Time series models with infinite-order partial copula dependence

Alexander J. McNeil

University of York

1st December 2021
Institute for Statistics and Mathematics, WU Vienna

joint work with Martin Bladt (Lausanne)

Overview

Overview

Summary

- In this talk we present:
  1. time series copula processes with infinite-order partial dependence;
  2. a parameterization of models using partial Kendall rank correlations;
  3. a generalization of the classical concepts of causality and invertibility for linear processes;
  4. non-Gaussian generalizations of classical Gaussian processes such as ARMA, seasonal ARMA, ARFIMA and FGN.
- With the added flexibility we can obtain superior statistical fits in real-world applications.

The copula approach to time series

- Given data $\{x_1, \ldots, x_n\}$ the idea is to find an appropriate strictly stationary stochastic process $(X_t)_{t\in\mathbb{Z}}$ consisting of:
  1. a continuous marginal distribution F_X;
  2. a copula process $(U_t)_{t\in\mathbb{Z}}$ satisfying $U_t = F_X(X_t)$ for all t.
- The latter is a process of standard uniform random variables.
- The main examples in the literature are first-order Markov copula processes (Chen and Fan, 2006; Ibragimov, 2009) and their higher-order d-vine generalizations (Smith, Min, Almeida, and Czado, 2010; Beare and Seo, 2015; Brechmann and Czado, 2015; Nagler, Krüger, and Min, 2020).
- These are based on pair copula decompositions described by Joe (1997) and Bedford and Cooke (2002), i.e. models constructed from bivariate copulas.

Overview

Notation for s-vines

- Let $(C_k)_{k\in\mathbb{N}}$ denote a sequence of bivariate copulas.
- Assume throughout that every $C_k \in \mathcal{C}^\infty$ and has density c_k which is strictly positive on $(0,1)^2$. Can probably be weakened.
- For $k \in \mathbb{N}$ let the forward and backward Rosenblatt functions

$$R_k^{(1)} : (0,1)^k \times (0,1) \to (0,1) \text{ and } R_k^{(2)} : (0,1)^k \times (0,1) \to (0,1)$$

be defined in a recursive, interlacing fashion by $R_1^{(1)}(u,x) = h_1^{(1)}(u,x)$, $R_1^{(2)}(u,x) = h_1^{(2)}(x,u)$ and, for $k \geq 2$,

$$R_k^{(1)}(\boldsymbol{u},x) = h_k^{(1)}\left(R_{k-1}^{(2)}(\boldsymbol{u}_{-1},u_1), R_{k-1}^{(1)}(\boldsymbol{u}_{-1},x)\right)$$
$$R_k^{(2)}(\boldsymbol{u},x) = h_k^{(2)}\left(R_{k-1}^{(2)}(\boldsymbol{u}_{-k},x), R_{k-1}^{(1)}(\boldsymbol{u}_{-k},u_k)\right)$$

where $h_k^{(i)}(u_1,u_2) = \frac{\partial}{\partial u_i} C_k(u_1,u_2)$ and $\boldsymbol{u}_{-i}$ indicates the vector $\boldsymbol{u}$ with ith component removed.

S-vine copulas

An n-dimensional s-vine copula $C_{(n)}$ has density of the form

$$c_{(n)}(u_1,\dots,u_n) = \prod_{k=1}^{n-1}\prod_{j=k+1}^{n} c_k\left(R^{(2)}_{k-1}(\boldsymbol{u}_{[j-k+1,j-1]},u_{j-k}), R^{(1)}_{k-1}(\boldsymbol{u}_{[j-k+1,j-1]},u_j)\right) \tag{1}$$

where $\boldsymbol{u}_{[j-k+1,j-1]} = (u_{j-k+1},\dots,u_{j-1})^\top$.

- This is a a d-vine copula subject to translation invariance conditions.
- A random vector $(U_1,\dots,U_n)$ following $C_{(n)}$ could be an excerpt from a stationary process.
- Moreover, for any $k \in \{1,\dots,n-1\}$ and $j \in \{k+1,\dots,n\}$,

$$R^{(1)}_k(\boldsymbol{u},x) = \mathbb{P}(U_j \le x \mid U_{j-k} = u_1,\dots,U_{j-1} = u_k)$$
$$R^{(2)}_k(\boldsymbol{u},x) = \mathbb{P}(U_{j-k} \le x \mid U_{j-k+1} = u_1,\dots,U_j = u_k).$$

- Where needed in formulas, $R^{(1)}_0(\cdot,x) = R^{(2)}_0(\cdot,x) = x$.

