Understanding the Sources of Macroeconomic Uncertainty

PRELIMINARY AND INCOMPLETE

Barbara Rossi* Tatevik Sekhposyan† Matthieu Soupre‡

February 14, 2016

Abstract

We propose a decomposition to distinguish between Knightian uncertainty (ambiguity) and risk, where the first measures the uncertainty on the probability distribution generating the data, while the second measures uncertainty about the odds of particular outcomes when the probability distribution is known. We use Survey of Professional Forecaster's (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 impact effect of the components of our decomposition in a structural model that features ambiguity and risk.

Keywords: Uncertainty, Risk, Knightian Uncertainty, Forecast Distributions, SPF, 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; tel.: +34-93-542-1655; e-mail: barbara.rossi@upf.edu

†Texas A&M University, 3060 Allen Building, 4228 TAMU, College Station, TX 77843, USA; tel.: +1-979-862-8857; e-mail: tsekhposyan@tamu.edu

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

1 Introduction

The recent financial crisis has renewed interest in measuring uncertainty and studying its macroeconomic effects. What is uncertainty? Rational agents face uncertainty when making their decisions, as the realization of the state of nature is not known in advance even if they can reasonably contemplate all possible states of nature and their likelihood. This situation is commonly known as risk. Frank Knight (1921) suggested a different definition of uncertainty, in which agents are uncertain either because they are unable to assign correct probabilities to future outcomes or because they disagree on those probabilities, even if they are correct. The literature has proposed several measures of uncertainty, but they do not distinguish between risk and Knightian uncertainty, nor explains how they relate to each other. In addition, while researchers report correlations between their proposed measure of uncertainty and the alternative ones considered in the literature, it is unclear how they are exactly related to each other.

This paper attempts to study uncertainty in a unified framework. In particular, we make the following contributions to the literature:

(i) We propose to decompose uncertainty into measures of aggregate Knightian uncertainty, risk and disagreement, as well as distinguish between ex-ante and ex-post uncertainty. Our approach will enable us to understand the sources of uncertainty, i.e. the various components that result in the general measure called uncertainty. Several of the components are of interest on their own. For example, Patton and Timmermann (2010) studied disagreement among professional forecasters, but did 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. Furthermore, in our context we will be able to distinguish between ex-ante and realized risk and understand how the mismatch between the two can change the role of the Knightian uncertainty.

(ii) We study how uncertainty and its sources evolve over time. For example, Patton and Timmermann (2010) study the resolution of disagreement over time; disagreement is only one of the components of uncertainty, and we investigate how important it is as a source of overall uncertainty over time.

(iii) We analyze how the various components in our decomposition are related to existing measures of uncertainty. This analysis will help our understanding of why the various measures differ from each other, and which one is more appropriate to use depending on the goals of the researcher.

(iv) We document the macroeconomic impact and transmission of this various sources of uncertainty.

(v) Lastly, we use a stylized macroeconomic model as a framework to discuss the interpretation

[1] **Acknowledgements:** We are grateful to Domenico Giannone for comments.

of the components of our decomposition in the presence of time-varying macroeconomic risk and ambiguity.

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. Thus, an important difference between this paper and others considered in the literature is that we use the probabilistic forecasts provided by Survey of Professional Forecasters (SPF) to measure and decompose uncertainty.

More in detail, our paper sheds light on the difference between uncertainty and risk in the economy, and their relationship. Risk is characterized by a situation where one knows the odds of the unknown, that is, one knows the probability distribution of stochastic events; Knightian uncertainty is characterized by a situation where even the odds or the outcomes are unknown (Knight, 1921). Furthermore, a large number of uncertainty measures (anything that depends on the realization of the data) considered in the literature are ex-post (such as the uncertainty measure recently proposed by Rossi and Sekhposyan, 2015, Jurado, Ludvigson and Ng, 2015, and Scotti, 2013), which arguably are hard to square with the notion of decision making of forward looking economic agents. In our framework, we will be able to distinguish between realized and ex-ante risk. Depending on the nature of the two, the resulting measures of Knightian uncertainty (which is a purely ex-post measure) will change. The advantage of our framework is that we are able to propose a measure that shares properties with a large body of uncertainty measures proposed in the literature, while at the same time, it enables us to obtain components that some might prefer from a decision-theoretic point of view.

Our paper is related to several existing contributions in the literature. Patton and Timmermann (2010) study the sources of disagreement among economic forecasters. To study the disagreement, they define a measure of cross-sectional dispersion among forecasters. They find that differences in forecasters' opinions are highest at the longer horizons, and persist over time; thus, they conclude, the sources of the disagreement is heterogeneity in priors or models, rather than different information sets. They also find that differences in opinions are countercyclical, and heterogeneity is more pronounced in recessions. Our paper has a different focus in that it studies uncertainty. However, among the various definitions of uncertainty considered in the literature, some are related to measures of cross-sectional dispersion (e.g. average disagreement among analysts, as in Bachmann et al., 2015), and Patton and Timmermann's (2010) measure will help understand how disagreement is related to alternative uncertainty measures available in the literature. Lahiri and Sheng (2010) focus on the first two moments of the distribution of density forecasts across individuals, that is the mean and the variance, and link the variance of the forecast error of the consensus forecast to the variance of the individual forecast errors. Focusing on the first two moments involves a loss of information relative to considering the whole density forecast, which we instead do. Our paper is also distantly related to our previous work (Rossi and Sekhposyan, 2015), where they used predictive densities of past forecast error realizations to construct uncertainty indices; in this paper, instead, we use predictive densities from the SPF, including individual predictive densities, and decompose uncertainty into meaningful components.

The paper is structured as follows. The next section presents our uncertainty measure based on survey density forecasts and the decompositions we investigate in this paper. Section 3 provides details on the estimation, and Section 4 discusses the data. Section 5 presents the empirical results, while section 6 analyzes the relationship between our indices and those existing in the literature. Section 7 interprets our decomposition through the lens of a DELETE{structural} macroeconomic model. Section 8 concludes.

2 An Uncertainty Index Based on Density Forecasts

The uncertainty index we propose in this paper measures the distance between the forecast distribution and the perfect forecast, where both are represented as cumulative distribution functions (CDFs).[2] We denote the CDF of 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: $x_{t+h}(r) \equiv 1(y_{t+h} < r)$. Note that the this CDF is defined over the support r, $r \in R$; by varying r, we can focus on different parts of the predictive distribution. Let $p_{s,t+h|t}(r)$ be the probability forecast of the outcome $x_{t+h}(r)$ made by forecaster s, $s = 1, ..., N$, i.e. $p_{s,t+h|t}(r) = P(x_{t+h}(r) = 1|\Omega_t)$, where Ω_t is the information set available at time t.[3] We measure the s-th forecaster's uncertainty as the MSFE of his/her forecast, i.e.:[4]

$$u_{s,t+h|t}(r) = E\left[\left(x_{t+h}(r) - p_{s,t+h|t}(r)\right)^2\right]. \tag{1}$$

Like Jurado, Ludvigson and Ng's (2015) measure, eq. (1) is a Mean Squared Forecast Error (MSFE); however, it is a MSFE applied to a forecast distribution. In fact, $u_{s,t+h|t}(r)$ compares the probability that forecaster s assigns to the different states of nature to the realization, while error-based measures à la Jurado, Ludvigson and Ng (2015) compare the point forecast to the realization.

[2] As we will explain later, our measure of uncertainty is similar to a Continuous Rank Probability Score (CRPS). In fact, it is the negative of the CRPS.

[3] To simplify notation, we assume in this section that R and N are fixed over time, although in the empirical application we will let them change over time.

[4] In the forecasting literature, this MSFE is known as the Brier score.

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

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

As mentioned above, by varying r we can explore measures of uncertainty in different parts of the predictive density. We will focus on an overall measure of uncertainty that averages MSFEs across all possible thresholds, that is,[5]

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

A graphical interpretation is provided in Figure 1. In the figure, the actual realization equals -1, denoted by a bar in the panel on the left; the predictive density is denoted by the Normal distribution. The panel on the right shows the CDF of the Normal distribution, as well as the perfect forecast, where the CDF of the perfect forecasts assumes zeros for values less than -1 and one otherwise. For any given r, the distance between the CDFs of the forecasted distribution and that of the perfect forecast is depicted by solid vertical lines. Our measure of uncertainty squares this measure and integrates it over the various values of r.

INSERT FIGURE 1

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

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

and the corresponding uncertainty measure for the aggregate predictive density as

$$u^A_{t+h}(r) \equiv E\left[\left(x_{t+h}(r) - p_{t+h|t}(r)\right)^2\right].$$

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

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

[5] Note that eq. (2) is the negative of the CRPS, as defined in Gneiting and Raftery (2007). In fact, the CRPS is an average of Brier scores (Hersbach, 2000, eq. 7).

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

$$U_{t+h|t} = \int_{-\infty}^{\infty} u_{t+h|t}^{A}(r)\,dr + \int_{-\infty}^{\infty} d_{t+h|t}(r)\,dr \tag{5}$$

$$= U_{t+h|t}^{A} + D_{t+h|t} \tag{6}$$

The component $u_{t+h|t}^{A}(r)$ can be further decomposed as:

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

where $V\left(x_{t+h}(r)\right) = E\left[\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right)\right)^2\right]$. Thus,

$$U_{t+h|t}^{A} = V_{t+h|t} + B_{t+h|t}, \tag{7}$$

where:

- $V_{t+h|t} = \int_{-\infty}^{\infty} V\left(x_{t+h}(r)\right)dr$ is the unconditional realized variance of the binary outcome, $V\left(x_{t+h}(r)\right)$, and thus stands for the inherent risk in the data;
- $B_{t+h|t} \equiv \int_{-\infty}^{\infty} E\left(\left[p_{t+h|t}(r) - E\left(x_{t+h}(r)\right)\right]^2\right)dr$ is the unconditional squared bias of the forecast distribution.

