKEW_{idio,t} = \left(1 - \frac{RV^2_{mkt,t}}{RV^2_{firm,t}}\right)^{-3/2} \left(RSKEW_{firm,t} - \left(\frac{RV^2_{mkt,t}}{RV^2_{firm,t}}\right)^{3/2} RKEW_{mkt,t}\right) \tag{11}$$

where $RV_{mkt,t}$ and $RV_{firm,t}$ are the market- and firm-level realized volatilities constructed from sample standard deviations of daily returns in month $t$.

To confirm that $RSKEW_{firm}$ is less noisy than what is obtained from cubed monthly returns, the bottom-right panel of figure A.1 plots it against

$$\frac{\sum_i mkt_{i,t}\left(r_{i,t} - \bar{r}_t\right)^3}{\left(\sum_i mkt_{i,t}\left(r_{i,t} - \bar{r}_t\right)^2\right)^{3/2}} \tag{12}$$

where $\bar{r}_t$ is the average return across all stocks in *month* $t$. The skewness from the monthly returns in (12) has a highly similar sample mean to $RSKEW_{firm,t}$ – -0.27 compared to -0.24 – but $RSKEW_{firm,t}$ is less volatile by a factor of 3.5. Neuberger (2012) similarly shows that $RSKEW_{S\&P,t}$ is much less noisy than calculating skewness based on realized cubed returns for the S&P 500.

The analysis here is entirely in terms of monthly skewness, both conditional and realized (i.e. a monthly horizon, not sampling frequency). I use a monthly horizon for comparability with the well known VIX, which is also a monthly measure, and also because most macroeconomic time series are reported at no higher than the monthly frequency. In addition, option markets are relatively liquid at shorter maturities. There is very little volume in 12-month options, for example, especially for individual stocks.

# 3 Validating implied skewness

This section validates the skewness measures by showing that they indeed are able to forecast realized skewness. Option-implied skewness at the firm, market, and idiosyncratic levels forecasts the skewness of realized returns, even after controlling for lagged realized skewness.

Implied and realized skewness is plotted in the left-hand panels of figure 1. In all three cases, implied and realized skewness are clearly correlated, but realized skewness has much more high-frequency variation, consistent with the presence of unpredictable surprises (since the realization should be equal to the expectation plus random noise). That effect is stronger for the S&P 500 than for the firm-level results because the averaging across firms reduces the noise.

Table 1 reports results of forecasting regressions for realized skewness. The first column in the top section shows that implied skewness has significant predictive power for realized skewness at the firm level. The second column shows that the predictive power holds even after controlling for lagged realized skewness. In other words, investors have information about the future skewness of firm-level economic outcomes, which they use to price options and which hence appears in our implied skewness measure. That information is independent of the information contained in lagged realized skewness. That result continues to hold if more lags of realized skewness or other controls are included and also in an ARMA(1,1) model.

The variation in conditional skewness implied by the first column of table 1 is substantial. The fitted values from the forecasting regression have a standard deviation of 0.15 and a 10th/90th percentile range of -0.46 to -0.08, implying considerable variation in conditional skewness.

Columns 3 and 4 show that similar results hold using realized skewness calculated from the less accurate method full monthly returns in equation (12). Lagged $SKEW_{firm,t}$ remains statistically significant, and its coefficient has a similar value, though the confidence bands are wider, consistent with the larger amount of noise in the monthly realized skew measure.

It is worth noting that the forecasting coefficient is substantially less than 1 in both cases. That fact is consistent with the presence of pure measurement error (e.g. due to mismeasurement of option prices or discretization error in the integrals) or time-varying risk premia in option-implied skewness. Likely both are at play. The constant in the regression is also not zero, which is visible from the difference in means in figure 1 and also implies the presence of risk premia or measurement error. Interestingly, there is little mean bias for $SKEW_{idio}$, implying that there is a skew risk premium at the aggregate but not the idiosyncratic level, as is also observed for volatility (Dew-Becker and Giglio (2020)).

As noted above, realized skewness has two components. Neuberger (2012) shows that the term involving the covariance of returns and changes in implied volatility is present when returns follow a diffusion, while the term involving cubed returns is driven by jumps (as the sampling frequency grows). The fifth and sixth columns of table 1 replace total skewness with those two components. The coefficients on $SKEW_{firm,t-1}$ in columns 5 and 6 are guaranteed to sum to the coefficient in column 1. Columns 5 and 6 show that the predictive power comes almost entirely from forecasting the diffusive component. That is for two reasons. First, the $R^2$ for the diffusive component is substantially higher than the jump component – 0.296 compared to 0.003. Second, the diffusive component is more volatile than the jump component by a factor of 2.6. That is, the diffusive component is both by far the largest component of overall realized skewness – accounting for almost 90 percent of its variance – and also much more predictable than the jump component.

The middle section of table 1 repeats the above regressions for the S&P 500. The results are similar. $SKEW_{S\&P}$ has significant forecasting power for realized skewness. The second column shows that realized skewness is driven out of the forecasting regression in this case: lagged realized skewness contains no marginal information for future skewness after controlling for option implied skewness. Since there is just a single observation, rather than a cross-section, monthly realized skewness is undefined in this case, so the third and fourth columns are empty. As with the firm-level returns, the predictive power for S&P 500 skewness is isolated to the diffusive component.

Finally, the bottom section of table 1 reports results for forecasting $RSKEW_{idio,t}$. The results are highly similar to those for $RSKEW_{firm,t}$, which we will see below is due to the fact that implied and realized skewness at the firm level are driven almost entirely by the idiosyncratic component of returns.

One additional robustness test I have run is to include a time trend in the regressions. That is not formally justified in the forecasting context, but it is worth checking in this case because realized and implied skewness both clearly trend down over the sample. For both the firm and S&P 500 regressions, adding a time trend does not drive implied skewness out of the regression, nor does it substantially reduce the magnitude of its coefficient (note also that the inclusion of lagged realized skewness in the regressions is similar to detrending (Hamilton (2018)) and can eliminate unit root problems).

# 4 The history of conditional skewness

This section examines the time-series properties of implied skewness. It first analyzes the basic properties and their relationship with other option-implied moments. Next, it analyzes the cyclicality of firm- and market-level skewness, showing that firm-level skewness is strongly procyclical while market skewness has little or no relationship with the business cycle. Finally, I show that credit spreads are significantly driven by firm-level skewness (and, most strongly, by the firm-level third moment).

## 4.1 Time series properties

The left-hand panels in figure 1 plot the time series of $SKEW_{firm,t}$, $SKEW_{S\&P,t}$, and $SKEW_{idio,t}$. All three have clear downward trends over time. Aggregate skewness has fallen from around 0 to -3, while firm skewness fell from 0 to -1. A significant fraction of the total decline in aggregate and firm skewness comes in two episodes in 1987 and 1989. In both cases, market skewness fell by 1.1, while firm skewness fell by 0.8 in 1987 and 0.4 in 1989. The 1987 episode is the October crash. In 1989, there was also a significant decline in the stock market associated with a large increase in volatility (which, as shown below, creates left skewness at the firm level). Prior to those events, both conditional and realized stock return skewness were near zero.

The correlation between firm and market skewness is, in some sense, surprisingly weak. They have an overall correlation of 0.64, but that is entirely due to their shared time trend. Once a linear trend is removed from both series, the correlation falls to 0.05. So while they share low-frequency variation, their business cycle and higher-frequency variation is essentially independent. The correlation between market and idiosyncratic skewness is similarly small.

Shocks to the three skewness series are also notably short-lived. Figure 2 plots the first 12 autocorrelations for Hodrick–Prescott detrended employment, industrial production, credit spreads, and skewness. The autocorrelations in the three skewness measures decay far more quickly than those for employment, IP, and credit spreads – to zero within a year, compared to approximately 0.4 for employment. Shocks to skewness thus appear to be relatively short-lived, which is also clear from simple inspection of figure 1.

Table 2 summarizes the results so far and gives statistics that can be used to guide calibrations of models driven by variation in idiosyncratic skewness. It reports the unconditional means and standard deviations of the three skewness measures along with the monthly, quarterly, and annual autocorrelations, both with and without removing a time trend. It should be noted that the 12th autocorrelation is much larger than the first autocorrelation raised to the 12th power, implying that skewness has transitory variation – such as an MA(1) term. If it were to be calibrated as an AR(1) process, matching the annual autocorrelation would therefore likely be most natural for most applications.

## 4.2 Aggregate and idiosyncratic contributions to firm skewness

Recalling equation (4), and rearranging the definition of $SKEW_{idio,t}$ in equation (6), we have

$$SKEW_{firm,t} = \left(\frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{3/2} SKEW_{mkt,t} + \left(1 - \frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{3/2} SKEW_{idio,t} \qquad (13)$$

The total skewness faced by individual firms, $SKEW_{firm,t}$, can therefore be thought of as having three sources: market skewness, $SKEW_{mkt,t}$, idiosyncratic skewness, $SKEW_{idio,t}$, and changes in their relative weights, through $IV_{mkt,t}/IV_{firm,t}$. Figure 4 illustrates the relative contributions of those three terms by plotting $SKEW_{firm,t}$ against the following three comparisons:

$$\text{Idio. skew only:} \quad \left(\frac{IV_{mkt}^2}{IV_{firm}^2}\right)^{3/2} SKEW_{mkt} + \left(1 - \frac{IV_{mkt}^2}{IV_{firm}^2}\right)^{3/2} \boxed{SKEW_{idio,t}} \qquad (14)$$

$$\text{Mkt. skew only:} \quad \left(\frac{IV_{mkt}^2}{IV_{firm}^2}\right)^{3/2} \boxed{SKEW_{mkt,t}} + \left(1 - \frac{IV_{mkt}^2}{IV_{firm}^2}\right)^{3/2} SKEW_{idio} \qquad (15)$$

$$\text{Volatility only:} \quad \left(\boxed{\frac{IV_{mkt,t}^2}{IV_{firm,t}^2}}\right)^{3/2} SKEW_{mkt} + \left(1 - \boxed{\frac{IV_{mkt,t}^2}{IV_{firm,t}^2}}\right)^{3/2} SKEW_{idio} \qquad (16)$$

In each case, only one of the three variables determining the value of $SKEW_{firm,t}$ is allowed to vary over time – the other two are held at their unconditional means over the sample (denoted by the elimination of the time subscript).

The top panel of figure 4 shows that when only $SKEW_{idio,t}$ is allowed to vary, the series looks nearly identical to total firm skewness, with a correlation of 80 percent. So variation over time in the total skewness faces by individual firms is almost entirely explained by variation in the skewness of their idiosyncratic shocks, consistent with models that emphasize variation in micro risk, such as Christiano, Motto, and Rostagno (2016) and Di Tella and Hall (2020).

