

Krüger, Fabian; Clark, Todd E.; Ravazzolo, Francesco

Conference Paper

Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2015: Ökonomische Entwicklung - Theorie und Politik - Session: Forecasting, No. C05-V2

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Krüger, Fabian; Clark, Todd E.; Ravazzolo, Francesco (2015) : Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2015: Ökonomische Entwicklung - Theorie und Politik - Session: Forecasting, No. C05-V2, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft

This Version is available at:

<http://hdl.handle.net/10419/113077>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts*

Fabian Krüger[†]

Todd E. Clark[‡]

Francesco Ravazzolo[§]

December 19, 2014

Abstract

This paper shows entropic tilting to be a flexible and powerful tool for combining medium-term forecasts from BVARs with short-term forecasts from other sources (nowcasts from either surveys or other models). Tilting systematically improves the accuracy of both point and density forecasts, and tilting the BVAR forecasts based on nowcast means and variances yields slightly greater gains in density accuracy than does just tilting based on the nowcast means. Hence entropic tilting can offer — more so for persistent variables than not-persistent variables — some benefits for accurately estimating the uncertainty of multi-step forecasts that incorporate nowcast information.

Keywords: Forecasting, Prediction, Bayesian Analysis

JEL classification code: E17, C11, C53

*We thank Frank Schorfheide, Ellis Tallman and the participants of the workshop on “Uncertainty and Economic Forecasting” (London, April 2014), the conference on “Uncertainty and Probabilistic Forecasting during the Financial and Economic Crisis” (Heidelberg, June 2014), the conference on “Advances in Applied Macro-Finance and Forecasting” (Istanbul, September 2014), the 8th CFE conference (Pisa, December 2014) and the 25th EC2 conference on “Advances in Forecasting” (Barcelona, December 2014) for helpful comments. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland, Federal Reserve System, or Norges Bank.

[†]Heidelberg Institute for Theoretical Studies, fabian.krueger@h-its.org

[‡]Federal Reserve Bank of Cleveland, todd.clark@clev.frb.org

[§]Norges Bank and BI Norwegian Business School, francesco.ravazzolo@norges-bank.no

1 Introduction

It is commonly known that models such as vector autoregressions (VARs) or dynamic stochastic general equilibrium (DSGE) models that are effective in medium-term macroeconomic forecasting are not as effective at short-horizon forecasting. As a result, VARs and DSGE models are often combined with current-quarter forecasts, or nowcasts, from another source. One such source is a judgmental forecast from a central bank or a survey of professional forecasters, motivated by evidence that such forecasts often provide useful information beyond that contained in econometric models (e.g. Ang, Bekaert, and Wei, 2007; Faust and Wright, 2013). Alternatively, relatively accurate short-horizon forecasts can be obtained from bridging equations or factor models, surveyed in Banbura, Giannone, and Reichlin (2013) and Banbura, Giannone, Modugno, and Reichlin (2013). Compared to medium-term forecasting models, these nowcasting approaches improve near-term forecast accuracy by better adding up information in data releases for the current quarter and require dealing with differences in data release dates within the quarter (what is known as the “ragged edge” of data).¹

A number of methods for combining (VAR or DSGE) medium-term forecasts with nowcasts from another source have been used in the recent literature. Faust and Wright (2009) use short-horizon forecasts from the Federal Reserve Board’s Greenbook as jumping-off points (treating them as data, appended to the actual data) for forecasts obtained from autoregressive and factor-augmented autoregressive models of GDP growth and inflation. Similarly, Faust and Wright (2013) use current-quarter forecasts from the Survey of Professional Forecasters as jumping-off points for inflation forecasts from a range of autoregressive, Phillips curve, and DSGE models. Schorfheide and Song (2013) and Wolters (2014) treat nowcasts from the Greenbook as data in forming forecasts at subsequent horizons from, respectively, a Bayesian VAR and DSGE models. Del Negro and Schorfheide (2013) combine current quarter Blue Chip Consensus forecasts of GDP growth, inflation, and interest rates with DSGE model forecasts by treating the Blue Chip forecasts as noisy data for the quarter.² Frey and Mokinski (2014) use survey nowcasts in estimating the parameters

¹While this discussion and our analysis focuses on combining forecasts from different sources, an alternative approach is to specify a single model in mixed frequency data (e.g., quarterly and monthly). For example, Schorfheide and Song (2013) and Giannone, Monti, and Reichlin (2014) develop mixed frequency BVAR and DSGE models, respectively.

²See section 5.3 of their paper for an explanation of the methodology, which relies on the Kalman filter, and some alternatives. Smets, Warne, and Wouters (2014) use similar methodology to combine medium-term survey forecasts, treated as noise in one case and news in another, with DSGE model forecasts. Monti (2010) develops a method for combining survey-based forecasts at a range of horizons with DSGE model forecasts.

of a VAR. Finally, while we are not aware of published examples, in practice it seems to be common to use conditional forecast methods (see e.g. Doan, Litterman, and Sims, 1984) to incorporate nowcast information into medium-term forecasts from BVARs.³

As this review suggests, there is no single, standard approach for combining forecasts from medium-term projection models with short-term forecasts from other sources, either surveys or nowcasting models. In this paper, we examine the effects of using entropic tilting to combine such forecasts. Robertson, Tallman, and Whiteman (2005) introduced entropic tilting into macroeconomic forecasting, using it to impose conditions on policy rates in a small BVAR forecasting model.⁴ Cogley and Sargent (2005) used tilting to produce BVAR forecasts that conditioned on information in the Bank of England’s forecast. More recently, Altavilla, Giacomini, and Ragusa (2013) use entropic tilting to combine survey-based forecasts of short-term interest rates with yield curve forecasts from econometric models. These studies primarily focused on point forecasts — not only tilting based on point forecasts but also measuring performance in terms of point forecast accuracy.

Compared to some other existing approaches for combining forecasts from multiple sources, tilting has some advantages. First, it is non-parametric and highly flexible. This flexibility is needed here; in particular, merging a full multi-step BVAR forecast density with an external nowcast is not a conventional density combination problem in the spirit of Hall and Mitchell (2007), Geweke and Amisano (2011) and others.⁵ Second, unlike simpler approaches such as treating the nowcast as additional data, entropic tilting permits the forecaster to properly capture uncertainty around the combined forecast.

Building on the aforementioned prior research, we use tilting to improve macroeconomic forecasts from BVARs by combining them with short-term forecasts (nowcasts) from surveys and specialized models that may be more effective in short-term forecasting. Extending past research, we

³As typically implemented assuming no variance around the nowcast condition, the Gaussian conditional forecasting approach yields forecasts that are the same as those obtained by treating the nowcasts as additional data. We return to this point in Section 4.4.

⁴For the purposes of forecasting tax revenues in Iowa, Lewis and Whiteman (2014) develop an entropically tilted prior that minimizes out-of-sample mean squared error subject to a Kullback-Leibler divergence constraint that the new prior not differ too much from the original.

⁵The density combination literature is concerned with merging a set of densities f_1, \dots, f_n , all of which refer to the same (univariate or multivariate) random variable. Our setting differs in two respects. First, the nowcast refers to a univariate random variable whereas the BVAR density is jointly for five horizons. Second, the nowcast information does not come as a full density but only as a set of moment conditions.

consider tilting the BVAR forecast distributions toward not just the means but also the variances of the nowcasts, and we consider the effects of tilting on the accuracy of not only point forecasts but also density forecasts. We also compare how proper combination of forecasts via tilting affects estimates of forecast uncertainty compared to cruder approaches that do not account for nowcast uncertainty.

In our implementation, we focus on forecasts of (U.S.) GDP growth, the unemployment rate, inflation in GDP price index, and the 3-month Treasury bill rate, all produced and evaluated with real-time data. A range of studies have considered similar variable sets (e.g. Clark, 2011; D’Agostino, Gambetti, and Giannone, 2013). We use forecasts from a BVAR with stochastic volatility as in Clark and Ravazzolo (2014).⁶ The survey-based forecasts we consider are taken from the Survey of Professional Forecasters (SPF). We also consider model-based nowcasts (current-quarter forecasts); for GDP and inflation, the model uses the Bayesian mixed frequency formulation of Carriero, Clark, and Marcellino (2014), while for the unemployment and T-bill rates, we use small VARs in monthly data (to construct quarterly nowcasts), detailed below.

Broadly, our results show entropic tilting to be a flexible, powerful, and effective tool for combining forecasts from BVARs with nowcasts from other sources, either a survey-based forecast or a model-based nowcast. We show that tilting, like other approaches to combining BVAR forecasts with nowcasts (e.g., Faust and Wright, 2013), systematically improves the accuracy of point forecasts of standard macroeconomic variables. Extending previous work, we also find that tilting based on nowcast means systematically improves the accuracy of density forecasts from our BVAR. We go on to show that tilting the BVAR forecasts based on not only nowcast means but also nowcast variances yields slightly greater gains in density accuracy than does just tilting based on the nowcast means. For less persistent variables such as GDP growth, the accuracy gains tend to die out as at the forecast horizon increases, but for unemployment and interest rates, the gains carry over to horizons as long as five quarters. Our results also show that tilting towards the nowcast mean and variance produces sharper forecast distributions than tilting towards the nowcast mean only. This is because the former approach incorporates the reduced variance of the nowcast – which uses intra-quarter information – whereas the latter approach implicitly conditions on the BVAR variance. Again, these effects are much more pronounced for the more persistent variables.

⁶We obtained very similar results with a BVAR including both stochastic volatility and a steady state prior used in such studies as Clark (2011) and Wright (2013).

As to the merits of the survey-based (SPF) nowcasts compared to the model-based nowcasts, for GDP and inflation, survey forecasts from the SPF are hard to beat, so the BVAR is improved more by tilting toward the SPF nowcast than the model-based nowcasts. But for the unemployment and T-bill rates, our model-based nowcasts are more accurate than their SPF counterparts, with corresponding effects on the tilted BVAR forecasts. In a comparison of tilting on a variable-by-variable basis to tilting jointly toward the nowcasts for all four variables of the BVAR, we find that the overall differences in forecast performance for the joint treatment of variables versus the individual treatment of variables are small.

The paper proceeds as follows. Sections 2 and 3 detail the data and models, respectively. Section 4 explains the implementation of tilting and provides simple examples. Section 5 provides our results. Section 6 concludes.

2 Data

2.1 Data for models

We use quarterly data to estimate BVAR models (detailed below) for growth of real GDP, inflation in the GDP price index or deflator (henceforth, GDP inflation), unemployment rate, and the 3-month Treasury bill rate. We compute GDP growth as 400 times the log difference of real GDP and inflation as 400 times the log difference of the GDP price index, to put them in units of annualized percentage point changes. The unemployment rate and interest rate are also defined in units of percentage points (annualized in the case of the interest rate), with quarterly rates formed as within-quarter averages of monthly rates.

In constructing model-based nowcasts of growth, inflation, unemployment, and the T-bill rate using models detailed in the next section, we rely on a small set of other indicators, for reasons we detail in the model section. For nowcasting GDP growth, we use two monthly coincident indicators taken from Carriero, Clark, and Marcellino (2014): employment growth and the Institute of Supply Management's production index for manufacturing. For nowcasting GDP inflation, we use monthly inflation rates of the CPI ex food and energy, the CPI for food, the CPI for energy, the PPI for capital goods, and the price deflator for new one-family houses under construction. We form nowcasts of unemployment using monthly data on not only unemployment but also growth in payroll employment and new claims for unemployment insurance. Finally, we construct

nowcasts of the T-bill rate using monthly data on the average (for the month) T-bill rate and the 3-month and 6-month T-bill rates on the 15th of the month.

In forming all of our model-based forecasts and nowcasts, for those indicators subject to significant revisions and for which we can easily obtain the needed data, we use real-time data from the Federal Reserve Bank of Philadelphia’s Real Time Dataset for Macroeconomists (RTDSM). The variables for which we use real time data are the following: GDP, GDP price index, unemployment, and employment. Note that, for simplicity, we use “GDP” and “GDP price index” to refer to the output and price series to be forecast, even though the measures are based on GNP and a fixed weight deflator for some of the sample. As described in Croushore and Stark (2001), the quarterly vintages of the RTDSM are dated to reflect the information available around the middle of each quarter. In vintage t , the available GDP and GDP price index data run through period $t - 1$. For all remaining variables, we use currently available data obtained from the FAME database of the Federal Reserve Board of Governors: the Institute of Supply Management’s production index for manufacturing, new claims for unemployment insurance, the CPI ex food and energy, the CPI for food, the CPI for energy, the PPI for capital goods, and the price index for new home construction.

