Time series with infinite-order partial copula dependence

Martin Bladt

Faculty of Business and Economics, University of Lausanne

Alexander J. McNeil

The York Management School, University of York

1 April 2022

Abstract

Stationary and ergodic time series can be constructed using an s-vine decomposition based on sets of bivariate copula functions. The extension of such processes to infinite copula sequences is considered and shown to yield a rich class of models that generalizes Gaussian ARMA and ARFIMA processes to allow both non-Gaussian marginal behaviour and a non-Gaussian description of the serial partial dependence structure. Extensions of classical causal and invertible representations of linear processes to general s-vine processes are proposed and investigated. A practical and parsimonious method for parameterizing s-vine processes using the Kendall partial autocorrelation function is developed. The potential of the resulting models to give improved statistical fits in many applications is indicated with an example using macroeconomic data.

AMS Subject Classification: 62M10; 62M05; 62H05; 60G10; 60G15; 60G22 Keywords: Time series; vine copulas; Gaussian processes; ARMA processes; ARFIMA processes.

1 Introduction

The principal aim of this paper is to show that the s-vine (or stationary d-vine) decomposition of a joint density provide a very natural vehicle for generalizing the class of stationary Gaussian time series to permit both non-Gaussian marginal behaviour and non-linear and non-Gaussian serial dependence behaviour. In particular, this approach provides a route to defining a rich class of tractable non-Gaussian ARMA and ARFIMA processes; the resulting models have the potential to offer improved statistical fits in any application where classical ARMA models or their long-memory ARFIMA extensions are used.

Vine models of dependence have been developed in a series of publications [24, 25, 6, 7, 8, 28, 1, 42]. There are a number of different configurations for vines but the most suitable one for longitudinal data applications is the d-vine, which is able to describe strict stationarity of a random vector under some additional translation-invariance restrictions on the vine structure. A recent paper by Nagler et al. [36] investigates the vine structures that can be used to construct stationary multivariate time series. The results of Nagler et al. imply that, for univariate applications, the d-vine is in fact the only structure for which translation-invariance restrictions are sufficient to guarantee stationarity; we follow them in referring to these restricted d-vines as stationary vines, or s-vines.

Vine models are best understood as copula models of dependence and there is now a large literature on copula models for time series. While the main focus of much of this literature has been on cross-sectional dependence between multiple time series, there is also a growing literature on modelling serial dependence within single series and lagged dependence across series. Firstorder Markov copula models [15, 12, 3, 17] are simple examples of s-vine processes. A number of authors have written on higher-order Markov extensions for univariate series or multivariate series [22, 43, 4, 10, 30, 36]. There is also literature showing how these models may be adapted to the particular requirements of time series showing stochastic volatility, including the mixture-copula approach of Louaiza-Maya et al. [30] and the v-transform approach of McNeil and Bladt [33, 9].

This paper makes the following novel contributions to the development of time series models based on vine copulas. First, we suggest how s-vine models may be generalized to infinite order and we propose accompanying generalizations of the classical concepts of causality and invertibility for linear processes that may be applied to s-vine processes. Second, we provide additional insight into the issues of stability and ergodicity for s-vine processes and we show how finite or infinite copula sequences may be used to develop non-linear filters of independent noise that generalize linear filters. Finally, we propose a practical and parsimonious approach to building s-vine processes in which copula sequences are parameterized by a function that we call the Kendall partial autocorrelation function; the latter may be borrowed from other well-known processes, such as Gaussian ARMA or ARFIMA processes, thus yielding natural non-Gaussian analogues of these models.

We believe that our approach may serve as a useful framework to faciliate further study in

the field. Several interesting theoretical questions remain, particularly relating to necessary and sufficient conditions for stability of models based on infinite copula sequences, as well as the interplay of copula sequences and long memory. However, on the practical side, the models are already eminently usable; methods exist for estimation and random number generation, and we suggest some new ideas for model validation using residuals. An example shows the benefits that may arise from using these models.

The paper is structured as follows. Section 2 sets out notation and basic concepts and makes the connection between s-vine copulas and s-vine processes; key objects in the development of processes are sequences of functions that we refer to as Rosenblatt functions. In Section 3 we show that finite-order s-vine processes are Markov chains belonging to the particular sub-category of non-linear state-space models. Section 4 explains why Gaussian processes form a sub-class of s-vine processes and shows how classical theory for linear processes may be reinterpreted as a theory of the behaviour of Rosenblatt functions. Section 5 uses the Gaussian analogy to suggest requirements for stable, infinite-order, non-Gaussian s-vine processes; a practical approach to model building is developed and illustrated with an application to macroeconomic data. Section 6 concludes. Proofs can be found in Appendix A while additional material on the Markov chain analysis of finite-order processes is collected in Appendix B.

2 S-vine processes

2.1 S-vine copulas

If a random vector (X_1, \ldots, X_n) admits a joint density $f(x_1, \ldots, x_n)$ then the latter may be decomposed as a d-vine. Writing f_{X_i} for the marginal density of X_i , the decomposition is

$$f(x_1, \dots, x_n) = \left(\prod_{i=1}^n f_{X_i}(x_i)\right) \prod_{k=1}^{n-1} \prod_{j=k+1}^n c_{j-k,j|S_{j-k,j}}(F_{j-k|S_{j-k,j}}(x_{j-k}), F_{j|S_{j-k,j}}(x_j))$$
(1)

where $S_{j-k,j} = \{j-k+1, \dots, j-1\}$ is the set of indices of the variables which lie between X_{j-k} and X_j , $c_{j-k,j|S_{j-k,j}}$ is the density of the bivariate copula $C_{j-k,j|S_{j-k,j}}$ of the joint distribution function (df) of X_{j-k} and X_j conditional on the intermediate variables $X_{j-k+1}, \dots, X_{j-1}$, and

$$F_{i|S_{j-k,j}}(x) = \mathbb{P}(X_i \leqslant x | X_{j-k+1} = x_{j-k+1}, \dots, X_{j-1} = x_{j-1}), \quad i \in \{j-k, j\}$$
(2)

denotes the conditional df of variable *i* conditional on these variables; note that $S_{j-1,j} = \emptyset$ and so the conditioning set is dropped in this case. The decomposition (1) implies a decomposition of the density $c(u_1, \ldots, u_n)$ of the unique copula of (X_1, \ldots, X_n) which is given implicitly by

$$c(F_1(x_1),\ldots,F_n(x_n)) = \prod_{k=1}^{n-1} \prod_{j=k+1}^n c_{j-k,j|S_{j-k,j}}(F_{j-k|S_{j-k,j}}(x_{j-k}),F_{j|S_{j-k,j}}(x_j)).$$
(3)

In practical modelling applications interest centres on models which admit the simplified d-vine decomposition in which the copula densities $c_{j-k,j|S_{j-k,j}}$ do not depend on the values of variables in the conditioning set $S_{j-k,j}$ and we can simply write $c_{j-k,j}$. Any set of copula densities $\{c_{j-k,j}: 1 \leq k \leq n-1, k+1 \leq j \leq n\}$ and any set of marginal densities f_{X_i} may be used in the simplified version of (1) to create a valid *n*-dimensional joint density. A number of papers have examined the limitations imposed by working with simplified vine copula models [20, 45, 44, 35]. In the last of these papers it is shown that the class of simplified vines is not dense in the space of copulas for a number of metrics including the one induced by total variation distance. These results may be interpreted as showing that there exist multivariate distributions that are difficult to approximate with simplified d-vines. However, the simplified d-vine construction still greatly enlarges the class of tractable densities for time series applications.

We are interested in strictly stationary stochastic processes whose higher-dimensional marginal distributions are simplified d-vines. As well as forcing $f_{X_1} = \cdots = f_{X_n}$, this requirement imposes translation-invariance conditions on the copula densities $c_{j-k,j}$ and conditional dfs $F_{\cdot|S_{j-k,j}}$ appearing in the simplified form of (1). It must be the case that $c_{j-k,j}$ is the same for all $j \in \{k+1,\ldots,n\}$ and so each pair copula density in the model can be associated with a lag k and we can write $c_k := c_{j-k,j}$ where c_k is the density of some bivariate copula C_k . The conditional dfs can be represented by two sets of functions $R_k^{(1)} : (0,1)^k \times (0,1) \to (0,1)$ and $R_k^{(2)} : (0,1)^k \times (0,1) \to (0,1)$ which are 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 $k \ge 2$,

$$R_{k}^{(1)}(\boldsymbol{u},x) = h_{k}^{(1)} \left(R_{k-1}^{(2)}(\boldsymbol{u}_{-1},u_{1}), R_{k-1}^{(1)}(\boldsymbol{u}_{-1},x) \right)$$

$$R_{k}^{(2)}(\boldsymbol{u},x) = h_{k}^{(2)} \left(R_{k-1}^{(2)}(\boldsymbol{u}_{-k},x), R_{k-1}^{(1)}(\boldsymbol{u}_{-k},u_{k}) \right)$$
(4)

where $h_k^{(i)}(u_1, u_2) = \frac{\partial}{\partial u_i} C_k(u_1, u_2)$ and u_{-i} indicates the vector u with *i*th component removed. Using this new notation, we obtain a simplified form of (1) in which the density of the copula c

4

in (3) takes the form

$$c_{(n)}(u_1,\ldots,u_n) = \prod_{k=1}^{n-1} \prod_{j=k+1}^n c_k \Big(R_{k-1}^{(2)}(\boldsymbol{u}_{[j-k+1,j-1]}, u_{j-k}), R_{k-1}^{(1)}(\boldsymbol{u}_{[j-k+1,j-1]}, u_j) \Big)$$
(5)

where $\boldsymbol{u}_{[j-k+1,j-1]} = (u_{j-k+1}, \dots, u_{j-1})^{\top}$. Note that, for simplicity of formulas, we abuse notation by including terms involving $R_0^{(1)}$ and $R_0^{(2)}$; these terms should be interpreted as $R_0^{(1)}(\cdot, u) = R_0^{(2)}(\cdot, u) = u$ for all u. Following Nagler et al. [36] we refer to a model with copula density of the form (5) as a stationary d-vine or *s-vine*.

