Black Swans and the Many Shades of Uncertainty

Nicholas Kozeniauskas, Anna Orlik and Laura Veldkamp
New York University and Federal Reserve Board [1]

July 14, 2014

[1]Preliminary draft. Please send comments to lveldkam@stern.nyu.edu. Thanks to Virgiliu Midrigan, Simon Mongey and Tom Sargent for useful discussions and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the position of the Board of Governors of the Federal Reserve or the Federal Reserve System.

Abstract

Various types of uncertainty shocks can explain many phenomena in macroeconomics and finance. But does this amount to throwing in a new, exogenous, unobserved shock to explain every challenging feature of business cycles? This paper explores the origin of micro uncertainty (uncertainty about firm-level shocks), macro uncertainty (uncertainty about aggregate shocks) and higher-order uncertainty shocks (disagreement) in a unified econometric framework. When agents use standard econometric techniques and real-time data to re-estimate parameters that govern the probability of black swans (unobserved extreme events), micro, macro and higher-order uncertainty covary just like their empirical counterparts. The results teach us that time-varying disaster risk and the many shades of uncertainty shocks are not distinct phenomena. All originate from using macro data to re-estimate the true probability distribution of economic outcomes.

A recent literature demonstrates that uncertainty shocks can explain business cycles, financial crises and asset price fluctuations with great success. But the way uncertainty shocks are measured and modeled varies from paper to paper. In some papers, an uncertainty shock means that an aggregate variable, such as GDP, becomes less predictable.[1] We refer to this as "macro uncertainty." In other papers, an uncertainty shock describes an increase in the variance of idiosyncratic shocks to firms, which manifests itself in an increase in the cross-sectional difference in firm outcomes.[2] We call this "micro uncertainty." Higher-order uncertainty, or the related term "confidence" describes the uncertainty about others' beliefs that arises when forecasts differ.[3] One reason that these phenomena are conflated is that they covary. At first glance, this paints a consistent picture of many uncertainty measures pointing to the same conclusion. But while macro uncertainty comes from aggregate shocks, micro uncertainty depends on firm-specific shocks, and higher-order uncertainty arises from private signal noise, all of which might well be independent. It is far from obvious then that these types of uncertainty should covary. This raises the question: What is the mechanism that generates and links the shades of uncertainty? If these belief shocks have no common underpinning, that is problematic for the whole literature because it amounts to introducing a new, unobserved, exogenous shock to explain every challenging feature of business cycles.

Macro uncertainty is the second moment of the distribution of a macro quantity (here, GDP growth) conditional on what an agent knows. Micro and higher-order uncertainty are cross-sectional variances that measure differences in firms' earnings or forecasts. A common justification for equating these measures is that when agents have greater prior uncertainty about macro outcomes, they respond more to their signals, which are heterogeneous. The greater role for heterogeneous signals causes the beliefs and outcomes among agents to be more heterogeneous as well. To investigate this story, we follow a standard approach.

[1]For macro uncertainty shocks, see e.g., Basu and Bundick (2012), Bianchi, Ilut, and Schneider (2012), (Schaal, 2012) on business cycles, and Bansal and Shaliastovich (2010), Pastor and Veronesi (2012) in the asset pricing literature.

[2]For micro uncertainty shocks, see e.g. Arellano, Bai, and Kehoe (2012), Christiano, Motto, and Rostagno (2014), Gilchrist, Sim, and Zakrajšek (2013), and (Bachmann and Bayer, 2012), who dispute the importance of these shocks for aggregate activity. Some papers such as Bloom (2009) and Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012) use both types of shocks.

[3]On the role of higher-order uncertainty see Angeletos and La'O (2014) and Angeletos, Collard, and Dellas (2014).

We measure fluctuations in macro uncertainty with a general autoregressive conditional heteroskedasticity (GARCH) model. Of course, one could add a stochastic process for the variance of firm-specific shocks, add another process for the dispersion of forecaster signals, and make all three variances covary. But that doesn't explain why these shocks covary or offer a unified framework for understanding where they come from. Instead, we build on a GARCH model, but hold the variance of the exogenous idiosyncratic shocks fixed so that all changes in micro and higher-order uncertainty are endogenous. We give agents signals about aggregate TFP, with a fixed amount of private signal noise, allow them to update beliefs with Bayes' law and then choose labor inputs in production. Quantitatively, this story flops. The problem is that the estimated macro uncertainty changes are small, which makes the effects on micro and higher-order uncertainties and their covariances miniscule. Allowing agents to re-estimate the forecasting parameters each period increases the size of all uncertainty shocks. But this still does not even come close to explaining the magnitude of any uncertainty shocks, their counter-cyclicality, or their high correlation with each other.

Next, we add one more ingredient to the forecasting model: learning about the probability of black swans. From Orlik and Veldkamp (2014), we know that when agents re-estimate parameters in a forecasting model with conditionally skewed outcomes, small changes in skewness estimates produce large fluctuations in the estimated probability of unobserved negative tail events (black swans). When we combine this mechanism with heterogeneous signals and a production economy, higher-order, macro and micro uncertainty all fluctuate more, covary more negatively with the business cycle, and covary more positively with each other. Thus the message is that, while shades of uncertainty are not obviously related, real-time estimation of the probability of black swans can produce fluctuations in the many shades of uncertainty, in a way that explains the uncertainty data.

Section 1 begins with measurement. It documents the covariance of micro, macro and higher-order uncertainty and addresses obvious alternative explanations. One possible explanation for the covariance is that all three uncertainties are driven by business cycle fluctuations. We show that, even after controlling for the business cycle, our uncertainty measures are related in a significant way. Another possible explanation has to do with binding constraints. When aggregate uncertainty rises, constraints are tighter, constrained firms have significantly different outcomes than unconstrained firms, and earnings dispersion rises. This may well be true. But this mechanism does not explain why both micro and macro uncertainty covary with higher-order uncertainty. An economic constraint bifurcates outcomes, but does not typically create differences in beliefs.

The benchmark production economy in Section 2 has productivity shocks with time-varying variance as the primary source of uncertainty fluctuations. The most common specification for time-varying variance is GARCH. Since GARCH variance estimates have only small fluctuations, we assume that firms do not know the GARCH parameters and use real-time data to re-estimate them each quarter. This real-time estimation amplifies macro uncertainty changes. The link between macro and micro uncertainty is updating with heterogeneous information. Each firm observes a private signal about TFP growth and then updates using Bayes' law. When past public information offers a less precise prediction of TFP and thus GDP growth, this represents high macro uncertainty. When forming beliefs in such uncertain times, firms weight the past public information less because its precision is low and weight the private signal more, in accordance with Bayes' law. A larger weight on the private, heterogeneous signal results in more heterogeneous beliefs. When firms have different beliefs about TFP, they choose different amounts of labor, causing output growth to become more dispersed. When we compare the model results with and without parameter learning, we see that learning about the parameters of the GARCH model can almost double the size of micro uncertainty shocks and increases fluctuations in macro uncertainty by an order of magnitude. The results do not support the hypothesis that large counter-cyclical uncertainty shocks create recessions. All the uncertainty fluctuations are quite small, acyclical, and have little covariance with each other. These results alone would suggest that uncertainty shocks are not a productive line of business cycle research.

What creates large uncertainty fluctuations with high correlation is using a forecasting model with conditional skewness (Section 3). The skewness of productivity shocks governs the probability of extreme events. Because GDP growth data has more extreme negative than positive outliers, our forecasters estimate skewness to be negative. When new pieces of data make estimated skewness more negative, left-tail events become more likely, the perceived risk of black swans rises and macro uncertainty increases. Because extreme events that have never been observed are more difficult to forecast accurately, a rise in disaster risk increases forecast dispersion. In high disaster-risk states, small differences in signals between agents about the state of the economy are amplified into large dispersion in forecasts. Divergent forecasts increase higher-order uncertainty, and generate larger differences in input choices and earnings, which is a spike in micro uncertainty. Thus, one of the key insights is that a forecasting model that reconciles higher-order, micro and macro uncertainty is also one that predicts a strong link between the risk of black swans and the many shades of of uncertainty shocks.

Related literature This paper builds on Orlik and Veldkamp (2014), in which a single representative agent estimates a forecasting model with conditional skewness and computes macro uncertainty. We add heterogeneous information and then embed the mechanism in a DSGE model to study effects on output. The new insight is that the degree of heterogeneity in firm output is endogenous and depends on macro uncertainty. This new mechanism can explain why different shades of uncertainty covary.

A few recent papers study the origins of uncertainty shocks. Some seek to explain why the cross-sectional dispersion of firm outcomes is countercyclical. Bachmann and Moscarini (2012) focus on the dispersion of prices and argue that recessions cause greater price dispersion because it is less costly for firms to experiment with their prices during bad economic times. Decker, D'Erasmo, and Moscoso Boedo (2013) argue that during recessions firms access fewer markets, which have idiosyncratic demand shocks, so they are less diversified and have more volatile outcomes. Others explain why macro uncertainty rises in recessions. In Nimark (2014), only outlier events are reported. Thus, the publication of a signal conveys both the signal content and information that the true event is far away from the mean, which increases macro uncertainty. In Van Nieuwerburgh and Veldkamp (2006) and Fajgelbaum, Schaal, and Taschereau-Dumouchel (2013), less economic activity generates less data, which increases uncertainty. Our paper differs because it connects dispersion across firms and forecasters to uncertainty about aggregate outcomes and shows why these shades of uncertainty covary, above and beyond what can be explained by the business cycle.

