

Federal Reserve Bank of New York
Staff Reports

Forecasting Macroeconomic Risks

Patrick A. Adams
Tobias Adrian
Nina Boyarchenko
Domenico Giannone

Staff Report No. 914
February 2020



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Forecasting Macroeconomic Risks

Patrick A. Adams, Tobias Adrian, Nina Boyarchenko, and Domenico Giannone

Federal Reserve Bank of New York Staff Reports, no. 914

February 2020

JEL classification: C22, E17, E37

Abstract

We construct risks around consensus forecasts of real GDP growth, unemployment, and inflation. We find that risks are time-varying, asymmetric, and partly predictable. Tight financial conditions forecast downside growth risk, upside unemployment risk, and increased uncertainty around the inflation forecast. Growth vulnerability arises as the conditional mean and conditional variance of GDP growth are negatively correlated: downside risks are driven by lower mean and higher variance when financial conditions tighten. Similarly, employment vulnerability arises as the conditional mean and conditional variance of unemployment are positively correlated, with tighter financial conditions corresponding to higher forecasted unemployment and higher variance around the consensus forecast.

Key words: macroeconomic uncertainty, quantile regressions, financial conditions

Boyarchenko: Federal Reserve Bank of New York and CEPR (email: nina.boyarchenko@ny.frb.org). Adams: MIT Sloan (email: patadams@mit.edu). Adrian: International Monetary Fund and CEPR (email: tadrian@imf.org). Giannone: Amazon.com and CEPR (email: giannond@amazon.com). Giannone's contributions to this work were done at the Federal Reserve Bank of New York, prior to his joining Amazon. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the International Monetary Fund, the Federal Reserve Bank of New York, or the Federal Reserve System.

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/staff_reports/sr914.html.

1 Introduction

Timely characterizations of risks to the economic outlook play an important role in both economic policy and private sector decisions. The Federal Open Market Committee (FOMC) and other central banks frequently discuss both upside and downside risks to growth, inflation, and unemployment in released statements and minutes. Financial institutions also closely monitor these risks,¹ and use measures such as value at risk to determine the susceptibility of their balance sheets to large losses. In this paper, we introduce a simple method to quantify time-varying risks around macroeconomic forecasts, and use this method to construct probabilistic forecasts for real GDP growth, unemployment, and inflation.

Adopting the methodology of [Adrian, Boyarchenko, and Giannone \(2019\)](#), we use quantile regressions to characterize upside and downside risks around the Survey of Professional Forecasters' (SPF) median consensus forecasts for each indicator, as a function of conditioning information available at the time of the forecasts (specifically, a broad-based index of financial conditions). Given the estimated quantiles obtained from these quantile regressions, we then fit a flexible smooth distribution function in order to obtain a full probability distribution. This method provides a forward-looking assessment of uncertainty, can capture asymmetries in risks over the course of the business cycle, and allows for the construction of informative measures of tail risks.

Studying the uncertainty around consensus point forecasts allows us to focus directly on how financial conditions shape the second and higher moments of the conditional predictive distributions of growth, unemployment, and inflation. Uncertainty around consensus point forecasts fluctuates substantially over time, and upside and downside risks do not vary one-for-one. In times of financial stress, risks around long-horizon forecasts for real GDP growth skew toward the downside, while risks around unemployment forecasts skew toward the

¹For example, see Goldman Sachs' "A Better Balance of Risks: 2018 Mid-Year Outlook" (https://www.gsam.com/content/dam/gsam/pdfs/common/en/public/articles/outlook/2018/GSAM_2018_Mid_Year_Investment_Outlook.pdf).

upside. Since these increases in uncertainty around consensus forecasts are accompanied by movements in the forecasts themselves as financial conditions tighten, decreasing for real GDP growth and increasing for unemployment, our probabilistic forecasts exhibit substantial variation over time in the lower quantiles for real GDP growth and in the upper quantiles for unemployment. In contrast, while risks around long-horizon forecasts for inflation also skew towards the upside during times of financial stress, in the post-Volcker disinflation era, these increases in uncertainty around the consensus inflation forecast are accompanied by decreases in the consensus forecast itself, leading to symmetric fluctuations of the lower and upper quantiles of inflation. Notably, prior to the Volcker disinflation, the upper quantiles of inflation exhibited more variation over time.

We find that, relative to forecasts constructed using the historical distribution of forecast errors, conditioning on financial conditions significantly improves the accuracy of out-of-sample forecasts for real GDP growth and unemployment and modestly for inflation. These findings indicate a potentially important connection between financial conditions and real business cycles, but a weaker connection with prices. Our empirical results linking tight financial conditions with increased uncertainty surrounding future real economic outcomes are consistent with macroeconomic models in which financial frictions generate endogenous fluctuations in the volatility of real variables. Models that achieve this result through frictions arising from within the financial intermediary sector include, among others, [Brunnermeier and Sannikov \(2014\)](#), [Adrian and Boyarchenko \(2015\)](#), and [Adrian and Duarte \(2017\)](#). However, while theory suggests that tightening financial conditions may exacerbate downside risks to inflation through the possibility of deflationary spirals ([Brunnermeier and Sannikov, 2016](#)), our in-sample results imply that post-1985 risks to inflation are fairly symmetric around the consensus forecast while pre-1985 upside risks to inflation vary more than downside risks over the course of the business cycle. [Gilchrist, Schoenle, Sim, and Zakrajšek \(2017\)](#) argue that the interaction of financial frictions with customer markets attenuates the response of inflation to the economic slack that emerges when financial conditions tighten.

Common approaches to assessing uncertainty around point forecasts adopt an “unconditional” perspective, using the distribution of historical forecast errors to construct estimates of uncertainty without incorporating additional information available at the time the forecasts are made. [Reifschneider and Tulip \(2019\)](#) use this approach to construct confidence bands around the median consensus forecasts from the FOMC’s Summary of Economic Projections, based on forecast errors within a twenty year rolling window. The use of rolling windows can capture low frequency changes in uncertainty, such as the decline in macroeconomic volatility associated with the Great Moderation beginning in 1985. However, this “unconditional” approach assumes that risks around consensus forecasts are unpredictable. In our out-of-sample evaluation, we compare our quantile regression-based density forecasts to a benchmark “unconditional” density forecast and find that conditioning on financial conditions yields statistically significant gains in forecast accuracy for real GDP growth and unemployment, indicating an important role for conditioning information in predicting macroeconomic risks.

Alternative approaches to modeling time-varying uncertainty around the consensus forecast include using information from survey-based density forecasts (as in e.g. [Andrade, Ghysels, and Idier, 2014](#); [Ganics, Rossi, and Sekhposyan, 2019](#)) or specifying an exogenous stochastic process for innovation volatilities (as in e.g. [Primiceri, 2005](#); [Cogley and Sargent, 2005](#); [Clark, 2012](#); [Clark, McCracken, and Mertens, 2018](#)). Both of these approaches have their own drawbacks. Since survey-based density forecasts are fixed-object forecasts (e.g. 2020 GDP growth) while consensus forecasts are fixed-horizon (e.g. four-quarter GDP growth), using density forecasts to model time-varying uncertainty around the consensus forecasts involves assumptions on the correspondence between fixed-object and fixed-horizon forecasts. Similarly, models in which uncertainty evolves exogenously can only infer increases in uncertainty after the realization of large prediction errors, and are thus less likely to detect fluctuations in risks at business cycle frequencies before they occur. In contrast, the quantile regression approach enables us to remain relatively agnostic about the relationship between

current financial conditions and uncertainty around the consensus forecast, allowing the data to inform us on that relationship instead.

Adrian, Boyarchenko, and Giannone (2019) show that downside risks to real GDP growth vary substantially over the business cycle as a function of financial conditions, while upside risks are stable over time. We extend these earlier findings along two directions. First, we show that these earlier results for real GDP growth hold even when we condition on consensus forecasts, which provide a more comprehensive summary of current and expected economic conditions than lagged GDP growth. Second, we contribute to the nascent literature on quantile regression approaches to measuring risks to inflation (Ghysels, Iania, and Striaukas, 2018) and unemployment (Kiley, 2018). As with real GDP growth, conditioning on the corresponding consensus forecasts arguably allows us to incorporate the most timely information on economic conditions available.

The paper also documents new facts about the SPF forecasts that complement other recent findings. Galbraith and van Norden (2018) show that the unconditional distribution of median SPF forecast errors for unemployment is positively skewed; we show that the degree of skewness in the conditional forecast error distribution varies significantly as a function of financial conditions. Barnes and Olivei (2017) show that financial variables are uninformative in predicting a common factor extracted from one-year-ahead consensus forecast errors for real GDP growth, unemployment, and CPI inflation. While they focus on predictability in terms of the mean of the conditional forecast error distribution, we focus on other features of this distribution and show that financial conditions do in fact provide considerable information about the *full distribution* of future forecast errors.

The rest of this paper is organized as follows. Section 2 describes the data used in our empirical analysis. Section 3 introduces our method for quantifying uncertainty around point forecasts, and presents both in-sample density forecasts and risk measures derived from these density forecasts. Section 4 presents out-of-sample forecasting results. Section 5 concludes.

2 Data

We use data on real-time survey forecasts for real output, unemployment, and inflation provided in the quarterly Survey of Professional Forecasters (SPF). Initially conducted by the American Statistical Association and the National Bureau of Economic Research in 1968, the SPF has been managed by the Federal Reserve Bank of Philadelphia since 1990Q3.² Professional forecasters participating in the survey provide their forecasts in the middle month of each quarter, and results are released to the public shortly after the submission deadline. For each variable, participants provide quarterly point forecasts for horizons ranging from the current quarter to four quarters ahead.³ We use the median forecasts for quarter-over-quarter real GDP growth, the quarterly average unemployment rate, and quarter-over-quarter GDP price index inflation.⁴ Our proposed method can be used to assess time-varying uncertainty and construct probabilistic forecasts for any point forecast with a sufficiently long history of available data. In this paper, we focus on SPF forecasts since these point forecasts have been studied extensively, are published regularly and are freely available. Other judgmental forecasts commonly used in the literature are either conducted less frequently (e.g. the Livingston Survey), available only through subscription (e.g. Blue Chip Economic Indicators or Consensus Forecasts), are released with a substantial lag (e.g. the Federal Reserve’s Greenbook forecasts), or refer to annual data frequencies (e.g. the IMF World Economic Outlook, the World Bank Global Economic Prospects, and the OECD Economic Outlook).

As an additional conditioning variable to construct density forecasts, we use the Federal Reserve Bank of Chicago’s National Financial Conditions Index (NFCI). The NFCI provides a weekly summary of U.S. financial conditions, using data on a broad set of 105 financial

²Historical forecasts, survey documentation, and other materials can be obtained from the Federal Reserve Bank of Philadelphia’s [website](#).

³For two quarters early in our evaluation period (1970Q1 and 1974Q3), four-quarter-ahead forecasts are not available. In these cases, we replace the missing median four-quarter-ahead forecasts with the available median three-quarter-ahead forecasts.

⁴SPF definitions for real output and prices have changed over time. From 1992 to 1995, the SPF collected forecasts for fixed-weighted real GDP and the GDP implicit price deflator. Prior to 1992, these forecasts were collected for GNP instead of GDP.

variables capturing risk premia, credit availability, and leverage. The index is constructed from a large dynamic factor model estimated using the quasi maximum likelihood estimator of Doz, Giannone, and Reichlin (2012); a complete description of the methodology is provided by Brave and Butters (2012). The NFCI is standardized to have an average value of zero and unit standard deviation over its full sample period. Positive readings of the index are indicative of tighter-than-average financial conditions, while negative readings are indicative of looser-than-average financial conditions.