Panels b and c essentially confirm the result from the top panel. Panel b shows that variation in the market skew explains essentially none of the variation in firm risk – their correlation is only 7 percent. Interestingly, the bottom panel shows that variation in volatility does matter somewhat. Recall from figure 1 that $SKEW_{mkt,t}$ is always more negative than $SKEW_{idio,t}$. That means that when market volatility rises relative to firm volatility, $IV_{mkt,t}/IV_{firm,t}$ rises, total firm skewness becomes more negative. In episodes when $IV_{mkt,t}/IV_{firm,t}$ rose significantly, such as the spikes in the late 1990's, the financial crisis, and the debt ceiling and Euro crises in 2010 and 2011, firm skewness became more negative. In those periods, though, there were typically also declines in the idiosyncratic component of skewness (with the exception of 1998, which seems to have been purely related to volatility changes). Overall, changes in volatility have a correlation of 55 percent with firm skewness. However, that component is much less volatile than the component due to idiosyncratic skewness. The idiosyncratic skewness component has a standard deviation three times higher than the component from volatility.

Because $SKEW_{firm,t}$ is a nonlinear function of the three components, there is no additive variance decomposition. Nevertheless, the ranking of importance, both in terms of ordering and magnitude, is clear.

Figure 4 thus yields two important conclusions. First, the vast majority of the variation in the skewness faced by firms is due to skewness induced by their idiosyncratic shocks, not skewness in aggregate shocks. Second, changes in skewness faced by firms can have subtle causes. In particular, even if the skewness of aggregate shocks stays fixed, when those shocks become more volatile, as in the various financial crises in the sample, they cause the total left skewness faced by firms to grow. In other words, volatility itself can cause skewness.

## 4.3 Relationship with other measures of skewness

Figure 3 plots $SKEW_{firm}$, minus an HP-filtered trend,[9] against three measures of cross-sectional skewness in annual growth rates: sales growth in Compustat, employment growth in the Census LBD, and the cross-section of income growth from the Social Security Administration.[10] These alternative measures are *not* detrended. The figure also displays the

[9]The figure uses a smoothing parameter of 129,600 on the monthly data and plots annual averages of detrended $SKEW_{firm,t}$ for comparability with the annual skewness series.

[10]I calculate Compustat sales and employment growth skewness directly. The LBD data is taken from Salgado, Guvenen, and Bloom (2020) and the Social Security Administration data is from Guvenen, Ozkan, and Song (2014). In all cases, Kelley skewness is $(p90 - 2 \times p50 + p10)/(p90 - p10)$, where $px$ denotes the

fraction of each year in which the economy was in a recession. Note that these are all measures of *realized* rather than conditional skewness. They measure the skewness of growth rates that actually were realized, as opposed to what agents expected the skewness so be.

In all three panels, there is a clear positive correlation between detrended $SKEW_{firm}$ and the alternative cross-sectional skewness measures. The overall correlations are all around 0.3, and the three series display similar patterns particularly since the late 1990's, all falling in the 2001 and 2008 recessions and displaying slight declines in 2016. On the other hand, it is notable that $SKEW_{firm}$ has a clear downward trend (see figure 1), as does realized stock return skewness, but that the alternative measures do not. That suggests that something has changed in the behavior of stock returns relative to employment and sales in the cross-section.

These alternative skewness series are only available at the annual frequency, yielding only 40 observations since 1980 and making identification of cyclicality difficult.

Another alternative to option-implied skewness, which has been explored in past work on the S&P 500 (e.g. Segal, Shaliastovich, and Yaron (2015)), is to measure upside and downside variance. Specifically, one can define, for either an individual firm or for an index

$$VIX_{up}^2 = E\left[r^2 \mid r > 0\right] \quad (17)$$

$$VIX_{down}^2 = E\left[r^2 \mid r < 0\right] \quad (18)$$

$$VIXA = \frac{VIX_{up} - VIX_{down}}{VIX} \quad (19)$$

that is, $VIX_{up}$ is the expected squared return conditional on returns being positive and $VIX_{down}$ is defined analogously. $VIXA$ is then a scaled measure of asymmetry, similar to skewness. The bottom-right panel of figure 3 plots $SKEW_{firm}$ along with $VIXA_{firm}$ (rescaled to have the same standard deviation as $SKEW_{firm}$). The two series have a correlation of 0.97, so they contain almost exactly the same information. A similar result holds for the case of S&P 500 skewness and VIX asymmetry. This result and those in figure A.1 jointly show that stock market skewness is highly similar regardless of the precise measure used.

# 5 Skewness over the business cycle

This section examines the cyclicality of skewness, both through simple correlations and also by studying the lead-lag structure. Firm-level skewness is procyclical and leads real activity. One proposed channel through which skewness could affect activity is its direct effect on

$x$th percentile of the cross-section of growth rates in each year.

credit spreads, skewness and spreads are in fact tightly linked.

## 5.1 Cyclicality

Figure 1, again, plots the three skewness series. Each panel includes gray bars indicating NBER-dated recessions. In all five recessions in the data there were declines in both total firm skewness and the idiosyncratic component. Market skewness, on the hand, shows no particular relationship with recessions – in many cases it actually rises. To the skewness of the shocks faced by firms is significantly procyclical, and that behavior is driven almost entirely by changes in the skewness of idiosyncratic shocks.

Table 3 reports correlations of the three skewness series, with and without linear trends removed, with various measures of real activity. In all cases, firm and idiosyncratic skewness are estimated to be significantly procyclical. They are positively correlated with both growth rates and detrended levels of employment and industrial production, the CBO output gap, and capacity utilization, and negatively correlated with the unemployment rate and an NBER recession indicator. Furthermore, the cyclicality is somewhat stronger for the idiosyncratic component.

An alternative way to quantify the cyclicality is to simply measure the average skewness in and out of recessions. After controlling for a time trend, firm skewness is on average lower by 0.17 in recessions, and idiosyncratic skewness is lower by 0.40, representing shifts of 0.67 and 0.74 standard deviations, respectively. There are thus clear declines in skewness in recessions, but, as is also visible in figure 1, recessions are also far from the only source of variation in skewness.

The correlations for $SKEW_{S\&P}$, on the other hand, are all close to zero and there is no consistent evidence of cyclicality – some correlations imply $SKEW_{S\&P}$ is procyclical and others imply it is countercyclical. During NBER-dated recessions, S&P 500 skewness is actually on average higher by 0.35, or 0.60 standard deviations. Overall, then, there is little evidence that aggregate skewness is procyclical – if anything is is countercyclical – while firm-level outcomes have consistently procyclical skewness, almost entirely due to variation in idiosyncratic skewness.

To further illustrate the procyclicality of firm skewness, the right-hand panels of figure 1 plot HP-filtered firm, market, and idiosyncratic skewness against HP-filtered unemployment, with the axis for the latter reversed. The HP filter here is used simply to remove low-frequency trends. In every recession, and especially 2001 and 1991, when unemployment initially rises there is a clear decline in both firm and idiosyncratic skewness. Skewness in many cases recovers quickly, and, unlike unemployment, it has substantial variation also outside recessions.

Again, for the S&P 500 there is no clear relationship with the business cycle. In fact, HP-filtered S&P 500 skewness is mildly positively correlated with HP-filtered unemployment (i.e. countercyclical).

## 5.2 Cross-correlations

The bottom four panels of figure 2 plots the cross-correlations of implied and realized skewness with industrial production and employment. The middle two panels plot cross-correlations with idiosyncratic implied and realized skewness, while the bottom two panels plot cross-correlations with S&P 500 implied and realized skewness. On the left are cross-correlations with employment, and on the right industrial production. As in Christiano, Motto, and Rostagno (2016), I apply the Hodrick–Prescott filter in order to focus on cyclical relationships, which means these correlations do not measure true forecasting relationships (see section 6).

$SKEW_{idio,t}$ has strong correlations with employment and industrial production, peaking around 0.40 for each. Furthermore, the correlations are highest between current levels of real activity and *lagged* implied and realized skewness. That is, skewness leads employment and IP. The correlations are higher for implied than realized skewness, which is consistent with the presence of greater noise in realized skewness.

The results for S&P 500 skewness are again drastically different from those for idiosyncratic skewness. Both implied and realized S&P 500 skewness are essentially uncorrelated with employment and IP. If anything, high current S&P 500 skewness actually appears to predict *low* future employment and IP.

Overall, then, we find that firm and idiosyncratic skewness are significantly procyclical, leading real activity. Skewness at the aggregate level, on the other hand, appears to have little relationship with the business cycle at all, regardless of leads or lags.

## 5.3 Credit spreads, skewness, and the third moment

There are a number of possible channels through which skewness and the business cycle could be related. A particularly prominent possibility is that they are linked through credit spreads: when skewness becomes more negative, that might increase credit spreads and hence borrowing costs for firms, reducing investment and GDP. That is the mechanism underlying the model of Christiano, Motto, and Rostagno (2016), though in that model the driving force is variation in volatility instead of skewness.

There is naturally a relationship between the skewness of stock returns and credit spreads. Corporate debt can be viewed as a short put option (Merton (1974)) – when firm value falls sufficiently far, the value of the debt is eventually impaired, whereas when firm value rises, there is a cap on the value of its debt. That implies that credit spreads should in general increase in the mass of the distribution of outcomes in the left tail. That mass can increase either due to the scale of the distribution – its variance – increasing, or due to a shift in its shape to be more left skewed. Those two effects can be summarized by simply hypothesizing that there should be a negative relationship between the third moment of the return distribution and credit spreads.

A number of structural models of the economy rely on such a relationship between credit spreads the distribution of shocks. Gourio (2013) shows that an increase in aggregate disaster risk – corresponding to an increase in the variance and left skewness of returns for the overall stock market (i.e. $IV_{S\&P}$ and $-SKEW_{S\&P}$) – can raise credit spreads and cause a recession. Christiano, Motto, and Rostagno (2014) study how an increase in the variance of *firm-specific* shocks (an increase in $IV_{firm,t}$ but not the $IV_{S\&P,t}$) can raise credit spreads and reduce GDP. One way to distinguish those two models is to evaluate whether aggregate or firm-specific risk is more important for driving variation in credit spreads.

To better understand the relationship, table 4 reports regressions of the Gilchrist–Zakrajsek (2012) measure of credit spreads on various option-implied moments. The first column confirms the negative relationship with firm skewness and the positive relationship with firm implied volatility, $IV_{firm,t}$. The second result is consistent with Christiano, Motto, and Rostagno (2016). The skewness result would presumably also hold in their model, but their shocks only affected volatility, not skewness, so it is technically outside the model.

The second column of table 4 replicates the regression from column 1, but replacing firm-level volatility and skewness with S&P 500 moments. Increased volatility and decreased skewness are again both associated with higher credit spreads, in this case consistent with Gourio (2013).

A simple way to summarize the positive relationship with volatility and the negative relationship with skewness is through the third moment of returns,

$$M3_{firm,t} = SKEW_{firm,t} \times VIX^3_{firm,t} \quad (20)$$

$$M3_{S\&P,t} = SKEW_{S\&P,t} \times VIX^3_{S\&P,t} \quad (21)$$

A decrease in skewness always decreases the third moment. When the skewness is negative, as is true for almost the entire sample, an increase in the volatility makes the third moment more negative. So we would expect the third moment to be negatively related overall to credit spreads and summarize the information in both skewness and implied volatility. The key question at that point is whether it is the firm- or aggregate-level third moment that is relevant for driving credit spreads.

