Asymmetric Forecast Densities for U.S. Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence

Michael S. Smith and Shaun P. Vahey

First Version: June 2013;

This Version: March 2015

Michael Smith is Chair of Management (Econometrics) at Melbourne Business School, University of Melbourne. Shaun Vahey is Professor at Warwick Business School, University of Warwick. Correspondence should be directed to Michael Smith at mike.smith@mbs.edu. The authors would like to thank the editorial team, participants at the 4th ESOBE and 8th CFE annual workshops and staff at the Narodowy Bank Polski for many helpful comments, as well as Todd Clark and Francesco Ravazzolo for providing their data and code, and also helpful suggestions. The authors are grateful to Tom Stark for helpful comments on the Survey of Professional Forecasters. This work was partially supported by Australian Research Council Future Fellowship FT110100729.

Asymmetric Forecast Densities for U.S. Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence

Abstract

Most existing reduced-form macroeconomic multivariate time series models employ elliptical disturbances, so that the forecast densities produced are symmetric. In this paper, we use a copula model with asymmetric margins to produce forecast densities with the scope for severe departures from symmetry. Empirical and skew t distributions are employed for the margins, and a high-dimensional Gaussian copula is used to jointly capture cross-sectional and (multivariate) serial dependence. The copula parameter matrix is given by the correlation matrix of a latent stationary and Markov vector autoregression (VAR). We show that the likelihood can be evaluated efficiently using the unique partial correlations, and estimate the copula using Bayesian methods. We examine the forecasting performance of the model for four U.S. macroeconomic variables between 1975:Q1 and 2011:Q2 using quarterly real-time data. We find that the point and density forecasts from the copula model are competitive with those from a Bayesian VAR. During the recent recession the forecast densities exhibit substantial asymmetry, avoiding some of the pitfalls of the symmetric forecast densities from the Bayesian VAR. We show that the asymmetries in the predictive distributions of GDP growth and inflation are similar to those found in the probabilistic forecasts from the Survey of Professional Forecasters. Last, we find that unlike the linear VAR model, our fitted Gaussian copula models exhibit nonlinear dependencies between some macroeconomic variables.

Key Words: Copula Multivariate Time Series Model; Gaussian Copula; Real-Time Density Forecasting; Survey of Professional Forecasters; Vine Copula.

1 Introduction

A communication issue arises with the forecast densities produced by conventional reduced-form macroeconomic models. As Cogley, Morozov and Sargent (2005) emphasize, vector autoregressions (VARs), with or without stochastic volatility, use elliptical disturbances, so that the forecast densities produced are either exactly or approximately symmetric. For policymakers and private sector forecasters alike, symmetric forecast densities periodically assign unpalatably high probability to extreme events such as deflation, nominal interest rates below the lower bound, and negative output growth. Some central banks, including the Bank of England and Norges Bank, explore the scope for asymmetric forecast densities as communication tools. Similarly, the predictive distributions for key U.S. macroeconomic variables published by the Survey of Professional Forecasters (SPF) periodically exhibit asymmetries, in particular during the recent recession. In this paper we propose using a copula model to construct asymmetric forecast densities for U.S. macroeconomic variables. In doing so, we aim to demonstrate the potential of the copula framework for modeling and density forecasting using multivariate time series in empirical macroeconomics.

Copula models allow for the construction of a multivariate distribution in a flexible fashion. The marginal distributions are modeled separately from the dependence structure, which is captured using a copula function; see McNeil et al. (2005) and Nelson (2006) for introductions. The end result is statistical models that can have richer dependence and distributional features than would otherwise be the case, yet can still be estimated using likelihood-based methods. Copulas have been used previously to model cross-sectional dependence in multivariate time series (Patton 2006; Rodriguez 2007; Patton 2012), and also to capture serial dependence in univariate series (Domma et al. 2009; Beare 2010; Smith et al. 2010). However, only limited attention has been given to the case we examine here, where a copula function is used to capture jointly both serial and cross-sectional dependence. This is because the copula required is of a high dimension, equal to the number of time points observed multiplied by the number of variables. To solve this problem Smith (2014) proposes employing a vine copula (Aas et al. 2009), which is parsimonious when the series is

Markov and stationary. Here, we employ a Gaussian copula (Song 2000) with a parsimonious parameter matrix for this problem.

While the Gaussian copula is a special case of the more general vine copula, a major advantage is that it can be represented using a latent Gaussian variable. For our multivariate time series model, we assume this latent variable arises from a stationary Markov VAR. The parameter matrix of the Gaussian copula is equal to the highly structured and parsimonious correlation matrix Ω of the latent VAR. Provided the marginal models are time invariant, the multivariate time series can be shown to also be (strongly) stationary and have the same Markov order as the latent process. In the special case when all variables have Gaussian marginal distributions, the copula model is simply a Gaussian VAR. In this case, the forecast densities for the individual macroeconomic variables are exactly the same as those from a VAR. However, the copula model provides a flexible extension whenever one or more of the margins are non-Gaussian— exactly the case that arises in our empirical analysis. To estimate the model we parameterize Ω in terms of its unique partial correlations. These can be used to compute the likelihood efficiently using the vine representation of the Gaussian copula. We also employ Bayesian ‘spike-and-slab’ priors for the partials, and undertake Bayesian selection and model averaging. Bayesian selection methods have been used in multivariate time series modeling of macroeconomic variables in a number of previous studies (George, Sun and Ni 2008; Korobilis 2010; Jochmann et al. 2013), and our study extends these.

Our real-time data consist of quarterly measurements on real GDP growth, inflation, the unemployment rate and the short-term interest rate. Clark and Ravazzolo (2014) document the real-time density forecasting performance of VARs for these macroeconomic variables; see also Jore, Mitchell and Vahey (2010) and Clark (2011) on the forecast accuracy of low-dimensional VARs with U.S. macroeconomic data. We consider quarterly vintages from 1975:Q1 to 2011:Q2, for which we show statistical tests reject normality for all four variables, except for a handful of early vintages of GDP growth and the interest rate. All four variables exhibit skew and/or heavy tails, so that we consider two alternative models for the margins: a skew t distribution (Azzalini and Capitanio 2003) and the empirical distribution. In each

case we compute the probability integral transformed (PIT) data, from which we fit the copula.

Even though dependence in the latent Gaussian process is strictly linear, the estimates of the serial and cross-sectional dependence structure of the fitted copula model exhibit nonlinearities. For example, the unemployment rate (which is a lagging indicator) is not only negatively related with output growth lagged 3 quarters, but with very different gradients during periods of contraction than during expansion. Using a quarterly expanding window, we construct forecast distributions of all four variables between 1975:Q1 and 2011:Q1, over a horizon of eight quarters. This is the same period considered by Clark and Ravazzolo (2014), and the forecasting component of our study provides an extension of their comprehensive comparison.

Overall, we find that our copula model provides improved forecast accuracy over a Bayesian VAR (BVAR). We show that the asymmetries of the fitted copula model can carry over into the out-of-sample predictive distributions, including at long horizons, and that these can be particularly useful for forecasting during periods of recession. For example, during the Great Recession, forecast distributions from BVAR models are symmetric and place high probabilities on substantial deflation and a negative interest rate, including a rate below -2 percent—events that are highly unlikely or impossible. In contrast, the predictive distributions for inflation and the interest rate from the copula model are sharper, positively skewed and put low or zero probability on substantial deflation or a negative interest rate. The improvement in forecasting performance over the BVAR during the recent recession offers evidence that the additional flexibility afforded by copula modeling matters for U.S. macroeconomic variables. Nevertheless, we find that the improvements in overall forecast accuracy achieved are modest relative to those obtained by augmenting the BVAR to account for stochastic volatility. The Gaussian copula does not allow for serial dependence in the second order moments. Finding a copula function that does so represents an important opportunity for further research— a point we return to in the discussion.

We also compare the predictive distributions from the copula model to the density forecasts produced for GDP growth and inflation by the SPF. These forecasts are created as

an ensemble from individual professional forecasters, and can be highly asymmetric and heavy-tailed. Croushore (1993) and Lahiri and Wang (2013) discuss the relative accuracy of SPF forecasts over various horizons. Hafer and Hein (1983), Faust and Wright (2013) and Kruger, Clark and Ravazzolo (2014) argue that SPF point forecasts perform well relative to existing time series models. We examine in detail the density forecasts for GDP growth over 2008-2009, and inflation over 2009-2010, over eight quarterly SPF reports. Using matching vintage data, we find the predictive distributions from the copula model exhibit a similar level of asymmetry and kurtosis to those produced in the quarterly reports, and differ substantially from those of the BVAR models. We find that there is a rapid shift in the predictive distributions of output growth from all the time series models with the onset of the Great Recession, although in the SPF reports the shift occurs one quarter earlier—probably reflecting the larger and more timely information set of the survey respondents. Nevertheless, we find that both during the Great Recession, and over a longer period of 18 years, at short horizons the predictive distributions from the SPF seem overly diffuse, and are less accurate than the time series models we consider.

The rest of the paper is organized as follows. Section 2 describes the data employed and the forecasting problem, including the forecasts provided by the SPF. Section 3 outlines the Gaussian copula model, and Section 4 provides an overview of how to undertake likelihood-based estimation and inference, including the construction of predictive distributions. Section 5 contains the empirical results and the forecasting study, while Section 6 concludes the paper. The Appendix gives details on computing the likelihood function, while a Web Appendix contains some properties of the model, and additional empirical results.

2 The Data and Forecasting Problem

As Croushore and Stark (2001) note, many macroeconomic variables are subject to considerable and badly-behaved data revisions. Failing to account for this by using heavily revised data can mask real-time predictive content. Therefore, we employ real-time quarterly data on four U.S. macroeconomic variables: output growth, inflation, the unemployment rate and the nominal short-term interest rate. Clark and Ravazzolo (2014) consider the same

variables in their comparison of different forecasting models. We summarize the data below, but refer the reader to Clark and Ravazzolo (2014) for a detailed discussion.

Output growth and inflation are measured as the logarithmic change from the previous quarter in real GDP and the GDP price level, respectively. The unemployment rate and the short-term interest rate are measured in percentage points. The data for output growth and inflation are observed in real-time, with quarterly vintage $T + 1$ containing data up to quarter T , which is the forecast origin. This reflects the time lag in the release of the national accounts. The quarter $T + 1$ vintage has the first release for the last time series observation, dated T , and the earlier observations, $t < T$, correspond to second or subsequent releases.

The quarterly data for GDP (in practice, GNP for some vintages) are taken from the Federal Reserve Bank of Philadelphia’s Real-Time Data Set for Macroeconomists. This is a collection of vintages of National Income and Production Accounts, with each vintage reflecting the information available around the middle of the respective quarter. Croushore and Stark (2001) provide a description of the database. The measure of inflation utilizes the analogous GDP price series, which like the GDP data, suffers from revisions. Monthly data on the unemployment rate and short-term interest rate (the 3-month Treasury bill rate) were obtained from the FAME database of the Federal Reserve Board of Governors. Both variables were transformed to quarterly frequency by averaging over the relevant quarter. Figure 1 plots the four data series from the 2011:Q2 vintage data in the left-hand panels.

Similar to Clark and Ravazzolo (2014), we utilize the real time data from 1954:Q1 until vintage $T + 1$ to estimate the model parameters. Forecasts are made using a quarterly expanding window with the first forecast origin at $T = 1974:Q4$ (so that the first vintage used for parameter estimation is 1975:Q1), and the last at $T = 2011:Q1$ (so that the last vintage used for parameter estimation is 2011:Q2). Forecasts are made over a horizon of up to 8 quarters for each window, and the out-of-sample forecast evaluation period runs from 1975:Q1 through 2011:Q2. Forecasts made one quarter ahead are for the values of the variables at the same quarter as that of the vintage, and are therefore usually called the ‘nowcast’.

To evaluate our density forecasts requires an additional assumption about which mea-

surements for GDP growth and inflation in the forecast horizon are to be considered the true value, commonly called the ‘outturn’. Corradi, Fernandez and Swanson (2009) discuss the implications of using alternative definitions of the target measurement for forecasting, as well as for parameter estimation. Here, we follow Romer and Romer (2000), Faust and Wright (2009), Clark and Ravazzolo (2014) and others, and use the second release as the outturn for forecast evaluation purposes.

Ideally, we would compare our forecast densities with those from the Federal Reserve. Unfortunately, this central bank does not publish the relevant complete forecast densities, and we instead compare our forecast densities with those produced by the Survey of Professional Forecasters (SPF). This is a quarterly survey aimed at professional forecasters, which began in 1968:Q4, but was taken over by the Philadelphia Federal Reserve in 1990:Q2. Croushore (1993) describes the SPF in detail and notes that roughly a third of respondents are employed in financial firms on Wall Street. The remainder are from other private sector firms and academia. Since 1990:Q2, the deadline for responses to the SPF questionnaire occurs in the second to third week of the middle month of each quarter. Respondents observe the first release of the national accounts data referring to the previous quarter before completing the survey. Because the deadline occurs mid-quarter, the respondents also observe up to half of the higher frequency data for the current quarter. For example, the report published on 12 February 2008 has a forecast origin of 2007:Q4 for national accounts data, although it is likely to be at least partially informed by additional information available in the first half of 2008:Q1.

We focus on the summary table labeled ‘Table Five: Mean Probabilities’ in the SPF reports. This records average probabilities across respondents for the ‘annual-average over annual-average’ for two variables: growth of real GDP and the corresponding measure of inflation. Probabilities are reported for bins of width one percentage point for the outcome of each variable. The individual responses are averaged within each bin—in effect, a linear opinion pool of professional forecasters’ views (Wallis 2005). The SPF request for probabilities from professional forecasters takes the following form:

“Please indicate what probabilities you would attach to the various possible percentage

changes (annual-average over annual-average) this year and the next in chain-weighted real GDP and the chain-weighted GDP price index.”

Respondents are asked to supply probabilities for the current year relative to the previous year, and also the following year relative to the current year. For example, the report dated 12 February 2008 asked for forecasts for 2007-2008 and 2008-2009, while the report dated 13 February 2009 asked for forecasts for 2008-2009 and 2009-2010.

We assume that the respondents produce their probabilities for GDP growth in the following manner, using the forecasts for 2008-2009 made in the 12 February 2008 report as an example. First, they produce multi-horizon forecasts for the value of GDP in each quarter of 2008, and then average them. The forecasts are for $GDP_{08:Q1}$, $GDP_{08:Q2}$, $GDP_{08:Q3}$ and $GDP_{08:Q4}$; the sum of which divided by four corresponds to the ‘annual average level of GDP for 2008’, labeled $\overline{GDP}_{08}$. The respondents repeat the process for 2009 to generate forecasts for $GDP_{09:Q1}$, $GDP_{09:Q2}$, $GDP_{09:Q3}$ and $GDP_{09:Q4}$, and the associated ‘annual average level of GDP for 2009’, labeled $\overline{GDP}_{09}$. Finally, the forecast ‘annual average over annual average’ growth rate reported in the SPF report ‘Table Five’ is constructed as $\overline{R}_{GDP,08-09} = \log(\overline{GDP}_{09}) - \log(\overline{GDP}_{08})$.

