

NOT-FOR-PUBLICATION APPENDIX: Macroeconomic Uncertainty Through the Lens of Professional Forecasters

Soojin Jo
and
Rodrigo Sekkel*

June 23, 2017

*Soojin Jo (soojin.jo@dal.frb.org): Research Department, Federal Reserve Bank of Dallas, Rodrigo Sekkel (rsekkel@bankofcanada.ca): Financial Markets Department, Bank of Canada.

This appendix provides a number of supplementary results and robustness checks related to our paper “Macroeconomic Uncertainty Through the Lens of Professional Forecasters”.

Section 1 provides additional evidence on the presence of a factor structure and co-movement in the survey forecast errors of four variables, i.e., GDP, IP, unemployment and housing starts, which are used in our baseline Factor Stochastic Volatility (FSV) model.

Section 2 discusses more on the Markov Chain Monte Carlo algorithm; we present details on prior distribution, starting values and the simulation algorithm; then we report several posterior convergence statistics that evaluate the Bayesian estimation algorithm of our FSV model.

Section 3 details an evaluation of our total volatility estimates for all four variables, and compares them to measures of uncertainty solely based on the disagreement among individual forecasters.

Section 4 presents more sensitivity analysis of our main FSV model: (i) estimates with an AR(1) model for the evolution of volatilities, and (ii) estimates using different consensus forecasts, i.e., different central tendency and the 25th and 75th percentiles of individual forecasters, to calculate forecast errors.

Section 5 compares our baseline macroeconomic uncertainty estimates with other commonly used proxies (VIX and forecast disagreement), and alternative estimates of macroeconomic uncertainty, as in Jurado et al. (2015).

Finally, Section 6 shows supplementary impulse response functions (IRFs) and variance decompositions with different identification choices in our baseline 11-variable VAR. Additionally, we show IRFs and variance decompositions with Bloom (2009) 8-variable VAR, for comparison purposes. Finally, we compare benchmark IRFs to our macroeconomic uncertainty (as documented in Section 5 of the main paper) with those to other widely-used uncertainty measures.

1 Evidence of a Factor Structure in our Data Set

It is worthwhile to discuss some details on our data set first. We use data from the Survey of Professional Forecasters (SPF), conducted by the Federal Reserve Bank of Philadelphia. This is the oldest quarterly survey of macroeconomic forecasts in the U.S. The National Bureau of Economic Research and the American Statistical Association initially introduced the survey in 1968, and then the Federal Reserve Bank of Philadelphia took it over in June 1990. For our benchmark index, we use four variables that have been part of the survey since its beginning: real GDP, unemployment rate, industrial production (IP) and housing starts.

The SPF surveys forecasters on a quarterly basis; the most recent quarter of data in their information set is the previous quarter. Survey questionnaires are sent after the Bureau of Economic Analysis (BEA) *advance report* of the national income and product accounts (NIPA), which contains the first estimates of the previous quarter's GDP. The survey's forecast submission deadline tends to occur around the second to third week of the middle month of a quarter. Thus, before the survey is submitted, forecasters may have access to the first month's of data for initial quarter of the survey. For instance, the Bureau of Labor Statistics releases the unemployment rate in the first week of a month; hence, it will be in the survey respondents' information set. By comparison, housing starts compiled by the Census Bureau have been released during the third week of each month in recent years. It then becomes less clear whether the first month's value in each quarter is included. Similarly, IP has been released near mid-month by the Federal Reserve System Board of Governors during our sample period, and, thus, the inclusion of the first month's value may not occur depending on particular dates in each month.

The calculation of forecast errors at any point in time is contingent on the realized value of the series, and thus, data revisions can ultimately affect our measurement of uncertainty through revisions to the forecast errors. For our analysis, we use the first release and revised data available in one and five quarters after the initial release, as well as the final 2016Q1 vintage data to compute four possible values of forecast errors. It is worthwhile to note here that the degree of the revisions may differ across variables. In addition, other than the first release, each variable from the same vintage may have been revised a different number

of times. For instance, BEA releases three vintages of the current quarterly estimate for GDP, while the Board of Governors of the Federal Reserve System announces each month a preliminary estimate of the previous month’s IP and revised estimates of five preceding months. In one quarter after the first release, GDP would be revised twice after the first “advance” estimate, while the first month value of IP in the previous quarter may have gone through four to five revisions. As such, some monthly indicators likely have undergone more number of revisions after a certain time since the initial release than quarterly one, i.e., GDP.

Now we provide more information about the presence of a factor structure in the forecast errors of the four variables used for our benchmark index. First, Table 1 summarizes the correlation coefficients among the forecast errors the four variables. The forecast errors are calculated as the difference between the median of individual forecasts and the initial data release. The correlations between GDP, IP and the unemployment rate forecast errors are in the range of 40 to 57% (negative correlations for the unemployment rate). The correlations are lower for housing, ranging from 8 to 28%. This smaller degree of comovement between housing and the remaining variables is reflected in our estimated factor loading (see Table 1 in the main text), as well as its higher share of idiosyncratic variance (see Figure 2 in the main text). More formally, we also estimate principal components and find that the first principal component explains 52% of the variability in forecasting errors. Hence, we conclude that there is a fair degree of co-movement across the forecasting errors of these four variables.

Table 1: Correlation Coefficients of Forecasting Errors

	GDP	IP	UR	HS
GDP	1			
IP	0.57	1		
UR	-0.40	-0.54	1	
HS	0.28	0.14	-0.09	1


In addition, we estimate a univariate Stochastic Volatility (SV) model separately for the four forecasting error series. Our factor model is constructed under the assumption that all comovement is captured by a factor and that innovations to the factor as well as

idiosyncratic errors are uncorrelated. Therefore, a factor structure in the levels of forecasting errors implies the same factor structure in volatilities of forecasting errors. Hence, we gauge if there is any similarity across the separately-estimated individual volatility series. We first report in Table 2 the correlation coefficients across the four individually-estimated SV series. The correlations range from 51 to 78%, reflecting close comovements as seen in Figure 1. Moreover, we plot in Figure 1 the average of individually estimated volatility series (in blue): This series shows very similar dynamics to our benchmark macroeconomic uncertainty from the FSV model (in red), with the correlation coefficient of 0.94. The similarity found in the four individually estimated SV series is again consistent with a factor structure imposed in our baseline model.

Table 2: Correlation of Volatility Series Estimated Separately from Univariate SV Models

	GDP	IP	UR	HS
GDP	1			
IP	0.78	1		
UR	0.59	0.65	1	
HS	0.70	0.68	0.51	1

Figure 1: Volatilities from Univariate Stochastic Volatility Models



Note: Figure shows volatility estimates from univariate SV models fitted to each forecasting error of the four macroeconomic indicators (i.e., GDP, IP, unemployment, and housing starts) used in our benchmark model. The blue line represents the average of the four univariate SV series, while the red line plots the benchmark macroeconomic uncertainty index from a FSV model.

2 More on the MCMC algorithm

We describe the setup of the priors as well as the MCMC posterior distribution simulation process in detail. In addition, we present several posterior convergence statistics evaluating the convergence of our estimation algorithm.

2.1 Prior Distributions and Starting Values

Our choice of prior distributions and their parameter values is very similar to Pitt and Shephard (1999). The prior for **factor loadings** is the Normal distribution, i.e., $\lambda_i \sim N(\lambda_0, \Lambda_0)$ with $\lambda_0 = 1$ and $\Lambda_0 = 25$, as in Pitt and Shephard (1999). The choice of large Λ_0 represents a fair degree of uncertainty around the factor loadings. The initial value of the

factor loadings is the OLS estimates of forecast errors on the first principal component as a proxy. Since the factor and loadings are not fully identified in a factor model, we set the loading of the first variable (i.e., GDP) equal to one, a common identification strategy for a factor model. A diffuse Normal prior is used as the prior for the **factor** conditional on $\{h_{f,t}\}_{t=1}^T$, consistent with Equation (3) in the main text (i.e., $f_t \sim N(\theta_0, \Theta_0)$ where $\theta_0 = 0$ and $\Theta_0 = h_{f,t}$).

