

NBER WORKING PAPER SERIES

CROSS-SECTIONAL UNCERTAINTY AND THE BUSINESS CYCLE:
EVIDENCE FROM 40 YEARS OF OPTIONS DATA

Ian Dew-Becker
Stefano Giglio

Working Paper 27864
<http://www.nber.org/papers/w27864>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2020

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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Ian Dew-Becker and Stefano Giglio. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Cross-Sectional Uncertainty and the Business Cycle: Evidence from 40 Years of Options Data
Ian Dew-Becker and Stefano Giglio
NBER Working Paper No. 27864
September 2020
JEL No. C58,D81,D84,E22,E30,E32,E37,G13

ABSTRACT

This paper presents a novel and unique measure of cross-sectional uncertainty constructed from stock options on individual firms. Cross-sectional uncertainty varied little between 1980 and 1995, and subsequently had three distinct peaks -- during the tech boom, the financial crisis, and the coronavirus epidemic. Cross-sectional uncertainty has had a mixed relationship with overall economic activity, and aggregate uncertainty is much more powerful for forecasting aggregate growth. The data and moments can be used to calibrate and test structural models of the effects of uncertainty shocks. In international data, we find similar dynamics and a strong common factor in cross-sectional uncertainty. The data is available on our websites. A companion paper [Dew-Becker and Giglio, "Real-time forward-looking skewness over the business cycle"] finds firm-level skewness is significantly procyclical.

Ian Dew-Becker
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
ian.dewbecker@gmail.com

Stefano Giglio
Yale School of Management
165 Whitney Avenue
New Haven, CT 06520
and NBER
stefano.giglio@yale.edu

A Data files is available at <http://www.dew-becker.org/#data>

1 Introduction

This paper reports a novel option-implied measure of cross-sectional uncertainty. Whereas the VIX, the most widely used option-implied uncertainty index, measures uncertainty about the state of the aggregate stock market (and, potentially, economy), we construct an index that tracks uncertainty about the cross-sectional distribution of firm outcomes. In many recent models and empirical analyses, it is precisely the cross-sectional component that is the critical driving force.¹

More formally, one might decompose the shock to a firm, $\eta_{i,t}$, into an aggregate component, μ_t , and an orthogonal component, $\varepsilon_{i,t}$ (which may be correlated across subsets of firms):

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

The *total uncertainty* a firm faces is measured by the conditional (time- t) variance of $\eta_{i,t+1}$. The VIX and other measures of *aggregate uncertainty* capture the conditional variance of μ_{t+1} . Finally, *cross-sectional uncertainty*, on which this paper focuses, is measured by the conditional variance of $\varepsilon_{i,t+1}$: it is the variance of the shocks faced by firms that are orthogonal to aggregate shocks. We measure cross-sectional uncertainty similarly to the VIX, using option-implied volatilities.

Our cross-sectional uncertainty measure is simple to construct: it is just average firm-level option-implied conditional variance minus market implied conditional variance ($\text{var}_t(\eta_{i,t+1}) - \text{var}_t(\mu_{t+1})$). 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. Because it is constructed from market prices, our measure is forward-looking, and is available continuously, in real time, making it particularly useful for policymakers. In addition, the measure is available for a long span of time (40 years), including six recessions. Past work has at most extended to 1996, observing only two business cycles.

In this paper, we document several empirical patterns on the relationship between our new measure of cross-sectional uncertainty and the economy. We focus on two types of patterns: the cyclical behavior of cross-sectional uncertainty and the forecasting power of cross-sectional uncertainty for future economic activity. We find that cross-sectional uncertainty has a mixed relationship with the state of the business cycle, rising during the tech boom of the late 1990's, but also during the financial crisis and coronavirus epidemic. The dark line in figure 1a plots cross-sectional uncertainty. From the start of our data, in 1980, up to 1995, there was surprisingly little variation. After 1995, firm-level uncertainty moves

¹Bewley (1986), Bloom (2009), Kaplan and Violante (2014), Christiano, Motto, and Rostagno (2016), and Ilut, Kehrig, and Schneider (2018).

much more (though still less than market uncertainty, in proportional terms), with three distinct increases, during the tech boom, the financial crisis, and the coronavirus epidemic. In the three episodes where uncertainty is elevated, it rapidly declines, returning to its long-run average by the trough of the recession. In a shorter sample, international data displays similar behavior and also has a very strong factor structure, implying that cross-sectional uncertainty is driven by global shocks.

Overall, the data appears to show that cross-sectional uncertainty is sometimes high in bad times, and sometimes high in good times. Two different classes of models exist that predict one or the other behavior for cross-sectional uncertainty, but not both. The financial crisis, with low activity and high uncertainty, is consistent with the models that emphasize countercyclical uncertainty, whether it is an endogenous response or an exogenous shock. Interestingly, though, if output tracked cross-sectional uncertainty over time, it would have recovered from the financial crisis by 2010 (when unemployment was still over 9 percent). In contrast to the financial crisis, the period of the late 1990's is consistent with models in which growth and innovation are associated with uncertainty, e.g. due to learning, creative destruction, or a risk/return trade-off in investment projects.²

Next, we examine the forecasting power of idiosyncratic uncertainty for aggregate output and employment, finding similarly mixed results. A key feature of the data is that it allows us to test whether aggregate or cross-sectional uncertainty is more relevant for forecasting, which represents a fruitful way to distinguish among classes of structural models and is also relevant for policymakers. We find strong evidence that it is aggregate rather than cross-sectional uncertainty that is most likely to be an important driver of the aggregate economy (though see Berger, Dew-Becker, and Giglio (2020) for questions about the effects of aggregate uncertainty).

We formally examine the cyclicity and forecasting regressions in two theoretical models of the macroeconomic effects of cross-sectional uncertainty shocks: Christiano, Motto, and Rostagno (2014) and Bloom et al. (2018). Both models predict that cross-sectional uncertainty should be clearly countercyclical and should be more tightly related to aggregate output than aggregate uncertainty or realized volatility, inconsistent with the data.

In addition to evaluating correlations and forecasts, we show that the data is also useful for giving a set of moments to aid in calibrating structural models. The data series, available on our websites, gives a direct measure of the underlying driving uncertainty process that needs to be parameterized, showing that many papers have used realistic amounts of variation in firm-specific uncertainty, while others require implausibly high quantities.

²See, for example, Acemoglu (2005), Imbs (2007), Comin and Mulani (2009), and Kogan et al. (2017).

