

THE COMMON ORIGIN OF UNCERTAINTY SHOCKS

Nicholas Kozeniauskas, Anna Orlik and Laura Veldkamp
New York University and Federal Reserve Board*

June 22, 2016

Abstract

Various types of uncertainty shocks can explain many phenomena in macroeconomics and finance. But does this just amount to inventing new, exogenous, unobserved shocks to explain challenging features of business cycles? This paper argues that three conceptually distinct fluctuations, all called uncertainty shocks, have a common origin. Specifically, we propose a mechanism that generates *micro uncertainty* (uncertainty about firm-level shocks), *macro uncertainty* (uncertainty about aggregate shocks) and *higher-order uncertainty* (disagreement) shocks from a common origin and causes them to covary, just as they do in the data. When agents use standard maximum likelihood techniques and real-time data to re-estimate parameters that govern the probability of disasters, the result is that micro, macro and higher-order uncertainty fluctuate and covary just like their empirical counterparts. Our findings suggest that time-varying disaster risk and the many types of uncertainty shocks are not distinct phenomena. They are outcomes of a quantitatively plausible belief updating process.

A recent literature starting with Bloom (2009) demonstrates that uncertainty shocks can explain business cycles, financial crises and asset price fluctuations with great success (e.g. Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012), Ordonez (2011) and Pastor and Veronesi (2012)). But what is the nature of these shocks and where do they come from? The nature of uncertainty shocks 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

*Please send comments to nic.j.koz@nyu.edu. We are grateful to Virgiliu Midrigan, Simon Mongey, and Tom Sargent for useful discussions and suggestions. We thank seminar participants at the SED meetings in Warsaw and the AEA meetings in San Francisco. 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.

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

itself in an increase in the cross-sectional difference in firm outcomes.² We call this *micro uncertainty*. *Higher-order uncertainty* describes the uncertainty about others' beliefs that arises when forecasts differ.³ These three types of uncertainty are conceptually different. While macro uncertainty comes from aggregate shocks, micro uncertainty depends on firm-specific shocks, and higher-order uncertainty arises from private signal noise. In principle all of these shocks could be independent. But they are not: these different types of uncertainty strongly covary, so much so that they are sometimes conflated. The fact that all three are not obviously related and yet covary strongly suggests that they may have a common cause. This paper is a search for the common cause of uncertainty shocks.

Our main argument is that if firms re-estimate a model with disaster risk each period before choosing inputs and producing, their macro, micro and higher-order uncertainty will covary in a realistic way. The reason that disaster risk can affect all three uncertainty measures originates from the fact that disasters arise infrequently, making their probability difficult to assess and the scope for disagreement large. When the risk of disasters rises, these errors and disagreements are amplified. Less accurate forecasts raise macro uncertainty. Divergent forecasts create higher-order uncertainty. Firms with divergent forecasts choose different inputs and obtain different outputs. This rise in the dispersion of firm output shows up as higher micro uncertainty. All three forms of uncertainty and their covariance can be explained in a unified framework that ties all three to disaster risk. This unification brings us one step closer to understanding what causes business cycle fluctuations.

Section 1 begins with measurement. It defines our measures of micro, macro and higher-order uncertainty and documents the strong positive correlation between them. Then it

²For micro uncertainty shocks see, for example, Arellano, Bai, and Kehoe (2012), Christiano, Motto, and Rostagno (2014), Gilchrist, Sim, and Zakrajsek (2013), Schaal (2012), Bachmann and Bayer (2013) and Bachmann and Bayer (2012). The last two papers 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 micro and macro uncertainty shocks.

³On the role of higher-order uncertainty see Angeletos and La'O (2014), Angeletos, Collard, and Dellas (2014) and Benhabib, Wang, and Wen (2015).

investigates the hypothesis that these types of uncertainty covary just because they are all countercyclical. We find that there is more to it than this because the correlation between them holds up even after controlling for the business cycle.

To explain why re-assessing tail risk can cause many forms of uncertainty to comove, Section 2 presents a production economy with firms that are uncertain about aggregate productivity (TFP). The firms face exogenous changes in productivity and its variance, which changes all three types of uncertainty endogenously. The model has three key features: First, stochastic volatility of TFP growth is essential because it is the source of our uncertainty fluctuations. Second, heterogeneous signals create heterogeneity amongst firms which is necessary to generate higher-order uncertainty and micro uncertainty. We can either assume heterogeneous beliefs or provide heterogeneous information. The latter generates more insight and endogenous fluctuations. The third ingredient is non-normal innovations. This is what creates a role for time-varying disaster risk. Normal distributions have thin tails, which makes disasters incredibly unlikely. In order to have meaningful changes in tail probabilities, we use a stochastic process that allows TFP to have conditional skewness and thick tails.

Each period our agents observe past data, re-estimate the parameters of this non-normal model and then receive private signals to update their beliefs about TFP growth. On the basis of their beliefs firms choose their labor input for production and forecast GDP growth. Since TFP growth has stochastic volatility and agents must re-estimate the model each period, the precision of agents' prior beliefs, and macro uncertainty, varies over time. This causes the weight that they place on their heterogeneous signals to vary causing changes in the dispersion of their forecasts and labor choices. This shows up as micro and high-order uncertainty shocks. This mechanism generates correlation between the three types of uncertainty, but can't explain their countercyclicality or magnitude. This is where disaster risk comes in. Because disaster probabilities are difficult to assess, a rise in disaster risk creates both uncertainty about aggregate outcomes (macro uncertainty) and disagreement.

When firms disagree, they make different input choices and have heterogeneous outcomes (micro uncertainty). They also make different forecasts (higher-order uncertainty). Thus, it is that fact that disasters are rare and difficult to predict that drives all three forms of economic uncertainty.

The strength of the disaster risk mechanism comes from how it creates disagreement about tail risk.⁴ Specifically, because tail probabilities are so sensitive to changes in skewness, small differences in skewness estimates produce large disagreement about disaster risk. Thus accounting for the skewness we see in the data causes higher-order, macro and micro uncertainty to fluctuate more, covary more negatively with the business cycle, and covary more positively with each other. Thus the message is that, while types of uncertainty are not obviously related, real-time estimation of disaster risk can produce fluctuations in the many types of uncertainty, in a way that explains the uncertainty data.

Related literature There is a growing literature studying the importance of uncertainty shocks for understanding business cycles and financial markets. In his seminal paper Bloom (2009) showed that various measures of uncertainty are countercyclical and studied the ability of uncertainty shocks to explain business cycle fluctuations. Several other papers have further investigated uncertainty shocks as a driving force for business cycles including Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012), Basu and Bundick (2012), Bianchi, Ilut, and Schneider (2012), Arellano, Bai, and Kehoe (2012), Christiano, Motto, and Rostagno (2014), Gilchrist, Sim, and Zakrajšek (2013), Schaal (2012), Bachmann and Bayer (2012) and Bachmann and Bayer (2013). A related strand of literature studies the impact of uncertainty shocks on asset prices (see, for example, Bansal and Shaliastovich (2010) and Pastor and Veronesi (2012)). Our paper complements this literature by asking what are the origins of the uncertainty shocks that most of these papers take as exogenous are.

⁴We use the term disagreement here to mean differing Bayesian estimates. It does not mean differences in priors.

A few recent papers also 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’Erasmus, 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 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. Benhabib, Liu, and Wang (2016) show how endogenous information acquisition can generate countercyclical uncertainty. In Van Nieuwerburgh and Veldkamp (2006) and Fajgelbaum, Schaal, and Taschereau-Dumouchel (2013), less economic activity generates less data, which increases uncertainty. Ludvigson, Ma, and Ng (2016) use statistical projection methods to argue that output fluctuations can cause uncertainty fluctuations or the other way around, depending on the type of uncertainty. Our paper differs because it explains why uncertainty can be explained by the business cycle and why dispersion across firms and forecasters is connected to uncertainty about aggregate outcomes.

A handful of papers have explored the importance of learning from past observations and its impact on second moments (Adam, Marcet, and Nicolini, 2016). While the idea of learning about tail risk has its origins in Orlik and Veldkamp (2014), this paper focuses on heterogeneity, a topic that the previous, representative agent forecasting model could not address. The model in this paper pushes that idea further: It adds heterogeneous information and then builds a production economy with these dispersed beliefs at its core. The new insight is that time-varying disaster risk, when paired with heterogeneous signals, can explain fluctuations in belief heterogeneity and firm output dispersion as well. With its macroeconomic predictions and realistic forecast dispersion, this model is more amenable

to empirical evaluation and can explain why different types of uncertainty covary.

Our work on cyclical uncertainty fluctuations complements the literature studying long run trends in aggregate and firm level volatility. In the past 30 years, aggregate volatility has declined while firm volatility has increased (Comin and Philippon, 2005). Comin and Mulani (2006) argue that these long-run trends come from a shift from general-purpose technology to more specific technology development. Our paper explains the cyclical fluctuations around these long-run trends.

1 The Empirical Puzzle

Since the objective of the paper is to uncover the link between the different types of uncertainty, we start in this section by defining our three measures of uncertainty and documenting some of their empirical properties. We document that the three types of uncertainty are positively correlated, that the relationships are statistically significant and that their covariance exceeds that which the business cycle alone can explain.