If a random vector (U_1, \ldots, U_n) follows the copula $C_{(n)}$ with density $c_{(n)}$ in (5) then for any $k \in \{1, \ldots, n-1\}$ and $j \in \{k+1, \ldots, n\}$, we have

$$R_k^{(1)}(\boldsymbol{u}, x) = \mathbb{P}(U_j \leq x \mid U_{j-k} = u_1, \dots, U_{j-1} = u_k)$$

$$R_k^{(2)}(\boldsymbol{u}, x) = \mathbb{P}(U_{j-k} \leq x \mid U_{j-k+1} = u_1, \dots, U_j = u_k).$$
(6)

and we refer to the conditional distribution functions $R_k^{(1)}$ and $R_k^{(2)}$ as forward and backward Rosenblatt functions. Henceforth we will often drop the superscript from the forward function and simply write $R_k = R_k^{(1)}$ to obtain less notationally cumbersome expressions. The conditional densities corresponding to the Rosenblatt functions may be derived from (5). Writing f_k for the density of the forward Rosenblatt functions we obtain $f_1(u, x) = c_{(2)}(u, x) = c_1(u, x)$ and, for k > 1

$$f_k(\boldsymbol{u}, x) = \frac{c_{(k+1)}(u_1, \dots, u_k, x)}{c_{(k)}(u_1, \dots, u_k)} = \prod_{j=1}^k c_j \left(R_{j-1}^{(2)}(\boldsymbol{u}_{[k-j+2,k]}, u_{k-j+1}), R_{j-1}(\boldsymbol{u}_{[k-j+2,k]}, x) \right).$$
(7)

The following assumption will be in force throughout the remainder of the paper.

Assumption 1. All copulas C_k used in the construction of s-vine models belong to the class C^{∞} of smooth functions with continuous partial derivatives of all orders. Moreover their densities c_k are strictly positive on $(0, 1)^2$.

This assumption applies to all the standard pair copulas that are used in vine copula models (e.g. Gauss, Clayton, Gumbel, Frank, Joe and t), as well as non-exchangeable extensions [29] or mixtures of copulas [30]. It ensures, among other things, that for fixed \boldsymbol{u} , the Rosenblatt functions are bijections on (0, 1) with well-defined inverses. Let us write $R_k^{-1}(\boldsymbol{u}, z)$ for the *inverses of the Rosenblatt functions*, satisfying $R_k^{-1}(\boldsymbol{u}, z) = x$ if and only if $R_k(\boldsymbol{u}, x) = z$. Inverses can also be defined for the Rosenblatt backward functions but will not be explicitly needed

In the sequel we refer to the copulas C_k as partial copulas. They should be distinguished from the bivariate marginal copulas given by $C^{(k)}(u, v) = \mathbb{P}(U_{j-k} \leq u, U_j \leq v)$ for any $j \in \{k+1, \ldots, n\}$. The two copulas are related by the formula

$$C^{(k)}(v_{1}, v_{2}) = \mathbb{E}\left(\mathbb{P}(U_{j-k} \leqslant v_{1}, U_{j} \leqslant v_{2} \mid U_{j-k+1}, \dots, U_{j-1})\right)$$

= $\mathbb{E}\left(C_{k}\left(R_{k-1}^{(2)}((U_{j-k+1}, \dots, U_{j-1})^{\top}, v_{1}), R_{k-1}((U_{j-k+1}, \dots, U_{j-1})^{\top}, v_{2})\right)\right)$
= $\int_{0}^{1} \cdots \int_{0}^{1} C_{k}\left(R_{k-1}^{(2)}(\boldsymbol{u}, v_{1}), R_{k-1}(\boldsymbol{u}, v_{2})\right) c_{(k-1)}(\boldsymbol{u}) \mathrm{d} u_{1} \cdots \mathrm{d} u_{k-1}.$ (8)

2.2 S-vine processes

We use the following general definition for an s-vine process.

Definition 1 (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 $n \ge 2$ 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 (5). An s-vine process $(U_t)_{t\in\mathbb{Z}}$ is an *s-vine copula process* if its univariate marginal distribution is standard uniform.

Our aim is to construct processes that conform to this definition and investigate their properties and practical application. Since s-vine processes can be endowed with any continuous univariate marginal distribution f_X , we will mostly investigate the properties of s-vine copula processes.

2.3 A note on reversibility

It is particularly common in applications of vine copulas to confine interest to standard exchangeable copulas C_k . In this case the resulting s-vine processes have the property of *reversibility*. For any $\boldsymbol{u} = (u_1, \ldots, u_n)^\top \in (0, 1)^n$ let us write $\overline{\boldsymbol{u}} = (u_n, \ldots, u_1)^\top$ for the reversed vector.

Definition 2. An s-vine copula process is reversible if for any $n \ge 2$ the higher dimensional marginal copulas satisfy $C_{(n)}(\boldsymbol{u}) = C_{(n)}(\overline{\boldsymbol{u}})$.

This is equivalent to saying that, for any $t, s \in \mathbb{Z}$ and any n > 2, the set of consecutive variables $(U_{t+1}, \ldots, U_{t+n})$ from the process has the same distribution as the reversed vector $(U_{s+n}, \ldots, U_{s+1})$. The process evolves forwards and backwards in a similar fashion, which may not be ideal for phenomena in which there is a clear temporal notion of causality; however, as soon as non-exchangeable copulas are included, the reversibility is broken. In summary we have following simple result. **Proposition 1.** If a copula sequence $(C_k)_{k\in\mathbb{N}}$ consists of exchangeable copulas then (i) the Rosenblatt forward and backward functions satisfy $R_k^{(2)}(\overline{\boldsymbol{u}}, x) = R_k(\boldsymbol{u}, x)$ for all $(\boldsymbol{u}, x) \in (0, 1)^k \times (0, 1)$ and (ii) the resulting s-vine copula process is reversible.

3 S-vine processes of finite order

3.1 Markov construction

The first class of processes we consider are s-vine copula processes of finite order p which are constructed from a set of copulas $\{C_1, \ldots, C_p\}$ using the Markov approach described by Joe [27, page 145]. Starting from a series of iid uniform innovation variables $(Z_k)_{k\in\mathbb{N}}$ we can set $U_1 = Z_1$ and

$$U_{k} = R_{k-1}^{-1} \Big((U_{1}, \dots, U_{k-1})^{\top}, Z_{k} \Big), \quad k \ge 2.$$
(9)

By using the inverses of the Rosenblatt forward functions we obtain, for any n, a random vector (U_1, \ldots, U_n) which forms a finite realization from an s-vine process $(U_t)_{t \in \mathbb{Z}}$. The copula $C_{(n)}$ of (U_1, \ldots, U_n) has density $c_{(n)}$ in (5) but the copula densities c_k appearing in this expression satisfy $c_k(u, v) = 1$ for k > p and the s-vine is said to be truncated at order p. Moreover, since $h_k^{(1)}(u, v) = v$ for k > p it follows from (4) that $R_k(u, x) = R_{k-1}(u_{-1}, x) = \cdots = R_p(u_{[k-p+1,k]}, x)$ and the updating equation (9) satisfies

$$U_{k} = R_{p}^{-1} \Big((U_{k-p}, \dots, U_{k-1})^{\top}, Z_{k} \Big), \quad k > p,$$
(10)

showing the Markovian character of the finite-order process.

The recursive nature of the construction (9) means that there is an implied set of functions that we will label $S_k : (0,1)^k \times (0,1) \to (0,1)$ for $k \in \mathbb{N}$ such that

$$U_k = S_{k-1}((Z_1, \dots, Z_{k-1})^\top, Z_k), \quad k \ge 2.$$
(11)

The functions $(S_k)_{k\in\mathbb{N}}$ satisfy $S_1(z_1, x) = R_1^{-1}(z_1, x)$ and

$$S_k(\boldsymbol{z}, x) = R_k^{-1} \Big(\big(z_1, S_1(z_1, z_2), \dots, S_{k-1}(\boldsymbol{z}_{[1,k-1]}, z_k) \big), x \Big), \quad k \ge 2.$$
(12)

The identity (11) can be thought of as a *causal* representation of the process while the complementary identity $Z_k = R_{k-1}((U_1, \ldots, U_{k-1})^{\top}, U_k)$ implied by (9) can be thought of as an *invertible* representation. We refer to the functions $(S_k)_{k \in \mathbb{N}}$ as Rosenblatt inverse functions; they should be distinguished from the inverses of the Rosenblatt forward functions

3.2 Non-linear state space model

The s-vine process of order p can be viewed as a p-dimensional Markov chain with state space $\mathcal{X} = (0,1)^p$. It is standard to treat Markov chains as being indexed by the natural numbers. To that end, for $t \in \mathbb{N}$, we introduce the vector-valued process $U_t = (U_t, \ldots, U_{t+p-1})^\top$, starting at $U_1 = (U_1, \ldots, U_p)^\top$, defined by the updating equation $U_t = F(U_{t-1}, Z_t)$ where

$$F: (0,1)^p \times (0,1) \to (0,1)^p, \quad F(\boldsymbol{u},z) = \left(u_2, \dots, u_p, R_p^{-1}(\boldsymbol{u},z)\right).$$
(13)

The Markov chain described by (13) defines a non-linear state space (NSS) model conforming exactly to the assumptions imposed in Meyn and Tweedie [34, Section 2.2.2]: under Assumption 1 the updating function F is a smooth (\mathcal{C}^{∞}) function; the state space $\mathcal{X} = (0, 1)^p$ is an open subset of \mathbb{R}^p ; the uniform distribution of innovations (Z_t) will be taken to be supported on the open set (0, 1).