A large literature has studied the relationship between uncertainty and the real economy. However this literature has either focused on aggregate uncertainty, or, if it has looked at individual firms, it has not used forward-looking measures of uncertainty (like ours), but backward-looking measures (realized volatility) that do not map into what uncertainty is in our models.³ This paper shows that that distinction changes the conclusions one draws from the data. This is the first work to deliver a long time-series of *forward-looking, cross-sectional* uncertainty. Only a few papers have similar forward-looking measures of firm-level uncertainty, primarily surveys, but in those cases it is not possible to disentangle the cross-sectional and aggregate components, whereas in the case of stock returns it is straightforward.⁴ This paper’s novelty is in developing an *ex ante* measure of idiosyncratic uncertainty that more directly maps into the shock processes driving structural models and has a long empirical sample.

2 Data

We obtain options price data from the Berkeley Options Database (BODB) for 1/1980–6/1995, Optionmetrics for 1/1996–12/2019, and Bloomberg for 1/2020–5/2020.

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, we use at-the-money Black–Scholes implied volatility. The latter requires only observing a single option price and is 99.5 percent correlated with the VIX.⁵ Since implied volatilities come from asset prices, they embed risk premia, meaning they are not errorless measures of investor beliefs. Nevertheless, they represent the most common measure of uncertainty studied in the literature.

Denote firm i ’s implied volatility in month t as $\sigma_{i,t}$ and implied volatility for the aggregate stock market as $\sigma_{mkt,t}$. One can always theoretically construct the linear projection of the return on stock i , $r_{i,t}$, on the market, $r_{mkt,t}$, as

$$r_{i,t} = \alpha_{i,t} + \beta_{i,t}r_{mkt,t} + \varepsilon_{i,t} \quad (2)$$

with $\varepsilon_{i,t} \perp r_{mkt,t}$ by construction. (2) is just a theoretical representation – it is not directly

³Specifically, Campbell et al. (2001), Bloom (2009), Herskovic et al. (2017), and Bloom et al. (2018) all examine measures of realized dispersion rather than conditional variances. Senga (2018) studies both realized volatility and total firm implied volatility (mixing aggregate and idiosyncratic components) since 1996.

⁴See Guiso and Parigi (1999), Ben-David et al. (2013), Bachmann, Elstner, and Sims (2013), and Bachmann et al. (2018).

⁵The model-free implied volatility requires a continuum of strikes, not usually available for individual stocks.

estimable since the parameters can change on every date, nor is it structural. We follow Campbell et al. (2001) in defining cross-sectional uncertainty as

$$\sigma_{\varepsilon,t}^2 = \sum_i w_{i,t} \sigma_{i,t}^2 - \sigma_{mkt,t}^2 \quad (3)$$

where $w_{i,t}$ are market capitalization weights; this equation is accurate when $\beta_{i,t} \approx 1$. We discuss robustness to this choice below.

Since $\varepsilon_{i,t}$ is only orthogonal to the market return, it can in general be correlated across firms, e.g. due to industry effects. Changes in the volatilities of cross-sectional factors will appear in $\sigma_{\varepsilon,t}^2$ so we refer to $\sigma_{\varepsilon,t}^2$ as *cross-sectional uncertainty*.

We measure $\sigma_{mkt,t}^2$ with S&P 500 option-implied volatility, and all volatilities are interpolated to a maturity of thirty days.

Figure A.1 in the Appendix plots 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 of those were listed on the CBOE. In 1996, when Optionmetrics becomes available, coverage jumps to 63 percent and then rises to 98 percent by the end of the sample.

To keep the sample consistent over time, our main results calculate cross-sectional uncertainty only for the 200 largest firms in the economy over the full sample; we show below that this choice is innocuous.⁶ Since we weight firms by market capitalization, and in any case only have data on public firms, our results necessarily apply to the largest firms in the economy. These firms account for a large fraction of total economic activity, though, and to the extent that idiosyncratic shocks affect the state of the economy, many theories imply it will be the largest firms whose shocks pass through to the aggregate economy (e.g. Gabaix (2011) and Acemoglu et al. (2012)).

3 Time-series behavior of cross-sectional uncertainty

This section reports the basic properties of cross-sectional uncertainty and examines its comovement with measures of real activity and financial stress.

⁶For BODB, tickers must be matched by hand to CRSP to obtain underlying stock prices. We did that only for the top 200 firms by size.

3.1 Univariate behavior and cyclicalities

3.1.1 Variability

Panel a of figure 1 plots the time series of cross-sectional uncertainty, $\sigma_{\varepsilon,t}$. In the first half of the sample, there is remarkably little variation: its standard deviation is only 9 percent of its mean for the period 1980–1997. But since 1997, it rose by nearly a factor of four to 37 percent of its mean.

$SD(\sigma)/E(\sigma)$:

	$\sigma_{mkt,t}$	$\sigma_{\varepsilon,t}$
Full sample:	0.40	0.29
1980–1997	0.30	0.09
1998–2020	0.43	0.37

Figure 1a also reports the implied volatility for the overall stock market. Relative to its mean, aggregate uncertainty is substantially more variable than cross-sectional uncertainty. The standard deviation of σ_{mkt} is 40 percent of its mean overall, compared to only 29 percent for cross-sectional volatility.

The variability of $\sigma_{mkt,t}$ is also much less isolated in time. Whereas the variation in cross-sectional uncertainty is driven primarily by just three episodes, there are numerous substantial jumps in market-level uncertainty, associated with the 1987 stock market crash, the first Gulf War, various events between 1998 and 2002, the debt ceiling, the Euro crisis, etc.

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

$$\begin{aligned} \text{var}\left(\sum_i w_{i,t} \sigma_{i,t}^2\right) &= \text{var}\left(\sigma_{mkt,t}^2\right) + \text{var}\left(\sigma_{\varepsilon,t}^2\right) + 2 \text{cov}\left(\sigma_{i,t}^2, \sigma_{mkt,t}^2\right) \\ 1.43 \times 10^{-2} &= \underbrace{0.45 \times 10^{-2}}_{32\%} + \underbrace{0.46 \times 10^{-2}}_{32\%} + \underbrace{0.52 \times 10^{-2}}_{36\%} \end{aligned}$$

Over the sample, the variation in the total uncertainty that firms face is equally driven by variation in aggregate and cross-sectional uncertainty.

The moments reported here on the volatilities of aggregate and cross-sectional uncertainty are useful for calibrating structural models of uncertainty shocks. We return to this point below.

3.1.2 Cyclicalities

Panel b of figure 1 plots cross-sectional uncertainty against the detrended level of the CRSP total stock market index. The periods of high cross-sectional uncertainty are all associated

with large changes in stock prices, but in different directions. During the dot-com boom, cross-sectional uncertainty tracks the rise of the stock market. They peak in almost exactly the same month, and uncertainty declines with the market. It follows the opposite pattern during the financial crisis: it is exactly when the stock market begins to decline that cross-sectional uncertainty rises. So uncertainty appears to be procyclical in the late 1990s and early 2000s, countercyclical in the financial crisis and in the recent coronavirus episode, and acyclical otherwise.