Throughout the paper we measure micro uncertainty with the cross-sectional interquartile range (IQR) of firm sales growth.⁵ The growth rate of firm i in quarter t is computed as

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

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

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

The data comes from Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012) and covers 1962Q1–2009Q3. It is based on observations for all public firms with at least 100

⁵We use the IQR in order to make contact with an existing literature, e.g., Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012), which uses the IQR to measure micro uncertainty.

quarters of data in Compustat between 1962 and 2010. The sample contains 2,465 firms. Since Comin and Mulani (2006) have already explained the long-run uncertainty trends, we focus on cyclical fluctuations by removing trends. We detrend micro uncertainty using a HP filter with smoothing parameter equal to the standard value for quarterly data (1600) and compute percentage deviations of the data from trend. The data is presented in Figure 1.

We measure higher-order uncertainty with the cross-sectional standard deviation of real GDP growth forecasts from the Survey of Professional Forecasters. These are one-period-ahead forecasts of real GDP: forecasts of quarter $t + 1$ real GDP made after the end of quarter t . From here on, we use the term GDP to refer to real GDP. Let $Q_t \equiv \int Q_{it} di$ denote aggregate GDP in quarter t . Let \mathcal{I}_{it} be the information set of agent i at the end of quarter t , and define GDP growth as $\Delta q_t = 400(\log Q_t - \log Q_{t-1})$ so that the units are annualized percentage growth. We compute approximate forecasts of GDP growth using the GDP forecast data as follows:

$$E[\Delta q_{t+1} | \mathcal{I}_{it}] \equiv 400(\log(E[Q_{t+1} | \mathcal{I}_{it}]) - \log Q_t). \quad (3)$$

We compute approximate forecasts because the data only provides forecasts of the level of GDP, not the growth rate of GDP. Higher-order uncertainty for quarter t is measured with the cross-sectional standard deviation of the forecasts in (3):

$$HU_t \equiv \sqrt{\frac{1}{N_t} \sum_{i \in \Omega_t} (E[\Delta q_{t+1} | \mathcal{I}_{it}] - \bar{E}_t)^2} \quad (4)$$

where $\bar{E}_t = 1/N_t \sum_i E[\Delta q_{t+1} | \mathcal{I}_{it}]$ is the average growth forecast, i indexes forecasters, Ω_t is the set of forecasters at the end of period t and N_t is the number of forecasters in this set.⁶ The higher-order uncertainty series covers 1968Q3–2011Q3. We detrend this series using the same method as we used for the micro uncertainty series in order to

⁶The average number of forecasters in SPF in a quarter is 41, with a standard deviation of 17

remove variation in the series at frequencies that we do not seek to explain. When we want to assess the relationship between higher-order uncertainty and micro uncertainty we need to make an adjustment for the fact that micro uncertainty is based on four quarter growth rates while higher-order uncertainty is based on one quarter growth rates. The adjustment we make is to average higher-order uncertainty over four quarters: $\widehat{HU}_t \equiv \frac{1}{4} \sum_{i=0}^3 HU_{t+i}$. \widehat{HU}_t is comparable in terms of timing to MiU_t . We also focus on the percentage deviation from trend for this series, which is presented in Figure 1. When evaluating the relationship between micro uncertainty and GDP growth we make a similar adjustment by using $\widehat{\Delta q}_{t+1} \equiv \frac{1}{4} \sum_{i=1}^4 \Delta q_{t+i}$ for GDP growth. Our real GDP data is from the BEA.

Conceptually we think of macro uncertainty of agent i at the end of period t as the standard deviation of that agents' beliefs about GDP growth in the next period: $\sqrt{V[\Delta q_{t+1} | \mathcal{I}_{it}]}$. To get an aggregate measure of macro uncertainty we average macro uncertainty across agents:

$$MaU_t \equiv \frac{1}{N_t} \sum_{i \in \Omega_t} \sqrt{V[\Delta q_{t+1} | \mathcal{I}_{it}]} \quad (5)$$

As for higher-order uncertainty, when we're comparing macro uncertainty to micro uncertainty we need to average it over four quarters. Let $\widehat{MaU}_t \equiv \frac{1}{4} \sum_{i=0}^3 MaU_{t+i}$ so that \widehat{MaU}_t is comparable to MiU_t in terms of timing. A difficulty with using this definition of macro uncertainty is that we don't have a good measure of it in the data. To measure this directly we would need detailed information on the probability that forecasters in the economy place on various outcomes for GDP growth, but such information is not available. Therefore to make progress empirically we use a proxy for macro uncertainty that we take from Bloom (2009). This proxy is based on the CBOE's VXO which utilizes the implied volatility of the S&P100 from options prices.⁷ This quarterly series provides our empirical

⁷This proxy is constructed as follows. For 1986 onwards the series is the CBOE's VXO. This series is the expected variance of the S&P100 over the next 30 days, as implied by options prices. Since the VXO isn't available prior to 1986 a normalized version of the realized volatility of the S&P500 is used for the pre-1986 period. Specifically Bloom takes the monthly standard deviation of the daily S&P500 index and normalizes the series so that for the period in which it overlaps with the VXO (1986 onwards) the two series

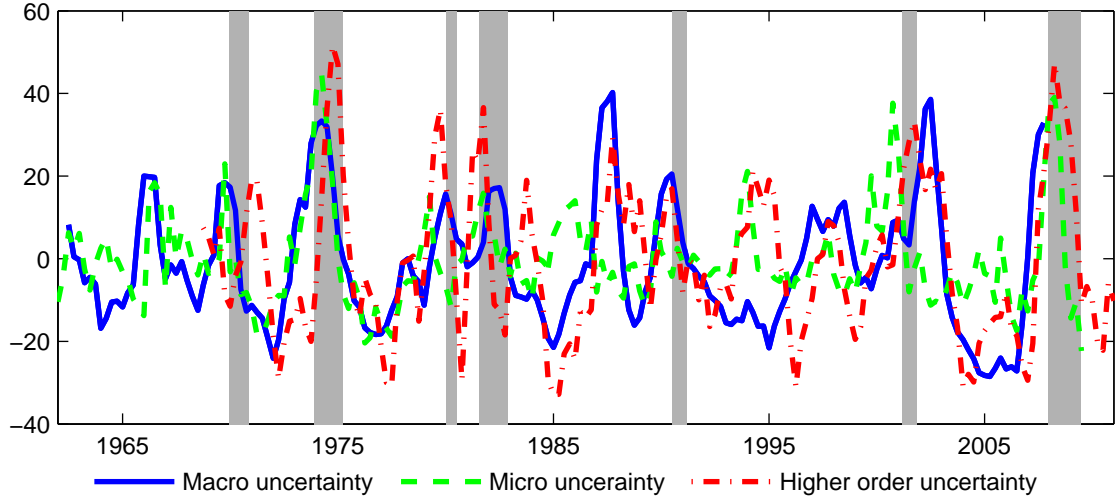


Figure 1: **Shades of uncertainty.** *Micro uncertainty* is MiU_t , *higher-order uncertainty* is \widehat{HU}_t and the *macro uncertainty* series is the empirical proxy for \widehat{MaU}_t . All series are presented as the percentage deviation from trend.

proxy for MaU_t from which we can compute \widehat{MaU}_t as described above. The MaU_t series covers 1962Q2–2008Q2. As for the other uncertainty series we focus on the percentage deviation from trend. The percentage deviation from trend of the empirical proxy for \widehat{MaU}_t is presented in Figure 1.

Having defined the three types of uncertainty, we now turn to discussing the relationship between them. 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 which is significant at the 1% level—see column (1) of the first panel of Table 1. The coefficient of 0.281 means that when higher-order uncertainty deviates from trend by one additional percentage point, micro uncertainty deviates from trend by an additional 0.281 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. The second and third panels of Table 1 shows that similar results hold for the relationships between macro uncertainty and higher-order uncertainty, and between

have the same mean and variance. This produces a monthly series and we follow Bloom in averaging the series across the months of each quarter to get a quarterly series.

<i>Micro Uncertainty</i>	(1)	(2)
Higher-order uncertainty	0.281*** (0.047)	0.134*** (0.046)
GDP growth		-2.549*** (0.357)
Obs	165	165
Sample	1968Q3–2009Q3	
<i>Higher-order uncertainty</i>	(1)	(2)
Macro uncertainty	0.350*** (0.110)	0.277** (0.111)
GDP growth		-1.949*** (0.706)
Obs	160	160
Period	1968Q3–2008Q2	
<i>Macro uncertainty</i>	(1)	(2)
Micro uncertainty	0.461*** (0.102)	0.252** (0.108)
GDP growth		-2.301*** (0.533)
Obs	182	182
Period	1962Q2–2007Q3	

Table 1: **Uncertainty shocks are correlated in the data, even after controlling for business cycles.** Micro uncertainty is MiU_t . Higher-order uncertainty is \widehat{HU}_t in the first panel and HU_t in the second panel. Macro uncertainty is MaU_t in the second panel and \widehat{MaU}_t in the third panel. GDP growth is Δq_t in the second panel and $\widehat{\Delta q}_t$ in the other panels. We multiple GDP growth by 400 so that its units are the annualized percentage growth rate. Standard errors are in parentheses. ***, ** and * denote significance at the 1%, 5% and 10% levels respectively.

macro and micro uncertainty. Column (1) shows that there is a positive and statistically significant relationship between both of these pairs of uncertainty at the 1% level. Macro uncertainty is also countercyclical, with a correlation with GDP growth of -0.26 . Several other papers starting with Bloom (2009) have documented similar facts.