The prior for the **variability of volatilities** is the inverse Gamma, i.e., σ_f^2 and $\sigma_i^2 \sim IG(\frac{v_0}{2}, \frac{\delta_0}{2})$. We set $v_0 = 1$ and $\delta_0 = 1$, which makes the conditional prior distribution flatter than that in Pitt and Shephard (1999) and more so than those in other recent studies applying SV (see, e.g., Primiceri 2005; Baumeister et al. 2013) to allow for a large time variation for volatilities *a priori*. Compared with the above studies, the total number of parameters is substantially smaller in our case. Thus, we use a more diffuse prior and put a larger weight on data. The prior of each **time-varying volatility** is the log-Normal. For the initial period's stochastic volatility, we have $\log h_0 \sim N(\mu_0^h, V_0^h)$, where $\mu_0^h = 1$ and $V_0^h = 10$ to allow a good chance for the data to determine the posterior distribution.

2.2 Posterior Distribution Simulation

The MCMC algorithm for the estimation of the joint posterior distribution of an FSV model closely follows Pitt and Shephard (1999). We divide the parameters into four blocks: (a) the factor loadings (λ); (b) the time series of the factor ($\{f_t\}_{t=1}^T$); (c) the hyperparameters for volatilities (σ_f ; and σ_i for all i); and (d) the volatility states ($\{h_{f,t}\}_{t=1}^T$ and $\{h_{i,t}\}_{t=1}^T$ for all i). The conditional independence across t and i as well as of f_t and $u_{i,t}$'s further breaks down most steps in the Gibbs sampler, i.e., drawing factor loadings $\{\lambda\}$, volatility states $\{h\}$ and the variance of volatilities $\{\sigma\}$ to drawing from a univariate process. Finally, the volatility states are drawn via Metropolis methods within the Gibbs sampler. Denoting by z^T the time series of a variable z from $t = 1$ to T , we describe the algorithm below.

Factor loadings. Conditional on all other parameters, this step becomes a simple Bayesian regression of forecast errors on the factor with known heteroskedastic error structures. Moreover, because all correlations are by definition captured by the factor, this step further decomposes into the n sub-steps of drawing each i -th loading sep-

arately from $\lambda_i|\varepsilon^T, f^T, \sigma_f, \sigma_i, h_f^T, h_i^T \sim N(\lambda_1, \Lambda_1)$, given the history of variable i , where $\Lambda_1 = (\Lambda_0^{-1} + \sum_{t=1}^T f_t^2/(h_{i,t}))^{-1}$ and $\lambda_1 = \Lambda_1(\Lambda_0^{-1}\lambda_0 + \sum_{t=1}^T f_t \cdot \varepsilon_{i,t}/(h_{i,t}))$.

Factor. Conditional independence also simplifies this step. Given all other parameter values, this step again becomes a Bayesian regression of available forecast errors on factor loadings with known heteroskedasticity for each period t . That is, $f_t|\varepsilon^T, \lambda_i, \sigma_f, \sigma_i, h_f^T, h_i^T \sim N(\theta_1, \Theta_1)$, where $\Theta_1 = \{h_{f,t}^{-1} + \sum_{i=1}^n \lambda_i^2(h_{i,t})\}^{-1}$, and $\theta_1 = \Theta_1(\sum_{i=1}^n \lambda_i \cdot \varepsilon_{i,t}/(h_{f,t}))$.

Innovation variance of volatilities. Since we model each stochastic volatility to follow a unit-root process without a drift, σ^2 is drawn from $\sigma|\varepsilon^T, f^T, h_f^T, h_i^T \sim IG(\frac{v_1}{2}, \frac{\delta_1}{2})$, where $v_1 = v_0 + T$, and $\delta_1 = \delta_0 + \sum_{t=1}^T (h_{i,t} - h_{i,t-1})^2$.

Volatility states. This step further decomposes to the $n + 1$ sub-steps of univariate stochastic volatility draws, based on the Markovian property of stochastic volatility. Following the algorithm by Jacquier et al. (2002) as used in Cogley and Sargent (2005), for each volatility series of an idiosyncratic error i or of the factor, we draw the exponential of volatility ($h_{i,t}^2$) one by one for each $t = 1, \dots, T$, based on $f(h_{i,t}|h_{i,t-1}, h_{i,t+1}, y_i^T, \lambda, f^T, \sigma)$.

Before sampling the states, we first transform forecast errors to be $\varepsilon_{i,t}^* = \varepsilon_{i,t} - \lambda_{i,t}f_t$. Such transformation is unnecessary for the factor. Then, we apply the Jacquier et al. (2002) algorithm for each date, i.e.,

$$\begin{aligned} f(h_{i,t}|(h_i)_{-t}^T, y_i^{*T}, \sigma) &= f(h_{i,t}|h_{i,t-1}, h_{i,t+1}, y_i^{*T}, \sigma) \\ &\propto f(y_{i,t}^*|h_{i,t})f(h_{i,t}|h_{i,t-1})f(h_{i,t+1}|h_{i,t}) \\ &= (h_{i,t})^{-1.5} \exp\left(\frac{-y_{i,t}^*}{2h_{i,t}}\right) \exp\left(\frac{-(\log h_{i,t} - \mu_t)^2}{2\sigma_c^2}\right), \end{aligned}$$

where μ_t and σ_c^2 are the conditional mean and variance of $\log h_{i,t}$, respectively. Under the unit-root specification of this paper, they can be calculated as

$$\begin{aligned} \mu_t &= \frac{(\log h_{i,t-1} + \log h_{i,t+1})}{2}, \\ \sigma_c^2 &= \frac{\sigma_i^2}{2}, \end{aligned}$$

for $t = 1, \dots, T - 1$. Thus, a trial value of $\log h_{i,t}$ is drawn from the Normal distribution with mean μ_t and variance σ_c^2 . For the beginning and end periods of each series i , the

following conditional mean and variance are used instead:

$$\begin{aligned} t = 0 : \quad & \sigma_c^2 = \frac{\sigma_i V_0^h}{\sigma_i + V_0^h}, & \mu &= \sigma_c^2 \left(\frac{\mu_0^h}{V_0^h} + \frac{\log h_{i,t+1}}{\sigma_i^2} \right), \\ t = T : \quad & \sigma_c^2 = \sigma_i^2, & \mu &= \log h_{i,T-1}. \end{aligned}$$

After obtaining a draw, the acceptance probability is evaluate using the conditional likelihood $f(y_{i,t}^* | h_{i,t})$, completing a Metropolis step (see Cogley and Sargent 2005 for details).

We can summarize the estimation procedure as follows:

1. Assign initial values for λ , f^T , σ_f , σ_i for all i , h_f^T , and h_i^T for all i .
2. Draw λ from $p(\lambda | \varepsilon^T, f^T, \sigma_f, \sigma_i, h_f^T, h_i^T)$.
3. Draw f^T from $p(f^T | \varepsilon^T, \lambda, \sigma_f, \sigma_i, h_f^T, h_i^T)$.
4. Draw σ_f and σ_i s from $p(\sigma | \varepsilon^T, f^T, h_f^T, h_i^T)$.
5. Draw h_f^T , and h_i^T s from $p(h | \varepsilon^T, f^T, \sigma_f, \sigma_i)$.
6. Go to step 2.

We iterate over the Metropolis-within-Gibbs sampler a total of 80,000 times, discarding the first 30,000 draws. Then we store every 10th draw in order to avoid potential autocorrelation across draws, and finally obtain 5,000 draws from the joint posterior distribution.

2.3 Convergence of the Markov Chain Monte Carlo Algorithm

Before we begin, it is important to note that different initial values, changes in the number of “burn-in” samples and that of posterior draws do not alter our baseline results. To conserve space, we do not report all results in this Appendix.

For the next two convergence statistics, we put together collected posterior draws in the following order: factor loadings (the second to the fourth element of λ , corresponding to points 1 to 3 in the x-axis of the Figure 2), the time series of a factor ($\{f\}_{t=1}^T$, corresponding to points 4 to 193), the time series of indicator-specific and macroeconomic volatilities ($\{h\}_{t=1}^T$ for all h_i 's and h_f , corresponding to points 194 to 1,143 representing 5 series for 190 periods), and the variability of the volatilities (σ_i 's and σ_f , corresponding to points 1144 to 1148).

