Real-time forward-looking skewness over the business cycle

Ian Dew-Becker*

August 16, 2022

Abstract

This paper measures option-implied skewness for individual firms and the overall stock market between 1980 and 2021, giving real-time measures of conditional micro and macro skewness. There are three key results: 1. Micro skewness is significantly procyclical, while macro skewness is acyclical; 2. Micro skewness leads the business cycle and is strongly linked to credit spreads, suggesting one potential causal channel; 3. Micro skewness is significantly, and not mechanically, correlated with macro 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 at the aggregate level, including time-varying disaster risk, along with 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, often (though not exclusively) at just an annual frequency. For example, research has measured the skewness of the distribution of annual firm-level employment growth and

*Northwestern University and NBER. I appreciate helpful comments from Stefano Giglio, Simon Oh, Laura Veldkamp, Matthias Kehrig, Hugues Langlois, Thiago Ferreira, Niels Gormsen, and participants at the NBER Summer Institute, Canadian Derivatives Institute, and Copenhagen Busines School. The Berkeley Options Database was saved by Stewart Mayhew, without whom this research would not be possible.

[1]On 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).

annual individual income growth. There is also a literature on realized skewness of firm stock returns at the monthly level.[2]

But in models in which time-varying risk – whether through volatility or skewness – is a driving force (e.g. Salgado, Guvenen, and Bloom (2020)), it is typically agents' ex ante beliefs about *future* skewness that matter. Forward-looking decisions, such as firm- and household-level investment, depend (by definition) on beliefs about the moments of future shocks, rather than the moments of the shocks that were already realized.[3] And realizations are always equal to expectations plus an error, meaning that they are noisy relative to the true conditional expectation. The transitory or "measurement error" component of cross-sectional stock return skewness in fact accounts for over half of its variance.

Beyond that, 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.[4] Policymakers also benefit from having measures available in real time that can be used for decision-making, which explains the wide use of the VIX for measuring real-time uncertainty, even though it is suboptimal from a theoretical perspective since it does not measure fundamental economic uncertainty.

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 significantly benefits from high-frequency, real-time measures of conditional moments.

Finally, past work has not distinguished the importance of the skewness of idiosyncratic versus aggregate shocks. There are independent analyses of both aggregate (macro) and firm- and household-level (micro) skewness, but nothing that analyzes them jointly. To see the macro/micro distinction, consider a simple specification,

$$x_{i,t+1} = \mu_{t+1} + \varepsilon_{i,t+1} \tag{1}$$

where $x_{i,t+1}$ is some firm-level outcome, μ_{t+1} is an aggregate shock, and $\varepsilon_{i,t+1}$ is an idiosyn-

[2]For 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), Salgado, Guvenen, and Bloom (2020), Gomez, Haddad, and Loualiche (2020), and Stockl and Kaiser (2021) study skewness. The evidence on stock returns in those papers is the exception to the general pattern that research has typically looked at low frequencies. However, the work on stock returns is on skewness in the cross-section of realized returns, not conditional skewness at either the aggregate or firm level.

[3]This is not to say that realized moments cannot matter – they will, for example, any time reallocation has real effects, for example if firms or households face borrowing constraints so that the cross-sectional distributions of firm equity and household wealth become relevant.

[4]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).

cratic shock. We then have that the date-t third moment (assuming the shocks are mean-zero and independent) is

$$\underbrace{E_t\left[x_{i,t+1}^3\right]}_{\text{Total}} = \underbrace{E_t\left[\mu_{t+1}^3\right]}_{\text{Macro}} + \underbrace{E_t\left[\varepsilon_{i,t+1}^3\right]}_{\text{Micro}} \tag{2}$$

Skewness for a firm, $E_t\left[x_{i,t+1}^3\right]$, can come either from macro or micro contributions, and part of this paper's aim is to disentangle them in order to give guidance for how to model time-varying skewness.

Contribution

This paper's fundamental contribution is to develop real-time forward-looking measures of the conditional skewness of both idiosyncratic and aggregate shocks (ε and μ, respectively) from option prices.[5] 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 both micro *and* macro 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

The paper's key finding is that skewness at the *firm level* – micro skewness – is significantly procyclical. In each of the past six recessions, firm-level skewness became noticeably more negative – falling by one standard deviation on average. Skewness of the overall stock market – macro skewness – is, on the other hand, actually slightly less negative than usual during NBER-dated recessions and acyclical overall. Those results distinguish among theories of skewness, implying that to the extent that exogenous skewness shocks are important drivers of the business cycle, it is skewness in firm-specific, or micro, rather than macro shocks that matters. Past work, such as Salgado, Guvenen, and Bloom (2020) has not been able to distinguish between fluctuations in micro and macro skewness skewness.

The paper's second result is that macro volatility and micro 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

[5] While 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. Iseringhausen and Theodoridis (2021) also develop a measure of quarterly aggregate skewness.

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 macro volatility and micro 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 Kozienauskas, 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.

Salgado, Guvenen, and Bloom (2020) also provide evidence that cross-sectional (micro) skewness is procyclical. The marginal contribution of this paper is, again, in looking at a real-time, conditional measure, one that is available at the monthly instead of annual frequency used there, and that can distinguish between shocks to macro and micro skewness. The higher frequency allows for a much more detailed investigation not just of cyclicality but also of the lead-lag structure of the data along with the ability of the conditional measure to forecast realized outcomes. Distinguishing the behavior of firm and aggregate skewness is useful for testing different models – e.g. those driven by macro risk shocks like time-varying disaster risk, versus those driven by micro behavior, like Ilut, Kehrig, and Schneider (2018).

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 tight relationship in the data between credit spreads and the conditional third moment faced by firms – they have a correlation of over 80 percent. However, Gourio's (2013) model is about macro skewness, while Christiano, Motto, and Rostagno (2014) is driven by micro risk, so the latter is more consistent with the results of this paper.

Credit spreads are 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. As is well known in macroeconomics in general, and in the literature on uncertainty shocks in particular, identifying the direction of causation is a subtle affair.[6] This paper has no particular contribution to make in terms of identification. Its contribution is entirely in

[6] There is a long-running debate in the uncertainty literature on the direction of causation between output and uncertainty. See, e.g., Berger, Dew-Becker, and Giglio (2021) and Ludvigson, Ma, and Ng (2021). Ilut, Kehrig, and Schneider (2018) discuss the same issue with regard to skewness.

measurement.

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 (2018) 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.[7]

Since the paper's measure of skewness comes from option prices, it is influenced by risk premia (just like the VIX). I therefore replicate the analysis with a skewness measure that is robust to the presence of time-varying risk premia and show that the results are unaffected.

Additional related literature

Beyond the work described above, this paper is also related to research using stock options to estimate investor views about conditional distributions, such as CBOE's VIX and SKEW indexes, which measure the conditional volatility and skewness of the S&P 500 (see also Kozhan, Neuberger, and Schneider (2013) and Kozlowski, Veldkamp, and Venkateswaran (2020)). Those are both aggregate moments, whereas this paper's novel contribution to measurement is in the firm-level skewness measure.

There is also a broader finance literature studying firm-level skewness, but almost entirely in the context of option pricing or return forecasting. Xing, Zhang, and Zhao (2010), Conrad, Dittmar, and Ghysels (2013), and An et al. (2014) study forecasts based on firm-level option-implied skewness. See also Bakshi, Kapadia, and Madan (2003) and Dennis and Mayhew (2002). Bollerslev and Todorov (2011) study measures of the extreme tails of the return distribution and their relationship with risk premia. This paper, instead of being focused on forecasting returns at either the firm or market level, is instead aimed at disentangling aggregate and idiosyncratic skewness and measuring their relationship with the real economy.

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) and Plagborg-Møller et al. (2020)), discussed in section 4.4.

[7]See also Kozienauskas, 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.

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.[8]

This paper measures skewness throughout as the scaled third moment. For a generic random variable x_{t+1},

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

where E_t denotes the statistical expectation conditional on information available at time t. The reason to use the skewness coefficient as opposed to the third moment is that the skewness coefficient is not mechanically driven by volatility. It measures purely the shape of the distribution of shocks, while the scale (which has been extensively studied elsewhere) can be analyzed separately.

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 (see section 4.5 and appendix A.2.2).

I measure implied skewness using the formula of Kozhan, Neuberger, and Schneider (2013). Option *I*mplied *S*kewness is IS_t, where

$$IS_t \equiv 3\frac{\int_0^\infty \frac{K-F_t}{K^2 F_t} O_t\left(K\right) dK}{\left[\int_0^\infty \frac{1}{K^2} O_t\left(K\right) dK\right]^{3/2}} \tag{4}$$

$O_t\left(K\right)$ is the price of an out-of-the-money option at strike K on date t and F_t is the forward

[8] Past work also studies skewness of endogenous outcomes, such as income (Guvenen et al. (2014); Busch et al. (2021)) and employment (Ilut et al. (2018); Decker et al. (2016)).

price of the underlying asset. Intuitively, skewness is measured as the difference of prices of options with strikes above the current forward price from those below. The numerator is the conditional third moment and the denominator is the conditional second moment, just like in equation (3), but under the risk-neutral probability measure, which adjusts for state prices.[9]

2.2 Data

I calculate implied skewness at the market level using S&P 500 option prices from the CME (formerly Chicago Mercantile Exchange) for 1983–1995 and Cboe (formerly Chicago Board Options Exchange) for 1996–2021. The analysis here is entirely in terms of monthly skewness – using options with a one-month average maturity. That yields comparability with the VIX 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.[10] I apply standard filters to isolate relatively liquid options (appendix A.1), and section A.4.3 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 1/1996–12/2021, 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)).

It is important to emphasize that the results here are for the largest firms in the economy – none of the firms in this paper's data are "small" by almost any definition. That will be more appropriate for testing 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)). Figure A.1 plots the fractions of aggregate market capitalization and employment covered by the sample. The BODB sample covers about 30% of market capitalization, while the

[9] If M_{t+1} is the stochastic discount factor, then the risk-neutral expectation operator is $E_t^Q[x_{t+1}] = E\left[\frac{M_{t+1}}{E_t[M_{t+1}]}x_{t+1}\right]$ for any random variable x_{t+1}.

[10] Dew-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 smooths out distributions).

main set of firms in the Optionmetrics sample is similar, though somewhat larger, covering about 50%. Coverage of aggregate employment rises from about 6% in the early part of the sample to 11% by the end. That smaller fraction is consistent with the fact that only about 1/3 of total employment is in publicly traded corporations (Davis et al. (2006), Dew-Becker and Giglio (2021)).

While data through 2021 is available, the events of the spring of 2020 in particular are sufficiently extreme as to dominate certain parts of the analysis. I therefore plot and discuss the univariate properties of skewness through the end of 2021, but all comparisons with macro variables use data only through the end of 2019.[11]

Since the analysis uses data on options for individual stocks, it applies to the largest corporations in the US, and, as noted above, the analysis is focused on the biggest of the big. In addition, the baseline skewness index will be weighted by stock market capitalization, similar to how stock indexes like the S&P 500 are constructed (which then also naturally influences the construction of the VIX, which measures S&P 500 uncertainty). Stock market capitalization differs from other measures of importance, like employment or sales shares (which are sufficient statistics for importance in network production models, for example). Section 3.4 shows that the skewness series is highly similar when it is modified with weights designed to make it more representative of the broad economy, implying that the exact composition and weighting are not major drivers of the results.

Finally, an important practical issue is that implied skewness is measured theoretically as an integral across strikes, but in the data only discrete strikes are available. That is in general addressed either through interpolation of observed option prices or the use of rectangular or trapezoidal numerical integration rules. Appendix A.1 describes the methods used here, which involve interpolation via a Gaussian process regression. That method has the added benefit that it helps reduce the influence of any outliers or errors in the data.[12]

2.3 Skewness indexes

Denote firm i's implied skewness on date t by $IS_{i,t}$. Implied skewness is calculated using equation (4), with the integral constructed using methods described in appendix A.1.

[11] Prior 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.

[12] A more common way to address outliers is to use Kelley skewness – based on quantiles of the distribution – instead of the scaled third moment. The drawback of Kelley skewness in this setting is that it is relatively difficult to calculate from option prices, and in fact can be more sensitive to data issues in option prices than the third moment. I have calculated Kelly skewness in the cross-section and find it to be highly similar to the scaled third moment, but significantly more noisy in practice. Kelley skewness is also less straightforward to compare to a realized outcome, making it more difficult to validate (see section 3.2).

Average firm-level skewness on date t is then

$$IS_{firm,t} \equiv \frac{\sum_i IS_{i,t} mkt_{i,t}}{\sum_i mkt_{i,t}} \tag{5}$$

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 (4) and S&P 500 option prices.

I calculate monthly values for skewness by simply averaging observations within each month.