Using standard arguments, the NSS model associated to (13) can be shown to be a ϕ -irreducible, aperiodic Harris recurrent Markov chain and to admit an invariant probability measure π which is the measure implied by the density $c_{(p)}$ given by (5); we summarise the arguments in Appendix B. This in turn allows the ergodic theorem for Harris chains to be applied [34, Theorem 13.3.3] to conclude that for any initial measure λ the Markov transition kernel $\mathsf{P}(\mathbf{x}, \cdot)$ satisfies

$$\left\|\int \lambda(\mathrm{d}\boldsymbol{x})\mathsf{P}^{n}(\boldsymbol{x},\cdot)-\pi(\cdot)\right\|\to 0, \quad n\to\infty$$

where $\|\cdot\|$ denotes the total variation norm. This is also sufficient for the strong law of large numbers (SLLN) to hold [34, Theorem 17.0.1]: for a function $g : \mathbb{R}^p \to \mathbb{R}$, if we define $S_n(g) = \sum_{k=1}^n g(U_k)$ and $\pi(g) = \int g(u)c_{(p)}(u) du$, then $\lim_{n\to\infty} n^{-1}S_n(g) = \pi(g)$, almost surely, provided $\pi(|g|) < \infty$.

Although the Markov models are ergodic, we caution that they can exhibit some very extreme behaviour, albeit for copula choices that we are unlikely to encounter in practice. Figure 1 shows a realisation of 10000 simulated values from a process of order p = 3 in which C_1 is a 180-degree rotated Clayton copula with parameter $\theta = 2$, C_2 is a Clayton copula with $\theta = 2$ and C_3 is a rotated Clayton copula with $\theta = 4$. Since the Clayton copula is well known to have lower tail dependence [25, 27], this means that C_1 and C_3 have upper tail dependence and C_3 is more strongly dependent than C_1 and C_2 . This increasing pattern of partial dependence, coupled with the strong upper tail dependence of C_3 , leads to a period of over 1500 successive values which are all greater than 0.6. An observer of this process who plots a histogram of the values in this period would have difficulty believing that the marginal distribution is uniform.



Figure 1: Realisation of 10000 simulated values from a process of order k = 3 in which C_1 is a 180-degree rotated Clayton copula with parameter $\theta = 2$, C_2 is a Clayton copula with $\theta = 2$ and C_2 is a rotated Clayton copula with $\theta = 4$.

This phenomenon is connected to rates of mixing behaviour and ergodic convergence for Markov processes. There is some literature for the case p = 1 in which these rates are shown to vary with the choice of copula and, in particular, its behaviour in joint tail regions [12, 13, 3, 5, 31]. For some results relevant to the case where p > 1, see Rémillard et al. [39].

4 Gaussian processes

Gaussian processes are processes whose finite-dimensional marginal distributions are multivariate Gaussian. We will identify the term Gaussian processes with *non-singular* Gaussian processes throughout; i.e. we assume that the finite-dimensional marginal distributions of Gaussian processes have invertible covariance matrices and admit joint densities. Such processes represent a subclass of the s-vine processes.

Proposition 2. 1. Every stationary Gaussian process is an s-vine process.

 Every s-vine process in which the pair copulas of the sequence (C_k)_{k∈N} are Gaussian and the marginal distribution F_X is Gaussian, is a Gaussian process.

4.1 S-vine representations of Gaussian processes

The first implication of Proposition 2 is that every Gaussian process has a unique s-vine-copula representation. This insight offers methods for constructing or simulating such processes as generic s-vine processes using (9) and estimating them using a likelihood based on (5).

Let $(X_t)_{t\in\mathbb{N}}$ be a stationary Gaussian process with mean μ_X , variance σ_X^2 and autocorrelation function (acf) $(\rho_k)_{k\in\mathbb{N}}$; these three quantities uniquely determine a Gaussian process. We assume the following:

Assumption 2. The acf $(\rho_k)_{k \in \mathbb{N}}$ satisfies $\rho_k \to 0$ as $k \to \infty$.

It is well known that this is a necessary and sufficient condition for a Gaussian process (X_t) to be a mixing process and therefore ergodic [32, 14].

The acf uniquely determines the partial autocorrelation function (pacf) $(\alpha_k)_{k\in\mathbb{N}}$ through a oneto-one transformation [2, 38]. Since the partial autocorrelation of a Gaussian process is the correlation of the conditional distribution of (X_{t-k}, X_t) given the intervening variables, the pair copulas in the s-vine copula representation are given by $C_k = C_{\alpha_k}^{\text{Ga}}$.

For $k \in \mathbb{N}$ let $\rho_k = (\rho_1, \dots, \rho_k)^{\top}$ and let P_k denote the correlation matrix of (X_1, \dots, X_k) . Clearly $P_1 = 1$ and, for k > 1, P_k is a symmetric Toeplitz matrix whose diagonals are filled by the first k-1 elements of ρ_k ; moreover, P_k is non-singular for all k under Assumption 2 [11, Proposition 4]. The one-to-one series of recursive transformations relating $(\alpha_k)_{k \in \mathbb{N}}$ to $(\rho_k)_{k \in \mathbb{N}}$ is $\alpha_1 = \rho_1$ and, for k > 1,

$$\alpha_{k} = \frac{\rho_{k} - \boldsymbol{\rho}_{k-1}^{\top} P_{k-1}^{-1} \overline{\boldsymbol{\rho}}_{k-1}}{1 - \boldsymbol{\rho}_{k-1}^{\top} P_{k-1}^{-1} \rho_{k-1}}, \quad \rho_{k} = \alpha_{k} \left(1 - \boldsymbol{\rho}_{k-1}^{\top} P_{k-1}^{-1} \rho_{k-1} \right) + \boldsymbol{\rho}_{k-1}^{\top} P_{k-1}^{-1} \overline{\boldsymbol{\rho}}_{k-1}; \quad (14)$$

see, for example, Joe [26] or the Durbin-Levinson Algorithm [11, Proposition 5.2.1].

Remark 1. Note that the restriction to non-singular Gaussian processes ensures that $|\rho_k| < 1$ and $|\alpha_k| < 1$, for all $k \in \mathbb{N}$, and this is henceforth always assumed.

We review three examples of well-known Gaussian processes from the point of view of s-vine processes.

Example 1 (Gaussian ARMA models). Any causal Gaussian ARMA(p,q) model may be represented as an s-vine process and full maximum likelihood estimation can be carried out using a joint density based on (5). If $\boldsymbol{\phi} = (\phi_1, \dots, \psi_p)^{\top}$ and $\boldsymbol{\psi} = (\psi_1, \dots, \psi_q)^{\top}$ denote the AR and MA parameters and $\rho_k(\phi, \psi)$ the acf, then we can use the transformation (14) to parameterize (5) in terms of ϕ and ψ using Gaussian pair copulas $C_k = C_{\alpha_k(\phi,\psi)}^{\text{Ga}}$. In practice, this approach is more of theoretical interest since standard estimation methods are generally much faster.

Example 2 (Fractional Gaussian noise (FGN)). This process has acf given by

$$\rho_k(H) = \frac{1}{2} \left((k+1)^{2H} + (k-1)^{2H} - 2k^{2H} \right), \quad 0 < H < 1,$$

where H is the Hurst exponent [41]. Thus the transformation (14) may be used to parameterize (5) in terms of H using Gaussian pair copulas $C_k = C_{\alpha_k(H)}^{\text{Ga}}$ and the FGN model may be fitted to data as an s-vine process and H may be estimated.

Example 3 (Gaussian ARFIMA models). The ARFIMA(p,d,q) model with -1/2 < d < 1/2 can be handled in a similar way to the ARMA(p,q) model, of which it is a generalization. In the case where p = q = 0, it has been shown [21] that

$$\alpha_k = \frac{d}{k-d}, \quad k \in \mathbb{N}; \tag{15}$$

see also Brockwell and Davis [11, Theorem 13.2.1]. The simple closed-form expression for the pacf means that the ARFIMA(0,d,0) model is even more convenient to treat as an s-vine than FGN; the two models are in fact very similar in behaviour although not identical. It is interesting to note that the pacf is not summable and similar behaviour holds for some other ARFIMA processes. For example, for $p, q \in \mathbb{N} \cup \{0\}$ and 0 < d < 1/2 the pacf satisfies $|\alpha_k| \sim d/k$ as $k \to \infty$ [23].

4.2 New Gaussian processes from s-vines

A further implication of Proposition 2 is that it shows how we can create and estimate some new stationary and ergodic Gaussian processes without setting them up in the classical way using recurrence equations, lag operators and Gaussian innovations. Instead we choose sequences of Gaussian pair copulas (C_k) parameterized by sequencies of partial correlations (α_k).

As in the previous section, we can begin with a parametric form for the acf $\rho_k(\boldsymbol{\theta})$ such that $\rho_k(\boldsymbol{\theta}) \to 0$ as $k \to \infty$ and build the model using pair copulas parameterized by the parameters $\boldsymbol{\theta}$ of the implied pacf $\alpha_k(\boldsymbol{\theta})$. Alternatively we can choose a parametric form for the pacf $\alpha_k(\boldsymbol{\theta})$ directly.

Any finite set of values $\{\alpha_1, \ldots, \alpha_p\}$ yields an AR(p) model which is a special case of the finiteorder s-vine models of Section 3. However, infinite-order processes that satisfy Assumption 2 are more delicate to specify. A necessary condition is that the sequence (α_k) satisfies $\alpha_k \to 0$ as $k \to 0$, but this is not sufficient. To see this, note that if $\alpha_k = (k+1)^{-1}$, the relationship (14) implies that $\rho_k = 0.5$ for all k which violates Assumption 2. A sufficient condition follows from a result of Debowski [16], although, in view of Example 3, it is not a necessary condition:

Assumption 3. The partial acf $(\alpha_k)_{k\in\mathbb{N}}$ satisfies $\sum_{k=1}^{\infty} |\alpha_k| < \infty$.