One piece of evidence that runs counter to the predictions of our model is the long-run decline in aggregate volatility and the concurrent increase in firm volatility over the last 30 years (Comin and Philippon, 2005). Comin and Mulani (2006) argue that this long-run trend comes from a shift from general-purpose technology to more specific technology development. Our paper has nothing to say about the long run trend in micro and macro uncertainty. It only explains the cyclical fluctuations. As such, it is a complement to this long-run theory.

1 The Empirical Puzzle

An important feature of this paper's explanation for micro uncertainty shocks is that they are closely related to macro uncertainty shocks. In this section, we first define micro and macro uncertainty and show how they differ conceptually. Then, we document the correlation between the two, argue that it is a statistically significant relationship and show that it is not because both measures are counter-cyclical.

In all quantitative exercises in this paper we will measure micro uncertainty with the cross-sectional interquartile range (IQR) of firm sales growth.[4] The growth rate of firm i in quarter t is computed as

$$g_{it}^{q} \equiv \frac{Q_{i,t+4} - Q_{i,t}}{\frac{1}{2}(Q_{i,t+4} - Q_{i,t})}. \tag{1}$$

where Q_{it} is the output of firm i in quarter t. Micro uncertainty in quarter t is the cross-sectional IQR of these growth rates:

$$MiU_t \equiv \text{IQR}(g_{it}^{q}). \tag{2}$$

In the theory sections of the paper we will discuss how this is an approximate measure of the uncertainty that firms have about their own growth rates due to idiosyncratic shocks. The data is from Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012) and covers 1962Q1–2009Q3. It is based on observations for all public firms with 100 quarters or more data in Compustat between 1962 and 2010. The sample contains 2,465 firms. We detrend the micro uncertainty series using a HP filter with smoothing parameter equal to the standard value for quarterly data (1600) and scale (divide) the detrended series by the trend.[5] The data is presented in Figure 1. We detrend the series to remove fluctuations in

[4]We use IQR in order to make contact with an existing literature, e.g., Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012), which uses IQR to measure micro uncertainty.

[5]The formula for detrending is (data$_t$ - trend$_t$)/trend$_t$.

micro uncertainty with different frequencies to those that we are explaining. The long-run trend is addressed by Comin and Mulani (2006) and Comin and Philippon (2005).

We measure higher-order uncertainty with the cross-sectional standard deviation of real GDP growth forecasts. The data is from the Survey of Professional Forecasters and is for 1968Q4–2011Q4. The forecast data provides one period ahead forecasts of real GDP (that is, forecasts of quarter t real GDP made at the end of quarter $t-1$). Throughout the paper we will be concerned with real rather than nominal GDP, so from here on we will refer to it as GDP. Let Q_t denote GDP in period t and define period t GDP growth to be $\Delta q_t = \log Q_t - \log Q_{t-1}$. We compute approximate forecasts of GDP growth using the GDP forecast data as follows:

$$E_{i,t-1}[\Delta q_t] \equiv \log(E_{i,t-1}[Q_t]) - \log Q_{t-1} \tag{3}$$

Higher-order uncertainty in period $t-1$ is measured with the cross-sectional standard deviation of the forecasts in (3):

$$HU_{t-1} \equiv \sqrt{\frac{1}{N_{t-1}} \sum_{i \in I_{t-1}} (E_{i,t-1}[\Delta q_t] - \bar{E}_{t-1})^2} \tag{4}$$

where $\bar{E}_{t-1} = 1/N_{t-1} \sum_i E_{i,t-1}[\Delta q_t]$ is the average growth forecast, i indexes forecasters, I_{t-1} is the set of forecasters at the end of period $t-1$ and N_{t-1} is the number of forecasters in this set. The average number of forecasters in a quarter is 41, with a standard deviation of 17. We detrend this data series using the same method as we used for the micro uncertainty data in order to remove variation in the series at frequencies that we do not seek to explain. The data is presented in Figure 2.

The correlation of micro and higher-order uncertainty is 0.43 and regressing micro uncertainty on higher-order uncertainty shows that they have a positive relationship that's significant at the 1% level—see column (1) in Table 1. The coefficient of 0.2807 means that when higher-order uncertainty deviates from trend by an additional one percentage point, micro uncertainty deviates from trend by an additional 0.2807 percentage points. Both types of uncertainty are also countercyclical: micro and higher-order uncertainty have correlations of -0.52 and -0.28, respectively, with GDP growth. Several other papers

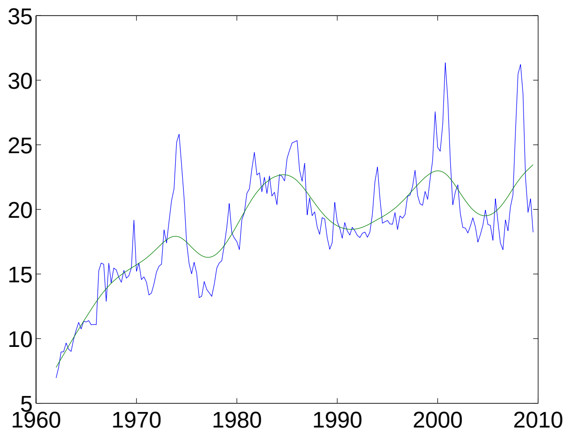

(a) Data and trend

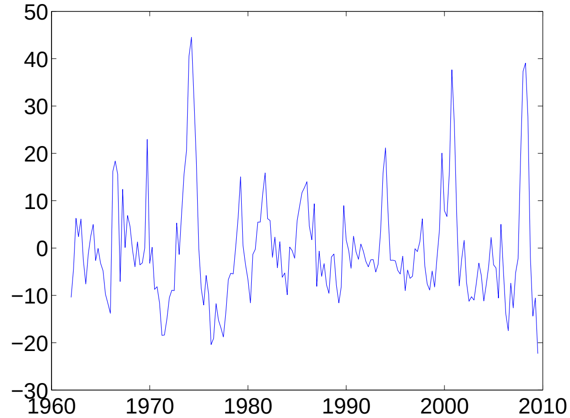

(b) Detrended data

Figure 1: **Micro uncertainty: IQR of firm sales growth.** Panel (a) presents the data and its trend. The scale is the annualized percentage. The trend is computed using a Hodrick-Prescott filter with the smoothing parameter equal to 1600. In panel (b) the data is detrended and scaled (divided) by the trend. The scale is the percentage deviation from trend.

starting with Bloom (2009) have documented similar facts.

One reason why shades of uncertainty could comove is that they are all driven by the business cycle. We test this hypothesis by regressing micro uncertainty on higher-order uncertainty, controlling for the business cycle. To control for the business cycle we use two variables, GDP growth and a recession variable which measures the number of months over the relevant period that were in recessions according to the NBER. The GDP growth data is from the BEA. The results are presented in columns (2) and (3) of Table 1. The main result is that micro and higher-order uncertainty have a significant positive relationship even after controlling for the business cycle.

Next, we regress higher-order uncertainty on macro uncertainty. Macro uncertainty is not as easy to observe as micro and higher-order uncertainty. It is the conditional variance of of next period's GDP growth: $Var_{i,t-1}[\Delta q_t]$. This is equivalent to the expected squared forecast error. The re-stating reveals that macro uncertainty depends on a forecasting model. We will discuss at length what this model should be. But a volatility option on the stock market increases in value as aggregate uncertainty rises. We follow Bloom (2009) and use the VXO, a equity volatility index, as a proxy for macro uncertainty. The results in

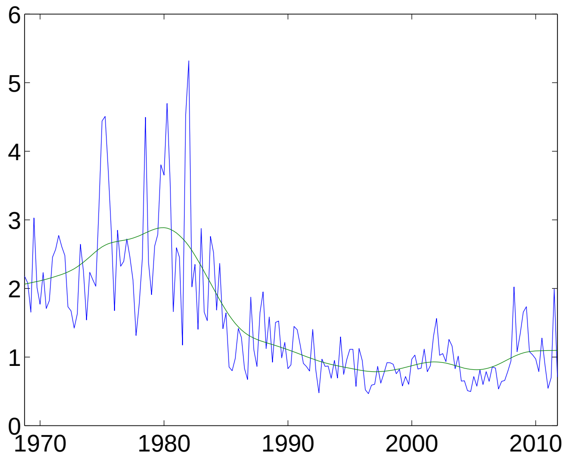

(a) Data and trend

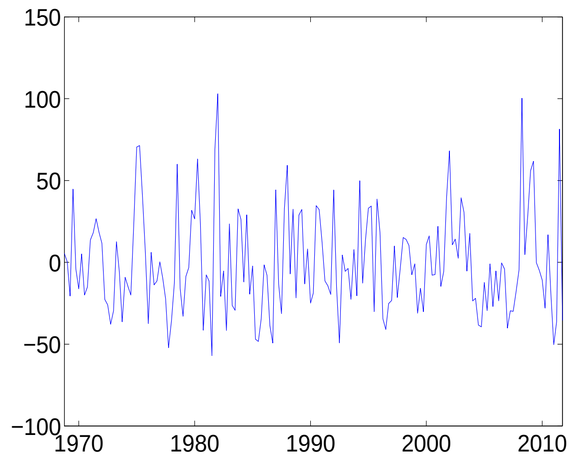

(b) Detrended data

Figure 2: **Higher-order uncertainty: GDP growth forecast dispersion.** GDP growth forecast dispersion is the cross-sectional standard deviation of GDP growth forecasts. The units are the annualized percentage. Panel (a) presents the data and its trend. The trend is computed using a Hodrick-Prescott filter with the smoothing parameter equal to 1600. In panel (b) the data is detrended and scaled (divided) by the trend. The scale is the percentage deviation from trend.

table 2 reveal that higher-order uncertainty and macro uncertainty are highly related, and that this relationship extends well beyond business cycle fluctuations. Taken together, these findings suggests a deeper connection between micro, macro and higher-order uncertainty.

