Understanding Uncertainty Shocks and the Role of Black Swans

Anna Orlik and Laura Veldkamp [*]

September 23, 2013

Abstract

A key question in macroeconomics is what shocks drive business cycles. A recent literature explores the effect of uncertainty shocks. But where do uncertainty shocks come from? Researchers often estimate stochastic volatility, using all available data, and equate it with uncertainty. If we define uncertainty to be the conditional standard deviation of a forecast error, then this stochastic volatility corresponds to the uncertainty of an agent who knows his model's parameters, well before they could ever be estimated, and is certain they are true. We show that a Bayesian forecaster who revises model parameters in real time and accounts for estimation uncertainty can experience large, counter-cyclical uncertainty shocks, even if his model has constant-volatility innovations. The exercise teaches us that large uncertainty shocks need not come from exogenous changes in variance. They may also come from "black swans": Events that are unlikely under the previous period's estimated model and cause agents to significantly revise their beliefs about the probability distribution of future outcomes.

Some times feel like uncertain times for the aggregate economy. At other times, events appear to be predictable, volatility is low, confidence is high. An active emerging literature argues that changes in uncertainty can explain asset pricing, banking crises, business cycle fluctuations, and the 2007 recession in particular. Uncertainty shocks are typically measured as innovations in a GARCH or stochastic volatility model, forecast dispersion or a price of a volatility option. But none of these measures captures an essential source

[*]Preliminary draft. Please send comments to anna.a.orlik@frb.gov and lveldkam@stern.nyu.edu. We are grateful to the seminar participants at the NBER Summer Institute, NYU Alumni Conference, North American meetings of the Econometric Society, SED meetings in Seoul, AEA meetings, Becker-Friedman Institute Policy Uncertainty Workshop, the NYU macro lunch, Tom Stark, and our AEA discussant, Lars Hansen. Many thanks also to Isaac Baley, David Johnson, Callum Jones and Nic Kozeniauskas for outstanding research assistance. We acknowledge grant assistance from the Stern Center for Global Economy and Business. 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.

of uncertainty, that no one knows the true model of the world. People in 1950 could not not estimate a model on a 65-year post-war data sample, as an econometrician does when inferring 1950 volatility today. In 2006, financial analysts' models told them that a nationwide decline in house prices was nearly impossible. Those models proved incorrect. Models are selected and estimated with error. We use real GDP data and Bayesian methods to measure uncertainty of a forecaster who considers estimation errors to be a source of uncertainty. Even when this forecaster uses a model which rules out the possibility of changes in the volatility of its innovations, uncertainty shocks still arise from the learning process itself.

Allowing for model uncertainty uncovers a new source of uncertainty shocks. It reveals that the largest shocks to uncertainty often come from events we dub "black swans." Under the probability distribution estimated based on the previous period's information set, these events would be highly improbable. Upon observing such an event, Bayes' law leads our forecaster to place a lower probability weight on his previous forecasting model and more weight on models that would make the observed event less unlikely. This shift of probability weight among various models raises model uncertainty. Even if the forecaster is certain that the true economic model has innovations with constant variance, the black swan causes our forecaster to be less certain about what stochastic process generates the data he sees, which is an uncertainty shock.

We define macroeconomic uncertainty as the variance of next-period GDP growth y_{t+1}, conditional on all information observed through time-t: $Std[y_{t+1}|\mathcal{I}_t]$. We use this definition because in most models, this is the theoretically-relevant moment. When there is an option value of waiting, beliefs with a higher conditional variance (imprecise beliefs) raise the value of waiting to observe additional information. Thus, it is uncertainty in the form of a higher conditional variance that typically delays consumption or investment and thus depresses economic activity.

A key message of our paper is that understanding and measuring economic uncertainty requires relaxing the full-information assumptions of rational expectations econometrics. In econometrics, rational expectations typically describes the assumption that agents know the true law of motion of the economy. If agents have rational expectations, they have no uncertainty about their model. Their only uncertainty is about realizations of model innovations. To measure the uncertainty of such a forecaster, it makes sense to estimate a model on as much data as possible, take the parameters as given, and estimate the conditional standard deviation of model innovations. This is what stochastic volatility estimates typically are. But in reality, the macroeconomy is not governed by a simple, known model and we surely don't know its parameters. Instead, our forecast data (from the Survey of Professional Forecasters or SPF) suggests that forecasters estimate simple models to approximate complex processes and constantly use new data to update our beliefs. Forecasters are not irrational. They simply lack rational expectations because they do not know the economy's true data-generating process. In such a setting, uncertainty and volatility can behave quite differently.

To build up intuition for how estimation uncertainty behaves, we start with one of the simplest forecasting models: a linear model of GDP growth with a hidden Markov state and homoskedastic innovations. We hold the volatility of the innovations fixed so that we can isolate the effect of estimation uncertainty shocks. Each period t, our forecaster observes time-t GDP growth and uses the complete history of GDP data to estimate her model and forecast GDP growth in $t+1$. We compare the uncertainty of our forecaster to volatility. Volatility is equivalent to the uncertainty of an agent who estimated the same model on the full sample of data and took those estimates to be truth when forming forecasts and computing conditional variance. The agent with known model parameters faces no uncertainty shocks because the estimated model is a homoskedastic one. But the forecaster with model uncertainty experiences changes in the conditional variance of her forecasts. These are uncertainty shocks.

However, a linear model with normally distributed shocks generates only small uncertainty shocks and it badly misses a key feature of the data: The average forecast is significantly lower than the average GDP growth realization. For an unbiased forecaster with a linear model, this should not be the case. But if GDP growth is a concave transformation of a linear-normal underlying variable, expected values can be systematically lower than the realizations. Adding extra uncertainty, such as parameter uncertainty, amplifies this effect. Therefore, we explore a non-linear forecasting model and find that combining non-linear forecasting with parameter uncertainty generates medium-sized shocks to uncertainty, particularly in recessions. Finally, we allow agents to learn about the non-linearity in their forecasting model. This generates the largest uncertainty shocks.

We compare our model-based uncertainty series to commonly-used uncertainty proxies and find that it is less variable, but more persistent than the proxy variables. The most highly correlated proxies are the price of a volatility option (VIX) or forecast dispersion. But neither achieves more than a 50% correlation with our uncertainty measure.

Related literature A new and growing literature uses uncertainty shocks as a driving process to explain business cycles (e.g., Bloom, Floetotto, Jaimovich, Sapora-Eksten, and Terry (2012), Basu and Bundick (2012) Christiano, Motto, and Rostagno (2012), Bianchi, Ilut, and Schneider (2012)), to explain investment dynamics Bachmann and Bayer (2012a), to explain asset prices (e.g., Bansal and Shaliastovich (2010), Pastor and Veronesi (2012)), and to explain banking panics (Bruno and Shin, 2012). These papers are complementary to ours. We explain where uncertainty shocks come from, while these papers trace out the economic and financial consequences of the shocks.[1]

More similar to our exercise is a set of papers that measure uncertainty shocks in various ways. Bloom (2009), Baker, Bloom, and Davis (2012), Stock and Watson (2012), Jurado, Ludvigson, and Ng (2013) Fernandez-Villaverde, Kuester, Guerron, and Rubio-Ramirez (2012), Fernandez-Villaverde, Guerron-Quintana, Rubio-Ramirez, and Uribe (2011), Nakamura, Sergeyev, and Steinsson (2012), Justiniano and Primiceri (2008) and Born, Peter, and Pfeifer (2011) document the properties of uncertainty shocks in the U.S. and in emerging economies, while Bachmann, Elstner, and Sims (2012) use forecaster data to measure ex-ante and ex-post uncertainty in Germany. While our paper also engages in a measurement exercise, what we add to this literature is a quantitative model of how and why such shocks arise.

The theoretical part of our paper grows out of an existing literature that estimates Bayesian forecasting models with model uncertainty. Cogley and Sargent (2005) use such a model to understand the behavior of monetary policy, while Johannes, Lochstoer, and Mou (2011) estimate a similar type of model on consumption data to capture properties of asset prices. While the mechanics of model estimation are similar, the focus on non-linear filtering, uncertainty shocks and uncertainty proxies distinguish our paper. Nimark (2012)

[1] In contrast, Bachmann and Bayer (2012b) argue that there is little impact of uncertainty on economic activity.

also generates increases in uncertainty by assuming that only outlier events are reported. Thus, the publication of a signal conveys both the signal content and information that the true event is far away from the mean. Such signals can increase agents' uncertainty. But that paper does not attempt to quantitatively explain the fluctuations in uncertainty measures. Our paper is in the spirit of Hansen's Ely Lecture (Hansen, 2007) and Chen, Dou, and Kogan (2013) which advocate putting agents in the model on equal footing with an econometrician who is learning about his environment over time and struggling with model selection.

