What Are Uncertainty Shocks?

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

July 12, 2017

Abstract

One of the primary innovations in modern business cycle research is the idea that uncertainty shocks drive aggregate fluctuations. But changes in stock prices (VIX), disagreement among macro forecasters, and the cross-sectional dispersion in firms' earnings, while all used to measure uncertainty, are not the same, either conceptually or statistically. Are these really measuring the same phenomenon and not just a collection of counter-cyclical second moments? If so, what is this shock that has such diverse impacts on the economy? Statistically, there is some rationale for naming all uncertainty shocks. There exists a subset of commonly-used uncertainty measures that comove significantly, above and beyond what the cycle alone could explain. Therefore, we explore a mechanism that generates *micro dispersion* (cross-sectional variance of firm-level outcomes), *higher-order uncertainty* (disagreement) and *macro uncertainty* (uncertainty about macro outcomes) from a change in macro volatility. The mechanism succeeds quantitatively, causing uncertainty measures to covary, just as they do in the data. If we want to continue the practice of naming these changes all "uncertainty shocks," these results provide guidance about what such a shock might actually entail.

*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.

One of the primary innovations in modern business cycle research is the idea that uncertainty shocks drive aggregate fluctuations. 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 et al. (2016), Ordonez (2011) and Pastor and Veronesi (2012)). But the measures of uncertainty are wide-ranging. Changes in the volatility of stock prices (VIX), disagreement among macro forecasters, and the cross-sectional dispersion in firms' earnings growth, while all used as measures of uncertainty, are not the same. Comparing VIX and firm earnings growth dispersion is like comparing business cycle volatility and income growth inequality. One measures aggregate changes in the time-series and the other differences in a cross-section. Are these disparate measures really capturing a common underlying shock? If so, what is it? Uncertainty is not exogenous. People do not spontaneously become uncertain, for no good reason. One person might. But a whole economy changing its beliefs, unprompted, is collective mania. Instead, people become uncertain after observing an event that makes them question future outcomes. That raises the question: What sorts of events can make agents uncertain in a way that shows up in all these disparate measures? Uncovering the answer to this question opens the door to understanding what this uncertainty shock is and why the aggregate economy fluctuates.

This paper contributes to answering these questions in the following ways. First it shows that the various measures of uncertainty are statistically distinct. While most measures of uncertainty are positively correlated after controlling for the business cycle, even the most highly correlated measures have correlations that are far from unity and some measures have correlations close to zero. Thus it's not obvious that these various measures of uncertainty are measuring the same shock to the economy. Using a model we show that, depending on the type of shock, different types of uncertainty can covary positively or negatively. The fact that these distinct measures are conflated in the literature is troubling because it means that there is not one uncertainty shock that explains the various aggregate outcomes linked to uncertainty. The discovery of many different shocks that explain many different outcomes is not the unified theory of fluctuations one would hope for. To unify uncertainty measures we use the model to identify a type of shock that can generate comovement in the different types of uncertainty that is consistent with the data. We find that changes in macroeconomic volatility are a quantitatively plausible explanation.

Section 1 starts with a statistical exploration of various uncertainty measures. We organize measures of uncertainty into three categories: measures of uncertainty about macroeconomic outcomes (*macro uncertainty*); measures of the dispersion of firm outcomes (*micro dispersion*); and measures of the uncertainty that people have about what others believe (*higher-order uncertainty*). We show that some measures are statistically unrelated to others, except for the fact that all are counter-cyclical. However, the data offers some hope for resuscitating uncertainty shocks as a unified phenomenon. Statistically there is a rationale for the practice of assigning a common name to some of these time-series, cross-sectional, output, price and forecast measures: A set of measures of these three types of uncertainty and dispersion have some common fluctuations at their core. This is not just a business cycle effect. These measures comove significantly above and beyond what the cycle alone can explain.

To understand these correlations and attempt to identify what kind of shock could generate them, we build a model (Section 2). In the model agents observe events, receive some private information, update beliefs with Bayes' law, form expectations and choose economic inputs in production. In this framework we formally define and solve for the three types of uncertainty and dispersion. The model has three possible second moment shocks that could generate fluctuations in uncertainty: changes in public signal noise, changes in private signal noise and changes in macro volatility. An increase in signal noise represents the idea that some information sources, such as news, ratings, or insights from friends might be less reliable or more open to interpretation.

Section 3 investigates the implications of the three possible shocks to the economy for the covariance of the three types of uncertainty and dispersion. We find that it is not given that the three types of uncertainty and dispersion are positively correlated. Their correlations can be negative. This shows that the different types of uncertainty are theoretically distinct. Therefore if we want to think of the various uncertainty shocks as a unified phenomenon, then we need a common origin for them. The negative correlations can arise when there are shocks to signal noise. For example, when signals are noisier, they convey less information, leaving agents with more macro uncertainty. At the same time, noisier signals get less weight in agents' beliefs. Since differences in signals are the source of agents' disagreement, weighting them less reduces disagreement, which results in less dispersed firm decisions (lower micro dispersion) and less dispersed forecasts (lower higher-order uncertainty).

In contrast, macro volatility fluctuations are a reliable common cause of the disparate collection of changes referred to as uncertainty shocks. Macro volatility creates macro uncertainty by making prior macro outcomes less accurate predictors of future outcomes. They create dispersion because when agents' prior information is less accurate, they weight that prior information less and weight signals more. This change in signal weighting generates greater differences in beliefs. Prior realizations are public information: Everyone saw the same GDP yesterday. But signals are heterogeneous. While one firm may incorporate their firm's sales numbers, another will examine its competitors' prices, and yet another will purchase a forecast from one of many providers. When macro uncertainty rises and prior realizations are weighted less, these heterogeneous sources of information are weighted more, driving beliefs apart.

Divergent beliefs (forecasts) create higher-order uncertainty and micro dispersion. When forecasts differ, and the difference is based on information others did not observe, another person's forecast becomes harder to predict. This is higher-order uncertainty. Firms with divergent forecasts also choose different inputs and obtain different outputs. This is more micro dispersion. All three forms of uncertainty and their covariance can be explained in a unified framework that brings us one step closer to understanding what causes business cycle fluctuations.

While our model points us to plausible sources of uncertainty measure comovement, it misses a mechanism to make uncertainty counter-cyclical. Of course we could assume that macro volatility rises in recessions, as many theories do. But since our goal is to uncover sources of fluctuations, it makes sense to ask why. To explain why uncertainty is counter-cyclical, we need to add one new ingredient to the model: disaster risk. We incorporate disaster risk by allowing the TFP growth process to have non-normal innovations. Normal distributions have thin tails, which makes disasters incredibly unlikely, and are symmetric so that disasters and miracles are equally likely. This is not what the data looks like. GDP growth is negatively skewed and our extension allows the model to have this feature. Disaster risk is important for understanding uncertainty because disaster probabilities are difficult to assess, so a rise in disaster risk creates both uncertainty about aggregate outcomes (macro uncertainty) and disagreement. And this is especially so in recessions when disasters are more likely.

Section 4 explores whether our model is quantitatively plausible. The simple model presented in Section 2 generates half of the fluctuations and most of the correlations of the various uncertainty measures. Adding disaster risk makes these uncertainty measures counter-cyclical. It also amplifies uncertainty fluctuations. The reason is that disasters are rare and difficult to predict. When outcomes are difficult to predict, firms disagree (higher-order uncertainty); they make different input choices and have heterogeneous outcomes (micro dispersion). With the learning and disaster risk mechanisms operating together, the model is able to generate over two thirds of the fluctuations in the various uncertainty measures, and have all of those uncertainty measures comove appropriately with the business cycle and each other.

Related literature 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. Since then, many other papers have further investigated uncertainty shocks as a driving force for business cycles.[1] A related strand of literature studies the impact of uncertainty shocks on asset prices.[2] Our paper complements this literature by investigating the nature and origins of their exogenous uncertainty shocks.

A few recent papers also question the origins of uncertainty shocks. Some propose reasons for macro uncertainty to fluctuate.[3] Others explain why micro dispersion is countercyclical.[4] Ludvigson et al. (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 not just statistically, but also economically, why dispersion across firms and forecasters is connected to uncertainty about aggregate outcomes, beyond what the business cycle can explain.

Finally, the tail risk mechanism that amplifies uncertainty changes in our quantitative exercise (Section 4) is also used in Orlik and Veldkamp (2014) to explain why macro uncertainty fluctuates and in Kozlowski et al. (2017) to explain business cycle persistence. This paper takes such macro changes as given and uses tail risk to amplify micro dispersion and higher-order uncertainty covariance. Without any heterogeneity in beliefs, those models cannot possibly address the central question of this paper: the distinction and connections between aggregate outcome uncertainty, micro dispersion and belief heterogeneity.

[1] e.g., Bloom et al. (2016), Basu and Bundick (2012), Bianchi et al. (2012), Arellano et al. (2012), Christiano et al. (2014), Gilchrist et al. (2013), Schaal (2012). Bachmann and Bayer (2013) dispute the effect of uncertainty on aggregate activity.

[2] e.g., Bansal and Shaliastovich (2010) and Pastor and Veronesi (2012).

[3] In Nimark (2014), the publication of a signal conveys that the true event is far away from the mean, which increases macro uncertainty. Benhabib et al. (2016) consider endogenous information acquisition. In Van Nieuwerburgh and Veldkamp (2006) and Fajgelbaum et al. (2016), less economic activity generates less data, which increases uncertainty.

[4] Bachmann and Moscarini (2012) argue that price dispersion rises in recessions because it is less costly for firms to experiment with their prices then. Decker et al. (2013) argue that firms have more volatile outcomes in recessions because they can access fewer markets and so diversify less.

1 Uncertainty Measures and Uncertainty Facts

This section makes two points. First, different types of uncertainty are statistically different—they have correlations that are far less than one—and therefore should be treated as distinct phenomena. Second, some measures have a significant positive relationship above and beyond the business cycle, which suggests that there is some force that links them. We begin with measurement and definitions. We discuss three types of uncertainty that have been discussed in the literature and introduce various measures of them. While it is well known that these measures are countercyclical, less is know about the relationship *between* the different types of uncertainty, which is our focus.

Conceptually there are three types of uncertainty that have been used in existing research. In some papers an uncertainty shock means that an aggregate variable, such as GDP, becomes less predictable.[5] We refer to this as *macro uncertainty.* In other papers an uncertainty shock describes an increase in the uncertainty that firms have about their own outcomes due to changes in idiosyncratic variables. We call this *micro uncertainty.*[6] *Higher-order uncertainty* describes the uncertainty that people have about others' beliefs, which usually arises when forecasts differ.[7]

To measure macro uncertainty, we would ideally like to know the variance (or confidence bounds) of peoples' beliefs about future macro outcomes. A common proxy for this is the VIX, which is a measure of the future volatility of the stock market, implied by options prices. To the extent that macro outcomes are reflected in stock prices and we accept the assumptions underlying options pricing formulas, this is a measure of the unpredictability of future aggregate outcomes, or macro uncertainty. Bloom (2009) constructs a series for macro uncertainty based on this and extends it back in time using the actual volatility of

[5]For macro uncertainty shocks see, for example, Basu and Bundick (2012) and Bianchi et al. (2012) on business cycles, and Bansal and Shaliastovich (2010), Pastor and Veronesi (2012) in the asset pricing literature.