2 A Baseline Model

This model illustrates the first of two mechanisms that can jointly explain the shades uncertainty. Firms form beliefs about aggregate productivity by combining prior information that is common to all firms with their own private signals. When the precision of prior beliefs is low, the common prior information that firms see is less informative about next period's productivity. This will generate high uncertainty about GDP growth. It will also cause firms to weight their private signals more when forming their posterior beliefs. Weighting private signals more acts like an increase in the variance of an idiosyncratic shock because it causes firms to have more dispersed beliefs about productivity and, as a result, choose more heterogeneous labor inputs. This causes output growth dispersion, our

Dependent variable: Micro uncertainty	(1)	(2)	(3)
Higher-order uncertainty	0.2807***	0.1338***	0.1264***
	(0.047)	(0.046)	(0.004)
GDP growth		−2.5489***	
		(0.357)	
Recession			1.8496***
			(0.226)
Obs	165	165	165
Sample		1968Q4–2009Q3	

Table 1: **Uncertainty shocks are correlated, even after controlling for business cycles.** *Micro uncertainty* is the IQR of firm sales growth, defined in equation (2). *Higher-order uncertainty* is the standard deviation of GDP growth forecasts, defined in equation (4). The units for both series are the percentage deviation from trend. Since micro uncertainty is based on four-quarter growth rates and macro uncertainty is based on forecasts of quarterly growth rates, we regress MiU_t on the average of HU_t, HU_{t+1}, HU_{t+2} and HU_{t+3} so that the periods of the variables are comparable. Similarly *GDP growth* is 400 times the average of Δq_{t+1}, Δq_{t+2}, Δq_{t+3} and Δq_{t+1}. The units are the approximate annualized percentage. *Recession* is the number of months in quarters $t+1$ to $t+4$ which were in recessions according to the NBER. Standard errors are in parentheses. ***, ** and * denote significance at the 1%, 5% and 10% levels respectively.

measure of micro uncertainty, to rise. At the end of the section, we calibrate the model to see whether the model can explain micro, macro and higher-order uncertainty, and their relationships to the business cycle.

2.1 Model

Time is discrete and starts in period 0. There is a unit mass of firms in the economy with each firm comprised of a representative agent who can decide how much to work. Agent i's utility in period t depends on his output Q_{it} and the effort cost of his labor L_{it}:

$$U_{it} = Q_{it} - L_{it}^{\gamma} \tag{5}$$

for some $\gamma > 1$. Output depends on labor effort and productivity A_t:

$$Q_{it} = A_t L_{it}. \tag{6}$$

	Dependent variable: Higher-order Uncertainty			
	(1)	(2)	(3)	(4)
Macro uncertainty	0.069***	0.053**	0.058**	0.055**
	(0.023)	(0.025)	(0.023)	(0.024)
Recession		0.046*		0.011
		(0.024)		(0.028)
GDP growth			−0.021***	−0.019**
			(0.007)	(0.008)
Period	1968Q4–2008Q3			

Table 2: **Regression results for the data** *Higher-order uncertainty* is the standard deviation of GDP growth forecasts, defined in equation (4). *Macro uncertainty* is the stock market volatility series constructed by Bloom (2009), normalized by (divided by) its standard deviation. For 1986 onwards this series is the CBOE's VXO. Since the VXO isn't available prior to 1986, Bloom uses 'the monthly standard deviation of the daily S&P500 index normalized to the same mean and variance as the VXO index when they overlap from 1986 onward.' We average over the months of each quarter to get a quarterly series. *Recession* is the number of months in each quarter which were part of recessions according to the NBER.

Aggregate output is $Q_t \equiv \int Q_{it} di$ and GDP growth is $\Delta q_t \equiv \log Q_t - \log Q_{t-1}$. The growth rate of productivity at time t, $\Delta a_t \equiv \log(A_t) - \log(A_{t-1})$, is normally distributed with a constant mean and a variance that evolves according to a GARCH(1,1) process:

$$\Delta a_t = \alpha_0 + \sigma_t \epsilon_t, \tag{7}$$

$$\sigma_t^2 = \alpha_1 + \rho \sigma_{t-1}^2 + \phi \epsilon_{t-1}^2, \tag{8}$$

where $\epsilon_t \sim N(0,1)$.

Agent i makes his labor choice L_{it} at the end of period $t-1$. His objective is to maximize expected period t utility.[6] The agent makes this decision at the end of period $t-1$, which means that the agent does not know productivity A_t. However, at the end of period $t-1$ he observes an unbiased signal about the growth of productivity from $t-1$ to t:

$$z_{i,t-1} = \Delta a_t + \eta_{t-1} + \psi_{i,t-1}, \tag{9}$$

[6]Decisions at time t have no effect on future utility so the agent is also maximizing expected discounted utility.

where $\eta_{t-1} \sim N(0, \sigma_\eta^2)$ and $\psi_{i,t-1} \sim N(0, \sigma_\psi^2)$. The public noise shock η_{t-1} is *i.i.d* over time and the private noise shock $\psi_{i,t-1}$ is i.i.d. across agents and over time.[7] Note that there is common and idiosyncratic signal noise.

2.2 Solution to the firm's problem

The first-order condition for agent i's choice of period t labor is:

$$L_{it} = \left(\frac{E_{i,t-1}[A_t]}{\gamma} \right)^{1/(\gamma-1)}. \tag{10}$$

In order to make his choice of labor the agent forecasts productivity. He forms a prior belief about TFP growth and then updates using his idiosyncratic signal. He knows that TFP growth follows the process specified in equations (7) and (8) but does not know the four parameters of that process. To form his prior about Δa_t at the end of period $t-1$ he uses the data $A^{t-1} \equiv \{A_0, A_1, \dots, A_{t-1}\}$ to estimate the TFP growth process. This prior is normally distributed with mean $E[\Delta a_t | A^{t-1}]$ and variance $v_{t-1} = V[\Delta a_t | A^{t-1}]$. When the agent receives his signal at the end of period $t-1$ he updates his beliefs according to Bayes' law:

$$E_{i,t-1}[\Delta a_t] = \frac{v_{t-1}^{-1} E[\Delta a_t | A^{t-1}] + (\sigma_\eta^2 + \sigma_\psi^2)^{-1} z_{i,t-1}}{v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}}.$$

Note that $(\sigma_\eta^2 + \sigma_\psi^2)^{-1}$ is the precision of that signal. Let the posterior variance be denoted $V_{i,t-1}[\Delta a_t] \equiv [v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}]^{-1}$. Note that this variance is common across agents because all agents receive signals with the same precision. If we define $\omega_{t-1} \equiv [(\sigma_\eta^2 + \sigma_\psi^2)(v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1})]^{-1}$, we can write agent i's forecast of Δa_t as a weighted sum of prior beliefs and the signal:

$$E_{i,t-1}[\Delta a_t] = (1 - \omega_{t-1}) E[\Delta a_t | A^{t-1}] + \omega_{t-1} z_{i,t-1}. \tag{11}$$

[7] We have assumed that there is no labor market, which means that there is not a wage which agents can use to learn about Δa_t. While a perfectly competitive labor market which everyone participates in could perfectly reveal Δa_t, there are many other labor market structures with frictions in which wages would provide no signal, or a noisy signal, about Δa_t. An additional noisy public signal would not provide much additional insight since we already allow for public noise in the signals that agents receive. It would however add complexity to model, so we close this learning channel down. Also note that if agents traded their output, prices would not provide a useful signal about TFP growth because once production has occurred, agents know TFP exactly.

Once the agent has beliefs about TFP growth, he computes his expected value of TFP using the fact that $A_t = A_{t-1}\exp(\Delta a_t)$:

$$E_{i,t-1}[A_t] = A_{t-1}\exp\left(E_{i,t-1}[\Delta a_t] + \frac{1}{2}V_{i,t-1}[\Delta a_t]\right) \tag{12}$$

and makes his labor choice according to equation (10).

2.3 Mapping productivity forecasts into output forecasts

In the quantitative section of the paper we will make use of GDP forecast data, so we derive these forecasts here. When making his GDP forecast an agent must consider his own beliefs about TFP, which will inform his output choice, and his beliefs about other agents' beliefs about TFP, as these will inform their output choices. Individual higher order beliefs don't matter because agents are measure zero.

From equations (6) and (10), firm i's level of output and aggregate output in period t are

$$Q_{it} = A_t\left(\frac{E_{i,t-1}[A_t]}{\gamma}\right)^{1/(\gamma-1)} \tag{13}$$

$$Q_t = \Gamma_{t-1}\exp[f(\Delta a_t, \eta_{t-1})], \tag{14}$$

where $f(\Delta a_t, \eta_{t-1}) \equiv \Delta a_t + (\frac{\omega_{t-1}}{\gamma-1})(\Delta a_t + \eta_{t-1})$ and

$$\Gamma_{t-1} \equiv \gamma^{1/(1-\gamma)}A_{t-1}\exp\left[\frac{1}{\gamma-1}\left(\log A_{t-1} + (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \frac{d_{t-1}^2}{2(\gamma-1)} + \frac{V_{i,t-1}[\Delta a_t]}{2}\right)\right].$$

The following lemma provides firm i's GDP growth forecast. The proof is in the appendix.