The third column of table 4 reports results from this regression. The $R^2$, at 0.65, is slightly higher than in the first and second columns, showing that the option-implied third moment summarizes the information available in the skewness and standard deviation. The key result in this column is that the coefficient on the firm-level third moment is strongly negative, while the coefficient on the market third moment is actually positive, though statistically insignificant. That result is consistent with the idea that credit spreads reflect firm risk that is also measured by third moments. Conditional on a measure of firm-level risk, there is no reason why aggregate risk should affect credit spreads (except perhaps through risk premia).

Overall then, consistent with the intuition from models driven by variation in both aggregate and cross-sectional risk, movements in equity volatility and skewness feed into credit spreads. That information is well summarized by the conditional third moment of returns.

# 6 VAR analysis

This section examines the dynamic relationship between skewness and other features of the economy in the context of a simple vector autoregression (VAR). The correlations reported above showed both that firm-level skewness is procyclical and that it leads the business cycle. A VAR is useful for quantifying the extent to which shocks to skewness can potentially explain variation in real activity.

An additional question is whether it is the implied or realized part of skewness that is the relevant driver of activity (at least in a simple forecasting sense). Berger, Dew-Becker, and Giglio (2020) ask the same question in the context of aggregate volatility. It is well known that aggregate volatility is countercyclical, but that is true of both realized and implied volatility. Berger, Dew-Becker, and Giglio (2020) show that realized volatility drives implied volatility out of forecasting regressions and VARs for IP and employment. After controlling for realized volatility, implied volatility has no forecasting power for activity, even though it does still forecast realized volatility (i.e. it represents a component of uncertainty). They propose a model based on skewed productivity shocks that can generate such behavior (Dew-Becker, Tahbaz-Salehi, and Vedolin (2020) suggest another explanation based on a nonlinear network production model). This section asks the analogous question for the case of skewness.

## 6.1 Regressions

Before estimating a full VAR, I examine a simple forecasting regression. As above, I focus on the third moment, as opposed to just skewness, since the third moment captures the total risk faced by firms. I use the idiosyncratic third moment, based on the results above that the market component of skewness is acyclical. Using the total firm third moment gives similar, though slightly weaker, results (since it includes aggregate skewness, which is acyclical). We have

$$M3_{idio,t} = M3_{firm,t} - M3_{S\&P,t} \tag{22}$$

$$R3_{idio,t} = \left(RSKEW_{firm,t} \times RSD^3_{firm,t}\right) - \left(RSKEW_{S\&P,t} \times RSD^3_{S\&P,t}\right) \tag{23}$$

For the VAR, I take the cube root of the third moments, so that they have the same units as returns (similar to using a standard deviation instead of a variance), which also reduces the influence of extreme observations.

Table 5 reports results from regressions of employment and IP growth on four of their own lags and four lags of $M3_{idio,t}$ and $R3_{idio,t}$. Including four monthly lags accounts for the delayed effects of skewness on aggregate output. The table reports the sum of the coefficients here for simplicity (the VAR results will more fully describe the dynamics). The top panel reports results for employment growth and the bottom panel IP growth.

The first and third columns of table 5 show that in regressions of employment and IP growth on lagged $M3_{idio}$, the conditional third moment predicts significantly higher growth rates (I leave interpretation of the magnitudes of the coefficients to the VAR where we can construct a forecast error variance decomposition). The second and fourth columns in table 5 report results including both the implied and realized third moment, $M3_{idio}$ and $R3_{idio}$. If it is shocks to expectations that are the relevant channel for skewness to affect the economy, then one would expect $M3_{idio}$ to have a positive coefficient and for the coefficient on $R3_{idio}$ to be no different from zero.

Table 5 shows instead that it is realized skewness that has a significantly positive coefficient, while the coefficient on implied volatility actually turns slightly negative. Furthermore, the coefficients on the realized third moment are three to four times larger than the coefficients on the implied third moment in columns 1 and 3, which is surprising given that the two variables have almost exactly the same volatility. The realized moment absorbs all of the predictive power of $M3_{idio}$, and in fact predicts far larger movements in output than the implied third moment. Those results are particularly surprising as we will see below that $M3_{idio}$ forecasts $R3_{idio}$.

The regressions in table 5 give the simplest evidence that third moments help forecast economic growth, but that it is really the realized rather than the implied third moment that drives the relationship.

## 6.2 VAR

The variables included in the VAR are the same as those in Berger, Dew-Becker, and Giglio (2020), but replacing volatility with skewness. Here I use a less sophisticated identification scheme and simply order $R3_{idio}$ first and $M3_{idio}$ second in a Cholesky factorization. The VAR also includes employment, IP, and the Fed Funds rate. The ordering of the latter three variables is irrelevant since the goal is just to estimate responses to shocks to the first two variables. I also estimate a separate VAR excluding $R3_{idio}$.

While the timing assumption for realized and implied third moments will certainly be invalid in many interesting models, it will be correct in models where variation in firm risk (i.e. in the conditional third moment) is an exogenous shock. Berger, Dew-Becker, and Giglio (2020) discuss this point in detail. In past work, the exogenous variation in risk has typically been in the second moment, but it could also plausibly be in the third moment.

The left-hand panels of figure 5 plot responses to shocks to $M3_{idio}$ in a VAR that, like the first column of table 5, includes just $M3_{idio}$ and not $R3_{idio}$. The shaded regions indicate 90-percent confidence bands. Increases in the implied third moment are persistent, lasting about two years on average, though falling by half relatively quickly, in two to four months. Even with the quick initial decline, though, they are followed by highly persistent increases in employment and IP, showing no significant reversion even after two years.

The left-hand panels of figure 6 plot forecast error variance decompositions. The fraction of the variance of employment and IP accounted for by third moment shocks is economically significant, but not enormous. At two years ahead, the point estimate is 14 percent for employment and the confidence interval has an upper bound of 32 percent. For IP, the point estimate is only 7 percent with an upper bound of 23 percent.

The right hand panels of figures 5 and 6 report results when both the realized and implied third moments are included. Rather than the usual scaling in standard deviation units, I scale the two sets of IRFs in the right-hand panels so that they correspond to equivalently sized third moment shocks. In particular, they are scaled to have the same cumulative impact on expected $R3_{idio}$ over the next 24 months. The magnitude of the impact of the two shocks on expectations of future realized third moments is therefore the same by construction.

The results in the right-hand panels are substantially different from the previous results in a number of ways. First, shocks to the implied third moment no longer have statistically or economically significant effects on output – the impulse responses differ little from zero, and the confidence bands only barely even contain the point estimates from the model with $M3_{idio}$ alone. Second, shocks to the *realized* third moment have significantly positive effects, and the confidence bands for the IRFs for $M3_{idio}$ and $R3_{idio}$ do not even overlap for the first 12 to 18 months. Moreover, the effects of $R3_{idio}$ are substantially larger than those observed in the left-hand column for $M3_{idio}$ in the previous VAR. Consistent with the regressions in the previous section, shocks to realized third moments not only drive out the implied third moment but actually have significantly larger effects on employment and IP.

The larger impulse responses translate to a larger fraction of variance explained in figure 6. At the point estimates for two years ahead, $R3_{idio}$ explains 36 percent of the variance of employment and 33 percent of the variance of IP. The upper bounds for the confidence intervals are 50 and 52 percent, respectively. Consistent with the IRFs, the variance explained by the implied third moment is now indistinguishable from zero.

All of that said, it is not the case in this VAR that the shocks to $M3_{idio}$ play no role at all. They still explain the majority of the forecast variance for $M3_{idio}$ itself. Furthermore, they forecast increases in the realized third moment (again, the IRFs are scaled so that the cumulative impact on $R3_{idio}$ is the same for the two shocks). The bottom panel in figure 6 plots the fraction of the variance of expected future $R3_{idio}$, defined as $E_t\left[\sum_{j=1}^{24} R3_{idio,t+j}\right]$, explained by shocks to $R3_{idio}$ and $M3_{idio}$ (note that the variance decompositions are unaffected by the choice of the scaling in the IRFs in figure 5). The two shocks explain approximately equal fractions of the total variation in expected future third moments. In other words, they represent essentially equivalent shocks to conditional third moments.

That fact makes it all the more surprising that the IRFs for employment and IP are so different. Essentially what the VAR results show is that we have two shocks that have the same effect on expected future third moments – so represent equivalent shocks to economic risk going forward – but that have very different effects on the real economy. That result implies that the effects of realized third moments on economic activity are not due purely to their effects on risk, but to something else.

# 7 Comparing and contrasting volatility and skewness

This last section briefly compares and contrasts the cyclical behavior of volatility and skewness at both the firm and aggregate level. In a pair of other papers – Berger, Dew-Becker, and Giglio (2020) and Dew-Becker and Giglio (2020) – I analyze with coauthors the behavior of aggregate and firm-specific stock return volatility. This section compares the results from those papers to the results obtained here for aggregate and firm-specific stock return skewness and third moments.

First, the various volatility and skewness series are not particularly jointly correlated. Their correlation matrix is reported below.

**Correlations between volatility and skewness measures**

| | Agg. vol. | Agg. skew | Idio. vol. | Idio. skew |
|---|---|---|---|---|
| Agg. vol. | 1 | | | |
| Agg. skew | 0.33 | 1 | | |
| Idio. vol. | 0.56 | 0.18 | 1 | |
| Idio. skew | -0.44 | -0.25 | -0.26 | 1 |

Note: time trends removed from the skewness indexes.

The largest correlation is between aggregate and idiosyncratic volatility, but that is still only 0.56, indicating that only 32 percent of their variation is shared. The next largest correlation is between aggregate volatility and idiosyncratic skewness, at -0.44. That is consistent with the fact that both are cyclical (and in opposing directions). Aggregate and idiosyncratic skewness, on the other hand, are actually negatively correlated. Overall, the first principal component explains 50 percent of the variation in the four series. So while they share some common variation, half remains independent of that common factor.

Even though they are not very strongly correlated, the series are similar in that they all have relatively weak autocorrelations. Shocks to volatility and skewness at the aggregate and firm levels both die out within about one year.

As to simple cyclicality, aggregate volatility is procyclical, while firm-specific volatility is acyclical. That result stands in contrast to skewness, where it is the firm-specific component that is significantly cyclical (procyclical, in this case), while the aggregate component is acyclical. Cyclical variation in risk at the aggregate level is therefore characterized by changes in the second moment, while cyclical variation in risk at the firm level is due to changes in the third moment.

The VAR results obtained here are very similar to those in Berger, Dew-Becker, and Giglio (2020). That paper finds that in a VAR with option-implied second moments only, second moment shocks appear to have a negative effect, but that disappears once the *realized* second moment is included, just as here. In both cases, shocks to option-implied moments are important for forecasting future realized moments – they are true risk shocks – but those shocks have only minimal effects on the business cycle. So for both of the cyclical patterns in risk – aggregate second moments and firm-specific third moments – it is actually the *realized* moments that are relevant drivers of the economy.

