Inversion Copulas from Nonlinear State Space Models

Michael Stanley Smith* and Worapree Maneesoonthorn

Melbourne Business School, University of Melbourne

NOT FOR CIRCULATION WITHOUT PERMISSION OF AUTHORS

First Version 9 November 2015

* Corresponding author; address for correspondence: Professor Michael Smith, Melbourne Business School, 200 Leicester Street, Carlton, VIC, 3053, Australia. Email: mike.smith@mbs.edu.

Inversion Copulas from Nonlinear State Space Models

NOT FOR CIRCULATION WITHOUT PERMISSION OF AUTHORS

Abstract

While copulas constructed from inverting latent elliptical, or skew-elliptical, distributions are popular, they can be inadequate models of serial dependence in time series. As an alternative, we propose an approach to construct copulas from the inversion of latent nonlinear state space models. This allows for new time series copula models that have the same serial dependence structure as a state space model, yet have an arbitrary marginal distribution—something that is difficult to achieve using other time series models. We examine the time series properties of the copula models, outline measures of serial dependence, and show how to use Bayesian methods to estimate the models. To illustrate the breadth of new copulas that can be constructed using our approach, we consider three example latent state space models: a stochastic volatility model with an unobserved component, a Markov-switching autoregression, and the Gaussian linear case. We use all three inversion copulas to model and forecast quarterly U.S. inflation data. We show how combining the serial dependence structure of the state space models, with flexible asymmetric and heavy-tailed margins, improves the accuracy of the fit substantially in every case. Last, we outline how the approach can be readily extended to construct copulas from the inversion of multivariate nonlinear state space models, thereby also allowing for new and flexible multivariate time series models.

Keywords: Copula Model; Nonlinear Time Series; Nonlinear State Space Models; Inflation

Forecasting.

1 Introduction

Parametric copulas constructed through the inversion of a latent multivariate distribution (Nelsen 2006, sec. 3.1) are popular for the analysis of high-dimensional dependence. For example, Gaussian (Song 2000), t (Embrechts, McNeil & Straumann 2001), skew t (Demarta & McNeil 2005; Smith, Gan & Kohn 2012) and other elliptical (Fang & Fang 2002) distributions are popular choices to form such 'inversion copulas'. However, these existing copulas cannot capture accurately the serial dependence exhibited by many time series. As an alternative, we instead propose a broad new class of inversion copulas, formed by inverting parametric nonlinear state space models. Even though the dimension of such a copula is high, it is parsimonious because its parameters are those of the underlying latent state space model. The copula also has the same serial dependence structure as the state space model. But when such copulas are combined with an arbitrary marginal distribution for the data, they allow for the construction of new time series models. These models are substantially more flexible than the underlying state space models themselves, because the latter usually have rigid margins that are often inconsistent with that observed empirically.

In general, the likelihood function of a nonlinear state space model cannot be expressed in closed form. Similarly, neither can the density of the corresponding inversion copula. Nevertheless, we show how Bayesian techniques can be used to compute posterior estimates of the copula model parameters and associated measures of serial dependence. We employ a Markov chain Monte Carlo (MCMC) sampler, where the wide range of existing methods for efficiently sampling the states of a nonlinear state space model can be employed directly. We also study the time series properties of the copula models. When the latent state space model is non-stationary, the resulting copula model for the data is also non-stationary, but with time invariant univariate margins. Alternatively, when the state space model is stationary and Markov, so is the resulting copula model for the data. In this case we also show how the copula density can be expressed as a product of lower-dimensional 'marginal copulas', which fully characterize the dependence structure. While our focus is on univariate time series, we also extend our approach to evaluate the copula obtained from inverting a multivariate nonlinear state space model.

Copula models are used extensively to model cross-sectional dependence, including between multiple time series; see Patton (2012) for a review. However, their use to capture serial dependence is much more limited. Joe (1997, pp.243-280), Lambert & Vandenhende (2002), Frees & Wang (2006), Domma, Giordano & Perri (2009), Beare (2010) and Smith, Min, Almeida & Czado (2010) use Archimedean, elliptical or decomposable vine copulas to capture serial dependence in univariate time series. While the likelihood is available in closed form for these copulas, they cannot capture the wide range of complex serial dependence structures that the inversion copulas proposed here can. This is also true when using existing copulas to capture serial dependence in multivariate time series. For example, Biller & Nelson (2003) and Smith & Vahey (2015) use the Gaussian copula to capture multivariate serial dependence. However, as shown in Smith & Vahey (2015), this fails to adequately capture features such as conditional heteroscedasticity in the series. Biller (2009), Smith (2015) and Beare & Seo (2015) use vine copulas to capture multivariate serial dependence. While these are more flexible than the Gaussian copula, it is difficult to choose appropriate pair-copula components. Moreover, evaluating the likelihood can become computationally prohibitive as the length or dimension of the time series increases. In comparison to previous copulas employed to capture serial dependence, the inversion copulas suggested here are more general and flexible, simple to specify, and often more parsimonious and easier to estimate.

To highlight the broad range of new copulas that can be formed using our approach, we consider three in detail. They are formed by inversion of three stationary latent state space models. The first is a stochastic volatility model with an unobserved autoregressive mean component, the second is a Markov switching autoregression, and the third is the linear Gaussian state space model. In each case, we consider issues of parameter identification and estimation. A key computational consideration when estimating the first and second copulas, is the repeated evaluation of the quantile function of the univariate margin of the latent state space model. We show how this can be computed quickly and accurately using spline interpolation. We show that the third inversion copula is a Gaussian copula, with correlation matrix equal to that of the latent linear state space model.

We employ each copula to model and forecast quarterly U.S. inflation. This is a long-standing problem in empirical macroeconomics, on which there is a large literature (Faust & Wright 2013). A wide range of univariate time series models have been used previously— see Amisano & Fagan (2013), Clark & Ravazzolo (2015) and Chan (2015) for some recent choices— including our latent state space models. However, all three state space models have marginal distributions that are inconsistent with that observed empirically for inflation, which exhibits strong positive skew and heavy tails. Moreover, the predictive distributions from the models are either exactly or approximately symmetric– a feature that places unpalatably high probability on unlikely events, such as severe deflation. In comparison, the inversion copula models can employ the same serial dependence structure as the latent state space models, but also incorporate much more accurate asymmetric marginal distributions. We show that this not only improves the fit of the time series models, but that it has a large effect on the predictive distributions, including allowing for a substantial degree of asymmetry.

The rest of the paper is organized as follows. In Section 2 we first outline the basic idea of our proposed inversion copula. We then discuss the copula model, estimation method, time series properties, measures of serial dependence and prediction in more detail. In Section 3 we discuss the three inversion copulas that we use as examples, including their application to model and forecast U.S. inflation. In Section 4 we outline how our approach to constructing inversion copulas can be readily extended to multivariate nonlinear time series, while Section 5 concludes.

2 Time Series Inversion Copulas

2.1 The Basic Idea

Consider a T-dimensional copula function $C(\boldsymbol{u})$, with corresponding density $c(\boldsymbol{u}) = \frac{d}{d\boldsymbol{u}}C(\boldsymbol{u})$, for $\boldsymbol{u} = (u_1, \ldots, u_T)$. A popular way to construct such a copula is by transformation from a latent continuous-valued random vector $\boldsymbol{z} = (z_1, \ldots, z_T)$ with a specified distribution. Let $\boldsymbol{z}$ have joint distribution function $F_z(\boldsymbol{z})$, and marginal distribution functions $F_t(z_t)$. Then, by setting $u_t = F_t(z_t)$, the vector $\boldsymbol{u} = (u_1, \ldots, u_T)$ has distribution function

$$C(\boldsymbol{u}) = F_z(F_1^{-1}(u_1), \ldots, F_T^{-1}(u_T)),$$

and density function

$$c(\boldsymbol{u}) = \frac{f_z(\boldsymbol{z})}{\prod_{t=1}^{T} f_t(z_t)}, \tag{2.1}$$

where $f_z(\boldsymbol{z}) = \frac{d}{d\boldsymbol{z}} F_z(\boldsymbol{z})$, $f_t(z_t) = \frac{d}{dz_t} F_t(z_t)$ and $z_t = F_t^{-1}(u_t)$. The transformation ensures that each u_t is marginally uniformly distributed on $[0, 1]$, so that C meets the conditions of a copula function (Nelsen 2006, p.45).

This approach to copula construction is called inversion (Nelsen 2006, sec. 3.1), and we label such a copula an 'inversion copula'. Previous choices for F_z include elliptical (especially the Gaussian and t) and skew t distributions. The dependence properties of the resulting inversion copulas are inherited from those of F_z, although all location, scale and other marginal properties of $F_t(z_t)$ are lost in the transformation. In this paper we propose to construct an inversion copula from a latent nonlinear state space model for $\boldsymbol{z}$. In doing so, we aim to construct new copulas that inherit the rich range of serial dependence structures that state space models allow.

The nonlinear state space model we consider is given by

$$z_t \sim H_t(z_t | \boldsymbol{x}_t; \boldsymbol{\psi}) \tag{2.2}$$

$$\boldsymbol{x}_t \sim K_t(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}; \boldsymbol{\psi}). \tag{2.3}$$

Here, H_t is the distribution function of z_t, conditional on a vector of additional latent states $\boldsymbol{x}_t$, while K_t is the distribution function of $\boldsymbol{x}_t$ conditional on $\boldsymbol{x}_{t-1}$. Parametric distributions are almost always adopted for H_t and K_t, and we do so here with parameters we denote collectively as $\boldsymbol{\psi}$. In the time series literature Equations (2.2) and (2.3) are called the observation and transition distributions, although in the copula context $\boldsymbol{z}$ are not direct observations, but are also latent.

A key requirement in constructing the inversion copula at Equation (2.1) is the evaluation of both the marginal distribution and density functions of z_t, which are

$$\begin{aligned} F_t(z_t|\boldsymbol{\psi}) &= \int H_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) f(\boldsymbol{x}_t|\boldsymbol{\psi}) \mathrm{d}\boldsymbol{x}_t \\ f_t(z_t|\boldsymbol{\psi}) &= \int h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) f(\boldsymbol{x}_t|\boldsymbol{\psi}) \mathrm{d}\boldsymbol{x}_t \,. \end{aligned} \tag{2.4}$$

Here, $h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) = \frac{d}{dz_t} H_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi})$, and $f(\boldsymbol{x}_t|\boldsymbol{\psi})$ is the marginal density of the state variable $\boldsymbol{x}_t$ which can be derived analytically from the transition distribution for most state space models used in practice. Evaluation of the integrals in Equation (2.4) is typically straightforward either analytically or numerically; for example, we show later how these margins can be computed efficiently for three different state space models.

A more challenging problem is the evaluation of the numerator in Equation (2.1). To compute this, the state vector $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T)$ needs to be integrated out, with

$$\begin{aligned} f_z(\boldsymbol{z}|\boldsymbol{\psi}) &= \int f(\boldsymbol{z}|\boldsymbol{x},\boldsymbol{\psi}) f(\boldsymbol{x}|\boldsymbol{\psi}) \mathrm{d}\boldsymbol{x} \\ &= \int \prod_{t=1}^{T} h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) \prod_{t=2}^{T} k_t(\boldsymbol{x}_t|\boldsymbol{x}_{t-1};\boldsymbol{\psi}) f(\boldsymbol{x}_1;\boldsymbol{\psi}) \mathrm{d}\boldsymbol{x} \,, \end{aligned}$$

where $k_t(\boldsymbol{x}_t|\boldsymbol{x}_{t-1};\boldsymbol{\psi}) = \frac{d}{d\boldsymbol{x}_t} K_t(\boldsymbol{x}_t|\boldsymbol{x}_{t-1};\boldsymbol{\psi})$. There are many methods suggested for evaluating this high-dimensional integral; for example, see Shephard & Pitt (1997), Durbin & Koopman (2002), Stroud, Müller & Polson (2003), Godsill, Doucet & West (2004), Jungbacker & Koopman (2007), Scharth & Kohn (2013) and references therein. In this paper we show how popular Markov chain Monte Carlo (MCMC) methods for solving this problem can also be used to estimate the inversion copula.

Because all features of the marginal distribution of z_t— including the marginal moments— are lost when forming the copula, parameters in a state space model that uniquely affect these are unidentified and can be excluded from $\boldsymbol{\psi}$. Moreover, where possible we also normalize H_t at Equation (2.2) so that the marginal distribution of z_t has zero mean and unit variance. While this has no effect on the copula function C, we find that this aids identification of the posterior distribution of the states $\boldsymbol{x}$.

2.2 Copula Model and Estimation

While inversion copulas can be used equally to model dependence in either continuous or discrete-valued time series data, we focus here on the former case. Let $\boldsymbol{y} = (y_1, \ldots, y_T)$ be a vector of time series observations, then its joint density is decomposed in a copula model as

$$f(\boldsymbol{y}) = c(\boldsymbol{u}) \prod_{t=1}^{T} g(y_t)\,. \tag{2.5}$$

Here, g is the marginal density of each observation, which we assume to be time invariant, and c is a T-dimensional copula density. The elements of $\boldsymbol{u}$ are $u_t = G(y_t)$, with $G(y) = \int_{-\infty}^{y} g(s)\mathrm{d}s$ being the marginal distribution function of the observations (not to be confused with the marginal distribution function $F_t(z_t)$ of the latent state space model). All serial dependence in the data is captured by the copula. All marginal features of the data are captured by G, which can be modeled separately, and either parametrically or non-parametrically.