S-vine process

Definition (S-vine process)

A strictly stationary time series $(X_t)_{t\in\mathbb{Z}}$ is an s-vine process if for every $t \in \mathbb{Z}$ and $n \geq 2$ the n-dimensional marginal distribution of the vector $(X_t, \ldots, X_{t+n-1})$ is absolutely continuous and admits a unique copula $C_{(n)}$ with a joint density $c_{(n)}$ of the form (1). An s-vine process $(U_t)_{t\in\mathbb{Z}}$ is an s-vine copula process if its univariate marginal distribution is standard uniform.

- We refer to C_k as the kth partial copula of the process - copula of conditional distribution of (U_{t-k}, U_t) given intervening variables.
- Should be distinguished from the bivariate marginal copula $C^{(k)}$ of (U_{t-k}, U_t). They are related by:

$$C^{(k)}(v_1, v_2) = \int_{[0,1]^{k-1}} C_k\left(R^{(2)}_{k-1}(\boldsymbol{u}, v_1), R^{(1)}_{k-1}(\boldsymbol{u}, v_2)\right) c_{(k-1)}(\boldsymbol{u})\mathrm{d}\boldsymbol{u}.$$

Overview

Construction of finite-order process

- Let $\{C_1, \ldots, C_p\}$ be a finite set of copulas. Think of these as first p terms of sequence $(C_k)_{k\in\mathbb{N}}$ where $C_k = C^{\perp}$ (independence) for $k > p$.
- Write the forward Rosenblatt functions as $R_k = R_k^{(1)}$ and note that they have unique inverses satisfying $R_k^{-1}(\boldsymbol{u}, z) = x \iff R_k(\boldsymbol{u}, x) = z$.
- Let $(Z_k)_{k\in\mathbb{N}}$ be a sequence of iid uniform innovations.
- Set $U_1 = Z_1$ and

$$U_k = R_{k-1}^{-1}\Big((U_1, \ldots, U_{k-1})^\top, Z_k\Big), \quad k \geq 2. \tag{2}$$

- $(U_1, \ldots, U_n)$ has copula density $c_{(n)}$ in (1) with $c_k(u, v) \equiv 1$ for $k > p$.
- Thus $(U_1, \ldots, U_n)$ is a realization from an s-vine process $(U_t)_{t\in\mathbb{Z}}$.
- The construction (2) is found in Joe (2015, page 145).

Finite-order process as Markov process

- For $k > p$ we find that $R_k(\boldsymbol{u}, x) = R_p(\boldsymbol{u}_{[k-p+1,k]}, x)$.
- The recursive equation (2) defining the process satisfies

$$U_k = R_p^{-1}\Big((U_{k-p}, \dots, U_{k-1})^\top, Z_k\Big), \quad k > p.$$

- Thus the process is pth order Markov and can be treated as a Markov process on the state space $(0, 1)^p$.
- It is an example of the non-linear state (NSS) model of Meyn and Tweedie (2009) and is a ϕ-irreducible, aperiodic, Harris-recurrent Markov chain.
- It satisfies the ergodic theorem for Harris chains (Meyn and Tweedie, 2009, Theorem 13.3.3) but many questions remain concerning rates of mixing and ergodic convergence for different sets of copulas $C_1, \dots, C_p$.

Rosenblatt inverse functions

- There is an implied set of functions $S_k : (0,1)^k \times (0,1) \to (0,1)$ such that

$$U_k = R_{k-1}^{-1}\left((U_1,\dots,U_{k-1})^\top, Z_k\right) = S_{k-1}((Z_1,\dots,Z_{k-1})^\top, Z_k), \quad k \geq 2.$$

- These functions satisfy $S_1(z_1, x) = R_1^{-1}(z_1, x)$ and

$$S_k(\boldsymbol{z}, x) = R_k^{-1}\left((z_1, S_1(z_1, z_2), \dots, S_{k-1}(\boldsymbol{z}_{[1,k-1]}, z_k)), x\right), \quad k \geq 2.$$