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, consider a single-factor specification,

$$r_{i,t} = \beta_{i,t} r_{mkt,t} + \varepsilon_{i,t} \tag{6}$$

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 $r_{mkt,t}$. $\beta_{i,t}$ is firm i's loading on the market on date t. The total skewness at the firm level is then, formally, the sum of two components

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

The micro contribution to firm skewness, 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, $\varepsilon^2_{i,t}$, and the return on the overall market (the example in the introduction assumed away the latter effect via independence).[13]

Following Campbell et al. (2001), I use the approximation $\beta_{i,t} \approx 1$ in calculating idiosyncratic skewness because it reduces the variance of the estimator (see also Dew-Becker and

[13] Formally, the $\varepsilon_{i,t}$ need not be independent across firms, so "idiosyncratic" here just means uncorrelated with aggregate returns.

Giglio (2021)). Option-implied idiosyncratic skewness is then measured empirically as

$$IS_{idio,t} = \left(1 - \frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{-3/2} \left(IS_{firm,t} - \left(\frac{IV_{mkt,t}^2}{IV_{firm,t}^2}\right)^{3/2} IS_{mkt,t}\right) \tag{8}$$

where $IV_{mkt,t}$ and $IV_{firm,t}$ are option-implied volatilities (measured as in the denominator of (4)) for the S&P 500 and averaged across individual firms.[14] Estimates of skewness accounting for $\beta_{i,t}$ are highly similar to the baseline (with a correlation over 90%).

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 1980–2021, compared to 0 to -1 for firm skewness. A significant fraction of the total declines in aggregate and firm skewness come in two episodes: the October, 1987 crash, which had the largest single-day decline in the history of the S&P 500, and a second large fall in 1989. 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 cyclicality 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. That said, much of the variation in skewness happens outside of recessions. Such behavior is also observed for stock market volatility as measured by the VIX. While the VIX typically rises in recessions, similar to firm skewness, both series also have significant variation outside recessions associated with shocks to financial markets, e.g. the various aftershocks to the financial crisis between 2010

[14]Note 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. The average of firm level idiosyncratic skewness is highly similar to the baseline (over 90% correlated), but noisier and has problems with large outliers in the early part of the sample when there is less data.

and 2012.

Firm and market skewness both share a downward trend over time. The remainder of the paper therefore reports results both for the raw data and after detrending using a backward-looking local linear trend model (Harvey (2006)). The same detrending method is used on the macro time series below (see appendix A.3 for details and robustness to alternative choices, including no detrending). Detrended variables are denoted by a circumflex, e.g. $\widehat{IS}_{firm,t}$. Appendix A.3 discusses the interpretation of the downward trend. It is shared between both implied and realized skewness for all three measures (Figure 2), but bigger for implied, consistent with a changing risk premium.

The correlation between firm and market skewness is surprisingly weak. They have an overall correlation of 0.64, but that is entirely due to their shared time trend. Once the trend is removed from both series, the correlation falls to 0.01. So while both have trended down, their business cycle and higher-frequency variation is independent.

Shocks to implied skewness are notably short-lived. Figure 5 shows that the autocorrelations in the three skewness measures decay far more quickly than those for employment, industrial production (IP), and credit spreads – to zero within a year, compared to 0.4–0.5 for the macro variables.

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 (most relevant for macro models), and annual autocorrelations, both with and without removing a time trend.

3.2 Validating implied skewness

Since skewness is a scaled moment, true conditional skewness does not exactly equal the statistical expectation of the realized skewness coefficient in equation (3) 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. I measure realized skewness here using methods in Neuberger (2012) (see appendix A.2). Figure 2 plots implied and realized skewness, showing that while they follow similar paths, realized skewness appears much noisier, as is expected for any realized moment (thus underlining the value of using the option-implied measure).

Table 1b reports results of forecasting regressions for realized skewness (realized skewness here naturally being measured for the same sample of firms used to calculate implied 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 that is independent of the information contained in lagged realized skewness. The remaining four columns show that similar results hold for market-level and idiosyncratic skewness.

Table 1b does not include R^2s because there is no reason to expect a conditional moment to explain all or even the majority of the variation in a realized moment. The next section quantifies the unpredictable variation in realized skewness and discusses how it can be interpreted.

Section 4.5 discusses an additional validation concern, the effects of time-varying risk premia.

3.3 Relationship with other measures of skewness

3.3.1 Realized stock return skewness

As discussed in the introduction, there is past work studying the skewness of realized stock returns. In addition to being theoretically suboptimal in that many models are driven by conditional rather than realized returns, realized skewness also has the drawback that it is contaminated with transitory "measurement error". That is, any sample moment is equal to the true conditional moment plus noise, and Figure 2 shows that component appears to be large.[15] The previous section shows that, as one would hope for a conditional moment, implied forecasts realized skewness. In addition, though, it is possible to estimate the magnitude of the transitory "measurement error" in realized skewness.

There are two ways to do that. First, one could take implied skewness as the true conditional moment and then regress realized on implied skewness. In that case, the fraction of the variation in realized skewness that represents noise would be $1 - R^2$ in the regression. In table 1b, that value is 79 percent at the firm level (i.e. the R^2 is 0.21), indicating that the vast majority of the variation in realized skewness is just transitory noise. That is likely an overestimate of the importance of measurement error, though.

An alternative way to measure the amount of noise in monthly realized skewness, which does not require assuming that option-implied skewness is a true conditional moment, is to look at its autocovariances. Under the assumption that true skewness is close to a random

[15] While one might appeal to the law of large numbers when studying the cross-section of stock returns to argue that the sample skewness is the "true" skewness, that would require the assumption that returns are i.i.d., which is a poor description of the data: stock returns are neither identically distributed not are they cross-sectionally independent (see, e.g., the large literature on factor pricing models).

walk at short horizons – i.e. its first difference is approximately serially uncorrelated – the variance of the "measurement error" is equal to minus the first autocovariance. Such a calculation makes it possible to get an estimate for measurement error in both implied and realized skewness. Alternatively, under the assumption that true skewness follows an AR(1) process, realized skewness will follow an ARMA(1,1) process, from which the variance of the transitory component can be estimated. The latter method reduces bias by not imposing the random walk assumption, but it is less precise since it requires the estimation of an ARMA model.

The table below reports results from the random-walk and AR(1) models for the amount of purely transitory variation in both implied and realized skewness. In all cases, at least half the variation in realized skewness at both the firm and idiosyncratic level is estimated to be purely transitory, compared to less than 10 percent, and often less than 5 percent, for implied skewness (and the differences between the random-walk and AR(1) models are small). These results are consistent with what one might expect from simple visual inspection of Figure 2.

Fraction of variance due to transitory variation (measurement error)

	Firm		Market		Idio.	
	Realized	Impl.	Realized	Impl.	Realized	Impl.
AR(1)	0.51	0.03	0.38	0.02	0.55	0.06
Random-walk	0.54	0.07	0.39	0.06	0.63	0.15

Overall, the results in this section help quantify the benefit from using conditional moments. For example, if a policymaker were using real-time data to evaluate the tail risk faced by individual firms, their most recently monthly observation would be about half noise, whereas the implied skewness measure would be only 3 percent noise.

3.3.2 Other cross-sectional skewness concepts

The top-left panel of Figure 3 plots $\widehat{IS}_{firm,t}$ and the skewness of media news about firms from the Ravenpack database of news articles (reported quarterly to reduce noise; see appendix A.4.1). Briefly, it is constructed as the cross-sectional skewness of Ravenpack's "Event Sentiment Scores", which are designed to measure the extent to which a given news event is associated with positive or negative sentiment for a firm. The text-based measure, which is novel to this paper and represents a new way to test the validity of option-implied moments, 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 two series are clearly strongly related, with an overall correlation of 0.52. Table 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.

The bottom panels on the left-hand side of Figure 3 plot $\widehat{IS}_{firm,t}$ against two alternative measures of cross-sectional skewness in annual growth rates: sales growth in Compustat and employment growth in the Census LBD.[16] These are again measures of realized rather than conditional skewness.

In both panels, there is a clear positive correlation – both about 0.5 – between $\widehat{IS}_{firm,t}$ 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.4.2 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)).

3.4 Robustness to design choices

The basic characteristics of the skewness index described so far are robust to a number of variations in the calculation method. None of these are major changes (like switching to looking at skewness over an outcome different from stock returns), but they show that the features of the data that the paper focuses on are not overly sensitive to small choices in the design of the index.

One question about the index is how representative it is. As described above, for the sake of consistency over the sample, the analysis focuses on the top 200 firms ranked by stock market capitalization. There are two concerns with that. First, smaller firms might behave differently. To shed light on that possibility, the top-left panel of Figure 4 plots skewness separately for the top 100 firms and those ranked 101–200 by size (just for the Optionmetrics sample). While that may seem like a small difference, due to the power law distribution of firm size there are in fact significant differences in size across those ranks – the top 100 firms are larger than those ranked 101–200 by a factor of 5 on average. Figure 4 shows that the skewness index is essentially unaffected by this change.[17]

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

[17]Ideally one might want to construct a skewness index for all publicly traded firms. As one goes further down in the size distribution, though, liquidity for options gets much worse (and also becomes endogenous

The second alternative is to look at skewness without weighting by market capitalization. The top-right panel of Figure 4 shows that the results are again quantitatively highly similar. Inevitably any measure based on publicly traded firms will be based on large corporations, but these results imply that it is not heavily influenced by value weighting among public firms.

An additional concern with weighting by market value is that market valuations may differ significantly from more fundamental measures of size like sales or employees. Furthermore, the industry composition of the stock market is not the same as the industry composition of the overall economy (Byun et al. (2022)). The middle-left panel of Figure 4 shows that when the firms in the sample are reweighted by industry to match observed industry employment shares from the BLS, the results are again minimally affected.[18] The middle-right panel of Figure 4 weights instead by firm employment reported by Compustat and again comes to similar conclusions.

The fact that the results are insensitive to the various weighting choices implies that either there is not much cross-sectional variation in implied skewness or it is not correlated with any of these factors. To see that, the bottom-left panel of Figure 4 plots the 10th, 50th, and 90th percentiles of the cross-sectional distribution of firm-level implied skewness (weighted by market capitalization). Because the series are so similar, I plot them quarterly to make the differences more visible. While there is clear cross-sectional variation, the quantiles move together extremely closely, and none of the conclusions of the paper would be affected by looking at any of them (the median, for example) instead of the mean.

The results in this section are also consistent with those reported in Dew-Becker and Giglio (2022), who find that average option-implied volatility is similarly unaffected by similar weighting and sample choices. None of this is to say that all possible measures of skewness are the same – the previous section shows they are far from identical. Rather, the details of how skewness is averaged across firms are just relatively unimportant.

4 Cyclicality 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 backward-looking detrending, with various measures of

to the distribution of returns), with few quotes and even fewer trades. In addition, calculating skewness for thousads of firms on thousands of days would be, while not necessarily infeasible, extremely time-consuming.

[18]Specifically, firms are weighted by market capitalization within three-digit NAICS industries, and then employment shares obtained from the BLS Current Employment Survey are used to weight across sectors.

real activity for the sample up to 12/2019. In all cases, firm and idiosyncratic skewness are estimated to be procyclical, with most of the correlations both statistically and economically significant, especially for detrended skewness and for the idiosyncratic component. After controlling for a time trend, firm skewness is on average lower by 0.23 in recessions, and idiosyncratic skewness is lower by 0.63, representing shifts of 0.98 and 1.25 standard deviations, respectively.

The correlations for IS_{mkt}, on the other hand, are all close to zero and there is no consistent evidence of cyclicality – some correlations imply IS_{mkt} is procyclical and others imply it is countercyclical. S&P 500 skewness is actually on average higher by 0.26 (0.61 standard deviations) in recessions. Overall, then, aggregate skewness has been slightly countercyclical over the past 40 years, while firm skewness is consistently procyclical.

To further illustrate the cyclicality of firm-level skewness, the bottom-right panel of Figure 3 plot $\widehat{IS}_{firm,t}$ against the Philadelphia Fed's Leading Indicators index. Each of the declines in the leading indicators index is associated with a decline in skewness.

4.1 Cross-correlations

The bottom four panels of Figure 5 plots the cross-correlations of implied skewness at the market and firm levels with employment, IP, the Chicago Fed National Activity index (CFNAI), and the Philadelphia Fed's Leading Indicator index. The shaded regions represent 90-percent confidence bands.

$\widehat{IS}_{firm,t}$ 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* $\widehat{IS}_{firm,t}$: skewness leads the cycle, and may even lead the leading indicators (and since the detrending method is purely backward-looking, these correlations indicate forecasting relationships).

The results for market-level skewness are again drastically different. In all four cases, the correlations are close to zero, and $\widehat{IS}_{mkt,t}$ again appears if anything to be countercyclical.

These results are robust to a number of changes to the specification, discussed in sections 4.5 and A.3 and plotted in Figures A.5 to A.8.

4.2 Aggregate and idiosyncratic contributions to firm skewness

Rearranging the definition of $IS_{idio,t}$ in equation (8), firm skewness is a time-varying mixture of market and idiosyncratic skewness,

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

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 $IV_{mkt,t}/IV_{firm,t}$. Since $IS_{mkt,t}$ tends to be more negative than $IS_{idio,t}$, when aggregate shocks have a larger scale ($IV_{mkt,t}/IV_{firm,t}$ grows), firm skewness will become more negative. The question is what effect that 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:

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

In each case, only one of the three variables determining the value of $IS_{firm,t}$ varies over time (outlined in red), 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 highly similar to total firm skewness, with a correlation of 79 percent and similar volatility. 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 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 39 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 (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 0.62 with firm skewness.

