Real-Time Macroeconomic Forecasting with a Heteroskedastic Inversion Copula

Rubén Loaiza-Maya and Michael S. Smith

First Version: July 2017

This Version: 23 July 2018

Rubén Loaiza-Maya is a PhD student and Michael Smith is Chair of Management (Econometrics), both at Melbourne Business School, University of Melbourne. We thank Joshua Chan, Tomasz Wozniak and participants at the 11th Annual RCEA Bayesian Econometric Workshop in Melbourne, Australia, for useful comments. We also thank the Editor Todd Clark and three referees for helpful remarks that have improved the manuscript substantially. This includes correction of an error concerning the treatment of missing observations in an earlier version of the paper, and making us aware of the previous work of Adrian et al. (2017). Correspondence should be directed to Michael Smith at mike.smith@mbs.edu.

Real-Time Macroeconomic Forecasting with a Heteroskedastic Inversion Copula

Abstract

There is a growing interest in allowing for asymmetry in the density forecasts of macroeconomic variables. In multivariate time series, this can be achieved with a copula model, where both serial and cross-sectional dependence is captured by a copula function, and the margins are nonparametric. Yet most existing copulas cannot capture heteroskedasticity well, which is a feature of many economic and financial time series. To do so, we propose a new copula created by the inversion of a multivariate unobserved component stochastic volatility model, and show how to estimate it using Bayesian methods. We fit the copula model to real-time data on five quarterly U.S. economic and financial variables. The copula model captures heteroskedasticity, dependence in the level, time-variation in higher moments, bounds on variables and other features. Over the window 1975Q1 – 2016Q2, the real-time density forecasts of all the macroeconomic variables exhibit time-varying asymmetry. In particular, forecasts of GDP growth have increased negative skew during recessions. The point and density forecasts from the copula model are competitive with those from benchmark models – particularly for inflation, a short term interest rate and current quarter GDP growth.

Key Words: Asymmetric Density Forecasting; Time Series Copula; Downside Risk; Macroeconomic Uncertainty.

1 Introduction

There is growing evidence that accounting for time-varying asymmetries and heavy tails in the real-time predictive distributions of macroeconomic variables improves their accuracy. Density forecasts constructed from expert surveys — such as the Survey of Professional Forecasters (SPF) in the U.S. — can exhibit strong asymmetries and heavy tails (Krüger and Nolte, 2015) and perform strongly for short horizons (Krüger et al., 2017). Quantile regression estimates of the distribution of U.S. GDP growth, conditional on economic and financial conditions, exhibit time-varying skew that turns negative during recessions (Adrian et al., 2017). Giglio et al. (2016) find similar results for several measures of U.S. and European real activity. Time series models that incorporate stochastic volatility produce time-varying heavy-tailed forecast densities that are among the most accurate from time series models (D'Agostino et al., 2013, Clark and Ravazzolo, 2015), and there is recent evidence (Chiu et al., 2016) that their extension to also allow for asymmetries can increase accuracy further. Moreover, in the post-recessionary period U.S. interest rates effectively reached their zero lower bound, during which time their predictive densities were necessarily positively skewed (Johannsen and Mertens, 2016). Despite this, only a few existing multivariate time series models produce predictive distributions that can feature strong asymmetries and heavy tails. To do so here we propose a new copula model, and employ it to model and forecast key U.S. macroeconomic variables in real time.

Copulas allow the margins and dependence structure of a multivariate distribution to be modeled separately, a feature that has made them popular in many areas (McNeil et al., 2005, Nelsen, 2006). In time series analysis, they are mostly used to model conditional cross-sectional dependence between multiple series; see Patton (2012) for a review. The correlation structure can be dynamic, but time-changing variance (i.e. heteroskedasticity) is modeled through the margins, which also account for any asymmetries in the distributions of the variables; see De Lira Salvatierra and Patton (2015) and Creal and Tsay (2015) for examples. However, a second way a copula can be employed in multivariate time series is to use it to capture both cross-sectional and serial dependence jointly. A major advantage of this approach is that the marginal distribution of each variable can be modeled directly, including as nonparametric. Smith and Vahey (2016) show that this is attractive for macroeconomic variables because their marginal distributions are often complex, and we follow this second approach here.

The main challenge in constructing such a model is the selection of an appropriate copula. Smith and Vahey (2016) employ a Gaussian copula, but this does not capture heteroskedasticity well, which is a key feature of GDP growth, inflation and other variables (Clark, 2011, Clark and Ravazzolo, 2015). Another choice is a vine copula (Brechmann and Czado, 2015, Smith, 2015, Beare and Seo, 2015), which can be used to capture heteroskedasticity when specific pair-copula components are employed (Loaiza-Maya et al., 2018). In this paper we suggest a new copula that can also do so as an alternative. This is an 'inversion copula' that is formed by a process of inverting (Nelsen, 2006, p.52) the data distribution from an existing multivariate time series model, which we call a 'pseudo time series' model. The inversion copula exactly replicates the dependence structure of the pseudo time series. However, when it is combined with new margins, the copula model forms a new and more flexible time series model. If the margins are the same as those of the pseudo time series model, then the copula model simply reproduces it. But by adopting nonparametric margins, the copula model can deviate substantially from the pseudo model, and has the potential to be much more accurate.

We construct our inversion copula from a multivariate unobserved component stochastic volatility (UCSV) model, where the trend component follows a stationary vector autoregresssion (VAR). Stochastic volatility is widely used to capture heteroskedasticity in macroeconomic series (Clark, 2011, D'Agostino et al., 2013, Carriero et al., 2016), while UCSV models are popular for modeling inflation (Stock and Watson, 2007, Mertens, 2016) joint with unemployment and an interest rate (Cogley and Sargent, 2005). By extracting the copula of this UCSV model we adopt a dependence structure that is established by previous studies as well-suited to these series. However, the margins of the UCSV model are symmetric and (as we show here) inconsistent with those observed empirically for most macroeconomic variables, so that adopting the copula model increases its accuracy.

Because the UCSV model does not have a closed form likelihood, its inversion copula cannot be written in closed form either. However, following Smith and Maneesoonthorn (2018), we show how methods for estimating state space models can also be used to estimate this copula. In particular, we develop a Markov chain Monte Carlo (MCMC) scheme to do so, although importance sampling methods (Jungbacker and Koopman, 2007, Scharth and Kohn, 2016) can also be used. The parameters of the underlying UCSV model become the copula parameters, for which constraints are derived that identify the likelihood. An advantage of inversion copulas is that they are fast and simple to simulate from. We show how measures of dependence, filtered and predictive distributions, and other inference can all be computed via simulation. Existing high-dimensional inversion copulas include elliptical and skew-elliptical copulas (Demarta and McNeil, 2005, Smith et al., 2012) and those constructed from univariate state space Smith and Maneesoonthorn (2018) or static factor models (Oh and Patton, 2017). However, as far as we are aware, ours is the first copula constructed by inversion of a multivariate nonlinear time series model.

The copula model is applied to quarterly U.S. real-time data on GDP growth, inflation, unemployment, a short term interest rate and stock market volatility. Adaptive kernel density estimates are used for the margins, and capture substantial skew and heavy tails for each variable and at most vintages, multi-modality for unemployment and the interest rate, and zero lower bounds for the interest rate and stock market volatility. When fitted to the 2016Q2 vintage data, strong serial dependence within and across series is captured by the copula. This includes negative correlation between stock market volatility and GDP growth next quarter, and strong heteroskedasticity in the time series – including volatility 'spillover'. The filtered distributions of the variables exhibit variance similar to that from a VAR with stochastic volatility (VARSV), but very different from that obtained by fitting an UCSV model directly to the data. Thus, incorporating more accurate margins changes the estimated serial dependence structure of the UCSV model substantially. The filtered distributions from the copula model also feature strong time-varying skew and kurtosis, unlike the necessarily symmetric distributions from the VARSV and UCSV models. For example, they are most heavily negatively skewed for GDP Growth prior to 1982 and during the subsequent recessions.

We examine the real-time out-of-sample density forecasts from the copula model for 1975Q1 to 2016Q2. These show that the features of the nonparametric margins are not simply transferred to the predictive distributions. Instead, they have moments (including higher order moments) that vary over time, and are — as one would expect — much sharper than the margins. In particular, the forecasts of GDP growth have lower quantiles that decrease during recessionary periods, whereas the upper quantiles remain stable over time, which is consistent with the findings of Adrian et al. (2017). The opposite is found for inflation, where the upper quantiles vary more over time than the lower quantiles, representing asymmetry in inflation uncertainty. Such features cannot be captured by VARSV and UCSV models with symmetric disturbances. The real-time forecasts for the four economic variables from our copula model are compared to those from a VAR, VARSV, UCSV and a Gaussian copula model as in Smith and Vahey (2016). The point and density forecasts from our copula model are more accurate for inflation and the interest rate (the two most skewed variables) than those from the benchmark models. Notably, the current quarter forecasts of GDP growth are the most accurate, including when compared to those from the VARSV model, which is a high-performing forecasting model for this variable (Clark and Ravazzolo, 2015). We also construct bivariate real-time forecast densities for the change in the unemployment rate and GDP growth, which exhibit asymmetries and non-linearities. In line with Okun's rule, Spearman correlations between the two variables are negative, but show a severe weakening of the relationship during the Great Recession.

Our model is not the only time series model that can produce asymmetric density forecasts. For example, a zero lower bound on the interest rate can be captured by adopting a shadow rate (Bauer and Rudebusch, 2016, Johannsen and Mertens, 2016), producing positively skewed density forecasts for this variable. Another approach is to extend existing time series models using finite or infinite mixture models (Kalli and Griffin, 2018, Jensen and Maheu, 2010, Chiu et al., 2016). However, an advantage of the copula model is that the margins can be modeled directly in an arbitrary manner at an initial stage, and the dependence separately in a second stage.

The rest of the paper is organized as follows. The heteroskedastic inversion copula is presented in Section 2, while Section 3 outlines the approach for its estimation and the computation of filtered and predictive distributions. Section 4 presents the empirical analysis. This includes the data, existing evidence of asymmetry in these series, the results of fitting the copula model to the 2016Q2 vintage, and the real-time forecast study. Section 5 discusses implications and extensions.

2 Copula Model

2.1 Copula Time Series Model

Consider a continuous-valued multivariate time series $\{\boldsymbol{Y}_t\}$, with $\boldsymbol{Y}_t = (Y_{1,t}, \ldots, Y_{r,t})'$ having r elements. Then Sklar's theorem (Nelsen, 2006, p.46) states that the joint distribution function F_Y of $\boldsymbol{Y} = (\boldsymbol{Y}_1', \ldots, \boldsymbol{Y}_T')'$ can be written as

$$F_Y(\boldsymbol{y}) = C(\boldsymbol{u})\,, \tag{1}$$

where $\boldsymbol{y} = (\boldsymbol{y}_1', \ldots, \boldsymbol{y}_T')'$, $\boldsymbol{y}_t = (y_{1,t}, \ldots, y_{r,t})'$, $\boldsymbol{u} = (\boldsymbol{u}_1', \ldots, \boldsymbol{u}_T')'$, $u_t = (u_{1,t}, \ldots, u_{r,t})'$, $u_{i,t} = F_{Y_{i,t}}(y_{i,t})$, and $F_{Y_{i,t}}$ is the marginal distribution function of $Y_{i,t}$. The function C is called the copula function, and is itself a joint distribution function on $[0,1]^{rT}$ with strictly uniform margins (Nelsen, 2006, p.47). It captures the entire dependence structure of $\boldsymbol{Y}$, which here is both the cross-sectional and temporal dependence structure of the time series. The major advantage of this decomposition is that it allows the margins $F_{Y_{i,t}}$ and dependence structure to be modeled separately.

Throughout this paper we assume (strong) stationarity, so that the margins of the series are time-invariant, and $F_{Y_{i,t}} = G_i$ for $i = 1, \ldots, r$. In this case, the density of the data is

$$f_Y(\boldsymbol{y}) = \frac{\partial^{rT}}{\partial y_{1,1} \cdots \partial y_{r,T}} F_Y(\boldsymbol{y}) = c(\boldsymbol{u}) \prod_{i=1}^{r} \prod_{t=1}^{T} g_i(y_{i,t})\,, \tag{2}$$

where $g_i(y) = \frac{d}{dy} G_i(y)$, and $c(\boldsymbol{u}) = \frac{\partial^{rT}}{\partial u_{1,1} \cdots \partial u_{r,T}} C(\boldsymbol{u})$ is called the 'copula density'. The main challenge of this copula model is the selection of an appropriate copula function C (or, equivalently, density c). For a univariate time series, Smith and Maneesoonthorn (2018) suggest constructing a copula by inverting a state space model, while for a multivariate time series Smith and Vahey (2016) construct a copula by inverting a VAR. We extend these ideas below, and propose using a new copula constructed from the inversion of a multivariate UCSV model.

2.2 Heteroskedastic Inversion Copula

Consider a second time series $\{\boldsymbol{Z}_t\}$, with $\boldsymbol{Z}_t = (Z_{1,t}, \ldots, Z_{r,t})'$ having r elements. This follows a stationary unobserved component model with stochastic volatility disturbances (UCSV) given by:

$$\boldsymbol{Z}_t \;=\; \boldsymbol{\mu}_t + \Lambda_t^{0.5} \boldsymbol{\varepsilon}_t\,,$$

$$
\begin{aligned}
\boldsymbol{\mu}_t &= \sum_{k=1}^{p} B_k \boldsymbol{\mu}_{t-k} + \boldsymbol{\epsilon}_t \,, \\
h_{i,t} &= \bar{h}_i + a_i(h_{i,t-1} - \bar{h}_i) + \eta_{i,t} \,, \text{ for } i = 1, \ldots, r \,,
\end{aligned} \tag{3}
$$

$\Lambda_t = \text{diag}\left(e^{h_{1,t}}, \ldots, e^{h_{r,t}}\right)$, $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, I_r)$, $\boldsymbol{\eta}_t = (\eta_{1,t}, \ldots, \eta_{r,t})' \sim N(\mathbf{0}, \Sigma_\eta)$, and $\boldsymbol{\epsilon}_t \sim N(\mathbf{0}, \Sigma_\epsilon)$. In this model, $\boldsymbol{Z}_t$ is decomposed into the sum of two zero mean stationary components: a level component $\boldsymbol{\mu}_t$ which follows a VAR(p) model, and a variance component which has r correlated elements that each follow stochastic volatility models, with $Z_{i,t}|\mu_{i,t}, h_{i,t} \sim N(\mu_{i,t}, e^{h_{i,t}})$. Eq. (3) defines a nonlinear state space model, with state vector $(\boldsymbol{\mu}_t', \boldsymbol{h}_t')'$, where $\boldsymbol{h}_t = (h_{1,t}, \ldots, h_{r,t})'$ is the log-volatility vector. The autoregressive coefficients $\boldsymbol{a} = (a_1, \ldots, a_r)'$ are bounded so that $-1 < a_i < 1$. We refer to $\{\boldsymbol{Z}_t\}$ as a 'pseudo time series', because it is not observed directly. We extract its dependence structure by constructing the Tr-dimensional copula density of $\boldsymbol{Z} = (\boldsymbol{Z}_1', \ldots, \boldsymbol{Z}_T')'$, and use it for c in Eq. (2).

Let Ψ be the parameters of the UCSV model at Eq. (3), and $\boldsymbol{Z}$ have distribution function $F_Z(\boldsymbol{z}|\Psi)$. Because the pseudo time series is stationary, the marginal distribution function $F_i(z_{i,t}|\Psi)$ of $Z_{i,t}$ is time invariant. Following Nelsen (2006, Sec.3.1), the copula function and density of this distribution can be obtained by inversion of the usual expression of Sklar's theorem, to give

$$
\begin{aligned}
C(\boldsymbol{u}|\Psi) &= F_Z\left(F_1^{-1}(u_{1,1}|\Psi), F_2^{-1}(u_{2,1}|\Psi), \ldots, F_r^{-1}(u_{r,T}|\Psi)|\Psi\right) \,, \text{ and} \\
c(\boldsymbol{u}|\Psi) &= \frac{f_Z\left(\boldsymbol{z}|\Psi\right)}{\prod_{i=1}^{r}\prod_{t=1}^{T} f_i(z_{i,t}|\Psi)} \,,
\end{aligned} \tag{4}
$$

where $z_{i,t} = F_i^{-1}(u_{i,t}|\Psi)$, $\boldsymbol{z}_t = (z_{1,t}, \ldots, z_{r,t})'$, $\boldsymbol{z} = (\boldsymbol{z}_1', \ldots, \boldsymbol{z}_T')'$, $f_Z(\boldsymbol{z}|\Psi) = \frac{\partial^{rT}}{\partial z_{1,1}\cdots\partial z_{r,T}} F_Z(\boldsymbol{z}|\Psi)$ and $f_i(z_{i,t}|\Psi) = \frac{d}{dz_{i,t}} F_i(z_{i,t}|\Psi)$. We call this an 'inversion copula' (it is also called an 'implicit copula' by McNeil et al. (2005, p.190)). However, for the data distribution of the UCSV model, f_Z and f_i cannot be expressed in closed form, so that neither can c in Eq. (4). Nevertheless, the likelihood at Eq. (2) can still be computed by expressing it conditional on $\boldsymbol{\mu} = (\boldsymbol{\mu}_1', \ldots, \boldsymbol{\mu}_T')'$ and $\boldsymbol{h} = (\boldsymbol{h}_1', \ldots, \boldsymbol{h}_T')'$, and integrating $\boldsymbol{\mu}$ and $\boldsymbol{h}$ out using Bayesian methods, which is the approach we employ.

The conditional likelihood can be obtained by a change of variables from $\boldsymbol{Z}$ to $\boldsymbol{Y}$, as

$$
f\left(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi\right) = f(\boldsymbol{z}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) J_{Z\to Y} = \prod_{t=1}^{T}\left(\phi_r\left(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \Lambda_t\right)\prod_{i=1}^{r}\frac{g_i(y_{i,t})}{f_i(z_{i,t}|\Psi)}\right) \,, \tag{5}
$$

where $J_{Z\to Y} = \prod_{t=1}^{T}\prod_{i=1}^{r}\frac{g_i(y_{i,t})}{f_i(z_{i,t}|\Psi)}$ is the the Jacobian of the transformation, and $\phi_r\left(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \Lambda_t\right)$ is the density of a $N(\boldsymbol{\mu}_t, \Lambda_t)$ distribution evaluated at $\boldsymbol{z}_t$. To compute Eq. (5), the most difficult task is evaluating the marginal densities f_i, and the corresponding quantile functions F_i^{-1}. The latter is needed to compute $z_{i,t} = F_i^{-1}(u_{i,t}|\Psi)$ from $u_{i,t} = G_i(y_{i,t})$. The density

$$f_i(z_{i,t}|\Psi) = \int \phi_1(z_{i,t}; \mu_{i,t}, h_{i,t}) f(\mu_{i,t}|\Psi) f(h_{i,t}|\Psi) \mathrm{d}(\mu_{i,t}, h_{i,t}) \,. \tag{6}$$

The state $\mu_{i,t}|\Psi \sim N(0, d_i)$, where $\boldsymbol{d} = (d_1, \ldots, d_r)$ can be computed from $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$ (Lütkepohl, 2005, pp.28-29), although (as we discuss below) we parameterize the vector autoregression directly in terms of $\boldsymbol{d}$ instead. The state $h_{i,t}|\Psi \sim N(\bar{h}_i, s_i^2)$, where $s_i^2 = \bar{\sigma}_{i,i}^2/(1 - a_i^2)$, with $\bar{\sigma}_{i,i}^2$ the ith leading diagonal element of Σ_η. Substituting these two state densities into Eq. (6), and integrating out $\mu_{i,t}$ analytically as a Gaussian, gives the expressions

$$\begin{aligned} f_i(z_{i,t}|\Psi) &= \int \phi_1\left(z_{i,t}; 0, d_i + e^h\right) \phi_1\left(h; \bar{h}_i, s_i^2\right) \mathrm{d}h \\ F_i(z_{i,t}|\Psi) &= \int \Phi_1\left(z_{i,t}; 0, d_i + e^h\right) \phi_1\left(h; \bar{h}_i, s_i^2\right) \mathrm{d}h \,, \end{aligned} \tag{7}$$

where $\Phi_r(\boldsymbol{z}; \boldsymbol{m}, \Omega)$ is the distribution function of a $N(\boldsymbol{m}, \Omega)$ evaluated at point $\boldsymbol{z}$, and the integrals over h are computed numerically. Finally, to evaluate F_i^{-1} from F_i, we use the fast numerical method proposed by Smith and Maneesoonthorn (2018, Appendix A).

Tab. A1 in the Online Appendix depicts the transformations underlying the inversion copula time series model, and tabulates the various distribution and density functions.

2.3 Parameterization and Identification

Let $\tilde{\boldsymbol{\mu}} = (\boldsymbol{\mu}_{p+1}', \ldots, \boldsymbol{\mu}_1')'$, then the stationary VAR(p) model for the level component $\{\boldsymbol{\mu}_t\}$ can be parameterized in terms of $\boldsymbol{d}$ and the unique partial correlations of the covariance matrix Var($\tilde{\boldsymbol{\mu}}$). If the partial correlation between series i and j, at lag k, is defined as

$$\begin{aligned} \gamma_{i,j}^0 &= \text{Corr}\left(\mu_{i,t}, \mu_{j,t}|\text{intervening elements of vector } \tilde{\boldsymbol{\mu}}\right), \text{if } k = 0 \,, \text{and} \\ \gamma_{i,j}^k &= \text{Corr}\left(\mu_{i,t-k}, \mu_{j,t}|\text{intervening elements of vector } \tilde{\boldsymbol{\mu}}\right), \text{for } 1 \leq k \leq p \,, \end{aligned}$$

then $\boldsymbol{\gamma} = \{\gamma_{i,j}^k; (i,j,k) \in \mathcal{P}\}$, for $\mathcal{P} = \{1 \leq j < i \leq r\,; k = 0\} \cup \{1 \leq i,j \leq r;\, 1 \leq k \leq p\}$, is the set of all unique partial correlations. There is a one-to-one relationship between $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$ and $\{\boldsymbol{\gamma}, \boldsymbol{d}\}$. Lütkepohl (2005, pp.29-31) outlines how to compute $\{\boldsymbol{d}, \boldsymbol{\gamma}\}$ from $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$, while

Appendix A shows how to undertake the reverse. In total there are $r^2p+r(r-1)/2$ partial correlations, so that there are 110 in our empirical analysis where $r=5$ and $p=4$. Adopting this parameterization simplifies our MCMC scheme in two ways. First, each $\gamma^k_{i,j}$ is unconstrained on $[-1,1]$ conditional on the other partials, simplifying generation and improving mixing. Second, the conditional posterior of $\boldsymbol{\gamma}$ is not a function of $F_1,\ldots,F_r$, so that generation of $\boldsymbol{\gamma}$ does not require re-computation of $\boldsymbol{z}_1,\ldots,\boldsymbol{z}_T$, which is slow. Following Daniels and Pourahmadi (2009), we also parameterize the $(r\times r)$ matrix Σ_η in terms of its variances and partial correlations.

When forming a copula via inversion, all information about the marginal location and scale of $Z_{i,t}$ is lost. To identify the parameters, constraints on these moments are required; for example, in the Gaussian copula $Z_{i,t}\sim N(0,1)$. We constrain the first two marginal moments of $Z_{i,t}$. First, $E(Z_{i,t})=0$ by definition of the model at Eq. (3). Second, the marginal variance is

$$\begin{aligned}\mathrm{Var}(Z_{i,t}) &= \mathrm{Var}(\mu_{i,t})+\mathrm{Var}(\exp(h_{i,t}/2)\epsilon_{i,t})\\ &= d_i+E(\exp(h_{i,t}))=d_i+\exp(\bar{h}_i+s_i^2/2)\,,\end{aligned}$$

where $s_i^2=\bar{\sigma}^2_{i,i}/(1-a_i^2)$. Setting $\mathrm{Var}(Z_{i,t})=1$, and solving for $\bar{h}_i$, gives the equality constraints $\bar{h}_i=\log(1-d_i)-s_i^2/2$, for $i=1,\ldots,r$. These are simple to enforce in the MCMC scheme by evaluating (and not generating) $\bar{\boldsymbol{h}}=(\bar{h}_1,\ldots,\bar{h}_r)'$ at each sweep. In addition, because $E(\exp(h_{i,t}))\geq 0$, the parameters $0\leq d_i\leq 1$, for $i=1,\ldots,r$. These inequality constaints are easily enforced in a Bayesian framework by adopting priors that constrain $d_i\in(0,1)$. With the parameterization and constraints outlined above, the copula parameters are now $\Psi=\{\boldsymbol{\gamma},\boldsymbol{d},\boldsymbol{a},\Sigma_\eta\}$.

It is important to note that the 'role' of each parameter in the UCSV model for $\{\boldsymbol{Z}_t\}$ does not translate to the same role for $\{\boldsymbol{Y}_t\}$, so that Ψ is simply a set of copula parameters. However, d_i can instead be interpreted as determining the relative importance of the two components for $Z_{i,t}$: serial dependence is captured by the homoskedastic component $\mu_{i,t}$ when $d_i\to 1$, and the heteroskedastic component when $d_i\to 0$. If all elements $\boldsymbol{d}=\boldsymbol{1}$, then $E(\exp(h_{i,t}))=0$ and $\boldsymbol{Z}_t=\boldsymbol{\mu}_t$. In this case, our proposed copula reduces to the inversion copula of a Gaussian VAR(p), which is the Gaussian copula discussed by Smith and Vahey (2016), but without the D-vine representation employed by these authors. If $\boldsymbol{d}=\boldsymbol{0}$, then $\boldsymbol{\mu}_t=\boldsymbol{0}$ and $\boldsymbol{Z}_t=\Lambda_t^{1/2}\boldsymbol{\varepsilon}_t$, so that C is the inversion copula of a multivariate stochastic volatility model only.

3 Estimation and Inference

3.1 Priors

All elements of $\boldsymbol{\gamma}, \boldsymbol{d}, \boldsymbol{a}$ are assumed independent a priori, and constrained to their admissible domains. Each cross-sectional partial correlation $\gamma_{i,j}^0$ has a $N(0, 100)$ prior constrained to $(-1, 1)$, which is close to a flat prior. Each partial $\gamma_{i,j}^k$ for $1 \leq k \leq p$ has a $N(0, \frac{0.04}{k^2})$ prior constrained to $(-1, 1)$, so that the longer the lag of the partial correlation, the tighter is its prior density around zero. This is similar to a Minnesota prior (Litterman, 1986), where higher levels of shrinkage are applied to autoregressive coefficients of longer lags. The prior for d_i is uniform on $[0, 1]$, while $\pi(a_i, \bar{\sigma}_{i,i}^2) \propto \mathrm{I}(0 < s_i^2 < 3\,;\, |a_i| < 0.99)$ which ensures that $\{a_i, \bar{\sigma}_{i,i}^2\}$ remain in a numerically stable region. Finally, uniform priors on $(-1, 1)$ are used for the $r(r-1)/2$ partial correlations of the matrix Σ_η, which is a standard uninformative prior on a correlation matrix (Daniels and Pourahmadi, 2009).

3.2 Posterior

Estimation of Ψ requires evaluation of the copula density at Eq. (4), which is the likelihood conditional on the margins $(G_1, \ldots, G_r)$. Even though the copula is high-dimensional, its parameters can still be estimated, as we now discuss. By exploiting the state space formulation of $\{\boldsymbol{Z}_t\}$ in Eq (3), the numerator of the copula density $f_Z(\boldsymbol{z}|\Psi)$ can be computed using a number of existing methods in the state space literature. We adopt a Bayesian MCMC approach and explicitly generate the states $\{\boldsymbol{\mu}, \boldsymbol{h}\}$ as part of the following sampling scheme.

Sampling Scheme

Step 1: Generate from $\boldsymbol{\mu}|\Psi, \boldsymbol{h}, \boldsymbol{y}$.

Step 2: Generate from $\boldsymbol{h}|\Psi, \boldsymbol{\mu}, \boldsymbol{y}$.

Step 3: Generate from $\boldsymbol{\gamma}|\{\Psi\backslash\boldsymbol{\gamma}\}\boldsymbol{\mu}, \boldsymbol{h}, \boldsymbol{y}$.

Step 4: Generate from $d_i|\{\Psi\backslash d_i\}, \boldsymbol{\mu}, \boldsymbol{h}, \boldsymbol{y}$, for $i = 1, \ldots, r$.

Step 5: Generate from $a_i|\{\Psi\backslash a_i\}, \boldsymbol{h}, \boldsymbol{\mu}, \boldsymbol{y}$, for $i = 1, \ldots, r$.

Step 6: Generate from $\Sigma_\eta|\{\Psi\backslash\Sigma_\eta\}, \boldsymbol{h}, \boldsymbol{\mu}, \boldsymbol{y}$.

In Step 1, we compute $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$ from $\{\boldsymbol{d}, \boldsymbol{\gamma}\}$ as shown in the Appendix A, and then generate $\boldsymbol{\mu}$ as a multivariate normal as in Chan and Jeliazkov (2009). To generate $\boldsymbol{h}$ in Step 2 we approximate the distribution of $\log(\varepsilon_{i,t}^2)$ for $i = 1, \ldots, r$, using a mixture of seven normals, and use the simulation smoother. In Step 3, the partial correlations $\boldsymbol{\gamma}$ are sampled as pairs of elements chosen at random, where each pair is generated conditional on the remaining elements of $\boldsymbol{\gamma}$ using adaptive random walk Metropolis-Hastings (MH). In Steps 4, 5 and 6 the elements of $\boldsymbol{d}$, $\boldsymbol{a}$, and the diagonal elements of Σ_η are all generated one at a time using MH with normal approximations to the conditional posteriors as the proposals. Finally, in Step 6 the partial correlations of the matrix Σ_η are generated one at a time using adaptive random walk MH. After convergence, we collect the Monte Carlo sample $\{\Psi^{[1]}, \ldots, \Psi^{[J]}\}$ from which posterior inference is computed in the standard Bayesian fashion; the Online Appendix gives a detailed description of the sampling scheme.