Figure 1d further emphasizes that point by plotting cross-sectional uncertainty against aggregate investment. Investment and uncertainty peak simultaneously in 2000, while uncertainty spikes and investment crashes in 2008. Figure 2a shows that uncertainty has a similarly mixed relationship with the unemployment rate.

To more formally quantify the cyclicity of cross-sectional uncertainty, panel a of table 1 reports the correlation of cross-sectional uncertainty with various measures of the state of the economy, over both the full sample and also pre- and post-1998. 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 correlations with detrended industrial production and capacity utilization are close to zero, and its correlation with credit spreads is positive, implying it is countercyclical. So relative to levels, it appears essentially acyclical. Furthermore, each of those correlations reverses sign between the first and second halves of the sample.

Relative to growth rates, uncertainty appears more consistently countercyclical, with negative correlations with all the cyclical measures. Those correlations also generally retain the same sign in the first and second halves of the sample.

Overall, to the extent that there is a clear cyclicity for cross-sectional uncertainty, it is that high uncertainty is associated with turning points. In levels, it has a weak relationship with the state of the economy. But relative to changes the correlations are much stronger. Cross-sectional uncertainty declines following its peaks in 2000 and 2008, emphasizing that it identifies change points.

3.2 Robustness and further results

In panels b–f, figure 2 plots five variations on the benchmark uncertainty series:

1. Using the 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 $\sigma_{\varepsilon,i,t}^2 = \sigma_{i,t}^2 - \beta_i^2 \sigma_{mkt,t}^2$.
3. Weighting by employment instead of market capitalization.
4. Using the full sample of options from Optionmetrics instead of just the largest 200 firms.
5. Controlling for industry effects by setting $\sigma_{\varepsilon,i,t}^2 = \sigma_{i,t}^2 - \sigma_{ind(i),t}^2$, where $\sigma_{ind(i),t}^2$ is the implied variance for the SPDR exchange traded fund covering stock i ’s sector.

The first test shows that the results are not driven just by the weighting by market capitalization – cross-sectional median uncertainty displays similar behavior. The main difference is a smaller increase in the late 1990’s.

The second test shows that the approximation where we treat the loadings on the market as all equal to 1, as in Campbell et al. (2001), has very little impact.

The third test – weighting by Compustat employment – is novel to this paper. Past work, in focusing on the S&P 500 index, puts weight on firms and their volatility based on their equity valuation. Differences in leverage and investor beliefs will then affect how uncertainty is aggregated, whereas employment weighting gives a more stable and fundamental measure. The results turn out to be highly similar. The largest effect is during the tech bubble, when tech stocks had high market capitalization relative to their share of the real economy and also relatively high implied volatilities.

Fourth, if we use the full Optionmetrics sample instead of just the 200 largest firms, the results are essentially unchanged, with just a slight level shift in the second half of the sample. Finally, removing the industry component of volatility also does not qualitatively change the conclusions. Its main effect is to dampen the cross-sectional uncertainty in the tech boom and financial crisis, consistent with those episodes being at least partially driven by sector shocks.

Appendix A.2 discusses the commonality in variation in uncertainty across firms. Consistent with Herskovic et al. (2017), 40–50 percent of cross-sectional variation in firm-specific uncertainty is captured by the common component.

Finally, in results available on request, we have also constructed firm-specific uncertainty using the so-called model-free implied volatility and obtain nearly identical results.

3.3 Uncertainty versus realized dispersion

Prior to this paper, the available data on forward-looking cross-sectional uncertainty covered at most only two recessions, making it uninformative about cyclicalit. For that reason, past studies have used realized dispersion – the cross-sectional standard deviation of the realizations of $\varepsilon_{i,t}$ – to proxy for $\sigma_{\varepsilon,t}^2$ and extend the sample (e.g. Bloom (2009)). The difference between the two is not innocuous.

Figure 1c plots $\sigma_{\varepsilon,t}^2$ against the realized cross-sectional standard deviation of the firm-specific residuals, $\varepsilon_{i,t}$. The realized standard deviation behaves substantially differently from the conditional standard deviation, $\sigma_{\varepsilon,t}$. It appears to have substantial high-frequency “noise”.

Recall that $\sigma_{\varepsilon,t-1}^2$ is the conditional expectation of the realized dispersion in returns, defined as

$$RD_t^{ret} = \sum_i w_{i,t} \text{var}(r_{i,t}) - \text{var}(r_{mkt,t}) \quad (4)$$

where $\text{var}(r_{i,t})$ and $\text{var}(r_{mkt,t})$ are calculated from *realized* returns within each month. Therefore,

$$RD_t^{ret} = \sigma_{\varepsilon,t-1} + \eta_t \quad (5)$$

where η_t is a mean-zero shock uncorrelated with $\sigma_{\varepsilon,t-1}$.

Equation (5) means that if one’s goal is to understand the behavior of *uncertainty* (the forward-looking $\sigma_{\varepsilon,t-1}$) – its variability, correlation with other variables, or its coefficient in forecasting regressions – then proxying for it with RD_t^{ret} will cause biases. The volatility of RD_t^{ret} is substantially higher than that of $\sigma_{\varepsilon,t-1}$, its correlation with other variables is lower, and in regressions there will be an attenuation bias even if η_t is exogenous. If η_t is correlated with outcomes of interest, that will further bias any regressions. Berger, Dew-Becker, and Giglio (2020), for example, show that when structural productivity shocks are skewed left (consistent with observed asymmetry in the business cycle), then realized volatility, η_t , will be negatively correlated with output, even if there is no structural effect of uncertainty on activity.

Furthermore, for a policymaker working in real time, the relative precision of $\sigma_{\varepsilon,t}$ is an added advantage. If one’s goal is to measure uncertainty in real time, the added noise in realized dispersion makes it less useful than the true uncertainty $\sigma_{\varepsilon,t}$.

4 Forecasting

This section examines the ability of cross-sectional uncertainty to forecast both future realized cross-sectional dispersion and also aggregate real activity.

4.1 Realized dispersion

The first question one must ask about the forecasting power of option-implied uncertainty is whether it forecasts future realized dispersion, as predicted by equation (5). Does it actually measure uncertainty?

