

Real-time forward-looking skewness over the business cycle

Ian Dew-Becker*

June 3, 2021

Abstract

This paper measures option-implied skewness at the firm and market levels between 1980 and 2020. It reports the first real-time measure of conditional micro skewness. There are three key results: 1. Firm skewness is significantly procyclical, while market skewness is acyclical; 2. Firm-level skewness leads the business cycle, and it is strongly linked to credit spreads, suggesting one potential causal channel; 3. Firm skewness is significantly, and not mechanically, correlated with aggregate volatility, implying that there is a common shock driving them both, which is also linked to the business cycle.

1 Introduction

Background

A growing literature has developed evidence that the conditional distributions of economic variables, at both the aggregate and micro levels, varies over time. There is work on shocks to volatility and uncertainty and the aggregate level, time-varying disaster risk, and estimates of cross-sectional variance and skewness of outcomes at the micro level.¹ The evidence on micro outcomes has come from analyses of the distribution of *realized* outcomes, typically at just an annual frequency. For example, research has measured the skewness

*Northwestern University and NBER. I appreciate helpful comments from Stefano Giglio, Simon Oh, Laura Veldkamp, Matthias Kehrig, and Thiago Ferreira.

¹On uncertainty, see, among many others, Bloom et al. (2018), Adrian, Boyarchenko, and Giannone (2019), and Berger, Dew-Becker, and Giglio (2020). On disasters and skewness, see Barro (2006), Gabaix (2012), Gourio (2012, 2013), and Wachter (2013). At the micro level, see also Storesletten, Telmer, and Yaron (2004), Guvenen, Ozkan, and Song (2014), and Kehrig (2015).

of the distribution of annual firm-level employment growth and annual individual income growth.²

But in models in which time-varying risk – whether through volatility or skewness – is a driving force (e.g. Christiano, Motto, and Rostagno (2016) and 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 moments of shocks that will be realized tomorrow, rather than the moments of the shocks that were already realized today. And realizations are always equal to expectations plus an error, meaning that they are noisy relative to the true conditional expectation. The use of annual data makes estimation and measurement difficult, since there are so few observations and they come at a low frequency relative to, for example, the length of the typical recession.³ Finally, policymakers benefit from having measures available in real time that can be used for actual decisionmaking. Estimation and testing of models based on time-varying risk, whether volatility or skewness, and also using measures of risk as inputs to policy, therefore requires high-frequency measures of conditional moments.

Contribution

This paper’s fundamental contribution is to develop real-time forward-looking measures of the *conditional* skewness of both firm-level and aggregate shocks from option prices.⁴ The measures are available since 1980 at up to the daily frequency (or even intraday) and without any reporting delay, making them particularly attractive to policymakers. The sample is of comparable length to those used in other studies and enables analysis of the behavior of micro and aggregate skewness across six business cycles.

To validate the skewness indexes, I show that they have significant predictive power for future realized skewness – i.e. the skewness of the shocks that are actually realized – for both stock returns and a novel text-based skewness measure from news articles. Those results hold even after controlling for lagged realized skewness. The option-implied skewness measures are also significantly positively correlated with annual measures of realized cross-sectional skewness of employment and income growth.

Results

²For variance, Campbell et al. (2001), Bloom (2009), Herskovic et al. (2016), and Bloom et al. (2018) all examine measures of realized dispersion rather than conditional variances. Guvenen, Ozkan, and Song (2014), Ferreira (2018), Jondeau, Zhang, and Zhu (2019), Oh and Wachter (2019), and Salgado, Guvenen, and Bloom (2020) study skewness. The evidence on stock returns is the exception to the general pattern that research has typically looked at low frequencies.

³See Guvenen, Ozkan, and Song (2014), Harmenberg and Sieversten (2017), Busch et al. (2018), Ilut, Kehrig, and Schneider (2018), and Salgado, Guvenen, and Bloom (2020).

⁴While the CBOE reports a market-level skewness index, it has received little attention, and this paper is novel for combining market skewness with firm-level information.

The paper’s key finding is that skewness at the *firm level* is significantly procyclical. In each of the past six recessions, firm-level skewness became noticeably more negative. Skewness of the overall stock market, on the other hand is actually slightly above average during NBER-dated recessions, and acyclical overall. Given that the total uncertainty faced by firms is a combination of aggregate and firm-specific shocks, the natural conclusion is that the skewness of firm-specific shocks must be procyclical, and I show formally that they are.

The paper’s second result is that aggregate volatility and firm-level skewness are significantly negatively correlated. Most or all of the jumps down in firm skewness are associated with jumps up in market-level volatility. A mechanical effect can, in principle, cause that link, but it turns out to be insufficient quantitatively (and the correlation between firm skewness and aggregate volatility is far from 1, inconsistent with a purely mechanical explanation). Instead, there appears to be a *common factor* affecting aggregate volatility and firm-level skewness, making both cyclical. That point is sharply identified here due to the high-frequency nature of the data, and cannot be seen in lower frequency realized measures. This result is reminiscent of the analysis in Kozienskis, Orlik, and Veldkamp (2018) showing how shocks to aggregate volatility can drive a range of measures of uncertainty and dispersion.

I next use cross-correlations to examine the lead-lag relationship between firm-level skewness and the business cycle. Skewness very clearly moves first, in the sense that it predicts measures of output better than output predicts skewness. Firm-level skewness also moves approximately contemporaneously with the Philadelphia Federal Reserve’s Leading Indicator index.

A number of papers have proposed channels through which changes in risk might drive the business cycle. In Gourio (2013) and Christiano, Motto, and Rostagno (2014), 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. This is certainly not the only channel, though, through which skewness and output could be related, and the paper does not provide causal evidence on this point.

In a related paper, Dew-Becker and Giglio (2021) show that firm-specific *volatility* has no consistent relationship with the business cycle. The results there and in this paper thus echo and extend those of Guvenen, Ozkan, and Song (2014), who argue that skewness in individual income growth is procyclical, while variance is acyclical. Ilut, Kehrig, and Schneider (2016) develop a model based on concave responses to firm-level shocks that generates

countercyclical aggregate volatility and procyclical cross-sectional skewness, consistent with the empirical results here.⁵

Many of this paper’s results rely on the high-frequency variation made newly available by using option data. Using annual data would not allow one to uncover lead/lag relationships or to see exactly how entries into recessions are associated with high skewness, and this is the only dataset that allows one to disentangle aggregate, firm, and idiosyncratic skewness.

Related literature

In addition to the work described above, this paper is also related to work using stock options to estimate investor views about conditional distributions, such as CBOE’s VIX and SKEW indexes, which measure the conditional volatility skewness of the S&P 500 (see also Kozhan, Neuberger, and Schneider (2013) and Kozlowski, Veldkamp, and Venkateswaran (2020)). Those are both aggregate moments, so Dew-Becker and Giglio (2021) provide a long time-series of firm-level implied volatility. This paper takes that work a step further by measuring firm-level 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.⁶ Finally, note that this paper focuses on skewness – measuring asymmetry – but total downside risk also depends on volatility (e.g. Adrian, Boyarchenko, and Giannone (2019)), discussed in section 4.4.

2 Data and methods

2.1 Option-implied skewness

The goal is to construct an estimate of the conditional skewness of economic outcomes at the aggregate and firm levels. I focus on stock returns as the outcome because stock options can be used to measure investor beliefs about their future distribution. Stock prices are economically relevant because they summarize investor beliefs about the future profitability of US corporations. It would arguably be preferable to have measure of skewness of some more fundamental concept, such as productivity, but productivity is not directly measurable, nor does it have an associated options market to help reveal agents’ probability distributions.⁷

This paper measures skewness throughout as the scaled third moment. For a generic

⁵See also Kozienskias, Orlik, and Veldkamp (2018) and Dew-Becker, Tahbaz-Salehi, and Vedolin (2021) for other models in which time-varying moments are generated through concave responses to structural shocks.

⁶Conrad, Dittmar, and Ghysels (2013) study firm-level option-implied skewness in the context of return forecasting. See also Bakshi, Kapadia, and Madan (2003).

⁷Past work also studies skewness of endogenous outcomes, such as income and employment.

random variable x_{t+1} ,

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

where E_t denotes the statistical expectation conditional on information available at time t .

Asset prices do not in general reveal statistical expectations, but instead “risk-neutral” expectations, which are affected by risk premia. This paper refers to “option-implied expectations”, by which it means those 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. This point is analyzed in detail in section 3.2 and appendix A.2.2.

More than one method of constructing option-implied skewness has been proposed in the literature. I use the formula of Kozhan, Neuberger, and Schneider (2013), denoting option Implied Skewness by IS_t ,

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

where $O_t(K)$ is the price of an out-of-the-money option at strike K on date t . 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.3 shows that the Kozhan, Neuberger, and Schneider (2013) and CBOE measures are highly similar, though not identical.

2.2 Data

I calculate implied skewness at the market level using S&P 500 option prices from the CME for 1983–1995 and the CBOE for 1996–2020. The analysis here is entirely in terms of monthly skewness – using options with a one-month average maturity. That yields comparability with the well known VIX index of monthly volatility and is consistent with the measurement frequency of macroeconomic data.

Option markets are also more liquid at shorter maturities. There is very little volume in 12-month options, for example, especially for individual stocks.⁸ As discussed in appendix

⁸Dew-Becker and Giglio (2021) extensively discuss the relationship between persistence of risk shocks and how they affect investment. Calculating a 12-month skewness is in principle feasible, but would be significantly more difficult (and the analysis here is already far from trivial) and would require making more assumptions and relying more strongly on interpolation and extrapolation. In addition, unlike variance, skewness is not directly comparable across horizons (intuitively, due to the force underlying the central limit theorem that smoothes out distributions).

A.1, I apply standard filters to isolate relatively more liquid options. Section A.3.4 examines differences in liquidity across options and shows they are very small and are unrelated to the patterns in skewness.

This paper’s innovation is to calculate a skewness index for *firm-level* outcomes using options on individual stocks. For 1/1996–12/2020, the data is from Optionmetrics, and for 1/1980–6/1995 from the Berkeley Options Database (note that there is a six-month gap). 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 (e.g. Mayhew and Mihov (2004)). Dew-Becker and Giglio (2021) provide an extensive analysis of the BODB data. It is important to emphasize that the results here are for the largest firms in the economy. That will be more appropriate for some models than others (for example, conditional skewness might be linked to the odds of receiving a granular shock to a large firm, à la Gabaix (2011)).

While data through 2020 is available, the events of the spring of that year in particular are sufficiently extreme as to dominate certain parts of the analysis. I therefore plot and discuss the univariate properties of uncertainty through the end of 2020, but all comparisons with macro variables use data only through the end of 2019.⁹

For the exercises forecasting realized skewness one month ahead, theory requires using implied skewness calculated on the last trading day of the month. For all other purposes, I average skewness over days in the month. The differences are minor.

2.3 Skewness indexes

Denote firm i ’s implied skewness on date t by $IS_{i,t}$. Implied skewness is calculated using equation (2), with the integral constructed using methods described in appendix A.1. Average firm-level skewness on date t is then

$$IS_{firm,t} \equiv \frac{\sum_i IS_{i,t} mkt_{i,t}}{\sum_i mkt_{i,t}} \quad (3)$$

where $mkt_{i,t}$ is the market capitalization of firm i on date t .

Skewness for the overall market – the S&P 500 – is denoted by $IS_{mkt,t}$, and is calculated using equation (2) and S&P 500 option prices.

⁹Prior to 2020, the largest monthly absolute growth rate in employment in the sample is 1.2 log points, while employment fell 14.7 log points in April, 2020.

2.3.1 Idiosyncratic skewness

Most of the analysis just focuses on firm and aggregate skewness, $IS_{firm,t}$ and $IS_{mkt,t}$, but it is also useful to understand the determinants of $IS_{firm,t}$. To do so, I consider the single-factor specification of Campbell et al. (2001),

$$r_{i,t} = r_{mkt,t} + \varepsilon_{i,t} \quad (4)$$

where $r_{i,t}$ is the return on stock i on date t , $r_{mkt,t}$ is the return on the market, and $\varepsilon_{i,t}$ is a residual uncorrelated with $r_{mkt,t}$. As in Campbell et al. (2001), the loading on the market is treated as approximately 1.