[6]For micro uncertainty shocks see, for example, Arellano et al. (2012), Christiano et al. (2014), Gilchrist et al. (2013), Schaal (2012).

[7]Angeletos and La'O (2014), Angeletos et al. (2014) and Benhabib et al. (2015) all use higher-order uncertainty.

stock prices for earlier periods in which the VIX is not available. We use this series and refer to it as *VIX*. Full details of this measure and all of the uncertainty measures that we use are provided in the appendix. A second proxy for macro uncertainty is the average absolute error of GDP growth forecasts, which we will call *forecast errors*. Assuming higher uncertainty is associated with more volatile future outcomes, forecast errors will be higher on average when uncertainty is higher. We construct this measure using data on real GDP growth forecasts from the Survey of Professional Forecasters (SPF). Our third measure of macro uncertainty comes from Jurado et al. (2015). These authors construct an econometric measure of macro uncertainty based on the variance of forecasts of macro variables made using a very rich dataset. We will refer to this measure as the *JLN* (Jurado, Ludvigson and Ng) uncertainty series.

True micro uncertainty is difficult to measure because data on firms' beliefs is rare. Dispersion of firm outcomes often proxies for micro uncertainty. We refer to these as measures of *micro dispersion*. In Section 3, we will discuss how closely related micro dispersion and micro uncertainty are in the context of our model. We use three measures of this which are constructed in Bloom et al. (2016). The first is the interquartile rage of firm sales growth for Compustat firms, which we call the *sales growth* uncertainty series. The second is the interquartile range of stock returns for public firms, which we call the *stock return* uncertainty series. Third is the interquartile range of manufacturing establishment TFP shocks, which is constructed using data from the Census of Manufacturers and the Annual Survey of Manufacturers. We call this the *TFP shocks* uncertainty series.

To measure higher-order uncertainty we use data from two forecasting datasets, the Survey of Professional Forecasters and Blue Chip Economic Indicators. Both datasets provide information on the forecasts of macro variables made by professional forecasters. We measure higher-order uncertainty using each dataset by computing the cross-sectional standard deviation of GDP growth forecasts. We call these two series *SPF forecasts* and *Blue Chip forecasts*.

The first question that we investigate is whether these different types of uncertainty are statistically distinct. So far the uncertainty shocks literature has focused on the fact that all types of uncertainty are countercyclical and therefore treated them as a single phenomenon. If they really are the same phenomenon then they should comove very closely. We put this idea to a simple test by computing the correlations between our uncertainty measures. We detrend all series using a HP filter and then take each measure of uncertainty and compute its correlation with all the measures of the other types of uncertainty (e.g. take a measure of macro uncertainty and correlate it with all the measures of micro dispersion and higher-order uncertainty). This gives us 42 correlations which we plot in Figure 1. A table of the individual correlations is provided in the appendix. The results show that the correlations for all measures of uncertainty are far from one. The maximum correlation is 0.62, the mean is 0.32 and several correlations are close to zero. Thus despite all three types of uncertainty being countercyclical, they each fluctuate in a distinct way.

The variation in the fluctuations of the three types of uncertainty raises the question of whether these are three independent phenomena that are all countercyclical, or whether they have a tighter link. To investigate this we assess whether there is a positive relationship between them that holds above and beyond the business cycle. We test this by regressing each measure of uncertainty on the measures of the other types of uncertainty controlling for the real GDP growth rate:

$$u_{1t} = \alpha + \beta u_{2t} + \gamma \Delta y_t + \varepsilon_t, \tag{1}$$

where u_{1t} and u_{2t} are two measures uncertainty for period t and Δy_t is real GDP growth for period t. Again, we detrend the data before performing the analysis and we put the data in percentage deviation from trend units. So $\beta = 1$ means that a 1% deviation in the rights hand side uncertainty measure is associated with a 1% deviation in the left hand side uncertainty measure. In Table 1 we report the β coefficients from these regressions.

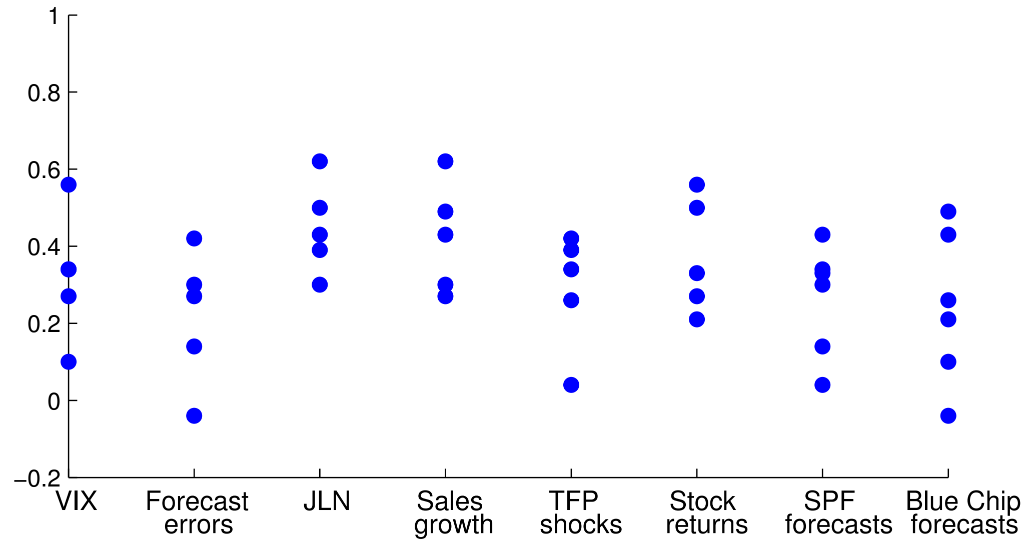

Figure 1: **Uncertainty correlations.** Correlation of each measure of uncertainty with the measures of the other types of uncertainty. VIX is a measure of uncertainty for 1962Q2–2008Q2 based on the volatility (realized and implied by options prices) of the stock market from Bloom (2009). Forecast errors is the average absolute error of GDP growth forecasts from the SPF for 1968Q3–2011Q3. JLN is the macro uncertainty measure from Jurado et al. (2015) for 1960Q4–2016Q4. Sales growth is the interquartile range of sales growth for Compustat firms for 1962Q1–2009Q3. TFP shocks is the interquartile range of TFP shocks for manufacturing establishments in the Census of Manufacturers and the Annual Survey of Manufacturers for 1972–2011 (annual data). Stock returns is the interquartile range of stocks returns for public firms for 1960Q1–2010Q3. SPF forecasts is the standard deviation of real GDP growth forecasts from the SPF for 1968Q3–2011Q3. Blue Chip forecasts is the standard deviation of real GDP growth forecasts from the Blue Chip Economic Indicators dataset for 1984Q3–2016Q4. For correlations with TFP shocks the other variables are averaged over the four quarters of each year. All uncertainty series are detrended. Additional details of the data are in the appendix.

The results show that most of the uncertainty measures have a positive and statistically significant relationship with the other measures. Aside from two series—the TFP shock measure of micro dispersion and the Blue Chip Economic Indicators measure of higher-order uncertainty—all of the series have a positive relationship with each other that's significant: for one pair of series the significance is at the 10% level, for two pairs of series it is at 5% and for all the others it is at 1%. We interpret this as evidence of some force in the economy beyond cyclical fluctuations that links these types of uncertainty.

(a) Macro uncertainty	VIX	Forecast errors	JLN
Micro dispersion			
Sales growth	0.33**	0.66***	0.18***
	(0.15)	(0.21)	(0.03)
TFP shock	0.94	2.27**	0.17*
	(0.60)	(0.87)	(0.10)
Stock returns	0.77***	1.14***	0.23***
	(0.10)	(0.26)	(0.02)
Higher-order uncertainty			
SPF forecasts	0.14**	0.23*	0.05***
	(0.06)	(0.14)	(0.02)
Blue Chip forecasts	0.08	−0.20	0.12***
	(0.10)	(0.26)	(0.02)
(b) Micro dispersion	Sales growth	TFP shock	Stock returns
Macro uncertainty			
VIX	0.08**	0.07	0.34***
	(0.04)	(0.05)	(0.04)
Forecast errors	0.09***	0.07**	0.09***
	(0.03)	(0.03)	(0.02)
JLN	0.91***	0.46*	1.38***
	(0.15)	(0.26)	(0.14)
Higher-order uncertainty			
SPF forecasts	0.13***	−0.04	0.14***
	(0.05)	(0.05)	(0.04)
Blue Chip forecasts	0.23***	0.05	0.14**
	(0.07)	(0.08)	(0.07)
(c) Higher-order uncertainty	SPF forecasts	Blue Chip forecasts	
Macro uncertainty			
VIX	0.28**	0.08	
	(0.11)	(0.10)	
Forecast errors	0.07*	−0.03	
	(0.04)	(0.04)	
JLN	1.14***	2.07***	
	(0.35)	(0.35)	
Micro dispersion			
Sales growth	0.37***	0.41***	
	(0.13)	(0.13)	
TFP shock	−0.42	0.34	
	(0.51)	(0.52)	
Stock returns	0.48***	0.30**	
	(0.14)	(0.14)	

Table 1: **Coefficients for uncertainty regressions controlling for the business cycles.** This table shows the value for β coefficient from the regressions specified in equation (1). Standard errors are in parentheses. ***, ** and * denote significance with respect to zero at the 1%, 5% and 10% levels respectively. All uncertainty series are detrended and put in percentage deviations from trend units. See the notes to Figure 1 and the appendix for details of the series.

2 Model

To make sense of the facts about macro uncertainty, higher-order uncertainty, and micro dispersion, we need a model with uncertainty about some aggregate production-relevant outcome, a source of belief differences, and firms that use their beliefs to make potentially heterogeneous production decisions. To keep our message as transparent as possible, we assume production is linear in labor alone. Since TFP is our key state variable, macro uncertainty comes from uncertainty about how productive current-period production will be (TFP). Belief differences arise from heterogeneous signals about that TFP. To capture the idea that some, but not all, of our information comes from common sources, the TFP signals have public and private signal noise. Finally, we need an exogenous shock that might plausibly affect all three uncertainties. We consider three possibilities. One is a time-varying variance of TFP innovations. The second and third are time-varying variances of public and private signal noise. This section formalizes these assumptions and describes the model's solution.

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}:

$$U_{it} = Q_{it} - L_{it}^{\gamma} \tag{2}$$

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

$$Q_{it} = A_t L_{it}. \tag{3}$$

Aggregate output is $Q_t \equiv \int Q_{it} di$ and GDP growth is $\Delta q_t \equiv \log Q_t - \log Q_{t-1}$.

The growth rate of productivity at time t, $\Delta a_t \equiv \log(A_t) - \log(A_{t-1})$, is:

$$\Delta a_t = \alpha_0 + \sigma_t \epsilon_t. \tag{4}$$

where $\epsilon_t \sim N(0,1)$ and draws are independent. The key feature of this equation is that the variance of TFP growth, σ_t^2, can be time-varying. This will be one of the potential sources of uncertainty shocks that we discuss in the next section.