Panel b of table 1 reports results of regressions of quarterly realized dispersion on the lag of $\sigma_{\varepsilon,t}$. In all cases, both the dependent and independent variables are standardized to have unit standard deviations. The first two columns show that $\sigma_{\varepsilon,t}$ has substantial forecasting power for RD_{t+1}^{ret} , in fact driving the lagged value of realized dispersion out of the regression. A unit standard deviation increase in cross-sectional uncertainty is associated with future cross-sectional uncertainty higher by 0.90 standard deviations, indicating that cross-sectional uncertainty is close to an unbiased predictor of changes in realized dispersion over time (the result is similar in absolute levels instead of standard deviations).⁷

The third and fourth columns of table 1b report results for two alternative measures of cross-sectional dispersion: the cross-sectional interquartile ranges of growth in industrial production (across sectors) and growth in sales (across Compustat firms), RD^{IP} and RD^{sales} . In both cases, $\sigma_{\varepsilon,t}$ again has substantial forecasting power. A unit standard deviation increase in uncertainty predicts about a 0.2 standard deviation increase in future dispersion. The relatively smaller magnitude is not surprising since $\sigma_{\varepsilon,t}$ measures uncertainty of stock returns, rather than IP or sales (and in this case, in results not reported here, $\sigma_{\varepsilon,t}$ is driven out of the regression by the lagged dependent variable).

4.2 Real activity

We study forecast of three monthly variables: the unemployment rate, non-farm private employment growth, and industrial production growth. All are again standardized to have unit variance. These regressions are valuable for two reasons, both independent interest and also as moments that can be used to test models.

The first column in the three sections of table 1c reports a regression of activity on lagged cross-sectional uncertainty. In all three cases, the coefficient implies that increases in

⁷The constant in the regression is not equal to zero, consistent with the presence of a risk premium.

uncertainty are followed by declines 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 about 0.14 standard deviations and an increase in unemployment of 0.09 standard deviations.

Sharper tests of theoretical models can be obtained by contrasting the forecasting power of cross-sectional uncertainty with that of other measures, like realized cross-sectional dispersion and aggregate uncertainty. Empirically, these measures are potentially correlated, but the existing literature has often offered sharp predictions about which type of volatility matters for the real economy.

Studying the effects of *aggregate* uncertainty, Berger, Dew-Becker, and Giglio (2020) show that uncertainty is driven out by realized volatility (the aggregate analog to realized dispersion). The second column of each section of table 1c shows that when realized dispersion is included in the forecasting regressions, it is associated with declines in both employment and industrial production and increases in unemployment, with larger coefficients than cross-sectional uncertainty.

Those results, together with those in the previous section, tell us that even though forward-looking cross-sectional uncertainty $\sigma_{\varepsilon,t}$ is a better measure of uncertainty than realized dispersion (lagged RD_t^{ret} is driven out by $\sigma_{\varepsilon,t}$ when predicting itself), realized dispersion has more forecasting power for real activity. Unanticipated shocks to realized dispersion – what appears in figure 1c to be just noise in RD_t^{ret} relative to $\sigma_{\varepsilon,t}$ – is actually more important for explaining economic fluctuations.

Similarly, table 1c also asks whether cross-sectional or aggregate uncertainty is most important for forecasting. In all three cases, market uncertainty drives cross-sectional uncertainty out of the regressions. Not only is market uncertainty dominant in relative terms, its coefficients are also very large in absolute terms, ranging in magnitude from 0.19 to 0.35, or twice as large as the coefficients on $\sigma_{\varepsilon,t}$ in the first column. The data therefore shows that to the extent that uncertainty is relevant for forecasting real activity, it is aggregate rather than cross-sectional uncertainty that matters.

Overall, while cross-sectional uncertainty does have some univariate forecasting power, it is delicate. Both realized dispersion and aggregate uncertainty are significantly stronger.

4.3 Summary

Section 3 showed that over the last 40 years, cross-sectional uncertainty has been relatively stable outside of three distinct episodes. Of those three peaks, two came during recessions and one during an expansion, showing that cross-sectional uncertainty has had a mixed rela-

tionship with the business cycle. This section takes that result a step further, showing that cross-sectional uncertainty overall has weak forecasting power for real activity. The results are difficult to reconcile with the view that firm-specific uncertainty is a major economic headwind.

Two alternative variables, realized dispersion and aggregate uncertainty, drive cross-sectional uncertainty out of the forecasting regressions. Between aggregate and cross-sectional uncertainty, aggregate uncertainty appears to be the relevant driver. Furthermore, to the extent that the cross-section matters, it is through the realization of shocks that generate dispersion or reallocation, not the expectation that such shocks might occur (even though realized dispersion, to the naked eye, simply appears to equal ex ante uncertainty plus significant noise).

The next section shows how these results can be used to calibrate and test structural models.

5 Calibrating and testing structural models

We now use $\sigma_{\varepsilon,t}$ to calibrate and testing structural models. We first estimate the time-series dynamics of uncertainty, which is a key input to calibrations. The regressions from table 1c are a useful test of structural models. Since most models are constructed at the quarterly frequency, the empirical results reported in this section are also estimated at the quarterly frequency.

We begin in section 5.1 by showing how moments from our dataset can be used to calibrate models. Sections 5.2 and 5.3 then focus on testing two specific models: the “really uncertain 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 Calibration moments

In the vast majority of models with time-varying uncertainty, uncertainty follows an AR(1) type process, which can be characterized by its standard deviation (which we scale relative to its mean) and autocorrelation. Panel a of table 2 reports those moments in the data, along with bootstrapped 95-percent confidence bands. These numbers are useful both as a guide for future calibrations and also for evaluating calibrations used in past work. As discussed above, the standard deviation of cross-sectional uncertainty relative to its mean is 29 percent over the full sample. Its quarterly autocorrelation is 0.906.

The next five columns of the table report analogous population moments in the calibrations of six recent structural models. In all five cases, the idiosyncratic uncertainty represents the volatility of firm-specific shocks to fundamentals (typically productivity). For the scaled standard deviation, the calibrations range from 0.09 to 0.71, lying inside the empirical confidence band in only a single case. For the autocorrelation, the results are similar. The calibrations range from 0.71 to 0.98 and are inside the empirical confidence bands in three of five cases. The table shows that there is little agreement in the literature on either the volatility or persistence of idiosyncratic risk. The data presented here can help resolve that disagreement.

5.2 Correlations

Table 2b reports raw correlations between major economic aggregates and uncertainty at the quarterly frequency. As in table 1, the correlations are weak and have mixed signs. The second and third columns report population correlations in the RUBC and CMR models. In both cases, with the exception of consumption growth, the correlations are substantially negative and well outside the empirical confidence bands. Both models focus on contractionary effects of uncertainty shocks. While the data presented here is consistent with the existence of such mechanisms during some episodes, in that uncertainty was high during the 2009 and 2020 recessions, the large increase in cross-sectional uncertainty during the boom of the late 1990's renders the overall correlation close to zero.

5.3 Forecasting regressions

A third way that the data can be used to evaluate models is to ask whether they can match the forecasting results obtained in table 1c. Recall that in addition to finding that cross-sectional uncertainty alone had some forecasting power, we also found that it was driven out by both realized cross-sectional dispersion and aggregate uncertainty.

