Cross-sectional uncertainty and the business cycle: evidence from 40 years of options data

Ian Dew-Becker and Stefano Giglio*

September 4, 2019

Abstract

This paper presents a novel and unique measure of cross-sectional uncertainty constructed from stock options on individual firms. We find that cross-sectional uncertainty varied little between 1980 and 1995, and subsequently had just two peaks – during the dot-com boom and the financial crisis. Overall, cross-sectional uncertainty is roughly acyclical and has little or no forecasting power for economic activity. Moreover, risk premia imply that investors view periods of high cross-sectional uncertainty as typically good states (as in the late 1990's). In general, cross-sectional uncertainty does not appear to represent a hindrance to economic growth.

1 Introduction

Background

The effects of uncertainty shocks are of interest to both policymakers and researchers. The most commonly studied real-time estimate of uncertainty about the economy is the VIX, which measures option-implied volatility for the S&P 500 index of large stocks. If one is interested in measuring uncertainty about the state of the aggregate economy – e.g. aggregate profits or output – the VIX is extremely valuable.

But in a wide range of models, uncertainty about cross-sectional or idiosyncratic shocks is the driving force. That could be because individual shocks cannot be perfectly insured or diversified (Bewley (1986)), because of nonlinearities in decision-making (Ilut, Kehrig,

*Dew-Becker: Northwestern University and the NBER. Giglio: Yale University. This paper would not have been possible without the version of the Berkeley Options Database preserved and shared by Stewart Mayhew. We thank Terry Hendershott and the Berkeley library for acquiring the data, converting it to a modern format, and making it available.

and Schneider (2018)), perhaps coming from fixed costs of adjustment (Bloom (2009) and Kaplan and Violante (2014)), or because of financial frictions, as in Christiano, Motto, and Rostagno (2016) and Gilchrist, Sim, and Zakrajsek (2014).

The contribution of this paper is to develop and study a novel and unique measure of cross-sectional uncertainty. It is analogous to the VIX in that it is also measured using option implied volatility, thus measuring real-time uncertainty from financial market data, but it measures uncertainty about cross-sectional shocks, rather than uncertainty about the level of the aggregate stock market. Whereas past work has measured the cross-sectional volatility of shocks that actually occur, this paper is unique in developing a real-time, forward-looking measure of cross-sectional uncertainty.

Methods

In many models, both theoretical and statistical, one can decompose the shock to a firm, $\eta_{i,t}$, into an aggregate component, μ_t, and a firm-specific component, $\varepsilon_{i,t}$,[1]

$$\eta_{i,t} = \mu_t + \varepsilon_{i,t} \tag{1}$$

The VIX and other measures of uncertainty about the state of the aggregate economy capture the conditional (time-t) variance of μ_{t+1}. Measures of total firm- or household-level uncertainty like the Michigan survey of consumers (used by Leduc and Liu (2015)) or the survey studied by Bachmann et al. (2018) apply to the total conditional variance of $\eta_{i,t+1}$. Measures of cross-sectional volatility, such as Campbell et al. (2001), measure the realized (or sample) cross-sectional standard deviation of $\varepsilon_{i,t}$. There are no studies that provide a measure of the conditional volatility on date t of the firm-specific component, $\varepsilon_{i,t+1}$. Part of our contribution is to fill that gap, and our paper is unique in giving a direct and forward-looking measure of that component of uncertainty.

Measuring and studying the uncertainty about the firm-specific or cross-sectional component, $\varepsilon_{i,t+1}$, is useful not just because it allows us to test models driven by that type of uncertainty, but also because it helps understand what drives variation in the total uncertainty faced by firms. Is it primarily due to aggregate factors, like uncertainty about government policy or some kind of technology shock, or about more idiosyncratic factors or policies that have disparate effects across the economy? And which is more relevant for aggregate activity?

[1]For example, Bloom (2009), Christiano, Motto, and Rostagno (2016), Bloom et al. (2018), and Ilut, Kehrig, and Schneider (2018) all use such a specification.

In general, we allow for cross-sectional correlation in $\varepsilon_{i,t}$. By subtracting market returns, the method implies that for any common factors in $\varepsilon_{i,t}$, the loadings in the cross-section sum to zero, so that the factors are purely cross-sectional, having no aggregate effects.

Our firm-specific implied volatility measure is simple to construct: it is just average firm-level option-implied volatility minus market implied volatility. Under general conditions, that gap measures the average variance of the residual from a regression of each stock's return on that of the market. It is an option-implied version of the volatility studied by Campbell et al. (2001), Gilchrist, Sim, and Zakrajsek (2014), and Herskovic et al. (2016).

The past work on idiosyncratic or cross-sectional volatility has exclusively studied the *realized* distribution of outcomes – i.e. the cross-sectional distribution of realized stock returns or sales growth.[2] We show that there are important differences in the behavior of realized and expected cross-sectional volatility.

Results

We obtain three key results about the behavior of cross-sectional uncertainty: it is acyclical, it does not forecast declines in real activity, and investors view periods of high cross-sectional uncertainty as good times.

The first part of the analysis examines the basic time series of cross-sectional uncertainty along with its raw correlations with other aggregate variables. Figure 1 plots cross-sectional uncertainty. For the start of our data, in 1980, up to 1995, there was surprisingly little variation. After 1995, firm-level uncertainty moves much more (though still less than market uncertainty), but with just two distinct increases, one during the dot-com boom and the other during the financial crisis.

Cross-sectional uncertainty over this period is also essentially acyclical. While it was high during the financial crisis, it was actually falling during the 2000-01 recession, and it shows no significant increase in any of the previous three recessions. Furthermore, in the two episodes where uncertainty is elevated, it rapidly declines, returning to its long-run average by the trough of the recession. So without very strong internal propagation, uncertainty cannot explain the length of time output stays below trend (see also Born, Breuer, and Elstner (2018)). If output tracked cross-sectional uncertainty over time, it would have recovered from the financial crisis by 2010.

The lack of cyclicality of cross-sectional uncertainty is surprising, since it is usually taken as a given in the literature that uncertainty is countercyclical. We show that cross-sectional *realized* volatility is more cyclical than uncertainty. That shows why using a forward-looking measure of uncertainty, rather than a backward-looking measure that quantifies how big were the shocks that actually occurred, is critical for understanding the cyclicality of uncertainty. The empirical acyclicality is inconsistent with many models, both those in which uncertainty is an exogenous driving force (Christiano, Motto, and Rostagno (2014) and Bloom et al.

[2]E.g. e.g. Campbell et al. (2001), Gilchrist, Sim, and Zakrajsek (2014), Herskovic et al. (2016), Bloom et al. (2018), Guvenen, Ozkan, and Song (2014), Bachmann, Elstner, and Sims (2013).

(2018)), and also those in which it is endogenously countercyclical, such as Bachmann and Moscarini (2012). We also show in the paper that these results are not unique to the US: they also hold in a sample of other countries for which we obtain analogous option prices: Switzerland, Germany, Spain, France, Great Britain, and the Netherlands.

We next examine the utility of idiosyncratic uncertainty for forecasting. Policymakers frequently cite uncertainty as a potential factor hindering growth. That might suggest that high uncertainty would forecast low future growth. At best, in univariate forecasting regressions idiosyncratic uncertainty has very weak forecasting power for aggregate output and employment. That result holds only for the second half of the sample, since 1997, though. In the first half of the sample, from 1980 to 1996, if anything high uncertainty forecasts *high* rather than low activity.

Cross-sectional uncertainty is also driven out in forecasting regressions by cross-sectional realized volatility, aggregate uncertainty, and, most strongly, aggregate realized volatility. Overall, then, we conclude that there is little evidence that cross-sectional uncertainty is relevant for forecasting or in general represents a major "headwind" that should be considered by policymakers.

As noted above, the acyclicality of cross-sectional uncertainty is inconsistent with the intuition behind some recent models of uncertainty shocks. We formally examine the cyclicality and forecasting regressions in two recent theoretical models: Christiano, Motto, and Rostagno (2014) and Bloom et al. (2018). In both cases, we find that the models predict that cross-sectional uncertainty should be countercyclical and should be more tightly related to aggregate output than aggregate uncertainty. Both of those predictions of the model are inconsistent with the data.

An alternative perspective from which to examine the relevance of idiosyncratic uncertainty shocks, advocated by Dew-Becker, Giglio, and Kelly (2019), is to ask how they are priced by investors. If idiosyncratic uncertainty shocks on average occur in bad states of the world – i.e. if they harm the economy and raise marginal utility – then theory predicts that they will carry a negative risk premium.

Our third contribution is therefore to measure the risk premium associated with idiosyncratic uncertainty shocks. Option portfolios have the advantage over other alternatives, such as portfolios of equities, that they give direct exposure to shocks to uncertainty. We show that shocks to uncertainty, both aggregate and idiosyncratic, carry *positive* risk premia. That is, rather than uncertainty being high in bad times, the risk premia suggest that investors view periods with positive uncertainty shocks as typically being good states of the world. That is consistent with the fact that the largest increase in cross-sectional uncertainty comes during the dot-com boom in the late 1990's.

The main takeaway from our analysis is that cross-sectional uncertainty – at least for large firms and as measured by option-implied volatility – is not in general something to fear: it is acyclical and it does not forecast future declines in activity. And in fact, investors do not appear to fear it.

In theory, there can be both contractionary and expansionary effects of cross-sectional variance. Such risk may be contractionary in the face of nonconvexities in adjustment costs. On the other hand, endogenous adjustment can make more volatile shocks beneficial. For example, if production can be shifted to the most productive firms, then the economy benefits from cross-sectional volatility, since aggregate output will depend roughly on the maximum of the cross-section of shocks. The so-called Oi–Hartmann–Abel effect is a similar expansionary mechanism, through which firms benefit from uncertainty because they can respond endogenously to it. Furthermore, in times of technological change, one might expect that there would be winners and losers, but that investors would not necessarily know which firms would benefit. That could cause cross-sectional uncertainty to be high as a result of good news about the overall economy (see also Pastor and Veronesi (2009)).

Related work

In addition to the literature discussed above, our work is also related to a literature in finance that examines the pricing of volatility at the firm level.[3] Most closely related in terms of asset pricing methods are Bakshi and Kapadia (2003) and Carr and Wu (2009), who both provide evidence that the idiosyncratic component of volatility is unpriced. Our analysis goes a step further by specifically distinguishing the pricing of *uncertainty*, which is typically what is most relevant to theoretical models.

The remainder of the paper is organized as follows. Section 2 describes the data and the method we use to construct firm-level uncertainty. Section 3 then examines its time-series properties and cyclicality. Section 4 studies forecasting regressions, while section 5 examines the ability of two leading theoretical models to match the empirical results. We then examine asset returns section 6 and international evidence in section 7. Finally, section 8 concludes.

2 Data

Much of the literature on uncertainty studies the VIX, which is an option-implied volatility for the S&P 500 index. It therefore measures uncertainty about aggregate stock prices. To measure firm-level uncertainty, we measure option-implied volatility for individual stocks.

[3]See Driessen, Maenhout, and Vilkov (2009), Goyal and Saretto (2009), Chen and Petkova (2012), Buraschi, Trojani, and Vedolin (2014), and Herskovic et al (2016).

We obtain data from two sources. For the period 1996–2017, the source is Optionmetrics (in particular, the implied volatility surface file). For the period 1980–1995, we use data from the Berkeley Options Database. That data has been used occasionally in the finance literature, but this is the first paper to use it to develop a time series of firm-level uncertainty back to 1980.[4]

The appendix describes the details of the construction of the implied volatilities. Whereas the VIX is measured using a so-called model-free implied volatility, here we just use at-the-money Black–Scholes implied volatility (the original definition of the VIX; the two series are 99.5 percent correlated). The model-free implied volatility requires a large (ideally infinite) number of strikes, which is often not available for individual stocks, especially in the early part of the sample. The Black–Scholes model's key assumption is that the underlying stock price follows a pure diffusion with constant volatility, which is not literally true empirically. Furthermore, since implied volatilities come from asset prices, they embed risk premia.[5] Option-implied volatilities are therefore not errorless measures of investor beliefs.

Denote firm i's implied volatility in month t as $\sigma_{i,t}$. Similarly, denote implied volatility for the aggregate stock market with $\sigma_{mkt,t}$. As long as all returns have finite first and second moments, one can theoretically construct the projection of the return on stock i, $r_{i,t}$, on the market, $r_{mkt,t}$, as

$$r_{i,t} = \alpha_{i,t} + \beta_{i,t-1} r_{mkt,t} + \varepsilon_{i,t} \tag{2}$$

where $\varepsilon_{i,t}$ and $r_{mkt,t}$ are conditionally uncorrelated by construction. Note that this is just a theoretical representation – it is not directly estimable since the parameters can in principle change on every date. Taking conditional variances of both sides, we have

$$\sigma^2_{i,t-1} = \beta^2_{i,t-1} \sigma^2_{mkt,t-1} + \sigma^2_{\varepsilon,i,t-1} \tag{3}$$

where $\sigma^2_{\varepsilon,i,t-1} = \text{var}_{t-1}(\varepsilon_{i,t})$, the variance conditional on information available on date $t-1$. Now consider a market value weighted average of firm-level implied variance, $\sigma^2_{i,t}$, minus

[4] While the Princeton University Library had the data from 1987–1995, the earlier part of the sample was nearly completely lost. Luckily, Stewart Mayhew, who as a PhD student was the last manager of the database, had kept tapes – carrying them in moves between various jobs for more than 20 years. It is entirely due to his foresight in keeping the data, and the generosity of Terry Hendershott and the UC Berkeley Library, that we are able to write this paper.