In the end, the results on realized skewness and volatility are consistent with models in which aggregate output is a concave function of micro shocks. Dew-Becker, Tahbaz-Salehi, and Vedolin (2020), motivated by Baqaee and Farhi (2020), analyze a model in which concavity arises from complementarity in a production network, while Ilut, Kehrig, and Schneider (2018) study a model where individual firms have concave responses to shocks. These types of models can match the procyclicality of skewness, the countercyclicality of volatility, and the fact that realized rather than implied moments are the driving force in VARs.[11]

# 8 Conclusion

This paper analyzes forward-looking skewness at the firm and aggregate level. The skewness of economic shocks has been of growing interest. Models of time-varying disaster risk emphasize time-varying aggregate skewness, while more recent work on time-varying micro dispersion emphasizes the cyclicality of skewness in cross-sectional shocks.

The innovation of this paper is to provide a real-time measure of conditional skewness that is available at high-frequency. Past work has measured skewness based on realized shocks and typically only at the annual frequency.

Aggregate skewness is essentially acyclical, but firm-level skewness, on the other hand, is strongly procyclical. In recessions, the left tail of outcomes becomes much longer. The business cycle also explains a large fraction of the variation in skewness – the correlation of firm skewness with capacity utilization is over 50 percent, for example. These results are consistent with models emphasizing cyclical variation in firm specific risk, such as those of Christiano, Motto, and Rostagno (2014) and Salgado, Guvenen, and Bloom (2020).

Finally, I estimate forecasting regressions and a VAR including both implied and realized third moments. In both cases, realized third moments drive out implied third moments, indicating that the relationship between skewness and output comes from the realization of skewed micro shocks, rather than agents' expectations about the future. That is consistent with recent models emphasizing endogenous aggregate skewness through a concave response of the economy to micro shocks.

[11]Kozlowski, Veldkamp, and Venkateswaran (2020) propose a model that would give another potential channel through which realized skewness could drive the business cycle: after a month with negative realized skewness, agents update their beliefs and act as though skewness will be lower going forward. That is still a model, though, in which it is fundamentally beliefs that are the driver (just in that case endogenous), so one would expect that it would be option-implied, not realized, skewness that would be the relevant shock.

# References


**Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan**, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *Review of Financial Studies*, 2003, *16* (1), 101–143.

**Barro, Robert J.**, "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 2006, *121(3)*, 823–866.

**Bollerslev, Tim and Viktor Todorov**, "Time-varying jump tails," *Journal of Econometrics*, 2014, *183* (2), 168–180.

**Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera**, "Asymmetric business-cycle risk and social insurance," 2018. Working paper.

**Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu**, "Have Individual Stocks Become More Volatile? an Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, 2001, *56* (1), 1–43.

**Christiano, Lawrence J., Roberto Motto, and Massimio Rostagno**, "Risk Shocks," *American Economic Review*, 2014, *104(1)*, 27–65.

**Conrad, Jennifer, Robert F Dittmar, and Eric Ghysels**, "Ex ante skewness and expected stock returns," *The Journal of Finance*, 2013, *68* (1), 85–124.

**Dew-Becker, Ian and Stefano Giglio**, "Cross-sectional uncertainty and the business cycle: Evidence from 40 years of options data," 2020. Working paper.

___ **and** ___ , "Cross-sectional uncertainty and the business cycle: Evidence from 40 years of options data," 2020. Working paper.

**Ferreira, Thiago R.T.**, "Stock market cross-sectional skewness and business cycle fluctuations," 2018. Working paper.

**Gabaix, Xavier**, "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 2012, *127(2)*, 645–700.

**Gilchrist, Simon and Egon Zakrajsek**, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102(4)*, 1692–1720.

**Gourieroux, Christian, Alain Monfort, Sarah Mouabbi, and Jean-Paul Renne**, "Disastrous Defaults," 2020. Working paper.

**Gourio, Francois**, "Disaster Risk and Business Cycles," *American Economic Review*, 2012, *102(6)*, 2734–2766.

___ , "Credit Risk and Disaster Risk," *American Economic Journal: Macroeconomics*, 2013, *5(3)*, 1–34. Working paper.

**Guvenen, Fatih, Serdar Ozkan, and Jae Song**, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621–660.

**Hamilton, James D**, "Why you should never use the Hodrick-Prescott filter," *Review of Economics and Statistics*, 2018, *100* (5), 831–843.

**Harmenberg, Karl and Hans Henrik Sievertsen**, "The Labor-Market Origins of Cyclical Income Risk," 2017. Working paper.

**Ilut, Cosmin, Matthias Kehrig, and Martin Schneider**, "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News," *Journal of Political Economy*, 2018, *126* (5), 2011–2071.

**Jondeau, Eric, Qunzi Zhang, and Xiaoneng Zhu**, "Average skewness matters," *Journal of Financial Economics*, 2019, *134* (1), 29–47.

**Kozhan, Roman, Anthony Neuberger, and Paul Schneider**, "The skew risk premium in the equity index market," *The Review of Financial Studies*, 2013, *26* (9), 2174–2203.

**Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran**, "The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation," 2020. Working paper.

**Merton, Robert C**, "On the pricing of corporate debt: The risk structure of interest rates," *The Journal of finance*, 1974, *29* (2), 449–470.

**Oh, Sangmin and Jessica A. Wachter**, "Cross-Sectional Skewness," 2019. Working paper.

**Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom**, "Skewed Business Cycles," 2020. Working paper.

**Segal, Gill, Ivan Shaliastovich, and Amir Yaron**, "Good and bad uncertainty: Macroeconomic and financial market implications," *Journal of Financial Economics*, 2015, *117* (2), 369–397.

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

Table 1: **Forecasting realized skewness**

| | Realized skew | Realized skew | Monthly skew | Monthly skew | Jumps | Diffusive skew |
|---|---|---|---|---|---|---|
| *Panel A: Firm-level* | | | | | | |
| $SKEW_{firm,t}$ | 0.27*** | 0.12*** | 0.31*** | 0.20*** | 0.01 | 0.25*** |
| | [0.02] | [0.02] | [0.09] | [0.09] | [0.10] | [0.02] |
| $RSKEW_{firm,t}$ | | 0.46*** | | | | |
| | | [0.04] | | | | |
| $RSKEW_{firm,t}$ (monthly) | | | | 0.25*** | | |
| | | | | [0.05] | | |
| Constant | −0.13*** | −0.08*** | −0.08 | −0.08 | 0.04*** | −0.17*** |
| | [0.02] | [0.02] | [0.06] | [0.06] | [0.01] | [0.01] |
| *Panel B: S&P 500* | | | | | | |
| $SKEW_{S\&P,t}$ | 0.39*** | 0.39*** | | | 0.03 | 0.36*** |
| | [0.06] | [0.06] | | | [0.05] | [0.02] |
| $RSKEW_{S\&P,t}$ | | 0.01 | | | | |
| | | [0.05] | | | | |
| Constant | 0.10 | 0.10 | | | 0.22** | −0.11*** |
| | [0.11] | [0.11] | | | [0.10] | [0.04] |
| *Panel C: Idiosyncratic* | | | | | | |
| $SKEW_{idio,t}$ | 0.22*** | 0.16*** | | | 0.04** | 0.19*** |
| | [0.02] | [0.03] | | | [0.17] | [0.02] |
| $RSKEW_{idio,t}$ | | 0.26*** | | | | |
| | | [0.05] | | | | |
| Constant | −0.20*** | −0.15*** | | | 0.08*** | −0.28*** |
| | [0.02] | [0.02] | | | [0.01] | [0.02] |

**Note:** Results from regressions of realized skewness on lagged implied and realized skewness. Each column in each section represents a different regression. Monthly skew is the cross-sectional skewness of firm-level monthly returns. Jumps are calculated from cubed daily returns, while diffusive skew is due to the product of returns with changes in implied volatility. * indicates significance at the 10% level, %% 5%, and *** 1%.

Table 2: **Calibration moments**

| | Mean | SD | AC(1) | AC(3) | AC(12) |
|---|---|---|---|---|---|
| *Panel A: Raw data* | | | | | |
| $SKEW_{firm,t}$ | -0.54 | 0.46 | 0.93 | 0.87 | 0.75 |
| | [0.09] | [0.04] | [0.01] | [0.03] | [0.08] |
| $SKEW_{S\&P,t}$ | -1.68 | 0.85 | 0.94 | 0.87 | 0.74 |
| | [0.18] | [0.12] | [0.02] | [0.05] | [0.11] |
| $SKEW_{idio,t}$ | -0.51 | 0.73 | 0.87 | 0.76 | 0.56 |
| | [0.13] | [0.09] | [0.03] | [0.06] | [0.11] |
| *Panel B: Time trend removed* | | | | | |
| $SKEW_{firm,t}$ | | 0.26 | 0.79 | 0.59 | 0.25 |
| | | [0.02] | [0.04] | [0.07] | [0.11] |
| $SKEW_{S\&P,t}$ | | 0.55 | 0.87 | 0.71 | 0.44 |
| | | [0.04] | [0.03] | [0.06] | [0.11] |
| $SKEW_{idio,t}$ | | 0.55 | 0.78 | 0.58 | 0.25 |
| | | [0.07] | [0.06] | [0.09] | [0.09] |

**Note:** Simple calibration moments for the three skewness series. The top panel reports statistics for the raw data, while the bottom panel removes a time trend from each. Standard errors reported in brackets are from a block bootstrap with two-year blocks.

Table 3: **Correlations with skewness indexes**

| | Firm | Firm detrended | S&P | S&P detrended | Idio. | Idio. detrended |
|---|---|---|---|---|---|---|
| Industrial prod. | 0.17 | 0.21 | -0.01 | -0.07 | 0.25 | 0.24 |
| Employment | 0.17 | 0.21 | 0.00 | 0.02 | 0.22 | 0.18 |
| Unemployment | -0.22 | -0.25 | 0.02 | 0.04 | -0.28 | -0.23 |
| CBO output gap | 0.19 | 0.22 | -0.12 | -0.05 | 0.29 | 0.26 |
| NBER rec. ind. | -0.12 | -0.14 | 0.09 | 0.15 | -0.17 | -0.17 |
| Capacity util. | 0.51 | 0.15 | 0.15 | -0.09 | 0.55 | 0.20 |
| Empl. growth | 0.17 | 0.09 | 0.02 | -0.16 | 0.22 | 0.17 |
| IP growth | 0.34 | 0.14 | 0.07 | -0.18 | 0.39 | .021 |

**Note:** Each column reports pairwise correlations with one of the skewness indexes, $SKEW_{firm,t}$ etc. Columns marked "detrended" have an HP-filtered trend removed from the skewness index, as to industrial production, employment, and unemployment.