Notice that as we move to the subsequent report, dated May 13, 2008, the information set shifts in time, and the first release of GDP for 2008:Q1 becomes known. Furthermore, the forecast horizon required to construct $\overline{R}_{GDP,08-09}$ varies with release date. The forecasts reported in May 13, 2008 require only forecasts 7 quarters ahead, not 8. By the report dated November 16, 2009, the only forecast horizon required is one quarter ahead (ie. the nowcast). The forecast for inflation $\overline{R}_{INF,08-09}$ is defined in the same fashion, but employing the price index rather than real GDP.

The predictive distribution of both $\overline{R}_{GDP,08-09}$ and $\overline{R}_{INF,08-09}$ resulting from a multivariate time series model can also be constructed as follows. Monte Carlo iterates of future quarterly growth and inflation are generated from the model and, combined with the observations already available in the vintage, used to compute Monte Carlo iterates of $\overline{R}_{GDP,08-09}$ and $\overline{R}_{INF,08-09}$. From these a histogram estimate of their distributions can be constructed, allowing direct comparison with the ensemble distributions published by the SPF.

3 Copula Model for Multivariate Time Series

In this section we outline the Gaussian copula model that we employ. We discuss how the partial correlations of a latent Gaussian process are an efficient parameterization of the Gaussian copula when it is represented as a vine copula. This representation is used later in Section 4 to compute the likelihood.

3.1 Model definition

Consider a continuous-valued random vector $\mathbf{Y}_t = (Y_{1,t}, \dots, Y_{m,t})'$ observed at times $t = 1, \dots, T$, with corresponding realization $\mathbf{y}_t = (y_{1,t}, \dots, y_{m,t})'$. Most multivariate time series models are defined via the transitional distributions of $\mathbf{Y}_t | \mathbf{Y}_{t-1}, \dots, \mathbf{Y}_1$. This restricts the form of the marginal distributions of each variable $Y_{i,t}$, which are evaluated either analytically for simple models or in a Monte Carlo fashion for most nonlinear time series. For example, in a VAR with Gaussian disturbances the margins are also Gaussian, and in a VAR with stochastic volatility (e.g. Primiceri 2005) the margins are usually a mixture of Gaussian or t distributions. In contrast, in a copula model of a multivariate distribution, the dependence structure is separated from the marginal distributions. This allows for any model to be adopted for the marginal distribution of each variable, including asymmetric, heavy-tailed, censored or even nonparametric distributions. These are then ‘glued together’ to form a multivariate model using a copula function, the form of which is not affected by the choice of margins. This feature is of great potential value in the time series modeling of macroeconomic variables, which often exhibit very varied marginal distributions empirically, and is something we exploit in our empirical work when computing forecast distributions.

To construct a copula model for the multivariate time series we first stack the elements into the vector $\mathbf{Y} = (\mathbf{Y}'_1, \dots, \mathbf{Y}'_T)'$, with $N = Tm$ elements. Then, by Sklar’s Theorem (Nelsen 2006; p.45), there always exists a copula function C such that the joint distribution function of $\mathbf{Y}$ is

$$F(\mathbf{y}) = C(\mathbf{u}), \tag{3.1}$$

where $\mathbf{u} = (\mathbf{u}'_1, \dots, \mathbf{u}'_T)'$ and $\mathbf{u}_t = (F(y_{1,t}), \dots, F(y_{m,t}))'$. The copula function C is itself a distribution function of $\mathbf{U} = (\mathbf{U}'_1, \dots, \mathbf{U}'_T)'$ on the unit cube $[0, 1]^N$, with $\mathbf{U}_t =$

$(F(Y_{1,t}), \dots, F(Y_{m,t}))'$ and all N margins being uniform distributions. All dependence between the elements of $\mathbf{Y}$ (both cross-sectional and serial) is captured by the copula function. In effect, once appropriate marginal distributions $F(Y_{j,t})$ have been selected for all variables, this approach transforms the data onto the unit cube. Our copula model then is a multivariate time series model for this transformed data, that is standardized to have realizations on $[0, 1]^m$ and uniform margins. This is something we illustrate later in our empirical work.

Differentiating the distribution function gives the density of $\mathbf{Y}$ as

$$f(\mathbf{y}) = c(\mathbf{u}) \prod_{t=1}^T \prod_{j=1}^m f(y_{j,t}). \quad (3.2)$$

Here, $f(y_{j,t})$ denotes the marginal density of $y_{j,t}$, $F(y_{j,t})$ is the corresponding distribution function, and $c(\mathbf{u}) = \frac{\partial}{\partial \mathbf{u}} C(\mathbf{u})$ is often called a ‘copula density’. Sklar’s Theorem reassures that such a copula always exists, although in practice it is unknown and has to be selected. In principle any copula can be used, although parametric copulas with densities that can be readily computed are preferred when undertaking likelihood-based estimation. This is the case for the Gaussian copula, which we employ here. If $\Phi(x)$ is the standard normal distribution function, then the Gaussian copula density is

$$c(\mathbf{u}) = c_{\text{Ga}}(\mathbf{u}; \Omega) = |\Omega|^{-1/2} \exp \left\{ -\frac{1}{2} \mathbf{w}' (\Omega^{-1} - I_N) \mathbf{w} \right\}, \quad (3.3)$$

with $\mathbf{u} = (u_1, \dots, u_N)'$, $\mathbf{w} = (\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_N))'$, and the correlation matrix Ω as dependence parameters. Song (2000) discusses the properties of a Gaussian copula, and it is straightforward to show that $\mathbf{w}$ is a realization from a latent variable $\mathbf{W} = (W_1, \dots, W_N)' \sim N(0, \Omega)$. We stress here that adopting a Gaussian copula does not mean that $\mathbf{Y}$ is distributed Gaussian. For example, in our empirical work we adopt highly non-Gaussian marginal distributions for the four macroeconomic variables in this study.

In general, the dimension of the copula $N = Tm$ is high, so that estimation of an unrestricted parameter matrix Ω would not be possible. However, we follow Biller and Nelson (2003) and Biller (2009) and assume that the latent Gaussian variable arises from a sta-

tionary Gaussian VAR(p), so that Ω is highly structured. To show this, note that here $\mathbf{W} = (\mathbf{W}'_1, \dots, \mathbf{W}'_T)'$ comprises a time series of multivariate vectors $\mathbf{W}_t = (\Phi^{-1}(U_{1,t}), \dots, \Phi^{-1}(U_{m,t}))'$. If the latent process $\{\mathbf{W}_t\}_{t=1}^T$ follows a stationary Gaussian VAR(p) with unit marginal variance (ie. $\text{Var}(\Phi^{-1}(U_{j,t})) = 1$ for all j and t), then Ω is the block Toeplitz correlation matrix of this process. That is, a matrix of $(T \times T)$ blocks, each of size $(m \times m)$, with the (s, t) th block being $\Omega(s - t) = \text{Corr}(\mathbf{W}_s, \mathbf{W}_t)$, where $\Omega(-t) = \Omega(t)'$ and $\Omega(0)$ is symmetric; see Lütkepohl (2006; p.30). For likelihood-based estimation, we parameterize Ω in terms of its partial correlations, and the copula in terms of its vine representation, both of which we discuss in more detail in Section 3.2.

If all the margins of $\mathbf{Y}$ are restricted to be Gaussian with time invariant means and variances, then $\mathbf{Y}$ is multivariate Gaussian (Song 2000) with correlation structure given by Ω , so that the model is a Gaussian VAR. However, the model extends the Gaussian VAR whenever one or more of the margins are non-Gaussian. Yet the dependence structure is still determined by the correlation structure of the latent stationary VAR(p). Some properties of the proposed time series model are given in the Web Appendix. In this we show that assuming a stationary VAR for the latent Gaussian process above ensures that the process $\{\mathbf{Y}_t\}_{t=1}^T$ is also (strongly) stationary, as long as the marginal models for elements $Y_{j,t}$ are invariant with respect to time t . This is the case in our empirical work, where the marginal models differ for each macroeconomic variable, but do not vary over time. Similarly, the time series has the same Markov order as that of the latent time series, which we set to $p = 4$ quarters in all our empirical work. Last, while in this study we only consider the special case of a Gaussian copula, by selecting different copula functions for C , different dependence structures can also be obtained, while retaining the same margins.

3.2 Partial correlations and vine representation

We parameterize the $N \times N$ correlation matrix Ω in terms of its (semi) partial correlations. If we denote the elements of $\mathbf{W} = (W_1, \dots, W_N)'$ with a single subscript, then these are

$$\phi_{i,j} = \text{Corr}(W_i, W_j | W_{j+1}, \dots, W_{i-1}),$$

for $i > j$, which are defined to be unconditional correlations when $i - j = 1$. Denoting the partial correlations as $\boldsymbol{\phi} = \{\phi_{i,j}; i = 2, \dots, N, j < i\}$, there is a one-to-one mapping between $\boldsymbol{\phi}$ and Ω that is widely attributed to Yule; for example, see Joe (2006) and Daniels and Pourahmadi (2009). Joe (2006) notes that an advantage of the partial correlations as a parameterization is that, conditional on all the other partials, $\phi_{i,j}$ remains unconstrained on $[0, 1]$, so that the likelihood is easier to analyze. However, for the multivariate time series employed here, we highlight two major additional advantages of this parameterization. The first is that it is highly parsimonious for stationary and/or Markov series, both of which we assume in our forecasting application. The second is that the partial correlations are also the parameters of a drawable vine (or D-vine for short) representation for the Gaussian copula. We discuss each point further below.

Let $a(t) = (t - 1)m + 1$ and $b(t) = tm$ denote the first and last elements of $\mathbf{Y}$ that correspond to the sub-vector $\mathbf{Y}_t$. Then, to facilitate a parsimonious parameterization, we group the partial correlations together into blocks as

$$\begin{aligned}\boldsymbol{\phi}_{(t,s)} &= \{\phi_{i,j}; a(t) \leq i \leq b(t), a(s) \leq j \leq b(s)\}, \text{ with } m^2 \text{ elements, and} \\ \boldsymbol{\phi}_{(t,t)} &= \{\phi_{i,j}; a(t) < i \leq b(t), a(t) \leq j < i\} \text{ with } m(m-1)/2 \text{ elements,}\end{aligned}$$

for $1 \leq s < t \leq T$. Then it is possible to show that Ω is a block Toeplitz matrix with blocks of size $(m \times m)$, if (and only if) the blocks of partial correlations $\boldsymbol{\phi}_{(t,t-k)} = \boldsymbol{\phi}_{(t',t'-k)} \equiv \boldsymbol{\phi}(k)$ for all t, t' and $k \geq 0$. This is exactly the condition established at Property 3 of the Web Appendix for $\{\mathbf{U}_t\}_{t=1}^T$ to be a strongly stationary series. If the margins are time invariant, then Property 4 establishes that $\{\mathbf{Y}_t\}_{t=1}^T$ is also a strongly stationary series. In addition, when the series is Markov of order p , then $\boldsymbol{\phi}(k) = \mathbf{0}$ for all $k > p$. In this case, the set of unique partial correlations $\tilde{\boldsymbol{\phi}} = \{\boldsymbol{\phi}(0), \dots, \boldsymbol{\phi}(p)\}$ provide a parameterization of the copula that is both parsimonious and unconstrained on the unit cube. For example, in our macroeconomic application $m = 4$ and we assume a stationary series with Markov order $p = 4$, which results in only $4(4^2) + 4(3)/2 = 70$ partials to be estimated, given the models for the margins.

Vine copulas are high-dimensional copulas which have densities equal to the product of a

sequence of component bivariate copula densities, called ‘pair-copulas’ (Aas et al. 2009). Let $c_{i,j}(u, v; \rho)$ denote the density of a bivariate Gaussian copula with a 2×2 correlation matrix with correlation coefficient ρ , which is the single dependence parameter of the copula. Then, Bedford and Cooke (2002) show that the Gaussian copula density at Equation (3.3) has the D-vine decomposition

$$c_{\text{Ga}}(\mathbf{u}; \Omega) = \prod_{i=2}^N \prod_{j=1}^{i-1} c_{i,j}(u_{i|j+1}, u_{j|i-1}; \phi_{i,j}). \quad (3.4)$$

The pair-copulas $c_{i,j}$ are bivariate Gaussian copulas, each of which has the partial correlation $\phi_{i,j}$ as its parameter. The arguments of each pair-copula, $u_{i|j+1}$ and $u_{j|i-1}$, can be computed from $\mathbf{u}$ and $\boldsymbol{\phi}$ using Algorithm 1 in Smith (2014), which we outline for the specific case of a Gaussian copula in the Appendix.

We use this D-vine representation of the Gaussian copula to compute the likelihood function. When $\phi_{i,j} = 0$, the bivariate pair-copula density $c_{i,j}(\cdot, \cdot; 0) = 1$, which is the density of the ‘independence copula’ (Nelsen 2006, p.25). For series with Markov order $p \ll T$, the vast majority of pair-copulas are independence copulas. This greatly speeds the computation of the copula density and likelihood using Equation (3.4). Moreover, as we discuss below, Bayesian estimators that shrink the remaining non-zero partial correlations in $\tilde{\boldsymbol{\phi}}$ towards zero through model averaging can further speed computation of the likelihood.

4 Estimation Methodology

4.1 Estimation

Copula models are usually estimated in two steps. The first step is the estimation of the marginal models, from which probability integral transformed values (PITs) are computed using the fitted marginal distribution functions. In the copula modeling literature, these are called the ‘copula data’ $\mathbf{u} = (\mathbf{u}'_1, \dots, \mathbf{u}'_T)'$. The second step involves the estimation of the copula parameters, which are the partial correlations here, using the copula data. The approach is often called ‘inference for margins’, and Joe (2005) shows it is only slightly

less efficient than full maximum likelihood. When estimating the copula parameters using Bayesian methods, it also simplifies the development of a Markov chain Monte Carlo (MCMC) sampling scheme substantially. In our empirical work we employ both skew t distributions and nonparametric approaches for estimating the marginal models, and computing the copula data. The unique copula parameters $\tilde{\phi}$ are then estimated using the Bayesian method proposed in Smith (2014) for D-vine copulas. We summarize this approach here in the special case of our Gaussian copula, but refer the reader to Smith (2014) for details.