The two components in eq. (7) have interesting interpretations. We view the variance component 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 as measuring how distant the predictive density is from the perfect prediction. Thus, we view the sum of bias and disagreement as measuring Knightian uncertainty. 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:

$$U_{t+h|t} = \underbrace{V_{t+h|t}}_{\text{"(Realized) Risk"}} + \underbrace{B_{t+h|t} + D_{t+h|t}}_{\text{"Knightian Uncertainty"}}$$

It is important to note that our proposed measure of uncertainty, $U_{t+h|t}$, is 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 ex-post realizations) and an ex-post component (which does). Also, one might wonder how the expected mean forecast and the ex-ante variance of the forecast distribution affect our measure of uncertainty. Let the aggregate predictive distribution made at time t for time $t+h$ be Normal with mean $\mu_{t+h|t}$ and variance $\sigma^2_{t+h|t}$; in that case, we have:

$$U^A_{t+h|t} = \underbrace{\left[2\sigma_{t+h|t}\phi\left(\frac{x_{t+h}-\mu_{t+h|t}}{\sigma_{t+h|t}}\right) + \left(x_{t+h}-\mu_{t+h|t}\right)\left(2\Phi\left(\frac{x_{t+h}-\mu_{t+h|t}}{\sigma_{t+h|t}}\right)-1\right)\right]}_{\text{"Ex-post"}} - \underbrace{\sigma_{t+h|t}/\sqrt{\pi}}_{\text{"Ex-ante"}} \tag{8}$$

where $\phi(.)$ and $\Phi(.)$ denote the PDF and the CDF of the Normal distribution, respectively. The proof is provided in Appendix A and follows Gneiting and Raftery (2007). Note that even if $U^A_{t+h|t}$ is the difference of two components, it is always positive; thus, the ex-post component is always bigger than ex-ante one. The rightmost component is the only component that is not affected by the realization, so we refer to it as the "ex-ante" measure of uncertainty. Moreover, it is a function of a standard deviation that is embedded in 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.[6] Our ex-ante component might be viewed as a measure of ex-ante risk. Note that, from eqs. (7) and (8), we have that $\textit{Ex-post} = B_{t+h|t} + (V_{t+h|t} - \textit{Ex-ante})$. Thus, the ex-post measure of uncertainty is a combination of a component of Knightian uncertainty, $B_{t+h|t}$, and the difference between realized risk (measured by the volatility in the economy, $V_{t+h|t}$) and ex-ante risk (measured by the variance of the predictive densities of the forecasters, *Ex-ante*). If the forecasters' ex-ante volatility measure equals the realized volatility, then the portion of Knightian uncertainty denoted by $B_{t+h|t}$ equals the ex-post measure of uncertainty.

3 Estimation

We propose to estimate the decomposition with its sample counterparts:

$$U_T = \int_{-\infty}^{+\infty} u_T(r)\, dr,$$

[6] For a Normal distribution, the inter-quantile range is 1.34σ.

where

$$u_T(r) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}u_{s,T}(r) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left[x_{t+h}(r) - p_{s,t+h|t}(r)\right]^2$$

and

$$U_T = \int_{-\infty}^{+\infty} B_T(r)\,dr + \int_{-\infty}^{+\infty} V_T(r)\,dr, \tag{9}$$

where the terms on the RHS of eq. (9) are as follows:

- $V(r)$ is an estimate of the variance of $x_{t+h}(r)$, which is a binary variable:

$$V_T(r) = \overline{x}_T(r)\left(1 - \overline{x}_T(r)\right);$$

- $B_T(r)$ is an estimate of the bias squared, estimated as:

$$B_T(r) = \frac{1}{T}\sum_{t=1}^{T}\left(p_{t+h|t}(r) - \overline{x}_T(r)\right)^2. \tag{10}$$

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.[7] We use the real GNP/GDP density forecasts to extract measures of macroeconomic uncertainty. We use real GNP/GDP as the measure of output. In the SPF data set, forecasters are asked to assign a probability value (over predefined intervals) of year-over-year inflation and output growth for the current and following (one-year-ahead) calendar years. The forecasters update the assigned probabilities for the current-year and the one-year-ahead forecasts on a quarterly basis.

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.

As noted, our uncertainty measure depends on realizations. The realized values of output growth we use are from the real-time data set for macroeconomists, also available from the Federal Reserve Bank of Philadelphia. We use the year-over-year growth rates of output and prices calculated from the first quarterly vintage of real GNP/GDP in each year to evaluate the density forecasts. For instance, in order to obtain the growth rate of real output for 1981, we take the 1982:I vintage of data and calculate the growth rate of the annual average GNP/GDP from 1980 to 1981. We

[7] The composition of the forecasters can be changing over time.

consider the annual-average over annual-average percent change (as opposed to fourth-quarter over forth-quarter percent change) in output to be consistent with the definition of the variables that SPF forecasters provide probabilistic predictions for.

The SPF density forecasts are fixed-event forecasts, which implies that the horizon changes over the quarter. Thus, we use the method proposed by Dovern et al. (2012) to construct fixed-horizon forecasts as a weighted average of the current-year and next-year forecasts. In detail, for each quarter $k = 1, 2, 3, 4$, the survey contains a pair of forecasts $x_{t+k|t}$ for the current-year and $x_{t+k+4|t}$ for the next-year forecasts.[8] In the first quarter of the year the current quarter forecast contains four 4-step-ahead forecasts, while the next year forecast contains zero. In the second quarter of the year, the current quarter forecast contains three 4-step-ahead forecasts, while the next year forecast contains only one, etc. The maximum number of 4-quarter-ahead forecasts that the two average forecasts jointly can contain is four. In order to construct fixed horizon forecasts from the fixed event forecasts, Dovern et al. (2012) weight the current year and next year forecasts by weights that count the number of four-step-ahead forecasts in each in the following way. Consider e.g. the 4-step-ahead forecasts written as $X_{t+4|t}$. For a particular value of k,

$$X_{t+4|t} = \frac{k}{4}\tilde{X}_{t+4|t} + \frac{4-k}{4}\tilde{X}_{t+8|t} \tag{11}$$

INSERT FIGURE 2 HERE

Figure 2 shows the evolution of the estimated density over time. The figure plots the mean of the distribution as well as several quantiles, together with the realization. The distribution has the typical features: there was higher ex-ante uncertainty (wider distribution) earlier in the sample. The forecasters appear to capture the great moderation starting mid-1980s. There appears to be no dramatic shift of the forecasted densities after the Great Recession.

5 The Dynamics of Uncertainty Over Time, and Its Sources

Figure 3, Panel A, shows the evolution of our estimated measure of uncertainty over time.[9] The figure highlights two interesting facts: disagreement is, in magnitude, only a small portion of the overall measure of uncertainty; in addition, it is roughly decreasing until the financial crisis of 2007; this is in sharp contrast with the overall measure of uncertainty, which has clear spikes in the early 1980s, early 2000s and the financial crisis. Thus, using disagreement as a measure of uncertainty may result in underestimating uncertainty in the economy. In addition, most would agree that early

[8] Note that this approach is not really essential for our analysis: we could alternatively construct uncertainty measures that are horizon-specific. Empirically, the differences between the latter approach and the one we pursue in the paper are negligible.

[9] We estimate the two components over time, and report the average measures over rolling windows of 4 quarters.

2007-2008 would be the most uncertain times in the latest decade; while disagreement increases during that period, it peaks only much later, after the end of the recession, in 2010. Thus, disagreement (i.e., the component of Knightian uncertainty due to disagreement among forecasters) may not be timely measure of uncertainty.

Panel B in Figure 3 depicts the decomposition of uncertainty into its bias and variance. The figure suggests that realized risk (measured by the variance component) was an important component of uncertainty throughout the last three decades, as it was Knightian uncertainty, measured by the bias component. Some differences between the two are important to note, however. The variance component was high during the latest financial crisis, and sharply decreased as soon as the recession was over; Knightian uncertainty (measured by the bias component) 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.

Turning to the ex-ante and ex-post components in uncertainty, depicted in panel C of Figure 3, as well as the overall uncertainty measure, it is very interesting to note that ex-ante uncertainty is quite constant over the 1980s and up to 2007, with a visibly sharp decline in 1992-1993, possibly due to changes in the survey questionnaire. Thus, movements in uncertainty during that period cannot be attributed to changes in ex-ante uncertainty affecting the economy. Ex-ante uncertainty does increase during the late recession, but only towards its end, and spikes much later than the peak of the financial crisis. This suggests that measures of volatility of the forecasters' predictive distribution are, themselves, not timely measures of uncertainty.

INSERT FIGURE 3 HERE

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 will be useful for practitioners to choose which measure is the most appropriate to use depending on their goals.

The top panel in Figure 4 plots Jurado, Ludvigson and Ng's (2014) uncertainty measure together with Bloom, Baker and Davis' (2013) index. 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 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 (2014) component does relative to the Bloom, Baker and Davis (2013) one. Thus, it seems that the Bloom, Baker and Davis (2013) uncertainty measure is driven more by ex-ante uncertainty, while the Jurado, Ludvigson and Ng (2014) uncertainty measure is clearly affected by ex-post uncertainty, namely uncertainty due to misspecification in the predictions.

INSERT FIGURE 4 HERE

To estimate the effects of the sources of uncertainty 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, the specific uncertainty index (either reliability, resolution or variance) and the overall uncertainty index. Note that our indices are calculated over a window of past data, thus there is no look-ahead bias in the measures and, thus, they can be included in the VAR. To better interpret the magnitude of the effects of the uncertainty indices, the uncertainty indices are standardized by their own mean and variance.

Panel A in Figure 5 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-D describe the effects of each of the components in the decomposition. Panel B focuses on the disagreement; it also decreases employment and the S&P500, although by a smaller magnitude; at the same time, it has a small and positive contemporaneous effect on GPD, while the medium term effect is again negative. The interest rate reacts marginally and negatively on impact, but the effects become quickly positive for several quarters. An increase in the bias (Panel C) and in the variance also have recessionary effects on all the macroeconomic variables; they are, however, smaller in magnitude than those of the overall index. Interestingly, the effects of the bias are much larger than those of the variance on GDP, employment and the stock market, while the effects on the interest rate are similar.

INSERT FIGURE 5 HERE

The effects of ex-ante and ex-post uncertainty on other macroeconomic variables are depicted in Figure 6. They both lead to decreases in employment, interest rates and the stock market 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.