2.2 SPF forecast data

We obtain quarterly SPF forecasts of growth, unemployment, inflation, and the T-bill rate from the website of the Federal Reserve of Philadelphia. At each forecast origin, the available forecasts span five quarterly horizons, from the current quarter through the next four quarters. We take the point forecast to be the median of the SPF responses. In some entropic tilting results, we also use a measure of forecast uncertainty. In the presented results, we consider what Clements (2014) refers to as an *ex post* measure: the variance of recent (real-time) forecast errors, which we compute over the previous 20 forecasts. Specifically, with $\hat{Y}_{t,h}$ denoting the (median) SPF forecast of Y_t at forecast horizon h (i.e., the forecast for t based on data up to $t - h$), our first h -period error measure is computed (for origin t) as

$$\hat{\sigma}_{t+h,h}^2 = \frac{1}{R} \sum_{r=0}^{R-1} (Y_{t-2-r} - \hat{Y}_{t-2-r,h})^2,$$

where $R = 20$. The end point of the window reflects our definition of forecast actuals, which we explain below.⁷

⁷We also considered a variance measure based on the cross sectional dispersion of point forecasts (see, e.g., Bomberger, 1996). The results with respect to entropic tilting were qualitatively very similar to the ones reported here.

2.3 Forecast evaluation sample

We evaluate forecasts from 1988:Q3 through 2013:Q2 (and over a pre-crisis sample of 1988:Q3-2007:Q4), which requires real-time data vintages from 1988 through 2013. The start date of 1988:Q3 marks the earliest possible for a common sample size across variables; SPF forecasts of the T-bill rate do not begin until 1981:Q3, and we require additional observations for computing the forecast error variance at all horizons. For each forecast origin t starting with 1988:Q3, we use the real-time data vintage t to estimate the forecast models and construct forecasts of quarterly values of all variables for periods t and beyond. Consistent with the availability of SPF forecasts, we report results for forecast horizons of 1-5 quarters ahead. In light of the time $t - 1$ information actually incorporated in the quarterly BVAR models used for forecasting at t , the 1-quarter ahead forecast is a current quarter (t) forecast, while the 2-quarter ahead forecast is a next quarter ($t + 1$) forecast, etc. For the BVAR used to forecast the four variables of interest, the starting point of the model estimation sample is 1955:Q1; we use data for the 1948-54 period to set the priors on some parameters, as detailed in the appendix. For the GDP and inflation nowcasting models, the starting point of model estimation is always 1970:Q2 and 1965:Q1, respectively.⁸ For the unemployment rate and T-bill nowcasting models, the estimation samples begin with January 1955 and January 1965, respectively, reflecting data availability.

As discussed in such sources as Romer and Romer (2000), Sims (2002), and Croushore (2006), evaluating the accuracy of real-time forecasts requires a difficult decision on what to take as the actual data in calculating forecast errors.⁹ We follow studies such as Romer and Romer (2000) and Faust and Wright (2009) and use the second available estimates of GDP/GNP and the GDP/GNP deflator as actuals in evaluating forecast accuracy. In the case of h -quarter ahead forecasts made for period $t + h$ with vintage t data ending in period $t - 1$, the second available estimate is taken from the vintage $t + h + 2$ data set. In light of our abstraction from real-time revisions in unemployment and interest rates, we use final vintage data for evaluating forecasts of these series.

In the case of the model-based nowcasts, we instead define the variance as the variance of the posterior distribution of forecasts for the period in question (see Section 5).

⁸The 1970 start date corresponds to the sample start used in Carriero, Clark, and Marcellino (2014), which was the earliest possible in the larger set of indicators they considered. For the inflation model, the 1965 start data is the earliest possible given the availability of the predictors we consider.

⁹The GDP data available today for, say, 1985, represent the best available estimates of output in 1985. However, output as defined and measured today is quite different from output as defined and measured in 1970. For example, today we have available chain-weighted GDP; in the 1980s, output in the U.S. was measured with fixed-weight GNP.

3 Models

This section provides the specifications of our models and an overview of the estimation methods. The priors and estimation algorithms are detailed in the appendix.

3.1 BVAR specification

We focus on forecasts from a BVAR with random walk stochastic volatility, the specification that Clark and Ravazzolo (2014) found to perform relatively well in a comparison of the forecasting performance (both point and density — stochastic volatility is particularly important for density accuracy) of a range of autoregressive models with and without time-varying volatility.

Let y_t denote the $k \times 1$ vector of model variables, $B_0 = a k \times 1$ vector of intercepts, and $B_i, i = 1, \dots, p$, denote a $k \times k$ matrix of coefficients on lag i . For our set of $k = 4$ variables, we consider a VAR(p) model with stochastic volatility, with a lag length of $p = 4$:

$$\begin{aligned}
 y_t &= B_0 + \sum_{i=1}^p B_i y_{t-i} + v_t \\
 v_t &= A^{-1} \Lambda_t^{0.5} \epsilon_t, \epsilon_t \sim N(0, I_k), \Lambda_t \equiv \text{diag}(\lambda_{1,t}, \dots, \lambda_{k,t}) \\
 \log(\lambda_{i,t}) &= \log(\lambda_{i,t-1}) + \nu_{i,t}, i = 1, k \\
 \nu_t &\equiv (\nu_{1,t}, \nu_{2,t}, \dots, \nu_{k,t})' \sim N(0, \Phi),
 \end{aligned} \tag{1}$$

where $A = a$ lower triangular matrix with ones on the diagonal and non-zero coefficients below the diagonal, and the diagonal matrix Λ_t contains the time-varying variances of underlying structural shocks. This model implies that the reduced form variance-covariance matrix of innovations to the VAR is $\text{var}(v_t) \equiv \Sigma_t = A^{-1} \Lambda_t A^{-1'}$. Note that, as in Primiceri's (2005) implementation, innovations to log volatility are allowed to be correlated across variables; Φ is not restricted to be diagonal.

To estimate this BVAR, we use a Gibbs sampler, detailed in the appendix. Stochastic volatility is estimated with the algorithm of Kim, Shephard, and Chib (1998), as detailed in Primiceri (2005).¹⁰ The VAR coefficients are drawn from a conditional posterior distribution that is multivariate normal, with a GLS-based mean and variance given in Clark (2011). All of our reported results are based on samples of 5000 posterior draws, obtained by retaining every 8th draw of a total sample

¹⁰However, we modify the algorithm of Primiceri (2005) to reflect the correction to the ordering of steps detailed in Del Negro and Primiceri (2014).

of 40,000 post-burn draws, with a burn period of 5000 draws.

The posterior distributions of forecasts reflect the uncertainty due to all parameters of the model and shocks occurring over the forecast horizon. To simulate the predictive density of the BVAR, from a forecast origin of period T , for each retained draw of the model parameters or latent states (B, A, Λ_t up through T , and Φ), we: (1) draw innovations to log volatilities for periods $T + 1$ through $T + H$ from a multivariate normal distribution with variance-covariance matrix Φ and use the random walk model of $\log \lambda_{t+h}$ to compute $\lambda_{T+1}, \dots, \lambda_{T+H}$; (2) draw innovations to y_{T+h} , $h = 1, \dots, H$, from a normal distribution with variance $\Sigma_{T+h} = A^{-1}\Lambda_{T+h}A^{-1'}$, and use the vector autoregressive structure of the model along with the coefficients B to obtain draws of y_{T+h} , $h = 1, \dots, H$. The draws of y_{T+h} are used to compute the forecast statistics of interest.

3.2 Nowcast model: GDP growth

To align with the typical timing of the Survey of Professional Forecasters, we use the Bayesian mixed frequency modeling approach of Carriero, Clark, and Marcellino (2014) to produce a current-quarter forecast of GDP growth with data available around the end of the first week of the second month of the quarter. More specifically, we forecast the quarterly growth rate of GDP in month two of the current quarter based on the regression:

$$\begin{aligned} y_t &= X_t' \beta + v_t \\ v_t &= \lambda_t^{0.5} \epsilon_t, \epsilon_t \sim i.i.d.N(0, 1) \\ \log(\lambda_t) &= \log(\lambda_{t-1}) + \nu_t, \nu_t \sim i.i.d.N(0, \phi), \end{aligned} \tag{2}$$

where t is measured in quarters and the vector X_t contains predictors available at the time the forecast is formed.

The specification of the regressor vector X_t is a function of the way the monthly variables are sampled. For the timing we follow in this analysis, the vector X_t contains variables available at about the end of the first week of month 2 of quarter t . Specifically, in our implementation, it contains a constant, GDP growth in quarter $t - 1$, employment growth in month 1 of quarter t , and the ISM index in month 1 of quarter t . We use employment and the ISM because, for our information timing, these are the two major coincident indicators that are available for forecasting GDP growth in the quarter. Our model with this small set of indicators performs comparably to models with the larger sets of indicators considered in Carriero, Clark, and Marcellino (2014).

3.3 Nowcast model: inflation in GDP price index

Our nowcasting model for inflation takes the same form as that described above for GDP growth, but with a different set of indicators included in X_t . While the information set of the typical SPF response has included just week 1 of month 2 of the quarter since the Philadelphia Fed took over the survey, prior to that time the information set (and survey response date) changed over time, and it was often later in the month. Accordingly, for simplicity, we construct nowcasts of GDP inflation using (inflation rates of) monthly price indexes released in the second half of month 2 of the quarter, for the CPI ex food and energy, the CPI for food, the CPI for energy, the PPI for capital goods, and the price index for new home construction. This set of indicators reflects major measures of consumption and investment prices, as typically available in the middle of the quarter

3.4 Nowcast model: unemployment rate

To align with current SPF timing, we obtain a nowcast of the quarterly average rate of unemployment by averaging the observed rate for month 1 of the quarter with forecasts for months 2 and 3. As noted above, the typical SPF response is based on an information set that includes labor market indicators for the first month of the quarter. We produce the forecasts of months 2 and 3 of the quarter using a BVAR(3) with stochastic volatility in monthly data, for the unemployment rate, growth in payroll employment, and new claims for unemployment insurance. We include unemployment claims in the model because they are commonly thought to be a leading indicator with some predictive content for the unemployment rate (e.g. Montgomery, Zarnowitz, Tsay, and Tiao, 1998), while employment is a major coincident indicator of the business cycle that might have predictive content for the unemployment rate, which has sometimes been considered to be a lagging indicator of the business cycle. This model takes the same basic form as the one detailed above, except in monthly rather than quarterly data.

3.5 Nowcast model: T-bill rate

To align with SPF timing, we obtain a nowcast of the quarterly average 3-month T-bill rate by averaging the observed rate for month 1 of the quarter with forecasts for months 2 and 3. As SPF timing has shifted over time and respondents have access to a wide range of financial indicators, we incorporate in our interest rate nowcasting model information through the 15th of month 2

of the quarter (in the event the 15th is not a business day, we use the preceding business day). Specifically, to forecast the monthly T-bill rate for months 2 and 3 of the quarter, we use a BVAR(3) with stochastic volatility in which the variable vector y_t is monthly and contains the average 3-month T-bill rate in t and the 3-month and 6-month T-bill rates on the 15th of month $t + 1$. We include the daily rates in the model as a way of capturing current information that would be available to a forecaster under our timing assumption. We include the 6-month rate because, under the expectations hypothesis, it should contain information about the expected future path of the 3-month rate. This BVAR takes the same form as the one detailed above, except in monthly rather than quarterly data.

4 Entropic tilting: methodology and examples

This section first details the general implementation of entropic tilting and then provides examples of our use.

4.1 General methodology

In using tilting to incorporate information from survey forecasts or model-based nowcasts into medium-term forecasts from a BVAR with stochastic volatility, our starting point is a “raw” sample of I (possibly vector-valued) MCMC forecast draws,

$$f := \{y_i\}_{i=1}^I,$$

where $y_i \in \mathbb{R}^p, p \geq 1$. In the following we interpret f as a discrete distribution with I possible outcomes, each of which has probability $1/I$.¹¹ For simplicity, at this stage we suppress dependence on a certain variable, forecast origin date and forecast horizon. We consider modifying the distribution f by imposing the moment condition

$$\mathbb{E}g(\mathbf{y}) = \bar{g},$$

where $g : \mathbb{R}^p \rightarrow \mathbb{R}^m$ and $\bar{g} \in \mathbb{R}^m, m \geq 1$. The following functional optimization problem is often called “entropic tilting”:

$$\min_{\tilde{f} \in \mathbb{F}} \text{KLIC}(\tilde{f}, f) \text{ subject to } \mathbb{E}_{\tilde{f}} g(\mathbf{y}) = \bar{g}. \quad (3)$$

¹¹We’re assuming that there are no ties, which is innocuous for our continuous setting.