[5] Bekaert, Hoerova, and Lo Duca (2013)**?** show that risk premia alone cannot account for the observed variation in implied volatilities.

market implied variance, $\sigma^2_{mkt,t}$:

$$\sigma^2_{\varepsilon,t} \equiv \sum_i w_{i,t}\sigma^2_{i,t} - \sigma^2_{mkt,t} \tag{4}$$

$$= \underbrace{\left(\sum_i w_{i,t}\beta^2_{i,t} - 1\right)}_{=\text{var}_w(\beta_{i,t})}\sigma^2_{mkt,t} + \sum_i w_{i,t}\sigma^2_{\varepsilon,i,t} \tag{5}$$

(where var_w denotes a variance weighted by market value). So the value-weighted mean of firm-level implied variance minus market variance is equal to the average of cross-sectional variance, $\sum_i w_{i,t}\sigma^2_{i,t}$, plus a second-order error that depends on the variance of the betas across firms. In general we would expect that term to be quantitatively small, and the appendix shows that in the data, $\text{var}_{w,t}(\beta_{i,t}) \approx 0.08$.

We therefore measure cross-sectional implied volatility as $\sigma^2_{\varepsilon,t} \equiv \sum_i w_{i,t}\sigma^2_{i,t} - \sigma^2_{mkt,t}$.[6] We focus on the definition of $\sigma^2_{\varepsilon,t}$ in (4) for the same reason Campbell et al. (2012) do: it is simple, transparent, and does not require estimation of the $\beta_{i,t}$. The construction of $\sigma^2_{\varepsilon,t}$ relies only on the assumptions that the return projection (6) exists, that option prices can yield a measure of implied volatility, and that $\text{var}_{w,t}(\beta_{i,t})$ is sufficiently small that the approximation error $\left(\sum_i w_{i,t}\beta^2_{i,t} - 1\right)\sigma^2_{mkt,t}$ is quantitatively small. The method here is similar to that of Campbell et al. (2001), who also weight by market capitalization and avoid estimating betas (implicitly setting them to 1).

We refer to $\sigma^2_{\varepsilon,t}$ variously as firm-specific, cross-sectional, or idiosyncratic uncertainty. It is idiosyncratic just in the sense that it is the part of total average firm uncertainty that comes from variation uncorrelated with the overall market. That is, it is total firm uncertainty minus the common component. Of course the $\varepsilon_{i,t}$ are correlated across firms, e.g. due to industry or location effects. Changes in the volatilities of those cross-sectional factors will appear in $\sigma^2_{\varepsilon,t}$. In that sense, the name "cross-sectional uncertainty" is most appropriate.[7]

We measure $\sigma^2_{mkt,t}$ as the at-the-money implied volatility for options on S&P 500 futures traded on the CME up to 1986, and then for CBOE S&P 500 options subsequently (when

[6] Note again that this relationship holds under both the physical and risk-neutral measures. That is, if we have the variances of individual firm and market returns under the risk-neutral measure, then $\sigma^2_{\varepsilon,t}$ is risk-neutral idiosyncratic variance.

[7] More formally, suppose the residuals can be decomposed as $\varepsilon_{i,t} = \sum_j \gamma_{i,j,t} f_{j,t} + \tilde{\varepsilon}_{i,t}$ for a set of factors $f_{j,t}$ and loadings $\gamma_{i,j,t}$, where $\tilde{\varepsilon}_{i,t}$ is uncorrelated across firms. Since $\varepsilon_{i,t}$ sums to zero across firms, the loadings $\gamma_{i,j,t}$ must also sum to zero (when weighting by market capitalization, $w_{i,t}$). The firm-specific variance is then $\sigma^2_{\varepsilon,t} = \text{var}_w(\gamma_{i,j,t})\,\text{var}(f_{j,t}) + \sum_i w_{i,t}\,\text{var}(\tilde{\varepsilon}_{i,t})$. It depends on the dispersion in loadings on the cross-sectional factors, the volatilities of the cross-sectional factors, and the volatility of the pure residual.

they overlap, the two series are essentially identical for our purposes).[8]

Throughout the analysis, we measure implied volatility interpolated to a maturity of thirty days, using the two nearest available maturities.

Figure A.1 in the appendix reports the fraction of total CRSP market capitalization for which we have implied volatilities in each month. For the period covered by the BODB, we have about one third, due to both the fact that not all firms had traded options and that only about half were listed on the CBOE (as opposed to other exchanges, like the NYSE or AMEX). Optionmetrics has data from more exchanges than just the CBOE. In 1996, when it becomes available, coverage jumps to 63 percent and then rises to 98 percent by the end of the sample.

3 Time-series behavior of cross-sectional uncertainty

3.1 Univariate behavior and cyclicality

3.1.1 Variability

The top panel of figure 1 plots the time series of cross-sectional uncertainty, $\sigma_{\varepsilon,t}$. The thin line plots realized cross-sectional realized volatility, constructed in the same way as $\sigma_{\varepsilon,t}$, but using squared returns instead of implied volatilities (see equation (??)). Gray bars indicate NBER-dated recessions. The implied volatility line has a gap between July and December, 1995, because the Berkeley options database ends in June, 1995, and Optionmetrics begins in January, 1996.

Outside the dot-com boom and the financial crisis, there is surprisingly little variation in cross-sectional volatility. In those two episodes, cross-sectional volatility rises above 60 percent per year, while in the rest of the sample its 5th and 95th percentiles are 19 and 26 percent. The overall standard deviation of cross-sectional volatility is 6.6 percentage points, but in the first half of the sample, through the end of 1997, it is smaller by a factor of three at only 2.1 percent.

The bottom panel of figure 1 compares cross-sectional uncertainty to implied volatility for the overall stock market (on a log scale). Aggregate uncertainty is substantially more variable than cross-sectional uncertainty. Its standard deviation is 38 percent of its mean,

[8] The S&P 500 options are not available until January, 1983. To construct firm-specific volatility in the earlier period, we impute the missing values for $\sigma^2_{mkt,t}$ by regressing it on $\sum_i w_{i,t}\sigma^2_{i,t}$, realized volatility on the aggregate stock market, the S&P 500 price/earnings ratio, and the Gilchrist–Zakrajsek excess bond premium (which yields an R^2 of 86 percent). We use the imputed values to inspect the basic behavior of uncertainty, but not for the forecasting regressions.

compared to only 27 percent for cross-sectional volatility (in absolute terms, the standard deviations are both 6.5 percent). The variability is also much less isolated. Whereas cross-sectional uncertainty has two major peaks, there are numerous clear peaks in market-level volatility, associated with the 1987 stock market crash, the first Gulf War, various events between 1998 and 2002, the 2008 financial crisis, and the debt ceiling and Euro crisis in 2010 and 2011.

The relative volatilities of $\sigma^2_{\varepsilon,t}$ and $\sigma^2_{mkt,t}$ can be used to construct a variance decomposition for the total variance faced by firms. Specifically,

$$\begin{array}{rcccccc} & \operatorname{var}\left(\sum_i w_{i,t}\sigma^2_{i,t}\right) = & \operatorname{var}\left(\sigma^2_{mkt,t}\right) & + & \operatorname{var}\left(\sigma^2_{\varepsilon,t}\right) & + & 2\operatorname{cov}\left(\sigma^2_{i,t},\sigma^2_{mkt,t}\right) \\ \text{Full sample:} & 4.31\times10^{-3} = & \underset{24\%}{1.02\times10^{-3}} & + & \underset{43\%}{1.86\times10^{-3}} & + & \underset{33\%}{1.43\times10^{-3}} \\ \begin{array}{r}\text{Outside '99-'02}\\ \text{and '08-'09}\end{array}: & 5.95\times10^{-4} = & \underset{61\%}{3.61\times10^{-4}} & + & \underset{15\%}{0.92\times10^{-4}} & + & \underset{24\%}{1.42\times10^{-4}} \end{array}$$

Over the full sample, 43 percent of the variation in total firm uncertainty comes from $\operatorname{var}\left(\sigma^2_{\varepsilon,t}\right)$ alone, not even counting the covariance term. When the financial crisis and dot-com boom are excluded, though, the variation in cross-sectional uncertainty falls by more than an order of magnitude, and it accounts for only 15 percent of the total variation in firm uncertainty (adding in half the covariance term raises the fraction to 27 percent). So for much of the sample, including the first three recessions, the majority of the variation in the total uncertainty firms face came from variation in uncertainty about the common component of returns.

The relatively low degree of variation in cross-sectional uncertainty is inconsistent with leading theoretical models of cross-sectional uncertainty. The standard deviation of cross-sectional uncertainty divided by its mean is 0.27 in the data. In the calibration of Bloom et al. (2018), it is 0.71, while in the estimation of Christiano, Motto, and Rostagno (2016) it is 0.58. So both models have a volatility for uncertainty that is more than twice as large as what we observe empirically. That gives another perspective from which the stability of cross-sectional implied volatility is surprising.

3.1.2 How much of the variation is common?

For most of the analysis, we follow the literature in studying the common component in cross-sectional uncertainty. It is worth asking, though, how much of the variation in firm-level uncertainty is driven by that common component. To do so, we use the law of total variance,

$$\underbrace{\mathrm{var}\left(x_{i,t}\right)}_{\text{Total variance}} = \underbrace{E\left[\underset{t}{\mathrm{var}}\left(x_{i,t}\right)\right]}_{\text{Average cross-sectional variance}} + \underbrace{\mathrm{var}\left[E_t\left(x_{i,t}\right)\right]}_{\text{Time-series variance of the average}} \quad (6)$$

where var_t and E_t refer to the cross-sectional variance and average on date t. The first term represents the residual variance after accounting for the cross-sectional average in each period, while the second term is the variance coming from that average. So the ratio of $\mathrm{var}\left[E_t\left(x_{i,t}\right)\right]$ to the total variance represents the fraction of the total variance explained by the cross-sectional mean in each period.

The variance decomposition identity (6) also holds with weights, so we weight by market capitalization as above (normalizing the sum of market capitalization to 1 on each date to give them equal weight overall). For $x_{i,t}$, we use total firm implied volatility measured here as

$$\sigma^2_{R,\varepsilon,i,t} \equiv \sigma^2_{i,t} - \beta^2_{i,R,t}\sigma^2_{mkt,t}$$

where $\sigma^2_{R,\varepsilon,i,t}$ is a rolling beta estimated using the previous 12 months of daily data. When we are just calculating the average of cross-sectional uncertainty across firms, the errors from setting $\beta_i \approx 1$ somewhat cancel out across firms. Here, though, those errors will affect the variance decomposition, so it is important to also examine what happens when we actually estimate β_i.

Variance decomposition for uncertainty measures

	Fraction from time-series component	
	Firm level	Sector level
Total firm uncertainty ($\sigma^2_{i,t}$)	0.50	0.75
Firm-specific uncertainty ($\sigma^2_{R,\varepsilon,i,t}$)	0.40	0.70

Depending on the measure, between 40 and 50 percent of the total variation in uncertainty is due to a common component (measured as the cross-sectional average), which is quantitatively consistent with the results in Herskovic et al. (2016). The fraction explained by a common component is greater for total firm uncertainty, which is natural since that includes market uncertainty, which affects all firms. The second column of the table above reports similar results for measures of uncertainty averaged within two-digit sectors (i.e. $\sum_{i\in S} w_{i,t}\sigma^2_{R,\varepsilon,i,t}$, where S represents the set of firms in some sector). These results therefore measure the extent to which cross-sectional uncertainty is similar across different sectors. In this case, the fraction of the variation explained by the time-series component is 1.5 times larger than in the firm-level case.

Overall, then, a surprisingly large amount of the total variation in cross-sectional or idiosyncratic risk is driven by a common component that hits all parts of the economy, which motivates us (and previous authors) to study a single common factor.

3.1.3 Cyclicality

The bottom panel of figure 1 adds further context by plotting cross-sectional uncertainty against market-level uncertainty and the detrended level of the NASDAQ index. The two periods of high cross-sectional uncertainty are both associated with large changes in stock prices, but in opposite directions. During the dot-com boom, cross-sectional uncertainty almost perfectly tracks the level of the NASDAQ. They peak in almost exactly the same month, and cross-sectional volatility declines in unison with the level of the NASDAQ. In other words, uncertainty is high in the late 1990's when stock prices are rising, and it falls as conditions deteriorate. It follows the opposite pattern during the financial crisis: it is exactly when the stock market index begins to decline that cross-sectional uncertainty begins to rise. So uncertainty appears to be procyclical in the late 1990s and early 2000s, countercyclical in the financial crisis, and acyclical otherwise.

Figure 2 further emphasizes that point by plotting cross-sectional uncertainty against measures of aggregate activity. The top panel plots the unemployment rate. $\sigma^2_{\varepsilon,t}$ peaks in 2000 at the same time that the unemployment rate reaches its trough. As unemployment subsequently rises, uncertainty falls. Conversely, unemployment and uncertainty rise simultaneously in 2007. The bottom panel shows similar behavior for private nonresidential fixed investment.

To more formally quantify the cyclicality of cross-sectional uncertainty, table 1 reports the correlation of cross-sectional uncertainty from figure 1 with various measures of the state of the economy. The signs of the variables are all set so that positive values indicate high activity (e.g. the sign of the unemployment rate is reversed). In terms of levels, cross-sectional uncertainty does not have a consistent correlation with economic indicators. It is positively correlated with the CBO output gap, detrended employment and minus the unemployment rate, implying it is procyclical. Its correlation with detrended industrial production is close to zero, and its correlation with the Gilchrist–Zakrajsek bond spread and an NBER recession indicator are both positive, implying it is countercyclical. So relative to levels, it appears essentially acyclical.

In first differences, the correlations are more consistent, with cross-sectional uncertainty associated with declines in industrial production, employment, and the CBO output gap, and with increases in unemployment. That said, even these results are not robust. The second and third columns of table 1 split the sample in the middle at the beginning of 1998.

Nine of the eleven change sign across the two subsamples, and by economically large amounts – e.g. from 0.19 to -0.26 for the unemployment rate. So, consistent with what is learned from visual inspection, at best we can say that cross-sectional uncertainty is high at the beginning of the last two recessions. Otherwise, though, it is not strongly cyclical, nor is it even particularly variable.