We compare 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 too. The alternative uncertainty indices that we explore (one-at-a-time) include: Bloom's (2009), labeled "VXO"; Baker et al.'s (2013) policy uncertainty index, labeled "BBD"; Jurado, Ludvigson and Ng's (2014), labeled "JLN"; and Scotti's (2013) macroeconomic surprise-based uncertainty index.

INSERT FIGURE 6 HERE

Panel A in Figure 7 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 (2014) 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 Scotti's index are qualitatively similar to those of the disagreement component, suggesting that disagreement might have similar effects to macroeconomic news shocks, which may affect forecasters differently. The effects of the variance component in our decomposition are more similar to those of the VXO.

INSERT FIGURE 7 HERE

7 The Macroeconomic Impact of the Various Sources of Uncertainty Through the Lens of a Structural Model

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

$$logZ_{t+1} = \rho_z logZ_t + \mu_t^* + u_{t+1} \tag{12}$$

u_{t+1} is i.i.d. $N(0, \sigma_u^2)$ and μ_t^* is a deterministic sequence the empirical sequence of which converges to an i.i.d. stochastic process $N(0, \sigma_z^2 - \sigma_u^2)$, such that observed values of $z_t \equiv logZ_{t+1} - \rho_z logZ_t$ look like realizations from an i.i.d. process with mean zero and variance σ_z^2. For all practical purposes, in the current context we treat μ_t^* as a realization from a stochastic process. Moreover, μ_t^* and u_t are independent. Thus, one could think about the GDP growth being driven by two sources of uncertainty in the economy: the first is unpredictable shocks, u_{t+1}; the second, μ_t^*, is a component that will be driven by ambiguity, as we discuss below.

We assume that the agents in this model know that the model generating GDP is autoregressive with persistence ρ_z and that there are two sources of uncertainty; however, they do not observe μ_t^*. They gather intangible information about μ_t^*, which sometimes makes them relatively confident that the correct forecast of future (log) GDP is $\rho_z logZ_t$, and sometimes makes them less confident, i.e. the signal is more ambiguous. The ambiguity is modeled by letting agents form their beliefs about GDP dynamics based on the following law of motion:

$$logZ_{t+1} = \rho_z logZ_t + \mu_t + u_{t+1}, \tag{13}$$

where $\mu_t \in [-a_t, -a_t + 2|a_t|]$ and u_{t+1} is i.i.d. Normal $(0, \sigma_u^2)$. The bounds on μ_t formalize the idea that sometimes agents are more ambiguous about the second source of disturbance to output growth: a bigger a_t implies a larger set of beliefs, and thus more ambiguity perceived at time t. Thus, we refer to a_t as the ambiguous component, or Knightian uncertainty.

Furthermore, agents learn about μ_t based on the process:

$$a_{t+1} - \bar{a} = \rho_a(a_t - \bar{a}) + \sigma_a \epsilon^a_{t+1}. \tag{14}$$

One can view ϵ^a_{t+1} as signals that the agents get about the ambiguity component. In some periods the signals can be such that the resulting a_t is higher. Consequently, 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 of the data generating process. Furthermore, there are certain parameter restrictions imposed to ensure that the average ambiguity is less than the total uncertainty about the process of $logZ_{t+1}$: these restrictions are that $\bar{a} = n\sigma_z$ and $\sigma_a = \sigma_n \sigma_z$ for $n \in (0, 1)$, where $\sigma^2_z = \sigma^2_\mu + \sigma^2_u$ and n and σ_n are parameters (one can think of $\bar{a}$ and σ^2_a as the mean and the variance of ambiguity).

In addition, when faced with ambiguity, modeled with eq. (14), the agents choose $\mu^{**}_t = min([-a_t, -a_t + 2|a_t|])$. Thus, the effective perceived law of motion becomes:

$$logZ_{t+1} = \rho_z logZ_t + \mu^{**}_t + u_{t+1}. \tag{15}$$

Note that when a_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 an output growth dynamics of a similar form. Thus, for simplicity, we directly model the dynamics of output growth and calibrate the parameters of the output growth process, ρ_z and σ^2_z, based on an AR(1) model estimated on the annualized US GDP growth. On the other hand, the parameters guiding ambiguity, i.e. ρ_a, n and σ_n are borrowed from the posterior mode estimates in Ilut and Schneider (2014).

We consider three scenarios:

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. In essence, the data is generated based on equation (12). The forecasters forecast based on the law of motion in equation (15). In this exercise we are changing the set of possible values that μ_t can take. As μ_t increases, both the conditional and unconditional means of the a_{t+1} increase (see equation (12)). Thus, the signals the agents get about this additional source of uncertainty become noisier. Given that the agents solve the issue with the multiple priors based on $\mu^{**}_t = min([-a_t, -a_t + 2|a_t|])$, this also makes the conditional and unconditional mean of the process $logZ_{t+1}$ biased downwards.

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 the eq. (12), the perceived law of motion is described by eq. (15), while learning under ambiguity occurs under eq. (14). In this case, increasing the level of uncertainty increases both the objective and perceived level of uncertainty. However, given that $\bar{a} = n\sigma_z$ 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) also increase. Thus, an increase in the 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 right model. That is, $\mu_t^{**} = \mu_t^*$. Thus, there is not ambiguity. In other words, the true model is still the one described by (12), while the model used for forecasting is not determined by equation (15), but instead by equation (12) itself. The design in this scenario intends to explore how the ex-post and overall uncertainty evolve when there is no ambiguity.

INSERT FIGURE 8 HERE

We simulate the model for 250 periods (using an additional 100 periods as the burn-in sample) for each of these scenarios, and use the simulated data to construct the components of our proposed two decompositions and plot them over time.

Panel A depicts the results for Scenario 1. The increase in ambiguity increases the Bias and the Ex-Post uncertainty components, 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 the Variance components, respectively. This follows from the fact that, as eqs. (12), (14) and (15) suggest, the overall level of uncertainty (both the true variance, σ_z, as well as the expected one, $\sigma_u^2 + \sigma_a^2 = \sigma_u^2 + \sigma_n^2 \sigma_z^2$) do 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 here that there is a feedback from the risk to ambiguity. As discussed in the description of Scenario 2, both the mean ($\overline{a}$) 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.

Lastly, 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, the comparison of the results suggests that, if there are any effects of ambiguity on uncertainty, then the total uncertainty should move proportionally more upon the increase of risk.

To summarize, our simulations show that the increase in ambiguity can increase the ex-post component, as well as the bias, thus resulting in an overall increase in uncertainty. The increase in the true volatility of the DGP will increase both the realized volatility as well as the ex-ante volatility measures. However, the increase in the overall uncertainty effects the ex-post volatility and bias as well. In the absence of ambiguity, the impact on the bias is negligible (it is more of a 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, our model has a 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.

8 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 its components and 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 a variance and a bias components, where the former measures risk and the latter measures Knightian uncertainty.

Given that our proposed uncertainty index is an ex-post measure of uncertainty, we propose also to decompose it into a component that only reflects ex-ante uncertainty, which we can relate to existing measures of uncertainty based on inter-quantile spreads, 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 (measured by our variance component) 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 due mostly 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 are 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 act as an amplifying mechanism for the increase in the level of risk.

9 References

Bachmann, R., S. Elstner and E. Sims (2013), "Uncertainty and Economic Activity: Evidence from Business Survey Data", *American Economic Journal: Macroeconomics* 5(2), 217-249.

Baker, S.R., N. Bloom, and S.J. Davis (2013), "Measuring Economic Policy Uncertainty," *mimeo.*

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

Boero, G., J. Smith and K. Wallis (2008), "Uncertainty and Disagreement in Economic Prediction: the Bank of England Survey of External Forecasters", *The Economic Journal* 118, 1107-1127.

Bomberger, W.A. (1996), "Disagreement as a Measure of Uncertainty", *Journal of Money, Credit and Banking* 28(3), 381-392

Gneiting, T. and A.E. Raftery (2007), "Strictly Proper Scoring Rules, Prediction, and Estimation", *Journal of the American Statistical Association* 102 (477), 359-378.

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.

Lahiri, K. and X. Sheng (2010), "Measuring Forecast Uncertainty by Disagreement: The Missing Link", *Journal of Applied Econometrics* 25, 514-538.

Newey, W.K. and K.D. West (1987), "A simple, Positive Semi-definite, Heteroskedasticity and

Autocorrelation Consistent Covariance Matrix", *Econometrica* 55(3): 703-708.

Patton A and A. Timmermann (2012), "Forecast Rationality Tests Based on Multi-horizon Bounds", *Journal of Business and Economic Statistics* 30(1), 1-17.

Rossi B. (2012), "Comment to: Forecast Rationality Tests Based on Multi-horizon Bounds", *Journal of Business and Economic Statistics* 30(1), 25-29.

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.

Scotti, C. (2013), "Surprise and Uncertainty Indexes: Real-time Aggregation of Real-Activity Macro Surprises," *Federal Reserve Board International Finance Discussion Paper* 1093.

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.

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

10 Appendix A. Proofs

The appendix provides the proofs for the results in the paper.