For c in Equation (2.5), we use an inversion copula constructed from a parametric state space model. Assuming a parametric margin $G(y_t;\boldsymbol{\theta})$ with parameters $\boldsymbol{\theta}$, the likelihood of the copula model is

$$f(\boldsymbol{y}|\boldsymbol{\psi},\boldsymbol{\theta}) = c(\boldsymbol{u};\boldsymbol{\psi}) \prod_{t=1}^{T} g(y_t;\boldsymbol{\theta}) = f_z(\boldsymbol{z}|\boldsymbol{\psi}) \prod_{t=1}^{T} \frac{g(y_t;\boldsymbol{\theta})}{f_t(z_t|\boldsymbol{\psi})}\,, \tag{2.6}$$

where the reliance of the copula density on $\boldsymbol{\psi}$ is made explicit. It is popular to estimate nonlinear state space models using MCMC methods where the latent states $\boldsymbol{x}$ are generated as a block. We adopt this approach to also estimate the inversion copula, although note that other methods based either on particle filtering or importance sampling also have the potential to be employed here. Conditional on the states, the likelihood is

$$f(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\psi},\boldsymbol{\theta}) = f(\boldsymbol{z}|\boldsymbol{x},\boldsymbol{\psi}) \prod_{t=1}^{T} \frac{g(y_t;\boldsymbol{\theta})}{f_t(z_t|\boldsymbol{\psi})} = \prod_{t=1}^{T} h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) \frac{g(y_t;\boldsymbol{\theta})}{f_t(z_t|\boldsymbol{\psi})}\,, \tag{2.7}$$

with $z_t = F_t^{-1}(G(y_t;\boldsymbol{\theta})|\boldsymbol{\psi})$. Adopting independent priors $\pi_\psi(\boldsymbol{\psi})$ and $\pi_\theta(\boldsymbol{\theta})$, estimation and inference from the model can be based on the sampler below.

Sampling Scheme

Step 1. Generate from $f(\boldsymbol{x}|\boldsymbol{\psi},\boldsymbol{\theta},\boldsymbol{y}) = f(\boldsymbol{x}|\boldsymbol{\psi},\boldsymbol{\theta},\boldsymbol{z}) \propto \prod_{t=1}^{T} h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) f(\boldsymbol{x}|\boldsymbol{\psi})$.

Step 2. Generate from $f(\boldsymbol{\psi}|\boldsymbol{x},\boldsymbol{\theta},\boldsymbol{y}) \propto \left(\prod_{t=1}^{T} h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi})/f_t(z_t|\boldsymbol{\psi})\right) f(\boldsymbol{x}|\boldsymbol{\psi})\pi_\psi(\boldsymbol{\psi})$.

Step 3. Generate from $f(\boldsymbol{\theta}|\boldsymbol{x},\boldsymbol{\psi},\boldsymbol{y}) \propto \left(\prod_{t=1}^{T} h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi}) g(y_t;\boldsymbol{\theta})/f_t(z_t|\boldsymbol{\psi})\right) \pi_\theta(\boldsymbol{\theta})$.

Crucially, once $\boldsymbol{z} = (z_1,\ldots,z_T)$ is computed, Step 1 is exactly the same as that for the underlying state space model, so that existing procedures for generating the latent states can be employed. Step 2 can be undertaken using Metropolis-Hastings, with a proposal based on a numerical or other approximation to the conditional posterior. However, for some state space models it can be more computationally efficient to generate sub-vectors of $\boldsymbol{\psi}$ from their conditional posteriors, using separate steps. For non-parametric marginal models it is often attractive to follow Shih & Louis (1995) and others, and employ two-stage estimation, so that Step 3 is not required. However, for parametric models, $\boldsymbol{\theta}$ can be generated at Step 3 using a Metropolis-Hastings step. It is initially appealing is to use the posterior of the marginal model as a proposal, with density $q(\boldsymbol{\theta}) \propto \prod_{t=1}^{T} g(y_t;\boldsymbol{\theta})\pi_\theta(\boldsymbol{\theta})$. However, this should be avoided because when there is strong dependence— precisely the circumstance where the copula model is most useful— this proposal can be a poor approximation to the conditional posterior.

We outline later key aspects of the samplers that are used to estimate the three specific inversion copulas considered. However, note that at Steps 2 and 3 the value of $\boldsymbol{z}$ needs updating, although not at the end of Step 1. When updating $z_t = F_t^{-1}(G(y_t;\boldsymbol{\theta})|\boldsymbol{\psi})$, evaluation of the quantile functions F_t^{-1} is often the most computationally demanding aspect of the sampling scheme. In the Appendix we outline how to achieve this quickly and accurately for a stationary nonlinear state space model using spline interpolation.

2.3 Time Series Properties

Inversion copulas at Equation (2.1) can be constructed from either stationary or non-stationary state space models for the latent series $\{z_t\}$. Here, stationarity refers to strong or strict stationarity, rather than weak or covariance stationarity; eg. see Brockwell &

Davis (1991, p.12). When a non-stationary state space model is used, the copula model at Equation (2.5) is a non-stationary time series model for $\{y_t\}$, but with a time invariant univariate marginal $G(y_t)$. (Note that the multivariate marginals $f(y_s, \ldots, y_t)$ for $s < t$ can be time-varying, however.) This is an uncommon variant of a non-stationary time series, although one with strong potential in applied analysis.

Conversely, when a stationary latent state space model is employed, then the time series model for the data is also stationary, as summarized below.

Lemma *(Stationarity)*

If $\{z_t\}$ is a stationary time series with marginal distribution function F_1, then:

(i) the time series $\{u_t\}$, with $u_t = F_1(z_t)$, is also stationary; and

(ii) the time series $\{y_t\}$, with $y_t = G^{-1}(F_1(z_t))$, is also stationary.

Proof of the above follows from the definition of strong stationarity and that both F_1 and $G^{-1} \circ F_1$ are one-to-one functions. It is also straightforward to show that the three series $\{z_t\}$, $\{u_t\}$ and $\{y_t\}$ share the same Markov order.

If the series is stationary with Markov order p, then an alternative representation of the copula density at Equation (2.1) can be derived as follows. First, denote the univariate marginal density of the latent process as $f_1(z_t) = \frac{d}{dz_t} F_1(z_t)$. Also, denote the (time invariant) r-dimensional marginal density of the latent process for $r \geq 2$ as

$$f^{(r)}(\boldsymbol{z}_{t-r+1:t}) = \int \cdots \int \left(\prod_{s=t-r+1}^{t} h_s(z_s|\boldsymbol{x}_s; \boldsymbol{\psi}) \right) f(\boldsymbol{x}_{t-r+1:t}|\boldsymbol{\psi}) \mathrm{d}\boldsymbol{x}_{t-r+1:t}\,, \tag{2.8}$$

where $t \geq r$ and $f(\boldsymbol{x}_{t-r+1:t}|\boldsymbol{\psi})$ is the r-dimensional marginal density of the states $\boldsymbol{x}_{t-r+1:t}$. Then the r-dimensional marginal copula density can be defined as

$$c^{(r)}(\boldsymbol{u}_{t-r+1:t}) = \frac{f^{(r)}(\boldsymbol{z}_{t-r+1:t})}{\prod_{s=t-r+1}^{t} f_1(z_s)}\,. \tag{2.9}$$

The $(p+1)$-dimensional marginal copula with density $c^{(p+1)}$ fully characterizes the serial dependence structure of the stationary Markov series. The copula density can be written in terms of both it, and lower dimensional copula margins, as

$$\begin{aligned} c(\boldsymbol{u}) &= \prod_{t=p+1}^{T} f(u_t|\boldsymbol{u}_{t-p:t-1}) \prod_{t=2}^{p} f(u_t|\boldsymbol{u}_{1:t-1}) \\ &= \prod_{t=p+1}^{T} \frac{c^{(p+1)}(\boldsymbol{u}_{t-p:t})}{c^{(p)}(\boldsymbol{u}_{t-p:t-1})} \prod_{t=2}^{p} \frac{c^{(t)}(\boldsymbol{u}_{1:t})}{c^{(t-1)}(\boldsymbol{u}_{1:t-1})}\,, \end{aligned} \tag{2.10}$$

where we define $c^{(r)} = 1$ whenever $r \leq 1$, and a product term to be equal to unity when its upper limit is less than its lower limit. For example, for a Markov order $p = 1$ series, $c(\boldsymbol{u}) = \prod_{t=2}^{T} c^{(2)}(\boldsymbol{u}_{t-1:t})$, with $c^{(2)}(u_1, u_2) = f^{(2)}(z_1, z_2)/f_1(z_1)f_1(z_2)$. For some state space models, employing Equation (2.10) can provide an expression for the likelihood that is more attractive computationally because it involves evaluation of integrals of p or lower dimension.

2.4 Serial Dependence

Measures of serial dependence at a given lag $l \geq 1$, can be computed from the inversion copula. These measures are time invariant when the latent state space model is stationary, and time varying otherwise. They include Kendall's tau, Spearman's rho, and measures of tail dependence; see Nelsen (2006, Ch. 5) for an introduction to such measures of concordance. These can be computed from the bivariate copula of the observations of the series at times s and $t = s + l$ as follows. If the density and distribution functions of (z_s, z_t) are denoted as

$$\begin{aligned} f^{s,t}(z_s, z_t|\boldsymbol{\psi}) &= \int\int h_s(z_s|\boldsymbol{x}_s;\boldsymbol{\psi})h_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi})f(\boldsymbol{x}_s,\boldsymbol{x}_t|\boldsymbol{\psi})\mathrm{d}\boldsymbol{x}_s\mathrm{d}\boldsymbol{x}_t\,, \\ F^{s,t}(z_s, z_t|\boldsymbol{\psi}) &= \int\int H_s(z_s|\boldsymbol{x}_s;\boldsymbol{\psi})H_t(z_t|\boldsymbol{x}_t;\boldsymbol{\psi})f(\boldsymbol{x}_s,\boldsymbol{x}_t|\boldsymbol{\psi})\mathrm{d}\boldsymbol{x}_s\mathrm{d}\boldsymbol{x}_t\,, \end{aligned} \tag{2.11}$$

then the bivariate copula function is

$$C^{s,t}(u_s, u_t|\boldsymbol{\psi}) = F^{s,t}(F_s^{-1}(u_s|\boldsymbol{\psi}), F_t^{-1}(u_t|\boldsymbol{\psi})|\boldsymbol{\psi})\,,$$

with corresponding density $c^{s,t}(u_s, u_t|\boldsymbol{\psi}) = f^{s,t}(z_s, z_t|\boldsymbol{\psi})/f_s(z_s|\boldsymbol{\psi})f_t(z_t|\boldsymbol{\psi})$. Kendall's tau, Spearman's rho and the lower tail dependence for quantile $0 < \alpha < 0.5$, are then

$$\begin{aligned}
\tau_{t,s} &= 4\int\int C^{s,t}(u,v)c^{s,t}(u,v)\mathrm{d}u\mathrm{d}v - 1 \\
r_{t,s} &= 12\int\int uvc^{s,t}(u,v)\mathrm{d}u\mathrm{d}v - 3 \\
\lambda_{t,s}^{--}(\alpha) &= \Pr(u_t < \alpha | u_s < \alpha) = C^{t,s}(\alpha,\alpha)/\alpha\,.
\end{aligned}$$

Tail dependencies in other quadrants, $\lambda_{t,s}^{++}(\alpha) = \Pr(u_t > 1-\alpha|u_s > 1-\alpha)$, $\lambda_{t,s}^{+-}(\alpha) = \Pr(u_t > 1-\alpha|u_s < \alpha)$ and $\lambda_{t,s}^{-+}(\alpha) = \Pr(u_t < \alpha|u_s > 1-\alpha)$ are computed similarly. When the latent state space model is stationary, we write the measures as $\tau_l, r_l, \lambda_l^{--}(\alpha), \lambda_l^{++}(\alpha), \lambda_l^{+-}(\alpha)$ and $\lambda_l^{-+}(\alpha)$, with $l = t - s$. Moreover, we note that for $l = 1$, the marginal copula $c^{(2)} = c^{s+1,s}$.

For some state space models, evaluation of the integrals at Equation (2.11) may be impractical. In this case, the dependence measures are readily calculated via simulation. For example, in the stationary case this involves generating iterates from the marginal copula with density $c^{(l+1)}$. To do this, simply generate $(z_1, \ldots, z_{l+1})$ from the state space model, and then transform each value to $u_t = F_t(z_t)$. If the iterates $(u_1^{[j]}, \ldots, u_{l+1}^{[j]})$, $j = 1, \ldots, J$, are generated in this way, then Spearman's rho for pairwise dependence between y_{t+l} and y_t is

$$r_l = 12E(u_{l+1}u_1) - 3 \approx \frac{12}{J}\sum_{j=1}^{J}\left(u_{l+1}^{[j]}u_1^{[j]}\right) - 3\,,$$

and the lower tail dependence for lag l is

$$\lambda_l^{--}(\alpha) \approx \frac{\sum_{j=1}^{J}\mathcal{I}(u_{l+1}^{[j]} < \alpha, u_1^{[j]} < \alpha)}{J\alpha}\,,$$

where $\mathcal{I}(A) = 1$ if A is true, and zero otherwise. The other measures are estimated similarly. In general, large Monte Carlo samples (e.g. $J = 50{,}000$) can be required for these estimates to be accurate. However, simulating from the marginal copula is both fast and can be undertaken in parallel, so that it is not a problem in practise.

