

Understanding the Sources of Macroeconomic Uncertainty

Barbara Rossi* Tatevik Sekhposyan[†] Matthieu Soupre[‡]

August 2, 2016

Abstract

We propose a decomposition to distinguish between Knightian uncertainty (ambiguity) and risk, where the first measures the uncertainty about the probability distribution generating the data, while the second measures uncertainty about the odds of the outcomes when the probability distribution is known. We use the Survey of Professional Forecasters (SPF) density forecasts to quantify overall uncertainty as well as the evolution of the different components of uncertainty over time and investigate their importance for macroeconomic fluctuations. We also study the behavior and evolution of the various components of our decomposition in a model that features ambiguity and risk.

Keywords: Uncertainty, Risk, Ambiguity, Knightian Uncertainty, Survey of Professional Forecasters, Predictive Densities.

J.E.L. Codes: C22, C52, C53.¹

*ICREA-University of Pompeu Fabra, Barcelona GSE and CREI, c/Ramon Trias Fargas 25/27, Barcelona 08005, Spain; e-mail: barbara.rossi@upf.edu

[†]Texas A&M University, 3060 Allen Building, 4228 TAMU, College Station, TX 77843, USA; e-mail: tsekposyan@tamu.edu

[‡]University of Pompeu Fabra, c/Ramon Trias Fargas 25/27, Barcelona 08005, Spain; e-mail: matthieu.soupre@upf.edu

¹**Acknowledgements:** We are grateful to Domenico Giannone and to seminar participants at the Fourth Inter-

1 Introduction

The recent financial crisis has renewed interest in measuring uncertainty and studying its macroeconomic effects. Stock and Watson (2012) suggest that liquidity-risk and uncertainty shocks are among the most important factors explaining the decline in U.S. GDP during the Great Recession, accounting for about two thirds of the GDP decline. Given that uncertainty is inherently unobserved, this has sparked a wide research agenda on various measures of uncertainty. However, as shown in Rossi and Sekhposyan (2015), the macroeconomic impact of the various uncertainty measures can be very different from each other. This naturally leads to the question of what exactly the uncertainty indices measure and how they differ from each other.

Typically the literature distinguishes between two types of uncertainty. The first type of uncertainty is the one that rational agents face when making their decisions, as the realization of the state of nature is not known in advance even if the agents can reasonably contemplate all possible states of nature and their likelihood. This situation is commonly known as risk. That is, risk is characterized by situations where one knows the odds of the unknown, that is, one knows the probability distribution of the stochastic events. Frank Knight (1921) suggested a different definition of uncertainty, in which agents cannot reasonably contemplate all the possible states of nature or characterize their probability distributions. Furthermore, even if they are able to characterize the distributions, they might be unable to assign correct probabilities to future outcomes. For example, disagreement on the probability distribution of future outcomes is a special case of Knightian uncertainty, since disagreeing on probability distributions automatically implies that the probability distributions are not correctly specified.

The empirical literature has proposed several measures of uncertainty, but does not distinguish between risk and Knightian uncertainty, nor explains how they relate to each other. In addition, while researchers routinely report correlations among various uncertainty measures or compare their macroeconomic effects, it is unclear how exactly they are related to each other or whether the difference in their macroeconomic effects depends on the type of uncertainty they measure.

This paper attempts to study uncertainty in a unified framework. To do so, we introduce a new measure of uncertainty that is based on forecast densities. Our new measure of uncertainty enables us to make two main contributions to the literature:

(i) The first main contribution is that we use our new measure of uncertainty to distinguish between Knightian uncertainty and risk, and their relationship. The use of forecast densities is key to provide a comprehensive measure of Knightian uncertainty because it allows to quantify

national Symposium in Computational Economics and Finance, the 1st Banque de France – Norges Bank Workshop in Empirical Macroeconomics, the 24th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics, the 9th ECB Workshop on Forecasting Techniques, the IAAE 2016 Annual Conference, the 2016 CEF Conference, the Chicago Fed, UCL, York, and Henan Universities for comments.

uncertainty pertaining to situations where the odds and outcomes are known, yet either one or both are characterized inaccurately, which is the definition of Knightian uncertainty we adhere to.²

(ii) The second main contribution is that we provide a decomposition of our uncertainty measure into several components that are related to the uncertainty measures used in the literature. This analysis sheds light on why the various measures of uncertainty differ from each other, and which one is more appropriate to use depending on the goals of the researcher. Again, the use of forecast densities is key to provide a comprehensive decomposition of uncertainty into its sources. In particular, we distinguish between disagreement and aggregate uncertainty. In this respect our contribution is similar to that of Lahiri and Sheng (2010), who consider the relationship between aggregate uncertainty and disagreement over the business cycle, yet measure it in terms of uncertainty and disagreement about the mean of the distribution, as opposed to the whole distribution. Our approach further enables us to distinguish between measures of realized volatility, ex-ante uncertainty and bias. These various components have all been used in the literature as measures of uncertainty. Our approach, on the other hand, enables us to distinguish among them and understand their relationship to each other.

Several of the components mentioned above have been of interest on their own. For example, Patton and Timmermann (2010) study disagreement among professional forecasters, but do not relate disagreement to measures of uncertainty, while Lahiri and Sheng (2010) consider the relationship of aggregate uncertainty and disagreement over the business cycle, yet they do not distinguish between risk and uncertainty. Jurado, Ludvigson and Ng (2015) use the forecast error variance as a measure of uncertainty, while D'Amico and Orphanides (2014) consider ex-ante measures of risk for inflation forecasting.

In addition to our main contributions, we also study how uncertainty and its sources resolve over time as the agents get closer in time to the event. For example, Patton and Timmermann (2010) study the resolution of disagreement over time; disagreement is only one of the components of uncertainty: we investigate both how important disagreement is as a source of overall uncertainty over time, as well as how the other components of uncertainty resolve over time. Furthermore, we document the macroeconomic impact and transmission of the various sources of uncertainty that we identify.

Lastly, we use a stylized macroeconomic model as a framework to discuss the interpretation of the components of our decomposition in the presence of macroeconomic risk and ambiguity. We show that the various components in our decompositions are indeed representative of sources of

²While we attempt to quantify Knightian uncertainty defined as the agents' inability to correctly characterize probability distributions or their disagreement on them, clearly we cannot quantify uncertainty associated with the agents inability to characterize all possible states of nature or situations where they have no opinions on the probability distributions associated with known states of the nature. Thus, one can think of our Knightian uncertainty measure as a lower bound on the actual Knightian uncertainty present in the economy.

uncertainty that the model implies.

It is important to note that the existing literature has focused mainly on quantifying and understanding uncertainty associated with point forecasts, for example by mapping uncertainty to forecasters' prediction errors. Though the individual point forecasts are on average consistent with the weighted mean of their predictive probability distributions (see Lambros and Zarnowitz, 1987), predictive distributions undoubtedly contain more information. Our goal is to take advantage of the richer information content of probabilistic forecasts to quantify Knightian uncertainty and distinguish among various sources of uncertainty. Thus, an important difference between this paper and the existing literature is that we use the probabilistic forecasts provided by the U.S. Survey of Professional Forecasters (SPF) to measure and decompose uncertainty.³ We focus mainly on output growth forecasts: since output growth is indicative of business cycle fluctuations, our analysis provides an overall measure of macroeconomic uncertainty; in addition, we also discuss inflation uncertainty measures that might help understand why monetary policy affects short and long term interest rates differently (Wright, 2011).

Furthermore, a large number of uncertainty measures considered in the literature are ex-post, since they depend on realizations (such as the uncertainty measures recently proposed by Jurado, Ludvigson and Ng, 2015; Rossi and Sekhposyan, 2015, 2016; and Scotti, 2016); such ex-post measures are arguably difficult to square with the notion of economic agents' forward looking decision making. In our framework, we are able to distinguish between ex-post measures of uncertainty (for instance, realized risk or bias) and ex-ante risk. The advantage of our framework is that we are able to propose an uncertainty measure that shares properties with a large body of uncertainty measures proposed in the literature, while at the same time, enables us to disentangle components that might be preferable from a decision-theoretic point of view.

The paper is structured as follows. The next two sections present our density-forecast-based uncertainty measures and the decompositions we investigate in this paper. Section 4 discusses the SPF data used for the empirical implementation, while Section 5 presents the empirical results. In Section 6 we analyze the macroeconomic impact of the various sources of uncertainty. Section 7 interprets our decomposition through the lens of a macroeconomic model. In Section 8 we extend our results to the analysis of inflation uncertainty, while Section 9 concludes.

³Our analysis can be done with any predictive density. We choose to use predictive densities from the SPF since they are produced by professional forecasters monitoring a wider range of indicators rather than a specific parametric model. Furthermore, the SPF is known for its superior forecasting performance from a point forecasting point of view, as shown in Giannone, Reichlin and Small (2008) and McCracken, Owyang and Sekhposyan (2015), among others.

2 An Uncertainty Index Based on Density Forecasts

The uncertainty index we propose in this paper measures the distance, on average across forecasters, between the forecast distribution provided by an individual forecaster and the perfect forecast corresponding to the realization, where both are represented by cumulative distribution functions (CDFs). We denote the perfect forecast by x_{t+h} , which formally is a random variable equal to one when the actual realization y_{t+h} is below some threshold r , and it is zero otherwise: $x_{t+h}(r) \equiv 1(y_{t+h} < r)$.⁴ Note that $x_{t+h}(r)$ is defined over the support r , $r \in \mathbb{R}$; by varying r , we can focus on different parts of the predictive distribution. Let $p_{s,t+h|t}(r)$ be the s -th forecaster's probability forecast of the event $x_{t+h}(r)$ equals one, i.e. $p_{s,t+h|t}(r) = P(x_{t+h}(r) = 1 | \Omega_{s,t})$, where $s = 1, \dots, N$ and $\Omega_{s,t}$ is the information set available to forecaster s at time t . We measure the s -th forecaster's uncertainty as the Mean Squared Forecast Error (MSFE) of his/her probabilistic forecast about a particular outcome, i.e.:⁵

$$u_{s,t+h|t}(r) = E \left[(x_{t+h}(r) - p_{s,t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right], \quad (1)$$

where \mathfrak{S}_{t-R}^t is the information set between time $t - R$ and time t .⁶ Note the difference between the two information sets $\Omega_{s,t}$ and \mathfrak{S}_{t-R}^t . $\Omega_{s,t}$ is the information set available to forecaster s when making its probability forecasts. On the other hand, \mathfrak{S}_{t-R}^t is the information set that we use to average squared errors over time.

Similarly to Jurado, Ludvigson and Ng's (2015) measure, eq. (1) is a MSFE; however, it is a MSFE applied to a forecast distribution for a given binary event. As such, it measures the unpredictable component associated with each possible value in the domain of the predictive distribution. In fact, $u_{s,t+h|t}(r)$ compares the probability that forecaster s assigns to the different states of nature with the realization, while error-based measures à la Jurado, Ludvigson and Ng (2015) compare the point forecast with the realization.⁷

The overall measure of uncertainty is then defined as the average of the individual uncertainty across forecasters:

$$u_{t+h|t}(r) = \frac{1}{N} \sum_{s=1}^N u_{s,t+h|t}(r) = \frac{1}{N} \sum_{s=1}^N E \left[(x_{t+h}(r) - p_{s,t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right].$$

As mentioned above, by varying r we can explore measures of uncertainty in different parts of the predictive density. We focus on an overall measure of uncertainty (which we label "Uncertainty")

⁴This notation is consistent with Hersbach (2000).

⁵In the forecasting literature, this MSFE is known as the Brier score.

⁶To simplify notation, we assume in this section that R and N are fixed over time, although in the empirical application we let N vary.

⁷In fact, if one associates the value $r \in \mathbb{R}$ with the corresponding quantile of the distribution, our uncertainty index measures an average squared error for that quantile.

that integrates the squared probability forecast errors over the whole domain of the distribution, that is:⁸

$$U_{t+h|t} = \int_{-\infty}^{+\infty} u_{t+h|t}(r) dr. \quad (2)$$

A graphical interpretation is provided in Figure 1. In the figure, the actual realization equals -2 , denoted by a vertical bar on the left panel; the predictive density is the Normal distribution. The panel on the right shows the CDF of the Normal distribution, as well as that of the perfect forecast, evaluated at the realization $y_{t+h} = -2$. Thus, the ideal CDF equals zero for $r < y_{t+h} (= -2)$ and one otherwise. At the realization, the distance between the CDFs of the perfect forecast and the forecasted distribution, $x_{t+h}(r) - p_{t+h|t}(r)$, is depicted by a hollow vertical bar. Our measure of uncertainty in eq. (2) squares this measure and integrates it over the various values of r (and averages it over time).

INSERT FIGURE 1

3 The Sources of Uncertainty

This section presents our main decompositions of uncertainty into its sources.

3.1 Aggregate Uncertainty and Disagreement

One of the goals of this paper is to link existing measures of uncertainty based on aggregate data with uncertainty measures based on disagreement among forecasters. To do so, we define an aggregate probability density ($\{p_{t+h|t}(r)\}_{r \in R}$), which is related to the individual ones ($\{p_{s,t+h|t}(r)\}_{r \in R}$) by:

$$p_{t+h|t}(r) = \frac{1}{N} \sum_{s=1}^N p_{s,t+h|t}(r). \quad (3)$$

The corresponding uncertainty measure for the aggregate predictive density is:

$$u_{t+h}^A(r) \equiv E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right].$$

⁸Our measure of uncertainty derives from the Continuous Rank Probability Score (CRPS), which is widely used to assess the quality of forecast distributions in statistics. In fact, the CRPS is the integral of Brier scores (Hersbach, 2000, eq. 7). In this particular case it could be viewed as an average CRPS across the forecasters. Note that eq. (2) is the negative of the CRPS, as defined in Gneiting and Raftery (2007). Moreover, it is a conditional average. In other words, for an individual forecaster, Gneiting and Raftery (2007) would define $CRPS_{s,t} = -\int_{-\infty}^{+\infty} (x_{t+h}(r) - p_{s,t+h|t}(r))^2 dr$, while our forecaster specific uncertainty measure would be $\int_{-\infty}^{+\infty} E[(x_{t+h}(r) - p_{s,t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t] dr$.

Appendix A shows that we can decompose the overall uncertainty measure as follows:

$$\begin{aligned}
u_{t+h|t}(r) &= E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right] + E \left[\frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right] \\
&= u_{t+h|t}^A(r) + d_{t+h|t}(r),
\end{aligned} \tag{4}$$

where $d_{t+h|t}(r) \equiv \frac{1}{N} \sum_{s=1}^N E \left[(p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 | \mathfrak{S}_{t-R}^t \right]$ measures the disagreement between individual forecast densities and the aggregate forecast density, and it is similar to the disagreement defined in Patton and Timmermann (2010) for point forecasts. Lahiri and Sheng (2010, eq. 18) discuss a similar decomposition for point forecasts.

Note that the decomposition in eq. (4) holds for a particular threshold r , thus it accounts for a forecast error associated with the binary outcome $1(y_{t+h} < r)$. The overall measure of uncertainty accounts for uncertainty at all possible values of r by considering the integral of the decomposition in eq. (4) over r . Thus, we have ‘‘Uncertainty’’ decomposed into ‘‘Aggregate Uncertainty’’ and ‘‘Disagreement’’:⁹

$$\begin{aligned}
U_{t+h|t} &\equiv \int_{-\infty}^{\infty} u_{t+h|t}(r) dr = \int_{-\infty}^{\infty} u_{t+h|t}^A(r) dr + \int_{-\infty}^{\infty} d_{t+h|t}(r) dr \\
&\equiv \underbrace{U_{t+h|t}^A}_{\text{‘‘Aggregate Uncertainty’’}} + \underbrace{D_{t+h|t}}_{\text{‘‘Disagreement’’}}
\end{aligned} \tag{5}$$

3.2 Knightian Uncertainty and Risk

As shown in Appendix A, we can further decompose the aggregate uncertainty, $U_{t+h|t}^A(r)$ into components that measure mean bias, dispersion of probability forecasts, realized risk and a covariance term between the forecasted and the ideal distribution as follows:

$$\begin{aligned}
u_{t+h}^A(r) &= \left([E(p_{t+h|t}(r) | \mathfrak{S}_{t-R}^t) - E(x_{t+h}(r) | \mathfrak{S}_{t-R}^t)]^2 \right) + V(p_{t+h|t}(r) | \mathfrak{S}_{t-R}^t) \\
&\quad + V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) - 2Cov(x_{t+h}(r) p_{t+h|t}(r) | \mathfrak{S}_{t-R}^t),
\end{aligned} \tag{6}$$

where $V(\cdot)$ denotes variance. Since the covariance term turns out to be rather small empirically, we summarize aggregate uncertainty with the following additive decomposition:

$$U_{t+h|t}^A \approx \underbrace{B_{t+h|t}}_{\text{‘‘Mean-Bias’’}} + \underbrace{V_{t+h|t}}_{\text{‘‘Dispersion’’}} + \underbrace{Vol_{t+h|t}}_{\text{‘‘(Realized) Risk’’}} \tag{7}$$

where:

⁹ A reason why the aggregate probability distribution, measured with a simple average of the individual probability distributions, is a good measure of aggregate uncertainty is the fact that, as in the context of point forecasts, combinations constructed by simple averages typically result in more accurately calibrated densities. Furthermore, the average of probability distributions is a measure widely used in a variety of central banks and policy institutions.

- $B_{t+h|t} \equiv \int_{-\infty}^{\infty} \left([E(p_{t+h|t}(r) | \mathfrak{S}_{t-R}^t) - E(x_{t+h}(r) | \mathfrak{S}_{t-R}^t)]^2 \right) dr$ is the mean squared bias of the forecast distribution;
- $V_{t+h|t} \equiv \int_{-\infty}^{\infty} V(p_{t+h|t}(r) | \mathfrak{S}_{t-R}^t) dr$ is the uncertainty about the ex-ante subjective probabilities in the aggregate distributional forecast;
- $Vol_{t+h|t} \equiv \int_{-\infty}^{\infty} V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) dr$ is the realized variance of the binary outcome, $x_{t+h}(r) \equiv 1(y_{t+h} < r)$, and thus stands for the inherent risk in the data.