To help focus in particular on the effects of uncertainty *shocks*, in both the model and data here we regress changes in macro aggregates on changes in uncertainty and realized dispersion. For CMR, then, the regression we estimate in the data and model is

$$\Delta y_t = b_0 + b_1 \Delta y_{t-1} + b_2 \Delta \sigma_{\varepsilon, t-1} + b_3 \Delta RD_{t-1}^{ret} + \eta_t \quad (6)$$

where η_t is a residual and y represents the log of GDP, consumption, investment, or hours worked. As in the analysis above, all variables are standardized to have unit standard deviation.

The first column in table 2c reports the estimates from the data, while the second column reports the (population) coefficients in CMR. In all four cases, there is again only a weak relationship in the data between economic activity and shocks to uncertainty – for investment, the relationship is actually positive – while realized dispersion has a statistically significantly negative relationship with real activity. In simulations of the CMR model, on the other hand, in all four cases it is uncertainty that is most important for forecasting real activity, rather than realized dispersion. As discussed above, realized dispersion is equal to uncertainty plus noise (the unexpected component of realized dispersion). In the CMR model, that noise does not have structural effects, so uncertainty dominates the forecasting regressions.

In the data, though, that “noise” – the gap between realized dispersion and its expectation – actually contains information. Models featuring concave responses to shocks, such as Ilut, Kehrig, and Schneider (2018) and Dew-Becker, Tahbaz-Salehi, and Vedolin (2020), are able to generate that effect. That is, in those models, realized dispersion does have effects on output, above and beyond the expected component encoded in uncertainty.

As discussed above, given that with our data we can construct a pair of matching measures of both aggregate and cross-sectional uncertainty, we can address the question of which of the two is more important for driving fluctuations. The second pair of columns in table 2c reports estimates from the regression

$$\Delta y_t = b_0 + b_1 \Delta y_{t-1} + b_2 \Delta \sigma_{\varepsilon, t-1} + b_3 \Delta \sigma_{mkt, t-t} + \eta_t \quad (7)$$

Compared to above, we have replaced RD_{t-1}^{ret} with $\sigma_{mkt, t-t}$. As in table 1c, table 2c shows that aggregate uncertainty drives cross-sectional uncertainty out of the empirical forecasting regressions. GDP, investment, and hours are all substantially more strongly driven by shocks to aggregate than cross-sectional uncertainty.

While CMR only has time-varying cross-sectional uncertainty, RUBC has fluctuations in both cross-sectional and aggregate uncertainty. In that paper’s benchmark calibration, cross-sectional and aggregate uncertainty are perfectly correlated. We run the simulation code three different times with calibrations where cross-sectional and aggregate uncertainty vary by different amounts, so that when the three calibrations are combined, the perfect correlation is broken.⁸

Table 2c shows that the coefficients on aggregate uncertainty in the forecasting regression

⁸In 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 $\{2.42, 2.76\}$ and $\{1.07, 6.21\}$ and append them to the baseline.

run on simulations of the RUBC model are, instead of being dominant, in all cases much smaller than the coefficients on cross-sectional uncertainty, by factors of three to four. In RUBC, it is primarily variation in cross-sectional rather than aggregate uncertainty that matters. That fact makes sense given that in the model, the vast majority of the variation in the total uncertainty faced by firms is from the cross-sectional component, but it is the opposite of what is observed empirically.

5.4 Implications

Table 2 makes two basic contributions. The top panel gives specific moments – volatility and autocorrelation – for calibrating structural models. The second contribution is to show how the raw correlations and forecasting regressions provide insights into aspects of the data that models can and cannot match. The data implies that realized dispersion is at least as important for forecasting as ex-ante uncertainty, if not more important. Similarly, it strongly implies that aggregate uncertainty is more important than cross-sectional uncertainty. Both of those features of the data are difficult for two leading models to match, so future work could use them as areas for improvement.

As discussed above, the targets in panels b and c of table 2 are difficult to match within a single model – they require being able to match the fact that uncertainty is sometimes good and sometimes bad, and also being able to explain why realized dispersion would have independent effects (which may require a nonlinear model). It is no criticism of RUBC and CMR that their benchmark calibrations do not match these new results. Rather, table 2 simply yields new empirical facts that structural models can be built or enriched to match.

6 International evidence

As in the US analysis, we obtain data from Optionmetrics Europe for 1/2002 to 12/2018 and from Bloomberg for 1/2019 to 5/2020. We have acceptable data for Switzerland, Germany, France, Great Britain, the Netherlands, and the constituents of the Euro Stoxx 50 index. 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.

Figure 3 plots cross-sectional uncertainty for each country against the US. In each case, cross-sectional uncertainty is clearly strongly correlated with that in the US, elevated in 2002, declining until the financial crisis, and then low and stable from 2010 until coronavirus in 2020 (the correlations are reported in the figure). That is true even though the path of

aggregate output in Europe over this period was very different from the US – a number of these countries went into recessions around 2012.

The table below summarizes the time-series behavior of market and cross-sectional uncertainty in the international data. The average time series standard deviation of market uncertainty is nearly twice that of cross-sectional uncertainty. 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 show that the finding of stability is not unique to the US: cross-sectional uncertainty has been similarly stable in other major developed economies. Second, there appears to be a strong international factor in cross-sectional uncertainty: uncertainty shocks have been global in nature over the last 18 years.

7 Conclusion

This paper reports a novel real-time index 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 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. 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 three episodes where it substantially grew, one a major economic expansion and the

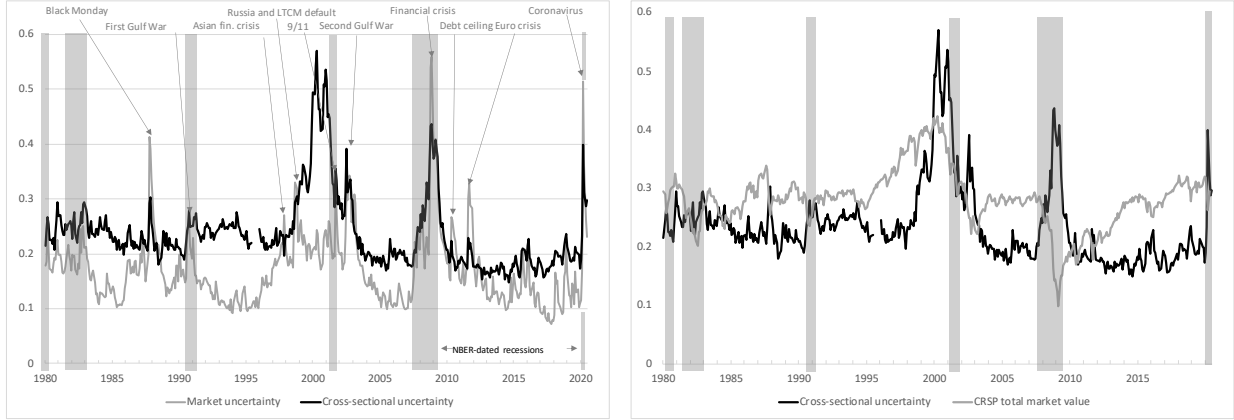
other two contractions. Studying raw correlations and forecasting regressions, we find that cross-sectional uncertainty is approximately acyclical overall and has little unconditional ability to forecast changes in future real activity. Overall, sometimes the data is consistent with models in which uncertainty shocks have causal negative effects on the economy, while in other periods it is consistent with models in which cross-sectional uncertainty is high following good shocks, perhaps due to a rise in creative destruction.