First, we compute the autocorrelation function of posterior draws, to see how well the

chain mixes. Following Primiceri (2005), we show the 20-th-order sample autocorrelation of the draws in the first panel of Figure 2. All parameters in the model have very low autocorrelations, i.e., below 0.005.

Second, we show that our MCMC algorithm converges to the ergodic distribution by calculating the Inefficiency Factors of posterior draws, as in Primiceri (2005) and Baumeister and Peersman (2013). That is, we compute the inverse of the relative numerical efficiency measure (RNE) as introduced in Geweke (1992):

$$RNE = (2\pi)^{-1} \frac{1}{S(0)} \int_{-\pi}^{\pi} S(\omega) d\omega, \quad (1)$$


where $S(\omega)$ denotes the spectral density of the collected draws of each parameters at frequency ω . The second panel of Figure 2 shows the inefficiency factors: Most parameters and state variables have the factor values lower than 0.2 with a few exceptions. However, the values do not go above 0.5 even in those cases, which is far below the value of 20 assigned as satisfactory in Primiceri (2005).

3 Evaluation of the Macro Uncertainty Series


It is difficult to evaluate the accuracy of an uncertainty measure, since uncertainty is unobservable. However, one indirect way of evaluating our baseline estimates of macroeconomic uncertainty is to exploit density forecasts for the four variables used in the FSV model. We define uncertainty as the time-varying conditional volatility of forecasting errors. Following ? and Clark (2012), we look at coverage rates based on the variances implied by our FSV model for each of the variables in our FSV model. That is, using the median posterior draws of the total volatility for the four variables in our model, we construct α -percent prediction intervals and report the frequency within which real-time GDP, IP, Unemployment and Housing Starts' initial releases fall. A frequency of more (less) than α -percent means that on average over our sample period, the forecast density based on our volatility estimates is too wide (narrow).

We choose α to be 50, unlike previous studies where 70-percent intervals were used. A 50-percent intervals makes it possible to compare our interval with ex-ante uncertainty

Figure 2: Convergence Diagnostics



(a) Autocorrelations




(b) Inefficiency Factors

Note: Figure plots the 20th-order sample autocorrelations of posterior draws (panel (a)) and inefficiency factors (panel (b)) of the collected posterior draws of all parameter and state variables.


bands constructed using SPF forecast disagreements (i.e., interquartile ranges) as in Section 2. Figure 6 reports forecast errors for each variable (dotted black line) along with the 50 percent intervals, i) one based on estimates from our FSV model (in red), and ii) the other obtained from forecast disagreements (in blue). The two prediction intervals have similar dynamics over time, while the ones using disagreements are usually narrower and more volatile than the ones based on our uncertainty estimates. Since the intervals based on the uncertainty estimates are wider, it is more likely that forecast errors would fall within these.

Nonetheless, it is still important to check statistically how close the frequency is to 50%. Table 3 reports t-test results for the frequencies that actual outcomes fall within the 50% bands using intervals based on estimates from our FSV model (FSV), and forecast disagreements (Disagreement). It also shows in parentheses p-values for the null hypothesis of correct coverage (i.e., empirical = nominal rate of 50 percent) based on t-statistics. In the last column, the table reports the p-values for the null hypothesis that the empirical frequencies obtained from the two intervals are the same.


Figure 3: Uncertainty Bands and Forecast Errors



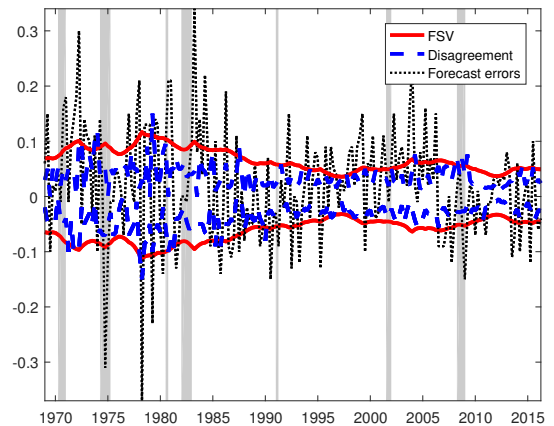
(a) GDP



(b) IP



(c) Unemployment



(d) Housing

Note: Figure plots prediction intervals based on conditional volatilities estimated from the baseline FSV model (in red solid lines) as well as cross-sectional interquartile ranges of survey forecasts (in blue dashed lines). The black dotted lines represent forecast errors measured as the differences between median SPF and initial data releases.

While frequencies associated with both intervals are all lower than 50%, the ones based on the uncertainty estimates from the FSV model are larger for all four variables. In addition, the p-values show coverage rates of GDP and housing starts are not statistically different from 50%. All considered, this supports the notion that on average, our index captures uncertainty reasonably well from a density forecasting perspective of the variables

Table 3: Forecast Coverage Rates, 1968Q4 - 2016Q1

	FSV	Disagreement	H0: same frequency
GDP	0.4789 (0.5630)	0.2684 (0.0000)	0.0000
IP	0.4263 (0.0419)	0.3263 (0.0000)	0.0100
Unemployment	0.4158 (0.0199)	0.3579 (0.0001)	0.0935
Housing	0.4526 (0.1924)	0.2895 (0.0000)	0.0000

Note: Table shows empirical coverage rates for actual outcomes falling within 50-percent prediction intervals based on i) the median posterior draws of total volatilities of each variable estimated from our FSV model (FSV) and ii) forecast disagreement of each variable (Disagreement). Parentheses show are p-values for the null hypothesis of the correct coverage, i.e., empirical = nominal rate of 50 percent, based on t-statistics. The last column reports the p-values for the null hypothesis that the empirical frequencies obtained from the two intervals are the same.

used to estimate our factor model.

4 Sensitivity Analysis

4.1 Alternative Model Specification: Autoregressive SV

Stochastic volatility has been adopted in macroeconomic literature in various specifications. While a unit root process has been widely imposed to the volatility process in studies such as in Clark (2012), Primiceri (2005), and ?, other papers have used an autoregressive (AR) SV process for volatility to evolve (see Mumtaz and Zanetti 2013 and Jo 2014, for example). Our baseline model assumes unit root processes for both indicator-specific and macroeconomic volatilities. However, we investigate how uncertainty estimates change when we use the AR SV specification instead of the random walk process for indicator-specific and macroeconomic volatilities, as follows:


$$\log h_t = \alpha + \beta \log h_{t-1} + \sigma \eta_t, \quad (2)$$

where cross-sectional subscripts are omitted for brevity. For both AR coefficients, we use the Normal distribution as a prior, with prior mean and standard deviation set as (-0.3, 0.4) and (0.8, 0.4) for α and β , respectively, for all common and indicator-specific volatility

series.

Table 4 summarizes the median posterior draws of the AR coefficients as well as the standard deviations of the draws. The median draws of β range from 0.48 to 0.86 with the macroeconomic stochastic volatility process having the highest slope coefficient. The resulting index of macroeconomic uncertainty is plotted in the blue dotted line in Figure 4. We find that switching from the unit-root to the AR process leads to a more volatile series, as expected from the posterior draws of the slope coefficient that is lower than 1, while overall dynamics remain close to our baseline index. As noted by Primiceri (2005), the unit root assumption impacts the estimated volatility series; also, the AR model specification captures many temporary shifts in uncertainty in addition to most significant increases.

Figure 4: Macroeconomic Uncertainty: Unit-Root vs AR SV models



Note: Figure plots macroeconomic uncertainty estimates obtained from an alternative specification of SV processes, i.e., autoregressive SV processes in the blue dotted line. The black solid line is our baseline uncertainty index with a unit-root SV specification.

However, one critical issue we find with the AR specification is that the parameter estimates depend to a great extent to the prior distribution specification. One place to see this is the median posterior draws of α , which do not vary much from the prior means (0.3) as shown in Table 4, reflecting high impacts of the priors. In addition, the volatility of the

macroeconomic uncertainty series changes dramatically to a small change in priors, while the three main peaks are preserved. This may have something to do with the fact that when shifting from the random walk assumption to the AR model, we are adding more parameters to be estimated without providing any new piece of information. This results lead us to conclude that the parsimonious unit-root specification is more suitable for our purpose.

