Markov Switching Regular Vine Copulas

Stöber, Jakob and Czado, Claudia
Lehrstuhl für Mathematische Statistik, Technische Universität München
Parkring 13
85748 Garching-Hochbrück, Germany
E-mail: stoeber@ma.tum.de, cczado@ma.tum.de

RÉSUMÉ (ABSTRACT)

Regular vine (R-vine) copulas, which are entirely constructed from bivariate copulas as building blocks, constitute a flexible class of high dimensional dependence models. In this paper, we study Markov switching R-vine copulas, which extend the flexibility of R-vine copulas by letting parameters and structures vary over time. Focusing on estimation methods, we give an extensive study of a modified Expectation Maximization (EM) algorithm and a Bayesian Gibbs sampling procedure. Using Bayesian methods, we can also assess the uncertainty in parameter estimates which is demonstrated in an application to exchange rate data.

Introduction

During the last years, new and increased computational capabilities have given birth to new and sophisticated models for statistical analysis. Using the theorem of Sklar (1959) to separate the description of marginal models from the dependence structure, one branch of multivariate statistics focuses on the study of high dimensional copula structures. Among these, Regular vine (R-vine) copulas, which are entirely constructed from bivariate copulas as building blocks, constitute one particular flexible class. Since Aas et al. (2009) introduced the R-vine copula structure in an inferential context, it has become popular for the accurate description of fixed dependence structures. However, it lacks the possibility to model dependence structures varying over time as it is displayed in international financial data sets. In this paper, we attempt to fill in this gap by introducing the so called Markov switching (MS) R-vine copula, generalizing work of Cholette et al. (2009). In particular, we study two methods for the estimation of parameters in these MS R-vine models, a modification of the Expectation Maximization (EM) Algorithm (Hamilton, 1990) and a Bayesian Gibbs sampler. This enables us to further assess the uncertainty in the estimation of parameters for a multivariate regime switching model, which was not possible beforehand. The remainder of this paper is structured as follows: Section 2 (Regular Vine Copulas) introduces R-vine copulas while Section 3 (Markov Switching Models) recapitulates MS models. Having introduced our main object we focus on estimation methods in Section 4 (Parameter Estimation), 5 (Expectation Maximization Algorithm) and 6 (Bayesian Gibbs Sampling) and present an application to exchange rate data in Section 7 (Application: Exchange Rates). Section 8 (Conclusion) concludes the paper and gives an outlook to possible areas of further research.

Regular Vine Copulas

Building on work of Joe (1996), Bedford and Cooke (2001) introduced the hierarchical R-vine structure in the area of multivariate statistics. Employing notation and methods from graph theory, they define an R-vine $\mathcal{V}$ in d variables as a sequence of connected trees (undirected, acyclic graphs) $T_1, \ldots, T_{d-1}$. A five-dimensional example is displayed in Figure 1. For a more detailed introduction to R-vine modeling we refer to Kurowicka and Cooke (2006), the computational treatment has been considered by Dißmann (2010).