Agent i makes his labor choice L_{it} at the end of period $t-1$. His objective is to maximize expected period t utility.[8] The agent makes this decision at the end of period $t-1$. At this point we assume that the agent knows σ_t but ϵ_t has not yet been realized so he does not know productivity A_t. This assumptions holds if σ_t^2 follows a GARCH process, for example, as it will later in the paper.

At the end of period $t-1$ each agent observes an unbiased signal about TFP growth which has both public (common) noise and private noise:

$$z_{i,t-1} = \Delta a_t + \eta_{t-1} + \psi_{i,t-1}, \tag{5}$$

where $\eta_{t-1} \sim N(0, \sigma^2_{\eta_{t-1}})$ and $\psi_{i,t-1} \sim N(0, \sigma^2_{\psi,t-1})$.[9]All draws of the public and private noise shocks, the η_{t-1}'s and $\psi_{i,t-1}$'s, are independent of each other.[10] We assume that the variances of public and private signal noises can be time-varying, so they are potential sources of uncertainty shocks. The information set of firm i at the end of period $t-1$ is

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

[9]These correlated signals also allow us to investigate the extreme cases of purely public and purely private signals. Pure public signals act just like a reduction in prior uncertainty. They can be created by setting private signal noise to zero $\sigma^2_{\psi,t-1} = 0$. Pure private signals are a special case where public signal noise (σ_η) is zero.

[10]As 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.

$\mathcal{I}_{i,t-1} = \{A^{t-1}, z_{i,t-1}\}$, where $A^{t-1} \equiv \{A_0, A_1, \ldots, A_{t-1}\}$. 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,[11] 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:

$$L_{it} = \left(\frac{E[A_t|\mathcal{I}_{i,t-1}]}{\gamma}\right)^{1/(\gamma-1)}. \tag{6}$$

In order to make his choice of labor the agent must forecast productivity. He forms a prior belief about TFP growth and then updates using his idiosyncratic signal. To form his prior belief he uses his knowledge of the mean of TFP growth in equation (4): $E[\Delta a_t|A^{t-1}] = \alpha_0$. The agent's prior belief that Δa_t is normally distributed with mean α_0 and variance $V[\Delta a_t|A^{t-1}] = \sigma_t^2$.

At the end of period $t-1$, the agent receives a signal with precision $(\sigma^2_{\eta,t-1} + \sigma^2_{\psi,t-1})^{-1}$ and updates his beliefs according to Bayes' law. The updated posterior forecast of TFP growth is a weighted sum of the prior belief and the signal:

$$E[\Delta a_t|\mathcal{I}_{i,t-1}] = (1-\omega_{t-1})\alpha_0 + \omega_{t-1} z_{i,t-1}, \tag{7}$$

$$\text{where} \quad \omega_{t-1} \equiv [(\sigma^2_{\eta,t-1} + \sigma^2_{\psi,t-1})(\sigma_t^{-2} + (\sigma^2_{\eta,t-1} + \sigma^2_{\psi,t-1})^{-1})]^{-1}, \tag{8}$$

The Bayesian weight on new information ω is also called the Kalman gain.

The posterior uncertainty, or conditional variance is common across agents because all agents receive signals with the same precision:

$$V_{t-1}[\Delta a_t] \equiv [\sigma_t^{-2} + (\sigma^2_{\eta,t-1} + \sigma^2_{\psi,t-1})^{-1}]^{-1} \tag{9}$$

[11] An agent that knows Q_{it} can back out X_t using the production function (3) and the productivity growth equation (4).

The agent's expected value of the level of TFP uses the fact that $A_t = A_{t-1}\exp(\Delta a_t)$:

$$E[A_t|\mathcal{I}_{i,t-1}] = A_{t-1}\exp\Big(E[\Delta a_t|\mathcal{I}_{i,t-1}] + \frac{1}{2}V_{t-1}[\Delta a_t]\Big) \qquad (10)$$

Given this TFP forecast, the agent makes his labor choice according to equation (6). The labor choice dictates the period t growth rate of firm i, $\Delta q_{it} \equiv \log Q_{it} - \log Q_{i,t-1}$, which is $\Delta q_{it} = \Delta a_t + \frac{1}{\gamma-1}\Big(\log(E[A_t|\mathcal{I}_{i,t-1}]) - \log(E[A_{t-1}|\mathcal{I}_{i,t-2}])\Big)$. Integrating over all firms' output delivers aggregate output:

$$Q_t = A_t \int \left(\frac{E[A_t|\mathcal{I}_{i,t-1}]}{\gamma}\right)^{1/(\gamma-1)} di. \qquad (11)$$

2.3 Uncertainty Measures in the Model

We derive macro, micro and higher-order uncertainty in the model, highlight the similarities and differences, and examine what forces make each one move.

Macro uncertainty For the model we define macro uncertainty to be the conditional variance of GDP growth forecasts, which is common for all agents:

$$\mathcal{U}_t \equiv V[\Delta q_t|\mathcal{I}_{i,t-1}] = \frac{(\gamma-1+\omega_{t-1})^2\sigma_t^2\sigma_{\psi,t-1}^2 + (\gamma-1)^2\sigma_t^2\sigma_{\eta,t-1}^2 + \omega_{t-1}^2\sigma_{\eta,t-1}^2\sigma_{\psi,t-1}^2}{(\gamma-1)^2(\sigma_t^2+\sigma_{\psi,t-1}^2+\sigma_{\eta,t-1}^2)}. \qquad (12)$$

If there is a prior belief about TFP with variance σ_t^2 and a signal with signal noise variance $\sigma_{\psi,t-1}^2 + \sigma_{\eta,t-1}^2$, then the variance of the posterior TFP belief is $\sigma_t^2(\sigma_{\psi,t-1}^2 + \sigma_{\eta,t-1}^2)/(\sigma_t^2 + \sigma_{\psi,t-1}^2 + \sigma_{\eta,t-1}^2)$. You can see this form showing up in the first two terns of the numerator and the denominator. The difference between $\mathcal{U}_t$ and the Bayesian posterior we just discussed is that $\mathcal{U}_t$ is the conditional variance of output, not of TFP. How TFP maps into output also depends on what other firms believe. Thus, the last term of the numerator $\omega_{t-1}^2\sigma_{\eta,t-1}^2\sigma_{\psi,t-1}^2$ represents uncertainty about how much other firms will produce.

Higher-order uncertainty We measure higher-order uncertainty with the cross-sectional variance of GDP growth forecasts (eq. 21):

$$\mathcal{H}_t \equiv V\big\{E[\Delta q_t|\mathcal{I}_{i,t-1}]\big\} = \sigma^2_{\psi,t-1}\omega^2_{t-1}\left(1+\frac{1}{(\gamma-1)[\sigma^2_{\psi,t-1}(\sigma^2_t+\sigma^2_{\eta,t-1})^{-1}+1]}\right)^2. \quad (13)$$

Higher-order uncertainty arises because there is private signal noise ($\sigma^2_{\psi,t-1} > 0$). The larger is the private signal noise, the more agents' signals differ. What also matters is how much they weight these signals in their beliefs (ω^2_{t-1}). If private signal noise is so large that the signal has almost no information, then ω^2_{t-1} becomes small, and beliefs converge again. The last term in parentheses is the rate of transformation of belief differences in TFP to belief differences in output.

We have constructed higher-order uncertainty to be two-sided and symmetric. Agent i's uncertainty about j is equal to j's uncertainty about i. In general, this need not be the case. For example, in sticky information (Mankiw and Reis, 2002) or inattentiveness (Reis, 2006) theories, information sets are nested. The agent who updated more recently knows exactly what the other believes. But the agent who updated in the past is uncertain about what the better-informed agent knows. If we average across all pairs of agents, then the average uncertainty about others' expectations would be the relevant measure of higher-order uncertainty. In such a model, belief dispersion and higher-order uncertainty are not identical, but move in lock-step. The only way agents can have heterogeneous beliefs, but not have any uncertainty about what others know, is if they agree to disagree. That doesn't happen in a Bayesian setting like this.

Micro dispersion Our measure of micro dispersion for the model is the cross-sectional variance of firm growth rates:

$$\mathcal{D}_t \equiv \int(\Delta q_{it}-\overline{\Delta q}_t)^2 di = \left(\frac{1}{\gamma-1}\right)^2(\omega^2_{t-1}\sigma^2_{\psi,t-1}+\omega^2_{t-2}\sigma^2_{\psi,t-2}), \quad (14)$$

where $\overline{\Delta q}_t \equiv \int \Delta q_{it} di$.[12] This expression shows that micro dispersion in period t depends on the variance of private signals ($\sigma^2_{\psi,t-1}$ and $\sigma^2_{\psi,t-2}$), the weights that agents place on their signals in periods $t-1$ and $t-2$, and the Frisch elasticity of labor supply ($1/(\gamma-1)$). Holding these weights fixed, when private signal noise increases, firms receive more heterogeneous signals. Beliefs about TFP growth, labor choices and output therefore become more dispersed as well. When agents place more weight on their signals, it also generates more of this dispersion. More weight on the $t-1$ signal increases dispersion in Q_{it} while more weight on the $t-2$ signal increases dispersion in $Q_{i,t-1}$. Both matter for earnings growth: $\Delta q_{it} = \log(Q_{it}) - \log(Q_{i,t-1})$. Micro dispersion depends on γ, the inverse elasticity of labor supply, because less elastic labor makes labor and output less sensitive to differences in signals.

Does dispersion measure firms' uncertainty? Micro dispersion is commonly interpreted as a measure of the uncertainty that firms have about their own economic outcomes. The uncertainty that firms have about their period t growth rates prior to receiving their signals in period $t-1$ is

$$V[\Delta q_{it}|A^{t-1}] = \sigma^2_t + \left(\frac{1}{\gamma-1}\right)^2 \omega^2_{t-1}(\sigma^2_{\eta,t-1} + \sigma^2_{\psi,t-1}). \tag{15}$$

This expression tells us that firms have uncertainty about their growth rates that comes from three sources. First there is uncertainty about TFP that shows up as σ^2_t. The firm also has uncertainty because it doesn't know what signal it will receive. This generates two additional sources of uncertainty: public signal noise ($\sigma^2_{\eta,t-1}$) and private signal noise ($\sigma^2_{\phi,t-1}$). The relevance of these depends on the weight that firms place on their signals (ω_{t-1}) and how sensitive a firm's labor choice is to changes in its beliefs, $1/(1-\gamma)$.

Comparing the uncertainty that a firm has about its growth (15) with the dispersion of firm growth (14) highlights two differences. First, dispersion only captures uncertainty

[12] We use a standard deviation rather than an IQR measure as is commonly used for the data since the standard deviation is much more analytically tractable.

due to private differences in information. Both the volatility of TFP σ_t and noise in signals that are publicly observed $\sigma_{\eta,t-1}$, make firms uncertain about their growth. But neither generates dispersion. Micro dispersion only captures uncertainty caused by idiosyncratic differences in information: $(\frac{1}{\gamma-1})^2\omega_{t-1}^2\sigma_{\psi,t-1}^2$. Second, micro dispersion systematically overestimates this component of a firm's uncertainty: $\mathcal{D}_t \geq (\frac{1}{\gamma-1})^2\omega_{t-1}^2\sigma_{\psi,t-1}^2$. The reason is that differences in levels of output yesterday and differences today both contribute to dispersion in firm growth rates. But firms are not uncertain about what they produced yesterday. It is only differences today that matter for uncertainty. This shows up mathematically as $\omega_{t-2}\sigma_{\psi,t-2}$ being in the expression for dispersion, but not in the expression for uncertainty. If $\omega_{t-1}\sigma_{\psi,t-1}$ and $\omega_{t-2}\sigma_{\psi,t-2}$ are closely correlated then changes in micro dispersion will be a good proxy for the part of uncertainty due to private information shocks. But it is possible to have dispersion, without there being any uncertainty.[13] Despite these issues, we focus on micro dispersion, as it allows us to speak to the existing literature.