But changes in volatility only generate a time series that is 39 percent as volatile as overall $IS_{firm,t}$. The variation in $IS_{firm,t}$ thus comes from the fact that market volatility and idiosyncratic skewness move *together*.

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 cyclicality 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), among others, 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 1980–2021 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 one 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)): it has a finite upside, but 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. Gourio (2013) focuses on aggregate skewness, while Christiano, Motto, and Rostagno (2014) develop a model driven by variation in 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 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 of credit spreads with volatility and the negative relationship with skewness is through the implied 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, consistent withe firm-level explanations as in Christiano, Motto, and Rostagno (2014).

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 $\widehat{IS}_{firm,t}$ 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. The correlation of $\widehat{IS}_{firm,t}$ with volatility is -0.63. Section 4.2 shows that that relationship is not purely mechanical. Also, the correlation of $\widehat{IS}_{idio,t}$, which removes the mechanical component, with $IV_{mkt,t}$, is -0.57, illustrating that the relationship comes from skewness in firm-specific shocks.

The bottom-right panel, shows, conversely, that there is no clear relationship between $\widehat{IS}_{firm,t}$ 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 $\widehat{IS}_{firm,t}$ and $IV_{idio,t}$ reflects that, at only -0.23 (about 1/3 of the correlation between $\widehat{IS}_{firm,t}$ and $IV_{mkt,t}$).

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.

4.5 Time-varying risk premia

The analysis up to here has ignored the potential presence of time-varying risk premia. To account for time-varying risk premia, I follow the usual method in the asset pricing literature and allow premia to depend on a set of state variables, the stock market price/earnings ratio, the Gilchrist-Zakrajsek credit spread, the CBO output gap, the unemployment rate, and the S&P 500 VIX (see also Dew-Becker and Giglio (2021)).

One can show that an unbiased estimate of the true conditional moments comes from just projecting realized third moments on implied skewness and also the state variables determining risk premia. The tradeoff in such an analysis is that any method that is robust to time-varying risk premia requires more estimation and modeling choices.

Appendix A.2 presents such an analysis in detail. Instead of calculating the expectations in the numerator and denominator of the skewness coefficient purely from option prices, it is done via linear regressions. The fitted skewness coefficient is then just the ratio of the third and second conditional moments from a statistical forecast.

Figure A.4 plots the baseline option-implied skewness measure against the alternative that is robust to time-varying risk premia. The characteristics of the two series are similar, with the only clear difference being a weaker downward trend in fitted skewness. Table A.1 in the appendix replicates the results on cyclicality from Table 1. The results are again similar. Figure A.8 replicates the cross-correlation results. Overall, while there is evidence of the presence of time-variation in risk premia, it does not qualitatively affect the paper's main findings.

In the end, the results here focus on the option-based measure of skewness because it is simple and it is in accord with the long tradition of measuring market-implied moments – volatility, in particular, from the VIX, but also, for example, expectations from futures prices as in Nakamura and Steinsson (2018) (among many others). Asset prices give simple measures of expectations that require no additional estimation or specification choices, but at the cost of contamination via risk premia.

5 Conclusion

This paper provides the first real-time, high-frequency measure of conditional skewness available jointly at the micro and macro levels. Its key finding is that macro skewness is essentially acyclical, but micro 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. That is, macro volatility is countercyclical and macro skewness acyclical, while micro volatility is acyclical and micro skewness procyclical.

References

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone, "Vulnerable growth," *American Economic Review*, 2019, *109* (4), 1263–89.

An, Byeong-Je, Andrew Ang, Turan G Bali, and Nusret Cakici, "The joint cross section of stocks and options," *The Journal of Finance*, 2014, *69* (5), 2279–2337.

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.

Berger, David, Ian Dew-Becker, and Stefano Giglio, "Uncertainty shocks as second-moment news shocks," *The Review of Economic Studies*, 2020, *87* (1), 40–76.

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.

Byun, Sung Je, Jonathan Loudis, and Lawrence Schmidt, "Market Factor Measurement Error and a Reevaluation of the Risk-Return Tradeoff."

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.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda, "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER Macroeconomics Annual*, 2006, *21*, 107–179.

Decker, Ryan, Pablo N. D'Erasmo, and Hernan Moscoso Boedo, "Market exposure and endogenous firm volatility over the business cycle," *American Economic Journal: Macroeconomics*, 2016, *8* (1), 148–198.

Dennis, Patrick and Stewart Mayhew, "Risk-neutral skewness: Evidence from stock options," *Journal of Financial and Quantitative Analysis*, 2002, *37* (3), 471–493.

Dew-Becker, Ian, Alireza Tahbaz-Salehi, and Andrea Vedolin, "Macro skewness and conditional second moments: evidence and theories," 2021. 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.

Gomez, Matthieu, Valentin Haddad, and Erik Loualiche, "The case of disappearing skewness," 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.

Hansen, Lars Peter and Robert J Hodrick, "Forward exchange rates as optimal predictors of future spot rates: An econometric analysis," *Journal of Political Economy*, 1980, *88* (5), 829–853.

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

Harvey, Andrew, "Forecasting with unobserved components time series models," *Handbook of economic forecasting*, 2006, *1*, 327–412.

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.

Iseringhausen, Martin and Konstantinos Theodoridis, "Aggregate Skewness and the Business Cycle," 2021. Working paper.

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.

Kozeniauskas, 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.

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

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.

Nakamura, Emi and Jón Steinsson, "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

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

Plagborg-Møller, Mikkel, Lucrezia Reichlin, Giovanni Ricco, and Thomas Hasenzagl, "When is growth at risk?," *Brookings Papers on Economic Activity*, 2020, *2020* (1), 167–229.

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.

Stöckl, Sebastian and Lars Kaiser, "Higher moments matter! Cross-sectional (higher) moments and the predictability of stock returns," *Review of Financial Economics*, 2021, *39* (4), 455–481.

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.

Xing, Yuhang, Xiaoyan Zhang, and Rui Zhao, "What does the individual option volatility smirk tell us about future equity returns?," *Journal of Financial and Quantitative Analysis*, 2010, *45* (3), 641–662.

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	0.23	0.59	0.42	0.19
	[0.07]	[0.04]	[0.01]	[0.03]	[0.07]	[0.02]	[0.05]	[0.06]	[0.08]
$IS_{mkt,t}$	-1.66	0.84	0.94	0.86	0.73	0.43	0.60	0.41	0.21
	[0.13]	[0.12]	[0.02]	[0.04]	[0.09]	[0.04]	[0.06]	[0.06]	[0.07]
$IS_{idio,t}$	-0.53	0.75	0.89	0.77	0.57	0.51	0.59	0.41	0.19
	[0.11]	[0.09]	[0.02]	[0.05]	[0.09]	[0.07]	[0.07]	[0.08]	[0.09]

(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]
$RS_{x,t}$		0.25		0.30		0.28
		[0.06]		[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.	0.26 [0.14]	0.32 [0.12]	0.06 [0.12]	-0.09 [0.09]	0.33 [0.15]	0.35 [0.10]
Employment	0.31 [0.14]	0.27 [0.10]	0.13 [0.13]	-0.09 [0.10]	0.35 [0.13]	0.28 [0.09]
Unemployment	-0.19 [0.15]	-0.33 [0.10]	-0.07 [0.11]	0.15 [0.11]	-0.26 [0.15]	-0.34 [0.10]
CBO output gap	0.08 [0.13]	0.12 [0.09]	-0.12 [0.10]	0.02 [0.10]	0.30 [0.10]	0.11 [0.08]
NBER rec. ind.	-0.03 [0.11]	-0.30 [0.11]	0.02 [0.09]	0.19 [0.11]	-0.21 [0.12]	-0.32 [0.13]
Capacity util.	0.14 [0.14]	0.33 [0.12]	0.07 [0.11]	-0.20 [0.09]	0.20 [0.16]	0.36 [0.11]
Empl. growth	0.14 [0.08]	0.24 [0.10]	-0.01 [0.05]	-0.16 [0.09]	0.25 [0.12]	0.32 [0.14]
IP growth	0.13 [0.08]	0.17 [0.10]	0.01 [0.03]	-0.12 [0.09]	0.21 [0.12]	0.23 [0.14]

(d) Credit spread regressions

Firm moments		Market moments		Market vs idiosyncratic	
$IS_{firm,t}$	-1.06	$IS_{mkt,t}$	-0.23	$I3_{idio,t}$	-22.15
	[0.17]		[0.07]		[2.65]
$IV_{firm,t}$	6.53	$IV_{mkt,t}$	11.00	$I3_{mkt,t}$	-13.00
	[1.68]		[2.43]		[3.71]
R^2	0.62		0.55		0.66

Notes: Standard errors, reported in brackets, are calculated using the Newey–West method with 12 lags (results are nearly identical using the method of Hansen and Hodrick (1980)). For details on detrending (panels a and c only) see section A.3. The sample in panels (c) and (d) stop in 12/2019.

Table 2: **Forecasting news skewness**

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

Notes: Forecasting regressions for news skewness from Ravenpack (see appendix A.4.1). 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 1: Implied skewness and its components

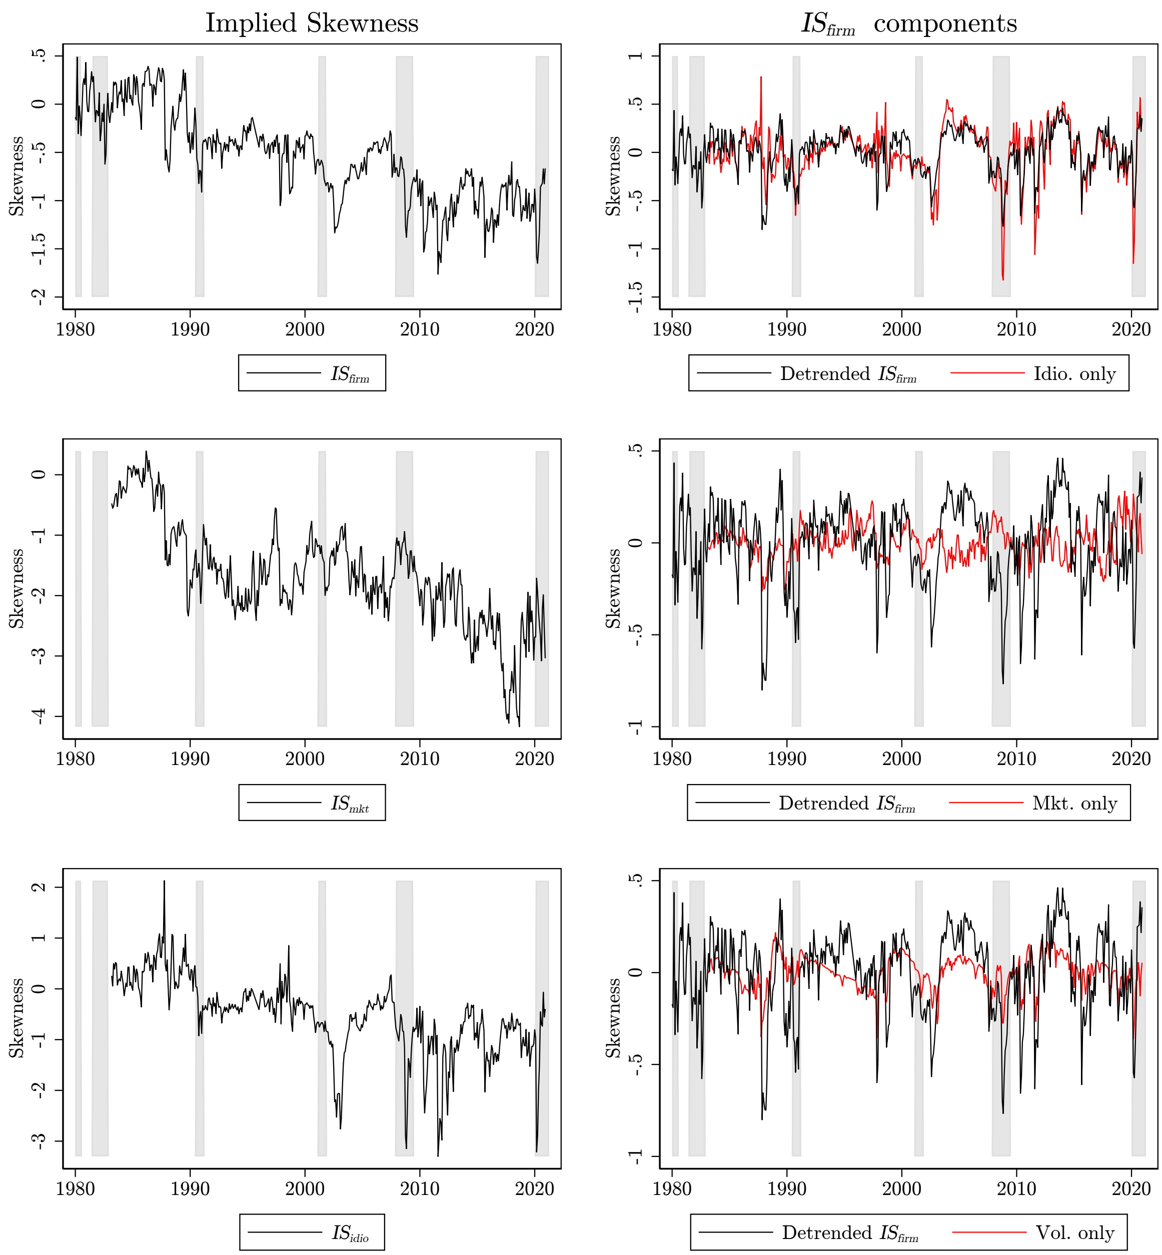

Note: The left-hand panels plot option-implied skewness at the firm, market, and idiosyncratic levels. The right-hand panels plot total firm skewness when just one of its three components varies over time, with the other two left at their unconditional means. The series in the right-hand panel are detrended (see section A.3). Gray bars indicate NBER-dated recessions.