3.3 Simulation and Forecasting

A major advantage of copulas constructed by inversion of a parametric model is that simulation of an iterate $\boldsymbol{U} \sim C(\cdot|\Psi)$ is both fast and simple. To do so, first simulate a series of values $\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_T$ from the UCSV model[1] at Eq. (3), and then compute $U_{i,t} = F_i(Z_{i,t}|\Psi)$ for $i = 1, \ldots, r$ and $t = 1, \ldots, T$. Monte Carlo estimates of a wide range of dependence measures — including slices of the copula density — can then be computed, as in Section 4.3.2.

The h-step ahead posterior predictive density of $\boldsymbol{Y}_{T+h}$, conditional on $\boldsymbol{Y}_{1:T} = \boldsymbol{y}$, is

$$f(\boldsymbol{y}_{T+h}|\boldsymbol{y}) = \int f(\boldsymbol{y}_{T+h}|\boldsymbol{y}, \boldsymbol{\mu}_{1:T}, \boldsymbol{h}_T, \Psi) f(\boldsymbol{\mu}_{1:T}, \boldsymbol{h}_T, \Psi|\boldsymbol{y}) \mathrm{d}(\boldsymbol{\mu}_{1:T}, \boldsymbol{h}_T, \Psi)\,.$$

The integration above cannot be computed analytically, but it is fast and simple to simulate from the distribution as follows. At each sweep of the sampler, simulate from the conditional distribution $\boldsymbol{Z}_{T+h}|\boldsymbol{\mu}_{1:T}, \boldsymbol{h}_T, \Psi$ by first simulating $\boldsymbol{\mu}_{T+h}$ and $\boldsymbol{h}_{T+h}$ using the transition equations, and then generating $\boldsymbol{Z}_{T+h}$ from the measurement equation, both given at Eq (3). Finally, compute $Y_{i,T+h} = G_i^{-1}(F_i(Z_{i,T+h}|\Psi))$ for $i = 1, \ldots, r$.

[1] To do so, we are careful to sample a pre-period to remove the effect of initial values of the VAR(p) process $\{\boldsymbol{\mu}_t\}$.

The analogous (in-sample) h-step-ahead predictive density of $\boldsymbol{Y}_{t+h}$, for $t = 1, \ldots, T-h$, is

$$f(\boldsymbol{y}_{t+h}|\boldsymbol{y}_{1:t}) \equiv \int f(\boldsymbol{y}_{t+h}|\boldsymbol{\mu}_{1:t}, \boldsymbol{h}_t, \Psi) f(\boldsymbol{\mu}_{1:t}, \boldsymbol{h}_t, \Psi|\boldsymbol{y}) \mathrm{d}(\boldsymbol{\mu}_{1:t}, \boldsymbol{h}_t, \Psi)\,,$$

which can be computed by simulating from $\boldsymbol{Z}_{t+h}|\boldsymbol{\mu}_{1:t}, \boldsymbol{h}_t, \Psi$ and then transforming, as above. For both densities, by simulating the iterates at each sweep of the sampler, Ψ is integrated out with respect to its posterior in the usual Bayesian fashion.

4 Empirical Analysis

4.1 U.S. Data and Asymmetry

We consider quarterly data on U.S. output growth, inflation, the unemployment rate, a nominal short-term interest rate and stock market volatility between 1954Q1 and 2016Q2. Output growth ($\Delta \ln(\mathrm{GDP})_t$) and inflation ($\mathrm{Infl}_t$) are the quarterly differences in the logarithms of real GDP and the GDP price deflator, respectively. These are observed in real time, and are obtained from the Real-Time Data Set for Macroeconomists (Croushore and Stark, 2001). The data are provided in quarterly vintages, where vintage $T+1$ contains observations up to quarter T, and reflects the lag in the national accounts. Each vintage also provides revisions for the most recent quarters, and sometimes also the entire series. The interest rate (IR_t) is the quarterly average of the monthly secondary market rate of the 3-Month T-Bill, and the unemployment rate (UR_t) is the quarterly average of the monthly seasonally-adjusted civilian unemployment rate; both obtained from the FRED database. Stock market volatility (VXO_t) is the quarterly average VXO volatility index sourced from the FRED database, and backfilled prior to 1986 using calibrated daily SP500 realized volatility as in Bloom (2009). We set $\boldsymbol{Y}_t = (\Delta \ln(\mathrm{GDP})_t, \mathrm{Infl}_t, \mathrm{UR}_t, \mathrm{IR}_t, \mathrm{VXO}_t)'$, and Fig. 1(f–j) plots the five series for vintage 2016Q2.

The four macroeconomic variables were considered by Clark and Ravazzolo (2015), Smith and Vahey (2016) and Krüger and Nolte (2015). Some vintages are missing early observations of output and the price deflator, and for these vintages we follow Clark and Ravazzolo (2015) and work only with the available information. For example, for the 1992Q1 vintage the 1959Q1 observations of output and the price deflator are the earliest values used. Adrian et al. (2017) use the National

Financial Conditions Index produced by the Federal Reserve Bank of Chicago as a measure of financial conditions. However, this only extends back to 1971, so that we instead use the VXO index, which Adrian et al. (2017) also found to be related to GDP growth.

There is strong evidence for asymmetry in the marginal distributions of these five variables. Shapiro-Wilk and Jarque-Bera tests (unreported) reject normality for all variables and vintages, except for some early vintages of the interest rate and unemployment. Using the nonparametric test of symmetry of Bai and Ng (2005), symmetry is rejected at the 5% level for GDP growth vintages between 1977Q4 and 1991Q4, inflation from vintage 1992Q2 onwards, and stock market volatility until 1987Q4 and between 2006Q1 and 2008Q4. However, this test can have low power, and even strong deviations from symmetry may be undetected. Fig. 1(a–e) plots histograms of the variables from vintage 2016Q2, and strong positive skew in inflation and market volatility is visible. Unemployment, market volatility and the interest rate are bounded on the left, with the latter having a zero lower bound which it effectively reached between 2008Q4 and 2015Q4. Fig. 1 also shows the adaptive kernel density estimates (KDEs) of Shimazaki and Shinomoto (2010), which are strongly asymmetric and heavy-tailed. They approximately capture lower bounds on unemployment and the interest rate, capture the smooth tails and asymmetry in the variables, and have multiple modes for unemployment and the interest rate. Except where otherwise stated, KDEs from each vintage are used for the margins $G_1, \ldots, G_5$ in our copula time series models.

Predictive distributions of growth can also exhibit asymmetry. To illustrate, we show that the predictive distributions for GNP/GDP reported in the Survey of Professional Forecasters (SPF) are often asymmetric. These are an ensemble of the survey respondents' predictive distributions for the "average annual over average annual" growth rate for both the same and next year; see Smith and Vahey (2016) for a discussion. These have the drawback that they mix over differing multi-horizon quarterly forecasts, and to address this we only consider the first quarter (Q1) reports. In the Q1 reports, the forecast distribution for the same year involves forecasts of GDP during all four quarters of the year of the report. In addition, the forecast distribution for the next year requires forecasts of GDP during the four quarters of the following year as well. Therefore, the same year forecasts involve a forecast horizon of four quarters, and the next year forecasts a horizon of eight quarters — horizons that remain the same in all Q1 reports. The SPF forecast distributions are a discrete piecewise specification, for which the Pearson skew coefficient can be uninformative. Therefore, we measure asymmetry by Bowley's skew coefficient Bow $= (q_3 + q_1 - 2q_2)/(q_3 - q_1)$, which uses the quantiles q_1, q_2 and q_3 and is more robust to the specification of the distribution. Fig. 2 plots this measure for the same and next year predictive distributions of annual growth in the Q1 reports, and sizable skew exists. Bootstrap standard errors show that this skew is statistically significant in many reports. Skew for the same year growth exhibits a high degree of persistence over years. Interestingly, the next year forecasts tend to be negatively skewed, which reflect beliefs about the long-run distribution of growth on the part of the SPF respondents. Later, we show that the real-time predictive distributions from our copula model also exhibit skew that varies over the vintages.

4.2 Benchmark Models

We label the heteroskedastic copula model as UCSV-C, and compare it to the following benchmarks:

(i) VAR: A fourth order Gaussian vector autoregression with a Minnesota prior.

(ii) VAR-C: The copula model of Smith and Vahey (2016), but with the prior $\pi(\boldsymbol{\gamma}, \boldsymbol{d})$ outlined in Section 3.1 and KDE margins, which we find improves the forecast accuracy slightly compared to the spike-and-slab prior these authors employ. This copula is equivalent to that in Section 2, but where $\boldsymbol{\varepsilon}_t = \mathbf{0}$, so that $\boldsymbol{Z}_t = \boldsymbol{\mu}_t$ and the pseudo time series follows a VAR(4).

(iii) VARSV: The fourth order Gaussian vector autoregression with stochastic volatility, implemented as outlined in Clark and Ravazzolo (2015).

(iv) UCSV: The unobserved component model with stochastic volatility outlined in Section 2.2, but applied directly to the data and without the parameter constraints required for identification of the copula (i.e. without the constraints on $\bar{\boldsymbol{h}}$ and $\boldsymbol{d}$).

The copula parameters are estimated using the Bayesian method outlined in Section 3. The VAR, VARSV and UCSV models are estimated using standard MCMC schemes; for example, see Koop et al. (2010). A total of 30,000 iterates were discarded for convergence, while a further 40,000 were collected as a Monte Carlo sample. These prove to be conservative choices, with the empirical results robust to variations in the starting states, and doubling or halving of the sampling and burnin periods.

The dependence structure of the heteroskedastic copula model fit to the 2016Q2 vintage data is discussed below first, followed by the model's real-time forecasting performance over the window 1975Q1–2016Q2.

4.3 Fitted Model for 2016Q2 Vintage

4.3.1 Parameter Estimates

Tab. 1 reports posterior means and 90% probability intervals of the copula parameters $\{\boldsymbol{d}, \boldsymbol{a}, \Sigma_\eta\}$. Posterior estimates of $\boldsymbol{\gamma}$, the states $\boldsymbol{\mu}, \boldsymbol{h}$ and the alternative VAR parameterization $\{B_1, \ldots, B_4, \Sigma_\epsilon\}$ are given in the Online Appendix. We briefly discuss their implication for the pseudo time series $\{\boldsymbol{Z}_t\}$. First, $\hat{d}_3 = 0.978$ and $\hat{d}_4 = 0.975$, so that the homoskedastic level components $\mu_{3,t}$ and $\mu_{4,t}$ explain a very high proportion of variation in $Z_{3,t}$ and $Z_{4,t}$, respectively. This suggests there is very little serial correlation in the variance of these two series. Second, $Z_{1,t}, Z_{2,t}$ and $Z_{5,t}$ have log-volatilities that exhibit positive serial correlation ($\hat{a}_1 = 0.763$, $\hat{a}_2 = 0.746$ and $\hat{a}_5 = 0.772$). This is consistent with time series plots of the posterior estimates of the states, which are the two conditional moments $\mu_{i,t} = E(Z_{i,t}|\boldsymbol{\mu}_t, \boldsymbol{h}_t)$ and $h_{i,t} = \log(\mathrm{Var}(Z_{i,t}|\boldsymbol{\mu}_t, \boldsymbol{h}_t))$. These are in the Online Appendix, along with estimates of $\boldsymbol{\gamma}$, Σ_η and Σ_ϵ, which indicate significant cross-sectional dependence between the elements of $\{\boldsymbol{Z}_t\}$; both contemporaneous and lagged.

4.3.2 Time Series Dependence Structure

To measure the estimated dependence structure of the actual time series $\{\boldsymbol{Y}_t\}$, we first compute the matrices of pairwise Spearman correlations $R(k) = \{\rho_{i,j}(k)\}_{i=1,\ldots,5;j=1,\ldots,5}$ at lag k, where

$$\rho_{i,j}(k) = \mathrm{Corr}(U_{i,t}, U_{j,t-k}) = 12E(U_{i,t}U_{j,t-k}) - 3\,,$$

and $U_{i,t} = G_i(Y_{i,t})$. The expectation above is with respect to the posterior distribution of Ψ. It is computed by averaging over draws from the copula $c(\boldsymbol{u}|\Psi)$ at the end of each sweep of the sampling scheme, integrating out Ψ with respect to its posterior. Tab. 2 presents the posterior means of $R(0)$, $R(1)$ and $R(4)$. For example, correlation between the unemployment rate and GDP growth lagged four quarters is -0.26, so that unemployment is a lagging indicator. Interestingly, correlation between GDP growth and stock market volatility lagged one quarter is -0.176, so that increases in stock market volatility leads to a reduction in output growth the following quarter, which is consistent with the quantile regression results of Adrian et al. (2017). Positive serial dependence exists in all five series, and it is strong for the unemployment rate, inflation, the interest rate and stock market volatility, but is weaker for GDP growth.

The Spearman correlations above measure dependence in the *level* of the time series. Yet the main advantage of our proposed copula is that it can also capture conditional heteroskedasticity. To illustrate, Fig. 3 plots four bivariate slices of the fitted copula density. These slices are in the variable pairs $(\Delta(\ln(\mathrm{GDP})_t, \Delta(\ln(\mathrm{GDP})_{t-1})$, $(\mathrm{Infl}_t, \mathrm{Infl}_{t-1})$, $(\mathrm{IR}_t, \mathrm{Infl}_{t-1})$ and $(\Delta(\ln(\mathrm{GDP})_t), \mathrm{VXO}_{t-1})$, and are constructed via simulation from the fitted copula. The first three slices have the majority of their mass concentrated along the leading diagonal (due to the positive level dependence between the variables) while the last slice has most mass on the off-diagonal (due to the negative level dependence between the variables). However, there is also a slight up-turn in the corners at (0,1) and (1,0) in panel (a), and to a less visible extent in panel (b). Loaiza-Maya et al. (2018) show that this characteristic is symptomatic of positive dependence between the *variances* of two variables. To measure this more precisely these authors propose computing the dependence in a variation process, which we also do here as outlined below.

Let $e_{i,t} = Y_{i,t} - E(Y_{i,t}|\boldsymbol{y}_{t-1}, \ldots, \boldsymbol{y}_{t-p})$ be the the innovation in the ith variable. Consider the transformation $V : \mathbb{R} \to \mathbb{R}^+$, such that (i) $V(a) = V(-a) \geq 0$, (ii) $V(0) = 0$, and (iii) $\frac{d}{da}V(a) > 0$ if $a > 0$.[2] Then Loaiza-Maya et al. (2018) show that the copula of any pair of transformed innovations $V(e_{i,t}), V(e_{j,s})$ is invariant to the specific choice of V. They propose measuring the volatility dependence between this pair using their Spearman correlation

$$\rho^V_{i,j}(k) = \mathrm{Corr}(\tilde{U}_{i,t}, \tilde{U}_{j,t-k}) = 12E(\tilde{U}_{i,t}, \tilde{U}_{i,t-k}) - 3\,,$$

where $\tilde{U}_{i,t} = F_{V_i}(V(e_{i,t}))$ and F_{V_i} is the distribution function of $V(e_{i,t})$. Computation of this metric can be undertaken via simulation for any time series model that has stationary innovations, including

[2] Examples of transformations that satisfy these conditions are simply $V(a) = |a|$ and $V(a) = a^2$.

all five time series models fit here. Tab. 3 reports the matrix $R^V(k) = \left\{\rho^V_{i,j}(k)\right\}_{i=1,\ldots,5;j=1,\ldots,5}$ for $k = 0, 1$ and both the VARSV and UCSV-C models. Both models capture heteroskedasticity (ie. $\rho^V_{i,i}(1) > 0$), and mostly positive volatility co-movements (ie. $\rho^V_{i,j}(0) > 0$ for $i \neq j$) and spillovers (ie. $\rho^V_{i,j}(1) > 0$ for $i \neq j$). The level of heteroskedasticity is similar for both models, while the volatility spillovers and co-movements are generally stronger for the VARSV model.

4.3.3 Predictive Distributions and Time-Varying Moments

The differences between the fitted UCSV-C, UCSV and VARSV models, are illustrated by the differences in their (in-sample) predictive distributions. For example, Fig. 4 plots the (a) mean, (b) standard deviation, (c) Pearson skew coefficient, and (d) kurtosis of the one-quarter-ahead predictive distributions of GDP growth from the three models, computed as in Section 3.3. Panel (b) shows that the UCSV-C model exhibits a much higher degree of heteroskedasticity than the UCSV model. Therefore, allowing for greater flexibility in the margins affects the estimated serial dependence structure substantially. The UCSV-C and VARSV models exhibit a similar degree of heteroskedasticity, which responds to macroeconomic events — such as the Great Moderation — similarly.

However, in panels (c) and (d), there is a large difference between the three models. The predictive distributions of the UCSV-C model exhibit negative skewness and higher kurtosis than the exactly symmetric distributions from the other two models. Moreover, even though only the first two moments of $\{\boldsymbol{Z}_t\}$ are time varying, when the inversion copula is combined with the five asymmetric margins, it also produces time variation in the third and fourth moments of $\{\boldsymbol{Y}_t\}$. For example, panel (c) shows that skew is sharply negative during each recession, reflecting increases in downside risk. Prior to the Great Moderation there is also a heavy lower tail (high kurtosis and negative skew), demonstrating a small, but persistent, probability of a contraction. Except during the three subsequent recessions, this heavy tail all but disappears after 1982. The ability to account for time-varying tail risk is a major feature of our proposed multivariate time series model. Equivalent plots for the other variables can be found in the Online Appendix, and show that the predictive distributions from the UCSV-C model also differ from those of the two benchmark models substantially.

Finally, as a measure of goodness-of-fit for each model, we compute the accuracy of the one-step-ahead predictive densities using the (mean) continuously ranked probability score (CRPS) of Gneiting and Raftery (2007). This metric is evaluated by integration of the quantile score as in Smith and Vahey (2016, Appendix A.2). Tab. 5 reports the values, and the UCSV-C model is more accurate than the benchmark models for every variable. As a guide, Harvey et al. (1997) small-sample adjusted Diebold-Mariano tests indicate that this dominance is statistically significant.

4.4 Real-Time Forecasts

We study the real-time density and point forecasts from our copula model, and their accuracy relative to those from the benchmark models. The models are fitted sequentially to each data vintage between 1975Q1 and 2016Q2, and the out-of-sample predictive densities $f(\boldsymbol{y}_{T+h}|\boldsymbol{y})$ computed for $h = 1, \ldots, 8$ and all vintages. The case where $h = 1$ corresponds to forecasting current quarter macroeconomic conditions, also called 'nowcasts' (Garratt et al., 2014).

4.4.1 Level of Asymmetry in Margins and Predictive Densities

Fig. 5(a-d) displays estimates of the marginals $G_1, \ldots, G_5$ obtained from fits to vintages 1985Q2, 1996Q4, 2006Q2 and 2016Q2, and we make three observations. First, the direction of skew does not vary over vintage, with GDP growth exhibiting negative skew, and the other variables positive skew. Second, the distribution of inflation in the post-Great Moderation 2006Q2 and 2016Q2 vintages is sharper with a lower mean and variance than the earlier two vintages. Last, the marginal distributions of unemployment and the interest rate are multi-modal, which is consistent with a small number of distinct regimes. It is possible to smooth away the modes by fitting a unimodal distribution, such as a skew t, but we retain them here and show that they do not inhibit forecast accuracy.

Fig. 5(e-h) plots the current quarter predictive densities for the same four vintages. These differ dramatically from the estimated margins, showing that the UCSV-C model does not simply transfer the features of the marginals to the predictive densities. For example, the predictive densities of unemployment and the interest rate are all unimodal, in contrast to their margins. The predictive densities are also much sharper than the marginals; for example, that for the interest rate at vintage 2016Q2 is concentrated between 0% and 1%, reflecting the monetary environment during that quarter.

In contrast, both the margins and predictive distributions of the VAR, VARSV and UCSV models are symmetric and unimodal. To illustrate the impact of this constraint on real-time forecasts, we examine the current quarter GDP growth predictive distributions. Fig. 6 plots the (a) standard deviation and (b) skew of these from both the UCSV-C and VARSV models. (We later show that those from the former model are more accurate). The predictive distributions from the VARSV are exactly symmetric, so have zero skew. In comparison, not only are the predictive distributions of the UCSV-C model skewed at times, but the skew varies depending on macroeconomic conditions. For example, forecasts are negatively skewed during the recessionary periods of the early 1980's, 2001 and the Great Recession, whereas post-2009Q2 they are largely positively skewed. Allowing for asymmetry in the predictive distributions also affects their standard deviation. For example, in the years prior to the Great Recession, the predictive standard deviation from the UCSV-C model was greater than that from the VARSV, but with upside risk due to positive skew. The converse is true during the second half of the Great Recession (ie. 2008Q4 – 2009Q2).

Adrian et al. (2017) characterize this asymmetry in real-time density forecasts using plots of the expected shortfall, and its upper tail equivalent, and we follow suit. Fig. 7 plots these at the 5% level for current quarter forecasts of GDP growth and inflation from the UCSV-C and VARSV models. From the UCSV-C model, the lower quantiles of GDP growth forecasts vary more over time than the upper quantiles. This is consistent with the empirical findings of Adrian et al. (2017), who use (skew t smoothed) quantile regression. For inflation, we find the opposite empirical result – the upper quantiles vary more over time than lower quantiles. In comparison, such features cannot appear in the VARSV forecasts. Similar conclusions can be drawn from forecasts at longer horizons, and equivalent plots when $h = 4$ are given in the Online Appendix.

To assess the impact of monetary policy after the Great Recession on real-time density forecasts of the interest rate, Fig. 8 plots the four moments of current quarter forecasts of this variable from 2000Q1 to 2016Q2. While all models show a post-recession mean close to zero in panel (a), both the VARSV and UCSV models overstate the standard deviation during this period in panel (b). In contrast, the UCSV-C model adapts to this unusual monetary policy well, with predictive distributions that have low variation, but high positive skew and kurtosis in panels (c) and (d) — a necessary feature of any distribution that is left-truncated with a mean close to the lower bound.

The Online Appendix provides plots of the standard deviation and skew (both Bowley and Pearson) of the real-time forecasts for all five variables and horizons $h = 1, 4$ from the UCSV-C model. Throughout, the forecasts exhibit time-varying asymmetry. The inflation forecasts are positively skewed at both horizons. Also given in the Online Appendix are rolling fan charts comparing the forecast densities from the VAR, VAR-C, VARSV and UCSV-C models.

4.4.2 Forecast Accuracy

To measure real-time forecast accuracy, we follow Romer and Romer (2000) and others and use GDP growth and inflation computed from the second revision of the national accounts as their true values, or 'outturns'. (Although the results are similar, and conclusions unchanged, when using the first revision — ie. second release — as in Smith and Vahey (2016).) For all models, the predictive means are the point forecasts, and their accuracy is measured using the root mean square error (RMSE). The accuracy of the density forecasts are measured by the mean CRPS metric. (The same conclusions are drawn from the logarithmic scores, which are reported in the Online Appendix.) For the UCSV-C model, we consider two marginal estimators in addition to the KDE. These are a skew t distribution estimated using maximum likelihood, and a Bayesian nonparametric constrained density estimator described in Appendix B and the Online Appendix. The latter enforces zero lower bounds on IR_t and VXO_t. To distinguish the results, we suffix UCSV-C with 'KDE', 'Skt' and 'BNP'.[3]

Tab. 4 reports the forecast accuracy for the four macroeconomic variables over the evaluation period 1975Q1 – 2016Q2, with the metrics for the copula models listed last. As an initial guide, results of a Harvey et al. (1997) small sample adjusted Diebold-Mariano two-sided test for the equality of the mean scores of the UCSV-C-KDE with each of the other models are reported. Rejection of the null hypothesis of equal means at the 10% or 5% level is denoted by '*' or '**' if favorable to the UCSV-C-KDE model, or by '+' or '++' if unfavorable. While no single multivariate time series

[3] For example, UCSV-C-Skt is the label for the copula model with skew t margins and the inversion copula constructed from the UCSV pseudo time series model.

model is uniformly more accurate for all four variables at all horizons, there are a number of clear differences, and we make five observations on these.

First, forecasts from the nonlinear models are more accurate than those from the linear VAR, except for unemployment. With this variable, forecasts from the nonlinear models are no more accurate than those from the VAR, except the density forecasts from the VARSV model at short horizons. Second, the results from the UCSV-C models are robust to the choice of margin (KDE, BNP or Skt), although with some exceptions. These include increased density forecast accuracy of GDP growth at long horizons ($h = 4, 8$) and inflation at short horizons ($h = 1, 2$) using the BNP margins. Moreover, the skew t margin does not account for the zero lower bound on the interest rate, and forecasts of this variable from the UCSV-C-Skt model are less accurate. Third, forecasts from the UCSV-C-KDE model are generally more accurate than those from the VAR-C-KDE model, and substantially so for inflation and GDP growth. This is consistent with the higher degree of heteroskedasticity in these series, which is better captured by the UCSV-C introduced here, than the VAR-C discussed by Smith and Vahey (2016). Fourth, the forecasts from the UCSV-C-KDE model are significantly more accurate than those from the UCSV model for GDP growth at $h = 1, 2$, and the interest rate at $h = 1, 2$. However, the reverse is never the case for any variable or horizon. Thus, combining nonparametric margins with the UCSV-C produces a more flexible time series model than the UCSV model, and improves forecast accuracy. Fifth, when compared to the VARSV model, forecasts of current quarter GDP growth and long-run ($h = 8$) inflation made from the UCSV-C-KDE model are significantly more accurate. Thus, the time-varying asymmetries in the current quarter GDP growth density forecasts documented above, translate into increased accuracy. This is a strong result, as the VARSV is known to be an accurate forecasting model for these series in real time (Clark, 2011, Clark and Ravazzolo, 2015). However, the VARSV dominates the UCSV-C-KDE for medium-run ($h = 4$) GDP growth and short-run ($h = 1$) unemployment forecasts.

Finally, we note that VAR-C model has greater relative forecast accuracy than that reported by Smith and Vahey (2016) for the same copula. This is because of differences between the studies in (i) the way that early missing observations were dealt with for some vintages, (ii) the priors employed, and (iii) the marginals employed.

4.4.3 Forecasting Economic Relationships in Real Time

Multivariate density forecasts from our copula model can also be constructed using iterates simulated from the predictive distribution, as in Section 3.3. To illustrate, we construct bivariate current quarter forecasts of $\Delta\mathrm{UR}_{T+1} = \mathrm{UR}_{T+1} - \mathrm{UR}_T$ and $\Delta\ln(\mathrm{GDP})_{T+1}$. Okun's rule suggests there should be a negative relationship between these two variables. The bivariate densities (see the Online Appendix) exhibit strong asymmetries, with a negative relationship between the two variables that can be nonlinear. In contrast, forecasts from the VARSV and UCSV models are necessarily elliptical densities with linear relationships. To illustrate how the relationship has changed over time, Fig. 9 plots the Spearman correlations of the current quarter real-time forecasts of these two variables between 1975Q1 and 2016Q2. The correlation is negative throughout, although it varies over time. It is strongest during the Great Moderation, and weakest during the Great Recession, which is consistent with existing evidence for this relationship (Ball et al., 2015).

5 Discussion

This paper makes two main contributions. The first is methodological, where we propose a new copula constructed by inversion of the data distribution from an UCSV model. When combined with nonparametric margins, it results in a multivariate time series model that has the same dependence structure as an UCSV model — including that due to heteroskedasticity — but is more flexible. Inversion copulas are attractive for the modeling of dependence in high dimensions, such as dimension Tr here. While a Gaussian copula is a popular choice of inversion copula its dependence structure is limited. For time series, drawable vine copulas are a viable alternative (Smith, 2015, Loaiza-Maya et al., 2018), but are less general and arguably more complex than the copula suggested here. Even though the density of our proposed copula cannot be written in closed form, the likelihood of the UCSV can be computed using Bayesian methods. A practical insight is that parameterizing the VAR for $\{\boldsymbol{\mu}_t\}$ in terms of its unique partial autocorrelations $\boldsymbol{\gamma}$ and variances $\boldsymbol{d}$, improves the efficiency of the MCMC scheme greatly. The prior for $\boldsymbol{\gamma}$ is similar to that induced by a Minnesota prior, and they can even be matched exactly by simulation.

The second contribution is to demonstrate that the copula time series model can improve real-time macroeconomic forecasts. The margins of many stationary time series models used in practice are inconsistent with those observed empirically. For the quarterly data examined here, the nonparametric margins in the copula model not only correct this inaccuracy, but also generate predictive distributions that have time-varying quantiles and higher order moments. Adrian et al. (2017) observe an increase in downside risk in GDP growth during recessions, conditional on financial and economic conditions. We observe the same feature in the real-time forecasts of GDP growth from our model, but these are joint with the other economic and financial variables. The current quarter GDP growth density forecasts are significantly more accurate than those from the benchmark models. Our empirical analysis also suggests asymmetries exist in the real-time forecast densities of other variables too. For example, there is upside risk in inflation, and post-recession forecasts of the interest rate have heavy positive upper tails – a necessary feature of a zero lower bound. While no one multivariate time series model dominates in forecasting accuracy for all variables and at all horizons, our model is the most accurate overall for inflation, the interest rate and current quarter GDP. The UCSV-C model also dominates the UCSV and VAR-C models, both of which it nests.