Lemma 1. *Given all information observed by the end of period t-1, firm i's forecast of period-t GDP is*

$$E_{i,t-1}[Q_t] = \Gamma_{t-1}\exp\left(E_{i,t-1}[f(\Delta a_t, \eta_{t-1})] + \frac{1}{2}V_{i,t-1}[f(\Delta a_t, \eta_{t-1})]\right). \tag{15}$$

where

$$E_{i,t-1}[f(\Delta a_t, \eta_{t-1})] = (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1} z_{i,t-1} \\ + \left(\frac{\omega_{t-1}}{\gamma-1}\right)\left(\frac{(v_{t-1}+\sigma_\eta^2)^{-1}E[\Delta a_t|A^{t-1}] + \sigma_\psi^{-2} z_{i,t-1}}{(v_{t-1}+\sigma_\eta^2)^{-1}+\sigma_\psi^{-2}}\right), \quad (16)$$

$$V_{i,t-1}[f(\Delta a_t, \eta_{t-1})] = \left[\left(1+\frac{\omega_{t-1}}{\gamma-1}\right) \quad \frac{\omega_{t-1}}{\gamma-1}\right] V_{i,t-1}[x_t] \left[\left(1+\frac{\omega_{t-1}}{\gamma-1}\right) \quad \frac{\omega_{t-1}}{\gamma-1}\right]', \quad (17)$$

$x_t \equiv [\Delta a_t - E[\Delta a_t|A^{t-1}], \ \ \eta_{t-1}]'$, *and*

$$V_{i,t-1}[x_t] = \begin{bmatrix} v_{t-1} & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix} - \begin{bmatrix} v_{t-1} \\ \sigma_\eta^2 \end{bmatrix} (v_{t-1}+\sigma_\psi^2+\sigma_\eta^2)^{-1} \begin{bmatrix} v_{t-1} & \sigma_\eta^2 \end{bmatrix}. \quad (18)$$

2.4 Sources of uncertainty

Why are macro, micro and higher-order uncertainty shocks related? When prior beliefs about GDP growth are imprecise (macro uncertainty), firms place a high weight on their idiosyncratic signals, generating dispersion in firm growth rates (micro uncertainty) and GDP growth forecasts (higher-order uncertainty). This section describes these effects in detail.

Let the log growth rate of firm i at time t be Δq_{it}. Using equation (13), this growth rate can be expressed as a function of beliefs about TFP:

$$\Delta q_{it} = \Delta a_t + \frac{1}{\gamma-1}\Big(\log(E_{i,t-1}[A_t]) - \log(E_{i,t-2}[A_{t-1}])\Big). \quad (19)$$

Making use of equations (11) and (12), the cross-sectional variance of firm growth rates in period t is therefore

$$V[\Delta q_{it}] = \left(\frac{1}{\gamma-1}\right)^2 \sigma_\psi^2(\omega_{t-1}^2+\omega_{t-2}^2). \quad (20)$$

Fixing ω_{t-2}, this tells us that shocks to measured micro uncertainty can come from two sources. The first is changes in private signal noise (σ_ψ). The second is the weight placed on signals (ω_{t-1}), which is determined by changes in the precision of prior beliefs relative to signals. Only the second channel operates in our model since σ_ψ is fixed over time. When prior beliefs are relatively imprecise, ω_{t-1} is relatively high. This means that agents place relatively high weight on their heterogeneous signals, which acts like an increase in the variance of an idiosyncratic shock. It amplifies dispersion in forecasts of TFP, choices of labor, levels of output and, ultimately, dispersion in growth rates.

How does measured micro uncertainty relate to the beliefs of firms about their own growth rates prior to receiving their idiosyncratic signals? Before receiving their idiosyncratic signals, firms are identical so they share the same beliefs. The variance of a firm's beliefs about its output growth prior to receiving its signal at the end of period $t-1$ is

$$V[\Delta q_{it}|A^{t-1}] = V[\Delta a_t|A^{t-1}] + \Big(\frac{1}{\gamma-1}\Big)^2 \omega_{t-1}^2(\sigma_\eta^2 + \sigma_\psi^2).$$

This can be derived using equations (11), (12) and (19). We can see that there are three sources of variance in firms' beliefs about their growth rates: variance in beliefs about aggregate TFP growth, variance due to public signal noise (σ_η) and variance due to private signal noise (σ_ψ). Thinking of micro uncertainty as the variance due to idiosyncratic shocks, micro uncertainty is $(1/(\gamma-1))^2\omega_{t-1}^2\sigma_\psi^2$. Comparing this to the measure of micro uncertainty that is observable—equation (20)—we see that measured micro uncertainty is the same except for one extra term, ω_{t-2}. This extra term arises because there is dispersion in firm output in period $t-1$ that does not contribute to a firm's uncertainty about its own growth rate, but does show up when when you measure the cross-sectional variance of growth rates. However, if ω_{t-2} has a reasonably high correlation with ω_{t-1} then measured micro uncertainty will be highly correlated with our theoretical notion of micro uncertainty.

Turning now to higher-order uncertainty, the measure that we use is the cross-sectional standard deviation of GDP growth forecasts, defined in equation (4). Using equations (31), (32) and (33), we can evaluate the cross-sectional variance:

$$V\big[\log(E_{i,t-1}[Q_t]) - \log Q_{t-1}\big] = \left[1 + \left(\frac{1}{(\gamma-1)[\sigma_\psi^2(v_{t-1}+\sigma_\eta^2)^{-1}+1]}\right)^2\right]\sigma_\psi^2\omega_{t-1}^2. \quad (21)$$

We want to understand why this cross-section variance covaries with macro uncertainty, the uncertainty for agent i is the variance of his beliefs at the end of period $t-1$ about GDP growth in period t: $V_{i,t-1}[\log(\Delta q_t)]$. Evaluating this expression for macro uncertainty, using equations (30), (33) and (34) gives

$$V_{i,t-1}[\log(\Delta q_t)] = \frac{(\gamma - 1 + \omega_{t-1})^2 v_{t-1}\sigma_\psi^2 + (\gamma-1)^2 v_{t-1}\sigma_\eta^2 + \omega_{t-1}^2\sigma_\eta^2\sigma_\psi^2}{(\gamma-1)^2(v_{t-1} + \sigma_\psi^2 + \sigma_\eta^2)}. \tag{22}$$

Note that the right hand side of the expression is independent of i. Since ω_{t-1} is increasing in v_{t-1} it is clear from equation (21) that the standard deviation of GDP growth forecasts is strictly increasing in v_{t-1}. Taking the derivative of (22) with respect to v_{t-1} and using the fact that $\gamma > 1$ gives the result that $V_{i,t-1}[\log(\Delta q_t)]$ is also strictly increasing in v_{t-1}. Thus higher-order uncertainty and macro uncertainty should comove.

What is the source of this comovement? Shocks to macro uncertainty in the GARCH model come from changes in the variance of the prior, v_{t-1}. Agents only have two sources of information, their prior and their signals, and the precision of their signals is the same over time. When agents have a less precise prior (v_{t-1} is larger) the precision of their beliefs about TFP growth are less precise, making them more uncertain about GDP growth. This shows up as an increase in the dispersion of GDP growth forecasts, because when agents have a less precise prior they place more weight on their idiosyncratic shocks. Consequently they have more dispersed forecasts of TFP growth and GDP growth. We can see this increase in the dispersion of GDP growth forecasts clearly in equation (21). When the prior is less precise, ω_{t-1} is larger, causing $V[\log(E_{i,t-1}[Q_t]) - \log Q_{t-1}]$ to increase. Note that the origin of higher-order shocks is exactly the same as the origin of micro uncertainty shocks, so the model will generate positive comovement between micro and higher-order uncertainty.

2.5 How well can the model explain uncertainty shocks?

The previous section shows that if this is an accurate description of the economic environment, then changes in macro uncertainty can generate changes in higher-order and micro uncertainty. But how large might this effect be? Can it explain the uncertainty relationships quantitatively? To answer this question, we calibrate the model to match moments of GDP growth and GDP growth forecasts and then assess how uncertainty shocks in the model compare to those in the data.

Calibration and model simulation The model has seven parameters: a parameter controlling the disutility of labor (γ), four parameters for the GARCH process (α_0, α_1, ρ and ϕ) and public and private signal noise (σ_η and σ_ψ respectively). We set $\gamma = 2$, which corresponds to a Frisch labor supply elasticity of one. This is at the lower end of the range of Frisch elasticities that Keane and Rogerson (2012) argue are reasonable at the macro level. The six remaining parameters are calibrated to target six moments of the data for 1968q4–2011q4: the mean, standard deviation, skewness and kurtosis of GDP growth;[8] the average cross-sectional standard deviation of GDP growth forecasts; and the average (over forecasters and over time) absolute error of GDP growth forecasts. Recall that all references to GDP are references to real GDP. We target moments of GDP growth rather than moments of TFP growth because our forecast data is for GDP growth, so we want the GDP growth process in the model to match the data, even if our production economy is simplistic. The cross-sectional standard deviation of GDP growth forecasts is defined in equation (4). The average absolute forecast error is

$$\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{i\in I_t}|E_{i,t-1}[\Delta q_t] - \Delta q_t|, \tag{23}$$

where T is the number of periods in the data.

Note that in calibrating the model we are only using one moment of the two uncertainty series that we are trying to explain as a target (the mean level of higher-order uncertainty). We are calibrating in a way that allows the time variation in all three types of uncertainty, their correlations with the business cycle and their correlations with each other to be endogenously determined. The only cyclical properties we calibrate to are the standard deviation, skewness and kurtosis of GDP growth.