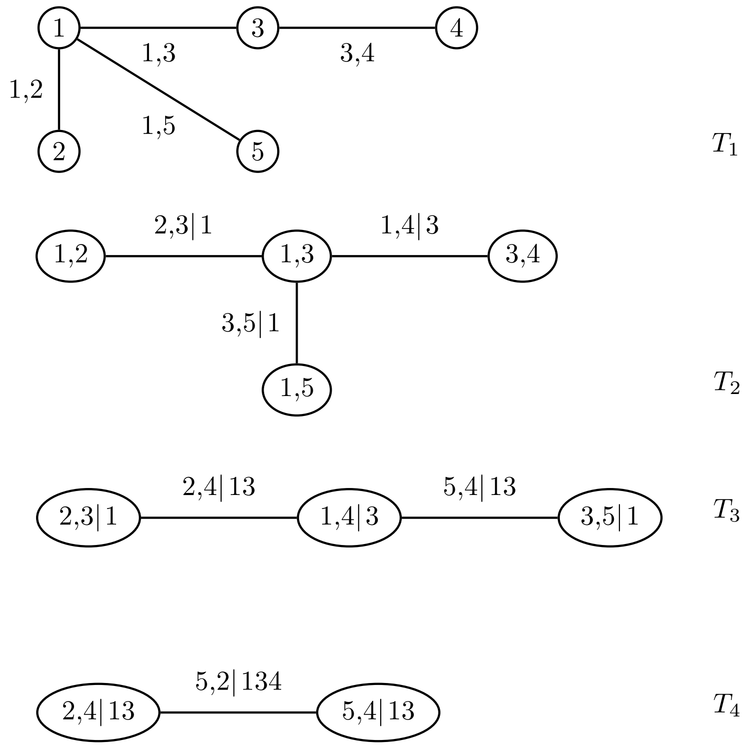

Figure 1: An R-vine tree sequence in five dimensions with corresponding edge indices in the notation of Czado (2010).

To satisfy the needs for statistical application, the nodes N_i and edges E_i, for $1 \leq i \leq d-1$, are required to satisfy the following properties:

1. T_1 has nodes $N_1 = 1, \ldots, d$.
2. For $i \geq 2$, T_i has nodes $N_i = E_{i-1}$ and a corresponding set of edges.
3. If two nodes are joined by an edge in T_{i+1}, the corresponding edges must share a common node in T_i.

This third technical condition is called proximity condition and it ensures that the densities of R-vine copulas can be represented in an integral free way. To build up a statistical model on this graph theoretic structure, we associate each edge $e \hat{=} j(e), k(e)|D(e)$ in the vine with a bivariate copula. This bivariate copula will be the copula corresponding to the bivariate conditional marginal distribution of $X_{j(e)}$ and $X_{k(e)}$ given $\mathbf{X}_{D(e)}$. For more details on our notation, we refer to the illustration in Figure 1 and Czado (2010). The density of an R-vine copula with a set $\mathbf{B}$ of bivariate copulas with parameters Θ associated to an R-vine $\mathbf{V}$ is denoted with $c(.|\mathcal{V}, \mathbf{B}, \mathbf{\Theta})$.

Markov Switching Models

In Markov switching models, which have been introduced into econometrics and statistical modeling by Hamilton (1989) different states of the world or the economy affect the development of a time series. For this, we assume a hidden Markov chain (S_t) with states 1,...,n, describing the progression of these states in time. For simplicity, we further assume that the Markov chain is homogeneous and of first order such that it can be completely characterized by its transition matrix $P(S_t = i|P_{t-1} = j) = P_{ij}$.

Many financial time series, like stock returns or exchange rates are influenced by external factors like the state of the economy or monetary policies which are not directly observable and which are thus included in the hidden state variable. Depending on these factors, different types of dependence are present in our data. This behavior is described by the specification of a Markov switching R-vine copula model for a time series $(\mathbf{u}_t)$ of a random vector with uniform marginal distributions as

$$c(\mathbf{u}_t|S_t, \mathcal{V}_{S_t}, \mathbf{B}_{S_t}, \mathbf{\Theta}_{S_t}) = \sum_{i=1}^{n} 1_{\{S_t=k\}} \cdot c(\mathbf{u}_t|\mathbf{V}_k, \mathbf{B}_k, \mathbf{\Theta}_k).$$

For the estimation problems considered in the upcoming section, we assume the specification of $\mathcal{V}_k$ and $\mathbf{B}_k$, for all $1 \leq k \leq n$, to be given such that only the sets of parameters $\mathbf{\Theta}_k$ are subject to estimation.

Parameter Estimation

Since our main interest lies in describing and estimating dependence structures, we will, for the remainder of this paper, assume that all random vectors have uniform marginal distributions such that they can directly be described by the means of a copula.