Last, we mention some directions for future research. While the MCMC estimator evaluates the exact posterior, the sampling scheme is slow for larger datasets. For these cases, alternatives such as simulated method of moments (Oh and Patton, 2013) and variational (Ormerod and Wand, 2010) estimators can be investigated. To extend our proposed copula to higher dimensions, factor structures for the pseudo time series can be considered, similar to the static factor model suggested by Oh and Patton (2017). Last, we note that the UCSV is Markov. The construction of inversion copulas from non-Markov stationary time series models has yet to be considered.

Appendix A

This appendix outlines how to compute $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$ from $\{\boldsymbol{\gamma}, \boldsymbol{d}\}$ for the stationary Gaussian VAR(p) model for $\{\boldsymbol{\mu}_t\}$. Let $\boldsymbol{\mu}_- = (\boldsymbol{\mu}'_{t-1}, \ldots, \boldsymbol{\mu}'_{t-p})'$, then $\{\boldsymbol{\gamma}, \boldsymbol{d}\}$ are the unique partial correlations and marginal variances of the $(p+1)r \times (p+1)r$ covariance matrix $V = \text{Var}((\boldsymbol{\mu}_t', \boldsymbol{\mu}_-')')$, from which V can be computed using the recursions of Yule; for example, see Daniels and Pourahmadi (2009). Then, if $V_1 = \text{Var}(\boldsymbol{\mu}_t)$, $V_{1,2} = \text{cov}(\boldsymbol{\mu}_t, \boldsymbol{\mu}_-)$ and $V_2 = \text{Var}(\boldsymbol{\mu}_-)$, the conditional distribution $\boldsymbol{\mu}_t|\boldsymbol{\mu}_- \sim N\left(V_{1,2}V_2^{-1}\boldsymbol{\mu}_-, V_1 - V_{1,2}V_2^{-1}V_{1,2}'\right)$. Matching these conditional moments with those from the usual specification of a VAR(p), gives $\Sigma_\epsilon = V_1 - V_{1,2}V_2^{-1}V_{1,2}'$ and $[B_1|\cdots|B_p] = V_{1,2}V_2^{-1}$.

Appendix B

This appendix outlines the Bayesian nonparametric constrained density estimator (labelled BNP). This represents an unknown density as an infinite mixture $f(x) = \sum_{j=1}^{\infty} w_j f_{\text{ker}}(x|\boldsymbol{\theta}_j)$. Here, w_j is the weight assigned to component j, and $f_{\text{ker}}(x|\boldsymbol{\theta}_j)$ is called the kernel density. We follow closely the approach of Kalli et al. (2011), but use a Gaussian density constrained to $A = \{x : x_L \leq x \leq x_U\}$ for the kernel density. For all five variables, the upper bound $x_U = \max(x) + \sigma_x$, where σ_x is the sample standard deviation of x. The lower bounds of the interest rate and stock market volatility are set to zero, while for the remaining variables $x_L = \min(x) - \sigma_x$. Estimation involves augmenting the mixture model with a 'slicing variable' ω, so that

$$f(x,\omega) = \sum_{j=1}^{\infty} \mathcal{I}(\omega < w_j)\mathcal{I}(x \in A)\phi(x;\mu_j,\sigma_j^2)/[\Phi(x_U;\mu_j,\sigma_j^2) - \Phi(x_L;\mu_j,\sigma_j^2)]\,,$$

and writing the weights as $w_j = \nu_j \prod_{k=1}^{j-1}(1-\nu_k)$, which is known as a 'stick breaking' representation. An efficient MCMC sampling scheme that we use to estimate the BNP constrained density estimator is outlined in Part B of the Online Appendix. As with the other two marginal models, we fit this estimator to each variable and vintage separately.

References

Adrian, T., Boyarchencko, N., and Domenico, G. (2017). Vunerable Growth. *Federal Research Bank of New York Staff Reports*, (794).

Bai, J. and Ng, S. (2005). Tests for Skewness, Kurtosis, and Normality for Time Series Data. *Journal of Business & Economic Statistics*, 23(1):49–60.

Ball, L., Jalles, J. T., and Loungani, P. (2015). Do forecasters believe in Okun's Law? An assessment of unemployment and output forecasts. *International Journal of Forecasting*, 31(1):176–184.

Bauer, M. D. and Rudebusch, G. D. (2016). Monetary Policy Expectations at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(7):1439–1465.

Beare, B. K. and Seo, J. (2015). Vine Copula Specifications for Stationary Multivariate Markov Chains. *Journal of Time Series Analysis*, 36(2):228–246.

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

Brechmann, E. C. and Czado, C. (2015). COPAR– multivariate time series modeling using the copula autoregressive model. *Applied Stochastic Models in Business and Industry*, 31(4):495–514.

Carriero, A., Clark, T. E., and Marcellino, M. (2016). Common Drifting Volatility in Large Bayesian VARs. *Journal of Business & Economic Statistics*, 34(3):375–390.

Chan, J. C. and Jeliazkov, I. (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2):101–120.

Chiu, C.-W. J., Mumtaz, H., and Pinter, G. (2016). Bayesian Vector Autoregressions with Non-Gaussian Shocks. Technical report, Centre for Research in Macroeconomics and Finance.

Clark, T. E. (2011). Real-Time Density Forecasts from Bayesian Vector Autoregressions with Stochastic Volatility. *Journal of Business & Economic Statistics*, 29(3):327–341.

Clark, T. E. and Ravazzolo, F. (2015). Macroeconomic Forecasting Performance under Alternative Specifications of Time-Varying Volatility. *Journal of Applied Econometrics*, 30(4):551–575.

Cogley, T. and Sargent, T. J. (2005). Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302.

Creal, D. D. and Tsay, R. S. (2015). High dimensional dynamic stochastic copula models. *Journal of Econometrics*, 189(2):335–345.

Croushore, D. and Stark, T. (2001). A real-time data set for macroeconomists. *Journal of Econometrics*, 105(1):111–130.

D'Agostino, A., Gambetti, L., and Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, 28(1):82–101.

Daniels, M. J. and Pourahmadi, M. (2009). Modeling covariance matrices via partial autocorrelations. *Journal of Multivariate Analysis*, 100(10):2352–2363.

De Lira Salvatierra, I. and Patton, A. J. (2015). Dynamic copula models and high frequency data. *Journal of Empirical Finance*, 30:120–135.

Demarta, S. and McNeil, A. J. (2005). The t Copula and Related Copulas. *International Statistical Review*, 73(1):111–129.

Garratt, A., Mitchell, J., and Vahey, S. P. (2014). Measuring output gap nowcast uncertainty. *International Journal of Forecasting*, 30(2):268–279.

Giglio, S., Kelly, B., and Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics*, 119(3):457–471.

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

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291.

Jensen, M. J. and Maheu, J. M. (2010). Bayesian semiparametric stochastic volatility modeling. *Journal of Econometrics*, 157(2):306–316.

Johannsen, B. K. and Mertens, E. (2016). A Time Series Model of Interest Rates with the Effective Lower Bound. *FEDS Working Paper*, 2016-033.

Jungbacker, B. and Koopman, S. J. (2007). Monte Carlo Estimation for Nonlinear Non-Gaussian State Space Models. *Biometrika*, 94(4):827–839.

Kalli, M. and Griffin, J. E. (2018). Bayesian nonparametric vector autoregressive models. *Journal of Econometrics*, 203(2):267–282.

Kalli, M., Griffin, J. E., and Walker, S. G. (2011). Slice sampling mixture models. *Statistics and computing*, 21(1):93–105.

Koop, G., Korobilis, D., et al. (2010). Bayesian Multivariate Time Series Methods for Empirical Macroeconomics. *Foundations and Trends® in Econometrics*, 3(4):267–358.

Krüger, F., Clark, T. E., and Ravazzolo, F. (2017). Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts. *Journal of Business & Economic Statistics*, 35(3):470–485.

Krüger, F. and Nolte, I. (2015). Disagreement versus uncertainty: Evidence from distribution forecasts. *Journal of Banking & Finance*, 72(Supplement):S5172–S5186.

Litterman, R. B. (1986). Forecasting With Bayesian Vector Autoregressions: Five Years of Experience. *Journal of Business & Economic Statistics*, 4(1):25–38.

Loaiza-Maya, R., Smith, M. S., and Maneesoonthorn, W. (2018). Time Series Copulas for Heteroskedastic Data. *Journal of Applied Econometrics*, 33(3):332–354.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis.* Springer Science & Business Media.

McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques and Tools.* Princeton Series in Finance.

Mertens, E. (2016). Measuring the Level and Uncertainty of Trend Inflation. *Review of Economics and Statistics*, 98(5):950–967.

Nelsen, R. B. (2006). *An Introduction to Copulas (Springer Series in Statistics).* Springer-Verlag New York, Inc., Secaucus, NJ, USA.

Oh, D. H. and Patton, A. J. (2013). Simulated method of moments estimation for copula-based multivariate models. *Journal of the American Statistical Association*, 108(502):689–700.

Oh, D. H. and Patton, A. J. (2017). Modeling Dependence in High Dimensions With Factor Copulas. *Journal of Business & Economic Statistics*, 35(1):139–154.

Ormerod, J. T. and Wand, M. P. (2010). Explaining variational approximations. *The American Statistician*, 64(2):140–153.

Patton, A. J. (2012). A review of copula models for economic time series. *Journal of Multivariate Analysis*, 110:4–18.

Romer, C. D. and Romer, D. H. (2000). Federal Reserve Information and the Behavior of Interest Rates. *The American Economic Review*, 90(3):429–457.

Scharth, M. and Kohn, R. (2016). Particle efficient importance sampling. *Journal of Econometrics*, 190(1):133–147.

Shimazaki, H. and Shinomoto, S. (2010). Kernel bandwidth optimization in spike rate estimation. *J. Comput. Neurosci.*, 29(1-2):171–182.

Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, 31(3):815 – 833.

Smith, M. S., Gan, Q., and Kohn, R. J. (2012). Modelling dependence using skew t copulas: Bayesian inference and applications. *Journal of Applied Econometrics*, 27(3):500–522.

Smith, M. S. and Maneesoonthorn, W. (2018). Inversion copulas from nonlinear state space models with an application to inflation forecasting. *International Journal of Forecasting*, 34(3):389–407.

Smith, M. S. and Vahey, S. P. (2016). Asymmetric Forecast Densities for U.S. Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence. *Journal of Business and Economic Statistics*, 34(3):416–434.

Stock, J. H. and Watson, M. W. (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and banking*, 39(s1):3–33.

				Σ_η				
	a_i	$\bar{h}_i$	d_i	$\Delta\ln$(GDP)	Infl	UR	IR	VXO
$\Delta\ln$(GDP)	0.763	-1.299	0.524	0.418	0.037	0.165	0.208	0.046
	(0.29,0.96)	(-1.93,-0.83)	(0.38,0.66)	(0.05,1.12)	(-0.22, 0.33)	(0.01,0.40)	(0.05,0.45)	(-0.13,0.25)
Infl	0.746	-2.408	0.821	0.037	0.526	0.067	0.006	0.237
	(0.43,0.94)	(-3.03,-1.87)	(0.70,0.90)	(-0.22,0.33)	(0.11, 1.20)	(-0.19,0.31)	(-0.19,0.18)	(0.02,0.64)
UR	0.829	-5.044	0.978	0.165	0.067	0.236	0.156	0.051
	(0.31,0.99)	(-5.74,-3.51)	(0.91,0.99)	(0.01,0.40)	(-0.19, 0.31)	(0.03,0.74)	(0.04,0.34)	(-0.07,0.21)
IR	0.928	-4.979	0.975	0.208	0.006	0.156	0.241	-0.009
	(0.83,0.98)	(-5.78,-3.93)	(0.93,0.99)	(0.05,0.45)	(-0.19, 0.18)	(0.04,0.34)	(0.07,0.51)	(-0.19,0.13)
VXO	0.772	-2.299	0.797	0.046	0.237	0.051	-0.009	0.489
	(0.33,0.97)	(-3.15,-1.67)	(0.67,0.90)	(-0.13,0.25)	(0.02, 0.64)	(-0.07,0.21)	(-0.19,0.13)	(0.04,1.44)

Table 1: **Estimates of the copula parameters from the 2016Q2 vintage data**. The posterior means and 90% probability intervals are reported for the copula parameters $\{\boldsymbol{d}, \boldsymbol{a}, \Sigma_\eta\}$, along with $\bar{\boldsymbol{h}}$, from the UCSV-C inversion copula model. Estimates of the partial correlations $\boldsymbol{\gamma}$, and alternative parameterization $\{B_1, \ldots, B_4, \Sigma_\epsilon\}$, can be found in the Online Appendix.

		$\Delta\ln(\text{GDP})_t$	Infl_t	UR_t	IR_t	VXO_t
$R(0)$	$\Delta\ln(\text{GDP})_t$	1.000	-0.099	-0.039	0.052	-0.180
	Infl_t	-0.099	1.000	0.041	0.568	-0.021
	UR_t	-0.039	0.041	1.000	-0.039	0.094
	IR_t	0.052	0.568	-0.039	1.000	0.040
	VXO_t	-0.180	-0.021	0.094	0.040	1.000
$R(1)$	$\Delta\ln(\text{GDP})_{t-1}$	0.327	-0.029	-0.148	0.100	-0.148
	Infl_{t-1}	-0.129	0.795	0.080	0.579	-0.033
	UR_{t-1}	0.065	0.015	0.958	-0.053	0.042
	IR_{t-1}	-0.001	0.547	-0.008	0.965	0.072
	VXO_{t-1}	-0.176	-0.023	0.151	0.006	0.709
$R(4)$	$\Delta\ln(\text{GDP})_{t-4}$	0.039	0.044	-0.260	0.159	-0.059
	Infl_{t-4}	-0.099	0.682	0.205	0.581	-0.059
	UR_{t-4}	0.167	0.008	0.769	-0.043	-0.065
	IR_{t-4}	-0.014	0.458	0.115	0.882	0.139
	VXO_{t-4}	-0.070	-0.025	0.253	-0.067	0.467

Table 2: **Estimates of level dependence from the 2016Q2 vintage data**. The posterior means of the matrices of Spearman pairwise correlations $R(0)$, $R(1)$ and $R(4)$ are reported for the fitted UCSV-C inversion copula. These measure dependence in the level of the time series $\{\boldsymbol{Y}_t\}$, and are computed by simulation. Estimates of $R(2)$, $R(3)$ can be found in the Online Appendix.

	VARSV					UCSV-C				
	$\Delta\ln(\text{GDP})_t$	Infl_t	UR_t	IR_t	VXO_t	$\Delta\ln(\text{GDP})_t$	Infl_t	UR_t	IR_t	VXO_t
	$R^V(0)$					$R^V(0)$				
$\Delta\ln(\text{GDP})_t$	1.000	0.074	0.247	0.174	0.115	1.000	0.010	0.092	0.053	0.015
Infl_t	0.074	1.000	0.181	0.206	-0.030	0.010	1.000	0.014	0.067	0.040
UR_t	0.247	0.181	1.000	0.226	0.024	0.092	0.014	1.000	0.057	0.027
IR_t	0.174	0.206	0.226	1.000	0.013	0.053	0.067	0.057	1.000	-0.006
VXO_t	0.115	-0.030	0.024	0.013	1.000	0.015	0.040	0.027	-0.006	1.000
	$R^V(1)$					$R^V(1)$				
$\Delta\ln(\text{GDP})_{t-1}$	0.120	0.060	0.088	0.109	0.096	0.069	0.004	0.022	0.035	0.006
Infl_{t-1}	0.056	0.177	0.143	0.158	-0.036	0.008	0.267	0.013	0.064	0.022
UR_{t-1}	0.084	0.143	0.152	0.149	0.015	0.026	0.010	0.171	0.043	0.025
IR_{t-1}	0.108	0.169	0.157	0.273	0.005	0.034	0.071	0.036	0.252	-0.003
VXO_{t-1}	0.083	-0.032	0.013	0.005	0.227	0.006	0.028	-0.003	0.006	0.192

Table 3: **Estimates of volatility dependence from the 2016Q2 vintage data.** Estimates of the matrices of pairwise Spearman correlations in the volatilities ($R^V(k)$) of the innovations. Results are reported separately for the VARSV on the left hand side, and the UCSV-C model on the right hand side. The top rows gives cross-sectional dependence ($k = 0$), and the bottom rows give first order serial dependence ($k = 1$). For example, for the UCSV-C model, the correlation between $V(\text{IR}_{t-1})$ and $V(\text{Infl}_t)$ is $\rho^V_{\text{Infl,IR}}(1) = 0.071$, which is positive volatility 'spillover'.

	Mean CRPS values for forecasts over 1975Q1 – 2016Q2							
	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$
	GDP Growth				Inflation			
Time Series Model								
VAR	0.4511**	0.4420**	0.4306	0.4240	0.1750	0.1954	0.2304*	0.3259*
VARSV	0.4014*	0.4003	0.3938^{++}	0.4022	0.1688	0.1867	0.2222	0.3245*
UCSV	0.4388**	0.4316**	0.4242	0.4159	0.1687	0.1859	0.2099	0.2910
VAR-C-KDE	0.3964*	0.4022	0.4126	0.4132	0.1767**	0.1931**	0.2078	0.2772
UCSV-C-BNP	0.3851	0.3999	0.4111^{++}	0.4033^{++}	0.1652^{++}	0.1797^{+}	0.1982	0.2651
UCSV-C-Skt	0.3841	0.3997	0.4188	0.4128	0.1660^{+}	0.1815	0.1964	0.2604
UCSV-C-KDE	*0.3845*	*0.4002*	*0.4158*	*0.4103*	*0.1685*	*0.1835*	*0.2010*	*0.2663*
	Unemployment				Interest Rate			
Time Series Model								
VAR	0.1485	0.2631	0.4632	0.6916	0.4020**	0.6450**	0.9643	1.5540
VARSV	0.1372^{+}	0.2466	0.4447	0.6924	0.3382	0.5835*	0.9358	1.5551
UCSV	0.1531	0.2654	0.4688	0.7110	0.4030**	0.6074**	0.9233	1.5178
VAR-C-KDE	0.1498	0.2675*	0.4805	0.7355	0.3128	0.5449	0.8597	1.4186
UCSV-C-BNP	0.1483	0.2633	0.4862*	0.7568	0.3354	0.5408	0.8603	1.4335
UCSV-C-Skt	0.1409^{++}	0.2490	0.4627	0.7258	0.3462**	0.5647**	0.9035**	1.4806**
UCSV-C-KDE	*0.1479*	*0.2606*	*0.4763*	*0.7279*	*0.3298*	*0.5384*	*0.8622*	*1.4348*

	RMSE of forecasts over 1975Q1 – 2016Q2							
	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$
	GDP Growth				Inflation			
Time Series Model								
VAR	0.8139**	0.7752**	0.7366	0.7203	0.3154	0.3666	0.4283	0.6202
VARSV	0.7496**	0.7360	0.7137^{++}	0.7307	0.3080	0.3532	0.4272	0.6233
UCSV	0.7871**	0.7568**	0.7305	0.7097	0.3032	0.3418	0.3780	0.5264
VAR-C-KDE	0.7272	0.7299	0.7393	0.7327	0.3121*	0.3513**	0.3824	0.5180
UCSV-C-BNP	0.7155	0.7287	0.7422	0.7242	0.2955^{+}	0.3281	0.3665	0.4939
UCSV-C-Skt	0.7108	0.7339	0.7601	0.7415	0.3016	0.3353	0.3690	0.4925
UCSV-C-KDE	*0.7065*	*0.7242*	*0.7426*	*0.7265*	*0.3029*	*0.3352*	*0.3746*	*0.4995*
	Unemployment				Interest Rate			
Time Series Model								
VAR	0.2759	0.5009	0.9069	1.3638	0.7936	1.2183	1.7640	2.7282
VARSV	0.2723	0.4956	0.8971	1.3663	0.7506	1.1538	1.6770	2.6030
UCSV	0.2836	0.5062	0.9188	1.3916	0.9787	1.2417	1.6664	3.1076
VAR-C-KDE	0.2715	0.4905	0.9078	1.3935	0.7287	1.1555	1.6572	2.5164
UCSV-C-BNP	0.2769	0.4977	0.9288**	1.4321	0.8141	1.1694	1.6757	2.5243
UCSV-C-Skt	0.2708	0.4847	0.9021	1.3823	0.8185	1.1959*	1.7129**	2.5629
UCSV-C-KDE	*0.2780*	*0.4901*	*0.9076*	*1.3860*	*0.7944*	*1.1585*	*1.6634*	*2.5112*

Table 4: **Real-time (out-of-sample) forecast accuracy of the multivariate time series models over 1975Q1 – 2016Q2.** The mean CRPS values of the density forecasts (top) and RMSE of the point forecasts (bottom) are given for the four macroeconomic variables and different models. Lower values indicate increased accuracy. Results for forecasts made $h = 1, 2, 4$ and 8 quarters ahead are given in columns. As a guide, results of a Harvey et al. (1997) small-sample adjustment of the Diebold-Mariano test are reported. This is a two-sided test of the equality of the mean score of the UCSV-C-KDE with each of the other models. Rejection of the null hypothesis of equal means at the 10% or 5% level is denoted by '*' or '**' if favorable to the UCSV-C-KDE model, or by '$^{+}$' or '$^{++}$' if unfavorable.

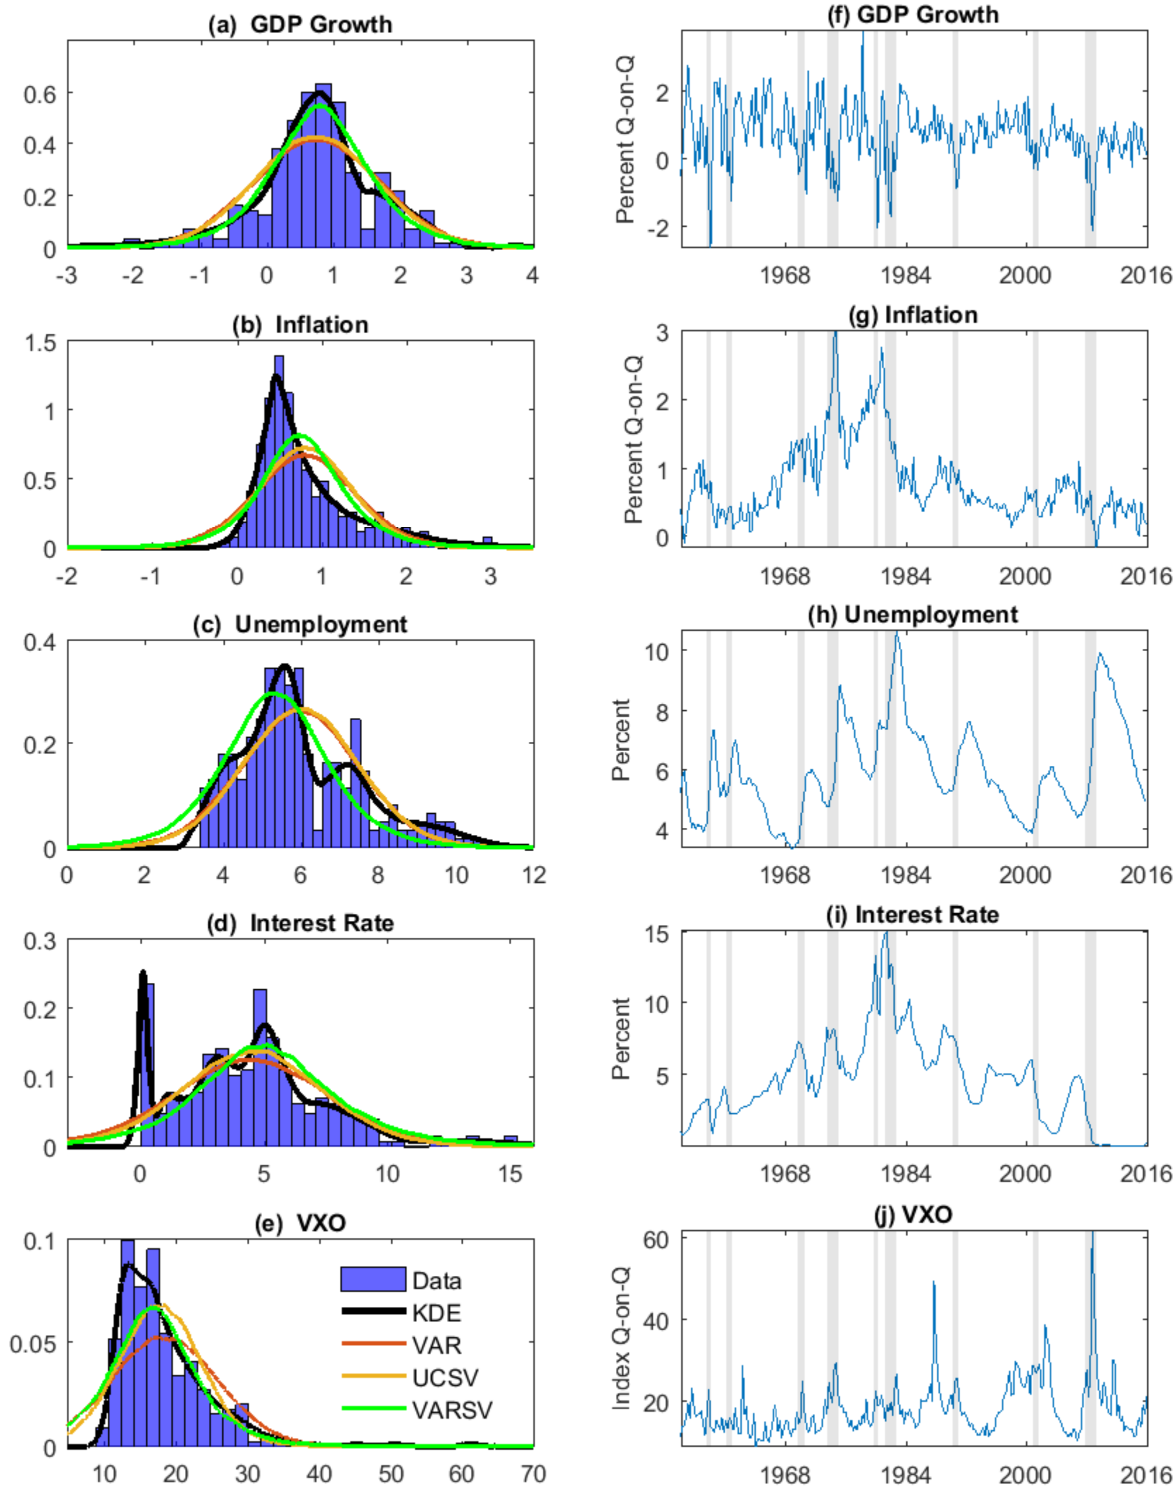

Figure 1: **Marginal densities and time series plots of the 2016Q2 vintage data.** Panels (a–e) give the histograms of the five variables, along with the adaptive kernel density estimates (black). The (necessarily symmetric) marginal densities from the fitted VARSV (green), VAR (red) and UCSV (yellow) are also plotted. Panels (f–j) give the time series plots of the data, where quarters that are classified as recessions by the NBER are shaded.

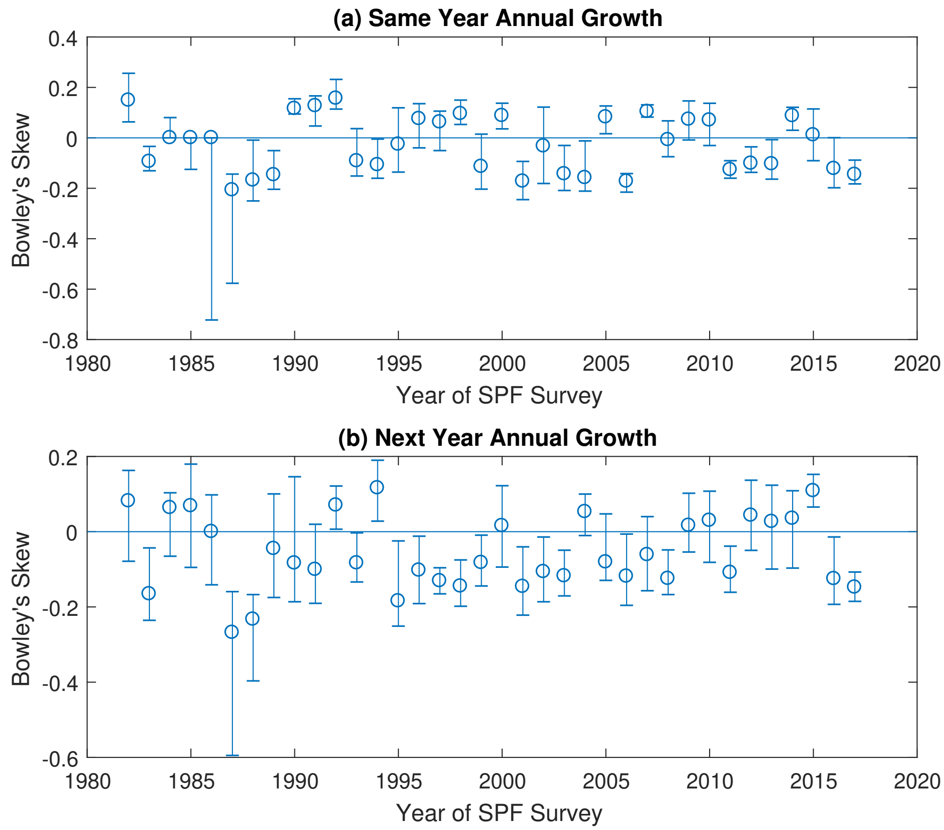

Figure 2: **Skew coefficients of the probabilistic GDP growth forecasts in the Q1 SPF reports.** The circles are Bowley's skew coefficients computed from the predictive distributions of "annual average over annual average" GNP/GDP growth, given in the first quarter SPF reports from 1982 to 2017. Panel (a) is for the distribution of growth in the same year as the report, and panel (b) in the next year. The vertical bands depict bootstrap 95% confidence intervals, computed by re-sampling the individual panelist probability distributions, which can be found on the Philadelphia Federal Reserve website.

