A hazard model of subfreezing temperatures in the United Kingdom using vine copulas

Symeon Koumoutsaris
Guy Carpenter, Tower Place, London, EC3R 5BU, UK
Correspondence to: Symeon Koumoutsaris (symeon.koumoutsaris@guycarp.com)

Abstract. Extreme cold weather events, such as the winters of 1962/63, the third coldest winter ever recorded to the Central England Temperature record, or more recently the winter of 2010/11, have significant consequences for the society and economy. This paper assesses the probability of such extreme cold weather across the United Kingdom. For that, a statistical model is developed in order to model the extremes of the Air Freezing Index (AFI), which is a common measure of magnitude and duration of freezing temperatures. A novel approach in the modelling of the spatial dependence of the hazard has been followed which takes advantage of the vine copula methodology. The method allows to model complex dependencies especially between the tails of the AFI distributions which is important to assess reliably the extreme behaviour of such events. The model suggests that the extreme winter 1962/63 has a return period of approximately once every 89 years, with 95% confidence intervals between 81 to 120 years. However, the relative short record length together with the unclear effects of anthropogenic forcing on the local climate add considerable uncertainty to this estimate. This model is used as part of a probabilistic catastrophe model for insured losses caused by the bursting of pipes.

1 Introduction

Extended periods of extreme cold weather can cause severe disruptions in human societies; on human health, by exacerbating previous medical conditions or due to reduction of food supply which can lead to famine and disease; agriculture, by devastating crops particularily if the freeze occurs early or late in the growing season; on infrastructure, e.g. severe disruptions in the transport system, burst of residential or system water pipes (Bowman et al., 2012). All these consequences lead to important economic losses.

Of particular interest for the insurance industry are the economical losses that originate as a result of bursting of pipes due to freeze events. Water pipes burst because the water inside them expands as it gets close to freezing which causes an increase in pressure inside the pipe. Whether a pipe will break or not, depends on the water temperature (and consequently on the air temperature), the freezing duration, the pipe diameter and composition, the wind chill effect (due to wind and air leakage on water pipes), and the presense of insulation (Gordon, 1996; McDonald et al., 2014).

Insurance losses from burst pipes have a significant impact on the UK insurance industry. They amount to more than £900 million in the last 10 years, representing around 10% of the total insured losses, mainly due to flood and windstorm, in the United Kingdom (UK) during the same period (ABI, 2017). Particular years can be very damaging, such as, for example, the

winter of 2010/2011 where losses from burst pipes have exceeded £300 million in UK making it the peril with the largest losses that year (ABI, 2017). Moreover, much more extreme cold winters have actually occurred in the UK in the last 100 years, such as the winters of 1946/47 and 1962/63. It is crucial for the insurance business to be able to anticipate the likelihood of occurrence of similar and even more extreme events so that they can adequately prepare for their financial impact (AIR, 2012). In fact, the capital requirements in (re)insurance is estimated in a 1 in 200 year return period (RP) loss basis, which is usually much larger than the available historical records.

Probabilistic catastrophe modelling is generally agreed to be the most appropriate method to analyze such problems. The main goal of catastrophe models is to estimate the full spectrum of probability of loss for a specific insurance portfolio (i.e. comprised by several residential, auto, commercial or industrial risks). This requires the ability to extrapolate the possible losses at each risk to high return periods (RP) which is usually achieved by simulating synthetic events that are likely to happen in the near future (typically a year). More importantly, it requires to consider also how all risks relate to each other and their potential synergy to create catastrophic losses. Such spatial dependence between risks can result from various sources, for example due to the spatial structure of the hazard (e.g. the footprint in a windstorm or the catchment area in a flood event) or due to similar building vulnerabilities between risks in the same geographical area (e.g. due to common building practices) (Bonazzi et al., 2012).

Modelling the spatial dependence of the hazard is usually achieved by taking advantage of certain characteristic properties of the hazard footprint, like for example the track path and the radius of maximum wind for and windstorms or the elevation in the case of floods. In the case of temperature, however, such a property cannot be easily defined; an alternative solution is to use multivariate copula models. Based on Sklar's theorem (Sklar, 1959), the joint distribution of all risk sources can be fully specified by the separate marginal distributions of the variables and by their copula, which defines the dependence structure between the variables.