In order to draw inference for a MS model, we need to overcome the challenge of having unobserved latent variables. To do so, let us consider the decomposition of the joint density of a time series of random vectors $\tilde{\mathbf{u}}_T = (\mathbf{u}_1, \ldots, \mathbf{u}_T)$ into conditional densities. In order to keep our notation short, we will suppress denoting the parameters of the MS R-vine model where they are not required.

$$\begin{aligned} f(\tilde{\mathbf{u}}_T) &= f(\mathbf{u}_1) \cdot \prod_{t=2}^{T} f(\mathbf{u}_t|\tilde{\mathbf{u}}_{t-1}) \\ &= \left[\sum_{i=1}^{n} f(\mathbf{u}_1|S_1 = i)P(S_1 = i)\right] \cdot \prod_{t=2}^{T} \left[\sum_{i=1}^{n} f(\mathbf{u}_t|S_t = i) \cdot P(S_t = i|\tilde{\mathbf{u}}_{t-1})\right]. \end{aligned}$$

The unconditional probabilities $P(S_1 = i)$ in this expression are known from the stationary distribution of the Markov chain, which we assume to exist, and we can apply the filter of Hamilton (1989) to obtain the probabilities $P(S_t = i|\tilde{\mathbf{u}}_{t-1})$. Being able to calculate all conditional probabilities in the equation, we can evaluate the likelihood for a given set of observations.

While this also enables us to estimate model parameters through maximization of the joint likelihood function, a direct approach is difficult since the problem is analytically not tractable and numerically difficult, because of the latent state variables. Furthermore, obtaining point estimates for parameters in a high dimensional setting is of little use without being able to assess the uncertainty in these estimates since inferred results might not be significant. Min and Czado (2010) demonstrated that estimating the covariance matrix of ML estimators for R-vines using the numerical Hessian often leads to non positive definite matrices and the effects are similar in the framework of MS R-vines. Because of long computation times also bootstrap methods are not a feasible way to examine the uncertainty of ML estimates. In the following, we discuss two possible approaches to overcome these challenges, tackling the problem of computation time with an approximative version of the EM Algorithm and demonstrating how to assess parameter uncertainty with a Bayesian approach.

Expectation Maximization Algorithm

The EM Algorithm, which has been proposed by Hamilton (1990) to overcome the problem of having unobserved state variables, constitutes an iterative procedure consisting of the following steps:

- Form the conditional expectation of unobserved variables and
- maximize the likelihood, replacing the latent state variables with their conditional expectation.

While this procedure has significant advantages in cases where it makes the maximization step analytically tractable, we would, in our case, still need to maximize the likelihood of a multivariate copula which is numerically difficult in high dimensions. To circumvent this problem, we can exchange the joint maximization with a stepwise maximization for each copula on the vine as it is similarly applied for model selection by Dißmann (2010).

We call this the *stepwise EM Algorithm*. Being a close approximation to the proper EM Algorithm, it leads to comparable results for most applications. While there are theoretical results considering the convergence of the EM Algorithm, c.f. Wu (1983), we certainly loose these properties with our approximation. All limit and convergence theorems however do rely on proper maximization at each step of the algorithm. This is almost impossible to guarantee in our case where we are faced with high dimensional numerical maximization problems. Therefore, also an implementation of the "proper" EM Algorithm has to be considered an approximation, which further justifies the use of our stepwise procedure.

Bayesian Gibbs Sampling

For a more accurate study and in order to assess the uncertainty in parameter estimation, we follow a Bayesian approach. To develop a Bayesian algorithm, two problems have to be adressed. We need to deal with the underlying Markov structure, and we need to estimate the parameters in an R-vine copula specification.

Since the conditional margins are unknown, the estimation of R-vine copula parameters is done in an MCMC framework utilizing the Metropolis-Hastings algorithm. Following Min and Czado (2010), we assume non informative priors for all copula parameters in our model, restricting the support to a finite interval in cases where the parameter range is not compact.

To deal with the latent variables, we follow the approach by Kim and Nelson (1998), such that the Gibbs sampler for the full model circles through the following steps:

1. Sample $(\tilde{S}_T|\tilde{\mathbf{u}}_T, \mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_n, P)$
2. Sample $(P|\tilde{S}_T, \tilde{\mathbf{u}}_T, \mathbf{\Theta}_1, \ldots, \mathbf{\Theta}_n) = (P|\tilde{S}_T)$
3. Sample $(\mathbf{\Theta}_i|\mathbf{u}_{\{t\in\{1,\ldots T\}|S_t=i\}})$ for $i = 1, \ldots, n$.