The three component decomposition in eq. (7) has an interesting interpretation. We view the realized volatility component $Vol_{t+h|t}$ as a measure of the underlying uncertainty in the data, and thus a measure of realized risk. On the other end, we view the bias component $B_{t+h|t}$ as a measure of how distant the predictive density is from the perfect prediction on average, while the dispersion, $V_{t+h|t}$, measures the variability in the predictive density. As we will show, $V_{t+h|t}$ is empirically small, so it can be ignored. Knightian uncertainty is measured, in our view, as the sum of bias and disagreement. In fact, Knightian uncertainty measures how uncertain agents were about events, either because they were unable to correctly assign probabilities to future outcomes even though they agreed to them, or because they disagreed on those probabilities. The realized variance or realized volatility, instead, is a measure of risk. To summarize, we have the following “Knightian uncertainty/(Realized) Risk” decomposition:

$$U_{t+h|t} \approx \underbrace{Vol_{t+h|t}}_{\text{“(Realized) Risk”}} + \underbrace{B_{t+h|t} + D_{t+h|t}}_{\text{“Knightian Uncertainty”}} .$$

3.3 Ex-ante Vs. Ex-post Uncertainty

It is important to note that our proposed measure of uncertainty, $U_{t+h|t}$, as well as aggregate uncertainty $U_{t+h|t}^A$, are constructed using ex-post realizations of the data. Thus, it is interesting to refine our measure by distinguishing between an ex-ante component (that does not include the realizations) and an ex-post component (which does). Also, one might wonder how the expected mean and the variance embedded in the forecast distribution affect our measure of uncertainty. Let the aggregate predictive distribution for the forecast of y_{t+h} made at time t be Normal with mean $\mu_{t+h|t}$ and variance $\sigma_{t+h|t}^2$ and the data be i.i.d.¹⁰ We have the following “Ex-ante/Ex-post” decomposition:

¹⁰Since the data are assumed to be i.i.d. in this sub-section, we could omit the time subscripts; however, we decided to keep them to make the notation consistent with the rest of the paper.

$$U_{t+h|t}^A = E \left\{ \underbrace{\left[2\sigma_{t+h|t}\phi\left(\frac{y_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}}\right) + (y_{t+h} - \mu_{t+h|t}) \left(2\Phi\left(\frac{y_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}}\right) - 1 \right) \right]}_{\text{“Ex-post”}} \right\} \quad (8)$$

$$- \underbrace{\frac{\sigma_{t+h|t}}{\sqrt{\pi}}}_{\text{“Ex-ante”}} \left| \mathfrak{S}_{t-R}^t \right\} \quad (9)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the PDF and the CDF of the Normal distribution, respectively. The proof is provided in Appendix A and follows Gneiting and Raftery (2007).¹¹

The rightmost component, $\sigma_{t+h|t}/\sqrt{\pi}$, is the only component that is not affected by the realization, so we refer to it as the “ex-ante” measure of uncertainty. In fact, as the proof suggests, this is the component that arises from the average distance of random draws from a given predictive distribution. Moreover, it is a function of the standard deviation of the forecaster’s density forecasts, and a common measure used in the uncertainty literature as a measure of ex-ante uncertainty. Note that the ex-ante measure of uncertainty is simply $\sigma_{t+h|t}/\sqrt{\pi}$, which, under Normality, is a monotone function of the width of the predictive distribution. Thus, the ex-ante measure is linked to the inter-quantile range measure proposed by Zarnowitz and Lambros (1987), among others.¹² Our ex-ante component might be viewed as a measure of ex-ante risk. Note that, from eqs. (7) and (8), we have that $Ex\text{-post} \approx B_{t+h|t} + V_{t+h|t} + Vol_{t+h|t} + Ex\text{-ante}$. Thus, the ex-post measure of aggregate uncertainty combines components of Knightian uncertainty, $B_{t+h|t}$, realized risk (measured by the volatility in the economy, $Vol_{t+h|t}$), ex-ante risk (measured by the variance of the predictive densities of the forecasters, $Ex\text{-ante}$) and dispersion, $V_{t+h|t}$. Note the difference between $V_{t+h|t}$ and $Ex\text{-ante}$: the first measures the variability of the probability distribution, while the second measures the width of the distribution at a particular point in time. Thus, if the aggregate density forecast does not change over time, $V_{t+h|t}$ would be zero. However, $Ex\text{-ante}$ will not be zero as long as the forecasters provide a distributional forecast.

We should note that there is a major difference between the two decompositions in that the “Ex-ante”/“Ex-post” decomposition is written in terms of the moments of the original predictive distribution, while the “Knightian Uncertainty/(Realized) Risk” decomposition is in terms of binary outcomes summarized by $x_{t+h}(r)$. As such, the latter decomposition could be applied to general situations (general forms of distribution and non-i.i.d. data), while the former one relies heavily on the assumption of Gaussianity and independence in the underlying predictive distribution. D’Amico

¹¹Note that even if $U_{t+h|t}^A$ is the difference of two components, it is always positive; thus, the ex-post component is always bigger than the ex-ante one.

¹²For a Gaussian distribution, the inter-quantile range is 1.34σ .

and Orphanides (2014) and Giordani and Soderlind (2003) provide empirical support in favor of Gaussianity for the Survey of Professional Forecasters, and the i.i.d. assumption would be satisfied for correctly calibrated conditional density forecasts.

A general note that applies to all proposed decompositions is that the resulting components are not orthogonal to each other. This is in line with the rest of the empirical literature which typically finds that a variety of uncertainty measures, constructed from different sources and measuring different aspects of uncertainty, are correlated with each other.

4 The Data

We use density forecasts from the Survey of Professional Forecasters (SPF) to calculate our uncertainty measures. The Federal Reserve Bank of Philadelphia provides the aggregate (mean probability distribution) forecasts, as well as the underlying disaggregate density forecasts of a panel of professional forecasters.¹³ We use the real GNP/GDP growth density forecasts to extract measures of macroeconomic uncertainty, as real GNP/GDP fluctuations are indicative of the state of the business cycle, and therefore are representative of macroeconomic uncertainty (Stock and Watson, 1999).

In the SPF data set, forecasters are asked to assign a probability value (over pre-defined intervals) to inflation and output growth for the current and the following (one-year-ahead) calendar years. The growth rate is defined as the rate of change in the average GDP from one year to another. The forecasters update the assigned probabilities for the current-year and the one-year-ahead forecasts on a quarterly basis. Thus, by construction, SPF forecasters provide four quarterly forecasts of the same target variable each year; this type of forecasts are typically referred to in the literature as “fixed-event” or “moving-horizon” forecasts. Being fixed-event forecasts, their horizon changes over the quarter. We use the method proposed by Doornik et al. (2012) to transform the SPF fixed-event forecasts into fixed-horizon forecasts by constructing a weighted average of the current-year and next-year forecasts. In detail, for each quarter the survey contains a pair of “fixed-event” density forecasts for the current-year, which we label $\hat{f}_{t+k|t}^{FE}$, and for the next-year, which we label $\hat{f}_{t+k+4|t}^{FE}$. The four-quarter-ahead (fixed-horizon) forecast at time t , which we label $\hat{f}_{t+4|t}^{FH}$, is calculated as the average of the two fixed event forecasts using weights that are proportional to their share of the overlap with the forecast horizon. Let k denote the number of quarters from time t until the end of the year. In quarter one, $k = 4$, while in quarter four, $k = 1$. Thus, for example, in the third quarter of the year, the four-step-ahead fixed-horizon forecast overlaps with the current year forecasts and next year forecasts 50% of the time, respectively. Thus, it would be the weighted average of the two-fixed event forecasts with weights equal to 2/4 and 2/4. Thus, in

¹³The composition of the forecasters can change over time.

general, for $k = 1, 2, 3, 4$.¹⁴

$$\widehat{f}_{t+4|t}^{FH} = \frac{k}{4} \widehat{f}_{t+k|t}^{FE} + \frac{4-k}{4} \widehat{f}_{t+k+4|t}^{FE}. \quad (10)$$

INSERT FIGURE 2 HERE

Panels A and B in Figure 2 show the evolution of the current and next year densities over time. The figures plot the mean as well as several quantiles of the distribution, together with the realization. Panel C, on the other hand shows the fixed horizon forecast, eq. (10). The fixed-horizon forecast is by construction less smooth than the fixed-event forecasts. However, both share the same feature that ex-ante uncertainty was higher earlier in the sample, in the sense that both density forecasts have a wider distribution prior to the mid-1980s relative to the later part of the sample; this suggests that forecasters noticed the Great Moderation starting mid-1980s. There appears to be no dramatic shift in the forecasted densities after the Great Recession. Some descriptive statistics on the SPF distributions is provided in Appendix B.

The analysis of SPF probability distributions is complicated since the SPF questionnaire has changed over time in various dimensions: there have been changes in the definition of the variables, the intervals over which probabilities have been assigned, as well as the time horizon for which forecasts have been made. To mitigate the impact of these problematic issues, we truncate the data set and consider only the period 1981:III-2014:II.¹⁵

As noted, our uncertainty measure depends on realizations. The realized values of output growth are from the real-time data set for macroeconomists, also available from the Federal Reserve Bank of Philadelphia. We use the four-quarter-ahead growth rates of output and prices calculated from the first release of the realization. For instance, in order to get the 4-quarter-ahead realization at the start of our sample, 1981:III, we calculate the growth rate between 1982:III and 1981:III using the 1982:IV vintage of the data.

5 The Dynamics of Uncertainty over Time, and Its Sources

Figure 3, Panel A, shows the evolution of our estimated measure of uncertainty and its components, aggregate uncertainty and disagreement, over time. The figure highlights two interesting facts: disagreement is, in magnitude, only a small portion of the overall measure of uncertainty;¹⁶ in

¹⁴The literature suggest alternatives to this weighting. For instance, Knueppel and Vladu (2016) propose weights for aggregating the fixed-event point forecasts which minimize the mean squared forecast error loss of the fixed-horizon forecast. This methodology requires a researcher to take a stand on the data generating process and is not directly applicable to our case of density weighting.

¹⁵We focus on quarterly data. See instead Ferrara and Guérin (2015) for a high-frequency analysis of uncertainty shocks.

¹⁶The magnitudes of $U_{h+h|t}$ and $U_{h+h|t}^A$ are reported on the y-axis on the left while that of disagreement is reported on the y-axis on the right. The magnitude of disagreement is small. This is due to the fact that, unlike the existing

addition, it is trending down until the financial crisis of 2007; this is in sharp contrast with the overall measure of uncertainty, as well as aggregate uncertainty, which have clear spikes in the early 1980s, early 2000s and the financial crisis. Thus, using disagreement as a measure of uncertainty may result in underestimating both the overall level of uncertainty in the economy as well as its fluctuations over time, as currently the level of disagreement is similar to what it was in the mid-1990s and lower than its value in the late 1980s. In addition, most would agree that the early 2007-2008 were probably the most uncertain times in the latest decades; while disagreement increases during that period, it peaks only much later, after the end of the recession, in 2009. Thus, disagreement (i.e., the component of Knightian uncertainty due to disagreement among forecasters) may not be a timely measure of uncertainty. Note that this result is not an artifact of constructing disagreement measures based on density forecasts: Sill (2014, Figure 1) shows a similar delay. In particular, Sill (2014) plots the dispersion of the mean one-year-ahead real GDP growth rate forecasts measured by the inter-quantile range: the first peak in the disagreement does not appear until the middle of the recession.

INSERT FIGURE 3 HERE

Panel B in Figure 3 depicts the decomposition of aggregate uncertainty into Knightian uncertainty and realized risk. The figure suggests that realized risk (measured by $Vol_{t+h|t}$) was an important component of uncertainty throughout the last three decades, as was Knightian uncertainty, measured by the mean bias component. Some differences between the two are important to note, however. The realized risk component was high during the latest financial crisis, and sharply decreased as soon as the recession was over; Knightian uncertainty (measured by the mean bias component, $B_{t+h|t}$) remained persistently high even after the end of the crisis. Thus, overall uncertainty remained persistently high after the end of the latest recession mostly because of forecasters' errors as opposed to risk being high. The role of dispersion in probability forecasts ($V_{t+h|t}$) as well as the co-movement between prediction and realization ($Cov_{t+h|t}$) are negligible for the cyclical dynamics of aggregate uncertainty.

Turning to the ex-ante and ex-post components, depicted in Panel C of Figure 3 together with the aggregate uncertainty measure ($U_{t+h|t}^A$), it is interesting to note that ex-ante uncertainty is quite constant in the 1980s and up to 2007. Thus, movements in uncertainty during that period cannot be attributed to changes in ex-ante uncertainty. Ex-ante uncertainty does increase during the latest recession, but only towards its end, and spikes much later than the peak of the recession. This suggests that measures of volatility in the forecasters' predictive distributions are, themselves, not timely measures of uncertainty.

measures of disagreement on point forecasts, we measure disagreement in probabilities, not in the mean forecast.

Finally, it is also of interest to investigate how the various components of uncertainty evolve as the forecasters get closer in time to the realization date, that is, as the forecast horizon becomes shorter. We separately consider forecasts for $h = 1, 2, \dots, 7, 8$ and compare them with the fixed-event realization. Both uncertainty as well as aggregate uncertainty decrease as the forecast horizon increases (Panel A in Figure 4, top left and right graphs). It may seem counter-intuitive that uncertainty decreases at longer horizons; to understand why, we examine its components. Clearly, disagreement decreases as forecasters get closer to the realization: in fact, disagreement decreases on average as the horizon decreases (cfr. bottom graph in Figure 4, Panel A). This finding is reassuring, as it is reminiscent of what Patton and Timmermann (2010) discovered for point forecasts, and our results show that similar results hold for disagreement calculated on density forecasts. The mean bias also decreases as the horizon decreases (Panel B in Figure 4). On the other hand, the dispersion of the density forecasts increases, thus increasing the aggregate uncertainty. The realized variance and covariance are constant over the horizons, and the latter hovers around zero.

INSERT FIGURE 4 HERE

The most striking patterns are displayed by ex-ante and ex-post uncertainty, depicted in Figure 4, Panel C. Clearly, ex-ante uncertainty decreases monotonically as the forecast horizon decreases; that is, forecasters' predictive densities become more spread out when the forecast horizon increases, thus reflecting more uncertainty in the economy when looking at events that are further in the future. However, there is no clear pattern in ex-post uncertainty. This means that, even though the forecasters' predictive densities become tighter as the realization gets closer in time, the uncertainty in the actual realizations does not diminish, as the size of the forecast errors does not diminish with the horizon.

Comparing the evolution of the ex-ante uncertainty in Panel C and the dispersion of the aggregate predictive density, $V_{t+h|t}$ in Panel B, we note that, although forecasters, on average, become less confident about the future as the forecast horizon increases, their views about uncertainty does not seem to be updated often for forecasts that are further in the future, thus resulting in the low variability of the predictive distribution over time. Moreover, as the distribution becomes more spread out with the forecast horizon, it has a higher chance of including the realization, thus resulting in a decline in the aggregate and overall uncertainty.

6 Understanding the Measures of Uncertainty in the Literature and Their Macroeconomic Effects

In this section, we use our decomposition to shed some light on why existing measures of uncertainty differ from each other. Understanding why they differ provides important insights on which measure

is the most appropriate for a particular analysis.

The top panel in Figure 5 plots Jurado, Ludvigson and Ng's (2015) uncertainty measure together with Baker, Bloom and Davis' (2016) index.¹⁷ Both indices are standardized for comparison. The figure shows that the former is overall smaller than the latter until 1995, then it becomes overall bigger, and in particular spikes up earlier than the latter during the latest financial crisis of 2007-2008. The lower panel plots the decomposition of our aggregate uncertainty index into ex-ante and ex-post components. The ex-post component is lower than the ex-ante component up to mid-1992, then it becomes systematically larger, and spikes up around 2007-2008, behaving similarly to how the Jurado, Ludvigson and Ng's (2015) behaves relative to Baker, Bloom and Davis (2016). Thus, it seems that the Baker, Bloom and Davis (2016) uncertainty measure is driven more by ex-ante uncertainty, while the Jurado, Ludvigson and Ng (2015) uncertainty measure is clearly affected by ex-post uncertainty, namely uncertainty due to misspecification in the predictions.

INSERT FIGURE 5 HERE

To estimate the effects of the uncertainty and its components on the economy, we estimate a Vector Autoregression (VAR) that includes (the log of) real GDP, (the log of) employment, the Federal Funds rate, (the log of) stock prices and the specific uncertainty indices one at a time. Identification is achieved via a Cholesky procedure, which follows the order in which the variables are listed. The VAR specification is the same as in Baker, Bloom and Davis (2016), although ours is at the quarterly frequency, and accordingly we use GDP instead of real industrial production. We order the variables as in Jurado et al.'s (2015) benchmark specification, i.e. from slow to fast moving. For completeness, we investigate the robustness of our results in a larger VAR in the Not-for-Publication Appendix. To better interpret and compare the magnitude of the effects of the uncertainty indices, the uncertainty indices are standardized by their own means and variances.

Panel A in Figure 6 shows the effects of our uncertainty index on the economy. Clearly, an increase in uncertainty has recessionary effects: both GDP and employment decrease, as well as the interest rate and the S&P 500. Panels B and C describe the effects of each of the components in the decomposition. Panel B shows the effects of a shock to aggregate uncertainty, which is in line with that of uncertainty since aggregate uncertainty is the main determinant of the total. Panel C focuses on disagreement; it also decreases employment although by a smaller magnitude; at the same time, it has no significant effects on the remaining variables.

INSERT FIGURE 6 HERE

Figure 7 shows the effects of uncertainty measured by mean bias, realized volatility and the dispersion in the probability forecasts. The mean bias and realized volatility appear to have recessionary effects (Panels B and D); dispersion in the density forecasts (Panel C) drives down

¹⁷We are using Jurado, Ludvigson and Ng's (2015) one-year-ahead uncertainty index.

employment, while it increases stock prices and output. It is important to note that, in magnitude, the mean bias and realized volatility have similar macroeconomic impact, though these effects are statistically significant for the first but not for the second.

INSERT FIGURE 7 HERE

The effects of ex-ante and ex-post uncertainty on other macroeconomic variables are depicted in Figure 8. They both lead to decreases in employment, interest rates and stock prices of similar magnitude; an increase in ex-ante uncertainty, however, has a small negative impact effect on GDP, while the medium run effect is positive and small, and the longer run effect is again negative; the effects of ex-post uncertainty on GDP are, instead, negative and large.

INSERT FIGURE 8 HERE

Figure 9 compares the results with those in the existing literature; the latter are also obtained by estimating VARs that include (the log of) real GDP, (the log of) employment, the Federal Funds rate, (the log of) stock prices, and the alternative uncertainty index, which is demeaned and standardized as well. The alternative uncertainty indices that we explore (one-at-a-time) include: Bloom (2009), labeled “VXO”; Baker et al.’s (2016) policy uncertainty index, labeled “BBD”; Jurado, Ludvigson and Ng (2015), labeled “JLN”; and Scotti’s (2016) macroeconomic surprise-based uncertainty index.

INSERT FIGURE 9 HERE

Panel A in Figure 9 shows that the VXO and BBD indices have similar effects on the economy, while an increase in uncertainty measured by the Jurado, Ludvigson and Ng’s (2015) index are qualitatively similar but much larger in magnitude, and, thus, are similar to the effects that we uncover for our ex-post index. The effects of Scotti’s index are again recessionary for GDP, employment and stock markets, and lead to an increase in the interest rate. The effects of this index are small and overall insignificant. The effects of our realized volatility measure are more similar to those of the VXO.