Table 4: **Credit spread regressions**

| | | | |
|---|---|---|---|
| $SKEW_{firm,t}$ | $-1.05^{***}$ | | |
| | [0.22] | | |
| $IV_{firm,t}$ | $6.49^{***}$ | | |
| | [2.10] | | |
| $SKEW_{S\&P,t}$ | | $-0.23^{***}$ | |
| | | [0.07] | |
| $IV_{S\&P,t}$ | | $10.99^{***}$ | |
| | | [2.93] | |
| $M3_{firm,t}$ | | | $-21.26^{***}$ |
| | | | [2.99] |
| $M3_{S\&P,t}$ | | | 7.39 |
| | | | [5.63] |
| $R^2$ | 0.60 | 0.54 | 0.65 |

**Note:** Estimates of regressions of corporate credit spreads on option-implied skewness and volatility. $M3_{firm,t}$ is the firm-level implied third moment. Standard errors are from a block bootstrap with two-year blocks.

Table 5: **Forecasting regressions**

| Dependent variable: | Employment growth | Employment growth | IP growth | IP growth |
|---|---|---|---|---|
| $M3_{idio}/1000$ | $0.10^{**}$ | $-0.04$ | $0.31^{*}$ | -0.28 |
| | [0.04] | [0.05] | [0.18] | [0.25] |
| $R3_{idio}/1000$ | | $0.32^{***}$ | | $1.21^{***}$ |
| | | [0.06] | | [0.37] |

**Note:** Estimates of regressions of employment and IP growth on lagged implied and realized third moments. Eah regression includes four lags of the dependent variable and four lags of the third moments. The numbers reported in the table correspond to the sum of the coefficients on the four lags.

Figure 1: Skewness time series

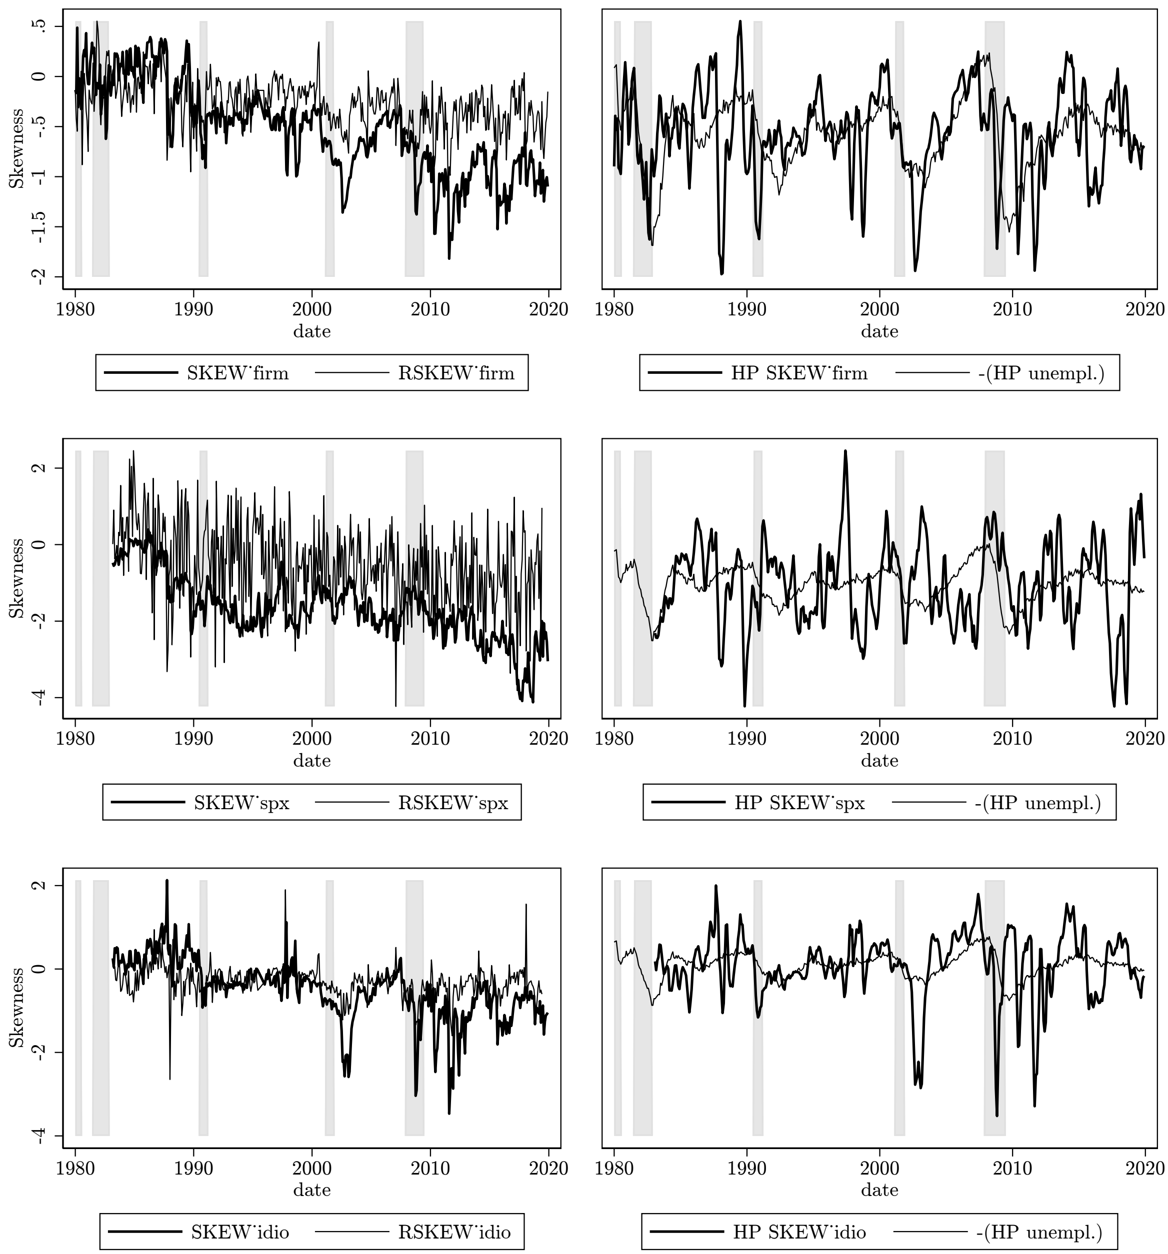


**Note:** The left-hand panels plot option-implied and realized skewness. The right-hand panels plot Hodrick–Prescott detrended skewness against detrended industrial production (both with a smoothing parameter of 129,600). Gray bars indicate NBER-dated recessions.

Figure 2: Auto- and cross-correlations

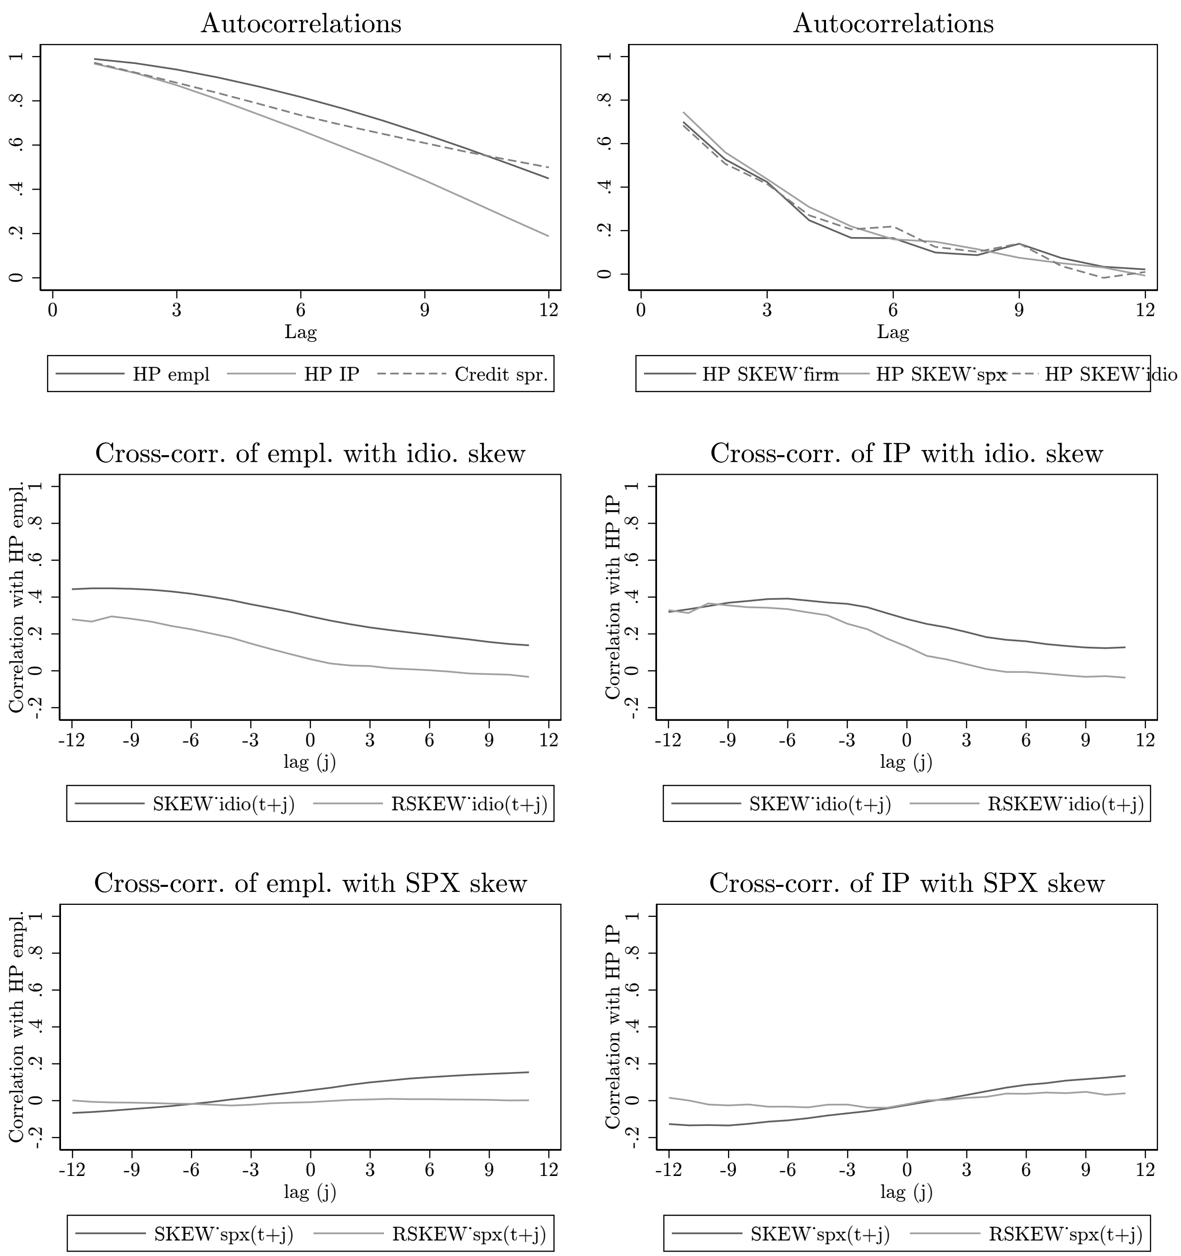


**Note:** The top two panels plot autocorrelations of various series. The remaining panels plot cross-correlations. The lines represent the correlation of employment or IP on date $t$ with skewness on date $t - j$, where $j$ is represented by the x-axis.