We make two assumptions to simplify computations. First, when agents estimate the TFP process and use the estimate to construct prior beliefs about TFP growth, they regard their parameter estimates as the truth. This ignores agents' parameter uncertainty. Second, we assume that agents do not use their signals when estimating TFP parameters.

[8] We attempted to match skewness and kurtosis and the reported parameter estimates represent our best attempt. But the model cannot match skewness or kurtosis of the GDP growth data. This reinforces our motive or using the skewed model in the next section. But it leaves us short two calibration targets. The next draft will find two new attainable calibration targets to identify the GARCH parameters.

Parameter	*Value*	*Target Moment*
γ	2	Frisch elasticity
α_0	0.0034	GDP growth mean
α_1	2.09e-5	GDP growth st. dev.
ρ	7.03e-15	TBD
ϕ	0.438	
σ_η	0.0059	Average abs. error for GDP growth forecasts
σ_ψ	0.0019	Average st. dev. of GDP growth forecasts

Table 3: **Parameter values for the GARCH model.** All moments of the data are for 1968q4–2011q4.

Because the data history is a long string of public signals and the private signal is only one noisy piece of information, parameters would differ only slightly across agents. We have done some small-scale experiments to verify that the heterogeneous parameter effect we ignore is small.

To calibrate the six parameters of the model excluding γ we use simulated method of moments with the identity weighting matrix. Details are in the appendix. The calibrated parameter values are reported in Table 3. In this table, as well as Table 6, this model is referred to as the *GARCH model.* One thing to note from the calibration is that $\rho \approx 0$ so the variance of TFP growth effectively follows an ARCH(1) process.

Our model does not match the average level of micro uncertainty (equation (2)) in the data. This will be the case for all of the models in the paper. This is unsurprising since there are many fixed sources of firm heterogeneity missing (e.g., size, age, industry, geography).[9] For now, we detrend the data and model output and compare the percentage deviations from trend.

Results Table 6 reports results for micro uncertainty, higher-order uncertainty, macro uncertainty and the calibration moments. As discussed in Section 1, we don't have a close data analog for macro uncertainty, so when comparing uncertainty in the model and the data we will focus on micro uncertainty and higher-order uncertainty. There are two columns in the table with results for the benchmark model. The column for which

[9] In the next revision, we plan to address this shortcoming by calibrating a distribution of firm fixed-effects.

parameter learning is "on" presents results for the model exactly as it has been described. In the other column, for which parameter learning is "off," we assume that agents know the parameters of the GARCH process. Examining the two columns in the results table without parameter learning highlights its important role. There are effectively no macro uncertainty shocks without parameter learning and therefore also no correlation with micro shocks. It is also an important source of micro uncertainty shocks, explaining nearly half of those. This shows us that learning about parameters amplifies changes in uncertainty and is a necessary feature of this model for generating non-trivial uncertainty shocks.

For comparing the model to the data the key moments are the standard deviation of micro and higher-order uncertainty, which measure the magnitude and frequency of shocks, the correlation of micro and higher-order uncertainty, and their correlations with GDP growth. Focus first on the results for when parameter learning is on. The main success of the model is that it generates positively correlated uncertainty shocks. However, relative to the data the shocks are small since the standard deviations of micro uncertainty and higher-order uncertainty are only about 30% and 20%, respectively, of their values in the data. Also, the correlation between micro and higher-order uncertainty is only just over one-third of its value in the data. Furthermore, all shades of uncertainty are acyclical in the model and countercyclical in the data. The lack of cyclicality in uncertainty is particularly problematic since much of what researchers want uncertainty to explain is the onset of recessions. Thus, while this model is a framework for thinking about many shades of uncertainty shocks, this mechanism is not quantitatively viable. This benchmark model makes clear that the relationship between the types of uncertainty is not obvious, mechanical, or easy to generate.

3 Adding Skewness to the Forecasting Model

Since the GARCH model cannot explain the size of uncertainty shocks, their correlation with each other, or their cyclicality, we explore a second forecasting model. One striking feature of the previous benchmark model is that it cannot replicate the negative skewness of the distribution of GDP growth: skewness is 0.01 in the model and -0.32 in the data. The model is missing the possibility of extreme negative events that might well explain high uncertainty. The next model explores whether relaxing the restriction that TFP growth is normally distributed makes a big difference. It turns out that it does. If you allow for non-normality in the GARCH model and calibrate it so that it matches the level of skewness in GDP growth, then the fluctuations in the various shades of uncertainty are large, countercyclical, and are correlated with each other. This non-normality is a way of capturing beliefs about rare, negative events. Thus, the model teaches us that re-estimating disaster risk can simultaneously produce fluctuations in many types of uncertainty.

3.1 Model

The preferences (5) and technology (6) are the same as in the GARCH model. The TFP growth process is a nonlinear transformation of X_t, which is normally distributed, with time-varying variance that follows a GARCH(1,1) process:

$$\Delta a_t = c + b\exp(X_t), \tag{24}$$

$$X_t = \alpha_0 + \sqrt{v_t}\epsilon_t, \tag{25}$$

$$v_t = \alpha_1 + \rho v_{t-1} + \phi\epsilon_{t-1}^2. \tag{26}$$

where $\epsilon_t \sim N(0,1)$ with draws being independent over time. The GARCH process for X_t is the same as the process that TFP growth followed in the GARCH model. As in that model we assume that agents know the structure of the process for X_t but have to estimate the parameter values.[10]

The point of the nonlinear transformation is to allow TFP growth to have a non-normal distribution. Specifically TFP growth is conditionally log-normal. We chose this transformation for several reasons. First, it allows TFP growth to be either positively or negatively skewed and does not dictate the degree of skewness, so the data can determine this. Second, the transformation takes the whole real line as its domain so the underlying random variable X can be normal. If $b < 0$ (as it is when the model is calibrated) the

[10]We assume that agents know the parameters of the transformation in equation (24). We make this assumption for simplicity. When we calibrate the model we need to simulate it each time we adjust parameters and this requires estimating the process for X_t 400,000 times. This can be done reasonably efficiently using maximum likelihood techniques so that it is feasible. Estimating b and c as well would increase the complexity of the problem significantly and make it difficult to calibrate the model in a reasonable period of time.

skewness of TFP growth equals

$$-\Big[\exp\left((\log|b|)^2\sigma_x^2\right)+2\Big]\sqrt{\exp\left((\log|b|)^2\sigma_x^2\right)-1},$$

so TFP growth is negatively skewed and the magnitude of skewness is increasing in $|b|$.

Agents now receive signals about X instead of Δa:

$$z_{i,t-1} = X_t + \eta_{t-1} + \psi_{i,t-1}. \tag{27}$$

As before, $\eta_{t-1} \sim N(0, \sigma_\eta^2)$ is common to all forecasters and i.i.d. over time, and $\psi_{i,t-1} \sim N(0, \sigma_\psi^2)$ is i.i.d. across forecasters and over time. Note that if we think of these as signals about TFP growth, then the distribution of signals is non-normal.

To forecast TFP in period t, agents first use the history of X, X^{t-1}, to estimate the process for X and form prior beliefs about X_t. These prior beliefs will have a normal distribution with mean $E[X_t|X^{t-1}]$ and variance $v_{t-1} = V[X_t|X^{t-1}]$. Agents update their beliefs with their signals. The mean and posterior variance of agent i's beliefs are:

$$E_{i,t-1}[X_t] = (1-\omega_{t-1})E[X_t|a^{t-1}] + \omega_{t-1}z_{i,t-1},$$
$$V_{i,t-1}[X_t] = [v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}]^{-1},$$

where ω_{t-1} has the same definition as it had for the GARCH model (see Section 2).

The equations for labor and firm output are the same as for the linear model (equations (10) and (13) respectively). In this setup we cannot solve for $E_{i,t-1}[A_t]$ and $E_{i,t-1}[Q_t]$ in closed form, but we can do so numerically.

3.2 Calibration and simulation

The calibration procedure is the same as for the GARCH model, except that the model has two extra parameters. As we did for the GARCH model, we set $\gamma = 2$. We can normalize α_0 to be 1 in this model. This leaves seven parameters to calibrate: three parameters of the GARCH process (α_1, ρ and ϕ), two parameters for the transformation (b and c), and two parameters for signal noise (σ_η and σ_ψ). To calibrate these parameters we use the same six parameters that we used to calibrate the GARCH model as well as the correlation

Parameter	*Value*	*Target Moment*
γ	2	Frisch elasticity
α_0	1	Normalization
c	0.0959	GDP growth mean
b	-0.0339	GDP growth st. dev., skewness, kurtosis and corr$(\Delta q_t, \lvert\Delta q_{t+1} - \Delta q_t\rvert)$
α_1	2.70e-4	
ρ	0.6678	
ϕ	0.2769	
σ_η	0.0696	Average abs. error for GDP growth forecasts
σ_ψ	0.0229	Average st. dev. of GDP growth forecasts

Table 4: **Parameter values for the full model.** All moments of the data are for 1968q4–2011q4.

between Δq_t and $|\Delta q_{t+1} - \Delta q_t|$.

The calibration procedure and the method for simulating the model to compute results is exactly the same as for the GARCH model. The calibration is presented in Table 4. In this table, as well as Table 6, this model is referred to as the *full model.*

3.3 Properties of uncertainty shocks in the full model

The results are reported in Table 6. There are two columns of results for the full model. The column for which parameter learning is "on" presents results for the model exactly as it has been described. In the other column, for which parameter learning is "off," we assume that agents know the parameters of the GARCH process (α_0, α_1, ρ and ϕ). Comparing the two columns tells us how important parameter uncertainty is.