2.5 Smoothed Estimates and Prediction

We also evaluate the following smoothed estimates of the first two moments of the time series. If we denote the posterior distribution augmented with the latent states $\boldsymbol{x}$ as $F_P(\boldsymbol{x}, \boldsymbol{\psi}, \boldsymbol{\theta}|\boldsymbol{y})$, then we compute the smoothed estimate of the mean

$$E_{F_P}(y_t) = \int E(y_t|\boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\psi})\mathrm{d}F_P \approx \frac{1}{J}\sum_{j=1}^{J} E(y_t|\boldsymbol{x}_t^{[j]}, \boldsymbol{\theta}^{[j]}, \boldsymbol{\psi}^{[j]})\,,$$

where $\{\boldsymbol{x}_t^{[j]}, \boldsymbol{\theta}^{[j]}, \boldsymbol{\psi}^{[j]}\}$, $j = 1, \ldots, J$, are the draws from the sampler. It is straightforward to show that for the inversion copula,

$$E(y_t|\boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\psi}) = \int G^{-1}(F_1(z_t|\boldsymbol{\psi}); \boldsymbol{\theta}) h_t(z_t|\boldsymbol{x}_t; \boldsymbol{\psi})\boldsymbol{d}z_t\,.$$

The univariate integral above is readily computed numerically at the end of each sweep of the MCMC scheme. The smoothed estimate of the standard deviation is $SD_{F_P}(y_t) = \left(E_{F_P}(y_t^2) - E_{F_P}(y_t)\right)^{1/2}$, where $E_{F_P}(y_t^2) = \int E(y_t^2|\boldsymbol{x}_t, \boldsymbol{\theta}, \boldsymbol{\psi})\mathrm{d}F_P$ is computed similarly. Interestingly, our later empirical work shows that even when the underlying latent state space model is homoscedastic, our copula model can exhibit some limited time variation in $SD_{F_P}(y_t)$ when the margin G is asymmetric.

In general, evaluation of the predictive distribution $F(y_{T+l}|\boldsymbol{y}, \boldsymbol{\psi}, \boldsymbol{\theta})$, $l \geq 1$, is difficult analytically, so that we instead evaluate it via simulation— a process which is both straightforward and fast. To simulate an iterate from the predictive distribution, we first simulate a ray of values $(z_{T+1}, \ldots, z_{T+l}) \sim F(z_{T+1}, \ldots, z_{T+l}|\boldsymbol{z}, \boldsymbol{\psi}, \boldsymbol{\theta})$ from the predictive distribution of the state space model. Then z_{T+l} is transformed to provide an iterate $y_{T+l} = G^{-1}(F_1(z_{T+l}|\boldsymbol{\psi})|\boldsymbol{\theta})$. Simulation can be undertaken conditional on either the point estimates of $(\boldsymbol{\psi}, \boldsymbol{\theta})$, or over the sample of parameter values from the posterior. The latter approach integrates out parameter uncertainty in the usual Bayesian fashion.

3 Three Inversion Copulas

We consider three time series inversion copulas in detail. The first two are constructed from two popular nonlinear state space models, while the last is constructed from the linear Gaussian state space model. In each case, we outline constraints required to identify the parameters when forming the copula by inversion, as well as how to implement the generic sampler in Section 2.2. While the first two copulas cannot be expressed in closed form, we show how the last is a Gaussian copula.

To illustrate, we fit each of the copula models to U.S. inflation from 1954:Q1 to 2013:Q4. We consider the differences $y_t = \log(P_t) - \log(P_{t-1})$ in the logarithm of the (seasonally adjusted) quarterly GDP price deflator P_t, sourced from the FRED database of the Federal Reserve Bank of Saint Louis. Figure 1(a) plots the time series of the $T = 240$ quarterly observations, while Figure 1(b) plots a histogram of the data. The distribution of inflation is far from symmetric, with sample skew 1.329 and kurtosis 4.515, and a Shapiro & Wilk (1965) test for normality is rejected at any meaningful significance level. A wide range of time series models have been fitted to quarterly inflation data previously (Faust & Wright 2013; Clark & Ravazzolo 2015), including the three state space models considered here. However, these three models— in fact, almost all time series models used previously— have margins that are inconsistent with that observed empirically. We show that combining the inversion copulas with more flexible margins produces time series models that have an improved fit. Moreover, we show that this also has a sizable effect on the predictive distributions produced during the onset of the recent recession; a period of particular interest to economists.

3.1 Stochastic Volatility Inversion Copula

The conditional variance of many financial and economic time series exhibit strong positive serial dependence. A popular model used to capture this is the stochastic volatility model, although a major limitation is that its marginal distribution is symmetric, which is inconsistent with most series. Our approach allows for the construction of time series models that have the same serial dependence as a stochastic volatility model, but also arbitrary marginals that can be asymmetric.

3.1.1 The Copula

We consider the stochastic volatility model with an unobserved autoregressive component, given by

$$
\begin{aligned}
z_t|\boldsymbol{x}_t, \boldsymbol{\psi} &\sim N(\mu_t, \exp(h_t)) \\
\mu_t|\boldsymbol{x}_{1:t}, \boldsymbol{\psi} &\sim N(\bar{\mu} + \rho_\mu(\mu_{t-1} - \bar{\mu}), \sigma_\mu^2) \\
h_t|\boldsymbol{x}_{1:t}, \boldsymbol{\psi} &\sim N(\bar{h} + \rho_h(h_{t-1} - \bar{h}), \sigma_h^2)\,,
\end{aligned}
\tag{3.1}
$$

where $\boldsymbol{x}_t = (\mu_t, h_t)$ is the state. We constrain $|\rho_\mu| < 1$ and $|\rho_h| < 1$, ensuring both $\boldsymbol{x}_t$ and z_t are (strongly) stationary first order Markov processes. The parameters $\bar{\mu}$ and $\bar{h}$ only affect the marginal mean and logarithm of the marginal variance of z_t, so that we set both to zero. With these constraints, $E(z_t) = 0$ and $\text{Var}(z_t) = \kappa^2 = s_\mu^2 + \exp(s_h^2/2)$, where $s_\mu^2 = \sigma_\mu^2/(1 - \rho_\mu^2)$ and $s_h^2 = \sigma_h^2/(1 - \rho_h^2)$. To normalize the marginal variance of z_t to one, we rewrite Equation (3.1) so that

$$
z_t|\boldsymbol{x}_t, \boldsymbol{\psi} \sim N\left(\frac{\mu_t}{\kappa}, \frac{\exp(h_t)}{\kappa^2}\right)\,,
\tag{3.2}
$$

The dependence parameters of the resulting inversion copula are therefore $\boldsymbol{\psi} = \{\rho_\mu, \rho_h, \sigma_\mu^2, \sigma_h^2\}$.

The marginal density at Equation (2.4) is

$$
f_1(z; \boldsymbol{\psi}) = \int\int \phi_1\left(z; \frac{\mu}{\kappa}, \frac{\exp(h)}{\kappa^2}\right) \phi_1(h; 0, s_h^2)\phi_1(\mu; 0, s_\mu^2) d\mu dh\,,
$$

where $\phi_1(z; a, b^2)$ is a univariate Gaussian density with mean a and variance b^2. The inner integral in μ can be computed analytically and (with a little algebra) the marginal density and distribution functions are

$$
\begin{aligned}
f_1(z; \boldsymbol{\psi}) &= \int \phi_1(z; 0, w(h)^2)\phi_1(h; 0, s_h^2) dh \\
F_1(z; \boldsymbol{\psi}) &= \int \Phi_1(z; 0, w(h)^2)\phi_1(h; 0, s_h^2) dh\,,
\end{aligned}
$$

with $w(h)^2 = (s_\mu^2 + \exp(h))/\kappa^2$. Computing the (log) copula density at Equation (2.1) requires evaluating $\log(f_1)$ and the quantile function F_1^{-1} at all T observations. The Appendix outlines how to compute these numerically using spline interpolation of both functions. In our empirical work we find these spline-based approximations to be accurate within 5 to 9 decimals places, and fast because they require direct evaluation of F_1^{-1} at only one point.

The Appendix also outlines how to compute the marginal copula density $c^{(2)}(u_t, u_{t-1}|\boldsymbol{\psi}) = f^{(2)}(z_t, z_{t-1}|\boldsymbol{\psi})/f_1(z_t|\boldsymbol{\psi})f_1(z_{t-1}|\boldsymbol{\psi})$ for this inversion copula. Because the time series has Markov order one, this bivariate copula characterizes the full serial dependence structure. For example, Figure 2(a) plots $c^{(2)}$ for the case when there is no unobserved mean component, $\rho_h = 0.952$ and $\sigma_h = 0.045$— typical values arising when fitting asset return data. The copula density is far from uniform, with high equally-valued tail dependence in all four quadrants ($\lambda_1^{++}(0.1) = 0.1428, \lambda_1^{++}(0.05) = 0.0964$ and $\lambda_1^{++}(0.01) = 0.0454$). This is a high level of first order serial dependence, yet $\tau_1 = r_1 = 0$. This is because τ_1 and r_1 measure 'level' dependence, whereas this copula instead captures bivariate dependence in variation. Most existing parametric copulas are not well-suited to represent such serial dependence arising from conditional heteroscedasticity in time series.

3.1.2 Estimation

We employ uniform priors for ρ_h and ρ_μ, and Inverse Gamma priors for σ_μ^2 and σ_h^2 with shape and scale parameters equal to 1.001. We outline here how to implement the first two steps of the sampling scheme in Section 2.2. In Step 1 we partition the state vector into $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_T)$ and $\boldsymbol{h} = (h_1, \ldots, h_T)$, and use the two separate steps:

Step 1a. Generate from $f(\boldsymbol{\mu}|\boldsymbol{\psi}, \boldsymbol{h}, \boldsymbol{\theta}, \boldsymbol{y}) \propto \prod_{t=1}^T \phi_1\left(z_t; \frac{1}{\kappa}\mu_t; \frac{1}{\kappa^2}\exp(h_t)\right) f(\boldsymbol{\mu}|\boldsymbol{\psi})$

Step 1b. Generate from $f(\boldsymbol{h}|\boldsymbol{\psi}, \boldsymbol{\mu}, \boldsymbol{\theta}, \boldsymbol{y}) \propto \prod_{t=1}^T \phi_1\left(z_t; \frac{1}{\kappa}\mu_t; \frac{1}{\kappa^2}\exp(h_t)\right) f(\boldsymbol{h}|\boldsymbol{\psi})$

The posterior of $\boldsymbol{\mu}$ in Step 1a can be recognized as normal with zero mean and a band 1 precision matrix, so that generation is both straightforward and fast. There are a number of efficient methods to generate $\boldsymbol{h}$ in Step 1b in the literature, and we employ the 'precision sampler' for the latent states outlined in Chan & Hsiao (2014). This is a fast sparse matrix implementation of the auxiliary mixture sampler (Kim, Shepherd & Chib 1998) that is known to mix well outside the copula context.

In Step 2 we generate each parameter one at a time, breaking generation into four separate Metropolis-Hastings steps. In each case normal approximations were used based on three Newton-Rapshon steps. Initial values for the optimization were the parameter means obtained from the density $q(\boldsymbol{\psi}) \propto f(\boldsymbol{\mu}|\boldsymbol{\psi})f(\boldsymbol{h}|\boldsymbol{\psi})\pi_\psi(\boldsymbol{\psi})$. The resulting four normal distributions proved to be accurate approximations to the conditional posteriors, with acceptance rates between 65.9% and 95.4% in our empirical work.

3.1.3 Modeling and Forecasting Inflation

In the copula literature it is popular to use the empirical distribution function as a nonparametric estimator of the margin G (Shih & Louis 1995; Tsukahara 2005). However, this can give inaccurate estimates of the tails of G, as illustrated by Smith & Vahey (2015) for quarterly macroeconomic time series. Therefore, we fit a kernel density estimator (KDE) using the locally adaptive bandwidth method of Shimazaki & Shinomoto (2010). The estimate is a smooth asymmetric and heavy-tailed distribution, and the density is also plotted on Figure 1(b). We use this as our estimate of $g(y)$, from which we compute the copula data $u_t = G(y_t)$. Figure 1(c) plots the resulting time series of the copula data, which can be seen to retain the serial dependence apparent in the original data.

The stochastic volatility model with an unobserved component is one of the most successful time series models of inflation (Clark & Ravazzolo 2015; Chan 2015). We therefore fit the corresponding inversion copula as discussed above. Table 1 reports the posterior estimates of $\boldsymbol{\psi}$, labeled as copula 'InvCop1'. There is strong positive correlation in both the level ($\hat{\rho}_\mu = 0.95$) and the log-volatilities ($\hat{\rho}_h = 0.71$) of z_t. This translates to strong positive dependence in y_t, and is apparent in the smoothed estimates of the mean and standard deviation of inflation, plotted in Figure 3(a,b).

Figure 4(a) plots the marginal copula density $c^{(2)}(u_t, u_{t-1}|\hat{\boldsymbol{\psi}})$ at the parameter point estimates. There are spikes in the density close to (0,0) and (1,1), so that the vertical axis is truncated at 7 to aid interpretation. The logarithm of the density is also plotted in panel (d).

The majority of mass is along the axis between (0,0) to (1,1), which is due to level dependence captured by the unobserved component. However, the conditional heteroscedasticity also affects the form of the copula. Table 2 reports measures of first order serial dependence for this copula. The unobserved mean component results in strong positive overall dependence in the series, with $\hat{r}_1 = 0.84$. There is high (symmetric) tail dependence $\hat{\lambda}_1^{++}(0.05) = \hat{\lambda}_1^{--}(0.05) = 0.56$, which is consistent with the copula density.

For comparison, we also fit the plain SVUC model in Equation (3.1), without any parameter constraints, directly to the inflation data. The marginal density for inflation, computed at the posterior mean of the parameters, is also shown in Figure 1(b). It is necessarily symmetric and inconsistent with that observed empirically. Table 3 provides the marginal likelihoods for both time series models, computed as outlined in the Appendix, showing that the copula model is a substantially better fit.

Last, to show that accounting for asymmetric margins can materially affect prediction, we refit the model to inflation from 1954:Q1 to 2009:Q1, and produce density forecasts over the horizon 2009:Q2 to 2011:Q2. This period coincides with the adoption of quantitative easing (QE1) by the Federal Reserve, and is of particular interest to macroeconomists. Figure 5(a,c,e,g) plots the predictive distributions 1, 2, 4 and 8 quarters ahead from both the copula model and the plain SVUC model. The forecasts from the copula model are highly skewed to the right. This corresponds to a small probability of a severe inflationary response to QE1, something that is also reported by the Survey of Professional Forecasters (February 2009) as discussed by Smith & Vahey (2015). In comparison, the predictive distributions from the plain SVUC model are necessarily symmetric, less sharp and place an unrealistically high probability on deflation. For example, the probability of deflation one quarter ahead is 70.4% for the SVUC model, compared to 4.0% for the copula model.