References

- Acemoglu, Daron**, “The Rise in Firm-Level Volatility: Causes and Consequences: Comment,” *NBER Macroeconomics Annual*, 2005, *20*, 203–215.
- , **Vasco M Carvalho**, **Asuman Ozdaglar**, and **Alireza Tahbaz-Salehi**, “The network origins of aggregate fluctuations,” *Econometrica*, 2012, *80* (5), 1977–2016.
- Ben-David, Itzhak**, **John R Graham**, and **Campbell R Harvey**, “Managerial mis-calibration,” *The Quarterly Journal of Economics*, 2013, *128* (4), 1547–1584.
- 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.
- Campbell, John Y.**, **Martin Lettau**, **Burton G. Malkiel**, and **Yexiao Xu**, “Have Individual Stocks Become More Volatile? an Empirical Exploration of Idiosyncratic Risk,” *Journal of Finance*, 2001, *56* (1), 1–43.
- Christiano, Lawrence J.**, **Roberto Motto**, and **Massimio Rostagno**, “Risk Shocks,” *American Economic Review*, 2014, *104*(1), 27–65.
- Comin, Diego** and **Sunil Mulani**, “A theory of growth and volatility at the aggregate and firm level,” *Journal of Monetary Economics*, 2009, *56* (8), 1023–1042.
- Dew-Becker, Ian**, **Alireza Tahbaz-Salehi**, and **Andrea Vedolin**, “Skewness and stochastic volatility in a network production model,” 2019. Working paper.
- Gabaix, Xavier**, “The granular origins of aggregate fluctuations,” *Econometrica*, 2011, *79* (3), 733–772.
- Gilchrist, Simon**, **Jae W. Sim**, and **Egon Zakrajsek**, “Uncertainty, Financial Frictions, and Investment Dynamics,” Working Paper 20038, National Bureau of Economic Research April 2014.

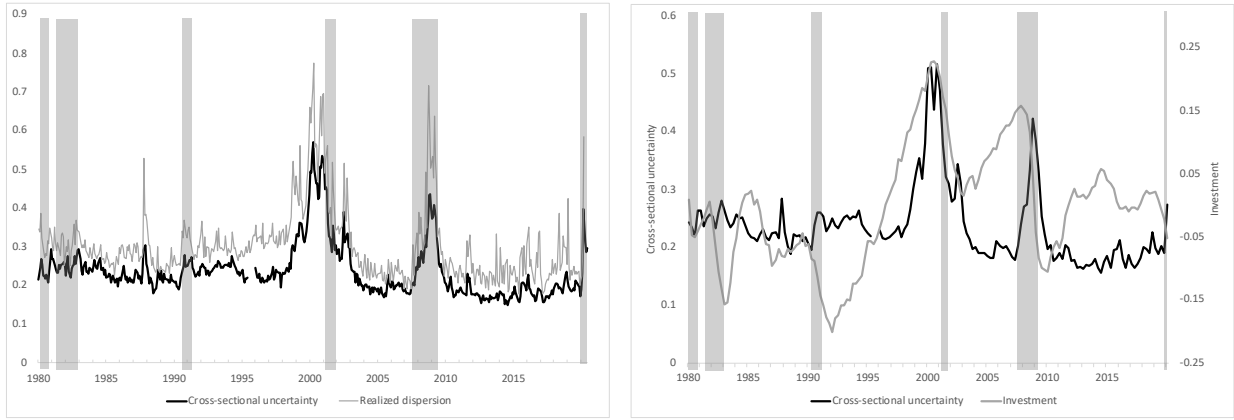
- Ilut, Cosmin, Matthias Kehrig, and Martin Schneider**, “Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News,” *Journal of Political Economy*, 2018, *126* (5), 2011–2071.
- Imbs, Jean**, “Growth and volatility,” *Journal of Monetary Economics*, 2007, *54* (7), 1848–1862.
- Kaplan, Greg and Giovanni L Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, *82* (4), 1199–1239.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, “Technological innovation, resource allocation, and growth,” *The Quarterly Journal of Economics*, 2017, *132* (2), 665–712.
- Schaal, Edouard**, “Uncertainty and unemployment,” *Econometrica*, 2017, *85* (6), 1675–1721.
- Senga, Tatsuro**, “A New Look at Uncertainty Shocks: Imperfect Information and Misallocation,” 2018. Working paper.
- Tella, Sebastial Di and Robert E. Hall**, “Risk Premium Shocks Can Create Inefficient Recessions,” 2020. Working paper.
- Tella, Sebastian Di**, “Uncertainty shocks and balance sheet recessions,” *Journal of Political Economy*, 2017, *125* (6), 2038–2081.

Figure 1: Time series of cross-sectional uncertainty



(a) Cross-sectional unc. and aggregate uncertainty

(b) Cross-sectional unc. and CRSP market value




(c) Cross-sectional unc. and realized dispersion

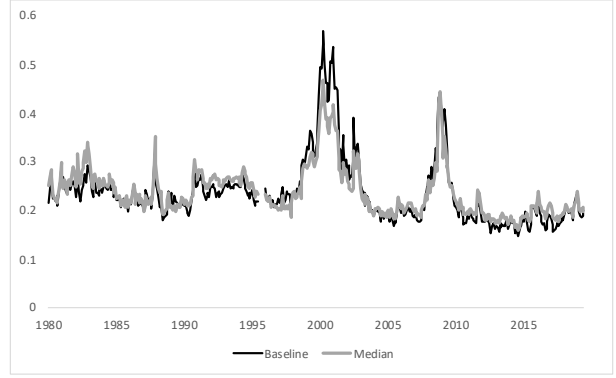
(d) Cross-sectional unc. and investment

Note: Each panel plots cross-sectional uncertainty (darker line) together with another time series: aggregate uncertainty (panel (a)), stock market value (panel (b)), realized dispersion (panel (c)), and investment (panel (d)). Options data before 1996 is from the Berkeley Options Dataset, for the period 1996/01-2019/06 is from Optionmetrics, and after 2019/06 is from Bloomberg. The VIX (market uncertainty) is obtained from CME options. Shaded areas are NBER recessions.