1 A Linear-Normal Forecasting Model

The purpose of the model is to explain why relaxing rational expectations and assuming that agents do not know the true distribution of outcomes with certainty opens up an additional source of uncertainty shocks. To isolate this new source of uncertainty shocks, we consider first a homoskedastic model. If an agent knew the true model (had rational expectations) would have uncertainty about random future shocks, but that uncertainty would be constant over time. There would be no uncertainty shocks.

We consider a forecaster who observes a real-time GDP growth data, in every quarter and forecasts the next period's growth. The agent knows that GDP growth comes from a linear model with a hidden Markov state, but does not know the parameters of this model. Each period, he starts with prior beliefs about these parameters and the current state, observes the new GDP data, and updates his beliefs using Bayes' law.

1.1 Definitions

A model, denoted $\mathcal{M}$, is a probability distribution over a sequence of outcomes. Let $y^t \equiv \{y_\tau\}_{\tau=0}^t$ denote a series of data available to the forecaster at time t. Models will differ in the information set available to forecasters and the driving process for y_t. In every model, the information set will include y^t, the history of y observations up to and including time t. The state S_t is never observed. Each model has a vector θ of parameters.

The agent, who we call a forecaster and index by i, is not faced with any economic choices. He simply uses Bayes' law to forecast future y outcomes. Specifically, at each date

t, the agent conditions on his information set $\mathcal{I}_{it}$ and forms beliefs about y_{t+1}. We call the expected value $E\left(y_{t+1}|\mathcal{I}_{it}\right)$ an agent i's *forecast* and the square root of the conditional variance $Var\left(y_{t+1}|\mathcal{I}_{it}\right)$ is what we call *uncertainty.* Forecasters' forecasts will differ from the realized growth rate. This difference is what we call a forecast error.

Definition 1. *An agent i's forecast error is the distance, in absolute value, between the forecast and the realized growth rate: $FE_{i,t+1} = |y_{t+1} - E[y_{t+1}|\mathcal{I}_{it}]|$.*

We date the forecast error $t+1$ because it depends on a variable y_{t+1} that is not observed at time t. Similarly, an average forecast error is

$$\bar{FE}_{t+1} = \frac{1}{N_t}\sum_{i=1}^{I_t} FE_{i,t+1}. \tag{1}$$

We define forecast errors and uncertainty over 1-period-ahead forecasts because that is the horizon we focus on in this paper. But future work could use these same tools to measure uncertainty at any horizon.

Definition 2. *Uncertainty is the standard deviation of the time-$(t+1)$ GDP growth, conditional on an agent's time-t information: $U_{it} = \sqrt{E\left[\left(y_{t+1} - E[y_{t+1}|\mathcal{I}_{it}]\right)^2 \middle| \mathcal{I}_{it}\right]}$.*

In settings where the forecaster's information set does not include the model or its parameters, we make a distinction between uncertainty and volatility. Volatility is the same standard deviation as before, but now conditional on the history y^t as well as the model $\mathcal{M}$ and the parameters θ:

Definition 3. *Volatility is the standard deviation of the unexpected innovations in y_{t+1}, taking the model and its parameters as given:*
$VOL_t = \sqrt{E\left[\left(y_{t+1} - E[y_{t+1}|y^t, \theta, \mathcal{M}]\right)^2 \middle| y^t, \theta, \mathcal{M}\right]}$.

Many papers equate volatility, uncertainty and squared forecast errors. These definitions allow us to understand the conditions under which these are equivalent. Volatility and uncertainty are both ex-ante measures because they are time-t expectations of $t+1$ outcomes, which are time-t measurable. However, forecast errors are an ex-post measure because it is not measurable at the time when the forecast is made. Substituting definition

1 into definition 3 reveals that $U_{it} = \sqrt{E[FE^2_{i,t+1}|\mathcal{I}_{it}]}$. So, uncertainty squared is the same as the expected squared forecast error. Of course, what people measure with forecast errors is typically not the *expected* squared forecast error. It is an average of *realized* squared forecast errors: $\sqrt{1/N_t \sum_i FE^2_{i,t+1}}$.

To compare volatility and uncertainty, we can examine definitions 3 and 2. When $\mathcal{I}_{it} = \{y^t, \theta, \mathcal{M}\}$, uncertainty and volatility are equivalent. Thus, when a forecaster knows his model and its parameters with certainty, volatility measures uncertainty. That leads us to define estimation uncertainty as the part of uncertainty that comes from relaxing rational expectations.

Definition 4. *The uncertainty generated by imperfect information about the model and its parameters (hereafter "estimation uncertainty") is* $\Delta = U - V$.

The difference between volatility and uncertainty is due to the difference between conditioning on an agent's information set, and conditioning on $\{y^t, \mathcal{M}, \theta\}$. If the agent's information set is $\mathcal{I}_{it} = \{y^t, \mathcal{M}, \theta\}$, then estimation uncertainty is zero. But if the agent is uncertain about the model or parameters, U and V will differ. Note that unlike U or V, Δ can be, and sometimes is, negative. For example, suppose the highest-probability model has high-variance innovations, but the probability of that model is just over 50%. Then accounting for model uncertainty will lead the agent to take some draws from the high-variance model and some from the low-variance model. The resulting set of outcomes may be less uncertain than they would be if the agent treated the high-variance model as if it was true.

1.2 The linear forecasting model

We begin by examining the following continuous-state hidden Markov process for y_t[2]

$$y_t = \alpha + s_t + \sigma\varepsilon_{y,t} \tag{2}$$

$$s_t = \rho s_{t-1} + \sigma_s \varepsilon_{s,t} \tag{3}$$

[2]We have also explored a version of the model with a discrete hidden state. The results are very similar and are reported in Appendix B.

where $\varepsilon_{y,t}$ and $\varepsilon_{s,t}$ are standard normal random variables independent of each other. In this model, the parameters are $\theta \equiv [\alpha, \rho, \sigma, \sigma_s]'$.

Information assumptions Each forecaster has an identical information set: $\mathcal{I}_{it} = \{y^t, \mathcal{M}\}$, $\forall i$. The model $\mathcal{M}$ is described by (2) and (3).

To compute the process for uncertainty, we use Bayesian updating to form the conditional variance in definition 2. When the parameters are known, (2) and (3) form the observation and state equations of a Kalman filtering system. The following recursive equations describe the conditional mean and variance of a Kalman system.

$$E[y_{t+1}|y^t, \theta, \mathcal{M}] = \rho E[y_t|y^{t-1}, \theta, \mathcal{M}] + \rho K_t y_t \tag{4}$$

where the term K_t is the Kalman gain

$$K_t = (I + \sigma^2 Var[y_t|y^{t-1}, \theta, \mathcal{M}])^{-1} \tag{5}$$

and the conditional variance of the estimate (the volatility) is

$$Var[y_{t+1}|y^t, \theta, \mathcal{M}] = \rho^2(I - K_t)(Var[y_t|y^{t-1}, \theta, \mathcal{M}] - \sigma^2 I) + (\sigma + \sigma_S^2)I. \tag{6}$$

Note that none of the parameters in the Kalman gain or conditional variance formulas are time-varying. That implies that conditional variance is constant. Setting $Var[y_{t+1}|y^t, \theta, \mathcal{M}] = Var[y_t|y^{t-1}, \theta, \mathcal{M}]$ and solving for K_t and variance yields

$$Var[y_{t+1}|y^t, \theta, \mathcal{M}] = \left((1 - \rho^2)\sigma^2 + \sigma_S^2\right) I \tag{7}$$

Notice that the square root of $Var[y_{t+1}|y^t, \theta, \mathcal{M}]$ is volatility. This teaches us that the volatility in this model is constant. That may or may not be a realistic feature of the data. But it is a helpful starting point because it will allow us to isolate the fluctuations in uncertainty that come from estimation uncertainty, rather than from changes in volatility. Later, we will re-introduce changes in volatility as well.

1.3 Data Description

There are two pieces of data that we use to evaluate and estimate our forecasting model. The first is real-time GDP data from the Philadelphia Federal Reserve. The variable we denote y_t is the growth rate of GDP. Specifically, it is the log-difference of the seasonally-adjusted real GDP series, times 400, so that it can be interpreted as an annualized percentage change. We use real-time data because we want to accurately assess what agents know at each date. Allowing them to observe final GDP estimates, that are not known until 2 years later, is not consistent with the goal. Therefore, y_t represents the estimate of GDP growth between then end of quarter $t-1$ and quarter t, based on the GDP estimates available at time t. Similarly, y^t is the history of GDP growth up to and including period t, based on the data available at time t.

