ARTICLE IN PRESS

International Journal of Forecasting ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Copula modelling of dependence in multivariate time series

Michael Stanley Smith *

Melbourne Business School, University of Melbourne, Australia

ARTICLE INFO

Keywords:
Copula model
Nonlinear multivariate time series
Bayesian model averaging
Multivariate stationarity

ABSTRACT

Almost all existing nonlinear multivariate time series models remain linear, conditional on a point in time or latent regime. Here, an alternative is proposed, where nonlinear serial and cross-sectional dependence is captured by a copula model. The copula defines a multivariate time series on the unit cube. A drawable vine copula is employed, along with a factorization which allows the marginal and transitional densities of the time series to be expressed analytically. The factorization also provides for simple conditions under which the series is stationary and/or Markov, as well as being parsimonious. A parallel algorithm for computing the likelihood is proposed, along with a Bayesian approach for computing inference based on model averages over parsimonious representations of the vine copula. The model average estimates are shown to be more accurate in a simulation study. Two five-dimensional time series from the Australian electricity market are examined. In both examples, the fitted copula captures a substantial level of asymmetric tail dependence, both over time and between elements in the series.

1. Introduction

Capturing cross-sectional and serial dependence jointly in multiple time series can improve empirical analyses in many applications. These include problems in areas as diverse as political science (Enders & Sandler, 1993), neuroimaging (Goebel, Roebroeck, Kim, & Formisano, 2003) and particularly macroeconomics (Sims, 1980). While the linear vector autoregression (VAR) is the most widely used model for such data, linearity can prove a restrictive assumption, and, increasingly, nonlinear models are preferred. The most popular of these are models with time-varying parameters (e.g. Cogley & Sargent, 2005; Koop, Leon-Gonzalez, & Strachen, 2009; Primiceri, 2005) and/or regime-switching (e.g. Hamilton, 1990; Sola & Driffill, 1994). However, such models remain linear, conditional upon a point in time or latent state. An alternative is to use a copula model. This allows for nonlinear and asymmetric cross-sectional and serial dependence, using a copula function with constant parameters and no latent regimes.

Copula models (Joe, 1997; Nelsen, 2006) are used widely for modelling dependence in cross-sectional data. They are attractive because they separate the modelling of the location and scale from any dependence, the latter of which is modelled on the unit cube. Patton (forthcoming) provides a recent overview of the literature on multivariate time series copula models. However, almost all previous work has used copulas to capture cross-sectional dependence only. For example, Patton (forthcoming) suggests using a copula to capture cross-sectional dependence conditional on the past. For univariate time series, Joe (1997), Beare (2010), Chen and Fan (2006), Domma, Giordano, and Perri (2009), Ibragimov (2009) and Smith, Min, Almeida, and Czado (2010) all consider using copulas to model serial dependence. The aim of this paper is to extend their use to enable the modelling of both cross-sectional and serial dependence in multivariate time series data.

* Correspondence to: Melbourne Business School, 200 Leicester Street, Carlton, VIC, 3053, Australia.
E-mail address: mike.smith@mbs.edu.

http://dx.doi.org/10.1016/j.ijforecast.2014.04.003
0169-2070/

A common approach to modelling multivariate time series is to select the form of the marginal distribution of the series at a point in time, along with a transitional density, from which the joint distribution of the series is derived. However, it is shown here that the analogous approach of selecting a copula to account for cross-sectional dependence, along with a transitional density on the unit cube, does not provide a well-defined copula model in general. Instead, the reverse approach is proposed, where a copula is selected to account for the dependence in the entire series, then the marginal distribution of the series and the transitional density are derived from this. In doing so, one difficulty is to select a flexible copula that provides closed form expressions for both of these. A copula called a drawable vine (or 'D-vine', see Aas, Czado, Frigessi, & Bakken, 2009; Bedford & Cooke, 2002) is used here to solve this problem.

Vines are copulas for higher dimensions that are constructed from a sequence of nested bivariate copula components called 'pair-copulas'. They are flexible because any combination of bivariate copulas can be used for the pair-copulas; see Haff, Aas, and Frigessi (2010) and Kurowicka and Cooke (2006) for recent overviews. The D-vine also has the property that the copula of any contiguous subvector remains a D-vine, with component pair-copulas that are subsets of those of the original vine. It is shown here how this 'closure' property makes the D-vine attractive for modelling the dependence in multivariate time series. Block functionals, which are equal to the product of subsets of the pair-copulas, are introduced to allow the marginal and transitional densities of the multivariate time series to be expressed concisely.