Debowski [16] showed that, if Assumption 3 holds, then the equality

$$1 + 2\sum_{k=1}^{\infty} \rho_k = \prod_{k=1}^{\infty} \frac{1 + \alpha_k}{1 - \alpha_k}.$$
 (16)

also holds. The rhs of (16) is a convergent product since absolute summability ensures that the sums $\sum_{k=1}^{\infty} \ln(1 \pm \alpha_k)$ converge. This implies the convergence of $\sum_{k=1}^{\infty} \rho_k$ which implies $\rho_k \to 0$ which in turn implies that Assumption 2 also holds, as we require.

Assumption 3 still allows some quite pathological processes, as noted by Debowski [16]. For example, even for a finite-order AR(p) process with $\alpha_k \ge a > 0$ for $k \in \{1, \ldots, p\}$ and $\alpha_k = 0$ for k > p it follows that $\sum_{k=1}^{\infty} \rho_k \ge 0.5(((1+a)/(1-a))^p - 1)$ and this grows exponentially with p leading to an exceptionally slow decay of the acf.

4.3 Rosenblatt functions for Gaussian processes

For Gaussian processes the Rosenblatt functions and inverse Rosenblatt functions take relatively tractable forms.

Proposition 3. Let $(C_k)_{k\in\mathbb{N}}$ be a sequence of Gaussian pair copulas with parameters $(\alpha_k)_{k\in\mathbb{N}}$ and assume that Assumption 2 holds. The forward Rosenblatt functions are given by

$$R_k(\boldsymbol{u}, x) = \Phi\left(\frac{\Phi^{-1}(x) - \sum_{j=1}^k \phi_j^{(k)} \Phi^{-1}(u_{k+1-j})}{\sigma_k}\right),\tag{17}$$

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

$$\phi_j^{(k)} = \begin{cases} \phi_j^{(k-1)} - \alpha_k \phi_{k-j}^{(k-1)}, & j \in \{1, \dots, k-1\}, \\ \alpha_k, & j = k. \end{cases}$$
(18)

The inverse Rosenblatt functions are given by

$$S_k(\boldsymbol{z}, x) = \Phi\left(\sigma_k \Phi^{-1}(x) + \sum_{j=1}^k \psi_j^{(k)} \Phi^{-1}(z_{k+1-j})\right),$$
(19)

where the coefficients $\psi_j^{(k)}$ are given recursively by

$$\psi_j^{(k)} = \sum_{i=1}^j \phi_i^{(k)} \psi_{j-i}^{(k-i)}, \quad j \in \{1, \dots, k\},$$
(20)

where $\psi_0^{(k)} = \sigma_k$ for $k \ge 1$ and $\psi_0^{(0)} = 1$.

We can analyse the behaviour of the Rosenblatt and inverse Rosenblatt functions as $k \to \infty$ in a number of different cases.

Gaussian processes of finite order. In the case of a Gaussian s-vine process of finite order p we have, for k > p, that $\alpha_k = 0$, $\sigma_k = \sigma_p$ and $\phi_j^{(k)} = \phi_j^{(p)}$. If $(U_k)_{k \in \mathbb{N}}$ is constructed from $(Z_k)_{k \in \mathbb{N}}$ using the algorithm described by (9), and if we make the substitutions $X_k = \Phi^{-1}(U_k)$ and $\epsilon_k = \Phi^{-1}(Z_k)$ as in the proof of Proposition 3, then it follows from (17) that $X_k = \sum_{j=1}^p \phi_j^{(p)} X_{k-j} + \sigma_p \epsilon_k$ for k > p, which is the classical recurrence equation that defines a Gaussian AR(p) process; from (11) and (19) we also have that $X_k = \sum_{j=1}^{k-1} \psi_j^{(k-1)} \epsilon_{k-j} + \sigma_p \epsilon_k$ for k > p. These two representations can be written in invertible and causal forms as

$$\epsilon_k = \sum_{j=0}^p \tilde{\phi}_j^{(p)} X_{k-j} \quad \text{and} \quad X_k = \sum_{j=0}^{k-1} \psi_j^{(k-1)} \epsilon_{k-j}, \quad k > p,$$
 (21)

where $\tilde{\phi}_{0}^{(p)} = 1/\sigma_{p}, \ \tilde{\phi}_{j}^{(p)} = -\phi_{j}^{(p)}/\sigma_{p} \text{ for } j > 1 \text{ and } \psi_{0}^{(k-1)} = \sigma_{p}.$

The first series in (21) is clearly a finite series while classical theory is concerned with conditions on the AR coefficients $\tilde{\phi}_j^{(p)}$ that allow us to pass to an infinite-order moving-average representation as $k \to \infty$ in the second series. In fact, by setting up our Gaussian models using partial autocorrelations, causality in the classical sense is guaranteed; this follows as a special case of Theorem 1 below.

Gaussian processes with absolutely summable partial autocorrelations. We next consider a more general case where the process may be of infinite order, but Assumption 3 holds. To consider infinite-order models we now consider a process $(U_t)_{t\in\mathbb{Z}}$ defined on the integers. The result that follows is effectively a restating of a result by Debowski [16] in the particular context of Gaussian s-vine copula processes.

Theorem 1. Let $(U_t)_{t\in\mathbb{Z}}$ be a Gaussian s-vine copula process for which the parameters $(\alpha_k)_{k\in\mathbb{N}}$ of the Gaussian pair copula sequence $(C_k)_{k\in\mathbb{N}}$ satisfy Assumption 3. Then, for all t, we have the almost sure limiting representations

$$U_{t} = \lim_{k \to \infty} S_{k}((Z_{t-k}, \dots, Z_{t-1})^{\top}, Z_{t})$$
(22)

$$Z_{t} = \lim_{k \to \infty} R_{k}((U_{t-k}, \dots, U_{t-1})^{\top}, U_{t})$$
(23)

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

Long-memory ARFIMA processes. As noted earlier, the pact of an ARFIMA(p,d,q) model with 0 < d < 0.5 is not absolutely summable [23] and so Theorem 1 does not apply in this case. Nevertheless, Brockwell and Davis [11, Section 13.2] show that the Gaussian process has a casual representation of the form $X_t = \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j}$ where convergence is now in mean square and the coefficients are square summable, i.e. $\sum_{j=0}^{\infty} \psi_j^2 < \infty$. Since convergence in mean square implies convergence in probability, the continuous mapping theorem implies that a representation of the form $U_t = \lim_{k\to\infty} S_k((Z_{t-k}, \dots, Z_{t-1})^\top, Z_t)$ at least holds under convergence in probability.

A non-causal and non-invertible case. If $\alpha_k = 1/(k+1)$ for all k, then $\rho_k = 0.5$ and both Assumptions 2 and 3 are violated. It can be verified (for example by induction) that the recursive formulas (18) and (20) imply that $\phi_j^{(k)} = 1/(k+1)$ and $\psi_j^{(k)} = \sigma_{k-j}/(k+2-j)$ for $j \ge 1$ (recall that $\psi_0^{(k)} = \sigma_k$). These coefficient sequences are unusual; the coefficients $\phi_j^{(k)}$ of the Rosenblatt function in (17) place equal weight on all past values $X_{k+1-j} = \Phi^{-1}(U_{k+1-j})$ while the coefficients $\psi_j^{(k)}$ of the inverse Rosenblatt function on the innovations in (19) place weight $\psi_k^{(k)} = 1/2$ on the first value $\epsilon_1 = \Phi^{-1}(Z_1)$ and decreasing weights on more recent values ϵ_j , j > 1.

As $k \to \infty$, we do have $\sigma_k^2 = \prod_{j=1}^k (1 - 1/(k+1)^2) \to 1/2$, but, for fixed $j \ge 1$, the terms $\psi_j^{(k)}$ and $\psi_j^{(k)}$ both converge to the trivial limiting value 0 and we do not obtain convergent limiting representations of the form (22) and (23) for Z_t or U_t in terms of past process values or past innovation values.

5 General s-vine processes

We now consider infinite-order s-vine copula processes constructed from general sequences $(C_k)_{k \in \mathbb{N}}$ of pair copulas.

5.1 Causality and invertibility

The key consideration for stability of an infinite-order process is whether it admits a convergent causal representation. A process $(U_t)_{t\in\mathbb{Z}}$ with such a representation is a convergent non-linear filter of independent noise. It will have the property that U_t and U_{t-k} are independent in the limit as $k \to \infty$, implying mixing behaviour and ergodicity. We suggest the following definition of the causality and invertibility properties for a general s-vine process.

Definition 3. Let $(C_k)_{k\in\mathbb{N}}$ be a sequence of pair copulas and let $(R_k)_{k\in\mathbb{N}}$ and $(S_k)_{k\in\mathbb{N}}$ be the corresponding Rosenblatt forward functions and Rosenblatt inverse functions defined by (4) and (12). An s-vine copula process $(U_t)_{t\in\mathbb{Z}}$ associated with the sequence $(C_k)_{k\in\mathbb{N}}$ is strongly causal if there exists a process of iid uniform random variables $(Z_t)_{t\in\mathbb{Z}}$ such that (22) holds almost surely for all t and it is strongly invertible if representation (23) holds almost surely for all t. If convergence in (22) and (23) only holds in probability, the process is weakly causal or weakly invertible.

We know that Gaussian ARMA processes defined as s-vine processes are always strongly causal (and invertible) and that the long-memory ARFIMA(p,d,q) process with 0 < d < 0.5 is weakly causal. When we consider sequences of Rosenblatt functions for sequences of non-Gaussian pair copulas, proving causality appears to be more challenging mathematically, since it is no longer a question of analysing the convergence of series. In the next section we use simulations to conjecture that causality holds for a class of processes defined via the Kendall correlations of the copula sequence.