The likelihood is equal to the copula density evaluated at $\mathbf{u}$, which we express using the D-vine representation at Equation (3.4). This can be further simplified for our time series copula as follows. It is possible to show that the pair-copulas with partial correlation parameters $\phi_{(t,s)}$, for $s < t$, are functions of only $\{\mathbf{u}_s, \dots, \mathbf{u}_t\}$ and not the full set of copula data. Therefore, the product of corresponding pair-copula terms can be denoted as the functional $K_{t,s}(\mathbf{u}_s, \dots, \mathbf{u}_t; \phi_{(t,s)})$. In addition, in our stationary Markov order p series, the partials $\phi_{(t,s)} = \phi(t-s)$, with $\phi(k) = \mathbf{0}$ for $k > p$, so that also $K_{t,s} = 1$ for $t-s > p$. This greatly reduces the number of pair-copula densities that need computing to evaluate the likelihood, which can be written as

$$\mathcal{L}(\phi(0), \dots, \phi(p); \mathbf{u}) = \prod_{t=1}^T \prod_{\substack{s=1 \\ t-s < p}}^t K_{t,s}(\mathbf{u}_s, \dots, \mathbf{u}_t; \phi(t-s)). \quad (4.1)$$

There is much empirical evidence that Bayesian shrinkage can increase the accuracy of forecasts of macroeconomic variables when using reduced-form models; for some recent examples, see De Mol et al. (2008), Banbura et al. (2010), Koop and Korobilis (2013) and Clark and Ravazzolo (2014). The vine representation allows for Bayesian shrinkage via a ‘spike-and-slab’ prior on the component pair-copulas, as we now outline. First, for each $k = 0, 1, \dots, p$, we denote the unique partial correlations as $\phi(k) = \{\phi_{l_1, l_2}^k\}$. Here, we distinguish between all $N(N-1)/2$ partial correlations, and the much smaller number of unique values in $\tilde{\phi}$ when the series is Markov and stationary. Then we introduce corresponding binary indicators $\gamma(k) = \{\gamma_{l_1, l_2}^k\}$, where $\gamma_{l_1, l_2}^k = 0$ if and only if $\phi_{l_1, l_2}^k = 0$ (i.e. the ‘spike’ or point mass), and $\gamma_{l_1, l_2}^k = 1$ otherwise. Notice that if $\gamma_{l_1, l_2}^k = 0$, then the corresponding

bivariate Gaussian pair-copulas are also equal to the independence copula with unit density. For the dependent pair-copulas, a uniform prior is placed on the value of Kendalls tau for that pair-copula. This prior is similar to that placed on elements of the inverse of Ω for a Gaussian copula by Pitt, Chan and Kohn (2006). However, such elements have complex bounds which are difficult to enforce when developing Bayesian estimation schemes; whereas here each (semi) partial correlation is unconstrained on $[0, 1]$. We employ the prior on $\tilde{\gamma} = \{\gamma(0), \dots, \gamma(p)\}$ given in Smith and Kohn (2002), which adjusts for the well-known problem of multiplicity; see Scott and Berger (2010) and Korobilis (2013).

Estimation is via a MCMC sampling scheme. This scheme randomly selects two pair-copulas and generates the two parameters and two corresponding binary indicators, jointly as a block. A Metropolis-Hastings step is employed to generate from the conditional posterior of this parameter block, within a wider Gibbs sampler. Computation of the likelihood can be undertaken using parallel processing, greatly speeding implementation in real time; see the Appendix. The end result is a Monte Carlo sample $\{(\tilde{\gamma}^{[1]}, \tilde{\phi}^{[1]}), \dots, (\tilde{\gamma}^{[J]}, \tilde{\phi}^{[J]})\}$ of size J from the posterior distribution of the parameters augmented with the binary variables $\tilde{\gamma}$. The posterior means can be used as point estimates of the parameters, or other measures of dependence, and approximated using the Monte Carlo iterates. Estimates are model averages over the distribution of $\tilde{\gamma}$, and selection of pair-copulas here extends the Bayesian selection and model averaging approaches proposed previously for a Gaussian VAR (George, Sun and Ni 2008; Koop and Korobilis 2010; Korobilis 2013).

4.2 Measuring dependence

At each sweep of the sampling scheme we compute $\{\Omega(0), \dots, \Omega(p)\}$ from the the current iterate of $\tilde{\phi}$. We do so by applying Yule's algorithm (Daniels and Pourahmadi 2009) to compute the sub-matrix of Ω comprised of only the first $m(p+1)$ rows and columns, and extract the component block matrices $\{\Omega(0), \dots, \Omega(p)\}$. If $\omega_{i,j}(k) = \text{Corr}(W_{i,t}, W_{j,t-k})$ is the (i, j) th element of $\Omega(k)$, then these values are measures of linear pairwise correlations for the latent Gaussian time series. However, the pairwise dependence between $Y_{i,t}$ and $Y_{j,t-k}$ may well be nonlinear, depending on the form of the margins. Therefore, we report Spearman's

rho, which is

$$\rho(Y_{i,t}, Y_{j,t-k}) = r_{i,j}(k) = \frac{2}{\pi} \arcsin(\omega_{i,j}(k)) ,$$

for lags $k = 0, 1, \dots, p$. If $R(k)$ is an $(m \times m)$ matrix with (i, j) th element $r_{i,j}(k)$, then $R(0)$ is a coherent measure of cross-sectional dependence of $\mathbf{Y}_t$, and $R(k)$ measures serial dependence at lags $1 \leq k \leq p$. Note that $R(0)$ is a symmetric matrix, whereas $R(k)$ is not for $k > 0$, because the autocorrelation matrices $\Omega(k)$ are also asymmetric. Again, the posterior means of $R(0), \dots, R(p)$ can be used as point estimates, and can be computed using the Monte Carlo iterates for $\Omega(0), \dots, \Omega(p)$. We note that it is also possible to compute $\Omega(k)$, for $k = p + 1, \dots, T - 1$, at each sweep of the sampler by employing Yule's algorithm to compute the full Ω matrix. However, we do not do so in our empirical analysis because they hold little information and their computation slows estimation.

4.3 Density forecasting

Predictive distributions are computed by simulating a 'ray' of future values from the copula model up to 8 quarters ahead, as discussed below. These rays are then transformed via the inverse distribution function for each marginal model to produce Monte Carlo iterates from the predictive distribution of $\mathbf{Y}_{T+8}, \dots, \mathbf{Y}_{T+1} | \mathbf{Y}_T, \dots, \mathbf{Y}_{T-p}$.

Simulation of values from the Gaussian copula model can be undertaken by generating an iterate from the Gaussian conditional distribution $\mathbf{W}_{T+8}, \dots, \mathbf{W}_{T+1} | \mathbf{W}_T, \dots, \mathbf{W}_{T-p}$ of the latent zero mean Gaussian process. Note that $\text{Var}(\mathbf{W}_{T-p}, \dots, \mathbf{W}_{T+8}) = \Sigma_f$ is a $m(p+8) \times m(p+8)$ block Toeplitz matrix with (s, t) th block given by $\Omega(s-t)$. Given Σ_f , it is straightforward to compute and simulate from the Gaussian conditional distribution. Moreover, by simulating one or more iterates at the end of each sweep of the MCMC scheme, uncertainty associated with the parameters of the copula model is accounted for in the usual Bayesian fashion. (An alternative approach to generate from the Gaussian copula model is proposed by Smith (2014), based on its vine copula representation and using only the partial correlations.)

5 Empirical Analysis

In this section we fit the Gaussian copula-based multivariate time series model to the four U.S. macroeconomic variables— output growth, inflation, unemployment and the interest rate. We discuss the fitted margins and the fitted copula separately, and to simplify presentation, with a focus on the 2011:Q2 vintage of data. We show that Gaussian margins are unrealistic for each variable, and therefore employ non-Gaussian models for the margins in our copula model. We also show how the resulting copula model allows for nonlinear dependence and multivariate asymmetry in the four variables, whereas a Gaussian VAR has necessarily linear correlation and symmetric multivariate margins. Next, we report on the forecasting performance of the copula model over our evaluation period using the real-time data from vintages 1975:Q1 to 2011:Q2. We highlight that the asymmetries evident in the fitted models carry over into the out-of-sample predictive distributions, and that the resulting forecasts are competitive with those from a Bayesian VAR. Last, we examine in detail the (ensemble) probabilistic forecasts from the SPF quarterly reports during the Great Recession, and compare them with the density forecasts from both our copula-based approach and the Gaussian VAR. We find a high degree of asymmetry in the SPF probabilistic forecasts, and that this is consistent with the asymmetry of the predictive distributions from the copula-based model, but inconsistent with the symmetric density forecasts from VAR based models.

5.1 Fitted model

Fitted margins, asymmetry and heavy tails

The Shapiro and Wilk (1965) test for normality is amongst the most powerful for the sample sizes employed here. Using this, normality is rejected at the 5% level for the four variables at all but a handful of early vintages of GDP growth and the interest rate. For example, for the 2011:Q2 vintage data, the p -value of the null hypothesis of normality for GDP growth is 0.0005, and is much lower for each of the other three variables. (A tabulation of the test results for all vintages is given in the Web Appendix.) We therefore consider fitting both a univariate skew t distribution (Azzalini and Capitanio 2003) and the empirical distribution

function (EDF) to each series. The former is a four parameter generalization of the t distribution, and is often employed to model data with skew and kurtosis (eg. see Walls 2006 and Panagiotelis and Smith 2008). The EDF is a nonparametric estimator and is widely advocated in the copula modeling literature (Shih and Louis 1995; Chen and Fan 2006).

Fig. 2(a), (c), (e) and (g) plot Gaussian, skew t and nonparametric density estimates for the vintage 2011:Q2 data. (For the nonparametric estimator we plot a kernel density estimate for visual presentation, simply because it is much smoother than the highly ‘spikey’ density arising from the interpolated EDF.) Both the skew t and nonparametric distributions capture strong positive skew in inflation, unemployment and the interest rate, along with modest heavy tails in GDP growth. Fig. 2(b), (d), (f) and (h) compare the densities of skew t distributions fit to four vintages that are approximately 50 quarters apart. They show that these characteristics are prevalent at other vintages as well. In addition, GDP growth also exhibits some negative skew at earlier vintages. Using both the fitted skew t distributions and the nonparametric EDFs we compute the copula data. Fig. 1 plots the four transformed copula data series for the skew t margins in the right-hand panels, and serial dependence is apparent visually; particularly in panels (d), (f) and (h). Last, we note that the fitted marginal distributions based on differing vintages of data in Fig. 2 should not be confused with the predictive distributions, which differ greatly as we show below.

Fitted copula and dependence structure

The Bayesian approach described in Section 4 is used to fit the Gaussian copulas. We use a burnin period of 5000 sweeps, and collect Monte Carlo samples during a further 10000 sweeps. Common measures suggest convergence to the posterior distribution appears to occur within 2000 sweeps, and the Monte Carlo estimates do not vary meaningfully with samples half or twice as large. To illustrate, in the rest of this subsection we report estimates obtained from the 2011:Q2 vintage data with skew t marginal models. Here, Bayesian selection identifies a high level of parsimony in the fitted model. There are 6 unique parameters in $\phi(0)$, and 16 in each of $\phi(1), \dots, \phi(4)$. The expected number of non-zero values in $\phi(0), \dots, \phi(4)$ is 5.7, 12.7, 6.5, 5.9 and 3.7, respectively, which suggests there is substantial shrinkage of the partial correlations to zero at longer lags.

Fig. 3 plots the posterior means of the matrices $R(0), \dots, R(4)$ defined in Section 4.2. In each matrix, the (i, j) th element of $R(k)$ contains the unconditional pairwise Spearman correlations between $Y_{j,t-k}$ and $Y_{i,t}$ from the fitted parametric model. That is, the lagged variables are the columns. The estimates confirm much of the usual cross-sectional and serial dependence structure for U.S. macroeconomic variables. For example, the element in the third row and first column of panel (e) displays a Spearman correlation of -0.28 between the unemployment rate at time t and GDP growth in time $t - 4$. That is, the unemployment rate is a lagging indicator. Tracking the analogous elements through the other matrices reveals that the Spearman correlation drops as the time lag shortens below three quarters. Similarly, the element in the bottom row and second column of panel (e) indicates a Spearman correlation of 0.64 between inflation at time $t - 4$ and the short-term interest rate in time t —the (nominal) interest rate rises with inflation at a time lag of four quarters, consistent with monetary control by the central bank. The sequence of corresponding elements in the other matrices shows that the correlation varies little as the lag length falls. A further finding is that the interest rate and the unemployment rate have the highest serial dependence of the four variables, whereas GDP growth exhibits the lowest level; observations that are consistent with the time series plots in Fig. 1.

Multivariate asymmetry and nonlinear dependence

Despite the latent Gaussian process having a linear dependence structure, our model not only captures strong asymmetry and heavy tails, it also allows for nonlinear dependence. To illustrate this, Fig. 4 provides density contour plots of two bivariate margins of the fitted copula model— one in the unemployment rate and GDP growth lagged three quarters, and another in the GDP growth and inflation lagged two quarters. These contours were estimated using bivariate kernel methods based on large Monte Carlo samples simulated from the fitted copula model using the approach in Smith (2014). Not only are the bivariate densities highly non-Gaussian (which is a direct result of the skew t univariate margins), but also show nonlinear dependence between the pairs of variables. To highlight this, smoothing splines were fit to the Monte Carlo iterates to provide estimates of the expectations of the non-lagged variable, conditional only on the lagged variable. These are plotted as thick

lines and are nonlinear. For example, the marginal impact of changes in GDP growth on the unemployment rate three quarters later, is greater during contractionary quarters (ie. when GDP growth is negative) than during expansionary quarters (ie. when GDP growth is positive). Scatterplots of the observed data overlay the contour plots, and are consistent with these copula model densities. The bivariate margins for a Gaussian VAR (not shown here) are constrained to be Gaussian with linear dependence by construction, and are inconsistent with the observed data. Further summaries of the fitted copula are given in the Web Appendix.

5.2 Forecasts

We consider the predictive performance over the evaluation period 1975:Q1 to 2011:Q2. An expanding window is used, with the in-sample period beginning at 1951:Q1, as outlined in Section 2. Two copulas models, one with EDF margins (labeled C-EDF) and the other with skew t margins (labeled C-Skt), are fit and large Monte Carlo samples obtained from the forecast distributions over a horizon of 8 quarters as outlined in Section 4.3. Both the margins and copula parameters are continually refit as the in-sample period is expanded quarter-by-quarter. We employ a Bayesian VAR with a Minnesota prior (labeled BVAR) as the benchmark model. Note that our copula model nests the Gaussian VAR, but has a shrinkage prior on the partial correlation parameters, which differs from a Minnesota prior. We also employ a Bayesian VAR with stochastic volatility models for the disturbances (labeled BVAR-SV), although we note that our copula model does not explicitly account for stochastic volatility. Implementation of these two models is outlined in Clark and Ravazzolo (2014). Our measure of point forecast performance is the Root Mean Squared Error (RMSE), while the Continuous Ranked Probability Score (CRPS) of Gneiting and Raftery (2007) is our measure of density forecast accuracy. We also compute a symmetric tail-weighted version of the CRPS proposed in Gneiting and Ranjan (2011) that rewards tail accuracy, which we label TW-CRPS. We compute the latter two metrics via numerical integration of the quantile score as outlined in the Appendix. This is a more accurate method of computation than using Monte Carlo methods as in Panagiotelis and Smith (2008) or Ravazzolo and Vahey (2014).