We use the second set of data, professional GDP forecasts, to evaluate our forecasting models. We describe below the four key moments that we use to make that assessment. The data come from the Survey of Professional Forecasters, released by the Philadelphia Federal Reserve. The data are a panel of individual forecaster predictions of real US output for both the current quarter and for one quarter ahead from quarterly surveys from 1968 Q4 to 2011 Q4. In each quarter, the number of forecasters varies from quarter-to-quarter, with an average of 40.5 forecasts per quarter.

Formally, $t \in \{1, 2, \ldots, T\}$ is the quarter in which the survey of professional forecasters is given. Let $i \in \{1, 2, \ldots, I\}$ index a forecaster and $I_t \subset \{1, 2, \ldots, I\}$ be the subset of forecasters who participate in a given quarter. Thus, the number of forecasts made at time t is $N_t = \sum_{i=1}^{I} \mathbb{I}(i \in I_t)$. Finally, let y_{t+1} denote the GDP growth rate over the course of period t. Thus, if GDP_t is the GDP at the end of period t, observed at the start of quarter $t+1$, then $y_{t+1} \equiv ln(GDP_t) - ln(GDP_{t-1})$. This timing convention may appear odd. But we date the growth $t+1$ because it is not known until the start of date $t+1$. The growth forecast is constructed as $E_{it}[y_{t+1}] = ln(E_{it}[GDP_t]) - ln(GDP_{t-1})$.

1.4 Estimation and results

Our forecaster starts with a prior distribution of parameters, $p(\theta)$, described in table 1. The prior means are chosen to match the mean, variance, and persistence of the GDP

Table 1: Prior assumptions

parameter	mean	stdev
ρ	0.5	0.08
σ^2	3	0.5
σ_s^2	5	1

growth data between 1968-2012.[3] Starting in quarter 4 of 1968, each period, the agent observes y_t, and updates his beliefs about parameters and states using Bayes' law. To compute posterior beliefs about parameters, we employ a Markov Chain Monte Carlo (MCMC) technique.[4] At each date t, the MCMC algorithm produces a sample of possible parameter vectors, $\left\{\theta^d\right\}_{d=1}^D$, such that the probability of any parameter vector θ^d being in the sample is equal to the posterior probability of those parameters, $p\left(\theta^d|y^t\right)$. Therefore, we can compute an approximation to any integral by averaging over sample draws: $\int f(\theta)p(\theta|y^t)d\theta \approx 1/D\sum_d f(\theta^d)$. Every parameter draw θ^d implies a probability distribution over states $p\left(S_t|y^t,\theta^d\right)$. Thus, the forecaster can construct his forecast as

$$E\left(y_{t+1}|y^t\right) = \int\int\int y_{t+1}p\left(y_{t+1}|\theta,S_{t+1},y^t\right)p\left(S_{t+1}|\theta,y^t\right)p\left(\theta|y^t\right)d\theta dS_{t+1}dy_{t+1} \quad (8)$$

$$\approx \frac{1}{D}\sum_{d=1}^{D} E[y_{t+1}|y^t,\theta^d,\mathcal{M}]. \quad (9)$$

where $E[y_{t+1}|y^t,\theta,\mathcal{M}]$ is computed using the Kalman filtering formula (4). To estimate uncertainty, we also calculate $E\left(y_{t+1}^2|y^t\right)$ in similar fashion, apply the variance formula $Var\left(y_{t+1}|y^t\right) = E\left(y_{t+1}^2|y^t\right) - E\left(y_{t+1}|y^t\right)^2$, and take the square root $U_t = \sqrt{Var\left(y_{t+1}|y^t\right)}$.

To understand what properties of our forecasts come from estimation uncertainty, we compare our forecasts to those of an agent who has rational expectations and thus takes the model parameters as given. These are the results in the column entitled 'θ known.' To keep the results comparable, the parameters that are used in the known-parameters model are the same as the mean prior beliefs in the parameter uncertainty mode.[5] The procedure

[3] The appendix explores priors estimated on data from 1947-68.

[4] More details are presented in the Appendix. Also, see Johannes, Lochstoer, and Mou (2011) for a recursive implementation of a similar discrete-state problem of sampling from the sequence of distributions.

[5] The appendix reports results where parameters come from matching moments of the data from 1947-

Moments	Data (SPF)	θ known	Model
Mean forecast	2.24%	2.68%	3.06%
Mean $\|FErr\|$	2.20%	2.38%	2.31%
Mean U_t	–	2.91%	3.40%
Stdev U_t	–	0	**0.20%**
Correl(U_t,GDP)	–	0	13%

Table 2: **Results from linear model** Column labeled 'θ known' uses equation (4) to forecast y_{t+1}. Columns labeled 'Model' uses equation (9) to forecast y_{t+1}.

for computing forecasts and uncertainty is as follows: We endow the agent with knowledge of the parameters and with the initial belief that $y_t = \alpha$, its long-run mean. Each period, the agent updates the current state belief using (4), (5) and (6).

The main takeaway from table 2 is that estimation uncertainty generates uncertainty shocks, albeit small ones. With known parameters, stdev(U_t) = 0. When parameters are updated every period, stdev(U_t) = 0.20. That illustrates how estimation can be a source of uncertainty, although the shocks generated by estimation uncertainty are quite small.

But the results also expose three dimensions along which this forecasting does not look very realistic. 1) Our forecasters' uncertainty is not counter-cyclical (Correl(U_t,GDP) = 13%). Every common proxy for uncertainty is counter-cyclical and most theories use uncertainty to explain the onset of a recession. So, a forecasting model that fails to deliver this feature is not very realistic. 2) The model does not explain the low average forecasts of GDP observed in the professional forecaster data. The true average of GDP growth over 1968:Q4-2012:Q4 is 2.68%. The average professional forecast of GDP growth is 2.24%, almost half a percentage point lower. This model fails to explain that gap.[6] 3) There is obviously no dispersion in forecasts. There is only one forecaster. Yet, forecast dispersion is a prominent and interesting feature of the data. Later, we will introduce heterogeneous signals to examine its role.

1968.

[6]This gap only arises in final GDP estimates. The average initial GDP announcement has 2.3% growth on average, in line with the forecasts. But if these initial announcements are themselves BEA forecasts of what the final GDP estimate will be, there is still a puzzle about why early estimates are systematically lower than final estimates.

In the next section, we examine a nonlinear forecasting model that remedies the first two problems: It generates counter-cyclical uncertainty and forecasts that are low on average. At the same time, the nonlinearity amplifies the uncertainty shocks and makes clear that ignoring estimation uncertainty may result in a significant under-estimation of uncertainty shocks.

2 A Non-Linear Forecasting Model

The linear model of the previous section was a starting point. It is a commonly-used framework where we shut down all sources of uncertainty shocks besides our estimation uncertainty and could see that mechanism at work. But that model misses important features of the forecast data and its uncertainty shocks are modest. Therefore, we turn to a non-linear model that both does a better job of matching features of forecast data and generates much larger uncertainty shocks.

Updating in non-linear models is typically challenging. We make this problem tractable by doing a change of measure. We transform GDP growth into variable $\widetilde{X}$, which is a continuous variable with a hidden persistent state.

$$\begin{aligned} \widetilde{X}_t &= \alpha + S_t + \sigma\epsilon_t \\ S_t &= \rho S_{t-1} + \sigma_S \varepsilon_t \end{aligned}$$

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

The fact that GDP growth is negatively skewed tells us that the mapping from $\widetilde{X}$ to y should be concave. A concave transformation of a normal variable puts more weight on very low realizations and makes very high realizations extremely unlikely. Since the $\widetilde{X}$ variable is normal, it can take on positive or negative values. So we need a concave mapping that is defined over positive and negative real numbers. One such function is a negative exponential:

$$y_t = c - exp(-\widetilde{X}_t) \tag{10}$$

One economic interpretation of this change of measure is to think of $\widetilde{X}_t$ as an economic fundamental condition. When the economy is functioning very well (high $\widetilde{X}_t$), then improving its efficiency results in a small increase in GDP. But if there is a high degree of dysfunction or inefficiency (low $\widetilde{X}$), then the economy can easily fall into a deep depression. Most modern macroeconomic models are not linear. Many generate exactly this type of effect through borrowing or collateral constraints, other financial accelerator mechanisms, matching frictions, or information frictions. Even a simple diminishing returns story could explain such a concave mapping.

What keeps the computation tractable is that we assume that the forecaster knows the non-linear mapping. He is uncertain about the $\widetilde{X}$ process and what parameters govern it. But he knows the relationship between the variable $\widetilde{X}$ and GDP y. That assumption allows us to take the GDP data, apply the inverse transform in (10) and convert our data into $\widetilde{X}_t$ data. The parameter c in equation (10) affects the skewness of the resulting distribution of GDP growth. Therefore, we choose $c = 24$ so that the skewness of y matches the skewness of GDP growth in the data from 1952-1968:Q3.[7] Then, we can use linear updating MCMC techniques to form beliefs about the $\widetilde{X}_t$ process. For each parameter draw θ_i from the MCMC algorithm, we compute $E[y_t|\mathcal{I}_t, \theta_i]$ and $E[y_t^2|\mathcal{I}_t, \theta_i]$. We average these expectations over all parameter draws and compute uncertainty as $U_t = E[y_t^2|\mathcal{I}_t] - E[y_t|\mathcal{I}_t]^2$.