Finally, it is noteworthy how quickly cross-sectional uncertainty declines following its peaks in 2000 and 2008. In both cases, cross-sectional uncertainty falls back to its long-run average much more quickly than market-level uncertainty does. This fact is a challenge to models arguing that it is uncertainty that drives the weak recoveries from the two recessions. If uncertainty that caused the weak recoveries, it must be uncertainty about the state of the aggregate economy, not cross-sectional uncertainty.[9]

Alternative measures

The last three columns of table 1 report correlations of measures of *realized* volatility with economic activity. As discussed above, previous work has not directly measured uncertainty about cross-sectional shocks, but rather has examined the realized distribution of shocks. We examine here and below the relationship between the previously used measures and our forward-looking uncertainty.

Table 1 reports results for cross-sectional realized stock return volatility (as in Campbell et al. (2001)), the monthly interquartile range of growth in industrial production (across NAICS sectors) and the quarterly interquartile range of firm-level sales growth from Compustat. With the exception of stock returns, realized volatility is much more strongly countercyclical than uncertainty, underscoring the importance of the distinction between the two.

3.2 ARCH effects and the response of cross-sectional uncertainty to macro events

Another notable feature of the top panel of figure 1 is that firm-level uncertainty responds relatively weakly to spikes in realized volatility. It is well known that large returns in the aggregate stock markets are associated with persistent increases in implied volatility (this is the foundation of the ARCH literature). Many of the spikes in cross-sectional realized volatility in figure 1, however, are not associated with changes in implied volatility. This is most visible in October, 1987 and the late 1990's.

[9] See, for example, Fajgelbaum, Schaal, and Taschereau-Dumouchel (2017) and Kozlowski, Veldkamp, and Venkateswaran (2019), who discuss the potential effects of persistent aggregate, rather than cross-sectional, uncertainty.

To formalize that result, we estimate a regression of cross-sectional uncertainty on its own lag and cross-sectional realized volatility (reported in appendix table A.1):

$$\sigma^2_{\varepsilon,t} = b_0 + b_1\sigma^2_{\varepsilon,t-1} + rv_{\varepsilon,t} + \eta_t \tag{7}$$

The marginal R^2 from including realized volatility in that regression is 0.049. Estimating the same regression but with market implied and realized volatility, the marginal R^2 from realized volatility is 0.142. So three times as much of the variation in market uncertainty is explained by realized volatility as for idiosyncratic uncertainty.

Given that $\sigma^2_{\varepsilon,t}$ is correlated with realized volatility, one might ask what fraction of the variance of the innovations in $\sigma^2_{\varepsilon,t}$ depends on innovations in realized volatility. That is, can these variables be separated at all, or are their shocks perfectly correlated? The top panel of figure 1 certainly makes one think not, given how much noisier realized volatility is than uncertainty. In a VAR(1) using $\sigma^2_{\varepsilon,t}$ and $rv_{\varepsilon,t}$, the correlation between the innovations to the two variables is 0.74, which means that 54 percent (i.e. $1\text{-}0.74^2$) of the variance of innovations to $\sigma^2_{\varepsilon,t}$ can be explained by innovations to $rv_{\varepsilon,t}$. The remaining 46 percent – essentially half the variation – is orthogonal to realized volatility. This result quantifies the degree to which the two series are different. They share essentially half of their variation, with the remaining half being independent.

Table A.1 also shows that idiosyncratic uncertainty also responds only very weakly to market events. In a regression of $\Delta\sigma^2_{\varepsilon,t}$ on market realized volatility, the coefficient is 0.06 and the R^2 is 0.05, compared to 0.16 and 0.21 for a regression of $\Delta\sigma^2_{mkt}$ on market realized volatility. So changes in idiosyncratic uncertainty are only weakly correlated with market realized volatility, unlike changes in aggregate uncertainty.

To further explore the relationship between events at the market level and changes in idiosyncratic uncertainty, major events associated with spikes in uncertainty, such as the September 11th attacks, the Asian financial crisis, etc., are labeled in the bottom panel of figure 1 (drawn mostly from Bloom (2009)). Those jumps in aggregate uncertainty, with the exception of the financial crisis, do not pass through to cross-sectional uncertainty.

To be clear, that result does not mean that macro uncertainty shocks are not associated with increases in the total uncertainty that firms face. Recall the decomposition from above,

$$\sigma^2_{i,t} = \beta^2_{i,t}\sigma^2_{mkt,t} + \sigma^2_{\varepsilon,i,t} \tag{8}$$

An increase in $\sigma^2_{mkt,t}$ will increase every firm's overall uncertainty, $\sigma^2_{i,t}$. What figure 1 and table A.1 show is that increases in $\sigma^2_{mkt,t}$ are not usually associated with increases in the cross-sectional part of uncertainty, $\sigma^2_{\varepsilon,i,t}$. Changes in $\sigma^2_{mkt,t}$ pass through nearly one-for-one into $\sigma^2_{i,t}$ (for $\beta \approx 1$), whereas if idiosyncratic uncertainty, $\sigma^2_{\varepsilon,i,t}$ responded to $\sigma^2_{mkt,t}$, then a unit shift in $\sigma^2_{mkt,t}$ would lead to a shift in $\sigma^2_{i,t}$ of more than that amount.

3.3 Robustness

Figure A.3 plots five variations on the benchmark uncertainty series:

1. Using the unweighted mean or median of implied volatility across firms (after taking out firm fixed effects), instead of weighting by market capitalization.

2. Correcting for each firm's loading on the market, by estimating β_i for each firm and setting $\sigma^2_{\varepsilon,i,t} = \beta_i^2 \sigma^2_{mkt,t} - \sigma^2_{i,t}$ (the results are highly similar using loadings estimated from rolling window regressions using $\beta_{i,R,t}$ from above).

3. Excluding tech stocks from the sample.

4. Controlling for industry shocks. This is done by replacing the market return in equation (6) with the return on the SPDR sector exchange traded fund (ETF) for the stock's sector. Equation (3) then uses the implied volatility for the sector ETF, instead of the market.

The first test shows that the results are not driven just by the weighting by market capitalization – they hold more broadly across our sample of stocks – in the sense that cross-sectional implied volatility is relatively stable over time outside the dot-com boom and financial crisis. Similarly, the second test shows that the approximation where we treat the loadings on the market as all equal to 1 has very little impact.

The third test, checking the importance of tech stocks, is obviously most important around the dot-com boom. Whereas in the full sample the dot-com boom has the largest peak in cross-sectional uncertainty, when tech stocks are excluded, uncertainty in the financial crisis becomes slightly higher than in the dot-com boom. None of the qualitative features of figure 1 are changed, however.

Finally, the fourth test shows that little of the cross-sectional uncertainty is driven by uncertainty shocks. Since options on sector ETFs did not exist until December, 1998, we cannot study the full sample in this case. In the available data, though, taking out the part of the variation explained by industry uncertainty has relatively minor effects. It always reduces cross-sectional uncertainty (since industries are one source of that uncertainty), with the largest effect during the financial crisis. The dynamics are again highly similar to the baseline, however.

In addition, section 7 reports estimates of cross-sectional uncertainty in six European countries and shows that they are qualitatively and quantitatively similar to the series for the US.

4 The predictive power of uncertainty for real activity

To further explore the relationship between cross-sectional uncertainty and the state of the economy, we now examine forecasting regressions to quantify whether shocks to cross-sectional uncertainty predict declines in real activity.

For monthly data, we examine three standard outcomes: the unemployment rate, non-farm private employment, and industrial production. Table 2 reports results from these forecasting regressions. To account for potential lagged effects, we regress activity on three-month moving averages of implied and realized volatility (results using individual monthly lags are similar, though statistically weaker). In all cases, we include three lags of the dependent variable, but do not report the coefficients in the table. All of the variables are standardized to have unit variance to aid interpretation. In each case, the dependent variable is entered as a first difference, so it is the changes that have unit variance.

4.1 Cross-sectional uncertainty alone

The first column in each of the three panels reports a regression of activity on lagged cross-sectional uncertainty. In all three cases, the coefficient implies that increases in uncertainty are followed by declines in real activity. The coefficients are significant at the 1-percent level for employment and the unemployment rate and at the 5-percent level for industrial production. So, consistent with the results in table 1, there is a relationship between cross-sectional uncertainty and changes in real activity. The magnitudes of the coefficients are similar, with a unit standard deviation increase in uncertainty being associated with declines in IP and employment of 0.09 standard deviations and an increase in unemployment of 0.15 standard deviations.

From the analysis above, the behavior of cross-sectional uncertainty appears rather different in the first and second halves of the sample. The second column of table 2 therefore estimates the forecasting regression allowing the coefficient on lagged uncertainty to change between the first and second halves of the sample. It shows that the forecasting power is isolated entirely in the second half. In all three regressions, the coefficients change sign between the two subsamples, and the changes are statistically significant for IP and employment. The coefficients in the second half are of the same magnitude as in the full sample, while in the first half their signs are reversed and the values are twice as large. So while cross-sectional implied volatility leads real activity in the simplest regression, that result is driven entirely by data from just half the sample – consistent with figure 1 – and thus it is not a strong stylized fact.

4.2 Cross-sectional uncertainty versus realized volatility

Table 1 shows that measures of realized cross-sectional volatility seem to be more strongly cyclical than option-implied uncertainty. One of the most useful features of our measure is that it allows us to directly compare uncertainty and realized volatility for the same underlying concept – stock returns. As in table 1, realized volatility here is calculated using the formula in (**??**). The third column of table 2 shows that when realized volatility is included in the forecasting regressions, it is associated with declines in both employment and industrial production and increases in unemployment, while implied volatility now has the wrong sign for employment and IP – being associated with expansions. The coefficients on realized volatility are also generally larger than those on uncertainty from column 1. These results are consistent with those of Campbell et al. (2001) and Gilchrist, Sim, and Zakrajsek (2014), who also find that realized cross-sectional volatility has forecasting power for real activity – our contribution in this column is to show that volatility actually drives out uncertainty.

4.3 Cross-sectional versus aggregate uncertainty

The last question about the regressions is whether it is cross-sectional or aggregate variation that is actually relevant. Overall, the correlation between S&P 500 uncertainty and cross-sectional uncertainty is 0.42 – cross-sectional uncertainty has historically moved with aggregate uncertainty (though this is also sample specific – the correlation is -0.29 in the first half and 0.55 in the second half). That implies that it is important to control for market uncertainty to isolate what is actually driving the results.

The fourth column does that and shows that market uncertainty drives cross-sectional uncertainty out of the regressions. In all three cases, the coefficient on cross-sectional uncertainty shrinks substantially (changing sign in one case) and becomes statistically insignificant, while the coefficient on aggregate uncertainty is significant in all three cases with the expected signs and with economically meaningful magnitudes.

So to the extent that uncertainty matters at all, it is aggregate uncertainty that is most important, rather than cross-sectional uncertainty. We show in section 5 that this result is useful in disciplining the type of model that can fit the data.

Similar to above, table 2 also asks whether it is actually aggregate uncertainty that drives the results, or if aggregate realized volatility (i.e. squares returns on the overall stock market) that is relevant. The fifth and final column of table 2 reports results from a full regression including both cross-sectional and market implied and realized volatility. Consistent with Berger, Dew-Becker, and Giglio (2019), market realized volatility is dominant and only its coefficient is statistically significant. It has significant forecasting power for two of the three of the outcomes, even after controlling for market implied volatility and cross-sectional realized and implied volatility. Moreover, its coefficient, ranging in magnitude between 0.16 and 0.35, is larger than we observed for any of the uncertainty. That is particularly surprising given that realized volatility appears to be noisy relative to uncertainty.

Finally, market and cross-sectional uncertainty are statistically insignificant and in fact have the wrong sign – the full regression shows that, after controlling for realized volatility, uncertainty is if anything expansionary.

Similar to table 1, then, these results show that it is critical to distinguish between uncertainty – which is about agents' conditional distributions for the future – and realized volatility, which measures the magnitude of shocks that have already occurred. It is the shocks that have already occurred that are relevant for forecasting, not the shocks that agents expect going forward. Moreover, it is realized volatility at the aggregate level that matters – neither cross-sectional uncertainty nor cross-sectional realized volatility ends up being relevant for forecasting.

4.4 Quarterly results

To further examine the forecasting power of uncertainty and realized volatility, table 3 reports results for quarterly regressions for output, consumption, investment, and hours worked. Consistent with table 2, in each case the right-hand side includes three lags of the dependent variable along with the previous quarter's value of the implied or realized volatility series. In the baseline regression just including cross-sectional uncertainty, none of the coefficients except for hours is statistically significant in this case. Column 4 shows that, similar to what is observed for monthly data, that the results are driven by realized volatility. In the full specification for investment – which is often seen as most likely to respond to uncertainty – both market and cross-sectional realized volatility have highly significant negative coefficients, while the coefficients on implied volatility are actually significantly positive, consistent with the results in table 2.

Overall, then, we conclude from this section that cross-sectional uncertainty has only extremely weak power for forecasting real activity – it works only in univariate regressions, and its statistical power is isolated to the second half of the sample. As soon as any controls are included in the regression, cross-sectional uncertainty is immediately driven out, and market realized volatility ends up dominant out of the variables we examine. So for policymakers, such as central banks, these results argue that cross-sectional uncertainty should not be used as an indicator of fundamental weakness that needs to be addressed, and it is not uncertainty, but rather the realization of large shocks that is best associated with future contractions.

5 Comparison to theoretical models

The analysis so far essentially reports moments – we make no claims about structural identification. We therefore now test two theoretical models by asking whether they match the results from the reduced-form regressions. Specifically, this section replicates the regressions above as closely as possible in two leading models of the effects of cross-sectional uncertainty shocks. The models are the "really uncertainty business cycles" (RUBC) model of Bloom et al. (2018), which is centered around a real options framework, and the model of financial frictions of Christiano, Motto, and Rostagno (CMR; 2016).

5.1 Christiano, Motto, and Rostagno (2016)

CMR is a New Keynesian business cycle model with financial frictions similar to those in Bernanke, Gertler, and Gilchrist (1999). There is a set of entrepreneurs who operate capital subject to an endogenous leverage constraint. When uncertainty is higher, the leverage constraint is tighter, forcing entrepreneurs to reduce investment, causing a recession. The volatilities of the aggregate shocks are constant, but cross-sectional uncertainty and realized volatility fluctuate over time.