4.2 Uncertainty Based on Different Measures of Central Tendency and Quantiles of Survey Forecasts

When calculating the forecasting errors for the baseline model, we use the consensus forecasts, i.e., the median across forecasters, to minimize potential influences from individual forecasting biases, as noted in Section 2 of the main paper. Capistrán and Timmermann (2009) provide support for the use of the median; they show that means of survey forecasts performs better than other methods of individual forecast combination when evaluated by root-mean-squared errors. In the case of SPF, the mean and median forecasts are very close to each other, as documented in Croushore (2010). In addition, Arai (2014) finds that the SPF median forecasts for GDP growth present no systematic biases. Hence, a number of previous studies consider the median as standard consensus projections.¹

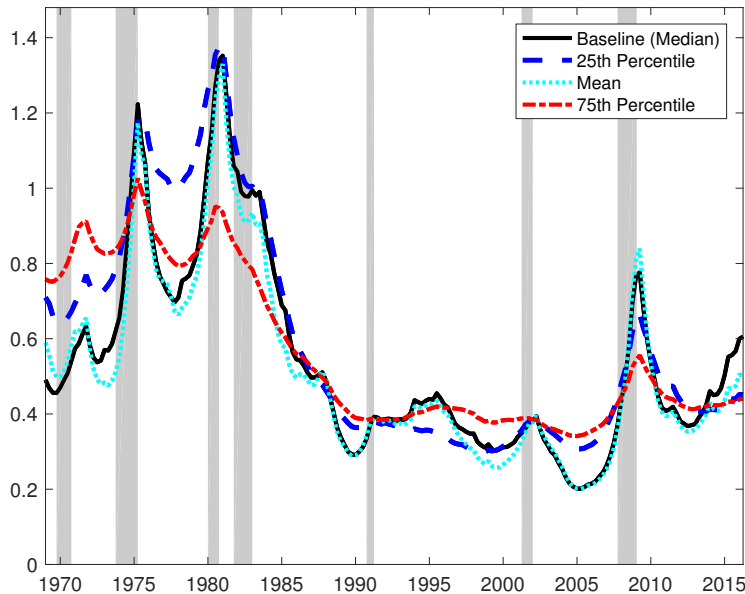
However, studies such as Gneiting (2011) show that quantiles might also be optimal forecasts when the loss function is generalized piecewise linear. Therefore, we re-estimate our FSV model using the 25th and 75th percentiles of the forecast error distribution, instead of the median as reported in the main text. In addition, we use the means of forecasts instead of the medians and reestimate the FSV model.

Figure 5 shows the resulting estimates. First, we find that the series based on the mean and our baseline index are very close to each other, with a correlation coefficient of 0.99. This is in line with a notion that the mean and median are very close to each other in the SPF. The uncertainty estimates based on the 25th and 75th percentiles of forecasts have similar dynamics with our baseline index, with the correlation coefficients amounting to 0.95 and 0.84, respectively. The differences in index appear to be most notable in the

¹See Engelberg et al. (2009) and Croushore (2010), for example.

earlier period of the sample, supporting the use of the median as consensus to remove idiosyncratic variations in individual forecasts.

Figure 5: Uncertainty Based on Different Optimal Forecasts

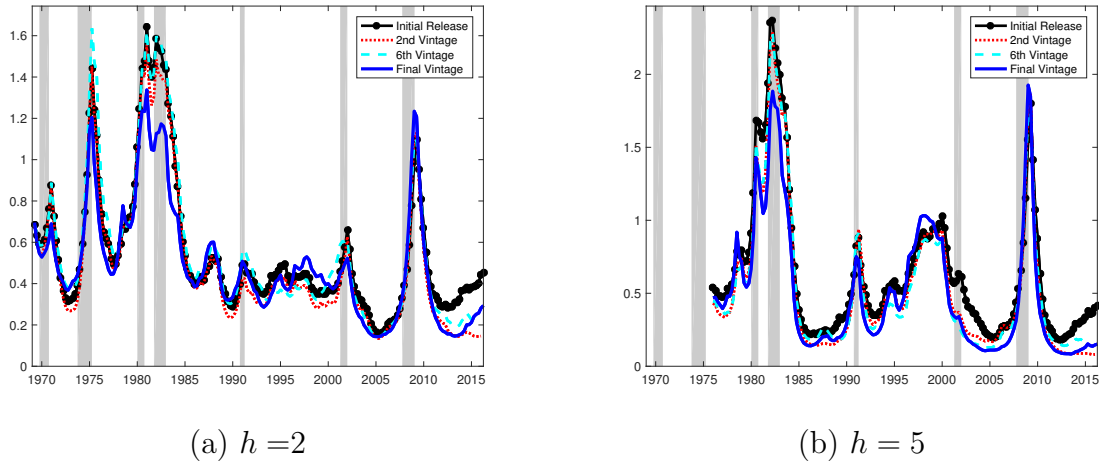


Note: Figure plots uncertainty estimates based on different quantiles in comparison with our baseline index that is based on median forecast errors. The blue, cyan, and red dashed lines represent the estimates based on the 25th percentile, mean, and 75th percentile forecasts, respectively.

4.3 Effects of Data Revisions on Uncertainties of Longer-Run Horizons

We further investigate how data revisions affect the long-horizon uncertainty estimates for two- and five-period-ahead forecasts. Similar to Figure 2 in the main text, Figure 6 shows the long-horizon macroeconomic uncertainty indexes based on the first-released data, the data revised in one and five quarter(s) after the initial release, and the final vintage data. We find that the size of differences across the indexes is much smaller for both longer-run horizons than nowcasts. Although the impacts of data revision is less pronounced, the key findings based on the nowcast uncertainty indexes still apply. In other words, the level of the second vintage-based index tends to be slightly lower than that of the first release-based

Figure 6: Effects of Data Revisions on Uncertainties of Longer-Run Horizons



Note: Figure compares estimated h -step-ahead macroeconomic uncertainty factor based on the first release with those based on data from different vintages, for $h = 2$, and 5. The black line with circles is the uncertainty index based on the initial release, the red dotted line is the one where forecast errors of each variable are computed using the revised data available in one quarter after the initial release, the cyan dashed line, the 6th vintage, and the blue solid line, the final vintage data.

index, and the relative size of the peaks changes depending on the data vintage.

5 Comparison with Other Common Measures of Macroeconomic Uncertainty

Various attempts have been made recently to quantify macroeconomic economic uncertainty: we compare our baseline measure with a measure based on forecast disagreements, the VIX, and finally, the one introduced by Jurado et al. (2015).

Before we offer a detailed comparison, a discussion regarding forecast horizons is in order. First, the horizons of forecast disagreements are identical to that of forecasts used in our baseline index: the disagreements are measured as interquartile ranges of the forecasts across professional forecasters, while our FSV model used the medians of the same forecasts.² Second, the VIX is implied volatility calculated by using the midpoint of real-time S&P 500 stock index option bid/ask quotes, representing the market's expectations about

²As noted in Section 2 of the main text, our baseline index is constructed using forecasting errors of one-step-ahead forecasts, namely nowcasts.

the volatility of the index over the next 30 days. Since it is the implied, expected volatility, it is different from ex-post or realized volatility, and, thus, the comparison we make in the paper is rather limited. Third, Jurado et al. (2015) provides three monthly indexes based on forecasts of different horizons, i.e., one-, three-, and twelve-month ahead. We construct a quarterly series by taking averages of the three-month-ahead uncertainty index. Again, forecast errors does not exactly align, but the focus of our discussion is to compare over all time-series dynamics of our uncertainty index and other widely-used measures.

Finally, we also compare our baseline index to stochastic volatility series estimated only using real GDP growth forecast errors. This will show clearly how other variables in the model contribute to the volatility of a factor, given the representative nature of real GDP growth for the current economic state.

Forecast disagreement

A widely used proxy for uncertainty based on survey forecasts is forecast disagreement, commonly measured as the interquartile range (e.g., Bachmann et al. 2013). Underlying this practice is the assumption that predictions of forecasters are more likely to be close to one another when economic uncertainty is low. However, forecast disagreement may just reflect heterogeneous, but not uncertain, beliefs.³ Figure 7 plots the forecast disagreements of the four economic series underlying our baseline model, along with the total variations of median forecast errors calculated as in equation (5) in the main text. The forecast disagreement is measured as the interquartile ranges, i.e., the 75th percentile minus the 25th percentile of individual forecasts.