Figure 2: Implied versus realized skewness

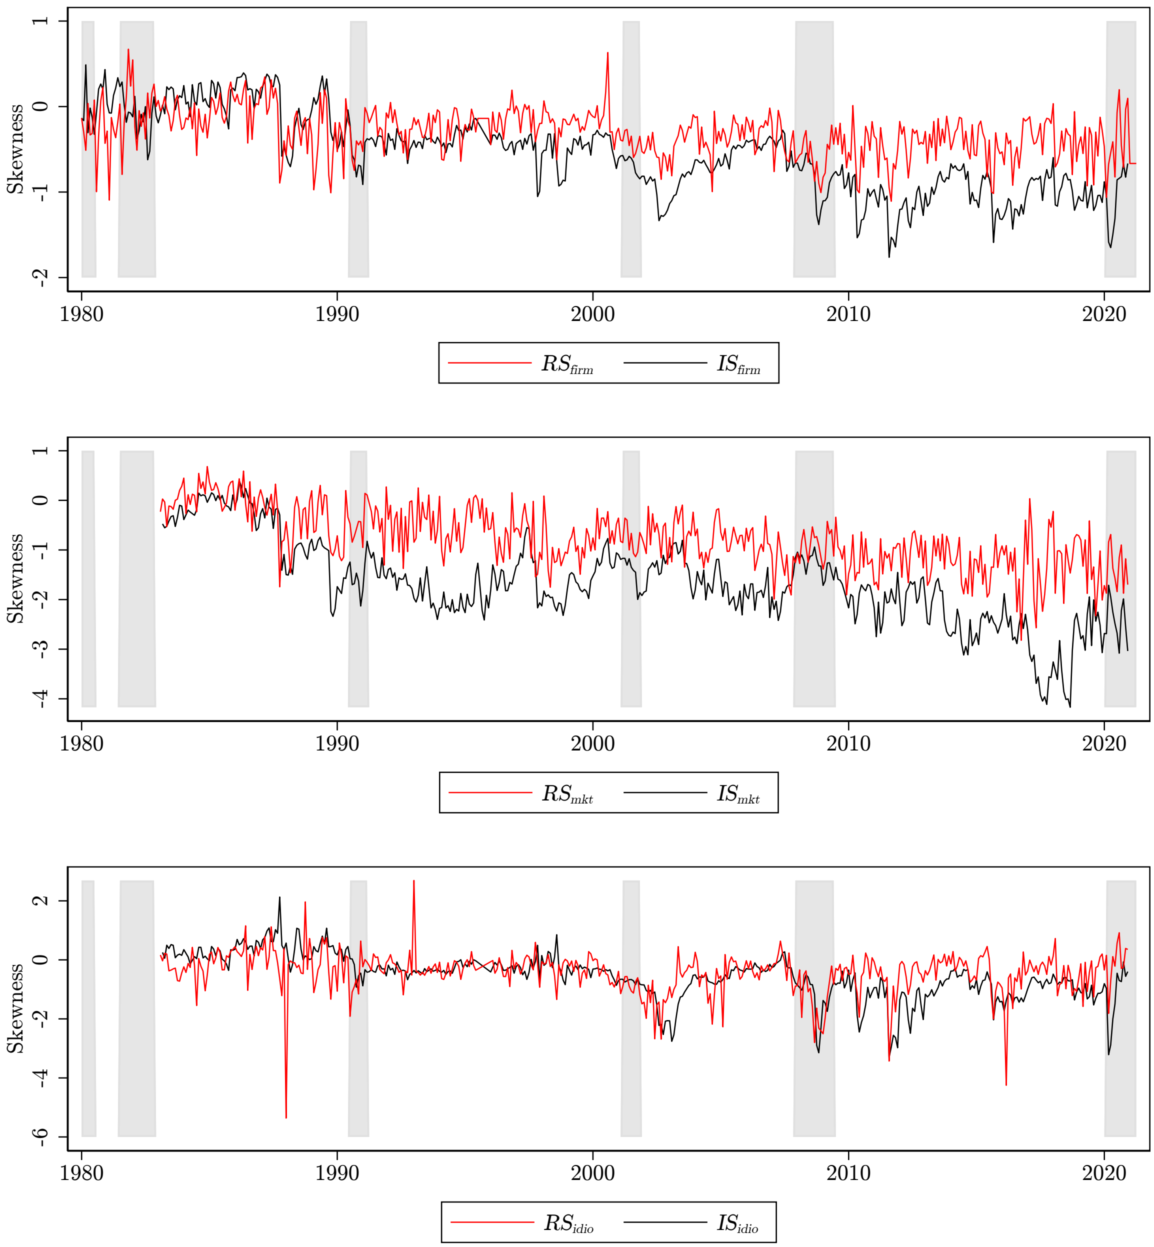

Note: Implied and realized skewness

Figure 3: Comparisons with other variables

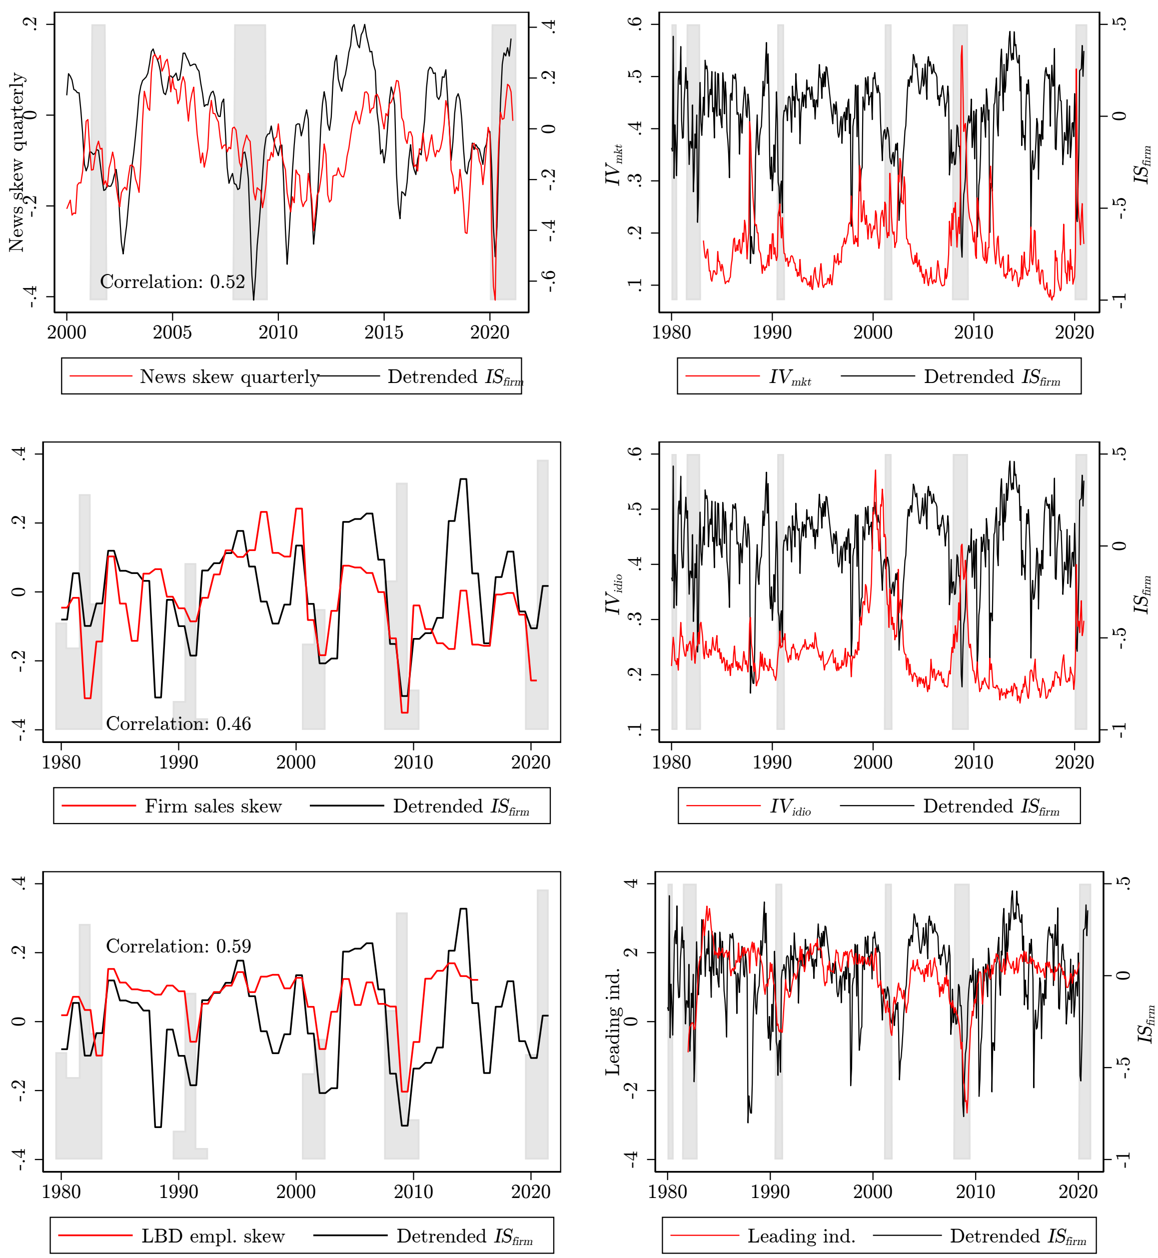

Note: The left-hand panels plot $\widehat{IS}_{firm,t}$ (detrended implied skewness) in black against three alternative measures of skewness in red. The scales for the alternatives are shifted to match $\widehat{IS}_{firm,t}$. In the bottom two panels on the left, the gray bars represent the fraction of months in each year that the economy was in recession. The right-hand panels plot $\widehat{IS}_{firm,t}$ in black against market and idiosyncratic implied volatility ($IV_{mkt,t}$ and $IV_{idio,t}$, respectively) and the Philadephia Fed's leading indicators index.