Here \mathbb{F} denotes the class of all discrete distributions that can be constructed by re-weighting the draws from f in an admissible way (such that the weights are positive and sum to one). Furthermore,

$$\begin{aligned} \text{KLIC}(\tilde{f}, f) &= \sum_{i=1}^I \tilde{\pi}_i \log(I\tilde{\pi}_i) \\ &= \log(I) + \sum_{i=1}^I \tilde{\pi}_i \log(\tilde{\pi}_i) \end{aligned} \quad (4)$$

is the Kullback-Leibler divergence between the candidate distribution \tilde{f} (which places weight $\tilde{\pi}_i$ on the i th MCMC draw) and f (which uses flat weights $1/I$). Finally,

$$\mathbb{E}_{\tilde{f}} g(\mathbf{y}) = \sum_{i=1}^I \tilde{\pi}_i g(\mathbf{y}_i)$$

is the expectation of \mathbf{y} under \tilde{f} . As discussed by Robertson, Tallman, and Whiteman (2005) and others, the tilting solution is given by setting

$$\pi_i^* = \frac{\exp(\gamma^{*'} g(\mathbf{y}_i))}{\sum_{i=1}^I \exp(\gamma^{*'} g(\mathbf{y}_i))} \quad (5)$$

$$\gamma^* = \arg \min_{\gamma} \sum_{i=1}^I \exp(\gamma' (g(\mathbf{y}_i) - \bar{g})). \quad (6)$$

The resulting (tilted) forecast distribution takes the form of an empirical distribution function (EDF).

We should note the following broad implications. First, the solution of the tilting problem comes in the form of a set of weights for the existing sample f . These can be used to either directly compute quantities of interest or simulate from the tilting distribution. Second, in practice tilting requires solving the minimization problem in (6), whose dimension equals the number of moment conditions (below we consider dimensions of one, two, four, and eight). This is often easy to do, given that the objective function is usually globally convex, and computing the gradient with respect to γ (and passing it to a numerical optimizer) is straightforward.¹² Third, equation (5) implies that the functional form of the tilting weights is determined by the choice of $g(\cdot)$; we explore this point in our first example below. Finally, it is possible to ensure some smoothness on

¹²In our implementation, we use the `optim` function of the R programming language (R Core Team, 2014), together with the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. In case the algorithm fails to converge, we impose a very small penalty on the L2 norm of the candidate parameter γ in order to regularize the problem. The corresponding R code is available from the first author.

the tilted forecast distribution by targeting a higher dimensional vector \bar{g} of moment conditions. We explore this below by experimenting with different sets of moment conditions.

4.2 Tilting variants considered in this paper

In the results to be presented below, we will consider the following variants of entropic tilting. First, for a given variable – indicated by the index (k) – we tilt the BVAR forecast distribution of the vector $\begin{bmatrix} y_{t+1}^{(k)} & \dots & y_{t+5}^{(k)} \end{bmatrix}$ to match a certain nowcast mean of variable k (dubbed “small m” below). Second, we tilt the same distribution to match a certain nowcast mean and variance for variable k (“small m/v”). Third, we consider the joint forecast distribution for the 20-dimensional vector $\begin{bmatrix} y_{t+1}^{(1)} & \dots & y_{t+5}^{(4)} \end{bmatrix}$ comprising four variables and five forecast horizons. We tilt this distribution to simultaneously match the nowcast means of all four variables (“big m”). Finally, we again consider the full 20-dimensional distribution and tilt it to simultaneously match the nowcast means and variances for all four variables (“big m/v”). To avoid clutter, we henceforth suppress the superindex (k) whenever we refer to a “representative” variable.

4.3 Example: Tilting the mean vs. mean and variance

In this section we illustrate how the forms of entropic tilting we will examine below are implemented and affect forecast distributions. In these examples, the model-based forecast comes from the BVAR-SV specification, and the forecast origin date is 2008:Q4. This origin date is interesting because it coincides with the recent recession becoming much more severe, which the SPF nowcasts pick up in real time but the BVAR in quarterly data by itself is slower to detect. The $p = 5$ variate vector of interest, $\mathbf{y}_{t:t+4} = [y_t, y_{t+1}, y_{t+2}, y_{t+3}, y_{t+4}]'$, contains the GDP growth rates from 2008:Q4 to 2009:Q4 (i.e., forecasts for GDP growth zero to four quarters ahead). The two panels of Figure 1 illustrate the following implementations of tilting:

- Targeting the SPF nowcast mean for GDP growth in 2008:Q4 (“small m”). This corresponds to setting

$$\begin{aligned} g(\mathbf{y}_{t:t+4}) &= y_t, \\ \bar{g} &= -2.94. \end{aligned}$$

As expected, the figure shows that the tilted distribution is located left of the raw one; this is

necessary to implement the SPF nowcast mean which is much smaller than the original one. In this case, the tilted density has a somewhat unconventional shape, featuring substantial probability mass at the lower end of its support.

- Targeting the SPF nowcast mean and variance (“small m/v”) corresponds to¹³

$$\begin{aligned} g(\mathbf{y}_{t:t+4}) &= [y_t, (y_t + 2.94)^2]' \\ \bar{g} &= [-2.94, 2.41]' . \end{aligned}$$

In this case, the tilted density again reaches the SPF nowcast mean, but the distribution is now bell-shaped. This is the result of targeting the nowcast variance in addition to the mean.

In these example, $\pi_{t,i}^*$ — the tilting weight on the i th MCMC draw — is a function of $y_{t,i}$ (the first element of the vector $\mathbf{y}_{t:t+4,i}$) alone; this follows from the specific choices of $g(\cdot)$ made here. For each example, the solutions to the tilting weights are given by the following.

- Small m:

$$\pi_{t,i}^* = \frac{\exp(-0.39 y_{t,i})}{\sum_{i=1}^I \exp(-0.39 y_{t,i})} .$$

- Small m/v:

$$\pi_{t,i}^* = \frac{\exp(-0.54 y_{t,i} - 0.21 (y_{t,i} + 2.94)^2)}{\sum_{i=1}^I \exp(-0.54 y_{t,i} - 0.21 (y_{t,i} + 2.94)^2)} .$$

Figure 2 plots the relationship for the two tilting variants. As the solutions and charts make clear, the choice of $g(\cdot)$ significantly affects the re-weighting of the draws in the tilted distribution. While the weight is a monotonic function of $y_{t,i}$ in the first variant, the relationship is bell-shaped in the second variant.

4.4 Example: Spillover effects on longer horizon forecasts

The results just presented demonstrate how tilting towards an external nowcast mean and variance yields a combined nowcast density. However, imposing moment conditions on the nowcast y_t also affects other elements of the vector of forecasts, $\mathbf{y}_{t:t+4} = [y_t, y_{t+1}, y_{t+2}, y_{t+3}, y_{t+4}]'$. While effects on other forecast horizons are difficult to see in the non-parametric solutions provided above

¹³As detailed above, the variance, intended to measure SPF forecast uncertainty, is computed as the variance of SPF forecast errors over the past 20 quarters




Figure 1: Histograms for raw and tilted samples. In each panel, the black vertical line shows the ex-post outcome of -6.55 .




Figure 2: Tilting weight $\pi_{t,i}^*$ as a function of the first element y_t of the vector $y_{t:t+4}$. The solid line corresponds to tilting towards the nowcast mean only; the dashed line corresponds to tilting towards the nowcast mean and variance.

in equations (5) and (6), we can use a Gaussian benchmark case (extending the example in Robertson, Tallman, and Whiteman, 2005) to provide some intuition.

Consider a five-variate vector $\mathbf{y}_{t:t+4}$, and suppose a forecaster uses a multivariate normal distribution $f = \mathcal{N}(\theta, \Sigma)$, where $\theta = [\theta_1 \dots \theta_5]'$ and Σ is a positive definite matrix with elements $\Sigma_{i,j}$.¹⁴ Consider the tilted density f^* which imposes that the first system variable have mean μ_1 and variance $\Omega_{1,1}$. Then, f^* is multivariate normal $\mathcal{N}(\mu, \Omega)$, with parameters

$$\mu_{2:5} = \theta_{2:5} + \Sigma_{1,1}^{-1} \Sigma_{1,2:5} (\mu_1 - \theta_1), \quad (7)$$

$$\Omega_{2:5,2:5} = \Sigma_{2:5,2:5} - \Sigma_{2:5,1} \Sigma_{1,1}^{-1} \Sigma_{1,2:5} \times \left(1 - \frac{\Omega_{1,1}}{\Sigma_{1,1}}\right), \quad (8)$$

$$\Omega_{2:5,1} = \Sigma_{2:5,1} \Sigma_{1,1}^{-1} \Omega_{1,1}, \quad (9)$$

where $A_{i:j, k:l}$ denotes the matrix consisting of rows $i : j$, columns $k : l$ of any matrix A . We write $A_{i:j, k}$ if the “matrix” is a column vector, and $A_{i, k:l}$ if it is a row vector. This Gaussian example yields the following implications.

- In the special case that y_t is fixed, such that $\Omega_{1,1} = 0$, we end up at the textbook formulas for conditioning in the multivariate normal distribution. That is, entropic tilting is exactly the same as conditional forecasting. It is also exactly the same as treating the nowcast as data or jumping-off points for forecasts at subsequent horizons (Faust and Wright, 2009, 2013). Appendix A presents a simple example which demonstrates this equivalence.
- The special case that $\Omega_{1,1} = \Sigma_{1,1}$ corresponds to a scenario in which the tilted variance for y_{t+1} is the same as the un-tilted variance. Interestingly, the same solution obtains when targeting a mean of θ_1 only, without making a tilting assumption about $\Omega_{1,1}$ – see e.g. Altavilla, Giacomini, and Ragusa (2013, Section 3.1). This implies that, at least in the Gaussian case, targeting the mean only is equivalent to targeting the mean *and* the original variance.
- If $\Omega_{1,1} < \Sigma_{1,1}$ and $\Sigma_{2:5,1} \neq 0$ in (8), tilting “reduces” (in a matrix sense) the variance of the forecasts at other horizons.
- The magnitude of the impact on $\theta_{2:5}$ and $\Omega_{2:5,2:5}$ mainly depends on $\Sigma_{2:5,1}$, the correlation of the nowcast with the other horizons. This correlation matrix reflects the persistence of the time series, as predicted by the BVAR.

¹⁴In practice, the parameters in θ and Σ will depend on past realizations of y . We suppress this here for simplicity.

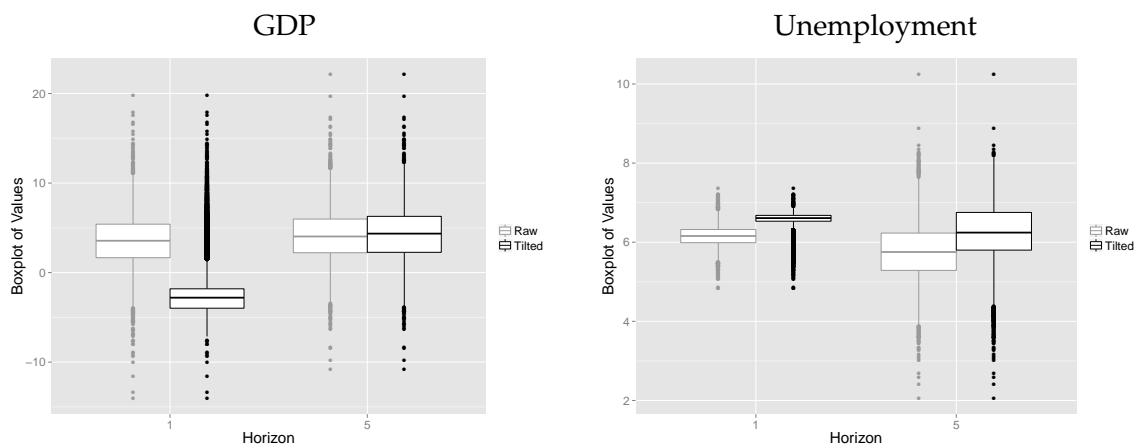


Figure 3: Vertical axis: Box plots of raw and tilted forecast distributions at origin date 2008:Q4. Horizontal axis: Forecast horizon h . Boxes range from the 25 % to the 75 % quantile of a forecast distribution; the end of the upper vertical line is the 75% quantile plus 1.5 times the interquartile range. Forecast draws exceeding that value are plotted as points.