7 The Dynamics of the Sources of Uncertainty Through the Lens of a Model

We consider a model of ambiguity which follows Ilut and Schneider (2014). The model is as follows. We assume that GDP growth, Z_{t+1} , evolves according to an autoregressive model with a time varying mean, μ_t^* :

$$Z_{t+1} = \rho_z Z_t + \mu_t^* + u_{t+1}, \tag{11}$$

where u_{t+1} is i.i.d. $N(0, \sigma_u^2)$ and μ_t^* is deterministic such that its empirical sequence converges to an i.i.d. stochastic process $N(0, \sigma_\mu^2)$, where $\sigma_\mu^2 = \sigma_z^2 - \sigma_u^2$. Consequently, the observed values of $z_t \equiv Z_{t+1} - \rho_z Z_t$ look like realizations from an i.i.d. process with mean zero and variance σ_z^2 . For all practical purposes, we treat μ_t^* as a realization from a stochastic process $N(0, \sigma_\mu^2)$. Moreover, μ_t^* and u_t are assumed to be independent. Thus, GDP growth is driven by two sources of uncertainty in the economy: the first source is the unpredictable shocks, u_{t+1} ; the second, μ_t^* , is a proxy for ambiguity, as we discuss below.

We assume that the agents in this model know that the data generating process for GDP growth is autoregressive with persistence ρ_z , and that there are two sources of uncertainty; however, they do not observe μ_t^* and u_{t+1} , even though they know the probability distribution of u_{t+1} . They gather intangible information about μ_t^* , which sometimes makes them relatively confident that the correct forecast of future GDP growth is $\rho_z Z_t$, and sometimes less confident, i.e. the signal is sometimes less and sometimes more ambiguous. One could think of a situation where the agents acquire poor quality information or conflicting news from newspapers or professional forecasters. The ambiguity is modeled by letting agents form their beliefs about GDP growth dynamics based on the following law of motion:

$$Z_{i,t+1} = \rho_z Z_{i,t} + \mu_{i,t} + u_{t+1}, \quad i = 1, 2, \dots, N \quad (12)$$

where $\mu_{it}^* \in [-a_{i,t}, -a_{i,t} + 2|a_{i,t}|]$, N is the total number of agents (equal to 100), and u_{t+1} is i.i.d. $N(0, \sigma_u^2)$. The bounds on μ_{it}^* formalize the idea that sometimes agents are more ambiguous regarding the second source of disturbance to output growth: those situations are associated, in the model, with a larger value of $a_{i,t}$, which implies a larger set of beliefs and more ambiguity perceived at time t by agent i . Thus, we refer to $a_{i,t}$ as the ambiguous component, or Knightian uncertainty.¹⁸

Furthermore, agents receive signals about $\mu_{i,t}^*$ from the process:

$$a_{i,t+1} - \bar{a}_i = \rho_{a,i}(a_{i,t} - \bar{a}_i) + \sigma_{a,i} \epsilon_{t+1}^a, \quad (13)$$

where ϵ_{t+1}^a is i.i.d. $N(0, 1)$. One can view ϵ_{t+1}^a as a signal that the agent gets about the ambiguity component, whose volatility depends on $\sigma_{a,i}$. In some periods the signal results in a higher $a_{i,t}$; in such cases, there is more ambiguity and the set of beliefs is larger. In other periods, depending on the received information, the set can be smaller, thus the agents are less ambiguous about the stochastic disturbances in the data generating process. Furthermore, we impose parameter restrictions to ensure that the average ambiguity is less than the total uncertainty about the process of Z_{t+1} : these restrictions are such that $\bar{a}_i = n_i \sigma_z$ and $\sigma_{a,i} = \sigma_n \sigma_z$ for $n \in (0, 1)$ for every agent

¹⁸This notion of Knightian uncertainty is similar to that of Ilut and Schneider (2014). We should note, however, that they assume that the total factor productivity shocks are ambiguous, while we do the same for output growth.

i , where n_i and σ_n are parameters (one can think of \bar{a}_i as the unconditional mean and $\sigma_{a,i}^2$ as the unconditional variance of the shock to perceived ambiguity). In particular, $n_i \sim iidN\left(n, \sigma_{n,I}^2\right)$, where $\sigma_{n,I}^2$ controls the cross sectional variability of N .¹⁹

Finally, when faced with ambiguity, modeled with eq. (13), the agents choose:

$$\mu_{i,t}^{**} = \min([-a_{i,t}, -a_{i,t} + 2|a_{i,t}|]). \quad (14)$$

Thus, the effective perceived law of motion for agent i becomes:

$$Z_{i,t+1} = \rho_z Z_{i,t} + \mu_{i,t}^{**} + u_{t+1}. \quad (15)$$

Note that when $a_{i,t}$ is bigger, ambiguity is higher, the set of beliefs is bigger, and the wider interval implies a lower worst case mean that the agents choose.

Our model is a simplification of Ilut and Schneider (2014): to be precise, they model ambiguity and risk about the technology process. However, under the assumption of fixed inputs, this would directly translate into a similar output growth dynamics. Thus, for simplicity, we directly model the dynamics of output growth and calibrate the parameters of the output growth process, ρ_z and σ_z , based on an AR(1) model estimated on the quarterly growth rate for the U.S. GDP. On the other hand, the ambiguity parameters, i.e. ρ_a , n and σ_n , are borrowed from their posterior mode estimates with the caveat that their estimates apply to the ambiguity in total factor productivity rather than output growth. Table 1 summarizes our baseline parameter values. Since μ_t and u_{t+1} can not be identified separately, the values for their respective variances are assigned arbitrarily. We let $\sigma_\mu = 0.5$, while σ_u is assigned a value to match the total conditional volatility in the output growth observed in the data, σ_z .

INSERT TABLE 1 HERE

We consider four scenarios. In the first three scenarios, there is no cross-sectional heterogeneity in ambiguity, i.e. $\sigma_{n,I}^2 = 0$ and $n_i = n$ for every agent; in the fourth scenario we consider heterogeneity by letting $n_i \neq n$.

Scenario 1: Ambiguity. We increase the level of ambiguity in the model, i.e. the level of n . We consider shifting the value of n from 0.2 to 0.8. While the data is generated by equation (11), the agents forecast output growth using the law of motion in equation (15). In this exercise we are changing the set of possible values that μ_t can take: as n increases, both the conditional and unconditional means of a_{t+1} increase – see equation (11), and the signals the agents get about the additional source of uncertainty, denoted by the set $[-a_t, -a_t + 2|a_t|]$, become noisier.

¹⁹Alternatively, one could model the level of ambiguity to be uniformly distributed across the forecasters. This would attenuate ambiguity.

Scenario 2: Risk and ambiguity. We increase the level of risk by increasing the value of σ_u from 0.3 to 1. In this experiment the model is still described by eq. (11), the perceived law of motion is described by eq. (15), while learning under ambiguity occurs under eq. (13). In this case, increasing the level of uncertainty increases both the objective and perceived level of uncertainty. However, given that $n_i = n$, $\bar{a} = n\sigma_z$ and $\sigma_a = \sigma_n\sigma_z$ for $n \in (0, 1)$, where $\sigma_z^2 = \sigma_\mu^2 + \sigma_u^2$, then both the level of ambiguity (\bar{a}) and the uncertainty about ambiguity (σ_a) increase. Thus, an increase in σ_u increases both risk and Knightian uncertainty in the model.

Scenario 3: Risk but no ambiguity. We increase the level of risk as in Scenario 2, yet change the model such that the agents are forecasting based on the true model: $\mu_t^{**} = \mu_t^*$. Thus, there is no ambiguity. In other words, the true model is still the one described by (11), while the model used for forecasting is not determined by equation (15), but instead by equation (11) itself. The design in this scenario intends to explore how the ex-post and overall uncertainty evolve when there is no ambiguity.

Scenario 4: Disagreement. We increase the variance of ambiguity across agents in the model, i.e. $\sigma_{n,I}$. We consider varying the value of $\sigma_{n,I}$ from 0.5 to 1 and letting $\rho_{a,i} \sim N(\rho_a, 0.01)$ be heterogeneous across agents. This way, agents differ both because the volatility of the signal they receive and its persistence. Note that, in this case, agents disagree on the level of ambiguity, although the aggregate level of ambiguity in the data is unchanged; that is, on average, \bar{a} equals $n\sigma_z$, which does not change as $\sigma_{n,I}$ increases.

INSERT FIGURE 10 HERE

We simulate the model for 254 periods for each of these scenarios (using an additional 100 periods as a burn-in sample); we use the simulated data to construct the components of our proposed decompositions and plot them over time.

Panel A in Figure 10 depicts the results for Scenario 1. The increase in ambiguity increases the Mean-Bias and the Ex-Post components of uncertainty, as well as the overall uncertainty. On the other hand, there is no change in either the perceived or the realized volatility, that is, the Ex-ante and Realized Risk components, respectively. This follows from the fact that: (i) the data generating process has not changed, and, thus, the realized variance (σ_z^2) has remained the same; and (ii) as eq. (15) suggests, the overall level of the ex-ante variance (σ_u^2) does not depend on n .

Panel B shows the simulation results for the second scenario. Here the increase in σ_u increases the measures of ex-ante (σ_u^2 itself) and ex-post risks. It is also important that there is feedback from risk to ambiguity. As discussed in the description of Scenario 2, both the mean (\bar{a}) and the variance (σ_a^2) of ambiguity are affected by the increase in the overall risk. Consequently, the overall measure of uncertainty increases due to both sources: increase in risk and increase in ambiguity.

Panel C shows the dynamics of uncertainty and its components when there is an increase in risk in a model with no ambiguity. In this setup it appears that both ex-ante and ex-post components of uncertainty increase. However, this increase is proportional such that the average level of overall uncertainty increases due to the upward shift in ex-ante uncertainty and its volatility mimics that of ex-post uncertainty (in the right panel). Thus, comparing Panels B and C suggests that, in the presence of ambiguity, uncertainty increases proportionally more than the increase of risk.

Finally, Panel D shows the dynamics when there is an increase in the cross-sectional dispersion of ambiguity while the overall level of ambiguity remains unchanged. Note that the component that is most largely affected by the increase in the cross-sectional dispersion in ambiguity is disagreement, as we would expect.

To summarize, our simulations show that the increase in ambiguity can increase the ex-post component, as well as the mean-bias, thus resulting in an overall increase in uncertainty. The increase in the true volatility of the DGP increases both the realized volatility as well as the ex-ante volatility measures. However, the increase in the overall uncertainty affects the ex-post volatility and mean-bias as well. In the absence of ambiguity, the impact on the bias is negligible (it is more similar to noise), thus the increase in the aggregate uncertainty reflects the increase in the ex-post volatility. On the other hand, the increase in the ex-post uncertainty is twice as much the increase in the ex-ante uncertainty, such that the resulting measure of aggregate uncertainty still reflects the increase in the ex-ante uncertainty. Now, in the presence of ambiguity, on the other hand, the bias goes up and the ex-post uncertainty goes up proportionally more, such that the aggregate uncertainty reflects the increase in all sources of uncertainty.

It appears that our simulation results could be reconciled with our empirical findings. The proposed model has a potential to generate an ex-ante uncertainty measure that is smoother than the realized variance. Moreover, the model has potential to generate relatively volatile measures of bias, as well as ex-post uncertainty. Our simulation results also suggest the existence of ambiguity in the empirical setup as the aggregate uncertainty does not move proportionally with the variance: in fact, the predominant sources of aggregate uncertainty are the Knightian measures.²⁰

8 Inflation Uncertainty

In this last section, we focus on inflation uncertainty. Understanding inflation uncertainty is important for several reasons. High uncertainty about future inflation, possibly spurred by high inflation itself, may have effects on real variables (Ball, 1990). For example, Gurkaynak and Wright (2012) and Wright (2011) have argued that inflation uncertainty matters because it might help explain

²⁰Note that it is possible that we underestimate the effect of the Knightian uncertainty, since it is possible that the data generating process, thus realized volatility, can also change in response to ambiguity.

the behavior of bond risk premia, and therefore help economists understand why monetary policy differently affects short term rates (the instrument of monetary policy) and the long term rate (the rate that is of interest for investors and consumers). In fact, Wright (2011) has found a positive and strong relationship between long-term inflation uncertainty and bond term premia in a large cross-section of countries. The important policy implication of Wright’s (2011) findings is the possibility that eliminating long-run inflation uncertainty might facilitate the transmission of monetary policy to the economy. Also, D’Amico and Orphanides (2014) consider ex-ante measures of risk for inflation forecasting and Caporale et al. (2012) have shown that inflation uncertainty has decreased in the Euro area, possibly due to the fact that inflation decreased steadily since the beginning of the Euro.

Figure 11 depicts our measure of uncertainty (Panel A) and the decompositions (Panels B,C). Inflation uncertainty was high in the early 1980s, possibly due to oil price shocks, and decreased substantially afterwards; typically, it tends to be high around recessions. The behavior over time of uncertainty is very different from that of disagreement, which instead does not necessarily peaks around recession times. While the volatility component is pretty constant over time, the majority of the fluctuations in aggregate inflation uncertainty are associated with the bias component and the ex-post components; interestingly, ex-ante inflation uncertainty seems to have decreased monotonically since the early 1980s.

Our empirical results suggest that the most effective policies to decrease inflation uncertainty are those that influence ex-post uncertainty. In other words, policies should aim at ensuring that ex-post realizations of inflation are in line with the average expected inflation (for example, by minimizing shocks to inflation), not those that decrease the agents’ ex-ante uncertainty (i.e. not those that affect the agents’ expectation formation process), although the latter can also be effective.

INSERT FIGURE 11 HERE

9 Conclusion

This paper proposes an alternative measure of uncertainty based on survey density forecasts. The new measure has the advantage that it can be used to decompose uncertainty into components that can help researchers understand what existing uncertainty indices measure. In particular, our measure of uncertainty can be decomposed into aggregate uncertainty and disagreement, and aggregate uncertainty can itself be decomposed into Knightian uncertainty and realized risk. The latter inherently measure different things, have specific business cycle dynamics and different macroeconomic impact. Moreover, these sources of uncertainty resolve differently across prediction horizons.

Given that our proposed uncertainty index is an ex-post measure of uncertainty, we also decompose it into a component that only reflects ex-ante uncertainty, which we can relate to existing

measures of uncertainty based on the inter-quantile spread of the forecast distribution, and a component that measures ex-post uncertainty. Our analysis uncovers that some existing measures of uncertainty capture ex-ante uncertainty (such as existing measures of uncertainty based on policy uncertainty), while others capture ex-post uncertainty.

We also investigate the effects of the sources of uncertainty on the macroeconomy. We find that, while an increase in overall uncertainty has recessionary effects, the effects of the various components of uncertainty differ. For example, disagreement is only a small portion of the overall uncertainty, and may both underestimate and lag the actual degree of uncertainty in the economy; thus it may not be a timely measure of uncertainty. In addition, both realized risk and Knightian uncertainty were important components of uncertainty over the last three decades, although the former sharply decreased as soon as the financial recession of 2007-2008 ended while the latter remained high even after the end of the crisis. This suggests that the high overall uncertainty that persisted after the end of the latest recession was mostly due to agents' being unable to assign the correct probability to the economic outcomes and disagreeing on them, rather than because risk was high. Simulation results from a stylized macroeconomic model suggest that the behavior of uncertainty and its components is largely reconcilable with a macroeconomic model with ambiguity. Ambiguity can be a source of its own in increasing the overall level of uncertainty; alternatively, it can also act as an amplifying mechanism for the increase in the level of risk.

References

- Baker, S.R., N. Bloom, and S.J. Davis (2016), “Measuring Economic Policy Uncertainty,” *Quarterly Journal of Economics*, forthcoming.
- Baringhaus, L. and C. Franz (2004), “On a New Multivariate Two-Sample Test,” *Journal of Multivariate Analysis* 88, 190–206.
- Ball, L. (1990), “Why Does High Inflation Raise Inflation Uncertainty?” *Journal of Monetary Economics*, 371-388.
- Bloom, N. (2009), “The Impact of Uncertainty Shocks,” *Econometrica* 77(3), 623-685.
- Caporale, G., L. Onorante and P. Paesani (2012), “Inflation and Inflation Uncertainty in the Euro Area,” *Empirical Economics* 43(2), 597-615.
- D’Amico, S. and A. Orphanides (2014), “Uncertainty and Disagreement in Bond Risk Premia,” *Federal Reserve Bank of Chicago Working Paper* 2014-24.
- Dovern, J., U. Fritsche and J. Slacalek (2012), “Disagreement Among Forecasters in G7 Countries,” *Review of Economics and Statistics* 94(4), 1081-1096.
- Ferrara, L. and P. Guérin (2015), “What Are The Macroeconomic Effects of High-Frequency Uncertainty Shocks?,” *EconomiX Working Papers* 2015-12, University of Paris West - Nanterre la Défense, EconomiX.
- Giannone, D., L. Reichlin and D. Small (2008), “Nowcasting: The Real-time Informational Content of Macroeconomic Data,” *Journal of Monetary Economics*, 55(4), 665-676.
- Giordani, P. and P. Soderlind (2003), “Inflation Forecast Uncertainty,” *European Economic Review* 47, 1037-1059.
- Gneiting, T. and A.E. Raftery (2007), “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association* 102 (477), 359-378.
- Gurkaynak, R. and J.H. Wright (2012), “Macroeconomics and the Term Structure,” *Journal of Economic Literature* 50(2), 331-67.
- Hersbach, H. (2000), “Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems,” *Weather and Forecasting* 15, 559-570.
- Ilut, C.L. and M. Schneider (2014), “Ambiguous Business Cycles,” *American Economic Review* 104(8), 2368-99.
- Jurado, K., S. Ludvigson and S. Ng (2015), “Measuring Uncertainty,” *American Economic Review* 105 (3), 1177-1216.
- Knight, F.H. (1921), *Uncertainty and Profit*. Boston: Houghton Mifflin.
- Knueppel, M. and A. Vladu (2016), “Approximating Fixed-Horizon Forecasts Using Fixed-Event Forecasts,” *mimeo*.
- Lahiri, K. and X. Sheng (2010), “Measuring Forecast Uncertainty by Disagreement: The Missing Link,” *Journal of Applied Econometrics* 25, 514-538.

McCracken, M.W., M.T. Owyang and T. Sekhposyan (2015), “Real-time Forecasting with a Large, Mixed Frequency, Bayesian VAR,” *St. Louis Fed Working Paper* 2015-030A.

Patton A. and A. Timmermann (2010), “Why Do Forecasters Disagree? Lessons From the Term Structure of Cross-Sectional Dispersion,” *Journal of Monetary Economics* 57(7), 803-820.

Rossi B. and T. Sekhposyan (2015), “Macroeconomic Uncertainty Indices Based on Nowcast and Forecast Error Distributions,” *American Economic Review Papers & Proceedings* 105(5), 650-55.

Rossi B. and T. Sekhposyan (2016), “A Macroeconomic Uncertainty Index for the Euro Area,” *mimeo*.

Scotti, C. (2016), “Surprise and Uncertainty Indexes: Real-time Aggregation of Real-Activity Macro Surprises,” *Journal of Monetary Economics* 82, 1-19.

Sill, K. (2014), “Forecast Disagreement in the Survey of Professional Forecasters,” *Philadelphia Fed Business Review* Q2-2014, 15-24.

Stock J.H. and M.W. Watson (1999), “Business Cycle Fluctuations in U.S. Macroeconomic Time Series”. In: J. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, North Holland, Vol. 1A, 3-64.

