

Modeling cold-related excess deaths via stationary vine copulas

Han Li^[a], Thomas Nagler^[b], and Claudia Czado^[c]

^[a]Centre for Actuarial Studies, Department of Economics, University of Melbourne

^[b]Department of Statistics, LMU Munich & Munich Center for Machine Learning

^[c]Department of Mathematics and Munich Data Science Institute, Technical University of Munich

March 25, 2026

Abstract

Extreme cold temperature events have long been associated with excess mortality via many different causes of death. Climate change is expected to intensify the frequency and severity of these extreme temperature events. To quantify and model cold-related excess deaths and, in turn, to better understand the potential impact of climate change on future mortality levels, we propose a new approach based on the state-of-the-art stationary vine copulas. We adopt the S-vine model for the first time in the context of climate-driven mortality risk, and introduce a special case of the model to aid model comparison and enhance interpretability of the results. This model is referred to as a (stationary) centrally connected C-vine (CCC-vine). Three types of dependence are captured by the proposed models, which are temporal dependence, contemporaneous cross-sectional dependence, and non-contemporaneous cross-sectional dependence. We fit the CCC-vine model to the US regional cause-specific death data over the period 1999–2018 and conclude that the model outperforms various benchmark models including the Gaussian copula model and the VAR model. Based on the fitted models, we generate several temperature scenarios and assess cause-specific excess deaths and overall excess deaths due to extreme cold temperatures. We also analyze and compare the geographical differences in cold-related excess deaths across six continental US regions. The results from our study can help public health interventions during extreme cold events to reduce temperature-driven excess deaths.

Keywords: Cause-of-death; Stationary vine; Extreme temperature; Vine copulas; Dependence modeling.

1 Introduction

A greater understanding of the link between climate change and mortality has a number of profound implications including for public policy planning and life insurance pricing. While there is no cause-of-death referred to as climate change, it is globally recognised that climate change can pose a serious threat to human lives (McMichael, 2011; Forzieri *et al.*, 2017). According to the World Health Organization (WHO), between 2030 and 2050, climate change is expected to cause approximately 250,000 additional deaths per year due to intense short-term temperature fluctuations and climate-sensitive diseases (WHO, 2014). A recent study found that approximately 9 out of every 100 deaths in the world during 2000–2019 were due to extreme cold temperatures (Zhao *et al.*, 2021). Extreme temperature events, including extreme cold weather, are likely to become more frequent and severe as a consequence of climate change (Kim *et al.*, 2017; Cohen *et al.*, 2018).

Previous studies found that ambient cold can lead to substantial short-term increases in mortality from multiple causes such as cardiovascular and respiratory diseases (Donaldson and Keatinge, 2002; Dushoff *et al.*, 2006; Arbuthnott *et al.*, 2018). Moreover, these extreme temperature events are likely to have a greater impact on vulnerable segments of the population, particularly the “oldest old”¹. Wan *et al.* (2022) examined the impact of low and high temperatures on death count using a quasi-Poisson regression model, with daily mean temperature as the predictor. They found that the elderly (ages 75+) in Scotland are more sensitive to both extreme cold and heat. By modeling joint extremes in temperature and mortality in the US, Li and Tang (2022) concluded that the extremes in cold temperatures and old-age (ages 85+) death counts exhibit the strongest level of dependence. On the other hand, they found that the relationship between deaths and extreme heat is small when compared to the relationship between deaths and extreme cold in the US.² In line with these findings we focus on the impact of extreme cold temperature on excess deaths. We consider US monthly cause-specific death count for people aged 85+ during 1999–2018, and a cold temperature index for the same period. Six continental regions are included in the study, namely Central West Pacific (CWP), Southwest Pacific (SWP), Southern Plains (SPL), Midwest (MID), Southeast Atlantic (SEA), and Central East Atlantic (CEA).

In this research, we jointly model time series data on cause-specific death counts and an underlying temperature index. Our approach differs from existing studies as we model the relationship between temperature and death count over their entire distributions, rather than the extreme values only. Moreover, instead of focusing on one or two causes of death, we look at the impact of cold weather on excess deaths due to several major causes including *Diabetes*, *External*, *Neoplasm*, *Respiratory*, and *Vascular*, as well as on the aggregated deaths. In this way, we take into account the dependence structure across different causes of death. Finally, the proposed model enables us to quantify any lagged effects of extreme temperature on death count, such as how cold temperatures in the previous period can affect excess deaths in the current period.

¹The American Geriatric Society and the World Health Organization define the oldest old as individuals aged over 80 years; the British Geriatrics Society uses 85 years as a threshold.

²These results are consistent with the conclusions made by the Society of Actuaries, one of the largest actuarial and insurance professional bodies in the world, in a recent industry report (Serre, 2022).

Three types of dependence need to be considered between the underlying temperature index and death count series, namely temporal dependence, contemporaneous cross-sectional dependence, and non-contemporaneous cross-sectional dependence. First, serial correlations in temperature and death count series over time need to be taken into account, for example, a higher rate of influenza deaths in one period may lead to more influenza deaths in the following period. Second, we consider cross-sectional dependence across multiple causes of death and temperature in the same period, where the magnitude, direction, and tail behaviour of the paired dependence are generally different. Third, there could be a lagged effect of cold temperatures on increased mortality risk from different causes. In other words, we would like to investigate if extreme cold weather in the previous period will have an impact on excess deaths in the current period. Therefore, non-contemporaneous cross-sectional dependence is another important component in our modeling framework.

To address these considerations, we propose a vine copula-based approach to model cause-specific excess deaths associated with extreme cold temperatures. The key idea behind vine copulas is to construct a flexible joint dependence structure using pair-copulas as bivariate building blocks (Bedford and Cooke, 2001, 2002; Aas *et al.*, 2009; Joe, 2014). Following this idea, Nagler *et al.* (2022) introduced the class of *stationary vine* (S-vine) models that guarantee stationarity under simple equality constraints on the pair-copulas. This model is ideal for simultaneously capturing cross-sectional dependence and temporal dependence, in contrast to the traditional two-step approach. In this paper, we adopt the S-vine model for the first time in the context of climate-driven mortality risk, over different geographical regions. To aid model comparison and enhance interpretability, we further propose a unified S-vine structure across all regions, based on preliminary analyses of the data. This new model is referred to as a (*stationary*) *centrally connected C-vine* (CCC-vine) model and it is introduced in Section 3.4.

The proposed S-vine and CCC-vine models can reflect many non-linear distributional effects that would be impossible to capture with conventional models (e.g., Analitis *et al.*, 2008), such as the Gaussian copula model. Additionally, a key advantage of stationary vine copula models in actuarial practice is their ability to efficiently perform complex scenario-based analyses conditioned on past or current events. The S-vine and CCC-vine models provide a single, generative model for the joint evolution of multiple causes of deaths. This enables practitioners to address a wide range of questions with any arbitrary conditioning on past or current events in a coherent and consistent manner, which is difficult to achieve using simpler alternative approaches such as empirical pairwise correlation coefficients.

Based on the fitted CCC-vine models, conditional Monte-Carlo simulations are used to generate several temperature scenarios, under which we assess the impact of extreme cold temperature on excess death from different causes as well as on the aggregated excess deaths. We also investigate the impact of extreme cold temperatures given that deaths from *Respiratory* are already at a high level. The empirical results show that across six US continental regions, *Vascular* and *Respiratory* are the two major causes of death (excluding all other causes combined) acting as the key “hubs” in the dependence structure, having relatively strong dependence with all other variables. On the other hand, *Diabetes* and *Neoplasms*

have the weakest level of cross-sectional dependence. We compare the performance of the standard S-vine model, the CCC-vine model, the Gaussian copula model, and the vector autoregression (VAR) model and find that CCC-vine models give the best overall performance. Based on the fitted CCC-vine models, we analyze the distribution of excess deaths under different climate scenarios. We find that two consecutive months of cold temperatures lead to the highest number of excess deaths. We also find that the impact of extreme cold temperatures on deaths due to *Diabetes*, *External*, *Neoplasms* is much smaller and even negligible. Among the six regions, it is concluded that CEA, SEA, and SWP have the highest number of cold-related excess deaths, while CWP and SPL have the smallest number of cold-related excess deaths. Across the six regions, MID seems to be least affected by extreme cold temperatures in terms of excess deaths.

The proposed modeling approach contributes to the existing literature in three ways. First, we are the first to utilize novel S-vine models in mortality modeling. The key strength of this new model is that it allows to generate scenario-based simulations that realistically reflect many non-linear distributional effects, such as tail risks, competing-risk, and harvesting effects³. Second, based on the S-vine models, we quantify the effect of extreme cold temperatures on excess deaths, both at the aggregate level and the cause-specific level, which makes important contributions to mortality risk modeling and management. Third, we propose the CCC-vine model which gives a unified structure across different data subgroups, focusing on relative importance of variables. This new modeling approach shows superior overall performance and enables straightforward comparison of the modeling results across different subgroups (in our case, different geographical regions). It also shows that preliminary data analysis and information pooling can enhance our understanding of the data, assist the vine copula model selection process, and in turn, improve the performance of our model. This innovation makes a valuable contribution to vine-based modeling that could be applied in other empirical settings.

The remainder of the paper is organized as follows. In Section 2, we describe and visualize the climate and cause-specific data for US from 1999–2018. Section 3 introduces S-vine copula modeling framework to model the cross-sectional and temporal dependence across temperature and causes of death. We propose novel Goodness-of-Fit tests and compare model performance in Section 4. We then analyze and discuss results from the scenario-based analysis in Section 5. Finally, Section 6 concludes. All analyses in this paper were carried out using R (R Core Team, 2022), and code is available in a GitHub repository which can be accessed at <https://github.com/tnagler/cold-mortality-svines>.

2 Data

In this section, we describe and visualize the datasets used in our empirical studies. We consider US monthly climate and mortality data over 20 years from 1999–2018.

³Harvesting effect is also referred to as “mortality displacement”, it describes the phenomenon where a compensatory decrease in mortality rates was observed in the subsequent weeks after an extreme event, suggesting that such events affect vulnerable individuals that they would have died in the short term anyway.

2.1 Actuaries Climate Index

For measurement of extreme cold temperatures, we collect the monthly T_{10} index from the Actuaries Climate Index (ACI). The ACI is developed and compiled by several actuarial professions in North America including the American Academy of Actuaries, the Casualty Actuarial Society, the Canadian Institute of Actuaries, and the Society of Actuaries. It measures the level of extreme climate and consists of six components: T_{10} (frequency of temperatures below the 10th percentile), T_{90} (frequency of temperatures above the 90th percentile), P (maximum rainfall per month in five consecutive days), D (annual maximum consecutive dry days), and W (frequency of wind speed above the 90th percentile). In particular, T_{10} is defined as

$$\frac{x_{10}}{x} \times 100,$$

where x represents the number of days in a given month and x_{10} denotes the number of days where the minimum temperature is below the 10th percentile of that particular month. To calculate relevant percentiles, the probability density function of the minimum temperature is estimated using a reference period from 1961 to 1990.

The ACI provides climate information on six continental US regions, which are illustrated in Figure 1. For each continental region, details of the states included can be found in Appendix 3 of the summary report (Actuaries Climate Index Executive Summary, 2018).

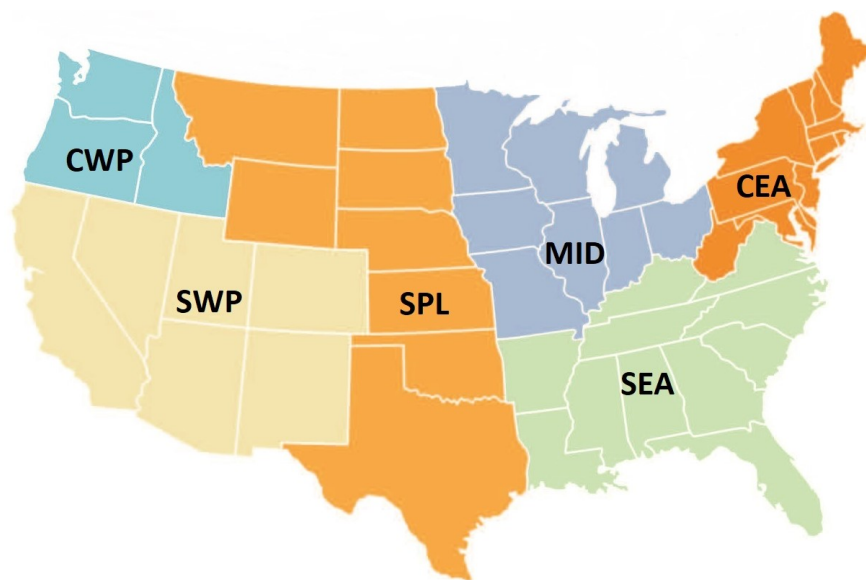


Figure 1: Six Continental US Regions. Source: Actuaries Climate Index Executive Summary, page 4 (Actuaries Climate Index Executive Summary, 2018).

Figure 2 plots the monthly T_{10} index for the six continental US regions over the period 2000–2018.⁴ We have added smooth trend lines to better visualize any underlying trends in

⁴Since we remove the trend and seasonality in monthly death series via seasonal differencing, our sample is reduced by 12 observations and thus the investigation period for the corresponding temperature index becomes 2000–2018.

the data. For each region, there is no apparent trend, seasonality, or heterogeneity observed in the T_{10} index. To assess the stationarity condition, we conduct the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test (Kwiatkowski *et al.*, 1992), the Augmented Dickey-Fuller (ADF) test (Dickey and Fuller, 1979), and the Mann-Kendall (MK) test (Mann, 1945; Kendall, 2022). We find that for all six regions, the T_{10} index fails to reject the null hypothesis of stationarity in both the KPSS test and MK test, but rejects the null hypothesis of non-stationarity in the ADF test, all at the 1% level of significance. Additionally, we conduct a Quasi-Seasonality (QS) test (Maravall, 2011; Ollech and Webel, 2020) for seasonality to confirm that there is no significant evidence of residual seasonal correlation. These results align with our visual examination of the data.