The results we obtain for our examples using the non-parametric entropic solution are broadly consistent with the implications of the Gaussian specification. In the interest of brevity, we suppress the details and provide a simple example here. As the forecast horizon increases from the current quarter through the following four quarters, tilting has more persistent effects on the forecasts of the unemployment and T-bill rates (the most persistent variables) than the forecasts of GDP growth and inflation. That is, tilting based on the forecast for 2008:Q4 has larger effects on the 2009:Q4 forecasts for the unemployment and T-bill rates than on the 2009:Q4 forecasts for growth and inflation. Figure 3 illustrates these points for GDP and unemployment forecasts. The figure uses box plots to visualize the raw and tilted distributions, for the nowcast (2008:Q4) and the one-year-ahead forecast (2009:Q4). For GDP (left panel), tilting leads to a strong downward revision of the nowcast. At the same time, the mean and variance of the one-year-ahead forecast are not strongly revised. For unemployment (right panel), the mean of the one-year-ahead forecast is revised almost as strongly as that of the nowcast itself, with the revision pointing to higher unemployment rates in both cases. The stronger spillover effects for unemployment (compared to GDP) are due to larger entries in $\Sigma_{2;5,1}$, which represents the persistence of the series (as predicted by the BVAR). For example, the predicted first-order autocorrelation is only 0.20 in the case of GDP, but 0.76 in the case of unemployment.

4.5 Example: Joint tilting across variables and forecast horizons

In the examples above, we performed tilting on a variable-by-variable basis. However, the methodology allows us to directly consider the full multivariate forecast distribution comprising all variables and forecast horizons. Forecasts tilted based on the joint set of nowcasts should be conceptually preferable. In practice one is likely to have available and be interested in using nowcasts for all model variables. Tilting based on the set of nowcasts together yields a single set of forecasts of all variables that reflects the BVAR-captured historical relationships among the variables. To see the logic of it, consider the very simple (and parametric) approach of incorporating nowcasts through Gaussian conditional forecasting. One could first condition on the nowcast for variable 1 and produce BVAR forecasts for all variables, then condition on the nowcast for variable 2 and produce BVAR forecasts for all variables, etc. This would of course produce an entire set of alternative forecasts for each variable, reflecting conditions imposed one at a time. In practice, it is more likely the case that the entire set of nowcast conditions would be imposed at once, to obtain a single set of forecasts that reflects the joint set of conditions. The reasoning is the same for a joint approach to entropic tilting.

Interestingly, “big” tilting turns out to be a more stringent version of the four “small” problems (one variable at a time). To see this, denote by f the full (20 dimensional) empirical MCMC distribution for all variables and horizons, by $f^{(k)}$ the distribution for variable k (five dimensions = forecast horizons), and by $\mathcal{C}^{(k)}$ the set of moment conditions imposed on variable k . Then, big tilting solves

$$\min_{\tilde{f}} \text{KLIC}(\tilde{f}, f) \quad \text{subject to} \quad \mathcal{C}^{(1)} \cup \dots \cup \mathcal{C}^{(4)}. \quad (10)$$

Small tilting for variable k solves

$$\min_{\tilde{f}^{(k)}} \text{KLIC}(\tilde{f}^{(k)}, f^{(k)}) \quad \text{subject to} \quad \mathcal{C}^{(k)}. \quad (11)$$

Notice that the candidate distributions \tilde{f} from (10) and $\tilde{f}^{(k)}$ from (11) are both characterized by a weight vector of dimension 5000 (the number of MCMC draws), and the raw distributions f and $f^{(k)}$ both feature flat weights. Hence, for a given weight vector, we have that $\text{KLIC}(\tilde{f}, f) = \text{KLIC}(\tilde{f}^{(k)}, f^{(k)})$, and thus (11) is equivalent to solving

$$\min_{\tilde{f}} \text{KLIC}(\tilde{f}, f) \quad \text{subject to} \quad \mathcal{C}^{(k)}. \quad (12)$$

Hence the minimization problem (12) for small tilting is a variant of the problem (10) for big tilting, featuring a less stringent set of constraints. This implies that big tilting will typically entail a more drastic move away from the baseline distribution compared to small tilting. In order to illustrate this point, we again consider the 2008:Q4 example, and tilting based on the SPF nowcast means *and* variances (big m/v, small m/v). Table 1 illustrates the logical necessity that the four small tilting approaches are KLIC-closer to raw MCMC than the big tilting approach (although, in this steep recession example, all approaches are fairly far away from the equal weights of raw MCMC, because the nowcast of growth is so different from the BVAR forecast). Similarly, Figure 4 presents “Lorenz curves” for the observation weights resulting from big versus small tilting. The figure shows that the weights for big tilting are highly unequal, with a small number of dominant MCMC draws (for example, the ten largest weights sum up to 0.5548). The weights for the small tilting problems are much more equal, i.e. the Lorenz curves in Figure 4 are left of the one for big tilting. Note that the small tilting method for inflation generates by far the most uniform weights (leftmost Lorenz curve), which is in line with the fact that in 2008:Q4 the current quarter MCMC forecast (2.74) is already close to the SPF mean nowcast (2.6).

In this example, the reliance of the “big” tilting distribution on a few Monte Carlo draws is striking. However, in our sample, it does not seem harmful in terms of forecasting performance. In particular, we experimented with importance sampling approaches which deliberately over-sample the tails of the BVAR distribution. While producing much more balanced tilting weights, these approaches did not result in more accurate distribution forecasts in terms of the Cumulative Ranked Probability Score (see Section 5 below). In the interest of simplicity, we hence focus on standard sampling approaches in the following.

Method	Variable	KLIC divergence from equal weights
big m/v	(all)	4995.603
	GDP	4994.113
small m/v	UNE	4993.483
	INF	4991.631
	TBI	4992.903

Table 1: KLIC divergences for big versus small tilting in the 2008Q4 example. See equation (4) for the underlying formula.




Figure 4: Lorenz curve of weights in the 2008Q4 example. *Reading example:* The topmost line indicates that for inflation, the 25 % smallest weights (horizontal axis) add up to roughly 0.07 (vertical axis).

5 Forecast results

We first consider the accuracy of point forecasts (defined as posterior means), using root mean square errors (RMSEs). We then consider density forecasts, using the average continuous ranked probability score (CRPS). Studies such as Gneiting and Raftery (2007) and Gneiting and Ranjan (2011) discuss the advantages of the CRPS over other measures. The CRPS, defined such that a lower number is a better score, is given by

$$\text{CRPS}_t(y_{t+h}^o) = \int_{-\infty}^{\infty} (F(z) - 1_{\{y_{t+h}^o \leq z\}})^2 dz, \quad (13)$$

where y_{t+h}^o denotes the observed outcome, F denotes the cumulative distribution function associated with the (posterior) predictive density f , and $1_{\{y_{t+h}^o \leq z\}}$ denotes an indicator function taking value 1 if $y_{t+h}^o \leq z$ and 0 otherwise. In our analysis, F takes the form of an empirical distribution function, whereby the observation weights are equal in the case of raw MCMC but not in the case of tilted distributions. We employ the algorithm by Hersbach (2000, Section 4) – which allows for non-equal weights – to compute the CRPS in both cases.

In order to test the statistical significance of differences in predictive performance, we consider

pairwise tests of equal predictive accuracy (henceforth, EPA; Diebold and Mariano, 1995; West, 1996) in terms of either RMSE or CRPS. All EPA tests we conduct compare the raw BVAR forecasts against a given variant of entropic tilting, using two sided tests and standard normal critical values. Based on simulation evidence in Clark and McCracken (2013), in computing the variance estimator which enters the test statistic, we employ a rectangular kernel truncated at lag $h-1$ and incorporate the finite sample correction due to Harvey, Leybourne, and Newbold (1997).¹⁵

Our use of EPA tests based on normal critical values may be viewed as an approximation that simplifies an inference problem that, in our context, features many complexities — possible nesting of forecasts and tilting that bears similarities to conditional forecasting — not necessarily easily dealt with in the forecast evaluation literature (see e.g. Clark and McCracken, 2013, 2014).¹⁶ Under the asymptotics of Giacomini and White (2006), a test of a null of equal forecast accuracy in the finite sample (at estimated model parameters) is generally normally distributed, subject to a requirement that the model parameters be estimated with a rolling sample of data. While we have not estimated the BVAR with a rolling sample of data, Monte Carlo evidence in Clark and McCracken (2013) indicates that, with nested models estimated with an expanding data sample (the approach we have used with our BVAR), EPA tests compared against normal critical values can be viewed as a somewhat conservative (modestly under-rejecting compared to nominal size) test for equal accuracy in the finite sample.

As noted in Section (4), we consider the following variants of entropic tilting: small m , in which we tilt the BVAR forecast distribution of the vector $\begin{bmatrix} y_{t+1}^{(k)} & \dots & y_{t+5}^{(k)} \end{bmatrix}$ to match a certain nowcast mean of variable k ; small m/v , in which we tilt the forecast distribution of variable k to match the nowcast mean and variance; big m , in which we tilt the entire 20 element vector of variables and horizons to simultaneously match the nowcast means of all four variables; and big m/v , in which we tilt the entire 20 element vector of variables and horizons to simultaneously match the nowcast means and variances of all four variables. We separately apply each of the tilting variants to nowcasts from the SPF and from the nowcasting models described in Section 3. In the case of

¹⁵In the rare cases in which the rectangular kernel yields a negative variance estimate, we resort to Bartlett kernel weights (Newey and West, 1987) in order to ensure positivity. In these cases, we use the automatic bandwidth selection procedure of Newey and West (1994) as implemented in R's `sandwich` package (Zeileis, 2004).

¹⁶At the one-step horizon, the tilted forecasts are, by construction, essentially the nowcasts, so the benchmark BVAR forecast and each tilted forecast are not nested, in which case the application of the EPA test is valid. At longer horizons, the picture is less clear; the tilted forecasts are functions of the nowcasts and the underlying BVAR forecasts. Under some conditions, at horizons of 2 or more periods, the tilted and BVAR forecasts could be seen as nested under a null of equal accuracy. Regardless, the multi-step tilted forecasts bear similarities to conditional forecasts; Clark and McCracken (2014) propose a modified test of EPA necessary for application to conditional forecasts.

the SPF-based results, the variances used in tilting are computed (using real time data available at each forecast origin) as described in section 2.2; for model-based nowcasts, the variances used in tilting are defined as the variance of the posterior distribution of BVAR forecasts for the period in question.

In light of the potentially large effects of the Great Recession of 2007-2009 and ensuing slow recovery, we provide results for a sample that ends in 2007:Q4, the business cycle peak as dated by the National Bureau of Economic Research, and a full sample that ends in 2013:Q2.

5.1 Comparison of current-quarter forecasts

Before examining the effects of entropic tilting of BVAR forecasts toward different nowcasts, it is useful to compare the accuracy of current quarter forecasts from the BVAR, the SPF, and the nowcasting models. Table 1 provides the RMSEs and CRPS scores of each current quarter forecast (except that we don't provide CRPS scores for the SPF forecasts because the SPF does not include the forecast density information needed to compute the CRPS over our sample). These results yield the following findings.

- Consistent with previous studies, current-quarter forecasts from the SPF and the models designed for nowcasting are generally more accurate than the current quarter forecasts from the BVAR. For example, in the case of GDP growth over the pre-crisis sample, the SPF and mixed frequency nowcasting models have RMSEs of 1.580 and 1.682, respectively, compared to the BVAR's RMSE of 1.974. The differences (for GDP growth and unemployment) are even larger in the full sample than in the pre-crisis sample.
- Compared to SPF, some of the nowcasting models yield better accuracy, while others yield less accuracy. For GDP growth, the mixed frequency nowcasting model is almost as accurate as SPF in the pre-crisis sample and modestly less accurate in the full sample, reflecting the better job the SPF did in picking up the sharp downturn of the Great Recession (see the discussion in Carriero, Clark, and Marcellino (2014)). For unemployment and the T-bill rate, the model-based nowcasts are at least somewhat more accurate than the SPF forecasts.¹⁷ For instance, in the pre-crisis sample, the model-based nowcast of the T-bill rate has an RMSE of

¹⁷Montgomery, Zarnowitz, Tsay, and Tiao (1998) also find that quarterly forecasts of the unemployment rate are greatly improved when the forecasting model takes account of the unemployment rate in the first month of the quarter.