Historical data for the NFCI are available starting in January 1971, and so our evaluation period begins in 1971Q1 and ends in 2018Q4. Since the SPF is conducted in the first week of the middle month of each quarter, throughout our empirical analysis we use the value of the NFCI as of the last Friday of the first month of the quarter in which each density forecast is generated, in order to avoid exploiting data that are not available at the time when forecasters are surveyed.

3 Methodology

To construct quarterly predictive distributions for real GDP growth, unemployment, and inflation, we use conditioning information available at the time of each SPF survey (specifically financial conditions, as measured by the NFCI) to determine the distribution of future forecast errors. To model this distribution, we use the two-step quantile regression methodology developed by Adrian, Boyarchenko, and Giannone (2019). We then use these distributions to construct measures of downside and upside risks for each variable and forecast horizon.

3.1 Quantile Regressions

Denote by y_{t+h} the value of the target variable of interest in quarter $t + h$. For real GDP growth or inflation, y_{t+h} represents the annualized average growth rate of real GDP or the GDP price index (respectively) between quarter t and quarter $t+h$; for unemployment, y_{t+h} is

the average unemployment rate in quarter $t+h$. Additionally, denote by $\hat{y}_{t+h|t}^{SPF}$ the h -quarter-ahead median SPF forecast for y_{t+h} , and the associated forecast error by $e_{t+h|t}^{SPF} \equiv y_{t+h} - \hat{y}_{t+h|t}^{SPF}$.

We first estimate quantile regressions (Koenker and Bassett, 1978) of the median SPF forecast errors $e_{t+h|t}^{SPF}$ on conditioning variables available at the time of the quarter t SPF survey. These conditioning variables are collected in the vector x_t , which also includes a constant. Given $\tau \in (0, 1)$, we wish to estimate the τ -quantile of the h -quarter-ahead forecast error distribution conditional on x_t , denoted by $F_{e_{t+h|t}^{SPF}|x_t}$. The τ -quantile is defined as

$$Q_{e_{t+h|t}^{SPF}|x_t}(\tau|x_t) \equiv \inf\{q \in \mathbb{R} | F_{e_{t+h|t}^{SPF}|x_t}(q|x_t) \geq \tau\}, \quad (1)$$

The quantile regression coefficients β_τ are chosen to minimize the sum of quantile-weighted absolute residuals:

$$\hat{\beta}_\tau \equiv \operatorname{argmin}_{\beta_\tau \in \mathbb{R}^k} \sum_{t=1}^{T-h} \left(\tau \cdot \mathbb{1}_{\{e_{t+h|t}^{SPF} > x_t' \beta_\tau\}} |e_{t+h|t}^{SPF} - x_t' \beta_\tau| + (1 - \tau) \cdot \mathbb{1}_{\{e_{t+h|t}^{SPF} < x_t' \beta_\tau\}} |e_{t+h|t}^{SPF} - x_t' \beta_\tau| \right) \quad (2)$$

The predicted value from the quantile regression,

$$\hat{Q}_{e_{t+h|t}^{SPF}|x_t}(\tau|x_t) \equiv x_t' \hat{\beta}_\tau, \quad (3)$$

provides a linear estimate of the τ -quantile of $e_{t+h|t}^{SPF}$ conditional on x_t . Under the assumption that the forecast error $e_{t+h|t}^{SPF}$ and forecast $\hat{y}_{t+h|t}^{SPF}$ are independent conditional on x_t , we can then obtain the implied quantiles of y_{t+h} by adding back the median SPF forecast $\hat{y}_{t+h|t}^{SPF}$ to obtain:⁵

$$Q_{y_{t+h|t}|x_t}(\tau|x_t) \equiv Q_{e_{t+h|t}^{SPF}|x_t}(\tau|x_t) + \hat{y}_{t+h|t}^{SPF} \quad (4)$$

Throughout our empirical analysis, the vector of conditioning variables x_t contains the

⁵Given the assumption of independence of $e_{t+h|t}^{SPF}$ and $\hat{y}_{t+h|t}^{SPF}$ conditional on x_t , our procedure can be viewed as a restricted quantile regression of y_{t+h} on both x_t and $\hat{y}_{t+h|t}^{SPF}$, in which the coefficients on the SPF forecast are restricted to be equal to 1 for all values of $\tau \in (0, 1)$. In unreported robustness checks, we estimate an unrestricted version of this quantile regression and find that out-of-sample forecasting performance deteriorates.

quarter t value of the NFCI (using the dating convention described in Section 2) and a constant. Figure 1 plots the observed one- and four-quarter-ahead SPF forecast errors against the value of the NFCI at the time of each SPF forecast. The colored lines depict the estimated 5th, 50th (median), and 95th quantiles of the forecast error distribution as a function of the NFCI (obtained via quantile regression), as well as the ordinary least squares regression line. All of these plots depict a strong asymmetry across quantiles in the relationship between future forecast errors and financial conditions at the time of each forecast. For real GDP growth, the slope of the conditional 95th quantile function is slightly steeper than that of the conditional 5th quantile function for one-quarter-ahead forecast errors. This relationship reverses for four-quarter-ahead forecast errors, for which the conditional 95th quantile does not appear to depend on financial conditions at all while the conditional 5th quantile decreases sharply as financial conditions tighten. As a result, short-horizon GDP forecast errors exhibit positive skewness during times of financial stress (which are indicated by large positive values of the NFCI), while long-horizon GDP forecast errors exhibit negative skewness. For unemployment, at both horizons the 5th conditional quantile of the forecast error distribution is essentially unaffected by financial conditions while the 95th conditional quantile rises substantially as financial conditions tighten, indicating that the forecast error distributions exhibits positive skewness in times of financial stress. For inflation, the asymmetry in the 5th and 95th conditional quantiles is also qualitatively similar for both short and long forecast horizons: tight financial conditions increase both downside and upside risks around inflation forecasts, with the latter rising more than the former.

As a result of these asymmetries in the relationship between financial conditions and uncertainty, the shape of the conditional forecast error distribution changes substantially as a function of prevailing financial conditions at the time of the forecast. When financial conditions are broadly consistent with historical averages (as indicated by an NFCI value of zero), downside and upside risks to forecasts are roughly balanced, and the distribution of future forecast errors is relatively symmetric. As financial conditions tighten, these distributions

become highly skewed, toward the left for long-term GDP growth forecasts and toward the right for short-term GDP growth forecasts, unemployment, and inflation. However, this does not necessarily imply that the SPF point forecasts fail to incorporate financial conditions: for GDP growth and inflation, the OLS regression lines are nearly flat, as would be expected if the median SPF forecasts represent conditional expectations based on information sets that include contemporaneous financial conditions.⁶ Even if fully optimal point forecasts based on information available at the time of each forecast were observed, there is no guarantee that the associated forecast errors would be homoskedastic, or that higher moments of the forecast error distribution would not vary over time.

Figure 2 plots the estimated coefficients on the NFCI from these quantile regressions.⁷ In all cases the estimated coefficients exhibit an upward-sloping pattern across quantiles, indicating that tightening financial conditions shift either one or both tails of the forecast error distribution outward, leading to increased uncertainty around the median forecasts. Moreover, many of the coefficients fall outside of the estimated 95% confidence bands, indicating that these nonlinearities are statistically significant.

Figures 3 plots the estimated predictive distributions for real GDP growth, inflation, and unemployment over time. As shown in Equation 4, these distributions are obtained by shifting the estimated forecast error quantiles by the median forecast for each variable. At each date, we plot the realization of each target variable, along with the median SPF forecast and the estimated quantiles from either one or four quarters prior. For real GDP growth,

⁶For GDP growth and inflation, in regressions of forecast errors on the NFCI and a constant we cannot reject the hypothesis that the coefficient on the NFCI is equal to zero at even the 10% level for either the one- or four-quarter-ahead horizons. However, for the unemployment forecast errors we can strongly reject this hypothesis, with associated t -statistics greater than 3 at both the one- and four-quarter-ahead forecast horizons. This indicates that the median SPF unemployment forecasts may fail to adequately incorporate information about financial conditions. The probabilistic forecasts we construct adjust for this fact by shifting the mean of the forecast error distribution away from zero in times when the value of the NFCI differs from its historical average of zero.

⁷Confidence bands are computed via bootstrapping under the assumption that the data are generated by a flexible linear model. Specifically, for each target variable of interest we estimate a three-variable vector autoregression (VAR) including the target variable, the median SPF forecast for the target variable, and the NFCI, using four lags and assuming i.i.d. Gaussian innovations. Given the estimated VAR parameters, we then simulate 1000 bootstrap samples to determine the distribution of the estimated quantile regression coefficients in the absence of any nonlinearities.

the lower quantiles of the growth distribution vary substantially over time and decrease sharply in periods of financial stress, as documented by [Adrian, Boyarchenko, and Giannone \(2019\)](#). Similar patterns arise in the estimated quantiles for inflation and unemployment. For unemployment, the lower quantiles of the predictive distribution shift one-for-one with the median SPF forecast over time, while the upper quantiles shift more than one-for-one as both the median forecast and estimated upside risks to the forecast rise during periods of financial stress. Uncertainty around unemployment forecasts is also much greater at the four-quarter-ahead horizon than at the one-quarter-ahead horizon. For inflation, both upside and downside risks to the median forecast fluctuate over time, and the lower quantiles of the predictive distribution for inflation are generally more stable than the upper quantiles.

To provide a clearer view of how our estimated measures of uncertainty behave over the business cycle, [Figure 4](#) plots the estimated interquartile range of the forecast error distribution (computed as the difference between the 75th and 25th conditional quantiles) against the point forecasts. For GDP growth and unemployment, this measure of uncertainty moves countercyclically, rising as forecasts for GDP decline and forecasts for unemployment increase. For GDP growth, the negative correlation between the median forecast and uncertainty leads to substantial instability in the left tail of our estimated predictive distributions, since shifts in the mean and dispersion move the 5th quantile of the conditional growth distribution in the same direction. For unemployment, the positive correlation between the median forecast and uncertainty leads to large shifts in the 95th quantile, and thus the right tails of our predictive distributions for unemployment vary substantially over time. [Kiley \(2018\)](#) finds similar asymmetries in the predictive distribution of the unemployment rate at long forecast horizons.

For inflation, the relationship between the median forecast and uncertainty differs before and after 1985. In the period before 1985 (represented by the red circles) there is a strong positive relationship between the level of expected inflation and inflation uncertainty, leading to the large and striking shifts in the upper quantiles of the conditional distribution during

this period which are shown in Figure 3. In contrast, from 1985 onward (represented by the blue diamonds) there is no clear relationship between the level of expected inflation and uncertainty. Stock and Watson (2007) and Cogley, Sargent, and Primiceri (2010) document changes in inflation dynamics - including its persistence and volatility - across these two periods.

3.2 Predictive Distributions

Our quantile regressions provide estimates of a finite set of conditional quantiles for each target variable. In order to construct a full conditional probability distribution from these estimates, we follow Adrian, Boyarchenko, and Giannone (2019) and fit a smooth quantile function from a flexible class of probability distributions to the estimated conditional quantiles. We consider probability distributions from the four-parameter skew t -family of Azzalini and Capitanio (2003), with probability density function given by

$$f(y; \mu, \sigma, \alpha, \nu) = \frac{2}{\sigma} t\left(\frac{y - \mu}{\sigma}; \nu\right) T\left(\alpha \left(\frac{y - \mu}{\sigma}\right) \sqrt{\frac{\nu + 1}{\nu + \frac{y - \mu}{\sigma}}}; \nu + 1\right) \quad (5)$$

Here $t(\cdot; n)$ and $T(\cdot; n)$ denote the probability density function and cumulative distribution function (respectively) of the standard student's t -distribution with n degrees of freedom. The skew t -distribution is specified by its location $\mu \in \mathbb{R}$, scale $\sigma \in \mathbb{R}_{++}$, shape $\alpha \in \mathbb{R}$, and degrees of freedom $\nu \in \mathbb{R}_{++}$. This family of distributions is quite general and allows to capture fat tails and skewness. However, it does not allow for other important features such as multimodality. These limitations are due to the necessity of parsimony, which is stringent since we fit a different distribution every time we make a new forecast.