We also see from Figure 2 that there are some natural variations in the T_{10} index across the six regions. Table 1 summarizes some key statistics of T_{10} . Overall, MID tends to have the largest mean T_{10} with the highest variance and 90th percentile, which is not surprising as the Midwest has experienced quite frequent winter storms and cold waves in recent years. On the other hand, SWP has relatively stable values in T_{10} and the index rarely exceeds 20 in value, which means that normally there are less than 20% of days in a month where the minimum temperature is below the 10th percentile).

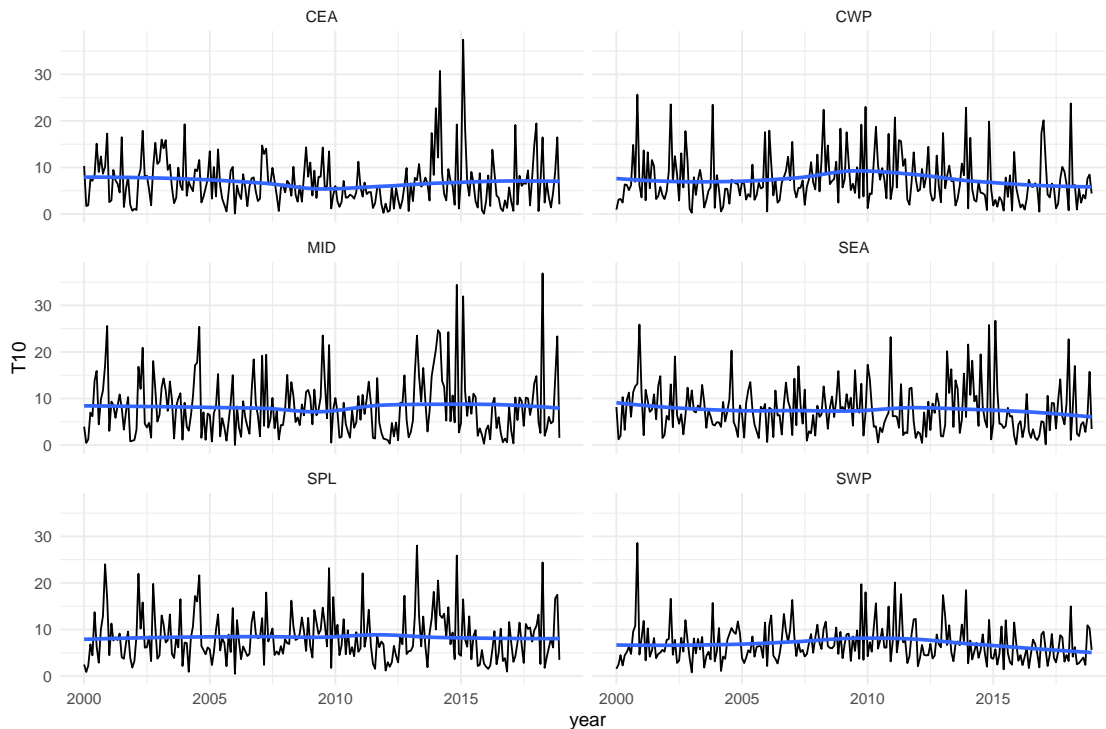


Figure 2: T_{10} index for six US regions over 2000–2018.

2.2 Cause-specific death data

We consider cause-of-death data from the CDC WONDER online query system⁵. The CDC WONDER database provides rich public health data and vital statistics in the US. We

⁵See more detail at <https://wonder.cdc.gov/ucd-icd10.html>.

Table 1: Summary statistics of T_{10} .

Region	Mean	Standard deviation	90th pctl
CEA	6.87	5.27	14.06
CWP	7.33	5.32	15.53
MID	8.24	6.51	16.64
SEA	7.53	5.16	13.97
SPL	8.31	5.17	15.10
SWP	6.89	3.99	11.68

collect monthly death counts from five major causes of death across six continental regions to be consistent with the climate data. The five major causes include *Diabetes*, *External*, *Respiratory*, *Neoplasms*, and *Vascular*, and they are classified based on the International Classification of Diseases, Version 10 (ICD-10). We also include one additional category for all remaining causes in ICD-10, which is denoted as *Other*. The detailed codifications are provided in Table 2 below.

Table 2: Codification of five major causes of death.

Cause of death	ICD-10 code
Diabetes	E10–E14
External	V01–Y89
Respiratory	J09–J98
Neoplasms	C00–D48
Vascular	I00–I78

It is widely acknowledged that elderly people are particularly vulnerable to the negative impact of climate change on health. For the US population mortality, a previous study found that cold temperatures and deaths at age 85+ exhibit the highest level of tail dependence (Li and Tang, 2022). Our research also focuses on cold-related excess deaths for the population aged 85+, or the “oldest-old”, and we further investigate the impact of extreme cold temperatures on different causes of death.⁶

In Figure 3, we plot the monthly cause-specific death counts in the six US continental regions for ages 85+. It can be seen that these death count series differ considerably in size and seasonal patterns. Except for *Vascular*, all causes show an overall upward trend over the investigation period.⁷ For *Vascular*, there is a slight downward trend in most regions in the beginning, which later stabilizes. Moreover, we observe strong seasonality in deaths due to vascular, respiratory, and other, and moderate seasonality deaths due to external causes and diabetes, where peaks in death counts usually happen in winter months (November, December, and January). On the other hand, neoplasm deaths do not seem to show strong

⁶We have also conducted the proposed modeling based on ages 75–84, and find those results to be in line with the results presented in Section 4. These additional results are available upon request.

⁷Since the number of deaths is determined by both mortality rate and population size, with decreasing mortality rate but growing population size, death counts can be increasing over time. This is particularly true for ages 85+ due to the rapidly aging population in the US

seasonal patterns compared to other causes. Across the six regions, the top 3 causes of death are *Vascular*, *Other*, and *Neoplasms*.

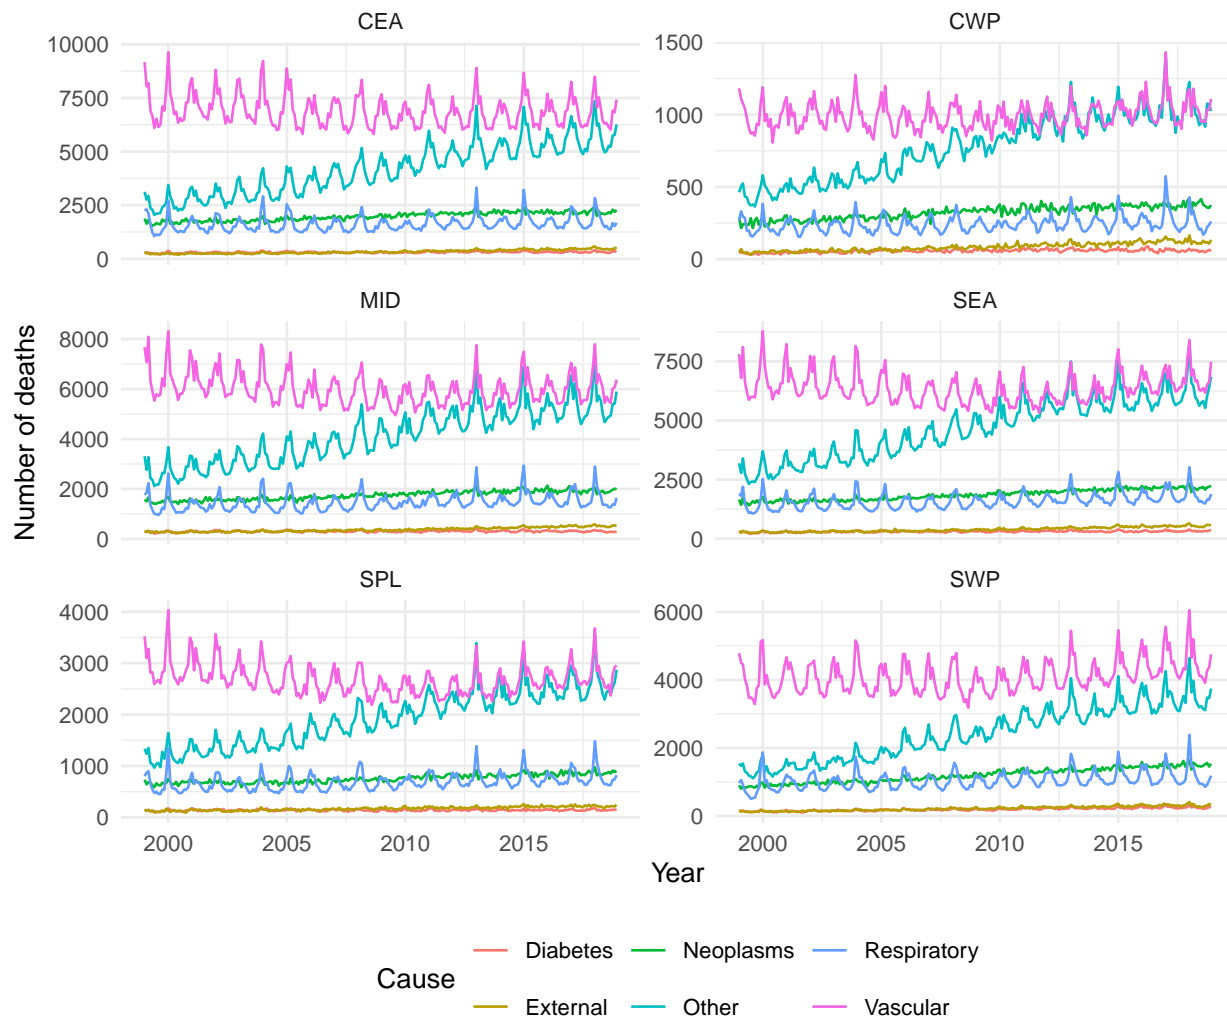


Figure 3: Monthly death counts for ages 85+ in six US regions.

The monthly death count series are not stationary. To remove the trend and seasonality in the data, we perform seasonal differencing at a lag of 12. It should be noted that the only data processing used in our research is seasonal differencing. The deseasonalized cause-specific series for CEA are plotted in Figure 4.⁸ Upon visual examination, there is no significant evidence of non-stationarity in the deseasonalized data. Similar to the T_{10} index, we conduct the three stationarity tests along with the QS test and find consistent results supporting the stationarity assumption. From Figure 4, we can see some co-movements in these deseasonalized series, particularly for *Respiratory*, *Vascular*, and *Other*, indicating the existence of tail dependence across the series. These simultaneous spikes observed across multiple causes are likely to be driven by common extreme weather and influenza events. As an example, simultaneous positive spikes can be observed during the 2012–2013 and 2014–2015 flu seasons in causes *Respiratory*, *Vascular*, and *Other*. According to the CDC Morbidity

⁸Plots for the other five continental regions can be found in the supplementary material, Figures S.1–S.5.

and Mortality Weekly Report (MMWR, Brammer *et al.*, 2013), the 2012–13 flu season recorded the highest number of influenza-associated pediatric deaths reported to the CDC since data collection began in 2004 (Excluding the 2009 H1N1 pandemic). During the 2014–15 influenza season, the proportion of deaths attributed to pneumonia and influenza remained above the epidemic threshold for eight consecutive weeks, from January 3 to February 21, 2015 (MMWR, Appiah *et al.*, 2015). These reports align with published influenza statistics for the United States.⁹ It should be noted that influenza epidemics are likely to coincide with periods of extreme cold weather, as viruses spread more easily at lower temperatures, and exposure to cold air can weaken the immune system (Donaldson and Keatinge, 2002; Dushoff *et al.*, 2006). Consequently, extreme cold temperatures can lead to substantial short-term increases in mortality not only from *Respiratory* but also from other causes, such as *Vascular* and *External*. The original monthly death counts plotted in Figure 3 also confirm the presence of genuine simultaneous spikes across causes, indicating tail dependence.