0.066, compared to the SPF RMSE of 0.133.

- The CRPS scores move closely in line with the RMSEs, both qualitatively and in terms of the magnitude of improvements of nowcasts over current-quarter forecasts from the BVAR.

		<i>Pre Crisis (88Q3 – 07Q4)</i>				<i>Complete (88Q3 – 13Q2)</i>			
		GDP	UNE	INF	TBI	GDP	UNE	INF	TBI
<i>RMSE</i>	<i>SPF</i>	1.580	0.125	0.767	0.133	1.591	0.151	0.809	0.133
	<i>BMF</i>	1.682	0.095	0.861	0.066	1.899	0.095	0.985	0.072
	<i>BVAR</i>	1.973	0.157	0.877	0.406	2.396	0.235	0.939	0.406
<i>CRPS</i>	<i>BMF</i>	0.960	0.053	0.495	0.037	1.048	0.053	0.554	0.035
	<i>BVAR</i>	1.123	0.089	0.506	0.214	1.277	0.117	0.541	0.214

Table 2: Root Mean Squared Errors and correlations for different nowcasts (SPF = Survey of Professional Forecasters, BMF = Bayesian Mixed Frequency, BVAR = Bayesian VAR with stochastic volatility). SPF and BMF use data up to daily frequency; BVAR is based on quarterly data.

5.2 Main results

We now consider tilting longer-horizon forecasts based on just current-quarter forecasts. As noted above, while it is common to condition longer-horizon model-based forecasts on nowcasts from other sources (judgment, models, etc.), tilting may be seen as a more flexible generalization. Importantly, it permits the forecaster to estimate the uncertainty around longer-horizon forecasts to properly reflect uncertainty surrounding the nowcast. The results are presented in Table 3. In light of the common central bank practice of reporting growth and inflation rates that are averages over four quarters, the table provides results for (annualized) quarterly forecasts four and five quarters ahead and for four-quarter averages four and five quarters ahead (in the columns “4*” and “5*”, respectively). These results yield the following key take-aways.

- In all cases, tilting forecasts based on just the nowcast (point or point and variance) from either the SPF or the nowcasting models improves the accuracy of point and density forecasts at horizons of two and three quarters.¹⁸ For example, in the full sample results for GDP growth at the three quarters-ahead horizon, under the small m approach, tilting toward the nowcast from the mixed frequency model lowers the RMSE of the BVAR forecast from 2.663

¹⁸In addition, tilting the nowcast (= one quarter ahead forecast) improves the forecast at that horizon itself. This result is to be expected from the preliminary findings in Table 2.

to 2.575; the difference is significant at the 1% level (two-sided test). For the same sample and horizon, tilting the T-bill forecasts toward the model-based nowcasts (small m approach) lowers the RMSE of the BVAR from 1.060 to 0.820 (difference significant at 1% level). Tilting has quantitatively similar effects on density forecast accuracy as measured by the CRPS.

- At forecast horizons of four and five quarters, the performance of forecasts tilted toward nowcasts is more mixed. At these horizons, tilting has relatively little benefit for forecasts of GDP growth and inflation. But it has some benefit for forecasts of the more persistent variables, the unemployment and T-bill rates. As an example, at the five step horizon, tilting the T-bill forecasts toward the model-based nowcasts lowers the full-sample RMSE of the BVAR from 1.594 to 1.428 (difference significant at 1% level). Again, tilting has quantitatively similar effects on density forecast accuracy as measured by the CRPS. These patterns align with the observations drawn in the illustration of Section 4.4.
- Tilting the BVAR forecasts toward both the mean and variance of nowcasts (“ m/v ”) — rather than just the mean or point nowcast (“ m ”)— yields small additional gains in density forecast accuracy. This pattern is very robust: In all scenarios (variables, subsamples and forecast horizons) covered by Table 3, the CRPS score of the best “ m/v ” specification is smaller than that of the best “ m ” specification. For example, in the case of the unemployment rate, $h = 2$ and the full sample, the best specification based on the mean only (“BMF big m ”) attains a CRPS of 0.145, whereas the best mean/variance specification (“BMF small m/v ”) attains a CRPS of 0.137. By comparison, the CRPS of the raw BVAR distribution is 0.217.
- Jointly considering the nowcasts of all four variables (“big m/v ”) versus considering all variables separately (“small m/v ”) tends to perform similarly well, with each approach outperforming the other in a number of scenarios. While one interpretation might be that joint treatment offers little overall advantage, an alternative interpretation might be that it is conceptually preferable for imposing tilting at a system level and does so at little (if any) cost, in terms of forecast accuracy.
- For GDP and inflation, both ‘tilting targets’ (survey and model nowcasts) perform comparably before the crisis, whereas survey nowcasts yield better results in the full sample. This likely reflects the superiority (noted above) of the survey in picking up the Great Recession on a timely basis. By contrast, this effect cannot be observed for unemployment and T-bill, where tilting toward model nowcasts performs better in both periods.

		Pre Crisis (88Q3 – 07Q4)							Complete (88Q3 – 13Q2)							
Horizon		1	2	3	4	5	4*	5*	1	2	3	4	5	4*	5*	
GDP																
RMSE	raw	1.973	2.040	2.015	1.992	1.969	1.441	1.452	2.396	2.593	2.663	2.636	2.612	2.086	2.140	
	SPF small m	1.580*	1.986	1.960*	2.059	2.012	1.239*	1.457	1.591**	2.454**	2.534**	2.672	2.633	1.712*	2.080	
	SPF small m/v	1.580*	1.963	1.959*	2.031*	2.006	1.229*	1.440	1.591**	2.385*	2.526**	2.642	2.622	1.692*	2.054	
	SPF big m	1.580*	1.963	1.979	2.042	2.051	1.234	1.441	1.591**	2.403*	2.563*	2.661	2.685	1.710*	2.083	
	SPF big m/v	1.578*	1.990	1.992	2.007	1.975	1.216	1.417	1.591**	2.388*	2.525	2.619	2.712	1.671*	2.039	
	BMF small m	1.682	1.991	1.969*	2.048	2.010	1.258	1.454	1.899**	2.492*	2.575**	2.666	2.631	1.841*	2.101	
	BMF small m/v	1.682	1.963	1.968	2.024	2.002	1.250	1.441	1.899**	2.469**	2.572*	2.643	2.620	1.832*	2.086	
	BMF big m	1.681	1.991	2.015	2.124*	2.068	1.295	1.507	1.903**	2.505	2.764	2.767*	2.709	1.944	2.194	
	BMF big m/v	1.686	2.007	2.098	2.094	1.990	1.290	1.491	1.895**	2.537	2.716	2.704	2.684	1.884*	2.154	
CRPS	raw	1.123	1.158	1.147	1.136	1.128	0.826	0.832	1.277	1.382	1.420	1.408	1.399	1.123	1.158	
	SPF small m	0.924**	1.131	1.122*	1.162*	1.143	0.722	0.831	0.966**	1.315**	1.358**	1.423	1.409	0.928**	1.124	
	SPF small m/v	0.901**	1.119	1.115*	1.156*	1.145	0.711	0.825	0.907**	1.282**	1.344**	1.414	1.408	0.915**	1.114	
	SPF big m	0.959*	1.141	1.125	1.160	1.139	0.730	0.824	0.995**	1.308*	1.364*	1.426	1.421	0.937**	1.124	
	SPF big m/v	0.899**	1.146	1.127	1.154	1.133	0.707	0.815	0.907**	1.304*	1.344*	1.409	1.453	0.907**	1.112	
	BMF small m	0.972**	1.132	1.128*	1.154	1.143	0.733	0.829	1.064**	1.335**	1.379**	1.419	1.408	0.989**	1.136	
	BMF small m/v	0.958*	1.120	1.123*	1.150	1.143	0.723	0.824	1.044**	1.321**	1.372**	1.413	1.406	0.979**	1.130	
	BMF big m	1.020	1.150	1.149	1.203**	1.176	0.758	0.856	1.109**	1.349	1.452	1.502*	1.469	1.030	1.182	
	BMF big m/v	0.961*	1.159	1.192	1.195	1.155	0.748	0.853	1.040**	1.353	1.465	1.466*	1.448	1.013*	1.170	
UNE																
RMSE	raw	0.157	0.274	0.397	0.511	0.605			0.235	0.463	0.706	0.940	1.147			
	SPF small m	0.125*	0.238	0.350	0.466	0.566			0.151*	0.352	0.577	0.817	1.038			
	SPF small m/v	0.125*	0.239	0.352	0.468	0.567			0.151*	0.350	0.573	0.810	1.028			
	SPF big m	0.125*	0.227*	0.329*	0.443	0.550			0.151*	0.329	0.550	0.780	1.003			
	SPF big m/v	0.122*	0.226*	0.330*	0.441	0.543			0.148*	0.328	0.535	0.760	0.987			
	BMF small m	0.095**	0.192*	0.296*	0.407*	0.511*			0.095**	0.255	0.468	0.720	0.969			
	BMF small m/v	0.096**	0.191*	0.295*	0.406*	0.512*			0.097**	0.264	0.499	0.754	1.007			
	BMF big m	0.095**	0.177*	0.269**	0.385*	0.503*			0.094**	0.252	0.497	0.760	1.009			
	BMF big m/v	0.095**	0.179*	0.277*	0.390*	0.498*			0.093**	0.272	0.507	0.757	0.994			
CRPS	raw	0.089	0.152	0.222	0.290	0.351			0.117	0.217	0.333	0.453	0.569			
	SPF small m	0.074**	0.137*	0.200	0.268	0.328			0.087**	0.178	0.280*	0.396	0.512			
	SPF small m/v	0.070**	0.133*	0.197	0.267	0.329			0.083**	0.176*	0.276*	0.392	0.507			
	SPF big m	0.075**	0.134*	0.192*	0.259	0.322			0.088**	0.171*	0.268*	0.378*	0.494			
	SPF big m/v	0.069**	0.126**	0.183*	0.250*	0.314			0.081**	0.166*	0.256*	0.367*	0.486*			
	BMF small m	0.064**	0.122*	0.179*	0.241*	0.298*			0.072**	0.147*	0.241	0.359*	0.479*			
	BMF small m/v	0.054**	0.110*	0.170*	0.233*	0.294*			0.054**	0.137*	0.241*	0.358*	0.482*			
	BMF big m	0.066**	0.117*	0.169**	0.232*	0.299*			0.073**	0.145*	0.247*	0.364*	0.489*			
	BMF big m/v	0.055**	0.106**	0.160**	0.224**	0.288*			0.054**	0.140**	0.237**	0.354**	0.477*			

Table 3: **Empirical results for entropic tilting.** “RMSE” rows contain root mean squared errors. “CRPS” rows contain mean cumulative ranked probability scores. raw – MCMC output of BVAR-SV model. Alternative tilting targets: SPF small m – SPF mean nowcast for the same variable. SPF small m/v – SPF nowcast mean and variance for the same variable. SPF big m – SPF nowcast means for all four variables. SPF big m/v – SPF nowcast means and variances for all four variables. BMF small m, BMF small m/v, BMF big m and BMF big m/v are defined analogously. One and two stars indicate rejections of equal predictive ability at the five and one percent level (two sided tests; implementation details described in the beginning of Section 5).