Application: Exchange Rates

For a demonstration of our procedures, we consider a data set consisting of exchange rates against the US dollar for nine different currencies, namely Euro, British pound, Canadian dollar, Australian dollar, Brazilian real, Japanese yen, Chinese yuan, Swiss franc and Indian rupee. This data set has also been studied by Czado et al. (2010) and we refer to their paper for more details on the data set which has 1007 daily observations.

We assume two regimes to be present in the data set for which we do first need to identify suitable R-vine structures $\mathcal{V}$ and sets of copulas $\mathbf{B}$ before we can turn to parameter estimation. To do so, we conduct a rolling window analysis to identify a regime which describes dependencies in times of crisis and we choose a second regime describing the average dependence over the whole data set. For this we use the structure selection methods of Dißmann (2010) and we refer to his work for more details on the employed methods. Considering the sets of copulas $\mathbf{B}$, we select Gaussian copulas to correspond to the conditional distributions of the first and second tree of the R-vine in the "normal" regime and of the second tree in the "crisis" regime. Rotated Gumbel copulas are assigned to the first tree in the crisis regime in order to account for possible tail dependencies.

Given these structures we iterate the approximative EM Algorithm until convergence and use the obtained point estimates as starting values for the Bayesian Gibbs sampling procedure. The resulting estimates for the smoothed probabilities of being in the "crisis" regime as obtained with the Gibbs sampling approach are plotted in Figure 2. Figure 3 shows histogram plots of the posterior distribution of the copula parameters in the "crisis" regime using a sample size of 6000 after appropriate thinning.

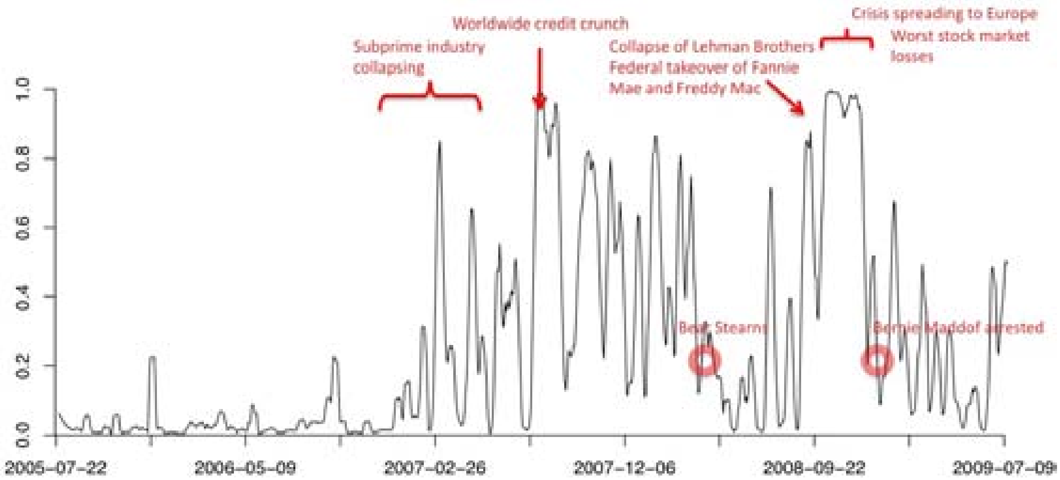

Figure 2: Smoothed probability that the latent state variable indicate the presence of the crisis regime plotted against the time. The annotations indicate important events during the financial crisis.

As Figure 2 demonstrates, times where our model indicates a high probability of the "crisis" structure to be present correspond to important events during the financial crisis. This confirms our conjecture that a dependence structure which is different from the long term average is present during these times. The distribution estimates in Figure 3 show that the estimated confidence bands are sufficiently narrow to draw conclusions from different posterior mode or posterior mean estimates. For the parameter $\theta_{6,5}$ many samples are close to the value 1 for which the Clayton copula corresponds to the independence copula, yielding the hypothesis that the exchange rates of the Swiss franc and the Chinese yuan to US dollar are independent during times of crisis.