However, one important difficulty is the limited choice of adequate copulas for more than two dimensions. For example, standard multivariate copula models such as the elliptical and Archimedean copulas do not allow for different dependency models between pairs of variables. Vine copulas provide a flexible solution to this problem based on a pairwise decomposition of a multivariate model into bivariate copulas. This approach is very flexible, as the bivariate copulas can be selected independently for each pair, from a wide range of parametric families, which enables modelling of a wide range of complex dependencies (Czado, 2010; Dißmann et al., 2013)

In this paper, the vine copula methodology is used in a novel application to develop a catastrophe model on insurance losses due to pipe bursts resulting from freeze events in the United Kingdom. The focus here is on the hazard component (Sect. 2) which is modeled using the Air Freezing Index (AFI), an index which takes account both the magnitude and duration of air temperature below freezing, calculated from temperature data from the last 51 years. The methods used are described in Sect. 3. Extreme value analysis is performed on the historical AFI values in order to extrapolate to longer return periods (Sect. 3.2). Stochastic winter-seasons are simulated taking into account the correlation of the hazard between all pair-cells with the help of regular vine copulas (Sect. 3.3). The resulting return periods of extreme cold winters in UK, including the underlying uncertainties, are discussed in Sect. 4. Concluding remarks are found in Sect. 5.

2 Temperature data

The hazard component of the catastrophe model is based on the gridded dataset of observed daily average temperature developed from the UK Met Office (Perry et al., 2009). The dataset covers the entire UK for the period from 1960 to 2011 at 5km x 5km resolution and georeferenced in the British National Grid projection. It is based on temperature data retreived from 540 stations accross UK with an average station density of 21 x 21 km^2 (Perry and Hollis, 2005; Perry et al., 2009). The data are rigorously quality-checked and interpolated to a regular grid using inverse-distance weighting, as described in Perry et al. (2009). The dependencies across cells, may, thus, be partially due to the interpolation itself.

For computational reasons, the data are regridded to a lower resolution of 50 km x 50km, which leads to a total of 170 cells over land. The use of a coarser horizontal resolution is expected to have relatively small influence in most cases given that winter climate anomalies are often coherent across large parts of the UK as they are primarily associated with large-scale atmospheric circulation patterns (Scaife and Knight, 2008). Nevertheless, local temperature may be subtly different in certain micro-climates, such as upland and urban regions. In particular over urban regions, which are most important from an insurance perspective, lower resolution may lead to temperatures that are biased towards lower values, leading though to a conservative view on the severity of extreme freeze events. In upland regions, on the other hand, extreme cold temperatures are most probably underestimated, although it is reasonable to expect that their damaging effects are somewhat mitigated from increased protection levels. For example, water pipes in properties located in mountaineous regions are usually better protected against cold spells.

3 Methods

3.1 Air-Freezing Index and historical events

The daily temperature gridded data are used to compute the AFI as the sum of the absolute average daily temperatures of all days with below 0°C temperatures during the freezing period (1). The freezing period in this study is defined from first of June to end of May of the following year, in order to include the entire winter season. Because AFI accounts both for the magnitude and duration of the freezing period, it is commonly used for determining the freezing severity of the winter season (Frauenfeld et al., 2007; Bilotta et al., 2015).

$$AFI = \sum_{day=1/6/(Year)}^{31/5/(Year+1)} |T_{day}| \text{ if } T_{day} < 0°C \quad (1)$$

Figure 1a shows a map of the AFI values for the season 1962/1963 (i.e. season starting from 1st June 1962 to 31st of May 1963), which was one of the coldest on the record in the United Kingdom (Walsh et al., 2001). The "Big Freeze of 1962/63", as it is also known, began on the 26 of December 1962 with heavy snowfall and went on for nearly three months until March 1963. The cause of the cold conditions has been the development of a large "blocking" anticyclone over Scandinavia and north-western Russia. Easterly winds on the southern edge of this system transported cold continental air westwards, displacing the

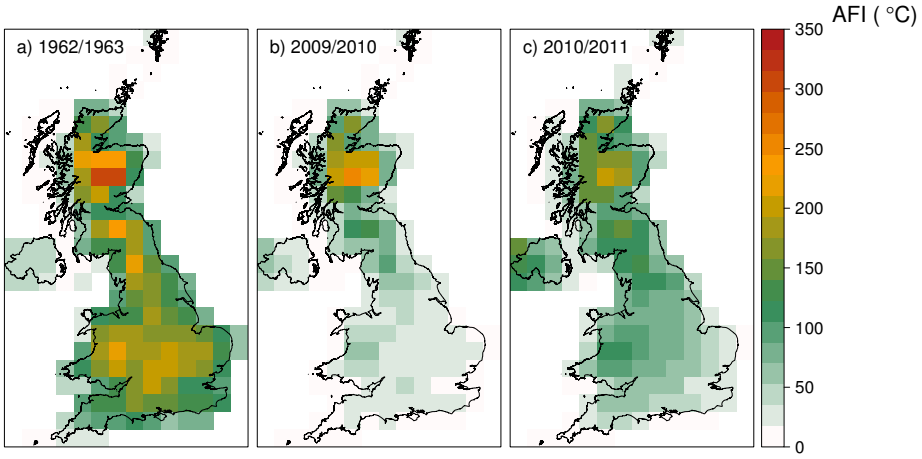

Figure 1. Map of AFI values (in °C) for the the winter-seasons of a) 1962/63, b) 2009/10, and c) 2010/11.