Stock J.H. and M.W. Watson (2012), “Disentangling the Channels of the 2007-2009 Recession,” *Brookings Papers on Economic Activity* Spring, 81-135.

Wright. J. H. (2011), “Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset,” *American Economic Review* 101 (4), 1514-34.

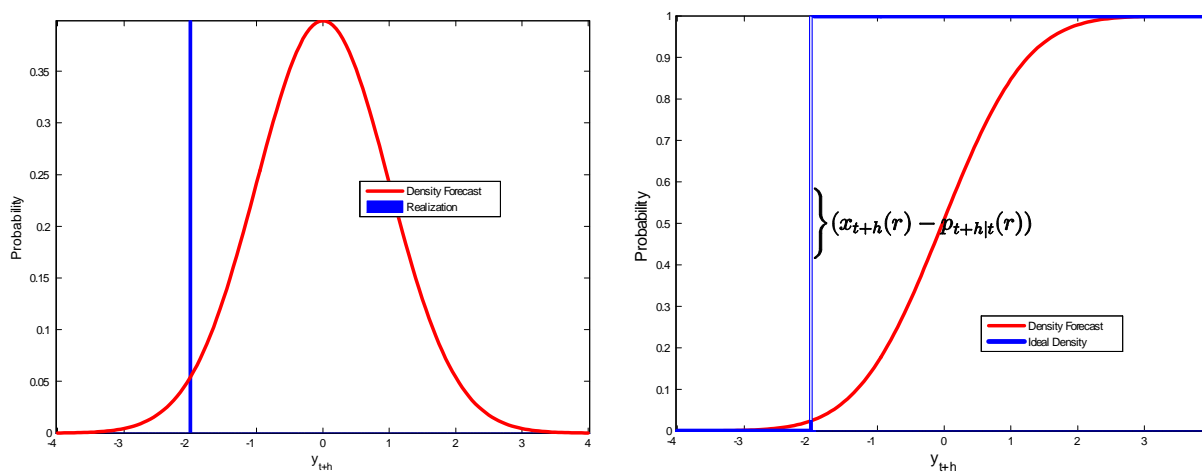
Zarnowitz, V. and L.A. Lambros (1987), “Consensus and Uncertainty in Economic Prediction,” *Journal of Political Economy* 95(3), 591-621.

10 Tables and Figures

Table 1: Baseline Parameter Values		
ρ_z	0.625	Estimated from an AR(1) model fitted to GDP growth
ρ_a	0.887	Ilut and Schneider's (2014) mode
n	0.995	Ilut and Schneider's (2014) mode
σ_u	0.780	Estimated from an AR(1) model fitted to GDP growth (σ_z)
σ_μ	0.500	Arbitrary, as the parameter is not separately identified
σ_n	0.134	Ilut and Schneider (2014) mode

Note: The table reports the parameter values used in the benchmark simulations.

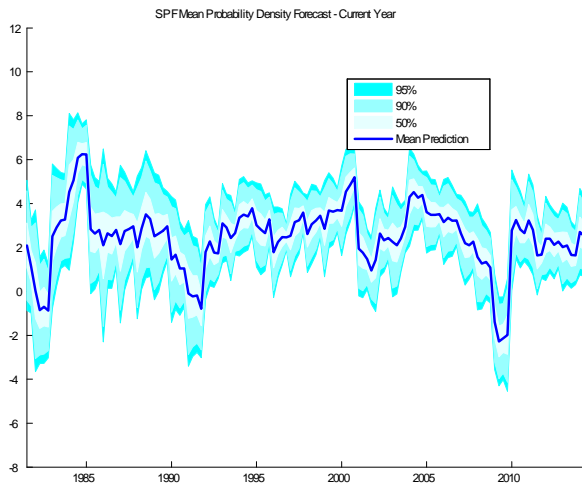
Figure 1: Brier Score Illustration



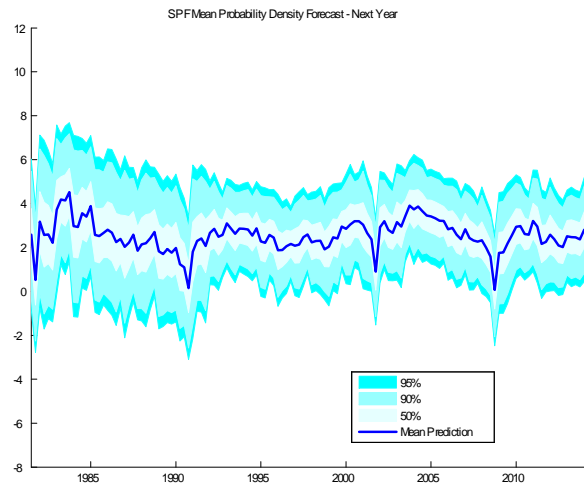
Note: The figure on the left shows the pdf of the predicted distribution together with the realization, $y_{t+h} = -2$. On the right we have the CDFs of the predicted and ideal distribution for the realization. The area between them (denoted by the hollow vertical bar) highlights the distance between the two.

Figure 2. The Survey of Professional Forecasters Data: GDP Growth

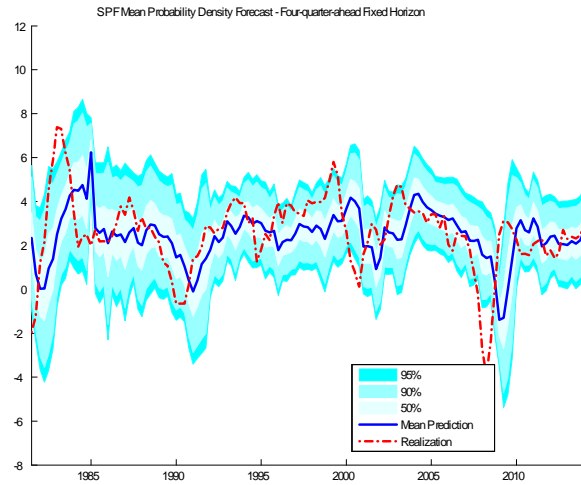
Panel A: Current Year Forecasts



Panel B: Next Year Forecasts



Panel C: Fixed-horizon Forecasts



Note: The figure shows the means and quantiles of the SPF's current year and next year predictive densities, as well as the constructed fixed horizon four-step-ahead predictive density. The four-step-ahead density is constructed from the SPF current year and next year density forecasts based on eq. (10). Panel C also shows the realized value of the GDP growth.

Figure 3: Decomposing Uncertainty in GDP

Panel A: Uncertainty, Aggregate Uncertainty and Disagreement

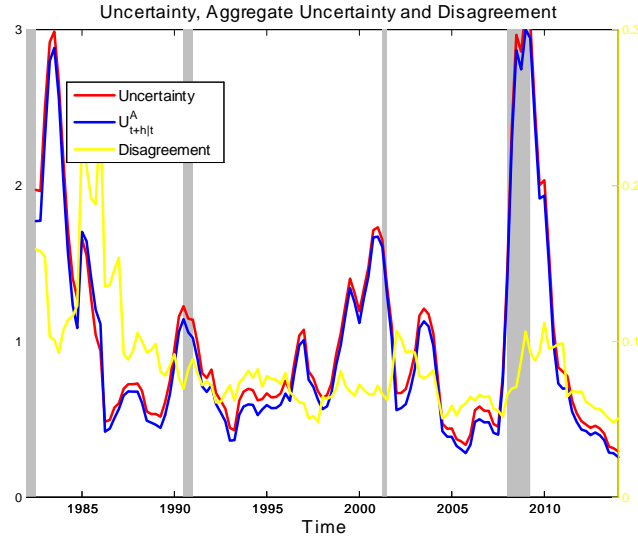
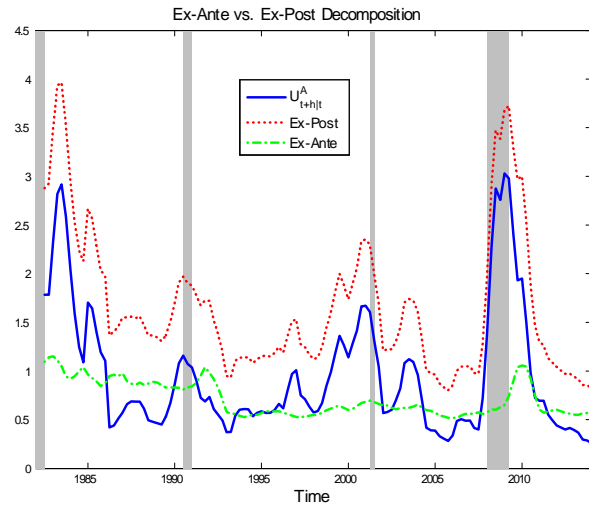
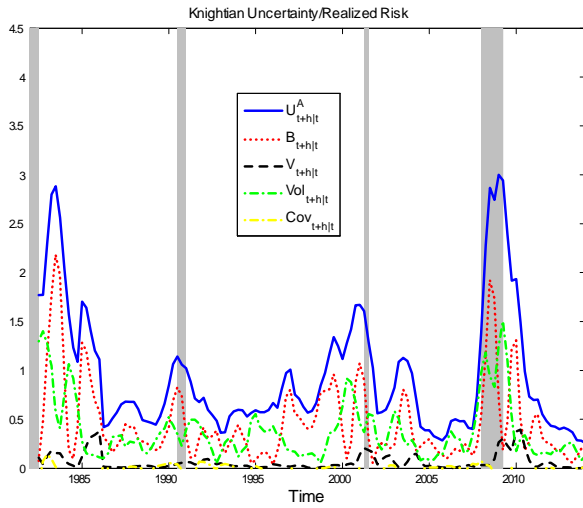


Figure 3: Decomposing Aggregate Uncertainty: GDP

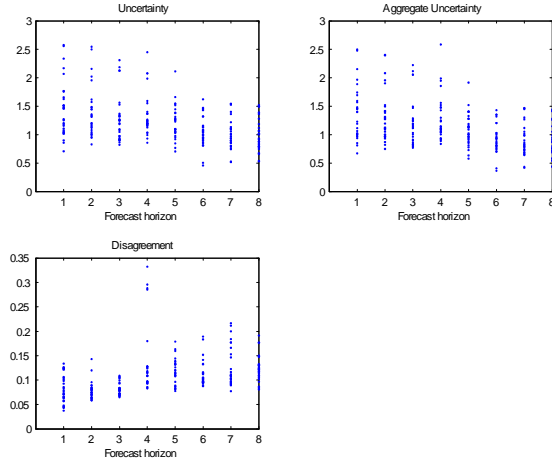
Panel B: Knightian Uncertainty Vs. Risk

Panel C: Ex-Ante Vs. Ex-Post

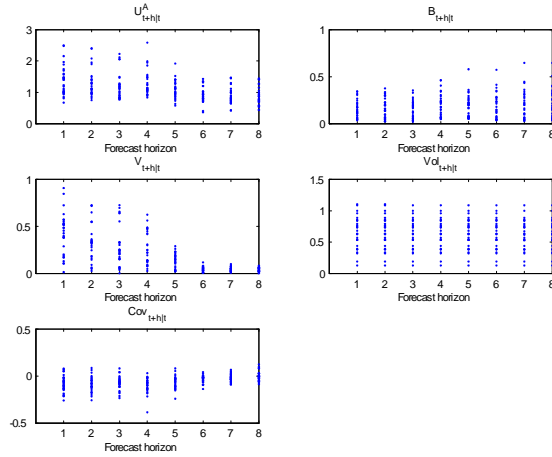


Note: Panel A of Figure 3 depicts the evolution of uncertainty, aggregate uncertainty and disagreement (eq. 5) over time. Panels B and C show the evolution of the components of aggregate uncertainty based on eq. (7) and eq. (8), respectively.

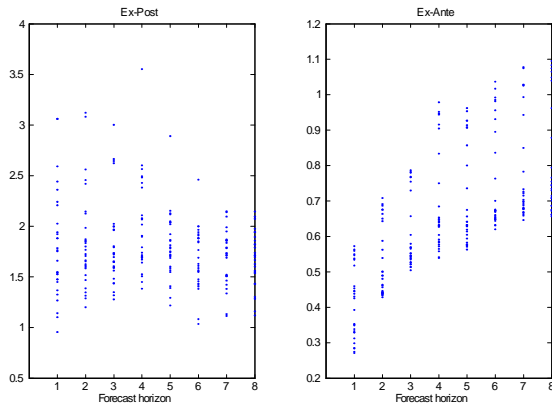
Figure 4: Decomposition of Uncertainty Across Horizons
Panel A: Uncertainty, Aggregate Uncertainty, Disagreement



Panel B: Knightian Uncertainty Vs. Risk

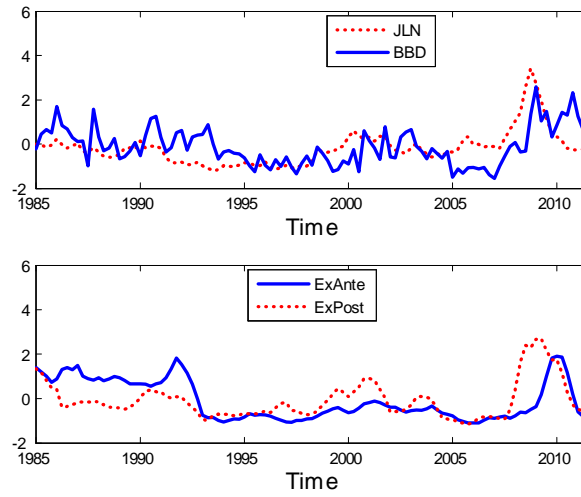


Panel C: Ex-Ante Vs. Ex-Post



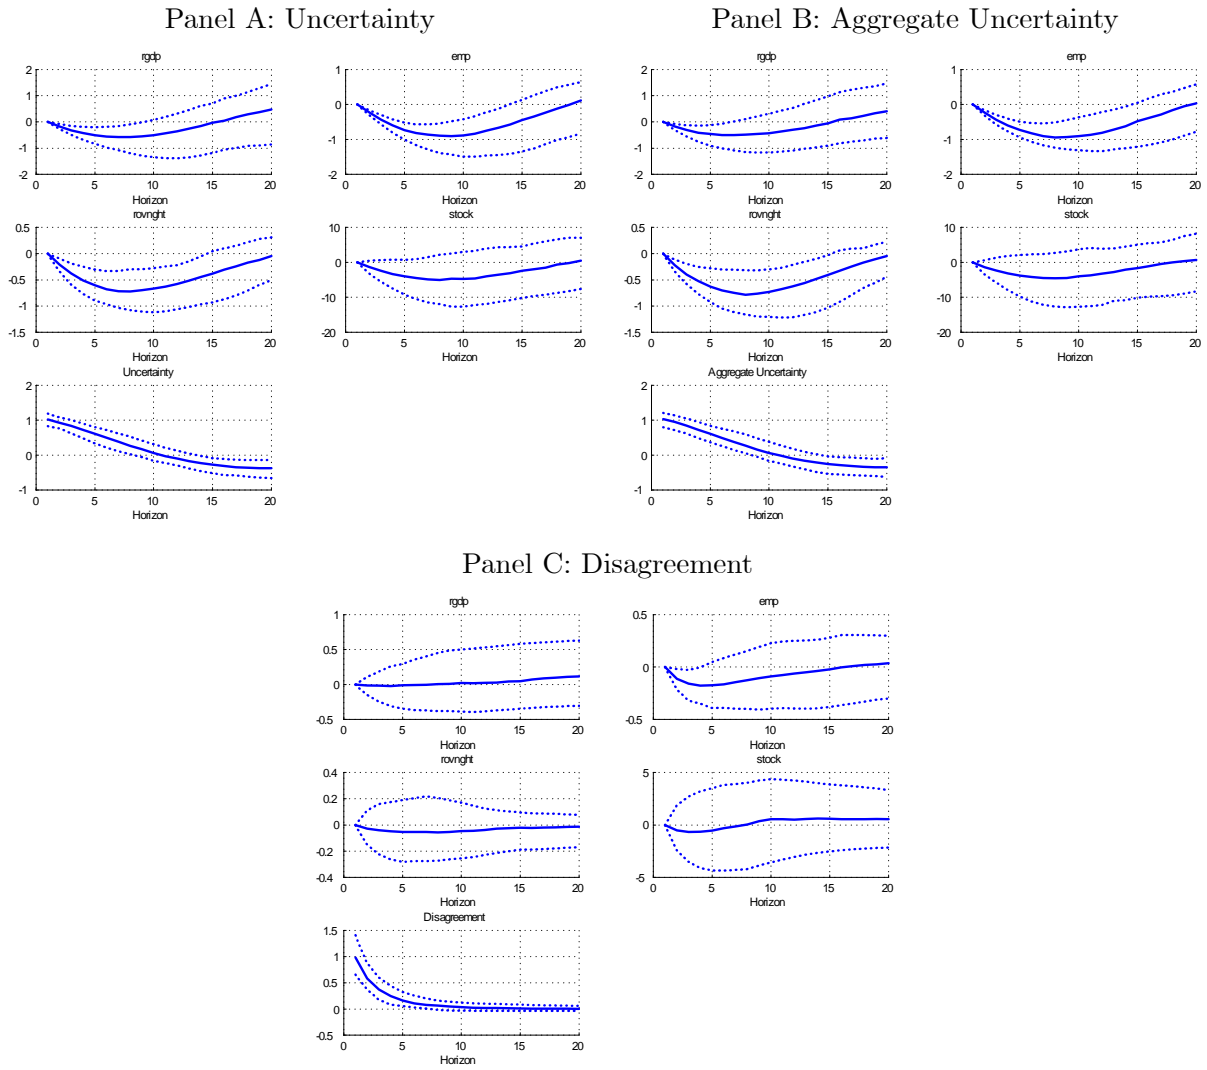
Note: Panel A shows uncertainty, aggregate uncertainty and disagreement over time. Panels B and C show the components in decompositions in eqs. (7) and (8), respectively.