One reason why the three types of uncertainty could comove is that they are all driven by the business cycle. We test this hypothesis by running the three regressions controlling for the business cycle with real GDP growth. The results are presented in column (2) of Table 1. The main result is that the types of uncertainty have a significant positive relationship even after controlling for the business cycle. Taken together, these findings suggest that micro, macro and higher-order uncertainty share some other common source of variation.

2 Model

We want to explore why micro, macro and high-order uncertainty comove. To do this, we feed in exogenous acyclical stochastic volatility shocks that mechanically move macro uncertainty and then look to see how and why micro and higher-order uncertainty respond. Since micro uncertainty is measured in the data using firm growth rates, we have firms making production decisions in the model. Since higher-order uncertainty is about differences in beliefs, we need to give these firms heterogeneous information on which to form different beliefs. This dispersed information is the source (and the only source) of firm heterogeneity. That is a limitation of the model that we'll return to.

In order for the transmission of macro to micro and higher-order uncertainty to be strong enough to resemble the data, we need firms to use their heterogeneous information to re-evaluate disaster risk each period. This amplifies changes in uncertainty in part because uncertainty is a variance and variance is very sensitive to changes in the probability of events far away from the mean. Re-evaluating disaster risk requires a non-normal model because normal distributions have infinitesimal probabilities of disasters (thin tails). We

introduce non-normality by taking a GARCH process for the state and then subjecting it to a simple non-linear transformation. We introduce learning about this distribution by allowing firms to re-estimate the parameters of the distribution each period. In what follows, we set up this model formally, describe its solution, and then describe a similar model without skewness. We use that model, which we'll call the *normal model*, later to isolate the role that the skewed distribution plays in the results.

A different way to introduce disaster risk would be to have a discrete disaster state and a regime-switching process. But this type of setup won't work for our purposes. For disaster risk to constantly fluctuate, we need for events in normal times to cause a reassessment of disaster risk. Thus normal events must be informative about disaster probabilities. Normal events can only be informative if the same parameters that govern ordinary events also govern disaster events. Discrete state disaster models have separate parameters governing ordinary and disaster states. In such a model disaster probability and severity estimates only experience shocks when disasters are realized. Thus a theory that encompasses time-varying disaster risk can either assume that such risk varies exogenously, or can link it to non-disaster outcomes with a non-normal distribution that covers the whole state space. To better understand why disaster risk varies we take the latter approach.

2.1 Environment

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} :

$$U_{it} = Q_{it} - L_{it}^\gamma \tag{6}$$

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

$$Q_{it} = A_t L_{it}. \tag{7}$$

Aggregate output is $Q_t \equiv \int Q_{it} di$ and GDP growth is $\Delta q_t \equiv 400(\log Q_t - \log Q_{t-1})$. We define GDP growth in this way in anticipation of the fact that we will calibrate to quarterly data so that the units are the annualized percentage.

A key feature of our model is that the growth rate of productivity at time t , $\Delta a_t \equiv \log(A_t) - \log(A_{t-1})$ has stochastic volatility. This stochastic volatility process is the source of our uncertainty shocks and takes the form of a general autoregressive conditional heteroskedasticity (GARCH) model. We use GARCH because it is simple to estimate and makes it feasible to add more richness to the rest of the model. Then we use an exponential transformation of this GARCH process to introduce non-normality and disaster risk:

$$\Delta a_t = c + b \exp(-X_t), \quad (8)$$

$$X_t = \sigma_t \epsilon_t, \quad (9)$$

$$\sigma_t^2 = \alpha + \rho \sigma_{t-1}^2 + \phi \sigma_{t-1}^2 \epsilon_{t-1}^2. \quad (10)$$

where $\epsilon_t \sim N(0, 1)$ with draws being independent from each other and over time. Note that in equation (8) we multiply X_t by -1 so that when $b < 0$, as it will be under the calibration, Δa_t is increasing in X_t .

We assume that agents know the GARCH structure of the process for X_t , equations (9)-(10), and the parameters b and c , but have to estimate the parameter values α , ρ and ϕ .⁸ Notice that the fundamental state of the economy X_t which is conditionally normally distributed is mapped into aggregate productivity growth. This change-of-variable procedure allows our forecasters to consider a family of non-normal distributions of TFP growth and convert each one into a linear-normal filtering problem with unknown parameters that

⁸We make this assumption about b and c 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. When b and c are known, this can be done efficiently using maximum likelihood techniques. Assuming agents do not know b and their beliefs about it change over time creates an additional source of time-variation in disaster risk that amplifies uncertainty shocks. Results available on request.

can be estimated jointly. The structural form of the mapping in (8) is dictated by a couple of observations. First, it is a simple (that is, computationally feasible) yet flexible formulation that allows us to focus our attention on conditionally skewed distributions. Note that skewness in this model is most sensitive to b because that parameter governs the curvature of the transformation (8) of the normal variable. Any function with similar curvature, such as a polynomial or sine function, would deliver a similar mechanism. Second, the historical distribution of GDP growth is negatively skewed which can be achieved by setting $b < 0$. Third, Orlik and Veldkamp (2014) show how a similar formulation reproduces important properties of the GDP growth forecasts in the Survey of Professional Forecasters (SPF).

Agent i makes his labor choice L_{it} at the end of period $t - 1$. His objective is to maximize expected period t utility.⁹ The agent makes this decision at the end of period $t - 1$, which means that he does not know productivity A_t . However, at the end of period $t - 1$ he observes an unbiased signal about the fundamental state X_t :

$$z_{i,t-1} = X_t + \eta_{t-1} + \psi_{i,t-1}, \quad (11)$$

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.¹⁰ Note that there is common and idiosyncratic signal noise. The information set of firm i at the end of period $t - 1$ is $\mathcal{I}_{i,t-1} = \{A^{t-1}, z_{i,t-1}\}$, where $A^{t-1} \equiv \{A_0, A_1, \dots, A_{t-1}\}$. Agents know the history of their private signals as well, but since signals are about X_t which is revealed after production at the end of each period (an agent that knows Q_{it} can back out X_t using

⁹Decisions at time t have no effect on future utility so the agent is also maximizing expected discounted utility.

¹⁰As in Lucas (1972), we assume that there is no labor market, which means that there is not a wage which agents can use to learn about X_t or 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 (e.g. a search market in which workers and firms Nash bargain over wages). 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.

the production function (7) and the productivity growth equation (8)), past signals contain no additional relevant information.

2.2 Solution to the firm's problem

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

$$L_{it} = \left(\frac{E[A_t | \mathcal{I}_{i,t-1}]}{\gamma} \right)^{1/(\gamma-1)}. \quad (12)$$

In order to make his choice of labor the agent must forecast productivity. He forms a prior belief about the state X_t that governs TFP growth and then updates using his idiosyncratic signal. He knows that X_t follows the process specified in equations (9)–(10) but does not know the three parameters of that process. To form his prior about X_t at the end of period $t-1$ he uses the data A^{t-1} to estimate the parameters of the process by maximum likelihood. This prior is normally distributed with mean $E[X_t | A^{t-1}]$ and variance $v_{t-1} = V[X_t | A^{t-1}]$.

When the agent receives his signal at the end of period $t-1$ he updates his beliefs according to Bayes' law:

$$E[X_t | \mathcal{I}_{i,t-1}] = \frac{v_{t-1}^{-1} E[X_t | A^{t-1}] + (\sigma_\eta^2 + \sigma_\psi^2)^{-1} z_{i,t-1}}{v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}}.$$

Note that $(\sigma_\eta^2 + \sigma_\psi^2)^{-1}$ is the precision of that signal. Let the posterior variance be denoted $V_{t-1}[X_t] \equiv [v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}]^{-1}$. 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 X_t as a weighted sum of prior beliefs and the signal:

$$E_{i,t-1}[X_t] = (1 - \omega_{t-1})E[X_t | A^{t-1}] + \omega_{t-1}z_{i,t-1}. \quad (13)$$

Once the agent has beliefs about the fundamental state X_t and TFP growth, he computes

his expected value of TFP using the fact that $A_t = A_{t-1} \exp(\Delta a_t)$:

$$E[A_t | \mathcal{I}_{i,t-1}] = A_{t-1} E \left[\exp(c + b \exp(-X_t)) | \mathcal{I}_{i,t-1} \right] \quad (14)$$

and makes his labor choice according to equation (12). The right hand side of this equation cannot be evaluated analytically so for quantitative results we will compute it numerically.

2.3 Beliefs about GDP

In the quantitative section of the paper we will make use of GDP forecast data, so we present the equations needed to compute these forecasts here. Using equations (7) and (12) firm i 's output in period t is

$$Q_{it} = A_t \left(\frac{E[A_t | \mathcal{I}_{i,t-1}]}{\gamma} \right)^{1/(\gamma-1)} \quad (15)$$

and aggregate output is

$$Q_t = A_t \int \left(\frac{E[A_t | \mathcal{I}_{i,t-1}]}{\gamma} \right)^{1/(\gamma-1)} di. \quad (16)$$

Using this equation we can take an individual agent's expectation to compute his GDP forecast and then his GDP growth forecast is computed according to equation (3). Since each agent's TFP forecast must be computed numerically, we also must also compute GDP growth forecasts numerically.