In a finite-order process the copula sequence for any lag k greater than the order p consists of independence copulas; it seems intuitively clear that, to obtain an infinite-order process with a convergent causal representation, the partial copula sequence $(C_k)_{k\in\mathbb{N}}$ should converge to the independence copula C^{\perp} as $k \to \infty$. However, in view of Example 4.3, this is not a sufficient condition and the speed of convergence of the copula sequence is also important. Ideally we require conditions on the speed of convergence $C_k \to C^{\perp}$ so that the marginal copula $C^{(k)}$ in (8) also tends to C^{\perp} ; in that case the variables U_t and U_{t-k} are asymptotically independent as $k \to \infty$ and mixing behaviour follows.

5.2 A practical approach to non-Gaussian s-vines

Suppose we take a sequence of pair copulas $(C_k)_{k\in\mathbb{N}}$ from some parametric family and parameterize them in such a way that (i) the copulas converge uniformly to the independence copula as $k \to \infty$ and (ii) the *level of dependence* of each copula C_k is identical to that of a Gaussian pair copula sequence that gives rise to an ergodic Gaussian process. The intuition here is that by sticking close to the pattern of decay of dependence in a well-behaved Gaussian process, we might hope to construct a stable causal process that is both mixing and ergodic.

A natural way of making 'level of dependence' concrete is to consider the Kendall rank correlation function of the copula sequence, defined in the following way.

Definition 4. The Kendall partial autocorrelation function (kpacf) $(\tau_k)_{k\in\mathbb{N}}$ associated with a copula sequence $(C_k)_{k\in\mathbb{N}}$ is given by $\tau_k = \tau(C_k)$, for $k \in \mathbb{N}$, where $\tau(C)$ denotes the Kendall's tau coefficient for a copula C.

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

$$\tau_k = \frac{2}{\pi} \arcsin(\alpha_k). \tag{24}$$

As in Section 4.2, suppose that $(\alpha_k(\boldsymbol{\theta}))_{k\in\mathbb{N}}$ is the pact of a stationary and ergodic model Gaussian process parametrized by the parameters $\boldsymbol{\theta}$, such as an ARMA or ARFIMA model; this implies a parametric form for the kpacf $(\tau_k(\boldsymbol{\theta}))_{k\in\mathbb{N}}$. The idea is to choose a sequence of non-Gaussian pair copulas that shares this kpacf.

A practical problem that may arise is that $\tau_k = \tau_k(\boldsymbol{\theta})$ can in theory take any value in (-1, 1); only certain copula families, such as Gauss and Frank, are said to be *comprehensive* and yield any value for τ_k . If we wish to use, for example, a sequence of Gumbel copulas to build our model then we need to find a solution for negative values of Kendall's tau. One possibility is to allow 90 or 270 degree rotations of the copula at negative values of τ_k and another is to substitute a comprehensive copula at any position k in the sequence where τ_k is negative.

Remark 2. Note that the assumption that the pair copulas C_k converge to the independence copula has implications for using t copulas $C_{\nu,\alpha}^t$ in this approach. The terms of the copula sequence $C_k = C_{\nu_k,\alpha_k}^t$ would have to satisfy $\nu_k \to \infty$ and $\alpha_k \to 0$ as $k \to \infty$; the sequence given by $C_k = C_{\nu,\alpha_k}^t$ for fixed ν does not converge to the independence copula as $\alpha_k \to 0$. While the sequence $(\alpha_k)_{k\in\mathbb{N}}$ can be connected to the kpacf by the same formula (24), the sequence $(\nu_k)_{k\in\mathbb{N}}$ is not fixed by the kpacf. It is simpler in this approach to work with copula families with a single parameter so that there is a one-to-one relationship between Kendall's tau and the copula parameter.

To compare the speed of convergence of the copula filter for different copula sequences sharing the same kpacf, we conduct some simulation experiments. For fixed n and for a fixed realization z_1, \ldots, z_n of independent uniform noise we plot the points $(k, S_k(\mathbf{z}_{[n-k,n-1]}, z_n))$ 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. When the copula sequence consists of Clayton copulas we will refer to the model as a Clayton copula filter; similarly Gumbel copulas yield a Gumbel copula filter; and so on. The following examples suggest that there are some differences in the convergence rates of the copula filters. This appears to relate to the tail dependence characteristics of the copulas [25, 27]. We recall that the Gumbel and Joe copulas are upper tail dependent while the Clayton copula is lower tail dependent; the Gauss and Frank copulas are tail independent. The filters based on sequences of tail-dependent copulas generally show slower convergence.

Example 4 (Non-Gaussian ARMA(1,1) models). In this example we consider s-vine copula processes sharing the kpacf of the ARMA(1,1) model with autoregressive parameter 0.95 and movingaverage parameter -0.85. Fixing n = 201 we obtain Figure 2. Convergence appears to be fastest for the Gaussian and Frank copula filters and slowest for the Clayton filter, followed by the Joe filter; the Gumbel filter is an intermediate case. We can also discern a tendency for jumps in the value of $S_k(\mathbf{z}_{[n-k,n-1]}, z_n))$ to be upward for the upper tail-dependent Gumbel and Joe copulas and downward for the lower tail-dependent Clayton copula.

Example 5 (Non-Gaussian ARFIMA(1,d,1) models). In this example we consider s-vine copula processes sharing the kpacf of the ARFIMA(1,d,1) model with autoregressive parameter 0.95, movingaverage parameter -0.85 and fractional differencing parameter d = 0.02. The latter implies that the pacf of the Gaussian process satisfies $|\alpha_k| \sim 0.02/k$ as $k \to \infty$ [23]. The lack of absolute summability means that the Gaussian copula process does not satisfy the conditions of Theorem 1. It is an unresolved question whether any of these processes is causal. Fixing n = 701 we obtain Figure 3. For the realized series of innovations used in the picture, convergence appears to take place, but is extremely slow. The tail-dependent Clayton and Joe copulas appear to take longest to settle down.

An obvious practical solution that circumvents the issue of whether the infinite-order process has a convergent causal representation is to truncate the copula sequence $(C_k)_{k\in\mathbb{N}}$ so that $C_k = C^{\perp}$



Figure 2: Plots of $(k, S_k(\boldsymbol{z}_{[n-k,n-1]}, z_n))$ for $k \in \{1, \ldots, n-1\}$ for the copula filters of ARMA(1,1) models; see Example 4. Horizontal lines show ultimate values $S_{n-1}(\boldsymbol{z}_{[1,n-1]}, z_n)$.

for k > p for some relatively large but fixed value p. This places us back in the setting of ergodic Markov chains but, by parameterizing models through the kpacf, we preserve the advantages of parsimony.

5.3 An example with real data

For this example we have used data on the US CPI (consumer price index) taken from the OECD webpage. We analyse the log-differenced time series of quarterly CPI values from the first quarter of 1960 to the 4th quarter of 2020, which can be interpreted as measuring the rate of inflation [46, Sections 14.2–14.4]. The inflation data are shown in the upper-left panel of Figure 4; there are n = 244 observations.

To establish a baseline model we use an automatic ARMA selection algorithm and this selects an ARMA(5,1) model. We first address the issue of whether the implied Gaussian copula sequence in an ARMA(5,1) model can be replaced by Gumbel, Clayton, Frank or Joe copula sequences (or 180 degree rotations thereof); for any lag k at which the estimated kpacf τ_k is negative we retain a



Figure 3: Plots of $(k, S_k(\boldsymbol{z}_{[n-k,n-1]}, z_n))$ for $k \in \{1, \ldots, n-1\}$ for the copula filters of ARFIMA(1,d,1) models; see Example 5. Horizontal lines show ultimate values $S_{n-1}(\boldsymbol{z}_{[1,n-1]}, z_n)$.

Gaussian copula and so the non-Gaussian copula sequences are actually hybrid sequences with some Gaussian terms. The data (x_1, \ldots, x_n) are transformed to pseudo-observations (u_1, \ldots, u_n) on the copula scale using the empirical distribution function and the s-vine copula process is estimated by maximum-likelihood; this is the commonly used pseudo-maximum-likelihood method [19, 12].

The best model results from replacing Gaussian copulas with Gumbel copulas and the improvements in AIC and BIC are shown in the upper panel of Table 1; the improvement in fit is strikingly large. While the presented results relate to infinite-order processes, we note that very similar result (not tabulated) are obtained by fitting s-vine copula processes of finite order where the kpacf is truncated at lag 30. Parameter estimates for the infinite-order models are given in Table 2.

The residual QQ-plots in the middle row of Figure 4 give further insight into the improved fit of the process with Gumbel copulas. In the usual manner, residuals are reconstructions of the unobserved innovation variables. If $(\hat{R}_k)_{k\in\mathbb{N}}$ denotes the sequence of estimated Rosenblatt forward functions, implied by the sequence $(\hat{C}_k)_{k\in\mathbb{N}}$ of estimated copulas, then residuals (z_1, \ldots, z_n) are constructed by setting $z_1 = u_1$ and $z_t = \hat{R}_{t-1}(\boldsymbol{u}_{[1,t-1]}, u_t)$ for t > 1. To facilitate graphical analysis these are transformed onto the standard normal scale so that the QQ-plots in the middle row of Figure 4 relate to the values $(\Phi^{-1}(z_1), \ldots, \Phi^{-1}(z_n))$ and are against a standard normal reference distribution. The residuals from the baseline Gaussian copula appear to deviate from normality whereas the residuals from the Gumbel copula model are much better behaved; the latter pass a Shapiro-Wilk test of normality (p-value = 0.97) whereas the former do not (p-value = 0.01).

	No.	pars	AIC	BIC
Gaussian copula process		6	-184.62	-163.64
Gumbel copula process		6	-209.28	-188.30
Gaussian process		8	372.73	400.71
Gaussian copula process $+$ skewed Student margin		10	352.50	387.47
Gumbel copula process $+$ skewed Student margin		10	319.17	354.14

Table 1: Comparison of models by AIC and BIC: top panel relates to models for the pseudo-copula data (u_1, \ldots, u_n) while the lower panel relates to full models of the original data (x_1, \ldots, x_n) .