Figure 3: Annual measures of realized skewness

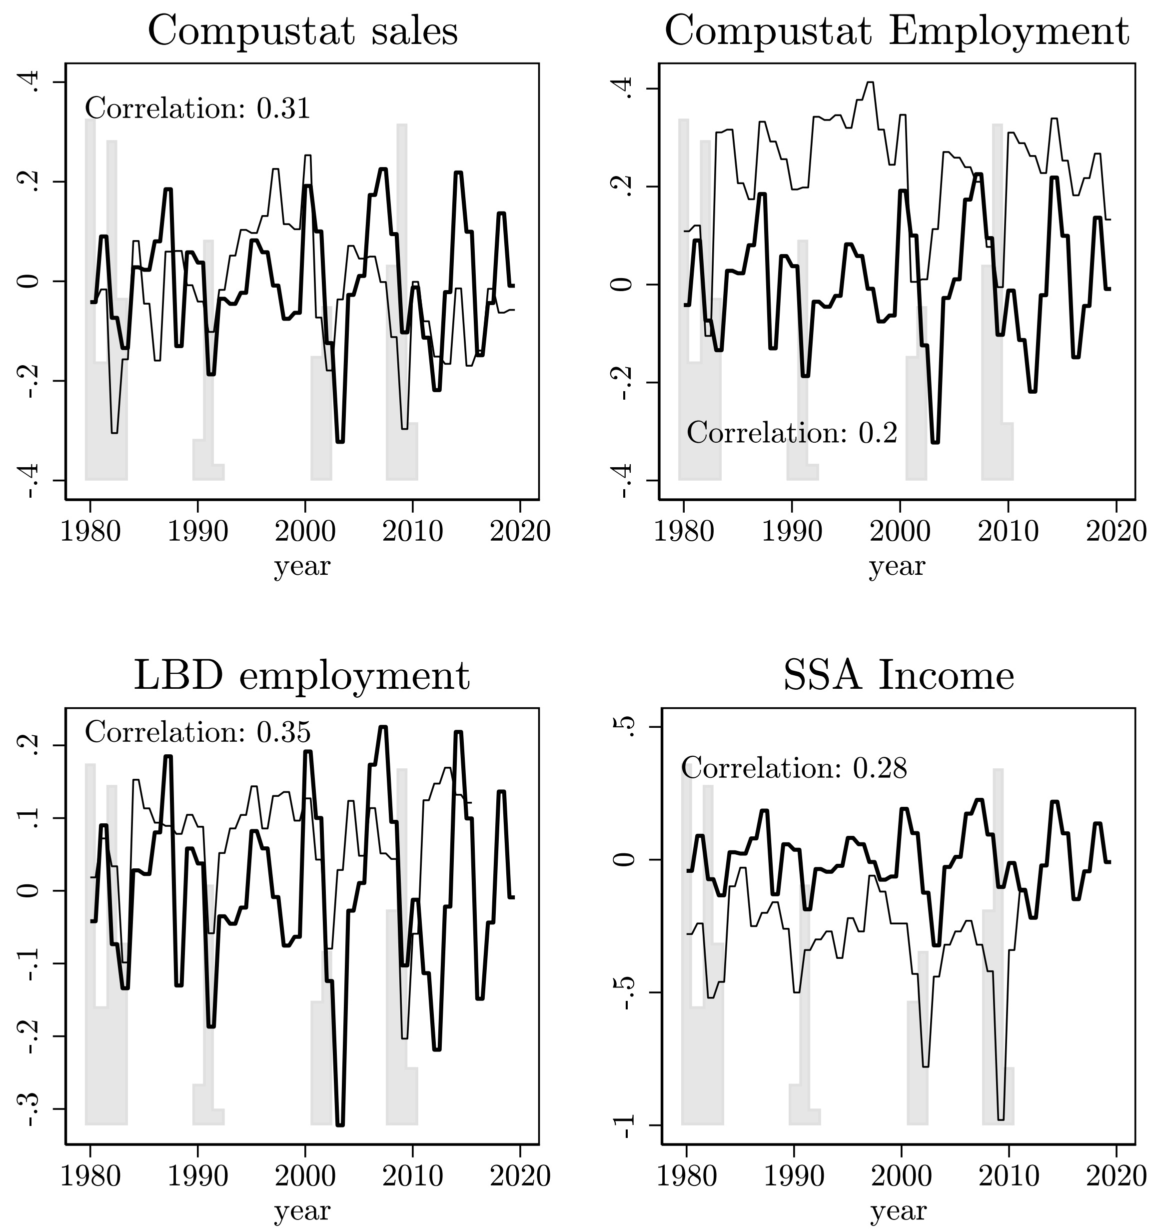


**Note:** The first three panels plot HP-detrended firm skewness (the heavy line) against a measure of annual skewness (the lighter line). Gray bars indicate the fraction of months in each year were in recessions, according to the NBER. The bottom-right panel plots firm skewness against firm VIX asymmetry, as defined in the text. VIX asymmetry is rescaled so that it has the same standard deviation as firm skewness.

Figure 4: Contributions to total firm skewness

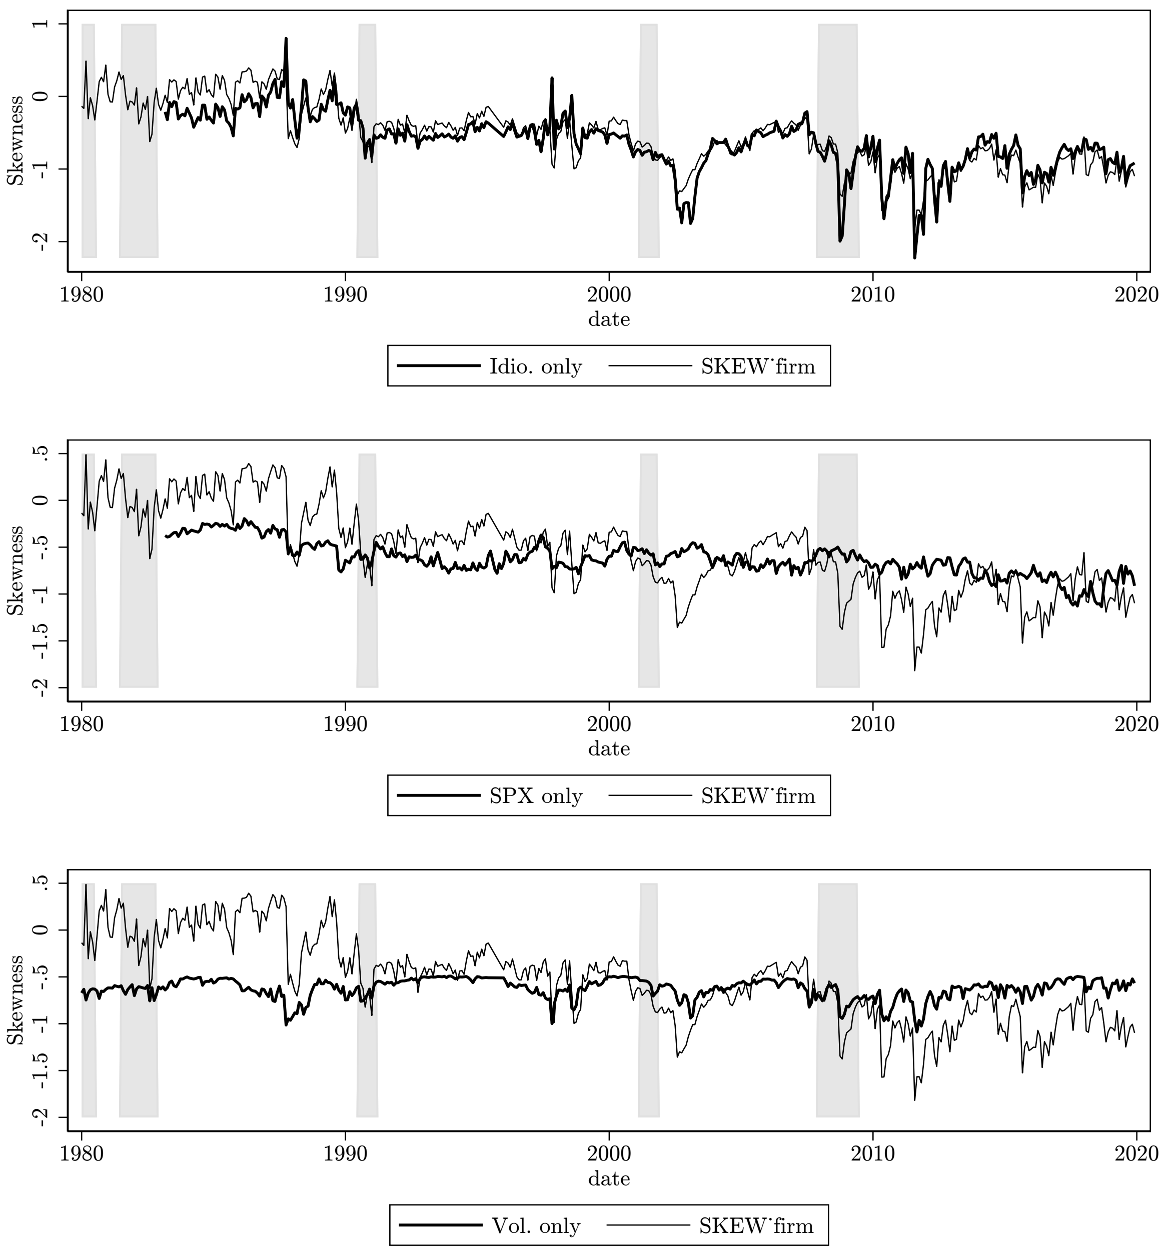


**Note:** Each panel plots total firm skewness when just one of its three components varies over time, with the other three left at their unconditional means.