2.1 Results: Uncertainty Shocks

Figure 1 compares the time series of uncertainty to the time series of volatility in each of our models.[8] In the model with known parameters, volatility and uncertainty are identical because the agent's information set includes the model and its parameters. In the other three models, U_t and V_t differ. In the linear model with parameter uncertainty, the two series are still quite similar. At the start of the sample, the presence of unknown parameters raises uncertainty above what the model with state uncertainty alone predicts. It does so directly, but also indirectly, as the unknown parameters make the state harder to infer. But by the end of the sample, beliefs about parameters have largely converged and the uncertainty levels are more similar to the other models. The nonlinear model generates the largest level of uncertainty. The magnitude of the shocks is large, but smaller relative

[7]We choose this time period because the agents would know this data before the first forecast is made in 1968. If we instead use the 1968-2012 sample, the results are nearly identical.

[8]In each period t, volatility is computed as standard deviation of one-period ahead expectation of GDP growth conditional on a parameter vector (and a model) associated with a maximum likelihood of the data through to date t from the posterior parameter distribution under each model.

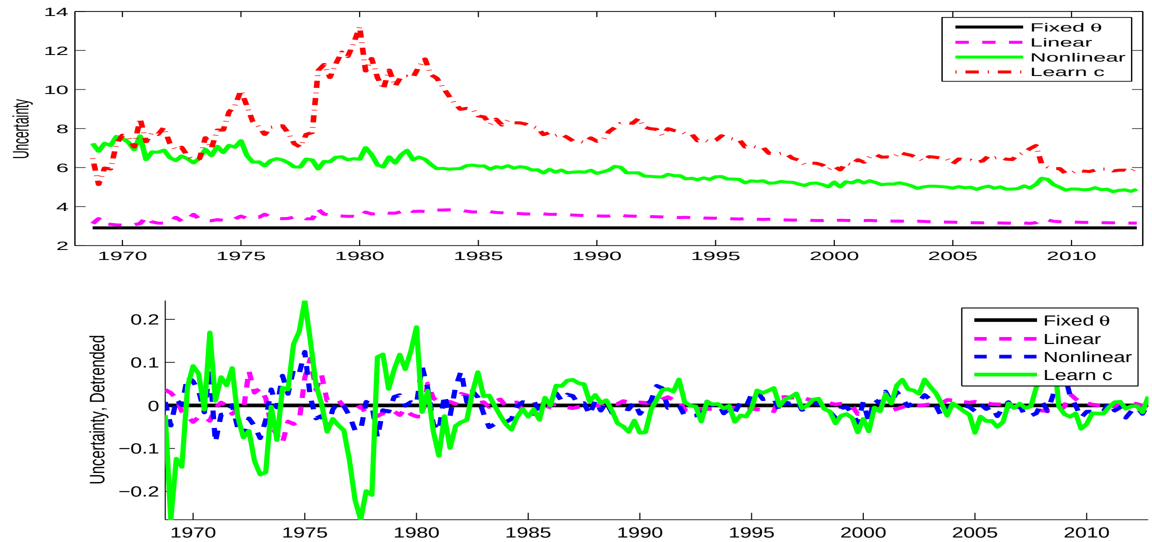

Figure 1: Uncertainty (U_t) implied by each of the models. The top panel plots the raw series, while the second shows the same uncertainty series in percentage

to the trend level than the other models. An agent who believes that GDP growth is not normally distributed has a lot more scope for uncertainty shocks.

Uncertainty is a very persistent process, with both low-frequency changes and fluctuations at the business cycle frequency. Since using growth rates of GDP is a form of trend-removal, it only makes sense to correlate a stationary series with another stationary series. Therefore, we detrend volatility and uncertainty in order to discern that nature of their cyclical components. We remove the low-frequency changes in uncertainty, using a bandpass filter to filter out frequencies lower than once every 32 quarters. Then, we compute a log deviation from this long-run trend.

$$\tilde{U}_t \equiv ln(U_t) - ln(U_t^{trend}) \tag{11}$$

The resulting series, plotted in the bottom two panels of figure 1, reveals large, highly counter-cyclical uncertainty shocks. In each of the recessions since 1968, uncertainty has risen sharply, 10-20% above trend. The summary statistics in table 3 summarize these findings.

model		linear (1)	nonlinear (2)	learn c (3)	signals (4)
Mean	U_t	3.38%	5.79%	7.65%	2.11%
	V_t	2.91%	6.82%	6.82%	2.73%
Std deviation	U_t	0.21%	0.71%	1.60%	0.05%
	V_t	0%	0.37%	0.37%	0.21%
Autocorrelation	U_t	0.95	0.95	0.95	0.95
	V_t	0	0.47	0.47	0
detrended data moments					
Corr($\tilde{U}_t$, y_t)		0.28	-0.90	-0.34	0.27
Corr($\tilde{V}_t$, y_t)		0	-0.92	-0.92	0
Corr($\tilde{U}_t$, y_{t+1})		0.24	-0.22	-0.10	0.24
Corr($\tilde{V}_t$, y_{t+1})		0	-0.23	-0.23	0

Table 3: **Properties of model uncertainty series.** Columns 1, 2 and 3 use equations (9), (10) and (18) to forecast y_{t+1}. Volatilities are computing assuming that the true parameters θ are the maximum likelihood estimates, using all gdp data available at time t.

2.2 Why Does the Nonlinear Forecast Model Perform So Differently?

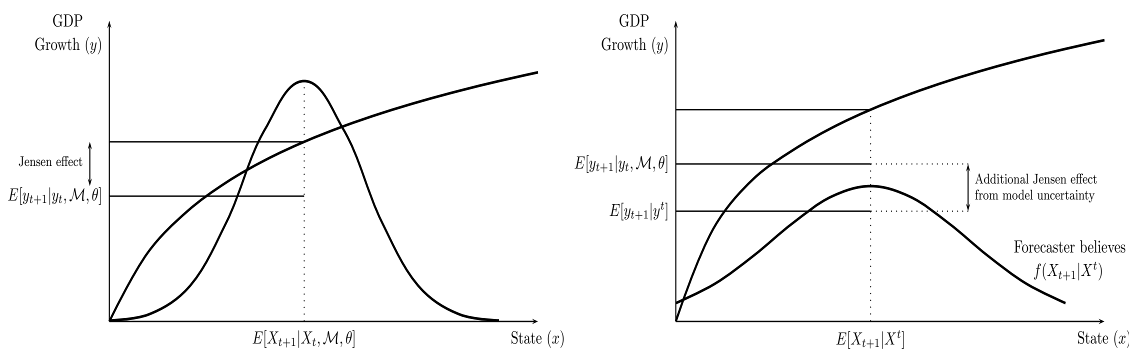

Figure 2: Explaining why average forecasts are lower than mean GDP growth. The result has two key ingredients: The forecaster faces more uncertainty than he would if he knew the true distribution of outcomes, and a Jensen inequality effect from the concave change of measure.

Aside from generating larger uncertainty shocks, the nonlinear model also explain the low GDP growth forecasts in the professional forecaster data. The average forecast is 2.2% in the model and 2.2% in the forecaster (SPF) data. Thus, the first puzzle is why the nonlinear model produces these average forecasts. If GDP growth is a concave transformation of a linear-normal underlying variable, Jensen's inequality tells us that expected values will be systematically lower than the median realization. But by itself, Jensen's inequality does not explain the forecast bias because the expected GDP growth and the mean GDP growth should both be lowered by the concave transformation (see figure 2, left panel). It must be that there is some additional uncertainty in expectations, making the Jensen effect larger for forecasts than it is for the unconditional mean of the true distribution (see figure 2, right panel). This would explain why our results tell us that most of the time, $E[y_{t+1}|\theta] > E[y_{t+1}|y^t]$. If the agent knew the true parameters, he would have less uncertainty about y_{t+1}. Less uncertainty would make the Jensen effect smaller and raise his estimate of y_{t+1}, on average. Thus, it is the combination of parameter uncertainty and the non-linear updating model that can explain the forecast bias.[9]

The same concave change of measure that can explain why low GDP forecasts can also generate large uncertainty shocks. The concave line in Figure 3 is a mapping from x into GDP growth, y. The slope of this curve is a Radon-Nikodym derivative. A given amount of uncertainty is like a band of possible x's. (If x was uniform, the band would represent the positive-probability set and the width of the band would measure uncertainty about x.) If that band is projected on to the y-space, the implied amount of uncertainty about y depends on the state x. When x is high, the mapping is flat, and the resulting with of the band projected on the y-axis (y uncertainty) is small. When x is low, the band projected on the y axis is larger and uncertainty is high.