Figure 4: Alternative design choices

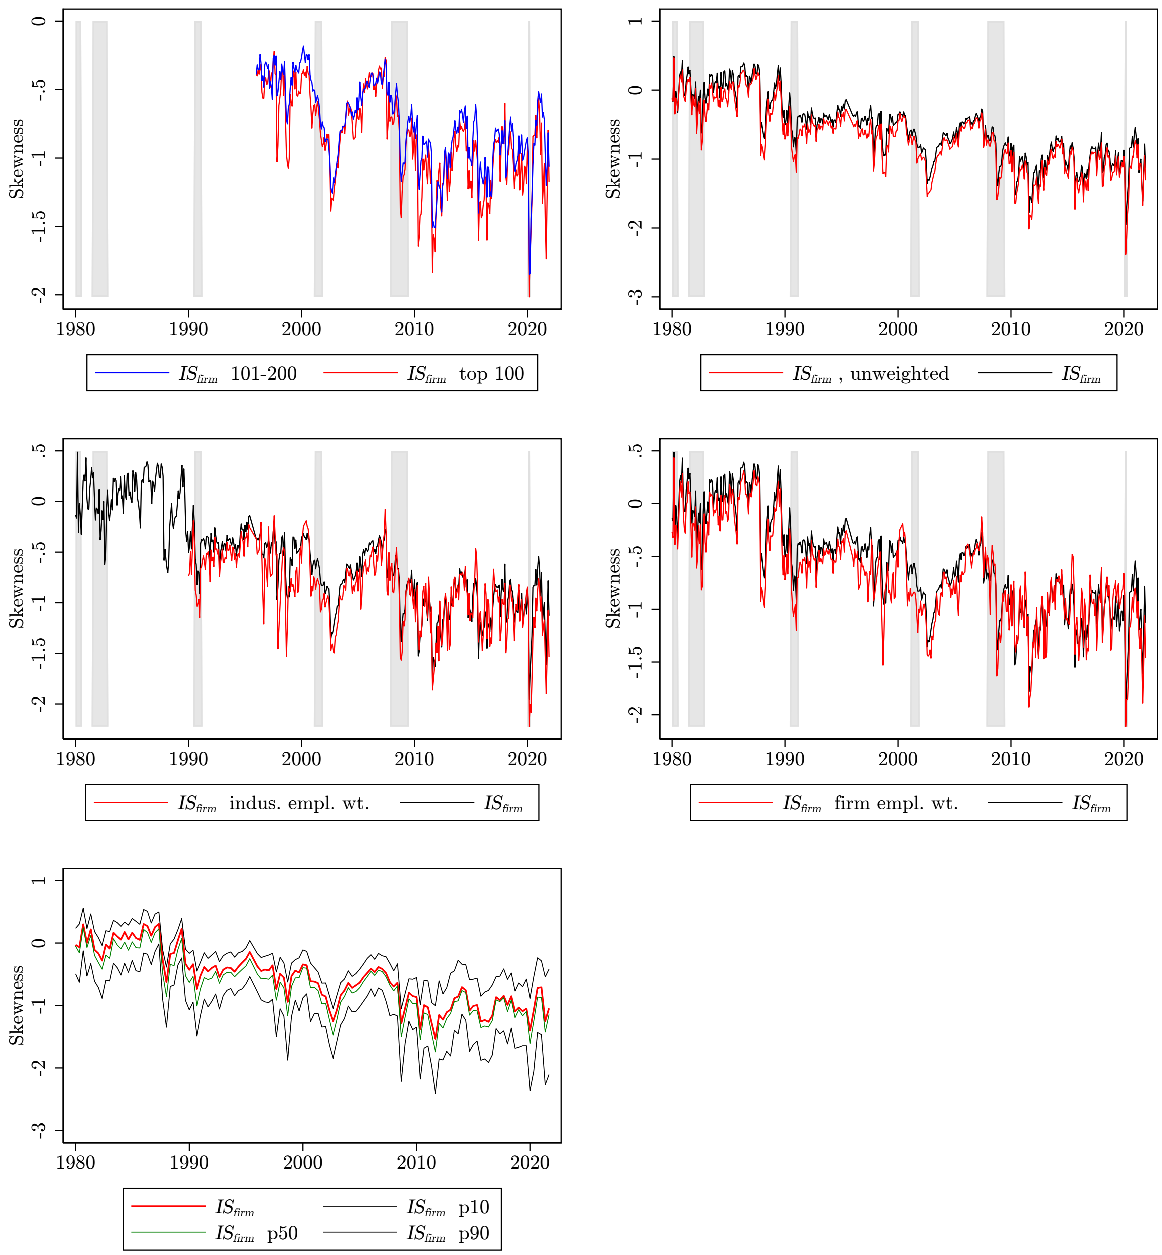

Note: Top-left: $IS_{firm,t}$ just for the top 100 or 101st to 200th ranked firms in the Optionmetrics sample instead of the top 200. Top-right: all firms weighted equally instead of by market capitalization. Middle-left: market capitalization weights within industries, BLS sector employment weights across industries. Middle-right: employment weights (from Compustat) instead of market capitalization. Bottom-left: percentiles of the cross-sectional distribution of implied skewness (weighted by market capitalization).

Figure 5: Auto- and cross-correlations

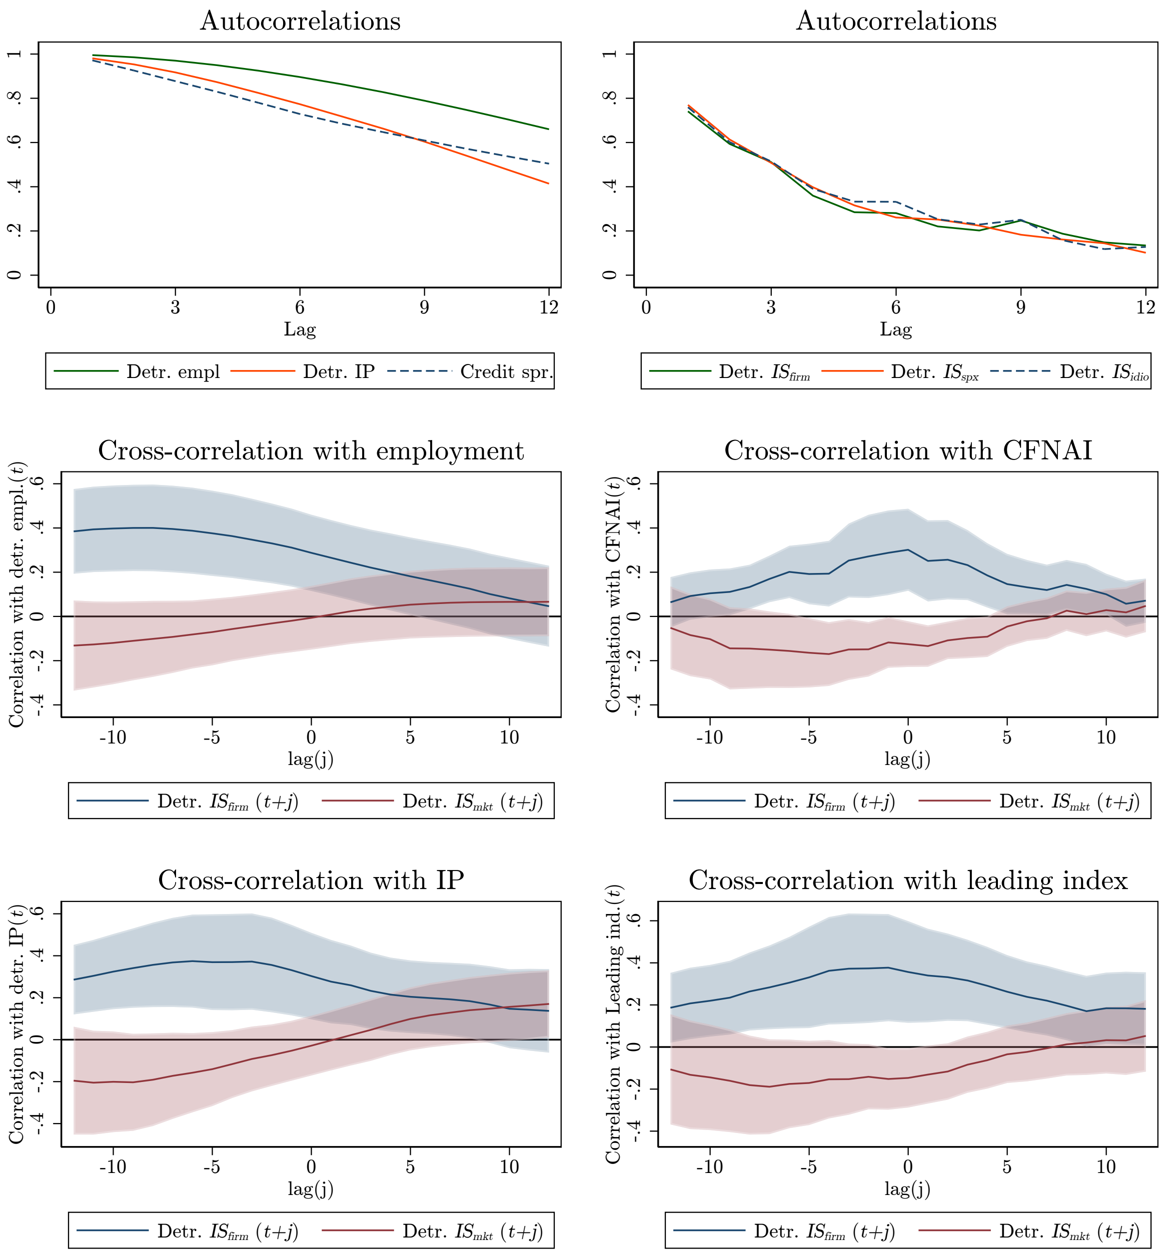

Note: The top two panels plot autocorrelations of various series. The remaining panels plot cross-correlations of detrended implied skewness at the firm and market levels with other variables. Shaded regions represent 90-percent confidence intervals.

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). The price of an out-of-the-money option on stock i at log strike k can always be written as

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

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. The assumption that $p_{i,k}$ has a linear asymptote is imposed by assuming that $f_i(k) \to 0$ as $k \to \pm\infty$. b_i^+ controls the difference between the slope as $k \to \infty$ and $k \to -\infty$.

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

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

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

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

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

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

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

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

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 $f_i(k) \to 0$ 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,

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

where $L(k) \equiv \frac{\exp(sk)}{\exp(sk)+1}$ is a scaled logistic function and s, σ_B^2, σ_D^2, $\sigma_{\hat{B}}^2$, and $\sigma_{\hat{D}}^2$ 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 trough 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: σ_a^2, σ_b^2, $\sigma_{b^+}^2$, $\sigma_{\hat{B}}^2$, and $\sigma_{\hat{D}}^2$. 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 (4) 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 in general appear to be 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 and only 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

The top panels of Figure A.3 plot the baseline versions of IS_{firm} and IS_{mkt} against an alternative using formula used by the CBOE. The results are essentially identical in both cases.

The middle-left panel compares the results using data from CME S&P 500 futures options to CBOE SPX options as reported by Optionmetrics. The differences are minor, implying that the exact data source does not make a difference.

The middle-right panel compares IS_{mkt} to the CBOE's reported SKEW index 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.

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 plotted in the bottom-left panel of Figure A.3.

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),

$$R3_{i,t} = \sum_{days \in t} 3\Delta p^E_{i,daily} \left(\exp\left(r_{i,daily}\right) - 1\right) + 6\left(r_{i,daily} \exp\left(r_{i,daily}\right) - 2\exp\left(r_{i,daily}\right) + r_{i,daily} + 2\right) \tag{A.11}$$

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. $\Delta p^E_{i,daily}$ 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.

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

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

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 (8) 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, $E_t\left[R3_{i,t+1}/RV^{3/2}_{i,t+1}\right]$, where RV is realized volatility,

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

That is,

$$IS_{i,t} \neq E_t \left[RS_{i,t+1} \right] \tag{A.14}$$

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

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

$$R3_{firm,t} = \beta_{3,firm}' X_{t-1} + \varepsilon_{3,firm,t} \tag{A.15}$$

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 the S&P 500 price/earnings ratio, the CBO output gap, the unemployment rate, the S&P 500 VIX, and the Gilchrist–Zakrajsek credit spread in all forecasting regressions. When forecasting the third moment, X_t includes the average option-implied third moment, and when forecasting the second moment it includes the option-implied second moment (in all cases, firm, market, and idiosyncratic, the option-implied moment matches the outcome being forecasted).

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

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

$$R2_{mkt,t} \equiv \sum_{days \in t} r^2_{mkt,daily} \tag{A.17}$$

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,

$$P3_{firm,t} \equiv \hat{\beta}'_{3,firm} X_t \tag{A.18}$$

$$P2_{firm,t} \equiv \hat{\beta}'_{2,firm} X_t \tag{A.19}$$

$$P3_{mkt,t} \equiv \hat{\beta}'_{3,mkt} X_t \tag{A.20}$$

$$P2_{mkt,t} \equiv \hat{\beta}'_{2,mkt} X_t \tag{A.21}$$

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

$$PS_{firm,t} \equiv \frac{P3_{firm,t}}{P2^{3/2}_{firm,t}} \tag{A.22}$$

$$PS_{mkt,t} \equiv \frac{P3_{mkt,t}}{P2^{3/2}_{mkt,t}} \tag{A.23}$$

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).

Figure A.4 plots the projected skewness series and shows that they are highly similar to the implied skewness series, especially after detrending. The reason is simply that in the forecasting regressions, the option-implied moments are dominant, so that the fitted value is very similar to the implied moments. Dew-Becker and Giglio (2021) report similar findings in the case of firm-level volatility – there is a risk premium, but its variation is not so large as to significantly change the behavior of the implied moments. that said, the trend in the implied and projected series is clearly different, implying that much of the long-term decline in skewness has been due to a change in the risk premium.

Table A.1 reports the correlations from Table 1c, but using projected instead of implied skewness. The results are highly similar to the benchmark results. Similarly, Figure A.8 shows that the cross-correlations are similar for IS and PS.

A.3 Detrending

A number of the variables in the behavior display unit root type behavior. In measuring economic activity, employment and industrial production have very strong unit roots, while the skewness indexes also display long-term downward trends. This section discusses how those trends are addressed.

The downward trend in skewness is evidence of a risk premium that increases over time if realized skewness does not also trend down. Looking at Figure 2, it is clear that realized skewness has a downward trend, but it is smaller than for implied skewness, implying that the risk premium has been trending larger. In Figure A.4, the version of implied skewness robust to risk premia has a downward trend (consistent with realized skewness). So accounting for a trend in skewness is important not only for eliminating time-varying risk premia, but also because there is a very persistent component in actual physical skewness.

Most of the results remove the trend. The main text focuses on a purely backward-looking trend. I use the local linear trend model (Harvey (2006)). Specifically, given a trending time series x_t, the filtered trend is $\bar{x}_t$, with

$$\bar{x}_t = \lambda_0 x_t + (1-\lambda_0)(x_{t-1} + b_{t-1}) \tag{A.24}$$

$$b_t = \lambda_1 (\bar{x}_t - \bar{x}_{t-1}) + (1-\lambda_1) b_{t-1} \tag{A.25}$$

and the detrended series is then $\hat{x}_t = x_t - \bar{x}_t$. The filter is motivated by a model in which both the level and growth rate of x_t have unit-root components.