For each quarter t , given the estimated conditional quantiles $\hat{Q}_{y_{t+h}|x_t}(\tau|x_t)$ ⁸ of y_{t+h} , we fit a skew t -distribution by choosing the parameters $\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h}\}$ to minimize the

⁸In case the estimated quantiles are not monotonically increasing, the uncrossing procedure of Chernozhukov, Fernández-Val, and Galichon (2010) can be applied in order to obtain a sequence of estimated quantiles that is monotonically increasing.

squared differences between the skew t -implied quantiles and our quantile regression estimates for $\tau = 0.05, 0.25, 0.75$, and 0.95 :

$$\{\hat{\mu}_{t+h|t}, \hat{\sigma}_{t+h|t}, \hat{\alpha}_{t+h|t}, \hat{\nu}_{t+h|t}\} = \underset{\mu, \sigma, \alpha, \nu}{\operatorname{argmin}} \sum_{\tau=0.05, 0.25, 0.75, 0.95} \left(\hat{Q}_{y_{t+h}|x_t}(\tau|x_t) - F^{-1}(\tau; \mu, \sigma, \alpha, \nu) \right)^2 \quad (6)$$

Here $F^{-1}(\tau; \mu, \sigma, \alpha, \nu)$ denotes the quantile function of the skew t -distribution.

In addition to constructing smooth probability distributions given the estimated quantiles $\hat{Q}_{y_{t+h}|x_t}(\tau|x_t)$ conditional on the NFCI, we also construct alternative predictive distributions based only on the current SPF forecast and the distribution of historical forecast errors. Following [Reifschneider and Tulip \(2019\)](#), we compute the unconditional quantiles of the forecast error distribution⁹, center these estimated quantiles around the current SPF forecast $y_{t+h|t}^{SPF}$, then fit a skew t -distribution to the implied quantiles of y_{t+h} . This alternative predictive distribution does not capture any changes in the conditional distribution of future forecast errors over time and thus represents an appropriate benchmark against which to compare our predictive distributions that incorporate information from financial conditions.

Figure 5 displays the predictive densities generated by our method in two particular quarters: 2008Q3 and 2017Q4 (the last date in our sample period for which we can compare the four-quarter-ahead forecasts against realized values). The 2008Q3 SPF round was conducted in early August 2008. Although the survey took place roughly one month before the collapse of Lehman Brothers, financial conditions were already relatively tight, as indicated by values of the NFCI 0.4 to 0.5 standard deviations above the index's historical average. In contrast, 2017Q4 was a period of relative stability and accommodative financial conditions, with the NFCI hovering 0.7 to 0.8 standard deviations below its historical average. Each chart plots the four quarter ahead predictive distribution obtained from the quantile regression-based model that conditions on financial conditions. For comparison, we also plot the unconditional distributions based only on the median SPF forecast, computed using the method

⁹This can be implemented by estimating the quantile regression described in Equation 2 with only a constant included in the set of conditioning variables.

described in the previous paragraph. The vertical lines represent the median SPF forecast at each date, which is used in the construction of both densities, and the realized outcome for the target variable (either average annualized quarterly GDP growth/inflation over the next four quarters, or the average quarterly unemployment rate in four quarters).

Inspection of these predictive densities reveals that financial conditions provide useful information about risks around the median SPF forecasts not only when financial conditions are relatively tight, but also when they are accommodative. During periods of financial stress like 2008Q3, uncertainty around the SPF point forecasts increases relative to the average level depicted by the unconditional densities. In contrast, during periods of accommodative financial conditions like 2017Q4, uncertainty decreases and the predictive density is concentrated near the SPF point forecasts. In out-of-sample density forecasting results presented in Section 4, we show that these latter periods are an important driver of the gains in predictive accuracy reaped by conditioning on financial conditions, since the unconditional predictive densities overstate uncertainty during these times and thus suffer from poor precision.

Figure 5 also highlights the asymmetry of shifts in risks to real activity over the business cycle. For GDP growth, the differences between the two predictive densities at each forecast date appear in the left tail and center of the distributions, with nearly identical right tails. The opposite is true of the predictive densities for unemployment, for which differences arise primarily in the right (rather than left) tails. These patterns again point to a role for financial conditions to provide information about downside - but not upside - risks around point forecasts for growth and employment, as emphasized in the discussion of the coefficient estimates presented in Figure 2.

3.3 Downside and Upside Risk Measures

Using our estimated predictive densities, we can construct informative measures of downside and upside risks around consensus forecasts. In this paper, we focus on the 5% expected shortfall and 95% expected longrise measures. These two measures capture the expected

severity of an event that occurs in either the left tail (for expected shortfall) or right tail (for expected longrise) of the predictive distribution. Specifically, these measures are defined by averaging the fitted quantile function $\hat{F}_{y_{t+h}|x_t}^{-1}(\tau|x_t)$ of the predictive distribution over the left and right tail regions (respectively):

$$SF_{t+h|t} = \frac{1}{0.05} \int_0^{0.05} \hat{F}_{y_{t+h}|x_t}^{-1}(\tau|x_t) d\tau, \quad LR_{t+h|t} = \frac{1}{0.05} \int_{0.95}^1 \hat{F}_{y_{t+h}|x_t}^{-1}(\tau|x_t) d\tau \quad (7)$$

Expected shortfall represents the average realization drawn from below the 5th quantile of the predictive distribution, while expected longrise represents the average realization drawn from above the 95th quantile of the predictive distribution. Both measures capture the expected severity of extreme tail events, *conditional* on their occurrence.

Figure 6 plots the expected shortfall and longrise of our predictive distributions over time. To illustrate the contribution of financial conditions to these tail risk measures, we compute them for both the predictive densities that incorporate financial conditions (solid lines) and the unconditional predictive densities that do not (dashed lines). Similar to the pattern observed in the estimated quantiles, the comovement of the median SPF forecast with the estimated uncertainty around the forecast leads to strong asymmetries between upside and downside tail risks over the course of the business cycle. For GDP growth at both forecasting horizons, the expected shortfall fluctuates substantially throughout our sample period while the expected longrise is more stable. For unemployment, the four-quarter-ahead expected longrise varies much more than the expected shortfall, which moves roughly one-for-one with the median SPF forecast shown in Figure 3. For inflation, the expected longrise also fluctuates more than the expected shortfall at both horizons, and is most volatile during the pre-1985 portion of our sample period.

Moreover, comparing these tail risk estimates for both densities sheds light on which risks financial conditions are or are not informative about. For example, substantial differences arise between expected longrise estimates for four-quarter-ahead unemployment depending

on whether financial conditions are taken into account, and the unconditional distribution appears to overestimate the risk of large increases in unemployment during times of accommodative financial conditions but underestimate this risk during times of financial stress. In contrast, incorporating financial conditions into the forecast seems to have no effect on the expected shortfall of unemployment. A similar pattern emerges for GDP growth at the four-quarter-ahead horizon, where incorporating financial conditions substantially changes the expected shortfall but not the expected longrise.

4 Out-of-sample Evaluation

To determine the importance of accounting for conditional uncertainty around the median SPF forecasts over the course of our sample period, we conduct an out-of-sample density forecasting exercise starting in 1992Q1, when twenty years of four-quarter-ahead forecast errors are first available. For each quarter, we re-estimate the quantile regression described in Equation 2 using forecast error and NFCI observations available through the previous quarter. We then use the median SPF forecast and the value of the NFCI for the given quarter to construct h -period ahead out-of-sample predictive distributions using the procedure described in Section 3.2. In the same manner, we construct the “unconditional” predictive distribution based only on the current quarter median SPF forecast and the historical distribution of SPF forecast errors available through the previous quarter. Additionally, to capture potential long-term trends in macroeconomic volatility, we follow [Reifschneider and Tulip \(2019\)](#) and also report results obtained using a twenty-year rolling window to estimate the unconditional predictive distribution (rather than all observations available at the date of each forecast). While our out-of-sample exercise replicates the timing when data become available to professional forecasters in real time, we use the latest available revised data rather than the real-time data published at each forecasting date.

To compare the accuracy of these out-of-sample density forecasts, we compute predictive

scores. For a given h -period ahead predictive density $\hat{f}_{y_{t+h}|\mathcal{I}_t}(\cdot)$, the log predictive score is calculated by evaluating the predictive density at the realized value of the target variable, denoted by y_{t+h}^o :

$$PS_{\hat{f}_{y_{t+h}|\mathcal{I}_t}}(y_{t+h}^o) \equiv \hat{f}_{y_{t+h}|\mathcal{I}_t}(y_{t+h}^o) \quad (8)$$

To compare two given forecasts $\hat{f}_{y_{t+h}|\mathcal{I}_t}(\cdot)$ and $\hat{g}_{y_{t+h}|\mathcal{I}_t}(\cdot)$, we compute the average difference in log predictive scores

$$\frac{1}{T - h - t_{1992Q1}} \sum_{t=t_{1992Q1}}^{T-h} (\log PS_{\hat{f}_{y_{t+h}|\mathcal{I}_t}}(y_{t+h}^o) - \log PS_{\hat{g}_{y_{t+h}|\mathcal{I}_t}}(y_{t+h}^o)) \quad (9)$$

over our out-of-sample evaluation period.

Additionally, to separately evaluate the calibration of the predictive distributions, we compute probability integral transforms (PITs), obtained by evaluating the estimated cumulative distribution functions $\hat{F}_{y_{t+h}|\mathcal{I}_t}(\cdot)$ at the realized value of the target variable:

$$PIT_{\hat{F}_{y_{t+h}|\mathcal{I}_t}}(y_{t+h}^o) \equiv \hat{F}_{y_{t+h}|\mathcal{I}_t}(y_{t+h}^o) \quad (10)$$

If the predictive distributions $\hat{F}_{y_{t+h}|\mathcal{I}_t}(\cdot)$ are correctly calibrated, then the PITs will be uniformly distributed. We assess the validity of this hypothesis by analyzing the empirical distribution of the PITs over our evaluation period.

Figure 7 shows the out-of-sample predictive scores for the financial conditions-based and unconditional predictive densities (with the latter density estimated using an expanding window). Accuracy gains from incorporating information from financial conditions are large, especially in normal times when accommodative financial conditions lead to sharper predictions by assigning higher probability around the modal outcomes and lower probability to the extreme outcomes observed during crises. In periods of crisis accuracy gains are less pronounced since the absolute probability of tail outcomes is low under both predictive distributions, and hence differences in relative performance are less visible.

These gains in predictive accuracy are quantified in Table 1, which presents differences in average log predictive scores. Positive values indicating superior average forecasting performance of the financial conditions-based density relative to the unconditional density. Following Diebold and Mariano (1995), we also report heteroskedasticity- and autocorrelation-robust standard errors for each difference in means.¹⁰ For all three variables and both forecasting horizons, average log predictive scores are larger for the financial conditions-based density, regardless of whether the unconditional density is estimated using an expanding window (top panel) or rolling window (bottom panel) of past forecast errors. The difference in predictive accuracy is largest for unemployment at the four-quarter-ahead horizon, for which we documented particularly large and persistent differences in upside risk estimates between the two distributions in Section 3.3 and Figure 6. For GDP growth and inflation, gains in predictive accuracy are instead larger at the one-quarter-ahead horizon rather than the four-quarter-ahead horizon.¹¹ While the differences in average log scores for inflation forecasts are large in absolute terms, the standard errors are relatively large due to the persistence of the difference in log scores.