Uncertainty is much more persistent. It is highest at the start of the sample, when data is most scarce, and then slowly decays over the rest of the sample. As noted by Collin-Dufresne, Johannes, and Lochstoer (2013), the persistent uncertainty process comes from the nature of learning: A single large shock to GDP growth results in a quick reevaluation of the parameter and model probabilities. These revisions in beliefs act as permanent, non-stationary shocks even when the underlying shock is transitory.

[9] When we use first-release GDP data, this forecast bias disappears. But that is consistent with an interpretation of this first release as itself a forecast. The first release may be lower, on average because it is more uncertain and the added uncertainty, combined with the non-linear model, lowers the average initial announcement.

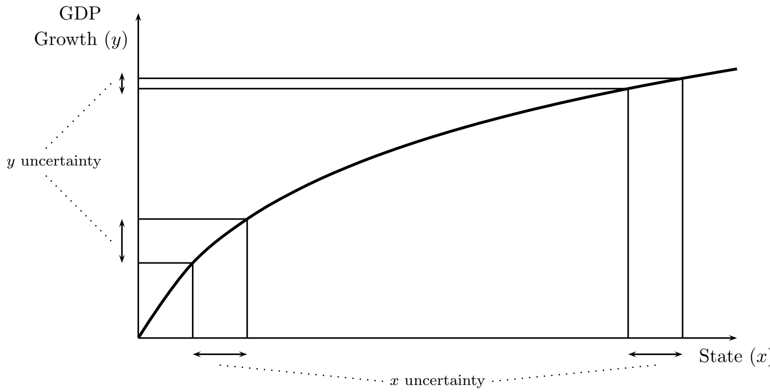

Figure 3: Nonlinear change of measure and state-dependent uncertainty. A given amount of uncertainty about x creates more uncertainty about y when x is low than it does when x is high.

2.3 The Role of the Black Swan

The previous discussion hopefully illuminated the mechanics of the nonlinear model. But it still leaves open the question: What generates an uncertainty shock? We explore the hypothesis that events which the forecasting model deems unlikely (black swan events) trigger uncertainty shocks. To see the relationship between uncertainty and unlikely events, we compute a black swan core for every GDP growth observation. The score is simply the number of standard deviations the realized GDP growth is from its forecasted level last period:

$$Black\,Swan\,Score_t = \frac{|y_t - E[y_t|y^{t-1}]|}{U_{t-1}}. \tag{12}$$

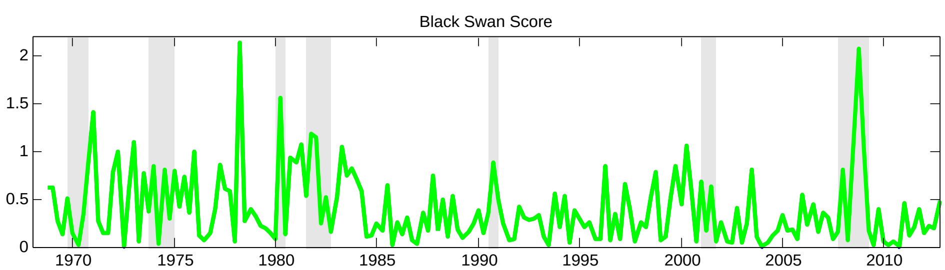

Figure 4: The largest uncertainty shocks correspond to black swan events. *Black Swan Score* is defined in (12).

Figure 4 uses the non-linear model 4 to compute the expectation and U_t and report this score. Comparing the black swan score to the uncertainty measure in the plot above it reveals two regularities. First, most but not all spikes in uncertainty coincide with an unusual event. Second, some black swan events do not generate uncertainty shocks. In particular, quarter 2 of 1978 registered 15.4% annualized growth. Such a high growth rate was unanticipated by the model forecasts (and by most forecasters at the time). Thus, it registers as a high black swan score. However, while most unusual events are low-growth events, this is a high-growth outlier. Such a positive black swan events seems to generate much less uncertainty than the negative black swan. The asymmetric change of measure causes positive and negative black swan events to have very different consequences.

2.4 Learning About Model Nonlinearity

There is no good reason why agents should be certain about the distribution of outcomes, but not the parameters, as the previous section assumed. One tractable way to allow agents to learn about the distribution is to allow them to learn about the parameters governing the change of measure. Since our change of measure function (10) has only one parameter, this suggests that we should allow the forecaster to learn about the parameter c.

The difficulty is that each value of c implies a different data sample $\{X^t\}$ on which to estimate the linear model with parameter uncertainty. Since computing the probability distribution of parameters with one set of data is already computationally intensive, considering many candidate $\{X^t\}$ samples, in each period, is too time-intensive to be feasible. Instead, we allow the agent to choose a value of c that matches the skewness of GDP growth, each period. As new GDP data points are observed, sample skewness changes and the agent adjusts c. We show that even without considering any uncertainty about c, the simple act of updating c generates large uncertainty fluctuations.

Column (3) of Table 3 shows that updating beliefs about the distribution and its non-linearity or non-normality have a large effect on uncertainty. Such learning increases the average level of uncertainty by 32%. More strikingly, it more than doubles the size of uncertainty shocks. The standard devitation of the uncertainty series was 0.71% without c updating and is 1.60% with updating beliefs about the distribution. One can interpret the magnitude of this standard deviation relative to the mean. A 1-standard deviation shock to uncertainty raises uncertainty 21% above its mean. That is quite a volatile process and offers quite a contrast to the relatively modest changes in volatility typically measured.

A third feature of uncertainty with distribution learning is that the uncertainty is significantly less counter-cyclical. The reason for this is that when GDP is high, it increases the estimated mean of the GDP process. This makes the negative outliers in the sample lie further away from the mean and increases the sample skewness. Since the beliefs about c are tied to model skewness, this causes a deline in c, which increases the slope of the change of measure function and as a result, increases uncertainty.

3 Data Used to Proxy for Uncertainty

Our model generates an endogenous uncertainty series. Next, we'd like to compare our measure to some of the empirical proxies for uncertainty that are commonly used, including the VIX, forecast dispersion, squared forecast errors, and GARCH volatility estimates. Although these measures are all supposed to serve as proxies for uncertainty, they have different properties from each other and from the model-generated uncertainty series.

3.1 Time-varying volatility models

One common procedure for estimating the size of volatility shocks is to estimate an ARMA process that allows for stochastic volatility. In order to compare such a volatility measure to our uncertainty series, we estimate a GARCH model of GDP growth, that allows for time-variation in the variance of the innovations. All data is quarterly and all of these series are non-stationary. To obtain stationary series, we use annualized growth rates.

The GARCH process that generates the best fit is one with an AR(1) process for GDP growth (y) and a GARCH(1) process for volatility, which includes 1 lagged variance:

$$y_{t+1} = 3.38 + 0.41 y_t + \epsilon_{t+1}, \quad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma_{t+1}^2) \tag{13}$$

$$\sigma_{t+1}^2 = 0.52 + 0.76\sigma_t^2 + 0.24\epsilon_t^2 \tag{14}$$

We assume that the errors in our models, ϵ_t, are Gaussian and estimate the process using our full sample of data (1947-2012). The complete set of estimates, stationarity tests, as well as more detail about the model selection process are reported in appendix B. We also estimate homoskedastic models for each of the time series to test the hypothesis of homoskedasticity. We use the ARCH-LM test to do this. This estimated GARCH process produces an estimate of the shock variance σ_{t+1}^2 in each period. This is what we call the volatility process.

The evidence for time-varying volatility is weak. The log likelihood of the highest-likelihood heteroskedastic model is only 3% higher than the best-fitting homoskedastic model. With an ARCH-LM test, one cannot reject the null hypothesis of homoskedasticity (pvalue is 0.22). We come to a similar conclusion if we estimate the model, starting in 1947, use a stochastic volatility model, or relax the distributional assumptions. Our baseline analysis assumes that errors ϵ_t have a Gaussian distribution. Using distributions with fatter tails (student-t) yields no difference in estimations or significance. Furthermore, we explore further lags of all variables; either coefficients were not significantly different from zero or the log-likelihood was reduced. Finally, we included different lags of linear terms for ϵ_t and variances σ_t^2 in the GARCH specification. Again, the estimated parameters were not significant or the log-likelihood was reduced.