Figure 5. Comparison of Uncertainty Measures



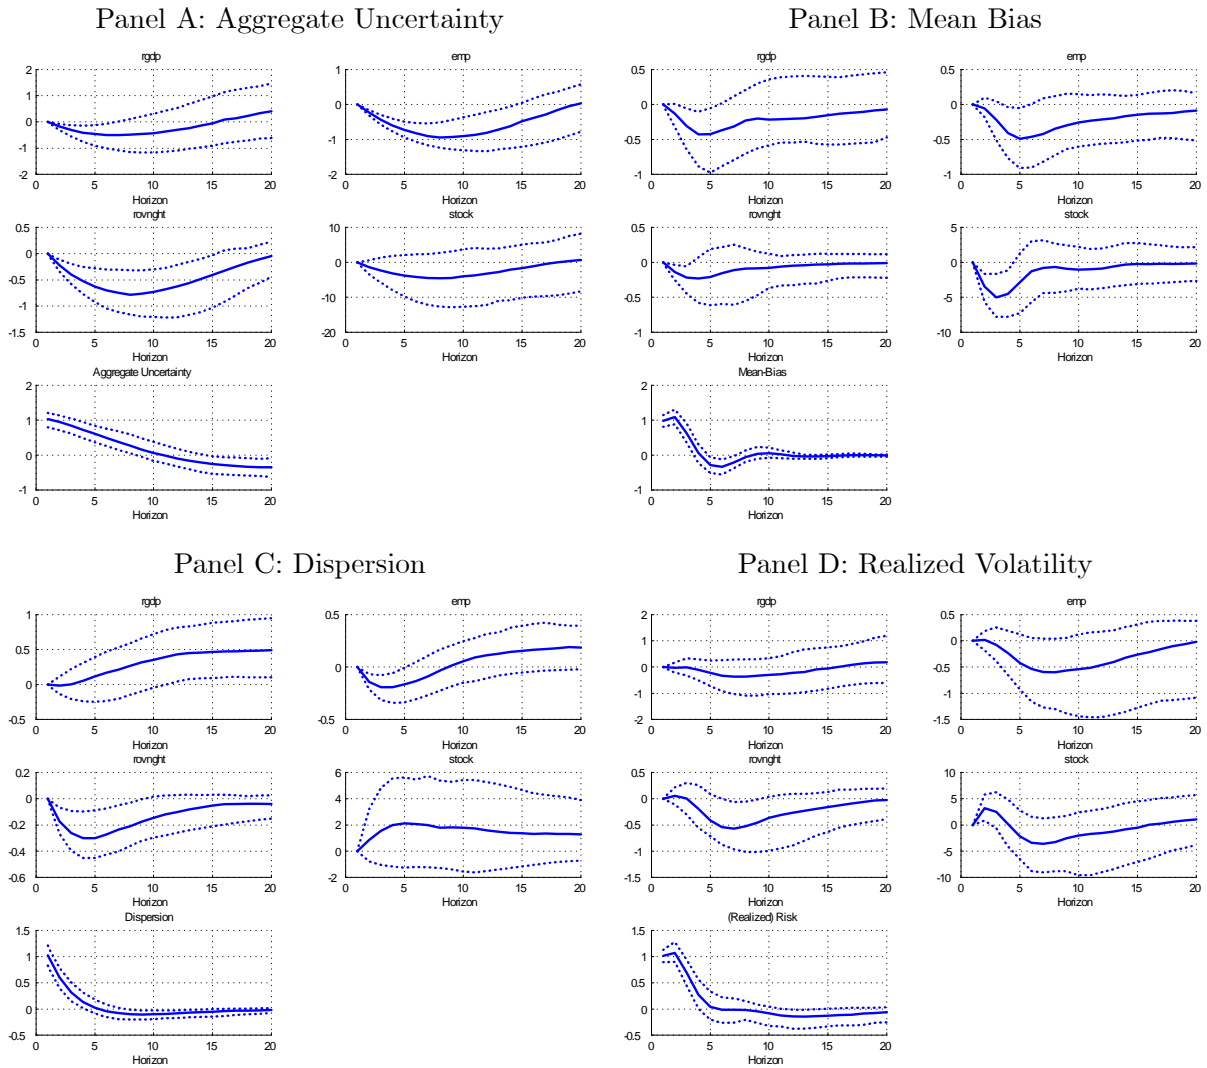
Note: The figure compares the Jurado, Ludvigson and Ng (2015) and Baker, Bloom and Davis (2016) uncertainty indices (top panel) with the ex-ante and ex-post components of our uncertainty measure, eq. (8), depicted in the bottom panel.

Figure 6. The Effects of Uncertainty on the Economy



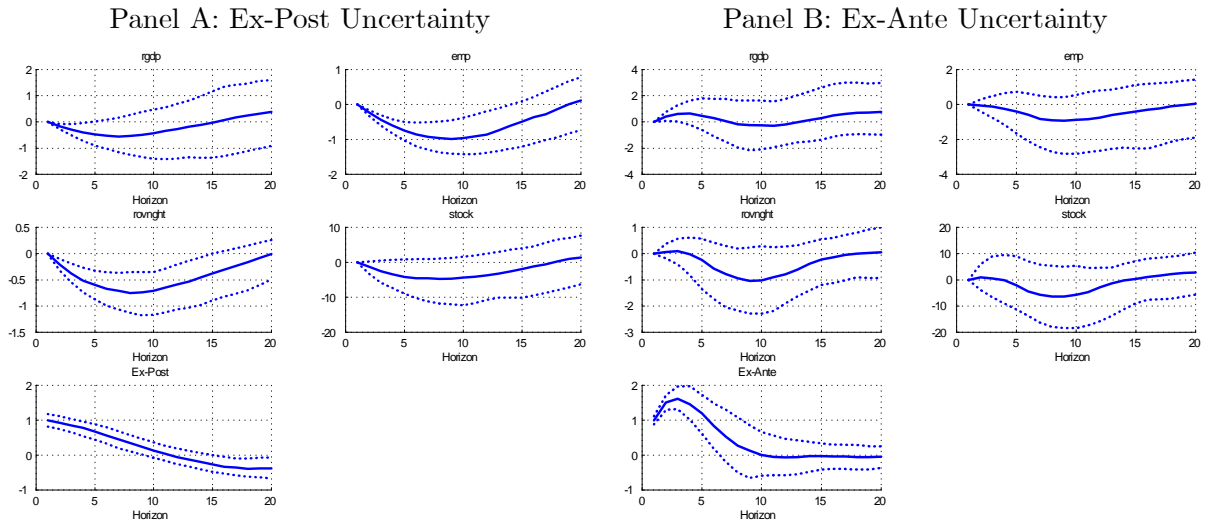
Note: The figure shows the impulse responses of uncertainty, aggregate uncertainty and disagreement shocks. The components are calculated as in eq. (5) All uncertainty measures have been standardized.

Figure 7. The Effects of Uncertainty on the Economy



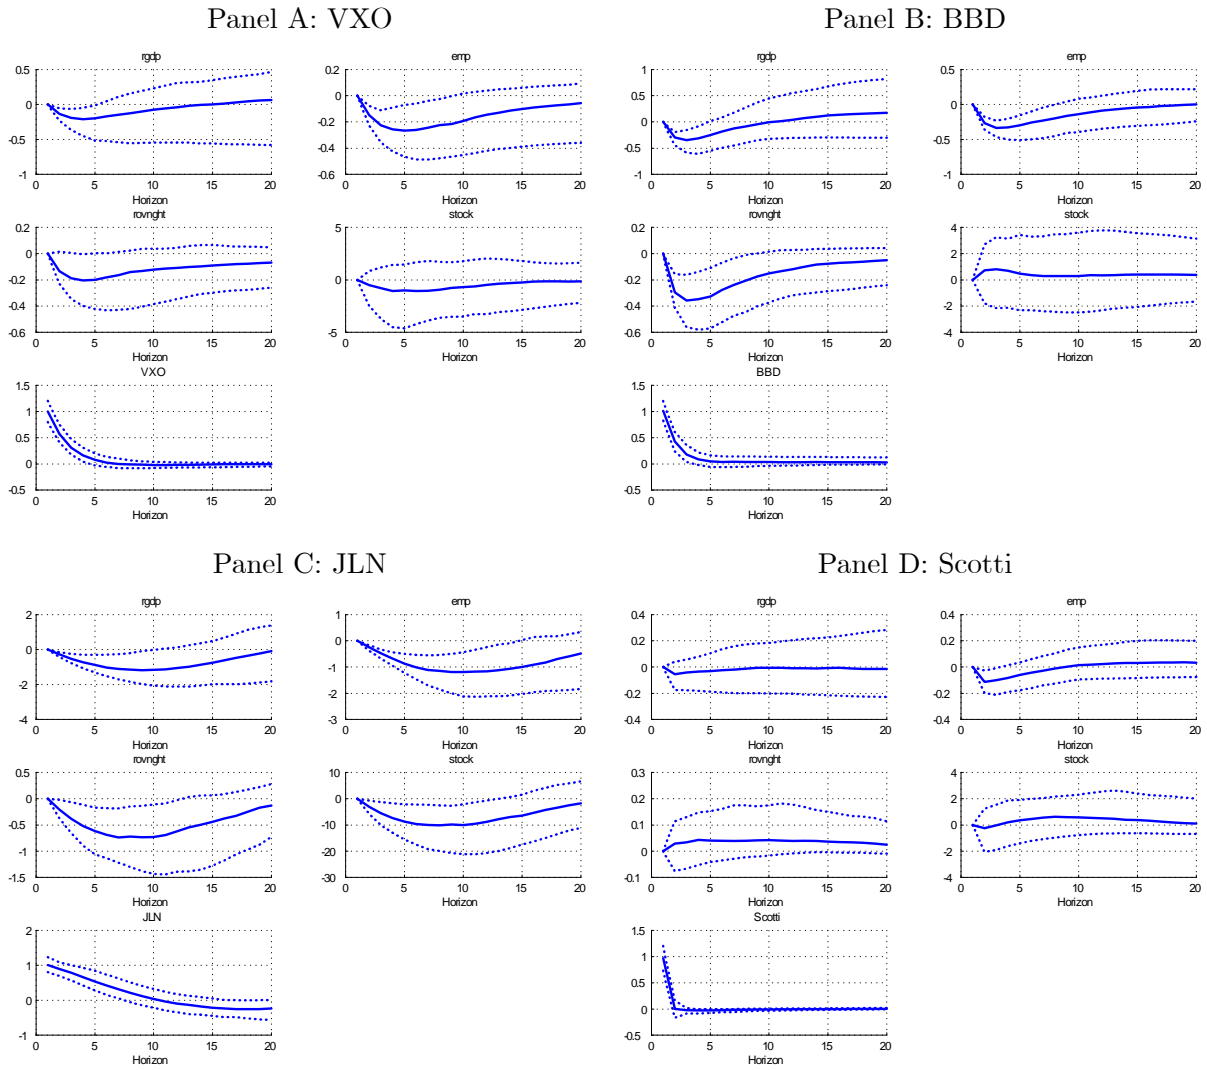
Note: The figure shows the impulse responses of the aggregate uncertainty and its components (mean bias, dispersion and realized risk measures) based on eq. (7). All uncertainty measures have been standardized.

Figure 8. The Effects of Uncertainty on the Economy



Note: The figure shows the impulse responses to a one standard deviation shock in ex-ante and ex-post measures of uncertainty based on eq. (8). All uncertainty measures have been standardized.

Figure 9. The Effects of Uncertainty on the Economy - Alternative Measures



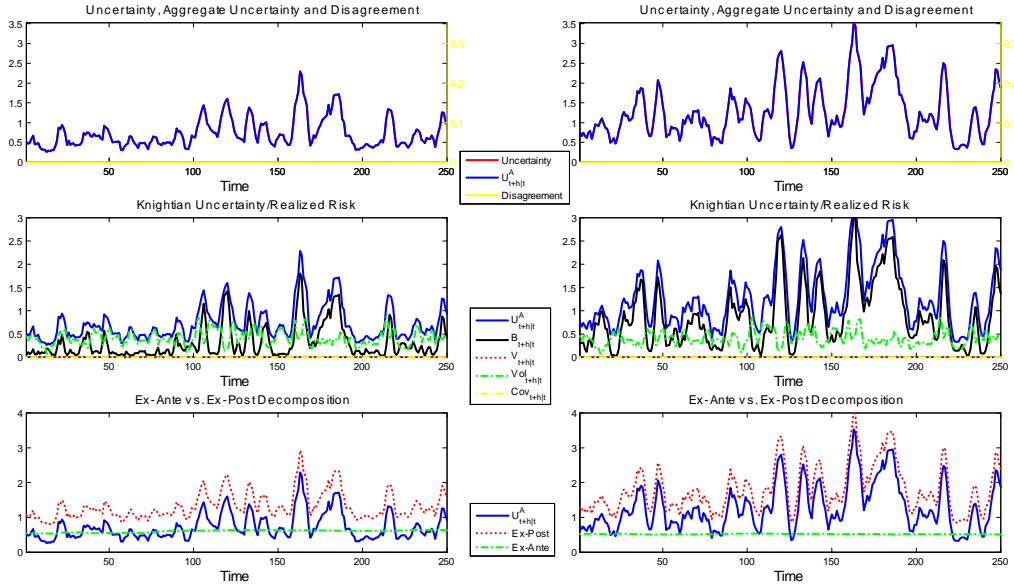
Note: The figure shows the impulse responses for the following uncertainty measures: VXO, JLN, BBD and Scotti. All uncertainty measures have been standardized.

Figure 10: Simulation Results

Panel A: Scenario 1 - Changing the Level of Ambiguity

$n = 0.2$

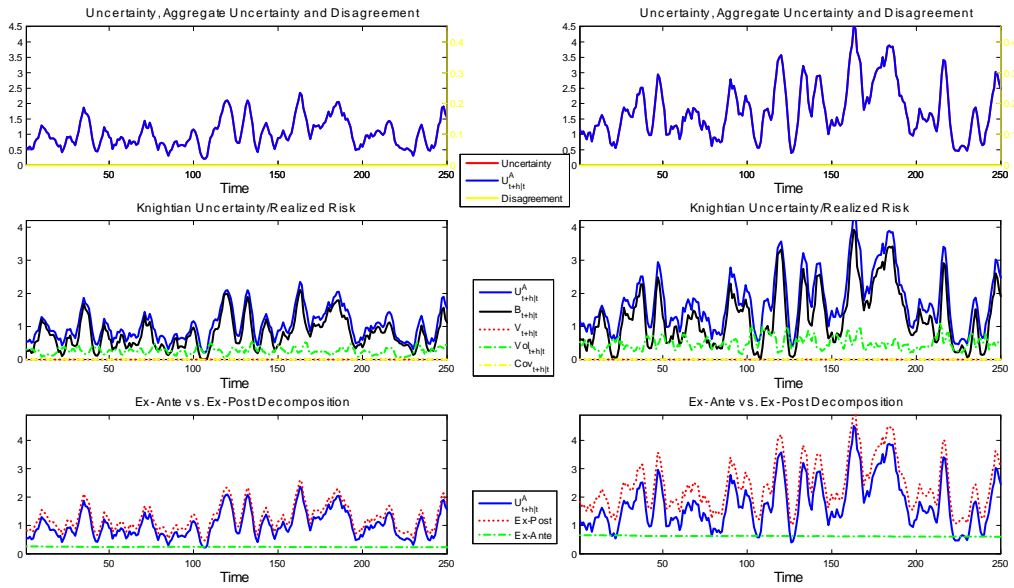
$n = 0.8$



Panel B: Scenario 2 - Changing the Level of Risk in the Model with Ambiguity

$\sigma_u = 0.3$

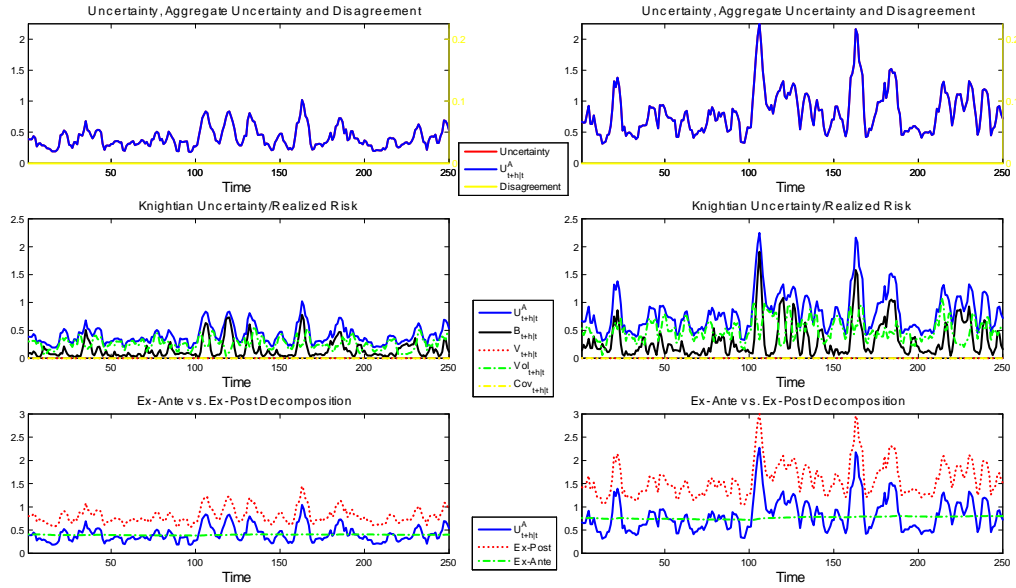
$\sigma_u = 1$



Panel C: Scenario 3 - Changing the Level of Risk in the Model with no Ambiguity

$\sigma_u = 0.3$

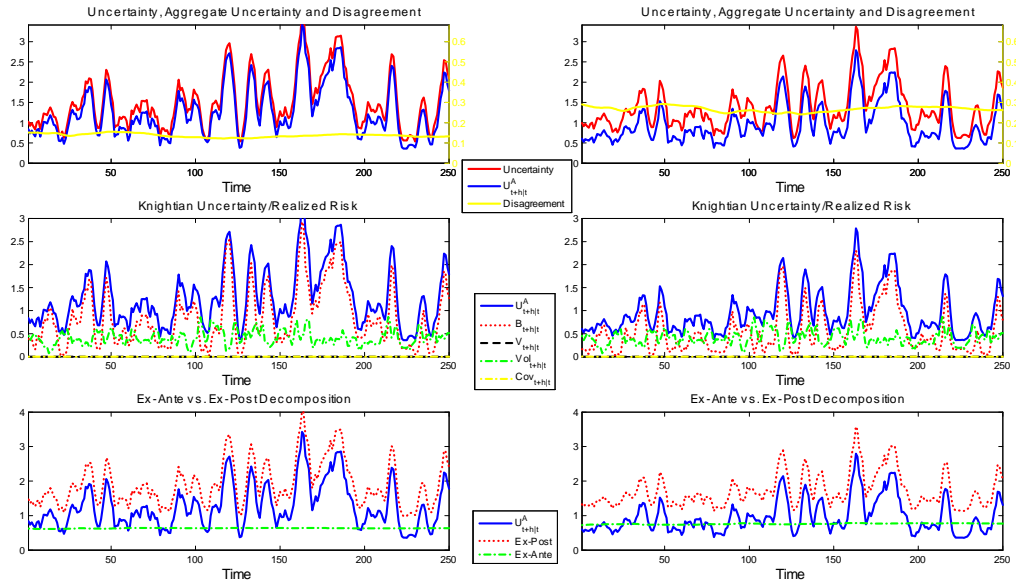
$\sigma_u = 1$



Panel D: Scenario 4 - Increasing Cross-Sectional Dispersion in Ambiguity

$\sigma_{n,I} = 0.5$

$\sigma_{n,I} = 1$



Note: The figure shows the simulation results described in more detail in Section 7; Panels A, B, C and D corresponding respectively to Scenarios 1, 2, 3 and 4, respectively.