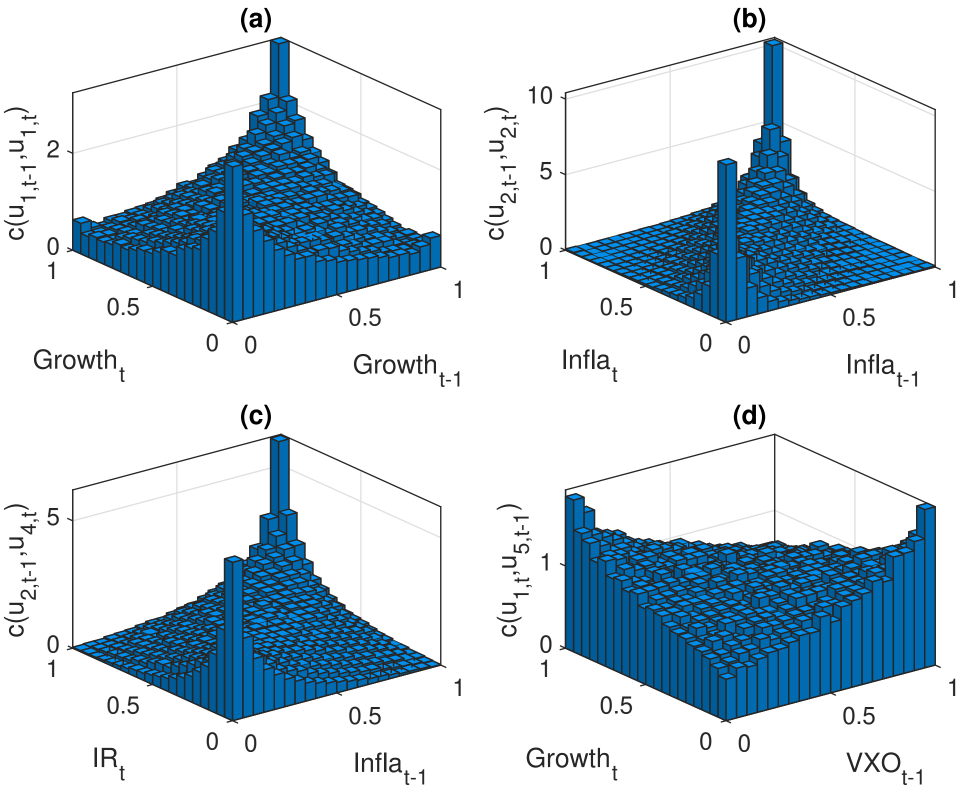

Figure 3: **Bivariate slices of the UCSV-C inversion copula fitted to the 2016Q2 vintage data.** The four panels correspond to the variable pairs (a) $(\Delta(\ln(GDP)_t, \Delta(\ln(GDP)_{t-1})$, (b) $(\text{Infl}_t, \text{Infl}_{t-1})$, (c) $(\text{IR}_t, \text{Infl}_{t-1})$ and (d) $(\Delta(\ln(GDP)_t), \text{VXO}_{t-1})$. The slices are constructed via simulation from the fitted copula, and presented as bivariate histograms.

	GDP Growth	Inflation	Unemployment	Interest Rate	Stock Market Volatility
VAR	0.4356	0.1290	0.1401	0.3261	2.2344
VARSV	0.4151	0.1234	0.1321	0.2718	2.0020
UCSV	0.4226	0.1093	0.1239	0.2391	1.8798
VAR-C	0.4117	0.1293	0.1314	0.2696	2.0331
UCSV-C	0.3778	0.0980	0.0854	0.1945	1.7005

Table 5: **A 'goodness-of-fit' measure for the UCSV-C and benchmark models fit to the 2016Q2 vintage data.** The mean of the cumulative rank probability score (CRPS) values computed from the one-step-ahead predictive distributions is reported for each variable and multivariate time series model. Lower values indicate increased goodness of fit. If the multivariate time series models provided perfect forecast densities, this measure would be zero. As a guide, Harvey et al. (1997) small-sample adjusted Diebold-Mariano tests indicate that the predictive densities from the UCSV-C model are significantly more accurate than those from all other models at the 5% level for all variables. (This is a two-sided test of the equality of the mean score of the UCSV-C model with each of the other models.)

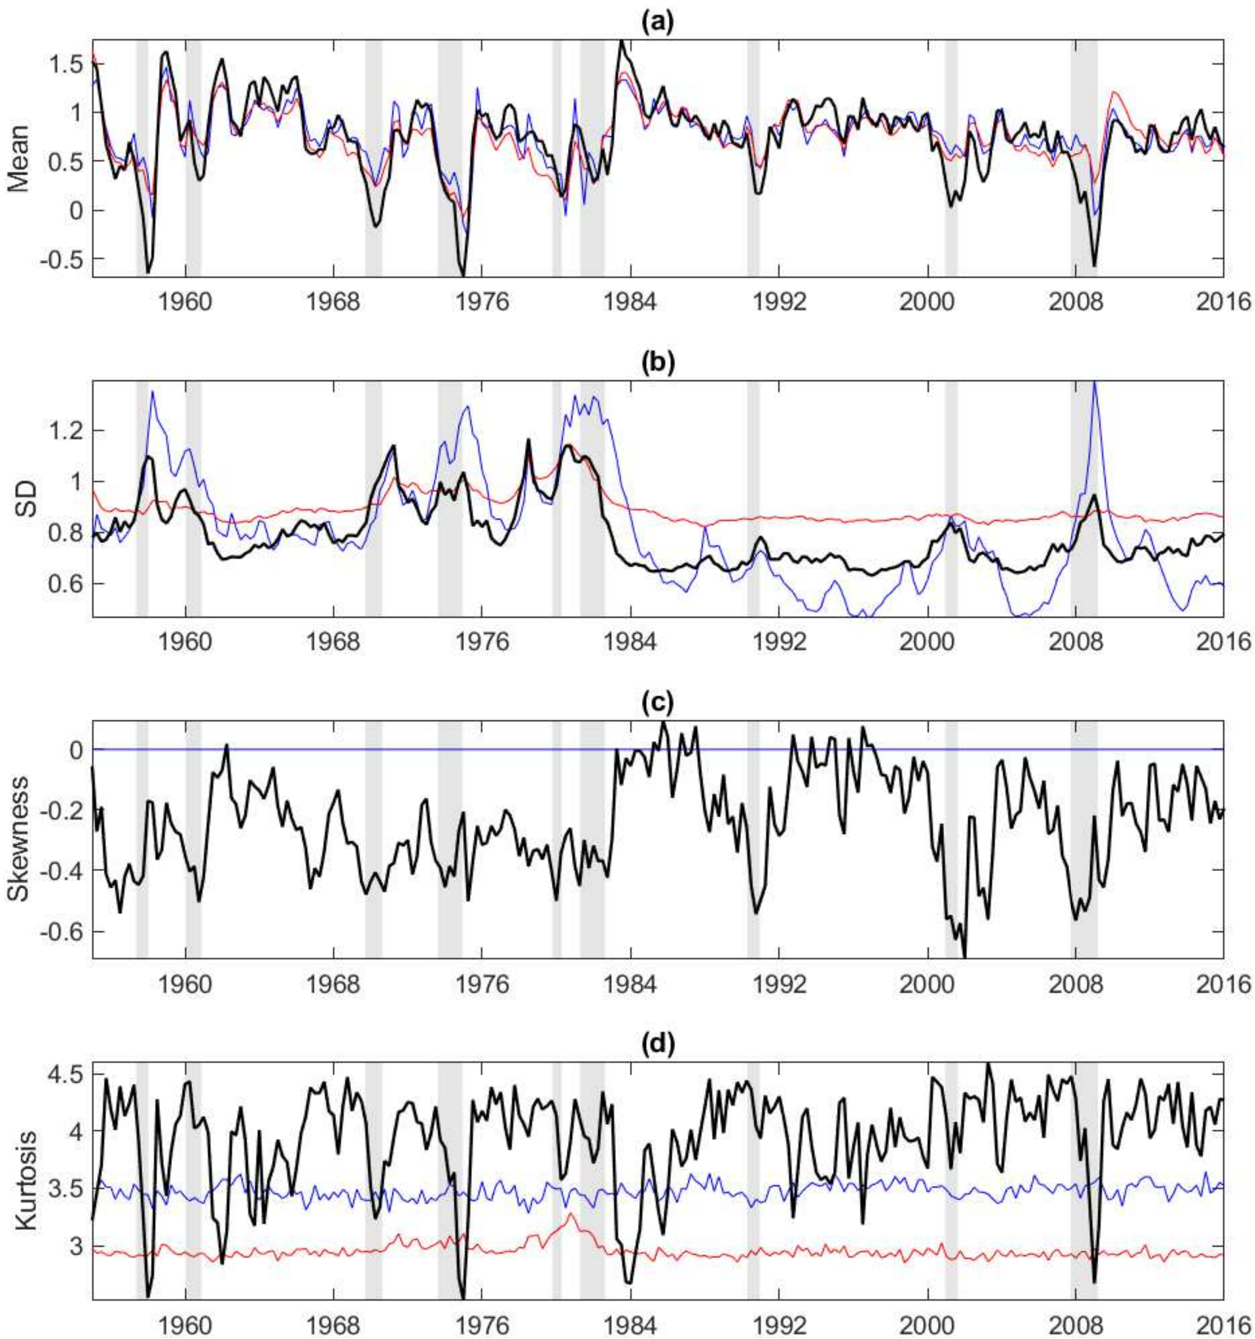

Figure 4: **Moments of the predictive distributions of GDP Growth from the 2016Q2 vintage data.** Panels (a–d) plot the mean, standard deviation, Pearson's skew coefficient and kurtosis from the one quarter ahead (ie. $h = 1$) in-sample predictive distributions of GDP growth. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

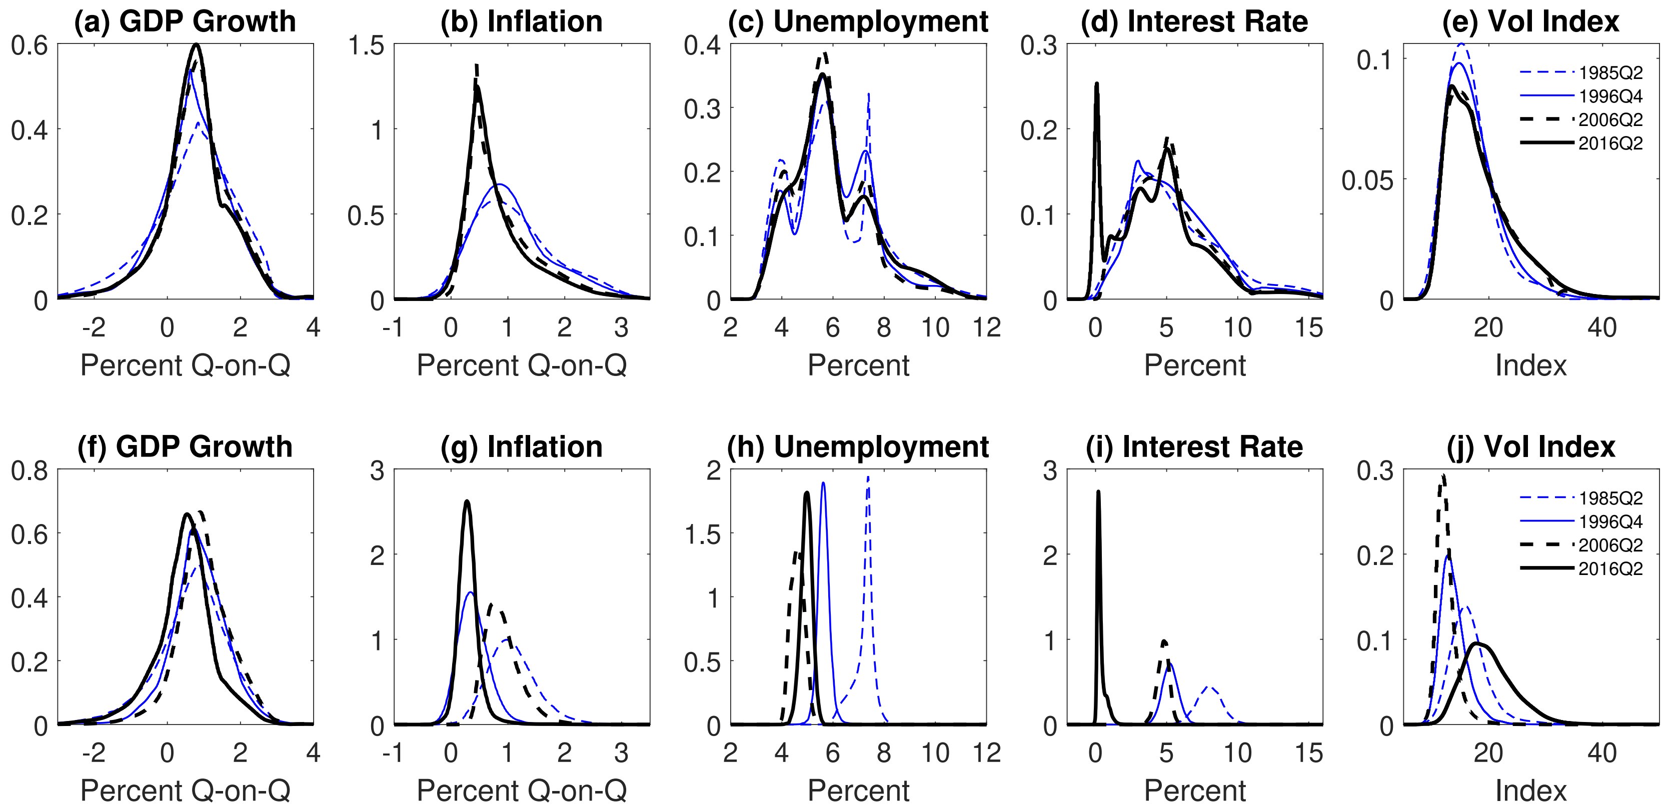

Figure 5: **Estimated marginal and current quarter real-time predictive densities for four vintages.** Panels (a–e) present the adaptive kernel density estimates for the five variables, and the four vintages 1985Q2 (thin blue dashed line), 1996Q4 (thin blue solid line), 2006Q2 (thick black dashed line) and 2016Q2 (thick black solid line). Panels (f–j) present the real-time (out-of-sample) current quarter predictive densities (ie. where $h = 1$) from the same four vintages from the UCSV-C multivariate time series model.

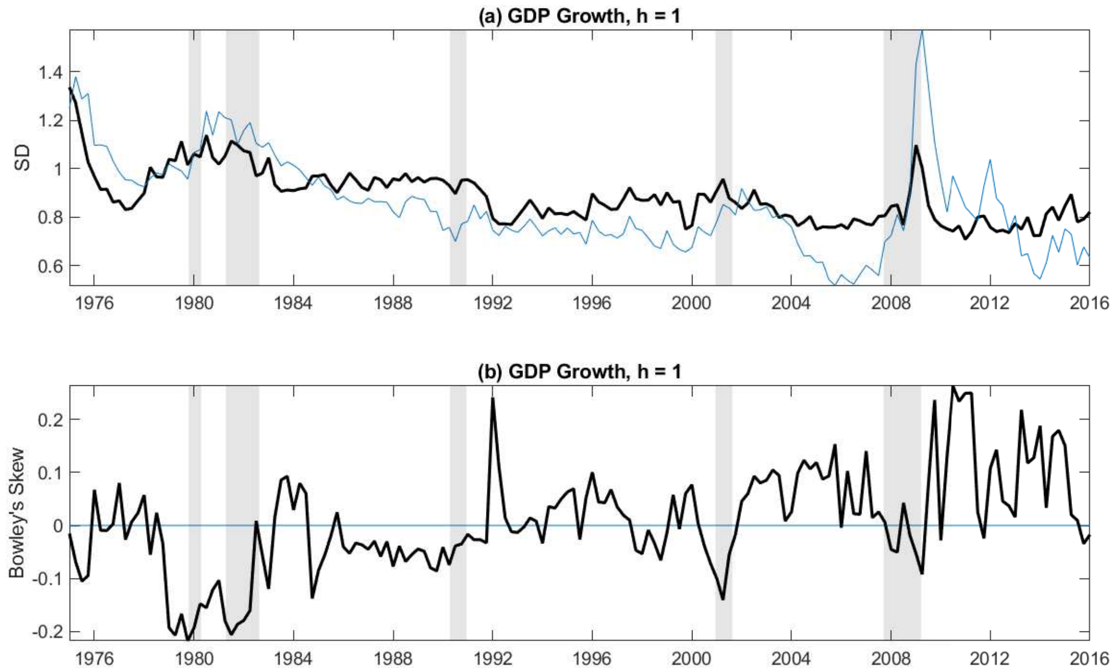

Figure 6: **Uncertainty and asymmetry in real-time predictions of GDP growth.** Panel (a) gives the standard deviation, and panel (b) Bowley's skew coefficient, for the current quarter (ie. $h = 1$) real-time (out-of-sample) predictive distributions of GDP growth from the UCSV-C (thick black line) and VARSV (thin blue line) models. The plots are for vintages 1975Q1 – 2016Q2.

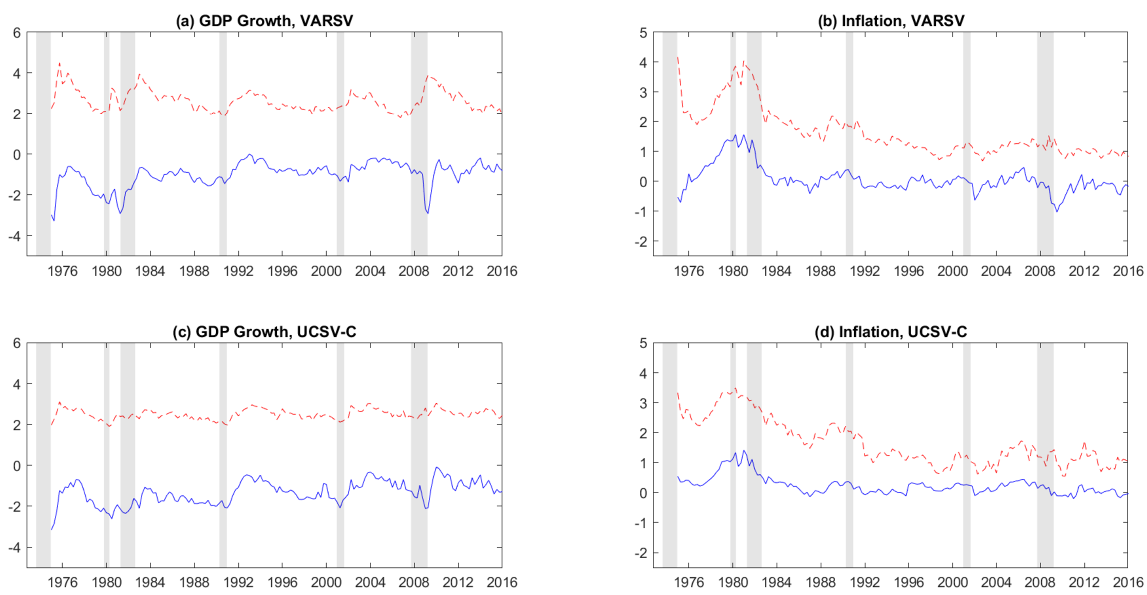

Figure 7: **Asymmetry in real-time inflation and GDP growth predictions.** Real-time (out-of-sample) current quarter ($h = 1$) expected shortfall (blue line) and upper tail equivalent (red line), for (a) GDP growth & VARSV, (b) inflation & VARSV, (c) GDP growth & UCSV-C, and (d) inflation & UCSV-C.

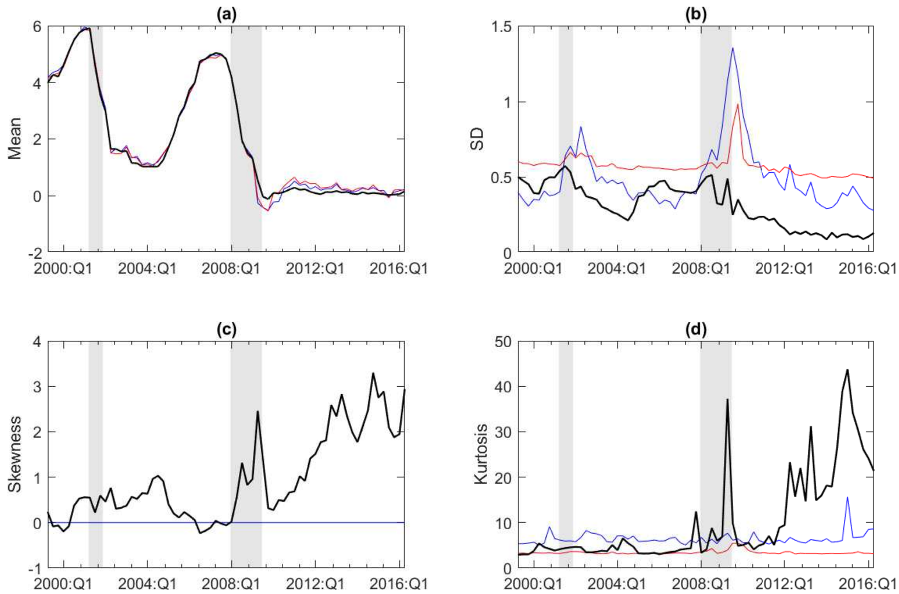

Figure 8: **Impact of the Great Recession on real-time interest rate predictions.** Panels (a–d) give the first four moments of the real-time (out-of-sample) current quarter (ie. $h = 1$) predictive distributions of the interest rate between 2000Q1 and 2016Q2. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

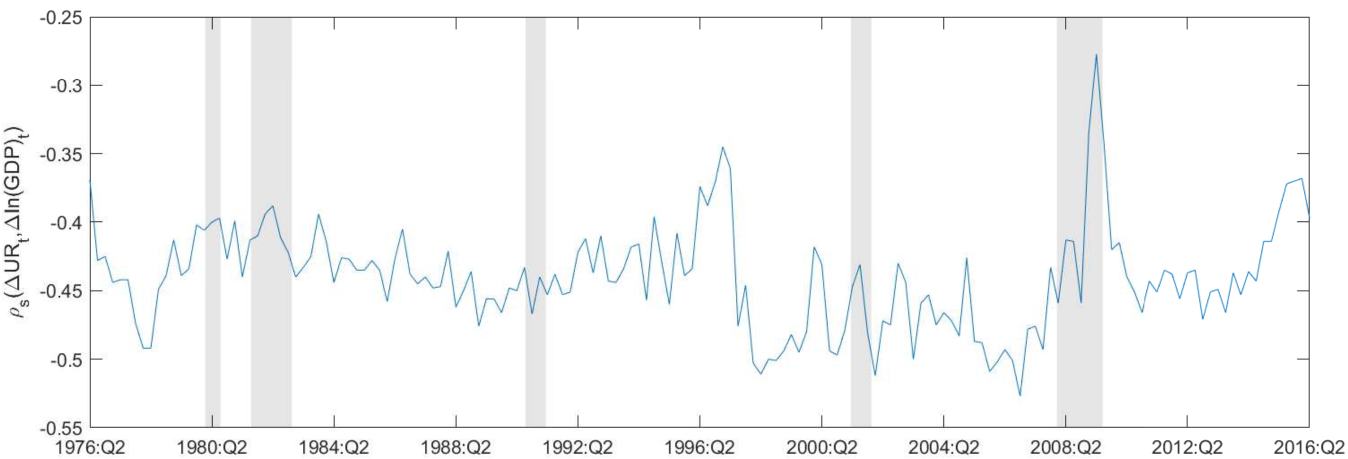

Figure 9: **Real-time current quarter forecasts of Okun's (Spearman) correlations.** Plot of the Spearman correlations between real-time (out-of-sample) current quarter forecasts of $\Delta \mathrm{UR}_{T+1} = \mathrm{UR}_{T+1} - \mathrm{UR}_T$ and $\Delta \ln(\mathrm{GDP})_{T+1}$, for vintages 1975Q1 to 2016Q2. Each correlation is the sample Spearman correlation coefficient of a large number of Monte Carlo iterates generated from the current quarter predictive distributions. Quarters classified as recessions by the NBER are shaded. Plots of the bivariate densities for 8 vintages can be found in the Online Appendix.

Online Appendix for 'Real-Time Macroeconomic Forecasting with a Heteroskedastic Inversion Copula Model'

This Online Appendix has four parts:

Part A: A technical appendix that outlines the MCMC sampling scheme used to estimate the copula parameters Ψ of the UCSV-C model in greater detail.

Part B: A technical appendix that outlines the MCMC sampling scheme used to estimate the Bayesian nonparametric constrained density (BNP) in Appendix B of the paper.

Part C: Provides additional empirical results for UCSV-C model fit to the 2016Q2 vintage data.

Table A2: Estimates of the all the partial correlations $\boldsymbol{\gamma}$.

Table A3: Estimates of Σ_ϵ.

Table A4: Estimates of $B_1, \ldots, B_4$.

Table A5: Estimates of level dependence $R(0), \ldots, R(4)$.

Figure A1: Time series plots of the estimates of $\boldsymbol{\mu}, \boldsymbol{h}$.

Figures A2 – A5: First four moments of the one-quarter-ahead (in-sample) predictive distributions from the UCSV-C, VARSV and UCSV.

Part D: Provides additional empirical results for the real-time out-of-sample forecasting study.

Table A6: Log-scores results for the real-time forecasting comparison.

Figure A6: Asymmetry in inflation and GDP growth prediction for $h = 4$.

Figure A7: Bowley's skew coefficient for the current quarter ($h = 1$) forecast densities.

Figure A8: Pearson's skew coefficient for the current quarter ($h = 1$) forecast densities.

Figure A9: Bowley's skew coefficient for the forecast densities with $h = 4$.

Figure A10: Pearson's skew coefficient for the forecast densities with $h = 4$.

Figures A11 – A12: Standard deviation of the $h = 1$ and $h = 4$ forecast densities.

Figures A13 – A17: Rolling fan charts for real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for all five variables.

Figure A18: Contour plots of the bivariate current quarter density forecasts for the change in unemployment and GDP growth. Results are given for 8 vintages.

Part A: Technical Appendix

	Data	$U_{i,t} = G_i(Y_{i,t})$	Copula	$Z_{i,t} = F_i^{-1}(U_{i,t})$	Pseudo
Time Series	$\{\boldsymbol{Y}_t\}_{t=1}^T$	$\longrightarrow$	$\{\boldsymbol{U}_t\}_{t=1}^T$	$\longrightarrow$	$\{\boldsymbol{Z}_t\}_{t=1}^T$
Domain	$\boldsymbol{Y}_t \in \mathbb{R}^r$	$\longrightarrow$	$\boldsymbol{U}_t \in [0,1]^r$	$\longrightarrow$	$\boldsymbol{Z}_t \in \mathbb{R}^r$
Joint CDF	$F_Y(\boldsymbol{y})$	$\longrightarrow$	$C(\boldsymbol{u}\|\Psi)$	$\longrightarrow$	$F_Z(\boldsymbol{z}\|\Psi)$
Joint PDF	$f_Y(\boldsymbol{y}) =$ $c(\boldsymbol{u}) \prod_{t=1}^T \prod_{i=1}^r g_i(y_{i,t})$	$\longrightarrow$	$c(\boldsymbol{u}\|\Psi) = \frac{f_Z(\boldsymbol{z}\|\Psi)}{\prod_{i=1}^r \prod_{t=1}^T f_i(z_{i,t}\|\Psi)}$	$\longrightarrow$	$f_Z(\boldsymbol{z}\|\Psi) =$ $\int \prod_{t=1}^T \phi_r(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \boldsymbol{h}_t) \mathrm{d}(\boldsymbol{\mu}, \boldsymbol{h})$
Marginal CDF	$G_i(y_{i,t})$	$\longrightarrow$	Uniform	$\longrightarrow$	$F_i(z_{i,t}\|\Psi) = \int \Phi_1(z_{i,t}; 0, d_i + e^h)$ $\times \phi\left(h; \bar{h}, \frac{\bar{\sigma}_{i,i}^2}{1-a_i^2}\right) \mathrm{d}h$
Marginal PDF	$g_i(y_{i,t})$	$\longrightarrow$	Uniform	$\longrightarrow$	$f_i(z_{i,t}\|\Psi) = \int \phi_1(z_{i,t}; 0, d_i + e^h)$ $\times \phi\left(h; \bar{h}, \frac{\bar{\sigma}_{i,i}^2}{1-a_i^2}\right) \mathrm{d}h$

Table A1: **Depiction of the one-to-one transformations underlying the time series inversion copula model.** These transformations are between the multivariate time series processes on the data domain $\{\boldsymbol{Y}_t\}$, copula domain $\{\boldsymbol{U}_t\}$ and pseudo domain $\{\boldsymbol{Z}_t\}$. The marginal and joint density/distribution functions are given for each of the three time series.