2.4 Simplified models for comparison

Our results in Section 4 will focus on the role of two features of our model: learning about the state X_t with imperfect information and the skewed shocks to TFP growth. We will refer to the first feature as the *learning mechanism* and the second feature as the *disaster risk mechanism*. In order to understand the roles of these features we will present results

for two simplified versions of the model. From now on we will refer to the main model as the *disaster risk model*. In the simplest version of the model, which we call the *no learning model*, we turn off both the learning and disaster risk mechanisms. To do this we make the TFP shocks normally distributed, eliminating disaster risk, and assume that agents know X_t at the end of period $t - 1$ when making their labor decisions and forecasts so that they have perfect information. Formally there are two changes to the disaster risk model. We change the TFP growth process by dropping the change of measure in equation (8) so that the TFP growth process is now:

$$\Delta a_t = c + X_t, \tag{17}$$

where X_t follows the same process as in equations (9) and (10). To endow agents with perfect knowledge of X_t we set both signal noises to zero: $\sigma_\eta = 0$ and $\sigma_\psi = 0$.

For the second simplified model we turn the learning mechanism on and leave the disaster risk mechanism off. Formally we assume that Δa_t follows the process in equation (17) and in this model there will be signal noise: we allow $\sigma_\eta, \sigma_\phi > 0$. We call this model the *normal model*. Comparing results for the no learning and normal models will demonstrate the role of the learning mechanism, while comparing results for the normal and disaster risk models will show what disaster risk does.

The solution methods for the no learning and normal models are exactly the same as for the disaster risk model. For these models GDP growth forecasts can be solved analytically. For the normal model we present these in the appendix. In the no learning model they are trivial since there is perfect information.

3 Calibration and Simulation

The disaster risk model has eight parameters: a parameter controlling the disutility of labor (γ), a parameter for the level of TFP growth (c), a parameter controlling the curvature of the change of measure function (b), three parameters for the GARCH process (α, 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 within the range of Frisch elasticities that Keane and Rogerson (2012) argue are reasonable at the macro level. The remaining parameters are calibrated to target seven moments of the data for 1968Q4–2011Q4: the mean, standard deviation, skewness and kurtosis of GDP growth (all real GDP); the standard deviation of the absolute difference between GDP growth and its mean, $\text{std}(|\Delta q_t - \overline{\Delta q}|)$ where $\overline{\Delta q}$ is average GDP growth for the sample; the average cross-sectional standard deviation of GDP growth forecasts (defined in (4)); and the average (over forecasters and over time) absolute error of GDP growth forecasts, which is

$$\frac{1}{T} \sum_{t=1}^T \frac{1}{N_t} \sum_{i \in I_t} |E[\Delta q_t | \mathcal{I}_{i,t-1}] - \Delta q_t|, \quad (18)$$

where T is the number of periods in the data.

To calibrate the parameters of the model excluding γ we use simulated method of moments. We assume that agents in the economy have access to data going back to 1947Q2, so when simulating the model we provide the agents with the same number of observations of TFP as they would have had in the real economy under this assumption. Full details of the calibration procedure are in the appendix. The calibrated parameter values are reported in Table 2 and the values of the target moments are reported in column (2) of Table 3, along with the corresponding values for the data in column (3) (calibration targets are marked with †'s). The model does a good job of getting close to the calibration targets.¹¹ For computing results we simulate the model 2000 times for a number of quarters

¹¹The model is not able to hit the calibration targets exactly. In particular the standard deviation of GDP growth is a little high while the kurtosis of GDP growth, the average standard deviation of GDP growth forecasts (higher-order uncertainty mean) and the average absolute error of these forecasts are a little low. This is because the Bayesian learning mechanism limits what the model can attain. To illustrate this, if you hold the weight that agents place on their signals (ω_{t-1}) fixed, then increasing σ_η and σ_ϕ will increase the variance and average absolute error of forecasts. But increasing these parameters reduces the weight that agents place on their signals, offsetting these effects. Alternatively you could adjust parameters of the GARCH process so that when agents estimate this process they have less precise prior beliefs about TFP growth and place more weight on their signals. But the moments of GDP growth that are targeted limit how much this can be done. The calibration balances these tradeoffs.

<i>Parameter</i>		<i>Normal model</i>	<i>Disaster risk model</i>
non-normality	b	n/a	-0.0767
TFP growth level	c	0.0034	0.0804
mean volatility	α	2.03e-5	4.36e-4
volatility persistence	ρ	0.2951	0.7419
volatility innovation var.	ϕ	0.1789	0.1937
private signal noise	σ_ψ	0.0090	0.1106
public signal noise	σ_η	0	0

Table 2: **Parameter values for the disaster risk and normal models**

equivalent to 1947Q2 to 2011Q4, compute moments of the simulated data over periods equivalent to those used for computing moments of the actual data and then average each moment across the simulations. Further details are in the appendix.

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 simple. Matching the forecast dispersion ensures that we have the right mean level of higher-order uncertainty, but leaves fluctuations in higher-order uncertainty as a free moment. Likewise, the average absolute forecast error is related to the average amount of macro uncertainty, but does not discipline the uncertainty shocks. Thus our calibration strategy 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 but not the cyclical properties of the different uncertainty time series which we seek to explain.

For the calibration of the normal model there is one less parameter to calibrate. Relative to the disaster risk model we don't have to calibrate b which controls the curvature of the change of measure function. In terms of calibration targets we drop the skewness and kurtosis of GDP growth and add the autocorrelation of the absolute difference between GDP growth and its mean, $ac(|\Delta q_t - \overline{\Delta q}|)$. We drop skewness and kurtosis because without

	Models		Data
	Normal (1)	Disaster risk (2)	(3)
GDP growth			
Mean ^{*†}	2.72	2.71	2.71
Std. ^{*†}	3.68	3.76	3.42
Skewness [†]	0.00	-0.33	-0.32
Kurtosis [†]	3.61	4.68	5.08
std($ \Delta q_t - \overline{\Delta q} $) ^{*†}	2.30	2.48	2.41
ac($ \Delta q_t - \overline{\Delta q} $) [*]	0.23	0.30	0.23
Higher-order uncertainty			
Mean ^{*†}	1.50	1.45	1.54
GDP growth forecasts			
Av. absolute error ^{*†}	2.14	2.01	2.23

Table 3: **Calibration targets.** Calibration targets for the normal model are marked with *'s. Calibration targets for the disaster risk model are marked with †'s. Higher-order uncertainty, GDP growth forecasts and GDP growth are defined in Section 1. The period for the GDP growth data and GDP growth forecasts data is 1968Q4–2011Q4. The period for the higher-order uncertainty data is 1968Q3–2011Q3. The model samples are simulated to be the same length, as described in the appendix. The two models are described in Section 2. The mean of higher-order uncertainty is computed using the HU_t series before detrending.

the non-normal distribution the model is not capable of matching these features of the data.

We follow the same simulated method of moments procedure for calibration that we used for the disaster risk model. The parameter values are presented in Table 2 and the values of the target moments are in column (1) of Table 3 (calibration targets are marked with *'s). The model is able to match the target moments well.¹² We do not do a separate calibration for the no learning model. It is sufficient for our purposes to use the parameter values of the normal model and change the signal noises to $\sigma_\eta = \sigma_\psi = 0$.

Our numerical solution method uses two approximations to keep the problem tractable. First, when agents estimate the TFP process and use the estimate to construct prior beliefs about TFP growth, they would optimally consider the whole distribution of possible parameter values, conditional on the data set. For this they would require Bayesian Monte Carlo methods to estimate the parameters. For the number of parameters we have to estimate each period this procedure is quite slow. Therefore we approximate the solution

¹²The normal model is not able to hit the calibration targets exactly for the same reasons that have been discussed for the disaster risk model.

by replacing the distribution of each parameter with its mean—the maximum likelihood estimate. This understates uncertainty slightly. However, when we experiment with lower-parameter versions of our model we find that uncertainty fluctuations are quite similar in the true and approximate model. Second, we put zero weight on heterogeneous signals when estimating parameters. We do this because if estimated distributions differ across agents then keeping track of the distribution of distributions is too memory-intensive. The quantitative effect is tiny because the data history is a long string of public signals and the private signal is only one noisy piece of information. The weights agents place on the heterogeneous signals would typically be less than one percent and thus create very small differences in beliefs. Experiments with fewer agents and periods confirm that the heterogeneous parameter effect we neglect is negligible.