The block functionals extend the idea of traditional partial correlation matrices to the case of nonlinear time series, and also provide simple conditions under which the series is strongly stationary and/or Markovian. In these latter cases, it is shown that the D-vine is also a parsimonious representation of dependence, with equality between many of the block functionals. Parametric copulas that allow for bi-directional dependence (i.e., both positive and negative associations) should be employed for the pair-copulas. Here, a mixture of rotated Archimedean copulas that allows for bi-directional dependence is proposed. When Gaussian pair-copulas are used, the D-vine can be shown to be a Gaussian copula (Haff et al., 2010). In this special case, the model nests those employed for multivariate time series by Biller and Nelson (2003), Lambert and Vandenhende (2002) and Smith and Vahey (2013).

The dimension of the copula is large at $N = Tm$, where m is the dimension of the multivariate vector and T is the number of observations. For the D-vine copula model, Aas et al. (2009) and Smith et al. (2010) compute the likelihood using an $O(N^2)$ recursive serial algorithm. This can involve too many computations to be employed in practice here. However, we show here how the algorithm can be re-ordered to allow the computations to be undertaken efficiently in parallel, thus greatly speeding execution of the algorithm in real time. When the series is Markov of order p, it is also shown how the absolute number of computations can be reduced to $O(Tpm^2)$. In addition, in independent work, Brechmann and Czado (2012) propose the use of D-vines for the modeling of serial dependence in multivariate time series. However, without exploiting these computational insights, they restrict their attention to the bivariate case where $m = 2$. Moreover, they do not employ or exploit the block notation, nor do they consider the properties of the time series, such as stationarity. Rémillard, Papageorgiou, and Soustra (2012) also consider using a copula to model the dependence of a multivariate series on the unit cube. They suggest using either an elliptical copula or a N-dimensional Archimedean copula. However, the latter is an unrealistic choice beyond the bivariate case, because it characterises all dependence using a single parameter, nor allows for bi-directional dependence. In comparison, it is shown here that a D-vine copula for a stationary Markov series has a similar number of pair-copulas as there are parameters in a VAR.

A Bayesian approach for estimating the copula model is proposed for stationary series. Bayesian inference has become increasingly popular for multivariate time series models, where point priors and model averaging can improve estimates and predictions; see Garrat, Koop, Mise, and Vahey (2009), George, Sun, and Ni (2008), Jochmann, Koop, Leon-Gonzalez, and Stachan (2013) and Korobilis (2013) for examples. Here, a two-level prior is suggested that places point masses on the blocks for Markov order selection, and also on each pair-copula, so as to be equal to the independence copula for parsimony. The posterior distribution is evaluated using a computationally efficient Markov chain Monte Carlo (MCMC) sampling scheme. Posterior inference on the dependence structure is obtained via simulation from the fitted copula. This includes estimates of both partial and marginal measures of pairwise serial and cross-sectional dependence, predictive distributions and estimates of generalized impulse-response functions.

The effectiveness of the approach is demonstrated using a simulation study. This shows that the selection method can improve the accuracy of the estimates of both the conditional and unconditional dependence structures. To illustrate that the methodology is effective in practice, two five-dimensional time series are examined and a multivariate analysis is found to be beneficial. One is the daily maxima of electricity demand in the five regions of the Australian National Electricity Market (NEM), and the other is daily spot prices of electricity in the same five regions. In both examples, the copula model captures significant levels of asymmetric and heavy-tailed dependence. It is well known that accounting for such properties in cross-sectional dependence can be beneficial (Patton, 2006), and the examples here demonstrate that this is also the case for multivariate serial dependence. The first example extends copula models for multivariate extremes (e.g. Favre, El Adlouni, Perreault, Theiémonge, & Bobée, 2004) to also account for serial dependence. The second demonstrates the flexibility of the copula model when applied to a series that economic theory suggests may exhibit extensive nonlinear dependence.

The rest of the paper is organized as follows. Section 2 outlines the copula model for multivariate time series, including block notation and the conditions for the series to be stationary and/or Markov. Estimation using fast parallel algorithms for computing the likelihood, priors

Please cite this article in press as: Smith, M. S., Copula modelling of dependence in multivariate time series. International Journal of Forecasting (2014), http://dx.doi.org/10.1016/j.ijforecast.2014.04.003