Figure 11: Decomposing Inflation Uncertainty

Panel A: Uncertainty, Aggregate Uncertainty and Disagreement

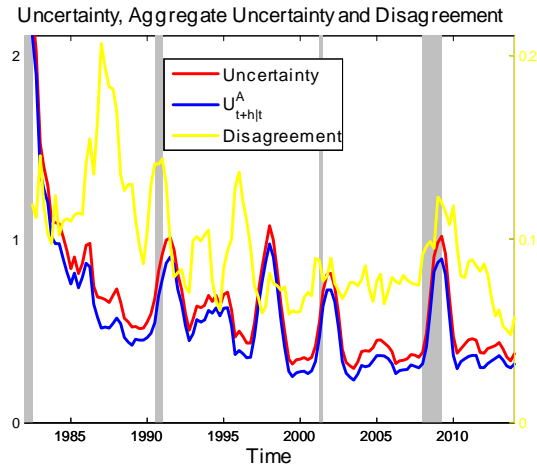
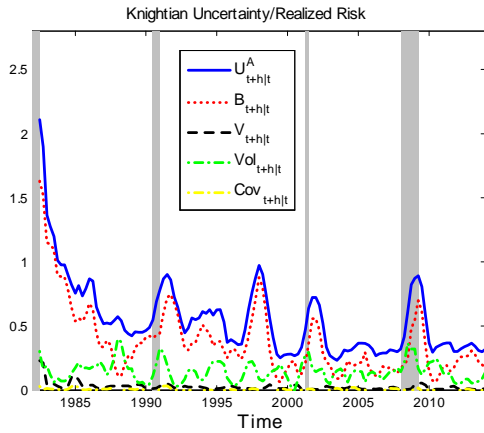
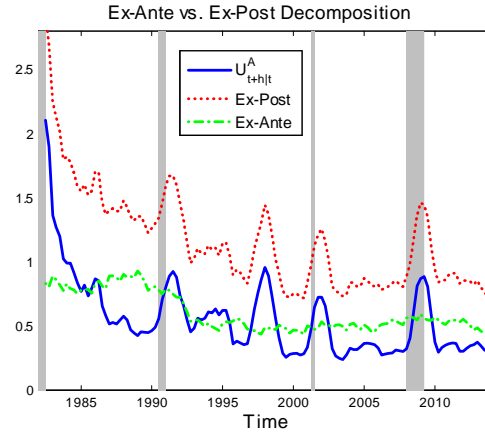


Figure 11: Decomposing Aggregate Inflation Uncertainty

Panel B: Knightian Uncertainty Vs. Risk



Panel C: Ex-Ante Vs. Ex-Post



Note: Panel A depicts the evolution of uncertainty, aggregate uncertainty, as well as disagreement (eq. 5) over time. Panels B and C show the evolution of the components of uncertainty based on eq. (7) and eq. (8), respectively.

Appendix A. Proofs

The appendix provides the proofs for the results in the paper. For simplicity, we write the proof for the unconditional expectation.

Proof of Eq. (4).

$$\begin{aligned}
& E \left[\frac{1}{N} \sum_{s=1}^N [x_{t+h}(r) - p_{s,t+h|t}(r)]^2 \right] \\
&= E \left[\frac{1}{N} \sum_{s=1}^N [x_{t+h}(r) - p_{t+h|t}(r) + p_{t+h|t}(r) - p_{s,t+h|t}(r)]^2 \right] \\
&= E \left(\frac{1}{N} \sum_{s=1}^N \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 + 2(x_{t+h}(r) - p_{t+h|t}(r))(p_{t+h|t}(r) - p_{s,t+h|t}(r)) \right] \right) \\
&+ E \left(\frac{1}{N} \sum_{s=1}^N \left[(p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 \right] \right) \\
&= E \left[\frac{1}{N} \sum_{s=1}^N (x_{t+h}(r) - p_{t+h|t}(r))^2 \right] + 2E \left[\frac{1}{N} \sum_{s=1}^N (x_{t+h}(r) - p_{t+h|t}(r))(p_{t+h|t}(r) - p_{s,t+h|t}(r)) \right] \\
&+ E \left[\frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 \right] \\
&= E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 \right] + 2E \left[(x_{t+h}(r) - p_{t+h|t}(r)) \frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r)) \right] \\
&+ E \left[\frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 \right] \\
&= E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 \right] + 2E \left[(x_{t+h}(r) - p_{t+h|t}(r)) \left(p_{t+h|t}(r) - \frac{1}{N} \sum_{s=1}^N p_{s,t+h|t}(r) \right) \right] \\
&+ E \left[\frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 \right] \\
&= E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 \right] + 0 + E \left[\frac{1}{N} \sum_{s=1}^N (p_{t+h|t}(r) - p_{s,t+h|t}(r))^2 \right].
\end{aligned}$$

■

Proof of Eq. (6).

$$\begin{aligned}
u_{t+h}^A(r) &\equiv E \left[(x_{t+h}(r) - p_{t+h|t}(r))^2 \right] = E \left[(x_{t+h}(r) - E(x_{t+h}(r)) + E(x_{t+h}(r)) - p_{t+h|t}(r))^2 \right] \\
&= E \left[(x_{t+h}(r) - E(x_{t+h}(r)))^2 \right] + E \left[(E(x_{t+h}(r)) - p_{t+h|t}(r))^2 \right] \\
&+ E \left[2(x_{t+h}(r) - E(x_{t+h}(r)))(E(x_{t+h}(r)) - p_{t+h|t}(r)) \right] \\
&= E \left([p_{t+h|t}(r) - E(x_{t+h}(r))]^2 \right) + V(x_{t+h}(r)) - 2Cov(x_{t+h}(r) p_{t+h|t}(r)),
\end{aligned}$$

where the last line follows from the fact that

$$\begin{aligned}
& E \left[2 \left(x_{t+h}(r) - E(x_{t+h}(r)) \right) \left(E(x_{t+h}(r)) - p_{t+h|t}(r) \right) \right] \\
&= 2E \left[\left(x_{t+h}(r) - E(x_{t+h}(r)) \right) \left(E(x_{t+h}(r)) - p_{t+h|t}(r) \right) \right] \\
&= 2E \left[x_{t+h}(r) E(x_{t+h}(r)) - x_{t+h}(r) p_{t+h|t}(r) - E(x_{t+h}(r))^2 + p_{t+h|t}(r) E x_{t+h}(r) \right] \\
&= 2E \left[p_{t+h|t}(r) E x_{t+h}(r) - x_{t+h}(r) p_{t+h|t}(r) \right] = 2 \left[E p_{t+h|t}(r) E x_{t+h}(r) - E(x_{t+h}(r) p_{t+h|t}(r)) \right] \\
&= -2Cov(x_{t+h}(r) p_{t+h|t}(r)).
\end{aligned}$$

Furthermore, note that $E \left([p_{t+h|t}(r) - E(x_{t+h}(r))]^2 \right) = \left([E(p_{t+h|t}(r)) - E(x_{t+h}(r))]^2 \right) + V(p_{t+h|t}(r))$.

$$\begin{aligned}
\text{In more detail, } E \left([p_{t+h|t}(r) - E(x_{t+h}(r))]^2 \right) &= E \left([p_{t+h|t}(r) - E(p_{t+h|t}(r)) + E(p_{t+h|t}(r)) - E(x_{t+h}(r))]^2 \right) \\
&= Var(p_{t+h|t}(r)) + E(E(p_{t+h|t}(r)) - E(x_{t+h}(r)))^2 \\
&\quad + 2E \left[(p_{t+h|t}(r) - E(p_{t+h|t}(r))) (E(p_{t+h|t}(r)) - E(x_{t+h}(r))) \right] \\
&= Var(p_{t+h|t}(r)) + E(E(p_{t+h|t}(r))). \quad \blacksquare
\end{aligned}$$

Proof of Eq. (8). Our measure of uncertainty is the negative of the expectation of the CRPS (Gneiting and Raftery, 2007). Note that $CRPS(F, y_{t+h}) = -\int_{-\infty}^{\infty} (F(r) - 1\{y_{t+h} < r\})^2 dr = -\bar{U}_{t+h}^A$, where $F(r)$ is the aggregate predictive distribution. Let $G(r)$ denote the ideal distribution, i.e. $G(r) = 1\{y_{t+h} < r\}$. The distributions are conditional distributions, but for notational simplicity we omit the conditionality. Then by Lemma 2.2 of Baringhaus and Franz (2004), we have $\bar{U}_{t+h}^A = \int_{-\infty}^{\infty} (F(r) - 1\{y_{t+h} < r\})^2 dy = E|Y_{1,t+h} - y_{1,t+h}| - \frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}| - \frac{1}{2}E|y_{1,t+h} - y_{2,t+h}|$, where $Y_{1,t+h}$ and $Y_{2,t+h}$ are i.i.d draws from F , while $y_{1,t+h}$ and $y_{2,t+h}$ are i.i.d. draws from $G(r)$, and both of these variables have finite expectations. Given Lemma 2.1 of Baringhaus and Franz (2004), $E|y_{1,t+h} - Y_{1,t+h}| = \int_{-\infty}^{\infty} F(r)(1-G(r))dr + \int_{-\infty}^{\infty} G(r)(1-F(r))dr$. Now, for $y_{1,t+h}$ and $y_{2,t+h}$, we have $E|y_{1,t+h} - y_{2,t+h}| = 2 \int_{-\infty}^{\infty} G(r)(1-G(r))dy = 2 \int_{-\infty}^{\infty} 1\{y_{t+h} < r\}(1 - 1\{y_{t+h} < r\})dy = 0$, where the last equality follows from the fact that, for a particular value of r , either $1\{y_{t+h} < r\}$ or $1 - 1\{y_{t+h} < r\}$ will be zero, and, thus, the product will be zero. Therefore,

$$\bar{U}_{t+h}^A = \int_{-\infty}^{\infty} (F(r) - 1\{y_{t+h} < r\})^2 dr = E|Y_{1,t+h} - y_{1,t+h}| - \frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}|. \quad (16)$$

This means we can rewrite aggregate uncertainty as the sum of expected absolute distance measures of random variables coming from the predictive distribution, and that coming from the predictive distribution and the true distribution which generates the realization. If $F(r)$ is the Gaussian distribution, i.e. if $Y_{t+h} \sim iidN(\mu_{t+h}, \sigma_{t+h}^2)$, then by the property that the difference of the i.i.d normal random variables is distributed normally (in this case centered around zero with a variance of $2\sigma_{t+h}^2$), and the fact that the absolute value of a mean zero random Normal variable has a half-normal distribution with mean $\frac{2\sigma_{t+h}}{\sqrt{\pi}}$, we have

$$\frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}| = \frac{\sigma_{t+h}}{\sqrt{\pi}}. \quad (17)$$

To obtain $E|Y_{t+h} - y_{t+h}|$, we use the properties of Dirac delta function. We denote the PDF of y_{t+h} by a Dirac delta function $\delta(r - y_{t+h})$. From the properties of the Dirac function, $E(y_{t+h}) = y_{t+h}$ and $Var(y_{t+h}) = 0$. Then, $Y_{1,t+h} - y_{1,t+h} \sim N(\mu_{t+h} - y_{t+h}, \sigma_{t+h}^2)$. By property of the folded Normal distribution, we have:

$$E|Y_{t+h} - y_{t+h}| = \sigma_{t+h} 2\varphi\left(-\frac{\mu_{t+h} - y_{t+h}}{\sigma_{t+h}}\right) + (\mu_{t+h} - y_{1,t+h}) \left(1 - 2\Phi\left(-\frac{\mu_{t+h} - y_{t+h}}{\sigma_{t+h}}\right)\right). \quad (18)$$

Substituting (18) and (17) into (16), we get the result:

$$\bar{U}_{t+h}^A = \underbrace{\left[2\sigma_{t+h}\phi\left(\frac{y_{t+h} - \mu_{t+h}}{\sigma_{t+h}}\right) + (y_{t+h} - \mu_{t+h}) \left(2\Phi\left(\frac{y_{t+h} - \mu_{t+h}}{\sigma_{t+h}}\right) - 1\right)\right]}_{\text{“Ex-post”}} - \underbrace{\frac{\sigma_{t+h}}{\sqrt{\pi}}}_{\text{“Ex-ante”}}, \quad (19)$$

and $U_{t+h}^A = E\left(\bar{U}_{t+h}^A | \mathfrak{S}_{t-R}^t\right)$. ■

Appendix B. Data

As the main text indicates, the fixed-horizon forecasts are constructed as a weighted average of the current and next year forecasts. Figure A1 shows the number of forecasters that provided forecasts for both, current year and next year, as well as the number of forecasters that have provided forecasts for either one of the years, but not both. As the figure shows, the latter group is not large. By limiting our attention to forecasters that provide forecasts for both years we lose 10% of the total number of observations. The maximum per period loss amounts to 30% of forecaster observations. In our sample we have 239 unique forecasters. Out of those, 108 have been providing forecasts more than twelve times. The sample has 31 forecasters that have provided density forecast for 8 or more but less than 12 times, while 37 of them provided forecasts for 4 times and more, but less than 8 times. Thus, the majority of the forecasters in our sample are repeated forecasters, which increases the confidence that our results are not driven by outliers.

D’Amico and Orphanides (2014) highlight the role of approximations in individual predictive distributions. The idea is that many forecasters tend to put a lot of weight on a few bins and zeros on other bins. D’Amico and Orphanides (2014) argue that this could be forecasters’ true perceived uncertainty. However, as they suggest, it is also possible that forecasters just use approximations and lump small probabilities into the adjacent bins. To shed some light on this issue, in Figure A2, Panel 1, we show the percentage of forecasters that put probabilities into one bin, two bins and three bins. In general, forecasters with all the probabilities on one bin and two bins are few. However, a non-negligible proportion of forecasters puts all the probabilities on three bins. The proportion of these forecasters is higher prior to 1992:I. This is explained by the structure of the bins at that point. In our sample period, the bin structure for GDP/GNP growth has changed three times. Between 1981:III and 1991:IV there were 6 bins covering [<-2 -2 to -0.1 0 to 1.9 2 to 3.9 4 to 5.9 6+], between 1992:I and 2009:I the bins were [<-2 -2 to -1.1 -1 to -0.1 0

to 0.9 1 to 1.9 2 to 2.9 3 to 3.9 4 to 4.9 5 to 5.9 6+], while since 2009:II the bins have been covering the following intervals [<-3 -3 to -2.1 -2 to -1.1 -1 to -0.1 0 to 0.9 1 to 1.9 2 to 2.9 3 to 3.9 4 to 4.9 5 to 5.9 $6+$]. Note that in the beginning of the sample the bins have been fairly wide, not giving forecasters opportunities to differentiate among bins. Given that we use a Gaussian approximation, in order to strive for accuracy we adjusted the bins in the period between 1981:III and 1991:IV. The modified grid doubles the number of bins in that period, splitting the original probabilities in each bin uniformly over the newly created ones. Effectively the grid structure in that period becomes the same as in the period between 1992:I and 2009:I. The summary of the number of forecasts with probabilities on one, two and three bins with the modified grid specification is provided in Figure A2, Panel B. The figure shows that, by construction, there are not many forecasts with probabilities on less than or equal to three bins left in the period prior to 1992:I. Moreover, we discarded densities that put all the probability mass on one bin in the calculations. The second source of approximation error that arises when working with SPF probability forecast histograms is the open ended nature of the first and last intervals. In practice, we close these intervals. We assume that the open intervals have the same length as the rest of the intervals in the respective grids. Panels C in Figure A2 shows the proportion of forecasters assigning a probability value on the leftmost and rightmost bins in the survey. On the one hand, these proportions are not negligible, and our choice of dealing with the leftmost and rightmost intervals might have some impact on the overall results. On the other hand, Panel D suggests that the probability value associated with these open intervals is small. Thus, closing the open intervals should not induce a large approximation error. Lastly, since we approximate the histograms with a Gaussian distribution, we use the mid-point approach: when fitting a Gaussian kernel we associate all the probability mass with the midpoint of the interval.

Figure A1: Forecasters with Current Year and Next Year Forecasts

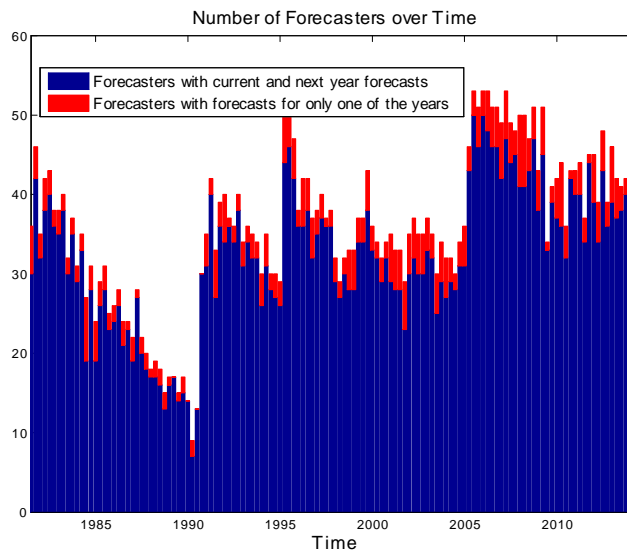
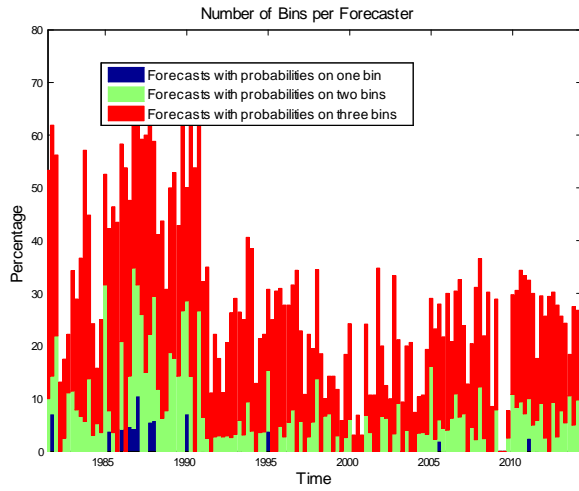
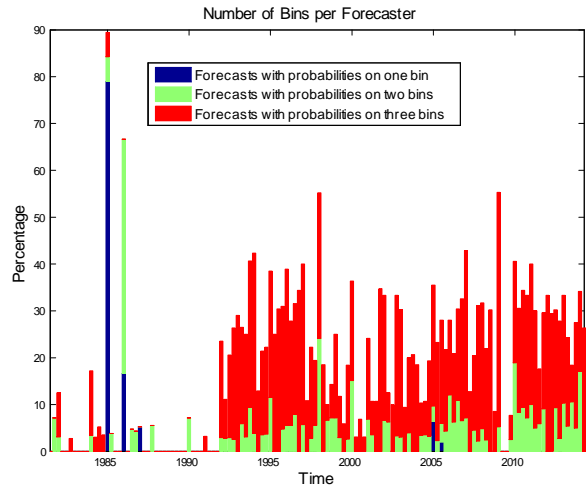