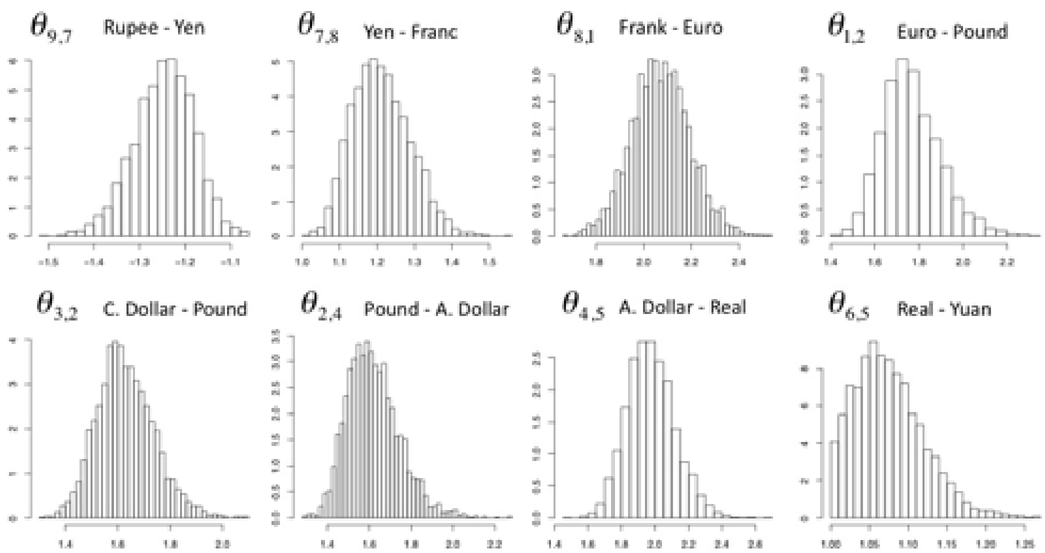

Figure 3: Histogram plots of our samples from the posterior distributions of the parameters of the "crisis" vine structure.

Conclusion

The model presented in this paper allows for a efficient parameter estimation for time-varying copula structures and to assess the uncertainty in such estimates. Building the model on truncated R-vines allows for highly flexible structures while keeping a low number of parameters, minimizing the risk of overestimating the model. Applying our procedures to exchange rate data, we find that the dependence structures where similar during important events of the financial crisis, which supports the use of MS models to account for behavioral changes during such times.

REFERENCES

Aas, K., Czado, C., Frigessi, A., and Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics* **44** (2), 182-198.

Bedford, T., and Cooke, R. M. (2001). Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence* **32**, 245-268.

Cholette, L., Heinen, A. and Valdesogo, A. (2009). Modeling International Financial Returns with a Multivariate Regime-Switching Copula. *Journal of Financial Econometrics* **7** (4), 437-480.

Czado, C. (2010) Pair copula constructions of multivariate copulas. In *Copula Theory and Its Applications*, Jaworski, P., Durante, F., Härdle, W., and Rychlik, T. (Eds.). Springer, Berlin.

Czado, C., Schepsmeier, U., and Min, A. (2010). Maximum likelihood estimation of mixed C-vines with application to exchange rates. To appear in *Statistical Modeling.*

Dißmann, J. (2010). Statistical inference for regular vines and application. *Diploma thesis, Technische Universität München, Germany.*

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* **57**, 357-384.

Hamilton, J. D. (1990). Analysis of Time Series Subject to Regime Changes. *Journal of Econometrics* **45**, 39-70.

Joe, H. (1996). Families of m-variate distributions with given margins and m(m-1)/2 bivariate dependence parameters. In *Distributions with Fixed Marginals and Related Topics*, Rüschendorf, L., Schweizer, B. and Taylor, M. D. (Eds.)

Kim, C.-J., and Nelson, C. R. (1998). Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime switching. *The Review of Economics and Statistics* **80** (2), 188-201.

Kurowicka, D., and Cooke, R. M. (2006). *Uncertainty Analysis with High Dimensional Dependence Modeling.* Wiley, Chichester.

Min, A., and Czado, C. (2010). Bayesian inference for multivariate copulas using pair-copula constructions. *Journal of Financial Econometrics* **8** (4), 511-546.

Sklar, A. (1959) Fonctions de repartition a n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris* **8**, 229-231.

Wu, C. F. (1983). On the convergence properties of the EM Algorithm. *Annals of Statistics* **11** (1), 95-103.