The initial value of $\bar{x}_t$ is set to the average value in the first year of data, while the initial value of b_t is zero. I set λ_0 and λ_1 so that the frequency-domain gain of the filter matches that of the HP filter (with smoothing parameter set to 129,600) as closely as possible. That yields $\lambda_0 = 0.0592$ and $\lambda_1 = 0.0315$.

Figure A.7 reports the cross-correlations using the more typical HP filter to detrend instead of the local linear trend model, while Table A.2 replicates the correlations from

Table 1c using the HP filter.

Figure A.5 reports the cross-correlations after linearly detrending the skewness series. Finally, Figure A.6 reports results without any detrending, using the raw skewness indexes and growth rates for employment and IP.

A.4 Other robustness checks

A.4.1 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 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 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.4.2 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

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

$VIXA_t$ 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.4.3 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

$$
\begin{aligned}
spr_{i,j,t} \quad = \quad & b_0 + a_i + b_1 \log\left(\frac{strike_{i,j,t}}{spot_{i,t}}\right)_{i,j,t} + b_2 \log\left(\frac{strike_{i,j,t}}{spot_{i,t}}\right)_{i,j,t} \times 1\left\{\log\left(\frac{strike_{i,j,t}}{spot_{i,t}}\right)_{i,j,t} > 0\right\} \\
& + b_3 1\left\{\log\left(\frac{strike_{i,j,t}}{spot_{i,t}}\right)_{i,j,t} > 0\right\} + \eta_{i,j,t} \qquad \text{(A.26)}
\end{aligned}
$$

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.9 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, the variation appears to be unrelated to short-term variation in implied skewness, especially at high frequencies.

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.9 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.9 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 robust to risk premia**

	$PS_{firm,t}$	$PS_{firm,t}$ detrended	$PS_{mkt,t}$	$PS_{mkt,t}$ detrended	$PS_{idio,t}$	$PS_{idio,t}$ detrended
Industrial prod.	0.35 [0.13]	0.09 [0.10]	0.10 [0.12]	-0.13 [0.08]	0.43 [0.13]	0.14 [0.12]
Employment	0.28 [0.13]	-0.10 [0.09]	0.08 [0.16]	-0.18 [0.08]	0.33 [0.12]	-0.03 [0.12]
Unemployment	-0.19 [0.13]	0.02 [0.09]	-0.01 [0.14]	0.21 [0.08]	-0.29 [0.13]	-0.04 [0.12]
CBO output gap	-0.13 [0.14]	-0.12 [0.08]	-0.15 [0.11]	0.02 [0.10]	0.03 [0.13]	-0.05 [0.10]
NBER rec. ind.	-0.08 [0.13]	-0.18 [0.10]	0.02 [0.05]	0.15 [0.10]	-0.29 [0.16]	-0.25 [0.14]
Capacity util.	0.19 [0.14]	0.14 [0.10]	0.06 [0.12]	-0.18 [0.08]	0.25 [0.15]	0.21 [0.12]
Empl. growth	0.21 [0.08]	0.25 [0.08]	0.05 [0.07]	-0.14 [0.06]	0.34 [0.11]	0.27 [0.10]
IP growth	0.14 [0.07]	0.09 [0.07]	0.16 [0.09]	-0.15 [0.07]	0.43 [0.13]	0.14 [0.12]

Notes: PS is the measure of conditional skewness based on direct forecasts of conditional moments that is robust to time-varying risk premia. See section 4.5. This table replicates panel c of table 1 and is identical other than the change in the skewness measure.

Table A.2: **Correlations using HP filter for detrending**

	$IS_{firm,t}$	$IS_{firm,t}$ HP detr.	$IS_{mkt,t}$	$IS_{mkt,t}$ HP detr.	$IS_{idio,t}$	$IS_{idio,t}$ HP detr.
Industrial prod.	0.16 [0.11]	0.25 [0.09]	0.00 [0.10]	-0.08 [0.09]	0.25 [0.10]	0.24 [0.09]
Employment	0.17 [0.13]	0.25 [0.09]	-0.01 [0.09]	0.02 [0.10]	0.22 [0.11]	0.19 [0.10]
Unemployment	-0.20 [0.12]	-0.29 [0.09]	0.02 [0.09]	0.04 [0.09]	-0.28 [0.10]	-0.24 [0.10]
CBO output gap	0.08 [0.13]	0.27 [0.09]	-0.12 [0.10]	-0.12 [0.09]	0.30 [0.10]	0.29 [0.07]
NBER rec. ind.	-0.03 [0.11]	-0.19 [0.09]	0.02 [0.09]	0.19 [0.10]	-0.21 [0.12]	-0.18 [0.11]
Capacity util.	0.16 [0.12]	0.25 [0.09]	-0.04 [0.10]	-0.12 [0.08]	0.25 [0.12]	0.23 [0.10]
Empl. growth	0.14 [0.08]	0.19 [0.10]	-0.01 [0.05]	-0.18 [0.09]	0.25 [0.12]	0.28 [0.15]
IP growth	0.13 [0.08]	0.16 [0.11]	0.01 [0.03]	-0.12 [0.08]	0.21 [0.12]	0.22 [0.16]

Notes: This replicates table 1c from the main text, but using the HP filter instead of the purely backward-looking filter.

Figure A.1: Sample coverage

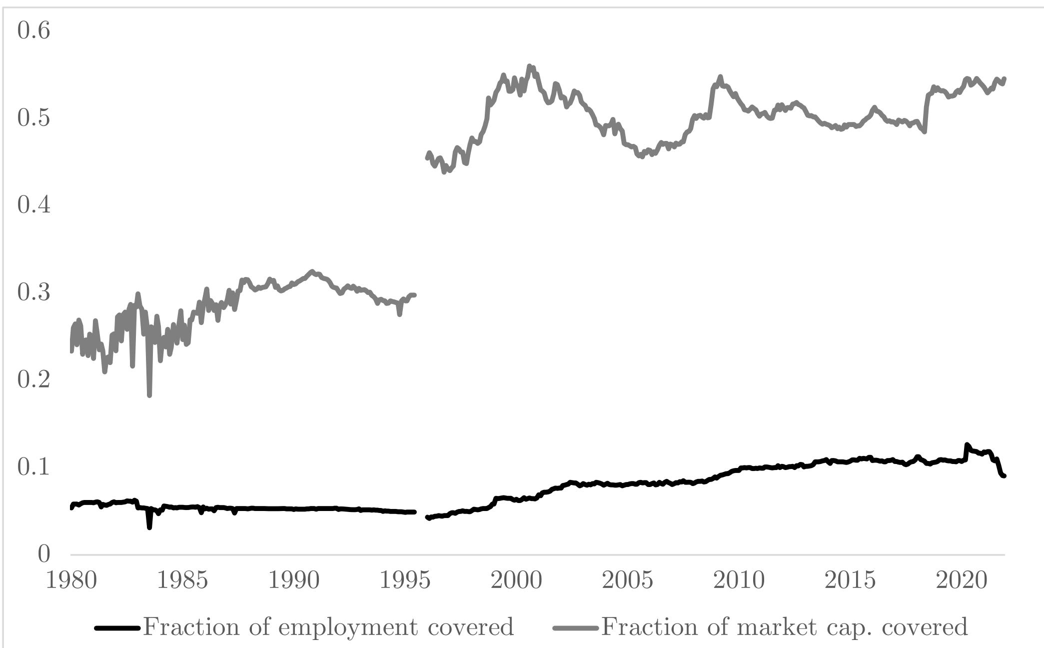

Note: The figure reports the fraction of CRSP total market capitalization accounted for the by firms used for calculating skewness, and the fraction of nonfarm payroll employment covered (from the BLS).

Figure A.2: Alternative skewness series

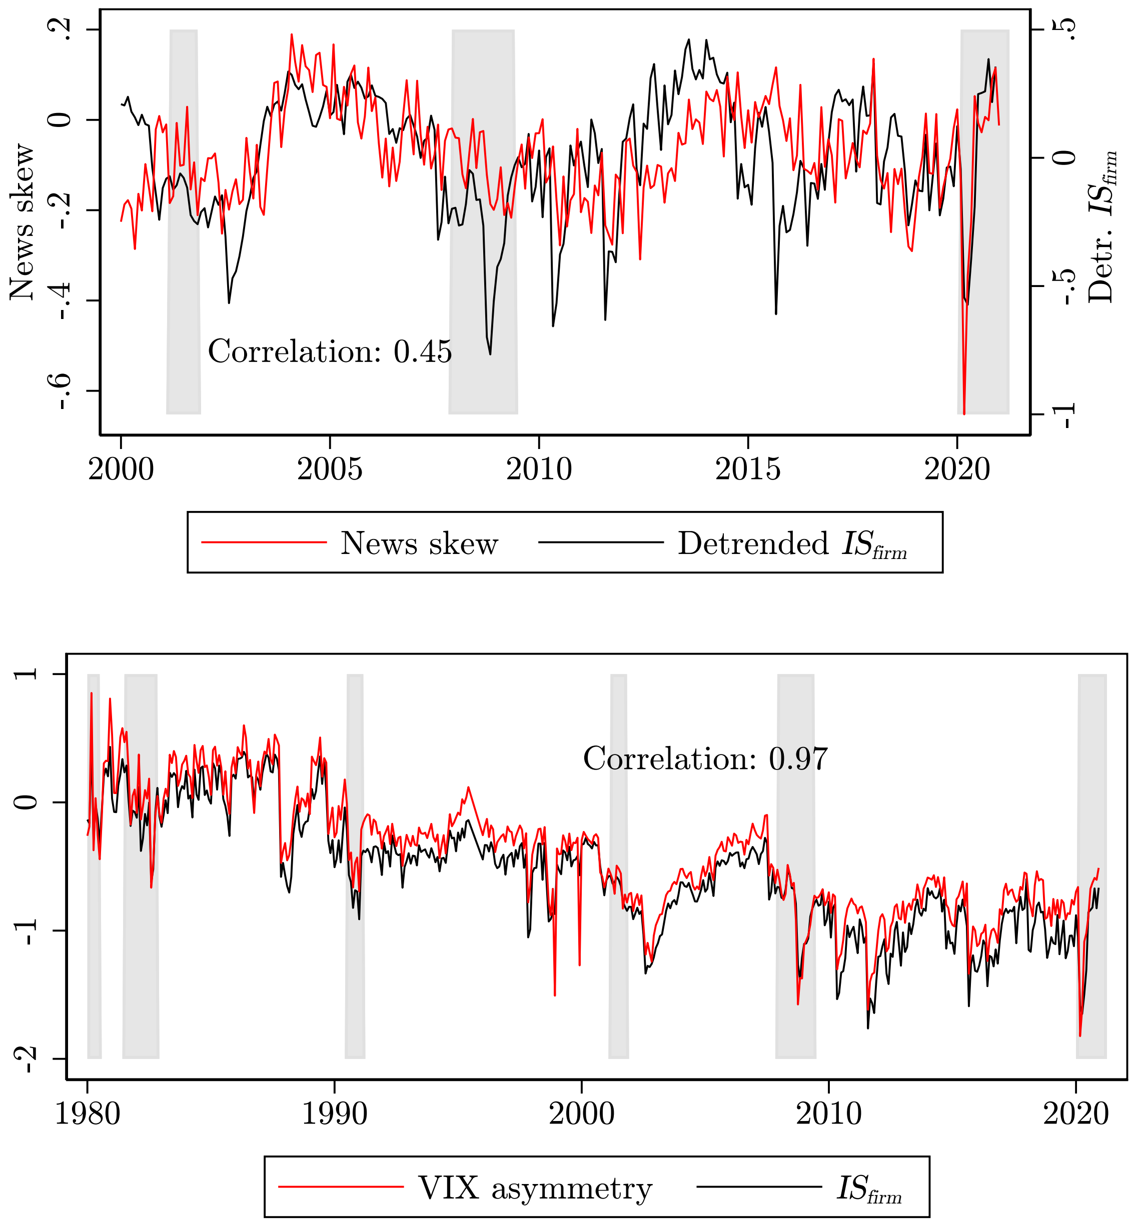

Note: The top panel replicates the top-left panel of figure 3, but with monthly instead of quarterly data. The bottom panel is the implied skew from VIX asymmetry (with both series normalized to have unit standard deviation). See sections A.4.1 and A.4.2, respectively.