Figure A2: Bin Statistics

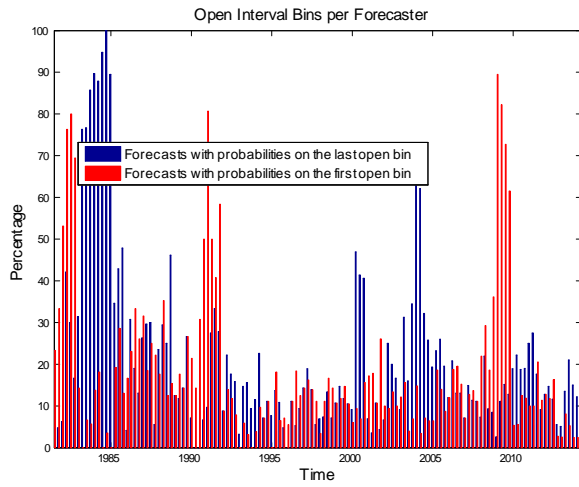
Panel A: Forecasts without Grid Adjustment



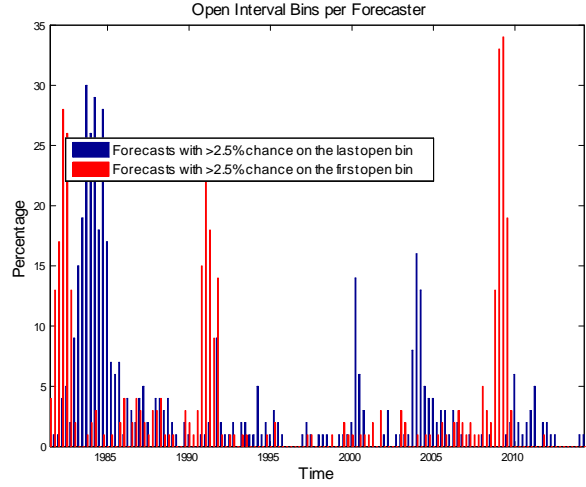
Panel B: Forecasts after Grid Adjustment



Panel C: Forecasts with Open Intervals



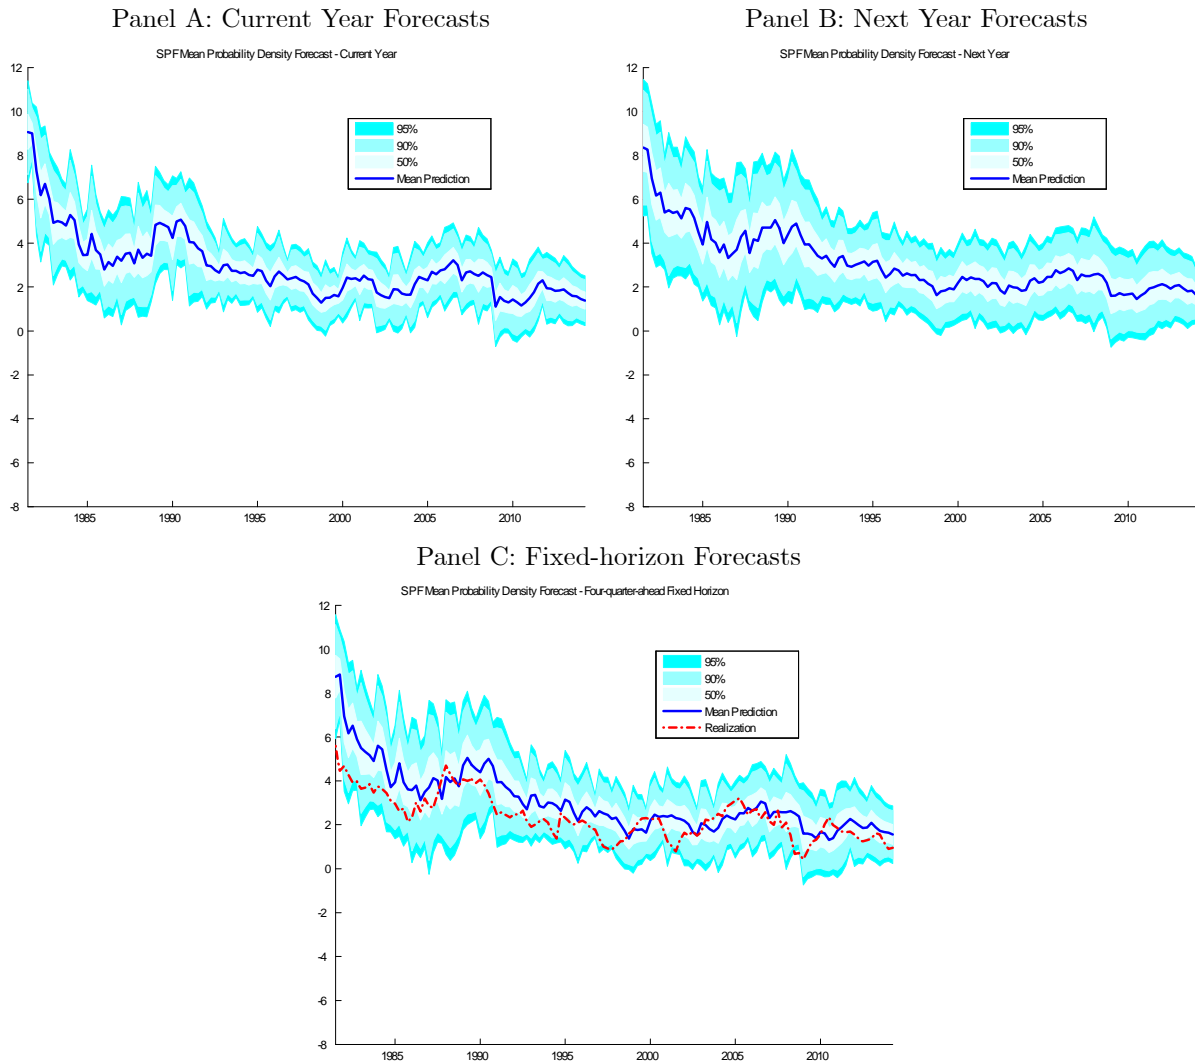
Panel D: Forecasts with >2.5% on Open Intervals



Not-for-Publication Appendix (For Online Publication Only)

A.1 Descriptive Analysis of Inflation Forecasts

Figure I. The Survey of Professional Forecasters Data: Inflation



Note: The figure shows the quantiles of the SPF four-step-ahead predictive density, its mean, as well as the realized value of inflation. The four-step-ahead density is constructed from SPF's current year and next year density forecasts based on eq. (10)

A.2 Reliability and Resolution Analysis

Note that an additional, interesting decomposition for $u_{t+h|t}^A(r)$ can be obtained following Murphy (1973):

$$u_{t+h|t}^A(r) \simeq REL_{t+h|t}(r) - RES_{t+h|t}(r) + V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) \quad (20)$$

where:

- $REL_{t+h|t}(r) \equiv E\left(\left[p_{t+h|t}(r) - E(x_{t+h}(r) | p_{t+h|t}(r))\right]^2 | \mathfrak{S}_{t-R}^t\right)$ measures the reliability of the forecast and scores the calibration of the forecast. A forecast is said “reliable” when the observed frequency is consistent with the probabilistic forecast made for a given event. For instance, forecasts that predict a probability of recession of 30 percent will be reliable if the economy effectively enters a recession 30 percent of the time every time such a forecast is made. Hence, reliability measures the unconditional (un)biasedness of the probabilistic forecasts. Because the term is expressed as a squared error, the smaller the calibration error, the better (i.e., the lower) the Brier score.
- $RES_{t+h|t}(r) \equiv E\left(\left[E(x_{t+h}(r) | p_{t+h|t}(r)) - E(x_{t+h}(r))\right]^2 | \mathfrak{S}_{t-R}^t\right)$ is the resolution, i.e. the average squared differential of the conditional and unconditional means of the observed outcomes. It captures the “decisiveness” of forecasts by comparing the forecast probability and the long-term average of the underlying process. The larger the term, the lower the Brier score.

As we show below, Eq. (20) holds up to an approximation error that involves within bin variation.

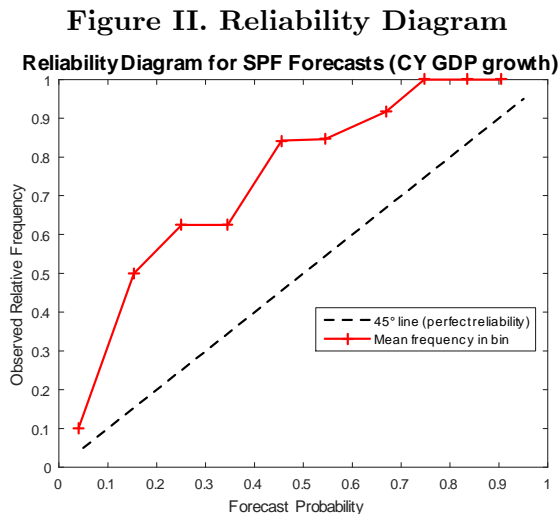
The decomposition can be estimated as follows.

Reliability is estimated as follows. For each t , determine which of the forecast bins $p_{t+h|t}(r)$ falls into. Let $\{p_{t+h|t}^{(k)}(r)\}$ be the collection of probabilities in the k -th bin and let $p_{t+h|t}^E(r)$ denote the unconditional expected value over the bin. We estimate $p_{t+h|t}^E(r)$ using a Uniform distribution over the bin, so that $p_{t+h|t}^E(r)$ is the midpoint of the bin.²¹ In addition, let the number of probabilities in each bin be n_k . Let \bar{x}_k be the average of the realizations conditional on the forecaster having made the probability forecast associated with the collection of probabilities in bin k , $\{p_{t+h|t}^{(k)}(r)\}$. Reliability is the average square calibration error, that is,

²¹In the 3-terms decomposition that we discuss here, we abstract from within bin variance and within bin covariance; thus, the unconditional expected value over the bin is indeed the midpoint of the bin and all forecasts in the bin are imposed to be equal to the midpoint (so their average is also the midpoint). We derive a 5-terms decomposition which includes within bin variance and within bin covariance (Stephenson, Coelho and Joliffe, 2008). In that case, the reliability will be calculated using the average forecast in the bin without imposing that all forecasts in the bin are equal. That is, $\bar{p}_{t+h|t}^{(k)}(r)$ (which is the average of the collection of probabilities in the k -th bin, $\{p_{t+h|t}^{(k)}(r)\}$), replaces $p_{t+h|t}^E(r)$ in eq. (21).

$$REL(r) = \frac{1}{T} \sum_{k=1}^K n_k \left(p_{t+h|t}^E(r) - \bar{x}_k(r) \right)^2. \quad (21)$$

Thus, reliability measures the squared deviation of the predicted probability from the observed outcome conditional probability of the event. This effectively tells the user how often (as a percentage) a forecast probability actually occurred. In theory, a perfect forecasting model will result in forecasts with a probability of $\alpha\%$ being consistent with the eventual outcome $\alpha\%$ of the time. Note that a forecast is reliable if the average square calibration error (REL) is small. Figure II provides intuition to understand reliability. The x-axis reports the forecast probability,²² while the y-axis reports the observed relative frequency. A reliable forecast would be the 45-degree line, where the observed frequency of realizations equals the forecast probability; the data clearly show departures from reliability in our sample.



Notes. The figure plots the reliability diagram for SPF forecasts of current year (CY) GDP growth.

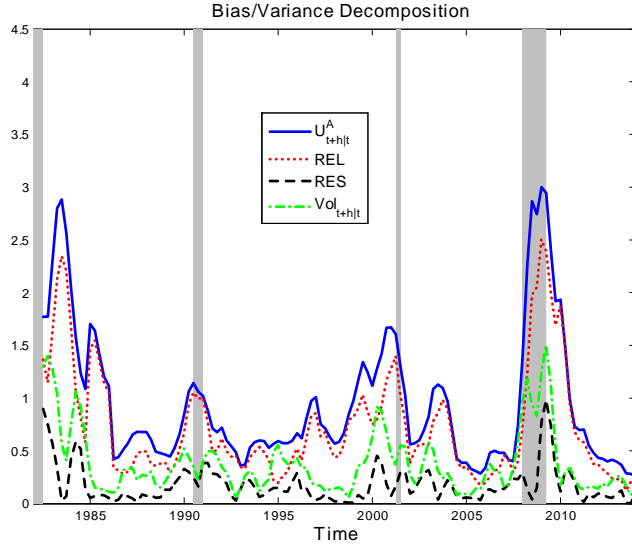
Resolution is the squared average difference between the conditional mean (given the forecast) and the unconditional mean: $RES(r) = \frac{1}{T} \sum_{k=1}^K n_k (\bar{x}_k(r) - \bar{x}(r))^2$. Note that good forecasts have high resolution.

Figure III shows the evolution of the components of the alternative decomposition over time.²³

²²The forecast probability is the mid-point of the bin in the forecast distribution.

²³Finally, note that the practical implementation of the Brier score involves “binning”. Binning smooths the data and makes them less noisy, as larger bins limit the “sparseness” problem (Stephenson et al., 2008). Some information is lost, however, by approximating continuous probability densities with a discrete number of bins.

Figure III. Aggregate Uncertainty, Reliability, Resolution and (Realized) Risk



Notes. The figure displays Aggregate Uncertainty, Reliability, Resolution and Realized Risk.

Proof of eq. (20). In practice, the Murphy decomposition requires partitioning the range of forecasts – essentially, the $[0,1]$ line – into K sub-segments. Let r be a number along the real line; let $\bar{p}^{(k)}$ denote the average probability in segment k ;²⁴ and let n_k denote the number of forecast probabilities that fall in the k -th sub-segment, for $k = 1, \dots, K$. Given all forecasts in the sample,

²⁴ Alternatively, one could consider $\bar{p}^{(k)}$ as the midpoint of the k -th segment

the Brier score can be broken down as follows:

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T [x_{t+h}(r) - p_{t+h|t}(r)]^2 &= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - p_{t+h|t}^{(j)}(r)]^2 \\
&= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r) + \bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r) + \bar{p}_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r)]^2 \\
&= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r)]^2 + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)]^2 \\
&\quad + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [\bar{p}_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r)]^2 \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r)] [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)] \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r)] [p_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r)] \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [p_{t+h|t}^{(j)}(r) - \bar{p}_{t+h|t}^{(k)}(r)] [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)] \\
&= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r)]^2 + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)]^2 \\
&\quad + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [\bar{p}_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r)]^2 \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r)] [p_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r)].
\end{aligned}$$

We can already recognize the reliability (REL) in the second term of this decomposition:

$$\begin{aligned}
REL(r) &= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)]^2 \\
&= \frac{1}{T} \sum_{k=1}^K n_k [\bar{x}_{t+h}^{(k)}(r) - \bar{p}_{t+h|t}^{(k)}(r)]^2.
\end{aligned} \tag{22}$$

The first term can be expressed as follows:

$$\begin{aligned}
\frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r) \right]^2 &= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[x_{t+h}^{(j)}(r) - \bar{x}(r) + \bar{x}(r) - \bar{x}_{t+h}^{(k)}(r) \right]^2 \\
&= \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[x_{t+h}^{(j)}(r) - \bar{x}(r) \right]^2 + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[\bar{x}(r) - \bar{x}_{t+h}^{(k)}(r) \right]^2 \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[x_{t+h}^{(j)}(r) - \bar{x}(r) \right] \left[\bar{x}(r) - \bar{x}_{t+h}^{(k)}(r) \right] \\
&= \frac{1}{T} \sum_{t=1}^T \left[x_{t+h}(r) - \bar{x}(r) \right]^2 - \frac{1}{T} \sum_{k=1}^K n_k \left[\bar{x}(r) - \bar{x}_{t+h}^{(k)}(r) \right]^2 \\
&\equiv V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) - RES(r).
\end{aligned}$$

Note that because the outcome variable x is binary, the uncertainty term can be expressed as $V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) = \bar{x}(r)(1 - \bar{x}(r))$. To summarize, we have decomposed the Brier score in the following way:

$$\begin{aligned}
\frac{1}{T} \sum_{t=1}^T \left[x_{t+h}(r) - p_{t+h|t}(r) \right]^2 &= V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) + REL(r) - RES(r) \\
&\quad + \frac{1}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[\bar{p}_{t+h|t}^{(k)}(r) - p_{t+h|t}^{(j)}(r) \right]^2 \\
&\quad + \frac{2}{T} \sum_{k=1}^K \sum_{j=1}^{n_k} \left[x_{t+h}^{(j)}(r) - \bar{x}_{t+h}^{(k)}(r) \right] \left[p_{t+h}^{(k)}(r) - p_{t+h|t}^{(j)}(r) \right].
\end{aligned}$$

The last two terms measure the variance of forecasts within the sub-segments and the co-movement between forecasts within a segment and their corresponding outcomes. The decomposition therefore writes:

$$\frac{1}{T} \sum_{t=1}^T \left[x_{t+h}(r) - p_{t+h|t}(r) \right]^2 = V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) + REL(r) - RES(r) + WSV(r) + WSC(r).$$

Remark that the last two terms equal zero when all forecasts within the same segment are assumed identical. Because $WSV(r)$ and $WSC(r)$ are quantitatively very small in the data, we will work under the simpler decomposition:

$$\frac{1}{T} \sum_{t=1}^T \left[x_{t+h}(r) - p_{t+h|t}(r) \right]^2 \simeq V(x_{t+h}(r) | \mathfrak{S}_{t-R}^t) + REL(r) - RES(r),$$

as per the definitions that we have written. ■

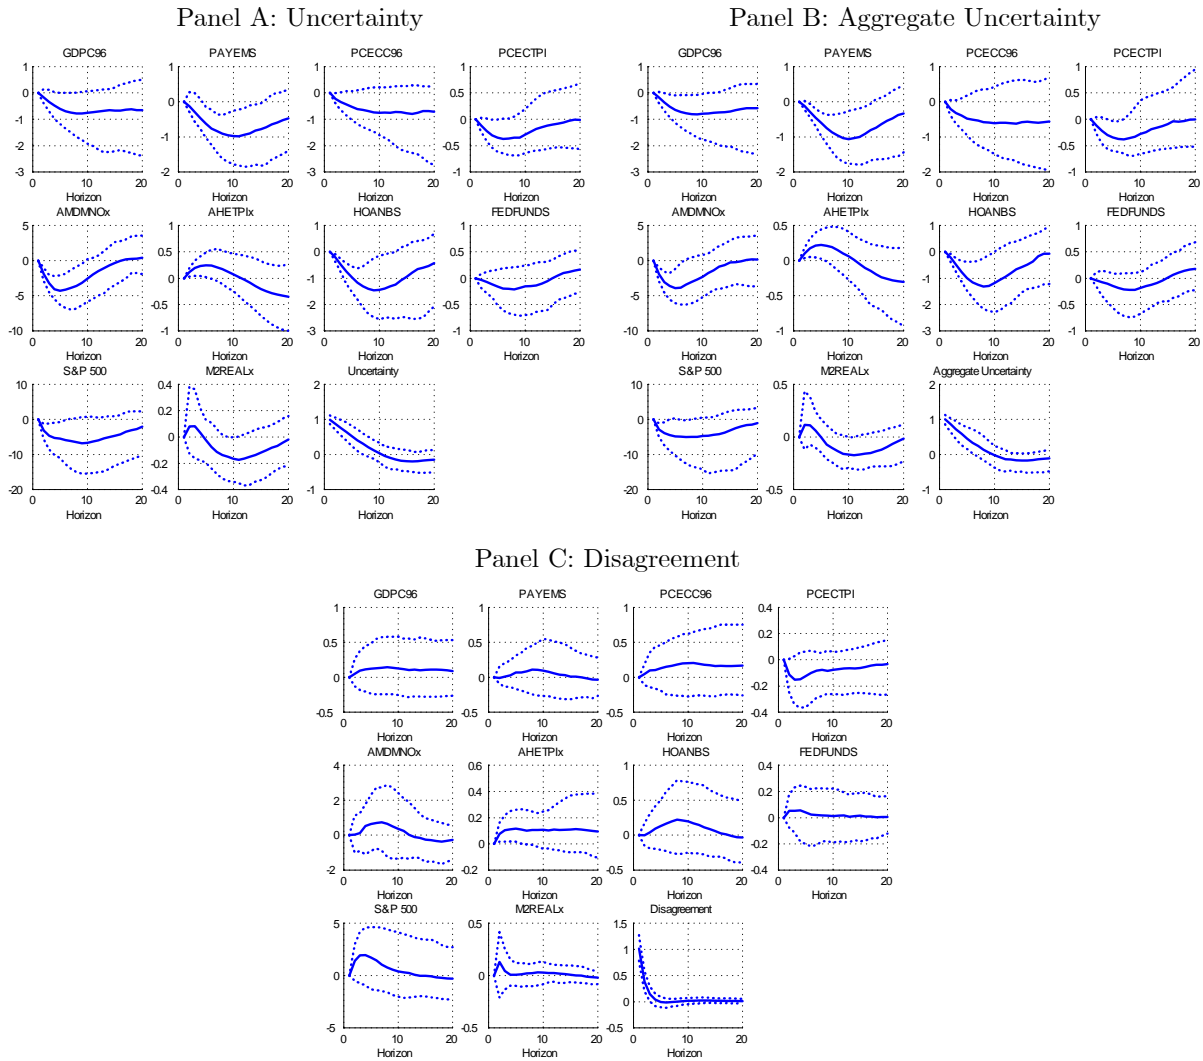
A.3 Results for the Large-Dimensional VAR

This section shows the robustness of our results to a mid-size 11 variable VAR as considered in Jurado, Ludvigson and Ng (2015) specified in the spirit of Christiano, Eichenbaum and Evans (2005). The VAR (11) is in the following variables: $\log(\text{real GDP})$, $\log(\text{employment})$, $\log(\text{real consumption})$, $\log(\text{PCE deflator})$, $\log(\text{real new order})$, $\log(\text{real wage})$, hours, federal funds rate, $\log(\text{S\&P 500 Index})$, growth rate of M2, and various uncertainty indices discussed in the paper. The variables are downloaded from the 2015-11 version of the FRED-QD (Quarterly Database for Macroeconomic Research) discussed in McCracken and Ng (2015). The labels on the impulse responses carry the mnemonics of the variables in the database described in Table I.