Proof of Eq. (8). Our measure of uncertainty is the negative of the CRPS. Note that $CRPS(F,x) = -\int_{-\infty}^{\infty}(F(y) - 1\{y \geq x\})^2 dy$. Let $G(y) = 1\{y \geq x\})$. Given Lemma 2.2 of Baringhaus and Franz (2004), $\int_{-\infty}^{\infty}(F(y) - 1\{y \geq x\})^2 dy = E|X_1 - Y_1| - \frac{1}{2}E|X_1 - X_2| - \frac{1}{2}E|Y_1 - Y_2|$, where X_1 and X_2 are i.i.d draws from F, while Y_1 and Y_2 are i.i.d. draws from G, which have finite expectations. Given Lemma 2.1 of Baringhaus and Franz (2004), $E|X - Y| = \int_{-\infty}^{\infty} F(y)(1 - G(y))dy + \int_{-\infty}^{\infty} G(y)(1 - F(y))dy$. Now for Y_1 and Y_2, we have $E|Y_1 - Y_2| = 2\int_{-\infty}^{\infty} G(y)(1 - G(y))dy = 2\int_{-\infty}^{\infty} 1\{y \geq x\}(1 - 1\{y \geq x\})dy = 0$. The last piece is true since for a particular value of y either $1\{y \geq x\}$ or $1 - 1\{y \geq x\}$ will be zero. Thus, the product will be zero. Thus, $\int_{-\infty}^{\infty}(F(y) - 1\{y \geq x\})^2 dy = E|X_1 - Y_1| - \frac{1}{2}E|X_1 - X_2|$. $CRPS(F,x) = -\int_{-\infty}^{\infty}(F(y) - 1\{y \geq x\})^2 dy = \frac{1}{2}E|X_1 - X_2| - E|X_1 - Y_1|$. If $F(.) = N(\mu, \sigma^2)$, then $X_1 - X_2 \sim (0, 2\sigma^2)$. Further $E|X_1 - X_2|$ will be a half-normal distribution and $E|X_1 - X_2| = \frac{\sqrt{2\sigma^2}\sqrt{2}}{\sqrt{\pi}} = \frac{2\sigma}{\sqrt{\pi}}$. Further, $\frac{1}{2}E|X_1 - X_2| = \frac{\sigma}{\sqrt{\pi}}$. Consider, $E|X_1 - Y_1|$, $X_1 \sim N(\mu, \sigma^2)$, while Y_1 is a draw from a distribution that has a CDF of $G(y) = 1\{y \geq x\})$. The PDF of Y_1 can be written as a Dirac delta function $\delta(y - x)$. Thus, $E(Y_1) = \int_{-\infty}^{\infty} y\delta(y - Y_1)dy = Y_1$, $E(Y_1^2) = \int_{-\infty}^{\infty} y^2\delta(y - Y_1)dy = Y_1^2$ and $Var(Y_1) = E(Y_1^2) - E(Y_1)^2 = 0$.Then, $X_1 - Y_1 \sim N(\mu - Y_1, \sigma^2)$. By the property of folded Normal distribution,

$$\begin{aligned} E|X_1 - Y_1| &= \sqrt{\sigma^2}\sqrt{\frac{2}{\pi}} exp\left(-\frac{(\mu - Y_1)^2}{2\sigma^2}\right) + (\mu - Y_1)\left(1 - 2\Phi\left(-\frac{\mu - Y_1}{\sqrt{\sigma^2}}\right)\right) \\ &= \sigma 2\varphi\left(-\frac{\mu - Y_1}{\sqrt{\sigma^2}}\right) + (\mu - Y_1)\left(1 - 2\Phi\left(-\frac{\mu - Y_1}{\sqrt{\sigma^2}}\right)\right) \end{aligned}$$

$$\begin{aligned} CRPS(F,x) &= \frac{\sigma}{\sqrt{\pi}} - \sigma 2\varphi\left(-\frac{\mu - x}{\sigma}\right) - (\mu - x)\left(1 - 2\Phi\left(-\frac{\mu - x}{\sigma}\right)\right) \\ &= \sigma\left(\frac{1}{\sqrt{\pi}} - 2\varphi\left(\frac{x - \mu}{\sigma}\right)\right) - \frac{x - \mu}{\sigma}\left(2\Phi\left(\frac{x - \mu}{\sigma}\right) - 1\right). \end{aligned}$$

■

Proof of Eq. (4).

$$\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left[x_{t+h}\left(r\right)-p_{s,t+h|t}\left(r\right)\right]^{2}=$$

$$=\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left[x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)+p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right]^{2}$$

$$=\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left[\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)^{2}+2\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)\right]$$

$$+\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left[\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)^{2}\right]$$

$$=\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)^{2}+2\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)$$

$$+\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)^{2}$$

$$=\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)^{2}+2\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)\frac{1}{N}\sum_{s=1}^{N}\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)$$

$$+\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)^{2}$$

$$=\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)^{2}+2\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)\left(p_{t+h|t}\left(r\right)-\frac{1}{N}\sum_{s=1}^{N}p_{s,t+h|t}\left(r\right)\right)$$

$$+\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)^{2}$$

$$=\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+h}\left(r\right)-p_{t+h|t}\left(r\right)\right)^{2}+0+\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N}\sum_{s=1}^{N}\left(p_{t+h|t}\left(r\right)-p_{s,t+h|t}\left(r\right)\right)^{2}.$$

■

Figures

Figure 1: Brier Score Illustration

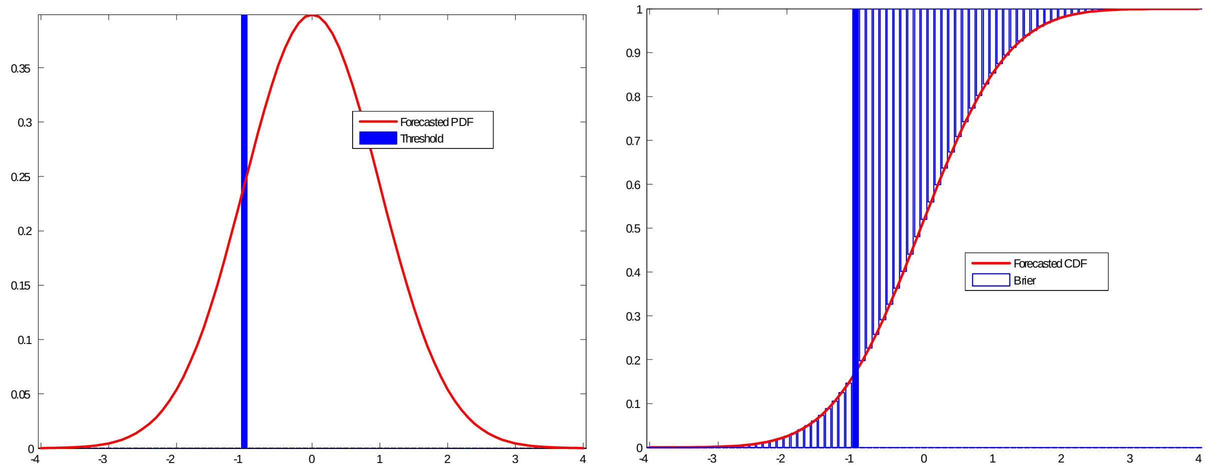

Note: The figure shows the pdf of the predicted and perfect distributions for a threshold $r = -1$ (on the left). On the right we have the CDF of the predicted and perfect distributions, with the area between them (in solid lines) depicting the distance between the two.

Figure 2. The Survey of Professional Forecasters Data

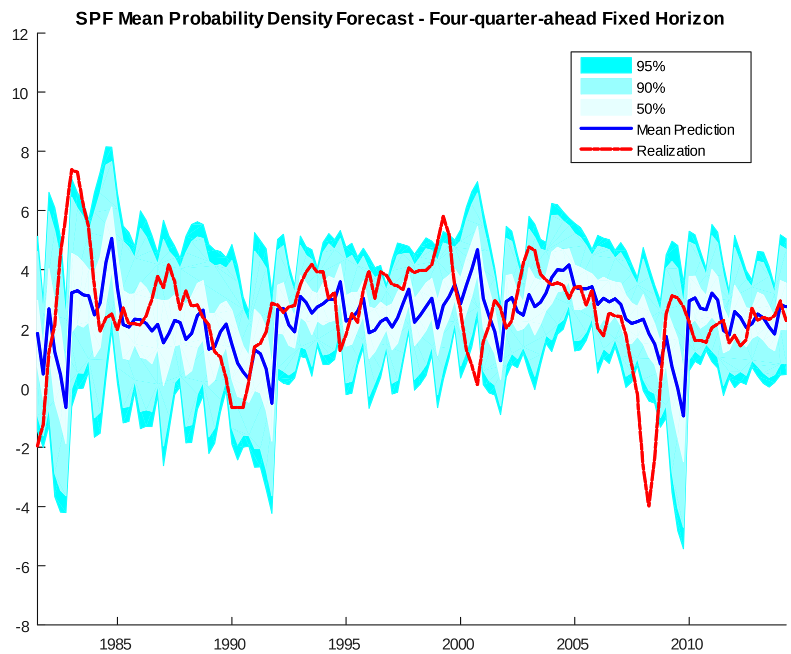

Note: The figure shows the quantiles of the SPF's four-step-ahead predictive density, its mean, as well as the realized GDP growth. The four-step-ahead density is constructed from SPF's current year and next year density forecasts based on eq. (11).