Figure 8 shows the empirical cumulative distribution function of the PITs for the two densities. Under the null hypothesis of perfect calibration of the predictive densities, the PITs are uniformly distributed and thus their empirical distributions should not deviate substantially from the 45-degree line. To assess the significance of deviations from uniformity, we construct confidence bands following Rossi and Sekhposyan (2019). These distributions provide evidence of good forecast calibration. The empirical distributions for the PITs of the financial conditions-based density fall outside of the confidence bands only for one-quarter-ahead forecasts of GDP growth and unemployment, and in the former case the same is true

¹⁰Inference based on these standard errors is asymptotically valid only for the predictions computed using the rolling window of 20 years. For the expanding window estimates, the standard errors should be taken as a general guidance.

¹¹We also compared the results of our quantile regression-based density forecasts to those generated by a conditionally Gaussian model for forecast errors, in which both the mean and log standard deviation of the conditional forecast error distribution are both linear functions of the NFCI. Both approaches achieve similar accuracy for predicting GDP growth and inflation forecast errors, but the conditionally Gaussian model performed significantly worse in predicting unemployment forecast errors.

for the unconditional density. In most other cases, the distribution of PITs for the financial conditions-based density are closer to the 45-degree line than for the unconditional density.¹²

Overall, our out-of-sample forecasting results show that our simple methodology for modeling time-varying risks around point forecasts as a function of conditioning information can improve substantially upon simple density forecasts which assume that uncertainty does not fluctuate over time, both in terms of forecast accuracy and calibration. We also confirm that the link between financial conditions and economic uncertainty is exploitable in constructing out-of-sample density forecasts even when we condition on rich information about macroeconomic expectations (median SPF forecasts).

5 Conclusion

In this paper, we present a simple method to construct probabilistic forecasts, using judgmental point forecasts and additional conditioning variables that can provide information about the uncertainty around these point forecasts. We use this method to construct probabilistic forecasts for real GDP growth, unemployment, and inflation. We document substantial variation in risks around the Survey of Professional Forecasters' median consensus forecasts over time, captured by changes in financial conditions. Incorporating information about financial conditions improves out-of-sample forecast accuracy noticeably for real GDP growth and unemployment, and mildly for inflation.

Our method can be easily adopted to quantify time-varying risks around any point forecasts. While this is of obvious use for judgmental point forecasts for which accompanying probability assessments are not provided, such as the Blue Chip or Federal Reserve Greenbook forecasts, it may also serve useful in characterizing uncertainty around model-based forecasts. For example, many policy institutions use dynamic factor models to produce

¹²Bands for one-quarter-ahead forecasts are based on critical values derived under the null of uniformity and independence of the PITs, while bands for four-quarter-ahead forecasts are computed by bootstrapping only assuming uniformity. The confidence bands should be taken as general guidance since they are derived for forecasts computed using a rolling window, while we use an expanding estimation window.

short-term forecasts of real GDP growth, and construct measures of uncertainty around these forecasts based on the models' historical forecast errors (Bok, Caratelli, Giannone, Sbordone, and Tambalotti, 2018). Our method can condition these measures of uncertainty on variables that may or may not be included in the model, and thus may provide a more convenient and robust alternative to incorporating stochastic volatility into the model.¹³ By including additional conditioning variables in the quantile regression step, our method can also be modified to incorporate information other than financial conditions that may serve as signals of time-varying macroeconomic risk, such as measures of economic policy uncertainty (Baker, Bloom, and Davis, 2016), geopolitical risk (Caldara and Iacoviello, 2018) or macroeconomic uncertainty (Jurado, Ludvigson, and Ng, 2015; Hengge, 2019).

We gauge financial conditions using a single summary index constructed from a large set of indicators. An important task for future research is to determine whether a one-dimensional index can in fact summarize the information content of financial conditions for predicting macroeconomic risks,¹⁴ and whether additional gains in forecast accuracy can be obtained by directly conditioning on the underlying financial variables. In this case, the standard quantile regression framework used in this paper must be modified in order to deal with the curse of dimensionality that arises in this setting.

References

- ADRIAN, T., AND N. BOYARCHENKO (2015): "Intermediary Leverage Cycles and Financial Stability," Staff Report No. 567, Federal Reserve Bank of New York.
- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): "Vulnerable Growth," *American Economic Review*, 109(4), 1263–89.

¹³Castelletti-Font, Diev, and Honvo (2019) use this approach to assess risks around GDP nowcasts using French data.

¹⁴Brunnermeier, Palia, Sastry, and Sims (2018) advocate for a multi-dimensional approach to measuring financial stress.

- ADRIAN, T., AND F. DUARTE (2017): “Financial Vulnerability and Monetary Policy,” Staff Report No. 804, Federal Reserve Bank of New York.
- ANDRADE, P., E. GHYSELS, AND J. IDIER (2014): “Inflation Risk Measures and Their Informational Content,” SSRN abstract N. 2439607.
- AZZALINI, A., AND A. CAPITANIO (2003): “Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t-Distribution,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 367–389.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring Economic Policy Uncertainty,” *Quarterly Journal of Economics*, 131, 1593–1636.
- BARNES, M., AND G. P. OLIVEI (2017): “Financial Variables and Macroeconomic Forecast Errors,” Working Paper No. 17-17, Federal Reserve Bank of Boston.
- BOK, B., D. CARATELLI, D. GIANNONE, A. M. SBORDONE, AND A. TAMBALOTTI (2018): “Macroeconomic Nowcasting and Forecasting with Big Data,” *Annual Review of Economics*, 10, 615–643.
- BRAVE, S., AND A. BUTTERS (2012): “Diagnosing the Financial System: Financial Conditions and Financial Stress,” *International Journal of Central Banking*, 8, 369–422.
- BRUNNERMEIER, M. K., D. PALIA, K. SASTRY, AND C. A. SIMS (2018): “Feedbacks: Financial Markets and Economic Activity,” Working paper.
- BRUNNERMEIER, M. K., AND Y. SANNIKOV (2014): “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, 104, 379–421.
- (2016): “The I Theory of Money,” Working paper.
- CALDARA, D., AND M. IACOVIELLO (2018): “Measuring Geopolitical Risk,” International Finance Discussion Paper 1222, Federal Reserve Board of Governors.
- CASTELLETTI-FONT, B., P. DIEV, AND W. HONVO (2019): “Are financial variables a useful complement for GDP nowcasting?,” Eco notepad blog post, Banque de France.
- CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND A. GALICHON (2010): “Quantile and Probability Curves Without Crossing,” *Econometrica*, 78, 1093–1125.

- CLARK, T. E. (2012): “Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility,” *Journal of Business and Economic Statistics*, 29, 327–341.
- CLARK, T. E., M. MCCracken, AND E. MERTENS (2018): “Modeling Time-Varying Uncertainty of Multiple-Horizon Forecast Errors,” Working Paper no. 17-15R, Federal Reserve Bank of Cleveland.
- COGLEY, T., AND T. J. SARGENT (2005): “Drifts and Volatilities: Monetary Policies and Outcomes in the post WWII US,” *Review of Economic Dynamics*, 8, 262–302.
- COGLEY, T., T. J. SARGENT, AND G. PRIMICERI (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2, 43–69.
- DIEBOLD, F. X., AND R. S. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13, 253–265.
- DOZ, C., D. GIANNONE, AND L. REICHLIN (2012): “A Quasi-Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models,” *Review of Economics and Statistics*, 94, 1014–1024.
- GALBRAITH, J. W., AND S. VAN NORDEN (2018): “Business Cycle Asymmetry and Unemployment Rate Forecasts,” Working paper.
- GANICS, G., B. ROSSI, AND T. SEKHPOSYAN (2019): “From fixed-event to fixed-horizon density forecasts: Obtaining measures of multi-horizon uncertainty from survey density forecasts,” Economics Working Papers 1689, Department of Economics and Business, Universitat Pompeu Fabra.
- GHYSELS, E., L. IANIA, AND J. STRIAUKAS (2018): “Quantile-based inflation risk models ,” Working Paper Research N. 349, National Bank of Belgium.
- GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJŠEK (2017): “Inflation Dynamics during the Financial Crisis,” *American Economic Review*, 107, 785–823.
- HENGGE, M. (2019): “Uncertainty as a Predictor of Economic Activity,” IHEID Working Papers 19-2019, Economics Section, The Graduate Institute of International Studies.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105(3), 1177–1216.
- KILEY, M. T. (2018): “Unemployment Risk,” Finance and Economics Discussion Series 2018-067, Federal Reserve Board of Governors.

KOENKER, R., AND G. BASSETT (1978): “Regression Quantiles,” *Econometrica*, 46, 33–50.

PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *Review of Economic Studies*, 72(3), 821–852.

REIFSCHEIDER, D., AND P. TULIP (2019): “Gauging the Uncertainty of the Economic Outlook Using Historical Forecasting Errors: The Federal Reserve’s Approach,” *International Journal of Forecasting*, 35, 1564–1582.

ROSSI, B., AND T. SEKHPOSYAN (2019): “Alternative Tests for Correct Specification of Conditional Predictive Densities,” *Journal of Econometrics*, 208, 638–657.

STOCK, J., AND M. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit, and Banking*, 39, 3–33.

Figure 1. SPF Forecast Errors and Financial Conditions. The figure shows quantile regression estimates of the conditional distributions of median SPF forecast errors, as a function of the NFCI value at the time of each SPF forecast. Results are reported for one quarter ahead (left column) and four quarter ahead (right column) forecasts of real GDP growth (top row), unemployment (middle row), and GDP price index inflation (bottom row).