3 What Is an Uncertainty Shock?

There are various, distinct measures of uncertainty. If we want to think of uncertainty as a unified concept that can explain many business cycle and financial facts, there needs to be a single shock that moves all these various measures. In this section, we explore three plausible candidates and show that one of them can explain the comovement of uncertainty measures in the data. In so doing, it provides a source of unification for theories based on uncertainty shocks.

In our model there are three potential sources of changes in uncertainty: changes in the amount of public signal noise ($\sigma_{\eta,t-1}$), changes in private signal noise ($\sigma_{\psi,t-1}$), and changes in the volatility of TFP shocks (σ_t). We consider what kind of comovement each one could

[13] If there is dispersion in firm output in period $t-1$ but agents have perfect information in period t (which will be the case if $\sigma_t^2 = 0$ so that $\omega_{t-1} = 0$) so that they all choose the same output then there will be positive micro dispersion in period t yet firms will have no uncertainty about their growth rates: $V[\Delta q_{it}|A^{t-1}] = 0$.

generate between the uncertainty and dispersion measures.

Public signal noise shocks. We first explore what happens to uncertainty measures when public signal noise rises.Public signal noise is sometimes referred to as "sender noise" because noise that originates with the sender affects receivers' signals all in the same way. The proofs of this and all subsequent results are in the appendix.

Result 1. ***Shocks to public signal noise can generate positive or negative covariances between any pair of the three types of uncertainty and dispersion:*** *Fix σ_t and $\sigma_{\psi,t-1}$ so that they do not vary over time.*

(a) If (25) *holds then $cov(\mathcal{D}_t, \mathcal{H}_t) > 0$ and otherwise $cov(\mathcal{D}_t, \mathcal{H}_t) \leq 0$;*

(b) If (27) *holds then $cov(\mathcal{D}_t, \mathcal{U}_t) < 0$ and otherwise $cov(\mathcal{D}_t, \mathcal{U}_t) \geq 0$;*

(c) If $\sigma_{\psi,t-1}$ is sufficiently small then $cov(\mathcal{H}_t, \mathcal{U}_t) < 0$ and if only one of conditions (25) *and* (27) *hold then $cov(\mathcal{H}_t, \mathcal{U}_t) > 0$.*

Public signal noise lowers micro dispersion. When public signal noise increases the signals that agents receive are less informative so they place less weight on them. This causes them to have less dispersed beliefs about TFP growth, which results in less dispersion in their labor input choices and less dispersion in their growth rates.

For higher-order uncertainty, public signal noise has two opposing effects. The direct effect comes from the same channel described above. A decrease in the dispersion of beliefs about TFP growth makes GDP growth forecasts less dispersed, which reduces higher-order uncertainty. The indirect effect arises because when agents are forecasting GDP growth, they need to forecast the labor input decisions of other agents, who produce that GDP. When public signal noise increases, the average signals of others ($\Delta a_t + \eta_{t-1}$) are more volatile. Because one's own signal is useful to predict others' signals, and those signals become more important to predict, agents to weight their own signals more. Greater weight on one's own signals causes GDP growth forecasts to be more dispersed and higher-order uncertainty to rise. If macro volatility is sufficiently high relative to signal noise, then the direct effect dominates because agents are mostly concerned about forecasting TFP growth, not others' signals. If macro volatility is low, others' signals are important to forecast, which makes the indirect effect stronger. Thus, public signal noise shocks can generate a positive or negative covariance between higher-order uncertainty and micro dispersion (Result 1a).

When public signal noise increases, there are also two effects on macro uncertainty. The direct effect is that less precise signals carry less information about TFP growth. This raises uncertainty about GDP growth. The indirect effect comes from agents needing to forecast the actions of others, in order to forecast their output (GDP). When public signal noise increases, agents weight their signals less in their production decisions. Since others' signals are unknown, less weight makes it easier for agents to forecast each others' decisions. This reduces uncertainty about GDP growth. Of the two opposing effects on macro uncertainty, either can dominate: If public signal noise is sufficiently high, then forecasting the actions of others will be very important. So the indirect effect will dominate and macro uncertainty will decrease. If private signal noise is sufficiently low, then all agents receive similar information, will be able to forecast each others' actions well and the direct effect will dominate. In this case macro uncertainty will increase. Thus micro dispersion and macro uncertainty can be positively or negatively correlated due to shocks to public signal noise (Result 1b).

The final part of Result 1 is about the covariance between higher-order uncertainty and macro uncertainty. These two types of uncertainty can have either a positive or negative covariance. There is a wedge between them because it is possible for uncertainty about TFP growth to increase and at the same time the dispersion of beliefs to decrease. This happens when private signal noise is low. In this case, an increase in public signal noise increases macro uncertainty because agents have more uncertainty about TFP growth. It decreases higher-order uncertainty because agents weight their signals less and have less dispersed beliefs about TFP growth.

Private signal noise shocks. A change in private signal noise can also generate positive or negative covariances.

Result 2. ***Shocks to private signal noise can generate positive or negative covariances between any pair of the three types of uncertainty and dispersion:*** *Fix σ_t and $\sigma_{\eta,t-1}$ so that they do not vary over time.*

(a) If $\sigma_{\psi,t-1}$ is sufficiently small then $cov(\mathcal{U}_t, \mathcal{H}_t) > 0$, $cov(\mathcal{U}_t, \mathcal{D}_t) > 0$ and $cov(\mathcal{H}_t, \mathcal{D}_t) > 0$;

(b) In only one of (29) and (30) holds then $cov(\mathcal{U}_t, \mathcal{H}_t) \geq 0$;

(c) In only one of (28) and (29) holds then $cov(\mathcal{U}_t, \mathcal{D}_t) \geq 0$;

(d) In only one of (28) and (30) holds then $cov(\mathcal{H}_t, \mathcal{D}_t) \geq 0$.

As with public noise, private signal has two competing effects on micro dispersion. There is more dispersion in the signals agents receive so if they hold the weight on their signals fixed they will have more dispersed beliefs about TFP growth, which will result in higher micro dispersion. But, because the signals are less informative agents weight them less and weight their prior beliefs more. Since agents have common prior beliefs, this decreases the dispersion in beliefs resulting in lower micro dispersion. Which of these forces is stronger depends on parameter values.

For higher-order uncertainty, these same two opposing forces are also at work. Recall from the discussion of public information shocks that when agents are forecasting GDP growth they need to forecast TFP growth as well as the actions of others. The discussion of micro dispersion tells us how an increase in private signal noise affects the dispersion of forecasts of TFP growth. For forecasts of other agents' actions the two forces also work in opposite directions. The fact that agents get signals that differ more from each other will mean that there will be greater differences in the forecasts that agents make of each others' actions. But since these signals are noisier agents will weight them less, which will bring their forecasts closer.

Private signal noise affects uncertainty about TFP growth and about the actions of others, both of which matter for macro uncertainty. The fact that agents have less precise information increases their uncertainty about TFP growth. In terms of forecasting the actions of others, agents will be more uncertain due to the fact that signals differ more across agents, but this is offset by the fact that agents are weighting their signals less.

When private signal noise is sufficiently small it is the increase in the dispersion of signals that is the dominant effect and the effects of changing signals weights are secondary. In this case all three types of uncertainty increase when private signal noise increases, so private signal noise shocks generate positive correlations between all three types of uncertainty and dispersion. This is part (a) of Result 2. This is not necessarily the case though. There are conditions under which the uncertainty and dispersion measures are negatively correlated, as provided by parts (b)–(d) of the result. There is a wedge between macro uncertainty and micro dispersion because an increase in private signal noise increase uncertainty about TFP growth but can increase or decrease the dispersion in TFP growth forecasts. This is also the cause of the wedge between macro uncertainty and higher-order uncertainty. There is a wedge between micro dispersion and higher-order uncertainty because agents weight their signals differently when TFP growth and when forecasting the actions of others. These wedges are why the different measures of uncertainty can react react in opposite ways to changes in private signal noise.

Macro volatility shocks. The third possible source of uncertainty shocks is changes in the volatility of TFP growth. Unlike the other potential sources of uncertainty shocks, this source generates positively correlated fluctuations in macro uncertainty, higher-order uncertainty and micro dispersion without additional conditions.

Result 3. ***Shocks to macro volatility generate positive covariances between all pairs of the three types of uncertainty and dispersion:*** *Fix* $\sigma_{\eta,t-1}$ *and* $\sigma_{\psi,t-1}$ *so that they do not vary over time. Then* $cov(\mathcal{U}_t, \mathcal{H}_t) > 0$*,* $cov(\mathcal{U}_t, \mathcal{D}_t) > 0$ *and* $cov(\mathcal{H}_t, \mathcal{D}_t) > 0$*.*

When macro volatility increases agents have less precise prior information about TFP growth and they therefore weight their signals more. Since those signals are heterogeneous this causes their beliefs about TFP growth to be more dispersed, which results in more dispersed production decisions and higher micro dispersion. In terms of macro uncertainty, the less precise information about TFP growth increases macro uncertainty. Agents weighting their signals more also makes it harder for agents to forecast the actions of others', which further increases macro uncertainty. Higher-order uncertainty also increases for two reasons. First the increase in dispersion of beliefs about TFP growth that results from agents weighting their signals more increases the differences between GDP growth forecasts. These differences also increase because agents weight their signals more when forecasting the actions of others, so they have more divergent forecasts of each others' actions.

There are two points to take from this section. The first is that the different types of uncertainty and dispersion are theoretically distinct. They are not mechanically linked and nor do they naturally fluctuate together. Only one of the possible sources of uncertainty shocks necessarily generates the positive correlation between all three types of uncertainty and dispersion that we see in the data. Therefore it is erroneous to treat these types of uncertainty and dispersion as single unified phenomenon, as the existing uncertainty shocks literature has tended to. If we want to unify these various shocks then we need a theory that ties them together. That's the second point: If we're after a common origin for the various uncertainty and dispersion shocks, then changes in macro volatility are a possible source. The next section evaluates whether macro volatility is quantitatively relevant for understanding uncertainty shocks.

4 Do Macro Volatility Shocks Generate Enough Uncertainty Comovement?

To develop a quantitatively viable theory, we enrich the model from Section 2 and calibrate it to the data. The augmented, calibrated model produces uncertainty shocks that are, in many respects, quantitatively similar to the data.