3.2 Markov Switching Inversion Copula

Another popular class of nonlinear state space models are regime switching models, which allow for structural changes in the dynamics of a series. In these models latent regime indicators usually follow an ergodic Markov chain, in which case the model is called a Markov switching model; see Hamilton (1994; Ch.22) for an introduction.

3.2.1 The Copula

We consider a two regime Markov switching AR(1) model given by

$$
\begin{aligned}
z_t &\sim N(a_{s_t} + \rho_{s_t} z_{t-1}, \sigma^2_{s_t}) \\
\Pr(s_t = j | s_{t-1} = i) &= p_{ij},
\end{aligned} \tag{3.3}
$$

for regimes $s_t \in \{1, 2\}$. This is a nonlinear state space model with state vector $\boldsymbol{x}_t = (z_{t-1}, s_t)$. We assume the Markov chain is ergodic, so that the marginal distribution of s_t is time invariant with $\Pr(s_t = 1) = \pi_1 = (1 - p_{22})/(2 - p_{11} - p_{22})$ and $\Pr(s_t = 2) = \pi_2 = 1 - \pi_2$. We also assume $0 < \rho_j < 1$, so that z_t follows a stationary first order Markov time series, with $E(z_t) = \bar{\mu} = \sum_{i=1,2} \pi_i \frac{a_i}{1-\rho_i}$ and $\text{Var}(z_t) = \kappa^2 = \sum_{i=1,2} \pi_i \frac{\sigma_i^2}{1-\rho_i^2}$. As before, we redefine the observation equation as

$$
z_t \sim N\left(\frac{1}{\kappa}\left(a_{s_t} - \bar{\mu} + \rho_{s_t} z_{t-1}\right), \frac{\sigma^2_{s_t}}{\kappa^2}\right), \tag{3.4}
$$

so that $E(z_t) = 0$ and $\text{Var}(z_t) = 1$. To identify the parameters in the copula, we also fix $\sigma_1^2 = 1$, so that σ_2^2 measures the magnitude of the variance of the second regime, relative to the first. The inversion copula parameters are therefore $\boldsymbol{\psi} = \{a_1, a_2, \rho_1, \rho_2, \sigma_2^2, p_{11}, p_{22}\}$.

The marginal distribution of z_t is a mixture of two Gaussians, with

$$
f_1(z|\boldsymbol{\psi}) = \sum_{i=1,2} \phi_1\left(z; \frac{1}{\kappa}\left(\mu_i - \bar{\mu}\right), \frac{s_i^2}{\kappa^2}\right)\pi_i, \tag{3.5}
$$

where $\mu_i = \frac{a_i}{1-\rho_i}$ and $s_i^2 = \frac{\sigma_i^2}{1-\rho_i^2}$. Following Aitken & Rubin (1985) and others, we identify the labels of the two components by assuming $\pi_1 < \pi_2$. The bivariate margin of any two contiguous elements of $\boldsymbol{z}$ is a mixture of four bivariate Gaussians, with density

$$
f^{(2)}(z_1, z_2|\boldsymbol{\psi}) = \sum_{i=1,2}\sum_{j=1,2} \phi_2\left((z_1, z_2); \frac{1}{\kappa}\boldsymbol{\mu}_{ij}, \frac{1}{\kappa^2}S_{ij}\right) p_{ij}\pi_i, \tag{3.6}
$$

where $\boldsymbol{\mu}_{ij} = (\mu_i - \bar{\mu}, \mu_j - \bar{\mu})$ and

$$S_{ij} = \begin{bmatrix} s_i^2 & \rho_j s_i^2 \\ \rho_j s_i^2 & s_j^2 \end{bmatrix}.$$

The marginal copula $c^{(2)}(u_1, u_2|\boldsymbol{\psi}) = f^{(2)}(z_1, z_2|\boldsymbol{\psi})/f_1(z_1|\boldsymbol{\psi})f_1(z_2|\boldsymbol{\psi})$ is therefore a copula constructed by inversion of a mixture of four Gaussians. This is very different than the more common 'mixture copula', which is built from a finite mixture of other copulas; for example, see Patton (2006). Unlike the other two inversion copulas examined, this copula can exhibit asymmetric first order serial dependence. To illustrate, Figure 2(b) plots the marginal copula $c^{(2)}$ when $p_{11} = p_{22} = 0.996$, $\sigma_2^2 = 4.66$, $\rho_1 = 0.288$, $\rho_2 = 0.047$, $c_1 = 0.031$ and $c_2 = -0.2$. In this case, $r_1 = 0.1248$, $\tau_1 = 0.0876$, tail dependence is asymmetric with $\lambda_1^{++}(\alpha) \neq \lambda_1^{--}(\alpha)$, and there is strong positive off-diagonal tail dependence; for example, $\lambda_1^{--}(0.1) = 0.188$, $\lambda_1^{++}(0.1) = 0.167$, and $\lambda_1^{+-}(0.1) = \lambda_1^{-+}(0.1) = 0.131$.

3.2.2 Estimation and Inflation

Amisano & Fagan (2013) employ the Markov switching AR(1) model at Equation (3.3) to model U.S. inflation, but only allow a_1 and a_2 to vary between regimes. We extend their study by using an inversion copula constructed from the latent Markov switching model, combined with the margin G given by the KDE in Figure 1(b). We also allow for a more general serial dependence structure, because the autoregressive parameter and (relative) variance of z_t also vary over regime. We employ the MCMC algorithm in Section 2.2 to estimate the model, where a forward filtering and backward sampling algorithm (Hamilton 1994, p.694) is used to sample the regime indicators $\boldsymbol{s} = (s_1, ..., s_T)$ in Step 1. Each element of $\boldsymbol{\psi}$ is sampled using truncated normal approximations obtained with five Newton-Raphson steps. The quantile function F_1^{-1} is computed using the spline interpolation method outlined in the Appendix.

Table 1 reports the posterior estimates of the copula parameters, ergodic probabilities π_i, and regime variances s_i^2, labeled as 'InvCop2'. There is a clear separation between the low (regime 1, with $\hat{\pi}_1 = 0.34$) and high (regime 2, with $\hat{\pi}_2 = 0.66$) persistence regimes. Regime 2 also has lower conditional volatility, although the posterior means of s_1^2 and s_2^2 are similar.

Table 2 reports the first order serial dependence metrics. As before there is high overall dependence with $\hat{r}_1 = 0.75$, although dependence is asymmetric with $\hat{\lambda}_1^{++}(0.05) = 0.57 > \hat{\lambda}_1^{--}(0.05) = 0.38$. This asymmetry is visable in $c^{(2)}(u_t, u_{t-1}|\hat{\boldsymbol{\psi}})$ plotted in Figure 4(b), and $\log(c^{(2)})$ plotted in Figure 4(e).

Figure 3(c,d) plots the smoothed means and standard deviations of inflation. The means $E_{F_P}(y_t)$ are less smooth in panel (c), than those found in panel (a). This is because InvCop1 is an inversion copula formed from a latent process with additive unobserved component for the mean of z_t. The smoothed variances in panel (d) exhibit a high degree of heteroscedasticity in inflation. However, the time series of smoothed variances is less persistent than that found in panel (b) for InvCop1. This is because the latter inversion copula is formed from a latent process with a smooth autoregression in the log-volatility. Table 3 reports the marginal likelihood for this copula model, along with that of the (non-copula) Markov switching model at Equation (3.3) fit with neither the constraint on σ_1^2 nor the normalization at Equation (3.4). Between these two models, the copula model provides a better fit, although the inversion copula InvCop2 does not capture the serial dependence structure as well as InvCop1.

Last, Figure 5(b,d,f,h) plots the predictive densities from this copula model, labeled 'InvCop2'. As before, the predictive densities are right skewed. However, the densities are bimodal for 2 and 4 quarters ahead, which can also be a feature of predictive densities produced from (plain) Markov switching AR(1) models.

3.3 Linear State Space Inversion Copula

The stationary discrete time linear Gaussian state space model encompasses many popular time series models; see Ljung (1999, Sec. 4.3) and Durbin & Koopman (2012, Part 1) for overviews. Despite its popularity, this model has Gaussian margins. As an alternative, we construct an inversion copula from a linear Gaussian state space model, and show that the result is a Gaussian copula with a parameter matrix that is a function of the state space parameters.

3.3.1 The Copula and Estimation

We consider the stationary linear state space model given by

$$\begin{aligned} z_t | \boldsymbol{x}_t, \boldsymbol{\psi} &\sim N(\boldsymbol{b}\boldsymbol{x}_t', h) \\ \boldsymbol{x}_t | \boldsymbol{x}_{t-1}, \boldsymbol{\psi} &\sim N(\boldsymbol{x}_{t-1}R', FQF'), \end{aligned} \tag{3.7}$$

where $\boldsymbol{x}_t$ is a $(1 \times k)$ vector of latent states, $\boldsymbol{b}$ is a $(1 \times k)$ vector, h is the disturbance variance, R is a $(k \times k)$ matrix of autoregressive coefficients, with absolute values of all eigenvalues less than one. The matrices F and Q are of sizes $(k \times q)$ and $(q \times q)$, respectively. We normalize the measurement equation as

$$z_t | \boldsymbol{x}_t, \boldsymbol{\psi} \sim N\left(\frac{1}{\kappa}\boldsymbol{b}\boldsymbol{x}_t', \frac{h}{\kappa^2}\right), \tag{3.8}$$

where $\kappa^2 = \boldsymbol{b}\Sigma_x\boldsymbol{b}' + h$ is the marginal variance of z_t. The matrix $\Sigma_x = Var(\boldsymbol{x}_t)$ is obtained as the solution of the equation $vec(\Sigma_x) = (I_{k^2} + R \otimes R)^{-1} vec(FQF')$. Note that even though the latent state vector $\boldsymbol{x}_t$ has Markov order one, z_t has an arbitrary Markov order; see Durbin & Koopman (2012) for properties of this model.

With the normalization at Equation (3.8), the margins of z_t are standard normal, so that $f_1(z; \boldsymbol{\psi}) = \phi(z; 0, 1)$ and $F_1(z; \boldsymbol{\psi}) = \Phi(z; 0, 1)$. This greatly simplifies construction and estimation of the copula, compared to those constructed from inversion of nonlinear state space models, where evaluation of f_1 and F_1 is dependent on the copula parameters $\boldsymbol{\psi}$. Moreover, $\boldsymbol{z} \sim N(0, \Omega_\psi)$, so that this inversion copula is a Gaussian copula. The correlation matrix is given by

$$\Omega_\psi = \begin{bmatrix} 1 & a_1/\kappa^2 & \dots & a_{T-1}/\kappa^2 \\ a_1/\kappa^2 & 1 & \dots & a_{T-2}/\kappa^2 \\ \vdots & \vdots & \ddots & \vdots \\ a_{T-1}/\kappa^2 & a_{T-2}/\kappa^2 & \dots & 1 \end{bmatrix},$$

where $a_l = \boldsymbol{b}\Gamma(l)\boldsymbol{b}'$, and $\Gamma(l) = \text{Cov}(\boldsymbol{x}_t, \boldsymbol{x}_{t-l})$ denotes the l^{th} autocovariance matrix of the state vectors, which can be computed using the multivariate Yule-Walker equations.

Not all parameters in the state space model are identified in the inversion copula. If $P(l) = D^{-1}\Gamma(l)D^{-1}$ is the lag l autocorrelation matrix of the state vector, then

$$\frac{a_l}{\kappa^2} = \frac{\boldsymbol{b}DP(l)D\boldsymbol{b}'}{\boldsymbol{b}DP(0)D\boldsymbol{b}' + h} = \frac{\boldsymbol{b}\bar{D}P(l)\bar{D}\boldsymbol{b}'}{\boldsymbol{b}\bar{D}P(0)\bar{D}\boldsymbol{b}' + 1},$$

where $\bar{D} = h^{-1/2}D$. Therefore, h is unidentified in Ω_ψ, so that we set $h = 1$ in Equation (3.7). The copula parameters are therefore $\boldsymbol{\psi} \subseteq \{\boldsymbol{b}, R, F, Q\}$, depending on the specific state space model adopted.

Estimation of $\boldsymbol{\psi}$ can employ the likelihood in Equation (2.6), where the copula density is

$$c_{Ga}(\boldsymbol{u};\boldsymbol{\psi}) = |\Omega_\psi|^{-1/2} \exp\left(-\frac{1}{2}\boldsymbol{z}(\Omega_\psi^{-1} - I)\boldsymbol{z}'\right).$$

However, we generate the states explicitly from a normal distribution using the Kalman filter, and estimate the model using the MCMC sampling scheme. Because the Gaussian copula is closed under marginalization, $c^{(r)}(\boldsymbol{u}_{t-r+1:t})$ is also a Gaussian copula density with a submatrix of Ω_ψ as parameters.

3.3.2 Example Copula and Inflation

We construct an inversion copula from a Gaussian unobserved component model that follows a stationary AR(4), where $z_t|\boldsymbol{x}_t, \boldsymbol{\psi} \sim N(\mu_t, 1)$ and $\mu_t = \sum_{j=1}^4 \rho_j \mu_{t-j} + \sigma_\mu^2$. This model can be written in state space form at Equation (3.7). We use this inversion copula to model the inflation data, along with a parametric skew t distribution (Azzalini & Capitanio 2003) for G. We use the MCMC scheme in Section 2.2 to compute the joint posterior of the marginal and copula parameters. As in Section 3.1.2, we use the precision sampler to sample the states at Step 1. The autoregression is parameterized by the partial autocorrelations $(\pi_1, \pi_2, \pi_3, \pi_4)$ using the Durbin-Levinson algorithm. In Step 2 of the scheme, these partials are sampled jointly from a normal approximation based on five Newton-Raphson steps, truncated to the unit cube, while σ_μ^2 is also generated using a truncated normal approximation. We use the location (ξ), scale (ω), skew (γ_1) and kurtosis (γ_2) coefficients as the skew t parameters, and generate $\boldsymbol{\theta} = (\xi, \omega, \gamma_1, \gamma_2)$ using a multivariate adaptive random walk proposal in Step 3.