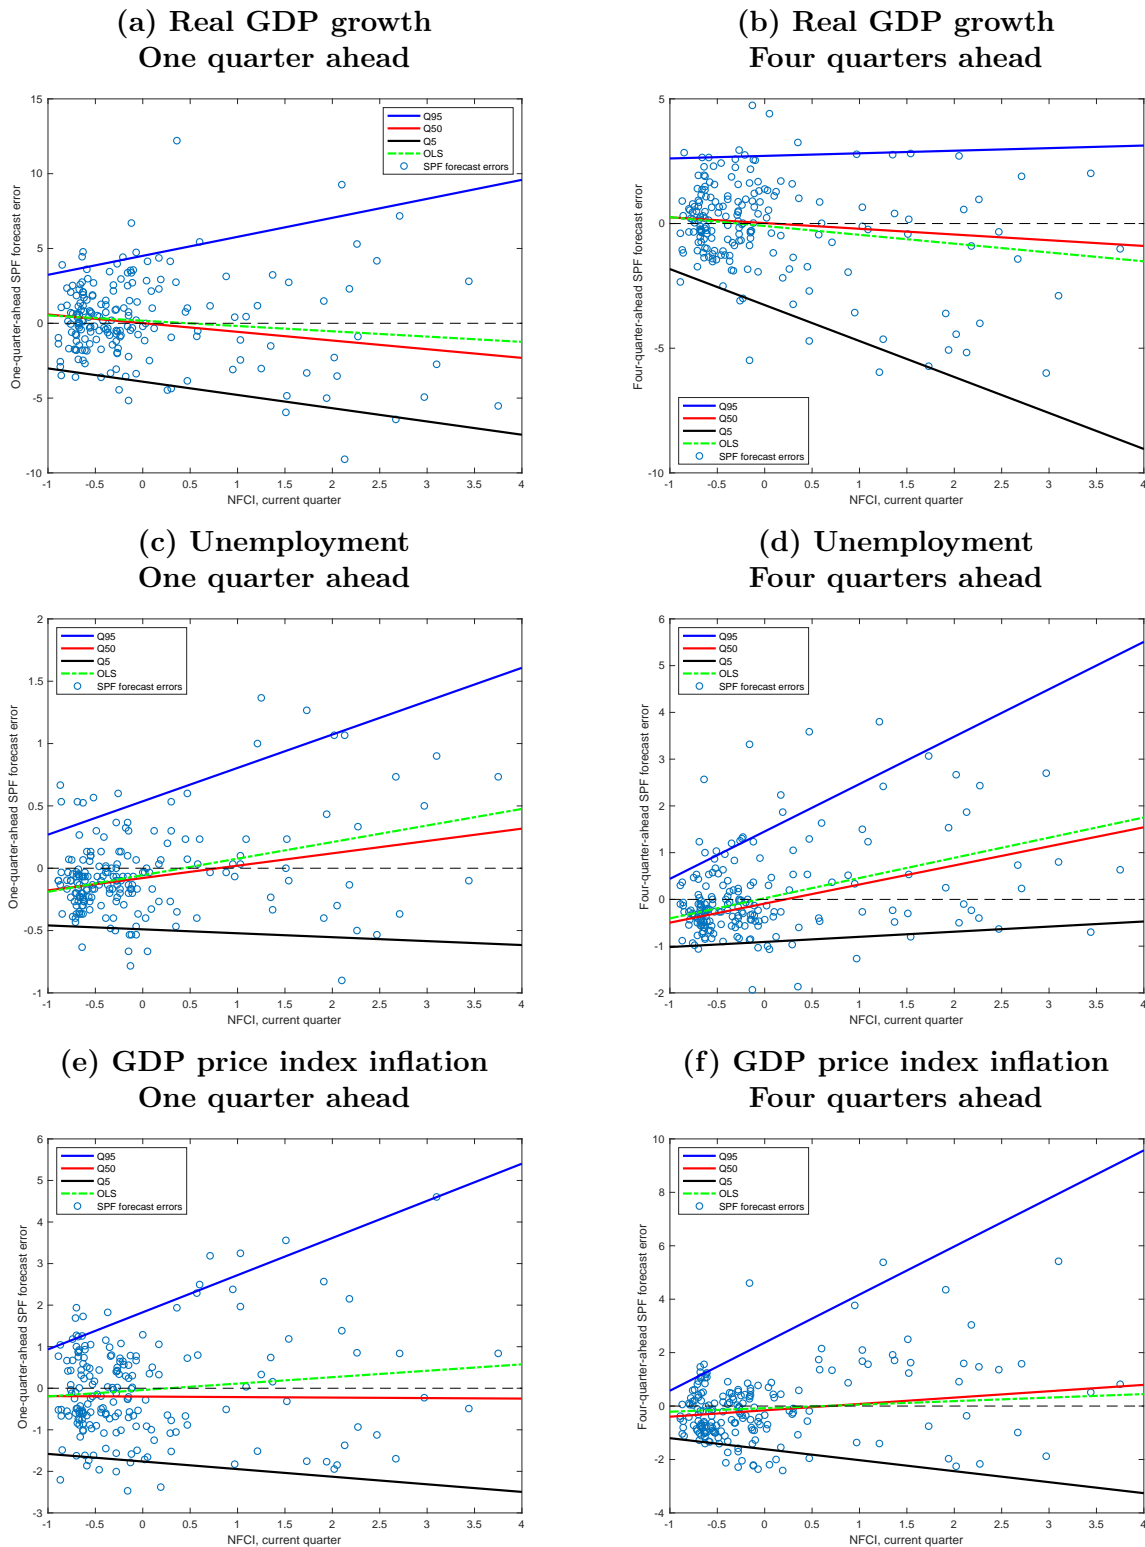


Figure 2. Estimated Quantile Regression Coefficients. The figure shows estimated coefficients from quantile regressions of median SPF forecast errors on the NFCI. Shaded bands represent 68%, 90%, and 95% confidence bounds computed under the null hypothesis that the true data generating process is a linear vector autoregression for the target variable, median SPF forecast, and NFCI, with i.i.d. Gaussian errors and four lags.

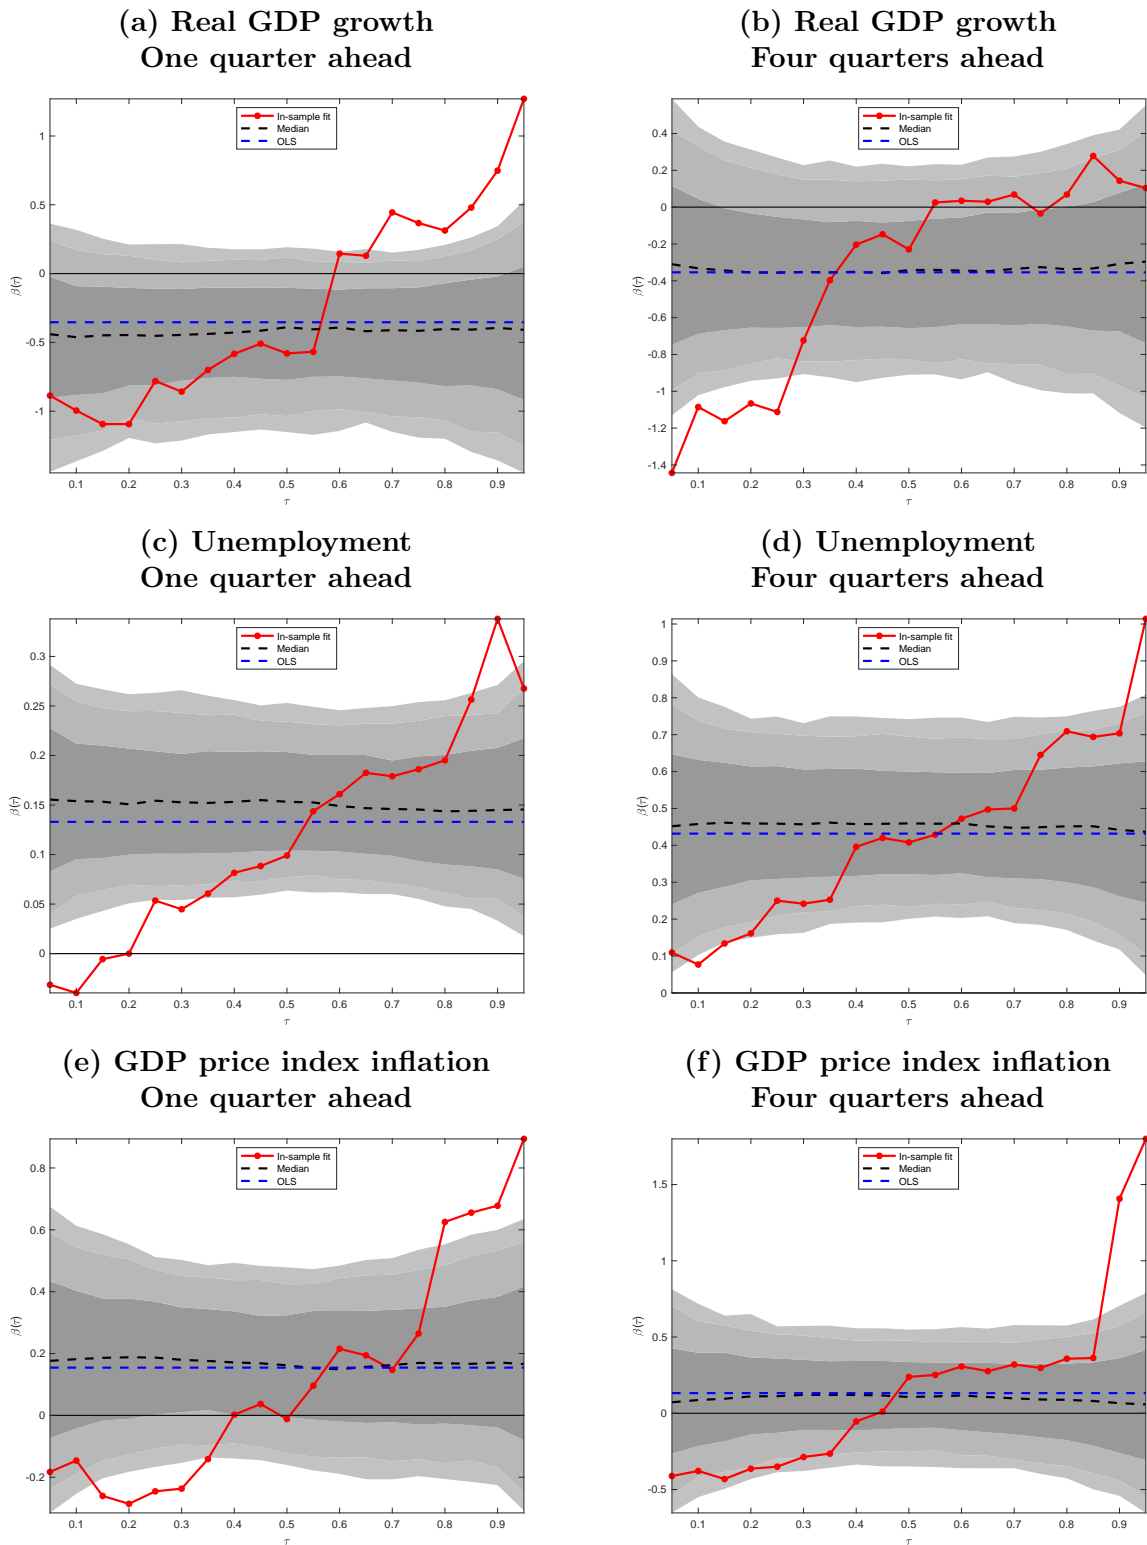


Figure 3. Estimated Quantiles. The figure shows the estimated quantiles of the predictive distributions over time. The shaded bands and black line depict the following quantiles: 5th, 10th, 25th, 50th (median, black line), 75th, 90th, 95th. The red dashed line depicts the median SPF forecast at each date, which is used in the construction of the predictive distributions.