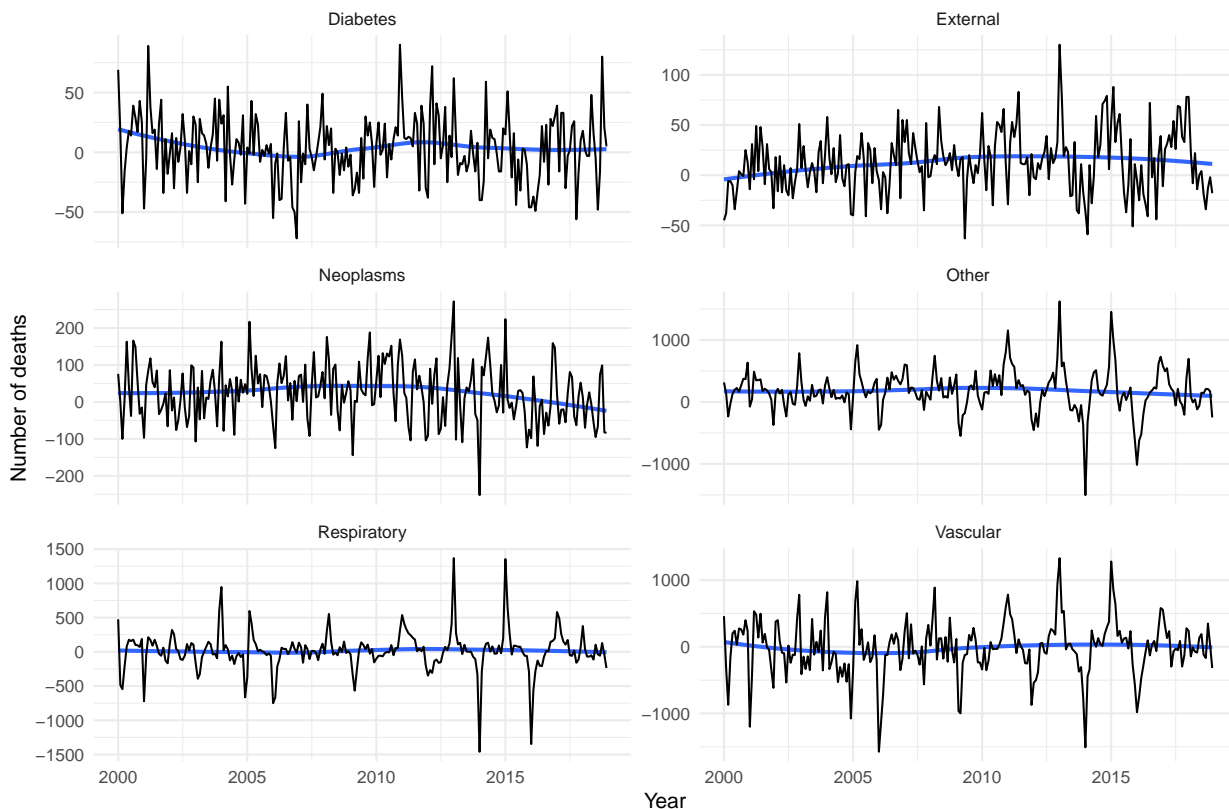


Figure 4: Deseasonalized cause-specific monthly death counts for ages 85+ in CEA.

The deseasonalized monthly death data will be used to produce results in Section 4. Since our copula model operates on the u -scale (*i.e.*, the unit interval $[0, 1]$), the deseasonalized time series must be transformed accordingly. Specifically, we transform the deseasonalized death series x_i using the marginal distribution function F_i , estimated via univariate kernel density estimation. This results in $u_i = \hat{F}_i(x_i)$ for each cause $i = 1, \dots, d$. We refer to u_i as our copula data in Section 3.

⁹see more details at https://en.wikipedia.org/wiki/United_States_influenza_statistics_by_flu_season.

3 Stationary vine copula models

3.1 Copulas and vines

The literature on copula-based modeling goes back to the late 1950s. Sklar (1959) proved that any multivariate distributions can be broken down into univariate marginal distributions and a copula, which describes the dependence structure between these univariate distributions. Mathematically, Sklar’s theorem can be described as follows.

Theorem 1 *Let $\mathbf{X} = (X_1, \dots, X_d)$ be an absolutely continuous random vector of dimension d with joint distribution function F and marginal distribution functions F_i , for $i = 1, \dots, d$, the joint distribution function can be expressed as*

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)), \quad (1)$$

with associated density function

$$f(x_1, \dots, x_d) = c(F_1(x_1), \dots, F_d(x_d)) \prod_{i=1}^d f_i(x_i), \quad (2)$$

where C is a copula function with density c , and f_i is the corresponding density for F_i . The copula C is the joint distribution function of the random vector

$$\mathbf{U} = (F_1(X_1), \dots, F_d(X_d)).$$

Parametrized d -dimensional copula functions that are frequently considered include the Gaussian copula, t copula, and Archimedean copulas (see more details in Czado, 2019). However, for high dimensional problems, these copulas suffer from a lack of flexibility and unrealistic symmetry constraints. In Figure 5, we illustrate the pairwise dependence structure of copula data based on temperature and death series in CEA. We can see that different causes of death generally exhibit different dependence structures. Overall, *Vascular* has the highest Kendall’s τ with other causes and the temperature index $T10$. On the other hand, the dependence structure of $T10$ with certain causes such as *Diabetes* and *External* is relatively weak. In the lower triangle, we plot the empirical pairwise contour plots of the copula data transformed to standard normal margins (Czado, 2019, Section 3.8). Although some of these contour plots look rather elliptical and, thus, compatible with the Gaussian copula, other plots suggest tail asymmetry and stronger tail dependence than a Gaussian copula. For example, the contour plot between *Vascular* and *Neoplasm* seems to suggest stronger upper tail dependence and weaker lower tail dependence. It should be noted that the pairwise plots in Figure 5 ignore the serial dependence in the data, which may also exhibit non-Gaussian dependence.

3.2 Vine copulas

To alleviate the issues of asymmetrical tail dependence and a lack of flexibility in high-dimensional copulas, Joe (1996) proposed a flexible construction of a multivariate copula using bivariate building blocks. Following this idea, Bedford and Cooke (2001, 2002) showed

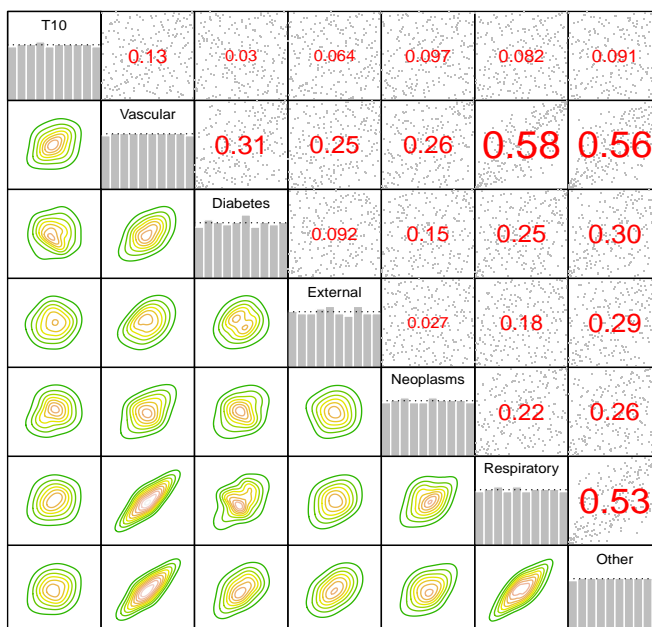


Figure 5: Pairwise copula data for CEA. Upper triangle: scatter plots of copula data with estimated Kendall’s τ ; Diagonal: marginal histograms of copula data; Lower triangle: empirical contour plots of marginally normalized copula data.

how a multivariate distribution can be decomposed into bivariate copulas (some corresponding to conditional dependence) and marginal densities. For organising all possible decompositions, they also developed a graphical model called *vine*. Unlike multivariate Archimedean copulas, vine copulas allow for a different direction and magnitude of dependence for each pair of variables. Unlike d -dimensional Gaussian and t copulas, vine copulas can assign a different type of dependence in each pair of causes. Vine copulas also allow for asymmetric tail dependence between variables. Vine copula-based modeling has been used in a wide range of areas such as finance (Aas, 2016), engineering (Schepsmeier and Czado, 2016; Coblenz *et al.*, 2020), and environmental sciences (Vernieuwe *et al.*, 2015; Kreuzer *et al.*, 2022). Given the empirical pairwise dependence structure illustrated in Figure 5, a vine copula-based approach is particularly suitable for our modeling purpose.

Definition 3.1 A regular vine (R -vine) is a sequence of trees $(T_k)_{k=1}^{d-1}$ with the following properties:

- (i) T_1 is a tree with vertices $V_1 = \{1, 2, \dots, d\}$ and edges E_1 ,
- (ii) for $k = 2, \dots, d - 1$, T_k is a tree with vertices $V_k = E_{k-1}$,
- (iii) (proximity condition) for $k = 2, \dots, d - 1$: if vertices $a, b \in V_k$ in T_k are connected by an edge $e \in E_k$, then the corresponding edges in T_{k-1} , $a = \{a_1, a_2\}$, $b = \{b_1, b_2\} \in E_{k-1}$, must share a common vertex: $|a \cap b| = 1$.

Each edge $e \in E_k$ in the graph is given a unique label $(a_e, b_e | D_e)$, where $a_e, b_e \in \{1, \dots, d\}$ and $D_e \subset \{1, \dots, d\} \setminus \{a_e, b_e\}$ with $|D_e| = k - 1$.

A vine copula model assigns a bivariate copula to each edge of this graph and constructs the copula density as

$$c(\mathbf{u}) = \prod_{k=1}^{d-1} \prod_{e \in E_k} c_{a_e, b_e; D_e}(u_{a_e | D_e}, u_{b_e | D_e} | \mathbf{u}_{D_e}),$$

where $u_{a_e | D_e} := C_{a_e | D_e}(u_{a_e} | \mathbf{u}_{D_e})$, $\mathbf{u}_{D_e} := (u_l)_{l \in D_e}$ is a subvector of $\mathbf{u} = (u_1, \dots, u_d) \in [0, 1]^d$ and, $C_{a_e | D_e}$ is the conditional distribution of U_{a_e} given \mathbf{U}_{D_e} . In this model, each *pair-copula* $c_{a_e, b_e; D_e}$ captures the dependence between variables X_{a_e} and X_{b_e} , conditional on the set $\{X_k, k \in D_e\}$, more precisely $c_{a_e, b_e; D_e}$ is the copula density associated with the bivariate conditional distribution $(X_{a_e}, X_{b_e}) | D_e$.

Figure 6 shows an example of an R-vine tree sequence on three variables. The nodes in the first tree represent the three variables A , B , and C . The edges are identified with bivariate pair-copulas, which describe the dependence between each pair of variables. In the second tree, the nodes are the edges of the first tree. The edges describe the dependence between node AB and BC conditional on B . The flexibility of vine copulas allows for a different direction and magnitude of dependence in each pair of variables. For example, AB could be assigned a copula with upper tail dependence (e.g., Gumbel), BC could be assigned a copula with lower tail dependence (e.g., Clayton) and $AC|B$ could be assigned a copula with no tail dependence (e.g., Gaussian).




Figure 6: Example of a 3-dimensional R-vine tree sequence.

There are two important subclasses of vine copula models, namely D-vine and C-vine models. In a D-vine copula, all trees in the vine are paths. This is appropriate when the variables have a natural ordering. In a C-vine copula, all trees are stars¹⁰, which is most natural when a certain variable drives the dependence among others. The sequence of root nodes of each tree level is called *order* of the C-vine. General R-vines allow for much more flexibility, however. For $d = 3$, any R-vine is both a D- and C-vine; differences only appear in higher dimension. Inference and simulation algorithms for such models were developed by Aas *et al.* (2009) and Dissmann *et al.* (2013).

¹⁰A “star” is a type of tree where one central node is directly connected to all other nodes, which are not connected to one another.

3.3 Stationary vine copula models

Brechmann and Czado (2015), Smith (2015), and Beare and Seo (2015) pioneered the use of vine copula models for multivariate time series. A multivariate time series is a sequence of random vectors $\mathbf{X}_1, \dots, \mathbf{X}_T \in \mathbb{R}^d$ that may exhibit both cross-sectional (within \mathbf{X}_t) and serial (across \mathbf{X}_t and \mathbf{X}_s) dependence. A vine copula model for this series consists of a large vine graph that treats each $X_{t,j}$ as a variable and assigns a bivariate copula to each edge. As is common in time series modeling, the working assumption is that the data are stationary, *i.e.*, the marginal and joint distributions do not change over the observation period. As a consequence, many of the pair-copulas in the model must be the same, which greatly reduces model complexity. For example, because the joint distribution of $(\mathbf{X}_t, \mathbf{X}_s)$ equals that of $(\mathbf{X}_{t+\tau}, \mathbf{X}_{s+\tau})$ for any $\tau \in N$, the copulas characterizing the dependence in $(\mathbf{X}_t, \mathbf{X}_s)$ must equal those characterizing $(\mathbf{X}_{t+\tau}, \mathbf{X}_{s+\tau})$. A second simplification arises from the assumption that the series is Markovian of order p : the behavior of \mathbf{X}_t only depends on the last p realizations $\mathbf{X}_{t-1}, \dots, \mathbf{X}_{t-p}$. In a corresponding vine copula model for $(\mathbf{X}_1, \dots, \mathbf{X}_T)$, this assumption induces many independence copulas, reducing complexity even further. In particular, all conditional copulas between variables $X_{t,j}$ and $X_{t+p+1,k}$ (conditional on the variables observed between times $t+1$ and $t+p$) must be independence copulas (Nagler *et al.*, 2022, Theorem 3).

Following these ideas, Nagler *et al.* (2022) proposed the class of *stationary vine* (S-vine) models that guarantee stationarity under simple equality constraints on the pair-copulas, which is ideal for capturing both cross-sectional dependence and time dependence. S-vines are constructed from a d -dimensional R-vine (called *cross-sectional structure*), connecting variables $X_{t,1}, \dots, X_{t,d}$ within a fixed time point. This structure is replicated at all time points and connected serially in a way that the overall graph remains a vine. Nagler *et al.* (2022) showed that this amounts to specifying an order in which individual margins are serially connected.

The models of Smith (2015) and Beare and Seo (2015) correspond to special S-vines, where the cross-sectional structure is a D-vine. For more details and related theoretical results, we refer to Nagler *et al.* (2022).

Definition 3.2 *An S-vine copula model consists of the following components:*

- *Models for the stationary marginal distributions $F_j(x) = P(X_{t,j} \leq x)$, $j = 1, \dots, d$.*
- *The S-vine graph, characterized by a cross-sectional vine structure (T_1, \dots, T_{d-1}) and an order for constructing serial connections.*
- *Bivariate pair-copulas assigned to the edges of the graph.*

Figure 7 provides an example of a 4-dimensional S-vine for three time points. Models for the marginal distributions and pair-copulas can be fitted using maximum-likelihood methods and selected using standard model selection criteria, see Nagler *et al.* (2022).