		<i>Pre Crisis (88Q3 – 07Q4)</i>							<i>Complete (88Q3 – 13Q2)</i>						
<i>Horizon</i>		1	2	3	4	5	4*	5*	1	2	3	4	5	4*	5*
		INF													
RMSE	raw	0.877	0.997	0.972	1.008	1.139	0.698	0.785	0.939	0.997	1.016	1.053	1.176	0.721	0.800
	SPF small m	0.767*	0.920**	0.943	0.997	1.115	0.634	0.746	0.809**	0.915**	0.977	1.031	1.146	0.636*	0.747
	SPF small m/v	0.767*	0.920**	0.944	0.997	1.117	0.634	0.748	0.809**	0.917**	0.977	1.031	1.149	0.637*	0.750
	SPF big m	0.767*	0.914**	0.929	0.986	1.111	0.617	0.735	0.809**	0.908**	0.949*	1.009	1.143	0.614**	0.730
	SPF big m/v	0.771*	0.911*	0.935	1.009	1.114	0.633	0.746	0.807**	0.903**	0.966	1.035	1.150	0.634*	0.746
	BMF small m	0.861	1.023	1.004	1.021	1.148	0.728	0.809	0.985	1.024	1.053	1.075	1.198	0.766	0.831
	BMF small m/v	0.861	1.029	1.004	1.021	1.149	0.729	0.810	0.985	1.030	1.055	1.076	1.200	0.769	0.833
	BMF big m	0.861	0.994	0.998	1.052	1.137	0.722	0.809	0.982	1.004	1.058	1.096	1.196	0.762	0.834
	BMF big m/v	0.869	1.020	1.028*	1.050	1.124	0.736	0.815	0.989	1.017	1.078*	1.094*	1.192	0.766	0.833
CRPS	raw	0.506	0.571	0.548	0.585	0.663	0.401	0.454	0.541	0.579	0.577	0.610	0.688	0.415	0.465
	SPF small m	0.443*	0.534**	0.541	0.579	0.653	0.374	0.440	0.470**	0.541**	0.565	0.602	0.676	0.379*	0.445
	SPF small m/v	0.436*	0.529**	0.538	0.576	0.651	0.367	0.436	0.464**	0.535**	0.560	0.597	0.673	0.370*	0.439
	SPF big m	0.451*	0.536*	0.539	0.577	0.654	0.371	0.436	0.480**	0.543**	0.559	0.598	0.677	0.378*	0.442
	SPF big m/v	0.441*	0.531*	0.536	0.581	0.656	0.366	0.434	0.465**	0.538**	0.554	0.597	0.681	0.368*	0.436
	BMF small m	0.491	0.581	0.560	0.589	0.667	0.411	0.465	0.546	0.591	0.591	0.619	0.698	0.430	0.478
	BMF small m/v	0.495	0.584	0.561	0.588	0.667	0.413	0.465	0.553	0.593	0.592	0.618	0.698	0.431	0.477
	BMF big m	0.496	0.567	0.565	0.603	0.665	0.414	0.466	0.549	0.583	0.597*	0.638	0.705	0.438	0.483
	BMF big m/v	0.499	0.590	0.580*	0.602	0.669	0.417	0.468	0.555	0.596	0.607*	0.631	0.702	0.432	0.480
		TBI													
RMSE	raw	0.406	0.751	1.029	1.271	1.482			0.406	0.757	1.060	1.339	1.594		
	SPF small m	0.133**	0.534**	0.848**	1.114**	1.362**			0.133**	0.536**	0.875**	1.181**	1.474**		
	SPF small m/v	0.132**	0.503**	0.823**	1.099**	1.345**			0.132**	0.512**	0.855**	1.164**	1.455**		
	SPF big m	0.133**	0.503**	0.819*	1.102*	1.345*			0.133**	0.521**	0.851**	1.156**	1.433**		
	SPF big m/v	0.134**	0.507**	0.826**	1.098**	1.338*			0.134**	0.516**	0.848**	1.148**	1.419**		
	BMF small m	0.066**	0.490**	0.822**	1.094**	1.346**			0.072**	0.480**	0.820**	1.132**	1.428**		
	BMF small m/v	0.067**	0.453**	0.794**	1.079**	1.327**			0.073**	0.453**	0.802**	1.113**	1.406**		
	BMF big m	0.066**	0.455**	0.772**	1.052**	1.307**			0.073**	0.485**	0.803**	1.122**	1.420**		
	BMF big m/v	0.069**	0.466**	0.802**	1.072**	1.300**			0.080**	0.472**	0.818**	1.124**	1.395**		
CRPS	raw	0.214	0.414	0.586	0.743	0.887			0.214	0.420	0.612	0.796	0.973		
	SPF small m	0.149**	0.326**	0.498**	0.658**	0.813**			0.154**	0.337**	0.526**	0.712**	0.898**		
	SPF small m/v	0.068**	0.267**	0.459**	0.636**	0.797**			0.070**	0.271**	0.478**	0.681**	0.875**		
	SPF big m	0.154**	0.324**	0.495**	0.655**	0.802**			0.158**	0.336**	0.519**	0.697**	0.871**		
	SPF big m/v	0.071**	0.270**	0.464**	0.635**	0.790**			0.072**	0.273**	0.478**	0.670**	0.851**		
	BMF small m	0.141**	0.311**	0.486**	0.648**	0.805**			0.148**	0.321**	0.506**	0.690**	0.875**		
	BMF small m/v	0.040**	0.238**	0.436**	0.617**	0.779**			0.041**	0.238**	0.444**	0.645**	0.840**		
	BMF big m	0.146**	0.309**	0.473**	0.628**	0.784**			0.158**	0.327**	0.502**	0.678**	0.864**		
	BMF big m/v	0.047**	0.247**	0.446**	0.616**	0.767**			0.048**	0.246**	0.460**	0.652**	0.832**		

Table 3: continued.

5.3 Entropic tilting and nowcast uncertainty

Table 3 implies that tilting towards the nowcast mean *and variance* consistently yields better CRPS scores than tilting towards the mean only. For $h = 1$, this effect is simply a consequence of the nowcast distributions being more accurate than the BVAR ones, which is well known in the literature. Much more interestingly, the result also holds for $h \geq 2$, which suggests that the “m/v” specification produces more favorable spillover effects on the horizons that are not directly affected by tilting. Table 4 investigates this result in more detail, by reporting the length and coverage of central prediction intervals obtained from both approaches (nominal level of 70%). In particular, we define length as the spread between the 15th and 85th percentiles of the forecast distribution and report the average length over time, and we measure coverage as the percent of actual outcomes of each variable falling within the 70% confidence band.

For all variables and forecast horizons, we observe that the “m/v” specifications produce shorter prediction intervals than the “m” specifications, which implies sharper (i.e., more concentrated) forecast distributions. This result is natural: The SPF and model nowcasts generally have lower variance than the current quarter forecasts produced by the BVAR (see Section 5.1). While the “m/v” tilting variant imposes this information, the “m” variant fails to do so. Instead, it penalizes the (KLIC) distance from the BVAR distribution, and thus implicitly targets the BVAR variance.¹⁹ These effects are clearest for the T-Bill and unemployment rates, where the “m/v” approaches produce prediction intervals whose average lengths (over time) are roughly 20-40 percent shorter than those of the “m” approaches. The differences are much smaller for GDP growth and inflation, where the average lengths of the prediction intervals typically differ by less than five percent.

Naturally, the reduced length of the “m/v” prediction intervals comes along with reduced coverage rates compared to the “m” variants. For GDP, unemployment and inflation, the coverage rates of “m/v” are mostly still above 60 percent (recall that the nominal level is 70 percent). A similar statement holds for the T-Bill rate and $h \in \{2, 3\}$. For the T-Bill rate and $h \in \{4, 5\}$, the coverage rates of the “m” approaches are already well below 70 percent, with the rates of “m/v” being even lower.

On balance, the increased sharpness of “m/v” appears to come at a small cost, in that the coverage rates are similarly close to (or far from) their nominal level as under the “m” approach. This assessment is consistent with the fact that the CRPS – which can be seen as a trade-off between sharpness and correct coverage, see e.g. Gneiting, Balabdaoui, and Raftery (2007) – consistently favors the “m/v” over the “m” approach.

¹⁹This effect can be seen most clearly in the Gaussian example of Section 4.4, where the “m” approach corresponds to targeting the original (BVAR) variance.

Horizon		Pre Crisis (88Q3 – 07Q4)							Complete (88Q3 – 13Q2)						
		1	2	3	4	5	4*	5*	1	2	3	4	5	4*	5*
		GDP													
Coverage	SPF m	0.846	0.705	0.718	0.718	0.718	0.667	0.641	0.870	0.700	0.700	0.700	0.660	0.640	0.560
	SPF m/v	0.731	0.679	0.692	0.718	0.718	0.628	0.641	0.710	0.680	0.680	0.700	0.660	0.590	0.560
	BMF m	0.795	0.692	0.718	0.718	0.718	0.692	0.641	0.790	0.680	0.700	0.700	0.660	0.640	0.560
	BMF m/v	0.705	0.692	0.692	0.705	0.705	0.667	0.641	0.710	0.680	0.680	0.690	0.650	0.610	0.560
Length	SPF m	4.756	4.460	4.495	4.561	4.563	2.671	2.652	5.242	4.877	4.914	4.971	4.969	2.978	2.952
	SPF m/v	3.377	4.304	4.385	4.442	4.486	2.432	2.601	3.332	4.643	4.732	4.785	4.813	2.640	2.857
	BMF m	4.577	4.437	4.458	4.528	4.538	2.630	2.634	4.927	4.812	4.844	4.899	4.907	2.913	2.919
	BMF m/v	3.798	4.323	4.377	4.450	4.485	2.488	2.600	3.979	4.667	4.736	4.798	4.815	2.726	2.864
CRPS	SPF m	0.924	1.131	1.122	1.162	1.143	0.722	0.831	0.966	1.315	1.358	1.423	1.409	0.928	1.124
	SPF m/v	0.901	1.119	1.115	1.156	1.145	0.711	0.825	0.907	1.282	1.344	1.414	1.408	0.915	1.114
	BMF m	0.972	1.132	1.128	1.154	1.143	0.733	0.829	1.064	1.335	1.379	1.419	1.408	0.989	1.136
	BMF m/v	0.958	1.120	1.123	1.150	1.143	0.723	0.824	1.044	1.321	1.372	1.413	1.406	0.979	1.130
		UNE													
Coverage	SPF m	0.859	0.833	0.756	0.731	0.667			0.860	0.820	0.750	0.710	0.650		
	SPF m/v	0.833	0.744	0.744	0.667	0.615			0.750	0.710	0.730	0.660	0.600		
	BMF m	0.987	0.923	0.872	0.808	0.744			0.990	0.900	0.830	0.760	0.670		
	BMF m/v	0.833	0.808	0.731	0.705	0.641			0.830	0.750	0.720	0.680	0.590		
Length	SPF m	0.389	0.623	0.818	0.977	1.090			0.451	0.715	0.932	1.114	1.238		
	SPF m/v	0.260	0.499	0.708	0.881	1.016			0.273	0.545	0.783	0.987	1.142		
	BMF m	0.441	0.692	0.887	1.039	1.148			0.505	0.786	1.004	1.179	1.312		
	BMF m/v	0.237	0.488	0.703	0.885	1.023			0.242	0.540	0.796	1.020	1.190		
CRPS	SPF m	0.074	0.137	0.200	0.268	0.328			0.087	0.178	0.280	0.396	0.512		
	SPF m/v	0.070	0.133	0.197	0.267	0.329			0.083	0.176	0.276	0.392	0.507		
	BMF m	0.064	0.122	0.179	0.241	0.298			0.072	0.147	0.241	0.359	0.479		
	BMF m/v	0.054	0.110	0.170	0.233	0.294			0.054	0.137	0.241	0.358	0.482		
		INF													
Coverage	SPF m	0.769	0.718	0.744	0.705	0.718	0.782	0.756	0.780	0.740	0.760	0.720	0.710	0.820	0.790
	SPF m/v	0.718	0.679	0.731	0.692	0.705	0.795	0.744	0.700	0.710	0.750	0.710	0.690	0.830	0.780
	BMF m	0.679	0.744	0.718	0.731	0.731	0.731	0.756	0.660	0.730	0.710	0.740	0.730	0.760	0.770
	BMF m/v	0.628	0.705	0.705	0.718	0.692	0.756	0.744	0.620	0.700	0.710	0.730	0.700	0.770	0.760
Length	SPF m	1.885	2.078	2.242	2.403	2.581	1.681	1.899	2.029	2.239	2.421	2.594	2.782	1.822	2.054
	SPF m/v	1.603	1.993	2.170	2.340	2.510	1.566	1.824	1.670	2.126	2.322	2.502	2.682	1.664	1.951
	BMF m	1.870	2.065	2.239	2.400	2.579	1.680	1.895	2.006	2.223	2.407	2.582	2.772	1.813	2.041
	BMF m/v	1.742	2.030	2.208	2.376	2.555	1.635	1.869	1.830	2.168	2.358	2.539	2.729	1.739	1.993
CRPS	SPF m	0.443	0.534	0.541	0.579	0.653	0.374	0.440	0.470	0.541	0.565	0.602	0.676	0.379	0.445
	SPF m/v	0.436	0.529	0.538	0.576	0.651	0.367	0.436	0.464	0.535	0.560	0.597	0.673	0.370	0.439
	BMF m	0.491	0.581	0.560	0.589	0.667	0.411	0.465	0.546	0.591	0.591	0.619	0.698	0.430	0.478
	BMF m/v	0.495	0.584	0.561	0.588	0.667	0.413	0.465	0.553	0.593	0.592	0.618	0.698	0.431	0.477
		TBI													
Coverage	SPF m	0.962	0.795	0.628	0.564	0.474			0.960	0.810	0.670	0.600	0.490		
	SPF m/v	0.795	0.692	0.551	0.449	0.423			0.810	0.730	0.610	0.490	0.440		
	BMF m	1.000	0.833	0.667	0.590	0.526			1.000	0.850	0.700	0.620	0.540		
	BMF m/v	0.949	0.718	0.603	0.462	0.462			0.940	0.760	0.640	0.500	0.460		
Length	SPF m	0.956	1.300	1.555	1.797	1.994			0.996	1.411	1.721	2.010	2.250		
	SPF m/v	0.263	0.810	1.175	1.457	1.692			0.299	0.887	1.298	1.627	1.897		
	BMF m	1.003	1.323	1.562	1.808	2.007			1.037	1.441	1.742	2.033	2.272		
	BMF m/v	0.238	0.806	1.171	1.462	1.693			0.241	0.859	1.280	1.617	1.889		
CRPS	SPF m	0.149	0.326	0.498	0.658	0.813			0.154	0.337	0.526	0.712	0.898		
	SPF m/v	0.068	0.267	0.459	0.636	0.797			0.070	0.271	0.478	0.681	0.875		
	BMF m	0.141	0.311	0.486	0.648	0.805			0.148	0.321	0.506	0.690	0.875		
	BMF m/v	0.040	0.238	0.436	0.617	0.779			0.041	0.238	0.444	0.645	0.840		