4.1 Quantitative Model

Since we are assessing the quantitative potential of changes in TFP growth volatility, to explain uncertainty shocks, we shut down time variation in signal noise: $\sigma_{\psi,t-1} = \sigma_\psi$ and $\sigma_{\eta,t-1} = \sigma_\eta$ for all t. To estimate TFP growth volatility, we need to impose some stochastic structure. Since GARCH is common and simple to estimate, we assume that σ_t^2 follows a GARCH process:

$$\sigma_t^2 = \alpha_1 + \rho\sigma_{t-1}^2 + \phi\sigma_{t-1}^2\epsilon_{t-1}^2. \tag{16}$$

where $\epsilon_t \sim N(0, 1)$, with draws being independent.

The final modification to the model is that we give TFP growth a negatively skewed distribution. This captures the idea of disaster risk. Disaster risk is a useful ingredient because it amplifies the uncertainty shocks and is gets their cyclicality correct. Disaster risk can amplify uncertainty during economic downturns because disasters are more likely during these periods and they extreme events whose exact nature is difficult to predict. This creates a lot of scope for uncertainty and disagreement.

We introduce non-normality into the model by assuming that the economy has an underlying state X_t which is subject to a non-linear transformation to generate TFP growth. Specifically, instead of Δa_t being determined by equation (4), we assume that

$$\Delta a_t = c + b\exp(-X_t), \tag{17}$$

$$X_t = \alpha_0 + \sigma_t\epsilon_t, \tag{18}$$

and σ_t^2 still follows equation (16). We have taken the TFP growth process from the baseline model, done an exponential transformation and linearly translated it. 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.[14] The structural form of the mapping in (17) is dictated by a couple of observations. First, it is a simple, computationally feasible 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 (17) 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.

We replace the signals in equation (5) with signals about X_t:

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

where where $\eta_{t-1} \sim N(0, \sigma_\eta^2)$ and $\psi_{i,t-1} \sim N(0, \sigma_\psi^2)$. The mechanics of learning are the same as in the baseline model, with two exceptions: (1) agents are now learning about X_t instead of Δa_t, and (2) once they form beliefs about X_t, they transform these into beliefs about Δa_t with equation (17).

When discussing the quantitative results, we will call the model presented in this section the *disaster risk model* and call the model presented in Section 2 the *normal model*. The difference between these models will demonstrate the role of the skewed TFP growth distribution. We will also present results for a version of the Section 2 model in which we assume that agents have perfect information ($\sigma_\eta = \sigma_\psi = 0$) about TFP growth, which

[14] It is also possible to allow the parameters of the model to be unknown, in which case agents need to estimate them each period. This version of the model with results is in the appendix. Adding parameter learning to the model modestly amplifies fluctuations in uncertainty.

we call the *perfect information model.* The difference between this model and the normal model will demonstrate the role of imperfect information and learning in the results.

4.2 Why Disaster Risk Makes Uncertainty Counter-Cyclical

Disasters are more likely in bad times. The heightened risk of disasters make bad states more uncertain times. Furthermore, disasters are hard to predict. They leave lots of scope for disagreement, dispersion of beliefs and dispersion of actions. These forces make all three uncertainty and dispersion measures counter-cyclical.

To formalize this argument, start with the change of variable function in equation (17), which introduced disaster risk. When we calibrate it, we find that the coefficient b is negative, meaning that the transformation is concave. 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.

Disaster risk and macro uncertainty. Figure 2 illustrates the effect of the concave change of variable on uncertainty. It plots a mapping from X into TFP growth, Δa. The slope of this curve is a Radon-Nikodym derivative. For illustrative purposes, suppose for a moment 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 it true: the band projected on the Δa axis is wider and uncertainty is higher.

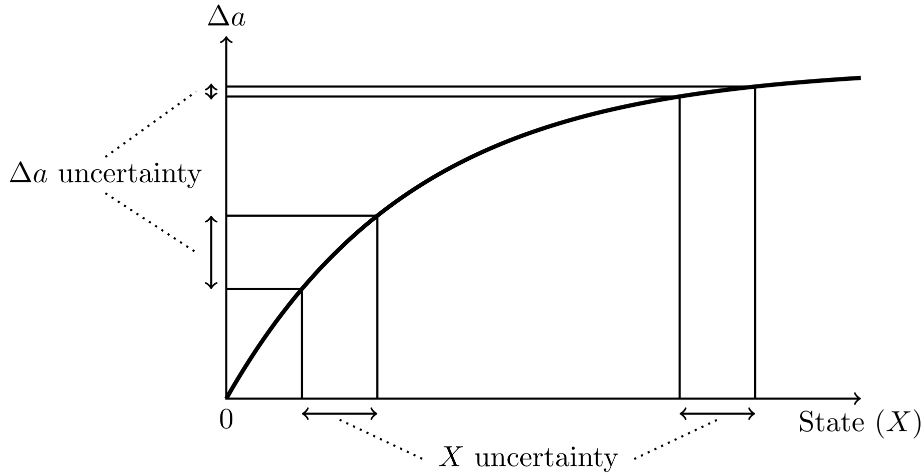

Figure 2: **Change of variable function and counter-cyclical forecast dispersion.** A given amount of uncertainty about X creates more uncertainty about TFP growth when X is low than it does when X is high.

Recall that posterior beliefs about X are actually assumed to be normally distributed. According to the formula for the variance of a log-normal distribution, the variance of beliefs about TFP growth is

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

Importantly, the source of counter-cyclical uncertainty appears: The agent's uncertainty about Δa_t (conditional variance) is decreasing in his expected value of X_t.

Counter-cyclical uncertainty arises because, in a skewed distribution, conditional variances are not independent of means. With normally distributed prior beliefs and signals, the variance of an agent's posterior beliefs would be independent of the value of his signal. Furthermore, the cross-sectional variance of the mean posterior beliefs of agents would be 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 a disaster.

Parameter		*Normal model*	*Disaster risk model*
Disutility of labor	γ	2	2
Non-normality	b	n/a	-0.0767
TFP growth level (disaster risk)	c	n/a	0.0804
TFP growth level (normal)	α_0	0.0034	0
Mean volatility	α_1	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**

Disaster risk, higher-order uncertainty and micro dispersion To understand the role of disaster risk in generating micro and higher-order 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). This means that 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. Greater disagreement about TFP growth generates more dispersion in labor choices and firm growth rates (micro dispersion), and more dispersion in GDP growth forecasts (higher-order uncertainty). Again, the intuition here is that when the economy is doing badly it is more sensitive to shocks so agents respond more to differences in information. Thus disaster risk generates fluctuations in micro dispersion and higher-order uncertainty beyond those coming from the learning mechanism and these fluctuations are countercyclical.

	Models		Data
	Normal	Disaster risk	
	(1)	(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
$\text{std}(\lvert\Delta q_t - \overline{\Delta q}\rvert)^{*\dagger}$	2.30	2.48	2.41
$\text{ac}(\lvert\Delta q_t - \overline{\Delta q}\rvert)^{*}$	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. GDP growth forecasts are from the Survey of Professional Forecasters. We use the measure of higher order uncertainty from this dataset that is discussed in Section 1 (the standard deviation of GDP growth forecasts). 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 the data, as described in the appendix.

4.3 Calibration and Simulation

We calibrate the normal and disaster risk models separately, but with the same procedure, at a quarterly frequency with data from 1968Q4–2011Q4. We use existing estimates for the disutility of labor γ. For the disaster risk model, α_0 can be normalized. The remaining parameters are determined using simulated method of moments. We calibrate the TFP growth process using moments of GDP growth data. We calibrate signal noises using moments of GDP growth forecasts (from the SPF dataset). The parameter values are reported in Table 2 and the calibration moments are in Table 3.

To evaluate the ability of the model to generate plausible uncertainty shocks, we pick one empirical measure of each type of uncertainty from Section 1 for comparison. For macro uncertainty we use the VIX, for micro dispersion, we use the sales growth dispersion and for higher-order uncertainty we use SPF forecast dispersion. Our model offers a direct analog to micro dispersion and higher-order uncertainty measures. But without a financial sector, we can't compute the VIX. Instead, we compare macro uncertainty (12) with the fluctuations in the VIX. We detrend all data so that we abstract from changes at frequencies that we are not attempting to explain.[15] We put both the real data and the output of the model in percentage deviation from trend units. To compute model output, we estimate moments or regression coefficients in each of the 2000 simulations, using the same number of observations for each series as in the data, and then report the mean (and/or standard deviation) of each moment or coefficient. Further calibration and simulation details are in the appendix.

4.4 Quantitative Results

The model can generate shocks to uncertainty and dispersion that are quantitatively similar to the data in terms of (i) their magnitude, (ii) their correlations with each other and (iii) their correlations with the business cycle. For each model, Table 4 reports the standard deviation of each uncertainty and dispersion measure, which is our measure of the size of uncertainty shocks, and the correlation of each measure with the other measures and GDP growth.[16]

As a benchmark, column (1) reports moments for the perfect information model. Because signals inform agents perfectly about the aggregate state X_t, they can forecast GDP growth perfectly and there is no macro uncertainty. Since all agents have the same beliefs, there is no higher-order uncertainty and no micro dispersion.

When we add learning and disaster risk to the model (column 3), the model generates plausible shocks to uncertainty and dispersion. Macro uncertainty shocks are about 65%

[15] We detrend using a HP filter exactly as we did in Section 1.

[16] We omit means of macro uncertainty and micro dispersion. The mean of macro uncertainty is 2.50 in the disaster risk model and the mean of the empirical proxy is 19.12. But a comparison between these numbers has no economic meaning because the data measure, the VIX, has units that are not comparable to the model. For the mean of micro dispersion, the problem is that size, age, industry, geography, all affect firm growth. Because the model's only source of firm heterogeneity is information, the model generates significantly less micro dispersion: 0.51 in the disaster risk model, compared to 18.59 in the data, on average. The appendix extends the model to match average dispersion. Because this additional complexity is not essential to our main point about uncertainty covariance, we suppress it in the main text. To compare other moments, we have detrended all series, as described in Section 4. Only the mean of higher-order uncertainty is reported, because it is a calibration target.

	Models			Data
	Perfect info.	Normal	Disaster risk	
	(1)	(2)	(3)	(4)
	(a) Macro uncertainty			
Std.	0	10.74	13.77	20.87
Corr. with GDP growth	0	0.00	−0.19	−0.26
Period	1962Q2–2008Q2			
	(b) Higher-order uncertainty			
Std.	0	18.49	22.69	31.13
Corr. with GDP growth	0	0.00	−0.12	−0.28
Corr. with Micro Unc.	0	0.42	0.75	0.43
Corr. with Macro Unc.	0	0.99	0.98	0.24
Period	1968Q3–2011Q3			
	(c) Micro dispersion			
Std.	0	10.29	13.83	11.58
Corr. with GDP growth	0	0.00	−0.06	−0.52
Corr. with Macro Unc.	0	0.43	0.76	0.32
Period	1962Q1–2009Q3			

Table 4: **Simulation results and data counterparts.** The three models are described in Sections 2 and 4. We use the VIX, sales growth and SPF measures of macro uncertainty, micro dispersion and higher-order uncertainty, resepctively, for the data. These are described in Section 1 and the appendix. All results are computed using the detrended series just as for the empirical analysis in Section 1. The periods are the periods for the data. The model samples are simulated to be the same length, as described in the appendix.

as large as in the data, higher-order uncertainty shocks are 70% as large and we get slightly more micro dispersion shocks than the data. All uncertainty series are also countercyclical and positively correlated with each other, with the magnitude of most of these correlations similar to their empirical counterparts.