S-vine copula models are particularly suited to model cause-of-death data. Extreme death counts are expected to exhibit strong and non-linear dependencies, which S-vines models

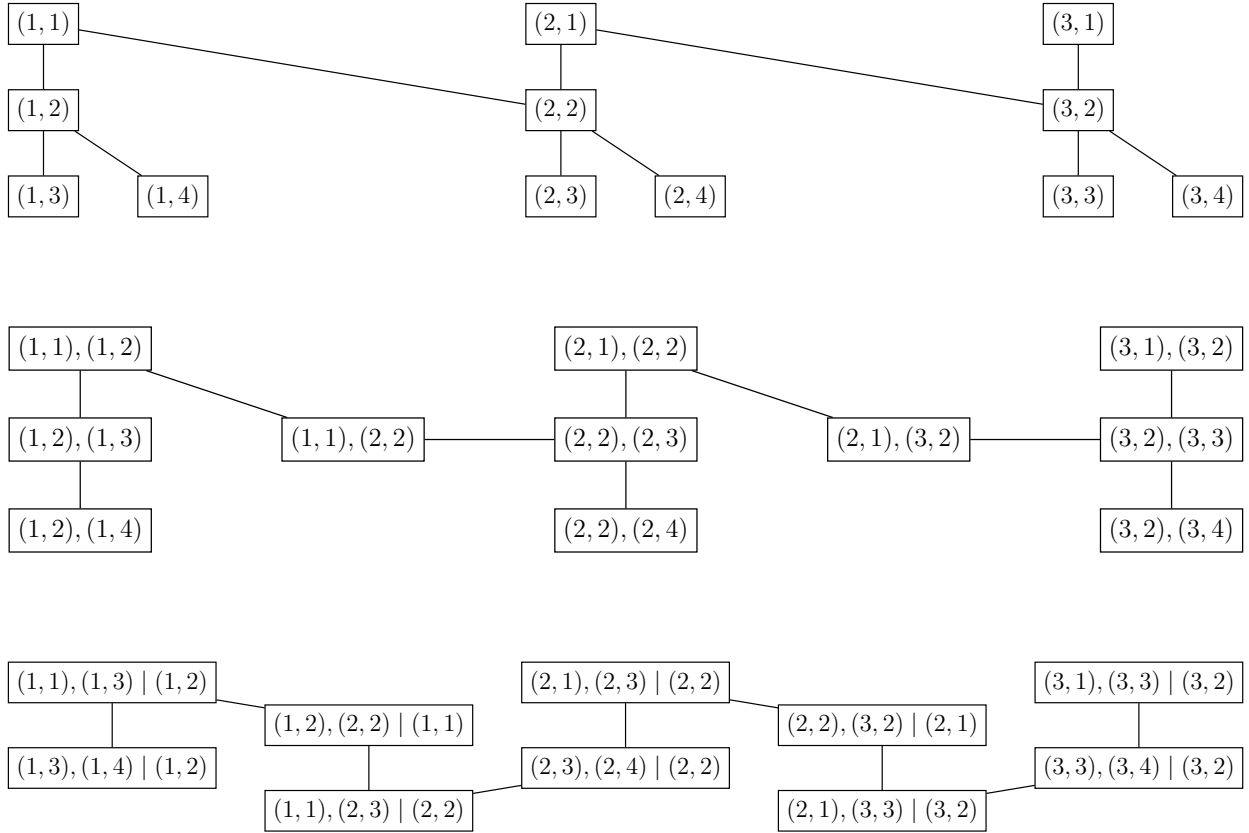


Figure 7: First three trees level of a four-dimensional S-vine on three time points. The first entry of (t, j) is the time index, while the second entry corresponds to the j th variable.

are designed to capture. We are also interested in their implications in various scenarios of extreme climate conditions. Such effects are easy to analyze via conditional Monte-Carlo simulation from S-vine models, see Nagler *et al.* (2022, Section S.3).

3.4 Stationary centrally connected C-vine models

A key ingredient to any vine copula model is the choice of the graph structure. Because of the huge number of options, an exhaustive search is computationally infeasible (see, Czado and Nagler, 2022, Section 4). Nagler *et al.* (2022) proposed a heuristic adapted from the cross-sectional vine model selection of Dissmann *et al.* (2013) that aims at maximizing dependence strength at low tree levels. In our empirical analysis, we can fit several S-vine copulas – one for each geographic region. The heuristic approach would then deliver models that differ in structure. This makes cross-regional comparison difficult.

To tackle this issue, we therefore propose a unified vine structure for all regions, based on domain knowledge and a preliminary analysis of the data. The proposed model is an S-vine model with centrally connected C-vines as cross-sectional structures, and thus is referred to as a (stationary) centrally connected C-vine (CCC-vine) model. Key features of the CCC-vine model are described in the following.

Definition 3.3 An CCC-vine copula model consists of the following components:

- The contemporaneous cross-sectional dependence structure is a C-vine.
- The C-vine is centrally connected, i.e. the central node of the C-vine is connected with its value of the previous time period.
- The order of the root nodes of the tree levels is selected based on the importance of variables.

For illustration purposes, in Figure 8, we plot the first three trees of a four-dimensional CCC-vine on two successive time points.

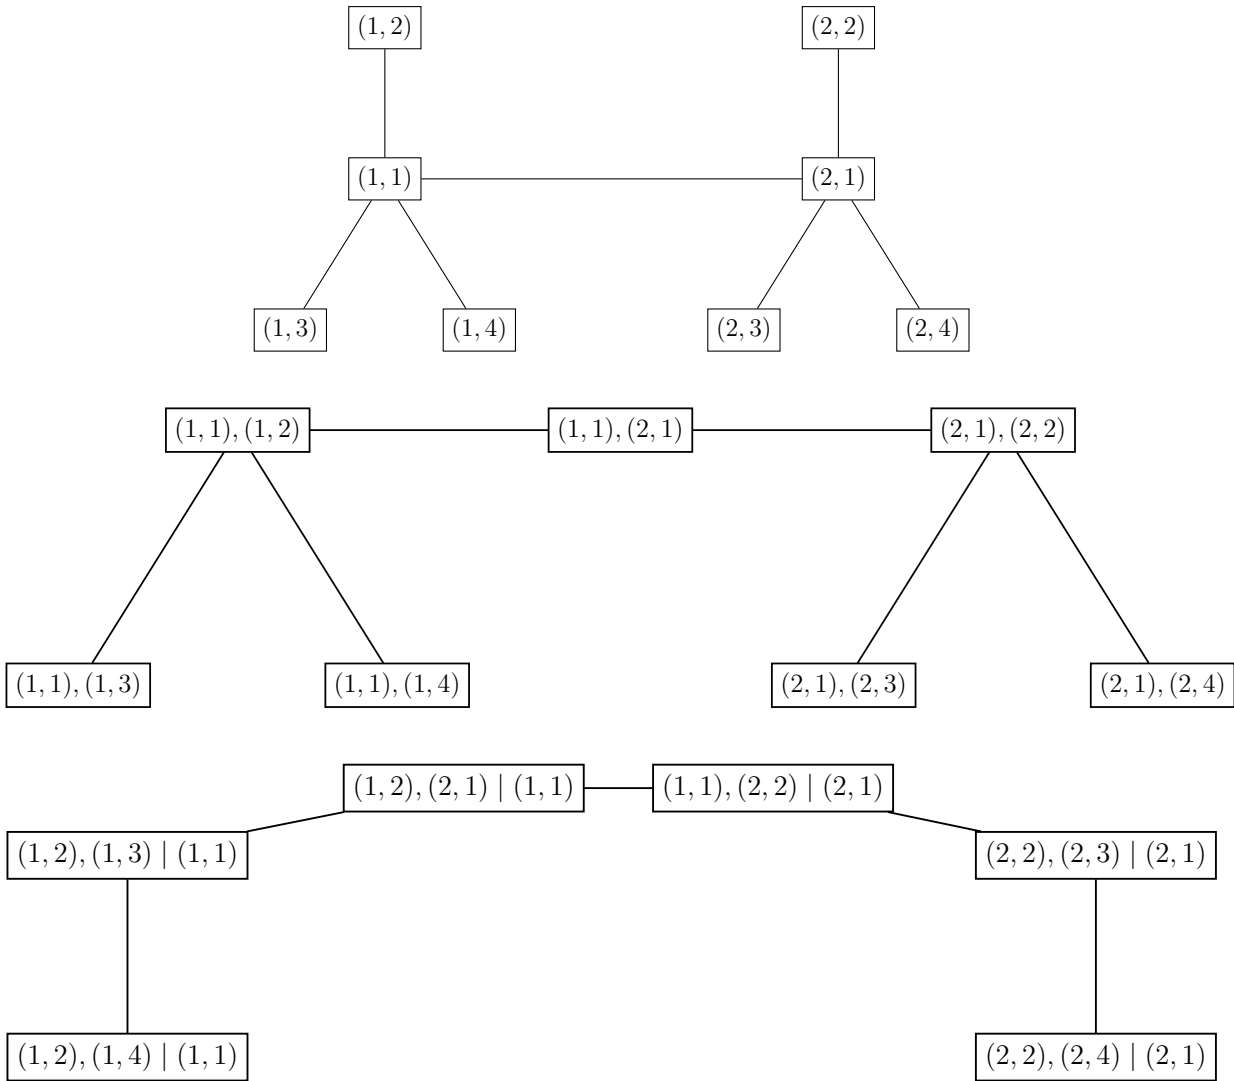


Figure 8: First three trees level of a four-dimensional S-vine on two time points with a cross-section C-vine. The first entry of (t, j) is the time index, while the second entry corresponds to the j th variable.

The CCC-vine is a special case of an S-vine model, so it inherits all of its benefits: the ability to capture complex non-linear dependencies as well as the ability to easily generate scenario-based simulations. Its distinguishing feature compared with other S-vines is its focus on relative importance of variables. In the first trees of the model, dependencies involving “important” variables are prioritized. In the exemplary CCC-vine in Figure 8, the most important variable would be $X_{t,1}$. The first tree contains copulas for all pair-wise dependencies between $X_{t,1}$ and $X_{t,j}$, $j = 2, \dots, 4$, as well as the temporal auto-correlation between $X_{t,1}$ and $X_{t+1,1}$. Importantly, no conditioning is involved, making the interpretation of pair-wise dependencies easy. With each further tree level, we must add one level of conditioning, so pair-copulas get increasingly difficult to interpret. This prioritization has a related secondary effect on the quality of model fit and inferences. Because pair-copulas are estimated sequentially tree-by-tree, dependencies involving the most important variables suffer less from error propagation.

The vine structure selection in the CCC-vine involves a simple function and does not lead to an increase in computation time. By pre-determining the vine copula structure in each tree, the calibration of the CCC-vine eliminates the need to search for alternative structures and therefore does not increase model complexity compared to the standard S-vine.

In our empirical studies, the cross-sectional C-vine structure is specified to reflect the importance (or “connectedness”) of different causes and the temperature index. To determine the ordering, in Table 3 we present the variable rankings according to aggregated strength of dependence for each region (measured by the summation of pair-wise Kendall’s τ across six causes of death and $T10$, illustrated in Figure 5 as an example for CEA). We then compute the overall rankings of each variable across the six regions. Therefore, the C-vine cross-sectional order is defined as $\{Vascular, Other, Respiratory, External, Diabetes, Neoplasms, T10\}$. Since this ordering reflects importance, it is then natural to also assign serial connections in the same order. It should be noted that the order of causes of death is also roughly in line with their relative frequency. The CCC-vine structure therefore prioritizes the most common causes, making the model more reliable when looking at aggregate death counts. Before fitting the S-vine copula, we select appropriate parametric models for marginal distributions based on AIC. The marginal models for each cause and region are provided in Table A.1.

Table 3: Variable rankings according to aggregated Kendall’s τ .

Region	$T10$	Diabetes	External	Neoplasms	Other	Respiratory	Vascular
CEA	7	4	6	5	2	3	1
CWP	6	5	4	7	2	1	3
MID	7	4	5	6	1	3	2
SEA	7	6	4	5	2	3	1
SPL	7	6	5	4	2	1	3
SWP	7	4	5	6	2	3	1
Total score (overall rankings)	41(7)	29(5)	29(4)	33(6)	11(2)	14(3)	11(1)

4 Goodness-of-Fit tests and model comparison

There are various ways to extend standard Goodness-of-Fit tests for copula models from the *iid* to a time series setting. Some methods have already appeared in the literature, but either do not match well with our model setup (Rémillard, 2017; Berghaus and Bücher, 2017). Instead we employ a strategy that exploits the Rosenblatt transform (Genest *et al.*, 2009, Section 4) and Markovian model structure of our models. The procedure is explained in more detail in the following.

Suppose $\mathbf{X}_1, \mathbf{X}_2, \dots \in \mathbb{R}^d$ is a stationary time series with Markov order 1 and suppose we have a model F_θ for the conditional distribution of \mathbf{X}_t given \mathbf{X}_{t-1} . Define the *pseudo-residual* $\mathbf{U}_t = (U_{t,1}, \dots, U_{t,d}) \in \mathbb{R}^d, t = 2, \dots, n$, via the conditional Rosenblatt transformation

$$\begin{aligned} U_{t,1} &= F_\theta(X_{t,1} \mid \mathbf{X}_{t-1}), \\ U_{t,j} &= F_\theta(X_{t,j} \mid \mathbf{X}_{t-1}, X_{t,1}, \dots, X_{t,j-1}), \quad j = 2, \dots, d. \end{aligned}$$

Note that the Rosenblatt transform is easily computed for S-vine models, see Nagler *et al.* (2022, Section S3.2). For the VAR model, it corresponds to the marginal probability integral transform of decorrelated residuals.