	$oldsymbol{ heta}^{(\mathrm{Ga})}$	s.e.	$oldsymbol{ heta}^{(\mathrm{Gu})}$	s.e.
ϕ_1	-0.381	0.104	-0.232	0.130
ϕ_2	0.144	0.081	0.136	0.094
ϕ_3	0.197	0.063	0.180	0.061
ϕ_4	0.462	0.075	0.410	0.077
ϕ_5	0.324	0.063	0.266	0.061
ψ_1	0.870	0.098	0.771	0.118

Table 2: Parameter estimates and standard errors for s-vine copula processes with Gaussian and Gumbel copula sequences fitted to the pseudo-copula data (u_1, \ldots, u_n) .

The picture of the kpacf in the top right panel of Figure 4 requires further comment. This plot attempts to show how well the kpacf of the fitted copula sequence matches the empirical Kendall partial autocorrelations of the data. The continuous line is the kpacf of the Gumbel/Gaussian copula sequence used in the best-fitting vine copula model of (u_1, \ldots, u_n) . The vertical bars show the empirical Kendall partial autocorrelations of the data at each lag k. However, the method should really be considered as 'semi-empirical' as it uses the fitted parametric copulas at lags $1, \ldots, k-1$ in order to construct the necessary data for lag k. The data used to estimate an empirical lag k rank correlation are the points

$$\left\{ \left(\widehat{R^{(2)}}_{k-1}(\boldsymbol{u}_{[j-k+1,j-1]}, u_{j-k}), \widehat{R}_{k-1}(\boldsymbol{u}_{[j-k+1,j-1]}, u_j)\right), \quad j = k+1, \dots, n \right\},\$$

where $\widehat{R_k}$ and $\widehat{R^{(2)}}_k$ denote the estimates of forward and backward Rosenblatt functions; it may be

noted that these data are precisely the points at which the copula density c_k is evaluated when the model likelihood based on $c_{(n)}$ in (5) is maximized.

The kpacf shows positive dependence between inflation rates at the first 5 lags; moreover the choice of Gumbel copula suggests asymmetry and upper tail dependence in the bivariate distribution of inflation rates at time points that are close together; in other words, large values of inflation are particularly strongly associated with large values of inflation in previous quarters, while low values are more weakly associated.

We next consider composite models for the original data (x_1, \ldots, x_n) consisting of a marginal distribution and an s-vine copula process. The baseline model is simply a Gaussian process with Gaussian copula sequence and Gaussian marginal distribution. We experimented with a number of alternatives to the normal marginal and obtained good results with the skewed Student distribution from the family of skewed distributions proposed by Fernandez and Steel [18]. Table 1 contains results for models which combine the Gaussian and Gumbel copula sequences with the skewed Student margin; the improvement obtained by using a Gumbel sequence with a skewed Student margin is clear from the AIC and BIC values. The QQ-plots of the data against the fitted marginal distributions in the bottom row of Figure 4 also show the superiority of the skewed Student to the Gaussian distribution for this dataset.

The fitting method used for the composite model results in Table 1 is the two-stage IFM (inference functions for margins) method [25] in which the margin is estimated first, the data are transformed to approximately uniform using the marginal model, and the copula process is estimated by ML in a second step.

The estimated values of the degree of freedom and skewness parameters in the skewed Student t marginal distribution are $\nu = 3.19$ and $\gamma = 1.47$ respectively. These suggest that inflation rates (changes in log CPI) follow a heavy tailed, infinite-kurtosis distribution (tail index = 3.19) that is skewed to the right.

6 Conclusion

The s-vine processes provide a class of tractable stationary models that can capture non-linear and non-Gaussian serial dependence behaviour as well as any continuous marginal behaviour. By defining models of infinite order and using the approach based on the Kendall partial autocorrelation function (kpacf), we obtain a very natural generalization of classical Gaussian processes, such as



Figure 4: Top row: log-differenced CPI data and estimated kpacf of s-vine copula process using Gumbel copula sequence. Middle row: QQ-plots for residuals from models based on Gaussian (left) and Gumbel (right) copula sequences. Bottom row: QQ-plots of the data against fitted normal (left) and skewed Student (right) marginal distributions.

Gaussian ARMA or ARFIMA.

The models are straightforward to apply. The parsimonious parametrization based on the kpacf makes maximum likelihood inference feasible. Analogues of many of the standard tools for time series analysis in the time domain are available, including estimation methods for the kpacf and residual plots that shed light on the quality of the fit of the copula model. By separating the issues of serial dependence and marginal modelling, we can obtain bespoke descriptions of both aspects that avoid the compromises of the more 'off-the-shelf' classical approach. The example of Section 5.3 indicates the kind of gains that can be obtained; it seems likely that many empirical applications of classical ARMA could be substantially improved by the use of models in the general s-vine class. In combination with v-transforms [33] s-vine models could also be used to model data showing stochastic volatility following the approach in Bladt and McNeil [9].

To increase the practical options for model building it would be of interest to consider how copulas with more than one parameter, such as the t copula or the symmetrized Joe-Clayton copula [37] could be incorporated into the methodology. The parameters would have to be allowed to change in a smooth parsimonious manner such that the partial copula sequence $(C_k)_{k\in\mathbb{N}}$ converged to the independence copula while the Kendall correlations $(\tau_k)_{k\in\mathbb{N}}$ followed the chosen form of kpacf for every k. This is a topic for further research.

The approach we have adopted should also be of interest to theoreticians as there are a number of challenging open questions to be addressed. While we have proposed definitions of causality and invertibility for general s-vine processes, we currently lack a mathematical methodology for checking convergence of causal and invertible representations for sequences of non-Gaussian pair copulas.

There are some very interesting questions to address about the relationship between the partial copula sequence $(C_k)_{k\in\mathbb{N}}$, the rate of convergence of causal representations and the rate of ergodic mixing of the resulting processes. The example of Figure 1 indicates that, even for a finite-order process, some very extreme models can be constructed that mix extremely slowly. Moreover, Example 5 suggests that non-Gaussian copula sequences serve to further elongate memory in long-memory processes and this raises questions about the effect of the tail dependence properties of the copula sequence on rates of convergence and length of memory.

It would also be of interest to confirm our conjecture that the pragmatic approach adopted in Section 5.2, in which the kpace of the (infinite) partial copula sequence $(C_k)_{k\in\mathbb{N}}$ is matched to that of a stationary and ergodic Gaussian process, always yields a stationary and ergodic s-vine model, regardless of the choice of copula sequence. However, for practical applications, the problem can be obviated by truncating the copula sequence at some large finite lag k, so that we are dealing with an ergodic Markov chain as in Section 3.

A Proofs

A.1 Proof of Proposition 1

In this proof we use the notation $(\boldsymbol{u})_i$ to denote the *i*th component of a vector \boldsymbol{u} and \boldsymbol{u}_{-i} to denote the vector \boldsymbol{u} with *i*th component removed. An exchangeable copula satisfies $C_k(u, v) = C_k(v, u)$ for all u, v and hence $h_k^{(2)}(u, v) = h_k^{(1)}(v, u)$. From this it follows that $R_1^{(2)}(u, x) = h_1^{(2)}(x, u) = h_1^{(1)}(u, x) = R_1(u, x)$. Part (i) follows by induction using the facts that for $\boldsymbol{u} = (u_1, \dots, u_k)^\top \in (0, 1)^k$ we have $(\overline{\boldsymbol{u}})_k = u_1, \, \overline{\boldsymbol{u}}_{-k} = (u_k, \dots, u_2)^\top$ and $\overline{\boldsymbol{u}}_{-k} = \boldsymbol{u}_{-1}$. We have that

$$R_{k}^{(2)}(\overline{\boldsymbol{u}},x) = h_{k}^{(2)} \left(R_{k-1}^{(2)}(\overline{\boldsymbol{u}}_{-k},x), R_{k-1}(\overline{\boldsymbol{u}}_{-k},(\overline{\boldsymbol{u}})_{k}) \right)$$
$$= h_{k}^{(1)} \left(R_{k-1}(\overline{\boldsymbol{u}}_{-k},u_{1}), R_{k-1}^{(2)}(\overline{\boldsymbol{u}}_{-k},x) \right) = h_{k}^{(1)} \left(R_{k-1}^{(2)}(\boldsymbol{u}_{-1},u_{1}), R_{k-1}(\boldsymbol{u}_{-1},x) \right) = R_{k}(\boldsymbol{u},x).$$

For part (ii) we observe that for any k > 1 the implication of part(i) is that for any $t, s \in \mathbb{N}$ the conditional distribution of $U_{t+k} \mid U_{t+1} = u_1, \ldots, U_{t+k-1} = u_{k-1}$ is the same as that of $U_{s+1} \mid U_{s+2} = u_{k-1}, \ldots, U_{s+k} = u_1$. It easily follows that $(U_{t+1}, \ldots, U_{t+k}) \stackrel{d}{=} (U_{s+k}, \ldots, U_{s+1})$ which proves reversibility of the process.

A.2 Proof of Proposition 2

If (X_t) is a Gaussian process its marginal distributions of all orders are multivariate Gaussian. The general d-vine copula decomposition in (1) can be applied to each *n*-dimensional marginal density. Since the conditional distributions of pairs (X_{j-k}, X_j) given intermediate variables are bivariate Gaussian distributions with covariance matrices that do not depend on the conditioning variables, the simplifying assumption holds for each pair copula density in (1) [27, pages 106–108]. The stationarity assumption ensures that the joint density of the *n*-dimensional copula takes the form (5).

Conversely, an s-vine process with Gaussian marginal density and Gaussian pair copulas is a stationary process with n-dimensional marginal densities of the form (7). These are the densities of multivariate Gaussian distributions and the resulting process is a Gaussian process.