Table 4 reports statistics from a simulation of the CMR model compared to what is observed empirically. We use simulations generated by the replication code posted with the published paper. The top panel reports correlations of GDP, investment, consumption, and hours with cross-sectional uncertainty. The left-hand section reports correlations of growth rates of the real quantities with changes in uncertainty, while the right-hand section reports correlations of levels of real quantities (minus linear trends) with levels of uncertainty.

In both levels and growth rates, the correlations of the real variables with uncertainty in the model are substantially larger than those in the data. In growth rates, the average correlation in the model is -0.27, compared to -0.02 in the data. In levels, the average correlation in the model is even more negative at -0.47, while in the data it is actually positive at 0.25. The model predicts that cross-sectional uncertainty should be cyclical in both levels and growth rates, with the correlations larger for levels, whereas in the data uncertainty is only weakly countercyclical at best, and for levels it appears if anything to be procyclical.

The remaining panels of table 4 report results from forecasting regressions for the same four real variables using cross-sectional implied and realized volatility. In each case, we report results first for univariate forecasts using just cross-sectional uncertainty, and second for a regression that also includes cross-sectional realized volatility.[10] The regressions in this section are not exactly identical to the forecasting regressions above because in CMR volatility is much more persistent and real activity responds to uncertainty shocks immediately, rather than with a lag.[11]

In all four cases, in the model simulations the coefficient on the first difference of cross-sectional implied volatility is negative. For GDP and investment, a unit standard deviation increase in uncertainty is associated with a 0.24 standard deviation decline in activity, compared to only 0.06 in the data. So while idiosyncratic uncertainty has quantitatively and statistically weak effects on activity empirically, in the CMR model it has large effects.

The second column shows that, counter to what was observed empirically above, when idiosyncratic realized volatility is included in the regression with simulated data, the coefficient on implied volatility is unchanged, while the coefficient on realized volatility is uniformly positive. In the data, on the other hand, the coefficient on implied volatility actually switches sign in most cases, and it is realized volatility that has the strongly negative coefficient.

So while the data implies that idiosyncratic implied volatility is driven out or weakened by realized volatility, in the CMR model it dominates, retaining its forecasting power even after realized volatility is included. The CMR model thus appears to be inconsistent with the data in that it implies counterfactually that idiosyncratic uncertainty both is a strong univariate forecaster of real activity and is not driven out by other variables.

Finally, recall that in the results above we found that aggregate uncertainty varied substantially more than cross-sectional uncertainty, and that aggregate uncertainty drives cross-sectional uncertainty out of forecasting regressions. It is important to note that CMR has zero variation in aggregate uncertainty. Furthermore, it is a model that has no specific channel for aggregate uncertainty to have substantial real effects. It thus very likely has the

[10]Firm-specific realized volatility is a variable generated by the simulations of the model (in the code, called volEquity). For market realized volatility, we use the squared value of the innovation in entrepreneur net worth, which CMR link to the level of the aggregate stock market.

[11]Specifically, the regressions use first differences of implied and realized volatility. Furthermore, in the empirical regressions here we measure implied volatility at the very beginning of the quarter (as in the model) rather than averaging over the previous quarter. Realized volatility is still measured as the sum over the previous quarter, but the value from the current quarter can be used instead and the results are qualitatively unchanged and quantitatively similar (in fact stronger).

counterfactual prediction that cross-sectional uncertainty is more important for explaining real activity than aggregate uncertainty.

5.2 Bloom et al. (2018)

We next consider results for RUBC. The RUBC model has fixed adjustment costs in investment. In certain calibrations, that assumption can cause variation in the conditional variance of shocks to reduce investment and thus real activity (though for multiple reasons this is not a global result).

In the RUBC model, aggregate and cross-sectional uncertainty are assumed to be perfectly correlated, meaning that it is not possible in the baseline specification to disentangle them. We solve that problem by running simulations in which the volatilities of aggregate and cross-sectional uncertainty vary by different amounts. We then append those simulations into a single long sample that allows us to investigate the behavior of the model when aggregate and cross-sectional uncertainty are imperfectly correlated.

More specifically, in the baseline RUBC calibration, aggregate and cross-sectional uncertainty rise by a factors of 1.61 and 4.14, respectively, in the high-uncertainty state. We construct two additional simulations in which they rise by the factors $\{1.72, 3.82\}$ and $\{1.48, 4.45\}$. When the three simulations are appended, aggregate and cross-sectional uncertainty are still very highly correlated, but no longer perfectly. Intuitively, the simulations essentially allow us to ask whether volatility in aggregate or idiosyncratic uncertainty has a larger effect on output.

As with CMR, the RUBC model predicts the uncertainty is substantially more countercyclical than we observe in the data. The top panel of table 5 shows that both in levels and growth rates, the RUBC model predicts uncertainty to be strongly negatively correlated with real activity, while in the data that correlation is close to zero (or even positive in levels).

Table 5 also reports results of regressions of output and investment growth on quarterly differences of aggregate and cross-sectional uncertainty in the RUBC model and the data. This time, instead of comparing uncertainty and realized volatility, we compare cross-sectional and aggregate uncertainty (this is because it is not possible for us, using the basic replication files for RUBC, to calculate realized stock return volatility).

In the model simulations, cross-sectional uncertainty is by far the most important driver of the forecasting results. Conditional on cross-sectional uncertainty, aggregate uncertainty is useless for forecasting real activity. In the data, though, that result is reversed: in three out of four cases, the coefficient on aggregate uncertainty is larger than that on cross-sectional uncertainty. And recall that the results in table 2 were even stronger for monthly data –

aggregate uncertainty completely drove out cross-sectional uncertainty.

Overall, then, both in terms of forecasting regressions and raw correlations, the RUBC model does not seem to be able to replicate the empirical results that we obtain with our measure of cross-sectional uncertainty.

6 Asset returns

So far we have shown that uncertainty is at best weakly cyclical and that it has little or no forecasting power for real activity. An alternative way to answer the question whether idiosyncratic uncertainty has real effects is to examine the cost to hedge it. If increases in uncertainty are bad in the sense that they are associated with states of the world with high marginal utility, then people should be willing to pay for insurance against uncertainty shocks.

More formally, consider a risk factor X (e.g. uncertainty) and an asset return, R_X, that hedges it in the sense that R_X is positively correlated with X. Specifically, suppose they are related through

$$R_{X,t} = R_{f,t} + \bar{R}_X + aX_t + \varepsilon_t \tag{9}$$

for a constant $a > 0$ (simply as a normalization) and where $\bar{R}_X$ is the average excess return of the asset R_X (where X has been normalized to have conditional mean zero; i.e. it is an innovation in uncertainty), R_f is the risk-free rate, and ε is the component of R_X uncorrelated with X. An ideal hedge would have $\varepsilon = 0$ in all states of the world, but that may not be possible in practice.

Defining M_t to be the stochastic discount factor (i.e. the Arrow–Debreu state prices divided by state probabilities), we have

$$0 = E_{t-1}\left[\left(R_{X,t} - R_{f,t}\right) M_t\right] \tag{10}$$

$$\frac{\bar{R}_X}{R_{f,t}} = -\underset{t-1}{\operatorname{cov}}\left(M_t, aX_t\right) - \underset{t-1}{\operatorname{cov}}\left(M_t, \varepsilon_t\right) \tag{11}$$

The first term on the right-hand side says that when increases in X are bad, in the sense that states when X is above its mean are high marginal utility states, the average return on R_X will be negative. The second term says that if the residual is also priced – i.e. if the part of R_X uncorrelated with X is related to marginal utility, then that can also potentially affect the average return.

The goal of this section is to construct asset returns, R_X, that hedge uncertainty. Equation (11) shows that ideally we want those returns to be as highly correlated with innovations to uncertainty as possible so as to minimize the magnitude of the ε term. As long as the ε term is not too large – in the sense that its covariance with M is small – the sign of $-E[R_X - R_f]$ identifies the sign of $\text{cov}(M, aX)$.

6.1 Option portfolios

The implied volatilities studied in the previous sections are based on option prices, so innovations in implied volatility are simply innovations in option prices. The most natural way to hedge those shocks, then, is to use portfolios of the options themselves. As in Cremers, Halling, and Weinbaum (2015) and Dew-Becker, Giglio, and Kelly (2019), we study returns on straddles to give exposure to uncertainty. A straddle is a portfolio long a put and call option with the same strike, roughly at the money (technically, chosen so that the Black–Scholes delta is 1/2). The appendix describes the details of the construction of the portfolio returns. Importantly, the data includes prices for options with multiple maturities on any given date.

Using an analysis similar to those in Cremers, Halling, and Weinbaum (2015) and Dew-Becker, Giglio, and Kelly (2019), the return on an n-month straddle for stock i can be approximated as

$$r_{n,i,t} \approx \frac{1}{2} n^{-1} \left(\frac{r_{i,t}}{\sigma_{i,t-1}} \right)^2 + \frac{1}{2} \frac{\Delta\sigma_{i,t}^2}{\sigma_{i,t-1}^2} \tag{12}$$

where $r_{n,i,t}$ is the return on an n-month straddle, $r_{i,t}$ is the return on the underlying asset, and $\sigma_{i,t}$ is the implied volatility for the underlying asset. The approximation says straddles are exposed, with a coefficient of 1/2, to proportional shocks to implied variance. They also carry exposure to squared shocks to the underlying asset itself. Intuitively, since the final payoff of a straddle depends on the absolute value of the return on the underlying, it has zero local exposure to returns on the underlying, but it is affected by higher-order variation in returns. Extreme returns in either direction are good for a straddle.

Using two maturities (i.e. two values of n), we can then construct a pair of portfolios for each stock that yield exposure to realized volatility, r_i^2, or implied volatility, $\Delta\sigma_{i,t}^2$.

Specifically,[12]

$$rv_{i,t} \equiv \sigma^2_{i,t-1}\left(r_{n_1,t} - r_{n_2,t}\right)\frac{1}{n_1^{-1} - n_2^{-1}} \tag{13}$$

$$iv_{i,t} \equiv \sigma^2_{i,t-1}\left(n_1 r_{n_1,t} - n_2 r_{n_2,t}\right)\frac{1}{n_1 - n_2} \tag{14}$$

Since we now have portfolios exposed to the total realized and implied volatility at the firm level, by value-weighting and subtracting analogous portfolios for the total stock market, we can construct portfolios with exposure to idiosyncratic volatility.[13] Specifically,

$$rv_{idio,t} \equiv \sum_i w_{i,t-1} rv_{i,t} - rv_{mkt,t} = rv_{idio,t} \tag{15}$$

$$iv_{idio,t} \equiv \sum_i w_{i,t-1} iv_{i,t} - iv_{mkt,t} \approx \Delta\sigma^2_{idio,t} \tag{16}$$

(where the second line includes the approximation that $\sum_i w_{i,t-1}\Delta\sigma^2_{idio,i,t} \approx \sum_i w_{i,t}\sigma^2_{idio,i,t} - w_{i,t-1}\sigma^2_{idio,i,t-1} = \Delta\sigma^2_{idio,t}$)

$iv_{idio,t}$ is the main object of interest here. It is the return on a portfolio that gives exposure to changes in idiosyncratic uncertainty. $rv_{idio,t}$ is a natural companion, giving exposure not to changes in uncertainty, but to the realization of volatility. That is, one portfolio depends on changes to expected future squared returns, while the other depends on the squared returns that actually occur.

6.2 Results

Table 6 reports estimates of risk premia. The first column reports the risk premia the rv and iv portfolios for idiosyncratic and market risk. To help with interpretation, we report Sharpe ratios – average returns divided by standard deviations – because they are invariant to leverage. The Sharpe ratio represents the excess return per unit of risk (measured by standard deviation). For reference, over our sample the Sharpe ratio for the total stock market was 0.56.

The two iv portfolios, for idiosyncratic and market uncertainty, both have positive Sharpe ratios over the sample – iv_{idio} at 0.82 and iv_{mkt} at 0.15. The latter result is consistent with the results of Dew-Becker, Giglio, and Kelly (2019) and others. The result for iv_{idio} is

[12]These portfolios have the same name as a similar pair of portfolios in Dew-Becker, Giglio, and Kelly (2019), but are scaled by $\sigma^2_{i,t-1}$ to give the exposures we need here

[13]The multiplication by $\sigma^2_{i,t-1}$ is essentially a time-varying leverage choice (under the assumption that the time periods are sufficiently short that the risk-free return is negligable, which will hold in practice).

novel to this paper. It says that a portfolio constructed to hedge shocks to cross-sectional uncertainty earns *positive* returns. Given equation (11), that implies that investors view increases in idiosyncratic volatility as on average being associated with low marginal utility. That is consistent with high cross-sectional uncertainty being associated with high growth, such as in the late 1990's.

The obvious question is how these results could be consistent with the fact that uncertainty was also high during the financial crisis. Obviously uncertainty is high in both good times. What the asset returns imply is that on average investors view rises in uncertainty as coming during good times. That is, when options are priced, regimes like the late 1990's are apparently viewed as a more typical outcome than regimes like the financial crisis.

For the two rv portfolios, table 6 shows that the returns have the opposite sign. Exposure to realized volatility earns a substantially negative risk premium for both rv_{idio} and rv_{mkt}. Again, the negative premium, along with its magnitude, for rv_{mkt} is consistent with past work. rv_{idio} also carries a negative premium, but it is substantially smaller – at the point estimates by a factor of three. These results are consistent with the fact that realized volatility appears to be a relatively robust predictor of declines in real activity.

The middle and bottom sections of table 6 show that these results, similar to those above, are consistent over time. In both the first and second halves of the sample, the premium for cross-sectional uncertainty – iv_{idio} – is positive, while the premia for market and cross-sectional realized volatility are consistently negative.

What these results say is that increases in uncertainty – when volatility is expected to be high in the future – are associated with low marginal utility. Investors view them as generally happening in good times. Surprises in realized volatility on the other hand – i.e. when the cross-sectional dispersion in returns is higher than expected – are associated with bad times, or high marginal utility. Again, these statements are just averages. Obviously it may be the case that sometimes cross-sectional uncertainty is bad – certainly at least during the financial crisis, it was associated with high marginal utility. But overall, investors apparently view periods of high uncertainty as coming in good states.