Table 4: **Impact of accounting for nowcast uncertainty.** “Coverage” and “Length” refer to central prediction intervals with a nominal level of 70 % (reported length is on average over time). “CRPS” reports the continuous ranked probability score (best = lowest number printed in bold). Note that the CRPS scores are identical to those in Table 3, and are reprinted here for ease of reference.

6 Conclusion

This paper is concerned with the problem of combining forecasts from a BVAR with nowcasts from other sources. This combination problem is non-standard, in that the BVAR implies a joint forecast distribution for several forecast horizons, whereas the nowcast information is restricted to mean and variance predictions for the current quarter. We argue that entropic tilting is a powerful tool to tackle these challenges; unlike other methods proposed in the literature, it does not require restrictive assumptions such as joint normality of the VAR system or zero variance of the nowcast.

In our empirical analysis, tilting systematically improves the accuracy of both point and density forecasts, and tilting the BVAR forecasts based on nowcast means and variances yields slightly greater gains in density accuracy than does just tilting based on the nowcast means. In a comparison of tilting on a variable-by-variable basis to tilting jointly toward the nowcasts for all four variables of the BVAR, we find that the overall differences in forecast performance for the joint treatment of variables versus the individual treatment of variables are small.

References

- ALTAVILLA, C., R. GIACOMINI, AND G. RAGUSA (2013): "Anchoring the Yield Curve Using Survey Expectations," Working Paper, UC London.
- ANG, A., G. BEKAERT, AND M. WEI (2007): "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?," *Journal of Monetary Economics*, 54, 1163 – 1212.
- BANBURA, M., D. GIANNONE, M. MODUGNO, AND L. REICHLIN (2013): "Now-Casting and the Real-Time Data Flow," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2, pp. 195 – 237. Elsevier.
- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2013): "Nowcasting," in *Oxford Handbook of Economic Forecasting*, ed. by M. P. Clements, and D. F. Hendry, pp. 193 – 224. Oxford University Press.
- BOMBERGER, W. A. (1996): "Disagreement as a Measure of Uncertainty," *Journal of Money, Credit and Banking*, 28, 381–392.
- CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2014): "Real-Time Nowcasting with a Bayesian Mixed Frequency Model with Stochastic Volatility," *Journal of the Royal Statistical Society: Series A*, forthcoming.
- CARTER, C. K., AND R. KOHN (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541–553.
- CLARK, T. E. (2011): "Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility," *Journal of Business & Economic Statistics*, 29, 327–341.
- CLARK, T. E., AND M. W. MCCRACKEN (2013): "Advances in Forecast Evaluation," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2, pp. 1107–1201. Elsevier.
- (2014): "Evaluating Conditional Forecasts from Vector Autoregressions," Federal Reserve Bank of Cleveland Working Paper 14-13.
- CLARK, T. E., AND F. RAVAZZOLO (2014): "The Macroeconomic Forecasting Performance of Autoregressive Models with Alternative Specifications of Time-Varying Volatility," *Journal of Applied Econometrics*, forthcoming.
- 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.
- CROUSHORE, D. (2006): "Forecasting with Real-Time Macroeconomic Data," *Handbook of Economic Forecasting*, 1, 961–982.
- D'AGOSTINO, A., L. GAMBETTI, AND D. GIANNONE (2013): "Macroeconomic Forecasting and Structural Change," *Journal of Applied Econometrics*, 28, 82–101.
- DEL NEGRO, M., AND G. E. PRIMICERI (2014): "Time-Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum," <http://faculty.wcas.northwestern.edu/~gep575/ErrataFinal2.pdf>, Working Paper, Northwestern University, Accessed: 2014-02-13.

- DEL NEGRO, M., AND F. SCHORFHEIDE (2013): "DSGE Model-Based Forecasting," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2, pp. 57–140. Elsevier.
- DIEBOLD, F. X., AND R. S. MARIANO (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 13, 253–263.
- DOAN, T., R. LITTERMAN, AND C. SIMS (1984): "Forecasting and Conditional Projection using Realistic Prior Distributions," *Econometric Reviews*, 3, 1–100.
- DURBIN, J., AND S. J. KOOPMAN (2002): "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis," *Biometrika*, 89, 603–616.
- FAUST, J., AND J. H. WRIGHT (2009): "Comparing Greenbook and Reduced Form Forecasts Using a Large Realtime Dataset," *Journal of Business & Economic Statistics*, 27, 468–479.
- (2013): "Forecasting Inflation," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2, pp. 2–56. Elsevier.
- FREY, C., AND F. MOKINSKI (2014): "Forecasting with Bayesian Vector Autoregressions Estimated Using Professional Forecasts," Working Paper, University of Konstanz.
- GEWEKE, J., AND G. AMISANO (2011): "Optimal Prediction Pools," *Journal of Econometrics*, 164, 130–141.
- GIACOMINI, R., AND H. WHITE (2006): "Tests of Conditional Predictive Ability," *Econometrica*, 74, 1545–1578.
- GIANNONE, D., F. MONTI, AND L. REICHLIN (2014): "Exploiting the Monthly Data-flow in Structural Forecasting," Centre for Macroeconomics (CFM) Working Paper 1416.
- GNEITING, T., F. BALABDAOUI, AND A. E. RAFTERY (2007): "Probabilistic Forecasts, Calibration and Sharpness," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 243–268.
- GNEITING, T., AND A. E. RAFTERY (2007): "Strictly Proper Scoring Rules, Prediction, and Estimation," *Journal of the American Statistical Association*, 102, 359–378.
- GNEITING, T., AND R. RANJAN (2011): "Comparing Density Forecasts using Threshold-and Quantile-weighted Scoring Rules," *Journal of Business & Economic Statistics*, 29.
- HALL, S. G., AND J. MITCHELL (2007): "Combining Density Forecasts," *International Journal of Forecasting*, 23, 1–13.
- HARVEY, D., S. LEYBOURNE, AND P. NEWBOLD (1997): "Testing the Equality of Prediction Mean Squared Errors," *International Journal of Forecasting*, 13, 281–291.
- HERSBACH, H. (2000): "Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems," *Weather and Forecasting*, 15, 559–570.
- KIM, S., N. SHEPHARD, AND S. CHIB (1998): "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *The Review of Economic Studies*, 65, 361–393.
- LEWIS, K. F., AND C. H. WHITEMAN (2014): "Empirical Bayesian Density Forecasting in Iowa and Shrinkage for the Monte Carlo Era," *Journal of Forecasting*, forthcoming.

- LITTERMAN, R. B. (1986): "Forecasting with Bayesian Vector Autoregressions – Five Years of Experience," *Journal of Business & Economic Statistics*, 4, 25–38.
- MONTGOMERY, A. L., V. ZARNOWITZ, R. S. TSAY, AND G. C. TIAO (1998): "Forecasting the US Unemployment Rate," *Journal of the American Statistical Association*, 93, 478–493.
- MONTI, F. (2010): "Combining Judgment and Models," *Journal of Money, Credit and Banking*, 42, 1641–1662.
- NEWKEY, W. K., AND K. D. WEST (1987): "A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.
- NEWKEY, W. K., AND K. D. WEST (1994): "Automatic Lag Selection in Covariance Matrix Estimation," *The Review of Economic Studies*, 61, 631–653.
- OMORI, Y., S. CHIB, N. SHEPHARD, AND J. NAKAJIMA (2007): "Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference," *Journal of Econometrics*, 140, 425–449.
- PRIMICERI, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 72, 821–852.
- R CORE TEAM (2014): *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- ROBERTSON, J. C., E. W. TALLMAN, AND C. H. WHITEMAN (2005): "Forecasting Using Relative Entropy," *Journal of Money, Credit and Banking*, 37, 383–401.
- ROMER, C. D., AND D. H. ROMER (2000): "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 90, 429–457.
- SCHORFHEIDE, F., AND D. SONG (2013): "Real-Time Forecasting with a Mixed-Frequency VAR," *Journal of Business & Economic Statistics*, forthcoming.
- SIMS, C. A. (2002): "The Role of Models and Probabilities in the Monetary Policy Process," *Brookings Papers on Economic Activity*, 2002, 1–40.
- SMETS, F., A. WARNE, AND R. WOUTERS (2014): "Professional Forecasters and Real-Time Forecasting with a DSGE Model," *International Journal of Forecasting*, 30, 981–995.
- WEST, K. D. (1996): "Asymptotic Inference about Predictive Ability," *Econometrica*, 64, 1067–1084.
- WOLTERS, M. H. (2014): "Evaluating Point and Density Forecasts of DSGE Models," *Journal of Applied Econometrics*, forthcoming.
- WRIGHT, J. H. (2013): "Evaluating Real-Time Forecasts with an Informative Democratic Prior," *Journal of Applied Econometrics*, 28, 762–776.
- ZEILEIS, A. (2004): "Econometric Computing with HC and HAC Covariance Matrix Estimators," *Journal of Statistical Software*, 11, 1–17.

Appendix

This appendix presents a simple demonstration of the equivalence between (1) Gaussian conditional forecasting with no uncertainty around the nowcast and (2) the Faust and Wright (2009) approach, and details the priors and algorithms used in estimation of the BVAR models and other nowcasting models.