A.3 Proof of Proposition 3

Let $(Z_k)_{k\in\mathbb{N}}$ be a sequence of iid standard uniform variables and $(U_k)_{k\in\mathbb{N}}$ a sequence of uniform random variables generated by setting $U_1 = Z_1$ and $U_k = R_{k-1}^{-1}((U_1, \ldots, U_{k-1})^{\top}, Z_k)$ for k >1 where $(R_k)_{k\in\mathbb{N}}$ denotes the sequence of Rosenblatt functions associated with the sequence of Gaussian pair copulas $(C_k)_{k\in\mathbb{N}}$. Moreover, let $(X_k)_{k\in\mathbb{N}}$ be a sequence of standard Gaussian variables defined by setting $X_k = \Phi^{-1}(U_k)$ for all k.

It follows that, for any $k \ge 1$, $(X_1, \ldots, X_{k+1}) \sim N_{k+1}(\mathbf{0}, P_{k+1})$ where P_{k+1} is the (k+1)dimensional correlation matrix implied by the acf $(\rho_i)_{i\in\mathbb{N}}$ of $(X_i)_{i\in\mathbb{N}}$ as in (14). The standard result for the conditional distribution of a multivariate normal implies that

$$X_{k+1} \mid X_1 = x_1, \dots, X_k = x_k \sim N\Big(\overline{\boldsymbol{\rho}}_k^\top P_k^{-1}(x_1, \dots, x_k)^\top, \ 1 - \boldsymbol{\rho}_k^\top P_k^{-1} \boldsymbol{\rho}_k\Big),$$

where $\boldsymbol{\rho}_k = (\rho_1, \dots, \rho_k)^{\top}$ as in (14) and $\overline{\boldsymbol{\rho}}_k$ is the reversed vector. The mean of the conditional distribution is the best linear predictor of X_{k+1} and the variance of the conditional distribution is the mean squared prediction error; let us write the former as $\sum_{j=1}^k \phi_j^{(k)} x_{k+1-j}$, where $\phi_j^{(k)} = (P_k^{-1} \boldsymbol{\rho}_k)_j$, and the latter as σ_k^2 . We then have

$$R_{k}(\boldsymbol{u}, x) = \mathbb{P}(U_{k+1} \leq x \mid U_{1} = u_{1}, \dots, U_{k} = u_{k})$$

$$= \mathbb{P}(X_{k+1} \leq \Phi^{-1}(x) \mid X_{1} = \Phi^{-1}(u_{1}), \dots, X_{k} = \Phi^{-1}(u_{k}))$$

$$= \Phi\left(\frac{\Phi^{-1}(x) - \sum_{j=1}^{k} \phi_{j}^{(k)} \Phi^{-1}(u_{k+1-j})}{\sigma_{k}}\right).$$

The expression $\sigma_k^2 = \prod_{j=1}^k (1 - \alpha_j^2)$ and the recursive formula (18) for the coefficients $\phi_j^{(k)}$ follow from the Durbin–Levinson Algorithm [11, Proposition 5.2.1].

It follows from (17) that, for k > 0,

$$U_{k+1} = R_k^{-1}((U_1, \dots, U_k)^{\top}, Z_{k+1}) = \Phi\left(\sigma_k \Phi^{-1}(Z_{k+1}) + \sum_{j=1}^k \phi_j^{(k)} \Phi^{-1}(U_{k+1-j})\right)$$
(25)

which may be written in terms of the variables $(X_k)_{k\in\mathbb{N}}$ as $X_{k+1} = \sigma_k \epsilon_{k+1} + \sum_{j=1}^k \phi_j^{(k)} X_{k+1-j}$ where we introduce the further notation $\epsilon_k = \Phi^{-1}(Z_k)$ for all $k \in \mathbb{N}$. An inductive argument then shows that this may be written in the form $X_{k+1} = \sigma_k \epsilon_{k+1} + \sum_{j=1}^k \psi_j^{(k)} \epsilon_{k+1-j}$ with the coefficients $\psi_j^{(k)}$ as defined in (20). Equation (19) then follows easily from (11).

A.4 Proof of Theorem 1

As in the proof of Proposition 3 we introduce the notation $(X_t)_{t\in\mathbb{Z}}$ and $(\epsilon_t)_{t\in\mathbb{Z}}$ where $X_t = \Phi^{-1}(U_t)$ and $\epsilon_t = \Phi^{-1}(Z_t)$. For fixed k the formulas $U_t = S_k((Z_{t-k}, \ldots, Z_{t-1})^\top, Z_t)$ and $Z_t = R_k((U_{t-k}, \ldots, U_{t-1})^\top, U_t)$ translate to

$$X_t = \sum_{j=0}^k \psi_j^{(k)} \epsilon_{t-j} \quad \text{and} \quad \epsilon_t = \sum_{j=0}^k \tilde{\phi}_j^{(k)} X_{t-j}$$

where $\tilde{\phi}_0^{(k)} = 1/\sigma_k$ and $\tilde{\phi}_j^{(k)} = -\phi_j^{(k)}/\sigma_k$ for j > 1.

Debowaski [16, Theorem 6] shows that under Assumption 3 the limiting representations

$$X_t = \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j} \quad \text{and} \quad \epsilon_t = \sum_{j=0}^{\infty} \tilde{\phi}_j X_{t-j}$$
(26)

hold where $(\tilde{\phi}_j)_{j\in\mathbb{Z}_+}$ and $(\psi_j)_{j\in\mathbb{Z}_+}$ are sequences such that $\lim_{k\to\infty} \tilde{\phi}_j^{(k)} = \tilde{\phi}_j$ and $\lim_{k\to\infty} \psi_j^{(k)} = \psi_j$ and such that $\psi_j = \operatorname{cov}(X_t, Z_{t-j})$, $\tilde{\phi}_0 = 1/\psi_0$ and $\sum_{j=0}^k \tilde{\phi}_j \psi_{k-j} = 0$ for k > 0; the series in the rhs of (26) converge absolutely, almost surely. Moreover, under Assumption 3 we also have that the terms $\sigma_k^2 = \prod_{j=1}^k (1 - \alpha_j^2)$ converge to a finite limit σ and so we can introduce a sequence $(\phi_j)_{j\in\mathbb{N}}$ such that $\phi_j = -\sigma \tilde{\phi}_j$ and write

$$X_t = \sigma \epsilon_t + \sum_{j=1}^{\infty} \psi_j \epsilon_{t-j} \quad \text{and} \quad \epsilon_t = \sigma^{-1} X_t - \sigma^{-1} \sum_{j=1}^{\infty} \phi_j X_{t-j}.$$

Finally, by Proposition 3, equations (22) and (23) are seen to be a restatement of the latter formulas in terms of the Rosenblatt functions.

B Markov chain analysis

The Markov chain specified by (16) under Assumption 1 is a well-behaved example of a chain on a general state space. The properties of the process can be verified by standard arguments which are collected here for completeness.

Invariance. The transition kernel of the Markov chain is given by

$$\mathsf{P}(\boldsymbol{u}, A) = \mathbb{P}(\boldsymbol{U}_{2} \in A \mid \boldsymbol{U}_{1} = \boldsymbol{u}) = \mathbb{P}(F(\boldsymbol{U}_{1}, Z) \in A \mid \boldsymbol{U}_{1} = \boldsymbol{u}) = \prod_{i=2}^{p} I_{\{u_{i} \in A_{i-1}\}} \int_{A_{p}} f_{p}(\boldsymbol{u}, x) dx$$

for a set $A = A_1 \times \cdots \times A_p \subseteq \mathcal{X}$. Writing π for the probability measure implied by $C_{(p)}$ and using (7), we have that

$$\int_{(0,1)^{p}} \pi(\mathrm{d}\boldsymbol{u}) \mathsf{P}(\boldsymbol{u}, A) = \int_{(0,1)^{p}} c_{(p)}(u_{1}, \dots, u_{p}) \prod_{i=2}^{p} I_{\{u_{i} \in A_{i-1}\}} \left(\int_{A_{p}} f_{p}(\boldsymbol{u}, x) \mathrm{d}x \right) \mathrm{d}\boldsymbol{u}$$

$$= \int_{0}^{1} \int_{A_{1}} \cdots \int_{A_{p-1}} c_{(p)}(u_{1}, \dots, u_{p}) \left(\int_{A_{p}} f_{p}(\boldsymbol{u}, x) \mathrm{d}x \right) \mathrm{d}u_{1} \cdots \mathrm{d}u_{p}$$

$$= \int_{0}^{1} \int_{A_{1}} \cdots \int_{A_{p}} c_{(p+1)}(u_{1}, \dots, u_{p}, u_{p+1}) \mathrm{d}u_{1} \cdots \mathrm{d}u_{p+1}$$

$$= \int_{A_{1}} \cdots \int_{A_{p}} c_{(p)}(u_{2}, \dots, u_{p}, u_{p+1}) \mathrm{d}u_{2} \cdots \mathrm{d}u_{p+1} = \pi(A)$$

showing that π is an invariant measure.