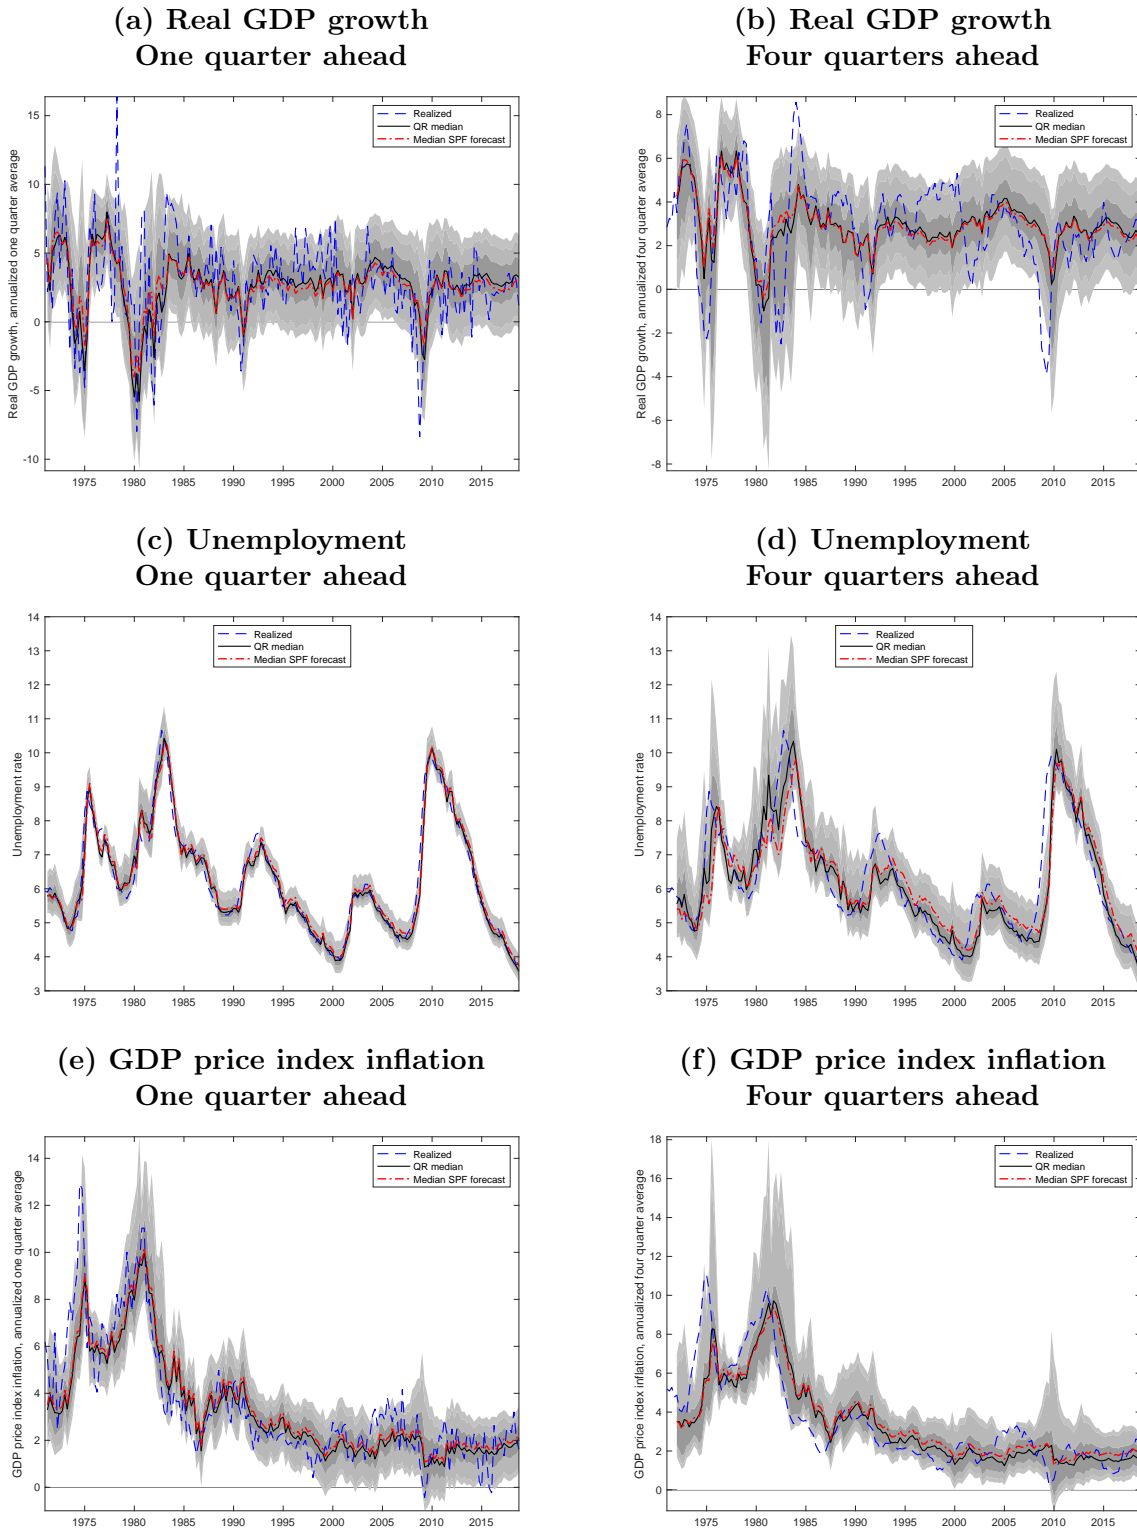


Figure 4. Predicted Forecast Error Interquartile Range vs. Median SPF Forecasts. This figure shows scatter plots of the estimated interquartile ranges (Q75-Q25) of the predictive distributions for SPF forecast errors against the median SPF forecasts. For inflation, we use different markers to differentiate between observations before 1985Q1 and after 1985Q1.

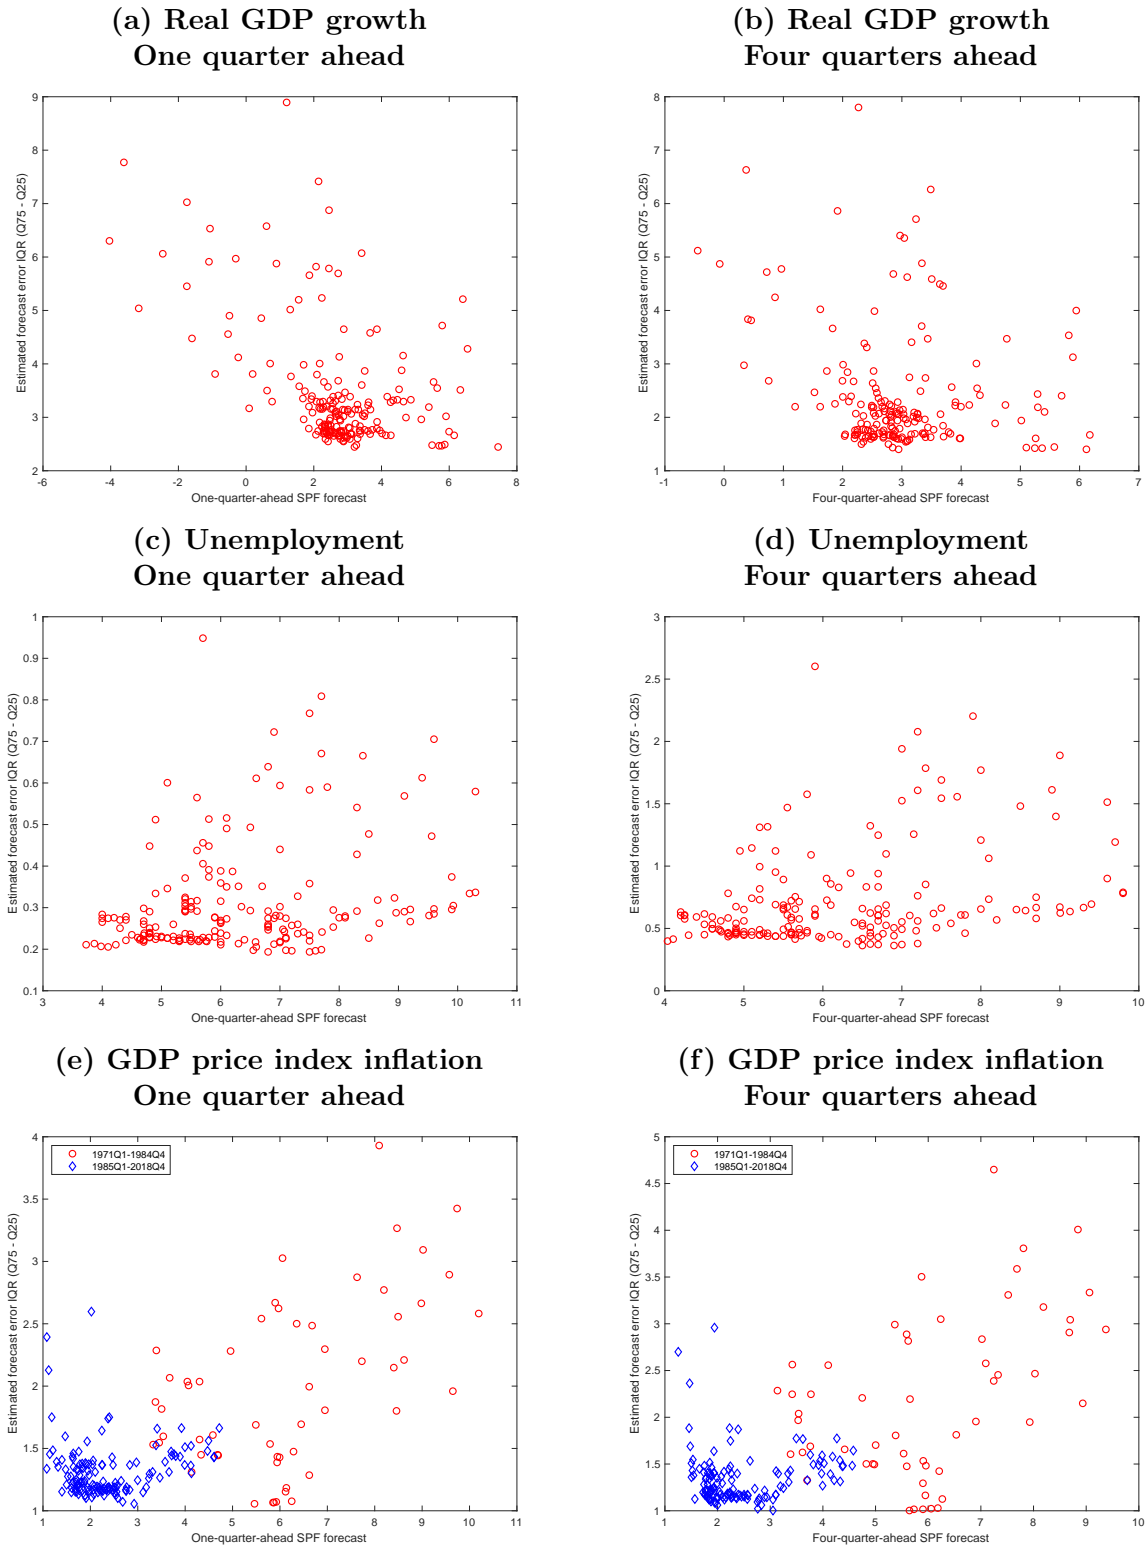


Figure 5. Predictive Densities. This figure shows estimated four quarter ahead predictive densities. The solid blue lines represent the predictive densities that condition on both the median SPF forecast and financial conditions, while the dashed orange lines represent the “unconditional” predictive densities computed from the distribution of historical forecast errors (see [Reifschneider and Tulip, 2019](#)). The vertical solid gray lines represents the median SPF forecast used in the construction of both the conditional and unconditional densities, while the red dotted lines represent realized values of the target variables.