- We refer to them as Rosenblatt inverse functions.
- We thus have two sets of equations expressing relationship between $(Z_k)_{k\in\mathbb{N}}$ and $(U_k)_{k\in\mathbb{N}}$:

$$\begin{aligned} U_k &= S_{k-1}((Z_1,\dots,Z_{k-1})^\top, Z_k) && \text{(causality)} \\ Z_k &= R_{k-1}((U_1,\dots,U_{k-1})^\top, U_k) && \text{(invertibility)}. \end{aligned}$$

Overview

Gaussian processes and s-vines

By Gaussian processes we refer to processes whose finite-dimensional marginal distributions are multivariate Gaussian distributions with non-singular covariance matrices.

Theorem

1. *Every stationary Gaussian process is an s-vine process.*
2. *Every s-vine process in which the pair copulas of the sequence* $(C_k)_{k \in N}$ *are Gaussian and the marginal distribution* F_X *is Gaussian, is a Gaussian process.*

- The ideas behind the proof are in Joe (2015).
- One implication is that every stationary Gaussian process can be treated as an s-vine; this offers generic (though not necessarily efficient) methods for simulation and estimation.

Gaussian processes as s-vines

For instance, the following Gaussian processes can be easily recast as s-vines:

- ARMA
- Seasonal ARMA (SARMA)
- Fractional Gaussian Noise (FGN)
- ARFIMA

The key logical steps are:

1. Calculation of the acf $(\rho_k)_{k \in \mathbb{N}}$.
2. Calculation of the pacf (partial autocorrelation function) $(\alpha_k)_{k \in \mathbb{N}}$ using well-known one-to-one mapping between acf and pacf.
3. Construction of the copula sequence $(C_k)_{k \in \mathbb{N}}$ by setting C_k to be a Gauss copula with parameter α_k.

New Gaussian processes from s-vines

- It is natural to ask what are the constraints on the sequence of Gauss copula parameters $(\alpha_k)_{k\in\mathbb{N}}$ to obtain a well-behaved Gaussian process.
- Obviously, we require $|\alpha_k| < 1$. The s-vine process will be stationary by construction, but not necessarily ergodic.
- A well known necessary and sufficient condition for a Gaussian process to be mixing is that the acf satisfies $\rho_k \to 0$ as $k \to \infty$ (Maruyama, 1970; Cornfeld, Fomin, and Sinai, 1982).
- Mixing implies ergodicity of the process.
- However, $\alpha_k \to 0$ is not sufficient for mixing behaviour. Counterexample given by sequence $\alpha_k = (k+1)^{-1}$ which yields $\rho_k = 0.5$, for all k.
- A sufficient (but not necessary) condition for mixing is that $\sum_{k=1}^{\infty} |\alpha_k| < \infty$ (Debowski, 2007).

Rosenblatt functions for Gaussian processes

The forward Rosenblatt functions for a mixing Gaussian process with pacf $(\alpha_k)_{k\in\mathbb{N}}$ can be calculated to be

$$R_k(\boldsymbol{u}, x) = \Phi\left(\frac{\Phi^{-1}(x) - \sum_{j=1}^{k} \phi_j^{(k)} \Phi^{-1}(u_{k+1-j})}{\sigma_k}\right),$$

where $\sigma_k^2 = \prod_{j=1}^{i}(1-\alpha_j^2)$ and the coefficients $\phi_j^{(k)}$ are given recursively by

$$\phi_j^{(k)} = \begin{cases} \phi_j^{(k-1)} - \alpha_k \phi_{k-j}^{(k-1)}, & j \in \{1,\dots,k-1\}, \\ \alpha_k, & j = k. \end{cases}$$

The inverse Rosenblatt functions can be calculated to be

$$S_k(\boldsymbol{z}, x) = \Phi\left(\sigma_k \Phi^{-1}(x) + \sum_{j=1}^{k} \psi_j^{(k)} \Phi^{-1}(z_{k+1-j})\right),$$

where the coefficients $\psi_j^{(k)}$ are given recursively by $\psi_j^{(k)} = \sum_{i=1}^{j} \phi_i^{(k)} \psi_{j-i}^{(k-i)}$ for $j \in \{1,\dots,k\}$ where $\psi_0^{(k)} = \sigma_k$ for $k \geq 1$ and $\psi_0^{(0)} = 1$.