3.2 Forecast dispersion

Some authors[10] use forecast dispersion (D_t in equation 17) as a measure of uncertainty. One of the advantages of this measure is that it is typically regarded as "model-free." It turns out that dispersion is only equivalent to uncertainty in a particular class of models. In that sense, it is not really model-free.

Whether dispersion accurately reflects uncertainty depends on private or public nature of information. Imprecise private information generates forecast dispersion, while imprecise public information typically does not. To understand the importance of this distinction, we decompose these forecast errors into their public and private components. Any unbiased forecast can be written as the difference between the true variable being forecast and some forecast noise that is orthogonal to the forecast. In other words,

$$y_{t+1} = E[y_{t+1}|\mathcal{I}_t] + \eta_t + e_{it} \tag{15}$$

where the forecast error $(\eta_t + e_{it})$ is mean-zero and orthogonal to the forecast. We can

[10]See e.g. Baker, Bloom, and Davis (2012) or Diether, Malloy, and Scherbina (2002), or Johnson (2004).

further decompose any forecast error into a component that is common to all forecasters η_t and a component that is the idiosyncratic error e_{it} of forecaster i. Using definition (2), we can write uncertainty as

$$U_{it} = E[(\eta_t + e_{it})^2|\mathcal{I}_t]$$

Since η and e are independent, mean-zero variables, this is

$$U_{it} = var(\eta_t) + var(e_{it})$$

Dispersion depends on the squared difference of each forecast from the average forecast. We can write each forecast as $y_{t+1} - \eta_t - e_{it}$. Then, with a large number of forecasters, we can apply the law of large numbers, set the average e_{it} to 0 and write the average forecast as $\bar{E}[y_{t+1}|\mathcal{I}_t] = y_{t+1} - \eta_t$. Thus,

$$D_t = \frac{1}{N}\sum_i (E[y_{t+1}|\mathcal{I}_t] - \bar{E}[y_{t+1}|\mathcal{I}_t])^2 = \frac{1}{N}\sum_i e_{it}^2$$

If there is a large number of forecasters, each with the same variance of their private forecast noise $var(e_t)$, and we apply the law of large numbers, then

$$D_t = var(e_t).$$

Comparing the formulae for U_{it} and D_t reveals that they are the same under two conditions: 1) there is no common component in forecast errors ($var(\eta_t) = 0$), and 2) private forecast errors have the same variance ($var(e_{it} = var(e_t), \forall i$). Thus, equating forecast dispersion and uncertainty is implicitly using a forecasting model that has these two properties. Property 2 may be violated and then dispersion may still be a good measure of average uncertainty. If property 1 is violated, but the variance of public signal noise does not change over time ($var(\eta_t) = \bar{v}, \forall t$), then dispersion will covary perfectly with uncertainty. While one can make such an assumption, it does raise the question why private information about the macroeconomy varies in its precision when public sources of information have a precision that is constant. As we will see in the next section, the forecast data does not support such an interpretation.

A Model with Heterogeneous Signals One feature of the forecaster data that the model so far does not speak to is the fact that there is heterogeneity in forecasts. A second feature is that forecasters obviously have access to more data than just past realizations of GDP. Additional data includes leading indicators or firm announcements of future hiring and investment plans.

To model this additional information and the heterogeneity of forecasts, we consider a setting where forecasters update beliefs as in the previous section. But each period, each forecaster i observes an additional signal z_{it} that is the next period's GDP growth, with common signal noise and idiosyncratic signal noise:

$$z_{it} = y_{t+1} + \eta_t + \epsilon_{it} \tag{16}$$

where $\eta_t \sim N(0, \sigma_\eta^2)$ is common to all forecasters and $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ is i.i.d. across forecasters.

We calibrate the two signal noise variances σ_η^2 and σ_ϵ^2 to match two moments. The first is the average dispersion of forecasts. Time-t forecast dispersion is

$$D_t = \sqrt{\frac{1}{N_t} \sum_{i \in I_t} (E[y_{t+1}|\mathcal{I}_{it}] - \bar{E}_t)^2} \tag{17}$$

where $\bar{E}_t = 1/N_t \sum_i E_{it}[y_{t+1}]$ is the average time-t growth forecast and N_t is the number of forecasters making forecasts at time t. The timer-series average of D_t is 1.6%.

The second moment is the average forecast error, $1/T \sum_t FE_t$. In the data, this average forecast error is 1.91%. Note that this average error is larger than the dispersion. It tells us that signal noise is not exclusively private signal noise. We choose σ_η and σ_ϵ such that average dispersion and forecast errors are identical in the model and in the data. The resulting estimated signal variances are $\sigma_\eta = 1.42\%$ and $\sigma_\epsilon = 2.31\%$.

Our forecasters use prior GDP growth realizations to estimate the parameters of the model, just as in (9). Then, they use their additional, heterogeneous signal to update these prior beliefs using Bayes' law:

$$E\left(y_{t+1}|y^t, z_{it}\right) = \int y_{t+1} \frac{f(z_{it}|y^{t+1}) f(y_{t+1}|y^t)}{\int f(z_{it}|y', y^t) f(y'|y^t) dy'} dy_{t+1}. \tag{18}$$

Comparing signal model to forecast data While the model with heterogeneous signal allows us to talk about the part of forecast errors that come from forecast dispersion, the amount of dispersion does not drastically alter our predictions. In a model with much less dispersion (0.4%), we find almost identical levels of uncertainty and correlations of uncertainty with GDP, but larger shocks to uncertainty.

But the fact that forecasters have signals about GDP growth makes a big difference. This model does a better job than any of the others in matching the average forecast error. But recall that we calibrated signal noise in order to match this moment of the data. Reducing the average error also brings down the standard deviation of forecast errors, to a level just slightly above that in the data.

The biggest success of the signal model is that it allows forecasts to be more highly correlated with future GDP growth. The correlation of 0.74 in this model is almost equal to the correlation of 0.72 in the data. While the other forecasting models do not have enormous signal errors because they are close to the true GDP growth number on average, they miss lots of the little ups and downs in GDP and therefore achieve low correlation. With a signal about future GDP, forecasters in this model are more likely to revise their forecasts slightly up when growth will be high and down when it will be low, achieving a higher correlation between forecast and GDP growth. But the bottom line is that building forecast heterogeneity in the model, similar to what we see in the data, is not a mechanism for generating large uncertainty shocks.

Should signal precision vary over time? One potential source of uncertainty shocks could be changes in the precision of signal. Here, we argue that this is an unlikely source of uncertainty shocks because it suggests other features of the data that are counter-factual. Suppose that in periods where the variance of the noise σ_η or σ_ϵ is high, z_t is a relatively poor predictor of y_{t+1}. Since agents' signal about y_{t+1} are low-precision, their uncertainty about y_{t+1}, conditional on this signal, will be high.

Of course, if all else is equal, a more volatile e_t would mean a more volatile y_t, and this story of forecasting variable precision changes would be the same as a story where uncertainty shocks come from volatility shocks. But it is possible that a fall in signal precision (increase in $\text{var}(e_t)$) could generate an uncertainty shock without a volatility shock to the GDP process y_t. For example, if z and e are independent, then $var(y_t) = var(z_t) + var(e_t)$, and a negative relationship between $var(z_t)$ and $var(e_t)$ could leave $var(y_t)$ unchanged. But this structure would imply that more uncertainty from less informative signals (high $var(e_t)$) is associated with lower macro volatility (low $var(z_t)$). There is no such inverse relationship between uncertainty and the volatility of forecasting variables in the data.

There is another way that $var(y_t)$ could be constant, despite a shock to signal precision. If e_t and z_t themselves are negatively correlated, then $var(z_t)$ and $var(e_t)$ can both rise. Since $var(y_t) = var(z_t) + var(e_t) - 2cov(y_t, e_t)$, then a sufficiently high covariance will allow signal precision $1/var(e_t)$ to change, resulting in an uncertainty shock, without a volatility shock. But the problem with this explanation is that then z_t is no longer an unbiased forecast of y_t. If we transform z_t to make it an unbiased signal, then this negative correlation with the estimation error would disappear. Exploring these possibilities makes the point that it is hard to see how changes in the precision of forecast variables can explain uncertainty shocks, in a way that is consistent with the data.

3.3 Mean-squared forecast errors

A measure related to forecast dispersion that captures both private and common forecast errors is the forecast mean-squared error.