In general, for each indicator, both disagreement and estimated total variations from our model exhibit similar dynamics over time. However, disagreements are notably more volatile than total variations. In addition, all variables show episodes where the disagreement and total variation deviate substantially from each other.

We further investigate the relationship between our macroeconomic uncertainty measure and an analogous, disagreement-based proxy of macroeconomic uncertainty. We take the averages of disagreement of the four indicators to construct a comparable series. Panel

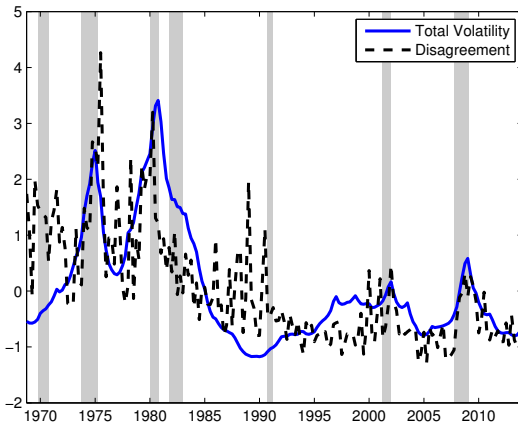
³See, for example, Mankiw et al. (2004), Lahiri and Sheng (2010), and Sill (2012) for more detailed discussions of measuring uncertainty using forecast disagreement.

(a) in Figure 8 presents the resulting average disagreement series along with our baseline measure of macroeconomic uncertainty. The average disagreement series evolves closer to our macroeconomic index, compared to the dynamics of individual disagreements; the three major spikes in disagreement coincide with the three main episodes of uncertainty increases in our baseline index.


VIX

Since employed in Bloom (2009), the VIX index has been one of the most commonly used measures of uncertainty. The VIX index quantifies implied volatility using a weighted average of the S&P 500 call-and-put option prices for the next 30 days and is available on a daily basis. Thus, it mainly captures the financial volatility. Panel (b) in Figure 8 plots the quarterly averages of the VIX along with our baseline measure of macroeconomic uncertainty after standardization for comparison. The VIX is more volatile and has more peaks, some of which are specifically due to financial market episodes that may not necessarily spill over to the real economy. More importantly, the VIX collectively reflects both stock market uncertainty and risk aversion, and, as a result, some jumps are mainly due to increases in risk aversion rather than uncertainty, as shown in Bekaert et al. (2013).


Figure 7: Stochastic Volatility and Dispersions




(a) GDP



(b) IP



(c) Unemployment



(d) Housing Starts

Note: Figure compares forecast disagreement from the SPF to total variation of underlying individual series calculated as equation (5) in the main text. The forecast disagreement is measured as the interquartile ranges of individual forecasts. All series are standardized for comparison.

Jurado et al. (2015)'s macroeconomic uncertainty index

The macroeconomic uncertainty measure constructed in Jurado et al. (2015) is based on the same definition of uncertainty: It is the time-varying conditional second moment of unexpected variations that affects forecasting errors of a number of economic indicators

simultaneously. Our approach and Jurado et al.'s measure share an additional similarity, as both employ a stochastic volatility model as well as a factor structure. However, Jurado et al. use the factor structure mainly for forecast generation: A large (dynamic) factor model is built encompassing 132 macroeconomic and 147 financial series, where the forecasts for all series are computed. These forecasts are used to calculate forecast errors of each series, after which a univariate stochastic volatility process is fitted to 132 macroeconomic series. Finally, the estimated volatilities are averaged to form a macroeconomic uncertainty index.

Thus, our approach contrasts with that of Jurado et al. (2015) in that, in our approach, (i) forecasts are simply taken from a survey and (ii) macroeconomic and indicator-specific uncertainties are jointly modeled and estimated in one step. Additionally, the number of underlying economic indicators is drastically different: While our measure is based on four indicators, Jurado et al. take into account close to 300. Despite the differences, it is surprising that both measure show very similar dynamics and identify the same three main episodes of high uncertainty, as illustrated in Panel (c) in Figure 8. This reflects the extensive information content in survey forecasts: Zarnowitz and Braun (1993) document that professional forecasters in the SPF generate forecasts based on various econometric models, leading indicators and surveys as well as a number of available macroeconomic and financial data, integrating information from a wide range of sources. Consequently, our approach removes predictable variations very effectively, providing a parsimonious yet efficient tool for the estimation of uncertainty.


Stochastic Volatility of Real GDP Growth Forecast Errors

One other way of employing survey forecasts to gauge economic uncertainty is to use a panel of density forecasts of individual forecasters and aggregate those. For instance, Lahiri and Liu (2006) and Lahiri and Sheng (2010) derive uncertainty from the density forecasts of the real GDP growth and inflation from the SPF, and investigate the driving factors as well as relationship to other uncertainty measures. Clements (2014) investigates the relationship between ex-ante density forecast and ex-post forecast errors and finds that forecasters are overconfident at longer horizons, but overstate the uncertainty around their short term forecasts. Unfortunately our baseline uncertainty index cannot be directly compared to


the ones based on the density forecasts due to data deficiency: SPF provides density forecasts only for the output and inflation. Instead, given the representative nature of real GDP growth for the current economic state, we estimate a univariate SV model for GDP forecast errors from the SPF and compare it to our baseline index.

Panel (d) of Figure 8 shows the SV estimates of GDP forecast errors. Although GDP is a representative economic indicator that contains important information regarding economic condition, the Figure illustrates that macroeconomic uncertainty may differ from those captured solely by unexpected variations in GDP. The factor loading estimates in Table 1 in the main text further support this view: The factor loadings for IP and unemployment rate forecast errors are of similar size as that of GDP, suggesting the importance of including other variables for the measurement of macroeconomic uncertainty.


Figure 8: Comparison with Other Uncertainty Measures



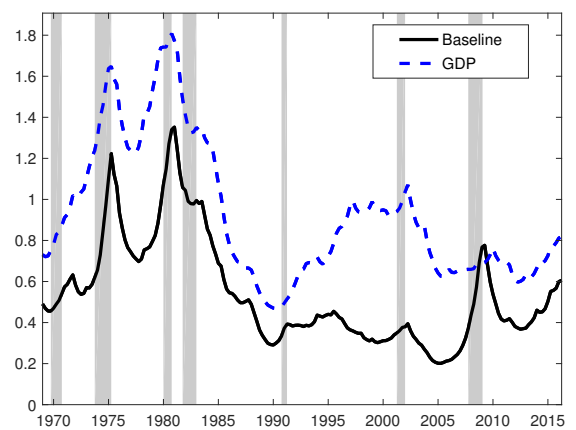
(a) Disagreement



(b) VIX



(c) JNL



(d) GDP SV

Note: Figure compares our baseline macroeconomic series to other widely used indexes of uncertainty. Panel (a) plots the averages of forecast disagreement of the four economic indicators; (b) the VIX; (c) the Jurado et al. (2015) uncertainty series; and (d) univariate SV estimated using GDP forecasting errors, along with our baseline measure, respectively. All series are standardized.

6 Innovations to Uncertainty and Macro Dynamics

In Section 5 of the main text, we show how innovations to our macroeconomic uncertainty estimates impact GDP and employment in the 11-variable VAR proposed by Christiano

et al. (2005). Consistent to the VAR in Jurado et al. (2015), our benchmark specification in the main text places the macroeconomic uncertainty index last. In this section, we examine how our main results change, when a different identification scheme is adopted. Namely, here we show IRFs and variance decompositions when our measure of macroeconomic uncertainty is ordered first in the same VAR. We also report below IRFs and variance decomposition in the 8-variable VAR proposed by Bloom (2009). To facilitate comparison, we lastly report how our IRFs compare when we substitute our uncertainty measure for other commonly used macroeconomic uncertainty estimates and proxies, such as the VIX, forecast disagreement or the Jurado et al. (2015) macroeconomic uncertainty estimates.