To outline the sampling scheme to evaluate the augmented posterior with density $f(\Psi, \boldsymbol{\mu}, \boldsymbol{h}|\boldsymbol{y})$, first notice from Section 2, that the conditional likelihood is

$$f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) = \prod_{t=1}^{T} \phi_r\left(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \boldsymbol{h}_t\right) \prod_{i=1}^{r} \frac{g_i(y_{t,i})}{f_i(z_{i,t}|\Psi)} .$$

The density

$$f(\boldsymbol{\mu}|\Psi) = f(\boldsymbol{\mu}|\boldsymbol{d}, \boldsymbol{\gamma}) = \prod_{t=2}^{T} f(\boldsymbol{\mu}_t|\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_{t-1}, \boldsymbol{d}, \boldsymbol{\gamma}) f(\boldsymbol{\mu}_1|\boldsymbol{d}, \boldsymbol{\gamma}) ,$$

is the likelihood of a VAR(p) model, which can be computed by first evaluating $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$ from $\{\boldsymbol{d}, \boldsymbol{\gamma}\}$ as in Appendix A. The density

$$f(\boldsymbol{h}|\Psi) = f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) = \prod_{t=2}^{T} f(\boldsymbol{h}_t|\boldsymbol{h}_{t-1}, \boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) f(\boldsymbol{h}_1|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) ,$$

is that of a (sparse) VAR(1) model, and readily computed from $\{\boldsymbol{a}, \bar{\boldsymbol{h}}, \Sigma_\eta\}$, where $\bar{\boldsymbol{h}}$ is a function of $\boldsymbol{d}$ as outlined in the manuscript. The steps of the sampling scheme are outlined below.

Step 1: The conditional posterior

$$\begin{aligned} f(\boldsymbol{\mu}|\Psi, \boldsymbol{y}, \boldsymbol{h}) &\propto f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) f(\boldsymbol{\mu}|\boldsymbol{d}, \boldsymbol{\gamma}) \\ &\propto \prod_{t=1}^{T} \phi_r\left(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \boldsymbol{h}_t\right) f(\boldsymbol{\mu}|\boldsymbol{d}, \boldsymbol{\gamma}) . \end{aligned}$$

To generate from the density above, first map the parameters $\{\boldsymbol{d}, \boldsymbol{\gamma}\}$ to $\{B_1, \ldots, B_p, \Sigma_\epsilon\}$, then efficiently draw $\boldsymbol{\mu}$ using the precision sampler in Chan and Jeliazkov (2009).

Step 2: The conditional posterior

$$\begin{aligned} f(\boldsymbol{h}|\Psi, \boldsymbol{y}, \boldsymbol{\mu}) &\propto f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) \\ &\propto \prod_{t=1}^{T} \phi_r\left(\boldsymbol{z}_t; \boldsymbol{\mu}_t, \boldsymbol{h}_t\right) f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) . \end{aligned}$$

To sample $\boldsymbol{h}$ from the density above we employ the following approximating state space form:

$$\begin{aligned} \boldsymbol{z}_t^* &= \boldsymbol{h}_t + \boldsymbol{e}_t , \\ \boldsymbol{h}_t &= \bar{\boldsymbol{h}} + A(\boldsymbol{h}_{t-1} - \bar{\boldsymbol{h}}) + \boldsymbol{\eta}_t , \end{aligned}$$

with $\boldsymbol{z}_t^* = \left(z_{1,t}^*, \ldots, z_{r,t}^*\right)'$, $z_{i,t}^* = \log((z_{i,t} - \mu_{i,t})^2 + \bar{c})$, $\bar{c}$ is set to 0.001, $\boldsymbol{e}_t = \left(\log(\varepsilon_{1,t}^2), \ldots, \log(\varepsilon_{r,t}^2)\right)'$,

$\bar{\boldsymbol{h}} = \left(\bar{h}_1, \ldots, \bar{h}_r\right)'$, and $A = \text{diag}(a_1, \ldots, a_r)$. We proceed exactly as in Primiceri (2005), using the mixture of seven normals to approximate the distributions of the elements of $\boldsymbol{e}_t$, and then the method in Carter and Kohn (1994) to generate the volatility states.

Step 3: If we denote the parameter prior densities by 'π', then the conditional posterior is

$$f(\boldsymbol{\gamma}|\boldsymbol{\mu}, \boldsymbol{h}, \{\Psi\backslash\boldsymbol{\gamma}\}, \boldsymbol{y}) \propto f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) f(\boldsymbol{\mu}|\Psi)\pi(\Psi) \propto f(\boldsymbol{\mu}|\boldsymbol{d}, \boldsymbol{\gamma})\pi(\boldsymbol{\gamma}) \,.$$

The partial correlations $\boldsymbol{\gamma}$ are sampled as pairs of elements selected at random, conditional on the remaining elements of $\boldsymbol{\gamma}$ using a Metropolis-Hastings (MH) step. The proposal density is a bivariate independent normal, constrained to the admissible region $(-1, 1)^2$. The variances of the proposal densities are selected adaptively to target an acceptance rate between 7% and 57%.

Notice that sampling of $\boldsymbol{\gamma}$ does not require re-computation of $f_1, \ldots, f_r$ and $F_1, \ldots, F_r$, because these functions depend only on $\boldsymbol{a}, \boldsymbol{d}$ and the diagonal elements of Σ_η, and not $\boldsymbol{\gamma}$; see Section 2.2.

Step 4: We generate each element of $\boldsymbol{d}$ one-at-a-time from its conditional density

$$\begin{aligned} f(d_i|\{\Psi\backslash d_i\}, \boldsymbol{\mu}, \boldsymbol{h}, \boldsymbol{y}) &\propto f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) f(\boldsymbol{\mu}|\Psi) f(\boldsymbol{h}|\Psi)\pi(d_i) \\ &\propto f(\boldsymbol{\mu}|\boldsymbol{d}, \boldsymbol{\gamma}) f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) \left(\prod_{t=1}^{T} \frac{1}{f_{i,t}(z_{i,t}|\Psi)}\right) \pi(d_i) \,. \end{aligned}$$

Each d_i is generated using MH with a truncated normal proposal density, centered around the mode of the conditional posterior. The mode is computed using 5 Newton-Raphson steps, where the functions F_i, its inverse and f_i are not updated to speed up optimization.

Step 5: We generate each element of $\boldsymbol{a}$ one-at-a-time from its conditional density

$$\begin{aligned} f(a_i|\{\Psi\backslash a_i\}, \boldsymbol{\mu}, \boldsymbol{h}, \boldsymbol{y}) &\propto f(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{h}, \Psi) f(\boldsymbol{h}|\Psi)\pi(a_i|\Sigma_\eta) \\ &\propto \pi(a_i|\Sigma_\eta) f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) \prod_{t=1}^{T} \frac{1}{f_{i,t}(z_{i,t}|\Psi)} \,, \end{aligned}$$

in the same fashion as the elements of $\boldsymbol{d}$.

Step 6: The conditional posterior

$$f(\Sigma_\eta|\{\Psi\backslash\Sigma_\eta\}, \boldsymbol{\mu}, \boldsymbol{h}, \boldsymbol{y}) \propto \pi(\Sigma_\eta) f(\boldsymbol{h}|\boldsymbol{a}, \boldsymbol{d}, \Sigma_\eta) \prod_{t=1}^{T} \prod_{i=1}^{k} \frac{1}{f_{i,t}(z_{i,t}|\Psi)}$$

We decompose $\Sigma_\eta = \text{diag}(\bar{\sigma}^2_{1,1}, \ldots, \bar{\sigma}^2_{r,r})^{1/2} \Omega_\eta \text{diag}(\bar{\sigma}^2_{1,1}, \ldots, \bar{\sigma}^2_{r,r})^{1/2}$. The parameters $\bar{\sigma}^2_{1,1}, \ldots, \bar{\sigma}^2_{r,r}$ are sampled one at a time in a similar manner as the elements of $\boldsymbol{a}$ and $\boldsymbol{d}$. Using the recursive relations due to Yule (see the discussion in Daniels and Pourhamadi (2009)) the unique partial correlations of the correlation matrix Ω_η are generated using adaptive random walk MH.

Additional References

Carter, C. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3): 541–553.

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

Part B: BNP Density Estimator

This appendix outlines the Bayesian nonparametric constrained density estimator (labelled BNP). The unknown density of a scalar random variable X is represented as the infinite mixture

$$f(x) = \sum_{j=1}^{\infty} w_j f_{\text{ker}}(x|\boldsymbol{\theta}_j)\,.$$

Here, w_j is the weight assigned to component j, and $f_{\text{ker}}(x|\boldsymbol{\theta}_j)$ is called the kernel density. We follow closely the slice sampling approach of Kalli, Griffin and Walker (2011), but use a Gaussian density constrained to $A = \{x : x_L \leq x \leq x_U\}$ for the kernel.

The method consists of augmenting the mixture model with a slicing variable ω, so that

$$f(x,\omega) = \sum_{j=1}^{\infty} \mathcal{I}(\omega < w_j)\mathcal{I}(x \in A)\phi(x;\mu_j,\tau_j)/[\Phi(x_U;\mu_j,\tau_j) - \Phi(x_L;\mu_j,\tau_j)]$$

The weights are transformed to their stick-breaking representation as $w_j = \nu_j \prod_{k=1}^{j-1}(1-\nu_k)$ for $j \geq 1$, and $\phi(x;\mu_j,\frac{1}{\tau_j})$ is a Gaussian probability density function with mean μ_j and precision τ_j. We approximate the infinite mixture with a finite sum over its first 100 values by setting $v_{100} = 1$, so that $\sum_{j=1}^{100} w_j = 1$, although we typically find that in excess of 99% of the density mass is explained by the first twenty terms.

The prior distributions for the component parameters are selected as $\tau_j|\beta \sim \text{Gamma}(\alpha,\beta)$, and $\mu_j|c \sim \text{N}(\varepsilon, \frac{R^2}{c})$. Following Richardson and Green (1997, p. 738), we employ a gamma hyperprior for β. We also allow c to have a gamma hyperprior. Denote as $\boldsymbol{x} = (x_1,\ldots,x_n)'$ a vector of n observed values of X, and as $\boldsymbol{\delta} = (\delta_1,\ldots,\delta_n)'$ a vector of corresponding mixture component allocation integers.

The MCMC sampling scheme to estimate the density is an extension of that in Kalli, Griffin and Walker (2011), and generates repeatedly from the following densities:

1. $f\left(\mu_j|\boldsymbol{x},\boldsymbol{\delta},\tau_j\beta,c\right) \propto \pi\left(\mu_j|c\right)\prod_{\delta_i=j}\phi\left(x_i;\mu_j,\frac{1}{\tau_j}\right)/[\Phi\left(x_U;\mu_j,\frac{1}{\tau_j}\right) - \Phi\left(x_L;\mu_j,\frac{1}{\tau_j}\right)]$
2. $f\left(\tau_j|\boldsymbol{x},\boldsymbol{\delta},\mu_j,\beta,c\right) \propto \pi\left(\tau_j|\beta\right)\prod_{\delta_i=j}\phi\left(x_i;\mu_j,\frac{1}{\tau_j}\right)/[\Phi\left(x_U;\mu_j,\frac{1}{\tau_j}\right) - \Phi\left(x_L;\mu_j,\frac{1}{\tau_j}\right)]$
3. $f\left(c|\boldsymbol{\mu}\right) \propto \pi(c)\prod_{j=1}^{J}\pi\left(\mu_j|c\right)$
4. $f\left(\beta|\boldsymbol{\tau}\right) \propto \pi(\beta)\prod_{j=1}^{J}\pi\left(\tau_j|\beta\right)$
5. $f(\nu_j|\boldsymbol{\delta}) \propto \text{Be}\left(\nu_j;a_j,b_j\right)$, where $a_j = a_j^0 + \sum_{i=1}^{n}\mathcal{I}\left(\delta_i = j\right)$, and $b_j = b_j^0 + \sum_{i=1}^{n}\mathcal{I}\left(\delta_i > j\right)$, for $j = 1,\ldots,100$.
6. $f(\omega_i|\boldsymbol{w},\boldsymbol{\delta}) \propto \mathcal{I}\left(0 < \omega_i < w_{\delta_i}\right)$ for $i = 1,\ldots,n$.
7. $\Pr(\delta_i = j) \propto \mathcal{I}\left(w_j > \omega_i\right)\phi(x_i;\mu_j,\frac{1}{\tau_j})/[\Phi(x_U;\mu_j,\frac{1}{\tau_j}) - \Phi(x_L;\mu_j,\frac{1}{\tau_j})]$ for $i = 1,\ldots,n$.

Here, the indicator function $\mathcal{I}(A) = 1$ if A is true, and zero otherwise, while Be denotes the density of a beta distribution. To generate from the densities at 1. and 2. above, we use MH with normal and gamma proposals, respectively. The posteriors at 3. and 4. are both recognizable as gamma distributions. The density at 6. is a uniform, while that at 7. is a multinomial with mass function computed by evaluating the terms up to proportionality, and then normalizing by their sum over j.

In our empirical applications we set $R = \text{Range}(x)$, $\varepsilon = \frac{\text{man}(x)+\min(x)}{2}$ and $\alpha = 2$. The hyperpriors are set as $\beta \sim \text{Gamma}(g, h)$, and $c \sim \text{Gamma}(\alpha_c, \beta_c)$, with $g = 0.2$, $h = \frac{10}{R^2}$, $\alpha_c = 2.34$ and $\beta_c = 0.065$. For $\boldsymbol{\nu}$, we use the prior $\nu_j \sim \text{Beta}(a_j^0, b_j^0)$, with $a_j^0 = \epsilon\kappa_j$, $b_j^0 = \epsilon(1-\kappa_j)$, $\kappa_j = \left(1 - \sum_{k<j} \xi_k\right)^{-1} \xi_j$, $\xi_j = E(w_j) = \frac{\Gamma\left(\tilde{a}+\tilde{b}\right)\Gamma(\tilde{a}+1)\Gamma\left(\tilde{b}+j-1\right)}{\Gamma\left(\tilde{\tilde{b}}\right)\Gamma\left(\tilde{b}\right)\Gamma\left(\tilde{a}+\tilde{b}+j\right)}$, $\tilde{a} = 2$, $\tilde{b} = 8$ and $\epsilon = 20$; see Kalli, Griffin and Walker (2011) for a discussion.

Additional References

Richardson, S., & Green, P. J. (1997). On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society: series B (statistical methodology)*, 59(4), 731-792.

Part C: Additional Empirical Results for the fit to Vintage 2016Q2

			$\Delta\ln(\text{GDP})_t$ ($j=1$)	Infl_t ($j=2$)	UR_t ($j=3$)	IR_t ($j=4$)	VXO_t ($j=5$)
$\gamma^0_{i,j}$	Infl_t	$(i=2)$	-0.158	1.000	-	-	-
			(-0.337, 0.000)	-	-	-	-
	UR_t	$(i=3)$	-0.055	0.056	1.000	-	-
			(-0.177, 0.042)	(-0.160,0.235)	-	-	-
	IR_t	$(i=4)$	0.218	0.652	-0.029	1.000	-
			(0.050, 0.365)	(0.435,0.789)	(-0.219, 0.169)	-	-
	VXO_t	$(i=5)$	-0.313	-0.077	0.111	0.050	1.000
			(-0.487,-0.102)	(-0.394,0.249)	(-0.097, 0.317)	(-0.267, 0.316)	-
$\gamma^1_{i,j}$	$\Delta\ln(\text{GDP})_{t-1}$	$(i=1)$	0.529	-0.179	-0.195	0.143	-0.228
			(0.341, 0.687)	(-0.350,0.007)	(-0.306,-0.088)	(-0.044, 0.281)	(-0.405,-0.047)
	Infl_{t-1}	$(i=2)$	0.050	0.942	0.121	0.706	-0.079
			(-0.230, 0.333)	(0.858,0.983)	(-0.121, 0.369)	(0.528, 0.815)	(-0.311, 0.312)
	UR_{t-1}	$(i=3)$	0.644	-0.044	0.993	-0.122	0.048
			(0.425, 0.876)	(-0.348,0.241)	(0.989, 0.995)	(-0.345, 0.071)	(-0.184, 0.264)
	IR_{t-1}	$(i=4)$	-0.096	-0.124	0.313	0.984	0.211
			(-0.326, 0.141)	(-0.351,0.103)	(0.087, 0.505)	(0.974, 0.990)	(-0.097, 0.463)
	VXO_{t-1}	$(i=5)$	0.088	-0.037	-0.009	-0.069	0.867
			(-0.185, 0.344)	(-0.311,0.245)	(-0.249, 0.218)	(-0.257, 0.125)	(0.788, 0.930)
$\gamma^2_{i,j}$	$\Delta\ln(\text{GDP})_{t-2}$	$(i=1)$	-0.062	-0.001	-0.138	0.008	0.026
			(-0.239, 0.102)	(-0.185,0.171)	(-0.309, 0.050)	(-0.143, 0.157)	(-0.142, 0.189)
	Infl_{t-2}	$(i=2)$	-0.029	0.008	0.014	-0.019	-0.003
			(-0.206, 0.145)	(-0.188,0.198)	(-0.171, 0.189)	(-0.193, 0.146)	(-0.178, 0.172)
	UR_{t-2}	$(i=3)$	-0.148	-0.042	-0.058	0.053	0.039
			(-0.347, 0.069)	(-0.229,0.141)	(-0.262, 0.161)	(-0.128, 0.220)	(-0.131, 0.214)
	IR_{t-2}	$(i=4)$	0.085	-0.034	-0.026	-0.281	-0.015
			(-0.100, 0.262)	(-0.210,0.142)	(-0.200, 0.141)	(-0.419,-0.142)	(-0.164, 0.128)
	VXO_{t-2}	$(i=5)$	0.059	0.009	-0.045	0.057	0.004
			(-0.117, 0.236)	(-0.176,0.187)	(-0.212, 0.131)	(-0.111, 0.217)	(-0.179, 0.176)
$\gamma^3_{i,j}$	$\Delta\ln(\text{GDP})_{t-3}$	$(i=1)$	-0.067	-0.005	0.011	0.028	-0.030
			(-0.194, 0.059)	(-0.137,0.126)	(-0.111, 0.131)	(-0.086, 0.143)	(-0.153, 0.087)
	Infl_{t-3}	$(i=2)$	-0.002	0.014	-0.019	-0.031	-0.021
			(-0.132, 0.122)	(-0.117,0.146)	(-0.145, 0.106)	(-0.153, 0.092)	(-0.150, 0.109)
	UR_{t-3}	$(i=3)$	-0.042	-0.033	-0.032	-0.021	0.033
			(-0.177, 0.095)	(-0.159,0.094)	(-0.162, 0.100)	(-0.148, 0.106)	(-0.090, 0.158)
	IR_{t-3}	$(i=4)$	0.054	-0.032	-0.004	0.028	0.042
			(-0.069, 0.176)	(-0.159,0.098)	(-0.125, 0.117)	(-0.087, 0.145)	(-0.078, 0.162)
	VXO_{t-3}	$(i=5)$	0.017	-0.023	0.014	0.002	0.001
			(-0.107, 0.140)	(-0.151,0.106)	(-0.105, 0.138)	(-0.117, 0.121)	(-0.126, 0.127)
$\gamma^4_{i,j}$	$\Delta\ln(\text{GDP})_{t-4}$	$(i=1)$	-0.009	0.004	0.007	0.017	-0.019
			(-0.107, 0.091)	(-0.094,0.100)	(-0.088, 0.101)	(-0.074, 0.108)	(-0.111, 0.073)
	Infl_{t-4}	$(i=2)$	0.004	0.021	-0.023	-0.013	0.016
			(-0.094, 0.103)	(-0.081,0.123)	(-0.120, 0.074)	(-0.109, 0.081)	(-0.081, 0.114)
	UR_{t-4}	$(i=3)$	-0.020	-0.013	-0.023	0.010	0.022
			(-0.119, 0.080)	(-0.115,0.090)	(-0.125, 0.078)	(-0.088, 0.105)	(-0.075, 0.120)
	IR_{t-4}	$(i=4)$	0.040	-0.004	-0.017	-0.012	0.036
			(-0.055, 0.133)	(-0.100,0.094)	(-0.112, 0.078)	(-0.102, 0.078)	(-0.055, 0.128)
	VXO_{t-4}	$(i=5)$	-0.001	-0.004	-0.004	-0.010	0.009
			(-0.098, 0.098)	(-0.103,0.095)	(-0.101, 0.092)	(-0.105, 0.081)	(-0.088, 0.106)

Table A2: **Estimates of all the partial correlations from the 2016Q2 vintage data**. The posterior means and 95% intervals are reported for the partial correlations $\gamma^0_{i,j}$ and $\gamma^1_{i,j}$ for the UCSV-C inversion copula.

	$\Delta \ln(\text{GDP})$	Infl	Σ_ϵ UR	IR	VXO
$\Delta \ln(\text{GDP})$	0.2860	-0.0162	-0.0372	0.0171	-0.0340
	(0.1867, 0.4068)	(-0.0421,0.0085)	(-0.0547,-0.0221)	(0.0069, 0.0283)	(-0.0764,0.0086)
Infl	-0.0162	0.0428	0.0012	0.0015	0.0112
	(-0.0421, 0.0085)	(0.0159,0.0829)	(-0.0042, 0.0066)	(-0.0028, 0.0061)	(-0.0067,0.0301)
UR	-0.0372	0.0012	0.0116	-0.0032	0.0042
	(-0.0547,-0.0221)	(-0.0042,0.0066)	(0.0081, 0.0162)	(-0.0058,-0.0007)	(-0.0040,0.0128)
IR	0.0171	0.0015	-0.0032	0.0120	0.0003
	(0.0069, 0.0283)	(-0.0028,0.0061)	(-0.0058,-0.0007)	(0.0087, 0.0156)	(-0.0071,0.0078)
VXO	-0.0340	0.0112	0.0042	0.0003	0.1524
	(-0.0764, 0.0086)	(-0.0067,0.0301)	(-0.0040, 0.0128)	(-0.0071, 0.0078)	(0.0862,0.2287)

Table A3: **Estimates of Σ_ϵ from the 2016Q2 vintage data**. The posterior means and 90% posterior intervals for Σ_ϵ from the inversion copula model with nonparametric margins fitted to the 2016Q2 vintage data.

B_1	$\Delta\ln(\text{GDP})_{t-1}$	Infl_{t-1}	UR_{t-1}	IR_{t-1}	VXO_{t-1}
$\Delta\ln(\text{GDP})_t$	0.542	-0.030	0.471	0.172	-0.241
	(0.29, 0.80)	(-0.50,0.46)	(-0.52,1.52)	(-0.49, 0.84)	(-0.49,-0.01)
Infl_t	0.171	0.999	0.087	0.067	0.041
	(0.06, 0.30)	(0.80,1.19)	(-0.33,0.55)	(-0.21, 0.38)	(-0.05, 0.14)
UR_t	-0.186	-0.017	1.031	-0.020	0.038
	(-0.27,-0.11)	(-0.12,0.08)	(0.78,1.26)	(-0.17, 0.13)	(-0.01, 0.09)
IR_t	0.083	0.053	-0.015	1.264	-0.028
	(0.03, 0.14)	(-0.04,0.15)	(-0.22,0.19)	(1.12, 1.41)	(-0.07, 0.02)
VXO_t	-0.047	-0.069	-0.223	0.193	0.879
	(-0.24, 0.13)	(-0.42,0.27)	(-1.01,0.54)	(-0.30, 0.71)	(0.71, 1.07)
B_2	$\Delta\ln(\text{GDP})_{t-2}$	Infl_{t-2}	UR_{t-2}	IR_{t-2}	VXO_{t-2}
$\Delta\ln(\text{GDP})_t$	0.109	-0.123	-0.308	-0.298	0.112
	(-0.16, 0.37)	(-0.79,0.48)	(-1.78,1.07)	(-1.32, 0.68)	(-0.15, 0.39)
Infl_t	0.024	0.005	-0.010	-0.067	0.015
	(-0.08, 0.14)	(-0.23,0.22)	(-0.58,0.55)	(-0.48, 0.34)	(-0.09, 0.12)
UR_t	-0.011	0.022	-0.037	0.066	-0.021
	(-0.07, 0.04)	(-0.10,0.16)	(-0.34,0.27)	(-0.14, 0.27)	(-0.08, 0.03)
IR_t	-0.013	0.016	0.031	-0.341	0.015
	(-0.07, 0.04)	(-0.11,0.14)	(-0.25,0.31)	(-0.55,-0.14)	(-0.04, 0.07)
VXO_t	-0.021	0.034	-0.020	-0.210	0.000
	(-0.22, 0.18)	(-0.42,0.49)	(-1.06,1.01)	(-0.98, 0.52)	(-0.21, 0.18)
B_3	$\Delta\ln(\text{GDP})_{t-3}$	Infl_{t-3}	UR_{t-3}	IR_{t-3}	VXO_{t-3}
$\Delta\ln(\text{GDP})_t$	-0.021	-0.011	-0.090	0.040	0.014
	(-0.22, 0.17)	(-0.44,0.44)	(-1.10,0.94)	(-0.72, 0.81)	(-0.18, 0.21)
Infl_t	0.012	-0.010	-0.070	-0.050	-0.010
	(-0.06, 0.09)	(-0.17,0.15)	(-0.53,0.33)	(-0.38, 0.25)	(-0.09, 0.07)
UR_t	0.007	0.003	0.006	0.003	0.003
	(-0.03, 0.05)	(-0.09,0.09)	(-0.20,0.22)	(-0.16, 0.16)	(-0.04, 0.04)
IR_t	0.003	-0.015	-0.034	0.045	0.003
	(-0.04, 0.04)	(-0.11,0.08)	(-0.24,0.17)	(-0.11, 0.21)	(-0.04, 0.04)
VXO_t	-0.062	-0.085	0.057	-0.002	-0.006
	(-0.21, 0.08)	(-0.42,0.24)	(-0.69,0.82)	(-0.57, 0.57)	(-0.14, 0.13)
B_4	$\Delta\ln(\text{GDP})_{t-4}$	Infl_{t-4}	UR_{t-4}	IR_{t-4}	VXO_{t-4}
$\Delta\ln(\text{GDP})_t$	0.003	0.016	0.030	0.156	-0.003
	(-0.13, 0.13)	(-0.25,0.29)	(-0.52,0.62)	(-0.23, 0.56)	(-0.12, 0.12)
Infl_t	0.010	0.024	-0.028	-0.001	-0.002
	(-0.04, 0.06)	(-0.07,0.12)	(-0.27,0.21)	(-0.16, 0.16)	(-0.05, 0.05)
UR_t	0.002	-0.010	-0.033	-0.010	0.000
	(-0.02, 0.03)	(-0.07,0.04)	(-0.15,0.08)	(-0.09, 0.07)	(-0.02, 0.02)
IR_t	0.000	-0.008	0.011	-0.011	-0.003
	(-0.03, 0.03)	(-0.06,0.05)	(-0.11,0.13)	(-0.09, 0.07)	(-0.03, 0.02)
VXO_t	-0.020	0.030	0.139	0.134	0.010
	(-0.12, 0.08)	(-0.17,0.24)	(-0.27,0.57)	(-0.16, 0.44)	(-0.08, 0.09)

Table A4: Posterior means and 90% posterior intervals for the autoregressive coefficient matrices B_1, B_2, B_3, B_4 from the inversion copula model with nonparametric margins fitted to the 2016Q2 vintage data.

	$\Delta\ln(\text{GDP})_t$	Infl_t	UR_t	IR_t	VXO_t
	$R(0)$				
$\Delta\ln(\text{GDP})_t$	1.000	-0.099	-0.039	0.052	-0.180
Infl_t	-0.099	1.000	0.041	0.568	-0.021
UR_t	-0.039	0.041	1.000	-0.039	0.094
IR_t	0.052	0.568	-0.039	1.000	0.040
VXO_t	-0.180	-0.021	0.094	0.040	1.000
	$R(1)$				
$\Delta\ln(\text{GDP})_{t-1}$	0.327	-0.029	-0.148	0.100	-0.148
Infl_{t-1}	-0.129	0.795	0.080	0.579	-0.033
UR_{t-1}	0.065	0.015	0.958	-0.053	0.042
IR_{t-1}	-0.001	0.547	-0.008	0.965	0.072
VXO_{t-1}	-0.176	-0.023	0.151	0.006	0.709
	$R(2)$				
$\Delta\ln(\text{GDP})_{t-2}$	0.205	0.010	-0.216	0.128	-0.107
Infl_{t-2}	-0.131	0.755	0.123	0.584	-0.041
UR_{t-2}	0.123	0.005	0.908	-0.057	0.000
IR_{t-2}	-0.019	0.518	0.031	0.941	0.100
VXO_{t-2}	-0.135	-0.029	0.196	-0.025	0.620
	$R(3)$				
$\Delta\ln(\text{GDP})_{t-3}$	0.099	0.035	-0.250	0.147	-0.082
Infl_{t-3}	-0.116	0.719	0.166	0.583	-0.051
UR_{t-3}	0.160	0.004	0.842	-0.052	-0.036
IR_{t-3}	-0.019	0.488	0.074	0.913	0.122
VXO_{t-3}	-0.095	-0.027	0.231	-0.049	0.537
	$R(4)$				
$\Delta\ln(\text{GDP})_{t-4}$	0.039	0.044	-0.260	0.159	-0.059
Infl_{t-4}	-0.099	0.682	0.205	0.581	-0.059
UR_{t-4}	0.167	0.008	0.769	-0.043	-0.065
IR_{t-4}	-0.014	0.458	0.115	0.882	0.139
VXO_{t-4}	-0.070	-0.025	0.253	-0.067	0.467

Table A5: **Estimates of level dependence from the 2016Q2 vintage data**. The posterior means of the matrices of Spearman unconditional pairwise correlations $R(k)$, for $k = 0, 1, \ldots, 4$, are reported for the fitted UCSV-C inversion copula. These measure dependence in the level of the time series $\{\boldsymbol{Y}_t\}$, and are computed by simulation.