Figure 2: Robustness




(a) Cross-sectional unc. and the unemployment rate




(b) Baseline vs. median IV




(c) Baseline vs. beta-adjusted



(d) Baseline vs. employment weights



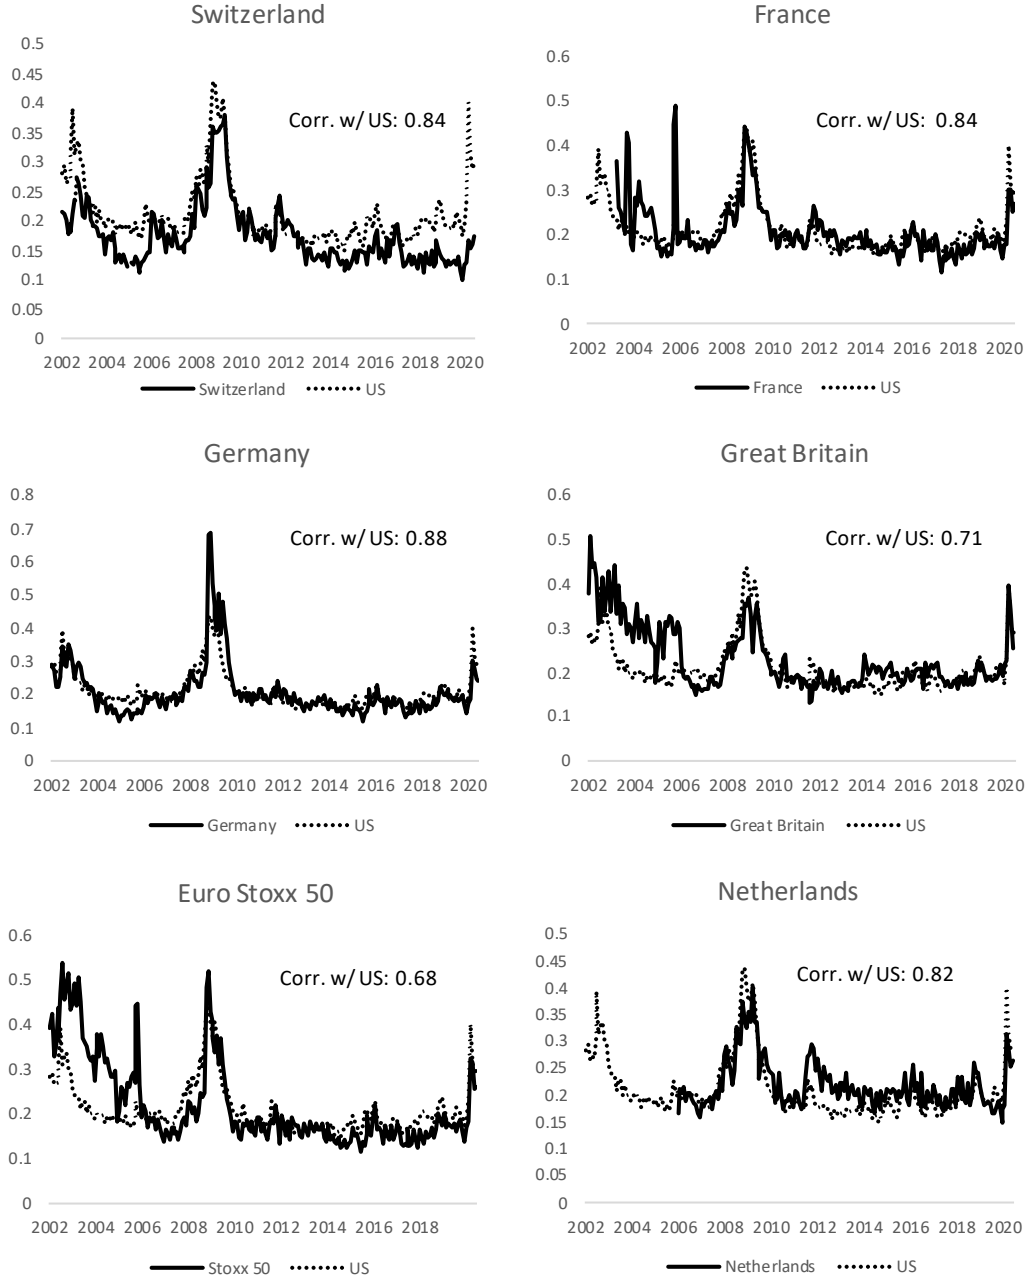
(e) Baseline vs. all Optionmetrics



(f) Baseline vs. industry-adjusted

Note: Panel (a) plots cross-sectional uncertainty together with the unemployment rate. Panel (b)-(f) plot our baseline measure of cross-sectional uncertainty together with alternative measures built in different ways. Specifically: in panel (b), the alternative measure uses the median of implied volatility across firms, instead of the weighted average by market cap; in panel (c), the alternative adjusts for betas; in panel (d), it uses employment weights; in panel (e), it uses all options in Optionmetrics instead of the largest 200; in panel (f), it controls for industry effects.

Figure 3: Cross-sectional uncertainty across countries



Note: Cross-sectional uncertainty from option data in European markets (solid line) against the one for the US (dotted line). Data is from Optionmetrics until 2018, and from Bloomberg since 2019.

Table 1: Cyclicity of cross-sectional regressions and forecasting results

(a) Correlations

	Detrended IP	Detrended empl.	Unemployment rate	CBO output gap	NBER recession	Capacity utilization	GZ bond spread	IP growth	Employment growth	Change in unempl. rate	Change in output gap
Full sample	0.01	0.22	-0.22	0.23	0.36	0.03	0.48	-0.20	-0.30	0.26	-0.22
Pre-1998	-0.34	-0.44	0.42	-0.43	0.28	-0.24	-0.16	-0.02	-0.15	0.06	0.07
Post-1998	0.03	0.31	-0.26	0.29	0.39	0.11	0.53	-0.23	-0.34	0.31	-0.28

(b) Forecasting realized dispersion

$$RD_t^x = b_0 + b_1 RD_{t-1}^x + b_2 \sigma_{\epsilon,t-1} + \eta_t$$

x :	Stock returns	Stock returns	IP growth IQR	Sales growth IQR
$\sigma_{\epsilon,t-1}$	0.91 [0.82,0.99]	0.69 [0.42,0.96]	0.17 [0.06,0.29]	0.25 [0.15,0.35]
RD_{t-1}^x		0.23 [-0.02,0.49]		
#obs.	143	143	142	141