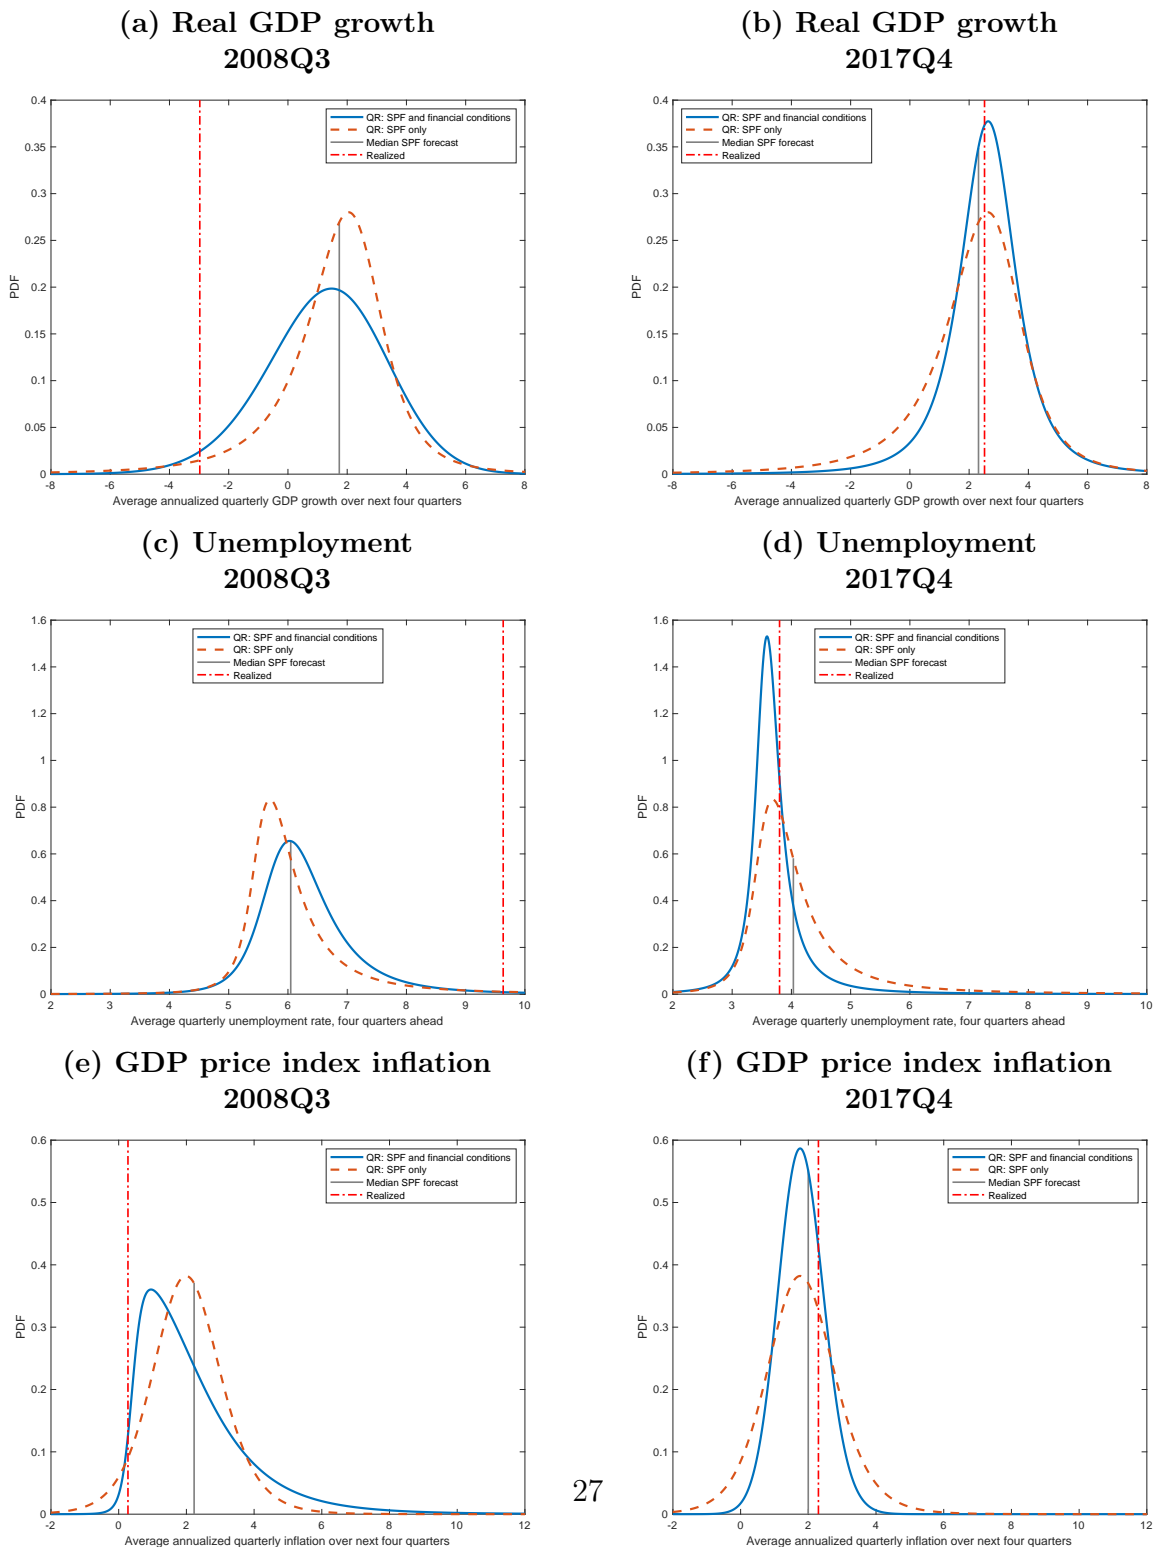


Figure 6. Expected Shortfall and Longrise. This figure shows the estimated 5% expected shortfall and 95% expected longrise of the predictive distributions. The solid blue and yellow lines represent the shortfall and longrise (respectively) for the predictive densities that condition on both the median SPF forecast and financial conditions, while the dashed orange and green lines represent the shortfall and longrise for the “unconditional” predictive densities computed from the distribution of historical forecast errors (see Reifschneider and Tulip, 2019). Gray bars denote recessions.

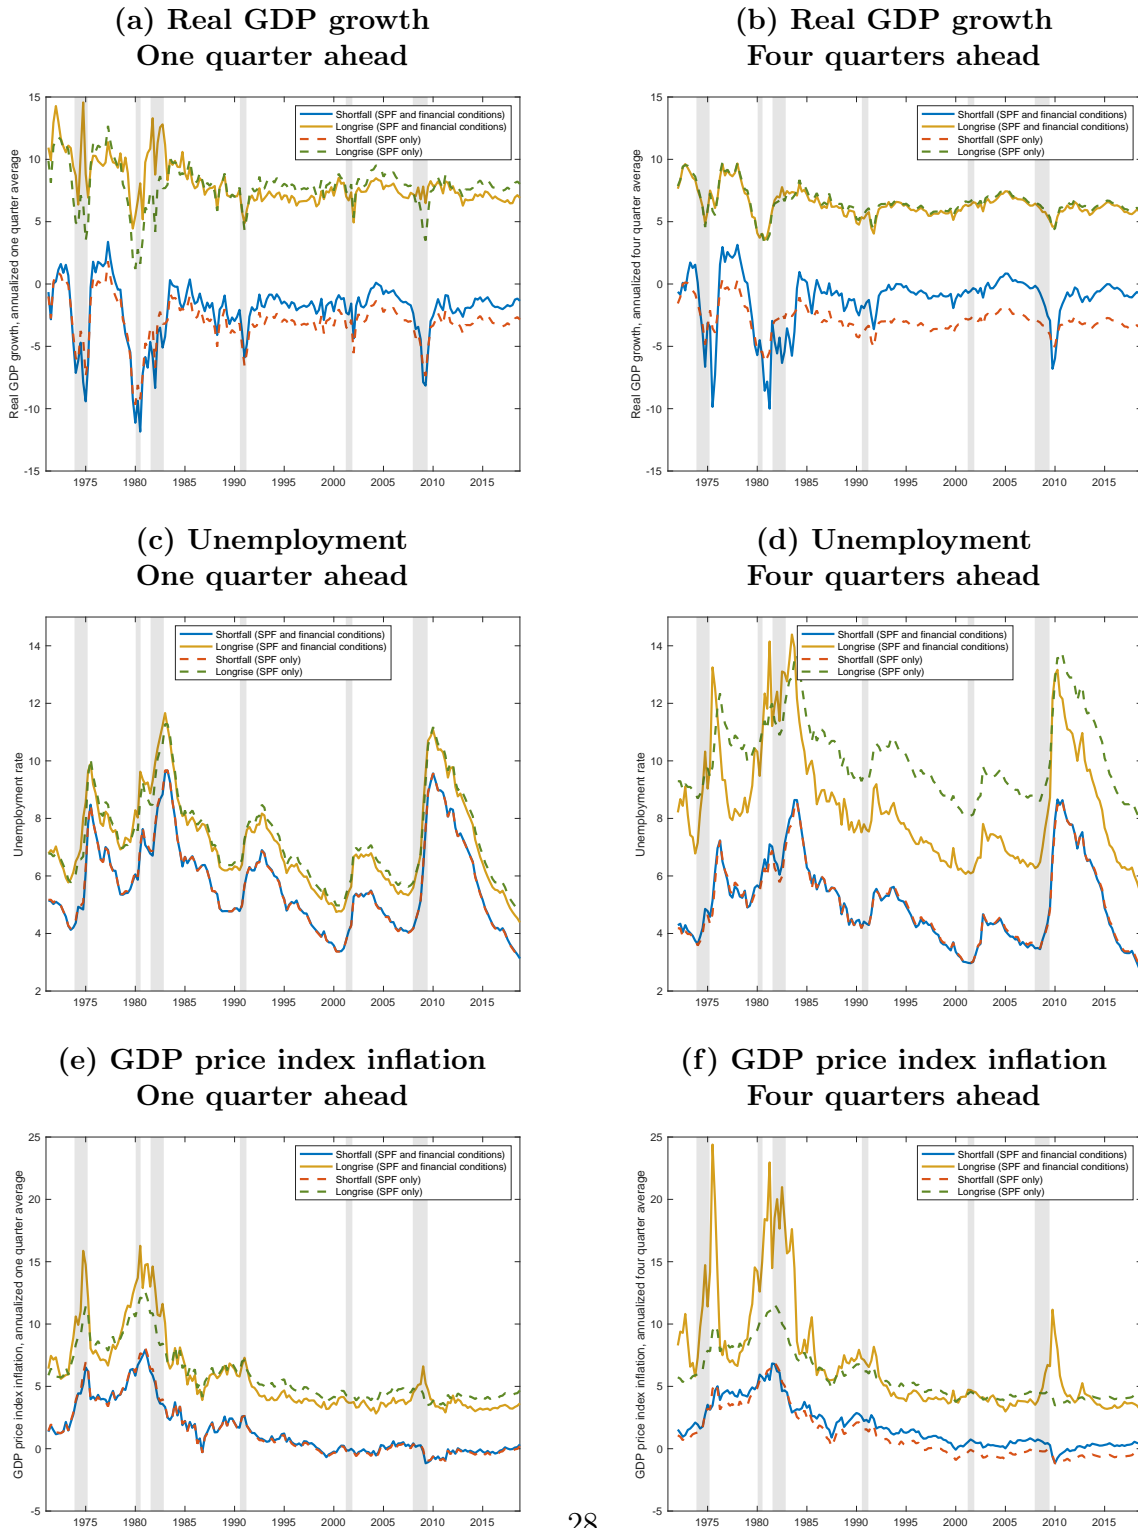


Figure 7. Out-of-Sample Predictive Scores. This figure shows predictive scores for out-of-sample density forecasts. The solid blue lines represent scores for the predictive densities that condition on both the median SPF forecast and financial conditions, while the dashed orange lines represent scores for the “unconditional” predictive density computed from the distribution of historical forecast errors (see [Reifschneider and Tulip, 2019](#)). Gray bars denote recessions. The first out-of-sample forecasts are made in 1992Q1.