By comparing the results in columns (2) and (3), we can evaulate the separate roles that the learning and disaster risk play in the results. Column (2) shows that learning alone generates 70–80% of the fluctuations in uncertainty and dispersion of the full model and also generates most of the positive correlations between the uncertainty and dispersion series. Adding disaster risk to the model (going from column (2) to column (3)) amplifies the uncertainty shocks and makes them countercyclical.

How much of the relationship between uncertainty and dispersion measures can the

	Macro Uncertainty	Higher-order Uncertainty	Micro Dispersion
Macro uncertainty		1.64^{***}	0.98^{***}
		(0.05)	(0.07)
Higher-order uncertainty	0.59^{***}		0.58^{***}
	(0.02)		(0.045)
Micro Dispersion	0.59^{***}	0.97^{***}	
	(0.06)	(0.11)	

Table 5: **Uncertainty shocks are correlated in the model, even after controlling for the business cycle.** These results are summary statistics for β from equation (1). We simulated the model 2000 times and run the regressions for each simulation. 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). Each series uses a simulation period that is analogous to the periods that the corresponding data series covers as described in detail in the appendix.

model explain, after controlling for the business cycle? To answer this question, we estimate the same equation (1) from Section 1 on model output, as we did on real data. Table 5 compares regression coefficients estimated from the model to those from the data. The columns determine the left hand side variable and the rows are the right hand side variable. For example, the first number in the macro uncertainty column says that when higher-order uncertainty deviates from trend by 1 percentage point, macro uncertainty deviates by 0.58 percentage points. Taken together, the results show that the model generates a positive and significant relationship between all three types of uncertainty and dispersion, above and beyond the business cycle, just as we found for the data.

It appears that our mechanisms, learning and disaster risk, are more than strong enough to explain the empirical comovement of uncertainty measures. In fact weakening the relationship between higher-order and micro dispersion, by adding more unpredictable variation in micro dispersion, might offer an even better fit to the data.

5 Conclusions

Fluctuations in micro dispersion, macro uncertainty and higher-order uncertainty have been used to explain recessions, asset price drops and financial crises. While all are dubbed "uncertainty shocks," they are distinct phenomena. 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.

We explore various possible causes and find that volatile macro outcomes create macro uncertainty, higher-order uncertainty and micro dispersion. The key to this mechanism is that the public information—past outcomes—become less informative predictors of the future, relative to private information. This makes agents put less weight on the public information, more weight on private information, and leads them to disagree more. Furthermore, when weak macro outcomes make highly-uncertainty disaster outcomes more likely, uncertainty of all types move in a correlated, volatile and countercyclical way. By offering a unified explanation for the origin of uncertainty shocks, our results provide insight into the nature of the shocks that drive business cycles.

References

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

Angeletos, M. and J. La'O (2014): "Sentiments," *Econometrica*, 81, 739–779.

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 (2013): "'Wait-and-See' Business Cycles?" *Journal of Monetary Economics*, 60, 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, 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, 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Sapora-Eksten, and S. Terry (2016): "Really Uncertain Business Cycles," Working paper.

Christiano, L., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

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 (2016): "Uncertainty Traps," *Quarterly Journal of Economics*, forthcoming.

Foster, L., J. Haltiwanger, and C. J. Krizan (2001): "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in *New Developments in Productivity Analysis*, ed. by C. Hulten, E. Dean, and M. Harper, University of Chicago Press, 303–372.

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

Jurado, K., S. Ludvigson, and S. Ng (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.

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

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2017): “The Tail that Wags the Economy: Beliefs and Persistent Stagnation,” NBER 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.

Mankiw, N. G. and R. Reis (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117, 1295–1328.

Nimark, K. (2014): “Man-Bites-Dog Business Cycles,” *American Economic Review*, 104, 2320,2367.

Ordonez, G. (2011): “The Asymmetric Effects of Financial Frictions,” *Journal of Political Economy*, 121, 844–895.

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*, 67, 1219–1264.

Reis, R. (2006): “Inattentive producers,” *Review of Economic Studies*, 73, 793–821.

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 Theoretical Appendix

A.1 GDP Forecasts

In this section we derive GDP and GDP growth forecasts, which we use in both Section 3 and for computing the quatitative results.

Using equation (10) and the fact that TFP growth forecasts are normally distributed in the cross section, it follows from equation (11) that:

$$Q_t = A_t \gamma^{1/(1-\gamma)} \exp\left[\frac{1}{\gamma-1}\Bigg(\log A_{t-1} + (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1}(\Delta a_t + \eta_{t-1}) + \frac{\omega_{t-1}^2\sigma_\psi^2}{2(\gamma-1)} + \frac{1}{2}V_{t-1}[\Delta a_t]\Bigg)\right]. \quad (19)$$

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 (19) that are known at the end of period $t-1$ from those that are unknown,

$$Q_t = \Gamma_{t-1}\exp[f(\Delta a_t, \eta_{t-1})], \quad (20)$$

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

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

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

$$E[Q_t|\mathcal{I}_{i,t-1}] = \Gamma_{t-1}\exp\Big(E[f(\Delta a_t, \eta_{t-1})|\mathcal{I}_{i,t-1}] + \frac{1}{2}V[f(\Delta a_t, \eta_{t-1})|\mathcal{I}_{i,t-1}]\Big). \quad (21)$$

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:

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

Therefore

$$E[f(\Delta a_t, \eta_{t-1})|\mathcal{I}_{i,t-1}] = (1-\omega_{t-1})E[\Delta a_t|A^{t-1}] + \omega_{t-1}z_{i,t-1} + \Big(\frac{\omega_{t-1}}{\gamma-1}\Big)\Bigg(\frac{(\sigma_t^2+\sigma_\eta^2)^{-1}E[\Delta a_t|A^{t-1}] + \sigma_\psi^{-2}z_{i,t-1}}{(\sigma_t^2+\sigma_\eta^2)^{-1} + \sigma_\psi^{-2}}\Bigg). \quad (22)$$

Let $\Lambda_t \equiv [\Delta a_t - E[\Delta a_t|A^{t-1}], \ \ \eta_{t-1}]'$. Then the variance term in equation (21) is

$$V[f(\Delta a_t, \eta_{t-1})|\mathcal{I}_{i,t-1}] = \left[\Big(1+\tfrac{\omega_{t-1}}{\gamma-1}\Big) \quad \tfrac{\omega_{t-1}}{\gamma-1}\right] V[\Lambda_t|\mathcal{I}_{i,t-1}] \left[\Big(1+\tfrac{\omega_{t-1}}{\gamma-1}\Big) \quad \tfrac{\omega_{t-1}}{\gamma-1}\right]' \quad (23)$$

where

$$V[\Lambda_t|\mathcal{I}_{i,t-1}] = \begin{bmatrix}\sigma_t^2 & 0\\ 0 & \sigma_\eta^2\end{bmatrix} - \begin{bmatrix}\sigma_t^2\\ \sigma_\eta^2\end{bmatrix}(\sigma_t^2+\sigma_\psi^2+\sigma_\eta^2)^{-1}\begin{bmatrix}\sigma_t^2 & \sigma_\eta^2\end{bmatrix}. \quad (24)$$

Together equations (21), (22), (23) and (24) define agent i's forecast of period t GDP.

The GDP growth forecast of agent i is

$$E[\Delta q_t|\mathcal{I}_{i,t-1}] = E[f(\Delta a_t, \eta_{t-1})|\mathcal{I}_{i,t-1}] + \log\Gamma_{t-1} + \log Q_{t-1}.$$

The expectation can be computed using equation (22).

A.2 Proofs of Results

Proof of Result 1 To prove this result we use the derivatives of $\mathcal{U}_t$, $\mathcal{H}_t$ and $\mathcal{D}_t$ with respect to $\sigma^2_{\eta,t-1}$ since these derivatives have the same sign as the derivative with respect to $\sigma_{\eta,t-1}$ and it simplifies the math. To show that $\partial\mathcal{D}_t/\partial\sigma_{\eta,t-1} < 0$ partial differentiate equation (14) with respect to $\sigma^2_{\eta,t-1}$ and use the fact that $\partial\omega_{t-1}/\partial\sigma^2_{\eta,t-1} < 0$ from equation (8).

For the effect of public signal noise on higher-order uncertainty start by taking the partial derivative $\partial\mathcal{H}_t/\partial\sigma^2_{\eta,t-1}$ and dropping terms that are not needed to sign the derivative. To simplify notation from now on I will drop the time subscripts from ω_{t-1}, $\sigma_{\eta,t-1}$ and $\sigma_{\psi,t-1}$. We get

$$\operatorname{sgn}\left\{\frac{\partial\mathcal{H}_t}{\partial\sigma^2_\eta}\right\} = \operatorname{sgn}\left\{\frac{\partial\omega}{\partial\sigma^2_\eta}\left(1+\frac{1}{(\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]}\right)+\frac{\omega(\gamma-1)\sigma^2_\psi}{(\sigma^2_t+\sigma^2_\eta)^2}\left(\frac{1}{(\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]}\right)^2\right\}$$

where $\operatorname{sgn}\{\cdot\}$ is the signum function that extracts the sign of a real number. Computing $\partial\omega/\partial\sigma^2_\eta$, substituting in the expressions for this and ω and dropping more terms that are not needed to sign the derivative gives

$$\begin{aligned}\operatorname{sgn}\left\{\frac{\partial\mathcal{H}_t}{\partial\sigma^2_\eta}\right\} &= \operatorname{sgn}\Bigg\{-\frac{\sigma^4_t}{\sigma^2_t+\sigma^2_\eta+\sigma^2_\psi}\left(1+\frac{1}{(\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]}\right)\\
&\qquad\qquad+\frac{\sigma^2_\psi}{(\sigma^2_t+\sigma^2_\eta)^2}\left(\frac{1}{(\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]^2}\right)\Bigg\}\\
&= \operatorname{sgn}\left\{-\frac{\sigma^4_t\big((\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]+1\big)}{\sigma^2_t+\sigma^2_\eta+\sigma^2_\psi}+\frac{\sigma^2_\psi}{(\sigma^2_t+\sigma^2_\eta)[\sigma^2_\psi+\sigma^2_t+\sigma^2_\eta]}\right\}\\
&= \operatorname{sgn}\left\{-\sigma^4_t\big((\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]+1\big)+\frac{\sigma^2_\psi}{\sigma^2_t+\sigma^2_\eta}\right\}.\end{aligned}$$

Therefore $\partial\mathcal{H}_t/\partial\sigma^2_\eta < 0$ if

$$\sigma^4_t\big((\gamma-1)[\sigma^2_\psi(\sigma^2_t+\sigma^2_\eta)^{-1}+1]+1\big) > \frac{\sigma^2_\psi}{\sigma^2_t+\sigma^2_\eta} \tag{25}$$

Part (a) of Result 1 follows from this condition. Since $\gamma > 1$ a sufficient condition for this to hold is

$$\sigma^4_t > \frac{\sigma^2_\psi}{\sigma^2_t+\sigma^2_\eta}. \tag{26}$$

For macro uncertainty, taking the partial derivative with respect to σ^2_η and dropping terms that are not needed to sign this derivative gives

$$\operatorname{sgn}\left\{\frac{\partial\mathcal{U}_t}{\partial\sigma^2_\eta}\right\} = \operatorname{sgn}\left\{(\gamma-1)^2\sigma^4_t+\omega^2\sigma^4_\psi-2\omega\sigma^2_t\sigma^2_\psi[(\gamma-1+\omega)\sigma^2_t(\sigma^2_t+\sigma^2_\psi+\sigma^2_\eta)+(\gamma-1)]\right\}.$$

Therefore $\partial\mathcal{U}_t/\partial\sigma^2_\eta > 0$ if

$$(\gamma-1)^2\sigma^4_t+\omega^2\sigma^4_\psi > 2\omega\sigma^2_t\sigma^2_\psi[(\gamma-1+\omega)\sigma^2_t(\sigma^2_t+\sigma^2_\psi+\sigma^2_\eta)+(\gamma-1)] \tag{27}$$

Part (b) of Result 1 follows from this condition. A sufficient condition for this to hold is that σ^2_ψ is sufficiently small:

$$\lim_{\sigma^2_\psi\to 0}\operatorname{sgn}\left\{\frac{\partial\mathcal{U}_t}{\partial\sigma^2_\eta}\right\} = \operatorname{sgn}\left\{(\gamma-1)^2\sigma^4_t\right\}.$$

Observe that in the limit as $\sigma^2_\psi \to 0$ the condition in equation (26) is also satisfied. Therefore in this case $\operatorname{cov}(\mathcal{U}_t,\mathcal{H}_t) < 0$, which proves the first part of Result 1(c). The second part of result 1(c) follows directly from conditions (25) and (27).