4 Results: Uncertainty Comovement

Our main result is that the combination of learning and disaster risk can explain why micro, macro and higher-order uncertainty fluctuate together. We begin by showing that the model generates uncertainty shocks of a similar size and with similar covariances to those in the data. We present results for both the normal and disaster risk models so that we separate the roles of learning and disaster risk in generating the results. We show that both mechanisms are important for generating uncertainty shocks of the size that we see in the data, that learning is important for getting the right correlation between the different uncertainty shocks and that disaster risk is critical for generating countercyclical uncertainty shocks. Getting the cyclicalities right is important because uncertainty shocks are typically used to explain downturns, not booms. We also replicate the estimation of the linear relationships between uncertainty measures, controlling for GDP growth, that we performed for the data in Table 1. We do this using simulated data from the model in order to understand whether the model answers the original motivating question: Why do uncertainty measures, that are conceptually distinct, covary above and beyond what the

	Models			Data
	No learning (1)	Normal (2)	Disaster risk (3)	(4)
<i>(a) Macro uncertainty</i>				
Std.	0	11.93	15.15	20.87
Corr. with GDP growth	0	0.00	-0.18	-0.26
Period	1962Q2–2008Q2			
<i>(b) Higher-order uncertainty</i>				
Std.	0	19.99	24.68	31.13
Corr. with GDP growth	0	0.00	-0.12	-0.28
Corr. with Micro Unc.	0	0.43	0.68	0.43
Corr. with Macro Unc.	0	0.99	0.98	0.24
Period	1968Q3–2011Q3			
<i>(c) Micro uncertainty</i>				
Std.	0	11.68	15.08	11.58
Corr. with GDP growth	0	0.00	-0.06	-0.52
Corr. with Macro Unc.	0	0.44	0.68	0.32
Period	1962Q1–2009Q3			

Table 4: **Simulation results and data counterparts.** The three models are described in Section 2. Macro uncertainty, higher-order uncertainty and micro uncertainty are defined in Section 1. All results are computed using the detrended series. As for the empirical results in Section 1 the correlation of micro uncertainty with higher-order uncertainty, macro uncertainty and GDP growth use \widehat{HU}_t , \widehat{MaU}_t and $\widehat{\Delta}q_{t+1}$ respectively. The periods are the periods for the data. The model samples are simulated to be the same length, as described in the appendix.

business cycle predicts? The fact that our model can explain these covariances is important because it supports the idea that the fluctuations we call uncertainty shocks arise in part from using real-time data to learn about the risk of disasters.

As a starting benchmark it is useful to consider the results for the no learning model, which are presented in column (1) of Table 4. Since agents receive signals that inform them perfectly about the aggregate state X_t agents all make the same decisions and can forecast GDP growth perfectly, so there is no micro, macro or higher-order uncertainty. We will now consider how the learning and disaster risk mechanisms generate each type of uncertainty shock, starting with macro uncertainty.

4.1 Macro Uncertainty

The results for macro uncertainty are presented in panel (a) of Table 4. Column (3) shows that the disaster risk model is able to generate uncertainty shocks that are about 75% of the size of those in the data (measured with the standard deviation of macro uncertainty), have a similar countercyclicality to in the data and are positively correlated with micro uncertainty as in the data. To understand how the model generates these results we will start with the role of the learning mechanism and then turn to the disaster risk mechanism.

The contribution of the learning mechanism is shown by comparing the results for the no learning and normal models in Table 4. This shows that the learning mechanism generates about half of the macro uncertainty shocks that we see in the data. This is because once agents must learn about the aggregate state X_t they have uncertainty about what GDP growth will be. There are two reasons that this uncertainty fluctuates over time. The main reason is that X_t has stochastic volatility so the prior beliefs of agents about X_t have varying precision. The precision of these prior beliefs will also vary because agents don't know the parameters of the X_t process. Their changing estimates of these parameters will affect the precision with which they think they can forecast X_t . In the appendix we show that it is stochastic volatility that generates the vast majority of the macro uncertainty shocks in the normal model. Changing uncertainty about X_t feeds directly through to changing uncertainty about GDP growth. To see this analytically we can use the closed form solution for the normal model presented in the appendix. We can express the uncertainty of agent i at the end of period $t - 1$ about GDP growth in period t as:

$$\sqrt{V[\Delta q_t | \mathcal{I}_{i,t-1}]} = 400 \sqrt{\frac{(\gamma - 1 + \omega_{t-1})^2 v_{t-1} \sigma_\psi^2 + (\gamma - 1)^2 v_{t-1} \sigma_\eta^2 + \omega_{t-1}^2 \sigma_\eta^2 \sigma_\psi^2}{(\gamma - 1)^2 (v_{t-1} + \sigma_\psi^2 + \sigma_\eta^2)}}. \quad (19)$$

This shows that an agent's uncertainty about GDP growth is a function of the parameters of the model, the variance of his prior beliefs about TFP growth v_{t-1} and the

weight that he places on his signals, ω_{t-1} , which is also a function of v_{t-1} . For $\gamma > 1$, $d\sqrt{V[\Delta q_t | \mathcal{I}_{i,t-1}]} / dv_{t-1} > 0$ showing that macro uncertainty is increasing in the variance of prior beliefs.

To the extent that these macro uncertainty shocks are caused by the stochastic volatility of X_t they are exogenous. However this does not mean that it was a given that this mechanism would generate uncertainty shocks with a magnitude similar to in the data. Recall that we did not calibrate the model to any moment of the data that ensured that we got time variation in uncertainty. Therefore the fact that the uncertainty shocks that have resulted are sizable relative to the data is evidence that our information mechanism is relevant for understanding uncertainty shocks.

Turning next to the correlation of macro uncertainty with GDP growth, the learning mechanism generates zero correlation. The reason is that the volatility of X_t is not correlated with its level, so macro uncertainty is just as likely to be high when the economy is booming as when it is in a bust. If this were the only mechanism in the model this result would be particularly problematic since much of what researchers want uncertainty to explain is the onset of recessions, crashes or crises. Next we'll show that it is the presence of disaster risk that firms have to learn about that is key to generating countercyclical macro uncertainty.

Disaster risk in the model comes from the change of variable function in equation (8). When we calibrate it we find that the coefficient b is negative, meaning that the transformation is concave. In part this is driven by the fact that GDP growth is negatively skewed in the data. 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 or, in other words, the level of disaster risk.

Figure 2 illustrates the effect of the concave change of variable on uncertainty. It plots




Figure 2: **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.

a mapping from X into TFP growth, Δa . The slope of this curve is a Radon-Nikodym derivative. 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 2. If that band is projected onto TFP growth (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 is true: the band projected on the Δa axis is wider and uncertainty is higher. 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

$$V[\Delta a_t | \mathcal{I}_{i,t-1}] = b^2 (\exp(V_{t-1}[X_t]) - 1) \exp(V_{t-1}[X_t] - 2E[X_t | \mathcal{I}_{i,t-1}]).$$

From this formula we can see that the agent's uncertainty about Δa_t is decreasing in his expected value of X_t .

The ability of disaster risk to generate countercyclical uncertainty is tied closely to the

fact that with a skewed distribution conditional variances are not independent of means. In the normal model agents have normally distributed prior beliefs about TFP growth and receive normally distributed signals. In this world the variance of an agent's posterior beliefs are independent of the value of the signal he receives and the cross-sectional variance of the mean posterior beliefs of agents are independent of the state X_t . This is no longer true when there is skewness. With negative skewness the variance of posterior beliefs is higher when the state is lower and there is a greater risk of disaster. The basic intuition is that when the economy is doing badly and is in the left tail, it is more sensitive to shocks and agents are more uncertain given a fixed amount of information. Thus the presence of disaster risk generates countercyclical fluctuations in macro uncertainty.

Quantitatively the disaster risk mechanism does a good job of bringing macro uncertainty in the model in line with the data. It increases the size of macro uncertainty shocks so that the model generates about three quarters of the shocks that are in the data and generates a negative correlation between macro uncertainty shocks and GDP growth that is about two-thirds as large as in the data.

Note that we have not reported results for the mean of macro uncertainty. The reason for this is that the model and the data are not comparable for this moment. Recall from Section 1 that in the model we measure macro uncertainty consistently with its theoretical notion as the average standard deviation of beliefs about quarterly GDP growth. We don't have an empirical counterpart for this so we are using a proxy based on the monthly volatility of aggregate stock market indexes. The construction of these measures is different so there is no reason that their means should be equal.¹³ To compare their standard deviations and correlations we have, as described in Section 1, detrended them and the series that we have worked with are their percentage deviations from trend.

¹³For completeness, the mean of macro uncertainty is 2.50 in the disaster risk model and the mean of the empirical proxy is 19.12.

4.2 Higher-order Uncertainty and Micro Uncertainty

The results for higher-order and micro uncertainty are presented in panels (b) and (c), respectively, of Table 4. The results for the disaster risk model in column (3) show that our model explains 80% of the magnitude of higher-order uncertainty fluctuations (standard deviation) and has micro uncertainty fluctuations that are slightly larger than their empirical counterparts. Both uncertainty series are countercyclical and the correlations between all three pairs of uncertainty are positive.

To understand these results first consider the contribution of the learning mechanism. For this we can focus on the normal model. Learning generates higher-order and micro uncertainty fluctuations as follows. The precision of the prior beliefs of agents vary over time for the reasons discussed in the explanation of the macro uncertainty results. When agents have imprecise prior beliefs (high v_{t-1}) they weight their heterogeneous signals more. With more weight on the heterogeneous signals (higher ω_{t-1}) there is more dispersion in beliefs, forecasts, labor choices and growth rates. That is, they generate more higher-order and micro uncertainty. Thus learning generates fluctuations in higher-order and macro uncertainty and these fluctuations are positively correlated with each other. They are also positively correlated with macro uncertainty because macro uncertainty also rises when the precision of prior beliefs increases.