Directly studying the iv portfolio is useful because it requires almost no estimation – we just need to measure its mean return. The drawback, though, is that we are relying on an approximation for its exposure, which one might worry could be inaccurate. As an alternative, the appendix describes results using a factor pricing model, which are qualitatively similar, though quantitatively more extreme. Overall, we view the simple returns as more economically plausible, and since they also are produced much more transparently, we focus on them here.

7 International evidence

The stability of cross-sectional uncertainty relative to aggregate uncertainty is a surprising result, so one might ask whether it also holds in other countries. Furthermore, the degree of comovement in cross-sectional uncertainty is informative for understanding the sources of the shocks that firms face. This section constructs market-level and cross-sectional uncertainty using similar methods to the analysis for the US above, using data for six European countries. It shows that not only does cross-sectional uncertainty display similar patterns in terms of overall volatility, but it is strikingly correlated across countries.

7.1 Data

Similar to the main analysis, we obtain data on individual firm and index-level implied volatilities from Optionmetrics using the volatility surface file, this time using the Europe database (there is also a database covering Asia, but its coverage is sparse until the late 2000's). The data in this case starts in 2002 and runs through the end of 2018. Data on market capitalization for individual firms is obtained from the Compustat Global database (matched to Optionmetrics by ISIN). To include a country in the sample, it must have option data for a high-level stock market index since at least 2008. We have acceptable data for Switzerland, Germany, Spain, France, Great Britain, and the Netherlands.[14] While all of the countries are from western Europe, the list includes countries with varying degrees of connection to the EU and countries on different currencies and with very different government fiscal states.

As in the main analysis, market-level uncertainty is calculated as at-the-money implied volatility for the stock market index. Cross-sectional uncertainty is (the square root of) average firm-level implied variance minus market-level implied variance, where the average is weighted by firm market capitalization.

7.2 Results

Figure 3 plots the time series of cross-sectional uncertainty for each country against that for the US. Across all six countries, cross-sectional uncertainty is clearly strongly correlated with that in the US. In all six countries, cross-sectional uncertainty also is elevated in 2002, declines until rising during the financial crisis, and then stays low and stable since 2010. That is true even though the path of aggregate output in Europe over this period was very different – a number of these countries went into recessions around 2012.

[14]The associated stock indexes are the SMI, DAX, IBEX 35, CAC 40, FTSE 100, and AEX.

Figure 4 plots the time series of aggregate uncertainty and shows that they are also highly similar across countries. While uncertainty in Spain was particularly elevated during the European sovereign debt crisis, the various series are otherwise highly correlated.

To formalize those results, the various figures report the correlation between each type of uncertainty in each country and the US. Across the six countries, the average correlation of cross-sectional uncertainty with that in the US is 0.80, while for market uncertainty the value is 0.91. Those numbers imply that US uncertainty explains on average 65 percent of the variation in cross-sectional uncertainty and 83 percent of the variation in market uncertainty in European countries. Furthermore, any measurement error in the uncertainty series will bias those values down.

The table below summarizes the various results for market and cross-sectional uncertainty. It shows that the average time series standard deviation of market uncertainty is nearly twice that of cross-sectional uncertainty, consistent with the findings for the US. Furthermore, it confirms the results on the fraction of the variation in uncertainty explained by a common factor – this time using the cross-sectional mean of uncertainty, rather than the US value. Finally, it reports the simple average of all the pairwise correlations across countries and shows the high degree of similarity for both types of uncertainty.

Statistics for cross-sectional and market uncertainty across countries

	Cross-sectional unc.	Market unc.
Avg. time-series s.d.	0.0474	0.0803
Avg. cross-sectional s.d.	0.0342	0.0322
Frac. of var. explained		
by cross-sectional mean	0.60	0.86
Avg. pairwise corr.	0.82	0.92

The results here are important for two reasons. First, they show that our findings are not unique to the US: cross-sectional uncertainty has been similarly stable in other major developed economies. Second, they show that there appears to be a very strong global (or at least international) factor driving cross-sectional uncertainty.

The latter result is particularly surprising. It is simple to envision a model in which there is a common component in market level uncertainty. That behavior will appear whenever there is a heteroskedastic shock that affects stock prices globally (e.g. news about global growth or changes in global risk appetite). A trade war, for example, might generally be expected to negatively impact all economies, so when there is more uncertainty about trade we would expect market uncertainty to rise everywhere.

But as we saw above, aggregate shocks typically do not pass into cross-sectional uncertainty. So the common variation in that type of uncertainty must come from common movements across countries in the volatility of shocks that affect relative performance across firms. In other words, whatever the shocks are that cause these cross-sectional movements – shifts in creative destruction, for example – they must occur across the various economies that we study here simultaneously.

8 Conclusion

This paper studies the behavior of cross-sectional implied volatility. A large literature studies the effects – both good and bad – of variation in the cross-sectional distribution of shocks that firms face. There is theoretical ambiguity about the effects of changes in cross-sectional uncertainty, but many policymakers appear to take the view that uncertainty represents a hindrance to economic growth. It is thus an important empirical question not just what the time series of firm-level of uncertainty has looked like, but also whether shocks to cross-sectional uncertainty are in fact contractionary.

We develop a novel index of cross-sectional uncertainty with data extending back to 1980. We show that the length of the sample is important – it is the data in the 1980's and early 1990's that emphasizes the extent to which the last two recessions have been anomalous. Prior to the late 1990's, there was little variation in cross-sectional uncertainty. Since then, there have been two episodes where it substantially grew, one a major economic expansion and the other a contraction. Studying raw correlations and forecasting regressions, we find that cross-sectional uncertainty is approximately acyclical and has little ability to forecast changes in future real activity. Furthermore, risk premia associated with portfolios that hedge cross-sectional uncertainty shocks imply that investors do not view periods of high uncertainty as bad.

Some recent work emphasizes models in which cross-sectional uncertainty is contractionary, but in many cases it can be expansionary. That is a well known result in models with fixed costs of adjustment (Oi (1961), Hartman (1972), and Abel (1983), for example). It is also natural in an economy where production can shift to the most efficient or productive firm. The optionality associated with optimization in general implies that more cross-sectional dispersion is a good thing. Our empirical results appear to be consistent with that view of the economy.

References

Abel, Andrew B., "Optimal Investment Under Uncertainty," *American Economic Review*, 1983, *73* (1), 228–233.

Bachmann, Rudiger and Giuseppe Moscarini, "Business Cycles and Endogenous Uncertainty," 2012. Working paper.

——, **Kai Carstensen, Stefan Lautenbacher, and Martin Schneider**, "Uncertainty and Change: Survey Evidence of Firms Subjective Beliefs," 2018. Working paper.

——, **Steffen Elstner, and Eric R Sims**, "Uncertainty and economic activity: Evidence from business survey data," *American Economic Journal: Macroeconomics*, 2013, *5* (2), 217–249.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Bakshi, Gurdip and Nikunj Kapadia, "Delta-Hedge Gains and the Negative Market Volatility Risk Premium," *The Review of Financial Studies*, 2003, *16(2)*, 527–566.

Basu, Susanto and Brent Bundick, "Uncertainty Shocks in a Model of Effective Demand," 2015. Working paper.

Bewley, Truman, "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in Werner Hildenbrand and Andreu Mas-Colell, eds., *Contributions to mathematical economics: In honor of Gerard Debreu*, New York: North-Holland, 1986.

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.

Born, Benjamin, Sebastian Breuer, and Steffen Elstner, "Uncertainty and the great recession," *Oxford Bulletin of Economics and Statistics*, 2018, *80* (5), 951–971.

Buraschi, Andrea, Fabio Trojani, and Andrea Vedolin, "When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia," *The Journal of Finance*, 2014, *69* (1), 101–137.

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.

Carr, Peter and Liuren Wu, "Variance Risk Premiums," *Review of Financial Studies*, 2009, *22(3)*, 1311–1341.

Chen, Zhanhui and Ralitsa Petkova, "Does idiosyncratic volatility proxy for risk exposure?," *The Review of Financial Studies*, 2012, *25* (9), 2745–2787.

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

Cremers, Martijn, Michael Halling, and David Weinbaum, "Aggregate Jump and Volatility Risk in the Cross-Section of Stock Returns," *The Journal of Finance*, 2015, *70* (2), 577–614.

Dew-Becker, Ian, Stefano Giglio, and Bryan Kelly, "Hedging macroeconomic and financial uncertainty and volatility," 2019. Working paper.

Driessen, Joost, Pascal J. Maenhout, and Grigory Vilkov, "The Price of Correlation Risk: Evidence from Equity Options," *The Journal of Finance*, 2009, *64(3)*, 1377–1406.

Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty traps," *The Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajsek, "Uncertainty, Financial Frictions, and Investment Dynamics," Working Paper 20038, National Bureau of Economic Research April 2014.

Goyal, Amit and Alessio Saretto, "Cross-section of option returns and volatility," *Journal of Financial Economics*, 2009, *94* (2), 310–326.

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

Hartman, Richard, "The effects of price and cost uncertainty on Investment," *Journal of Economic Theory*, 1972, *5* (2), 258–266.

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.

Jurado, Kyle, Sydney Ludvigson, and Serena Ng, "Measuring Uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Kaplan, Greg and Giovanni L Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, "The tail that keeps the riskless rate low," *NBER Macroeconomics Annual*, 2018, *33* (1), 253–283.

___ , ___ , and ___ , "The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation," 2019. Working paper.

Leduc, Sylvain and Zheng Liu, "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 2016, *82*, 20–35.

Oi, Walter, "The Desirability of Price Instability Under Perfect Competition," *Econometrica*, 1961, *29* (1), 58–64.

Pástor, L'uboš and Pietro Veronesi, "Technological Revolutions and Stock Prices," *The American Economic Review*, 2009, pp. 1451–1483.

Figure 1: Time series of cross-sectional uncertainty

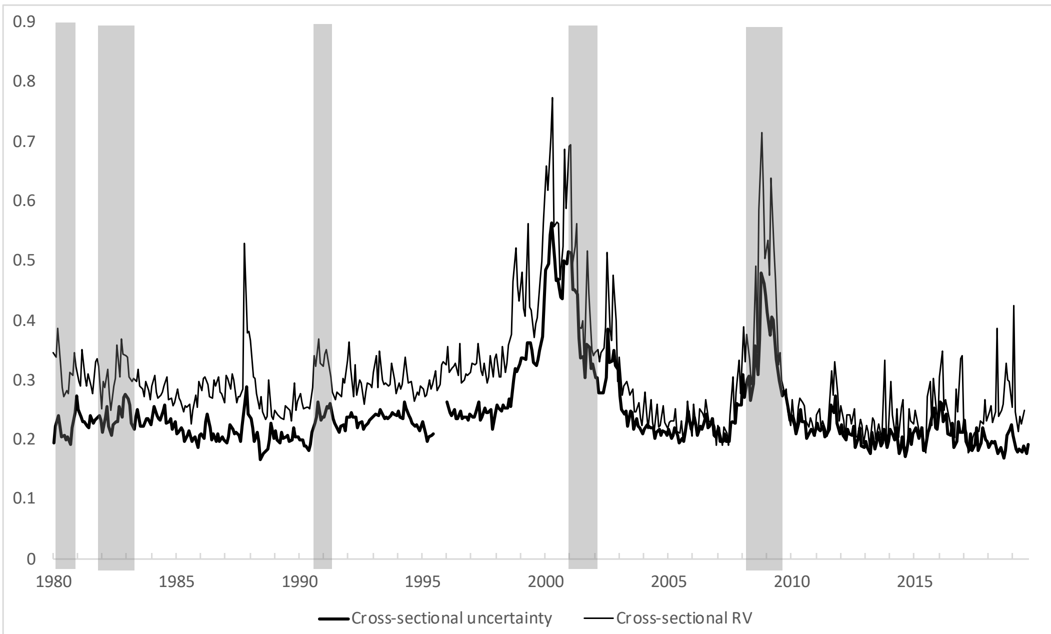

(a) Cross-sectional uncertainty and cross-sectional realized variance

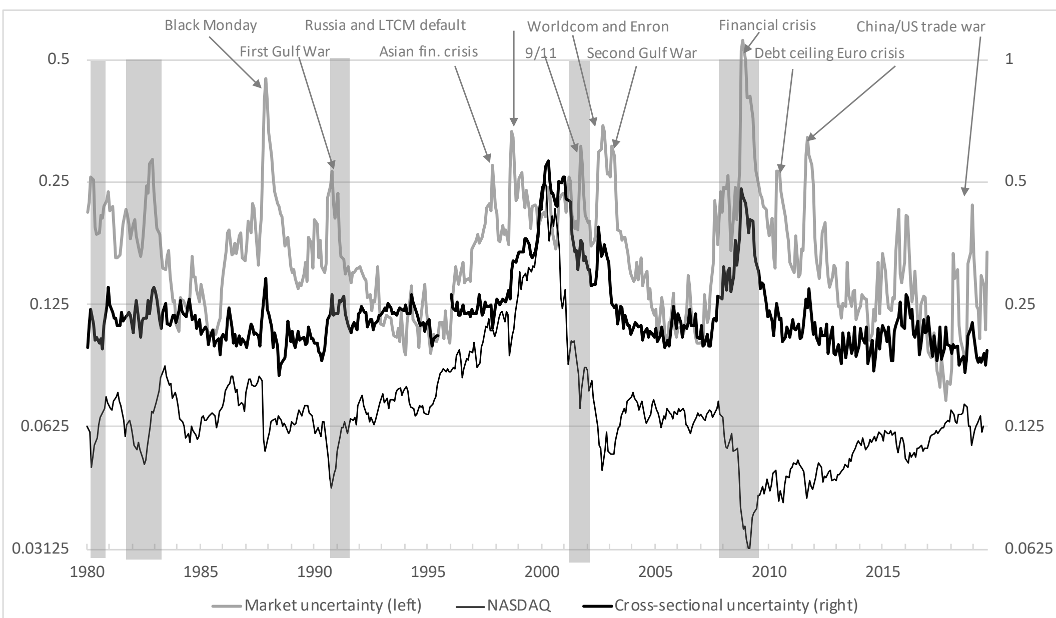

(b) Cross-sectional uncertainty and VIX

Note: The top panel reports the time series of cross-sectional uncertainty and cross-sectional realized variance. The bottom panel reports the time-series of cross-sectional uncertainty, the VIX, and the detrended cumulative value of the NASDAQ index. Options data before 1996 is from the Berkeley Options Dataset, and options data after 1996 is from Optionmetrics. The VIX is obtained from CME options. Shaded areas are NBER recessions.