Let F be the true distribution of the data. Under the null-hypothesis, $F = F_\theta$, each vector $\mathbf{V}_t = (\mathbf{U}_t, \mathbf{U}_{t+1})$ follows a uniform distribution on $[0, 1]^{2d}$. Testing for this is essentially equivalent to testing independence in a $2d$ -dimensional random vector. A slight complication is that the sequence $(\mathbf{V}_t)_{t=1}^{n-1}$ is not independent because \mathbf{U}_t is contained in both \mathbf{V}_{t-1} and \mathbf{V}_t . This issue will disappear after a data splitting step discussed later, so let us assume that $(\mathbf{V}_t)_{t=1}^n$ is independent under the null for the moment. We use the test of Genest *et al.* (2007) based on the Möbius transform of the empirical copula process. More specifically, for any $A \subset \{1, \dots, 2d\}$, define

$$\mathbb{G}_{A,n}(\mathbf{v}) = \frac{1}{\sqrt{n}} \sum_{t=1}^n \prod_{j \in A} [1\{V_{t,j} \leq v_j\} - v_j].$$

Under the null hypothesis of independence, the collection of $\mathbb{G}_{A,n}, A \subset \{1, \dots, 2d\}$, converges jointly to a certain Gaussian process. Each $\mathbb{G}_{A,n}$ can therefore be used to test independence in the components of $(\mathbf{V}_{t,j})_{j \in A}$. Moreover, Genest *et al.* (2007) have shown that $\mathbb{G}_{A,n}$ and $\mathbb{G}_{A',n}$ are asymptotically independent whenever $A \neq A'$. This makes it possible to aggregate p -values of the tests using Fisher's method (Littell and Folks, 1971), which relies on the independence between individual statistics. Since in our case, $2d = 14$ is relatively large, a full test for multivariate independence over all subsets would lack power (Genest *et al.*, 2007, Section 5). We therefore construct the statistic from all subsets of cardinality four or smaller.

Finally, we have to account for the fact that the parameter θ has been estimated from the same data that is used for testing. We employ the generic method of Braun (1980): first split the data into m subsets, then perform a test on each subset, and finally compute the aggregate p -value $p = 1 - (1 - \min(p_1, \dots, p_m))^m$. More precisely, we assign the vector \mathbf{V}_t to the k th subset whenever $1 + t \pmod{m} = k$, which ensures that no \mathbf{U}_t is contained in

more than one \mathbf{V}_i in each subset. The vectors in each subset are then indeed *iid* standard uniform under the null. The reported p -values all use $m = 5$, but our results were found to be relatively insensitive to this choice. The same adjustment is used for the Cramér-von-Mises tests. Note that for the Goodness-of-Fit tests, the marginal models used are those listed in Table A.1, with parameters estimated using Maximum Likelihood Estimation.

Using the proposed Goodness-of-Fit tests, we compare the model performance of the CCC-vine, a standard S-vine with heuristic structure selection, a Gaussian copula model, as well as the VAR model¹¹, all with Markov order equal to 1.¹² Table 4 shows p -values of goodness-of-fit tests for the four models. The values for the CCC-vine and S-vine models are generally quite high and would not lead to rejections at the 5% level in any region. The Gaussian copula is rejected only in CWP, while the VAR model would be rejected in all six regions.

Table 4: Goodness-of-fit p -values.

	CEA	CWP	MID	SEA	SPL	SWP
CCC-vine	0.90	0.07	0.91	0.88	0.99	0.26
S-vine	0.89	0.49	0.60	0.91	0.90	0.62
Gaussian copula	0.89	0.00	0.92	0.25	0.80	0.24
VAR	0.03	0.00	0.04	0.00	0.01	0.00

The superior performance of the CCC-vine is further confirmed in Table 5, where the smallest AIC and BIC values are highlighted in bold for each region. It should be noted that the AIC and BIC are measured on the original scale rather than on the copula data scale. We also include the number of parameters (“df”) used in each model. It can be seen that, in 4 out of 6 cases, the CCC-vine models outperform the other three methods in both AIC and BIC. Overall, the standard S-vine models only provide slightly worse results compared to the CCC-vine models, followed by the VAR models and the Gaussian copula models. This is not surprising as the VAR and Gaussian copula models do not account for conditional heteroscedasticity and tail behaviour of the data. Across the three copula models, the CCC-vine and the standard S-vine have a very similar number of parameters, whereas the Gaussian copula model requires noticeably more parameters. In contrast, the VAR model has the fewest parameters; however, it also exhibits worse AIC and BIC values compared to both the CCC-vine and standard S-vine models. These observations are consistent with our earlier conclusions, highlighting that the CCC-vine provides strong goodness-of-fit without increasing model complexity, and effectively captures the dependence structure in the data.

To better understand and visualize the dependence structure across multiple causes and extreme temperature index, we plot the first two trees of the CCC-vine models in Figures 9 and 10 for CEA.¹³ Fitted copula families and their strength of dependence (in terms of

¹¹For the VAR model, we use a probit transformation for $T10$ to deal with its bounded domain.

¹²The choice of Markov order 1 is motivated by the analysis in the following section. There, we always focus on two subsequent months. Because of the nested structure of stationary vines, including more than one lag in the overall model would not affect any results.

¹³The CCC-vine structures for the remaining regions are included in the supplementary materials, Figures S.6–S.15.

Table 5: AIC and BIC comparison across different models.

Region	Criterion	CCC-vine	Standard S-vine	Gaussian copula	VAR
CEA	AIC	17023	17039	17165	17181
	BIC	17236	17245	17460	17373
	df	62	60	86	56
CWP	AIC	14165	14146	14283	14221
	BIC	14353	14335	14578	14412
	df	55	55	86	56
MID	AIC	17180	17215	17322	17285
	BIC	17406	17427	17620	17477
	df	66	67	87	56
SEA	AIC	16997	17007	17164	17140
	BIC	17234	17240	17462	17332
	df	71	68	87	56
SPL	AIC	15770	15770	15914	15870
	BIC	15973	15959	16209	16061
	df	59	55	86	56
SWP	AIC	16098	16124	16280	16206
	BIC	16328	16357	16579	16398
	df	67	68	87	56

Kendall's τ) are shown as edge labels. For more information about the commonly available choices of pair-copula families, we refer to Joe (2014) and Czado (2019). In Figure 9, for the first tree of the CCC-vine, we observe strong dependence between causes *Vascular* and *Respiratory* (Kendall's τ of 0.58), and *Vascular* and *Other* (Kendall's τ of 0.56). On the other hand, the dependence is relatively weaker between *Vascular* and *Neoplasms* (Kendall's τ of 0.3), and *Vascular* and *External* (Kendall's τ of 0.25). These results are largely consistent with the pairwise copula data shown in Figure 5. However, it is important to note that Figure 5 does not account for the serial dependence in the variables. In terms of the choice of copula, the BB7 copula is selected for four pairs, indicating the need to account for complex forms of dependence. The BB7 copula (see Joe, 2014, p. 202) combines the behavior of Joe and Clayton copulas, and allows for various forms of asymmetric tail behavior. For example, we can see that the pairwise plot for *Vascular* and *Neoplasms* in Figure 5 exhibits positive tail dependence, as shown by the non-elliptical contour plot with a greater concentration in the upper corner. The contour plots of two fitted BB7 copulas (*Vascular-Neoplasms* and *Vascular-Other*) can be seen in Figure 11. Both copulas show small, but notable differences in tail behavior. The shapes of the upper and lower tail of *Vascular-Neoplasms* are quite different, although the tail-dependence coefficients are almost equal (lower: 0.29, upper: 0.28). For *Vascular-Other*, the shapes are more similar, but the upper tail-dependence coefficient for *Vascular-Other* is 0.67 and the lower tail-dependence coefficient is 0.57. The Gaussian copula is only selected for *Vascular-External*, and *Vascular-Diabetes*, which is also consistent with the elliptical shapes of the contour plots observed in Figure 5.

Tree 1

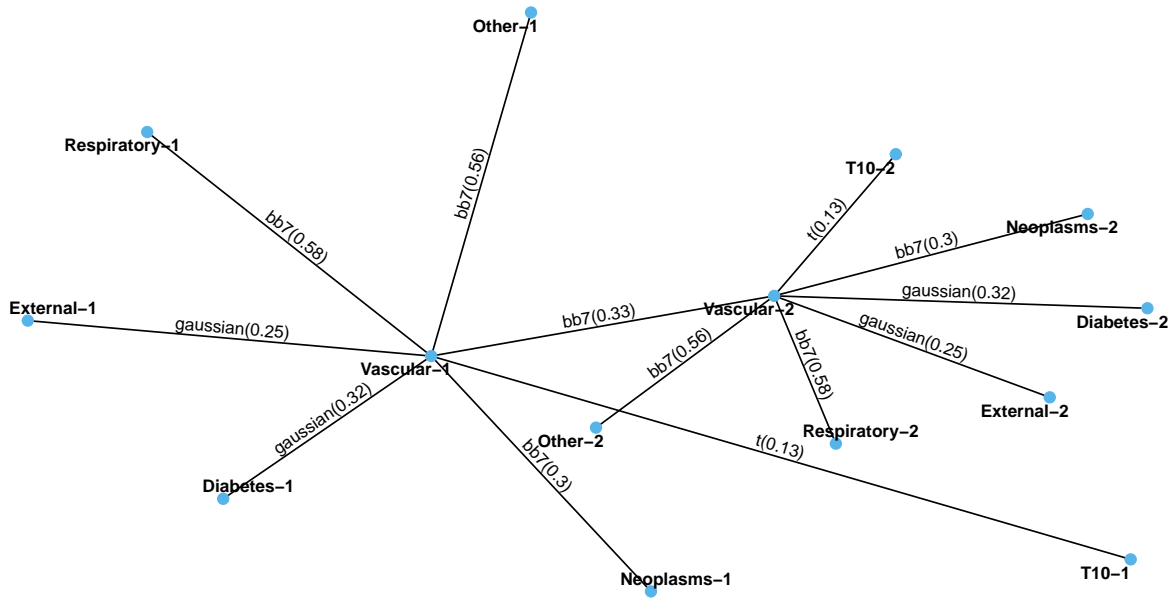


Figure 9: First tree of the selected stationary CCC-vine structure for CEA.

Tree 2

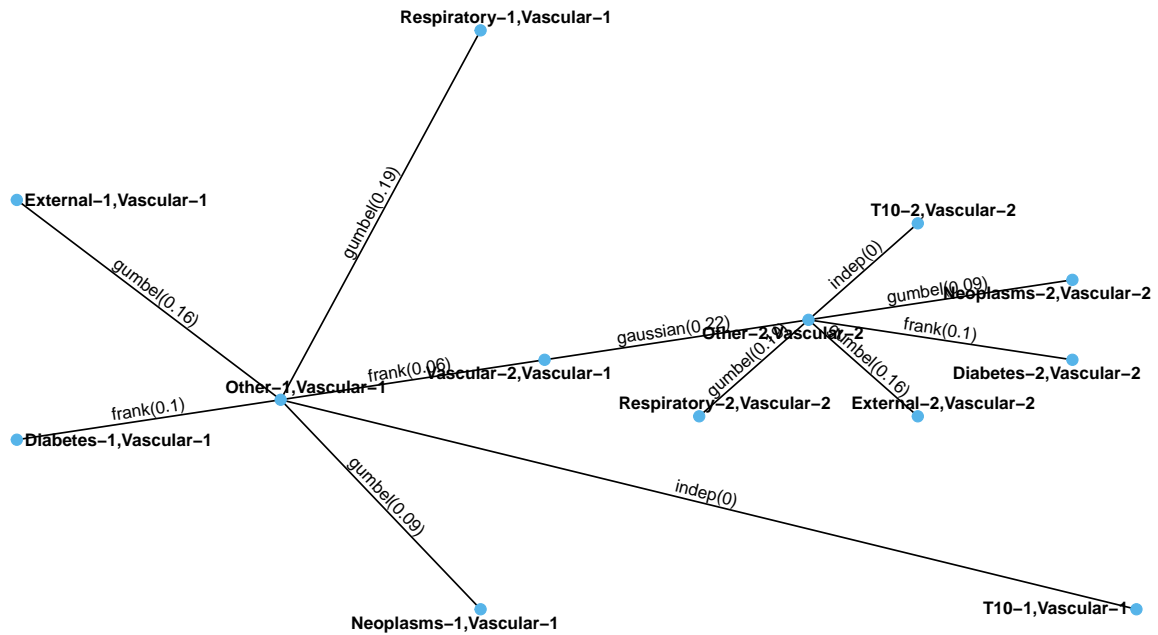


Figure 10: Second tree of the selected stationary CCC-vine structure for CEA.

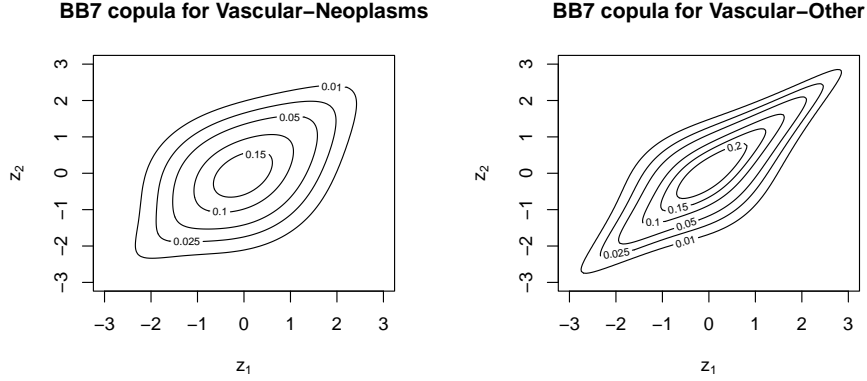


Figure 11: Two examples of fitted BB7 copulas in the CCC-vine model for region CEA.