We can see the relationship between learning and higher-order and micro uncertainty in the normal model analytically using results derived in the appendix. Higher-order uncertainty, the cross-sectional standard deviation of GDP growth forecasts, is

$$\sqrt{V[\log(E[Q_t|\mathcal{I}_{i,t-1}]) - \log Q_{t-1}]} = \sigma_\psi \omega_{t-1} \sqrt{1 + \left(\frac{1}{(\gamma - 1)[\sigma_\psi^2(v_{t-1} + \sigma_\eta^2)^{-1} + 1]} \right)^2}. \quad (20)$$

Let the log growth rate of firm i at time t be $\Delta q_{it} \equiv \log Q_{it} - \log Q_{i,t-1}$. An alternative measure of micro uncertainty (which we use here because it is more analytically tractable than the main measure we use) is the cross-sectional variance of firm growth rates in period

t :

$$\int (\Delta q_{it} - \overline{\Delta q_t})^2 di = \left(\frac{1}{\gamma - 1}\right)^2 \sigma_\psi^2 (\omega_{t-1}^2 + \omega_{t-2}^2), \quad (21)$$

where $\overline{\Delta q_t} \equiv \int \Delta q_{it} di$. These expressions show that when prior beliefs are relatively imprecise so that v_{t-1} and ω_{t-1} are relatively high, both micro and higher-order uncertainty increase. Agents weight their heterogeneous signals more, which amplifies the effect of the idiosyncratic signal noise (captured by the σ_ψ term).

Column (2) of Table 4 shows the contribution of the learning mechanism to explaining the micro and higher-order uncertainty data. It generates about two-thirds of the higher-order uncertainty shocks (standard deviation) and about the same amount of micro uncertainty shocks as we see in the data. It also generates a positive correlation between all three types of uncertainty. The problem with this mechanism is that it doesn't generate countercyclical higher-order and micro uncertainty shocks. The reasons for this are exactly the same as those for why learning does not generate countercyclical macro uncertainty shocks, which we discussed above. It is the presence of disaster risk that is key to being able to generate this countercyclicality.

To understand the role of disaster risk in generating higher-order and micro uncertainty shocks return to Figure 2. Because of idiosyncratic signals there will be a distribution of beliefs about the state of the economy each period. For getting intuition assume that this distribution is uniform and represent it with the bands on the horizontal axis in the figure. Due to the changing slope of the change of measure function a fixed dispersion in beliefs about the state will generate greater dispersion in beliefs about TFP growth when the economy is performing worse (lower X). When the economy is doing poorly small differences in information and beliefs about the state generate far more disagreement about TFP growth than when the economy is doing well because the distribution of TFP growth is much flatter. Greater disagreement about TFP growth generates more dispersion in labor choices and firm growth rates (micro uncertainty), and more dispersion in GDP growth forecasts (higher-order uncertainty). Thus disaster risk generates countercyclical fluctua-

tions in these two types of uncertainty in the same way as it does for macro uncertainty. Comparing the results in columns (2) and (3) of Table 4 shows the quantitative effect of adding disaster risk to the model. It increases the size of higher-order and micro uncertainty shocks so that the model generates about 80% of the higher-order uncertainty shocks and a little more micro uncertainty shocks than are in the data. It also makes uncertainty shocks countercyclical as they are in the data.

So far we have not reported results for the mean of micro uncertainty. This is because the model is not setup to match this moment of the data. There are many dimensions of heterogeneity amongst firms that we are not modeling (e.g., size, age, industry, geography) which are likely to affect the growth rates of firms. The only source of firm heterogeneity in the model is information. It is therefore expected that the model generates significantly less micro uncertainty on average than there is in the data. This is the case: the mean level of micro uncertainty in the disaster risk model is 0.51 compared to 18.59 in the data. In the appendix, we show that we can extend the model to match this moment of the data. We omit these extensions from our main model because they are not essential to our main point about uncertainty covariance and introduce significant additional complexity. As described in Section 1 we have detrended the uncertainty series and worked with the series for percentage deviations from trend to make the model and data comparable.

4.3 The Role of Parameter Updating

In the discussion of the macro uncertainty results we mentioned that one reason why the precision of prior beliefs varies over time, which is an integral part of the learning mechanism, is because agents don't know the parameters the process for X_t and must re-estimate them each period. Having agents who do not know the true model parameters and estimate them in real time is both realistic and helps modestly to amplify the uncertainty fluctuations. In the appendix we quantify the contribution of parameter updating for both the normal and disaster risk models by presenting results for modified versions of these

models under the assumption that agents know the parameter values for the X_t process. Agents updating their beliefs about these parameters over time increases the standard deviations of micro, higher-order and macro uncertainty by 8–14% in the two models.

4.4 Do Uncertainty Measures Covary Beyond the Business Cycle?

Our original question was why uncertainty measures varied above and beyond what the business cycle predicts. To determine whether our model offers a plausible explanation we regress micro uncertainty on higher-order uncertainty, higher-order on macro uncertainty and macro uncertainty on micro uncertainty in exactly the same way as we did for the data in Table 1. As in the regressions for the data we also control for the business cycle using GDP growth. We estimate regression coefficients in each of the 2000 simulations of the model and then report the mean and standard deviation of each coefficient. The first panel of Table 5 reports the relationship between micro uncertainty and higher-order uncertainty. When higher-order uncertainty deviates from trend by one additional percentage point, micro uncertainty deviates from trend by an additional 0.56 percentage points in the model. In the data it is 0.28 percentage points (see Table 1). When we control for the business cycle micro and higher-order uncertainty still have a significant positive relationship that is stronger in the model than in the data.

The second panel of Table 5 reports the relationship between higher-order uncertainty and macro uncertainty, and the third panel reports results for regressing macro uncertainty on micro uncertainty. We find that there is a significant positive relationship between both pairs of uncertainty and that this relationship survives controlling for the business cycle.

These results reveal that our mechanisms, learning and disaster risk, are more than strong enough to explain the empirical comovement of uncertainty measures, above what the business cycle can explain. In fact weakening the relationship between higher-order and micro uncertainty, by adding more unpredictable variation in micro uncertainty, might

<i>Micro Uncertainty</i>	(1)	(2)
Higher-order uncertainty	0.546*** (0.074)	0.544*** (0.076)
GDP growth		0.014 (0.373)
Obs	165	165
<i>Higher-order uncertainty</i>	(1)	(2)
Macro uncertainty	1.625*** (0.084)	1.648*** (0.070)
GDP growth		0.374** (0.180)
Obs	160	160
<i>Macro uncertainty</i>	(1)	(2)
Micro uncertainty	0.507*** (0.123)	0.496*** (0.123)
GDP growth		-0.400 (0.463)
Obs	182	182

Table 5: **Uncertainty shocks are correlated in the model, even after controlling for business cycles.** These regressions results are summary statistics from 2000 simulations of the model. For each simulation regressions were run to replicate those performed for the data. The reported coefficients are the averages across the simulations. The numbers in parentheses are the standard deviations of the coefficients across the simulations. Significance levels are computed using the fraction of simulations for which each coefficient is above zero. ***, ** and * denote significance at the 1%, 5% and 10% levels respectively (two-tail tests). Micro uncertainty is MiU_t . Higher-order uncertainty is \widehat{HU}_t in the first panel and HU_t in the second panel. Macro uncertainty is MaU_t in the second panel and \widehat{MaU}_t in the third panel. GDP growth is $\widehat{\Delta q}_t$ in the first and third panels and Δq_t in the second panel. Each series uses the simulation periods that are analogous to the periods that the corresponding data series covers, as discussed in detail in the appendix. The rows titled *Obs* provide the number of periods of simulated data used in each regression.

offer an even better fit to the data.

4.5 Does Micro Uncertainty Measure Uncertainty?

Micro uncertainty measured with earnings dispersion (equation (2)), or something similar, is commonly used as a measure of the uncertainty that firms have about their own economic outcomes. In this subsection we investigate how earnings dispersion relates to the uncertainty of firms about their growth rates. To do this we use the normal model because that allows us to derive uncertainty expressions in closed form. The variance of a firm's beliefs about its output growth in period t prior to receiving its signal at the end of period

$t - 1$, derived in the appendix, is

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

Note that all firms condition on the same information set so they have the same variance of beliefs. 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). Comparing the variance due to private signal noise, $(1/(\gamma - 1))^2 \omega_{t-1}^2 \sigma_\psi^2$, to the expression for measured micro uncertainty—equation (21)—we see that measured micro uncertainty is the same except for one extra term, ω_{t-2} . If you ignore this term for a moment, this tells us that measured micro uncertainty is capturing the uncertainty that firms have about their growth rates that is due to the private noise in their signal.

The reason that ω_{t-2} shows up in measured micro uncertainty but not in the variance of firms' beliefs about their growth rates is 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 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 the uncertainty of firms about their growth rates that arises because of private signal noise. Thus the model tells us to interpret the cross-sectional dispersion of firms growth rates as an approximate measure of the uncertainty that firms face because of idiosyncratic shocks.

5 Conclusions

Fluctuations in tail risk, micro, macro and higher-order uncertainty have been used to explain recessions, asset price drops and financial crises. Risk and uncertainty are moments based on beliefs. If we are going to use beliefs to explain these important phenomena

perhaps we should stop to ask: Why might beliefs change so much?

It is possible that each form of belief shock has its own distinct cause. However, the strong comovement of the various shocks suggests that something links them. It suggests that one should look for a common cause to these various forms of belief shocks. This paper explores such a common cause.