Figure A.3: Skewness comparisons

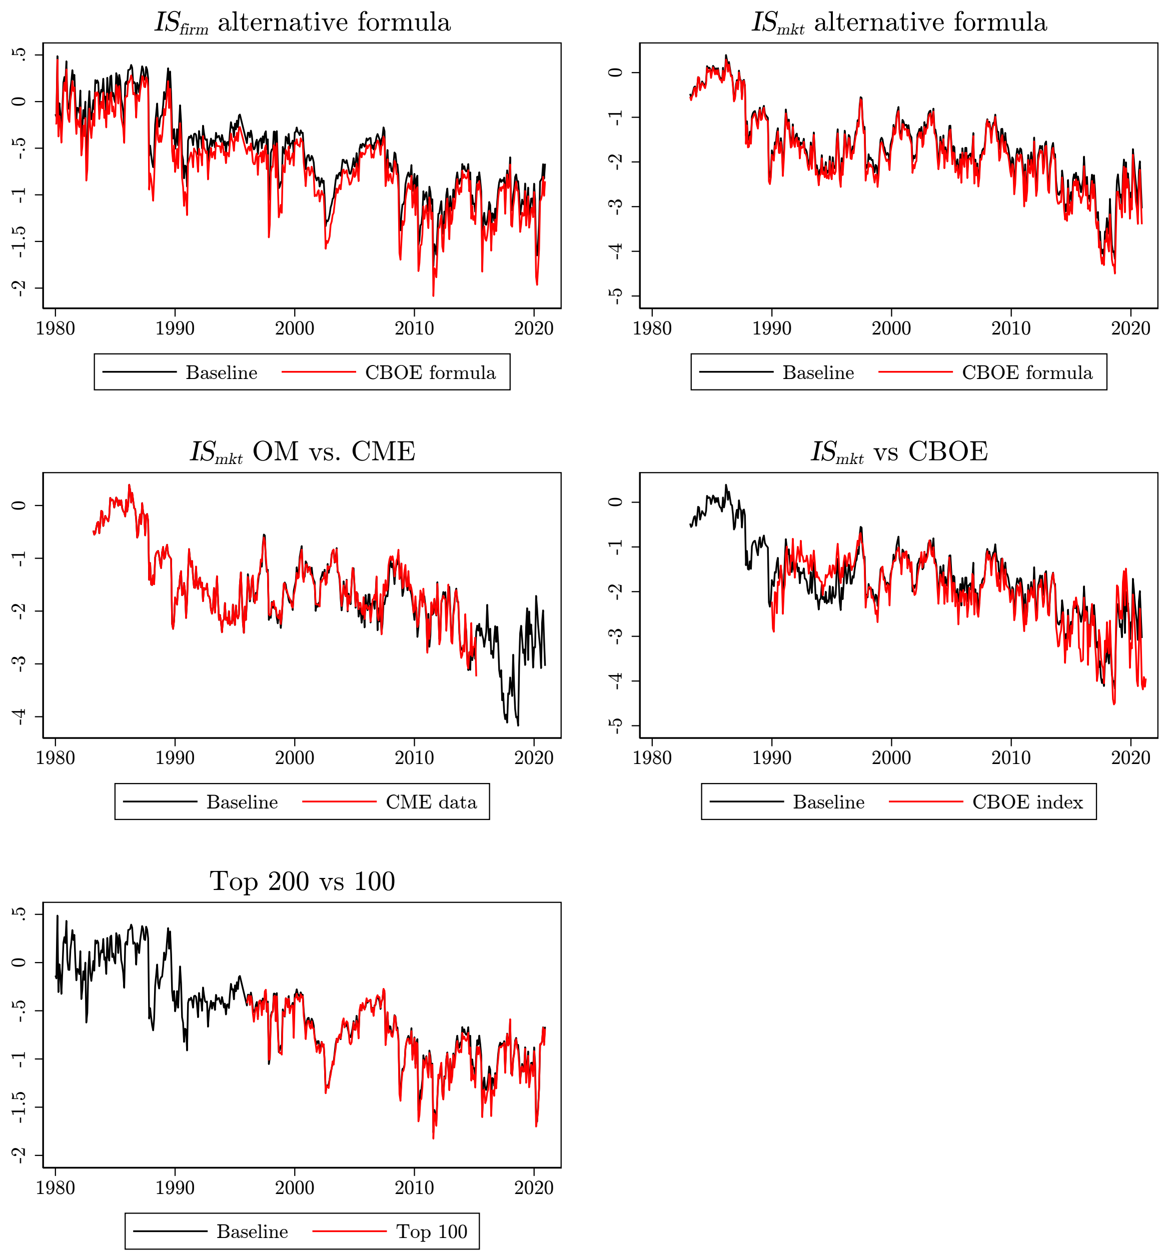

Note: Baseline skewness series and alternatives.

Figure A.4: Skewness robust to risk premia

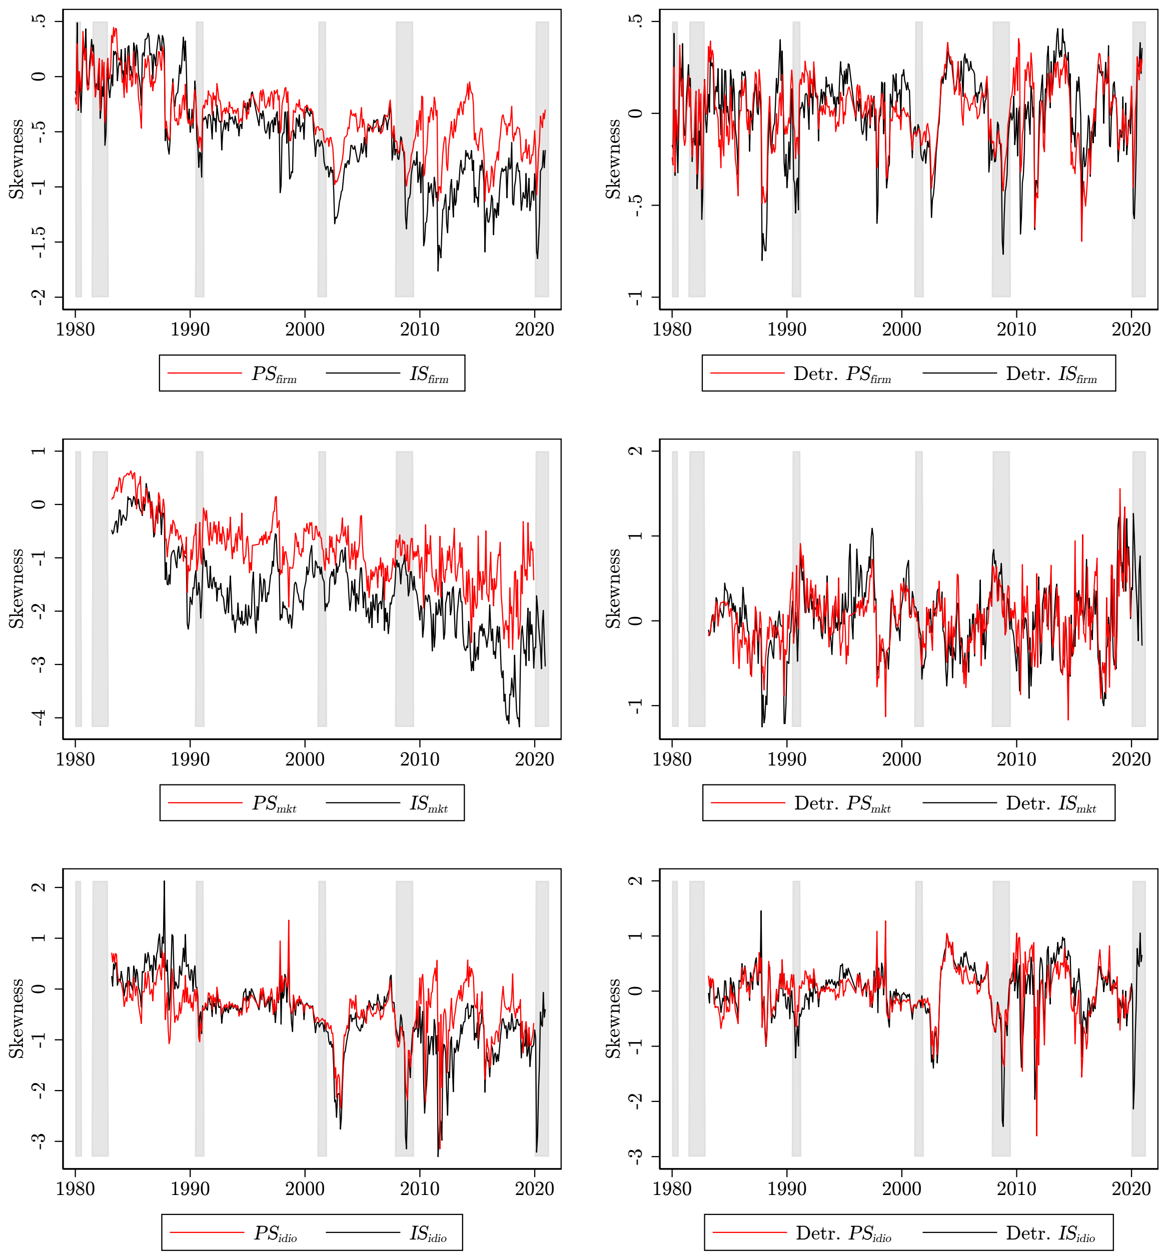

Note: Skewness constructed from forecasting regressions – PS, or projected (as opposed to implied) skewness – which is robust to the presence of time-varying risk premia. See section 4.5.

Figure A.5: Cross-correlations with linear detrending of implied skewness

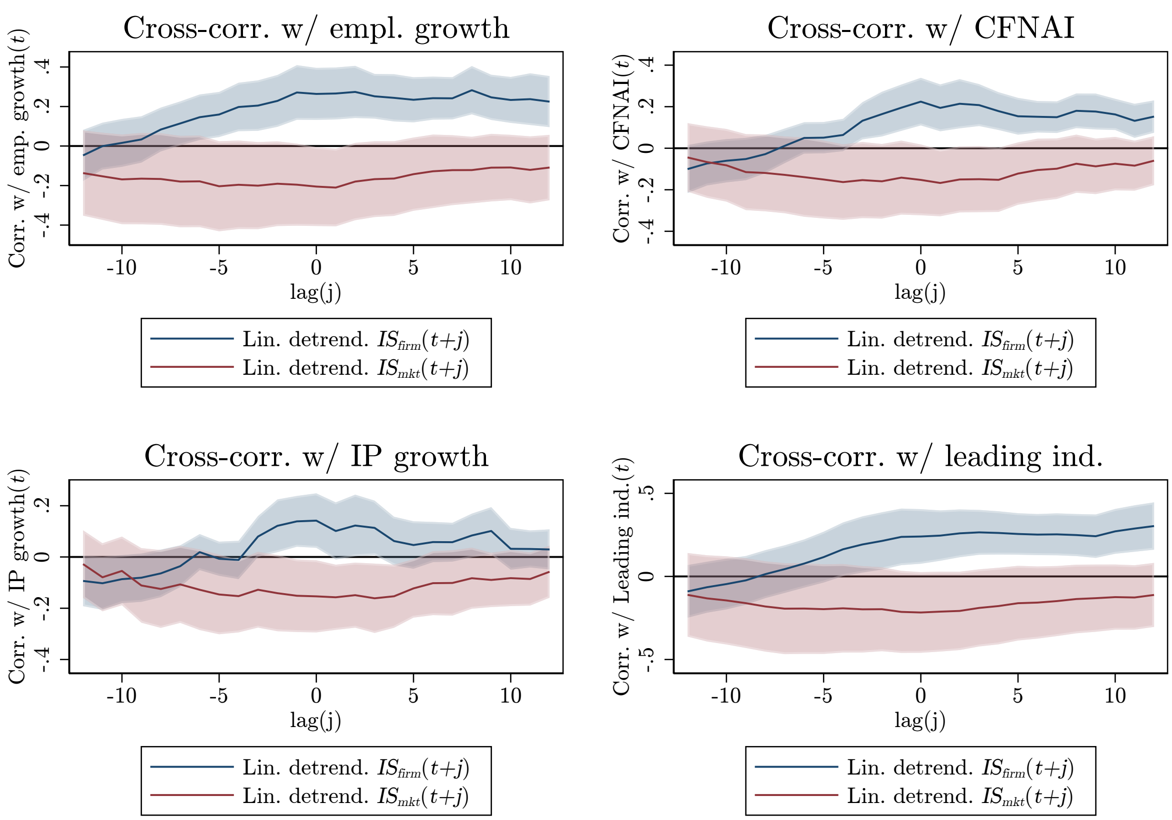

Note: This is an alternative version of the cross-correlations with two changes: the implied skewness series are linearly detrended instead of using the HP filter, and HP-filtered employment and IP are replaced with their growth rates.