Proof of Result 2 To prove this result we use the derivatives of $\mathcal{U}_t$, $\mathcal{H}_t$ and $\mathcal{D}_t$ with respect to $\sigma^2_{\psi,t-1}$ since these derivatives have the same sign as the derivative with respect to $\sigma_{\psi,t-1}$ and it simplifies the math. As for the proof of Result 1 we drop the $t-1$ subscripts on ω, σ_ψ and σ_η.

Taking the partial derivative of $\mathcal{D}_t$ with respect to $\sigma^2_{\psi,t-1}$ and dropping terms that are not need to sign the derivative gives

$$\text{sgn}\left\{\frac{\partial \mathcal{D}_t}{\partial \sigma^2_\psi}\right\} = \text{sgn}\left\{1 - \frac{2\sigma^2_\psi}{\sigma^2_\eta + \sigma^2_\psi + \sigma^2_t}\right\}.$$

Therefore $\partial \mathcal{D}_t / \partial \sigma^2_\psi > 0$ if

$$1 > \frac{2\sigma^2_\psi}{\sigma^2_\eta + \sigma^2_\psi + \sigma^2_t} \tag{28}$$

For macro uncertainty, taking the partial derivative and dropping terms that are not needed to sign the derivative gives

$$\begin{aligned}\text{sgn}\left\{\frac{\partial \mathcal{U}_t}{\partial \sigma^2_\psi}\right\} = \text{sgn}\Big\{&2(\gamma - 1 + \omega)\frac{\partial \omega}{\partial \sigma^2_\psi}\sigma^2_t\sigma^2_\psi(\sigma^2_t + \sigma^2_\psi + \sigma^2_\eta) + (\gamma-1)^2\sigma^4_t \\ &+ 2(\gamma-1)\omega\sigma^2_t(\sigma^2_t + \sigma^2_\eta) + \omega^2(\sigma^2_t + \sigma^2_\eta)^2\Big\}.\end{aligned}$$

Using the fact that $\partial\omega/\partial\sigma^2_\psi = -\sigma_t^{-2}\omega^2$ this simplifies to

$$\text{sgn}\left\{\frac{\partial \mathcal{U}_t}{\partial \sigma^2_\psi}\right\} = \text{sgn}\Big\{(\gamma-1)^2\sigma^4_t + 2(\gamma-1)\omega\sigma^2_t(\sigma^2_t + \sigma^2_\eta) + \omega^2(\sigma^2_t + \sigma^2_\eta)^2 - 2(\gamma - 1 + \omega)\omega\sigma^2_\psi(\sigma^2_t + \sigma^2_\psi + \sigma^2_\eta)\Big\}.$$

Therefore $\partial \mathcal{U}_t / \partial \sigma^2_\psi > 0$ if

$$(\gamma-1)^2\sigma^4_t + 2(\gamma-1)\omega\sigma^2_t(\sigma^2_t + \sigma^2_\eta) + \omega^2(\sigma^2_t + \sigma^2_\eta)^2 > 2(\gamma - 1 + \omega)\omega\sigma^2_\psi(\sigma^2_t + \sigma^2_\psi + \sigma^2_\eta) \tag{29}$$

Using the same approach for higher order uncertainty the sign of it's derivative with respect to $\sigma^2_{\psi,t-1}$ satisfies:

$$\text{sgn}\left\{\frac{\partial \mathcal{H}_t}{\partial \sigma^2_\psi}\right\} = \text{sgn}\left\{\left(1 - \frac{2\sigma^2_\psi\omega}{\sigma^2_t}\right)\left((\gamma-1)\left(\frac{\sigma^2_\psi}{\sigma^2_t + \sigma^2_\eta} + 1\right) + 1\right) - \frac{2\sigma^2_\psi}{\sigma^2_\psi + \sigma^2_\eta + \sigma^2_t}\right\}.$$

Therefore $\partial \mathcal{H}_t / \partial \sigma^2_\psi > 0$ if

$$\left(1 - \frac{2\sigma^2_\psi\omega}{\sigma^2_t}\right)\left((\gamma-1)\left(\frac{\sigma^2_\psi}{\sigma^2_t + \sigma^2_\eta} + 1\right) + 1\right) > \frac{2\sigma^2_\psi}{\sigma^2_\psi + \sigma^2_\eta + \sigma^2_t} \tag{30}$$

To prove part (a) of Result 2 observe that conditions (28), (29) and (30) are all satisfied in the limit as $\sigma^2_\psi \to 0$. For (28) and (30) this can be seen by inspection. For (29) the limit of the left hand side is zero while the limit of the right hand side is greater than $(\gamma-1)^2\sigma^4_t > 0$.

Parts (b)–(d) of Result 2 follow directly from conditions (28), (29) and (30).

Proof of Result 3 This result follows directing from partial differentiating equations (12), (13) and (14) with respect to σ_t and using the fact that $\partial\omega_{t-1}/\partial\sigma_t > 0$, which follows from equation (8).

B Additional Empirical Details

In this section we provide additional details of the data series that we use in Section 1 and how we compute the empirical results, and also provide some extra results. We start by describing the data sources and the construction of each uncertainty series.

VIX The VIX measure of macro uncertainty comes from Bloom (2009). 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 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. The series covers 1962Q2–2008Q2.

Forecast errors The forecast errors series is constructed using data from the Survey of Professional Forecasters. This dataset provides one period ahead forecasts of quarterly real GDP (forecasts of quarter $t+1$ real GDP made after the end of quarter t). We use this data to construct a real GDP growth forecast for each forecaster for each period and then compute the average absolute error of GDP growth forecasts, which is our forecast errors series.[17] The data allows us to construct this variable for 1968Q3–2011Q3. The average number of forecasters in a quarter is 41, with a standard deviations of 17.

JLN The JLN measure of macro uncertainty is taken from Jurado et al. (2015). That paper measures the h-month ahead uncertainty of a variable y_{jt} as

$$\mathcal{U}_{jt}^{y}(h) \equiv \sqrt{E\left[\left(y_{j,t+h} - E\left[y_{j,t+h}|I_t\right]\right)^2 \Big| I_t\right]},$$

where I_t is the information set available to agents at time t, which they construct using a very large set of data. The authors measure the uncertainty about many variables and then take a weighted average of them to construct their measure of macro uncertainty. We use the measures from this paper to construct a quarterly measure of macro uncertainty and an annual measure. For the quarterly measure we take the 3-month head measure of uncertainty and use the December, March, June and September observations for Q1, Q2, Q3 and Q4 measures of uncertainty, respectively. For the annual series we take the December observation in year $t-1$ of 12-month ahead uncertainty as the measure of year t uncertainty. Our quarterly series covers 1960Q4–2016Q4 and our annual series covers 1961–2016.

Micro dispersion measures Our three measures of micro dispersion—the sales growth, stock return and TFP shocks series—are all taken from Bloom et al. (2016). Full details of these series are provided in that paper, so we provide a brief summary here. The sales growth series is the interquartile range of firms sales growth computed using data on all public firms with at least 100 quarters of data in Compustat between 1962 and 2010. The sample contains 2,465 firms. The growth rate of firm i in quarter t is measured as

$$\frac{Q_{i,t+4} - Q_{it}}{\frac{1}{2}(Q_{i,t+4} + Q_{it})}, \tag{31}$$

where Q_{it} is the sales of firm i in quarter t. This series is quarterly and covers 1962Q1–2009Q3. The stock returns series is the interquartile range of monthly stock returns for all public firms with at least 25 years of data in CRSP between 1960 and 2010. We make the series quarterly by averaging it across the three months within each quarter. The series covers 1960Q1–2010Q3. The TFP shocks series is constructed using data on establishments with at least 25 years of observations in the Census of Manufacturers or Annual Survey of Manufacturers from the U.S. Census Bureau. There are 15,673 establishments in the sample and the data is annual covering 1972–2011. TFP pf establishment j ($\hat{z}_{j,t}$) is measured using the method from Foster et al. (2001) and TFP shocks are measured as the residuals from:

$$\log(\hat{z}_{j,t}) = \rho \log(\hat{z}_{j,t-1}) + \mu_j + \lambda_t + e_{j,t}.$$

μ_j and λ_t are establishment and time fixed effects, respectively. The TFP shocks uncertainty series is the interquartile range of these TFP shocks. We remove Census of Manufacturers year fixed effects from the series by regressing it on a constant and a dummy variable for the years in which these Censuses occur and subtracting the Census year effect from the observations for these years.

SPF forecasts The SPF forecasts series is constructed using the same data from the Survey of Professional Forecasters as the forecast errors series. For this series we compute the standard deviation of real GDP growth forecasts for each quarter. The forecast for quarter t GDP growth is made at the end of quarter $t-1$. The series covers 1968Q3–2011Q3.

[17] The quarter t GDP growth forecasts of forecaster i is $\log(E_i[Q_t]) - \log(Q_{t-1})$ where Q_t is quarter t real GDP.

(a) Macro uncertainty	VIX	Forecast errors	JLN
Micro dispersion			
Sales growth	0.27	0.30	0.62
TFP shock	0.34	0.42	0.39
Stock return	0.56	0.27	0.50
Higher-order uncertainty			
SPF forecasts	0.34	0.14	0.30
Blue Chip forecasts	0.10	−0.04	0.43
(b) Micro dispersion	Sales growth	TFP shocks	Stock returns
Macro uncertainty			
VIX	0.27	0.34	0.56
Forecast errors	0.30	0.42	0.27
JLN	0.62	0.39	0.50
Higher-order uncertainty			
SPF forecasts	0.43	0.04	0.33
Blue Chip forecasts	0.49	0.26	0.21
(c) Higher-order uncertainty	SPF forecasts	Blue Chip forecasts	
Macro uncertainty			
VIX	0.34	0.10	
Forecast errors	0.14	−0.04	
JLN	0.30	0.43	
Micro dispersion			
Sales growth	0.43	0.49	
TFP shock	0.04	0.26	
Stock return	0.33	0.21	

Table 6: **Correlations of uncertainty measures.** All uncertainty series are detrended.