Importantly, note that nothing about the calculation of $IS_{firm,t}$ and $IS_{mkt,t}$ relies on that approximation.

It is then straightforward, using the definition of skewness, to show that the contribution of firm-specific shocks to $IS_{firm,t}$ is

$$skew(\varepsilon_{i,t}) + 3 \frac{E[r_{m,t} \varepsilon_{i,t}^2]}{\sigma_{\varepsilon,t}^3} \quad (5)$$

That contribution, which I refer to as idiosyncratic skewness has two sources: skewness in the firm-specific shocks themselves, and covariance between the magnitude of the firm-level shocks, $\varepsilon_{i,t}^2$, and the return on the overall market.¹⁰ Option-implied idiosyncratic skewness is then

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

where $IV_{mkt,t}$ and $IV_{firm,t}$ are option-implied volatilities (measured as in the denominator of (2)) for the S&P 500 and averaged across individual firms.^{11,12}

¹⁰Formally, the $\varepsilon_{i,t}$ need not be independent across firms, so “idiosyncratic” here just means uncorrelated with aggregate returns.

¹¹Equation (6) comes from expanding the skewness formula for (4) to get

$$skew_t(r_{i,t}) = \left(\frac{\sigma_{m,t}^2}{\sigma_{i,t}^2}\right)^{3/2} skew_t(r_{m,t}) + \underbrace{\left(1 - \frac{\sigma_{m,t}^2}{\sigma_{i,t}^2}\right)^{3/2} \left(skew_t(\varepsilon_{i,t}) + 3 \frac{E[r_{m,t} \varepsilon_{i,t}^2]}{\sigma_{\varepsilon,t}^3} \right)}_{\text{Idiosyncratic skewness}} \quad (7)$$

where the $\sigma_{i,t}^2$ terms are conditional variances. Option-implied idiosyncratic skewness in (6) inserts option-implied moments for the conditional moments.

¹²Note that $IS_{idio,t}$ is not an average across firms, but is calculated from average moments across firms, reducing the impact of measurement error. Formally, $IS_{idio,t}$ is the contribution of idiosyncratic shocks to skewness for a hypothetical firm that has the average level of skewness and volatility on date t .

3 The time series of conditional skewness

3.1 Basic characteristics

The left-hand panels in figure 1 plot the time series of $IS_{firm,t}$, $IS_{mkt,t}$, and $IS_{idio,t}$. All three have clear downward trends over time. Aggregate skewness is now significantly more negative than firm skewness, having fallen from around 0 to -3 over 1980–2020, compared to 0 to -1 for firm skewness. A significant fraction of the total declines in aggregate and firm skewness come in two episodes in 1987 and 1989. The 1987 episode is the October crash, which had the largest single-day decline in the history of the S&P 500. In 1989, there was also a significant decline in the stock market associated with a large increase in volatility. Prior to those events, both conditional and realized stock return skewness were near zero. This link between the downward shifts and large jumps in the market is the first indication of a link between aggregate volatility and firm-level skewness.

Gray bars in the figure represent NBER-dated recessions. Section 4 studies the cyclicity of skewness in detail, but one can see in figure 1 that firm and idiosyncratic skewness have fallen in each recession in the sample, while market skewness displays no clear pattern, sometimes rising and sometimes falling in recessions. Perhaps most strikingly, in March and April of 2020, when firm skewness reached some of its most negative values and employment fell by 15 percent, market-level skewness actually rose toward 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 both have trended down, their business cycle and higher-frequency variation is essentially independent.

Shocks to $IS_{firm,t}$, $IS_{mkt,t}$, and $IS_{idio,t}$ are 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.

Table 1a 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.¹³

¹³Note that the 12th autocorrelation is much larger than the first autocorrelation raised to the 12th power, implying that skewness has transitory variation. If it were to be calibrated as an AR(1) process, matching the annual autocorrelation would likely be most natural.

3.2 Validating implied skewness

Since skewness is a scaled moment, the formal link between conditional and realized skewness is not simple. True conditional skewness does not equal the statistical expectation of the realized skewness coefficient in equation (1) since the latter is a nonlinear function of two realized moments. That said, one would still in general expect conditional skewness to predict realized skewness, and this section checks that. I measure realized skewness here using methods developed in Neuberger (2012) and described in detail in appendix A.2. Figure A.1 plots implied and realized skewness, showing that while they follow similar paths, realized skewness appears much noisier.

Table 1b reports results of forecasting regressions for realized skewness. The first column 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 the implied skewness measure. That information is independent of the information contained in lagged realized skewness. The remaining four columns in table 1b show that similar results hold for market-level and idiosyncratic skewness.

Appendix A.2 provides a more thorough analysis of the forecasting power of the conditional moments, treating the theoretical side more carefully – forecasting the numerator and denominator of the skewness coefficient separately – and also directly accounting for the possibility of time-varying risk premia in the option prices. It shows that the paper’s results are robust to those modifications. Appendix A.2 derives an alternative skewness measure based on direct projections of realized returns on state variables. The advantage of the paper’s focus on option prices is that it requires no significant modeling choices – skewness is simply read from option prices directly.

3.3 Relationship with other measures of skewness

The top-left panel of figure 3 plots $IS_{firm,t}$ and the skewness of media news about firms using event sentiment scores from the Ravenpack database of news articles (reported quarterly to reduce noise; see appendix A.3.1). The text-based measure is useful for giving a measure that is independent of asset prices and still available at high frequency, though the news measure is more about the skewness of realized shocks rather than a conditional skewness. The text-based skewness measure is novel to this paper and potentially of independent interest.

The two series are clearly strongly related, with an overall correlation of 0.39 and (0.47

for quarterly changes). Table A.2 reports results of regressions of news skewness on its own lag and lagged implied and realized skewness. Implied skewness has economically and statistically significant forecasting power for news-based skewness, showing that it not only predicts the behavior of stock returns but also the distribution of news events (consistent with the Ravenpack series measuring realized rather than conditional skewness).

The bottom panels on the left-hand side of figure 3 plot $IS_{firm,t}$, minus an HP-filtered trend, against two alternative measures of cross-sectional skewness in annual growth rates: sales growth in Compustat and employment growth in the Census LBD.¹⁴ These are measures of *realized* rather than conditional skewness, measuring the skewness of growth rates that actually occurred, as opposed to what agents expected the skewness so be. Furthermore, they are growth rates (of an annual flow, in one case) making the comparison to a monthly time series difficult.

In both panels, there is a clear positive correlation – both about 0.35 – between detrended $IS_{firm,t}$ and the alternative cross-sectional skewness measures. The three series display similar patterns particularly since the late 1990’s, all falling in the 2001 and 2008 recessions, for example. Since these series are only available at the annual frequency, forecasting regressions have insufficient power to test whether they are predicted by implied skewness.

Appendix A.3.2 shows that $IS_{firm,t}$ is also very similar to a measure of skewness based on the difference between upside and downside variances (e.g. Segal, Shaliastovich, and Yaron (2015)).

4 Cyclicity of skewness

As discussed above, one can see, based on the gray bars in figure 1, that firm skewness is procyclical and aggregate skewness is acyclical. Table 1c reports correlations of the three skewness series, with and without HP-filtering, with various measures of real activity for the sample up to 12/2019.¹⁵ In all cases, firm and idiosyncratic skewness are estimated to be significantly procyclical. Furthermore, the cyclicity is somewhat stronger for the idiosyncratic component.

An alternative way to quantify the cyclicity 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.

¹⁴The LBD data is taken from Salgado, Guvenen, and Bloom (2020). Both measures use the Kelley skewness coefficient to control outliers.

¹⁵Table A.1 reports results through 2020.

The correlations for IS_{mkt} , on the other hand, are all close to zero and there is no consistent evidence of cyclicity – some correlations imply IS_{mkt} is procyclical and others imply it is countercyclical. During NBER-dated recessions, S&P 500 skewness is actually on average higher by 0.35 (0.60 standard deviations). Overall, then, aggregate skewness has been acyclical or even countercyclical over the past 40 years, while firm-level outcomes have consistently procyclical skewness.

To further illustrate the cyclicity of firm-level skewness, the bottom-right panel of figure 3 plot detrended $IS_{firm,t}$ against the Philadelphia Fed’s Leading Indicators index. Each of the declines in the leading indicators index is associated with a decline in skewness (though the reverse is not true).

4.1 Cross-correlations

The bottom four panels of figure 2 plots the cross-correlations of implied skewness and the market and firm levels with employment, the unemployment rate, the Chicago Fed National Activity index (CFNAI), and the Philadelphia Fed’s Leading Indicator index. The shaded regions represent 90-percent confidence bands.

I apply the Hodrick–Prescott filter to both employment and skewness in order to focus on cyclical relationships, which means these correlations do not measure true forecasting relationships (appendix A.3.3 shows similar, though somewhat weaker, results hold without detrending).

$IS_{firm,t}$ has strong correlations with all four measures of activity, peaking at 0.2–0.4 for each. Furthermore, the correlations are highest between current levels of real activity and lagged $IS_{firm,t}$: skewness leads the cycle. Its correlations with the leading indicators index peak near lag zero, implying they move contemporaneously.

The results for market-level skewness are again drastically different from those for firm skewness. In all four cases, the correlations are close to zero, and in fact $IS_{mkt,t}$ again appears if anything to be countercyclical. Furthermore, the estimated correlations are statistically significantly different from those for IS_{firm} .

4.2 Aggregate and idiosyncratic contributions to firm skewness

Rearranging the definition of $IS_{idio,t}$ in equation (6),

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

where, again, IV represents option-implied volatility. The total skewness faced by individual firms, $IS_{firm,t}$, has three sources: market skewness, $IS_{mkt,t}$, idiosyncratic skewness, $IS_{idio,t}$, and changes in their relative weights, through $IV_{mkt,t}/IV_{firm,t}$. Intuitively, since $IS_{mkt,t}$ tends to be more negative than $IS_{idio,t}$, when $IV_{mkt,t}/IV_{firm,t}$ is larger, so that aggregate shocks have a larger scale, firm skewness will become more negative. The question is what effect that mechanism has quantitatively.

The right-hand panels of figure 1 illustrate the relative contributions of those three terms by plotting $IS_{firm,t}$ against the following three comparisons:

$$\begin{aligned}
\text{Idio. skew only:} & \quad \left(\frac{\overline{IV_{mkt}^2}}{\overline{IV_{firm}^2}} \right)^{3/2} \overline{IS_{mkt}} + \left(1 - \frac{\overline{IV_{mkt}^2}}{\overline{IV_{firm}^2}} \right)^{3/2} \boxed{IS_{idio,t}} \\
\text{Mkt. skew only:} & \quad \left(\frac{\overline{IV_{mkt}^2}}{\overline{IV_{firm}^2}} \right)^{3/2} \boxed{IS_{mkt,t}} + \left(1 - \frac{\overline{IV_{mkt}^2}}{\overline{IV_{firm}^2}} \right)^{3/2} \overline{IS_{idio}} \\
\text{Volatility only:} & \quad \left(\frac{\overline{IV_{mkt,t}^2}}{\overline{IV_{firm,t}^2}} \right)^{3/2} \overline{IS_{mkt}} + \left(1 - \frac{\overline{IV_{mkt,t}^2}}{\overline{IV_{firm,t}^2}} \right)^{3/2} \overline{IS_{idio}}
\end{aligned}$$

In each case, only one of the three variables determining the value of $IS_{firm,t}$ varies over time, while the other two are held at their unconditional means over the sample (denoted by overbars). The figure removes time trends from each to focus on higher frequency variation.

The top-right panel in figure 1 shows that when only $IS_{idio,t}$ varies, the series looks nearly identical to total firm skewness, with a correlation of 78 percent. The volatility of this series is also about the same as the volatility of $IS_{firm,t}$. So variation over time in the total skewness faced 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.

The next two panels essentially confirm the result from the top panel. The middle-right panel of figure 1 shows that variation in the market skew explains essentially none of the variation in firm risk – their correlation is -0.01, and the volatility of this series is only 38 percent of that of $IS_{idio,t}$.