Tables 1 and 2 report the mean values of RMSE and CRPS for all four time series models

over the full forecast evaluation period, and 9 contractionary quarters (1982:Q1–1982:Q2, 1991:Q1–1991:Q3 and 2008:Q4–2009:Q3). The last are the three post-1975 quarters where there were two or more consecutive quarters of negative GDP growth in the outturn data. Results are reported for horizons $h = 1, 2, 4$ and 8 quarters ahead, with $h = 1$ being the nowcast. As a rough guide, results of a two-sided test of the equality of the mean score of the BVAR with each of the other three models are reported for the full evaluation period. The test employed is the Diebold and Mariano (1995) test with the small sample adjustment suggested by Harvey, Leybourne and Newbold (1997). Table 1 of the Web Appendix reports results for the TW-CRPS metric, which we do not discuss here because they are very similar to those for the CRPS.

Overall, the results follow the finding of Clark and Ravazzolo (2014) that the BVAR-SV model is the best forecasting method during the full evaluation period. However, we also find that the C-Skt often has the second lowest metric values, and usually below those of the BVAR model. We also find that the use of EDF margins can seriously degrade the forecasts, so that the C-EDF model produces poor forecasts at times. This is because the nonparametric estimator is based on small samples, so that the marginal density estimates can be poor. In particular, the minimum and maximum possible forecast values are equal to the smallest and largest values observed previously, which can give tail probabilities that are too low. In contrast, C-Skt has smooth forecast densities with tails on the entire real line.

We note some additional features of the forecasts, mainly focusing on forecast density accuracy as measured by the CRPS in Table 2. First, BVAR-SV provides particularly accurate forecasts for GDP growth, which suggests that the capturing of persistence in quarters of subdued or inflated volatility in output growth is important. In contrast, this is not a feature of either copula model fit, and something we discuss further in Section 6. Second, C-Skt and BVAR-SV have almost identical forecast accuracy for inflation overall. Last, both C-Skt and BVAR-SV improve significantly over the BVAR in interest rate forecasting at shorter horizons, with the C-Skt giving the best long horizon forecasts.

During the contractionary periods the results lack statistical significance due to the low

sample size (9 quarters). Nevertheless, the forecast densities differ greatly between models, which we illustrate using forecasts of inflation and the interest rate during the Great Recession— a period that is of particular interest to monetary policymakers. Fig. 5 shows inflation forecast densities made using vintage 2009:Q2 data, which is during the recession. Forecasts from the BVAR and BVAR-SV models are Gaussian (up to integration over the parameter posterior distribution), with the latter having much higher variance because by this quarter the model has registered a large jump in volatility. Symmetry here seems unrealistic, and both models provide high probabilities of substantial deflation that we show below are inconsistent with those produced by the SPF. In contrast, both copula forecast densities are sharper (ie. more peaked) and have a heavy upper tail.

Fig. 6 shows fan charts for interest rate forecasts made using vintages 2007:Q4 through 2011:Q1 for the four time series models. The thick lines are the median of the predictive distributions from 1 to 8 quarters ahead. The circles denote the observed data, and the 10th and 90th percentiles of the predictive distributions are plotted as thin lines, thereby creating a series of ‘fans’. Two specific fans are highlighted in black. The first is based on vintage 2007:Q4 data, which is during the financial crisis that preceded the recession. Here, the BVAR-SV model has low variance estimates carrying over from the previous Great Moderation period, producing a fan that is narrower than those of the other three models. In contrast, by the end of the recession, using vintage 2009:Q4 data, forecasts from the BVAR-SV model have very high variance. This produces high probabilities of a negative interest rate, which is improbable. Forecast probabilities of a rate below -2% exceed 10% during 5 consecutive quarters in the forecast horizon. In comparison, the C-Skt model gives a much lower probability of negative rate and C-EDF none. In the latter case, this is because the lower bound of the margin is the minimum previously observed interest rate (which has to be positive). This highlights an attractive feature of our copula model— that bounds and other constraints can be readily be placed on the variable values via the choice of marginal models. Here, a zero lower bound could be imposed on the interest rate, which can be difficult in other time series frameworks, and important in the current monetary environment. Fan charts for the other three variables are provided in the Web Appendix.

5.3 Forecast comparison with SPF

Section 2 discusses how the probabilistic forecasts for both GDP growth and inflation from the ensemble of SPF survey respondents are reported as the distributions of ‘annual average over annual average’ values. We track the probabilistic forecasts for the 2008-2009 value of GDP growth (labeled $\overline{R}_{\text{GDP},08-09}$ in Section 2) and the 2009-2010 value of inflation (labeled $\overline{R}_{\text{INF},09-10}$) over the quarterly reports in which predictive distributions were provided for each value. These are given in the bottom sections of Tables 3 and 4, along with the date of each SPF report and the number of respondents from which the ensemble is constructed. Also reported are some summary statistics (mean, standard deviation, skew and kurtosis) computed from the histogram distributions, along with the CRPS calculated using the outturns for both annual values.

As discussed in Section 2, the information set changes with each report. For example, for the first report in Table 3 dated February 2008, no quarterly GDP growth figures that contribute to $\overline{R}_{\text{GDP},08-09}$ are known. However, by the last report, dated November 2009, national accounts data are available up to 2009:Q3 and the only forecast required is one-quarter ahead—the nowcast. We note that the reports that provide the predictive distribution of $\overline{R}_{\text{GDP},08-09}$ straddle the onset of the Great Recession, while those that detail the predictive distribution of $\overline{R}_{\text{INF},09-10}$ record inflation expectations during a period when the U.S. Federal Reserve adopted record low interest rates and embarked on its first round of quantitative easing.

There are three striking features of these SPF forecasts during the Great Recession. First, the predictive distributions are usually both asymmetric and heavy-tailed. For example, in Table 3 the report dated May 2009 gives a distribution for annual GDP growth with a skew of 2.18 and a kurtosis of 11.29. While in Table 4 the distribution for inflation in the report dated August 2009 has skew 1.13 and kurtosis 6.42. Previous research noting that the SPF forecast densities can be non-Gaussian include Croushore (1993) and Andrade, Ghysels and Idler (2014). Second, in Table 3, the professional forecasters adjust their views on GDP growth very rapidly at the start of 2009. In the third quarter of 2008 the mode of GDP

growth is 1 to 2 percent (mean 1.61 percent), but by the first quarter of 2009, the mode drops to -2 to -1 percent (mean -1.37 percent). A third striking feature is that despite the rapid adjustment to macroeconomic news, the professional forecasters do make mistakes. The outturn of $\bar{R}_{\text{GDP},08-09}$ is -2.46 percent. But by the last quarter of 2009, the probability of the appropriate bin is only 71.5%. For inflation, with an outturn value of 0.95%, in the final report dated November 2010 in Table 4, there is a probability of 12.6% of inflation greater than 2 percent, which is impossible in the absence of anything other than an extreme price shock during the fourth quarter of 2010. For both variables, the ensemble of forecast densities lack sharpness.

To compare the SPF forecasts with those from the time series models, the top three sections of Tables 3 and 4 report the comparable predictive distributions of $\bar{R}_{\text{GDP},08-09}$ and $\bar{R}_{\text{INF},09-10}$ for our proposed copula model with skew t margins, and also the BVAR and BVAR-SV. These are constructed for the eight data vintages that match the SPF report dates. For example, forecasts made using vintage 2009:Q2 data employ the first release of national accounts for 2009:Q1, as does the SPF report dated 15 May 2009. When using the time series models to create predictive distributions of $\bar{R}_{\text{GDP},08-09}$ or $\bar{R}_{\text{INF},09-10}$ for a given vintage, we employ all known national accounts data. Any remaining GDP and price index values that are required are computed using the iterates of quarterly GDP growth and inflation produced from the time series models. The number h of these quarterly periods that require forecasting is also given in Tables 3 and 4 for each data vintage employed.

A number of substantial differences can be observed when comparing the forecast distributions (and their summary statistics) from the time series models against each other, and also with those provided by the SPF. First, the forecast distributions produced using the BVAR exhibit a kurtosis close to 3 throughout. This is because the underlying predictive distributions of quarterly GDP growth and inflation are assumed Gaussian. In contrast, the forecast distributions of both other time series models exhibit excess kurtosis throughout, as do those provided by the SPF. Second, both the BVAR and BVAR-SV produce forecast distributions that are close to symmetric. Again, this is because the predictive distributions of underlying quarterly GDP growth and inflation are also symmetric (except for very minor

perturbations due to integration over the parameter space). In comparison, the forecast distributions from the copula model exhibit substantial skew, particularly for inflation in Table 4. In the latter case, this positive skew is consistent with that found in the SPF reports. Third, both the BVAR and BVAR-SV models provide very high probabilities of deflation (up to 58.7%) in Table 4 using data vintages 2009:Q1 to 2010:Q2. This is at odds with the SPF, which never report more than a 8.4% chance of deflation. In comparison, the copula model produces low probabilities of deflation, and its predictive distributions are both sharp and located around the ultimate outturn value of 0.95%. Last, we observe in Table 3 that the SPF reacts to the onset of the Great Recession one quarter earlier than the statistical models. This is likely due to the availability of higher frequency data in the first half of 2009:Q1, which is not utilized by the statistical models.

Finally, Table 5 extends the comparison between the SPF probabilistic forecasts, and those from the time series models, to cover all reports and matching data vintages between 1992:Q1 and 2009:Q4. We employ the SPF discrete ensemble distributions as density forecasts, and their means as the SPF point forecasts. The same three forecast metrics are employed, with the outturn values computed using the data vintage that had the second release of the final GDP and price index values required to compute $\bar{R}_{\text{GDP}}$ and $\bar{R}_{\text{INF}}$. Each SPF report provides forecasts for the current year Y and the next year $Y + 1$. For example, the May 2002 report provides forecasts for both $\bar{R}_{\text{GDP},01-02}$ and $\bar{R}_{\text{GDP},02-03}$. Notice that the earlier of these two values requires forecasting of GDP growth for $h = 3$ quarters (2002:Q2, 2002:Q3 and 2002:Q4), whereas the latter value requires forecasts for $h = 7$ quarters (2002:Q2 to 2003:Q4). Therefore, in Table 5 we report mean metric values broken down by the number of quarterly GDP growth (inflation) values that need forecasting, and those that are known.

We observe that at longer horizons the SPF respondents produce more accurate forecasts than the time series models for both variables. This is particularly the case with GDP growth, where their forecasts are the most accurate by all measures whenever $h \geq 3$ quarterly forecasts are required. However, the time series models are more accurate than the SPF at short horizons. This is particularly the case for inflation, where the distributions reported by

the SPF have poor tail estimates. For example, all time series models have a lower TW-CRPS metric than the SPF reports for inflation measures with $h = 1, 2, 3$ or 4 forecast quarters (i.e. forecasts for the current year Y). The poor performance of the SPF forecasts at short horizons is at least in part due to the discrete nature of the ensemble forecast distributions. At short horizons the forecast distributions are tight and the discretization into bins of width one percent is very coarse. Whereas, at longer horizons the forecast distributions are more diffuse and discretization has much less of an effect.

Croushore (1993) and Lahiri and Wang (2013) provide discussions of the relative accuracy of SPF forecasts over short and long horizons. Hafer and Hein (1983) and Faust and Wright (2013) argue that SPF forecasts for inflation represent a difficult benchmark for time series models to beat, as do Kruger, Clark and Ravazzolo (2014) for nowcasting of both quarterly inflation and GDP growth. However, these authors use the quarterly point forecasts provided by the SPF, which are the medians of the survey respondents' forecasts of the level of each variable. These are different data than the SPF ensemble density forecasts that we employ, and do not suffer from the discretization issues discussed above.

6 Discussion

The aim of this paper is to demonstrate the potential for copula-based multivariate time series models to improve density forecasts in empirical macroeconomics. All four of our macroeconomic variables have marginal distributions that are far from Gaussian, and exhibit strong skew and/or heavy tails throughout almost all the vintages of our real-time dataset. Copula models allow these to be modeled directly and separately from the dependence structure. While nonparametric models for the margins are popular in the copula literature (eg. Shih and Louis 1995), tails of distributions can be poorly estimated when the sample sizes are low, which is the case in many macroeconomic studies. Skew t distributions provide a smooth parametric alternative, although other distributions are also feasible.

By restricting our attention to the Gaussian copula we assume tail dependence is zero. While this is unrealistic for the modeling of high frequency financial returns (eg. see Pat-

ton 2006), it is not necessarily as restrictive in the analysis of low frequency macroeconomic time series. On the other hand, the increased parsimony and analytical tractability represent advantages, as does the equivalence with a latent Gaussian VAR (Biller and Nelson 2003; Biller 2009). With Gaussian margins, the copula model specializes to the familiar Gaussian VAR. But with one or more non-Gaussian margins, the copula model is a nonlinear multivariate time series model. Using the 2011:Q2 vintage data, this is starkly illustrated in Figure 4 by the nonlinear relationship between unemployment and output growth lagged three quarters.

By selecting time invariant margins, the time series model for the macroeconomic variables is stationary. This follows from a property in the Web Appendix. By further assuming the series also has a Markov order of four quarters, the dependence structure is captured parsimoniously using the partial correlations. Moreover, using the vine representation of the Gaussian copula, the Bayesian selection methodology of Smith (2014) allows for further parsimony to be exploited. By using the parallel algorithm outlined in the Appendix to evaluate the likelihood, estimation is also computationally feasible in practice. For example, for the 2011:Q2 vintage data we completed 1000 sweeps of the Bayesian sampling scheme in under 8 minutes using FORTRAN code on a standard desktop with 8 threads.

In our forecasting study, the copula model with skew t margins is competitive with the BVAR in terms of improved values of a number of performance metrics, although these were only statistically significant in a few cases. However, during the Great Recession the asymmetric and heavy-tailed predictive distributions produced avoid a number of pitfalls of the symmetric forecasts produced by both the BVAR and the BVAR-SV, and were much more consistent with those produced by the SPF.

Clark (2011) and Clark and Ravazzolo (2014) show that the inclusion of stochastic volatility improves density forecast accuracy for low-dimensional models. We find the BVAR-SV model outperforms the copula model overall. In theory, such persistence in the second moment can also be incorporated into a copula model by the choice of a suitable copula function. For example, one approach would be to employ a more general vine copula as in Smith (2014), but with component pair-copulas that allow for dependence in higher order

moments. There is no current research on this topic of which we are aware, although such copula models would have great potential, and present an opportunity for future research.