Blue Chip forecasts The Blue Chip forecasts series is constructed using data from the Blue Ship Economic Indicators dataset. Like the Survey of Professional Forecasters this dataset provides real GDP growth forecasts made by professional forecasters. We use this data to compute the standard deviation of forecasts of real GDP growth for each quarter, with the forecasts made at the start of the quarter. The series covers 1984Q3 to 2016Q4.

Additional details for empirical analysis In the paper we present results on the correlations between the uncertainty series and regressions of the series on each other controlling for GDP growth. The full set of correlation results are provided in Table 6. For both the correlation and regression analysis it is necessary to ensure that the various series are measuring uncertainty over the same horizon. For correlations and regressions that use pairs from the following uncertainty series no adjustments are required as they all measure uncertainty over a one quarter horizon: VIX, forecasts erros, JLN uncertainty, stock returns, SPF forecasts and Blue Chip forecasts. For correlations and regressions using one of these 6 series and the sales growth dispersion measure we make an adjustment because sales growth dispersion is computed using sales growth over four quarters. To make the other series match up with the horizon of sales growth we average them over the same four quarters that sales growth is measured over.

The TFP shocks series measures micro dispersion over an annual horizon. The make the horizon of the six quarterly series match up with this we average them over the 4 quarters of each year. To match the horizon of the sales growth series with this we use the observation for the last quarter of year $t-1$ (which uses growth over the subsequent four quarters) to measure dispersion for year t.

As noted in the text, for all of the regressions and correlations we detrend the series using a HP filter. For analysis at quarterly and annual frequencies we set the smoothing parameter equal to 1600 and 100, respectively.

C Calibration and Simulation Details

Disaster risk model calibration The disaster risk model has nine 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), four parameters for the X_t process (α_0, α_1, ρ and ϕ), and public and private signal noise (σ_η and σ_ψ respectively). We set $\gamma = 2$, which corresponds to a Frisch labor supply elasticity of one. This is within the range of Frisch elasticities that Keane and Rogerson (2012) argue are reasonable at the macro level. α_0 can be normalized and we set it equal to 0. 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; and the average (over forecasters and over time) absolute error of GDP growth forecasts, which is

$$\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{i\in I_t}|E[\Delta q_t|\mathcal{I}_{i,t-1}] - \Delta q_t|, \tag{32}$$

where T is the number of periods in the data. The forecast data for the calibration is taken from the Survey of Professional Forecasters.

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 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. 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. For each simulation we compute moments using the the same number of periods of observations as the data that we're matching has (173 periods for 1968Q4–2011Q4) so that model and data are comparable. Each simulation consists of: simulating X_t for an initial burn in period; simulating the X_t process for 173 periods that we will use for computing moments; computing prior beliefs for each period; 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 (11); computing the five moments of GDP growth that we use for calibration; computing higher-order uncertainty using the definition of GDP growth forecasts in footnote 17 (so that the model and data are consistent) and equation (21), and then computing its mean; and computing the average absolute error of GDP growth forecasts for each period using the definition of GDP growth forecasts in footnote 17 and equation (11), 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.

The 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.

Disaster risk model calibration For the normal model we no longer have the parameters b and c, but α_0 can no longer be normalized. So there is one less parameter to calibrate. For 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, $\mathrm{ac}(|\Delta q_t - \overline{\Delta q}|)$. We drop skewness and kurtosis because without 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 model is able to match the target moments well, but not exactly for the same reasons that have been discussed for the disaster risk model. We do not do a separate calibration for the perfect information model. It is sufficient for our purposes to use the parameter values of the normal model and change the signal noises to $\sigma_\eta = \sigma_\psi = 0$.

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 dispersion and macro uncertainty. Macro uncertainty can be computed by Gaussian quadrature using equation (12). The measure of micro dispersion cannot be computed by Gaussian quadrature. Instead for each simulated sample we compute sample paths for 2000 firms and compute micro dispersion using the same definition as we use for the data: the interquartile range of firm sales growth, with sales growth defined in equation (31). Computing sample paths requires drawing signals for each firm each period and computing their output decisions.

As for the calibration, for each simulation we compute moments of the model using the same number of periods as we have for the relevant data. The data that we are evaluating the model with covers 1962Q1 to 2011Q4 so this means generating 200 periods of data for each simulation. For each moment of the model we only use the observations that are analogous to the data. For example, the micro dispersion series covers 1962Q1 to 2009Q3 so we use periods 1 to 191 from the model. Just like for the calibration, we run 2000 simulations and average each moment across the them.

D Extension with Parameter Learning

In this section we describe how the results would change if the agents did not know the parameters of the GARCH process for TFP growth. We provide this extension since having agents who do not know the true model parameters and estimate them in real time is both realistic and amplifies uncertainty fluctuations.

The modification to the environment is as follows. Take the normal and disaster risk models that are presented in Sections 2 and 4 of the paper. Our baseline assumption is that agents know the parameters of the GARCH process for the variance of X_t: α_0, α_1, ρ and ϕ in equation (16). We now drop this assumption.[18] Agents therefore need to estimate these parameters in order to form beliefs about TFP growth. Specifically the process for forming beliefs at the end of period $t-1$ is as follows. Using the observed history of TFP, A^{t-1}, agents estimate the process for Δa_t (X_t for the disaster risk model) by maximum likelihood and based on the parameter estimates from this they form their prior beliefs about Δa_t (X_t). These beliefs are normally distributed with mean $E[\Delta a_t|A^{t-1}]$ ($E[X_t|A^{t-1}]$) and variance $v_{t-1} = V[\Delta a_t|A^{t-1}]$ ($v_{t-1} = V[X_t|A^{t-1}]$). Given these priors the agents then update their beliefs using their signals in the same way as in the model presented in the paper, that is according to equations (7) and (9).

Dropping the assumption that the parameters of the GARCH process are known adds additional scope for prior beliefs to fluctuate over time. Shocks to Δa_t (X_t) will now not only affect its future variance through the GARCH process, but will cause agents to re-estimate the parameters of the process. As the estimated parameters change agents beliefs about future values of Δa_t (X_t) will change, which is an additional source of uncertainty shocks. Note that conditional on the sequence of prior beliefs, the model with parameter learning is identical to the model used in the quantitative exercise in the paper. Thus uncertainty shocks due to changes in beliefs about parameters have the same qualitative effects as uncertainty shocks due to

[18] We still assume that agents the values of parameters b and c. We do this so that the model is computationally feasible. 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.

changes in σ_t int he model in which parameters are known. The question is how much parameter learning changes the magnitude of shocks to the economy.

To assess the quantitative contribution of parameter learning to the results we repeat the quantitative exercise from the body of the paper for this extension of the model. Before presenting the results there are a few comments to make on how we perform the quantitative exercise with parameter learning. When we solve this model numerically we make two simplifying assumptions. 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 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.

For the calibration of the models with parameter learning we use the same parameters as for the models in the body of the paper. We take this approach since the calibration moments of the models change negligibly when parameter learning is added to the model. For simulating the model to compute results it is important that the model is simulated for the same length of time as the data so that agents have access to the same amount of information for estimating parameters. 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 use for assessing the model cover 1962Q1–2011Q4. This means that for the first period in which we compute results we give agents the previous 59 periods of data (corresponding to 1947Q2–1961Q4) to estimate the GARCH process with, and each period after that they get one additional period of data.

The results for these models are presented in Table 7. We also reproduce the results for the models without parameter learning for comparison. First look at the results for the normal model, in columns (1) and (2) of the table. In column (1) parameter updating is off and in column (2) it is turned on. Comparing the results in these columns shows that the main function of parameter learning is to generate additional uncertainty shocks. The standard deviations of micro dispersion, higher-order uncertainty and macro uncertainty increase by 14%, 8% and 11%, respectively, when parameter updating is turned on. The results for the two versions of the disaster risk model, presented in columns (3) and (4) of Table 7, tell the same story. Turning parameter updating on increases the standard deviations of micro dispersion, higher-order uncertainty and macro uncertainty by 9%, 9% and 10%, respectively.

E Extension with Additional Firm Heterogeneity

In this section we address the inability of the models presented in the main text to match the mean level of micro dispersion 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 dispersion.

The change that we make is to modify the TFP growth process to allow for additional heterogeneity amongst firms. Specifically, TFP is assumed to be idiosyncratic and to follow the process:

$$\Delta a_{it} = \alpha_{it} + \theta_i \sigma_t \epsilon_t, \tag{33}$$

where $\alpha_{it} \sim N(\alpha_0, \sigma_\alpha^2)$ 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 σ_t follows the same process as in equation (16). This TFP growth 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 α_{it} is now a random variable rather than a constant captures the idea that firms experience transitory idiosyncratic shocks in addition

Model	Normal		Disaster risk	
Parameter updating	Off	On	Off	On
	(1)	(2)	(3)	(4)
(a) Macro uncertainty				
Std.	10.74	11.93	13.77	15.15
Corr. with GDP growth	0.00	0.00	−0.19	−0.18
Period	1962Q2–2008Q2			
(b) Higher-order uncertainty				
Std.	18.49	19.99	22.69	24.68
Corr. with GDP growth	0.00	0.00	−0.12	−0.12
Corr. with Micro Unc.	0.42	0.43	0.75	0.68
Corr. with Macro Unc.	0.99	0.99	0.98	0.98
Period	1968Q3–2011Q3			
(a) Micro dispersion				
Std.	10.29	11.68	13.83	15.08
Corr. with GDP growth	0.00	0.00	−0.06	−0.06
Corr. with Macro Unc.	0.43	0.44	0.76	0.68
Period	1962Q1–2009Q3			

Table 7: **Separating the roles of state uncertainty and parameter updating.** The three models are described in Sections 2 and 4. All results are computed using the detrended series just as for the empirical analysis in Section 1. Parameter updating off means that agents know the values of α_0, α_1, 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.

to being affected by aggregate shocks. To maintain the learning structure of the model signals are now

$$z_{i,t-1} = \sigma_t \epsilon_t + \eta_{t-1} + \psi_{i,t-1},$$

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_α^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 dispersion. 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 dispersion and do this in the same way as described in earlier in this appendix; 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 8 and the results in Table 9. The results show that the model is capable of matching the mean and standard deviation of micro dispersion 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 cyclicality can be brought closer to the data by adding the disaster risk mechanism to this model.

Parameter		*Value*
TFP growth level	α_0	0.0048
Std. transient idiosyncratic shocks	σ_α	0.1459
Std. sensitivity to agg. shocks	σ_θ	4.9414
mean volatility	α_1	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 8: **Parameter values for the extended model**

	Extended Model	Data
Micro dispersion		
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 9: **Simulation results for extended model and data counterparts.** Micro dispersion is measured with the sales growth series that is described in Section 1. The mean of micro dispersion is computed using the raw data. The rest of the results are computed using the detrended series. The periods are the periods for the data. The model samples are simulated to be the same length, as described in Appendix C.