6.1 Christiano et al. (2005) VAR with long-run uncertainty and a different ordering

Below we show in impulse response functions and variance decompositions in the same VAR specification, but with our baseline macroeconomic uncertainty index ordered first. Panel (a) in Figure 9 is similar to Figure 8 in the main paper, which plots the IRFs from the VAR with uncertainty ordered last, while panel (b) shows the IRFs from the VAR where the same uncertainty index is ordered first. Additionally, we also add the IRFs obtained by replacing our baseline uncertainty index (based on one-period-ahead, i.e., nowcast errors) by long-run uncertainty estimates (as shown in Section 4.1 of the main paper) for both specifications.⁴

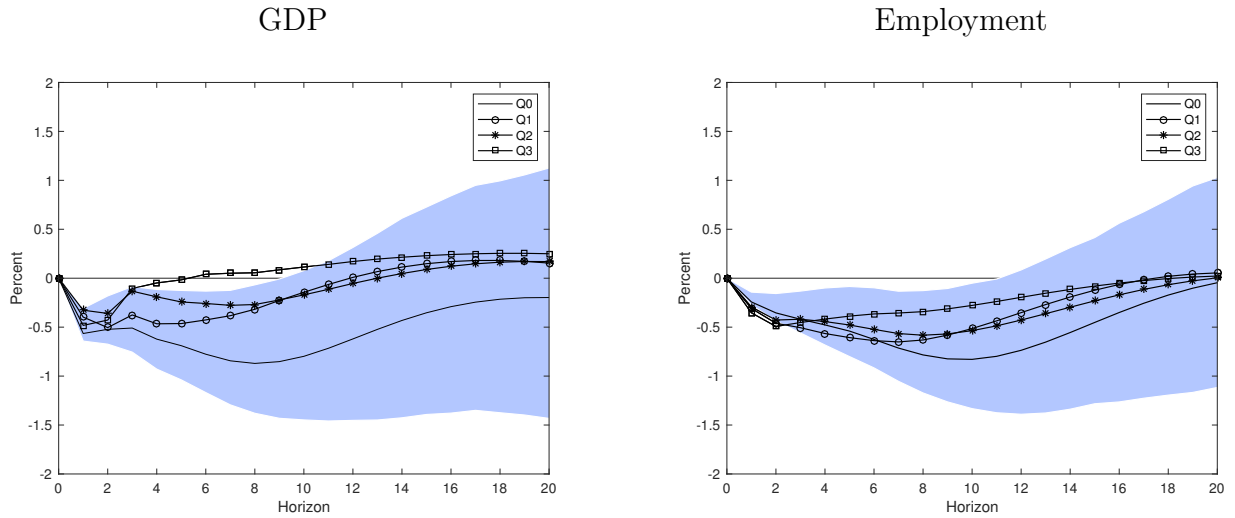
As one would expect, when the uncertainty series is placed first in the system, we find a larger impact of macroeconomic uncertainty shocks on GDP and employment than in our benchmark specification in the main text. For example, a four standard deviation shock to the baseline macroeconomic uncertainty index yields slightly smaller than a 1%-drop in GDP at its peak (panel (a)), whereas panel (b) shows that it falls by about 2.8% when macroeconomic uncertainty is ordered first in the same VAR. Longer-horizon uncertainties did not alter the IRFs significantly in both specifications.

The stronger impact of macroeconomic uncertainty in this alternative VAR ordering

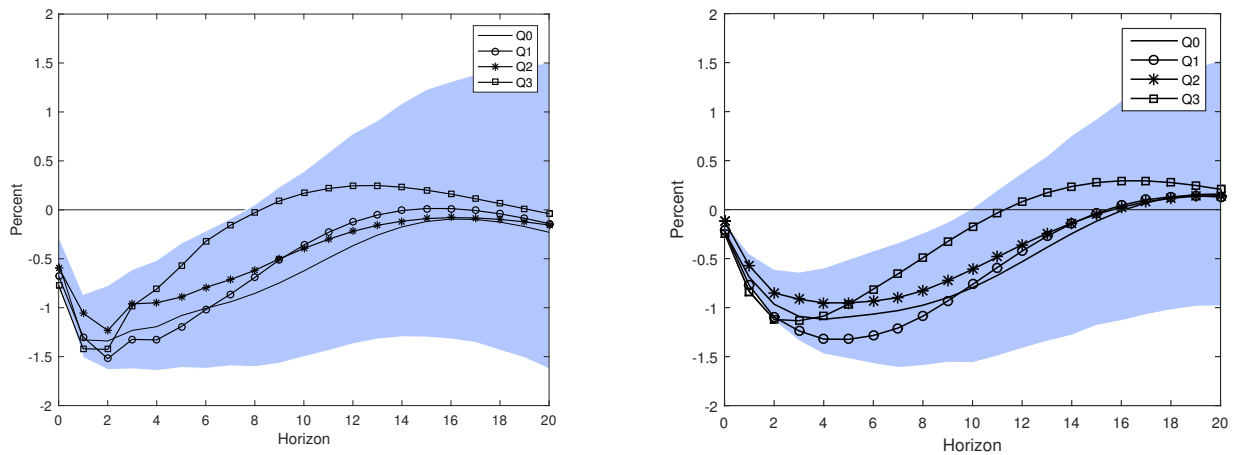
⁴We exclude five-step-ahead uncertainty estimates, as these are available only from 1975Q4, to facilitate comparison.

is also reflected in the variance decomposition estimates. Table 6 shows that the baseline macroeconomic uncertainty index accounts for a maximum of 16.5 and 15.7% of GDP and employment variability, respectively, when uncertainty is ordered first in the VAR. Across all horizons, the shares are larger than those presented in Table 5 using the VAR with uncertainty ordered last.

Figure 9: Responses to Innovations to Macroeconomic Uncertainty in the 11-variable VAR



(a) Uncertainty Ordered Last (Baseline)



(b) Uncertainty Ordered First

Note: Figure shows the response of GDP and Nonfarm Payroll Employment to a four standard deviation macroeconomic uncertainty shock in a 11-variable VAR with the following variables and ordering: Uncertainty, GDP, Employment, consumption, real private investment, wages, hours, Fed Funds rate, profits, and M2. Q0 to Q3 denote IRFs to innovations in uncertainty about nowcast, one-, two-, and three-quarter-ahead forecasts, respectively. The shaded area represents 68% confidence bands using Kilian (1998)'s bootstrap-after-bootstrap method.

Table 4: AR Coefficients Estimates

	GDP	IP	Unemployment	Housing	Factor
α	-0.307 (0.625)	-0.301 (0.635)	-0.306 (0.635)	-0.293 (0.636)	-0.314 (0.634)
β	0.652 (0.134)	0.855 (0.075)	0.750 (0.097)	0.475 (0.142)	0.858 (0.070)

Note: Table reports the median posterior draws of the AR coefficients of each AR SV series. Numbers in parenthesis are the standard deviations of the draws.

Table 5: Variance Decomposition for the 11-variable VAR with Uncertainty Ordered Last

Quarters	<i>GDP</i>				<i>Employment</i>			
	Q0	Q1	Q2	Q3	Q0	Q1	Q2	Q3
2	2.34	1.18	0.78	1.77	1.39	2.34	2.09	3.11
4	2.57	1.77	0.78	1.33	2.14	3.50	2.72	3.49
8	3.79	1.82	0.66	0.56	3.41	4.14	2.86	2.14
20	4.28	1.94	1.01	1.83	4.80	4.14	2.99	3.91

Note: Table shows forecast error variance decomposition for our baseline 11-variable VAR with following variables and ordering: GDP, Employment, consumption, real private investment, wages, hours, fed funds rate, profits, M2 and our macroeconomic uncertainty series. The VAR is estimated in levels with 3 lags. All variables, except the Fed Funds rate and our uncertainty measure, enter the VAR in logs. Q0 to Q3 denote uncertainty about nowcast, one-, two-, and three-quarter-ahead forecasts, respectively.