Causality and invertibility of Gaussian processes

Theorem

Let $(U_t)_{t\in\mathbb{Z}}$ be a Gaussian s-vine copula process with absolutely summable copula parameters $(\alpha_k)_{k\in\mathbb{N}}$. Then, almost surely, for all t,

$$U_t = \lim_{k\to\infty} S_k((Z_{t-k},\ldots,Z_{t-1})^\top, Z_t)$$
$$Z_t = \lim_{k\to\infty} R_k((U_{t-k},\ldots,U_{t-1})^\top, U_t)$$

for an iid uniform innovation process $(Z_t)_{t\in\mathbb{Z}}$.

- Proof is adaptation of result in Debowski (2007).
- Set $X_t = \Phi^{-1}(U_t)$ and $\epsilon_t = \Phi^{-1}(Z_t)$. Result reduces to familiar

$$X_t = \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j}, \quad \epsilon_t = \sum_{j=0}^{\infty} \phi_j X_{t-j}, \quad \psi_j = \lim_{k\to\infty} \psi_j^{(k)}, \quad \phi_j = \lim_{k\to\infty} \phi_j^{(k)}.$$

- Open issue: generalize to include mixing processes without absolutely summable $(\alpha_k)_{k\in\mathbb{N}}$, including some long-memory models.

Overview

Non-Gaussian copula sequences

- What are the conditions on a general copula sequence $(C_k)_{k\in\mathbb{N}}$ that enable us to construct processes $(U_t)_{t\in\mathbb{Z}}$ from uniform innovations $(Z_t)_{t\in\mathbb{Z}}$ such that we have convergent causal and invertible expressions

$$U_t = \lim_{k\to\infty} S_k((Z_{t-k},\dots,Z_{t-1})^\top, Z_t),$$
$$Z_t = \lim_{k\to\infty} R_k((U_{t-k},\dots,U_{t-1})^\top, U_t) \qquad ?$$

- It seems clear that $C_k \to C^\perp$ as $k \to \infty$.
- However, this is not sufficient (even in Gaussian case) and the speed of convergence is also important.
- Ideally we require conditions such that $C_k \to C^\perp$ also implies $C^{(k)} \to C^\perp$ (independence of U_t and U_{t-k} in the limit, mixing).
- This is more of a theoretical than practical issue as we can also view the models we simulate and fit as finite-order processes of very high order.

Parameterization via Kendall pacf

- Given parametric pair copulas $(C_k)_{k\in\mathbb{N}}$ we would like that:
  1. the copulas converge uniformly to the independence copula as $k \to \infty$;
  2. the level of dependence of each copula C_k is identical to ergodic Gaussian processes.
- To translate to non-Gaussian copulas, we use the Kendall partial autocorrelation function (kpacf) $(\tau_k)_{k\in\mathbb{N}}$ asssociated with a copula sequence $(C_k)_{k\in\mathbb{N}}$, given by

$$\tau_k = \tau(C_k),\ k \in \mathbb{N}.$$

- For a Gaussian copula sequence with $C_k = C^{\text{Ga}}_{\alpha_k}$ we have

$$\tau_k = \frac{2}{\pi}\arcsin(\alpha_k). \tag{3}$$

- For each pacf $(\alpha_k(\boldsymbol{\theta}))_{k\in\mathbb{N}}$ there is an implied kpacf $(\tau_k(\boldsymbol{\theta}))_{k\in\mathbb{N}}$. Idea: choose non-Gaussian pair copulas with this kpacf.
- Copulas should have Kendall correlation in the entire $(-1, 1)$. Otherwise rotations or replacements are needed.

Examples of copula filters

- Speed of convergence of different copula filters can be explored numerically, for the same kpacf.
- For fixed n and for a fixed realization $z_1, \ldots, z_n$ of independent uniform noise we plot the points $(k, S_k(\boldsymbol{z}_{[n-k,n-1]}, z_n))$ for $k \in \{1, \ldots, n-1\}$. We expect the points to converge to a fixed value as $k \to n-1$, provided we take a sufficiently large value of n.
- Non-Gaussian ARMA(1,1), parameters: 0.95, -0.85, $n = 201$.