Figure 5: Impulse response functions

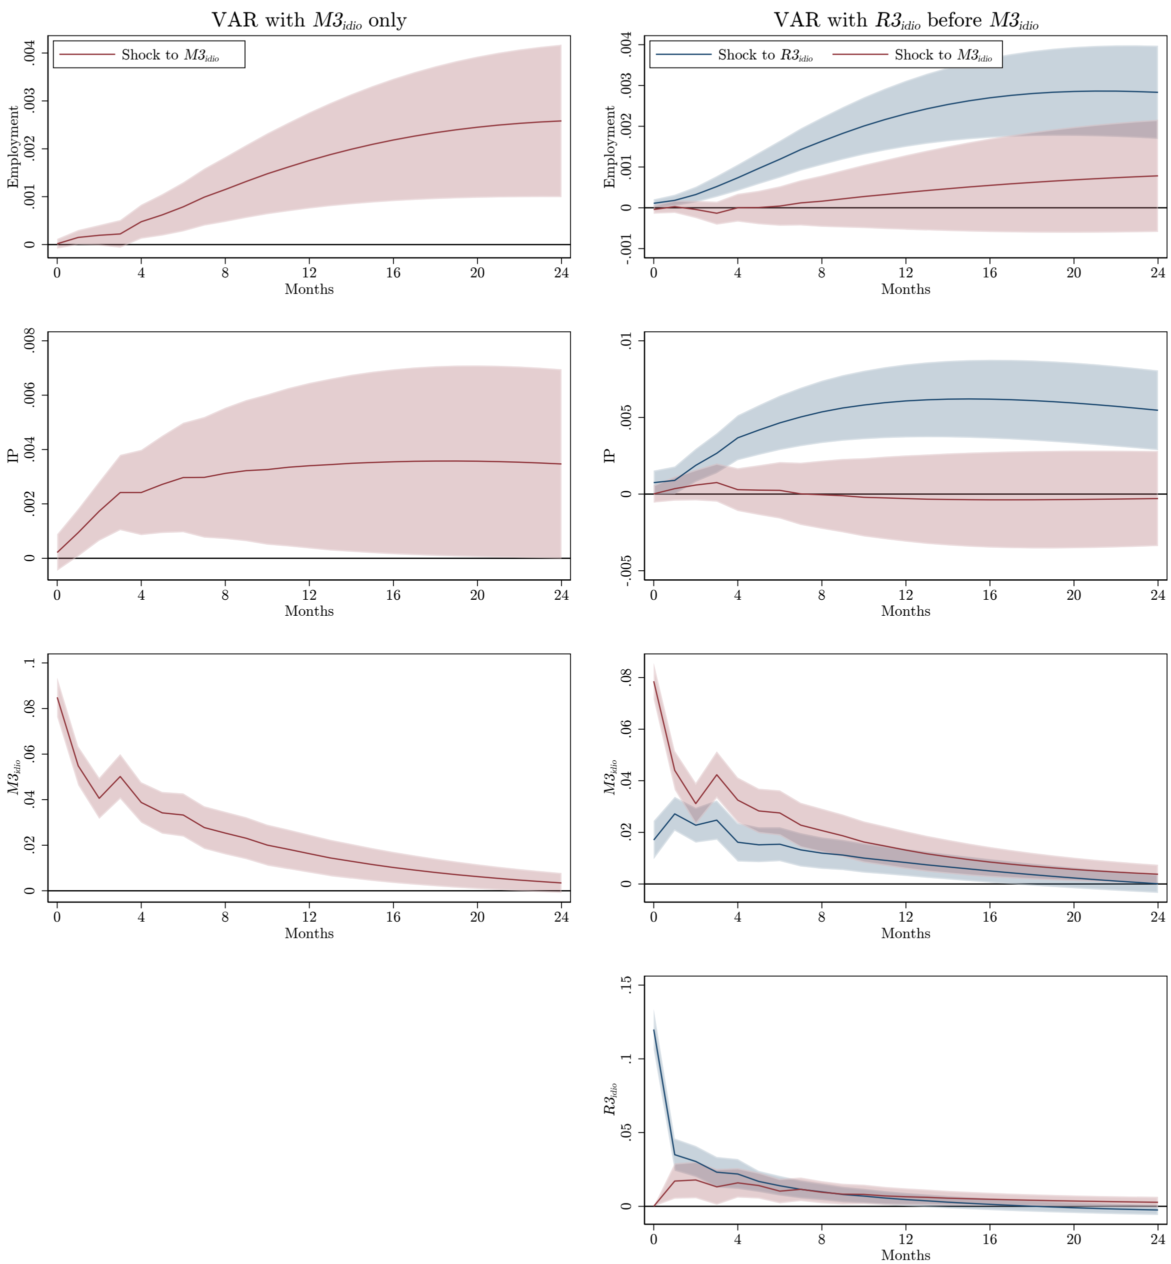


**Note:** Impulse response functions with 90-percent confidence bands. The two columns represent separate VARs. The left includes $M3_{idio}$ only (along with the macro variables), while the right includes both $M3_{idio}$ and $R3_{idio}$, with $R3_{idio}$ ordered first in the Cholesky factorization.

Figure 6: Forecast error variance decompositions

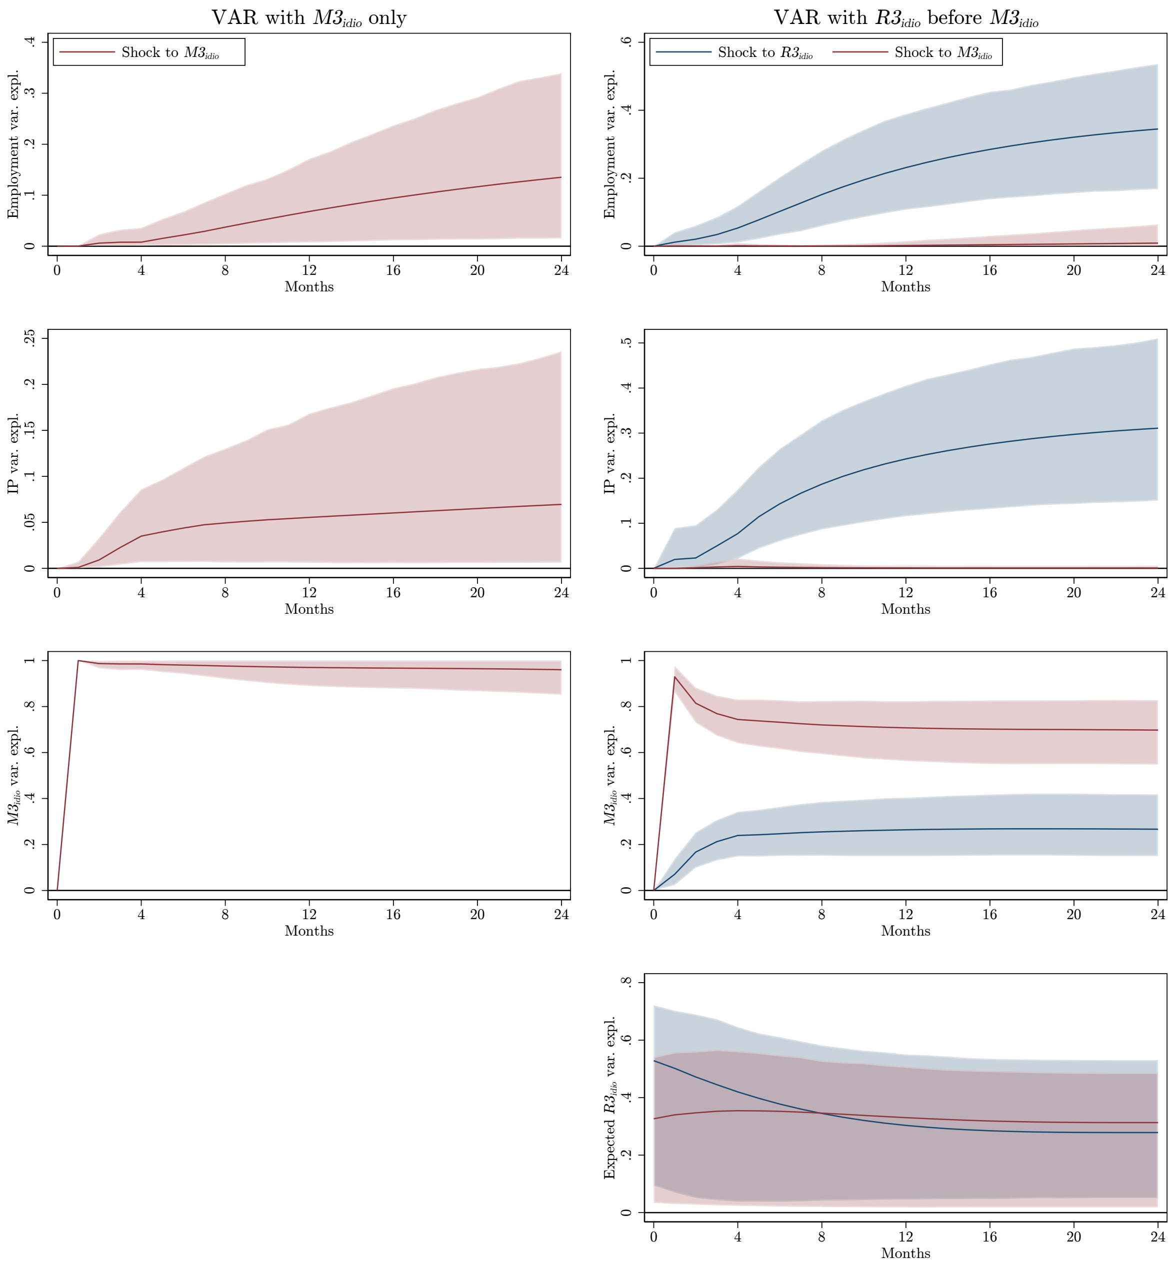


**Note:** Forecast error variance decompositions with 90-percent confidence bands. See figure 5. The decomposition for expected $R3_{idio}$ is based on the VAR-implied expectation of $R3_{idio}$ over the next 24 months.

## A.1 Fitting option prices to get risk-neutral moments

This section describes how we fit option prices in order to calculate risk-neutral moments at the firm and market level. We first describe the firm specification, which is more general. The method for the S&P 500 is a special case with restrictions imposed.

The key assumption is that the log price of a stock option approaches an asymptote that is linear in the log strike, as in Bollerslev and Todorov (2014). That is, the price of an out-of-the-money option on stock $i$ at log strike $k$ is

$$p_{i,k} = a_i + b_i |k| + b_i^+ \max(k, 0) + f_i(k) + \varepsilon_{i,k} \tag{A.1}$$

where $a_i$, $b_i$, and $b_i^+$ are coefficients. $f_i(k)$ is some function of the strike, and $\varepsilon_{i,k}$ is a measurement error that is uncorrelated across strikes and firms. $f_i(k)$ captures the deviation of the log option price from the linear asymptotes. The assumption that $p_{i,k}$ has a linear asymptote is imposed by assuming that $f_i(k) \to 0$ as $k \to \pm\infty$.

The key step is estimating the parameters $a_i$, $b_i$, and $b_i^+$ and the function $f_i$. We treat them as having a joint Normal distribution, so this is a standard Gaussian process regression. For the parameters $a_i$, $b_i$, and $b_i^+$, we assume that

$$a_i \sim N\left(a, \sigma_a^2\right) \tag{A.2}$$

$$b_i \sim N\left(b, \sigma_b^2\right) \tag{A.3}$$

$$b_i^+ \sim N\left(b^+, \sigma_{b^+}^2\right) \tag{A.4}$$

$a$, $\sigma_a^2$, etc. are hyperparameters that can be estimated by MLE. Conditional on $a$, $b$, and $b^+$, the parameters $a_i$, $b_i$, and $b_i^+$ are independent across firms.