Interestingly, the bottom-right panel shows that variation in volatility does matter somewhat. Since market volatility is countercyclical – and more so than idiosyncratic volatility (see Dew-Becker and Giglio (2021)) – that channel will drive firm skewness to be procyclical. 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 0.67 with firm skewness. However, changes in volatility only generate a time series that is 38 percent as volatile as overall $IS_{firm,t}$.¹⁶

Because $IS_{firm,t}$ is a nonlinear function of the three components, there is no additive variance decomposition. Nevertheless, the ranking of importance of the determinants of firm skewness, both in terms of ordering and magnitude, is clear: variation in idiosyncratic skewness is most important, variation in the relative volatility of aggregate and idiosyncratic shocks contributes to procyclicality but is small quantitatively, and variation in aggregate skewness is unimportant at the firm level.

The fact that the cyclicity of skewness is driven by the firm-specific component has important implications. Models in Gabaix (2012), Wachter (2013), Gourio (2013), and Kozlowski, Veldkamp, and Venkateswaran (2020) etc. are driven by variation in the conditional skewness of aggregate shocks. But the results here show that while aggregate skewness does vary, it is not related to the business cycle in the 1980–2020 sample. Instead, it is variation in firm-level risk that has been important, consistent with models that emphasize micro-level mechanisms, including Christiano, Motto, and Rostagno (2014), Ilut, Kehrig, and Schneider (2018), Salgado, Guvenen, and Bloom (2020), and Dew-Becker, Tahbaz-Salehi, and Vedolin (2021).

4.3 Credit spreads, skewness, and the third moment

There are a number of channels through which skewness and the business cycle could be related. A prominent channel is credit spreads: declines in skewness could raise firm borrowing costs, reducing investment and GDP.

Corporate debt is like a short put option (Merton (1974)) – when firm value falls sufficiently far, the value of the debt is eventually impaired. Credit spreads should thus 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 increasing – rising variance – or due to a shift in its shape to be more left skewed. A number of structural models of the economy rely on those effects. In Gourio (2013), the effect is due to variation in aggregate risk, while in Christiano, Motto, and Rostagno (2014) it is about firm-level risk.

Table 1d reports regressions of the Gilchrist–Zakrajsek (2012) credit spread on option-implied moments. The first column confirms the negative relationship with firm skewness and the positive relationship with firm implied volatility, $IV_{firm,t}$, consistent with Christiano, Motto, and Rostagno (2016). The second column replicates the regression from column

¹⁶The results are the same if the three components of firm skewness are shut off one at a time.

1, but replacing firm-level volatility and skewness with S&P 500 moments, showing that increased volatility and decreased skewness are again both associated with higher credit spreads, 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, $I3$:

$$\begin{aligned} I3_{idio,t} &= IS_{idio,t} \times IV_{idio,t}^3 \\ I3_{mkt,t} &= IS_{mkt,t} \times IV_{mkt,t}^3 \end{aligned}$$

The third column of table 1d reports results from a regression of credit spreads on $I3_{idio,t}$ and $I3_{mkt,t}$. The R^2 , at 0.66, 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. Both $I3_{idio,t}$ and $I3_{mkt,t}$ have significantly negative coefficients, though the coefficient on $I3_{idio,t}$ is far larger. Furthermore, its marginal R^2 is substantially larger – 14.0 compared to 3.1 percent for $I3_{mkt,t}$. Overall, then, fluctuations in both market and firm-specific risk affect credit spreads, but the firm-specific component is quantitatively far more important.¹⁷

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, and that information is summarized well by the third moment. Most of the variation in risk driving spreads comes from common shocks to firm-level rather than aggregate risk.

4.4 Comparing volatility and skewness

The results in this paper are closely related to work on time-variation in the conditional volatility of stock returns. Dew-Becker and Giglio (2021) show that market-level conditional volatility is countercyclical, while idiosyncratic conditional volatility is acyclical. That is the opposite of what is found here. So, when studying stock returns, cyclical variation in the aggregate contribution to risk is characterized by changes in the second moment, while cyclical variation in the firm-level contribution to risk is due to changes in the third moment.

To see that, the bottom two panels on the right-hand side of figure 3 plot $IS_{firm,t}$ relative to market and idiosyncratic volatility (i.e. market and idiosyncratic option-implied conditional standard deviations, which sum, in terms of squares, to total firm volatility). In every episode where market volatility jumps up, firm skewness jumps down. Their overall correlation, after removing a linear trend from $IS_{firm,t}$, is -0.47. Section 4.2 shows that that

¹⁷Table A.1 reports results for the same regressions on the sample through 2020.

relationship is not purely mechanical. Also, the correlation of $IS_{idio,t}$, which removes the mechanical component, with $IV_{mkt,t}$, is -0.38, illustrating that the relationship comes from skewness in firm-specific shocks.

The bottom-right panel, shows, conversely, that there is no clear relationship between $IS_{firm,t}$ and the idiosyncratic component of volatility, $IV_{idio,t}$. Idiosyncratic volatility rises in three main episodes, and skewness is low in two of them (2008 and 2020) and high in the other (the late 1990's). The correlation between $IS_{firm,t}$ and $IV_{idio,t}$ reflects that, at only -0.12.

There appears, then, to be a single shock that both drives firm-level skewness to become more negative and market-level volatility to become more positive. That common component in risk is also clearly cyclical, rising in every recession in the sample. This link has not been previously noticed since, up to now, there was not a high-frequency measure of conditional skewness – the common jumps in market volatility and firm skewness are not visible in annual data.

5 Conclusion

This paper provides the first real-time, high-frequency measure of conditional skewness. Its key finding is that aggregate skewness is essentially acyclical, but firm-level skewness is strongly procyclical, due to variation in the skewness of firm-specific shocks. Combining the results here with those in other work, recessions are characterized by greater dispersion in aggregate outcomes, but little or no change in asymmetry, while firm-level shocks become more tilted to the left but, conditionally, no more volatile.

The business cycle also potentially 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.

References

- Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone**, “Vulnerable growth,” *American Economic Review*, 2019, 109 (4), 1263–89.
- Bakshi, Gurdip and Charles Cao**, “Risk Neutral Kurtosis, Jumps, and Option Pricing: Evidence from 100 Most Actively Traded Firms on the CBOE,” 2003. Working paper.
- , **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.
- Bekaert, Geert, Eric Engstrom, and Andrey Ermolov**, “Bad environments, good environments: A non-Gaussian asymmetric volatility model,” *Journal of Econometrics*, 2015, *186* (1), 258–275.
- Bloom, Nicholas**, “The Impact of Uncertainty Shocks,” *Econometrica*, 2009, *77*(3), 623–685.
- , **Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry**, “Really uncertain business cycles,” *Econometrica*, 2018, *86* (3), 1031–1065.
- 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, Alireza Tahbaz-Salehi, and Andrea Vedolin**, “Macro skewness and conditional second moments: evidence and theories.” Working paper.
- and **Stefano Giglio**, “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**, “The granular origins of aggregate fluctuations,” *Econometrica*, 2011, *79* (3), 733–772.
- , “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.
- 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.
- Harmenberg, Karl and Hans Henrik Sievertsen**, “The Labor-Market Origins of Cyclical Income Risk,” 2017. Working paper.
- Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh**, “The common factor in idiosyncratic volatility: Quantitative asset pricing implications,” *Journal of Financial Economics*, 2016, 119 (2), 249–283.
- 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.
- Kehrig, Matthias**, “The Cyclical Nature of the Productivity Distribution,” 2015. Working paper.
- Kozeniasukas, Nicholas, Anna Orlik, and Laura Veldkamp**, “What are uncertainty shocks?,” *Journal of Monetary Economics*, 2018, 100, 1–15.
- 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.
- Mayhew, Stewart and Vassil Mihov**, “How do exchanges select stocks for option listing?,” *The Journal of Finance*, 2004, 59 (1), 447–471.
- Merton, Robert C**, “On the pricing of corporate debt: The risk structure of interest rates,” *The Journal of finance*, 1974, 29 (2), 449–470.
- Muravyev, Dmitriy and Neil D. Pearson**, “Option Trading Costs Are Lower Than You Think,” *Review of Financial Studies*, 2020. Forthcoming.
- 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.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 2004, 112 (3), 695–717.
- 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: Moments and regressions

(a) Calibration moments									
	Raw data					Detrended			
	Mean	SD	AC(1)	AC(3)	AC(12)	SD	AC(1)	AC(3)	AC(12)
$IS_{firm,t}$	-0.54	0.46	0.93	0.87	0.76	0.22	0.70	0.44	0.03
	[0.07]	[0.04]	[0.01]	[0.03]	[0.07]	[0.02]	[0.04]	[0.06]	[0.07]
$IS_{mkt,t}$	-1.66	0.84	0.94	0.86	0.73	0.40	0.73	0.43	0.02
	[0.13]	[0.12]	[0.02]	[0.04]	[0.09]	[0.04]	[0.03]	[0.06]	[0.09]
$IS_{idio,t}$	-0.53	0.75	0.89	0.77	0.57	0.48	0.72	0.43	0.02
	[0.11]	[0.09]	[0.02]	[0.05]	[0.09]	[0.07]	[0.05]	[0.09]	[0.06]

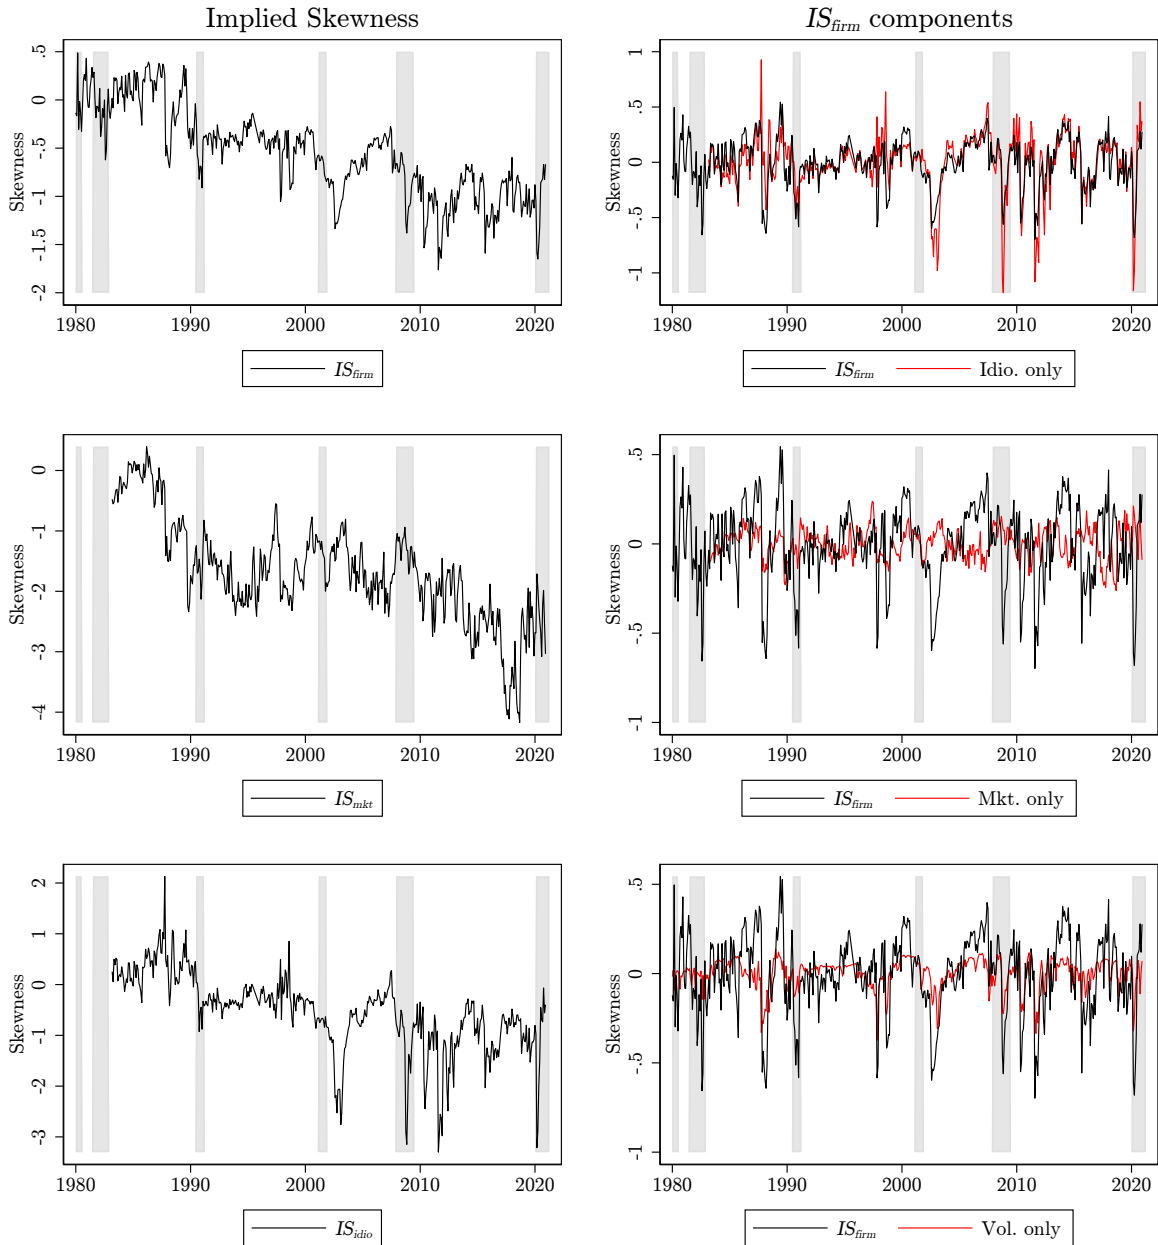
(b) Forecasting realized skewness						
	Firm		Market		Idiosyncratic	
$IS_{x,t}$	0.29	0.20	0.43	0.30	0.32	0.20
	[0.04]	[0.03]	[0.04]	[0.04]	[0.08]	[0.06]
$RS_{x,t}$		0.46		0.30		0.28
		[0.04]		[0.06]		[0.08]
Constant	-0.14	-0.11	-0.04	-0.04	-0.25	-0.19
	[0.03]	[0.02]	[0.07]	[0.05]	[0.05]	[0.04]