Table 1 reports the posterior estimates of $\boldsymbol{\psi}$ and $\boldsymbol{\theta}$ (labeled 'InvCop3'), while Table 2 reports the first order serial dependence metrics. The estimated skew t margin has high positive skew $\hat{\gamma}_1 = 1.59$, heavy tails $\hat{\gamma}_2 = 8.72$, and is similar to the KDE. The estimated partial correlations indicate that the unobserved component is Markov order two, and translates into positive first order serial dependence in inflation with $\hat{r}_1 = 0.84$. Figure 3(e,f) plots the smoothed mean and standard deviation of inflation from the fitted model. Interestingly, despite the state space model for z_t being homoscedastic, some heteroscedasticity can be seen in Figure 3(f). Figure 4(c,f) plots the marginal copula density, $c^{(2)}(u_t, u_{t-1}|\hat{\boldsymbol{\psi}})$ and its logarithm at the point estimates of the parameters. This is a bivariate Gaussian copula, and is therefore symmetric along the axes (0,1) to (1,0). Tail dependence is positive with $\hat{\lambda}_1^{++}(0.01) = 0.46$ and $\hat{\lambda}_1^{++}(0.05) = 0.56$; although $\lim_{\alpha \to 0} \lambda_1^{++}(\alpha) = 0$ because $c^{(2)}$ is a Gaussian copula. Table 3 reports the marginal likelihood value for the copula model, along with that for a (non-copula) Gaussian AR(4) unobserved component model. Again, the copula model dominates the state space model, although the serial dependence structure is better captured by InvCop1.

Last, we compute predictive densities for the horizon 2009:Q2 to 2011:Q2 as discussed in Section 3.1.3. Figure 5(b,d,f,h) plots these densities, labeled as 'InvCop3'. At 2, 4 and 8 quarters ahead, the predictive densities are positively skewed. However, all the predictive densities differ substantially from those produced using the copulas InvCop1 and InvCop2, and are much more diffuse.

4 Extension to Multivariate Time Series

The approach is readily extended to the multivariate case as follows. Let $\boldsymbol{y}_t = (y_{1,t}, \ldots, y_{m,t})$ be a vector of m time series observed at time t, for $t = 1, \ldots, T$. Then, a copula model employs a $N = Tm$-dimensional copula to capture both cross-sectional and serial dependence in $\boldsymbol{y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_T)$. The copula can be constructed from inverting a latent multivariate nonlinear state space model, with observation equation $\boldsymbol{z}_t = (z_{1,t}, \ldots, z_{m,t}) \sim H_t(\boldsymbol{z}_t|\boldsymbol{x}_t; \boldsymbol{\psi})$, and transitional distribution at Equation (2.3). The marginal distribution and density functions of $z_{j,t}$ are

$$
\begin{aligned}
F_{j,t}(z_{j,t}|\boldsymbol{\psi}) &= \int H_{j,t}(z_{j,t}|\boldsymbol{x}_t;\boldsymbol{\psi})f(\boldsymbol{x}_t|\boldsymbol{\psi})\mathrm{d}\boldsymbol{x}_t \\
f_{j,t}(z_{j,t}|\boldsymbol{\psi}) &= \int h_{j,t}(z_{j,t}|\boldsymbol{x}_t;\boldsymbol{\psi})f(\boldsymbol{x}_t|\boldsymbol{\psi})\mathrm{d}\boldsymbol{x}_t\,.
\end{aligned}
$$

Here, $H_{j,t}$ and $h_{j,t}$ are the univariate distribution and density functions of $z_{j,t}$, respectively. In practice, these are usually known in closed form from H_t. If $F_z(\cdot|\boldsymbol{\psi})$ and $f_z(\cdot|\boldsymbol{\psi})$ are the joint distribution and density functions of $\boldsymbol{z} = (z_{1,1}, \ldots, z_{m,T})$, respectively, then the inversion copula function and density are

$$
\begin{aligned}
C(\boldsymbol{u};\boldsymbol{\psi}) &= F_z\left(F_{1,1}^{-1}(u_{1,1}|\boldsymbol{\psi}), \ldots, F_{m,T}^{-1}(u_{m,T})|\boldsymbol{\psi}\right) \\
c(\boldsymbol{u};\boldsymbol{\psi}) &= f_z(\boldsymbol{z})/\prod_{t=1}^{T}\prod_{j=1}^{m} f_{j,t}(z_{j,t};\boldsymbol{\psi})\,,
\end{aligned}
$$

with $u_{j,t} = F_{j,t}(z_{j,t};\boldsymbol{\psi})$, and $\boldsymbol{u} = (u_{1,1}, \ldots, u_{m,T})$.

Univariate marginal distributions $G_j(y_{j,t};\boldsymbol{\theta})$ for each series $j = 1, \ldots, m$ are required to complete the copula model, which has likelihood

$$
f(\boldsymbol{y}|\boldsymbol{\psi},\boldsymbol{\theta}) = f_z(\boldsymbol{z}|\boldsymbol{\psi})\prod_{t=1}^{T}\prod_{j=1}^{m}\frac{g_j(y_{j,t};\boldsymbol{\theta})}{f_{j,t}(z_{j,t}|\boldsymbol{\theta})}\,,
$$

where $g_j(y_{j,t};\boldsymbol{\theta}) = \frac{d}{dy_{j,t}}G_j(y_{j,t};\boldsymbol{\theta})$. Posterior estimation can be undertaken using MCMC as outlined in Section 2.2, but with conditional likelihood

$$
f(\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\psi},\boldsymbol{\theta}) = \prod_{t=1}^{T} h_t(\boldsymbol{z}_t|\boldsymbol{x}_t;\boldsymbol{\psi})\prod_{j=1}^{m}\frac{g_j(y_{j,t};\boldsymbol{\theta})}{f_{j,t}(z_{j,t}|\boldsymbol{\psi})}\,,
$$

where $z_{j,t} = F_{j,t}^{-1}(G_j(y_{j,t};\boldsymbol{\theta})|\boldsymbol{\psi})$. Therefore, at Step 1 of the sampling scheme the states can be generated using existing methods for state space models, such as the Kalman filter for the linear Gaussian case, or the particle filter for the nonlinear case.

As before, it is straightforward to show that multivariate stationarity and Markov properties of $\boldsymbol{z}_t$ are inherited by the time series. Moreover, in this stationary case, $F_{j,1}^{-1}$ and

$\log(g_j)$ can be approximated using splines for each margin j as in the Appendix. Estimates of serial dependence, analogous to those in Section 2.4, can be computed using Monte Carlo methods based on iterates simulated from the copula. Last, if the series is Markov order p, then the serial dependence is fully characterized by the marginal copula $c^{(p+1)}$ where

$$c^{(r)}(\boldsymbol{u}_{t-r+1:t}) = \frac{f^{(r)}(\boldsymbol{z}_{t-r+1:t})}{\prod_{s=t-r+1}^{t} \prod_{j=1}^{m} f_{j,1}(z_{j,s})} \,,$$

$\boldsymbol{u}_t = (u_{1,t}, \ldots, u_{m,t})$, and $f^{(r)}$ is the obvious generalization of Equation (2.8).

5 Discussion

This paper proposes a new class of copulas for capturing serial dependence. They are constructed from inversion of a general nonlinear state space model, so that the potential range of dependence structures that they can produce is incredibly broad. A major insight is that such copulas can be very different than those that are widely used for capturing cross-sectional dependence. The latter include elliptical and vine copulas, which have also been used previously to capture serial dependence; see Joe (1997), Beare (2010) and Smith et al. (2010) for examples. Yet, the three inversion copulas studied in detail highlight the much wider serial dependence structures that can be captured by our proposed inversion copulas. For example, existing closed form parametric copulas cannot easily replicate the marginal copula densities found in either Figures 2 or 4(a,b,d,e).

As with the likelihood of the underlying state space model, in general the density of the corresponding inversion copula cannot be expressed in closed form. However, an important insight is that existing simulation methods for evaluating the likelihood of the latent state space model can also be employed in the copula context. While we generate the states $\boldsymbol{x}$ in an MCMC scheme, particle and importance sampling methods (e.g. DeJong et al. 2013; Koopman & Lucas 2015; Scharth & Kohn 2013) may also be used to compute the numerator of the copula density f_z. Either way, a major computational challenge when estimating the copula model is the repeated evaluation of the marginal quantile functions $z_t = F_t^{-1}(u_t; \boldsymbol{\psi})$ and densities $f_t(z_t; \boldsymbol{\psi})$ at the T observations. When the latent state space model is stationary, so that $F_t = F_1$ and $f_t = f_1$, we show how this can be achieved using spline approximations. These approximations prove highly accurate in our examples, are fast to derive, and can be employed with potentially large values of T in practice.

Recently, copulas with time-varying parameters have proven popular for the analysis of multivariate time series data; for example, see Almeida & Czado (2012), Hafner & Manner (2012), De Lira Salvatierra & Patton (2015) and Creal & Tsay (2015). However, these authors use copulas to account for the (conditional) cross-sectional dependence as in Patton (2006). This is completely different to our objective of constructing a T-dimensional copula for serial dependence. Moreover, when our approach is applied to a latent multivariate state space model, a Tm-dimensional inversion copula of the joint dependence structure is obtained– again, very different from the dynamic conditional copula models proposed previously. Semi- and nonparametric copula functions (Kauermann, Schellhase & Ruppert 2013; Tran et al. 2013) can also be used to model serial dependence. However, such an approach is better suited to longitudinal data, where there are repeated observations of the time series.

An interesting result is that the inversion copula of a Gaussian linear state space model is a Gaussian copula. Therefore, in this special case, the likelihood is available in closed form. Nevertheless, estimation using simulation methods can still prove efficient, just as it is for the latent state space model itself. While all three of our example inversion copulas are Markov and stationary, copulas can also be derived from non-stationary latent state space models. The resulting time series copula model is also non-stationary, but with a time invariant margin. Such copula models are an interesting topic for further study, although a new approach to computing F_t^{-1} and f_t efficiently is needed. Another interesting extension is to employ the proposed time series inversion copulas to capture serial dependence in discrete data. Here, the copula remains unchanged, but G would be a discrete distribution function. The model can be estimated using Bayesian data augmentation, as discussed in Pitt, Chan & Kohn (2006). This would require $\boldsymbol{z}$ to be generated as an additional step in the sampling scheme in Section 2.2.

To illustrate our methodology, we use it to model and forecast U.S. inflation data. This is an important application of nonlinear time series models in macroeconomics (Faust &

Wright 2013). We employ nonlinear state space models that are currently considered among the most useful or best performing for this series (Amisano & Fagan 2013; Clark & Ravazzolo 2015; Chan 2015). Our inversion copulas extract the serial dependence features of these time series models. When combined with highly asymmetric and heavy-tailed nonparametric or flexible marginal distributions, the fit of the resulting copula time series models can substantially improve on that of the state space models themselves. This is because these state space models have rigid margins, which are very far from that observed for inflation empirically.

Acknowledgments

The work of Michael Smith was supported by Australian Research Council Grant FT110100729.

Appendix

Part A

This part of the appendix outlines how to efficiently evaluate the quantile function F_1^{-1} and the logarithm of the marginal density $\log(f_1(z))$ for a stationary nonlinear state space model. For both, we use spline interpolations based on their values at N absciassae, where we set $N = 100$ in practise. The advantage of such spline-based approximations is that they are highly accurate (between 5 and 9 decimal places in our empirical work), yet are fast to compute at the T observations once the interpolation is complete– even for large values of T.

We use a uniform grid for the N quantile function values $\{q_1, \ldots, q_N\}$, which have correpsonding probability values $\{p_1, \ldots, p_N\}$. We set $p_1 = 0.0001$ and $p_N = 0.9999$, so that the function is approximated far into the tails of the distribution. The following steps obtain the points at which the interpolations are made.

1. Set $p_1 = 0.0001$ and $p_N = 0.9999$, and evaluate both $q_1 = F_1^{-1}(p_1)$ and $q_N = F_1^{-1}(p_N)$ using a root finding algorithm.
2. Set step size to $\delta = (q_N - q_1)/(N-1)$, and a construct uniform grid as $q_i = q_1 + (i-1)\delta$, for $i = 2, \ldots, N$.
3. For $i = 1, \ldots, N$ (in parallel):

   3a. Compute $p_i = F_1(q_i)$

   3b. Compute $b_i = \log(f_1(q_i))$

We then interpolate the points $\{(p_i, q_i); i = 1, \ldots, N\}$ and $\{(q_i, b_i); i = 1, \ldots, N\}$ using splines. We employ natural cubic smoothing splines using the (fast and efficient) spline toolbox in MATLAB, although other fast interpolating methods could also be employed. Notice that numerical inversion of F_1 is undertaken above only twice in Step 1. Moreover, F_1 and f_1 are evaluated only N times in step 3, something that can also be undertaken parallel.

If the marginal distribution F_1 is symmetric, as in Section 3.1, then $F_1^{-1}(1-u) = -F_1^{-1}(u)$ for $0 \leq u \leq 1/2$, $f_1(z) = f_1(-z)$ and $F_1(-z) = 1 - F_1(z)$ for $z \geq 0$. These identities can be exploited to reduce the number of computations at Steps 1 and 3 by one half, further speeding the algorithm.