Table I. Description of Variables Included in the VAR		
<i>Variable</i>	<i>Mnemonics</i>	<i>Description</i>
real GDP	GDPC96	Real Gross Domestic Product, 3 Decimal (Billions of Chained 2009 Dollars)
Employment	PAYEMS	All Employees: Total nonfarm (Thousands of Persons)
Real Consumption	PCECC96	Real Personal Consumption Expenditures (Billions of Chained 2009 Dollars)
PCE deflator	PCECTPI	Personal Consumption Expenditures: Chain-type Price Index (Index 2009=100)
real new order	AMDMNOx	Real Manufacturers' New Orders: Durable Goods (Millions of 2009 Dollars), deflated by Core PCE
real wage	AHETPIx	Real Average Hourly Earnings of Production and Nonsupervisory Employees: Total Private (2009 Dollars per Hour), deflated by Core PCE
hours	HOANBS	Nonfarm Business Sector: Hours of All Persons (Index 2009=100)
federal funds rate	FEDFUNDS	Effective Federal Funds Rate (Percent)
S&P 500 Index	S&P 500	S&P's Common Stock Price Index: Composite
M2	M2REALX	Real M2 Money Stock (Billions of 1982-84 Dollars)

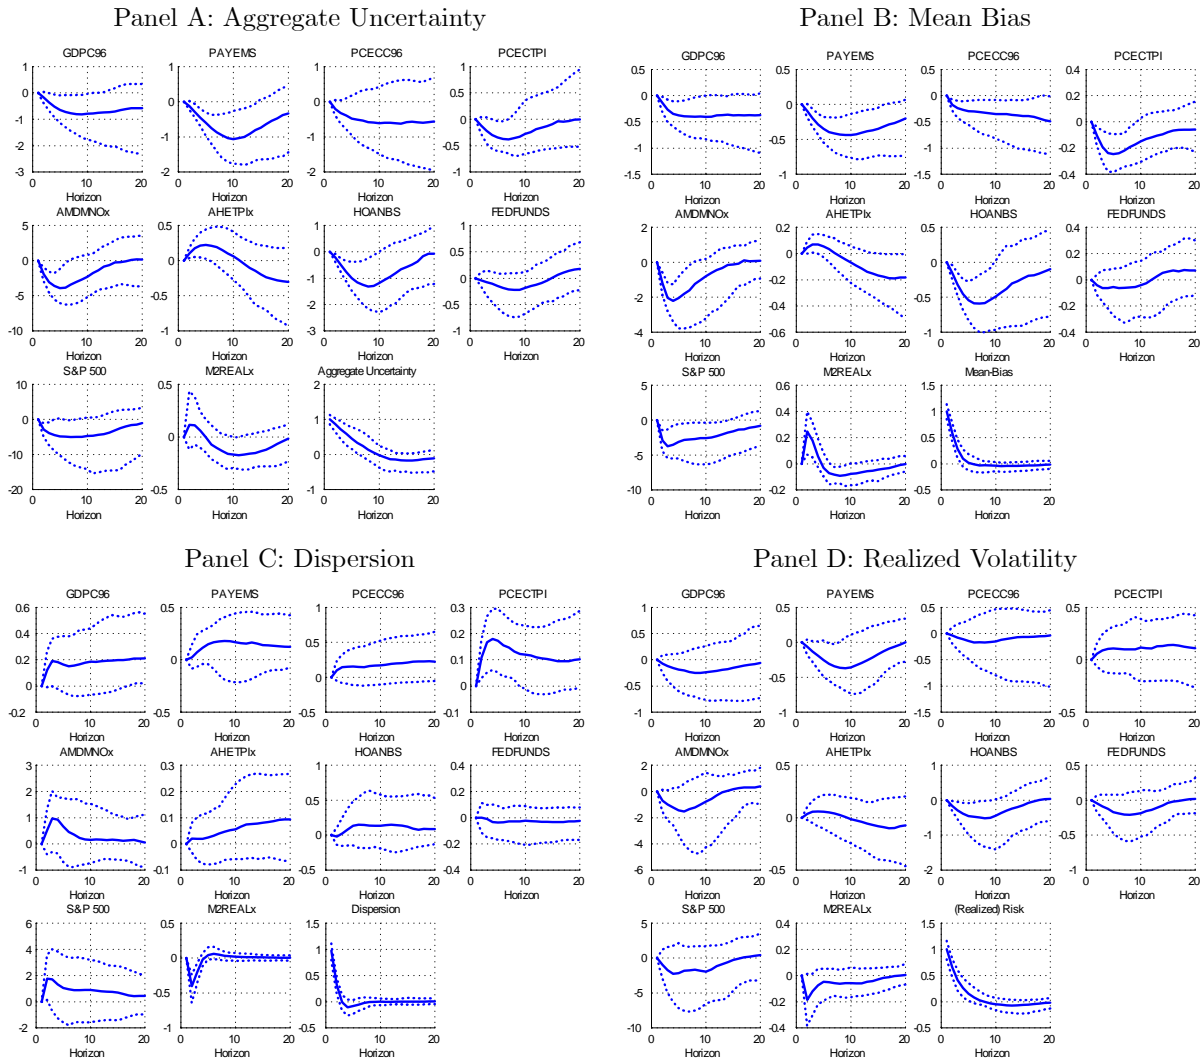
The impulse responses to shocks in the uncertainty indices are displayed in Figures IV-VII. The figures show that the findings are in general the same as those we report in the main text: all uncertainty measures are recessionary in nature. The ex-post measures, as well as realized volatility, have higher impact in magnitude than disagreement or ex-ante uncertainty. Also, even in the large VAR, GDP increases after a shock to dispersion.

Figure IV: Macroeconomic Impact of Uncertainty



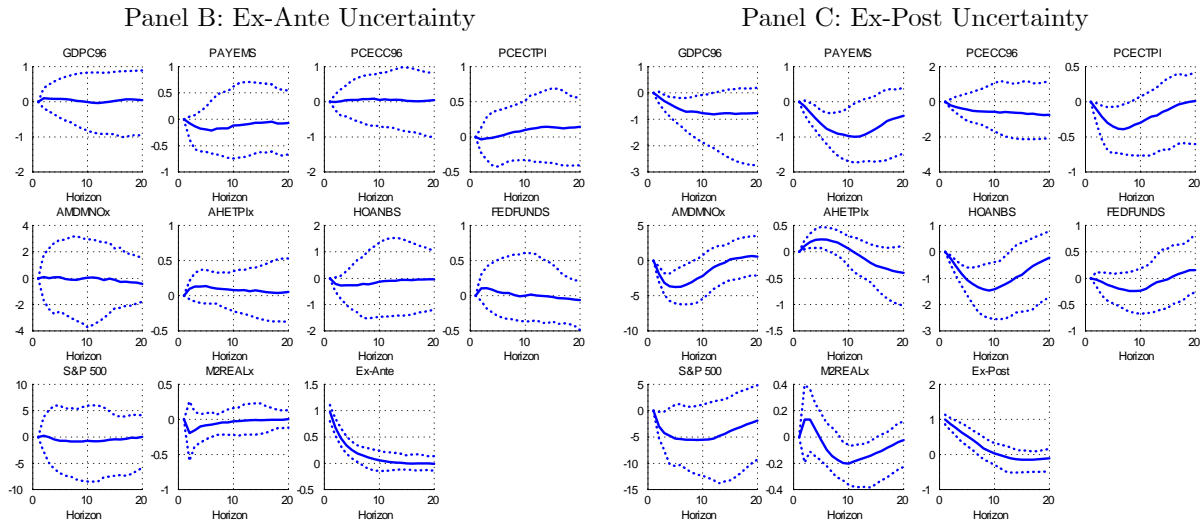
Note: The figure shows the impulse responses of uncertainty, aggregate uncertainty and disagreement based on eq. (5).

Figure V: Macroeconomic Effect of Uncertainty



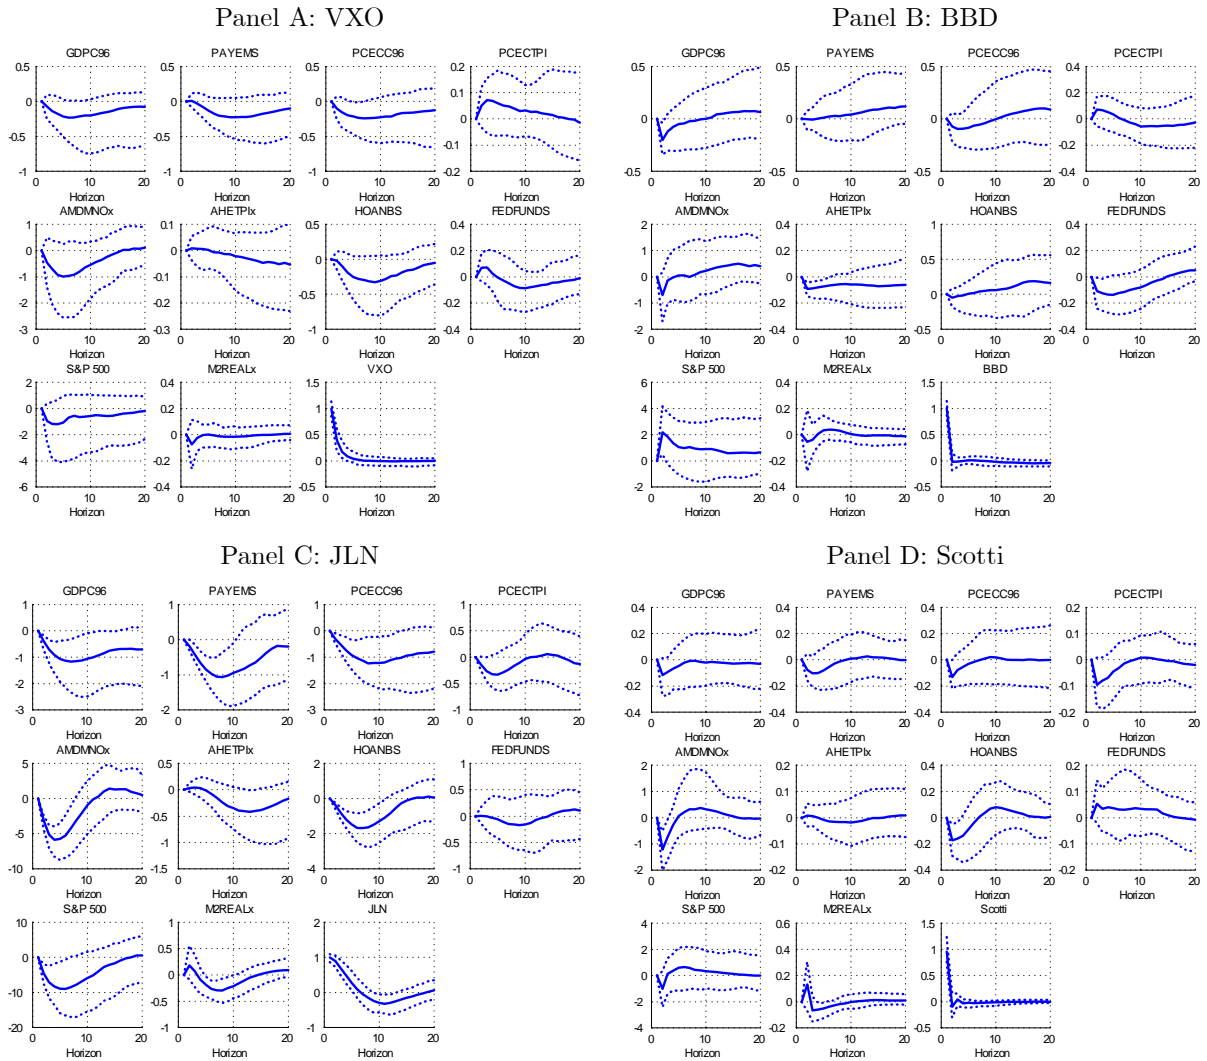
Note: The figure shows the impulse responses of the ex-ante and ex-post measures of uncertainty based on eq. (7). The uncertainty measures have been standardized.

Figure VI: Macroeconomic Impact of Uncertainty



Note: The figure shows the impulse responses of the ex-ante and ex-post measures of uncertainty based on eq. (8). The uncertainty measures have been standardized.

Figure VII: Macroeconomic Effect of Uncertainty - Alternative Measures



Note: The figure shows the impulse responses for the following uncertainty measures: VXO, JLN, BBD and Scotti. The uncertainty measures have been standardized.

A.4 Estimation

We propose to estimate the decomposition with its sample counterparts:

$$\widehat{U}_{t+h|t} = \int_{-\infty}^{+\infty} \widehat{u}_{t+h|t}(r) dr, \quad t = R, \dots, T$$

where R is the size of the rolling window,

$$\widehat{u}_{t+h|t}(r) = \frac{1}{R} \sum_{j=t-R+1}^t \frac{1}{N} \sum_{s=1}^N u_{s,j+h|j}(r) = \frac{1}{R} \sum_{j=t-R+1}^t \frac{1}{N} \sum_{s=1}^N [x_{t+h}(r) - p_{s,j+h|j}(r)]^2$$

and

$$\begin{aligned} \widehat{U}_{t+h|t}^A &= \int_{-\infty}^{+\infty} \left(\bar{p}_{t+h|t}(r) - \bar{x}_{t+h}(r) \right)^2 dr + \int_{-\infty}^{+\infty} \widehat{V}(p_{t+h|t}(r)) dr \\ &\quad + \int_{-\infty}^{+\infty} \widehat{V}ol_{t+h|t}(r) dr - 2 \int_{-\infty}^{+\infty} \widehat{Cov}(x_{t+h}(r), p_{t+h|t}(r)) dr, \end{aligned} \quad (23)$$

where the terms on the RHS of eq. (23) are as follows:

- $\bar{p}_{t+h|t}(r)$, $\bar{x}_{t+h}(r)$ are estimated by $\frac{1}{R} \sum_{j=t-R+1}^t p_{j+h|j}(r)$, $\frac{1}{R} \sum_{j=t-R+1}^t x_{j+h}(r)$;
- $\widehat{V}ol_{t+h}(x_{t+h}(r))$ is an estimate of the variance of $x_{t+h}(r)$, which is a binary variable, recursively over time:

$$\widehat{V}ol_{t+h}(x_{t+h}(r)) = \bar{x}_{t+h}(1 - \bar{x}_{t+h});$$

- $\widehat{V}_{t+h}(p_{t+h|t}(r))$ is an estimate of the variance of $p_{t+h|t}(r)$ recursively over time:

$$\widehat{V}_{t+h}(p_{t+h|t}(r)) = \frac{1}{R} \sum_{j=t-R+1}^t \left(p_{j+h|j}(r) - \bar{p}_{t+h|t}(r) \right)^2$$

- $\widehat{Cov}(x_{t+h}(r), p_{t+h|t}(r))$ is estimated as:

$$\widehat{Cov}(x_{t+h}(r), p_{t+h|t}(r)) = \frac{1}{R} \sum_{j=t-R+1}^t \left(p_{j+h|j}(r) - \bar{p}_{t+h|t}(r) \right) \left(x_{j+h}(r) - \bar{x}_{t+h}(r) \right)$$

While we do not need the Normality assumption to calculate the decomposition above, in practice we fit a Gaussian distribution to the predictive density. The main reason is to guarantee that the “Knightian uncertainty/(Realized) Risk” decomposition is consistent with the “Ex-ante”/“Ex-post”, since the latter is valid only under Normality. Furthermore, in the empirical implementation we let $R = 4$, which amounts to calculating 4-quarter-moving average of the various components of uncertainty, and we proxy the indefinite integrals with definite ones by treating the extrema of either the realization or the bins as integral bounds.

Additional References

Christiano, L.J., M. Eichenbaum and C.L. Evans (2005), “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy* 113(1), 1-45.

McCracken and Ng (2015), “FRED-MD: A Monthly Database for Macroeconomic Research,” *Federal Reserve Bank of St. Louis Working Paper* 2015-012B.

Murphy, A.H. (1986), “A New Decomposition of the Brier Score: Formulation and Interpretation,” *Monthly Weather Review* 114(12), 2671-2673.

Rossi, B. and T. Sekhposyan (2016), “Understanding the Sources of Macroeconomic Uncertainty,” *mimeo*.

Stephenson, D.B., C.A.S. Coelho and I.T. Joliffe (2008), “Two Extra Components in the Brier Score Decomposition,” *Weather and Forecasting* 23, 752-757.