(c) Correlations						
	$IS_{firm,t}$	$IS_{firm,t}$ detrended	$IS_{mkt,t}$	$IS_{mkt,t}$ detrended	$IS_{idio,t}$	$IS_{idio,t}$ detrended
Industrial prod.	0.17	0.21	-0.01	-0.07	0.25	0.24
Employment	0.17	0.22	-0.01	0.02	0.23	0.19
Unemployment	-0.22	-0.26	0.02	0.04	-0.28	-0.24
CBO output gap	0.15	0.21	-0.16	-0.08	0.27	0.27
NBER rec. ind.	-0.11	-0.14	0.10	0.15	-0.17	-0.16
Capacity util.	0.51	0.15	0.14	-0.09	0.56	0.20
Empl. growth	0.17	0.09	0.01	-0.17	0.23	0.18
IP growth	0.34	0.14	0.06	-0.19	0.40	0.22

(d) Credit spread regressions						
	Firm moments		Market moments		mkt vs. idio.	
$IS_{firm,t}$	-1.06		$IS_{mkt,t}$	-0.23	$I\mathcal{B}_{idio,t}$	-22.15
	[0.17]			[0.07]		[2.65]
$IV_{firm,t}$	6.53		$IV_{mkt,t}$	11.00	$I\mathcal{B}_{mkt,t}$	-13.00
	[1.68]			[2.43]		[3.71]
R^2	0.62			0.55		0.66

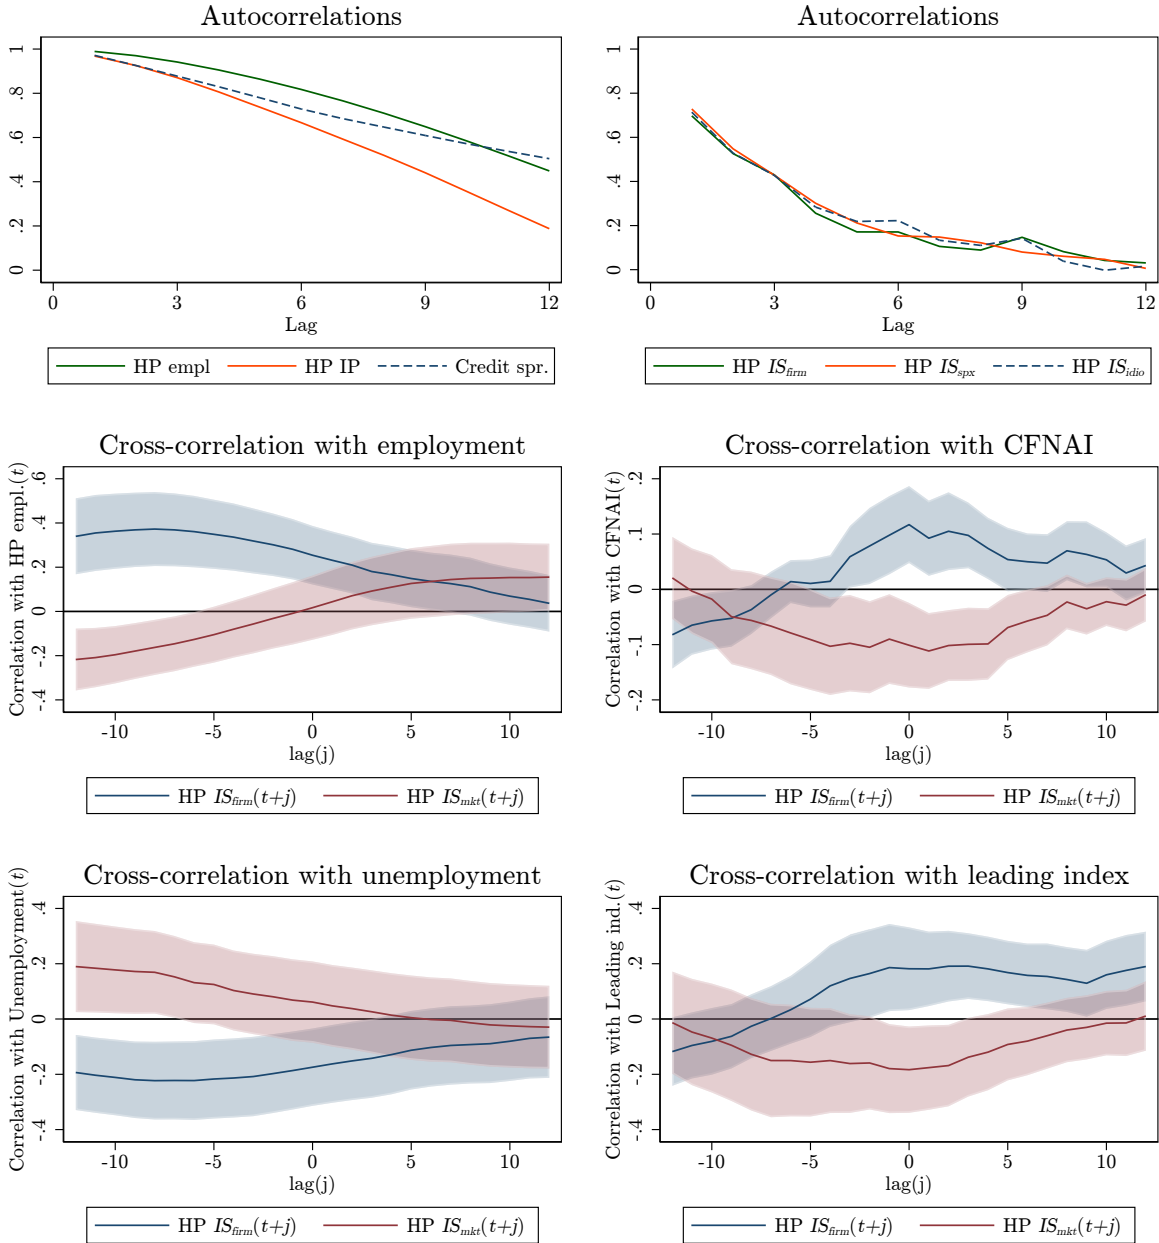
Notes: In panels (a), (b), and (d), standard errors, reported in brackets, are calculated using the Newey–West method with 12 lags. Panels (a) and (c) detrend using the HP filter with a smoothing parameter of 129,600. The levels of industrial production and employment are detrended in the same way. There is no detrending in panel (d). The sample in panels (c) and (d) stop in 12/2019.

Figure 1: Implied skewness and its components



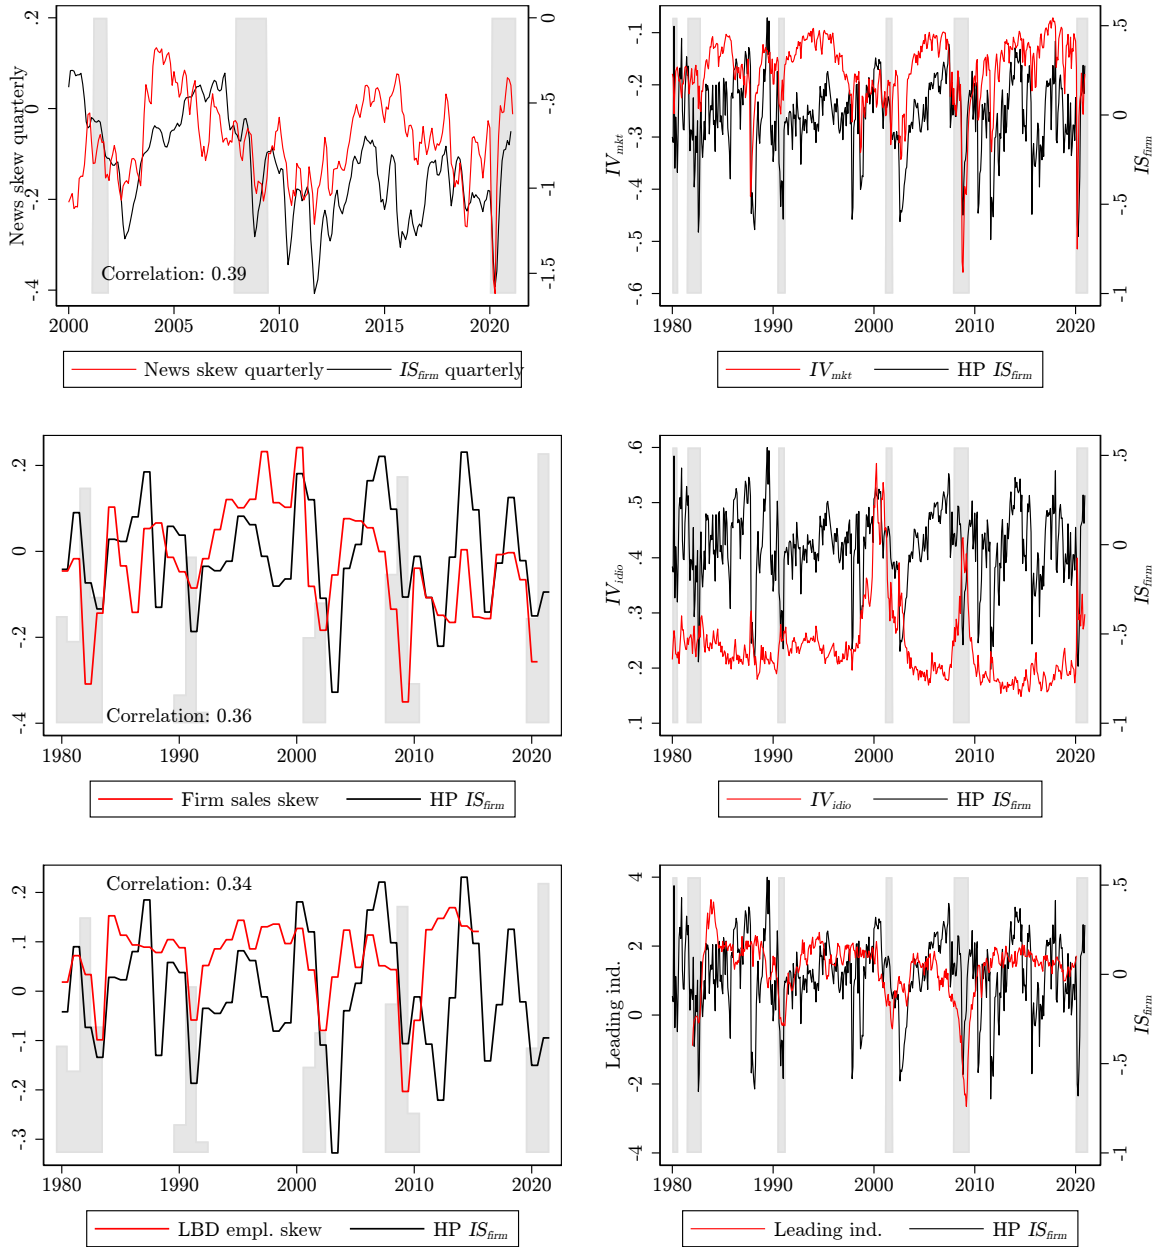
Note: The left-hand panels plot option-implied skewness at the firm, market, and idiosyncratic levels. The right-hand panels plot total firm skewness when just one of its three components varies over time, with the other two left at their unconditional means. The series in the right-hand panel are HP-filtered. 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 of HP-filtered implied skewness at the firm and market levels with other variables. Shaded regions represent 90-percent confidence intervals.