Figure A.6: Cross-correlations without detrending

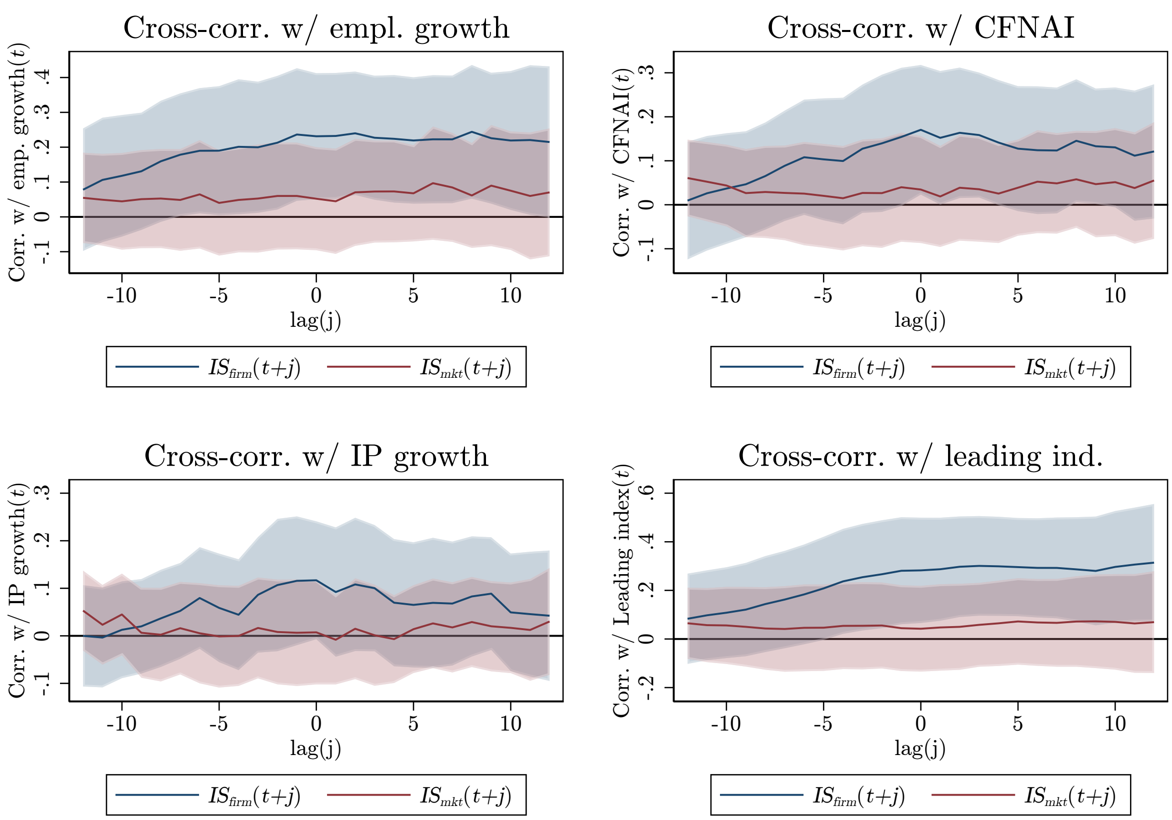

Note: Alternative version of the cross-correlations in which no series are detrended. Employment and IP are included as growth rates since their levels are nonstationary in the absence of detrending.

Figure A.7: Cross-correlations with HP filtering

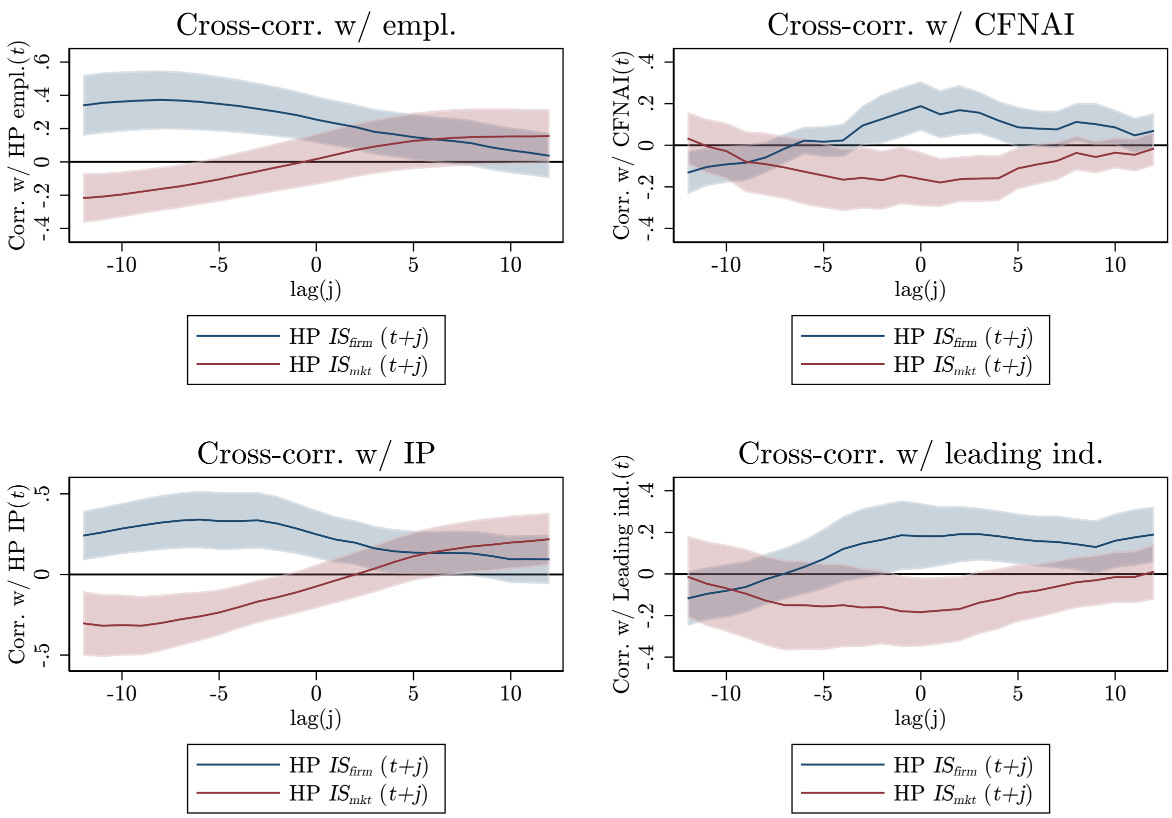

Note: Alternative version of the cross-correlations in which the HP filter with a coefficient of 129,600 is used for detrending.

Figure A.8: Cross-correlations for $\widehat{PS}$

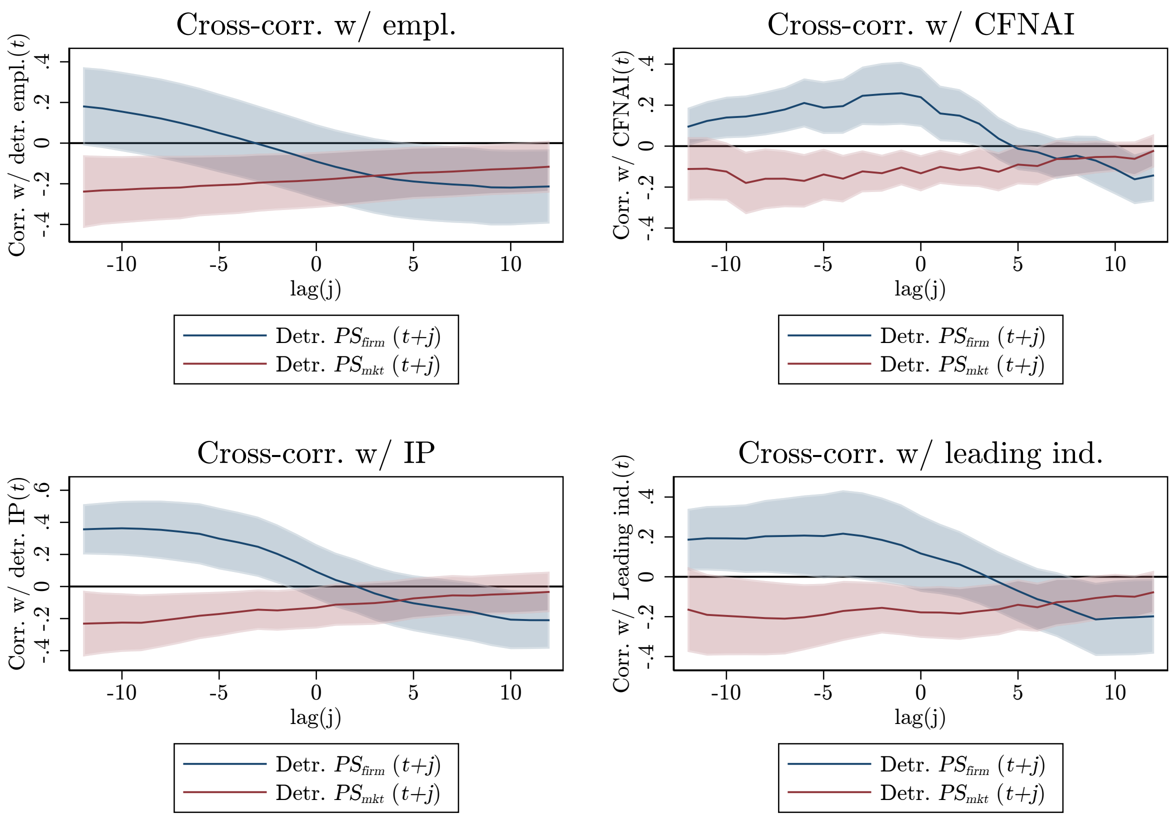

Note: Alternative version of the cross-correlations using $\widehat{PS}$ (the version robust to risk premia) instead of $\widehat{IS}$.

Figure A.9: Liquidity and skewness

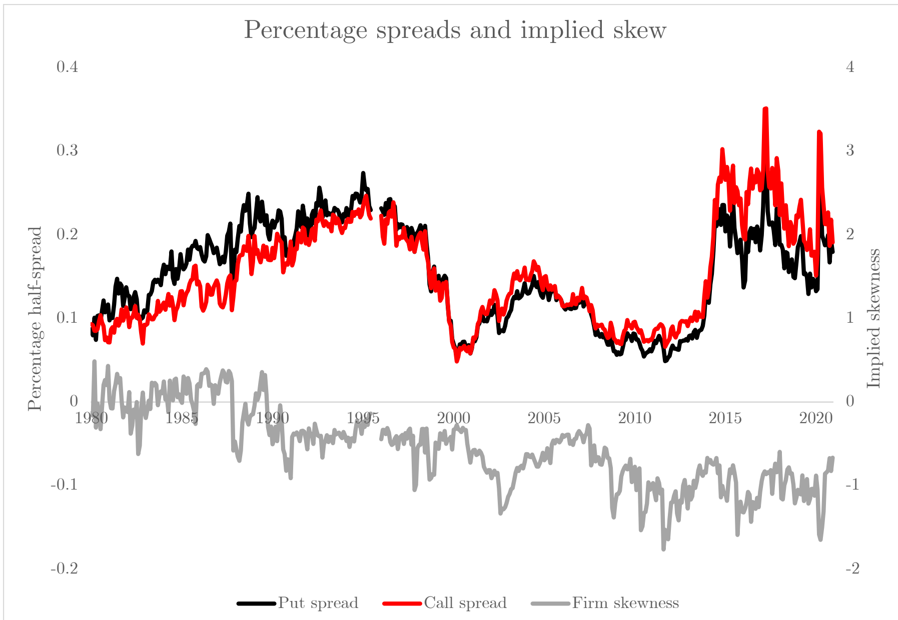

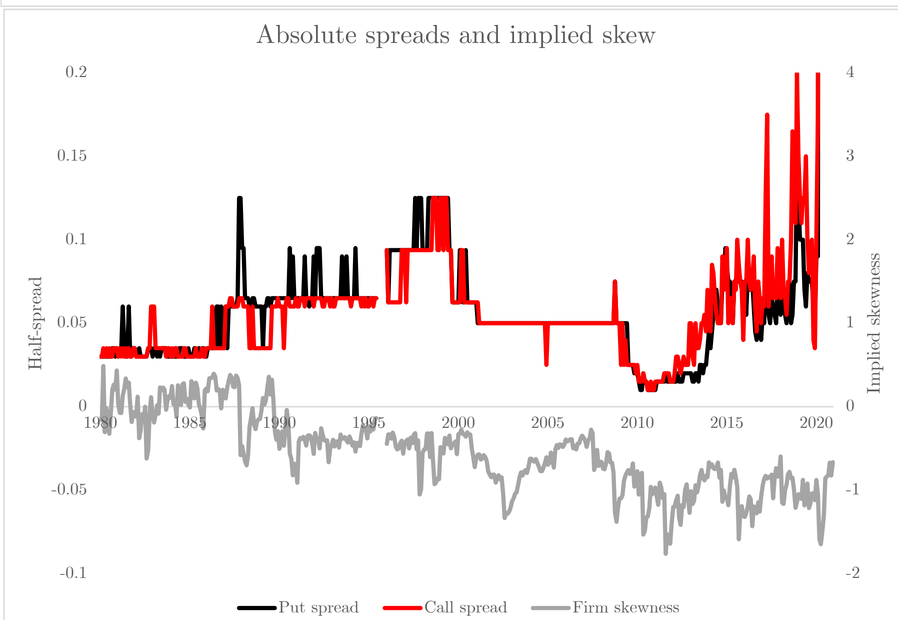

Note: In the top panel, the put spread is for a put option with strike/spot equal to 0.9, calculated as the fitted value from the regression A.26. The call spread is for a call option with strike/spot equal to 1.1. Both are scaled by the price of the option. The bottom panel reports the median raw bid/ask 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.