To illustrate the effectiveness of the method, we consider the approximations to F_1^{-1} and $\log(f_1)$ when $\boldsymbol{\psi}$ equals the posterior mean in the inflation study in Section 3.1.3. Plots of the approximations are visually indistinguishable from the true functions, which can be evaluated (slowly) using numerical methods. The integrated absolute error between the approximate and true functions are 7.9×10^{-6} and 3.6×10^{-6} for the quantile and log-density, respectively. Computation of both approximations, and their evaluation at the $T = 240$ datapoints, takes only 0.29s using MATLAB on a standard four core desktop.

Part B

This part of the appendix shows how to evaluate the bivariate marginal copula density

$$c^{(2)}(u_1, u_2; \boldsymbol{\psi}) = \frac{f^{(2)}(z_1, z_2|\boldsymbol{\psi})}{f_1(z_2|\boldsymbol{\psi}) f_1(z_1|\boldsymbol{\psi})},$$

for the latent stochastic volatility model with an unobserved (level) component given in Section 3.1. The univariate marginal density $f_1(z|\boldsymbol{\psi})$ can be computed readily as in Part A above. The bivariate density

$$f^{(2)}(z_1, z_2|\boldsymbol{\psi}) = \int\int\int\int \phi_2\left(\boldsymbol{z}; \frac{1}{\kappa}\boldsymbol{\mu}, \frac{1}{\kappa^2}H\right) \phi_2(\boldsymbol{\mu}; \mathbf{0}, S_\mu) \mathrm{d}\boldsymbol{\mu} \phi_2(\boldsymbol{h}; \mathbf{0}, S_h) \mathrm{d}\boldsymbol{h},$$

where $\boldsymbol{z} = (z_1, z_2)'$, $\boldsymbol{\mu} = (\mu_1, \mu_2)'$, $\boldsymbol{h} = (h_1, h_2)'$,

$$H = \begin{bmatrix} \exp(h_1) & 0 \\ 0 & \exp(h_2) \end{bmatrix}, \; S_h = s_h^2 \begin{bmatrix} 1 & \rho_h \\ \rho_h & 1 \end{bmatrix}, \; S_\mu = s_\mu^2 \begin{bmatrix} 1 & \rho_\mu \\ \rho_\mu & 1 \end{bmatrix},$$

and $\phi_2(\boldsymbol{x}; 0, \Omega)$ is a bivariate normal density with mean zero and variance Ω evaluated at point $\boldsymbol{x}$. The inner two integrals in $\boldsymbol{\mu}$ of this 4-dimensional integral can be computed analytically by recognising a bivariate normal. Then, by recognising a second bivariate normal in $\boldsymbol{z}$, the density can be written as:

$$f(z_1, z_2|\boldsymbol{\psi}) = \int\int \phi_2(\boldsymbol{z}; \boldsymbol{0}, W(\boldsymbol{h}))\phi_2(\boldsymbol{h}; \boldsymbol{0}, S_h)\mathrm{d}\boldsymbol{h}\,,$$

where $W(h)^{-1} = \kappa^2(S_\mu + H)^{-1}$. This bivariate integral can be computed numerically.

Part C

We briefly outline here how we compute the marginal likelihood for the copula model. The logarithm of the marginal likelihood is

$$\log m(\boldsymbol{y}) = \log f(\boldsymbol{y}|\boldsymbol{\psi}, \boldsymbol{\theta}) + \log \pi_\psi(\boldsymbol{\psi}) + \log \pi_\theta(\boldsymbol{\theta}) - \log f(\boldsymbol{\psi}, \boldsymbol{\theta}|\boldsymbol{y})\,,$$

which we compute at the posterior mean $\boldsymbol{\psi}^\star, \boldsymbol{\theta}^\star$. The likelihood is

$$f(\boldsymbol{y}|\boldsymbol{\psi}, \boldsymbol{\theta}) = f_z(\boldsymbol{z}|\boldsymbol{\psi}) \prod_{t=1}^{T} \frac{g(y_t; \boldsymbol{\theta})}{f_t(z_t|\boldsymbol{\psi})}\,.$$

In the case of our three stationary copula models, we compute f_1 using the spline approximation, and g using either the density of a skew t, or the KDE. For the Gaussian copula (InvCop3), $f_z(\boldsymbol{z}|\boldsymbol{\psi})$ is simply a Gaussian density easily computed using the Kalman filter. For InvCop1 we employ the bootstrap particle filter to integrate out $\boldsymbol{h}$, with the Kalman filter embedded to integrate out the unobserved mean component $\boldsymbol{\mu}$ analytically. For InvCop2 the Hamilton filter for discrete Markov states is applied. For other nonlinear state space models, computing f_z typically involves using either importance sampling or a particle filter. Last, we follow Chib & Jeliazkov (2001) and use iterates from reduced runs of the sampling schemes to compute the posterior density $f(\boldsymbol{\psi}^\star, \boldsymbol{\theta}^\star|\boldsymbol{y})$ component of the marginal likelihood.

References

Aitken, M. & D. Rubin, (1985). 'Estimate and hypothesis testing in finite mixture models', *Journal of the Royal Statistical Society*, Series B, 47, 67–75.

Almeida, C. and C. Czado, (2012). 'Bayesian inference for stochastic time-varying copula models', *Computational Statistics & Data Analysis*, 56, 1511–1527.

Amisano, G. & G. Fagan, (2013). 'Money growth and inflation: A regime switching approach', *Journal of International Money and Finance*, 33, 118–145.

Azzalini, A. & A. Capitanio, (2003). 'Distribution generated by perturbation of symmetry with emphasis on a multivariate skew t distribution', *Journal of the Royal Statistical Society*, Series B, 65, 367–389.

Beare, B. K., (2010). 'Copulas and temporal dependence', *Econometrica*, 78, 395–410.

Beare, B. K., & J. Seo (2015). 'Vine copula specifications for stationary multivariate Markov chains', *Journal of Time Series Analysis*, 36, 228–246.

Biller, B. (2009). 'Copula-Based Multivariate Input Models for Stochastic Simulation', *Operations Research*, 57, 4, 878–892.

Biller, B. & B. Nelson, (2003). 'Modeling and Generating Multivariate Time-Series Input Processes Using a Vector Autoregressive Technique', *ACM Transactions on modeling and Computer Simulation*, 13, 3, 211–237.

Brockwell, P.J. & R.A. Davis, (1991), *Time Series: Theory and Methods*, 2nd Ed., NY: Springer.

Chan, J.C.C (2015). 'The Stochastic Volatility in Mean Model with Time-Varying Parameters: An Application to Inflation Modeling', *Journal of Business and Economic Statistics*, to appear.

Chan, J.C.C. & C.Y.L. Hsiao, (2014). 'Estimation of Stochastic Volatility Models with Heavy Tails and Serial Dependence', in *Bayesian Inference in the Social Sciences*, Eds. I. Jeliazkov and X.-S. Yang, Wiley: NJ.

Chib, S. & I. Jeliazkov (2001). 'Marginal Likelihood From the Metropolis-Hastings Output', *Journal of the American Statistical Association*, 96, 270–281.

Clark, T. E. & F. Ravazzolo, (2015). 'Macroeconomic Forecasting Performance under Alternative Specifications of Time-varying Volatility', *Journal of Applied Econometrics*, 30, 551–575.

Creal, D.D. & R.S. Tsay, (2015). 'High dimensional dynamic stochastic copula models', *Journal of Econometrics*, in press.

DeJong, D.N., R. Liesenfeld, G.V. Moura, J.-F. Richard, H. Dharmarajan, (2013). 'Efficient Likelihood Evaluation of State-Space Representations', *Review of Economic Studies*, 30, 538–567.

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

Demarta, S. & A. McNeil, (2005). 'The t-copula and related copulas', *International Statistical Review*, 73, 111–129.

Domma, F., S. Giordano, & P. Perri, (2009). 'Statistical Modeling of Temporal Dependence in Financial Data via a Copula Function', *Communications in Statistics*, 38, 703–728.

Durbin, J. & S.J. Koopman, (2002). 'A simple and efficient simulation smoother for state space time series models', *Biometrika*, 89, 603–616.

Durbin, J. & S.J. Koopman, (2012). *Time Series Analysis by State Space Methods*, 2nd Ed., OUP.

Embrechts, P., A. McNeil & D. Straumann, (2001). 'Correlation and dependency in risk management: properties and pitfalls', in M. Dempster & H. Moffatt, (Eds.) *Risk Management: Value at Risk and Beyond* Cambridge University Press, 176–223.

Fang, H. & K. Fang, (2002). 'The meta-elliptical distributions with given marginals', *Journal of Multivariate Analysis*, 82, 1–16.

Faust, J. & J. Wright, (2013). 'Forecasting Inflation', in G. Elliot & A. Timmermann (eds.), *Handbooks in Economics*, Vol. 2A., 3–56.

Frees, E.W. & P. Wang, (2006). 'Copula credibility for aggregate loss models', *Insurance: Mathematics and Economics*, 38, 360–373.

Hafner, C.M. & H. Manner, (2012). 'Dynamic stochastic copula models: estimation, inference and applications', *Journal of Applied Econometrics*, 27, 269–295.

Hamilton, J.D., (1994). *Time Series Analysis*, Princeton University Press: NJ.

Godsill, S., A. Doucet & M. West (2004). 'Monte Carlo Smoothing for Nonlinear Time Series', *Journal of the American Statsitical Association*, 99, 156–168.

Joe, H., (1997). *Multivariate Models and Dependence Concepts*, Chapman and Hall.

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

Kauermann, G., C. Schellhase, D. Ruppert (2013). 'Flexible Copula Density Estimation with Penalized Hierarchical B-splines', *Scandinavian Journal of Statistics: Theory and Applications*, 40, 685–705.

Ljung, L. (1999). *System Identification: Theory for the User*, 2nd Ed., Prentice-Hall.

Kim, S., N. Shephard & S. Chib, (1998). 'Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models', *Review of Economic Studies*, 65, 361–393.

Lambert, P. & F. Vandenhende, (2002). 'A copula-based model for multivariate non-normal longitudinal data: analysis of a dose titration safety study on a new antidepressant', *Statistics in Medicine*, 21, 3197–3217.

Nelsen, R., (2006), *An Introduction to Copulas*. 2nd ed., New York, Springer.

Patton, A.J., (2006). 'Modelling Asymmetric Exchange Rate Dependence', *International Economic Review*, 47, 527–556.

Patton, A., (2012). 'A Review of Copula Models for Economic Time Series', *Journal of Multivariate Analysis*, 110, 4–18.

Pitt, M., D. Chan & R. Kohn (2006). 'Efficient Bayesian Inference for Gaussian Copula Regression Models', *Biometrika*, 93, 537–554.

Scharth, M. & Kohn, R., (2013). 'Particle Efficient Importance Sampling', *Journal of Econometrics*, to appear.

Shapiro, S. & M. Wilk, (1965). 'An Analysis of Variate Test for Normality (Complete Samples)', *Biometrika*, 52, 3-4, 591–611.

Shephard, N. & M. Pitt, (1997). 'Likelihood analysis of non-Gaussian measurement time series', *Biometrika*, 84:3, 653–667.

Shih, J.H. & T.A. Louis, (1995). 'Inferences on the association parameter in copula models for bivariate survival data', *Biometrics*, 51, 4, 1384–99.

Shimazaki, H. & S. Shinomoto, (2010). 'Kernel bandwidth optimization in spike rate estimation', *Journal of Computational Neuroscience*, 29, 171–182.

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

Smith, M.S., Q. Gan & R. Kohn (2012). 'Modelling Dependence using Skew t Copulas: Bayesian Inference and Applications', *Journal of Applied Econometrics*, 27, 500–522.

Smith, M., A. Min, C. Almeida & C. Czado, (2010). 'Modeling Longitudinal Data using a Pair-Copula Decomposition of Serial Dependence, *Journal of the American Statistical Association*, 105, 492, 1467-1479.

Smith, M.S. & S.P. Vahey, (2015). 'Asymmetric density forecasting of U.S. macroeconomic variables using a Gaussian copula model of cross-sectional and serial dependence', *Journal of Business and Economic Statistics*, in press.

Song, P. (2000). 'Multivariate Dispersion Models Generated from Gaussian Copula', *Scandinavian Journal of Statistics*, 27, 305–320.

Stroud, J.R., P. Müller & N.G. Polson, (2003). 'Nonlinear State-Space Models with State-Dependence Variances', *Journal of the American Statistical Association*, 98:462, 377–386.

Survey of Professional Forecasters, (2009). 'First Quarter 2009', Federal Reserve Bank of Philadephia: February 2009.

Tran, M., P. Giordani, X. Mun, R. Kohn & M. Pitt, (2014). 'Copula-Type Estimators for Flexible Multivariate Density Modeling Using Mixtures', *Journal of Computational and Graphical Statistics*, 23, 1163–78.

Tsukahara, H. (2005). 'Semiparametric estimation in copula models', *The Canadian Journal of Statistics*, 33, 357–375.

Copula	Parameters					
InvCop1	ρ_μ	σ^2_μ	ρ_h	σ^2_h		
	0.95	2.83	0.71	1.14		
	(0.93,0.97)	(0.77, 6.49)	(0.43,0.92)	(0.37,2.32)		
InvCop2	c_1	ρ_1	σ^2_1	s^2_1	p_{11}	π_1
	-0.27	0.43	1	1.28	0.85	0.34
	(-0.57,0.05)	(0.18,0.63)	–	(1.03,1.65)	(0.58,0.99)	(0.12,0.50)
	c_2	ρ_2	σ^2_2	s^2_2	p_{22}	π_2
	0.04	0.88	0.23	1.08	0.94	0.66
	(-0.01,0.18)	(0.78,0.93)	(0.15,0.35)	(0.82,1.35)	(0.84,0.99)	(0.50,0.88)
InvCop3	π_1	π_2	π_3	π_4	σ^2_μ	
	0.91	0.40	-0.09	0.25	2.54	
	(0.86,0.93)	(0.23,0.56)	(-0.20,0.19)	(-0.07,0.28)	(1.60,3.13)	
	ξ	ω	γ_1	γ_2		
	0.21	0.64	1.59	8.72		
	(0.10,0.27)	(0.47,0.92)	(1.49,1.61)	(7.53,8.77)		

Table 1: Posterior parameter estimates for the three inversion copulas fit to the U.S. inflation data. The three latent state space models are (i) InvCop1: a first order stochastic volatility model with an unobserved AR(1) component; (ii) InvCop2: a Markov-switching AR(1) model; and (iii) InvCop3: a Gaussian state space model with an unobserved AR(4) component. For InvCop3, the parameters of the jointly estimated skew t margin are also reported. For each parameter the posterior mean is reported, along with 90% posterior probability intervals below.