The model does well at generating higher-order, micro and macro uncertainty shocks. The model explains 75% of higher-order uncertainty shocks (the standard deviation of macro uncertainty) and all micro uncertainty shocks and a little more (the standard deviation of micro uncertainty). Note that the interaction of the mechanisms that generate micro and macro uncertainty shocks in the baseline and skewed models amplify each other in the full model: the standard deviations of micro and macro uncertainty in the full model are more than the sum of their values in the GARCH and homoskedastic skewed models (described in detail in Appendix C). The model also does reasonably well at replicating the relationship between micro and higher-order uncertainty and between higher-order uncertainty and the business cycle. It generates about half of the correlation between micro and higher-order uncertainty that is in the data and about two-thirds of the correlation between higher-order uncertainty and GDP growth. The weakest point in the uncertainty results is that micro uncertainty is not sufficiently countercyclical. Comparing results when parameter learning is on and off shows that the main role of parameter learning is to generate larger uncertainty shocks. Micro uncertainty varies less than half as much without parameter uncertainty and higher-order uncertainty varies by only about one-third as much.

The final part of the results is to test whether higher-order, micro and macro uncertainty have a significant relationship in the model, after controlling for the business cycle. To do this we regress micro uncertainty on higher-order uncertainty and GDP growth and then higher-order on macro uncertainty, in exactly the same way as we did for the data in Section 1. Each simulation of the model produces different series for the three variables in the regression, so we perform the regression for each of the 2000 simulations of the model and report the mean and standard deviation of each coefficient. The first set of results are presented in Table 5. For comparison we report the results for the data again. When we regress micro uncertainty on higher-order uncertainty without controlling for the business cycle, we find that the coefficient on higher-order uncertainty is very close to what we find for the data. When higher-order uncertainty deviates from trend my one additional percentage point, micro uncertainty deviates from trend by an additional 0.29 percentage points in the model and 0.28 percentage points in the data. When we control for GDP growth we find that micro and higher-order uncertainty still have a significant positive relationship in both the model and the data. The second set of results are in Table 5. The positive, significant coefficients on macro uncertainty reveal that the model generates a strong, positive relationship between macro and higher-order uncertainty, above and beyond what GDP and recessions can explain.

3.4 How Are Uncertainty Shocks Related to Disasters?

What really causes uncertainty to fluctuate in this model? One way to gain intuition about these fluctuations is to think through the effect of a change of variable from a normal to a skewed variable. When we calibrate this change of variable function (24), we find that the

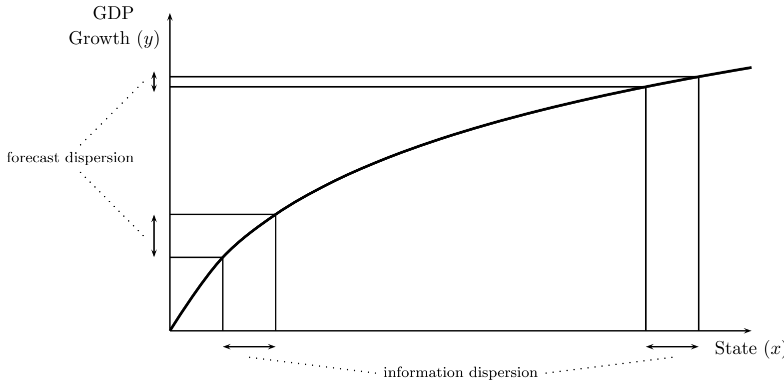

Figure 3: **Change of variable function and counter-cyclical forecast dispersion.** A given amount of uncertainty about X creates more uncertainty about TFP growth when $-X$ is low than it does when $-X$ is high. Likewise, a given amount of dispersion in beliefs about x creates more forecast dispersion in low TFP-growth times (low $-X$).

coefficient b is negative, meaning that the transformation is concave. A concave change of variable makes extreme, low realizations of TFP growth more likely and makes very high realizations less likely. In other words, a concave transformation creates a negatively-skewed variable. The concavity, and thus degree of negative skewness determines the probability of negative outlier events. So, one interpretation of this change of variable function is that it controls the level of disaster risk.

Figure 3 illustrates the effect of a concave change of variable on uncertainty and forecast dispersion. It plots a mapping from $-X$ into TFP growth, Δa. The slope of this curve is a Radon-Nikodym derivative. First consider macro uncertainty in this environment. For illustrative purposes, suppose that an agent has beliefs about X that are uniformly distributed. We can represent these beliefs by a band on the horizontal axis in Figure 3. If that band is projected onto the Δa-space, the implied uncertainty about (width of the band for) Δa depends on the state X. When $-X$ is high, the mapping is flat, and the resulting band projected on the Δa-axis is narrow. This means that uncertainty about TFP growth and therefore GDP growth is small, so macro uncertainty is small. When $-X$ is low the opposite it true: the band projected on the Δa axis is wider and uncertainty is high. If we now return to thinking about an agent with posterior beliefs about X that are normally distributed, using properties of the log-normal distribution we can express the variance of his beliefs about TFP growth as

$$V_{i,t-1}[\Delta a_t] = b^2(\exp(V_{i,t-1}[X_t]) - 1)\exp(2E_{i,t-1}[X_t] + V_{i,t-1}[X_t]).$$

From this formula we can see that the agents' uncertainty about Δa_t is increasing in his expected value of X_t, just as we illustrated for the case when the agent's beliefs are uniform.

This concave change of variable also explains why forecast dispersion and higher-order uncertainty vary. Suppose now that the bands on the x-axis represent the cross-sectional standard deviation of firms' forecasts of X. When $-X$ is high, the transformation in Figure 3 is flat so the dispersion of TFP growth forecasts is small. When $-X$ is low the opposite is true. Thus GDP growth forecast dispersion is high exactly when firms have a high degree of uncertainty about GDP growth and this occurs when GDP growth is low.

Now consider micro uncertainty. Recall that our measure of micro uncertainty is the IQR of firm output growth rates. When $-X$ is low and GDP growth is low, the dispersion of TFP growth forecasts is high. Dispersed forecasts create dispersion in labor choices, output and firm growth rates (micro uncertainty). In contrast, when $-X$ is high, forecast dispersion, output dispersion and thus micro uncertainty are all low. In sum, micro uncertainty is countercyclical.

In the benchmark model, we showed that our measure of micro uncertainty approximately measures the variance of a firm's beliefs about its own growth rate due its idiosyncratic signal (shock). We discussed that the measure is approximate because it picks up dispersion in output in period $t-1$, which is not a source of uncertainty for an individual firm. The link between our measure of micro uncertainty and the beliefs of an individual agent about his growth rate before receiving his signal is not as strong in this model. The reason is that in this model the effect of the idiosyncratic component of a firm's signal at time $t-1$ depends on the value of $X_t + \eta_{t-1}$: when $X_t + \eta_{t-1}$ is large the effect of the idiosyncratic component of the signal on firm output will be large because the transformation in Figure 3 is steep, and when $X_t + \eta_{t-1}$ is small the effect will be small. In this model our measure of micro uncertainty measures the variance in a firm's growth rate due to its idiosyncratic shock, conditional on $X_t + \eta_{t-1}$. This is different to the variance of the firm's beliefs about its growth rate before receiving its signal, because at that point in time $X_t + \eta_{t-1}$ is unknown.

This analysis teaches us that the level of disaster risk is closely related to the size of fluctuations in higher-order, micro and macro uncertainty. The higher disaster risk is, the more skewed the Δa distribution is, the more curvature the change of variable function has, and the larger are the fluctuations in all shades of uncertainty.

4 Conclusions

Is it plausible that micro, macro and higher-order uncertainty all rise in recessions, fuel the downturn, lower asset prices, and possibly trigger financial crisis? We argue that it is. When weak macro outcomes make agents re-assess their beliefs about the distribution of aggregate outcomes and in particular, about the skewness of that distribution, uncertainty of all shades are correlated with each other, volatile and counter-cyclical.

We started by working through standard arguments for why forecast dispersion and firm earnings dispersion might be related to macro uncertainty: When uncertainty is high, agents have imprecise prior beliefs, they weight their heterogeneous signals more. With more weight in beliefs, heterogeneous signals generate more dispersion in forecasts, actions and outcomes. That is, they generate higher-order and micro uncertainty. That mechanism is logical and straightforward, but it fails quantitatively.

When we allow TFP growth to have a conditional distribution that's skewed, uncertainty shocks are much larger, are countercyclical and are positively correlated with each other. Allowing for a skewed distribution allows the model to replicate the left tail of the empirical GDP growth distribution. Our results support the growing literature on uncertainty shocks and disaster risk by unifying it and by explaining why such shocks arise.

	Dependent variable: Micro uncertainty			
	Full Model		Data	
	(1)	(2)	(3)	(4)
Higher-order uncertainty	0.2900***	0.2756**	0.2807***	0.1338***
	(0.1003)	(0.1040)	(0.047)	(0.046)
GDP growth		−0.1842		−2.5489***
		(0.3533)		(0.357)
Obs	165	165	165	165
Sample	1968Q4–2009Q3			
	Dependent variable: Higher-order Uncertainty			
	(1)	(2)	(3)	(4)
Macro uncertainty	46.39***	47.20**	47.55***	47.68**
	(5.49)	(5.31)	(5.09)	(5.14)
Recession		−6.68		−3.22
		(4.41)		(4.57)
GDP growth			0.57**	0.47*
			(0.22)	(0.24)

Table 5: **Higher-order, micro and macro uncertainty shocks are correlated, after controlling for business cycles.** *Micro uncertainty* is the IQR of firm sales growth, defined in equation (2). *Higher-order uncertainty* is the standard deviation of GDP growth forecasts, defined in equation (4). The units for both series are the percentage deviation from trend. Since micro uncertainty is based on four-quarter growth rates and higher-order uncertainty is based on forecasts of quarterly growth rates, we regress MiU_t on the average of HU_t, HU_{t+1}, HU_{t+2} and HU_{t+3} so that the periods of the variables are comparable. Similarly *GDP growth* is 400 times the average of Δq_{t+1}, Δq_{t+2}, Δq_{t+3} and Δq_{t+1}. The units are the approximate annualized percentage. *Macro uncertainty* is av$(\sqrt{V_{i,t-1}[log(\Delta q_t)]})$. *Recession* is an indicator variables for the 5% of quarters with the lowest GDP growth rates. The coefficients for the model regressions are the averages across 2000 simulations and the numbers in parenthesis are the standard deviations of these estimates. For the data, standard errors are in parentheses. ***, ** and * denote significance at the 1%, 5% and 10% levels respectively. For the model significance levels are computed using the distribution of parameter estimates from the simulations.