more usual mild westerly influence from the Atlantic Ocean on the British Isles. Over the Christmas period, the Scandinavian High collapsed, but a new one formed near Iceland, bringing Northerly winds. The mean AFI value (mAFI) in the entire UK (i.e. average of AFI values across all gridcells) mounted up to 98.3°C, which represents four standard deviations larger than the average of the entire 51-year period (19.6°C). The event affected the entire country with peak AFI values exceeding 200°C both in the South and in the North of the country (Fig. 1).

After 1962/63, a long run of mild winters followed until late 1978 and early 1979 (Fig. 2). However, temperatures in 1978/79 were not as low and the cold weather was interrupted frequently by brief periods of thaw (Cawthorne and Marchant, 1980). The mAFI value of winter 1978/79 reached 59.2°C. The 1980s stands out as a decade with several cold spells in UK, with mAFI values above 40°C for the winters 1981/82, 1984/85, and 1985/86 (64.8, 43.9, and 50.6 °C, respectively). Finally, the winters of 2009/2010 and 2010/2011 brought frigid temperatures to parts of Europe and the UK (Guirguis et al., 2011; Osborn, 2011; Seager et al., 2010), with mAFI values across UK of 39.1 and 62.2°C (Fig. 1b and c). As mentioned previously, the latter one had a significant financial impact on the UK insurance industry. The relation between cold winter spells and the North Atlantic Oscillation (NAO), a large-scale mode of natural climate variability, is discussed in detail in Sect. 4.2.2.

3.2 Extreme value analysis

Since the historical data only extends for 51 years and our interest lies in very rare events (such as 1 in 200 years), it is necessary to extrapolate by fitting an extreme value distribution. The Generalized Extreme Value (GEV) family of distributions has been chosen, which includes the Gumbel, the Fréchet, and Weibull distributions. An additional term was included, the probability of no hazard (P0), in order to account for the cells, mainly on the south England coast, that have years with no

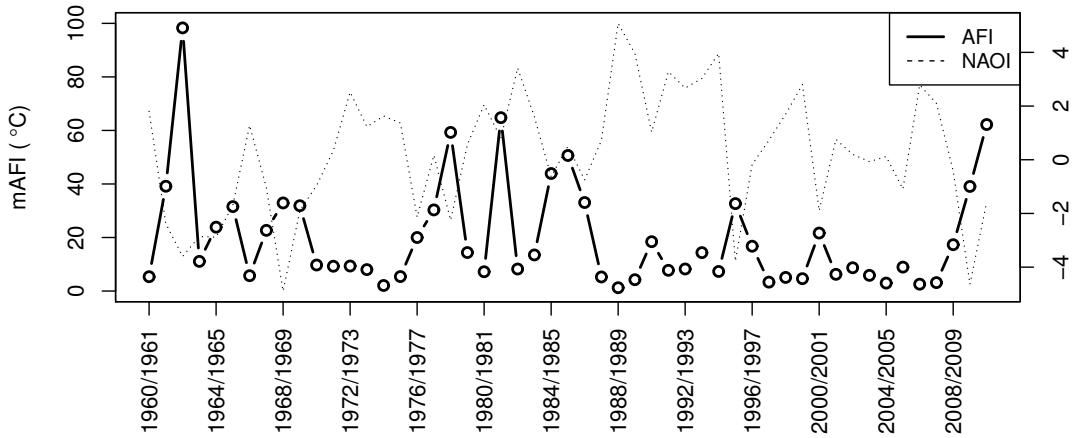

Figure 2. Interannual variation of UK average AFI over the study period. The North Atlantic Oscillation Index (NAOI) is also shown in dotted line.

negative temperatures at all. The probability therefore that the AFI value (X) inside a cell j will exceed a certain value (x) takes the form:

$$F(x) = P(X \leq x) = P0 + (1 - P0)exp\left\{-\left(1 + \xi\frac{x-\mu}{\sigma}\right)^{-\frac{1}{\xi}}\right\} \tag{2}$$

where μ, σ, and ξ represent the location, scale, and shape parameters of the distribution, respectively. $P0$ is computed as [total number of years with no hazard]/[total number of years + 1]. $F(x)$ is defined when $1 + \xi\frac{x-\mu}{\sigma} > 0$, $\mu \in \Re$, $\sigma > 0$, and $\xi \in \Re$. Its derivative, the GEV probability density function $f(x)$ is given by:

$$f(x) = \begin{cases} P0, & \text{if } x = 0 \\ (1 - P0)\frac{1}{\sigma}\left[1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right]^{-\frac{1}{\xi}-1} exp\left\{-\left[1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right]^{-\frac{1}{\xi}}\right\}, & \text{if } x > 0 \end{cases} \tag{3}$$