Table 6: Variance Decomposition for the 11-variable VAR with Uncertainty Ordered First

Quarters	<i>GDP</i>				<i>Employment</i>			
	Q0	Q1	Q2	Q3	Q0	Q1	Q2	Q3
2	15.15	16.04	10.73	19.81	12.43	15.90	8.34	18.86
4	16.22	19.86	11.96	17.53	15.66	21.03	11.62	20.12
8	12.35	15.48	9.01	8.69	13.96	19.85	10.53	12.01
20	16.46	19.86	12.76	20.47	15.66	21.59	11.62	21.40

Note: Table shows forecast error variance decomposition for our baseline 11-variable VAR with following variables and ordering: Uncertainty, GDP, employment, consumption, real private investment, wages, hours, fed funds rate, profits, and M2. Q0 to Q3 denote uncertainty about nowcast, one-, two-, and three-quarter-ahead forecasts, respectively.


6.2 IRFs with Macroeconomic Uncertainty Calculated with Different Data Vintages

Section 4.4 of the main paper showed how data revisions can affect the measurement of uncertainty, as using different data vintages may change the size and direction of forecast errors. Our baseline macroeconomic uncertainty measure is estimated with forecast errors computed using the first data release, but we also show in Section 4.4 of the main paper the uncertainty series using the forecast errors calculated using different vintages, i.e., the second, sixth, and the final vintages. Among those, we select the series based on the second and final vintages, and re-run the 11-variable VAR proposed by Christiano et al. (2005), with uncertainty ordered last.

Figure 10 plots the IRFs of GDP and employment following a 4-standard-deviation shock to different uncertainty series. While it is common across all measures that a shock to uncertainty generates negative responses to GDP and employment, the magnitudes of the responses differ; the uncertainty based on the second data vintage yields less negative responses of both GDP and employment, but the final vintage-based uncertainty generates more stronger and persistent responses relative to the baseline case.

These results shows a further implication of the impact of data revisions. Depending on a particular data vintage chosen, one likely sees differential effects of an uncertainty shocks, in addition to varying estimates of uncertainty, as shown in Section 4.4 of the main text.

Figure 10: Responses with Measures of Uncertainty Estimated with Different data Vintages



Note: Figure shows the response of GDP and employment to four-standard-deviation uncertainty shocks in a 11-variable VAR, but with our uncertainty measure estimated with forecast errors for the individual series calculated with different choices of data vintages, i.e., the second and sixth vintages. The shaded area represents 68% confidence bands for our baseline results using Kilian (1998)'s bootstrap-after-bootstrap method.


6.3 Comparison of IRF with Uncertainty Estimated Recursively

In Section 4.3 of the main paper, we report a recursively-estimated uncertainty series. That is, we run another set of estimation with expanding samples by adding one set of observations (i.e., real-time quarterly forecast errors for the four indicators) at a time after the first 20 quarters. We then constructed an uncertainty series by retaining the median draw of the volatility of a common factor for the latest period. Here, we run the 11-variable VAR proposed by Christiano et al. (2005), with uncertainty ordered last, using the recursively-estimate uncertainty measure.

Figure 11 plots the IRFs of GDP and employment following a 4-standard-deviation shock to uncertainty. The new set of IRFs differs most from those using the baseline uncertainty index in three to four quarters after the impact: while the baseline IRFs present immediate negative responses of GDP and employment, the new results suggest that it may take a while for an uncertainty shock to have negative impacts on economic activity, if we

measure macroeconomic uncertainty recursively by adding one set of observations at a time.

Figure 11: Responses with Uncertainty Estimated Recursively (Fig. 5 in the main paper)



Note: Figure shows the response of GDP and employment to four-standard-deviation uncertainty shocks in a 11-variable VAR, but with our uncertainty measure estimated recursively (see main text, Figure 5). Namely, the second and sixth releases. The shaded area represents 68% confidence bands using Kilian (1998)'s bootstrap-after-bootstrap method.

6.4 Bloom (2009) VAR

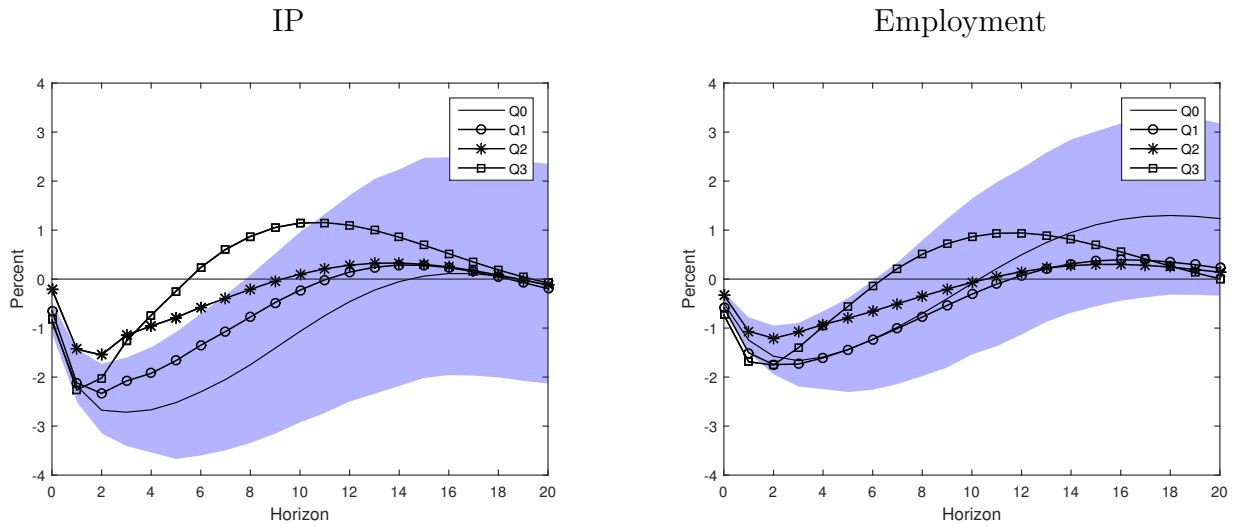
In this subsection, we estimate IRFs and variance decompositions from an alternative, smaller VAR as introduced in Bloom (2009). The VAR is comprised of the following 8 variables: S&P index, our baseline macroeconomic uncertainty index, the fed funds rate, wages, CPI, hours, employment and IP. Contrary to Bloom (2009) and in accordance to Jurado et al. (2015), we do not HP filter the variables before estimating the VAR. As in the previous section, we also examine IRFs and variance decompositions, replacing our baseline uncertainty index (based on one-period-ahead, i.e., nowcast errors) by two- to four-period-ahead forecast errors for both specification.

Figure 12 shows the IRFs of IP and employment following a 4-standard-deviation shock to our macroeconomic uncertainty series. The peak responses for IP and employment are similar to the ones obtained with the 11-variable VAR when macroeconomic uncertainty is

ordered first. Contrary to the previous VAR though, this peak impact occurs earlier, and is slightly less persistent for the IRFs to the longer horizon uncertainty shock. However, we still do not see evidence of a significant overshoot following the negative impacts, unlike 12.

Finally, Table 7 documents the share of variance attributed to macroeconomic uncertainty in this alternative 8-variable VAR. We find that our baseline macroeconomic uncertainty explains a maximum of 12 and 9% of IP and employment, respectively.

Figure 12: Responses to Innovations to Macroeconomic Uncertainty in an 8-variable VAR



Note: Figure shows the response of GDP and nonfarm payroll employment to uncertainty shocks in a 8-variable VAR with the following variables and ordering: S&P 500, uncertainty, the fed funds rate, wages, CPI, hours, employment and industrial production. Q0 to Q3 denote IRFs to innovations in uncertainty about nowcast, one-, two-, and three-quarter-ahead forecasts, respectively. The shaded area represents 68% confidence bands using Kilian (1998)'s bootstrap-after-bootstrap method.