We argue that beliefs change because the true distribution of outcomes is unknown. Each period, when a new piece of data is observed, agents re-estimate that distribution. When weak macro outcomes make agents re-assess their beliefs about the skewness of the shocks they face, uncertainty of all types move in a correlated, volatile and counter-cyclical way. The skewness of the distribution is a crucial part of the story because when we allow the distribution of TFP growth to be negatively skewed, the probability of disasters surges. Uncertainty shocks become amplified, countercyclical and positively correlated with each other. By offering a unified explanation for the origin of uncertainty shocks, our results support the growing literature that uses uncertainty to explain some of the most important recent events in the macroeconomy.

References

- ADAM, K., A. MARCET, AND J. NICOLINI (2016): “Stock Market Volatility and Learning,” *Journal of Finance*, 71(1), 33–82.
- 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 C. BAYER (2013): “‘Wait-and-See’ Business Cycles?,” *Journal of Monetary Economics*, 60(6), 704–719.
- 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.
- BENHABIB, J., X. LIU, AND P. WANG (2016): “Endogenous Information Acquisition and Countercyclical Uncertainty,” Working paper.
- BENHABIB, J., P. WANG, AND Y. WEN (2015): “Sentiments and Aggregate Demand Fluctuations,” *Econometrica*, 83(2), 549–585.
- 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.
- LUCAS, R. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4, 103–124.
- LUDVIGSON, S., S. MA, AND S. NG (2016): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?,” New York University working paper.
- NIMARK, K. (2014): “Man-Bites-Dog Business Cycles,” *American Economic Review*, forthcoming.
- ORDONEZ, G. (2011): “The Asymmetric Effects of Financial Frictions,” *Journal of Political Economy*.
- ORLIK, A., AND L. VELDKAMP (2014): “Understanding Uncertainty Shocks and the Role of the Black Swan,” NBER Working Paper No. 20445.
- 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 Solution of the normal model

In this section we present closed form solutions for GDP growth forecasts, micro uncertainty, higher-order uncertainty and macro uncertainty for the normal model.

GDP growth forecasts In the normal model the prior belief of agents about TFP growth are normally distributed and are summarized by $E[\Delta a_t|A^{t-1}]$ and $v_{t-1} = V[\Delta a_t|A^{t-1}]$. Agents receive signals $z_{i,t-1}$ about X_t which can be transformed into signals about Δa_t : $\tilde{z}_{i,t-1} \equiv c + z_{i,t-1}$. These signals are normally distributed. The mean and variance of the posterior beliefs of agent i about TFP growth are therefore

$$E[\Delta a_t|\mathcal{I}_{i,t-1}] = (1 - \omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1}\tilde{z}_{i,t-1} \quad (22)$$

and $V_{t-1}[\Delta a_t] \equiv [v_{t-1}^{-1} + (\sigma_\eta^2 + \sigma_\psi^2)^{-1}]^{-1}$, respectively, where ω_{t-1} has the same definition as for the disaster risk model. Agent i 's forecast of period t TFP growth in the normal model is then

$$E[A_t|\mathcal{I}_{i,t-1}] = A_{t-1}E[\exp(\Delta a_t)|\mathcal{I}_{i,t-1}] = A_{t-1}\exp\left(E[\Delta a_t|\mathcal{I}_{i,t-1}] + \frac{1}{2}V_{t-1}[\Delta a_t]\right). \quad (23)$$

Note that because signals are normally distributed in the cross section, TFP growth forecasts and TFP forecasts are also normally distributed.

GDP is given by the same expression as the for the disaster risk model, equation (16). Using equation (23) and the fact that TFP growth forecasts are normally distributed in the cross section, it follows from equation (16) that:

$$Q_t = A_t\gamma^{1/(1-\gamma)}\exp\left[\frac{1}{\gamma-1}\left(\log A_{t-1} + (1 - \omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1}(\Delta a_t + \eta_{t-1}) + \frac{\omega_{t-1}^2\sigma_\psi^2}{2(\gamma-1)} + \frac{1}{2}V_{t-1}[\Delta a_t]\right)\right]. \quad (24)$$

Note that $\omega_{t-1}^2 \sigma_\psi^2$ is the cross-sectional variance of firms forecasts of TFP growth in period t . Separating the terms in equation (24) that are known at the end of period $t - 1$ from those that are unknown,

$$Q_t = \Gamma_{t-1} \exp[f(\Delta a_t, \eta_{t-1})], \quad (25)$$

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

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

Now consider agent i 's forecast of GDP. 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

$$E[Q_t | \mathcal{I}_{i,t-1}] = \Gamma_{t-1} \exp \left(E[f(\Delta a_t, \eta_{t-1}) | \mathcal{I}_{i,t-1}] + \frac{1}{2} V[f(\Delta a_t, \eta_{t-1}) | \mathcal{I}_{i,t-1}] \right). \quad (26)$$

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

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

Therefore

$$\begin{aligned} E[f(\Delta a_t, \eta_{t-1}) | \mathcal{I}_{i,t-1}] &= (1 - \omega_{t-1}) E[\Delta a_t | A^{t-1}] + \omega_{t-1} \tilde{z}_{i,t-1} \\ &+ \left(\frac{\omega_{t-1}}{\gamma-1} \right) \left(\frac{(v_{t-1} + \sigma_\eta^2)^{-1} E[\Delta a_t | A^{t-1}] + \sigma_\psi^{-2} \tilde{z}_{i,t-1}}{(v_{t-1} + \sigma_\eta^2)^{-1} + \sigma_\psi^{-2}} \right). \end{aligned} \quad (27)$$

Let $\Lambda_t \equiv [\Delta a_t - E[\Delta a_t | A^{t-1}], \eta_{t-1}]'$. Then the variance term in equation (26) is

$$V[f(\Delta a_t, \eta_{t-1}) | \mathcal{I}_{i,t-1}] = \left[\left(1 + \frac{\omega_{t-1}}{\gamma-1} \right) \quad \frac{\omega_{t-1}}{\gamma-1} \right] V[\Lambda_t | \mathcal{I}_{i,t-1}] \left[\left(1 + \frac{\omega_{t-1}}{\gamma-1} \right) \quad \frac{\omega_{t-1}}{\gamma-1} \right]' \quad (28)$$

where

$$V[\Lambda_t|\mathcal{I}_{i,t-1}] = \begin{bmatrix} v_{t-1} & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix} - \begin{bmatrix} v_{t-1} \\ \sigma_\eta^2 \end{bmatrix} (v_{t-1} + \sigma_\psi^2 + \sigma_\eta^2)^{-1} \begin{bmatrix} v_{t-1} & \sigma_\eta^2 \end{bmatrix}. \quad (29)$$

Together equations (26), (27), (28) and (29) define agent i 's forecast of period t GDP.

Micro uncertainty Let the log growth rate of firm i in quarter t be $\Delta q_{it} \equiv \log Q_{it} - \log Q_{i,t-1}$. Using equation (15), firm i 's growth rate can be expressed as a function of i 's beliefs about TFP, $E[A_t|\mathcal{I}_{i,t-1}]$, the true change in productivity Δa_t and the preference parameter γ :

$$\Delta q_{it} = \Delta a_t + \frac{1}{\gamma - 1} \left(\log(E[A_t|\mathcal{I}_{i,t-1}]) - \log(E[A_{t-1}|\mathcal{I}_{i,t-2}]) \right).$$

Making use of equations (15), (22) and (23), the cross-sectional variance of firm growth rates in period t is

$$\int (\Delta q_{it} - \overline{\Delta q}_t)^2 di = \left(\frac{1}{\gamma - 1} \right)^2 \sigma_\psi^2 (\omega_{t-1}^2 + \omega_{t-2}^2),$$

where $\overline{\Delta q}_t \equiv \int \Delta q_{it} di$. This is an alternative measure of micro uncertainty to the IQR measure in equation (2). We use this for analytical work since it has a more tractable expression.

Higher-order uncertainty Using equations (26), (27) and (28) we can evaluate the cross-sectional variance of GDP growth forecasts::

$$V[\log(E[Q_t|\mathcal{I}_{i,t-1}]) - \log Q_{t-1}] = \left[1 + \left(\frac{1}{(\gamma - 1)[\sigma_\psi^2(v_{t-1} + \sigma_\eta^2)^{-1} + 1]} \right)^2 \right] \sigma_\psi^2 \omega_{t-1}^2.$$

Macro uncertainty Evaluating macro uncertainty, using equations (25), (28) and (29) gives

$$\sqrt{V[\Delta q_t | \mathcal{I}_{i,t-1}]} = 400 \sqrt{\frac{(\gamma - 1 + \omega_{t-1})^2 v_{t-1} \sigma_\psi^2 + (\gamma - 1)^2 v_{t-1} \sigma_\eta^2 + \omega_{t-1}^2 \sigma_\eta^2 \sigma_\psi^2}{(\gamma - 1)^2 (v_{t-1} + \sigma_\psi^2 + \sigma_\eta^2)}}.$$

Firm uncertainty about own growth rate The variance of a firm's beliefs about its output growth in period t prior to receiving its signal at the end of period $t - 1$ is $V[\Delta q_{it} | A^{t-1}]$. This can be evaluated using equations (15), (22) and (23) to be

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

B Calibration and Simulation Details

Calibration For the disaster risk model we have seven parameters to calibrate, excluding γ which we calibrate externally, and we use seven moments of the data to pin these down, as discussed in the main text. The calibration procedure is simulated method of moments. We divide the moment conditions by the value of the relevant moment of the data and use an identity weighting matrix so that we're minimizing the sum of squared percentage deviations of the model moments from their data counterparts. 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 moments of the data that we are using for calibration are based on GDP growth in 1968Q4 to 2011Q4, so agents who are making decisions for the first of these quarters (1968Q4) have 86 quarters of data to use when they are first estimating the X_t process. They get an extra quarter of data for each subsequent quarter. 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. Note that we calculate moments for periods 87–259 of the simulation so that the moments are analogous to the data series for GDP growth and higher-order uncertainty. Each simulation consists of: simulating the X_t process for 259 periods; estimating the X_t 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 (16) and then computing the five moments of GDP growth that we use for calibration; computing higher-order uncertainty using equations (3), (4) and (16) and then computing its mean; and computing the average absolute error of GDP growth forecasts for each period using equations (3), (16) and

$$\frac{1}{N_{t-1}} \sum_{i \in \Omega_{t-1}} |E[\Delta q_t | \mathcal{I}_{i,t-1}] - \Delta q_t|$$

and then averaging over the relevant periods.