A key advantage of the CCC-vine over the standard S-vine is its consistent and simpler model structure across all datasets based on preliminary analysis of the data, which enables easier interpretation of modeling results. For the CCC-vine models, the copula structure is the same at each tree level across the six regions, with differences only in the marginal models and copula types. Therefore, the CCC-vine model facilitates a clearer understanding of the output, capturing both serial dependence and cross-sectional dependence. For example, Figure 9 and Figure S.6 of the supplementary material show an identical structure in the CCC-vine model for the first tree for the CEA and CWP regions. Meanwhile, the constituent pair copulas and dependence parameters are allowed to differ across the two regions. The common structure shows that serial and cross-sectional dependence occur through *Vascular* in the first tree. For CEA, we observe strong dependence between *Vascular* and *Respiratory* (Kendall’s $\tau = 0.58$), and between *Vascular* and *Other* (Kendall’s $\tau = 0.56$). For CWP, the dependence is generally weaker, with the strongest relationship also between *Vascular* and *Respiratory* (Kendall’s $\tau = 0.25$), captured by the same copula type, BB7. We also note that in the first tree, the serial dependence of CEA is modeled by a BB7 copula, whereas for CWP it is captured by a Gaussian copula. In contrast, the S-vine copula model requires fitting separate copulas for each geographic region. Finding the “optimal” vine tree structure for each region results in different S-vine models, as demonstrated in Figure A.1 of the Appendix. This makes cross-regional comparisons difficult. As a result, no clear conclusions can be drawn from these differing S-vine structures. This simple example illustrates the advantage of having a unified and consistent copula tree structure across different datasets.

In summary, by comparing model complexity, interpretability, and goodness-of-fit and forecast accuracy¹⁴, we conclude that the CCC-vine offers a simpler and unified model structure, with results that are easier to interpret, without compromising short-term forecasting per-

¹⁴While short-term forecast accuracy is not the primary focus of our research, we conducted additional back-testing exercises. We found that the CCC-vine achieves a comparable level of accuracy to the standard S-vine model, while outperforming the VAR model. These additional results are presented in Figure S.20 of the supplementary material. However, it should be noted that the power of these tests is limited due to the small sample size: the out-of-sample evaluation is based on only 80 observations, with 160 used for fitting. Consequently, the results should be interpreted with caution and may be inconclusive, as the limited data constrains the strength of any statistical tests.

formance. Importantly, our study is the first to apply S-vine models in the context of quantifying the impact of extreme climate events on cause-specific mortality, thereby extending their use to a novel and practically significant domain.

5 Scenario-based analysis

The key contribution of our modeling approach, as well as the focus of the research, is scenario-based analysis using historical data, rather than producing short-term mortality forecasts. Accurate mortality forecasting that incorporates a climate component is challenging because it relies on precise predictions of climate variables, which are often unavailable or highly uncertain. In contrast, our goal is to sufficiently capture the temporal dependence, contemporaneous cross-sectional dependence, and non-contemporaneous cross-sectional dependence present in historical mortality and climate data, under a single copula framework. This allows us to generate scenarios beyond the observed data and simulate situations where temperatures extremes exceed certain thresholds.

From the perspective of actuarial practitioners and insurers, when assessing climate-driven mortality, the key question is often not the mean mortality forecast over the coming months, but rather the potential risks under extreme conditions, such as the “worst-case scenario” or the Value at Risk (VaR) at a given probability level. Our research therefore provides important insights into these questions and contributes to risk management, pricing, and capital allocation decisions in the presence of climate-driven mortality risk.

In this section, we simulate samples from different models to assess and quantify the impact of extreme cold temperatures on cause-specific deaths as well as the total number of deaths, using a scenario-based approach.¹⁵ In the scenario-based analysis, we investigate the impact of extreme cold temperatures on death distributions, both with and without making additional assumptions about the level of respiratory deaths being high. Based on these results, we investigate the regional differences in cold-related deaths and identify the most affected regions by extreme cold temperatures. Additionally, we break down total deaths by causes and identify the most temperature-sensitive cause for each region.

Based on the selected CCC-vine models, we simulate temperature and death data for 1 million pairs of two consecutive months and for all six regions. In our scenario-based analysis, we consider two conditioning events. First, we look at scenarios conditioning only on cold temperatures, *i.e.* $T10$. Second, we look at scenarios conditioning both on cold temperatures and respiratory deaths, *i.e.* $T10$ and *Respiratory*. In each case, we compute the monthly deseasonalized deaths by cause and in total under various scenarios. We also compare the prediction intervals computed from the CCC-vine models with those from the VAR models and find that VAR models tend to underestimate the upper tail risk of deseasonalized deaths, particularly when we experience two consecutive months of extreme cold temperatures.

¹⁵Note that the copula models considered in the simulation have identical margins. Therefore, evaluating the models after transforming the data to the original scale would lead to identical conclusions.

5.1 Scenarios conditioning only on cold temperatures

In this section, our baseline scenario does not impose any restrictions on the values of the temperature index $T10$. We consider the distribution of deseasonalized deaths at time t under three temperature scenarios as follows:

1. The $T10$ index at time t exceeds its 90th percentile.¹⁶
2. The $T10$ index at time $t - 1$ exceeds its 90th percentile.
3. The $T10$ index at both time t and $t - 1$ exceeds its 90th percentile.

The first scenario describes the circumstance where the number of extreme cold days in the current month t is higher than its 90th percentile threshold based on historical data. Therefore, it measures the contemporaneous impact of extreme cold temperature on excess deaths. The second scenario describes the circumstance where the number of extreme cold days in the previous month $t - 1$ is higher than its 90th percentile threshold. Therefore, it measures the lagged impact of extreme cold temperature in previous month on excess deaths of the current month. The third scenario describes the circumstance where the number of extreme cold days in both the current month t and the previous month $t - 1$ are exceeding its 90th percentile threshold. Therefore, it measures the impact of extreme cold temperature in two consecutive months on the excess deaths of current month. These three scenarios will be compared against a baseline scenario where there is no condition imposed on the value of temperature index $T10$. Roughly we have 10^5 pairs for Scenario 1, 10^5 pairs for Scenario 2, and slightly more than 10^4 pairs for Scenario 3.¹⁷


Prediction intervals of aggregated deaths by region We compute the prediction intervals of monthly deaths at the 10th and 90th percentiles under different scenarios, as well as the median level (50th percentile) of the monthly deseasonalized deaths. For insurance pricing and government planning, adverse mortality experience is far more critical than favourable mortality experience. Therefore, our focus is on the 90th percentile of the distribution. Figure 12 plots the prediction intervals for the six regions as well as the median of monthly deseasonalized deaths, for the CCC-vine models, the Gaussian copula models, and the VAR models. The prediction intervals obtained from the three models differ substantially. We observe that, although the median deaths estimated by the three approaches are relatively close, the 90th percentiles produced by the CCC-vine models are noticeably higher in some regions. The Gaussian copula models and VAR models yield prediction intervals that are largely symmetric. On the other hand, the prediction intervals from the CCC-vine models tend to be asymmetrical and skewed to the right, reflecting the commonly observed heavy-tailed behaviour of death count series. These results highlight the greater potential risk in climate-related mortality, demonstrating the practical value of the CCC-vine model

¹⁶The 90th percentile is computed based on the historical $T10$ data over the period 2000–2018.


¹⁷Given that we have a total of 1 million pairs of simulated data for two consecutive months (*i.e.*, $t - 1$ and t), under Scenario 1, where $T10$ in the current month exceeds its 90th percentile, we observe roughly 10^5 pairs of simulated data (10% of 1 million). Similarly, for Scenario 2, where $T10$ in the previous month exceeds its 90th percentile, we also observe roughly 10^5 pairs. Finally, for Scenario 3, where $T10$ exceeds its 90th percentile in both the previous and current months, we observe approximately 10^4 pairs.

in capturing tail events and extreme dependence, as well as in assessing extreme scenarios.

Across the six regions, we can see that CEA has the highest number of excess deaths under all scenarios based on the CCC-vine, closely followed by SEA. This result is not surprising as these two regions have large population exposures. Moreover, SEA has a relatively warm climate across the six regions (as shown in Table 1), so it may be less resilient against extreme cold temperatures. SWP appears to primarily experience excess deaths in Scenarios 1 and 3, but not in Scenario 2. For regions with relatively small populations such as CWP and SPL, the number of excess deaths are considerably lower.



(a) CCC-vine



(b) Gaussian copula

(c) VAR

Figure 12: Prediction intervals of monthly deseasonalized total deaths at 10th and 90th percentiles. The dot represents the median of the distribution. Baseline scenario does not impose any restrictions.


From Figure 12 Panel (a), we can see that for all regions, Scenario 1 shifts up the 10th and 90th percentiles of monthly death distribution, indicating a higher likelihood of extreme mortality events under extreme temperatures. Scenario 3 seems to have the most severe

impact on extreme mortality, where the shift in 90th percentile is the biggest across all six regions. This indicates that longer periods of unusually cold temperatures lead to higher excess mortality. Finally, Scenario 2 does not seem to have the same impact on the death distribution compared to the other two scenarios: for most regions, the shift in the level of the 90th percentile is much smaller compared to Scenarios 1 and 3. Nevertheless, it indicates that if extreme cold temperatures occurred in the previous month, excess deaths are likely to happen regardless of the temperature situation in the current month. In summary, two consecutive months of extreme cold temperatures are most likely to trigger extreme mortality events.

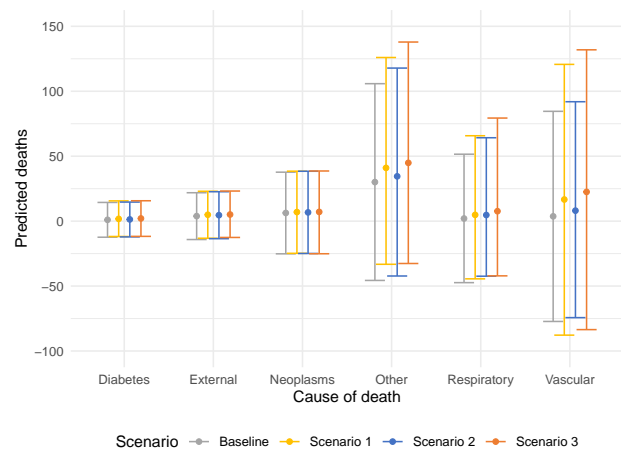
Based on results from the CCC-vine models, we compare the 90th percentiles of the predicted deaths under all scenarios across the six regions. For CEA, SEA, and SWP, we can see that the increase from the unconditional scenario to Scenario 3 is notably bigger than the increase in the cases of Scenarios 1 and 2. On the other hand, for CWP, MID, and SPL, it seems that extreme cold temperatures do not have a huge influence on the 90th percentile of the predicted deaths, indicating that the “worst” case scenario remains somewhat unchanged under different temperature scenarios. We suspect that this is also due to a higher level of adaptation to cold climates in these two regions. It is also interesting to note that in all regions but CWP, the 10th percentile of monthly deaths actually decreases when (additionally) conditioning on low temperatures in the previous month (*i.e.*, going from baseline to Scenario 2 and from Scenario 1 to Scenario 3). This observation can potentially be explained by the so-called “harvesting effect”, which suggests an increase in mortality level due to a short-term, acute environmental event is likely to be followed by a decrease in mortality in the preceding period.

Prediction intervals by region and cause To obtain further insights into how cause-specific deaths are affected by extreme temperatures, we break down the total deaths into the six major causes considered in our modeling, under three temperature scenarios. We plot the prediction intervals of monthly cause-specific deaths in Figure 13.¹⁸ We can see that *Vascular*, *Other*, and *Respiratory* have the top 3 largest contributions to cold-related deaths for the 90th percentile, under all scenarios across all regions. On the other hand, *Diabetes*, *External*, and *Neoplasms* seem largely unaffected by extreme cold temperatures. Therefore, we conclude that the main drivers of cold-related excess mortality are *Vascular*, *Other*, and *Respiratory*. Figure 13 also provides more insight into the harvesting effect observed in the 10th percent quantiles. Recall that the harvesting effect is indicated by a decrease in the 10th percentile of excess death when going from baseline to Scenario 2 and from Scenario 1 to Scenario 3, respectively. It appears across all causes in some regions (CEA, SEA, MID), but only for some of the causes in other regions (CWP, SPL, SWP). This again highlights regional differences in cold-related deaths.

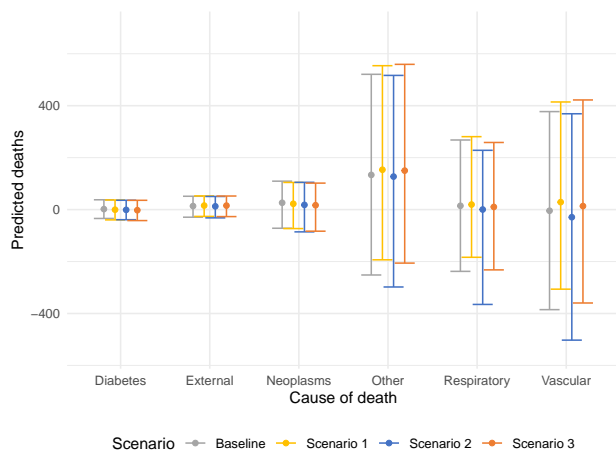
¹⁸The corresponding results and plots based on the Gaussian copula models and the VAR models are included in the supplementary materials, Figures S.16 and S.18.