Figure 2: Cross-sectional uncertainty and the real economy

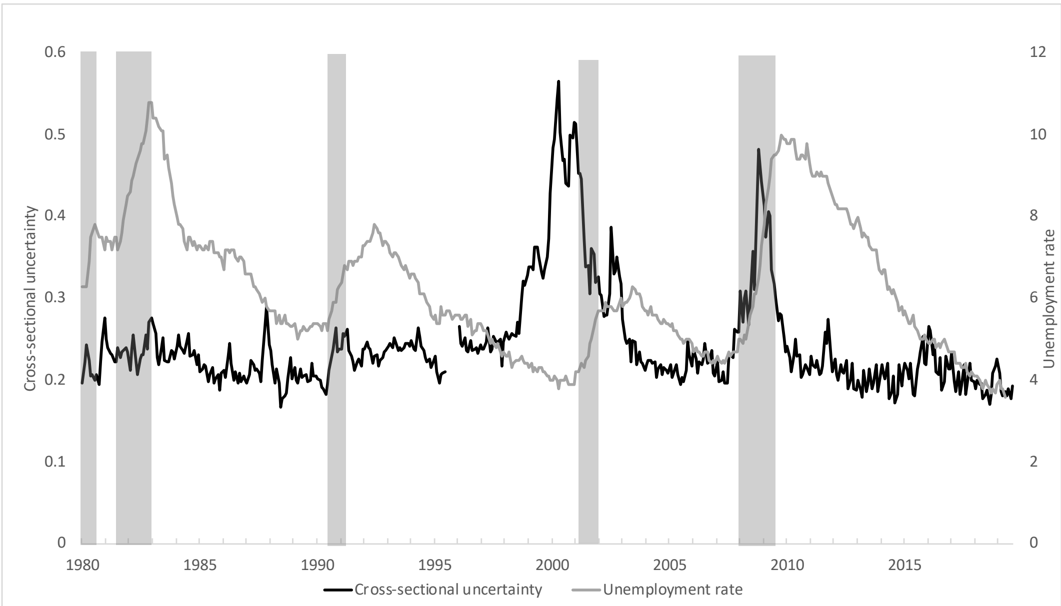

(a) Cross-sectional uncertainty and the unemployment rate

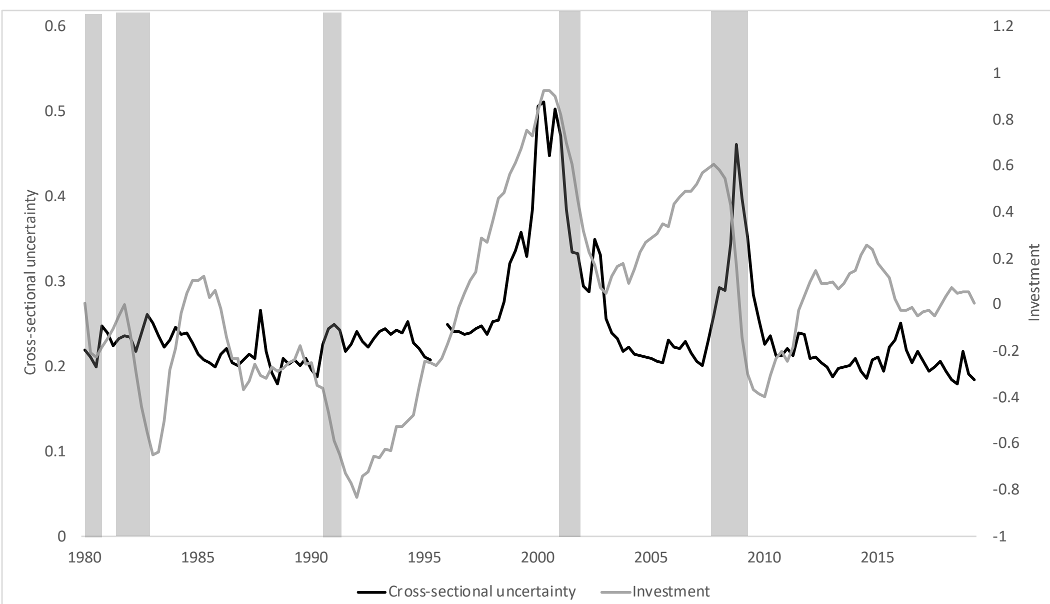

(b) Cross-sectional uncertainty and private nonresidential fixed investment

Note: The top panel plots the monthly unemployment rate and cross-sectional uncertainty. The bottom panel plots log quarterly real private nonresidential fixed investment minus a linear trend and the quarterly average of cross-sectional uncertainty. Shaded areas are NBER recessions.

Figure 3: Cross-sectional uncertainty across countries

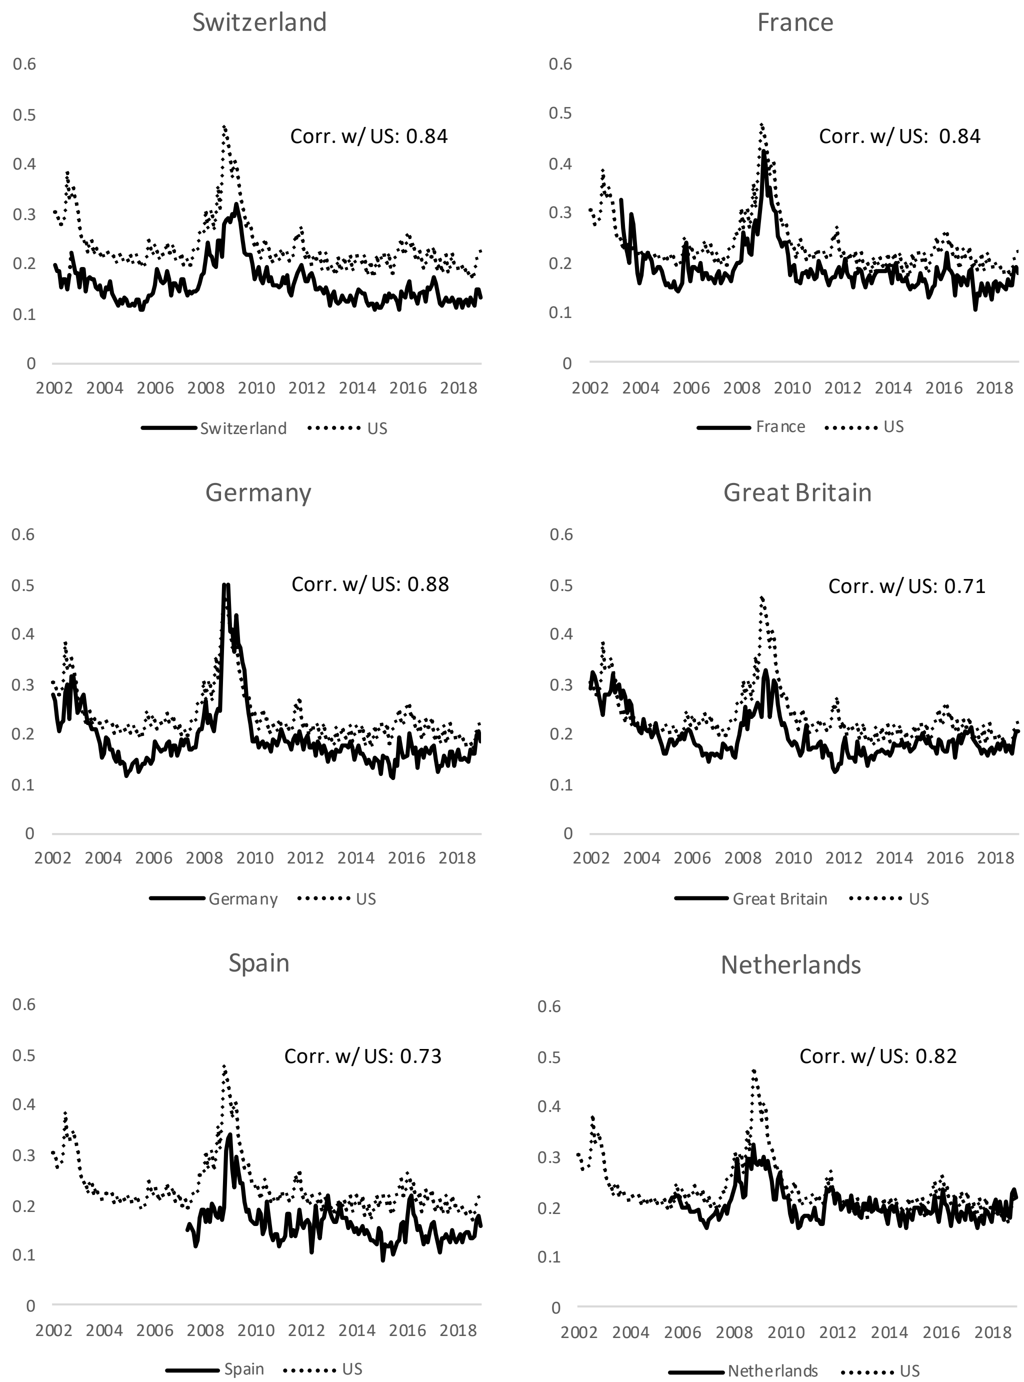

Note: Cross-sectional uncertainty from option data in European markets

Figure 4: Aggregate uncertainty across countries

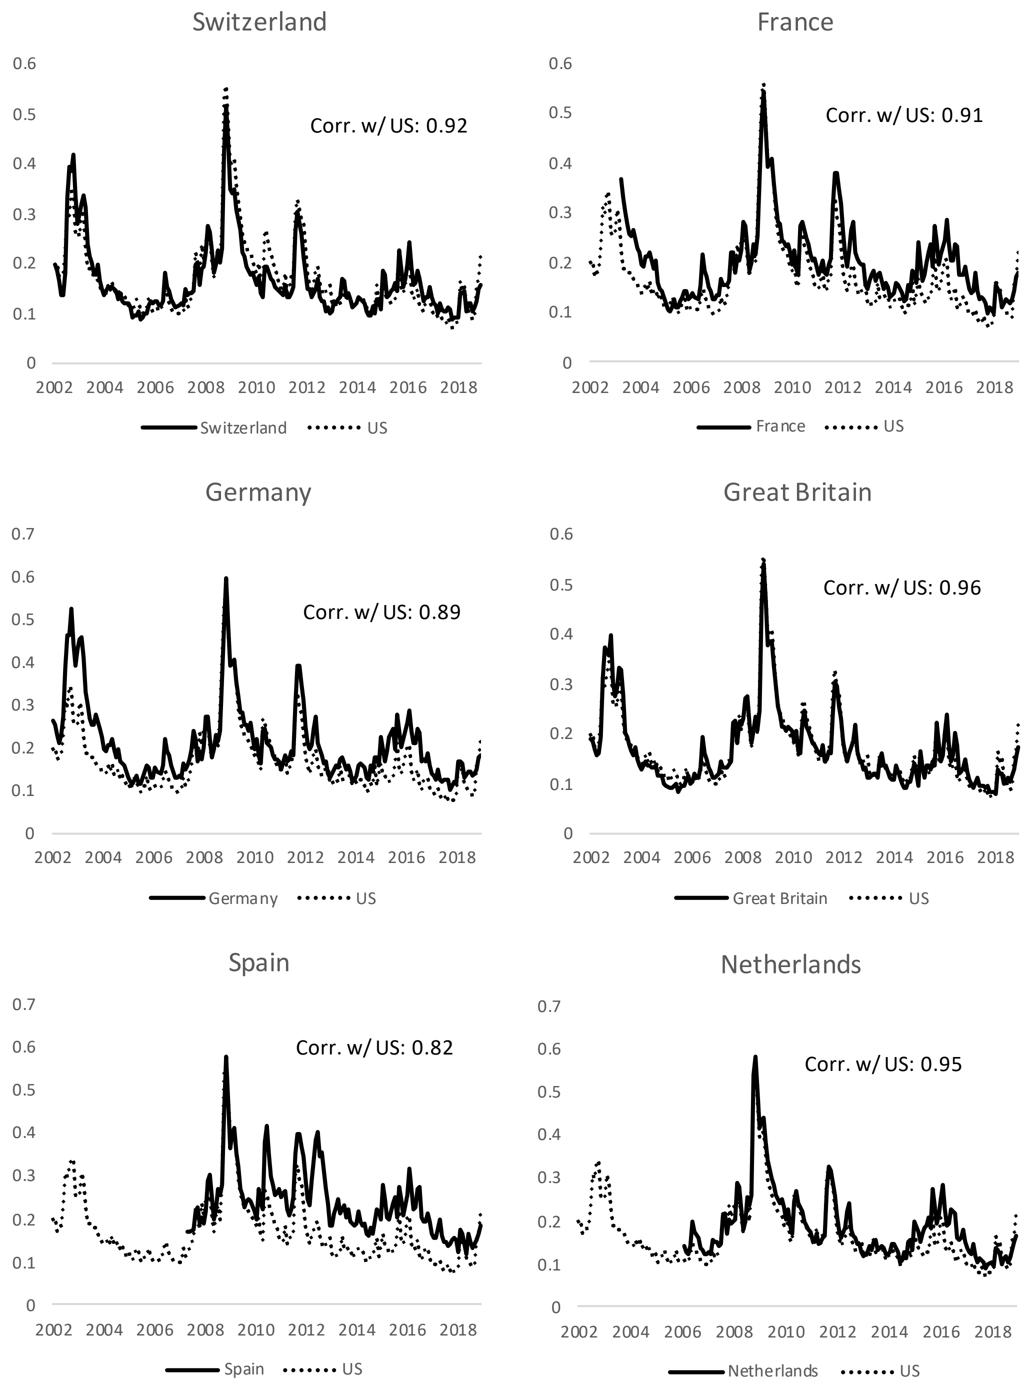

Note: Aggregate (market index) uncertainty from option data in European markets

Table 1: Raw correlations with cross-sectional uncertainty

	Cross-sectional uncertainty			Cross-sectional	Interquartile	Interquartile
	Full sample	Pre-1998	Post-1997	realized volatility	range, ΔIP	range, $\Delta Sales$
Detrended IP	-0.02	-0.39	0.00	-0.01	-0.23	-0.41
Detrended empl.	-0.21	0.49	-0.31	-0.22	-0.02	0.10
Unemployment rate	0.23	-0.16	0.26	0.24	-0.23	-0.11
CBO output gap	0.20	-0.14	0.26	0.22	-0.31	-0.08
NBER recession	-0.36	-0.18	-0.40	-0.36	-0.42	-0.37
Capacity utilitzation	-0.11	-0.05	0.08	0.02	-0.42	-0.20
GZ bond spread	-0.60	0.29	-0.59	-0.50	-0.60	-0.49
IP growth	-0.22	0.11	-0.23	-0.20	-0.32	-0.12
Employment growth	-0.35	-0.07	-0.36	-0.28	-0.57	-0.42
Change in unempl. rate	-0.26	-0.01	-0.33	-0.26	-0.31	-0.31
Change in output gap	-0.23	0.12	-0.29	-0.24	-0.18	-0.17

Note: Each row reports the raw correlation of cross-sectional uncertainty with a different outcome. The first three columns use cross-sectional uncertainty (for different time periods). The fourth column replaces cross-sectional uncertainty with cross-sectional realized volatility. The last two columns use interquartile range of IP growth and sales growth across NAICS industries and COMPUSTAT firms, respectively (see Bloom et al. (2018)).