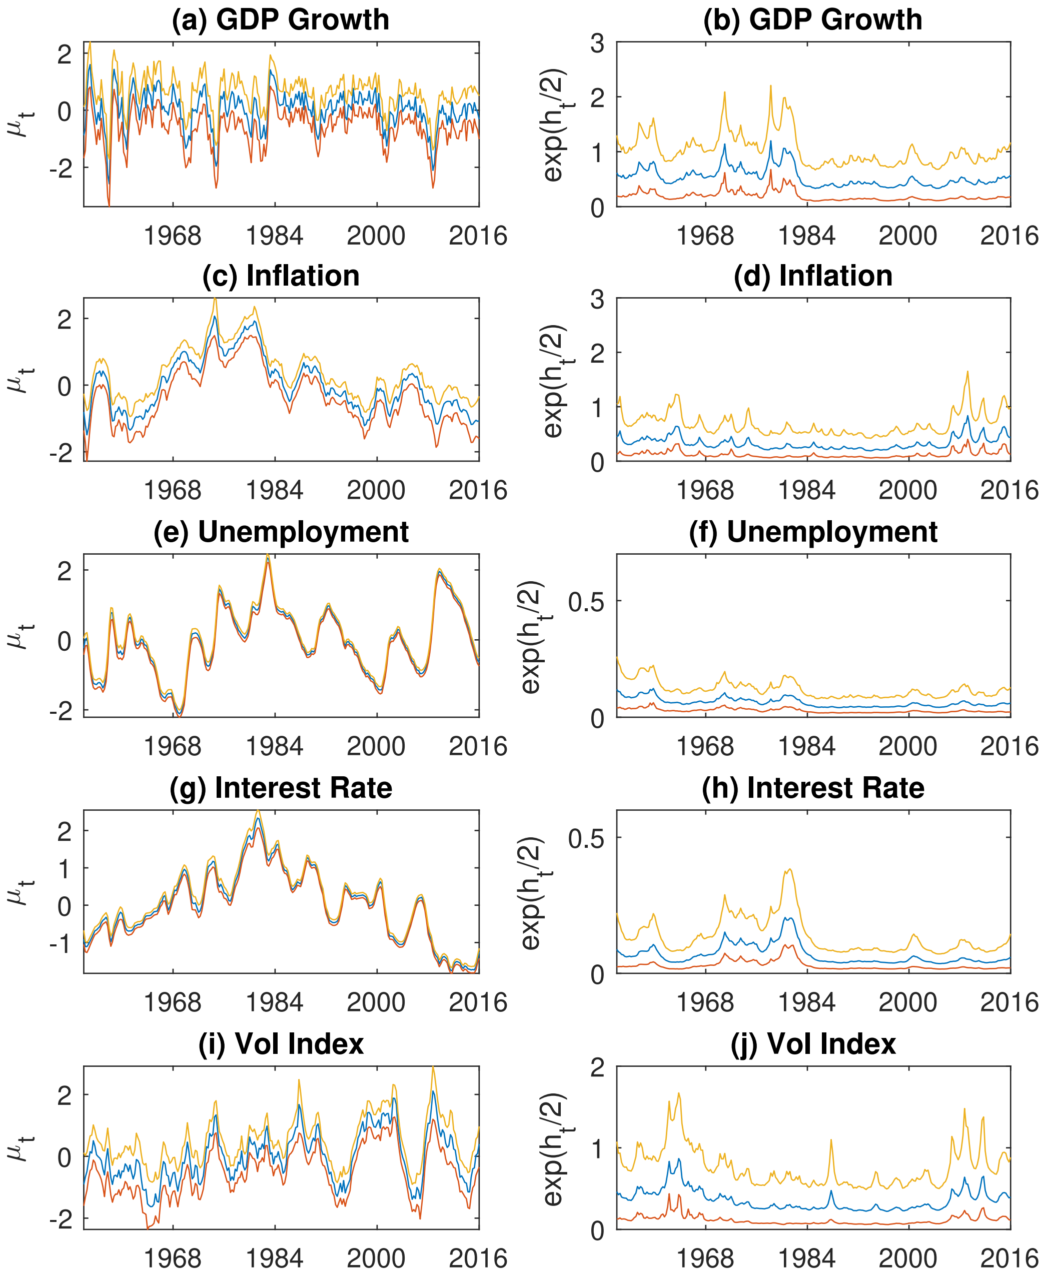

Figure A1: The posterior mean and 95% posterior intervals for the unobserved component ($\mu_{i,t}$) and volatility of the latent variable ($exp(h_{i,t}/2)$) from the inversion copula model are displayed in the first and second column respectively.

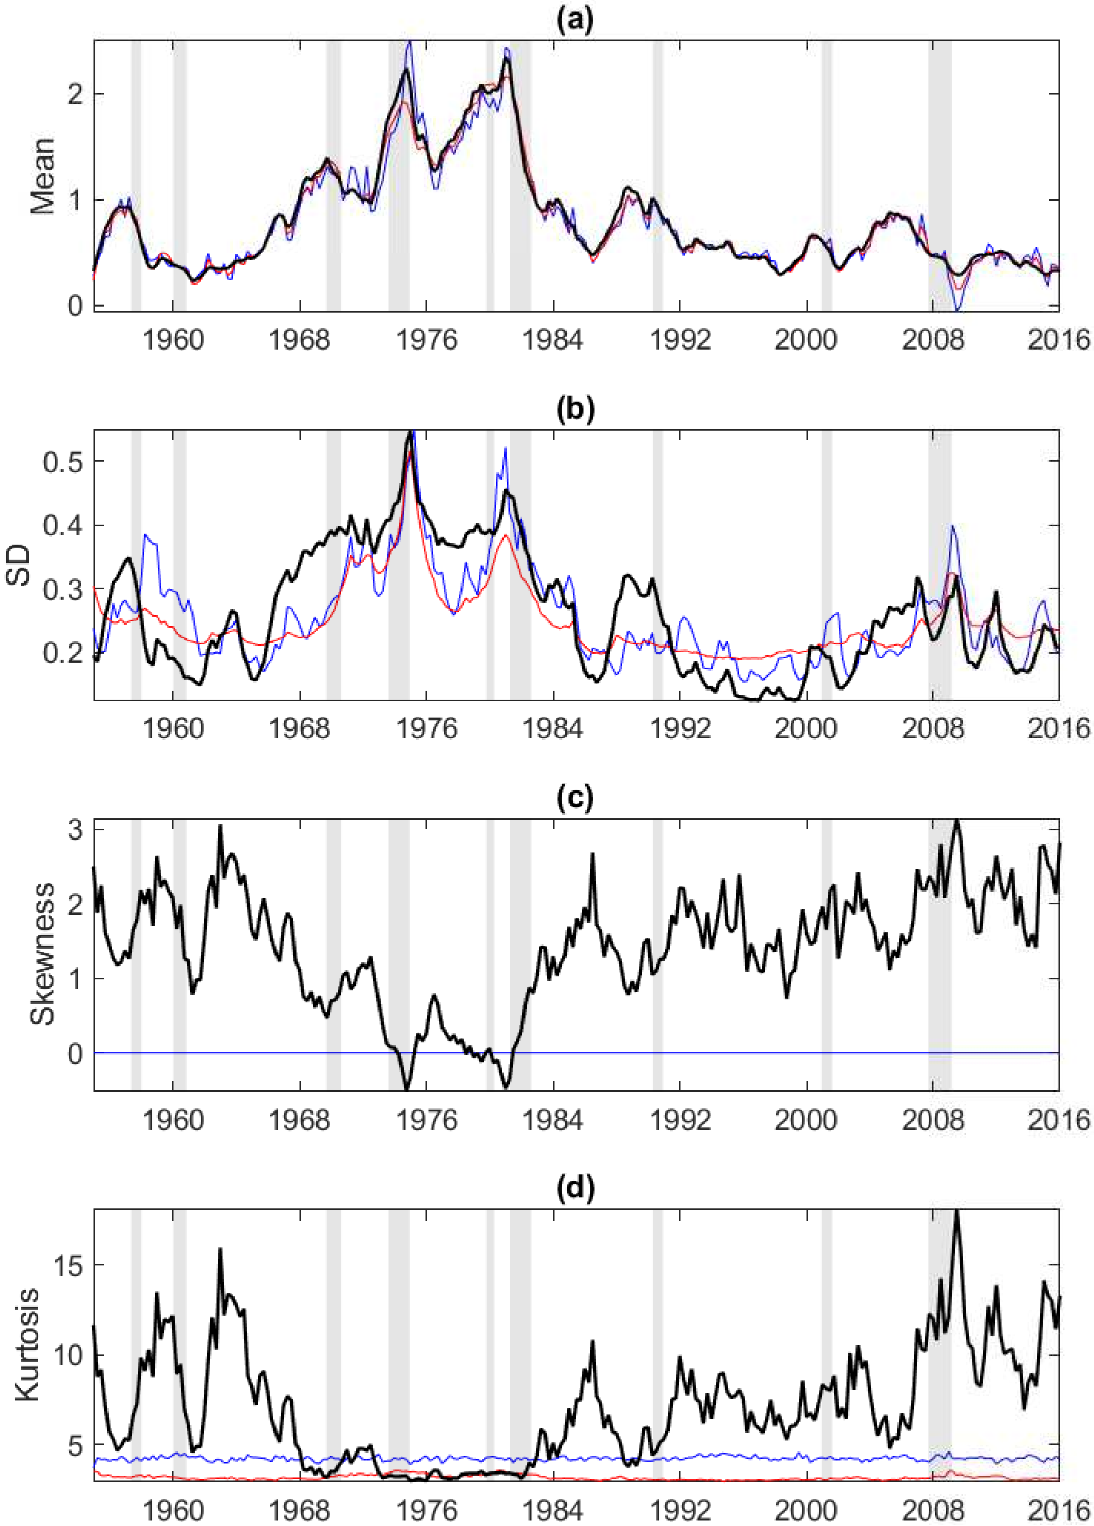

Figure A2: **Moments of the one-quarter-ahead predictive distributions of inflation from the 2016Q2 vintage data.** Panels (a–d) plot the mean, standard deviation, Pearson's skew coefficient and kurtosis from the one quarter ahead (ie. $h = 1$) in-sample predictive distributions of inflation. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

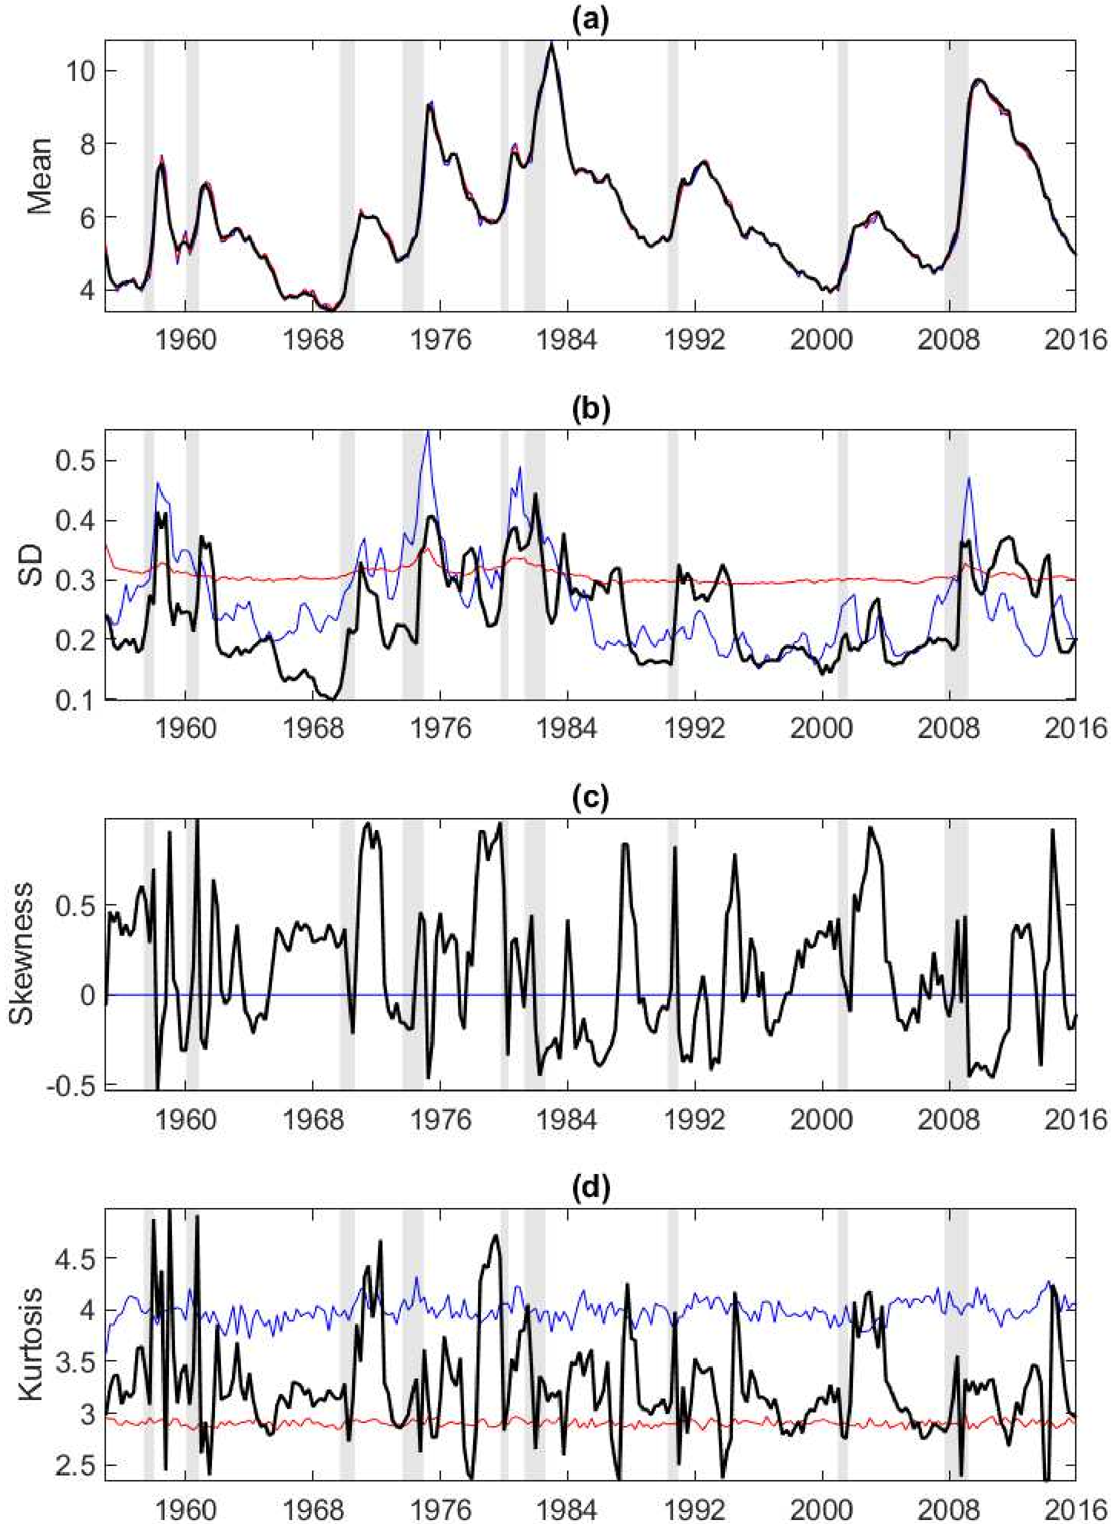

Figure A3: **Moments of the one-quarter-ahead predictive distributions of unemployment from the 2016Q2 vintage data.** Panels (a–d) plot the mean, standard deviation, Pearson's skew coefficient and kurtosis from the one quarter ahead (ie. $h = 1$) in-sample predictive distributions of unemployment. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

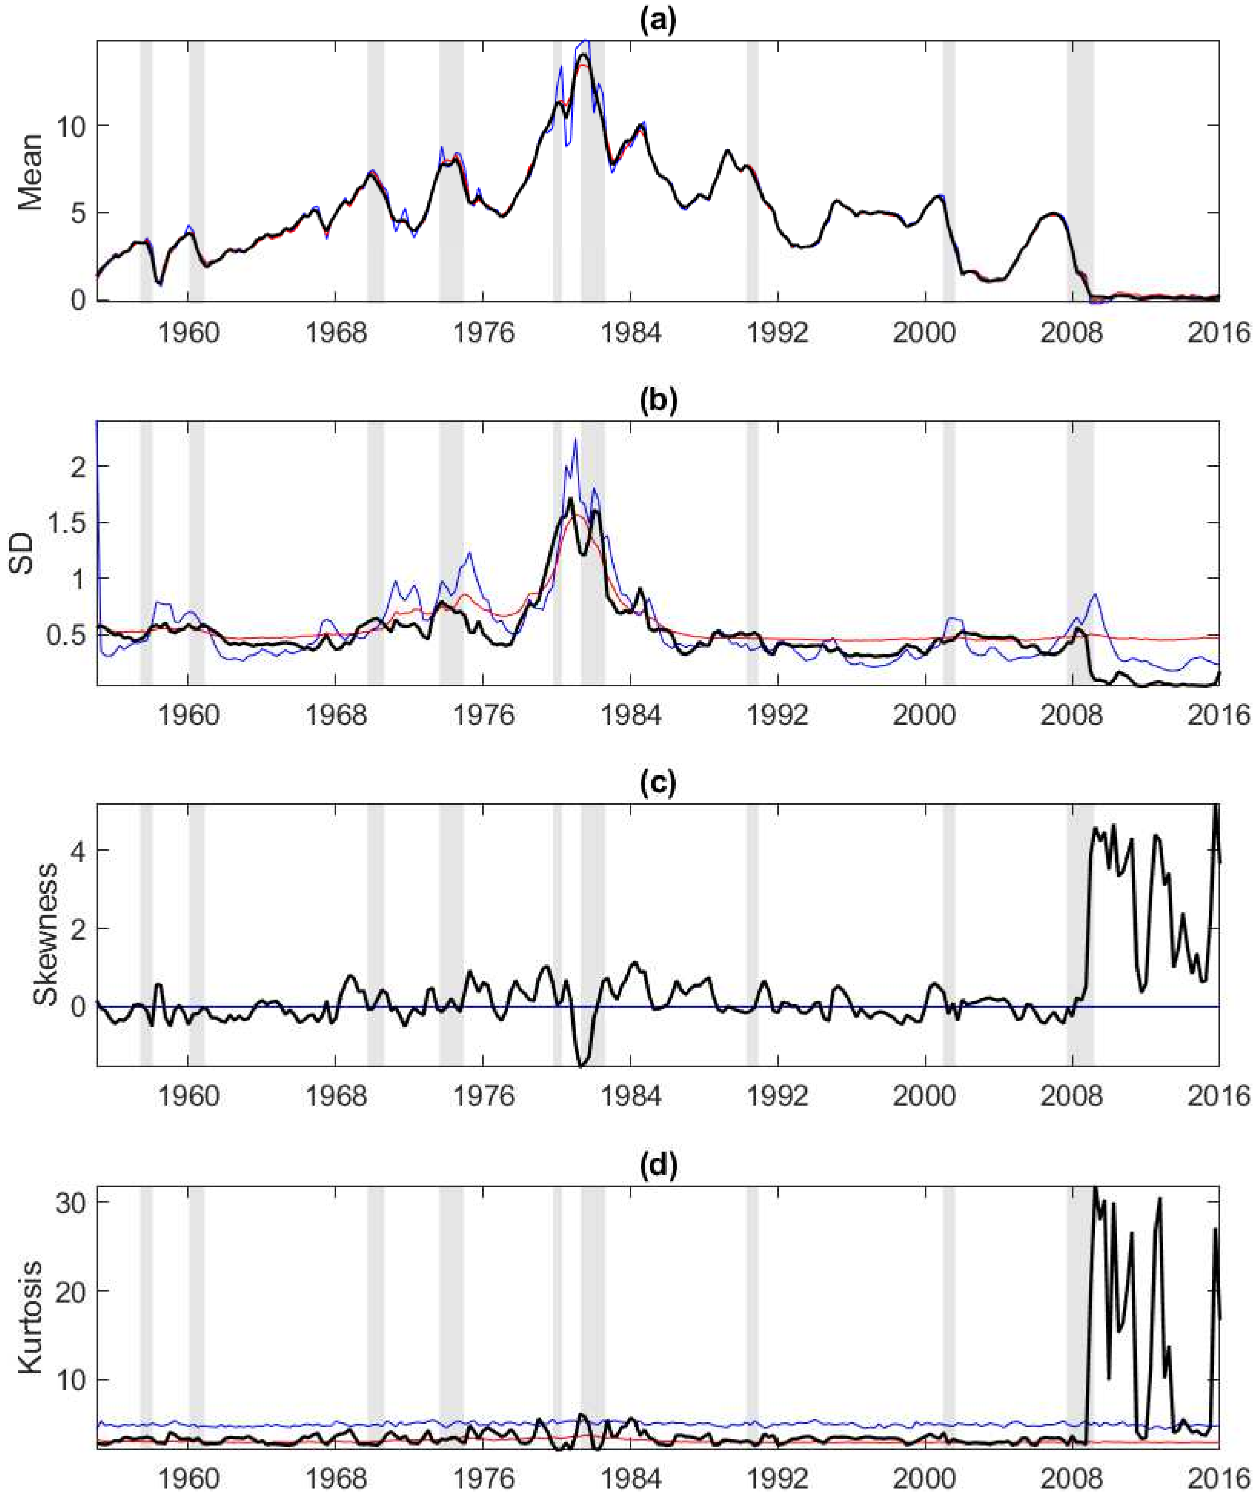

Figure A4: **Moments of the one-quarter-ahead predictive distributions of the interest rate from the 2016Q2 vintage data.** Panels (a–d) plot the mean, standard deviation, Pearson's skew coefficient and kurtosis from the one quarter ahead (ie. $h = 1$) in-sample predictive distributions of the interest rate. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

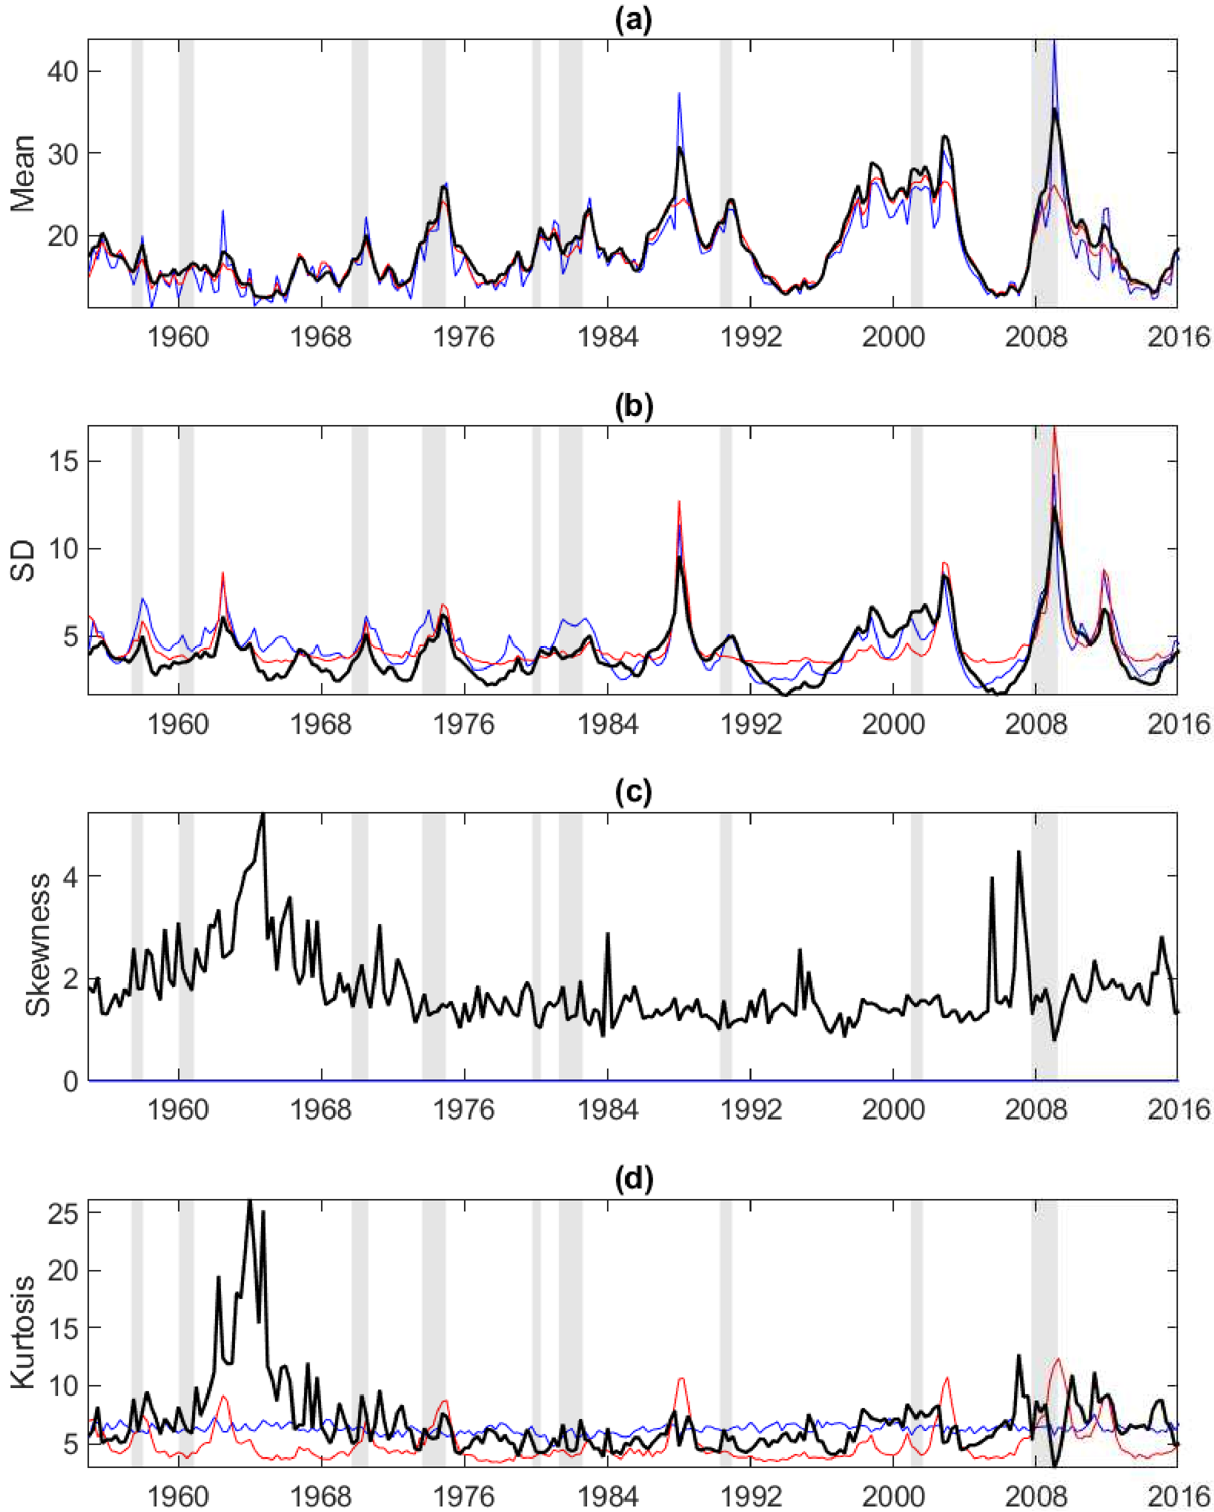

Figure A5: **Moments of the one-quarter-ahead predictive distributions of stock market volatility from the 2016Q2 vintage data.** Panels (a–d) plot the mean, standard deviation, Pearson's skew coefficient and kurtosis from the one quarter ahead (ie. $h = 1$) in-sample predictive distributions of stock market volatility. In each panel, moments from the fitted VARSV (blue line), UCSV (red line) and UCSV-C (black line) models are plotted. Quarters classified as recessions by the NBER are shaded.