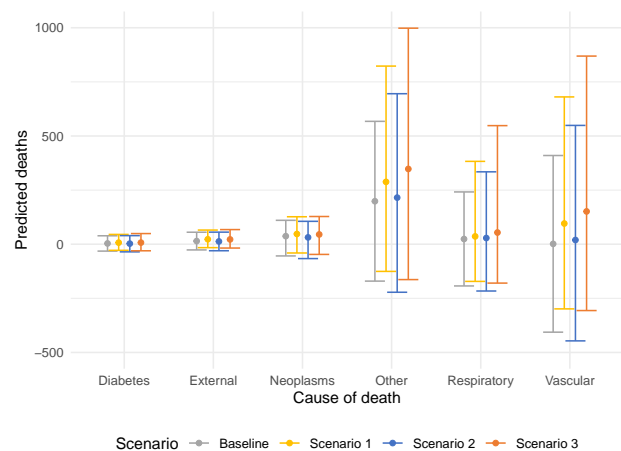
(a) CEA




(b) CWP




(c) MID



(d) SEA



(e) SPL



(f) SWP

Figure 13: Prediction intervals of monthly deseasonalized cause-specific deaths at 10th and 90th percentiles based on CCC-vine. The dot represents the median of the death distribution. Baseline scenario does not impose any restrictions.

5.2 Scenarios conditioning on both cold temperature and respiratory deaths

It is well-acknowledged that flu seasons are often associated with an elevated level of respiratory deaths. In light of this, we assess the impact of extreme cold temperatures conditioning on heightened respiratory death counts. In this section, our baseline scenario does not impose any restriction on the values of temperature index, but conditions on the death count of *Respiratory* in the current month (*i.e.* time t), requiring it to exceed its 90th percentile. Similarly, we consider the distribution of deseasonalized deaths at time t under three temperature scenarios as follows:

1. The $T10$ index at time t exceeds its 90th percentile, and *Respiratory* death at time t exceeds its 90th percentile.
2. The $T10$ index at time $t - 1$ exceeds its 90th percentile, and *Respiratory* death at time t exceeds its 90th percentile.
3. The $T10$ index at both time t and $t - 1$ exceeds its 90th percentile, and *Respiratory* death at time t exceeds its 90th percentile.

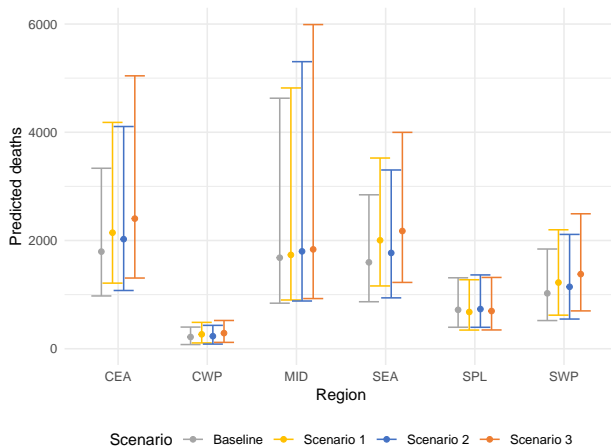
In terms of the sample size, roughly we have 10^5 for the baseline scenario, 3×10^4 pairs for Scenario 1, 2×10^4 pairs for Scenario 2, and slightly less than 5×10^3 pairs for Scenario 3. In Table 6, we report some key statistics of *Respiratory* deaths for the six regions. It can be seen that, the mean values are much smaller than the standard deviations, indicating that the distribution of *Respiratory* deaths is largely fat-tailed. The 90th percentile for each region is also presented in the table.

Table 6: Summary statistics of *Respiratory*.

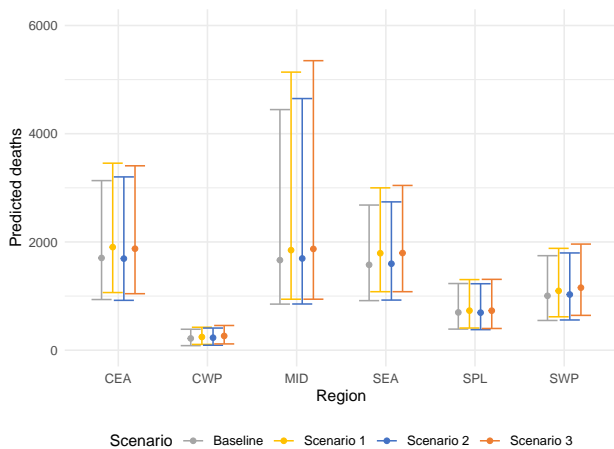
Region	Mean	Standard deviation	90th pctl
CEA	13.89	287.20	256.90
CWP	1.72	47.06	43.00
MID	13.79	291.20	246.60
SEA	22.53	232.31	243.70
SPL	9.68	127.35	129.20
SWP	21.96	167.78	165.30

Prediction intervals of aggregated deaths by region We first look at the 10th, 50th, and 90th percentiles of the prediction intervals of monthly deaths in Figure 14. Based on the CCC-vine models, compared to Figure 12, it is clearly shown that there is a substantial increase in the 90th percentiles of total predicted deaths by region. We suspect this increase in excess deaths is partly due to the positive dependence between *Respiratory* and other causes, particularly *Vascular*. Besides this increase, we observe that for all regions, the distributions of predicted deaths become more skewed to the right when we impose restrictions on a high level of *Respiratory* deaths. This shape better reflects tail mortality risk under

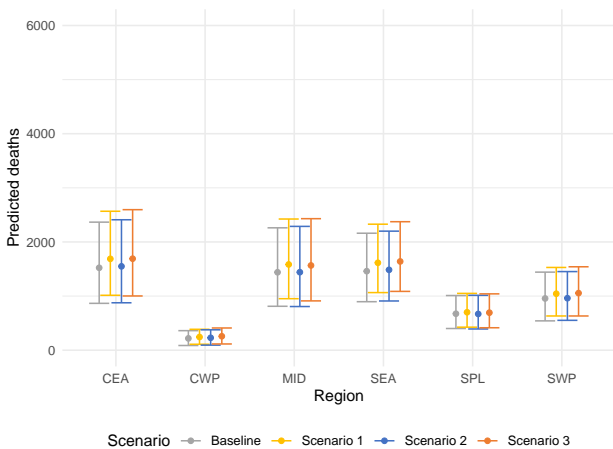
extreme events. Another interesting observation is that, in Figure 14, the 10th percentiles are always above 0 across all regions. While under the scenarios described in Section 5.1, the 10th percentiles are consistently negative, indicating less than expected number of deaths in the “best cases”. When comparing the results from the CCC-vine models and the two alternative models, again, we argue that the Gaussian model and the VAR model are likely to underestimate the extreme mortality risk due to extreme cold temperatures. For CEA, CWP, SEA, and SWP, Scenario 3 has the highest number of deaths for the 90th percentile (*i.e.* in the “worst case”). It should also be noted that for a majority of regions, the 10th percentiles do not change much from the baseline scenario to the other three scenarios.



(a) CCC-vine



(b) Gaussian copula



(c) VAR

Figure 14: Prediction intervals of monthly deseasonalized total deaths at 10th and 90th percentiles. The dot represents the median of the distribution. Baseline scenario imposes restrictions on high *Respiratory* deaths.

Prediction intervals by region and cause Finally, we decompose total deaths by the six causes and plot the prediction intervals of monthly cause-specific deaths in Figure 15.¹⁹ First, we can see that extreme cold temperatures still have very small impact on deaths from *Diabetes*, *External*, and *Neoplasms*. When conditioning on a high level of *Respiratory* deaths, we can see that extreme cold temperatures are still likely to cause excess deaths from *Vascular* and *Other*. Additionally, we should anticipate regional variations in the response of cause-specific mortality to extreme cold temperature events.

It should be noted that Figure 4 provides a benchmark for the CEA results presented in Figures 13 and 15.²⁰ Take *Respiratory* and *Vascular* deaths in the CEA region under the scenarios conditioning on both cold temperature and respiratory deaths as illustrative examples. In Figure 4, we observe two peaks in *Respiratory* death data during the winter months of 2013 and 2015, both at approximately 1,300. From the CEA plot in Figure 2, we also observe two extremely large values of $T10$ in the same time period. These align with the scenarios described in Section 5.2, where we impose conditions on both $T10$ and *Respiratory* deaths exceeding a high threshold. In panel (a) of Figure 15, we observe that for *Vascular* deaths, the upper limits of the prediction intervals are roughly between 1,250 and 1,500 over all scenarios, which is consistent with the observed peaks in death counts in Figure 4. Overall, we find that the prediction intervals are largely consistent with the observed data in Figure 4, which offers additional confidence in the performance and usefulness of our model.


6 Conclusions

To quantify and model cold-related excess deaths, this paper investigates the relationship between extreme cold temperatures and death counts from several major causes of death including *Diabetes*, *External*, *Neoplasm*, *Respiratory*, and *Vascular*. An innovative stationary CCC-vine model is proposed which allows for flexible dependence structures across variables while taking advantage of preliminary data analysis and information pooling. Besides the valuable contribution to the literature of mortality modeling, it should be noted that the vine copula-based approach introduced in this project is readily applicable to other areas such as joint risk modeling and portfolio management.


The empirical results illustrate good performance of our proposed model and provide new insights into an important research area. Older age groups have always faced higher risks of cold-related death. Understanding how and to what extent extreme cold affects mortality of this vulnerable segment of the population is an important step toward finding better solutions to protect the elderly against cold-related deaths, such as improved awareness and increased heating coverage. Our results show that geographical differences across regions also need to be taken into account when designing plans for cold weather on a national level. Public health policies and interventions should be tailored to different demographic and geographical factors.

¹⁹The corresponding results and plots based on the Gaussian copula models and VAR models are included in the supplementary materials, Figures S.17 and S.19.

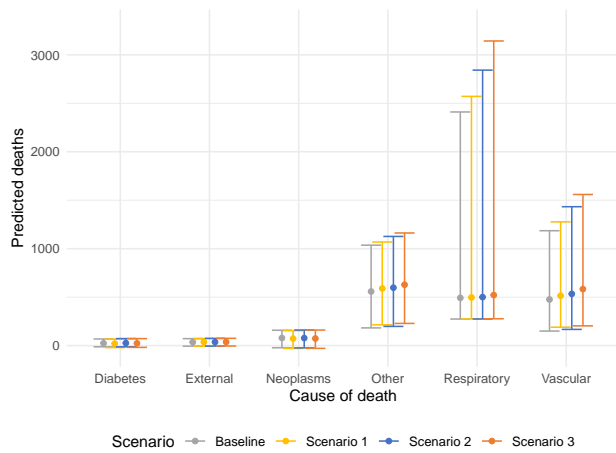
²⁰For the other 5 regions, comparisons can be made between Figures S.1–S.5 and Figures 13 and 15.




(a) CEA




(b) CWP



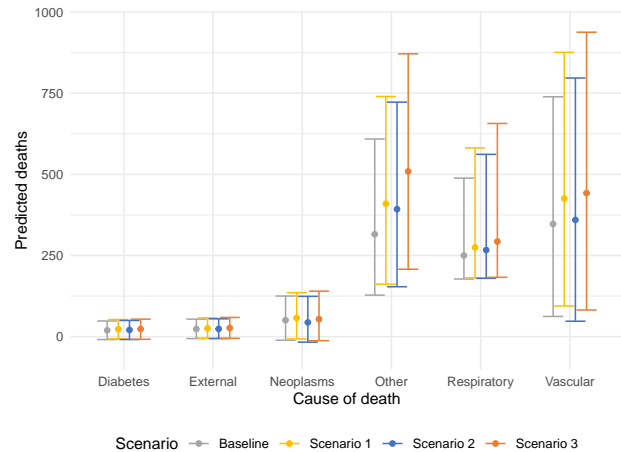
(c) MID



(d) SEA



(e) SPL



(f) SWP

Figure 15: Prediction intervals of monthly deseasonalized cause-specific deaths at 10th and 90th percentiles based on CCC-vine. The dot represents the median of the death distribution. Baseline scenario imposes restrictions on high *Respiratory* deaths.

References

- Aas, K. (2016). Pair-copula constructions for financial applications: A review. *Econometrics*, **4**(4), 43.
- Aas, K., Czado, C., Frigessi, A., and Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, **44**(2), 182–198.
- Actuaries Climate Index Executive Summary (2018). Available at: http://actuariesclimateindex.org/wp-content/uploads/2016/04/ACI_ExecutiveSummary5.18.pdf.
- Analitis, A., Katsouyanni, K., Biggeri, A., Baccini, M., Forsberg, B., Bisanti, L., Kirchmayer, U., Ballester, F., Cadum, E., Goodman, P. G., Hojs, A., Sunyer, J., Tiittanen, P., and Michelozzi, P. (2008). Effects of Cold Weather on Mortality: Results From 15 European Cities Within the PHEWE Project. *American Journal of Epidemiology*, **168**(12), 1397–1408.
- Appiah, G. D., Blanton, L., D’Mello, T., Kniss, K., Smith, S., Mustaquim, D., Steffens, C., Dhara, R., Cohen, J., Chaves, S. S., *et al.* (2015). Influenza activity—United States, 2014—15 season and composition of the 2015—16 influenza vaccine. *Morbidity and Mortality Weekly Report*, **64**(21), 583.
- Arbuthnott, K., Hajat, S., Heaviside, C., and Vardoulakis, S. (2018). What is cold-related mortality? A multi-disciplinary perspective to inform climate change impact assessments. *Environment International*, **121**, 119–129.
- Beare, B. K. and Seo, J. (2015). Vine copula specifications for stationary multivariate Markov chains. *Journal of Time Series Analysis*, **36**(2), 228–246.
- Bedford, T. and Cooke, R. (2001). Monte carlo simulation of vine dependent random variables for applications in uncertainty analysis. *ESREL 2003*.
- Bedford, T. and Cooke, R. M. (2002). Vines: A new graphical model for dependent random variables. *Annals of Statistics*, pages 1031–1068.
- Berghaus, B. and Bücher, A. (2017). Goodness-of-Fit tests for multivariate copula-based time series models”. *Econometric Theory*, **33**(2), 292–330.
- Brammer, L., Kniss, K., Epperson, S., Blanton, L., Mustaquim, D., Steffens, C., D’Mello, T., Perez, A., Dhara, R., Chaves, S. S., *et al.* (2013). Influenza activity—United States, 2012—13 season and composition of the 2013—14 influenza vaccine. *Morbidity and Mortality Weekly Report*, **62**(23), 473.
- Braun, H. (1980). A simple method for testing goodness of fit in the presence of nuisance parameters. *Journal of the Royal Statistical Society. Series B (Methodological)*, **42**(1), 53–63.
- Brechmann, E. C. and Czado, C. (2015). COPAR—multivariate time series modeling using the copula autoregressive model. *Applied Stochastic Models in Business and Industry*, **31**(4), 495–514.