References

Angeletos, M., F. Collard, and H. Dellas (2014): "Quantifying Confidence," MIT working paper.

Angeletos, M., and J. La'O (2014): "Sentiments," *Econometrica*, forthcoming.

Arellano, C., Y. Bai, and P. Kehoe (2012): "Financial Frictions and Fluctuations in Volatility," Federal Reserve Bank of Minneapolis working paper.

Bachmann, R., and C. Bayer (2012): "Investment Dispersion and the Business Cycle," University of Aachen working paper.

Bachmann, R., and G. Moscarini (2012): "Business Cycles and Endogenous Uncertainty," Yale University working paper.

Bansal, R., and I. Shaliastovich (2010): "Confidence Risk and Asset Prices," *American Economic Review*, 100(2), 537–41.

Basu, S., and B. Bundick (2012): "Uncertainty Shocks in a Model of Effective Demand," Boston College working paper.

Bianchi, F., C. Ilut, and M. Schneider (2012): "Ambiguous business cycles and uncertainty shocks," Stanford University working paper.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3).

Bloom, N., M. Floetotto, N. Jaimovich, I. Sapora-Eksten, and S. Terry (2012): "Really Uncertain Business Cycles," NBER working paper 13385.

Christiano, L., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*.

Comin, D., and S. Mulani (2006): "Diverging Trends in Aggregate and Firm Volatility," *The Review of Economics and Statistics*, 88(2), 374–383.

Comin, D., and T. Philippon (2005): "The Rise in Firm-Level Volatility: Causes and Consequences," *NBER Macroeconomics Annual*, pp. 167–201.

Decker, R., P. N. D'Erasmo, and H. Moscoso Boedo (2013): "Market Exposure and Endogenous Firm Volatility over the Business Cycle," University of Maryland working paper.

Fajgelbaum, P., E. Schaal, and M. Taschereau-Dumouchel (2013): "Uncertainty Traps," NYU working paper.

Gilchrist, S., J. Sim, and E. Zakrajšek (2013): "Uncertainty, financial frictions, and irreversible investment," Boston University and Federal Reserve Board working paper.

Keane, M., and R. Rogerson (2012): "Reconciling micro and macro labor supply elasticities: A structural perspective," Oxford working paper.

Nimark, K. (2014): "Man-Bites-Dog Business Cycles," *American Economic Review*, forthcoming.

Orlik, A., and L. Veldkamp (2014): "Understanding Uncertainty Shocks and the Role of the Black Swan," Working paper.

Pastor, L., and P. Veronesi (2012): "Uncertainty about Government Policy and Stock Prices," *Journal of Finance*, forthcoming.

Schaal, E. (2012): "Uncertainty, Productivity and Unemployment in the Great Recession," Federal Reserve Bank of Minneapolis working paper.

Van Nieuwerburgh, S., and L. Veldkamp (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53(4), 753–772.

A Proofs

Deriving forecasts Integrating over firm i's output (13), we get

$$Q_t = A_t \gamma^{1/(1-\gamma)} \int E_{i,t-1}[A_t]^{1/(\gamma-1)} di. \quad (28)$$

Using equation (12) and the fact that TFP growth forecasts are normally distributed, it follows from equation (28) that:

$$Q_t = A_t \gamma^{1/(1-\gamma)} \exp\left[\frac{1}{\gamma-1}\left(\log A_{t-1} + (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1}(\Delta a_t + \eta_{t-1}) + \frac{d_{t-1}^2}{2(\gamma-1)} + \frac{1}{2}V_{i,t-1}[\Delta a_t]\right)\right], \quad (29)$$

where d_{t-1}^2 is the cross-sectional variance of firms forecasts of TFP growth in period t, Δa_t. It follows from equation (11) that $d_{t-1}^2 = \omega_{t-1}^2 \sigma_\psi^2$. Separating the terms in equation (29) that are known at the end of period $t-1$ from those that are unknown,

$$Q_t = \Gamma_{t-1} \exp[f(\Delta a_t, \eta_{t-1})], \quad (30)$$

where $f(\Delta a_t, \eta_{t-1}) \equiv \Delta a_t + (\frac{\omega_{t-1}}{\gamma-1})(\Delta a_t + \eta_{t-1})$ and

$$\Gamma_{t-1} \equiv \gamma^{1/(1-\gamma)} A_{t-1} \exp\left[\frac{1}{\gamma-1}\left(\log A_{t-1} + (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \frac{d_{t-1}^2}{2(\gamma-1)} + \frac{V_{i,t-1}[\Delta a_t]}{2}\right)\right].$$

Under agent i's beliefs at the end of period $t-1$, $\Delta a_t + \eta_{t-1}$ is normally distributed. Therefore $f(\Delta a_t, \eta_{t-1})$ is normally distributed under these beliefs so we can express agent i's forecast of period t GDP as

$$E_{i,t-1}[Q_t] = \Gamma_{t-1} \exp\left(E_{i,t-1}[f(\Delta a_t, \eta_{t-1})] + \frac{1}{2}V_{i,t-1}[f(\Delta a_t, \eta_{t-1})]\right). \quad (31)$$

To evaluate $E_{i,t-1}[f(\Delta a_t, \eta_{t-1})]$ we need the mean of agent i's posterior belief about $\Delta a_t + \eta_{t-1}$. This can be computed by Bayes' law:

$$E_{i,t-1}[\Delta a_t + \eta_{t-1}] = \frac{(v_{t-1} + \sigma_\eta^2)^{-1} E[\Delta a_t|A^{t-1}] + \sigma_\psi^{-2} z_{i,t-1}}{(v_{t-1} + \sigma_\eta^2)^{-1} + \sigma_\psi^{-2}}.$$

Therefore

$$E_{i,t-1}[f(\Delta a_t, \eta_{t-1})] = (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1} z_{i,t-1} + \left(\frac{\omega_{t-1}}{\gamma-1}\right)\left(\frac{(v_{t-1} + \sigma_\eta^2)^{-1} E[\Delta a_t|A^{t-1}] + \sigma_\psi^{-2} z_{i,t-1}}{(v_{t-1} + \sigma_\eta^2)^{-1} + \sigma_\psi^{-2}}\right). \quad (32)$$

Let $x_t \equiv [\Delta a_t - E[\Delta a_t|A^{t-1}], \;\; \eta_{t-1}]'$. Then the variance term in equation (31) is

$$V_{i,t-1}[f(\Delta a_t, \eta_{t-1})] = \left[\left(1 + \tfrac{\omega_{t-1}}{\gamma-1}\right) \quad \tfrac{\omega_{t-1}}{\gamma-1}\right] V_{i,t-1}[x_t] \left[\left(1 + \tfrac{\omega_{t-1}}{\gamma-1}\right) \quad \tfrac{\omega_{t-1}}{\gamma-1}\right]' \tag{33}$$

where

$$V_{i,t-1}[x_t] = \begin{bmatrix} v_{t-1} & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix} - \begin{bmatrix} v_{t-1} \\ \sigma_\eta^2 \end{bmatrix} (v_{t-1} + \sigma_\psi^2 + \sigma_\eta^2)^{-1} \begin{bmatrix} v_{t-1} & \sigma_\eta^2 \end{bmatrix}. \tag{34}$$

Together equations (31), (32), (33) and (34) define agent i's forecast of period t GDP.

B Calibration and Simulation Details for the GARCH Model

Calibration To calibrate the six parameters of the model excluding γ we use simulated method of moments with the identity weighting matrix. Since agents in the model are assumed to not know the parameters of the GARCH process, it is important that the model is simulated for the same length of time as the data that agents would have access to in the economy. We assume that agents have access to post-war data so each simulation of the model is for 259 periods, corresponding to 1947Q2 to 2011Q4. The data that we're using for calibration starts in 1968Q4, so agents who are making decisions for this period have 86 periods of data to use when they are first estimating the TFP process. To calculate the moments of the model we simulate it 2000 times, calculate the moments of the model for each simulation and then average each moment across the simulations. Each simulation consists of: simulating the TFP process for 259 periods; estimating the TFP process for periods 87–259 using only the data available up to, but not including, the relevant period; using these estimations of the process to construct prior beliefs about TFP growth for periods 87–259; computing the path for TFP using the simulated path for TFP growth, normalizing TFP in the first period to be 1; computing GDP growth making use of equation (28); and computing GDP growth forecast dispersion and the average absolute error of GDP growth forecasts using equations (3), (4), (31) and (23).

In each period of the model, the only source of heterogeneity amongst agents is the realization of their idiosyncratic signal noise, ψ_{it}. Since this random variable is normally distributed and draws are i.i.d. across agents, we can use Gaussian quadrature to compute aggregate moments.