Irreducibility. A process is ϕ -irreducible if there is a measure ϕ on \mathcal{X} such that for every set $A \subseteq \mathcal{X}$ with $\phi(A) > 0$ and every $\boldsymbol{u} \in (0,1)^p$ there exists $n = n(\boldsymbol{u}, A) > 0$ such that $\mathsf{P}^n(\boldsymbol{u}, A) > 0$. In our case it suffices to take n = p, independent of \boldsymbol{u} and A, and ϕ to be Lebesgue measure. After *p*-fold iteration of the Markov updating scheme in (16) we obtain the random vector $\boldsymbol{U}_{p+1} = (U_{p+1}, \ldots, U_{2p})^{\top}$ and for a set $A = A_1 \times \cdots \times A_p$ with positive Lebesgue measure we have

$$\mathsf{P}^{p}(\boldsymbol{u}, A) = \mathbb{P}(\boldsymbol{U}_{p+1} \in A \mid \boldsymbol{U}_{1} = \boldsymbol{u})$$

= $\mathbb{P}(U_{p+1} \in A_{1}, \dots, U_{2p} \in A_{p} \mid U_{1} = u_{1}, \dots, U_{p} = u_{p})$
= $\int_{A_{1}} \cdots \int_{A_{p}} f_{U_{p+1},\dots,U_{2p}\mid U_{1},\dots,U_{p}}(x_{1},\dots,x_{p} \mid u_{1},\dots,u_{p}) \mathrm{d}x_{1} \cdots \mathrm{d}x_{p}$

and this probability is strictly positive if the integrand is strictly positive. Since the function in the integrand satisfies

$$f_{U_{p+1},\dots,U_{2p}|U_1,\dots,U_p}(u_{p+1},\dots,u_{2p} \mid u_1,\dots,u_p) = \prod_{j=1}^p f_p((u_j,\dots,u_{j+p-1})^\top,u_{j+p})$$

and the conditional densities $f_p(\boldsymbol{u}, x)$ as defined in (7) are products of strictly positive pair-copula densities, the result follows.

Recurrence. Since the Markov chain is ϕ -irreducible and admits an invariant probability measure, it is a positive recurrent chain. The absolute continuity of the transition kernel with respect to Lebesgue measure (exploited above) also means it is a Harris recurrent chain: for any point $\boldsymbol{x} \in \mathcal{X}$ and any set A with invariant measure $\pi(A) = 1$, either $\boldsymbol{x} \in A$ or, if not, $\mathsf{P}(\boldsymbol{x}, A) = 1$ so that it is certain that the time to entering A is finite and this is a condition for Harris recurrence [40, Theorem 6(v)].

Aperiodicity. A periodic chain would cycle through $m \ge 2$ disjoint subsets of the state space $\mathcal{X}_1, \ldots, \mathcal{X}_m$, each satisfying $\phi(\mathcal{X}_i) > 0$, in successive steps. If such behaviour were to occur we could find an $n \ge p$ such that $\mathsf{P}^n(\boldsymbol{u}, \mathcal{X}_2) = \mathbb{P}(\boldsymbol{U}_{n+1} \in \mathcal{X}_2 \mid \boldsymbol{U}_1 = \boldsymbol{u}) = 1$ for all $\boldsymbol{u} \in \mathcal{X}_1$, which would imply that $\mathsf{P}^n(\boldsymbol{u}, \mathcal{X}_1) = 0$ for all $\boldsymbol{u} \in \mathcal{X}_1$. However, since $\phi(\mathcal{X}_1) > 0$, the argument used to establish the ϕ -irreducibility of the process can be repeated to show that $\mathsf{P}^n(\boldsymbol{u}, \mathcal{X}_1) > 0$ for all $\boldsymbol{u} \in \mathcal{X}$, which yields a contradiction.

Conflict of interest statement. The authors declare no conflict of interest.

References

- Aas, K., C. Czado, A. Frigessi, and H. Bakken (2009). Pair-copula constructions of multiple dependence. *Insurance Math. Econom.* 44 (2), 182–198.
- [2] Barndorff-Nielsen, O. and G. Schou (1973). On the parametrization of autoregressive models by partial autocorrelation. J. Multivariate Anal. 3(4), 408–419.
- [3] Beare, B. (2010). Copulas and temporal dependence. *Econometrica* 78, 395–410.
- [4] Beare, B. and J. Seo (2015). Vine copula specifications for stationary multivariate Markov chains. J. Time Series Anal. 36, 228–246.
- [5] Beare, B. K. (2012). Archimedean copulas and temporal dependence. *Econometric Theory* 28(6), 1165–1185.
- [6] Bedford, T. and R. Cooke (2001a). Probabilistic Risk Analysis: Foundations and Methods. Cambridge: Cambridge University Press.
- [7] Bedford, T. and R. M. Cooke (2001b). Probability density decomposition for conditionally independent random variables modeled by vines. Ann. Math. Artif. Intell. 32, 245–268.

- [8] Bedford, T. and R. M. Cooke (2002, 08). Vines-a new graphical model for dependent random variables. Ann. Statist. 30(4), 1031–1068.
- [9] Bladt, M. and A. J. McNeil (2021). Time series copula models using d-vines and v-transforms. Econom. Stat.. Available online 23 July 2021.
- [10] Brechmann, E. C. and C. Czado (2015). COPAR—multivariate time series modeling using the copula autoregressive model. Appl. Stoch. Models Bus. Ind. 31(4), 495–514.
- [11] Brockwell, P. J. and R. A. Davis (1991). Time Series: Theory and Methods (2nd ed.). New York: Springer.
- [12] Chen, X. and Y. Fan (2006). Estimation of copula-based semiparametric time series models.
 J. Econometrics 130(2), 307–335.
- [13] Chen, X., W. B. Wu, and Y. Yi (2009). Efficient estimation of copula-based semiparametric markov models. Ann. Statist. 37(6B), 4214–4253.
- [14] Cornfeld, I., S. Fomin, and G. Sinai (1982). *Ergodic Theory*. Springer-Verlag.
- [15] Darsow, W., B. Nguyen, and E. Olsen (1992). Copulas and Markov processes. Illinois J. Math. 36(4), 600–642.
- [16] Debowski, L. (2007). On processes with summable partial autocorrelations. Statist. Probab. Lett. 77, 752–759.
- [17] Domma, F., S. Giordano, and P. F. Perri (2009). Statistical modeling of temporal dependence in financial data via a copula function. *Comm. Statist. Simulation Comput.* 38(4), 703–728.
- [18] Fernández, C. and M. Steel (1998). On Bayesian modeling of fat tails and skewness. J. Amer. Statist. Assoc. 93(441), 359–371.
- [19] Genest, C., K. Ghoudi, and L. Rivest (1995). A semi-parametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika* 82, 543–552.
- [20] Haff, H., K. Aas, and A. Frigessi (2010). On the simplified pair copula construction simply useful or too simplistic? J. Multivariate Anal. 101, 1296–1310.
- [21] Hosking, J. (1981). Fractional differencing. *Biometrika* 68(1), 165–176.

- [22] Ibragimov, R. (2009). Copula-based characterizations for higher order markov processes. *Econo*metric Theory 25(3), 819–846.
- [23] Inoue, A. (2002). Asymptotic behavior for partial autocorrelation functions of fractional ARIMA processes. Ann. Appl. Probab. 12(4), 1471–1491.
- [24] Joe, H. (1996). Families of m-variate distributions with given margins and m (m-1)/2 bivariate dependence parameters. *Lecture Notes-Monograph Series*, 120–141.
- [25] Joe, H. (1997). Multivariate Models and Dependence Concepts. London: Chapman & Hall.
- [26] Joe, H. (2006). Generating random correlation matrices based on partial correlations. J. Multivariate Anal. 97, 2177–2189.
- [27] Joe, H. (2015). Dependence Modeling with Copulas. CRC Press, Boca Raton.
- [28] Kurowicka, D. and R. Cooke (2006). Uncertainty Analysis with High Dimensional Dependence Modelling. Chichester: Wiley.
- [29] Liebscher, E. (2008). Construction of asymmetric multivariate copulas. J. Multivariate Anal. 99, 2234–2250.
- [30] Loaiza-Maya, R., M. Smith, and W. Maneesoonthorn (2018). Time series copulas for heteroskedastic data. J. Appl. Econometrics 33, 332–354.
- [31] Longla, M. and M. Peligrad (2012). Some aspects of modeling dependence in copula-based Markov chains. J. Multivariate Anal. 111, 234–240.
- [32] Maruyama, G. (1970). Infinitely divisible processes. Theory Probab. Appl. 15(1), 1–22.
- [33] McNeil, A. J. (2021). Modelling volatile time series with v-transforms and copulas. Risks 9(14).
- [34] Meyn, S. and R. L. Tweedie (2009). Markov chains and stochastic stability (2nd edition ed.). Cambridge University Press.
- [35] Mroz, T., S. Fuchs, and W. Trutschnig (2021). How simplifying and flexible is the simplifying assumption in pair copula constructions - analytic answers in dimension three and a glimpse beyond. *Electron. J. Stat.* 15(1), 1951–1992.
- [36] Nagler, T., D. Krüger, and A. Min (2022). Stationary vine copula models for multivariate time series. J. Econometrics 227(2), 305–324.

- [37] Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. Internat. Econom. Rev. 47(2), 527–556.
- [38] Ramsey, F. (1974). Characterization of the partial autocorrelation function. Ann. Statist. 2(6), 1296–1301.
- [39] Rémillard, B., N. Papageorgiou, and F. Soustra (2012). Copula-based semiparametric models for multivariate time series. J. Multivariate Anal. 110, 30–42.
- [40] Roberts, G. O. and J. S. Rosenthal (2006). Harris recurrence of Metropolis-within-Gibbs and trans-dimensional Markov chains. Ann. Appl. Probab. 16(4), 2123–2139.
- [41] Samorodnitsky, G. (2007). Long range dependence. Foundations and Trends in Stochastic Systems 1(3), 163–257.
- [42] Smith, M., A. Min, C. Almeida, and C. Czado (2010). Modeling Longitudinal Data Using a Pair-Copula Decomposition of Serial Dependence. J. Amer. Statist. Assoc. 105 (492), 1467–1479.
- [43] Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. Internat. J. Forecasting 31(3), 815–833.
- [44] Spanhel, F. and M. Kurz (2019). Simplified vine copula models: approximations based on the simplifying assumption. *Electron. J. Stat.* 13(1), 1254–1291.
- [45] Stöber, J., H. Joe, and C. Czado (2013). Simplified pair copula extensions—limitations and extensions. J. Multivariate Anal. 119, 101–118.
- [46] Stock, J. and M. Watson (2011). Introduction to Econometrics (3rd edition). Addison Wesley Longman.