- Coblenz, M., Holz, S., Bauer, H.-J., Grothe, O., Koch, R., *et al.* (2020). Modelling fuel injector spray characteristics in jet engines by using vine copulas. *Journal of the Royal Statistical Society Series C*, **69**(4), 863–886.
- Cohen, J., Pfeiffer, K., and Francis, J. A. (2018). Warm Arctic episodes linked with increased frequency of extreme winter weather in the United States. *Nature Communications*, **9**(1), 1–12.
- Czado, C. (2019). Analyzing dependent data with vine copulas. *Lecture Notes in Statistics*, Springer.
- Czado, C. and Nagler, T. (2022). Vine copula based modeling. *Annual Review of Statistics and Its Application*, **9**(1), 453–477.
- Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, **74**(366a), 427–431.
- Dissmann, J., Brechmann, E. C., Czado, C., and Kurowicka, D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics & Data Analysis*, **59**, 52–69.
- Donaldson, G. and Keatinge, W. (2002). Excess winter mortality: influenza or cold stress? Observational study. *BMJ*, **324**(7329), 89–90.
- Dushoff, J., Plotkin, J. B., Viboud, C., Earn, D. J., and Simonsen, L. (2006). Mortality due to influenza in the United States—an annualized regression approach using multiple-cause mortality data. *American Journal of Epidemiology*, **163**(2), 181–187.
- Forzieri, G., Cescatti, A., e Silva, F. B., and Feyen, L. (2017). Increasing risk over time of weather-related hazards to the European population: A data-driven prognostic study. *The Lancet Planetary Health*, **1**(5), e200–e208.
- Genest, C., Quessy, J.-F., and Rémillard, B. (2007). Asymptotic local efficiency of Cramér–von Mises tests for multivariate independence. *The Annals of Statistics*, **35**(1), 166 – 191.
- Genest, C., Rémillard, B., and Beaudoin, D. (2009). Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, **44**(2), 199–213.
- Joe, H. (1996). Families of m-variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. *Lecture Notes-Monograph Series*, pages 120–141.
- Joe, H. (2014). *Dependence modeling with copulas*. CRC press.
- Kendall, M. (2022). Rank correlation methods (4th edn.) Charles Griffin (1975). *Scientific Reports*, **12**, 8724.
- Kim, J.-S., Kug, J.-S., Jeong, S.-J., Huntzinger, D. N., Michalak, A. M., Schwalm, C. R., Wei, Y., and Schaefer, K. (2017). Reduced North American terrestrial primary productivity linked to anomalous Arctic warming. *Nature Geoscience*, **10**(8), 572–576.

- Kreuzer, A., Valle, L. D., and Czado, C. (2022). A bayesian non-linear state space copula model for air pollution in beijing. *Journal of the Royal Statistical Society Series C*, **71**(3), 613–638.
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of econometrics*, **54**(1-3), 159–178.
- Li, H. and Tang, Q. (2022). Joint extremes in temperature and mortality: A bivariate pot approach. *North American Actuarial Journal*, **26**(1), 43–63.
- Littell, R. C. and Folks, J. L. (1971). Asymptotic optimality of Fisher’s method of combining independent tests. *Journal of the American Statistical Association*, **66**(336), 802–806.
- Mann, H. B. (1945). Nonparametric tests against trend. *Econometrica: Journal of the econometric society*, pages 245–259.
- Maravall, A. (2011). Seasonality tests and automatic model identification in tramo-seats. *Bank of Spain: Madrid, Spain*.
- McMichael, T. (2011). Climate change: A dose of reality. *Nature*, **472**(7343), 292–293.
- Nagler, T., Krüger, D., and Min, A. (2022). Stationary vine copula models for multivariate time series. *Journal of Econometrics*, **227**(2), 305–324.
- Ollech, D. and Webel, K. (2020). A random forest-based approach to identifying the most informative seasonality tests.
- R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rémillard, B. (2017). Goodness-of-fit tests for copulas of multivariate time series. *Econometrics*, **5**(1).
- Schepsmeier, U. and Czado, C. (2016). Dependence modelling with regular vine copula models: a case-study for car crash simulation data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **3**(65), 415–429.
- Serre, D. (2022). Climate risk analysis for life and health insurance companies. *Research Report*.
- Sklar, M. (1959). Fonctions de repartition an dimensions et leurs marges. *Publ. inst. statist. univ. Paris*, **8**, 229–231.
- Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, **31**(3), 815–833.
- Vernieuwe, H., Vandenberghe, S., De Baets, B., and Verhoest, N. E. (2015). A continuous rainfall model based on vine copulas. *Hydrology and Earth System Sciences*, **19**(6), 2685–2699.

- Wan, K., Feng, Z., Hajat, S., and Doherty, R. M. (2022). Temperature-related mortality and associated vulnerabilities: evidence from scotland using extended time-series datasets. *Environmental Health*, **21**(1), 1–14.
- WHO (2014). *Quantitative risk assessment of the effects of climate change on selected causes of death, 2030s and 2050s*. World Health Organization.
- Zhao, Q., Guo, Y., Ye, T., Gasparrini, A., Tong, S., Overcenco, A., Urban, A., Schneider, A., Entezari, A., Vicedo-Cabrera, A. M., *et al.* (2021). Global, regional, and national burden of mortality associated with non-optimal ambient temperatures from 2000 to 2019: a three-stage modelling study. *The Lancet Planetary Health*, **5**(7), e415–e425.

Appendix

1. Marginal models

Table A.1: Marginal models for causes of death and temperature index selected by the AIC. Each cell contains the model family, parameter values, and p -value of a Cramér-von-Mises goodness-of-fit test.

Region	T10	Diabetes	External	Neoplasms	Other	Respiratory	Vascular
CEA	Gamma (1.7, 0.2) 0.19	Logistic (3.6, 15.5) 0.50	Logistic (12, 17) 0.86	Normal (25, 78) 0.54	G Normal (174, 339.1, 0.8) 0.94	Student-t (16, 4829, 2) 0.88	Laplace (-12, 294) 0.66
	Gamma (2, 0.3) 0.45	Normal (1, 10.4) 0.67	Normal (3.8, 14) 0.23	Logistic (6.3, 14.3) 0.11	G Normal (30, 63.6, 1.2) 0.38	G Normal (2.0, 46.9, 0.8) 0.96	Logistic (3.8, 36.8) 0.11
MID	Weibull (1.3, 8.9) 0.54	Logistic (1.8, 16.4) 0.98	S Normal (12.1, 31.5, 0.9) 0.98	S Normal (22, 70.8, 0.8) 0.65	Laplace (134, 240) 0.82	Cauchy (15, 82) 0.17	Student-t (-4.1, 581, 2.3) 0.84
	Weibull (1.5, 8.4) 0.26	Logistic (3.6, 16.1) 0.15	Logistic (15, 19) 0.98	S G Normal (31.9, 67.2, 1.4, 0.9) 0.64	Laplace (198, 229) 0.72	G Normal (24, 232.5, 0.6) 0.60	Laplace (1.5, 253.6) 0.48
SPL	Gamma (2.6, 0.3) 0.87	Normal (1.2, 16.9) 0.90	Normal (4.7, 22.4) 0.93	Normal (11, 44) 0.69	Student-t (73.8, 148.1, 3.4) 0.96	G Normal (10, 127.3, 0.7) 0.86	Laplace (1.5, 130) 0.63
	Lognormal (1.8, 0.6) 0.76	Normal (5.9, 22) 0.68	Normal (9.6, 22.6) 0.16	S Normal (33.2, 50.9, 1.3) 0.78	Laplace (130, 132) 0.38	S G Normal (22, 167.8, 0.6, 0.9) 0.97	Laplace (34, 187) 0.36

2. Standard S-vine structures

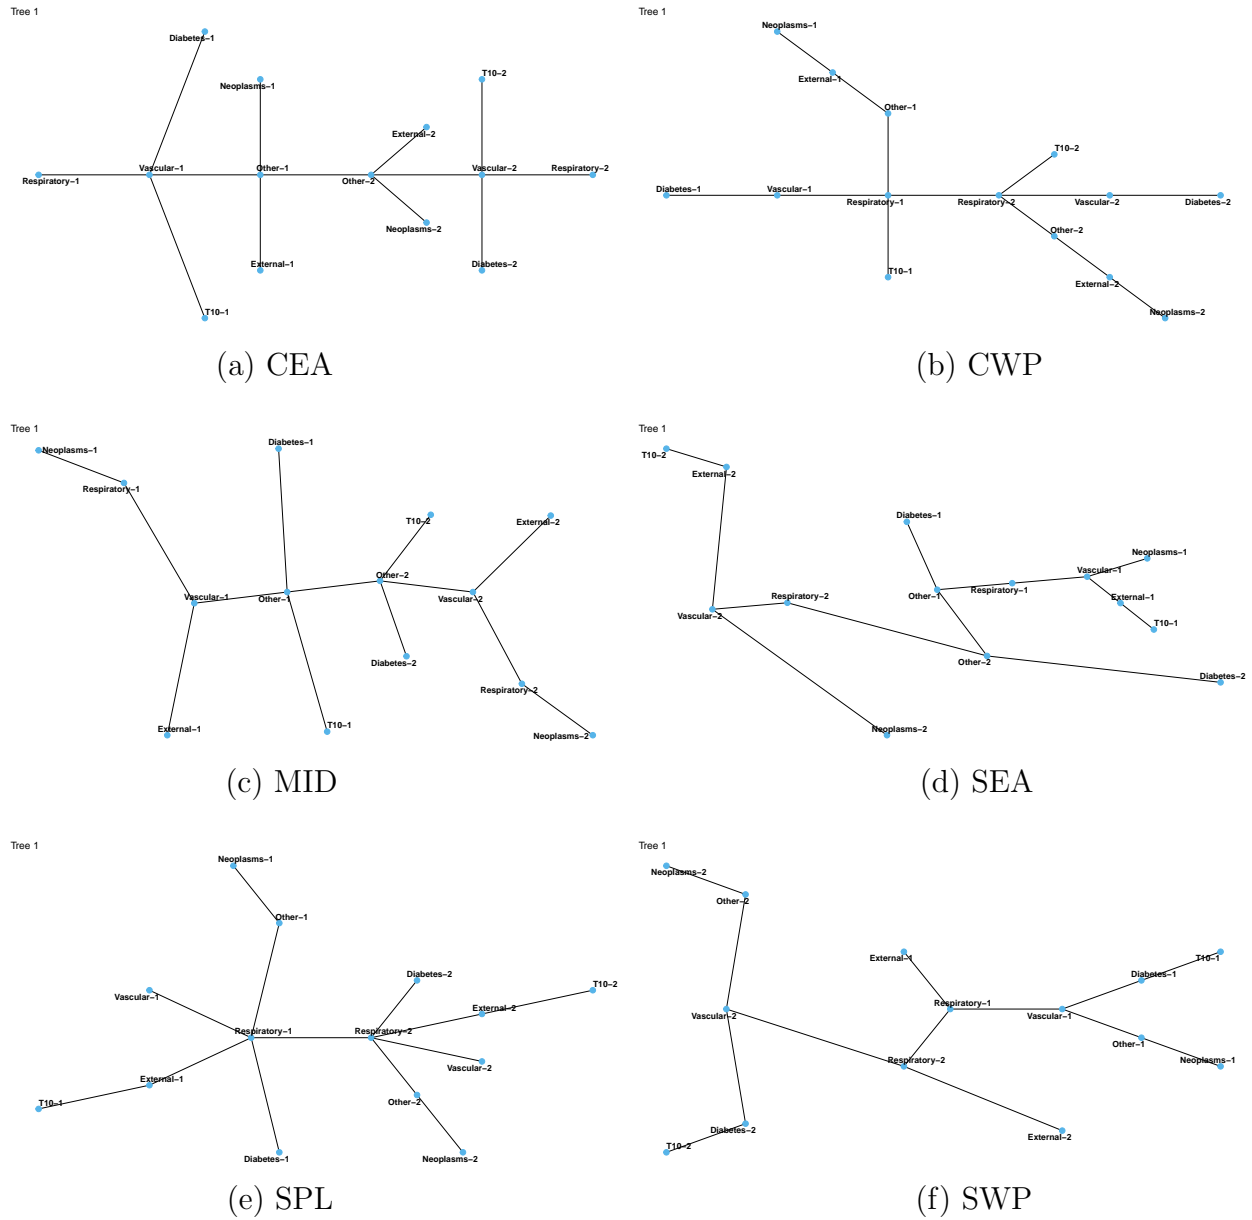


Figure A.1: First tree of the selected standard S-vine structures for the six regions.