Figure 3: Decomposing Uncertainty in GDP

Panel A: Uncertainty, Aggregate Uncertainty and Disagreement

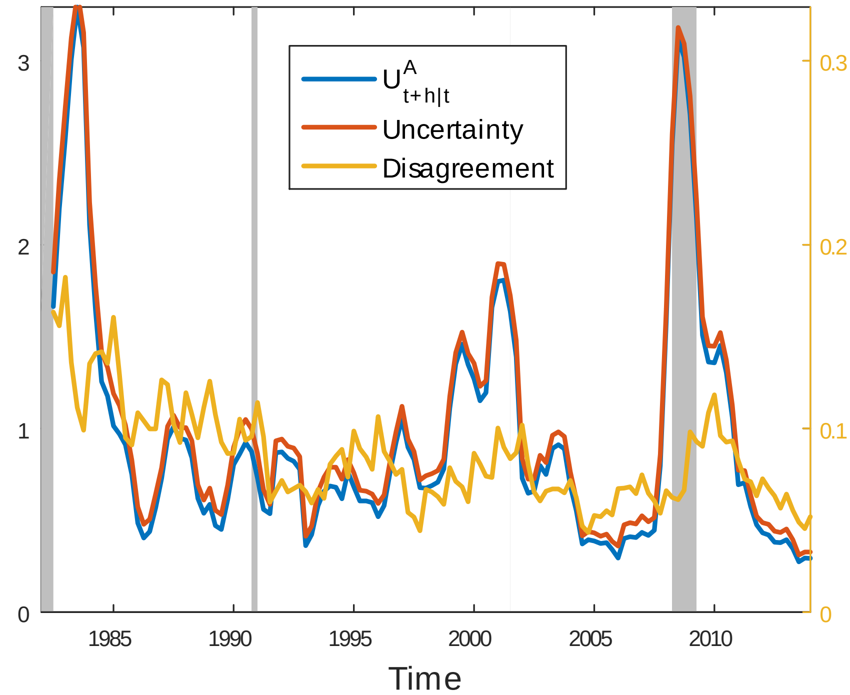

Figure 3: Decomposing Aggregate Uncertainty: GDP

Panel B: Knightian Uncertainty Vs. Risk

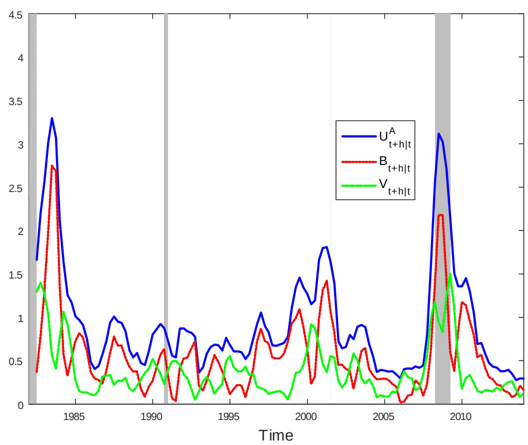

Panel C: Ex-Ante Vs. Ex-Post

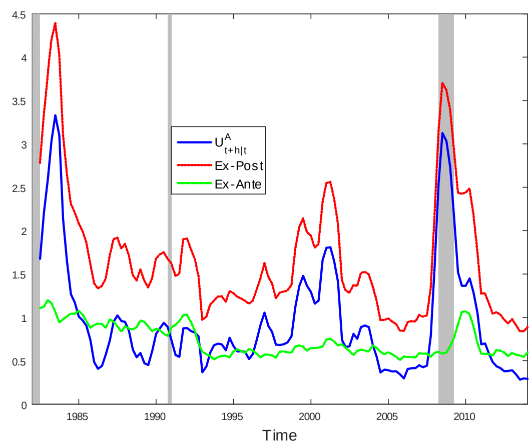

Note: Panel A of figure depicts the evolution of total uncertainty, aggregate uncertainty and disagreement (eq. 4) over time. Panels B and C show the evolution of the components of uncertainty based on eq. (7) and eq. (8), respectively.

Figure 4. Comparison of Uncertainty Measures

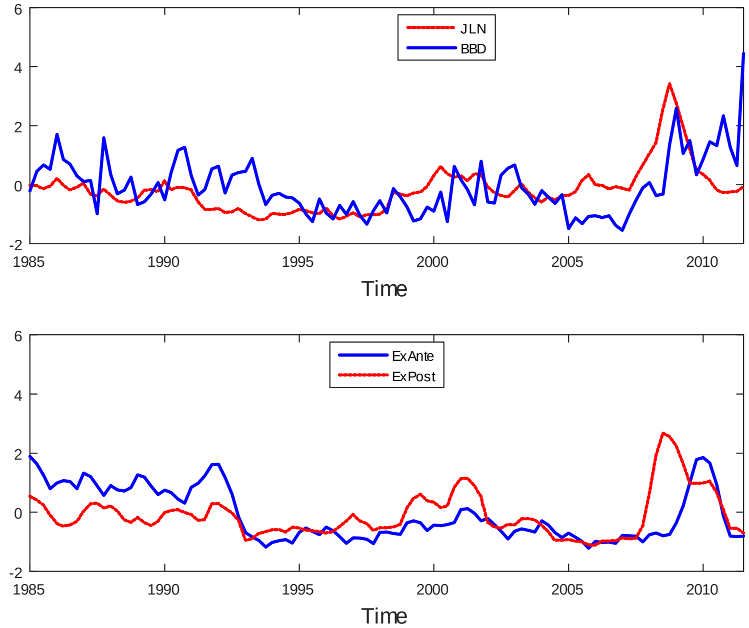

Note: The figure shows the ex-ante ex-post components of our uncertainty measure, obtained based on the decomposition in eq. (8).

Figure 5. The Effects of Uncertainty on the Economy:

Panel A. Our Uncertainty Index

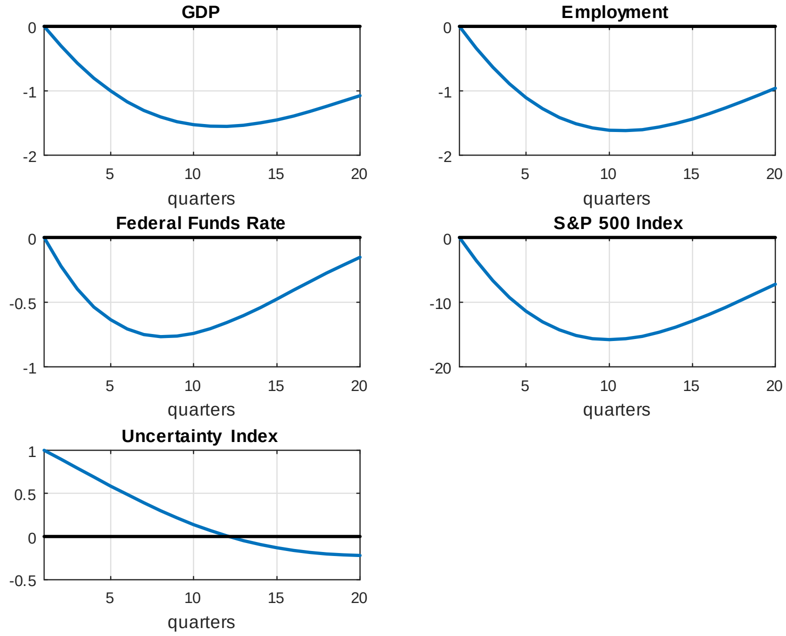

Panel B. Disagreement

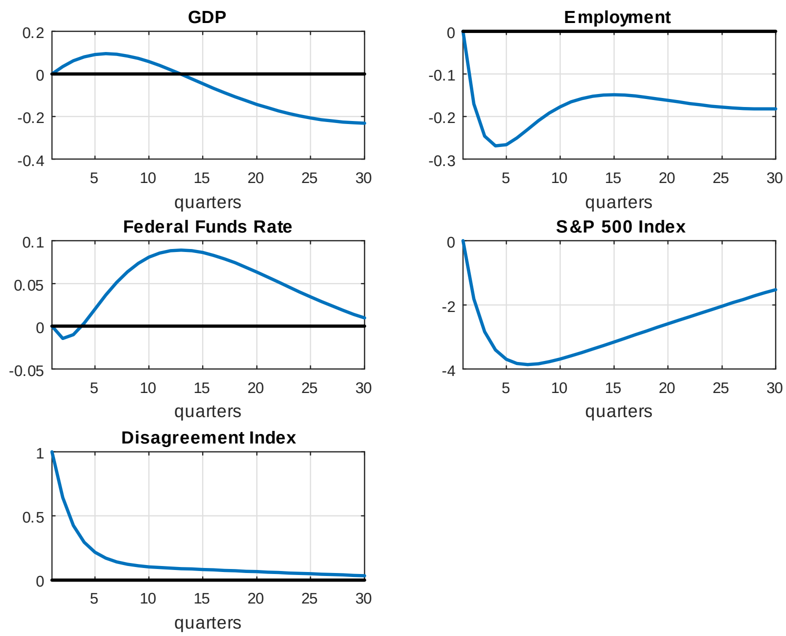

Panel C. Bias

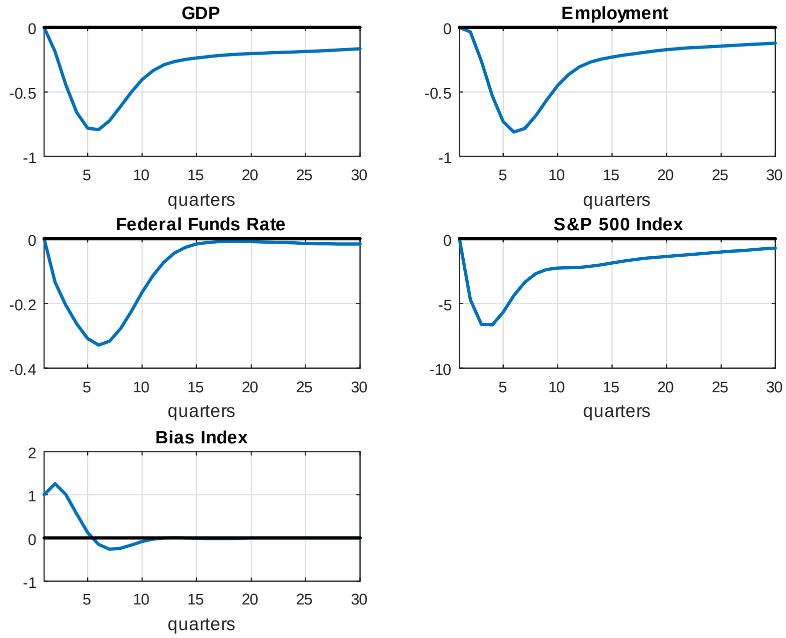

Panel D. Variance

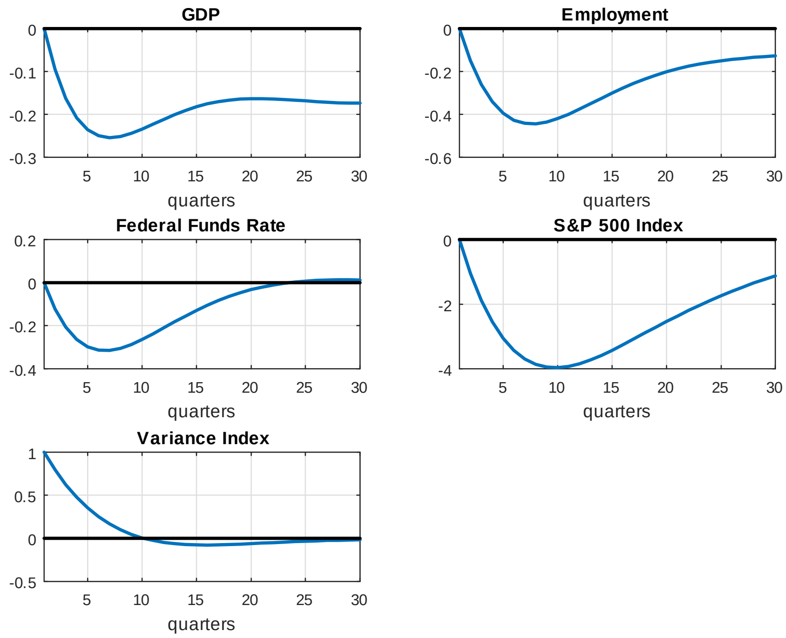

Note: The figure shows the impulse responses of the overall measure of uncertainty and disagreement based on eq. (4) as well as the various components of the aggregate uncertainty based on eq. (7) .