We define a forecast mean-squared error (MSE_{t+1}) of a forecast of y_{t+1} made in quarter t as the square root of the average squared distance between the forecast and the realized value

$$MSE_{t+1} = \sqrt{\frac{\sum_{i \in I_t} (E[y_{t+1}|\mathcal{I}_{it}] - y_{t+1})^2}{N_t}}. \tag{19}$$

If forecast errors were completely idiosyncratic, with no common component, then dispersion in forecasts and mean-squared forecasting errors would be equal. To see this, note that $FE_{jt}^2 = (E[y_{t+1}|\mathcal{I}_{jt}] - y_{t+1})^2$. We can split up FE_{jt}^2 into the sum $((E[y_{t+1}|\mathcal{I}_{jt}] - \bar{E}_t[y_{t+1}]) + (\bar{E}[y_{t+1}] - y_{t+1}))^2$, where $\bar{E}_t[y_{t+1}] = \int_j E[y_{t+1}|\mathcal{I}_{jt}]$ is the average forecast. If the first term in parentheses is orthogonal to the second, $1/N \sum_j FE_{jt}^2 = MSE_t^2$ is simply the sum of forecast dispersion and the squared error in the average forecast: $E[y_{t+1}|\mathcal{I}_{jt}] - \bar{E}_t[y_{t+1}])^2 + (\bar{E}_t[y_{t+1}] - y_{t+1})^2$.

We can then use this insight along with our forecast data to evaluate the extent to which variation in mean-squared errors (MSE) comes from changes in the accuracy of average forecasts and how much comes from changes in dispersion. If most of the variation comes from changes in disperion, then perhaps the assumption that the common components of forecast errors are zero or negligible is not a bad assumption. But if most of the fluctuation in MSE comes from changes in average forecast errors, then using forecast dispersion as a proxy for uncertainty will miss an important sourcee of variation.

Table 4 examine forecasts of four different macro variables covered by the survey of professional forecasters. For each one, at least half of the variation comes from average forecasts. This finding tells us that the model which equates forecast dispersion with uncertainty, which would imply that dispersion and mean-squared error would be identical, is a model which is not well supported by the data. In other words, forecast dispersion may have very different properties from uncertainty.

	GDP	FedSpend	SLSpend	IntRt
R^2	80%	54%	97%	65%

Table 4: **The fraction of MSE variation explained by average forecast errors.** This is the R^2 of a regression of the forecast mean-squared error, MSE^2, defined in (19) on $(\bar{E}_t[y_{t+1}] - y_{t+1})^2$. The remaining variation is due to changes in forecast dispersion.

Forecast bias Recent research argues that professional forecasts are biased. For example, Elliott and Timmermann (2008) argue that stock analysts over-estimate earnings growth and the Federal Reserve under-estimates GDP growth. But none of these findings suggest that the bias is volatile. In other words, forecast bias may explain why forecasts are further away from the true outcome. But uncertainty shocks come from fluctuations in the expected size of forecast errors. A fixed bias does not create such fluctuations.

3.4 VIX and confidence measures

We discuss the volatility and forecast-based measures in most detail because there are measures that we can relate explicitly to our theory. Other proxy variables for uncertainty are interesting but have a less clear connection to our model. The market volatility index (VIX) is a traded blend of options that measures expected percentage changes of the S&P500 in the next 30 days. It captures expected volatility of equity prices. But it would take a rich and complicated model to link macroeconomic uncertainty to precise movements in the VIX. Nevertheless, we can compare its statistical properties to those of the uncertainty measure in our model. Figure 5 does just this.

Another commonly cited measure of uncertainty is business or consumer confidence. The consumer confidence survey asks respondents whether their outlook on future business or employment conditions is "positive, negative or neutral." Likewise, the index of consumer sentiment asks respondents whether future business conditions and personal finances will be "better, worse or about the same." While these indices are indeed negatively correlated with the GARCH-implied volatility of GDP, they are not explicitly questions about uncertainty. Furthermore, we would like to use a measure that we can compare to the forecasts in our model. Since it is not clear what macro variable "business conditions" or "personal finances" corresponds to, it is not obvious what macro variable respondents are predicting.

3.5 Comparing uncertainty proxies to model-generated uncertainty

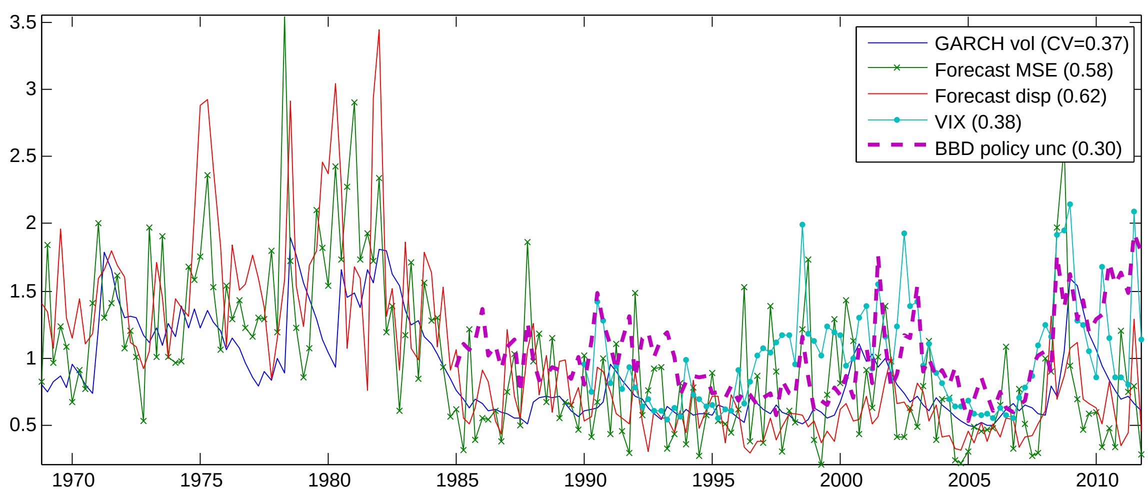

Figure 5: Comparing variables used to measure uncertainty in the literature.

Figure 5 plots each of the uncertainty proxies. There is considerable comovement, but also substantial variation in the dynamics of each process. These are clearly not measures of the same stochastic process, each with independent observation noise. Furthermore, they have properties that are quite different from our model-implied uncertainty metric.

	Mean	Standard deviation	autocorr	correlation with g_{t+1}	correlation with $\tilde{U}_t$
forecast MSE	2.64%	1.53	0.48	0.04	-3.02%
forecast dispersion	1.54%	0.95	0.74	-0.19	19.98%
GARCH volatility	3.65%	1.35	0.90	0.06	7.26%
GARCH real-time	3.45%	1.04	0.90	-0.07	8.32%
VIX	20.55	7.81	0.58	-0.41	35.80%
BBD policy uncertainty	105.95	31.79	0.65	-0.41	20.95%
model 4 uncertainty ($\tilde{U}_t$)	0	0.03	0.41	-0.23	100%

Table 5: **Properties of forecast errors and volatility series for macro variables.** Forecast MSE and dispersion are defined in (19) and (17) and use data from 1968q4-2011q4. Growth forecast is constructed as $ln(E_t(GDP_t)) - ln(E_t(GDP_{t-1}))$. GARCH volatility is the σ_{t+1} that comes from estimating (14), using data from 1947q2-2012q2. At each date t, GARCH real-time is the σ_{t+1} that comes from estimating (14) using only data from 1947:Q2 through date t. VIX_t is the Chicago Board Options Exchange Volatility Index closing price on the last day of quarter t, from 1990q1-2011q4. BBD policy uncertainty is the Baker, Bloom, and Davis (2012) economic policy uncertainty index for the last month of quarter t, from 1985q1-2011q4. $\tilde{U}_t$ is nonlinear model uncertainty, measured as the log deviation from trend (eq. 11).

Table 5 shows that our uncertainty metric is less volatile, moderately counter-cyclical, but the raw uncertainty series is more persistent than the proxy variables.

3.6 Inferring uncertainty from probability forecasts

One way to infer the uncertainty of an economic forecaster is to ask them about the probabilities of various events. The survey of professional forecasters does just that. They ask about the probability that GDP growth exceeds 6%, is between 5-5.9%, between 4-4.9%, . . . , between -1 and -2%, and below -2%. However, the survey only reports a single probability weight that is averaged across all forecasters.

Since this data does not completely describe the distribution of y_{t+1} beliefs, computing a variance requires some approximation. The most obvious approximation is to assume a discrete distribution. For example, when agents assign a probability to $1-2\%$ GDP growth, we treat this as if that is the probability placed on the outcome of 1.5% GDP growth. When the agent says that there is probability $p_{6.5}$ of growth above 6%, we treat this as probability $p_{6.5}$ placed on the outcome $y_{t+1} = 6.5\%$. And if the agent reports probability $p_{-2.5}$ of growth below -2%, we place probability of $p_{-2.5}$ on $y_{t+1} = -2.5\%$. Then the expected rate of GDP growth is $\bar{y} = \sum_{m \epsilon M} p_m m$ for $M = \{-2.5, -1.5, \ldots, 6.5\}$. Finally, the conditional variance of beliefs about GDP growth are $var[y|\mathcal{I}] = \sum_{m \epsilon M} p_m (m - \bar{y})^2$.