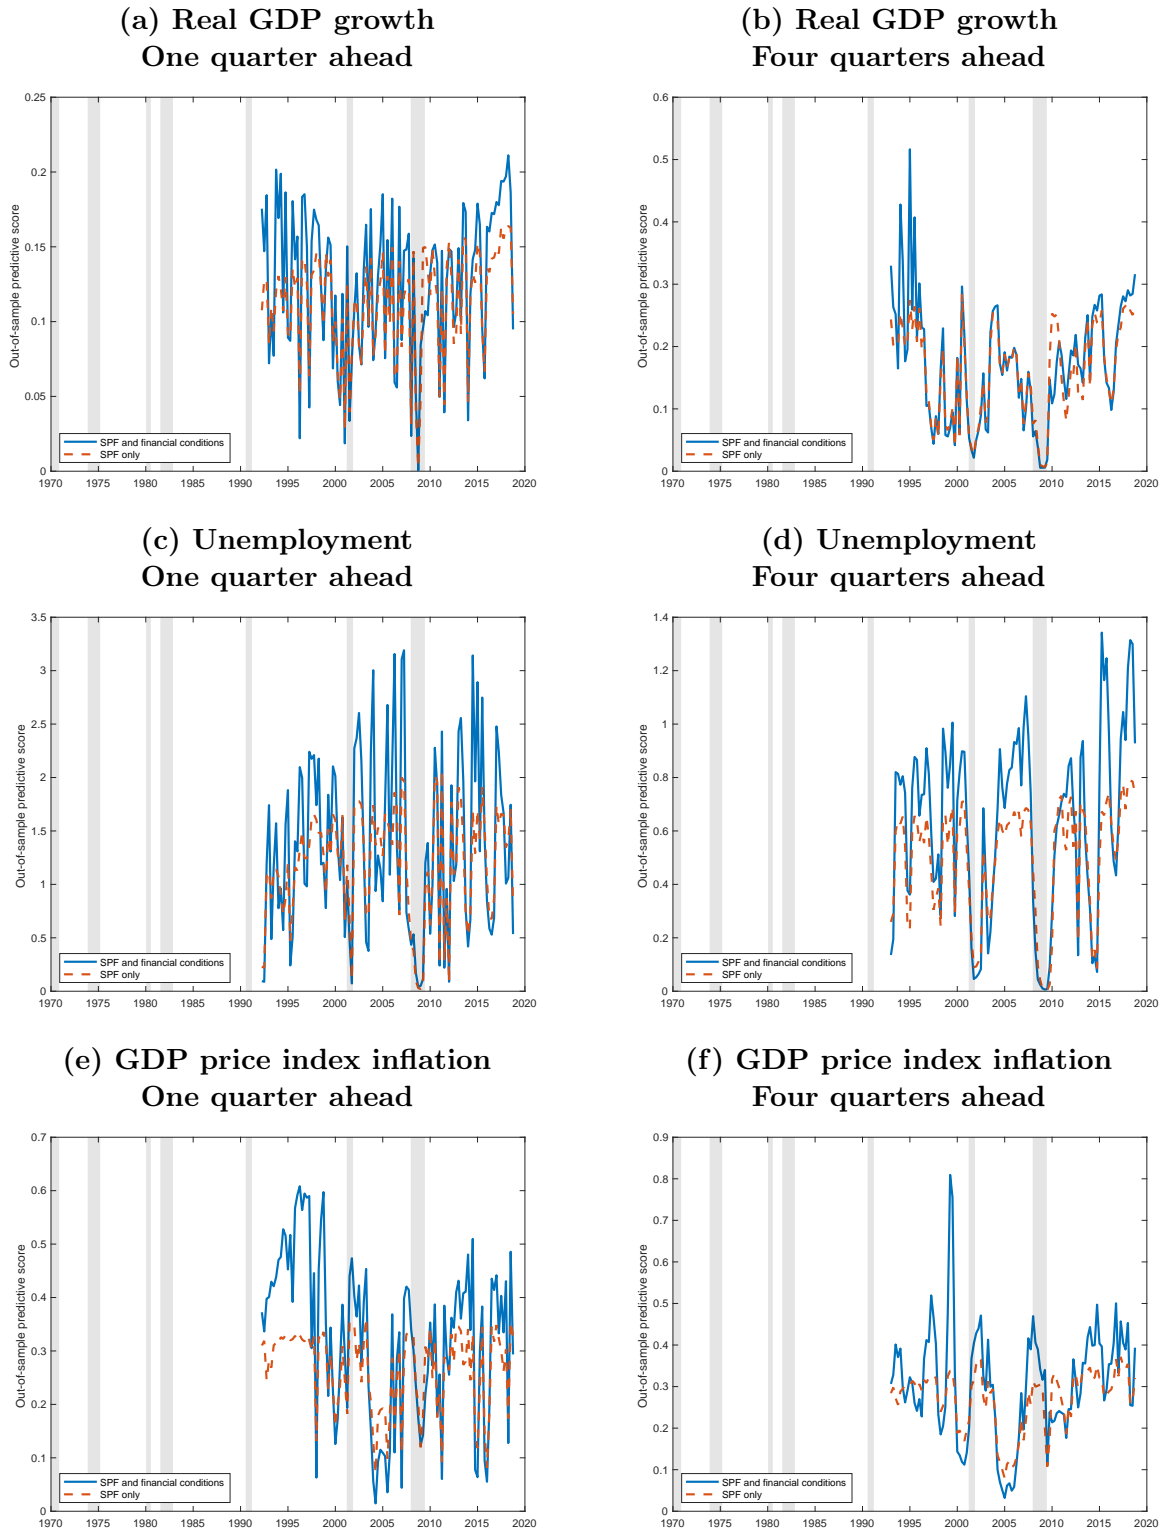


Figure 8. Out-of-Sample Probability Integral Transforms. This figure shows the empirical cumulative distribution of probability integral transforms (PITs) for out-of-sample density forecasts. The solid blue lines represent distributions for the predictive densities that condition on both the median SPF forecast and financial conditions, while the dashed orange lines represent distributions for the “unconditional” predictive densities computed from the distribution of historical forecast errors (see Reifschneider and Tulip, 2019). The first out-of-sample forecasts are made in 1992Q1. 95% confidence bands for tests of correct calibration are computed following Rossi and Sekhposyan (2019) and plotted parallel to the 45-degree line.

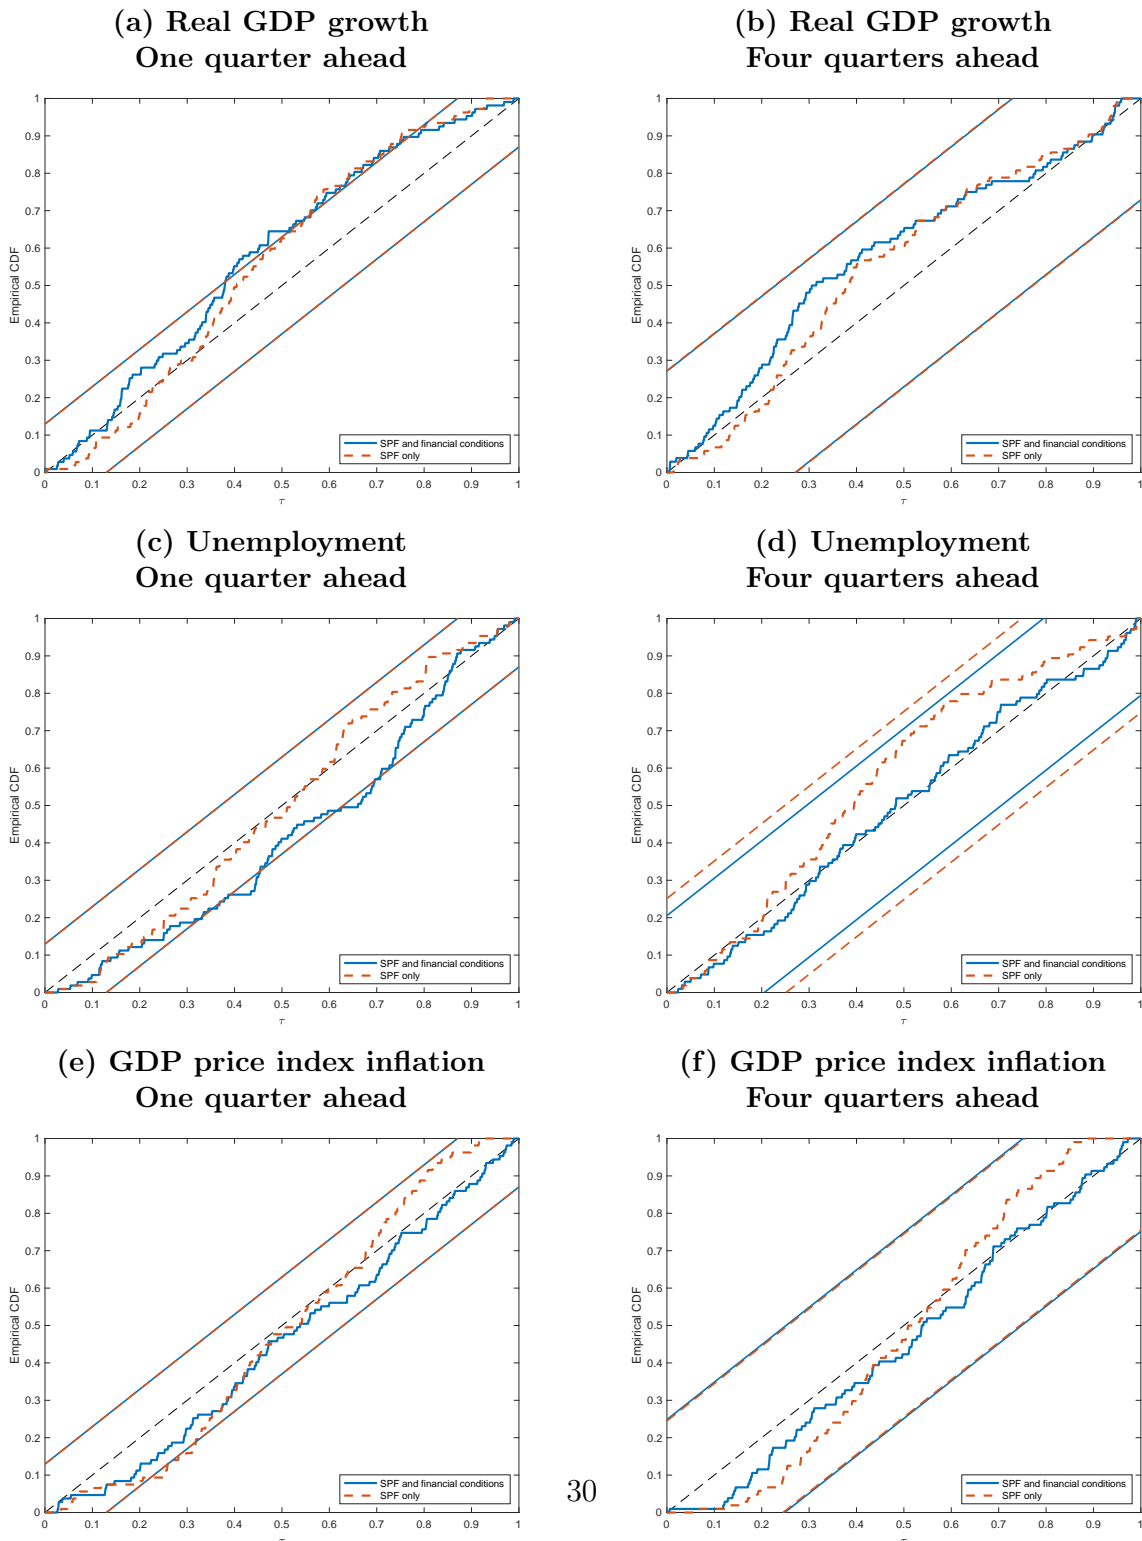


Table 1: Out-of-Sample Predictive Scores. This table reports *differences* in average out-of-sample log predictive scores between the predictive densities that condition on both the median SPF forecast and financial conditions, and the “unconditional” predictive densities computed from the distribution of historical forecast errors (see Reifschneider and Tulip, 2019). Positive values indicate superior average forecasting performance of the densities which incorporate financial conditions. The top panel reports results using expanding windows of past forecast errors to estimate the unconditional predictive densities, while the bottom panel reports results using 20-year rolling windows to estimate the unconditional predictive densities (the conditional distribution is always estimated using an expanding window). The first out-of-sample forecasts are made in 1992Q1. Heteroskedasticity- and autocorrelation-robust standard errors are reported in parentheses.

Average difference in log scores:			
SPF and financial conditions - SPF only (expanding window)			
	Real GDP growth	GDP price index inflation	Unemployment
$h = 1$	0.056	0.073	0.049
<i>(s.e.)</i>	(0.023)	(0.071)	(0.039)
$h = 4$	0.012	0.040	0.138
<i>(s.e.)</i>	(0.034)	(0.051)	(0.050)
Average difference in log scores:			
SPF and financial conditions - SPF only (20-year rolling window)			
	Real GDP growth	GDP price index inflation	Unemployment
$h = 1$	0.013	0.135	0.060
<i>(s.e.)</i>	(0.025)	(0.047)	(0.047)
$h = 4$	0.073	0.148	0.157
<i>(s.e.)</i>	(0.041)	(0.105)	(0.063)