Figure 6. The Effects of Uncertainty on the Economy:

A. Ex-Ante Uncertainty

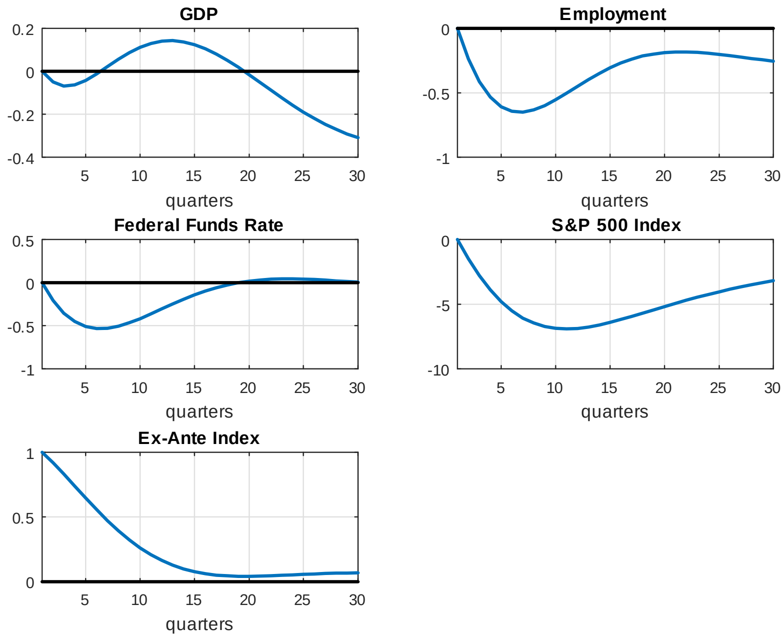

B. Ex-Post Uncertainty

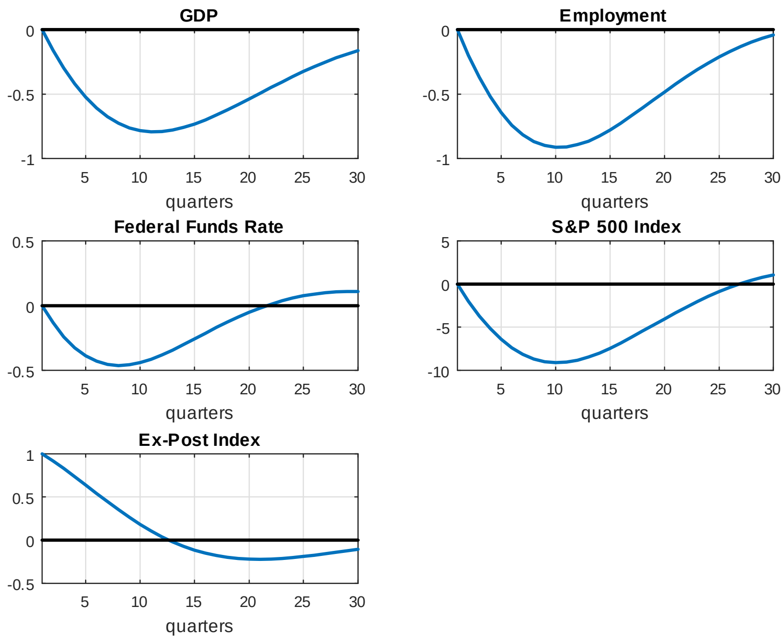

Note: The figure shows the impulse responses of the ex ante and ex post measures of uncertainty based on eq. (8) .

Figure 7. The Effects of Uncertainty on the Economy:

Panel A. VXO Uncertainty Index

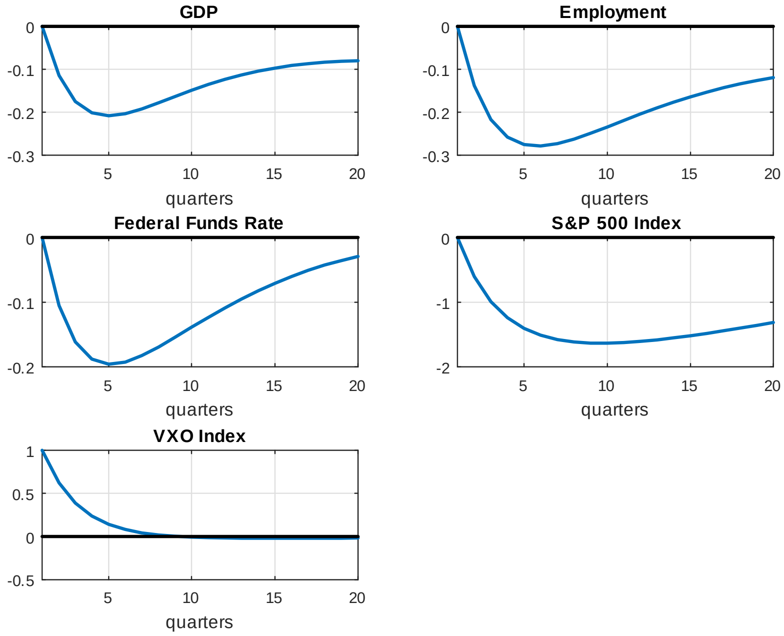

Panel B: JLN Uncertainty Index

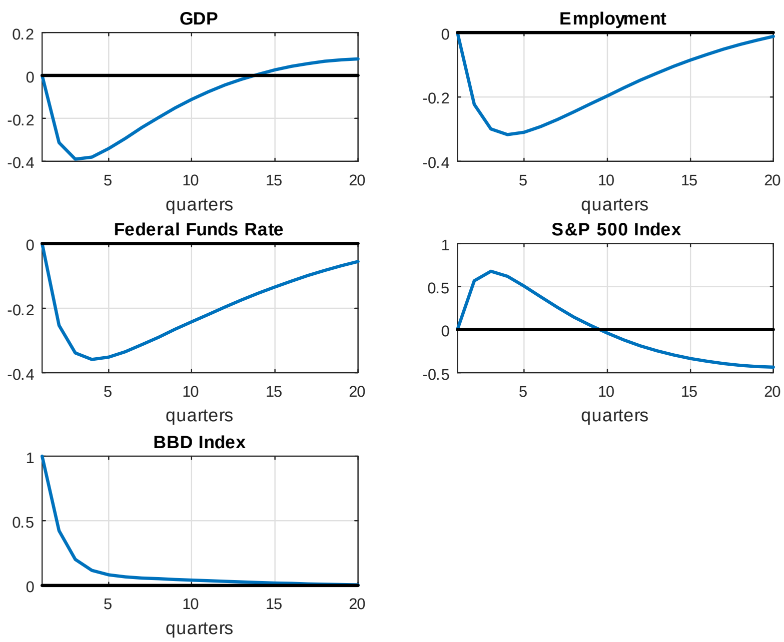

Panel C: BBD Uncertainty Index

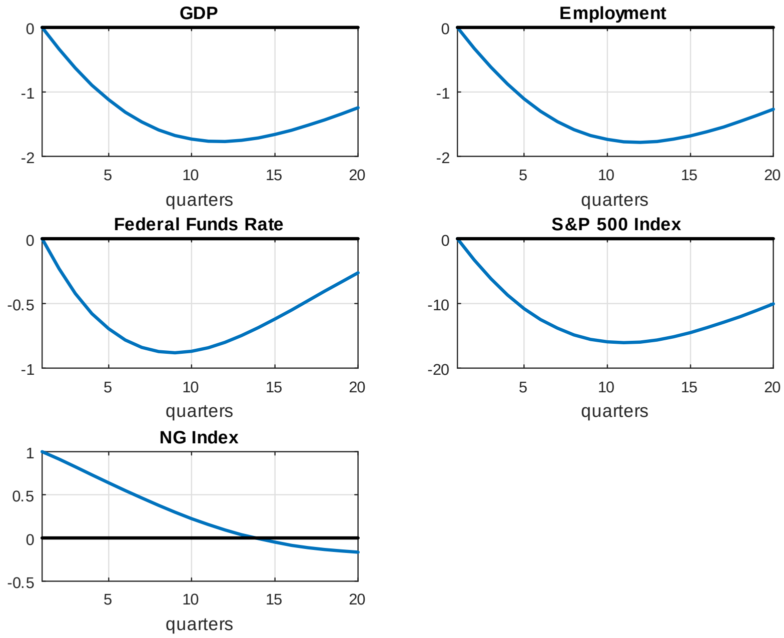

Panel D: Scotti Uncertainty Index

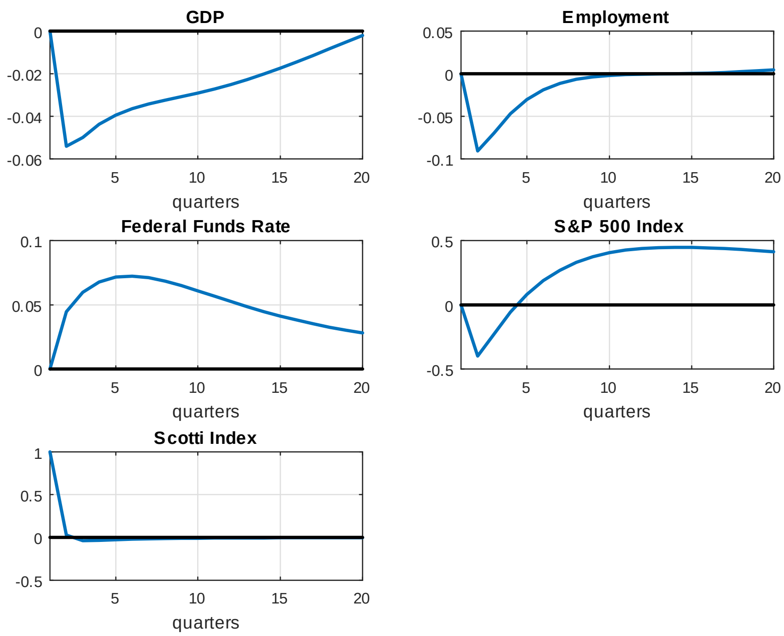

Note: The figure shows the impulse responses for the following uncertainty measures: VXO, JLN, BBD and Scotti. The measures have been standardized.