References

- Aas, K., C. Czado, A. Frigessi & H. Bakken, (2009). ‘Pair-copula Constructions of Multiple Dependence’, *Insurance: Mathematics and Economics*, 44, 2, 182–198.
- Andrade, P, E. Ghysels, & J. Idler, (2014). ‘Inflation Risk Measures and Their Informational Content’, <http://dx.doi.org/10.2139/ssrn.2439607>.
- Azzalini, A. & A. Capitanio, (2003). ‘Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t Distribution’, *Journal of the Royal Statistical Society, Series B*, 65, 367–389.
- Banbura, M., D. Giannone & L. Reichlin, (2010). ‘Large Bayesian Vector Auto Regressions’, *Journal of Applied Econometrics*, 25, 1, 71–92.
- Beare, B. K., (2010). ‘Copulas and Temporal Dependence’, *Econometrica*, 78, 395–410.
- Bedford, T. & R. Cooke, (2002). ‘Vines - a New Graphical Model for Dependent Random Variables’, *Annals of Statistics*, 30, 1031–1068.
- Biller, B., (2009). ‘Copula-Based Multivariate Input Models for Stochastic Simulation’, *Operations Research*, 57, 4, 878–892.
- Biller, B. & B. Nelson, (2003). ‘Modeling and Generating Multivariate Time-Series Input Processes Using a Vector Autoregressive Technique’, *ACM Transactions on Modeling and Computer Simulation*, 13, 3, 211–237.
- Chen, X. & Y. Fan, (2006). ‘Estimation and Model Selection of Semiparametric Copula-based Multivariate Dynamic Models under Copula Misspecification’, *Journal of Econometrics*, 135, 125–154.
- Clark, T. E., (2011). ‘Real-Time Density Forecasts From Bayesian Vector Autoregressions with Stochastic Volatility’, *Journal of Business and Economic Statistics*, 29, 327–341.
- Clark, T. E. & F. Ravazzolo, (2014). ‘Macroeconomic Forecasting Performance under Alternative Specifications of Time-varying Volatility’, *Journal of Applied Econometrics*, doi: 10.1002/jae.2379
- Cogley, T., S. Morozov & T. Sargent, (2005). ‘Bayesian Fan Charts for U.K. inflation: Forecasting and Sources of Uncertainty in an Evolving Monetary System’, *Journal of Economic Dynamics and Control* 29(11), 1893–1925.
- Corradi, V., A. Fernandez & N. R. Swanson, (2009). ‘Information in the Revision Process of Real-Time Data Sets’, *Journal of Business and Economic Statistics*, 27, 455–467.

- Croushore, D., (1993). ‘Introducing: the Survey of Professional Forecasters’, *Business Review*, Federal Reserve Bank of Philadelphia, Issue November, 3–15.
- Croushore, D. & T. Stark, (2001). ‘A Real-time Data Set for Macroeconomists’, *Journal of Econometrics*, 105, 111–130.
- D’Agostino, A., L. Gambetti & D. Giannone, (2013). ‘Macroeconomic Forecasting and Structural Change’, *Journal of Applied Econometrics*, 28, 1, 82–101.
- Daniels, M. & M. Pourahmadi, (2009). ‘Modeling Covariance Matrices via Partial Autocorrelations’, *Journal of Multivariate Analysis*, 100, 2352–2363.
- Diebold, F. X & R.S. Mariano, (1995), ‘Comparing Predictive Accuracy’, *Journal of Business and Economic Statistics*, 13, 253–263.
- Domma, F., S. Giordano, & P. Perri, (2009). ‘Statistical Modeling of Temporal Dependence in Financial Data via a Copula Function’, *Communications in Statistics*, 38, 703–728.
- De Mol C., D. Giannone, & L. Reichlin, (2008). ‘Forecasting using a Large Number of Predictors: is Bayesian Shrinkage a Valid Alternative to Principal Components?’, *Journal of Econometrics*, 146, 318–328.
- Faust, J. & J. H. Wright, (2009). ‘Comparing Greenbook and Reduced Form Forecasts using a Large Real-time Dataset’, *Journal of Business and Economic Statistics*, 27, 468–479.
- Faust, J. & J. H. Wright, (2013). ‘Forecasting Inflation’, *Handbook of Forecasting*, forthcoming.
- George, E., Sun, D. & S. Ni, (2008). ‘Bayesian Stochastic Search for VAR Model Restrictions’, *Journal of Econometrics*, 142, 553–580.
- Gneiting, T. & A. E. Raftery, (2007). ‘Strictly Proper Scoring Rules, Prediction and Estimation’, *Journal of the American Statistical Association*, 102, 359–378.
- Gneiting, T. & R. Ranjan, (2011). ‘Comparing Density Forecasts using Threshold- and Quantile-Weighted Scoring Rules’, *Journal of Business and Economic Statistics*, 29, 3, 411–422.
- Hafer, R.W. & S.E. Hein, (1983). ‘On the Accuracy of Time Series, Interest Rate and Survey Forecasts of Inflation’, *Journal of Business*, 58, 377–398.
- Harvey, D., S. Leybourne & P. Newbold, (1997). ‘Testing the Equality of Prediction Mean Squared Errors’, *International Journal of Forecasting*, 13, 281–291.
- Jochmann, M., G. Koop, R. Leon-Gonzalez & R. W. Strachan, (2013). ‘Stochastic Search Variable Selection in Vector Error Correction Models with an Application to a Model of the UK Macroeconomy’, *Journal of Applied Econometrics*, 28, 1, 62–81.
- Joe, H., (2005). ‘Asymptotic Efficiency of the Two-stage Estimation Method for Copula-based Models’, *Journal of Multivariate Analysis*, 94, 401–419.

- Joe, H., (2006). 'Generating Random Correlation Matrices Based on Partial Correlations', *Journal of Multivariate Analysis*, 71, 2177–2189.
- Jore, A.S., J. Mitchell & S. P. Vahey, (2010). 'Combining Forecast Densities from VARs with Uncertain Instabilities', *Journal of Applied Econometrics*, 25, 621–634.
- Koop, G. & D. Korobilis, (2010). *Bayesian Multivariate Time Series Methods for Empirical Macroeconomists, Foundations and Trends in Econometrics*, 3, 4, 267–358.
- Koop, G. & D. Korobilis, (2013). 'Large Time-Varying Parameter VARs', *Journal of Econometrics*, 177, 185–198.
- Korobilis, D., (2013). 'VAR Forecasting using Bayesian Variable Selection', *Journal of Applied Econometrics*, 28, 2, 204–230.
- Kruger, F., T. Clark, & F. Ravazzolo, (2014), 'Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts', working paper.
- Lahiri, K. & J.G. Wang, (2013). 'Evaluating Probability Forecasts for GDP Decline using Alternative Methodologies', *International Journal of Forecasting*, 29, 1, 175–190.
- Lütkepohl, H., (2006). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag: Berlin.
- McNeil A.J., R. Frey & R. Embrechts, (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton University Press: Princeton, NJ.
- Nelsen, R., (2006). *An Introduction to Copulas*. 2nd ed., New York: NY: Springer.
- Panagiotelis, A. & M. Smith, (2008). 'Bayesian Density Forecasting of Intraday Electricity Prices using Multivariate Skew t Distributions', *International Journal of Forecasting*, 24, 710–725.
- Patton, A., (2006). 'Modelling Asymmetric Exchange Rate Dependence', *International Economic Review*, 47, 527–556.
- Patton, A., (2012). 'A Review of Copula Models for Economic Time Series', *Journal of Multivariate Analysis*, 110, 4–18.
- Pitt, M., D. Chan & R. Kohn, (2006). 'Efficient Bayesian Inference for Gaussian Copula Regression Models', *Biometrika*, 93, 537–554.
- Primiceri, G. E., (2005). 'Time Varying Structural Vector Autoregressions and Monetary Policy', *Review of Economic Studies*, 72, 821–852.
- Ravazzolo, F. & S. P. Vahey, (2014). 'Forecast densities for economic aggregates from disaggregate ensembles', *Studies in Nonlinear Dynamics & Econometrics*, 18(4), 367–381.
- Rodriguez, J., (2007). 'Measuring Financial Contagion: A Copula Approach', *Journal of Empirical Finance*, 14, 401–423.

- Romer, C.D. & D.H. Romer, (2000). ‘Federal Reserve Information and the Behavior of Interest Rates’, *American Economic Review*, 90, 429–457.
- Scott, J.G. & J.O. Berger, (2010). ‘Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-selection Problem’, *Annals of Statistics*, 38, 5, 2587–2619.
- Shapiro, S. & M. Wilk, (1965). ‘An Analysis of Variate Test for Normality (Complete Samples)’, *Biometrika*, 52, 3-4, 591–611.
- Shih, J.H. & T.A. Louis, (1995). ‘Inferences on the Association Parameter in Copula Models for Bivariate Survival Data’, *Biometrics*, 51, 4, 1384–99.
- Smith, M. S., (2014). ‘Copula Modelling of Dependence in Multivariate Time Series’, *International Journal of Forecasting*, in press.
- Smith, M. & R. Kohn, (2002). ‘Parsimonious Covariance Matrix Estimation for Longitudinal Data’, *Journal of the American Statistical Association*, 91, 460, 1141–1153.
- Smith, M., A. Min, C. Almeida & C. Czado, (2010). ‘Modeling Longitudinal Data using a Pair-Copula Decomposition of Serial Dependence’, *Journal of the American Statistical Association*, 105, 492, 1467–1479.
- Song, P., (2000). ‘Multivariate Dispersion Models Generated from Gaussian Copula’, *Scandinavian Journal of Statistics*, 27, 305–320.
- Wallis, K.F., (2005). ‘Combining Density and Interval Forecasts: a Modest Proposal’, *Oxford Bulletin of Economics and Statistics*, 67, 983–994.
- Walls, W.D., (2006). ‘Modelling Heavy Tails and Skewness in Film Returns’, *Applied Financial Economics*, 15, 17, 1181–1188.

Appendix

Part A

We outline how to evaluate the likelihood $\mathcal{L}$ at Equation (4.1) for the copula, given the copula data $\mathbf{u} = (\mathbf{u}'_1, \dots, \mathbf{u}'_T)'$. This involves evaluating the functionals

$$\begin{aligned} K_{t,s}(\mathbf{u}_s, \dots, \mathbf{u}_t; \boldsymbol{\phi}(t-s)) &= \prod_{i=a(t)}^{b(t)} \prod_{j=a(s)}^{b(s)} c_{i,j}(u_{i|j+1}, u_{j|i-1}; \phi_{i,j}) \\ K_{t,t}(\mathbf{u}_t; \boldsymbol{\phi}(0)) &= \prod_{i=a(t)}^{b(t)} \prod_{j=a(t)}^{i-1} c_{i,j}(u_{i|j+1}, u_{j|i-1}; \phi_{i,j}), \end{aligned}$$

where $a(\cdot)$ and $b(\cdot)$ are defined in Section 3.2. For each $K_{t,s}$, the component pair-copulas have parameters equal to the partial correlations $\phi_{i,j}$. For our stationary series, these are elements of $\boldsymbol{\phi}(t-s)$, and are also denoted as ϕ_{l_1, l_2}^k , with $k = t-s$, $l_1 = i - m(t-1)$ and $l_2 = j - m(s-1)$. The pair-copulas themselves are bivariate Gaussian copulas, each of which has density

$$c_{Ga}(u, v; \rho) = (1 - \rho^2)^{-1/2} \exp \left\{ \frac{-\rho^2(\tilde{u}^2 + \tilde{v}^2) - 2\rho\tilde{u}\tilde{v}}{2(1 - \rho^2)} \right\},$$

where $\tilde{u} = \Phi^{-1}(u)$ and $\tilde{v} = \Phi^{-1}(v)$. The arguments of the pair-copula are computed using the algorithm given in Smith (2014), and given below for the Gaussian copula case. It is based on the algorithm of Aas, Czado, Frigessi and Bakken (2009), but has been adjusted to account for the high number of independence pair-copulas in our copula model, and also with the computations re-ordered so that they can be undertaken in parallel.

The algorithm involves repeated evaluations of the conditional bivariate Gaussian copula function, given by

$$h(u|v; \rho) = \Phi \left(\frac{\tilde{u} - \rho\tilde{v}}{(1 - \rho^2)^{1/2}} \right),$$

where $h(u|v; \rho = 0) = u$. The following algorithm evaluates the required arguments $u_{i|j}$ and $u_{j|i}$ ($i > j$) of all pair-copula densities required to compute the likelihood $\mathcal{L}$.

Algorithm

For $r = 1, \dots, (p+1)m - 1$:

For $i = r + 1, \dots, N$ (in parallel):

Step 1. Set $j = i - r$, $s = \lceil j/m \rceil$, $t = \lceil i/m \rceil$, $k = t - s$, $l_1 = i - m(t - 1)$, $l_2 = j - m(s - 1)$.

Step 2. If $k > p$ or $\gamma_{l_1, l_2}^k = 0$ then

Step 2.1. $u_{i|j} = u_{i|j+1}$ and $u_{j|i} = u_{j|i-1}$,

else

Step 2.2. $u_{i|j} = h(u_{i|j+1}|u_{j|i-1}; \phi_{l_1, l_2}^k)$ and $u_{j|i} = h(u_{j|i-1}|u_{i|j+1}; \phi_{l_1, l_2}^k)$.

In the above, $\lceil x \rceil$ is the integer ceiling function evaluated at x . The values $u_{i|j}$ and $u_{j|i}$ are stored in the upper and lower triangular elements of an $N \times N$ banded matrix $\mathcal{M}$. The algorithm fills in these elements in the upper and lower diagonal bands, from band $r = 1$ to band $r = (p + 1)m - 1$. The computations in the inner loop can be undertaken in parallel on up to $N - r$ different nodes.

Part B

To compute the CRPS and its tail-weighted version (denoted TW-CRPS) we follow Gneiting and Ranjan (2011) and express both metrics in terms of the quantile score, so that for forecast density $f(y)$ with observation y_{obs} , the metrics are given by the score function

$$S(f, y_{\text{obs}}, v(\alpha)) = \int_0^1 \text{QS}_\alpha(F(\alpha)^{-1}, y) v(\alpha) d\alpha.$$

Here, $\text{QS}_\alpha(q, y_{\text{obs}}) = 2(\mathbb{I}(y_{\text{obs}} < q) - \alpha)(q - y_{\text{obs}})$ is the quantile score for forecast quantile q at the level $0 < \alpha < 1$, F^{-1} is the inverse distribution function of the forecast density and $v(\alpha)$ is a weighting function. The CRPS results from using a uniform weighting function where $\text{CRPS}(f, y_{\text{obs}}) = S(f, y_{\text{obs}}, 1)$, and the tail-weighted CRPS value results from using a quadratic weighting function where $\text{TW-CRPS}(f, y_{\text{obs}}) = S(f, y_{\text{obs}}, (2\alpha - 1)^2)$. To compute these metrics in practise using the (large) Monte Carlo samples from the forecast distributions, we approximate F using the EDF, and compute the integral above numerically. In experiments we have found this to be a more accurate method of computation than expressing the CRPS as a difference of expectations, which are each approximated using Monte Carlo iterates.