(c) Uncertainty forecasting real activity

$$y_t = b_0 + b_1 y_{t-1} + b_2 \sigma_{\epsilon,t-1} + b_3 RD_{t-1}^{ret} + b_4 \sigma_{mkt,t-1} + \eta_t$$

	Unemployment			$\Delta \log(Empl.)$			$\Delta \log(IP)$		
$\sigma_{\epsilon,t-1}$	0.09 [0.04,0.14]	-0.01 [-0.05,0.04]	-0.02 [-0.08,0.04]	-0.15 [-0.25,-0.05]	-0.00 [-0.05,0.04]	0.02 [-0.09,0.14]	-0.13 [-0.22,-0.04]	-0.03 [-0.21,0.15]	-0.00 [-0.11,0.10]
RD_{t-1}^{ret}		0.04 [-0.01,0.08]			-0.02 [-0.06,0.03]			-0.04 [-0.21,0.14]	
$\sigma_{mkt,t-1}$			0.19 [0.13,0.25]			-0.35 [-0.48,-0.23]			-0.24 [-0.35,-0.13]
#obs.	442	434	442	442	434	442	442	434	442

Note: Panel (a) reports correlations between cross-sectional uncertainty and various macroeconomic variables. Panel (b) reports the results of a regression of three different measures of realized cross-sectional dispersion on lagged cross-sectional uncertainty and lagged realized dispersion. The three measures are: realized cross-sectional dispersion of stock returns, the cross-sectional interquartile ranges of growth in industrial production (across sectors), and growth in sales (across Compustat firms). Panel (c) reports forecasting regressions of real activity, in the three sections, respectively: unemployment, change in employment, and change in industrial production. In each section, the first column uses lagged cross-sectional uncertainty as predictor, the second column adds lagged realized cross-sectional dispersion, and the last column adds instead market-wide uncertainty.

Table 2: Comparing models to data moments

(a): Calibration moments

	Data	RUBC	CMR	Schaal	Di Tella	Gilchrist et al.
$SD[\sigma_{\varepsilon,t}]/E[\sigma_{\varepsilon,t}]$	0.29	0.71	0.58	0.14	0.20	0.09
$Corr[\sigma_{\varepsilon,t}, \sigma_{\varepsilon,t-1}]$	^[0.16,0.37] 0.91	0.91	0.98	0.94	0.71	0.90
	_[0.86,0.96]					

(b) Correlations of growth rates with $\Delta\sigma_{\varepsilon,t}$

	Data	RUBC	CMR
GDP	-0.04	-0.53	-0.39
Consumption	^[-0.17,0.12] 0.10	0.37	-0.05
Investment	^[-0.08,0.32] -0.01	-0.59	-0.43
Hours worked	^[-0.10,0.10] 0.10	-0.76	-0.21
	_[-0.09,0.25]		

(c) Forecasting regressions

$$CMR : y_t = b_0 + b_1 y_{t-1} + b_2 \sigma_{\varepsilon,t-1} + b_3 RD_{t-1}^{ret} + \eta_t$$

$$RUBC : y_t = b_0 + b_1 y_{t-1} + b_2 \sigma_{\varepsilon,t-1} + b_3 \sigma_{mkt,t-1} + \eta_t$$

y		Data	CMR		Data	RUBC
GDP:	$\sigma_{\varepsilon,t-1}$	-0.002	-0.24	$\sigma_{\varepsilon,t-1}$	-0.002	-0.049
		^[-0.013,0.009]			^[-0.005,0.002]	
	RD_{t-1}^{ret}	-0.002	0.05	$\sigma_{mkt,t-1}$	-0.003	-0.027
		^[-0.012,0.008]			^[-0.007,0.000]	
Consumption:	$\sigma_{\varepsilon,t-1}$	-0.005	-0.04	$\sigma_{\varepsilon,t-1}$	-0.000	-0.033
		^[-0.013,0.000]			^[-0.003,0.003]	
	RD_{t-1}^{ret}	0.004	0.03	$\sigma_{mkt,t-1}$	-0.002	-0.020
		^[-0.005,0.012]			^[-0.005,0.001]	
Investment:	$\sigma_{\varepsilon,t-1}$	0.017	-0.23	$\sigma_{\varepsilon,t-1}$	-0.008	-0.039
		^[-0.025,0.059]			^[-0.022,0.006]	
	RD_{t-1}^{ret}	-0.035	0.09	$\sigma_{mkt,t-1}$	-0.019	-0.016
		^[-0.076,0.005]			^[-0.033,-0.005]	
Hours:	$\sigma_{\varepsilon,t-1}$	-0.008	-0.04	$\sigma_{\varepsilon,t-1}$	-0.010	-0.023
		^[-0.038,0.022]			^[-0.020,0.001]	
	RD_{t-1}^{ret}	-0.016	0.00	$\sigma_{mkt,t-1}$	-0.022	-0.012
		^[-0.045,0.014]			^[-0.032,-0.013]	

Note: Panel (a) reports moments of the time-series of $\sigma_{\varepsilon,t}$ in the data and in various papers. Panel (b) reports correlations of $\sigma_{\varepsilon,t}$ with macroeconomic variables in the data and in the models. Panel (c) reports the results of forecasting regressions in the data and in the models. RUBC is Bloom et al. (2018); CMR is Christiano, Motto, and Rostagno (2016). All variables are at the quarterly frequency.

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.

In the period since 2009, there has been an increase in the seasonality of implied volatility around earnings announcement dates. For the period 2010–2020, we therefore estimate a nonlinear regression for average firm-level implied volatility that fits a sine curve to the data with precisely four cycles per year. That sine curve is then removed to yield the seasonally adjusted series.

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.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,

$$\underbrace{\text{var}(x_{i,t})}_{\text{Total variance}} = \underbrace{E\left[\text{var}_t(x_{i,t})\right]}_{\text{Average cross-sectional variance}} + \underbrace{\text{var}[E_t(x_{i,t})]}_{\text{Time-series variance of the average}} \quad (\text{A.1})$$

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 $\text{var}[E_t(x_{i,t})]$ to the total variance represents the fraction of the total variance explained by the cross-sectional mean in each period.

The variance decomposition identity (A.1) 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

$$\sigma_{R,\varepsilon,i,t}^2 \equiv \sigma_{i,t}^2 - \beta_{i,R,t}^2 \sigma_{mkt,t}^2 \quad (\text{A.2})$$

where $\sigma_{R,\varepsilon,i,t}^2$ 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 common component	
	Firm level	Sector level
Total firm uncertainty ($\sigma_{i,t}^2$)	0.50	0.75
Firm-specific uncertainty ($\sigma_{R,\varepsilon,i,t}^2$)	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. $\sum_{i \in S} w_{i,t} \sigma_{R,\varepsilon,i,t}^2$, 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.

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.