Key to being able to calibrate the model in a reasonable amount of time is that all of the calibration moments can be computed using Gaussian quadrature rather than by simulating the model for a large number of agents and aggregating their decisions. This can be done because 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.

For the normal model the calibration procedure is exactly the same. There is one less parameter to calibrate and the moments of GDP growth that we use as calibration targets are different, as discussed in the main text, but neither of these changes affect the calibration procedure.

Results simulation To simulate the model for computing results we follow the same procedure as for the simulations performed during the calibration. The difference is that in addition to computing GDP growth, higher-order uncertainty and the average absolute error of GDP growth forecasts, we also need to compute micro uncertainty and macro uncertainty. Macro uncertainty can be computed by Gaussian quadrature using equations (5) and (16). The measure of micro uncertainty cannot be computed by Gaussian quadrature. Instead for each simulated sample we compute sample paths for 2000 firms and compute micro uncertainty using equation (2). Computing sample paths requires drawing signals for each firm each period and computing their output decisions. The moments of macro uncertainty and micro uncertainty are computed using the periods of each sample path that are analogous to the periods for these series in the data. For example, the micro uncertainty series covers 1962Q1 to 2009Q3 in the data, so we use periods 60–250 of each simulation.

C Separating the roles of state uncertainty and parameter updating

This section separates the roles of state uncertainty in the GARCH process and parameter updating for the learning mechanism in our model. As discussed in the main text both of these elements of the model cause time variation in the variance of agents' prior beliefs, v_{t-1} , which affects the weight that agents place on their signals. We demonstrate the separate contributions of the GARCH process and parameter updating to this mechanism by presenting two versions of each of the normal and disaster risk models. In one version we turn parameter updating off by assuming that agents know α , ρ and ϕ . Under this assumption agents do not have to reestimate these parameters at the end of each period and instead form their prior beliefs about X_t using these known parameters and the structure of the X_t process given in equations (9) and (10). In the second versions parameter updating

is on and the models are exactly as described in Section 2 of the main text.

The results for these models are presented in Table 6. We also present the results for the no learning model in which agents know X_t at the end of period $t - 1$ for comparison. First look at the results for the normal model, in columns (2) and (3) of the table. In column (2) parameter updating is off and in column (3) it is turned on. Comparing the results in these columns shows that the main function of parameter updating is to generate additional uncertainty shocks. The standard deviations of micro uncertainty, higher-order uncertainty and macro uncertainty increase by 14%, 8% and 11%, respectively, when parameter updating is turned on. This shows that it is the GARCH process that generates most of the variation in the variance of prior beliefs and that parameter updating just provides some additional kick.

The results for the two versions of the disaster risk model, presented in columns (4) and (5) of Table 6, tell the same story. Turning parameter updating on increases the standard deviations of micro uncertainty, higher-order uncertainty and macro uncertainty by 9%, 9% and 10%, respectively.

D Extended Model with Firm Heterogeneity

In this section we address the inability of the models presented in the main text to match the mean level of micro uncertainty that is observed in the data. We use an extension of the normal model for this purpose. We use the normal model rather than the disaster risk model to reduce the burden of the calibration. While this means that the model will not match all of the uncertainty moments of interest it will be sufficient to demonstrate the main point of this section: that we can extend our model to match the mean level of micro uncertainty.

The change that we make is to modify the TFP growth process to allow for additional

Model	No learning	Normal		Disaster risk		
		Off	On	Off	On	
Parameter updating	n/a	(1)	(2)	(3)	(4)	(5)
<i>(a) Macro uncertainty</i>						
Std.	0	10.74	11.93	13.77	15.15	
Corr. with GDP growth	0	0.00	0.00	-0.19	-0.18	
Period	1962Q2–2008Q2					
<i>(b) Higher-order uncertainty</i>						
Std.	0	18.49	19.99	22.69	24.68	
Corr. with GDP growth	0	0.00	0.00	-0.12	-0.12	
Corr. with Micro Unc.	0	0.42	0.43	0.75	0.68	
Corr. with Macro Unc.	0	0.99	0.99	0.98	0.98	
Period	1968Q3–2011Q3					
<i>(a) Micro uncertainty</i>						
Std.	0	10.29	11.68	13.83	15.08	
Corr. with GDP growth	0	0.00	0.00	-0.06	-0.06	
Corr. with Macro Unc.	0	0.43	0.44	0.76	0.68	
Period	1962Q1–2009Q3					

Table 6: **Separating the roles of state uncertainty and parameter updating.** Macro uncertainty, higher-order uncertainty and micro uncertainty are defined in Section 1. All results for these series are computed using the detrended series. As for the empirical results in Section 1 the correlation of micro uncertainty with higher-order uncertainty, macro uncertainty and GDP growth use \widehat{HU}_t , \widehat{MaU}_t and $\widehat{\Delta q}_{t+1}$ respectively. Parameter updating off means that agents know the values of α , ρ and ϕ . When it is off agents don't know these parameters and estimate them each period. The periods are the periods for the data. The model samples are simulated to be the same length, as described in the appendix.

heterogeneity amongst firms. Specifically, TFP is assumed to be idiosyncratic and to follow the process:

$$\Delta a_{it} = c_{it} + \theta_i X_t, \tag{30}$$

where $c_{it} \sim N(c, \sigma_c^2)$ and is i.i.d. across agents and over times, $\theta_i \sim N(1, \sigma_\theta^2)$ and is i.i.d. across agents and X_t follows the same process as in equations (9) and (10). This process contains two changes relative to the normal model in the main text. θ_i captures the idea that firms differ in their sensitivity to aggregate shocks. This sensitivity is firm specific and fixed over time. The fact that c_{it} is now a random variable rather than a constant captures the idea that firms experience transitory idiosyncratic shocks in addition to being affected by aggregate shocks. The structure of signals is given by equation (11) and the remainder of the model is exactly the same as in Section 2.

This model has two more parameters to calibrate than the normal model. These are the variance of transitory idiosyncratic shocks, σ_c^2 , and the variance of firm sensitivity to aggregate shocks, σ_θ^2 . The targets that we use for these parameters are the mean and standard deviation of micro uncertainty. The rest of the calibration targets are the same as for the normal model. The calibration procedure is identical to that for the normal model with two exceptions: we need to compute micro uncertainty and do this in the same way as described in Appendix B; and since this model is much more computationally intensive than the standard model we use 20 simulated samples instead of 2000. Tests show that increasing the number of samples does not significantly affect the results.

The parameter values are presented in Table 7 and the results in Table 8. The results show that the model is capable of matching the mean and standard deviation of micro uncertainty and simultaneously generates a standard deviation of higher-order uncertainty that's very similar to the value for the data. Simulation experiments indicate that the correlations between the different types of uncertainty and their cyclicalities can be brought closer to the data by adding the disaster risk mechanism to this model.

<i>Parameter</i>		<i>Value</i>
TFP growth level	c	0.0048
Std. transient idiosyncratic shocks	σ_c	0.1459
Std. sensitivity to agg. shocks	σ_θ	4.9414
mean volatility	α	5.59e-7
volatility persistence	ρ	0.1179
volatility innovation var.	ϕ	0.0086
private signal noise	σ_ψ	0.0020
public signal noise	σ_η	0.0061

Table 7: **Parameter values for the extended model**

	Extended Model	Data
<i>Micro uncertainty</i>		
Mean	14.21	18.59
Std.	11.11	11.58
Corr. with GDP growth	0.02	-0.52
Corr. with Higher-order Unc.	0.02	0.43
Period	1962Q1–2009Q3	

Table 8: **Simulation results for extended model and data counterparts.** Micro uncertainty is defined in Section 1. The means of micro uncertainty is computed using the MiU_t . The rest of the results are computed using the detrended series. As for the empirical results in Section 1 the correlation of micro uncertainty with higher-order uncertainty and GDP growth use \widehat{HU}_t and $\widehat{\Delta q}_{t+1}$ respectively. The periods are the periods for the data. The model samples are simulated to be the same length, as described in Appendix B.