The resulting conditional variance series is not very informative. It hardly varies (range is $[0.0072, 0.0099]$). It does not rise substantially during the financial crisis. In fact, it suggests that uncertainty in 2008 was roughly the same as it was in 2003. The reason this measure does not detect high uncertainty in exteme events is that the growth rates are top- and bottom-coded. All extremely bad GDP events are grouped in the bin "growth less than 2%." All the uncertainty was about how bad this recession might be. Instead, what the probability bins reveal is a high probability weight on growth below 2%. Since most of the probability is concentrated in one bin, it makes the uncertainty look low. The bottom line is that while using surveys to ask about ex-ante probabilities of GDP events is a promising approach to measuring uncertainty, the available data that uses this approach does not seem useful for our purposes.

4 Conclusions

The data typically used to measure uncertainty offer a muddy picture of what uncertainty shocks look like: Survey data reveals large swings in the ability of agents to make accurate forecasts. Yet, there is weak evidence of time-varying variance of macro aggregates and the estimated magnitude of these shocks is much smaller than what would explain the survey data. Uncertainty shocks appear to come not from the properties of the data being forecast, but from some state of mind of the forecaster himself.

Our model reconciles these findings. It takes the actual series of GDP as an input and generates forecasts with properties similar to those in the forecast data. But with this model structure, we can compute actual uncertainty – the conditional variance of the forecast of next quarter growth. We find that model uncertainty is a source of large, counter-cyclical uncertainty shocks that covary imperfectly with standard uncertainty proxies.

The model also helps us to understand where uncertainty shocks come from. People use simple models to forecast complex economic processes. They have to. Any model that is as complicated as the economy itself would be intractable and not useful. Recognizing this, it is natural for agents to consider more than one possible model of the world. Of course, forecasters are not endowed with knowledge of the parameters of their model either.

They estimate them over time. When we estimate a forecasting model with these three ingredients: state uncertainty, model uncertainty and parameter uncertainty, the resulting uncertainty shocks resemble some of the uncertainty proxies, but are not identical to them.

A next step for this agenda is to explore firm-level earnings forecasts and firm-level shocks and determine whether our model can also explain uncertainty facts at the micro level.

References

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

——— (2012b): "Wait-and-See Business Cycles?," University of Aachen working paper.

Bachmann, R., S. Elstner, and E. Sims (2012): "Uncertainty and Economic Activity: Evidence from Business Survey Data," University of Aachen working paper.

Baker, S., N. Bloom, and S. Davis (2012): "Measuring Economic Policy Uncertainty," Stanford University working paper.

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

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

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

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

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

Born, B., A. Peter, and J. Pfeifer (2011): "Fiscal News and Macroeconomic Volatility," Bonn econ discussion papers, University of Bonn, Germany.

Bruno, V., and H. Shin (2012): "Capital Flows, Cross-Border Banking and Global Liquidity," Princeton University working paper.

Chen, H., W. Dou, and L. Kogan (2013): "Measuring the Dark Matter in Asset Pricing Models," MIT working paper.

Christiano, L., R. Motto, and M. Rostagno (2012): "Risk Shocks," Northwestern University working paper.

Cogley, T., and T. Sargent (2005): "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–563.

Collin-Dufresne, P., M. Johannes, and L. Lochstoer (2013): "Parameter Learning in General Equilibrium: The Asset Pricing Implications," Columbia University working paper.

DIETHER, K., C. MALLOY, AND A. SCHERBINA (2002): "Differences of Opinion and the Cross-Section of Stock Returns," *Journal of Finance*, 57, 2113–2141.

ELLIOTT, G., AND A. TIMMERMANN (2008): "Economic Forecasting," *Journal of Economic Literature*, 46(1), 3–56.

FERNANDEZ-VILLAVERDE, J., P. GUERRON-QUINTANA, J. F. RUBIO-RAMIREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101(6), 2530–61.

FERNANDEZ-VILLAVERDE, J., K. KUESTER, P. GUERRON, AND J. RUBIO-RAMIREZ (2012): "Fiscal Volatility Shocks and Economic Activity," University of Pennsylvania working paper.

HANSEN, L. (2007): "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," *American Economic Review*, 97(2), 1–30.

JOHANNES, M., L. LOCHSTOER, AND Y. MOU (2011): "Learning About Consumption Dynamics," Columbia University working paper.

JOHNSON, T. (2004): "Forecast Dispersion and the Cross Section of Expected Returns," *Journal of Finance*, 59, 1957–1978.

JURADO, K., S. LUDVIGSON, AND S. NG (2013): "Measuring Uncertainty," New York University working paper.

JUSTINIANO, A., AND G. PRIMICERI (2008): "Time-Varying Volatility of Macroeconomic Fluctuations," *American Economic Review*, 98(3), 604–641.

NAKAMURA, E., D. SERGEYEV, AND J. STEINSSON (2012): "Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence," Columbia University working paper.

NIMARK, K. (2012): "Man Bites Dog Business Cycles," Working paper.

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

STOCK, J., AND M. WATSON (2012): "Disentangling the Channels of the 2007-2009 Recession," NBER working paper 18094.

A Estimating the model

We estimate the parameters of the model using random-walk Metropolis Hastings algorithm. The likelihood function of a given vector of parameters, θ, (which we need for the accept/reject step of the algorithm) is given by

$$p\left(y^T|\theta\right) = \prod_{t=0}^{T-1} p\left(y_{t+1}|y^t,\theta\right)$$

In turn, the predictive distribution of the data, $p\left(y_{t+1}|y^t,\theta\right)$ can be obtained as an integral against the filtering distribution

$$p\left(y_{t+1}|y^t,\theta\right) = \int\int p\left(y_{t+1}|s_{t+1},\theta\right) p\left(s_{t+1}|s_t,\theta\right) p\left(s_t|y^t,\theta\right) ds_t ds_{t+1}$$

We apply Kalman filtering techniques to obtain the distribution of the filtered values of the hidden state, $p\left(s_t|y^t,\theta\right)$, which will be characterized by the following moments $\widehat{s}_t \equiv E\left[s_t|y^{t-1},\theta\right]$ and $Var_{s,t} \equiv E\left[\left(s_t-\widehat{s}_t\right)\left(s_t-\widehat{s}_t\right)'\right]$. The predictive distribution of the data is then given by the following moments $E\left[y_t|y^{t-1},\theta\right] = \widehat{s}_t$ and $Var\left[y_t|y^{t-1},\theta\right] = Var_{s,t}+\sigma^2$.

B Estimated ARCH/GARCH process for GDP

In this section we present the GARCH models that we estimate to infer volatility proxies for GDP. We estimate volatility using an ARMA model with ARCH or GARCH errors. The estimation procedure is maximum likelihood. We considered several models and chose the AR and MA orders based on the significance of additional variables and their effect on the log-likelihood. Similarly, we considered different lags of linear terms for ϵ_t and variances σ_t^2 in the GARCH specification and used the significance and effect of additional variables on the log-likelihood to inform the specification choice.

We use quarterly GDP growth rate data for 1947:Q2–2012:Q2.[11] The subsample of the GDP data which matches the uncertainty data is 1968:Q4 to 2011:Q4. The results of ADF tests indicate that the series for the full sample and matching subsample are stationary.

Best-fitting homoskedastic model Recall that we defined $y_{t+1} \equiv \ln(gdp_t) - \ln(gdp_{t-1})$. The best-fitting processes is an ARMA(1,0):

$$\Delta gdp_{t+1} = 3.25 + 0.37\Delta gdp_t + \epsilon_{t+1}, \quad \varepsilon_{t+1} \sim \mathcal{N}(0, 14.43)$$

[11]Data source: Bureau of Economic Analysis (http://www.bea.gov/national/index.htm#gdp). We are using the seasonally adjusted annual rate for the quarterly percentage change in real GDP. Note that this data is for the actual percentage change, not an approximation. The version of the data is 27 July 2012, downloaded 2 August 2012.

The log-likelihood is -713.27.

Best-fitting heteroskedastic model For heteroskedastic processes, the best specification is ARMA(1,0) for growth and GARCH(1) for variance:

$$\Delta gdp_{t+1} = 3.38 + 0.41\Delta gdp_t + \epsilon_{t+1}, \quad \varepsilon_{t+1} \sim \mathcal{N}(0, \sigma_{t+1}^2)$$

$$\sigma_{t+1}^2 = 0.52 + 0.76\sigma_t^2 + 0.24\epsilon_t^2$$

The log-likelihood is -693.90.

For each sample period of each variable we use Augmented Dickey-Fuller (ADF) tests to check that the series is stationary. This test is based on estimating an AR model for the data. For each sample we use the best-fitting AR model as the basis for the test. Each of the series we use passed the stationarity test.