To help absorb the cross-sectional variation in those parameters, we first normalize all underlyings to have a price of 1. That is, if we have an observation with an underlying price of $S_{i,t}$, a strike of $\tilde{K}$, and an option price of $\tilde{P}_{\tilde{K},i,t}$, then we simply divide through by $S_{i,t}$, to get

$$K_t = \tilde{K}_t / S_{i,t} \tag{A.5}$$

$$P_{K,i,t} = \tilde{P}_{\tilde{K},i,t} / S_{i,t} \tag{A.6}$$

This is simply a renormalization in terms of a different numeraire (units of the underlying instead of dollars), which will be valid under any standard model.

To fully specify the likelihood of the data, under the assumption that $f$ is Gaussian, we need to define its mean and covariance matrix (i.e. its covariance kernel). We model it as following a Brownian bridge stretched to cover the entire real line and with a jump at $k = 0$. The stretched Brownian bridge assumption implies that $f_i(k) \to 0$ almost surely as $k \to \pm\infty$. The assumption of a jump at zero allows the price function to be continuous at zero even though the linear part of the model, $a_i + b_i |k| + b_i^+ \max(k, 0)$ has a discontinuity (i.e. $f$ is able to smooth out the discontinuity with an offsetting jump).

More formally, similar to above,

$$
\begin{aligned}
E[f_i(k)] &= 0 \forall i, k && \text{(A.7)} \\
\operatorname{cov}(f_i(k), f_j(m)) &= \sigma_B^2 (\min(L(k), L(m)) - L(k) L(m)) && \text{(A.8)} \\
&\quad + \sigma_D^2 \delta_{k>0} \delta_{m>0} (1 - L(k))(1 - L(m)) && \text{(A.9)} \\
&\quad + \delta_{i=j} \begin{pmatrix} \sigma_{\hat{B}}^2 (\min(L(k), L(m)) - L(k) L(m)) \\ + \sigma_{\hat{D}}^2 \delta_{k>0} \delta_{m>0} (1 - L(k))(1 - L(m)) \end{pmatrix} && \text{(A.10)}
\end{aligned}
$$

where $L(k) \equiv \frac{\exp(sk)}{\exp(sk)+1}$ is a scaled logistic function and $s$, $\sigma_B^2$, $\sigma_D^2$, $\sigma_{\hat{B}}^2$, and $\sigma_{\hat{D}}^2$ are hyperparameters. $\delta_z$ is an indicator function equal to 1 if $z$ is true and 0 otherwise.

The specification for the distribution of $f_i$ is such that it is the sum of two Brownian bridges: one that is common to all firms, and a second that is specific to firm $i$. The logistic function $L$ is what stretches the Brownian bridge to cover the full real line, while the parameter $s$ determines the rate at which the variance of $f$ converges to zero as $|k|$ grows.

Intuitively, the method here incorporates information across firms on a given date. That information is shared through the parameters that are common to all firms: $a$, $b$, $b^+$, and also the covariance kernel of $f_i$. Firms also have scope for independent variation trough the random components in $a_i$, $b_i$, $b_i^+$, and $f_i(k)$.

For the S&P 500, we run the estimation separately from other firms. The $i$-specific hyperparameters above are therefore eliminated: $\sigma_a^2$, $\sigma_b^2$, $\sigma_{b+}^2$, $\sigma_{\hat{B}}^2$, and $\sigma_{\hat{D}}^2$. The only unknown parameters are then $a$, $b$, $b^+$, $s$, $\sigma_B^2$, and $\sigma_D^2$.

### A.1.1 Data filters and estimation

We estimate the model above using daily closing prices for options on individual stocks and the S&P 500 index. On each day, we take prices for out-of-the-money options that are at least three ticks above zero and where the strike is less than five at-the-money standard deviation units from the underlying price. The maturity must be greater than 7 days and less than 75. We weight observations in the estimation by the underlying firm's market value divided by the bid-ask spread for the log price (i.e. the log ask minus the log bid). When we do the estimation for the S&P 500, we require at least six total prices on each date, with at least two strikes above and below the underlying price.

For the S&P 500, we estimate all of the hyperparameters on each date. When using all firms, estimating the hyperparameters is slow and does not appear to have a substantial impact on the results. We therefore fix the variance hyperparameters at values that we found to be optimal on a representative date. The parameters $a$, $b$, and $b^+$ are re-estimated, as is the residual variance for the fitting errors ($\text{var}(\varepsilon_{i,k})$). The fact that the variance hyperparameters are fixed does not mean that the $f_i$ or $a_i$, etc. are fixed. They are fitted each day, but the variances are fixed. This is mathematically equivalent to running the Kalman filter with fixed variance parameters, just updating the state estimates. The estimation of the $f_i$ and $a_i$ is linear conditional on the variances, and hence has linear closed-form expressions.

### A.1.2 Extensions to the baseline specification

Figure A.1 plots our measure of skewness first against skewness calculated using the formula used by the CBOE. The results are essentially identical for both firms and the S&P 500. Next, it compares our skewness measure to the CBOE's reported values for the S&P 500. Finally, it compares the results using data from CME S&P 500 futures options to CBOE SPX options as reported by Optionmetrics. In all cases, the results are quantitatively and qualitatively similar.

To give further context to the data, the top panel of figure A.2 plots the average daily number of firms used within each month, while the bottom panel plots the faction of total US equity market capitalization covered by those firms. Both the number of firms and their market capitalization has trended up over time. The fraction of total equity market capitalization covered rises from about one fourth in the very earliest part of the sample up to 60 percent by the end.

The Optionmetrics sample is restricted to the top 100 firms by market capitalization on each date for the sake of consistency with the Berkeley Options Database, for which we only have the largest firms. To evaluate the impact of that restriction on the results, we also re-estimated the option-implied skewness using the top 200 firms. With 200 firms, skewness has a correlation with the baseline measure of 99.4 percent, confirming that any difference from including smaller firms, at least under value weighting, is quantitatively small. When we use 200 firms in the Optionmetrics sample, the average fraction of total market capitalization covered in the Optionmetrics period rises from 48 to 60 percent.

Figure A.1: Alternative skewness series

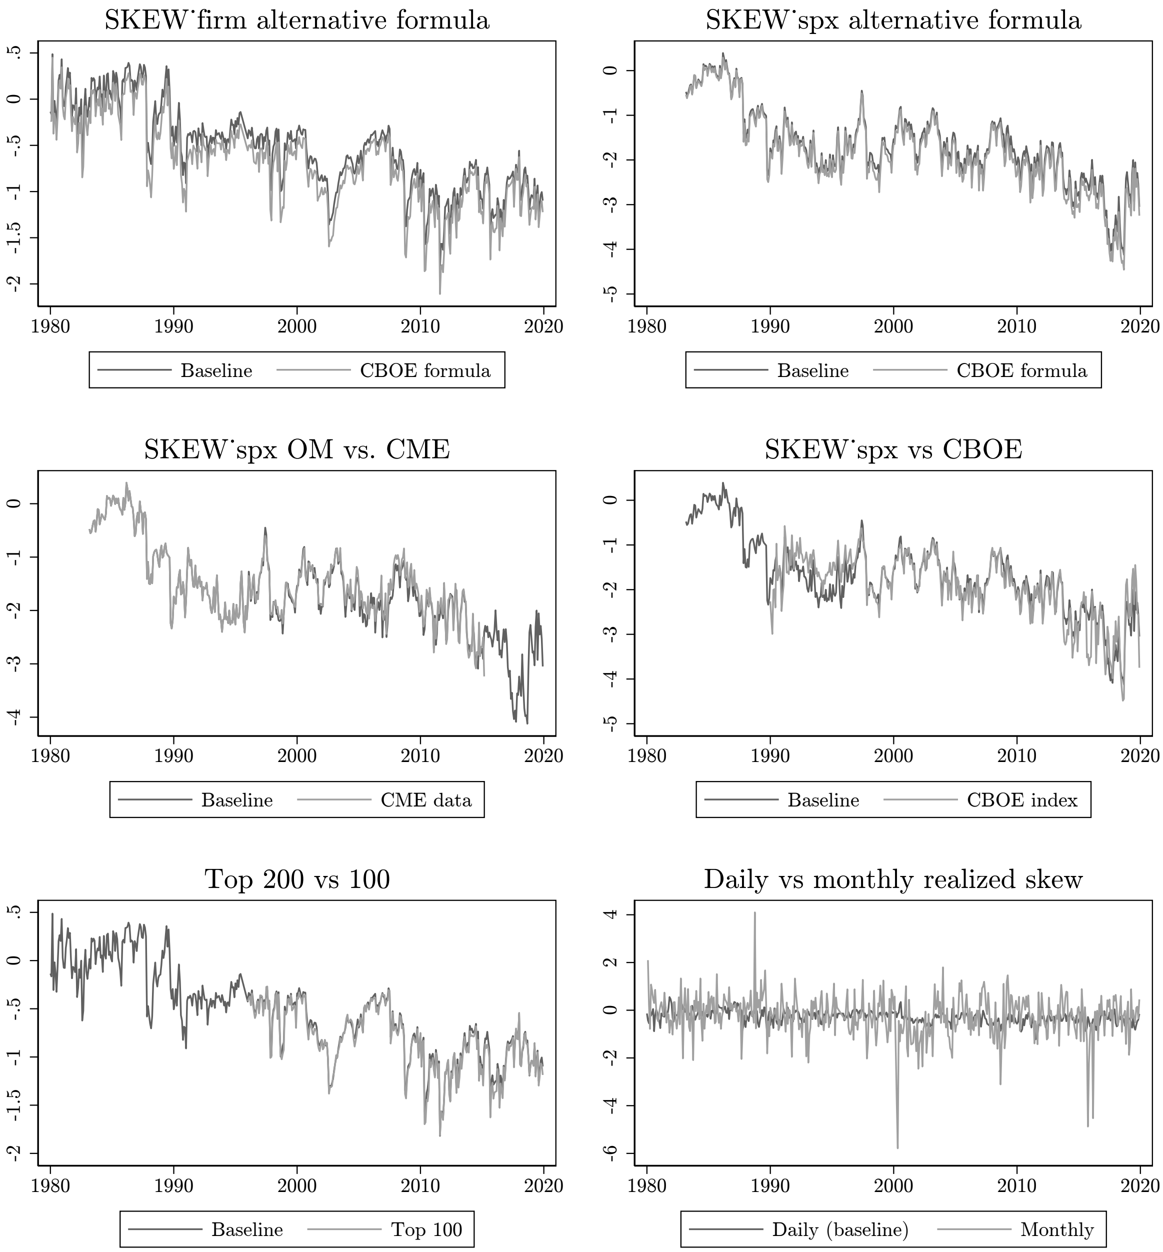


**Note:** Baseline skewness series and alternatives.