Figure 3: Comparisons with other variables



Note: The left-hand panels plot $IS_{firm,t}$ in black against three alternative measures of skewness in red. The scales for the alternatives are shifted to match $IS_{firm,t}$. For the top two panels, $IS_{firm,t}$ is detrended. The right-hand panels plot linearly detrended $IS_{firm,t}$ in black against market and idiosyncratic implied volatility ($IV_{mkt,t}$ and $IV_{idio,t}$, respectively) and the Philadelphia Fed's leading indicators index.

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

This section describes how I fit option prices in order to calculate risk-neutral moments at the firm and market level. I 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

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

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 $f_i(k) \rightarrow 0$ as $k \rightarrow \pm\infty$. b_i^+ controls the difference between the slope as $k \rightarrow \infty$ and $k \rightarrow -\infty$.

The main step is estimating the parameters a_i , b_i , and b_i^+ and the function f_i . I 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^+ , I assume that

$$a_i \sim N(a, \sigma_a^2) \quad (\text{A.2})$$

$$b_i \sim N(b, \sigma_b^2) \quad (\text{A.3})$$

$$b_i^+ \sim N(b^+, \sigma_{b^+}^2) \quad (\text{A.4})$$

The parameters governing the distributions, a , σ_a^2 , etc., are hyperparameters that can be estimated through 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, I first normalize all underlyings to have a price of 1. That is, given an observation with an underlying price of $S_{i,t}$, a strike of \tilde{K} , and an option price of $\tilde{P}_{\tilde{K},i,t}$, we divide through by $S_{i,t}$, to get

$$K_t = \tilde{K}_t / S_{i,t} \quad (\text{A.5})$$

$$P_{K,i,t} = \tilde{P}_{\tilde{K},i,t} / S_{i,t} \quad (\text{A.6})$$

This is simply a renormalization in terms of a different numeraire (units of the underlying instead of dollars).

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). I 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 $f_i(k) \rightarrow 0$ almost surely as $k \rightarrow \pm\infty$. 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,

$$E[f_i(k)] = 0 \forall i, k \quad (\text{A.7})$$

$$\text{cov}(f_i(k), f_j(m)) = \sigma_B^2 (\min(L(k), L(m)) - L(k)L(m)) \quad (\text{A.8})$$

$$+ \sigma_D^2 \delta_{k>0} \delta_{m>0} (1 - L(k))(1 - L(m)) \quad (\text{A.9})$$

$$+ \delta_{i=j} \left(\begin{array}{l} \sigma_B^2 (\min(L(k), L(m)) - L(k)L(m)) \\ + \sigma_D^2 \delta_{k>0} \delta_{m>0} (1 - L(k))(1 - L(m)) \end{array} \right) \quad (\text{A.10})$$

where $L(k) \equiv \frac{\exp(sk)}{\exp(sk)+1}$ is a scaled logistic function and $s, \sigma_B^2, \sigma_D^2, \sigma_b^2$, and σ_D^2 are hyperparameters. δ_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 $L(k)$ is what stretches the Brownian bridge to cover the full real line, while the parameter s in $L(k)$ 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 through the random components in a_i, b_i, b_i^+ , and $f_i(k)$.

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

The Gaussian process regression not only fits the option prices but also can be used to create an estimate of prices at any unobserved strike – specifically, one can get the expectation of the price at a given strike conditional on the other observed prices and the hyperparameters. The integral in equation (2) can then be calculated to arbitrary precision over a discrete grid of strikes.

A.1.1 Data filters and estimation

I estimate the model above using daily closing prices for options on individual stocks and the S&P 500 index. On each day, I take prices for out-of-the-money options where the bid is greater than \$0.10 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 64. I weight observations in the estimation by the underlying firm’s market value divided by the bid-ask spread for the log option price (i.e. the log ask minus the log bid for each option). I also drop any observations where the implied volatility is below 0.01 or above 5 (which are generally data errors). When I do the estimation for the S&P 500, I 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, I 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. I therefore fix the variance hyperparameters at values that were 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 fit for 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 the 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 the skewness measure to the CBOE’s reported values for the S&P 500. Those time series are highly similar. While skewness under the CBOE index is more negative in 2020 than in this paper’s series, the CBOE series rises towards zero almost monotonically between 1/1/2020 and 4/1/2020 – i.e. as the market was falling and volatility rising.

It compares the results using data from CME S&P 500 futures options to CBOE SPX options as reported by Optionmetrics. The results are highly similar.

The Optionmetrics sample is restricted to the top 200 firms by market capitalization on each date for the sake of consistency with the Berkeley Options Database, for which I only have the largest firms. To evaluate the impact of that restriction on the results, I also re-estimated the option-implied skewness using the top 100 firms. With 100 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 I 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. That is also plotted in figure A.1.

A.2 Realized skewness and risk premia

A.2.1 Realized skewness from the text

The results in the main text calculate the realized third moment using the method of Neuberger (2012) (see proposition 6),

$$R3_{i,t} = \sum_{days \in t} 3\Delta p_{i,daily}^E (\exp(r_{i,daily}) - 1) + 6(r_{i,daily} \exp(r_{i,daily}) - 2\exp(r_{i,daily}) + r_{i,daily} + 2) \quad (\text{A.11})$$

where $R3_{i,t}$ denotes the realized third moment for firm i in month t . The summation is taken over the days in the month. $\Delta p_{i,daily}^E$ is the daily change in the price of the entropy contract for firm i (equation (24) of Kozhan, Neuberger, and Schneider (2013)). $r_{i,daily}$ is the daily log return on firm i 's stock. In Neuberger's (2012) formulation, the maturity of the log contract changes over the course of the month. I just take it at the average maturity of 15 days for the sake of simplicity (this difference has minimal impact on the results).

Based on (A.11), and following Neuberger (2012), I calculate average *Realized Skewness* at the firm level in the main text as

$$RS_{firm,t} = \sum_i mkt_{i,t} \frac{\sum_{days \in t} r_{i,daily}^3 + 3r_{i,daily} \Delta IV_{i,daily}^2}{\left(\sum_{days \in t} r_{i,daily}^2\right)^{2/3}} \quad (\text{A.12})$$

and $RS_{mkt,t}$ is again defined analogously for the S&P 500 is used to represent the aggregate market. Idiosyncratic realized skewness, $RS_{idio,t}$, is then constructed using the formula (6) replacing implied with realized moments.

Note that even if $IS_{i,t}$ is the true conditional skewness coefficient, it will not be the case that $IS_{i,t}$ is equal to the expectation of the realized skewness coefficient, $E_t \left[R3_{i,t+1} / RV_{i,t+1}^{3/2} \right]$, where RV is realized volatility,

$$RV_{i,t} \equiv \left(\sum_{days \in t} r_{i,daily}^2 \right)^{1/2} \quad (\text{A.13})$$

That is,

$$IS_{i,t} \neq E_t [RS_{i,t+1}] \quad (\text{A.14})$$

Neuberger (2012) and Kozhan, Neuberger, and Schneider (2013) also discuss this issue. It is due to the fact that the skewness coefficient is a nonlinear function of moments. The next section discusses how to address it and also how to account for time-varying risk premia.

A.2.2 Accounting for risk premia and nonlinearity in the skewness coefficient

The main text takes $IS_{i,t}$ as an estimate of conditional skewness. An alternative to that is to simply directly calculate the conditional expectations – the moments that go into the skewness coefficient, $E_t r_{i,t+1}^3$ and $E_t r_{i,t+1}^2$ – from linear projections. The idea in that case is that there is some set of state variables determining the conditional moments and risk premia, so then the conditional expectations can be recovered from a projection of the realized second and third moments on those state variables.

To do that, I estimate regressions of the form

$$R3_{firm,t} = \beta'_{3,firm} X_{t-1} + \varepsilon_{3,firm,t} \quad (\text{A.15})$$

where $R3_{firm,t}$ is the average of the realized third moment, $R3_{i,t}$, across firms, weighting by market values as elsewhere. The same regression can be run for the average realized second moment, $R2_{firm,t}$, and the same for the overall stock market. X_t here is a vector of date- t state variables, and $\beta_{3,firm}$ is a vector of coefficients. $\varepsilon_{3,firm,t}$ is then a zero-mean residual.

For X_t , I include $IS_{firm,t}$, $IS_{mkt,t}$, $IV_{firm,t}$, $IV_{mkt,t}$, the S&P 500 price/earnings ratio, the CBO output gap, the unemployment rate, and the Gilchrist–Zakrajsek credit spread. The same X_t is used in all forecasting regressions.

In the data, the realized moments have distributions with heavy tails – they have kurtosis of over 200. That means that there are large outliers, so OLS performs poorly, and it is inefficient since the residuals are not Gaussian. Instead, I obtain the coefficients β through maximum likelihood estimation in which I assume that the residuals are t -distributed with four degrees of freedom (results are similar assuming a few more degrees of freedom or switching to a Cauchy distribution).

The regression is estimated for the second and third moments at both the market and firm levels – i.e. using $R3_{firm,t}$, $R2_{firm,t}$, $R3_{mkt,t}$, and $R2_{mkt,t}$ as the dependent variable,

where

$$R2_{firm,t} \equiv \frac{\sum_i \left(mkt_{i,t} \sum_{days \in t} r_{i,daily}^2 \right)}{\sum_i mkt_{i,t}} \quad (\text{A.16})$$

$$R2_{mkt,t} \equiv \sum_{days \in t} r_{mkt,daily}^2 \quad (\text{A.17})$$

The maximum likelihood estimation of the four regressions then yields estimated coefficients $\hat{\beta}_{3,firm}$, $\hat{\beta}_{2,firm}$, $\hat{\beta}_{3,mkt}$, and $\hat{\beta}_{2,mkt}$. That then yields estimates of the conditional moments based on *Projections*,

$$P3_{firm,t} \equiv \hat{\beta}'_{3,firm} X_t \quad (\text{A.18})$$

$$P2_{firm,t} \equiv \hat{\beta}'_{2,firm} X_t \quad (\text{A.19})$$

$$P3_{mkt,t} \equiv \hat{\beta}'_{3,mkt} X_t \quad (\text{A.20})$$

$$P2_{mkt,t} \equiv \hat{\beta}'_{2,mkt} X_t \quad (\text{A.21})$$

and the *Projected Skewness* at the firm and market levels is defined as

$$PS_{firm,t} \equiv \frac{P3_{firm,t}}{P2_{firm,t}^{3/2}} \quad (\text{A.22})$$

$$PS_{mkt,t} \equiv \frac{P3_{mkt,t}}{P2_{mkt,t}^{3/2}} \quad (\text{A.23})$$

In other words, *PS* is defined analogously to *IS*, but instead of the conditional moments coming from option prices, which may contain risk premia, they come from direct regressions of realized moments on observables. In principle, the analysis in the main text could have been done entirely in terms of *PS* instead of *IS* (that might have been natural in the absence of available option prices).

I now confirm that the key results from the text on implied skewness also hold for skewness from the projection. Specifically:

1. **Low correlation between firm and market skewness:** The correlation between $IS_{firm,t}$ and $IS_{mkt,t}$ after removing time trends is 0.04. For $PS_{firm,t}$ and $PS_{mkt,t}$ it is -0.06. So in both cases market and firm skewness are close to uncorrelated.

2. **Negative correlation between firm skewness and market volatility:** The correlation between $IS_{firm,t}$ and $IV_{mkt,t}$ after removing time trends is -0.52. For $PS_{firm,t}$ and $IV_{mkt,t}$ it is -0.62. So in both cases firm skewness is significantly negatively correlated with market volatility. For $PS_{firm,t}$ and $P2_{mkt,t}$, the correlation is -0.50.