Simulation The simulation method for computing the results is mostly the same as for the calibration simulations. One difference is that we can't use a Gaussian quadrature to compute micro uncertainty because it is not a moment that requires integrating over firms. Therefore we explicitly simulate the model for 2000 firms to compute the results for micro uncertainty. Again, we do this 2000 times and average the results over simulations.

C How well can skewness explain uncertainty shocks without stochastic volatility?

The preferences (5) and technology (6) are the same as in the GARCH model. The TFP growth process is now:

$$\Delta a_t = c + b\exp(X_t), \tag{35}$$

$$X_t \sim N(\mu_x, \sigma_x^2), \tag{36}$$

and draws of X_t are i.i.d. over time. The point of the transformation is to allow TFP growth to have a non-normal distribution. Specifically the TFP growth distribution is a linear transformation of a log-normal distribution. We have chosen this transformation for several reasons. First, it allows TFP growth to be either positively or negatively skewed and does not dictate the degree of skewness, so the data can determine this. Second, the transformation takes the whole real line as its domain so the underlying random variable X can be normal. If $b < 0$ (as it is when the model is calibrated) the skewness of TFP growth equals

$$-\left[\exp\left((\log|b|)^2\sigma_x^2\right) + 2\right]\sqrt{\exp\left((\log|b|)^2\sigma_x^2\right) - 1},$$

so TFP growth is negatively skewed and the magnitude of skewness is increasing in $|b|$.

Agents now receive signals about X instead of Δa:

$$z_{i,t-1} = X_t + \eta_{t-1} + \psi_{i,t-1}. \tag{37}$$

As before, $\eta_{t-1} \sim N(0, \sigma_\eta^2)$ is common to all forecasters and i.i.d. over times, and $\psi_{i,t-1} \sim N(0, \sigma_\psi^2)$ is i.i.d. across forecasters and over time. Note that if we think of these as signals about TFP growth, then the distribution of signals is non-normal. We assume that agents know the parameters of the distribution of X_t, so their prior belief about X_t will be that $X_t \sim N(\mu, \sigma_x^2)$ every period. By making priors homoskedastic we eliminate the effect that variation in the precision of the prior has on the dispersion of agents' beliefs about TFP growth. The mean of agent i's beliefs about X_t after receiving her signal at the end of period $t-1$ is

$$E_{i,t-1}[X_t] = (1-\omega_{t-1})\mu_x + \omega_{t-1}z_{i,t-1}. \tag{38}$$

where ω_{t-1} is defined as it was for the GARCH model and the variance of the prior belief is $v_{t-1} = \sigma_x^2$.

The equations for labor, firm output and aggregate output are the same as for the linear model (equations (10), (13) and (28) respectively). In this setup we cannot solve for $E_{i,t-1}[A_t]$ and $E_{i,t-1}[Q_t]$ in closed form, but we can do so numerically.

Calibration and simulation The calibration procedure for this model is similar to that used for the GARCH model. We again set $\gamma = 2$ so that the Frisch elasticity of labor supply is 1. The other parameters of the model are four parameters for the distribution of TFP growth (b, c, μ_x and σ_x) and the two parameters for public and private signal noise (σ_η and σ_ψ). We can normalize μ_x to 1 since for any distribution of TFP growth that can be written in the form of equations (35) and (36) there is a continuum of pairs $(|b|, \mu_x)$ that can generate that distribution. To see this note that we can write TFP growth as

$$\Delta a_t = c + \text{sign}(b) \times \exp(\log |b| \mu_x + \log |b| \sigma_x \epsilon_t)$$

where $\epsilon_t \sim N(0, 1)$. We calibrate the remaining five parameters of the model to match five moments of the data. These are the same moments as we used for calibrating the GARCH model, except that we drop the kurtosis of GDP growth.

Unlike for the GARCH model, we do not need to replicate the sample length of the data when calibrating this model. There are two reasons for this: agents know the parameters of the model so parameter learning is not relevant, and TFP growth is drawn i.i.d. each period so there is no time dependence in the model. Using Gaussian quadrature for numerical integration over the distribution of firms and the distribution of state variables for the economy, we can actually compute the moments of this model without simulating it. To choose parameter values we minimize the sum of squared percentage deviations of the calibration moments for the model from those moments of the data.

The calibration is presented in Table 7. In this table, as well as Table 6, this model is referred to as the *skewed model.* As indicated earlier, $b < 0$ so that the model can replicate the negative skewness of GDP growth in the data. Also note that there is no public signal noise under the calibration. This is because private signal noise alone generates enough errors in GDP growth forecasts that no public signal noise is needed to explain the average absolute error of GDP growth forecasts in the data.

We compute the same results for this model as we did for the GARCH model. We can compute most of these using Gaussian quadrature and avoid simulating the model. The exception to this is micro uncertainty. To compute this we use a simulation of 5000 firms for 10,000 periods. In this simulation we compute micro uncertainty each period, as well as GDP growth and macro uncertainty since we want to know the correlation of these moments with micro uncertainty.

Results We can see from the standard deviations that are reported in the panels (a) and (b) of Table 6 that this model generates uncertainty shocks, but they are small relative to the data. The model does however generate countercyclical uncertainty as was predicted, which the GARCH model was not able to do, and generates a correlation between micro and macro uncertainty that is nearly half of that in the data. Also note that this model can match the skewness of GDP growth and that our measure of macro uncertainty is a good measure of uncertainty that individual agents have about GDP growth (these moments are nearly perfectly correlated). Since both the skewed model and the GARCH model hit some moments of the data, but not others, the final model in the main text puts both mechanisms together to achieve an outcome that resembles the data.

	Benchmark model		Skew homskd.	Full model		Data
Parameter learning	Off	On	Off	Off	On	
(a) Micro uncertainty						
Mean	0.37%	0.36%	0.54%	0.37%	0.36%	18.59%
Std.	2.50%	4.66%	3.69%	6.50%	14.70%	11.58%
Corr. with GDP growth	−0.01	−0.00	−0.20	−0.14	−0.07	−0.52
Corr. with Higher-order Unc.	0.01	0.16	0.19	0.23	0.21	0.43
Period	1962Q1–2009Q3					
(b) Higher-order uncertainty						
Mean*	1.54%	1.53%	1.52%	1.54%	1.47%	1.54%
Std.	0.13%	5.88%	4.03%	8.61%	23.58	31.13%
Corr. with GDP growth	0.00	−0.00	−0.85	−0.43	−0.17	−0.28
Corr. with Macro Unc.	0.96	0.96	0.98	0.97	0.95	
Period	1968Q4–2011Q4					
(b) Macro uncertainty						
Mean	2.81%	2.79%	2.60%	2.92%	2.81%	
Std.	0.00%	0.10%	0.11%	0.26%	0.49%	
Corr. with GDP growth	−0.02	0.00	−0.90	−0.46	−0.26	
Corr. with Micro Unc.	−0.01	0.15	0.20	0.23	0.21	
(c) GDP growth forecasts						
Av. absolute error*	2.23%	2.22%	2.12%	2.23%	2.17%	2.23%
Period	1968Q4–2011Q4					
(d) GDP growth						
Mean*	2.71%	2.71%	2.71%	2.71%	2.71%	2.71%
Std.*	3.39%	3.39%	3.58%	3.39%	3.39%	3.39%
Skewness*	0.01	0.01	−0.32	−0.30	−0.30	−0.32
Kurtosis†	4.93	4.91	3.19	5.25	5.25	4.98
$\text{corr}(\Delta q_t, \lvert\Delta q_{t+1} - \Delta q_t\rvert)^{\ddagger}$	−1.92e-4	2.88e-4		−0.06	−0.06	−0.06
Period	1968q4–2011q4					

Table 6: **Simulation results and data counterparts.** *Micro uncertainty* (MiU) is the IQR of firm sales growth, defined in equation (2). *Higher-order uncertainty* is the standard deviation of GDP growth forecasts, defined in equation (4). The units for the means these series are the annualized percentage ($100 \times MiU$ and $400 \times MaU$). The standard deviations and correlations in panels (a) and (b) are computed using the detrended and scaled series for micro and higher-order uncertainty. The units for the standard deviations are the percentage deviation from trend. Since micro uncertainty is based on four-quarter growth rates and macro uncertainty is based on forecasts of quarterly growth rates, the correlation between them is computed as the correlation between MiU_t and the average of MaU_t, MaU_{t+1}, MaU_{t+2} and MaU_{t+3}. Similarly the correlation of micro uncertainty with GDP growth is the correlation of MiU_t with the average of Δq_{t+1}, Δq_{t+2}, Δq_{t+3} and Δq_{t+1}. The units for GDP growth are the approximate annualized percentage: $400\Delta q_t$. The units for GDP growth forecasts are the same. *Macro uncertainty* is $\text{av}(\sqrt{V_{i,t-1}[log(\Delta q_t)]})$ and its units are the annualized percentage. * denotes moments that are used as calibration targets for all models. † denotes moments that are used to calibrate the GARCH and full models only. ‡ denotes moments that are used to calibrate only the full model.

Parameter	*Value*	*Target Moment*
γ	2	Frisch elasticity
μ_x	1	Normalization
c	0.0439	GDP growth mean
b	-0.0147	GDP growth st. dev. and skewness
σ_x	0.1470	
σ_η	0.0000	Average abs. error for GDP growth forecasts
σ_ψ	0.2034	Average st. dev. of GDP growth forecasts

Table 7: **Parameter values for the skewed, homoskedastic model.** All moments of the data are for 1968q4–2011q4.