	Dependence Metric				
	r_1	$\lambda_1^{--}(0.01)$	$\lambda_1^{--}(0.05)$	$\lambda_1^{++}(0.01)$	$\lambda_1^{++}(0.05)$
Copula					
InvCop1	0.84	0.40	0.56	Sym	Sym
	(0.77,0.89)	(0.24,0.52)	(0.47,0.64)		
InvCop2	0.75	0.24	0.38	0.48	0.57
	(0.64,0.85)	(0.08,0.47)	(0.21,0.57)	(0.31,0.58)	(0.45,0.65)
InvCop3	0.84	0.46	0.56	Sym	Sym
	(0.78,0.89)	(0.37,0.53)	(0.51,0.63)		

Table 2: Posterior means of first order serial dependence measures for each of the three inversion copulas fit to the U.S. inflation data. The copulas InvCop1 and InvCop3 have symmetric tail dependence,while InvCop3 has asymmetric tail dependence. Also reported in parentheses are the 90% posterior probability intervals for each dependence metric.

Model	Marginal Likelihood (Logarithm)
Copula Time Series Models	
KDE Margin & InvCop1	-4.08
KDE Margin & InvCop2	-25.25
Skew t Margin & InvCop3	-20.70
State Space Models	
SV(1) with Unobserved AR(1) Component	-10.46
Markov Switching AR(1)	-48.83
Gaussian Unobserved AR(4) Component	-88.69

Table 3: Marginal likelihood values on the logarithmic scale for six models fit to the U.S. inflation data. A higher value indicates a better fit. The first three models employ inversion copulas constructed from latent state space models, along with an asymmetric margin. The bottom three are the same state space models fit directly to the data.

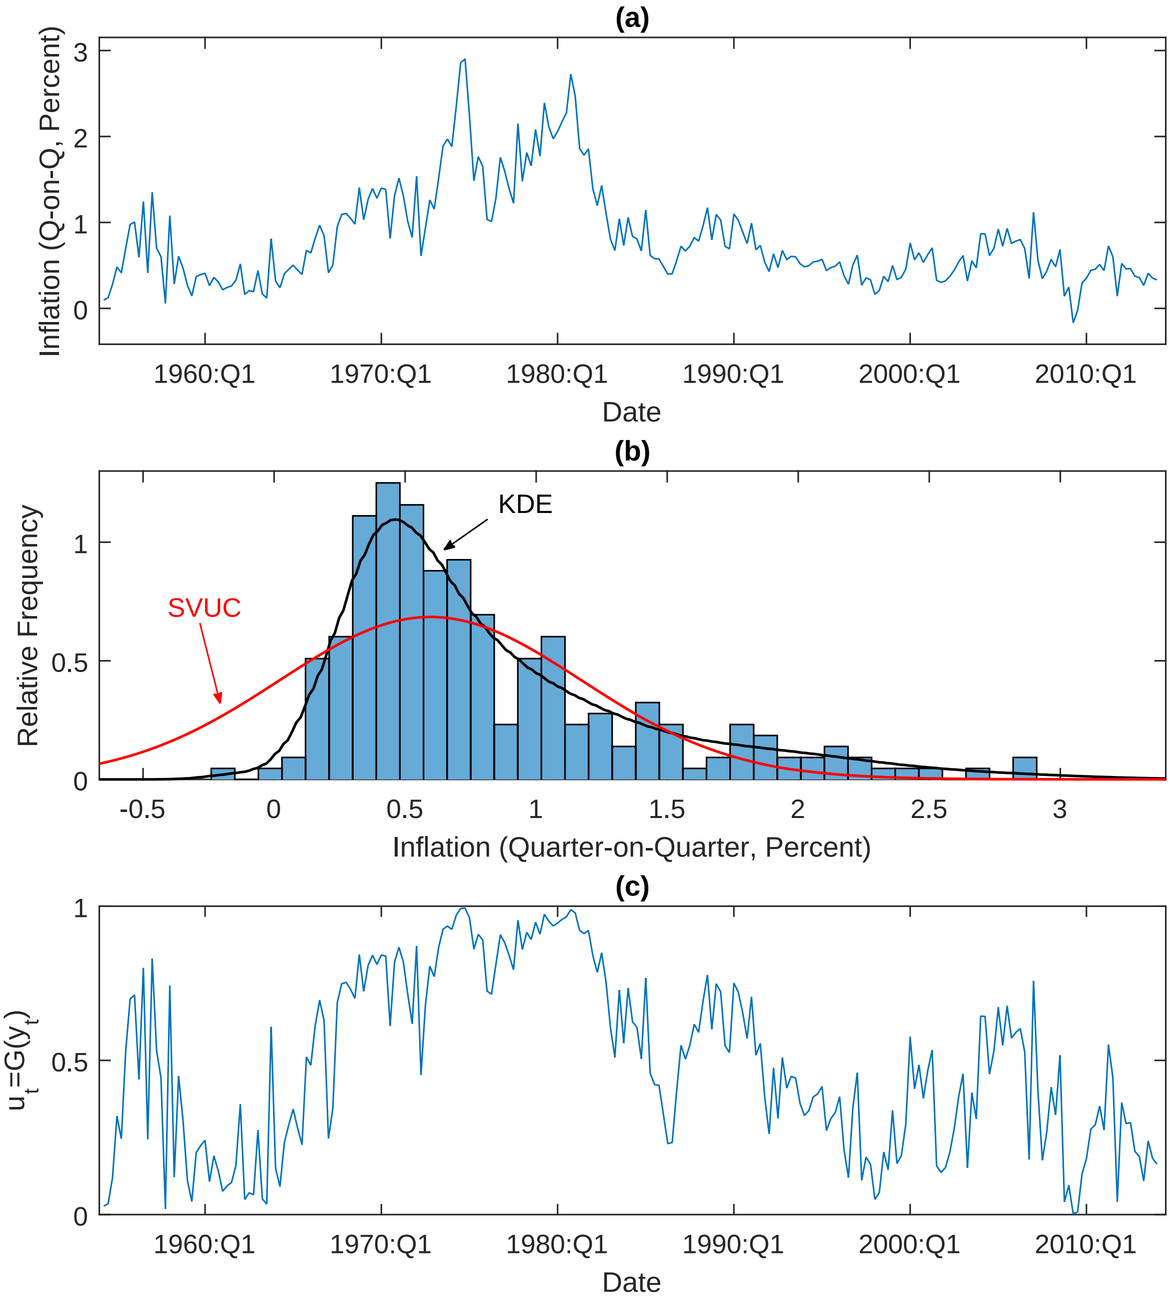

Figure 1: Panel (a) is a time series plot of the quarter-on-quarter U.S. inflation data. Panel (b) plots the (normalized) histogram of the inflation data. Also plotted as a black line is a kernel density estimate (KDE), computed using an adaptive bandwidth. Plotted as a red line for comparison is the marginal density of a plain stochastic volatility model with an unobserved component. Panel (c) is a time series plot of the copula data $u_t = G(y_t)$, where G is the distribution function computed from the KDE in panel (b).

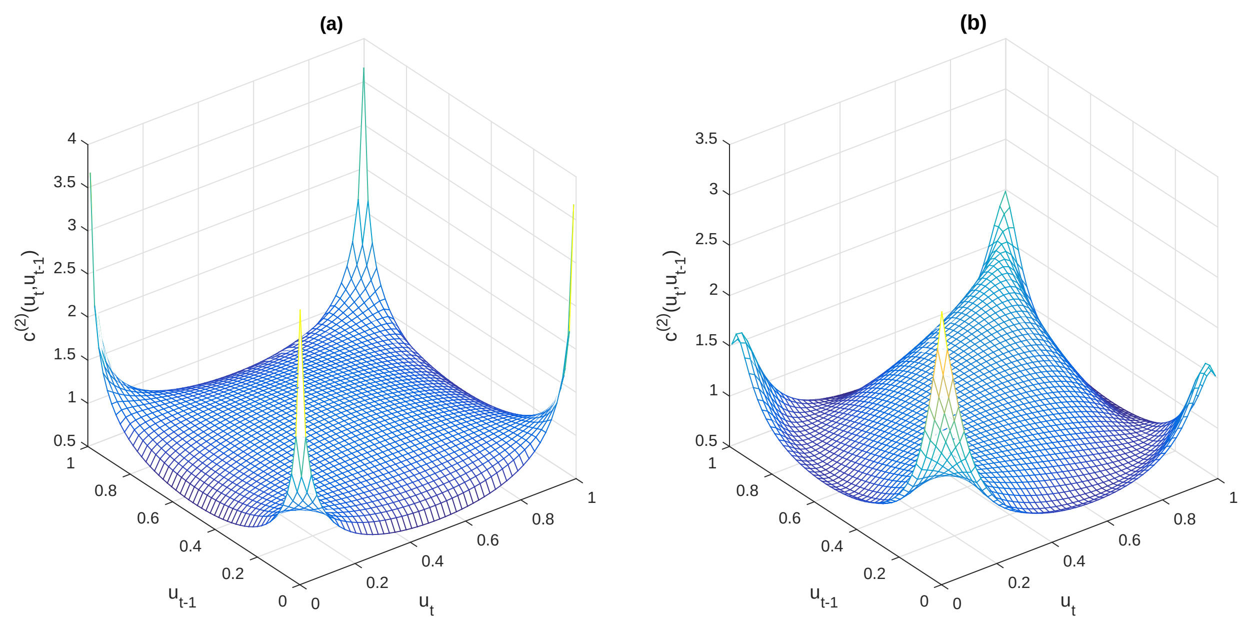

Figure 2: Bivariate marginal copula densities $c^{(2)}(u_t, u_{t-1}|\boldsymbol{\psi})$ of two copulas constructed by inversion of latent nonlinear state space models. Panel (a) is for a first order stochastic volatility model. Here, the values of overall 'level' dependence (ie. Kendall's tau or Spearman's rho) for this copula are exactly zero, yet the copula has high (equally-valued) tail dependence in all four quadrants. Panel (b) is for a Markov switching autoregression. Here, dependence is asymmetric, so that tail dependence differs in each of the four quadrants.