A Example to establish equivalence of data augmentation to conditional forecasting

As mentioned in the text, the data augmentation technique of Faust and Wright (2009) simply appends the nowcast to the data. Here we provide an example to demonstrate that this method is equivalent to Gaussian conditional forecasting with no uncertainty around the nowcast (so $\Omega_{11} = 0$; see Section 4.4). To this end, suppose y_t follows an AR(1) process, and we use it to forecast $h \geq 2$ steps ahead, with a one-step nowcast of value μ_1 . The model is

$$y_t = \phi y_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma^2).$$

A.1 Unconditional forecast without nowcast information

The usual h step ahead forecast for y_{t+h} is

$$\hat{y}_{t+h}^u = \phi^h y_t,$$

where the superscript u stands for “unconditional” (i.e., no nowcast information). The corresponding forecast error is

$$\hat{e}_{t+h}^u = \sum_{j=0}^{h-1} \phi^j \epsilon_{t+h-j},$$

and has variance

$$\Sigma_{hh} = \sigma^2 \sum_{j=0}^{h-1} \phi^{2j}.$$

Furthermore, the predicted covariance between the 1-step and h -step ahead forecast errors is given by

$$\Sigma_{1h} = \sigma^2 \phi^{h-1}.$$

A.2 Data augmentation approach to incorporating nowcast

Under the data augmentation approach, the nowcast of μ_1 is treated as data for period $t + 1$, and we form the forecast for period $t + h$ as an $(h - 1)$ -step ahead forecast using the pseudo-data for $t + 1$. In this case, using a superscript a to denote the data augmentation approach, we obtain a forecast for y_{t+h} of:

$$\hat{y}_{t+h}^a = \phi^{h-1} \mu_1.$$

With the nowcast treated as data for period $t + 1$, the forecast error variance for period $t + h$ is defined as just the $(h - 1)$ -step ahead error variance, which is given by

$$\sigma^2 \sum_{j=0}^{h-2} \phi^{2j}.$$

A.3 Conditional forecast approach to incorporating nowcast

Under the Gaussian conditional forecasting approach, the forecast for period $t + h$ is formed under the condition that y_{t+1} take the nowcast value of μ_1 . Using the more general conditional forecast solution given in Section 4.4, as well as the unconditional forecast quantities derived above, the point forecast is given by

$$\begin{aligned} \hat{y}_{t+h}^c &= y_{t+h}^u + \frac{\Sigma_{1h}}{\Sigma_{11}} (\mu_1 - y_{t+1}^u) \\ &= \phi^h y_t + \frac{\phi^{h-1} \sigma^2}{\sigma^2} (\mu_1 - \phi y_t) \\ &= \phi^{h-1} \mu_1. \end{aligned}$$

Given the nowcast condition on y_{t+1} (without uncertainty around it, so $\Omega_{11} = 0$), the general solution in Section 4.4 yields a conditional forecast error variance as follows:

$$\begin{aligned} \Omega_{hh} &= \Sigma_{hh} - \frac{\Sigma_{1h}^2}{\Sigma_{11}} \\ &= (1 + \phi^2 + \dots + \phi^{2h-2}) \sigma^2 - \frac{\phi^{2h-2} \sigma^4}{\sigma^2} \\ &= \sigma^2 \sum_{j=0}^{h-2} \phi^{2j}. \end{aligned}$$

Accordingly, in this simple example, it follows that the data augmentation and conditional forecast approaches to combining a nowcast with another forecast yield the same forecasts, as long as the conditional forecast is implemented assuming a zero variance around the nowcast.

B Prior for BVARs with stochastic volatility

We describe in this section the priors used with the BVAR-SV models to produce quarterly forecasts of the variables of interest and nowcasts of unemployment and the T-bill rate.

For the VAR coefficients, we use a conventional Minnesota prior, without cross-variable shrinkage (note that i and j refer to the row and column of B_l):

$$\underline{\mu}_B \text{ such that } E[B_l^{(ij)}] = 0 \quad \forall i, j, l \quad (14)$$

$$\underline{\Omega}_B \text{ such that } V[B_l^{(ij)}] = \begin{cases} \frac{\theta^2}{l^2} \frac{\sigma_i^2}{\sigma_j^2} & \text{for } l > 0 \\ \varepsilon^2 \sigma_i^2 & \text{for } l = 0 \end{cases} \quad (15)$$

Following common settings, we set $\theta = 0.2$, $\varepsilon = 1000$, and the scale parameters σ_i^2 at estimates of residual variances from $AR(p)$ models from the estimation sample. With all of the variables of our VAR models defined so that they should be stationary, we set the prior mean of all the VAR coefficients to 0.

In the prior for the volatility-related components of the model, we use an approach to setting them similar to that of such studies as Cogley and Sargent (2005), Primiceri (2005) and Clark (2011). The prior for A is uninformative, with a mean and variance for each row vector of $\underline{\mu}_{a,i} = 0$, $\underline{\Omega}_{a,i} = 1000^2 \cdot I_{i-1}$, $i = 2, \dots, k$. We make the priors on the volatility-related parameters loosely informative. The prior for Φ is inverted Wishart, with mean of $0.01 \times I_k$ and $k + 1$ degrees of freedom. For the initial value of the log volatility of each equation i , we use a mean of $\log \hat{\lambda}_{i,0,OLS}$ and variance of 4. To obtain $\log \hat{\lambda}_{i,0,OLS}$, we use the residuals from $AR(p)$ models estimated over a training sample preceding the estimation sample. For each $j = 2, \dots, k$, we regress the residual from the AR model for j on the residuals associated with variables 1 through $j - 1$ and compute the error variance $\hat{\sigma}_{i,0}^2$. We set the prior mean of log volatility in period 0 at $\log \hat{\lambda}_{i,0,OLS} = \log \hat{\sigma}_{i,0}^2$. For the quarterly model and the unemployment nowcasting model, the training sample is 1949-54; for the T-bill nowcasting model, the training sample is 1959-63. For the quarterly model and the unemployment nowcasting model, because a handful of the data vintages do not start until later than most others, we use the same prior mean on initial volatility for all vintages (forecast origins), computed using the last available vintage of data.

C Prior for Bayesian mixed frequency models

We describe in this section the priors used for mixed frequency nowcasting models for GDP growth and inflation. Since the form of the prior is not dependent on the month timing m , in spelling out the prior we drop the index m from the model parameters for notational simplicity.

In all cases, for the coefficient vector β , we use a prior distribution that is normal, with mean 0 (for all coefficients) and variance that takes a diagonal, Minnesota-style form. The prior variance is Minnesota style in the sense that shrinkage increases with the lag (with the quarter, not with the month within the quarter), and in the sense that we impose more shrinkage on the monthly predictors than on lags of GDP growth or inflation. The shrinkage is controlled by three hyperparameters (in all cases, a smaller number means more shrinkage): λ_1 , which controls the overall rate of shrinkage; λ_2 , which controls the rate of shrinkage on variables relative to GDP or GDP inflation; and λ_3 , which determines the rate of shrinkage associated with longer lags.

At each forecast origin, the prior standard deviation associated with the coefficient on variable $x_{i,j,t-l}$ of X_t , where i denotes the indicator (employment, etc.), j denotes the month within which the quarter at which the indicator has been sampled, and l denotes the lag in quarters (while we only consider a lag of 1 in this paper, Carriero, Clark, and Marcellino (2014) include results for models with a lag of 2), is specified as follows:

$$\text{sd}_{i,j,t-l} = \frac{\sigma_y}{\sigma_{i,j}} \frac{\lambda_1 \lambda_2}{l^{\lambda_3}}. \quad (16)$$

For coefficients on lag l of y , the prior standard deviation is

$$\text{sd}_l = \frac{\lambda_1}{l^{\lambda_3}}. \quad (17)$$

Finally, for the intercept, the prior is uninformative:

$$\text{sd}_{int} = 1000\sigma_y. \quad (18)$$

In setting these components of the prior, for σ_y and $\sigma_{i,j}$ we use standard deviations from AR(4) models for GDP growth or inflation and $x_{i,j,t}$ estimated with the available sample of data.

In all of our results, the hyperparameters are set at values that may be considered very common in Minnesota-type priors (e.g. Litterman, 1986): $\lambda_1 = 0.2$, $\lambda_2 = 0.2$, and $\lambda_3 = 1$.

Finally, in the prior for the volatility-related components of the model, our approach is similar to that used in such studies as Cogley and Sargent (2005), Primiceri (2005) and Clark (2011). For

the prior on ϕ , we use a mean of 0.035 and 5 degrees of freedom. For the period 0 value of volatility of each equation i , we use a prior of

$$\underline{\mu}_\lambda = \log \hat{\lambda}_{0,OLS}, \quad \underline{\Omega}_\lambda = 4. \quad (19)$$

To obtain $\log \hat{\lambda}_{0,OLS}$, we use a training sample of 40 observations preceding the estimation sample to fit an AR(4) model to GDP growth or inflation.

D Algorithm for BVAR with stochastic volatility

We estimate the BVAR-SV model with a five-step Gibbs sampling algorithm. Let X_t denote the collection of right-hand side variables of each equation of the VAR and B denote the vector of the system of VAR coefficients contained in B_i , $i = 0, \dots, p$, as defined in the paper's equation (1). go here

Step 1: Draw the VAR coefficients B conditional on the history of Λ_t , A , and Φ .

The vector of coefficients is sampled from a conditional posterior distribution that is multivariate normal with mean $\bar{\mu}_B$ and variance $\bar{\Omega}_B$, based on prior mean $\underline{\mu}_B$ and variance $\underline{\Omega}_B$. Letting $\Sigma_t = A^{-1}\Lambda_t A^{-1'}$, the posterior mean and variance are:

$$\text{vec}(\bar{\mu}_B) = \bar{\Omega}_B \left\{ \text{vec} \left(\sum_{t=1}^T X_t y_t' \Sigma_t^{-1} \right) + \underline{\Omega}_B^{-1} \text{vec}(\underline{\mu}_B) \right\} \quad (20)$$

$$\bar{\Omega}_B^{-1} = \underline{\Omega}_B^{-1} + \sum_{t=1}^T (\Sigma_t^{-1} \otimes X_t X_t'). \quad (21)$$

Step 2: Draw the elements of A conditional on B , the history of Λ_t , and Φ .

Following Cogley and Sargent (2005), rewrite the VAR as

$$A(y_t - X_t' B) = A \hat{y}_t \equiv \tilde{y}_t = \Lambda_t^{0.5} \epsilon_t, \quad (22)$$

where, conditional on B , \hat{y}_t is observable. This system simplifies to a set of $i = 2, \dots, k$ equations, with equation i having as dependent variable $\hat{y}_{i,t}$ and as independent variables $-1 \cdot \hat{y}_{j,t}$, $j = 1, \dots, i-1$, with coefficients a_{ij} . Multiplying equation i by $\lambda_{i,t}^{-0.5}$ eliminates the heteroskedasticity associated with stochastic volatility. Then, proceeding separately for each transformed equation i , draw the i 'th equation's vector of coefficients a_i (a vector containing a_{ij} for $j = 1, \dots, i-1$) from

a normal posterior distribution with the mean and variance implied by the posterior mean and variance computed in the usual way. See Cogley and Sargent (2005) for details.

Step 3: Draw the elements of the states for the mixture distribution used to approximate the χ^2 distribution under the Kim, Shephard, and Chib (1998) algorithm, conditional on B , A , the history of Λ_t , and Φ .

See Primiceri (2005) for details. However, we depart from Primiceri by using a 10 state approximation of the χ^2 distribution from Omori, Chib, Shephard, and Nakajima (2007) instead of the 7-state approximation from Kim, Shephard, and Chib (1998).

Step 4: Draw the elements of the variance matrix Λ_t conditional on B , A , Φ , and the mixture states.

Following Primiceri (2005), the VAR can be rewritten as

$$A(y_t - X_t' B) \equiv \tilde{y}_t = \Lambda_t^{0.5} \epsilon_t,$$

where $\epsilon_t \sim N(0, I_k)$. Taking logs of the squares yields

$$\log \tilde{y}_{i,t}^2 = \log \lambda_{i,t}^2 + \log \epsilon_{i,t}^2, \quad i = 1, \dots, k.$$

The conditional volatility process is

$$\log(\lambda_{i,t}^2) = \log(\lambda_{i,t-1}^2) + \nu_{i,t}, \quad i = 1, \dots, k.$$

The estimation of the time series of $\lambda_{i,t}^2$ uses the vector of the measured $\log \tilde{y}_{i,t}^2$ and Primiceri's version of the Kim, Shephard, and Chib (1998) algorithm; see Primiceri for further detail (we depart from his implementation by using the Durbin and Koopman (2002) simulation smoother instead of the one proposed by Carter and Kohn (1994)).

Step 5: Draw the variance matrix Φ , conditional on B , the history of Λ_t , and A .

Following Primiceri (2005), the sampling of Φ , the variance of innovations to the log variances, is based on inverse Wishart priors and posteriors. The scale matrix of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{\nu}_t \hat{\nu}_t'$, where $\hat{\nu}_t$ denotes the vector of innovations to the posterior draw of the volatilities for the set of variables.

E Algorithm for Bayesian mixed frequency model

The mixed frequency nowcasting model with stochastic volatility is estimated with a Metropolis-within-Gibbs algorithm, used in such studies as Clark (2011) and Carriero, Clark, and Marcellino (2014). The posterior mean and variance of the coefficient vector are given by

$$\bar{\mu}_\beta = \bar{\Omega}_\beta \left\{ \sum_{t=1}^T \lambda_t^{-1} X_t y_t + \underline{\Omega}_\beta^{-1} \underline{\mu}_\beta \right\} \quad (23)$$

$$\bar{\Omega}_\beta^{-1} = \underline{\Omega}_\beta^{-1} + \sum_{t=1}^T \lambda_t^{-1} X_t X_t', \quad (24)$$

where we again omit the m index from the parameters for notational simplicity.