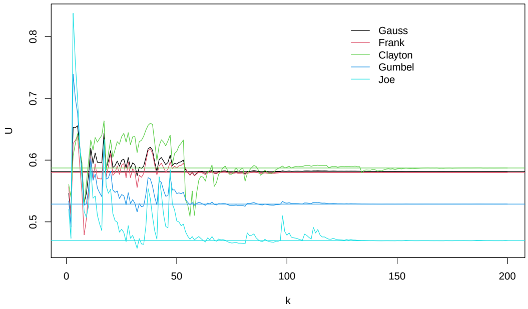

Examples of copula filters

- Speed of convergence of different copula filters can be explored numerically, for the same kpacf.
- For fixed n and for a fixed realization $z_1, \ldots, z_n$ of independent uniform noise we plot the points $(k, S_k(\boldsymbol{z}_{[n-k,n-1]}, z_n))$ for $k \in \{1, \ldots, n-1\}$. We expect the points to converge to a fixed value as $k \to n-1$, provided we take a sufficiently large value of n.
- Non-Gaussian ARFIMA(1,d,1) models, parameters: 0.95, -0.85, $d = 0.02$, and $n = 701$. In particular $|\alpha_k| \sim 0.02/k$ (not summable!)

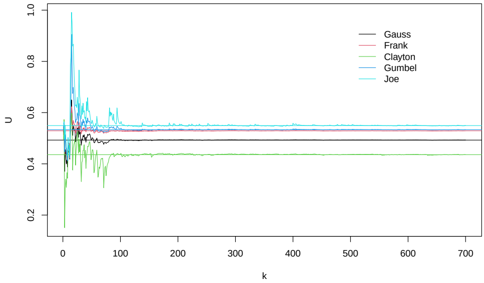

Overview

Real data

- We have written an R library to fit S-vines via their kapcf.
- Package `tscopula` (in particular using `rvinecopulib`).
- We open HTML for two examples.

Overview

Conclusions

- Our models give an interpretation to the idea of non-Gaussian ARMA, ARFIMA, etc. The process is named after the Gaussian process with which it shares a kpacf.
- We can generalize the idea of model residuals. We reconstruct the unobserved innovations $(Z_t)_{1 \leq t \leq n}$ using the invertibility formula for the fitted model.
- There is a need for parsimonious comprehensive bivariate copula families to give more options in fitting.
- Questions remain concerning the convergence of infinite copula filters.
- There are also statistical issues to resolve, such as consistency and asymptotic normality of parameter estimates in the pseudo-ML and full-ML estimation methods.

For further reading

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

Bedford, Tim, and Roger M. Cooke, 2002, Vines–a new graphical model for dependent random variables, *Annals of Statistics* 30, 1031–1068.

Brechmann, Eike Christian, and Claudia Czado, 2015, Copar - multivariate time series modeling using the copula autoregressive model, *Applied Stochastic Models in Business and Industry* 31, 495–514.

Chen, X., and Y. Fan, 2006, Estimation of copula-based semiparametric time series models, *Journal of Econometrics* 130, 307–335.

Cornfeld, I., S. Fomin, and G. Sinai, 1982, *Ergodic Theory* (Springer-Verlag).

Debowski, L., 2007, On processes with summable partial autocorrelations, *Statistics and Probability Letters* 77, 752–759.

Ibragimov, R., 2009, Copula-based characterizations for higher-order Markov processes, *Econometric Theory* 25.

Joe, H., 2015, *Dependence Modeling with Copulas* (CRC Press, Boca Raton).

For further reading (cont.)

Joe, Harry, 1997, *Multivariate Models and Dependence Concepts* (Chapman & Hall, London).

Maruyama, G., 1970, Infinitely divisible processes, *Theory of Probability and its Applications* 15, 1–22.

Meyn, S., and R. L. Tweedie, 2009, *Markov chains and stochastic stability*, second edition (Cambridge University Press).

Nagler, T., D. Krüger, and A. Min, 2020, Stationary vine copula models for multivariate series, Preprint.

Smith, Michael, Aleksey Min, Carlos Almeida, and Claudia Czado, 2010, Modeling Longitudinal Data Using a Pair-Copula Decomposition of Serial Dependence, *Journal of the American Statistical Association* 105, 1467–1479.