3. **Firm skewness is procyclical and market skewness is acyclical:** The table below reports correlations between $IS_{firm,t}$, $IS_{mkt,t}$, $PS_{firm,t}$, $PS_{mkt,t}$, and measures of real activity. Both $IS_{firm,t}$ and $PS_{firm,t}$ are consistently procyclical, while $PS_{mkt,t}$ has correlations close to zero, as with $IS_{mkt,t}$. As in the main text, the data for the table runs to the end of 2019.

	$IS_{firm,t}$	$IS_{mkt,t}$	$PS_{firm,t}$	$PS_{mkt,t}$
HP indus. prod.	0.17	-0.01	0.21	0.03
HP employment	0.16	-0.01	0.13	-0.03
Unemployment	0.01	0.31	-0.15	0.35
CBO output gap	0.15	-0.16	0.23	-0.18
NBER rec. ind.	-0.12	0.10	-0.26	0.00
Capacity util.	0.48	0.14	0.51	0.17
Empl. growth	0.32	0.06	0.40	0.17
IP growth	0.15	0.01	0.21	0.04

A.3 Other robustness checks

A.3.1 Ravenpack news skewness

To construct a measure of news skewness, I obtain news Event Sentiment Scores (ESS) from the Ravenpack database for the firms included in the implied skewness index. The firms are linked from Optionmetrics by finding matches on either CUSIP or ticker. For each firm×month observation, I calculate the skewness coefficient for the ESS across all events. I then average the skewness across firms within each month (the results are highly similar – with a correlation of 90 percent – with and without weighting by stock market capitalization) to get the overall skewness index.

Figure 3 in the main text plots quarterly moving averages of news skewness. The top panel of figure A.2 plots the original monthly data, where the transitory noise is clearly evident. The 1-month autocorrelation of changes in news skewness is -0.31 (compared to only -0.13 for IS_{firm}), consistent with the presence of transitory measurement error or other noise.

Table A.2 reports results from regressions of news skewness on lagged $IS_{firm,t}$ along with controls for lagged news skewness and lagged $RS_{firm,t}$. The dependent and independent variables are normalized to have unit variance to aid interpretation of the coefficients. In all three cases, the coefficient on lagged option implied skewness is statistically and economically

significant. Option prices thus have the ability to forecast the skewness of realized news events at the firm level, even after controlling for lagged news skewness and lagged realized stock market skewness. The regressions in table A.2 use the full data sample (2000–2020), but the results are almost identical if the sample is stopped at the end of 2019.

A.3.2 VIX asymmetry

Recent work has analyzed upside and downside variance (e.g. Bekaert, Engstrom, and Ermolov (2015) and Segal, Shaliastovich, and Yaron (2015)). To use that to give a measure of asymmetry, I define

$$VIXA_t = \frac{E_t [r_{t+1}^2 | r_{t+1} > 0]^{1/2} - E_t [r_{t+1}^2 | r_{t+1} < 0]^{1/2}}{E_t [r_{t+1}^2]^{1/2}}$$

$VIXA_t$ is a scaled measure of asymmetry, similar to skewness. The bottom panel of figure A.2 plots $IS_{firm,t}$ along with $VIXA_{firm,t}$ (rescaled to have the same standard deviation as $IS_{firm,t}$). The two series have a correlation of 0.97, and a similar result holds when comparing S&P 500 skewness to S&P 500 VIX asymmetry. The information contained in the option-implied skewness coefficient is thus highly similar to that in relative upside and downside variance.

A.3.3 Cross-correlations and alternative detrending

Figures A.4 and A.5 plot cross-correlations similar to those in the main text, but with two changes. Figure A.4 replaces the HP filtering of implied skewness with linear detrending, and also uses growth rates for employment and IP instead of HP-filtered levels. Figure A.5 eliminates all detrending. The correlations are somewhat smaller in these figures, which is natural since the trend absorbs some of the variance of implied skewness. Directionally, the results are almost all identical.

A.3.4 Liquidity

As discussed in the main text, this is not the first paper to measure skewness in individual stocks, and liquidity is discussed elsewhere in the literature. The filters described in section A.1.1 are similar to those used in past work to address liquidity concerns; see, e.g., Bakshi and Cao (2003).

The concern with liquidity is that the liquidity of options could affect their price. Differences in liquidity across strikes could then create implied skewness, even if investors believe returns to be symmetrically distributed. If those differences in skewness change systematically over time, e.g. correlated with the business cycle, then liquidity could drive some of the results.

A standard way to measure liquidity is to examine bid/ask spreads. It is important to note that in the case of options it is often possible to trade inside the quoted spread, so that the effective spread paid is smaller than the quoted spread, which at times can be large. Muravyev and Pearson (2020) show that the effective spreads that traders actually pay can be less than 40 percent of the quoted spreads.

To evaluate the possibility that liquidity drives the results, I calculate percentage bid-ask half-spreads (i.e. scaled by the mid-price of the option) for puts and calls over time. The half-spreads are calculated for a hypothetical put and call with strike/spot equal to 90% for the put and 110% for the call. They are obtained by taking the fitted values from the regression

$$\begin{aligned}
spr_{i,j,t} = & b_0 + a_i + b_1 \log \left(\frac{strike_{i,j,t}}{spot_{i,t}} \right)_{i,j,t} + b_2 \log \left(\frac{strike_{i,j,t}}{spot_{i,t}} \right)_{i,j,t} \times 1 \left\{ \log \left(\frac{strike_{i,j,t}}{spot_{i,t}} \right)_{i,j,t} > 0 \right\} \\
& + b_3 1 \left\{ \log \left(\frac{strike_{i,j,t}}{spot_{i,t}} \right)_{i,j,t} > 0 \right\} + \eta_{i,j,t}
\end{aligned} \tag{A.24}$$

where $spr_{i,j,t}$ is the bid/ask spread for option j of firm i on date t , a_i is a set of firm fixed effects normalized to have zero mean across the firms, $1 \{ \cdot \}$ represents an indicator function, $\eta_{i,j,t}$ is a residual, and the b 's are coefficients. The regression is run month-by-month to get monthly average fitted spreads. The results are highly similar taking the median bid/ask spread for all options with strike/spot within 2 percentage points of 90% and 110% or if the observations are weighted by the market capitalization of the firm.

The top panel of figure A.6 plots the average fitted spreads over time. The magnitude of the half-spreads is similar to what is reported in past work, e.g. Muravyev and Pearson (2020) (and, again, is an overestimate of actual trading costs by likely a factor of two). In addition, the spreads are almost identical for strikes above and below the spot price, which implies that there is no significant asymmetry in liquidity that might distort implied skewness.

In addition, while the spreads vary significantly over time (potentially partly due to variation in minimum tick size, e.g. related to the CBOE Penny Pilot program), the variation appears to be unrelated to short-term variation in implied skewness, especially at high fre-

quencies. Jumps down in skewness, as are often observed in recessions and when the S&P 500 VIX jumps up, do not appear to be related to shifts in liquidity.

At very low frequencies – i.e. comparing the very beginning of the sample to the very end – liquidity for low strikes has risen relative to liquidity for high strikes.

The conclusion from the top panel of figure A.6 is therefore that liquidity, based on bid/ask spreads, is similar at strikes above and below the underlying price, and while the spreads vary over time, they do not do so in a way that is correlated with business cycle frequency shifts in skewness.

To further examine liquidity, the bottom panel of figure A.6 plots the simple median of raw bid/ask spreads across options with moneyness between 88% and 92% and 108% and 112%. The raw spreads vary over time in a way that is generally similar to what is observed for the percentage spreads, and again without any substantial difference between spreads above and below the spot price and without a clear relationship with short-term movements in skewness. Again, at the lowest frequencies, the put spreads have fallen slightly relative to the call spreads as implied skewness has declines, but the change is small relative to the overall magnitude of spreads.

Table A.1: Correlations and regressions including 2020

(c) Correlations						
	$IS_{firm,t}$	$IS_{firm,t}$ detrended	$IS_{mkt,t}$	$IS_{mkt,t}$ detrended	$IS_{idio,t}$	$IS_{idio,t}$ detrended
Industrial prod.	0.18	0.24	-0.04	-0.12	0.26	0.28
Employment	0.15	0.22	-0.05	-0.04	0.20	0.19
Unemployment	-0.18	-0.26	0.06	0.09	-0.25	-0.25
CBO output gap	0.20	0.24	-0.11	-0.12	0.30	0.29
NBER rec. ind.	-0.17	-0.15	0.02	0.19	-0.21	-0.18
Capacity util.	0.53	0.18	0.17	-0.13	0.57	0.23
Empl. growth	0.15	0.15	0.00	-0.13	0.21	0.22
IP growth	0.17	0.17	-0.02	-0.18	0.25	0.28

(d) Credit spread regressions						
	Firm moments		Market moments		mkt vs. idio.	
$IS_{firm,t}$	-1.01 [0.17]		$IS_{mkt,t}$	-0.20 [0.07]	$I3_{idio,t}$	-22.68 [2.25]
$IV_{firm,t}$	6.21 [1.64]		$IV_{mkt,t}$	10.32 [2.36]	$I3_{mkt,t}$	-6.68 [5.87]
R^2	0.60			0.53		0.60

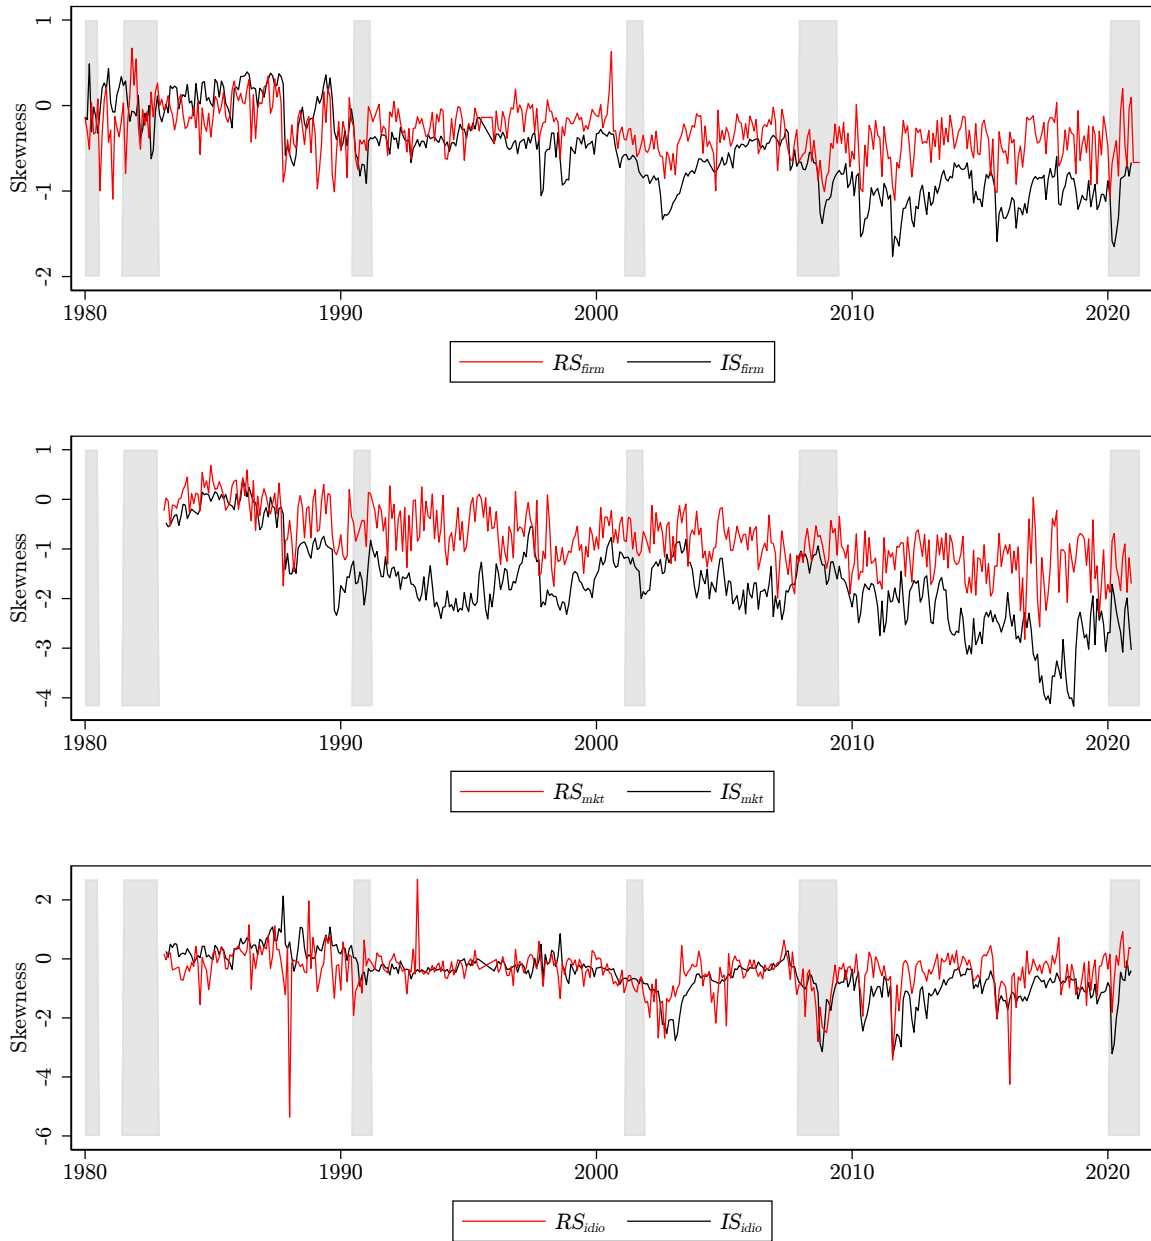
Notes: These results replicate panels (c) and (d) from table 1 but using the sample through the end of 2020.