Figure 8: Simulation Results

Panel A: Changing the Level of Ambiguity

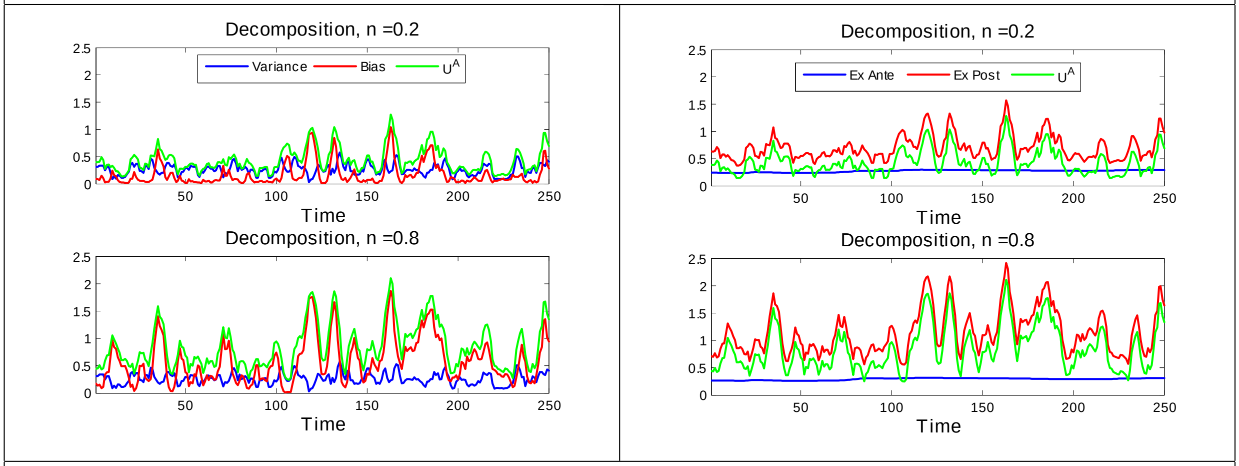

Panel B: Changing the Level of Risk

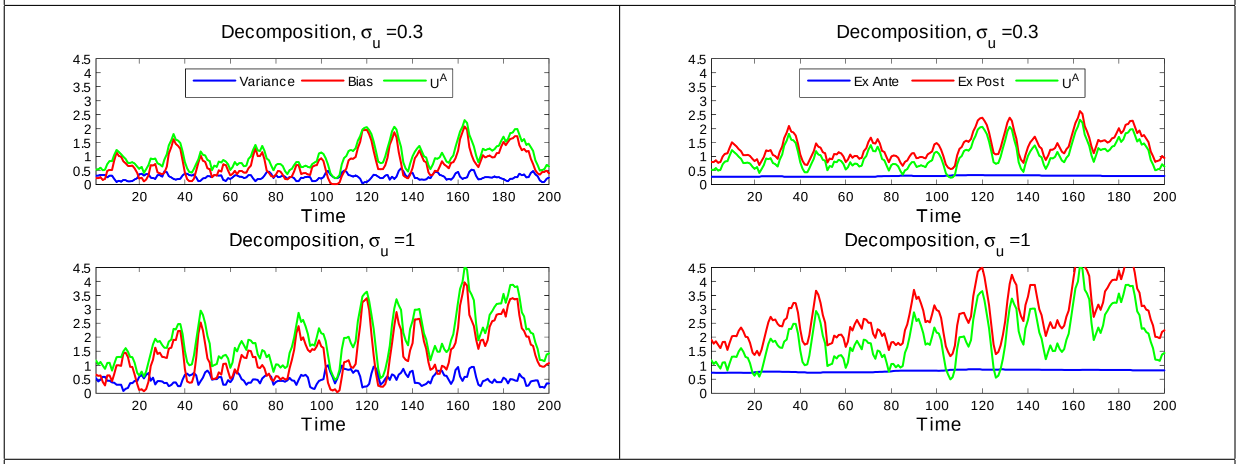

Panel C: Changing the Level of Risk in the Model with no Ambiguity

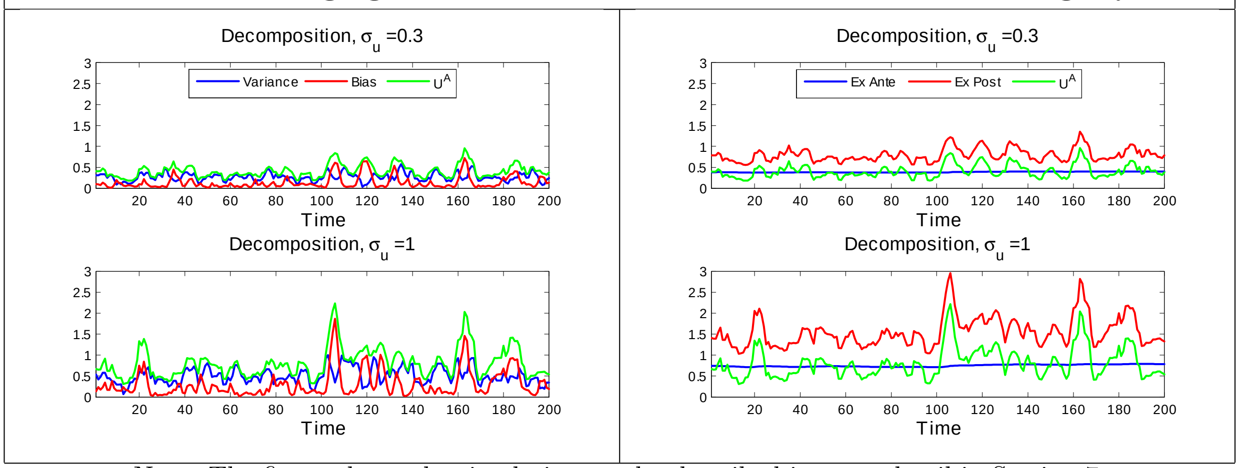

Note: The figure shows the simulation results described in more detail in Section 7.

11 Appendix B

Note that an additional, interesting decomposition of the bias component has been proposed by Murphy (1973), as follows:

$$Bias_{t+h|t}(r) = REL_{t+h|t}(r) - RES_{t+h|t}(r) \tag{16}$$

where:

- $REL_{t+h|t}(r) \equiv E\left(\left[p_{t+h|t}(r) - E\left(x_{t+h}(r) | p_{t+h|t}(r)\right)\right]^2\right)$ 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.

- $RES_{t+h|t}(r) \equiv E\left(\left[E\left(x_{t+h}(r) | p_{t+h|t}(r)\right) - E\left(x_{t+h}(r)\right)\right]^2\right)$ 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.

Eq. (16) holds up to an approximation error that involves within bin variation. The proof of eq. (16) as well as the exact form of the approximation error is provided in detail at the end of Appendix B.

Reliability measures the unbiasedness of the probability forecast and resolution measures the ability of the forecast to capture extreme outcomes. Thus, both are, broadly speaking, measuring how good the forecasts are, and thus can be viewed as measures of risk.

The decomposition of the bias squared into reliability and resolution can be estimated as follows. Note that $B_T(r) = REL_T(r) - RES_T(r)$, where:

- $REL(r)$ is obtained as follows. For each t, determine which of the forecast bins $p_{t+h|t}(r)$ falls into. Let $\left\{p^{(k)}_{t+h|t}(r)\right\}$ be the collection of probabilities in the k-th bin and let $p^E_{t+h|t}(r)$ denote the unconditional expected value over the bin. We estimate $p^E_{t+h|t}(r)$ using a Uniform distribution over the bin, so that $p^E_{t+h|t}(r)$ is the midpoint of the bin.[10] In addition, let the number of probabilities in each bin be n_k. Let $\overline{x}_k$ be the average of the realizations conditional on the forecaster having

[10] In 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

made the probability forecast associated with the collection of probabilities in bin k, $\left\{p^{(k)}_{t+h|t}(r)\right\}$. Reliability is the average square calibration error, that is,

$$REL(r) = \frac{1}{T}\sum_{k=1}^{K} n_k \left(p^{E}_{t+h|t}(r) - \overline{x}_k(r)\right)^2 . \tag{17}$$

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 B1 provides simple intuition to understand reliability. The x-axis reports the forecast probability[11], 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.

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

Figures B2 Panel A-C show the results of a VAR estimated with the components of the alternative decomposition.

The practical implementation of the Brier score involves "binning". In our case, the SPF forecast densities are defined over pre-specified bins, each of which corresponds to a given range of values for GDP growth or inflation.[12] 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. These bins are a feature of the SPF dataset, and we take them as given.

are imposed to be equal to the midpoint (so their average is also the midpoint). The appendix derives a 5-terms decomposition which includes within bin variance and within bin covariance. 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, $\overline{p}^{(k)}_{t+h|t}(r)$ (which is the average of the collection of probabilities in the k-th bin, $\left\{p^{(k)}_{t+h|t}(r)\right\}$), replaces $p^{E}_{t+h|t}(r)$ in eq. (17).

[11] The forecast probability is the mid-point of the bin in the forecast distribution.

[12] See the data appendix for a more precise description of the data set.