Table 2: Forecasting regressions: monthly

Panel A: Unemployment	(1)	(2)	(3)	(4)	(5)
Cross-sectional unc. (t-1)	0.151***		0.106	0.0421	0.130
	(0.0467)		(0.129)	(0.0554)	(0.129)
Cross-sectional unc. (t-1), before 1998		-0.215			
		(0.228)			
Cross-sectional unc. (t-1), after 1998		0.172***			
		(0.0491)			
Cross-sectional RV (t-1)			0.0472		-0.120
			(0.125)		(0.132)
Market uncertainty (t-1)				0.205***	0.121
				(0.0579)	(0.109)
Market RV (t-1)					0.121
					(0.103)
p-value difference		0.0434			

Panel B: Employment	(1)	(2)	(3)	(4)	(5)
Cross-sectional unc. (t-1)	-0.0899***		-0.00970	-0.0523	-0.0483
	(0.0318)		(0.0976)	(0.0371)	(0.0977)
Cross-sectional unc. (t-1), before 1998		0.277*			
		(0.157)			
Cross-sectional unc. (t-1), after 1998		-0.116***			
		(0.0337)			
Cross-sectional RV (t-1)			-0.0798		0.0253
			(0.0917)		(0.0983)
Market uncertainty (t-1)				-0.0798*	0.110
				(0.0408)	(0.0785)
Market RV (t-1)					-0.222***
					(0.0779)
p-value difference		0.0155			

Panel C: Industrial production	(1)	(2)	(3)	(4)	(5)
Cross-sectional unc. (t-1)	-0.0945**		0.0219	-0.0208	-0.0486
	(0.0428)		(0.138)	(0.0513)	(0.137)
Cross-sectional unc. (t-1), before 1998		0.354			
		(0.222)			
Cross-sectional unc. (t-1), after 1998		-0.128***			
		(0.0464)			
Cross-sectional RV (t-1)			-0.116		0.0736
			(0.131)		(0.138)
Market uncertainty (t-1)				-0.133**	0.153
				(0.0519)	(0.103)
Market RV (t-1)					-0.350***
					(0.108)
p-value difference		0.0373			

Note: Forecasting regressions of unemployment (panel A), employment (panel B), and industrial production (panel C) using lagged uncertainty and realized volatility measures. All regressors are 3-month moving averages of the underlying variables. The last row reports is the p-value for the difference of the two coefficients of specification (2).

Table 3: Forecasting regressions: quarterly

Panel A: Output	(1)	(2)	(3)	(4)
Cross-sectional uncertainty (t-1)	-0.00728	-0.00547	0.00830	0.00682
	(0.00684)	(0.00841)	(0.0235)	(0.0237)
Market uncertainty (t-1)		-0.00344		0.0214
		(0.00925)		(0.0198)
Cross-sectional RV (t-1)			-0.0118	-0.00744
			(0.0171)	(0.0184)
Market RV (t-1)				-0.0232
				(0.0174)
N	142	142	142	142

Panel B: Consumption	(1)	(2)	(3)	(4)
Cross-sectional uncertainty (t-1)	-0.00343	-0.00749	-0.0117	-0.0102
	(0.00563)	(0.00728)	(0.0199)	(0.0202)
Market uncertainty (t-1)		0.00703		0.0135
		(0.00800)		(0.0172)
Cross-sectional RV (t-1)			0.00624	0.00314
			(0.0145)	(0.0158)
Market RV (t-1)				-0.00706
				(0.0154)
N	142	142	142	142

Panel C: Investment	(1)	(2)	(3)	(4)
Cross-sectional uncertainty (t-1)	-0.0684*	-0.0182	0.200	0.172
	(0.0382)	(0.0461)	(0.128)	(0.124)
Market uncertainty (t-1)		-0.0944*		0.221**
		(0.0495)		(0.105)
Cross-sectional RV (t-1)			-0.203**	-0.124
			(0.0927)	(0.0947)
Market RV (t-1)				-0.283***
				(0.0905)
N	142	142	142	142

Panel D: Hours	(1)	(2)	(3)	(4)
Cross-sectional uncertainty (t-1)	-0.0167**	-0.0108	-0.00323	-0.00492
	(0.00644)	(0.00746)	(0.0201)	(0.0202)
Market uncertainty (t-1)		-0.0128		0.00901
		(0.00825)		(0.0173)
Cross-sectional RV (t-1)			-0.0101	-0.00275
			(0.0142)	(0.0150)
Market RV (t-1)				-0.0202
				(0.0143)
N	142	142	142	142

Note: Quarterly forecasting regressions of output (panel A), consumption (panel B), investment (panel C), and hours (panel D) using lagged uncertainty and realized volatility measures.

Table 4: CMR model vs. the data

(a) Cyclical correlations with cross-sectional uncertainty

	Growth rates		Levels	
	Model	Data	Model	Data
Investment	-0.43	-0.05	-0.46	0.24
Consumption	-0.05	-0.07	-0.21	0.26
GDP	-0.39	-0.04	-0.70	0.24
Hours	-0.21	0.08	-0.51	0.26

(b) Regression analysis

	Model		Data	
Investment				
ΔCross-sectional uncertainty	-0.236	-0.235	-0.057	0.067
			[0.053]	[0.065]
ΔCross-sectional RV		0.112		-0.211***
				[0.068]
Consumption				
ΔCross-sectional uncertainty	-0.037	-0.04	-0.094*	-0.083
			[0.048]	[0.062]
ΔCross-sectional RV		0.033		-0.017
				[0.064]
GDP				
ΔCross-sectional uncertainty	-0.235	-0.235	-0.07	0.006
			[0.051]	[0.062]
ΔCross-sectional RV		0.03		-0.139**
				[0.067]
Hours				
ΔCross-sectional uncertainty	-0.027	-0.028	-0.03	-0.006
			[0.046]	[0.058]
ΔCross-sectional RV		0.013		-0.043
				[0.061]

Note: Panel A of the table reports correlations between cross-sectional uncertainty and various macroeconomic series, using linearly detrended series; the table reports correlations in the CMR model and in the data, and reports them using both growth rates (left side) and levels (right side). Panel B reports the results of regressions of various macroeconomic variables onto first differences in cross-sectional uncertainty and realized volatility. Like Panel A, Panel B reports the results of the regressions in the model and in the data. For both model and data, the left column uses only cross-sectional uncertainty as regressor; the right column includes also cross-sectional RV.

Table 5: RUBC model vs. the data

(a) Cyclical correlations with cross-sectional uncertainty

	Growth rates		Levels	
	Model	Data	Model	Data
Investment	-0.50	-0.05	-0.39	0.24
Consumption	0.20	-0.07	-0.16	0.26
GDP	-0.52	-0.04	-0.32	0.24
Hours	-0.72	0.08	-0.30	0.26

(b) Regression analysis

	Model		Data	
Investment				
ΔCross-sectional uncertainty	-0.51	-0.5	-0.06	-0.01
			[0.05]	[0.06]
ΔAggregate uncertainty		-0.01		-0.09
				[0.06]
Consumption				
ΔCross-sectional uncertainty	0.2	0.17	-0.11**	-0.13**
			[0.05]	[0.06]
ΔAggregate uncertainty		0.03		0.03
				[0.06]
GDP				
ΔCross-sectional uncertainty	-0.52	-0.54	-0.07	-0.04
			[0.05]	[0.06]
ΔAggregate uncertainty		0.01		-0.05
				[0.06]
Hours				
ΔCross-sectional uncertainty	-0.73	-0.74	-0.01	0.02
			[0.05]	[0.05]
ΔAggregate uncertainty		0.01		-0.05
				[0.05]

Note: Panel A of the table reports correlations between cross-sectional uncertainty and various macroeconomic series, using linearly detrended series; the table reports correlations in the RUBC model and in the data, and reports them using both growth rates (left side) and levels (right side). Panel B reports the results of regressions of various macroeconomic variables onto first differences in cross-sectional and aggregate uncertainty. Like Panel A, Panel B reports the results of the regressions in the model and in the data. For both model and data, the left column uses only cross-sectional uncertainty as regressor; the right column also includes aggregate uncertainty.

Table 6: Sharpe ratios on the *iv* and *rv* portfolios

	iv portfolio	*rv portfolio*
Full sample		
Cross-sectional	0.82	-0.48
	[0.17]	[0.17]
Market	0.15	-1.00
	[0.17]	[0.17]
Pre-1998		
Cross-sectional	0.86	-0.76
	[0.25]	[0.25]
Market	0.72	-1.19
	[0.25]	[0.25]
Post-1997		
Cross-sectional	0.80	-0.33
	[0.22]	[0.22]
Market	-0.1	-0.97
	[0.22]	[0.22]

Note: The table reports annualized Sharpe ratios on the *iv* and *rv* portfolios, hedging respectively uncertainty and realized volatility. The table reports two versions of each portfolio, one hedging cross-sectional uncertainty or realized volatility, and one hedging market uncertainty or realized volatility. The top panel uses all available data, the bottom panels split the sample in 1998.

A.1 Data calculations

A.1.1 Constructing implied volatility

For the Optionmetrics sample, we obtain at-the-money implied volatilities as the delta=50 IVs with maturity of 30 days from the Optionmetrics surface file.

For the BODB, the steps are as follows:

1. We calculate closing bid and ask prices for each option as the average of the final value and any other values recorded in the last 15 minutes of trading.

2. For each date/maturity/ticker combination, we take the strike immediately above and below the underlying price, as long as it is within 20 percent of the underlying.

3. Option prices are calculated as the midpoint between the bid and ask.

4. We drop all options with maturity less than 7 days.

5. The BODB reports a spot price. We replace the spot price with the value implied by put-call parity with a dividend of zero if the put-call parity implied price differs from the reported spot by more than 20 percent (this is to eliminate some clear data errors).

6. Implied volatilities are constructed using the Black–Scholes formula for European options ignoring dividends. For the one-month maturity, early exercise has generally very small effects on prices. We experimented by using the same method on data from Optionmetrics and comparing it to the implied volatilities that they report (which use a model for dividends and also account for early exercise) and we found the differences were quantitatively small.

7. We interpolate between maturities – and extrapolate where necessary – to get 30-day implied volatilities. Firm-level implied volatilities are set to have a maximum of 200 percent annualized and a minimum of zero (the interpolated values are Winsorized).

8. The implied volatilities are then collapsed across firms weighting by market capitalization. We matched the tickers in the BODB to CRSP permco numbers to get market capitalization. In the large majority of cases, the BODB tickers are the same as the stock exchange tickers (they differ most for NASDAQ listings; the BODB manual, available online or on request from us, discusses this issue). The remainder are matched by hand where possible.

A.1.2 Calculating returns

For Optionmetrics, we directly use the data on closing bid and ask prices for options. We use the set of firms that was ever in the top 200 sorted by size during the Optionmetrics sample (1996–2017). For the BODB, the closing bid and ask are constructed as discussed above and we use the entire available sample (which is tilted towards large firms). From there, the construction of returns is the same.

1. We drop all observations where the bid/ask spread is greater than 20 percent. We also apply step 5 above to the BODB data.

2. The two-week return is calculated by looking forward 10 trading days.

3. Return observations are dropped if at initiation the bid or ask price is less than 10 cents, the bid or ask volume is zero, or the maturity is less than 21 calendar days (the price filters are applied only at initiation so as not to introduce look-ahead bias).

4. We take the straddle with strike immediately above and below the spot, as long as they are within 20 percent of the spot and interpolate by log strike to construct an approximately at-the-money straddle. This requires having a valid straddle both above and below the spot, which often is not available.

5. Returns are collapsed across firms weighting by market capitalization at initiation. Returns are then interpolated and extrapolated to monthly maturities. For the one-month maturity, there must be a straddle available with a maturity of less than 60 days, and for the five-month maturity there must be available a straddle with maturity of at least 120 days (that is, we do not extrapolate by more than 30 days).

The S&P 500 straddle returns are constructed in a similar manner on data from the CME for the S&P 500 futures options and on Optionmetrics for SPX index options. In the BODB sample, the CME option returns are used to represent the market return, while the CBOE SPX options are used in with the Optionmetrics data.