Table A.2: Forecasting news skewness

$IS_{firm,t-1}$	0.49 [0.17]	0.19 [0.09]	0.40 [0.16]
News skew $_{t-1}$		0.57 [0.06]	
$RS_{firm,t-1}$			0.14 [0.12]
Constant	0.35 [0.19]	0.14 [0.08]	0.33 [0.19]

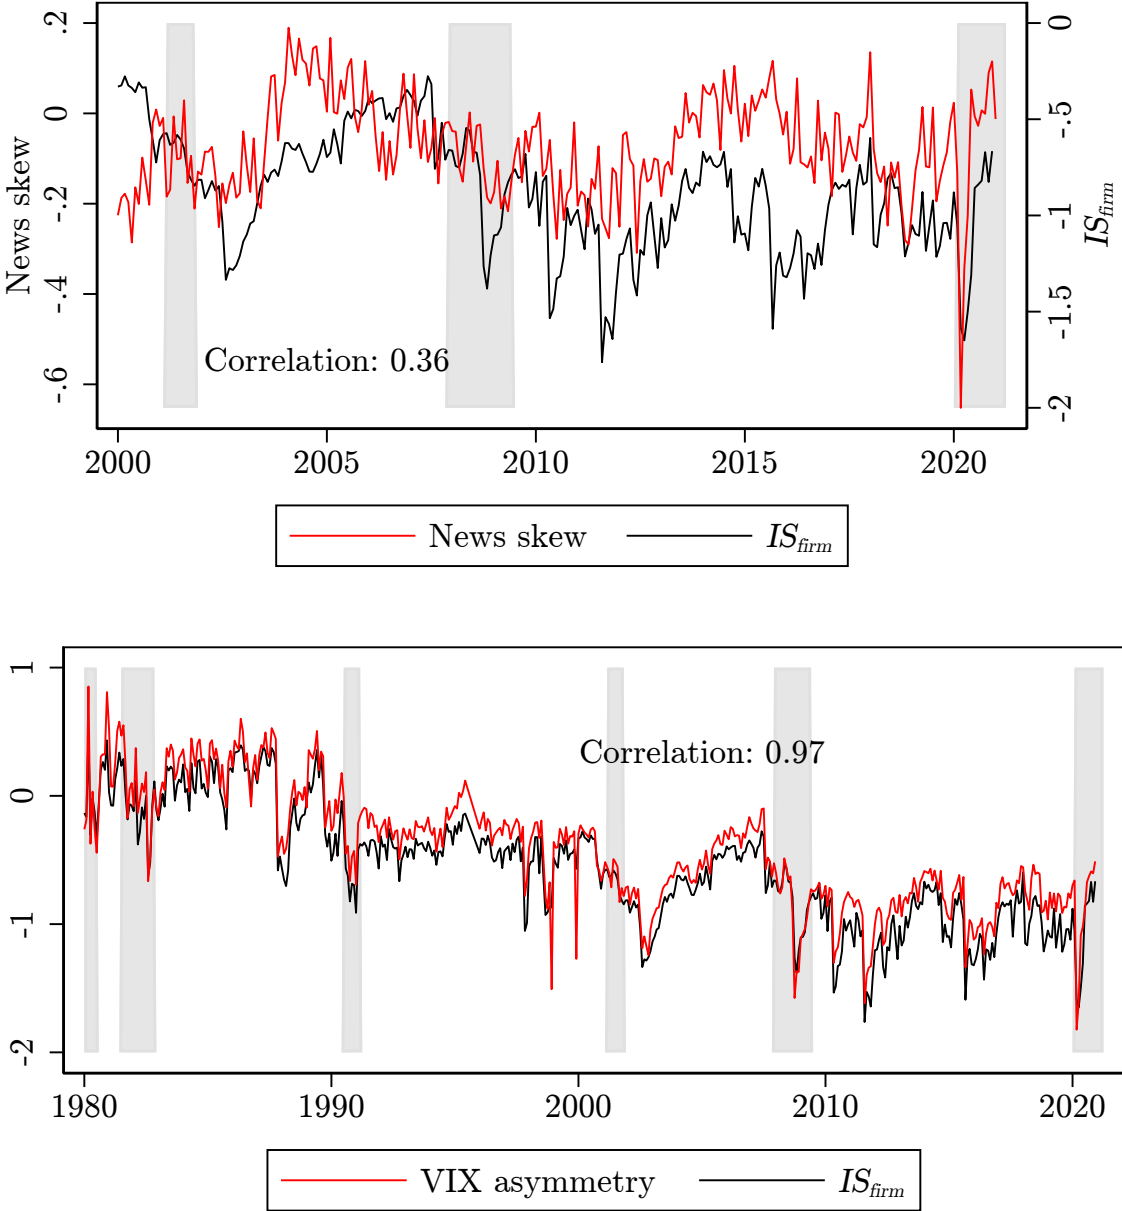
Notes: Forecasting regressions for news skewness from Ravenpack (see appendix A.3.1). All variables are normalized to have zero mean and unit variance. Standard errors calculated by the Newey–West method with six lags are reported in brackets. The data is monthly, 2000–2020.