Part D: Additional Empirical Results for the Real-Time Forecasting Exercise

	Log-scores							
	$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 1$	$h = 2$	$h = 4$	$h = 8$
	GDP Growth				Inflation			
Time Series Model								
VAR	-1.2357**	-1.2424**	-1.2487**	-1.2522**	-0.2516**	-0.3739**	-0.5330**	-0.8315**
VARSV	-1.0601	-1.0823	-1.0944^{++}	-1.1306	-0.2115	-0.3021	-0.4627	-0.8039
UCSV	-1.2230**	-1.2235**	-1.2321**	-1.2310**	-0.2250	-0.3274	-0.4514	-0.7463
VAR-C-KDE	-1.0776**	-1.1050	-1.1461	-1.1557**	-0.3601**	-0.3545*	-0.4227	-0.7309
UCSV-C-BNP	-1.0449	-1.0971	-1.1334^{++}	-1.1245^{++}	-0.2497	-0.2999	-0.3874	-0.7154
UCSV-C-Skt	-1.0481	-1.1045	-1.1451	-1.1465	-0.2153^{++}	-0.2899	-0.3623	-0.6021
UCSV-C-KDE	*-1.0423*	*-1.0971*	*-1.1508*	*-1.1479*	*-0.2485*	*-0.3080*	*-0.3944*	*-0.7072*
	Unemployment				Interest Rate			
Time Series Model								
VAR	-0.1534**	-0.6914**	-1.2196^{++}	-1.6435^{++}	-1.1435**	-1.5844**	-1.9852**	-2.4864**
VARSV	0.0439^{++}	-0.5179^{++}	-1.1300^{+}	-1.6271	-0.7317**	-1.3717	-1.9866	-2.6226
UCSV	-0.1901**	-0.7061	-1.2371	-1.6657^{+}	-0.9716**	-1.4645**	-2.0020	-2.8035
VAR-C-KDE	-0.1354**	-0.7269**	-1.3338	-1.7693	-0.5764	-1.2218	-1.8183	-2.3821
UCSV-C-BNP	-0.1230*	-0.6833	-1.2909	-1.8155	-0.5173^{++}	-1.0894^{++}	-1.6498^{+}	-2.2700
UCSV-C-Skt	-0.0071^{++}	-0.5871^{++}	-1.2102	-1.7166	-0.7318**	-1.3172	-1.9126	-2.6204**
UCSV-C-KDE	*-0.1044*	*-0.6809*	*-1.3208*	*-1.7621*	*-0.5885*	*-1.2050*	*-1.8401*	*-2.4322*

Table A6: Log-scores for the real-time (out-of-sample) predictive densities from the five multivariate time series models (higher values indicate increased accuracy). Separate results are reported for forecasts made $h = 1, 2, 4$ and 8 quarters ahead in columns. Each log-score is computed as the mean of the logarithms of the predictive densities evaluated at the observations, where each predictive density is computed as outlined in Section 3.3. As a rough guide, results of a small sample adjusted Diebold and Mariano (1995) test are reported. This is a two-sided test of the equality of the mean score of the UCSV-C-KDE with each of the other models. Rejection of the null hypothesis of equal means at the 10% or 5% level is denoted by '*' or '**' if favorable to the UCSV-C-KDE model, or by '$^{+}$' or '$^{++}$' if unfavorable.

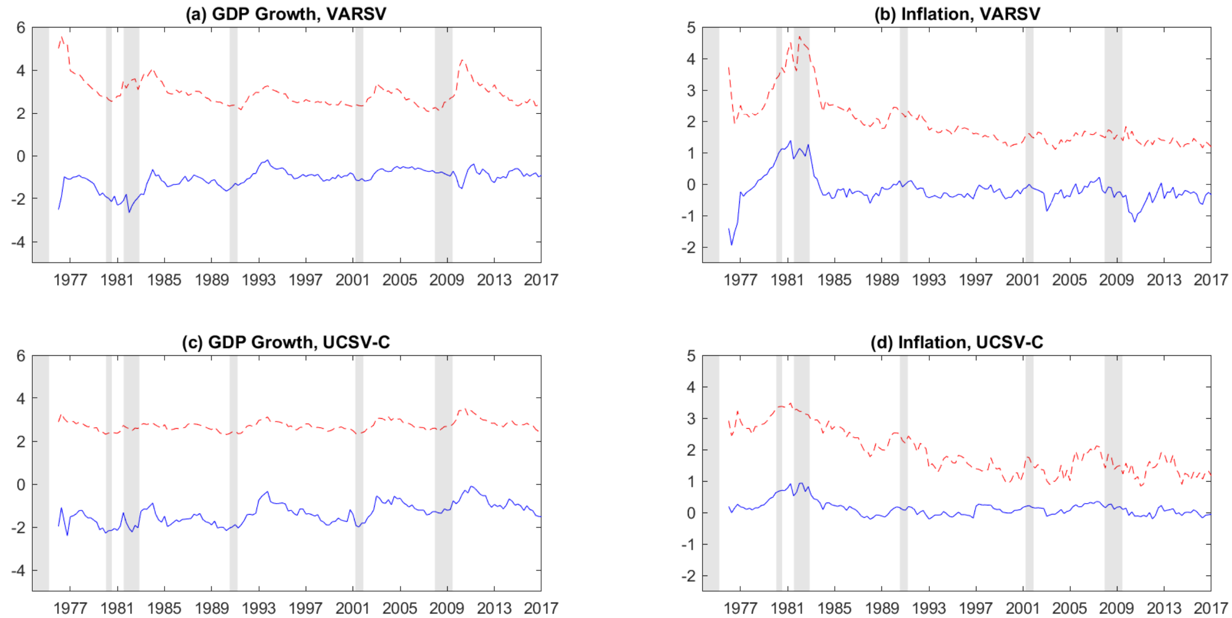

Figure A6: **Asymmetry in real-time inflation and GDP growth prediction for $h = 4$.** The real-time (out-of-sample) one-year-ahead expected shortfall (blue line) and upper tail equivalent (red line), from the VARSV and UCSV-C models. The first and second column present estimates for inflation and GDP Growth, respectively.

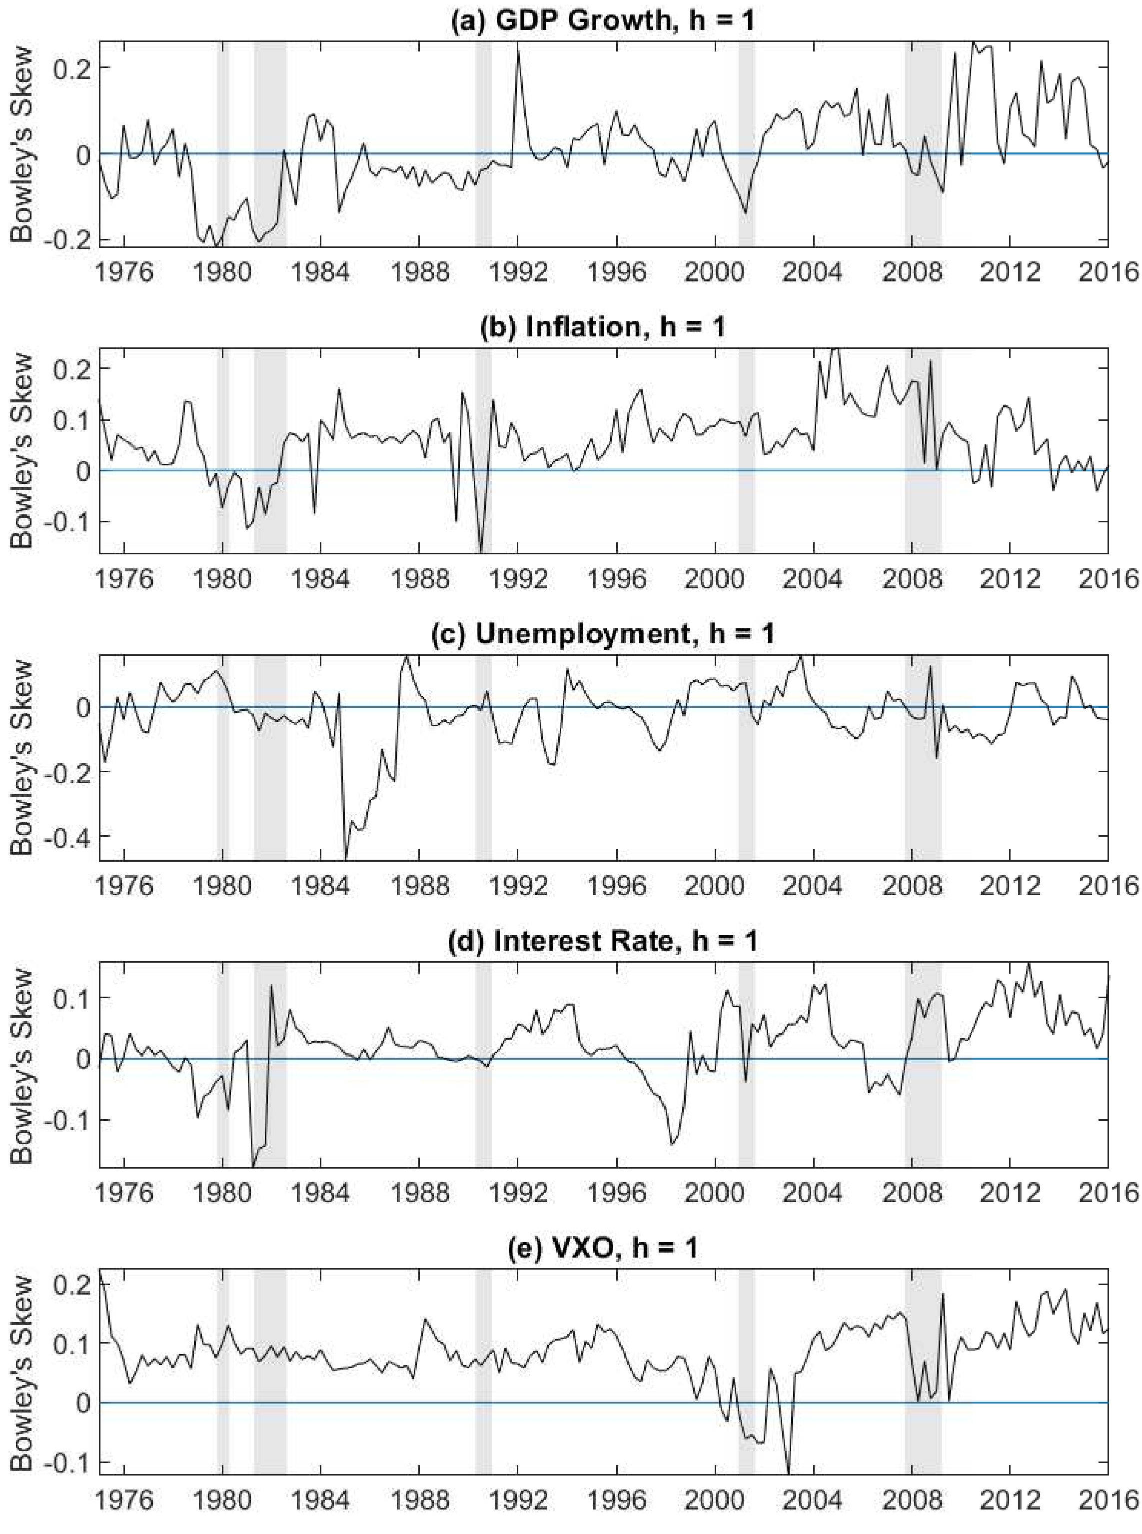

Figure A7: Bowley's skew coefficients of the current quarter real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

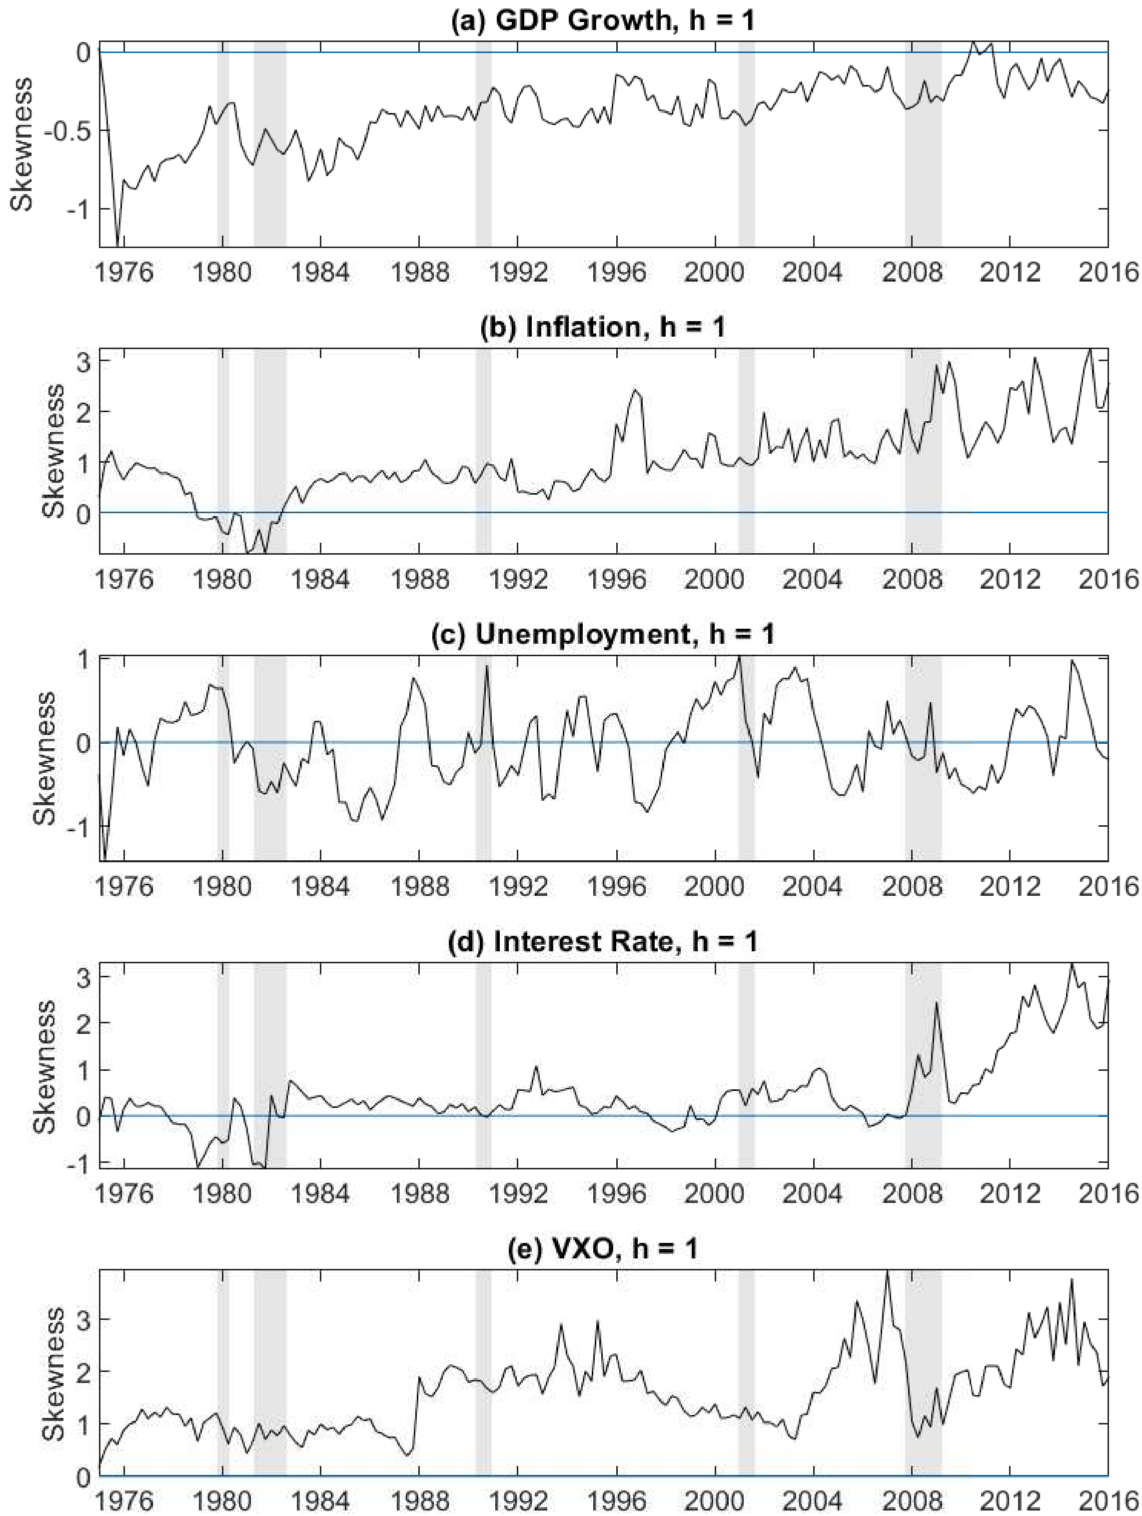

Figure A8: Pearson's skewness of the current quarter real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

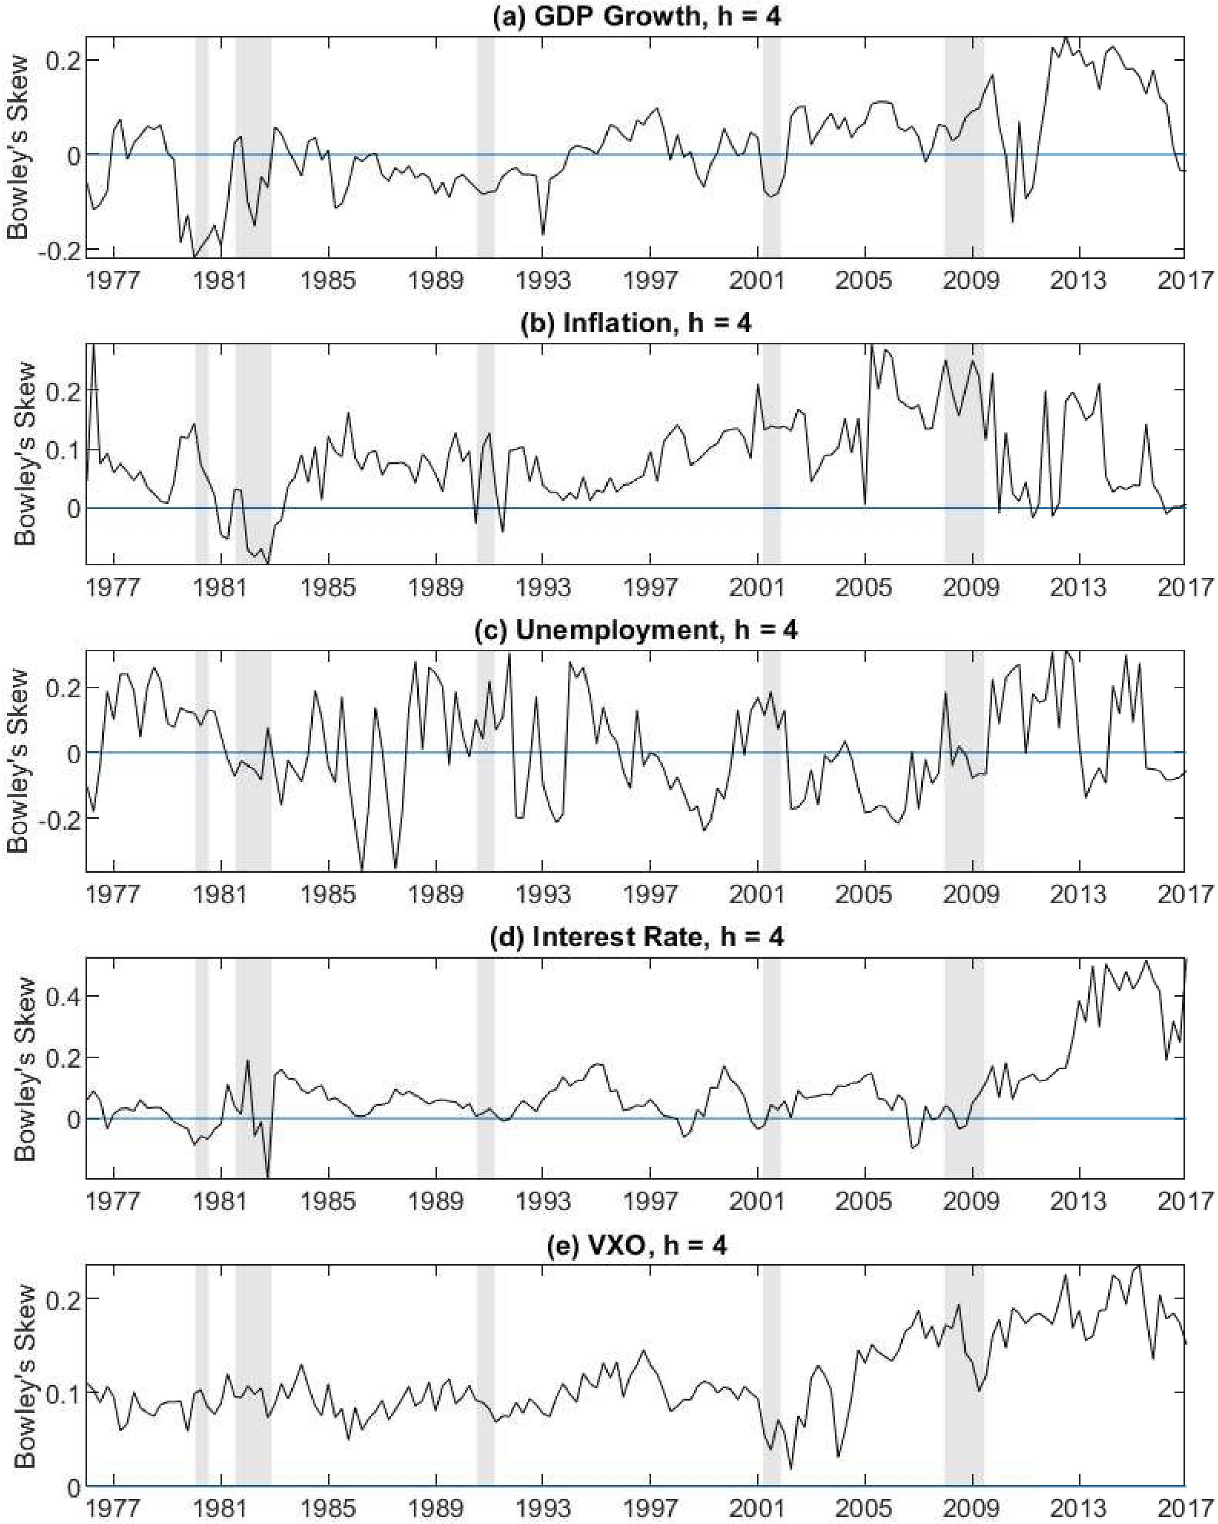

Figure A9: Bowley's skew coefficients of the $h = 4$ quarter ahead real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

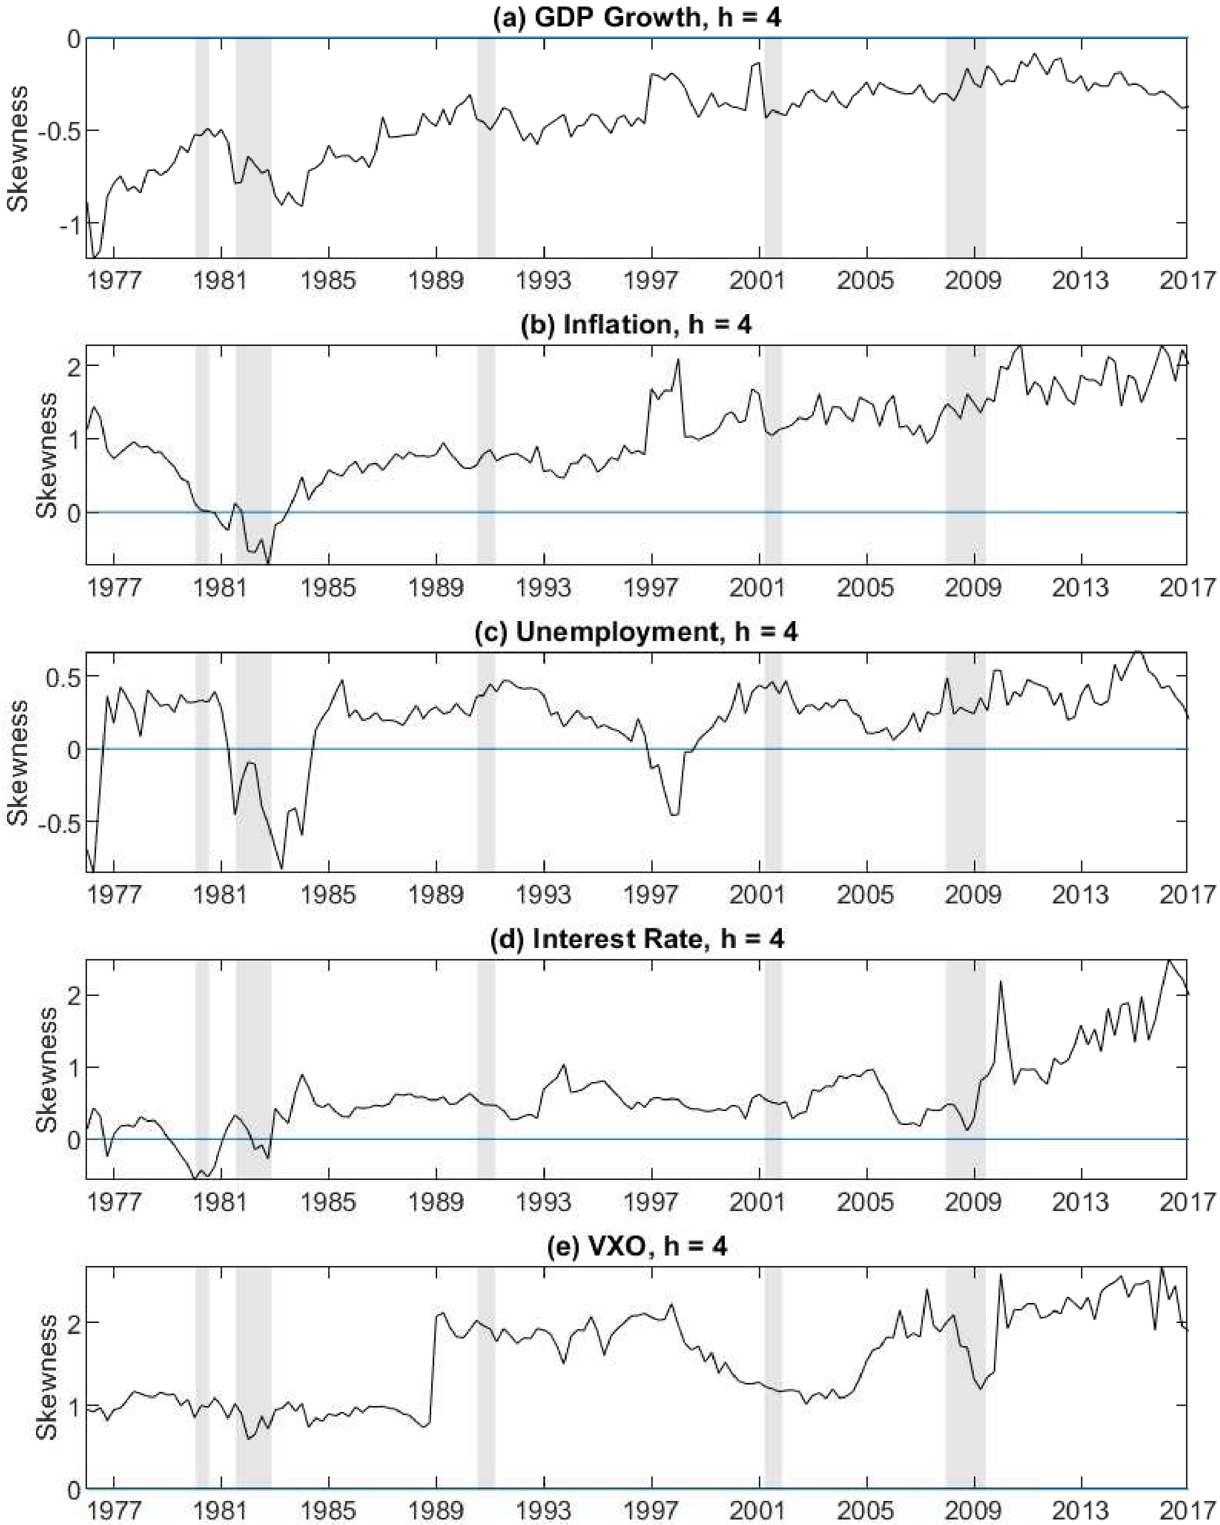

Figure A10: Pearson's skewness of the $h = 4$ quarter ahead real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