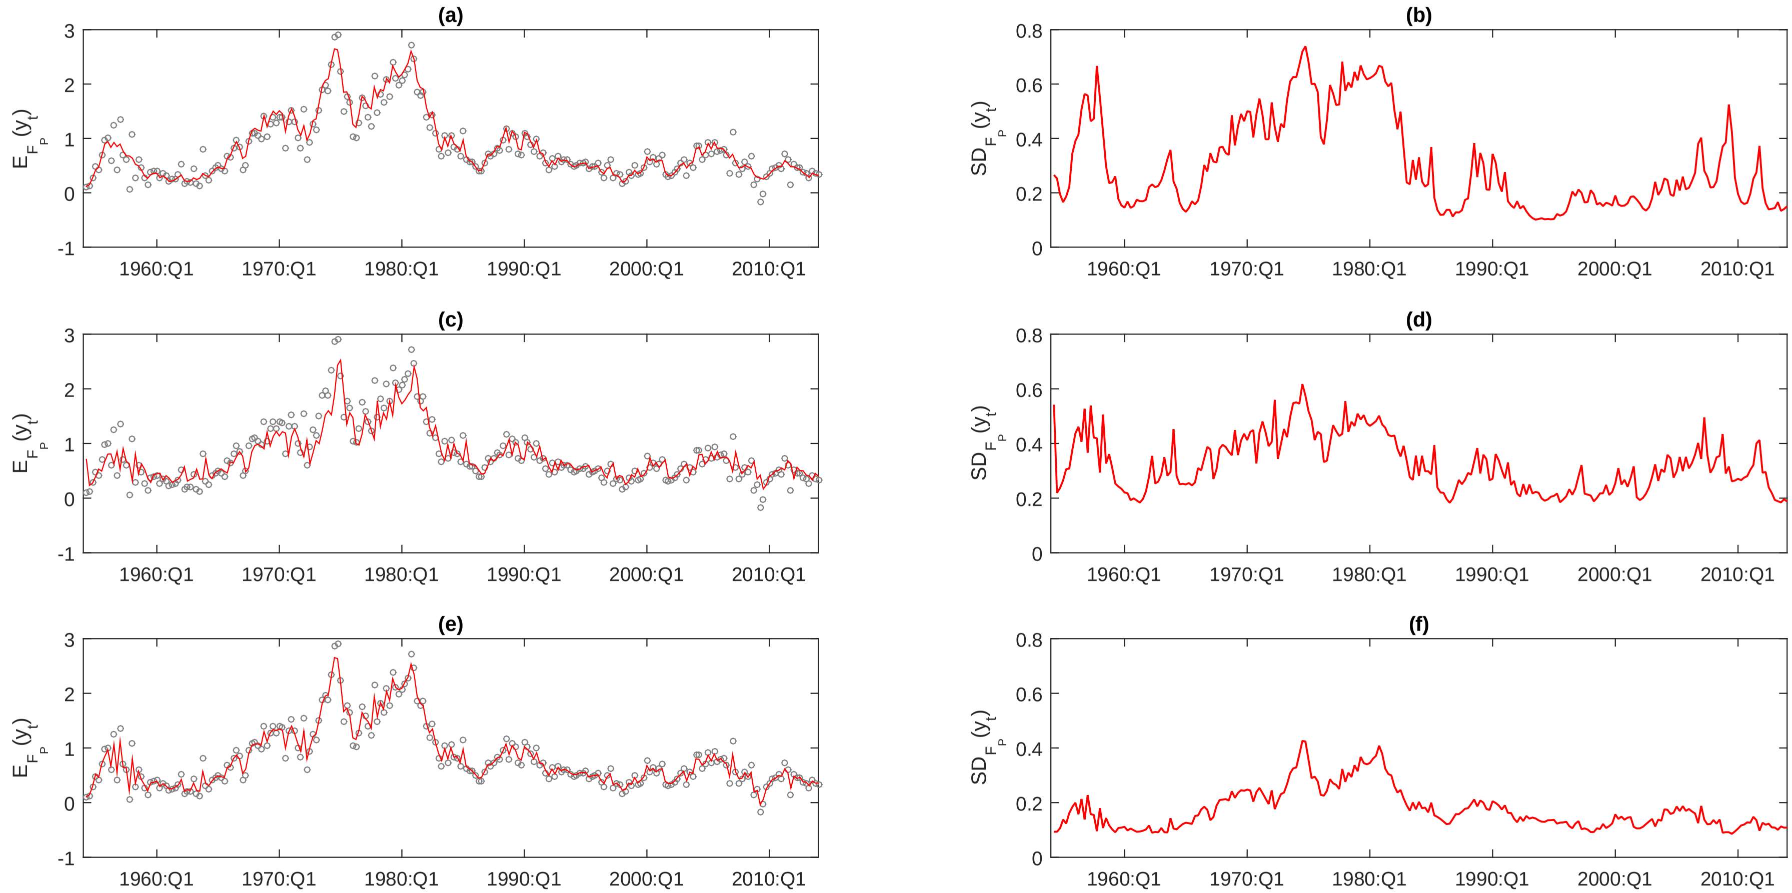

Figure 3: Smoothed estimates of the means $E_{F_P}(y_t)$ [panels (a,c,e)] and the standard deviations $SD_{F_P}(y_t)$ [panels (b,d,f)] of U.S. inflation for the three fitted copula models. Panels (a,b) are for the model outlined in Section 3.1, panels (c,d) are for the model outlined in Section 3.2, and panels (e,f) are for the model outlined in Section 3.3. The times series observations also plotted in the left hand panels as circles.

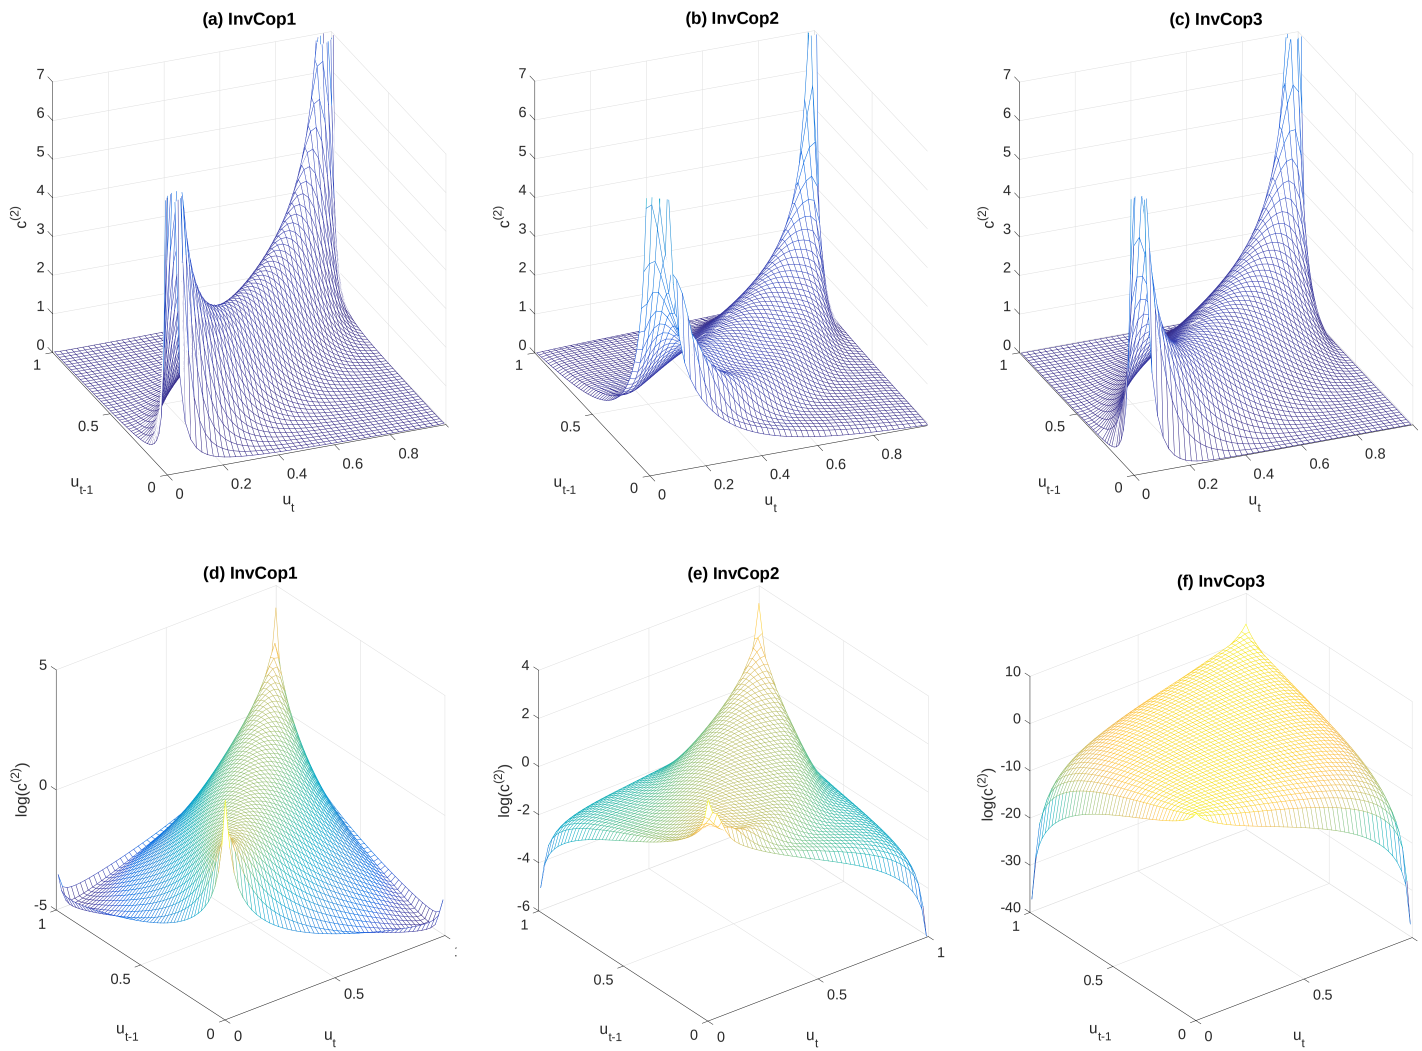

Figure 4: Marginal copula densities $c^{(2)}(u_t, u_{t-1}|\hat{\boldsymbol{\psi}})$ for each of the three inversion copula models fitted to the inflation data. Panels (a,b,c) plot the densities with a common vertical axis truncated at 7 for interpretation. Panels (d,e,f) plot the logarithm of the same three densities. Each density has been computed at the posterior mean $\hat{\boldsymbol{\psi}}$ of the copula parameters. The three latent state space models for the inversion copulas are: panels (a,d) the stochastic volatility model with unobserved component (InvCop1); panels (b,e) the Markov switching AR(1) model (InvCop2); and panels (c,f) the Gaussian unobserved AR(4) component model (InvCop3).

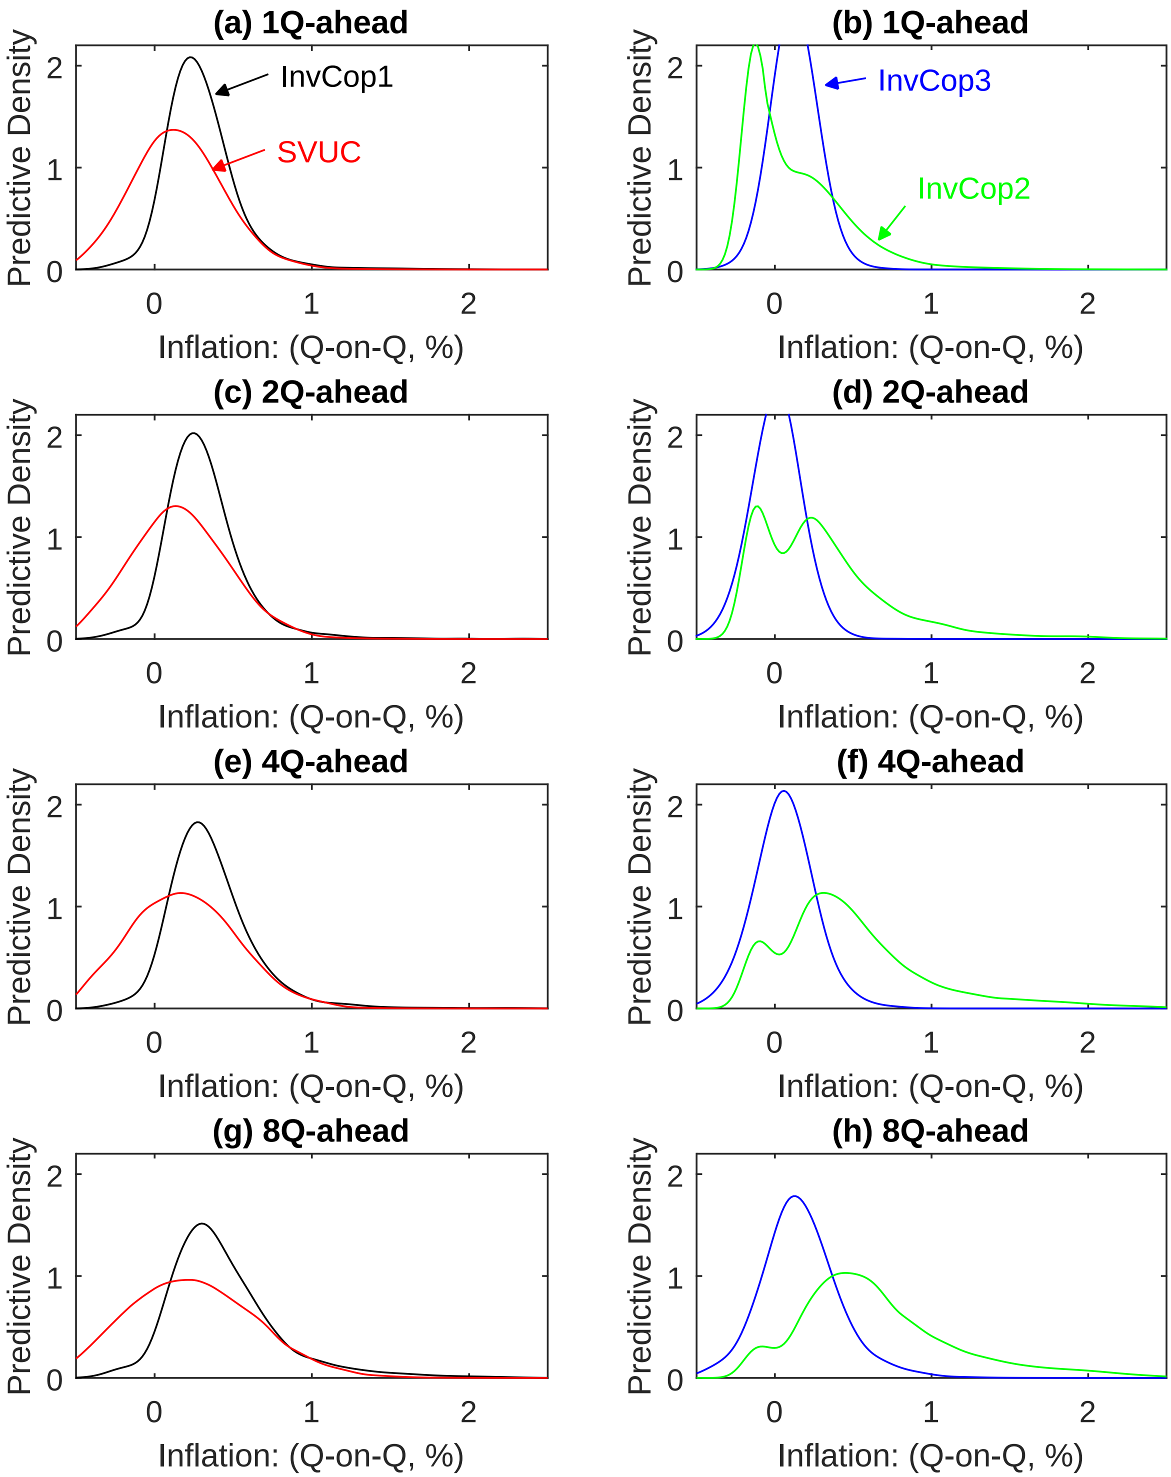

Figure 5: Predictive densities for U.S. inflation (quarter-on-quarter, in percent) using a forecast origin of 2009:Q1. Forecasts are made 1, 2, 4 and 8 quarters ahead. Panels (a,c,e,g) present the predictive densities for the inversion copula model in Section 3.1 (InvCop1), and for the benchmark SVUC model. Panels (b,d,f,h) present the predictive densities for the two inversion copula models in Section 3.2 (InvCop2), and Section 3.3 (InvCop3).