Figure 1: The left-hand panels plot the 2011:Q2 vintage data, where the series are (a) quarterly real GDP growth in percent, (c) quarterly inflation in percent, (e) unemployment rate in percent and (g) short-term interest rate in percent. The copula data (PIT values), computed using the fitted skew t marginal distributions, are on the right-hand side panels for the corresponding series.

Figure 2: Densities of marginal models fit to each of the four variables. The lefthand panels (a), (c), (e) and (g) each show three densities fit to the 2011:Q2 vintage data. These are a Gaussian (thin black line), a skew t (blue line) and a nonparametric kernel density (red line). The righthand panels (b), (d), (f) and (h) each show four skew t densities fit to 1975:Q1 (dashed red line), 1987:Q2 (red line), 1999:Q4 (dashed black line) and 2011:Q2 (black line) vintage data.

Figure 3: The posterior means of the matrices of Spearman unconditional pairwise correlations $R(k)$, for $k = 0, 1, \dots, 4$, from the copula model with skew t margins fit to the 2011:Q2 vintage data. The four variables are GDP growth (Grth), inflation (Infln), the unemployment rate (UR) and the interest rate (IR). The value in row i and column j of $R(k)$ is the Spearman correlation between $Y_{i,t}$ and $Y_{j,t-k}$. For example, in panel (e) GDP growth at time $t - 4$ and the unemployment rate at time t have a Spearman correlation of -0.28.

Figure 4: Contour plots of two bivariate margins of the multivariate copula model with skew t univariate margins fitted to the vintage 2011:Q2 data. Panel (a) shows the bivariate margin $f(Y_{1,t-3}, Y_{3,t})$ of GDP growth lagged 3 quarters and the unemployment rate, and panel (b) shows the bivariate margin $f(Y_{2,t-2}, Y_{1,t})$ of inflation lagged 2 quarters and GDP growth. The contours are constructed using bivariate kernel density estimates from large Monte Carlo samples simulated from the fitted parametric model. The contours are unlabeled for visual clarity, but show unimodal densities that are far from Gaussian or elliptical. Estimates of the conditional means $E(Y_{3,t}|Y_{1,t-1})$ and $E(Y_{1,t}|Y_{2,t-2})$ are given as thick red lines, and are constructed using smoothing splines fit to these Monte Carlo iterates. Strong nonlinear bivariate relationships are observed. Scatterplots of the data are overlaid to show that the bivariate margins of the fitted model are consistent with the observed data.

Figure 5: Forecast densities of inflation from using vintage 2009:Q2 data with forecast origin 2009:Q1. Panel (a) is for $h = 1$ quarter ahead (i.e. the nowcast), panel (b) for $h = 2$ quarters ahead (2009:Q3), panel (c) for $h = 4$ quarters ahead (2010:Q1), and panel (d) for $h = 8$ quarters ahead (2011:Q1). Forecast densities are from the BVAR (blue line), BVAR-SV (black dashed line), Coupla & EDF margins (blue dashed line) and Copula & Skew t margins (black line). The outturn values for each quarter are also reported.

Figure 6: Fan charts of density forecasts for the short term interest rate, with each panel corresponding to a different forecasting model. Predictions are made with forecast origins at 2007:Q3 to 2010:Q4 (vintages 2007:Q4 to 2011:Q1). The thick lines are the predictive medians between 1 and 8 quarters ahead, and the thin lines are the 10th and 90th percentiles of the predictive distributions. The forecast origins are plotted as circles, which are also the outturn values as this variable is not subject to revision in later vintages. The fans for forecasts made at two particular periods (vintages 2007:Q4 and 2009:Q4) are plotted in black to distinguish them, and discussed in the text.

Full Period: 1975:Q1 – 2011:Q2								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
	GDP Growth				Inflation			
<i>Time Series Model</i>								
BVAR	0.786	0.771	0.758	0.725	0.317	0.365	0.412	0.580
Copula & EDF	0.767	0.766	0.753	0.729	0.316	0.371	0.425	0.566
Copula & Skew t	0.767	0.754	0.769	0.740	0.311	0.359	0.403	0.542
BVAR-SV	0.736**	0.716*	0.720	0.732	0.311	0.352	0.415	0.593
	Unemployment				Interest Rate			
<i>Time Series Model</i>								
BVAR	0.287	0.526	0.967	1.394	0.818	1.272	1.789	2.645
Copula & EDF	0.327**	0.589	1.039	1.529	0.847	1.282	1.785	2.658
Copula & Skew t	0.288	0.512	0.966	1.517	0.759	1.282	1.807	2.701
BVAR-SV	0.282	0.505	0.918	1.386	0.789**	1.211**	1.729	2.585
Three Recessions: 1982:Q1–1982:Q2, 1991:Q1–1991:Q3 and 2008:Q4–2009:Q3								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
	GDP Growth				Inflation			
<i>Time Series Model</i>								
BVAR	1.474	1.592	1.477	1.417	0.510	0.647	0.764	0.829
Copula & EDF	1.473	1.704	1.536	1.461	0.406	0.515	0.628	0.616
Copula & Skew t	1.326	1.622	1.586	1.466	0.491	0.603	0.652	0.594
BVAR-SV	1.330	1.444	1.456	1.408	0.511	0.643	0.751	0.756
	Unemployment				Interest Rate			
<i>Time Series Model</i>								
BVAR	0.520	1.255	2.314	2.372	0.675	1.052	1.822	3.290
Copula & EDF	0.642	1.461	2.550	2.744	1.149	1.321	1.803	3.510
Copula & Skew t	0.407	1.173	2.330	2.658	0.939	0.970	1.693	3.766
BVAR-SV	0.500	1.192	2.256	2.512	0.714	0.966	1.645	3.302

Table 1: Mean values of RMSE computed over two evaluation periods. Separate results are reported for forecasts made $h = 1, 2, 4$ and 8 quarters ahead in columns, and each row corresponds to results from a different model. As a rough guide, results of a Harvey, Leybourne and Newbold (1997) small-sample adjustment of the Diebold and Mariano (1995) test are reported for the first evaluation period (the second has only 9 quarters). This is a two-sided test of the equality of the mean score of the BVAR with each of the other three models. For these three models, rejection of the null hypothesis of equal means at the 10% or 5% level is denoted with ‘*’ or ‘**’.

Full Period: 1975:Q1 – 2011:Q2								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
GDP Growth				Inflation				
<i>Time Series Model</i>								
BVAR	0.431	0.427	0.420	0.407	0.176	0.196	0.224	0.317
Copula & EDF	0.409	0.415	0.412	0.404	0.181	0.204	0.233	0.314
Copula & Skew t	0.411	0.414	0.421	0.412	0.174	0.189	0.216	0.289
BVAR-SV	0.387**	0.386**	0.391**	0.396	0.171	0.187	0.217	0.308
Unemployment				Interest Rate				
<i>Time Series Model</i>								
BVAR	0.152	0.274	0.492	0.699	0.410	0.675	0.985	1.529
Copula & EDF	0.180**	0.317*	0.547	0.814	0.398	0.658	0.986	1.564
Copula & Skew t	0.153	0.271	0.505	0.801	0.364**	0.643**	0.961	1.544
BVAR-SV	0.142**	0.253**	0.460	0.711	0.361**	0.616**	0.977	1.573
Three Recessions: 1982:Q1–1982:Q2, 1991:Q1–1991:Q3 and 2008:Q4–2009:Q3								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
GDP Growth				Inflation				
<i>Time Series Model</i>								
BVAR	0.892	0.985	0.894	0.851	0.303	0.348	0.441	0.470
Copula & EDF	0.910	1.109	0.975	0.912	0.237	0.283	0.372	0.354
Copula & Skew t	0.807	1.046	1.013	0.919	0.291	0.309	0.360	0.335
BVAR-SV	0.804	0.925	0.933	0.911	0.291	0.335	0.394	0.407
Unemployment				Interest Rate				
<i>Time Series Model</i>								
BVAR	0.295	0.791	1.477	1.449	0.387	0.618	1.110	2.142
Copula & EDF	0.410	1.020	1.877	1.936	0.495	0.793	1.255	2.315
Copula & Skew t	0.230	0.751	1.626	1.815	0.487	0.584	1.126	2.512
BVAR-SV	0.279	0.767	1.596	1.674	0.429	0.657	1.191	2.303

Table 2: Mean values of CRPS computed over two evaluation periods. Separate results are reported for forecasts made $h = 1, 2, 4$ and 8 quarters ahead in columns, and each row corresponds to results from a different model. As a rough guide, results are reported for a test equivalent to the Diebold and Mariano (1995) test (with the Harvey, Leybourne and Newbold (1997) small-sample adjustment) applied to the CRPS values. Test results are only reported for the first evaluation period (the second has only 9 quarters). This is a two-sided test of the equality of the mean CPRS of the BVAR with each of the other three models. For these three models, rejection of the null hypothesis of equal means at the 10% or 5% level is denoted with ‘*’ or ‘**’.

		GDP Growth Bands in Percent									Summary Statistics				
		< -2	-2-1	-1-0	0-1	1-2	2-3	3-4	4-5	>5	Mean	SD	Skew	Kurt.	CRPS
Vintage	h	Copula & Skew t													
08:Q1	8	0.2	0.6	1.8	4.6	11.9	20.3	25.1	18.6	16.9	3.42	1.71	0.03	3.91	4.94
08:Q2	7	0.1	0.3	0.8	3.2	8.5	18.3	24.2	22.5	22.0	3.79	1.64	0.03	3.63	5.34
08:Q3	6	0.0	0.1	0.4	2.0	6.6	15.5	25.1	25.1	25.1	4.04	1.59	0.23	3.70	5.61
08:Q4	5	0.1	0.3	1.3	4.7	13.4	22.5	26.1	19.4	12.2	3.28	1.55	-0.01	3.74	4.87
09:Q1	4	0.0	0.6	2.7	11.3	25.7	31.0	19.1	7.2	2.4	2.32	1.30	0.18	3.65	4.06
09:Q2	3	21.6	43.4	29.0	5.1	0.8	0.1	0.0	0.0	0.0	-1.34	0.92	-0.17	4.03	0.72
09:Q3	2	82.4	17.1	0.5	0.0	0.0	0.0	0.0	0.0	0.0	-2.45	0.53	-0.15	4.13	0.11
09:Q4	1	94.3	5.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-2.34	0.21	0.45	4.99	0.07
Vintage	h	BVAR													
08:Q1	8	0.0	0.4	1.7	4.6	11.7	20.5	24.0	20.2	16.8	3.45	1.63	-0.05	3.03	4.99
08:Q2	7	0.0	0.1	0.4	2.3	7.5	17.4	25.2	24.5	22.7	3.90	1.52	0.09	3.11	5.51
08:Q3	6	0.0	0.1	0.3	1.3	5.2	13.8	24.2	25.8	29.3	4.19	1.48	-0.04	3.17	5.82
08:Q4	5	0.0	0.0	0.2	1.7	5.3	15.0	25.1	26.5	26.2	4.08	1.43	-0.01	3.06	5.74
09:Q1	4	0.1	0.5	2.4	9.9	22.3	30.5	22.4	9.8	2.3	2.48	1.30	-0.12	3.02	4.20
09:Q2	3	9.1	30.3	38.5	18.2	3.5	0.4	0.0	0.0	0.0	-0.73	0.98	0.05	3.06	1.21
09:Q3	2	57.6	40.0	2.4	0.0	0.0	0.0	0.0	0.0	0.0	-2.11	0.56	-0.05	3.00	0.22
09:Q4	1	81.2	18.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-2.21	0.23	0.07	2.89	0.15
Vintage	h	BVAR-SV													
08:Q1	8	0.0	0.2	0.7	2.8	10.1	26.0	33.6	18.5	8.1	3.30	1.28	0.12	4.59	5.05
08:Q2	7	0.1	0.1	0.8	2.6	8.6	22.2	31.8	22.4	11.4	3.48	1.36	0.06	4.08	5.20
08:Q3	6	0.0	0.1	0.5	1.9	6.8	20.4	32.7	24.7	12.8	3.62	1.30	-0.04	3.88	5.36
08:Q4	5	0.2	0.4	1.3	4.1	13.2	26.0	28.8	17.0	9.1	3.16	1.44	-0.01	3.93	4.83
09:Q1	4	2.0	3.7	10.7	21.3	26.2	20.5	9.6	4.0	2.0	1.48	1.63	-0.04	3.91	3.06
09:Q2	3	38.9	28.7	19.2	8.9	3.1	0.8	0.3	0.1	0.0	-1.60	1.50	0.07	3.87	0.52
09:Q3	2	70.7	24.7	4.1	0.5	0.0	0.0	0.0	0.0	0.0	-2.40	0.82	0.08	4.03	0.18
09:Q4	1	88.2	11.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-2.33	0.29	0.05	4.29	0.09
Date	#R	SPF Reports													
Feb-08	44	0.4	1.0	2.5	7.2	23.9	40.9	18.6	3.9	1.5	2.29	1.24	-0.42	4.35	4.08
May-08	44	0.4	0.9	2.7	12.8	31.4	37.6	11.1	2.3	0.8	1.98	1.18	-0.23	4.07	3.79
Aug-08	43	0.7	1.0	3.2	18.6	43.1	26.0	6.1	0.9	0.3	1.61	1.11	-0.27	4.51	3.47
Nov-08	47	3.1	16.4	34.5	24.0	12.9	7.2	1.9	0.1	0.0	0.07	1.33	0.45	2.97	1.79
Feb-09	39	33.8	36.0	19.2	7.2	3.0	0.7	0.1	0.1	0.1	-1.37	1.16	1.07	4.72	0.56
May-09	46	45.9	19.9	7.4	2.0	0.7	0.2	0.1	0.1	0.1	-1.89	1.01	2.18	11.29	0.24
Aug-09	33	57.9	18.3	8.5	2.8	0.8	0.2	0.2	0.1	0.2	-1.92	1.07	2.44	12.61	0.22
Nov-09	40	71.5	13.2	5.0	1.0	0.8	3.8	0.8	0.0	0.1	-1.95	1.29	2.60	9.84	0.17

Table 3: Predictive distributions for $\bar{R}_{\text{GDP},08-09}$, which is the growth in annual average GDP from 2008 to 2009, as defined by the SPF. In each row, probabilities are reported for different output growth bands (columns), along with the first four moments of the distribution and the CRPS using the outturn value of -2.46%. Distributions are given for three time series models (Copula with skew t margins, BVAR and BVAR-SV), along with those published by the SPF. For the time series models, different rows correspond to prediction using different vintage datasets, and the number of quarters (h) of GDP growth that need to be forecast to compute $\bar{R}_{\text{GDP},08-09}$ is also noted. For the SPF predictions, the date of the report is given, along with the number of survey respondents (#R).