6.5 Comparison with IRF with other measures of uncertainty

We further compare IRFs to the innovation to our baseline uncertainty estimates with the ones obtained with other widely-used estimates and proxies for macroeconomic uncertainty. We use the 11-variable VAR similar to Christiano et al. (2005), where uncertainty is ordered last. Figure 13 plots responses to innovations in our baseline uncertainty series (FSV),

Table 7: Variance Decomposition for the 8-variable VAR

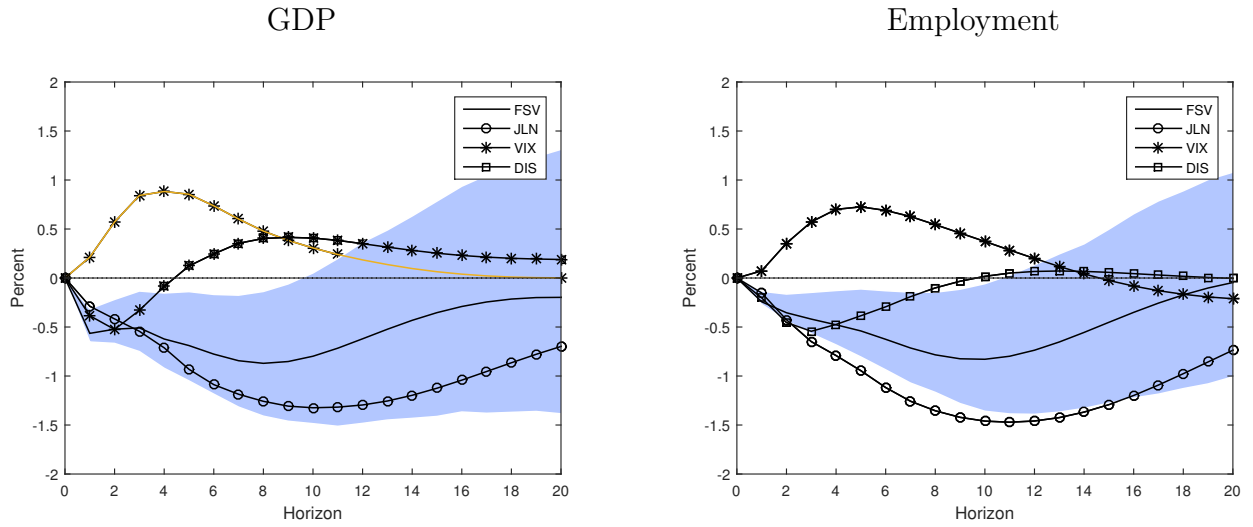
Quarters	<i>IP</i>				<i>Employment</i>			
	Q0	Q1	Q2	Q3	Q0	Q1	Q2	Q3
2	10.26	9.15	3.84	11.21	8.99	13.15	6.29	17.57
4	12.02	8.91	3.39	7.18	8.26	10.15	4.47	10.22
8	10.83	6.27	1.96	3.28	5.62	6.35	2.38	3.89
20	12.02	9.82	4.12	11.21	9.04	13.15	6.29	17.57

Note: Table shows forecast error variance decomposition in a 8-variable VAR with following variables and ordering: S&P 500, Uncertainty, the fed funds rate, wages, CPI, hours, employment and industrial production. Q0 to Q3 denote uncertainty about nowcast, one-, two-, and three-quarter-ahead forecasts, respectively.

as well as the Jurado et al. (2015) index (JLN), the average disagreement introduced in Section 2 (DIS) and the VIX.⁵ There is a clear divergence in results. Both innovations to our measures and the JLN measure of macroeconomic uncertainty result in large drops in economic activity. On the other hand, innovations to proxies of macroeconomic uncertainty, the disagreement and the VIX, result in very small and short-lived negative impacts on economic activity, followed by strong rebounds.

⁵We again estimate IRFs to a four-standard-deviation shock, in line with our analysis in the main text and Section 4.1.

Figure 13: Responses with Alternative Measures of Uncertainty



Note: Figure shows the response of GDP and employment to four-standard-deviation uncertainty shocks in a 11-variable VAR using alternative uncertainty indexes: FSV is our baseline measure of macroeconomic uncertainty as perceived by professional forecasters; JLN is the measure proposed by Jurado et al. (2015); VIX is the stock market implied volatility used by Bloom (2009); and DIS is the average disagreement as discussed in Section 2. The shaded area represents 68% confidence bands using Kilian (1998)'s bootstrap-after-bootstrap method.

References

- ARAI, N. (2014): “Using forecast evaluation to improve the accuracy of the Greenbook forecast,” *International Journal of Forecasting*, 30, 12–19.
- BACHMANN, R., S. ELSTNER, AND E. SIMS (2013): “Uncertainty and Economic Activity: Evidence from Business Survey Data,” *American Economic Journal: Macroeconomics*, 5, 217–49.
- BAUMEISTER, C., P. LIU, AND H. MUMTAZ (2013): “Changes in the Effects of Monetary Policy on Disaggregate Price Dynamics,” *Journal of Economic Dynamics and Control*, 37, 543–560.
- BAUMEISTER, C. AND G. PEERSMAN (2013): “Time-Varying Effects of Oil Supply Shocks on the US Economy,” *American Economic Journal: Macroeconomics*, 5, 1–28.
- BEKAERT, G., M. HOEROVA, AND M. L. DUCA (2013): “Risk, uncertainty and monetary policy,” *Journal of Monetary Economics*, 60, 771–788.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- CAPISTRÁN, C. AND A. TIMMERMANN (2009): “Forecast combination with entry and exit of experts,” *Journal of Business & Economic Statistics*, 27, 428–440.
- 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, 1–45.
- CLARK, T. E. (2012): “Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility,” *Journal of Business & Economic Statistics*.
- CLEMENTS, M. P. (2014): “Forecast Uncertainty – Ex Ante and Ex Post: US Inflation and Output Growth,” *Journal of Business & Economic Statistics*, 32, 206–216.
- COGLEY, T. AND T. J. SARGENT (2005): “Drifts and Volatilities: Monetary Policies and Outcomes in the Post-WWII U.S.” *Review of Economic Dynamics*, 8, 262–302.
- CROUSHORE, D. (2010): “An evaluation of inflation forecasts from surveys using real-time data,” *The BE Journal of Macroeconomics*, 10.
- ENGELBERG, J., C. F. MANSKI, AND J. WILLIAMS (2009): “Comparing the point predictions and subjective probability distributions of professional forecasters,” *Journal of Business & Economic Statistics*, 27, 30–41.
- GEWEKE, J. (1992): “Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments,” in *Bayesian Statistics, edited by Bernardo J. M., Berger J., Dawid A. P., Smith A. F. M.*, Oxford University Press, 169–193.

- GNEITING, T. (2011): “Quantiles as optimal point forecasts,” *International Journal of Forecasting*, 27, 197–207.
- JACQUIER, E., N. G. POLSON, AND P. E. ROSSI (2002): “Bayesian Analysis of Stochastic Volatility Models,” *Journal of Business & Economic Statistics*, 20, 69–87.
- JO, S. (2014): “The Effects of Oil Price Uncertainty on Global Real Economic Activity,” *Journal of Money, Credit and Banking*, 46, 1113–1135.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KILIAN, L. (1998): “Small-Sample Confidence Intervals for Impulse Response Functions,” *Review of Economics and Statistics*, 80, 218–230.
- LAHIRI, K. AND F. LIU (2006): “Modelling multi-period inflation uncertainty using a panel of density forecasts,” *Journal of Applied Econometrics*, 21, 1199–1219.
- LAHIRI, K. AND X. SHENG (2010): “Measuring Forecast Uncertainty By Disagreement: The Missing Ink,” *Journal of Applied Econometrics*, 25, 514–538.
- MANKIW, N. G., R. REIS, AND J. WOLFERS (2004): “Disagreement About Inflation Expectations,” in *NBER Macroeconomics Annual 2003, Volume 18*, The MIT Press, 209–270.
- MUMTAZ, H. AND F. ZANETTI (2013): “The impact of the volatility of monetary policy shocks,” *Journal of Money, Credit and Banking*, 45, 535–558.
- PITT, M. AND N. SHEPHARD (1999): “Time-Varying Covariances: A Factor Stochastic Volatility Approach,” *Bayesian Statistics*, 6, 547–570.
- PRIMICERI, G. E. (2005): “Time varying structural vector autoregressions and monetary policy,” *Review of Economic Studies*, 72, 821–852.
- SILL, K. (2012): “Measuring Economic Uncertainty Using the Survey of Professional Forecasters,” *Federal Reserve Bank of Philadelphia Business Review*, 4, 16–27.
- ZARNOWITZ, V. AND P. BRAUN (1993): “Twenty-two years of the NBER-ASA quarterly economic outlook surveys: Aspects and comparisons of forecasting performance,” in *Business cycles, indicators and forecasting*, University of Chicago Press, 11–94.