Figure A.1: Implied versus realized skewness



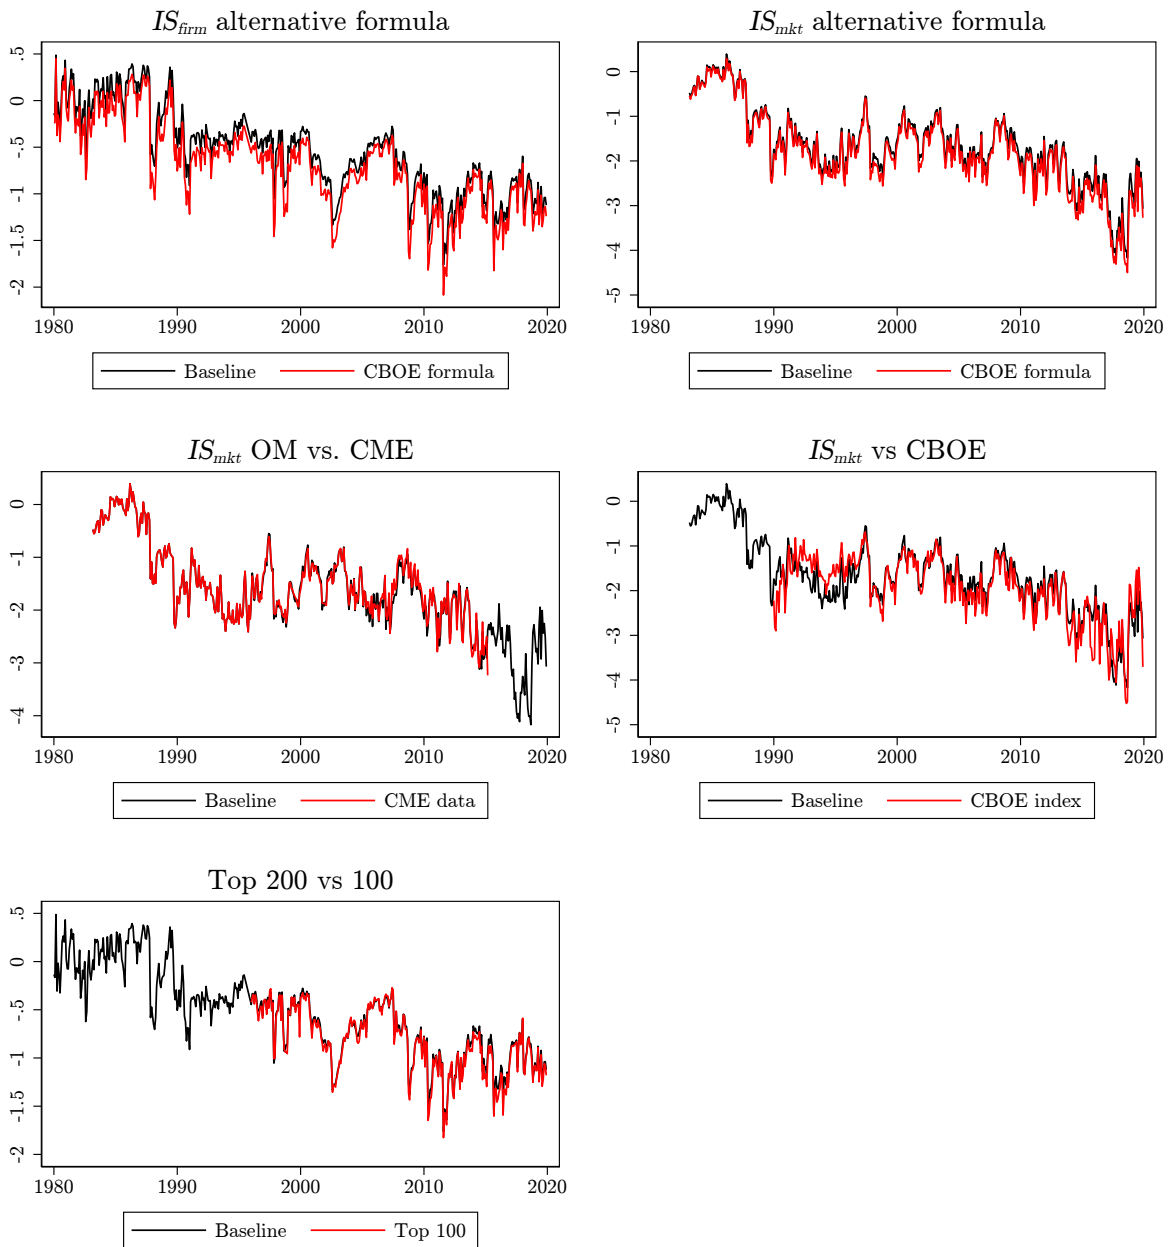
Note: Implied and realized skewness

Figure A.2: Alternative skewness series



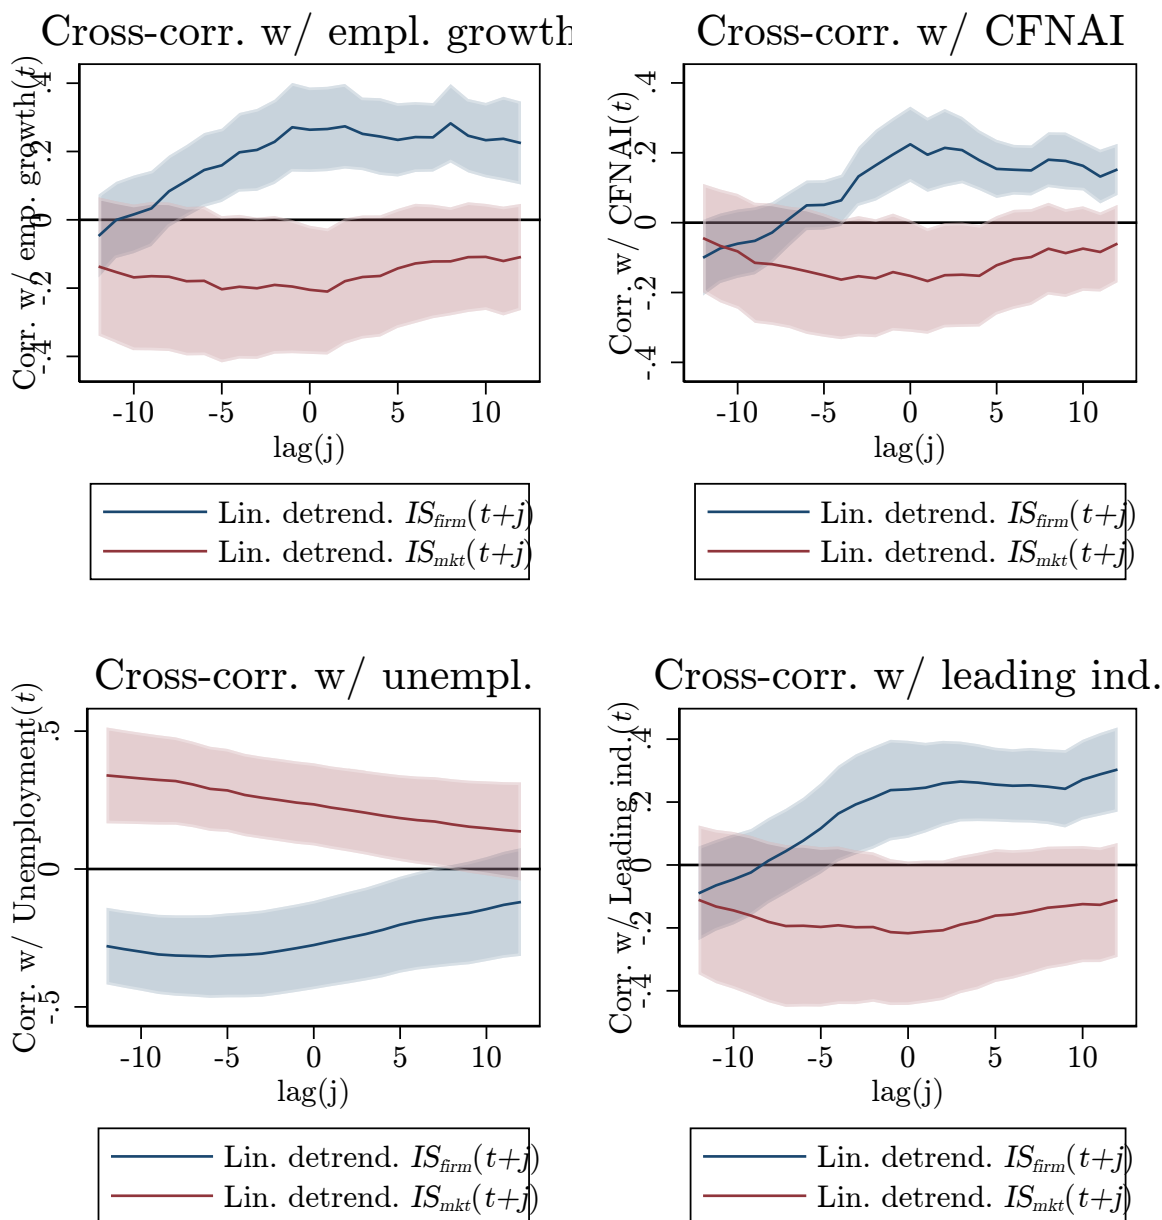
Note: The top panel replicates the top-left panel of figure 3, but with monthly instead of quarterly data. The bottom panel is the implied skew from VIX asymmetry (with both series normalized to have unit standard deviation). See sections A.3.1 and A.3.2, respectively.

Figure A.3: Skewness comparisons



Note: Baseline skewness series and alternatives.

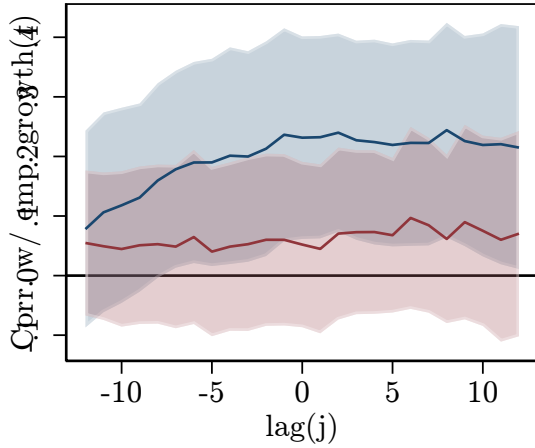
Figure A.4: Cross-correlations with linear detrending of implied skewness



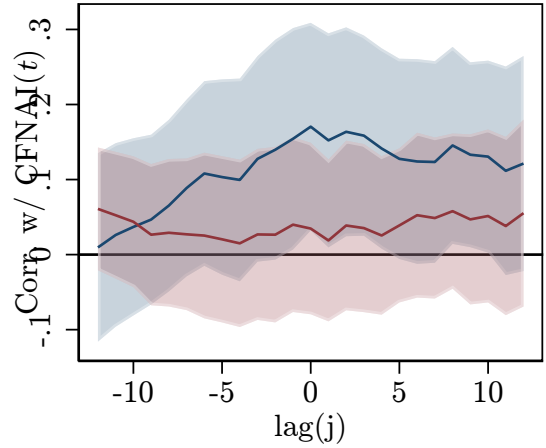
Note: This is an alternative version of the cross-correlations with two changes: the implied skewness series are linearly detrended instead of using the HP filter, and HP-filtered employment and IP are replaced with their growth rates.

Figure A.5: Cross-correlations without detrending

Cross-correlation w/ empl. grov Cross-correlation w/ CFNAI

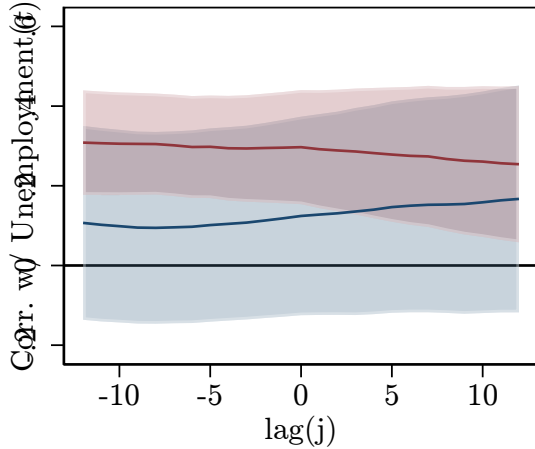


— $IS_{firm}(t+j)$ — $IS_{mkt}(t+j)$

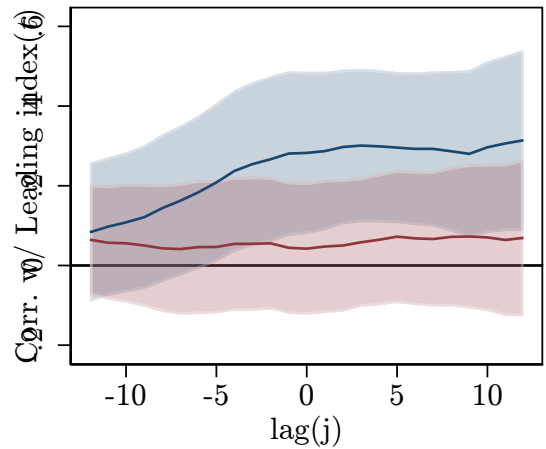


— $IS_{firm}(t+j)$ — $IS_{mkt}(t+j)$

Cross-correlation w/ unemployon Cross-correlation w/ leading inde



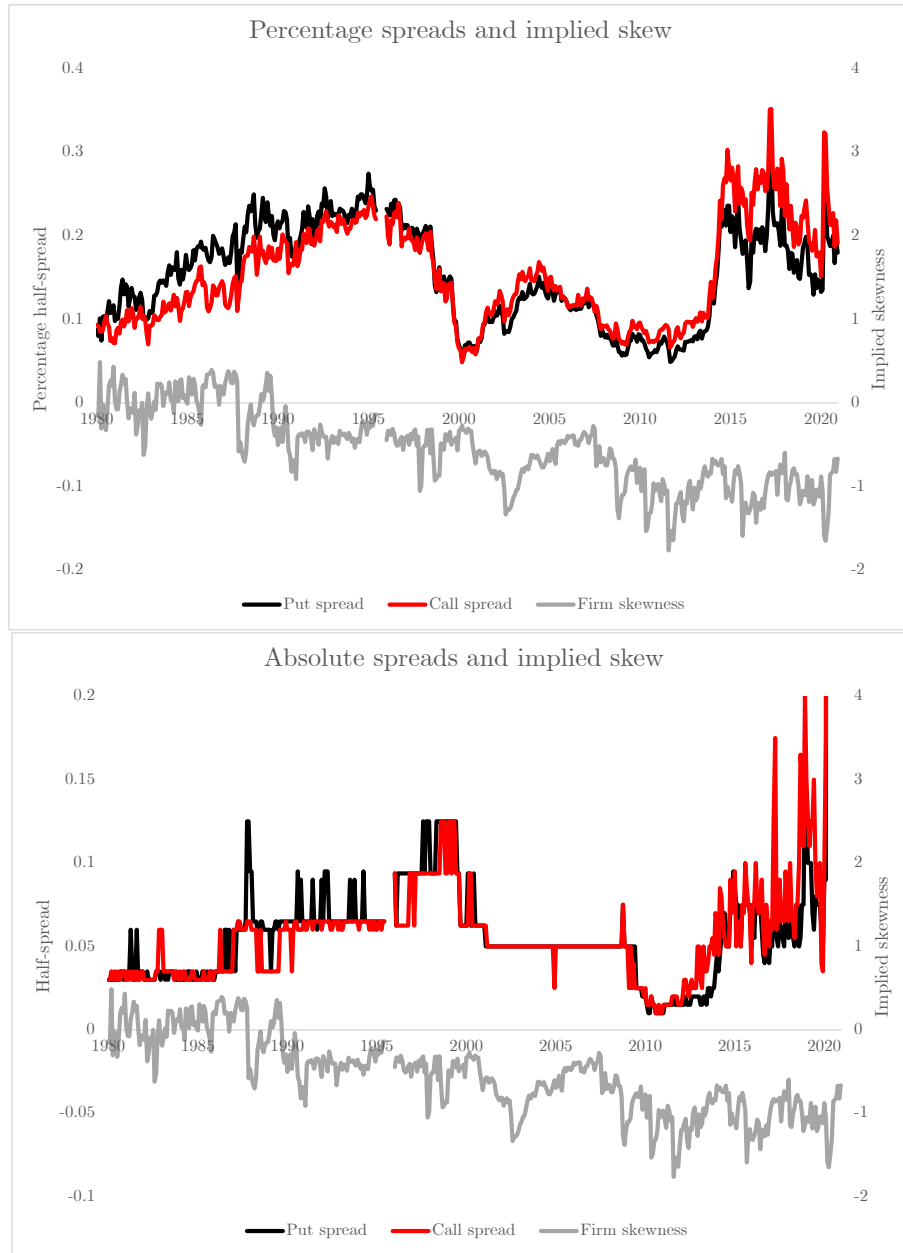
— $IS_{firm}(t+j)$ — $IS_{mkt}(t+j)$



— $IS_{firm}(t+j)$ — $IS_{mkt}(t+j)$

Note: Alternative version of the cross-correlations in which no series are detrended.

Figure A.6: Liquidity and skewness



Note: In the top panel, the put spread is for a put option with strike/spot equal to 0.9, calculated as the fitted value from the regression A.24. The call spread is for a call option with strike/spot equal to 1.1. Both are scaled by the price of the option. The bottom panel reports the median raw bid/ask spread for puts and calls with strike/spot between 0.88 and 0.92 for puts and 1.08 and 1.12 for calls. In both cases, the figure plots half-spreads.