		Inflation Bands in Percent										Summary Statistics				
		<0	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	>8	Mean	SD	Skew	Kurt.	CRPS
Vintage	h	Copula & Skew t														
09:Q1	8	4.4	50.4	39.0	5.4	0.6	0.1	0.0	0.0	0.0	0.0	0.98	0.64	0.69	5.15	0.13
09:Q2	7	0.1	16.3	57.8	21.4	3.8	0.6	0.0	0.0	0.0	0.0	1.64	0.70	0.90	4.39	0.37
09:Q3	6	1.9	54.6	40.8	2.5	0.1	0.0	0.0	0.0	0.0	0.0	0.96	0.50	0.51	4.40	0.11
09:Q4	5	2.0	73.3	24.3	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0.76	0.38	0.24	3.85	0.13
10:Q1	4	0.3	79.0	20.5	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.78	0.30	0.55	4.46	0.12
10:Q2	3	0.0	92.4	7.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.75	0.17	0.51	4.21	0.13
10:Q3	2	0.0	77.4	22.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.93	0.12	0.75	4.31	0.03
10:Q4	1	0.0	2.3	97.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.08	0.05	0.90	4.33	0.10
Vintage	h	BVAR														
09:Q1	8	25.5	27.0	25.9	15.0	5.4	1.2	0.1	0.0	0.0	0.0	0.90	1.39	-0.03	3.04	0.32
09:Q2	7	17.3	25.4	28.2	19.4	7.9	1.6	0.1	0.0	0.0	0.0	1.26	1.34	-0.07	2.98	0.34
09:Q3	6	58.7	27.1	11.2	2.7	0.3	0.0	0.0	0.0	0.0	0.0	-0.24	1.18	0.03	3.06	0.73
09:Q4	5	44.0	39.4	14.4	2.2	0.0	0.0	0.0	0.0	0.0	0.0	0.12	0.92	-0.02	2.96	0.49
10:Q1	4	38.1	51.5	10.1	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.19	0.64	0.00	3.13	0.48
10:Q2	3	9.0	81.2	9.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.51	0.38	0.02	3.10	0.28
10:Q3	2	0.0	72.7	27.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.88	0.20	0.01	3.15	0.06
10:Q4	1	0.0	21.3	78.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.06	0.07	-0.02	2.89	0.07
Vintage	h	BVAR-SV														
09:Q1	8	26.3	20.4	20.3	16.6	9.1	4.1	1.9	0.6	0.3	0.3	1.17	2.03	-0.09	4.63	0.45
09:Q2	7	22.9	17.6	20.1	17.9	11.7	5.5	2.6	1.0	0.5	0.3	1.44	2.11	-0.07	4.39	0.51
09:Q3	6	37.7	26.4	20.9	9.2	3.6	1.3	0.4	0.2	0.0	0.1	0.44	1.71	-0.12	5.56	0.42
09:Q4	5	30.1	34.6	23.9	8.5	2.1	0.6	0.1	0.0	0.0	0.0	0.58	1.25	0.10	4.71	0.32
10:Q1	4	24.6	50.0	21.7	3.3	0.2	0.1	0.0	0.0	0.0	0.0	0.51	0.82	0.03	4.33	0.27
10:Q2	3	7.0	78.1	14.7	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.60	0.42	-0.10	4.25	0.21
10:Q3	2	0.2	72.3	27.4	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.88	0.24	0.03	4.84	0.06
10:Q4	1	0.0	20.5	79.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.05	0.08	0.01	4.78	0.07
Date	#R	SPF Reports														
Feb-09	39	8.4	22.1	35.7	22.6	7.3	2.2	1.2	0.4	0.1	0.1	1.62	1.30	0.69	4.38	0.40
May-09	46	6.8	21.8	33.4	28.9	7.5	1.2	0.3	0.1	0.0	0.0	1.64	1.15	0.22	3.47	0.43
Aug-09	33	2.9	19.9	42.4	23.5	7.3	2.3	0.8	0.4	0.2	0.2	1.77	1.21	1.13	6.42	0.44
Nov-09	40	4.6	19.9	44.1	21.4	7.0	2.7	0.1	0.1	0.1	0.0	1.66	1.12	0.54	4.18	0.40
Feb-10	39	5.3	25.6	46.5	18.1	3.4	0.9	0.2	0.0	0.0	0.0	1.42	0.99	0.39	3.93	0.31
May-10	44	4.0	34.6	43.3	13.0	4.1	0.8	0.2	0.0	0.0	0.0	1.32	0.98	0.68	4.12	0.26
Aug-10	36	2.4	42.2	43.7	10.3	1.2	0.2	0.0	0.0	0.0	0.0	1.17	0.82	0.60	4.21	0.21
Nov-10	43	0.9	32.3	54.2	10.2	1.7	0.5	0.2	0.0	0.0	0.0	1.32	0.81	0.86	5.68	0.24

Table 4: Predictive distributions for $\bar{R}_{\text{INF},09-10}$, which is the growth in annual average prices from 2009 to 2010, as defined by the SPF. In each row, probabilities are reported for different inflation bands (columns), along with the first four moments of the distribution and the CRPS using the true outturn value of 0.95%. Distributions are given for three time series models (Copula with skew t margins, BVAR and BVAR-SV), along with those published by the SPF. For the time series models, different rows correspond to prediction using different vintage datasets, and the number of quarters (h) of inflation that need to be forecast to compute $\bar{R}_{\text{INF},09-10}$ is also noted. For the SPF predictions, the date of the report is given, along with the number of survey respondents (#R).

Forecast Year	Y	Y	Y	Y	Y+1	Y+1	Y+1	Y+1
Report/Vintage Quarter	Q4	Q3	Q2	Q1	Q4	Q3	Q2	Q1
# Quarters Observed	7	6	5	4	3	2	1	0
# Quarters Forecast	1	2	3	4	5	6	7	8
Output Growth $\overline{R}_{\text{GDP}}$								
	<i>RMSE</i>							
Copula & Skew t	0.369	0.425	0.652	1.416	1.879	2.080	2.058	2.022
BVAR	0.371	0.439	0.827	1.551	2.189	2.253	2.211	2.040
BVAR-SV	0.359	0.409	0.692	1.314	1.850	1.946	1.961	1.917
SPF	0.382	0.399	0.463	0.807	1.285	1.497	1.585	1.689
	<i>Mean CRPS</i>							
Copula & Skew t	0.172	0.241	0.378	0.722	0.997	1.087	1.088	1.082
BVAR	0.177	0.250	0.470	0.815	1.165	1.195	1.187	1.107
BVAR-SV	0.169	0.232	0.374	0.690	0.989	1.030	1.057	1.055
SPF	0.224	0.243	0.284	0.460	0.738	0.839	0.899	0.928
	<i>Mean TW-CRPS</i>							
Copula & Skew t	0.045	0.054	0.086	0.175	0.224	0.251	0.252	0.251
BVAR	0.047	0.055	0.102	0.187	0.270	0.278	0.277	0.250
BVAR-SV	0.045	0.053	0.084	0.152	0.224	0.239	0.243	0.239
SPF	0.056	0.062	0.074	0.098	0.160	0.192	0.206	0.211
Inflation $\overline{R}_{\text{INF}}$								
	<i>RMSE</i>							
Copula & Skew t	0.212	0.239	0.336	0.382	0.612	0.770	1.034	0.947
BVAR	0.211	0.237	0.323	0.397	0.666	0.792	0.959	0.913
BVAR-SV	0.210	0.228	0.297	0.388	0.619	0.662	0.845	0.807
SPF	0.256	0.310	0.401	0.438	0.603	0.745	0.854	0.855
	<i>Mean CRPS</i>							
Copula & Skew t	0.083	0.126	0.191	0.231	0.357	0.419	0.555	0.528
BVAR	0.080	0.129	0.187	0.245	0.388	0.472	0.566	0.555
BVAR-SV	0.078	0.126	0.168	0.235	0.338	0.387	0.494	0.481
SPF	0.192	0.217	0.252	0.281	0.349	0.419	0.482	0.479
	<i>Mean TW-CRPS</i>							
Copula & Skew t	0.023	0.029	0.043	0.054	0.082	0.091	0.122	0.115
BVAR	0.022	0.029	0.040	0.058	0.091	0.103	0.126	0.128
BVAR-SV	0.021	0.029	0.035	0.054	0.076	0.085	0.112	0.108
SPF	0.049	0.058	0.067	0.073	0.082	0.092	0.104	0.101

Table 5: Measures of accuracy of the predictive distributions of ‘annual average over annual average’ GDP growth $\overline{R}_{\text{GDP}}$ (top half) and inflation $\overline{R}_{\text{INF}}$ (bottom half). Distributions from the SPF reports from 1992:Q1 to 2009:Q4 are compared to those obtained from three time series models (BVAR, BVAR-SV and C-Skt) fitted to matching vintages of real-time data. In row 1, the results are broken down by forecasts for the current year Y and the next year $Y+1$, as stated in the SPF reports. In row 2, the results are further broken down by quarter in which the the report/vintage occurs. Rows 3 and 4 give the number of quarters already observed in the national accounts, and the number that require forecasting, to compute $\overline{R}_{\text{GDP}}$ and $\overline{R}_{\text{INF}}$.

**Web Appendix for ‘Asymmetric Forecast Densities for
U.S. Macroeconomic Variables from a Gaussian Copula
Model of Cross-Sectional and Serial Dependence’**

Part A: Time Series Properties

In Part A we outline some time series properties of the copula model, which establish the stationarity and Markov order of the multivariate time series.

1. First, it is straightforward (eg. see Song 2000) to show that all sub-vectors of $\mathbf{U}$ also have distributions that are Gaussian copulas. This type of ‘closure’ property for the copula function C in Equation (3.1) is required for the time series to be stationary (Smith 2014). Using this property, the joint marginal density of the contiguous vectors $(\mathbf{U}'_s, \dots, \mathbf{U}'_t)'$ with $t > s$, which we denote as

$$c_{(s,\dots,t)}(\mathbf{u}_s, \dots, \mathbf{u}_t) = \int c(\mathbf{u}) d\mathbf{u}_1 \dots d\mathbf{u}_{s-1} d\mathbf{u}_{t+1} \dots d\mathbf{u}_T = c_{\text{Ga}}(\mathbf{u}_s, \dots, \mathbf{u}_t; \Omega_{(s,t)}),$$

is also that of a Gaussian copula. It has a parameter matrix which we denote as $\Omega_{(s,t)}$, and is given by the $(t-s+1)m \times (t-s+1)m$ sub-matrix of Ω that comprises rows/columns $(s-1)m+1, \dots, tm$. This copula captures the dependence in the distribution of $(\mathbf{Y}_s, \dots, \mathbf{Y}_t)$, which has density

$$f(\mathbf{y}_s, \dots, \mathbf{y}_t) = c_{(s,\dots,t)}(\mathbf{u}_s, \dots, \mathbf{u}_t) \prod_{l=s}^t \prod_{j=1}^m f(y_{j,l}).$$

When $s = t$, this gives the copula density for $\mathbf{Y}_t$, which we denote as $c_{(t)}(\mathbf{u}_t) = c_{\text{Ga}}(\mathbf{u}_t; \Omega_{(t)})$, with $\Omega_{(t)}$ the $m \times m$ sub-matrix of Ω comprising rows/columns $(t-1)m+1, \dots, tm$. The matrix $\Omega_{(t)}$ captures any cross-sectional dependence between the elements of $\mathbf{Y}_t$.

2. Second, Property 1 and the definition of a conditional density, can be used to express the transitional densities of the time series model as

$$\begin{aligned} f(\mathbf{y}_t | \mathbf{y}_{t-1}, \dots, \mathbf{y}_s) &= \frac{c_{(s,\dots,t)}(\mathbf{u}_s, \dots, \mathbf{u}_t)}{c_{(s,\dots,t-1)}(\mathbf{u}_s, \dots, \mathbf{u}_{t-1})} \prod_{j=1}^m f(y_{j,t}) \\ &= f(\mathbf{u}_t | \mathbf{u}_{t-1}, \dots, \mathbf{u}_s) \prod_{j=1}^m f(y_{j,t}). \end{aligned}$$

The density $f(\mathbf{u}_t | \mathbf{u}_{t-1}, \dots, \mathbf{u}_s)$ is that of the conditional distribution of $\mathbf{U}_t | \mathbf{U}_{t-1}, \dots, \mathbf{U}_s$,

and it captures the dynamics in the multivariate time series. While this conditional density is equal to the ratio of two copula densities, it is not itself the density of a copula, and the elements of $\mathbf{U}_t$ are not marginally uniformly distributed conditional on the past values. This is important, because otherwise the predictive distribution for each variable $Y_{j,t}$ would simply be given by its marginal density $f(y_{j,t})$.

3. Third, the series $\{\mathbf{U}_t\}_{t=1}^T$ is also stationary when a stationary vector autoregression is adopted for the latent Gaussian process $\{\mathbf{W}_t\}_{t=1}^T$, as we now show. Stationarity here refers to strong stationarity, where

$$c_{(t-k, \dots, t)}(\mathbf{u}_{t-k}, \dots, \mathbf{u}_t) = c_{(t'-k, \dots, t')}(\mathbf{u}_{t'-k}, \dots, \mathbf{u}_{t'}), \text{ for all } t, t' \text{ and } k \geq 0.$$

From Property 1 this condition holds for a Gaussian copula if and only if $\Omega_{(t-k, t)} = \Omega_{(t'-k, t')}$. This is the case here because Ω is a block Toeplitz matrix when a stationary vector autoregression is employed for the latent Gaussian process. One implication of this property is that the marginal copulas $c_{(t)}(\mathbf{u}_t) = c_{\text{Ga}}(\mathbf{u}_t; \Omega_{(t)})$ that account for cross-sectional contemporaneous dependence in $\mathbf{Y}_t$, are time invariant with parameter matrix $\Omega_{(t)} = \Omega(0)$ for all t .

4. The last property we establish here is strong stationarity of the multivariate times series $\{\mathbf{Y}_t\}_{t=1}^T$, which is where

$$f(\mathbf{y}_{t-k}, \dots, \mathbf{y}_t) = f(\mathbf{y}_{t'-k}, \dots, \mathbf{y}_{t'}), \text{ for all } t, t' \text{ and } k \geq 0.$$

This follows immediately from Properties 1 and 3, as long as the marginal densities for $f(y_{j,t})$ are invariant with respect to t . However, they do not have to be invariant with respect to j , and are not so in this study.