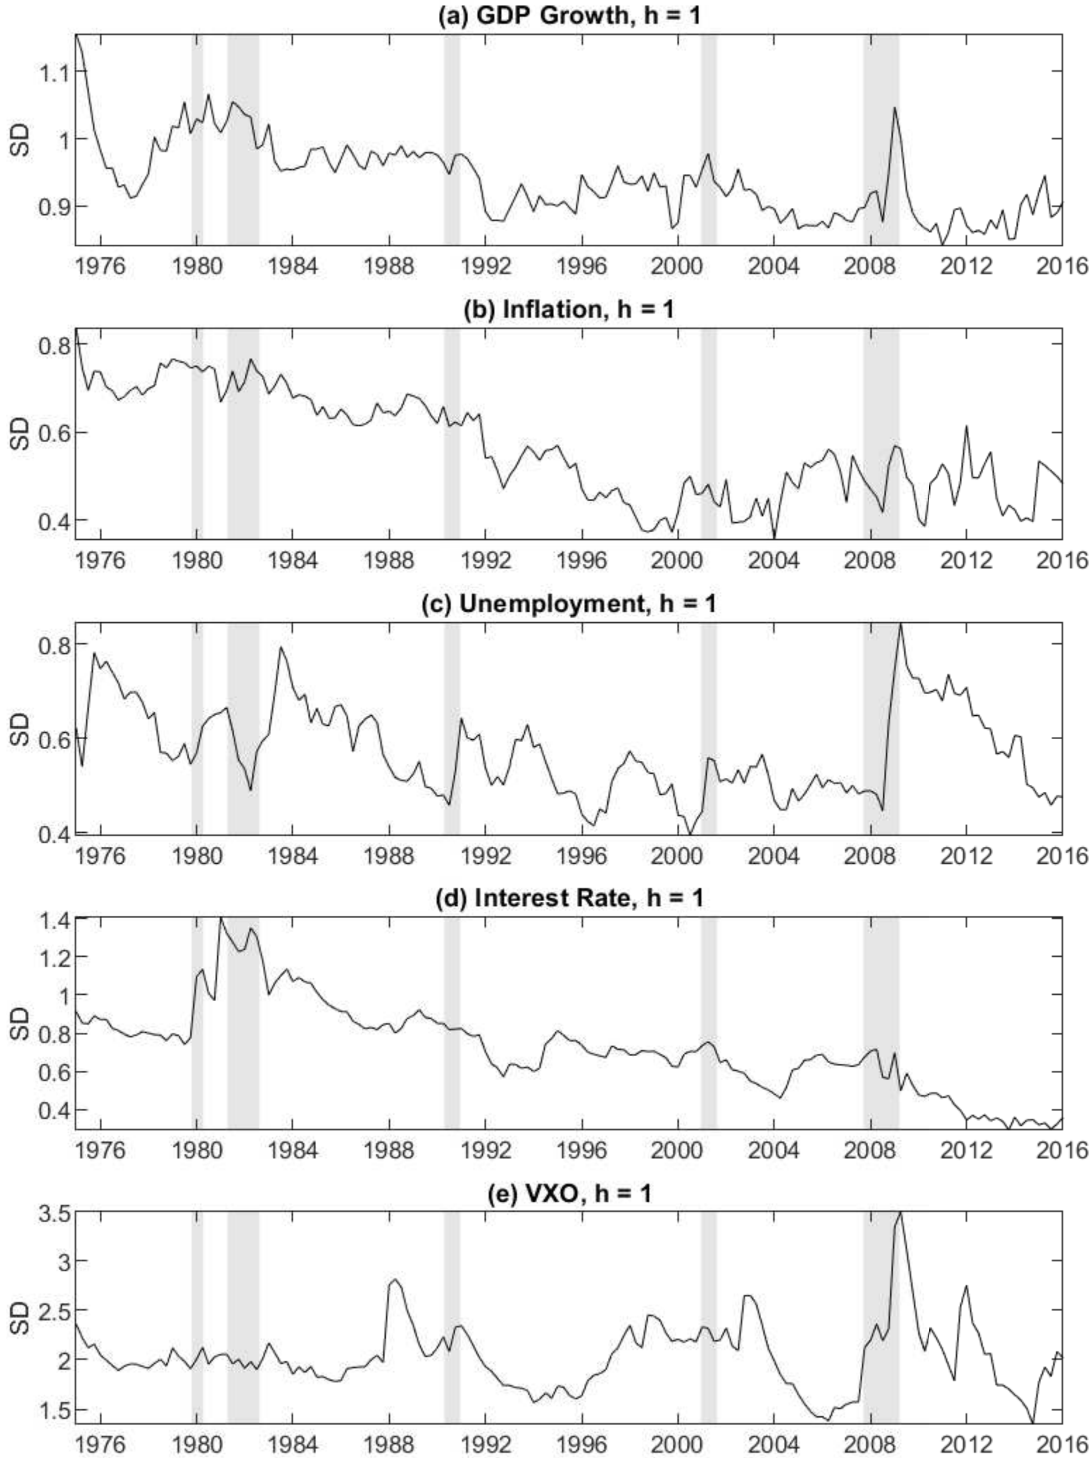

Figure A11: Standard deviation of the $h = 1$ quarter ahead real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

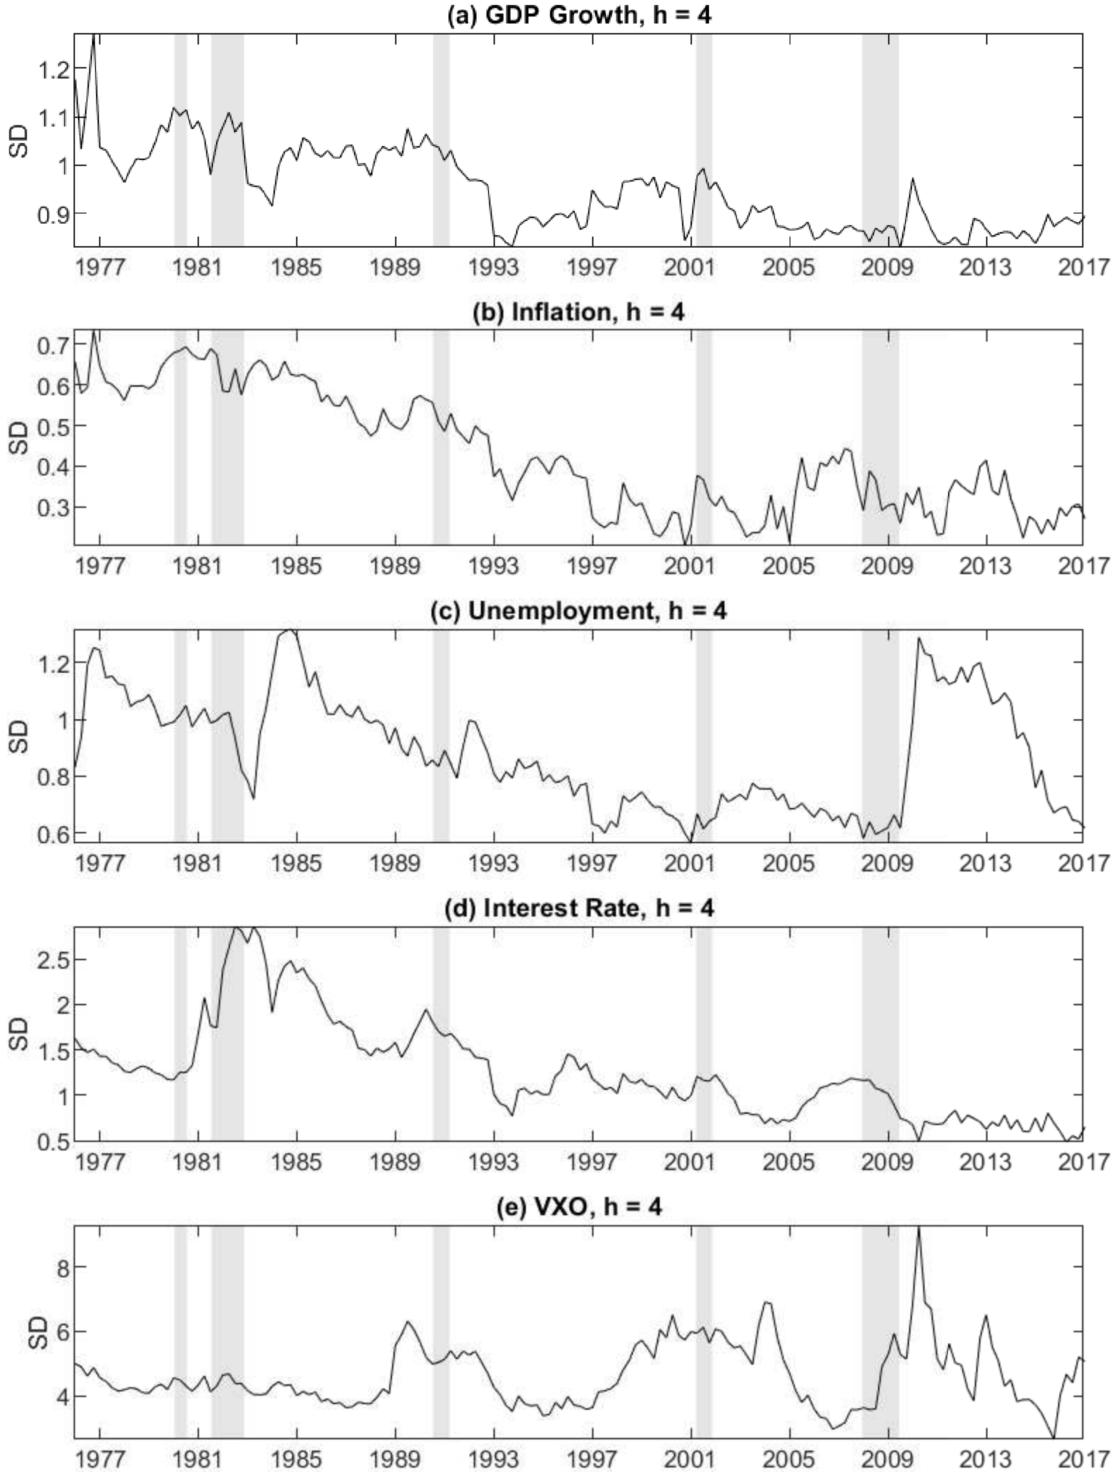

Figure A12: Standard deviation of the $h = 4$ quarter ahead real-time predictive densities from the UCSV-C model. The shaded areas are quarters that correspond to recessions, as defined by the NBER.

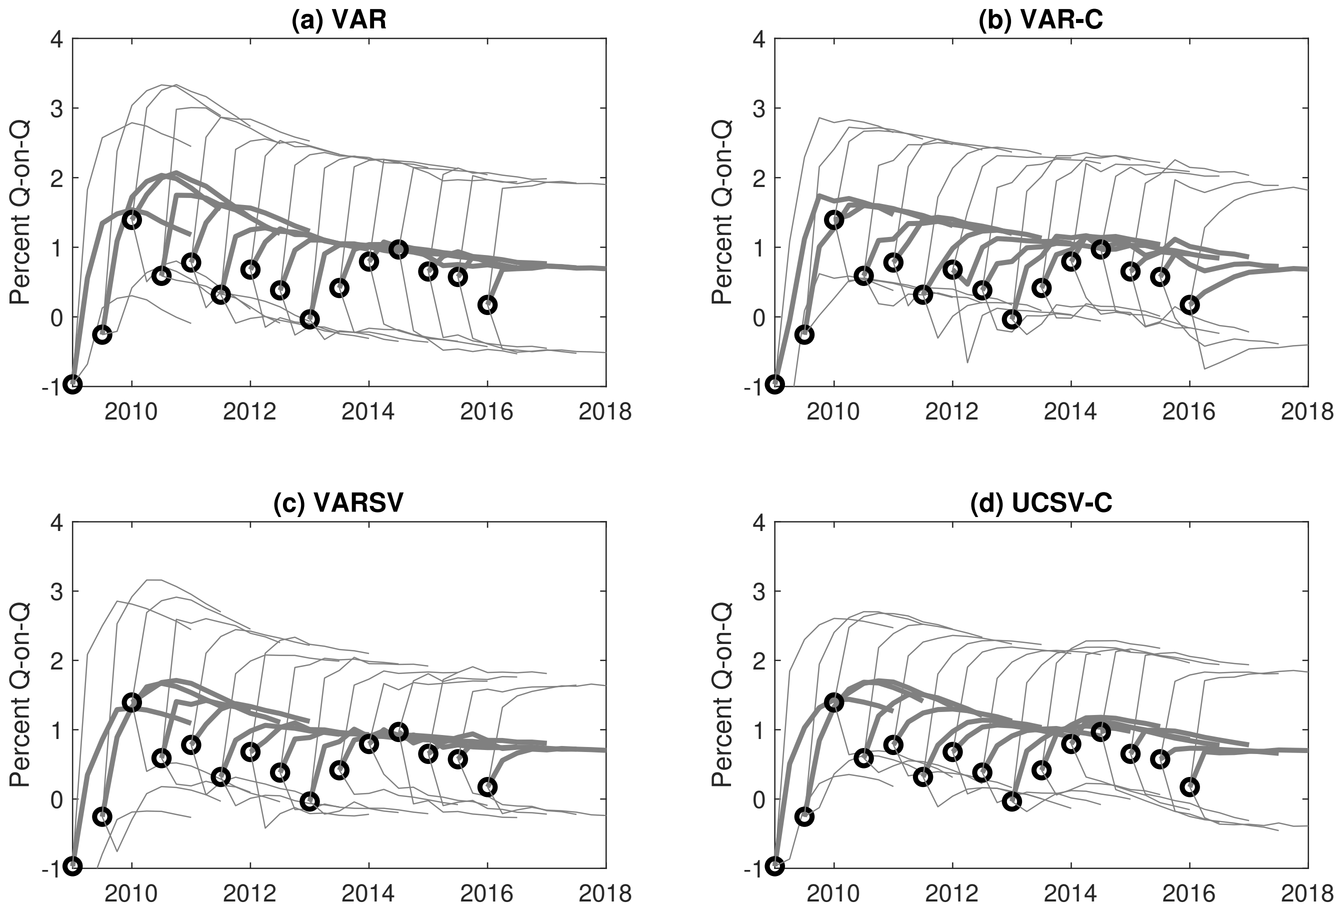

Figure A13: Rolling fan charts showing the real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for GDP Growth. The thick line displays the predictive mean, the thin lines the 10% and 90% predictive quantiles, and the circle shows the prediction origin.

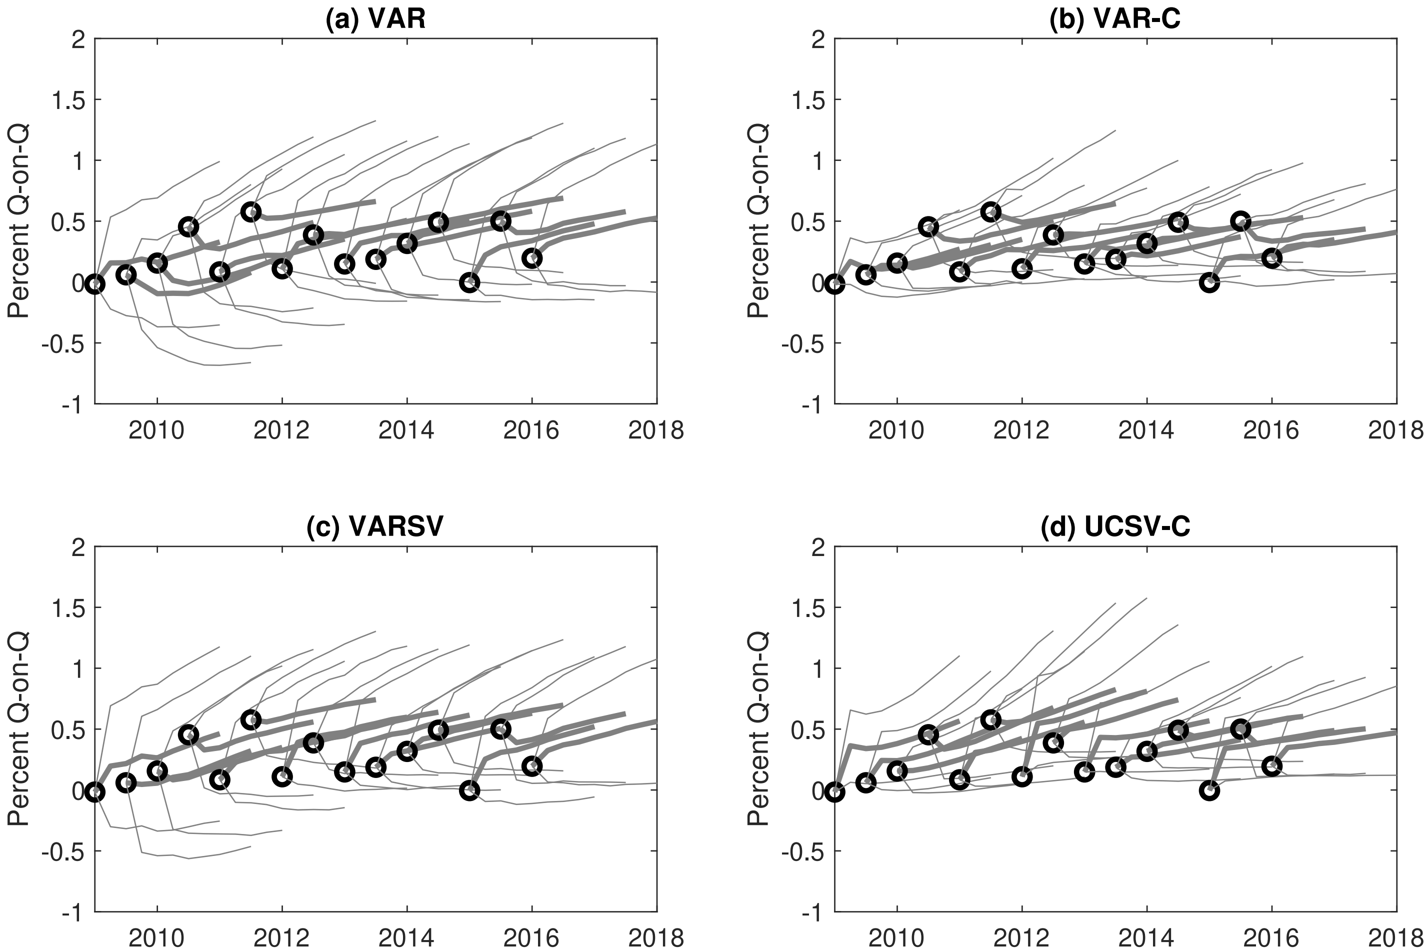

Figure A14: Rolling fan charts showing the real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for inflation. The thick line displays the predictive mean, the thin lines the 10% and 90% predictive quantiles, and the circle shows the prediction origin.

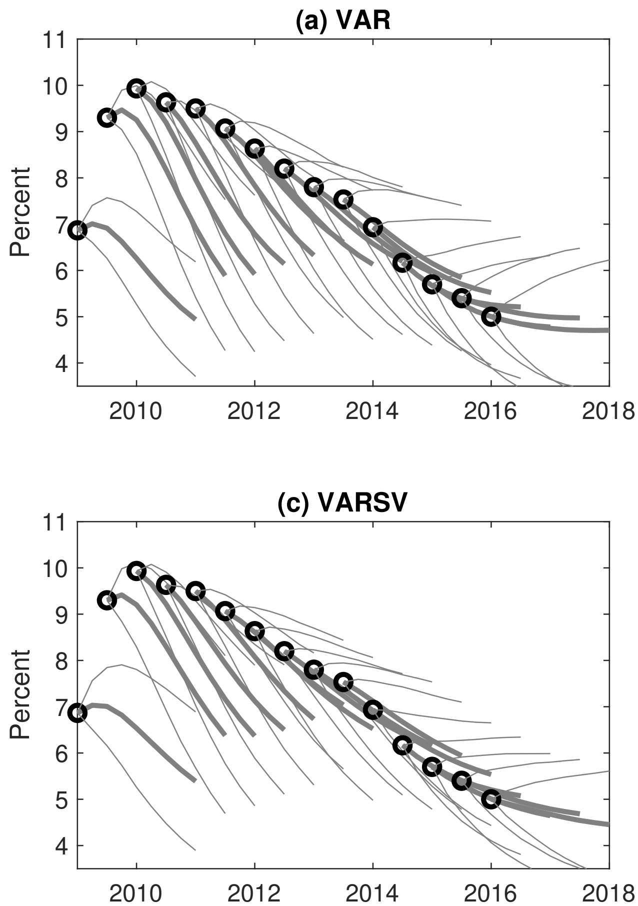

Figure A15: Rolling fan-charts showing the real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for unemployment. The thick line displays the predictive mean, the thin lines the 10% and 90% predictive quantiles, and the circle shows the prediction origin.

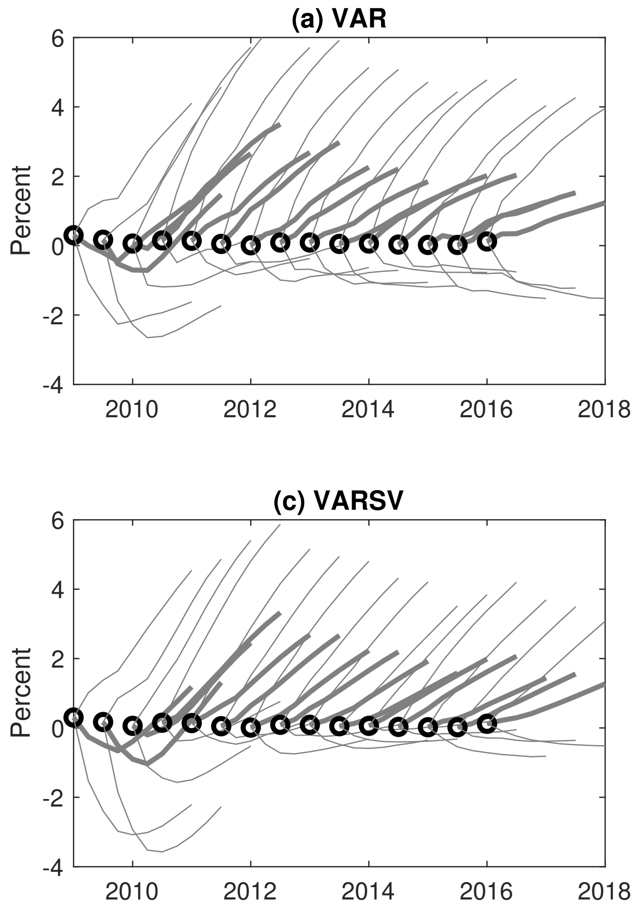

Figure A16: Rolling fan charts showing the real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for the interest rate. The thick line displays the predictive mean, the thin lines the 10% and 90% predictive quantiles, and the circle shows the prediction origin.

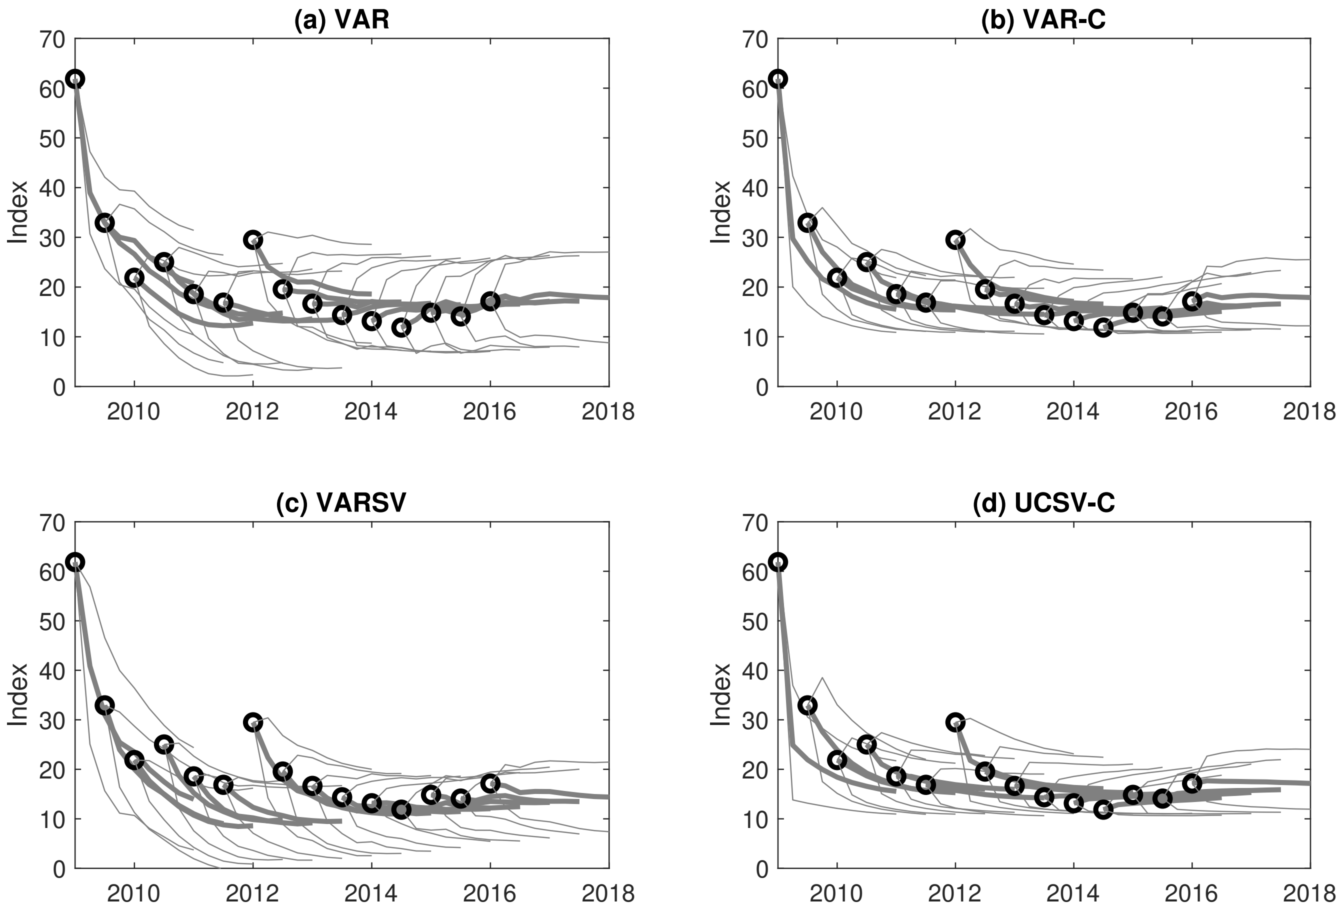

Figure A17: Rolling fan-charts showing the real-time predictions from the VAR, VAR-C, VARSV and UCSV-C models for VXO. The thick line displays the predictive mean, the thin lines the 10% and 90% predictive quantiles, and the circle shows the prediction origin.

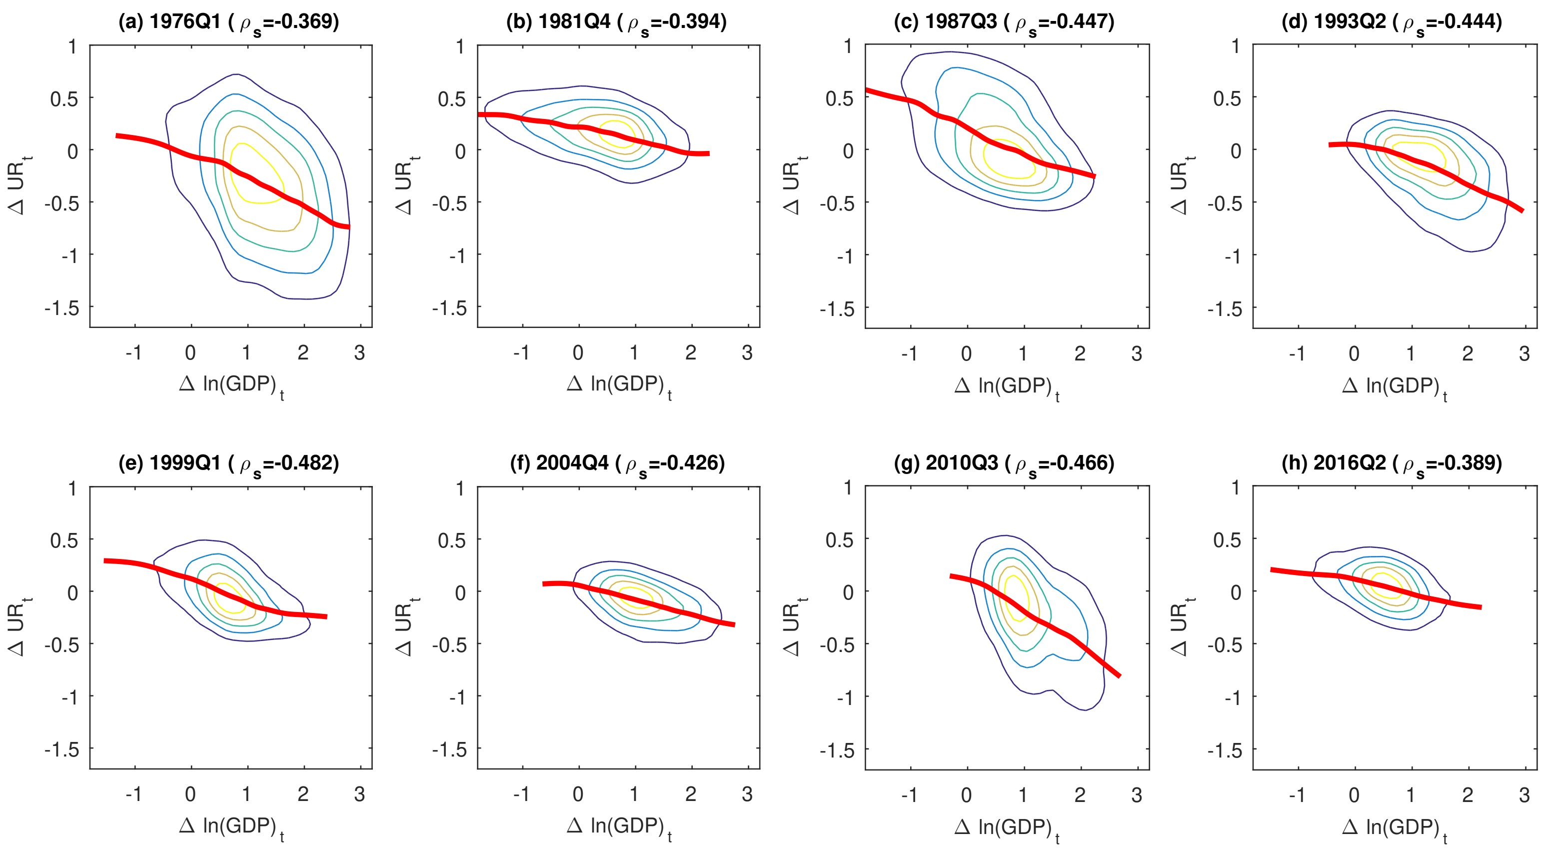

Figure A18: Contour plots of the bivariate current quarter real-time predictions of the change in unemployment and GDP growth, for eight equally-spaced quarters. Also plotted in red is an estimate of $E(\Delta \mathrm{UR}_t|\Delta \ln(\mathrm{GDP})_t)$. The Spearman correlations of these bivariate distributions are also reported in the panel titles in parenthesis.