A.2 Additional results

A.2.1 Alternative methods of calculating implied volatility

Figure A.1 plots four versions of the idiosyncratic implied volatility time series. The line labeled "mean" uses the unweighted mean of implied volatility across stocks, instead of weighting by market capitalization. We remove firm fixed effects to control for changes in composition over time and shift the level of the time series to that its mean is the same as in the baseline case. The line labeled "median" is the same, but uses the median across stocks instead of the mean. In both cases, the set of options used in the BODB sample is somewhat different from in the baseline case since for these two lines we can use the full sample instead of just those for which we have market capitalization data. Finally, the line labeled "beta" uses the formula $\sum_i w_{i,t}\sigma_{i,t}^2 - \left(\sum_i w_{i,t}\beta_i^2\right)\sigma_{mkt,t}^2$ for idiosyncratic volatility, where β_i is measured using the full sample of data for firm i. This is done separately for the

BODB and Optionmetrics because the stock prices have different sources – for the BODB we use CRSP, while for Optionmetrics we use the prices reported by Optionmetrics. In both cases, we only keep β_i for stocks with at least 60 months of data. β_i is calculated using monthly data with the market return obtained from Kenneth French's website.

A.2.2 Factor pricing model for returns

We estimate a standard factor model using the rv and iv portfolios as test assets. Specifically, the model's two equations are

$$r_t = \bar{r} + \beta_{rv,idio} RV_{j,t} + \beta_{iv,j} IV_{j,t} + \varepsilon_t \tag{A.1}$$

$$E[r_t] = \lambda_{rv,j} \beta_{rv,j} std(RV_{j,t}) + \lambda_{iv,idio} \beta_{iv,j} std(IV_{j,t}) + \alpha \tag{A.2}$$

for $j \in \{idio, mkt\}$, where

$$RV_{idio,t} \equiv \sum_i w_{i,t-1} r_{i,t}^2 - r_{mkt,t}^2 \qquad RV_{mkt,t} \equiv r_{mkt,t}^2$$
$$IV_{idio,t} \equiv \sum_i w_{i,t-1} \Delta\sigma_{idio,i,t}^2 \qquad IV_{mkt,t} \equiv \Delta\sigma_{mkt,t}^2$$

The first equation is a time-series model for returns, while the second equation measures the relationship between risk premia and factor exposures. Intuitively, the λ parameters measure the Sharpe ratio for a portfolio that loads with a coefficient of 1 on the associated factor and 0 on the others (multiplying the betas by the factor standard deviations ensures that the λ parameters represent Sharpe ratios). $\lambda_{iv,idio}$, for example, represents the Sharpe ratio earned for exposure to changes in idiosyncratic implied volatility. It is therefore our main parameter of interest, and the analysis above implies it should take a similar value to the mean return on $iv_{idio,t}$.

That method intuitively involves using all of the rv and iv portfolios as potential hedges for shocks to σ_{idio}^2, so that the R_X in (9) is then a combination of the rv and iv portfolios, which reduces the magnitude of the ε. The factor model does not require the approximation or any of the assumptions in the previous subsection. Furthermore, note that the factor model does not even really require using the rv and iv portfolios. Since they are just linear combinations of straddles, the results would be unchanged if the factor model was estimated using the original straddle returns (with the scaling by implied volatility) instead of the rv and iv portfolios.

The table below reports risk premia estimated from the factor model. The signs of the λ parameters – which, again, represent Sharpe ratios for exposures to the realized and implied volatility factors – are the same as for the average returns themselves.

Factor model estimates

$\lambda_{iv,idio}$	4.60	$\lambda_{iv,mkt}$	0.93
	[0.71]		[0.29]
$\lambda_{rv,idio}$	-2.53	$\lambda_{rv,mkt}$	-1.48
	[0.61]		[0.36]

The Sharpe ratios estimated in the factor model are more extreme than those reported for the rv and iv portfolios. The reason for that difference is that the R^2s in the first stage regression of the returns on the factors are much less than 1. Since the factor model assumes the fitting errors are unpriced, it implies that the excess return in those portfolios is earned for only a fraction of their total risk. We view the magnitudes of the Sharpe ratios in the factor model as economically fairly implausible, which is why the main text focuses on the actual tradable returns. The implied Sharpe ratios in the factor model are not anything that an investor can literally earn since the pricing factors are untraded.

Figure A.1: Fraction of market capitalization covered by the options data

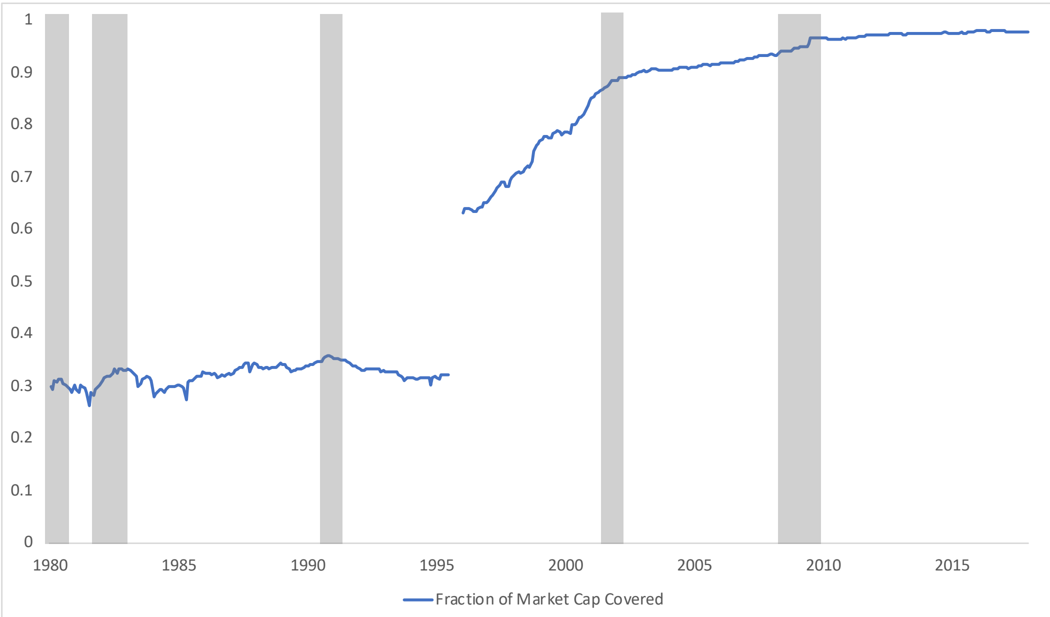

Note: The figure reports the ratio of total market capitalization for the firms for which we observe options data to the total market capitalization. Data before 1996 is from the Berkeley Options Dataset, and data after 1996 is from Optionmetrics. Shaded areas are NBER recessions.

Figure A.2: Average firm realized variance for different sets of firms

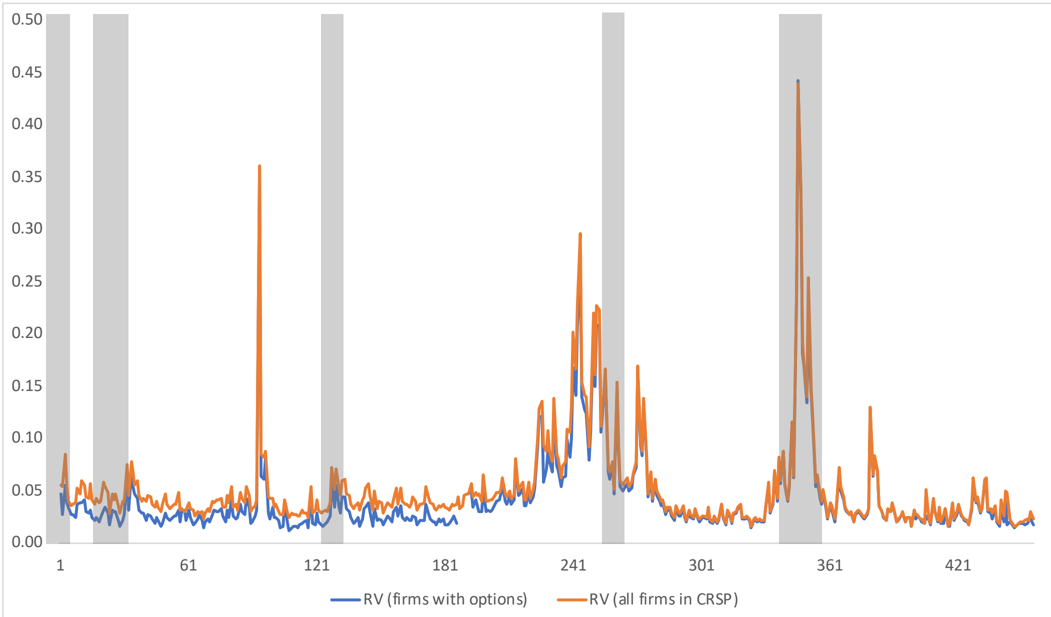

Note: The figure reports cross-sectional average realized variance for all firms in CRSP and for the firms for which we observe options. Option data before 1996 is from the Berkeley Options Dataset, and data after 1996 is from Optionmetrics. Shaded areas are NBER recessions.

Figure A.3: Robustness

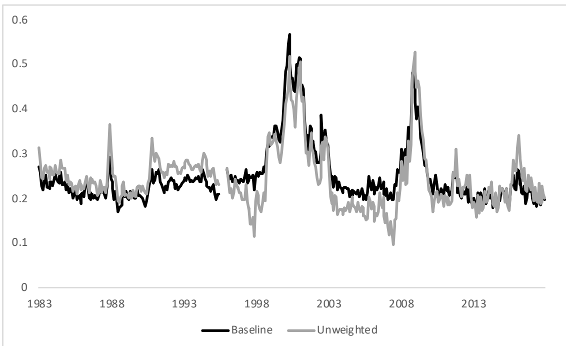

(a) Unweighted mean uncertainty

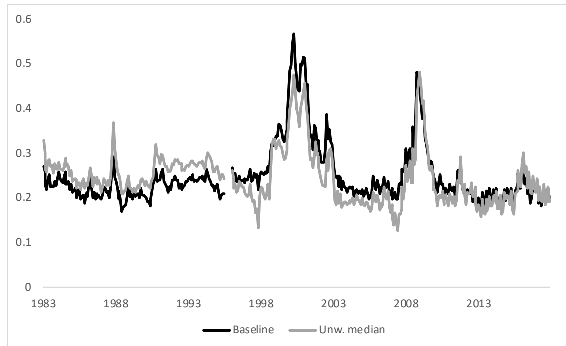

(b) Unweighted median uncertainty

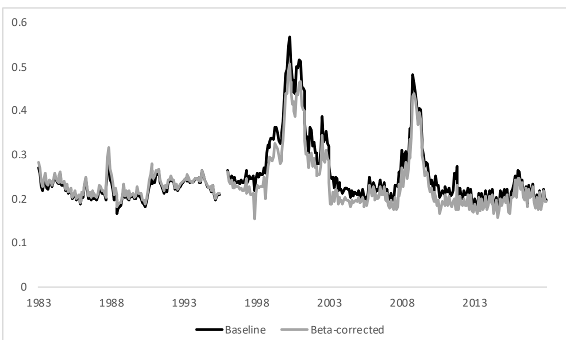

(c) Beta-corrected uncertainty

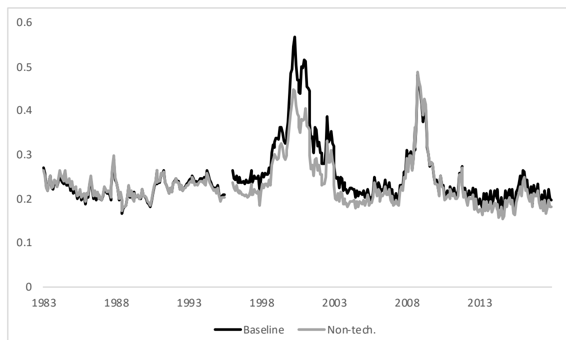

(d) Uncertainty for non-tech firms

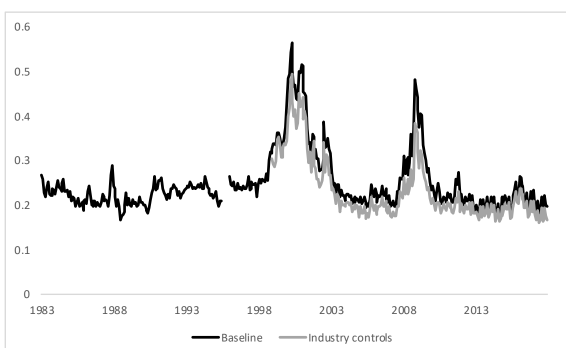

(e) Uncertainty controlling for industry shocks

Note: Each panel plots the baseline measure of cross-sectional uncertainty along with an alternative. Panels (a) and (b) weight all firms equally (i.e. ignoring market values) and use either the cross-sectional mean or median of firm implied volatility. Panel (c) uses the formula from the text that corrects for the cross-sectional variance of beta. Panel (d) drops all tech firms – with three-digit SIC code 357 or 737 or trading on the NASDAQ. Panel (e) reports residual uncertainty after controlling for the industry component (instead of the aggregate component).

Table A.1: Arch-type regressions

Uncertainty (dep. var.):	Cross-sec.	Cross-sec.	Mkt.	Mkt.	ΔCross-sec.	ΔMkt.
Cross-sec. unc. (t-1)	0.95***	0.59***				
	(0.02)	(0.02)				
Cross-sec. RV (t)		0.29***				
		(0.01)				
Mkt. unc. (t-1)			0.90***	0.57***		
			(0.02)	(0.02)		
Mkt. RV (t-1)				0.37***		
				(0.02)		
Mkt. RV (t)					0.056***	0.159***
					(0.011)	(0.015)
R^2	0.904	0.953	0.80	0.94	0.05	0.211

Note: Coefficients of monthly regressions of cross-sectional and market uncertainty onto contemporaneous and lagged values of cross-sectional and market uncertainty and RV.