Part B: Additional Empirical Results

Full Period: 1975:Q1 – 2011:Q2								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
GDP Growth					Inflation			
<i>Time Series Model</i>								
BVAR	0.972	0.971	0.982	0.975	0.381	0.441	0.503	0.684
Copula & EDF	0.949	0.963	0.955	0.949*	0.395	0.442	0.511	0.665
Copula & Skew t	0.941	0.947	0.963	0.948*	0.386	0.421	0.482	0.613
BVAR-SV	0.850**	0.871**	0.896**	0.921	0.373	0.416	0.474	0.645
Unemployment					Interest Rate			
<i>Time Series Model</i>								
BVAR	0.349	0.624	1.132	1.673	1.003	1.574	2.253	3.508
Copula & EDF	0.438**	0.771**	1.338	1.994	1.003	1.576	2.293	3.649
Copula & Skew t	0.346	0.627	1.167	1.899	0.836**	1.451**	2.164	3.568
BVAR-SV	0.317**	0.560**	1.061	1.738	0.826**	1.392	2.275	3.825
Three Recessions: 1982:Q1–1982:Q2, 1991:Q1–1991:Q3 and 2008:Q4–2009:Q3								
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
GDP Growth					Inflation			
<i>Time Series Model</i>								
BVAR	2.066	2.337	2.103	1.935	0.651	0.817	1.005	1.014
Copula & EDF	2.136	2.608	2.24	2.039	0.498	0.592	0.832	0.789
Copula & Skew t	1.865	2.475	2.331	2.042	0.625	0.631	0.805	0.772
BVAR-SV	1.924	2.319	2.354	2.304	0.604	0.717	0.799	0.833
Unemployment					Interest Rate			
<i>Time Series Model</i>								
BVAR	0.773	2.140	4.187	4.082	0.838	1.229	2.273	4.936
Copula & EDF	1.015	2.583	4.939	4.948	1.138	1.784	2.930	5.417
Copula & Skew t	0.549	1.813	4.080	4.483	1.096	1.305	2.540	5.952
BVAR-SV	0.698	2.021	4.546	4.881	1.022	1.516	2.805	5.808

Table 1: Mean values of $10 \times \text{TW-CRPS}$ computed over two evaluation periods. Separate results are reported for forecasts made $h = 1, 2, 4$ and 8 quarters ahead in columns, and each row corresponds to results from a different model. As a rough guide, results are reported for a test equivalent to the Diebold and Mariano (1995) test (with the Harvey, Leybourne and Newbold (1997) small-sample adjustment) applied to the TW-CRPS values. Test results are only reported for the first evaluation period (the second has only 9 quarters). This is a two-sided test of the equality of the mean TW-CRPS of the BVAR with each of the other three models. For these three models, rejection of the null hypothesis of equal means at the 10% or 5% level is denoted with ‘*’ or ‘**’.

Figure 1: Additional posterior summaries of the fitted Gaussian copula model for the vintage 2011:Q2 data with skew t margins. Estimates of the posterior means of the selection indicators $\gamma(0), \dots, \gamma(4)$ are given in the top row. White cells correspond to elements with $\Pr(\gamma_{l_1, l_2}^k = 1 | y) \approx 1$, while black cells correspond to elements with $\Pr(\gamma_{l_1, l_2}^k = 0 | y) \approx 1$. The next row gives estimates of the Kendall's tau for each unique pair copula in blocks $K_{t, t-k}$ for $k = 0, \dots, 4$. These are computed directly from the unique partial correlations $\phi(0), \dots, \phi(4)$. The last row gives the corresponding unconditional Spearman correlations. The variables are output growth (Grth), inflation (Infln), the unemployment rate (UR) and the interest rate (IR).

Figure 2: Forecasts from the BVAR made using a quarterly expanding window with data vintages 2007:Q4 to 2011:Q1. Panels (a) to (d) correspond to variables Y_1 to Y_4 . The solid blue lines are the predictive means between 1 and 8 quarters ahead, and the dashed blue lines are the the 10th and 90th percentiles of the predictive distributions. The second release of the actual values are plotted as a solid black line. The blue circles are the forecast origins (2007:Q3 to 2010:Q4), and correspond to the first release of the data.

Figure 3: Forecasts from the BVAR-SV made using a quarterly expanding window with data vintages 2007:Q4 to 2011:Q1. Panels (a) to (d) correspond to variables Y_1 to Y_4 . Details are given in the caption of Web Appendix Figure 2.

Figure 4: Forecasts from the copula model with EDF margins made using a quarterly expanding window with data vintages 2007:Q4 to 2011:Q1. Panels (a) to (d) correspond to variables Y_1 to Y_4 . Details are given in the caption of Web Appendix Figure 2.

Figure 5: Forecasts from the copula model with skew t margins made using a quarterly expanding window with data vintages 2007:Q4 to 2011:Q1. Panels (a) to (d) correspond to variables Y_1 to Y_4 . Details are given in the caption of Web Appendix Figure 2.

Vintage	Y_1	Y_2	Y_3	Y_4	Vintage	Y_1	Y_2	Y_3	Y_4
1975:Q1	0.0127	0.0000	0.0225	0.0067	1993:Q2	0.0107	0.0004	0.0027	0.0000
1975:Q2	0.0041	0.0000	0.0184	0.0098	1993:Q3	0.0111	0.0003	0.0030	0.0000
1975:Q3	0.0039	0.0000	0.0043	0.0136	1993:Q4	0.0078	0.0003	0.0033	0.0000
1975:Q4	0.0040	0.0000	0.0019	0.0172	1994:Q1	0.0071	0.0003	0.0037	0.0000
1976:Q1	0.0209	0.0002	0.0012	0.0231	1994:Q2	0.0066	0.0002	0.0040	0.0000
1976:Q2	0.0167	0.0002	0.0014	0.0297	1994:Q3	0.0054	0.0002	0.0042	0.0000
1976:Q3	0.0145	0.0001	0.0017	0.0379	1994:Q4	0.0047	0.0002	0.0041	0.0000
1976:Q4	0.0124	0.0001	0.0018	0.0470	1995:Q1	0.0043	0.0001	0.0037	0.0000
1977:Q1	0.0113	0.0001	0.0019	0.0564	1995:Q2	0.0037	0.0001	0.0032	0.0000
1977:Q2	0.0098	0.0001	0.0008	0.0659	1995:Q3	0.0040	0.0001	0.0029	0.0000
1977:Q3	0.0141	0.0003	0.0010	0.0771	1995:Q4	0.0033	0.0001	0.0026	0.0000
1977:Q4	0.0125	0.0003	0.0015	0.0941	1996:Q1	0.0000	0.0000	0.0023	0.0000
1978:Q1	0.0102	0.0004	0.0021	0.1146	1996:Q2	0.0000	0.0000	0.0020	0.0000
1978:Q2	0.0145	0.0006	0.0028	0.1297	1996:Q3	0.0000	0.0000	0.0017	0.0000
1978:Q3	0.0107	0.0007	0.0032	0.1398	1996:Q4	0.0000	0.0000	0.0015	0.0000
1978:Q4	0.0100	0.0006	0.0038	0.1141	1997:Q1	0.0000	0.0000	0.0013	0.0000
1979:Q1	0.0089	0.0010	0.0040	0.0941	1997:Q2	0.0067	0.0000	0.0011	0.0000
1979:Q2	0.0110	0.0014	0.0042	0.0978	1997:Q3	0.0058	0.0000	0.0010	0.0000
1979:Q3	0.0147	0.0014	0.0042	0.0527	1997:Q4	0.0051	0.0000	0.0009	0.0000
1979:Q4	0.0151	0.0022	0.0042	0.0285	1998:Q1	0.0044	0.0000	0.0009	0.0000
1980:Q1	0.0157	0.0026	0.0044	0.0097	1998:Q2	0.0038	0.0000	0.0008	0.0000
1980:Q2	0.0170	0.0029	0.0054	0.0004	1998:Q3	0.0029	0.0000	0.0008	0.0000
1980:Q3	0.0046	0.0011	0.0063	0.0005	1998:Q4	0.0023	0.0000	0.0007	0.0000
1980:Q4	0.0045	0.0011	0.0061	0.0006	1999:Q1	0.0023	0.0000	0.0006	0.0000
1981:Q1	0.0021	0.0025	0.0066	0.0001	1999:Q2	0.0022	0.0000	0.0006	0.0000
1981:Q2	0.0018	0.0031	0.0067	0.0000	1999:Q3	0.0018	0.0000	0.0005	0.0000
1981:Q3	0.0027	0.0026	0.0068	0.0000	1999:Q4	0.0000	0.0000	0.0004	0.0000
1981:Q4	0.0038	0.0025	0.0066	0.0000	2000:Q1	0.0000	0.0000	0.0004	0.0000
1982:Q1	0.0032	0.0025	0.0052	0.0000	2000:Q2	0.0060	0.0000	0.0003	0.0000
1982:Q2	0.0045	0.0017	0.0038	0.0000	2000:Q3	0.0053	0.0000	0.0003	0.0000
1982:Q3	0.0088	0.0015	0.0035	0.0000	2000:Q4	0.0048	0.0000	0.0002	0.0000
1982:Q4	0.0097	0.0016	0.0020	0.0000	2001:Q1	0.0049	0.0000	0.0002	0.0000
1983:Q1	0.0134	0.0016	0.0013	0.0000	2001:Q2	0.0046	0.0000	0.0002	0.0000
1983:Q2	0.0120	0.0017	0.0005	0.0000	2001:Q3	0.0054	0.0000	0.0002	0.0000
1983:Q3	0.0146	0.0017	0.0003	0.0000	2001:Q4	0.0072	0.0000	0.0001	0.0000
1983:Q4	0.0131	0.0012	0.0002	0.0000	2002:Q1	0.0089	0.0000	0.0001	0.0000
1984:Q1	0.0113	0.0011	0.0003	0.0000	2002:Q2	0.0072	0.0000	0.0001	0.0000
1984:Q2	0.0100	0.0011	0.0004	0.0000	2002:Q3	0.0103	0.0000	0.0001	0.0000
1984:Q3	0.0110	0.0007	0.0005	0.0000	2002:Q4	0.0087	0.0000	0.0001	0.0000
1984:Q4	0.0106	0.0006	0.0006	0.0000	2003:Q1	0.0094	0.0000	0.0001	0.0000
1985:Q1	0.0105	0.0005	0.0007	0.0000	2003:Q2	0.0081	0.0000	0.0001	0.0000
1985:Q2	0.0115	0.0006	0.0009	0.0000	2003:Q3	0.0072	0.0000	0.0001	0.0000
1985:Q3	0.0132	0.0005	0.0010	0.0000	2003:Q4	0.0071	0.0000	0.0001	0.0000
1985:Q4	0.0117	0.0004	0.0004	0.0000	2004:Q1	0.0048	0.0000	0.0001	0.0000
1986:Q1	0.0711	0.0265	0.0005	0.0000	2004:Q2	0.0043	0.0000	0.0001	0.0000
1986:Q2	0.0743	0.0206	0.0006	0.0000	2004:Q3	0.0032	0.0000	0.0001	0.0000
1986:Q3	0.0397	0.0144	0.0007	0.0000	2004:Q4	0.0028	0.0000	0.0001	0.0000
1986:Q4	0.0401	0.0119	0.0009	0.0000	2005:Q1	0.0024	0.0000	0.0001	0.0000
1987:Q1	0.0404	0.0109	0.0010	0.0000	2005:Q2	0.0020	0.0000	0.0001	0.0000
1987:Q2	0.0371	0.0094	0.0012	0.0000	2005:Q3	0.0019	0.0000	0.0000	0.0000
1987:Q3	0.0416	0.0114	0.0013	0.0000	2005:Q4	0.0016	0.0000	0.0000	0.0000
1987:Q4	0.0365	0.0088	0.0034	0.0000	2006:Q1	0.0017	0.0000	0.0000	0.0000
1988:Q1	0.0302	0.0069	0.0032	0.0000	2006:Q2	0.0014	0.0000	0.0000	0.0000
1988:Q2	0.0295	0.0053	0.0029	0.0000	2006:Q3	0.0014	0.0000	0.0000	0.0000
1988:Q3	0.0243	0.0047	0.0027	0.0000	2006:Q4	0.0013	0.0000	0.0000	0.0000
1988:Q4	0.0246	0.0055	0.0024	0.0000	2007:Q1	0.0011	0.0000	0.0000	0.0000
1989:Q1	0.0246	0.0058	0.0021	0.0000	2007:Q2	0.0011	0.0000	0.0000	0.0000
1989:Q2	0.0196	0.0057	0.0019	0.0000	2007:Q3	0.0011	0.0000	0.0000	0.0000
1989:Q3	0.0182	0.0040	0.0016	0.0001	2007:Q4	0.0010	0.0000	0.0000	0.0000
1989:Q4	0.0162	0.0031	0.0014	0.0001	2008:Q1	0.0011	0.0000	0.0000	0.0000
1990:Q1	0.0183	0.0028	0.0012	0.0001	2008:Q2	0.0012	0.0000	0.0000	0.0000
1990:Q2	0.0175	0.0030	0.0011	0.0001	2008:Q3	0.0013	0.0000	0.0000	0.0000
1990:Q3	0.0260	0.0026	0.0009	0.0001	2008:Q4	0.0015	0.0000	0.0000	0.0000
1990:Q4	0.0277	0.0022	0.0008	0.0001	2009:Q1	0.0020	0.0000	0.0000	0.0000
1991:Q1	0.0395	0.0018	0.0008	0.0001	2009:Q2	0.0010	0.0000	0.0000	0.0000
1991:Q2	0.0519	0.0021	0.0010	0.0001	2009:Q3	0.0010	0.0000	0.0000	0.0000
1991:Q3	0.0582	0.0018	0.0011	0.0001	2009:Q4	0.0008	0.0000	0.0000	0.0000
1991:Q4	0.0596	0.0016	0.0013	0.0001	2010:Q1	0.0009	0.0000	0.0000	0.0000
1992:Q1	0.0000	0.0000	0.0015	0.0001	2010:Q2	0.0007	0.0000	0.0000	0.0000
1992:Q2	0.0000	0.0000	0.0017	0.0001	2010:Q3	0.0007	0.0000	0.0000	0.0000
1992:Q3	0.0000	0.0000	0.0020	0.0001	2010:Q4	0.0006	0.0000	0.0000	0.0000
1992:Q4	0.0000	0.0000	0.0023	0.0001	2011:Q1	0.0005	0.0000	0.0000	0.0000
1993:Q1	0.0116	0.0005	0.0025	0.0000	2011:Q2	0.0005	0.0000	0.0000	0.0000

Table 2: P-values from Shapiro Wilk (1965) tests for normality applied to output growth (Y_1), inflation (Y_2), unemployment (Y_3) and the interest rate (Y_4). The null hypothesis is normality. Tests are performed on each of the 146 vintages of data from 1975:Q1 to 2011:Q2. P-values greater than 0.05 are given in bold.

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Smith, MS;Vahey, SP

Title:

Asymmetric Forecast Densities for US Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence

Date:

2016-07-01

Citation:

Smith, M. S. & Vahey, S. P. (2016). Asymmetric Forecast Densities for US Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence. JOURNAL OF BUSINESS & ECONOMIC STATISTICS, 34 (3), pp.416-434. <https://doi.org/10.1080/07350015.2015.1044533>.

Persistent Link:

<http://hdl.handle.net/11343/123420>