Figure B1.Reliability Diagram

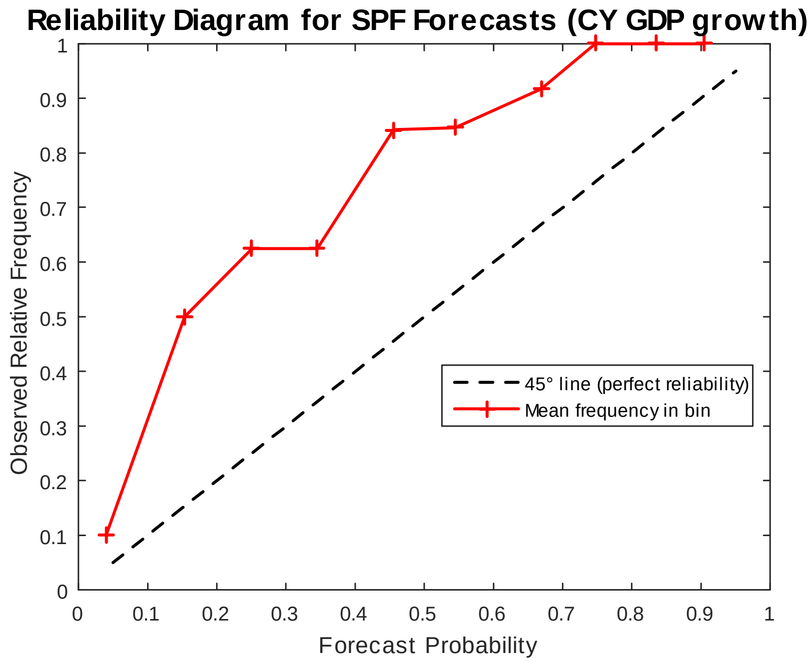

Figure B2. The Effects of Uncertainty on the Economy

Panel A: Reliability

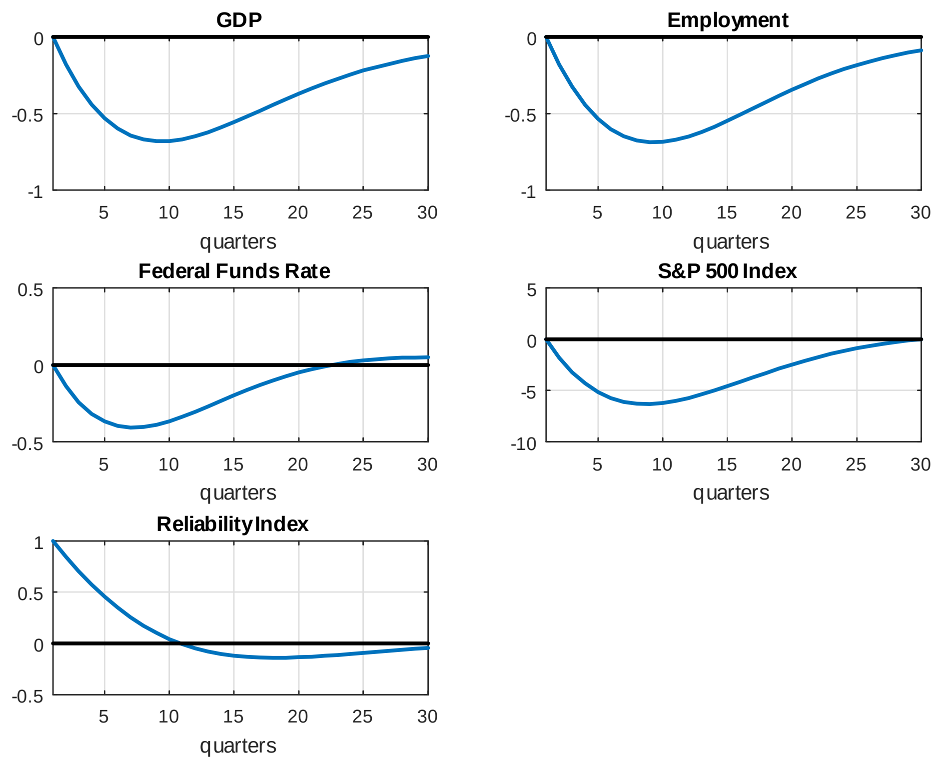

Panel B: Resolution

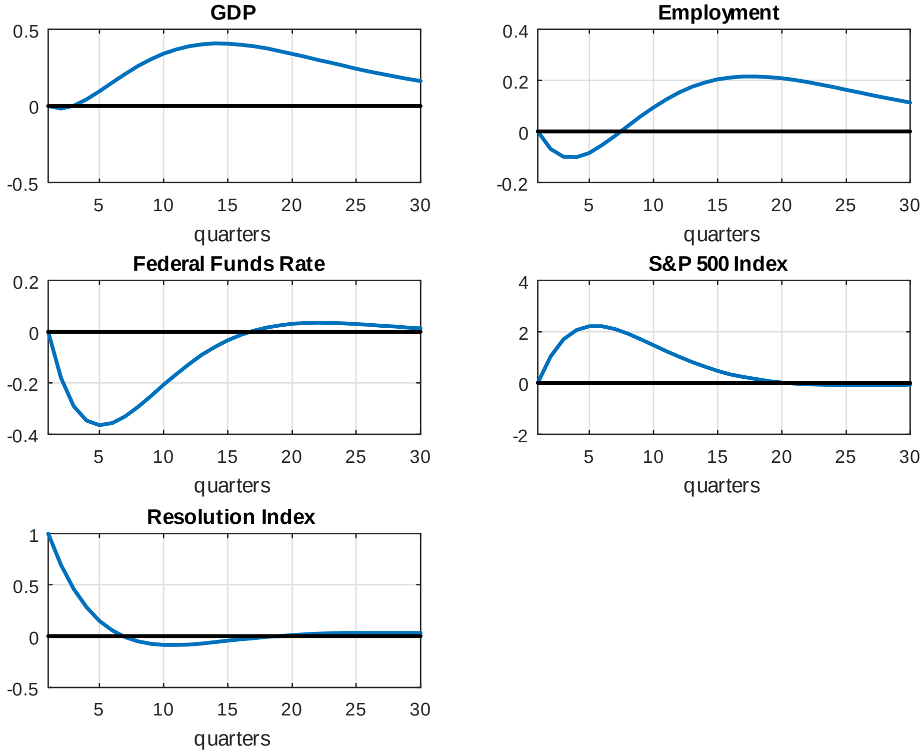

Panel C: Variance

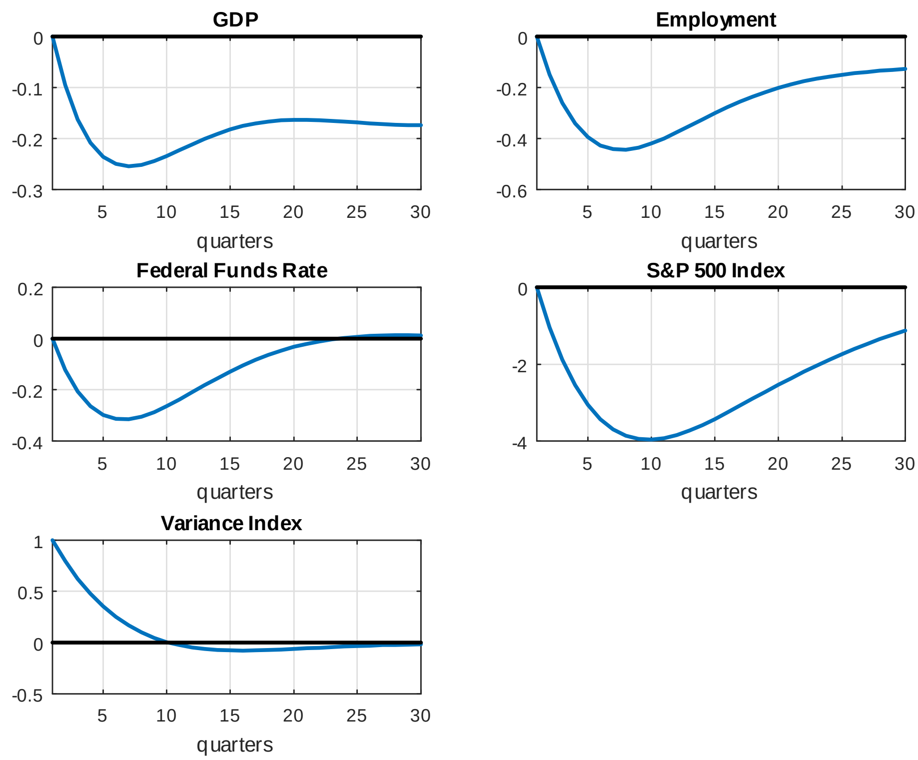

Note: The figure shows the impulse responses for the following uncertainty measures: VXO, JLN, BBD and Scotti. The measures have been standardized.

Proof of eq. (16).. In practice, the Murphy decomposition requires partitioning the range of forecasts – essentially, the [0,1] line – into M sub-segments. Let r be a number along the real line; let $\overline{p}_{\cdot}^{(k)}$ denote the average probability in segment k;[13] and let n_k denote the number of forecast probabilities that fall in the k-th sub-segment, for $k = 1, \ldots, M$. Given all forecasts in the sample, the Brier score can be broken down as follows:

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

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

$$
\begin{aligned}
REL(r) &= \frac{1}{T}\sum_{k=1}^{M}\sum_{j=1}^{n_k}\left[\overline{x}_{t+h}^{(k)}(r) - \overline{p}_{t+h|t}^{(k)}(r)\right]^2 \\
&= \frac{1}{T}\sum_{k=1}^{M} n_k \left[\overline{x}_{t+h}^{(k)}(r) - \overline{p}_{t+h|t}^{(k)}(r)\right]^2. \qquad (18)
\end{aligned}
$$

[13] Alternatively, one could consider $\overline{p}_{\cdot}^{(k)}$ as the midpoint of the k-th segment

The first term can be expressed as follows:

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

Note that because the outcome variable x is binary, the uncertainty term can be expressed as $V(r) = \overline{x}(r)(1-\overline{x}(r))$. To summarize, we have decomposed the Brier score in the following way:

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

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:

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

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:

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

as per the definitions that we have written. ■