There are various methods of parameter estimation for fitting the GEV distribution, such as least squares estimation, maximum likelihood estimation (MLE), probability weighted moments, and others. Traditional parameter estimation techniques give equal weight to every observation in the dataset. However, the focus in catastrophe modeling is mainly on the extreme outcomes and, thus, it is preferable to give more weight to the long return periods. I therefore use the Tail-Weighted Maximum Likelihood Estimation (TWMLE) method developed by (Kemp, 2016) to estimate the GEV parameters which introduces ranking dependended weights ($w_{(r)}$) in the maximum likelihood. The weights are defined for each cell based on the historical

Table 1. GEV parameters for a single cell over London (cellid = 32).

method	location	scale	shape	P0
MLE	3.53	5.12	1.14	0.12
TWMLE + geographical smoothing	5.33	12.05	0.28	0.12

winter-season AFI values, i.e. the lowest historical AFI value in the cell (rank r=1 out of n observations) has the lowest weight, while the largest historical AFI value (rank r=n) has the largest weight, as follows:

$$w_{(r)} = AFI_{(r)} / \sum_{r=1}^{n} AFI_{(r)} \tag{4}$$

Along with the TWMLE method described above, a second modification has been implemented in order to geographically smooth the GEV parameters. The smoothing is incorporated into the fitting process by minimizing the local (ranked) log-likelihood. More precisely, the log-likelihood at each grid cell i is calculated using all grid points but weighted by their distance d_{ij}:

$$LogL_i = \sum_{j=1}^{170} (k_{ij} * LogL_j) \tag{5}$$

where $k_{ij} = \frac{1}{\sqrt{2\pi}} e^{-\frac{d_{ij}^2}{2L^2}}$, L is the smoothing parameter, and $LogL_j$ is the ranked log-likelihood for cell j.

Because the historical gridded data are already geographically smoothed, a small length scale parameter L of 15 km has been used (in comparison to the 50km grid size). In general, the increase of the sample size at each grid point allows for a more precise estimation of the parameters, especially for the shape parameter which is highly influential in estimating the hazard levels and high return periods.

As an example, the GEV fit for a single cell over London is shown in Fig. 3. The curve fitted as described above (black line) is closer to the empirical estimates (black circles, computed as described in Sect. 4.1) in comparison with the GEV fit with no weighting applied (grey line). As shown in table 1, for both fits the shape parameter is positive (i.e. both fits correspond to the Fréchet distribution), but for the approach followed here (TWMLE + geographical smoothing), the shape parameter is smaller leading to a shorter tail and a curve that is nearer to the empirical estimate. The largest AFI empirical point in this figure represents the 1962/63 exceptional winter which is estimated to be a more rare event than what the historical data suggests (i.e. larger than 1 in 52 years), as further discussed in the following sections.

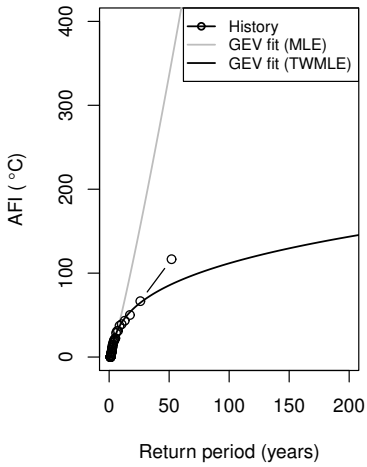

Figure 3. AFI return period curves for a single cell over London:,(black circles) empirical fit, (grey line) GEV fitted with MLE, and (black line) GEV fitted with TWMLE and geographical smoothing.

3.3 Copulas and vine copulas

The stochastic behaviour of the hazard (i.e. AFI) at each cell is fully described by its corresponding GEV probability distribution, as described in Sect. 3.2. However, insurance portfolio loss analysis requires the calculation of the combined stochastic behaviour of the hazard across all the model domain (i.e. all cells). This is described by the joint distribution of the hazard which, according to Sklar's theorem, can be fully specified by the separate marginal GEV distributions and by their (d-dimensional) copula, which models the hazard dependence between the cells.

More precisely, consider a vector of $X = (X_1, ..., X_d)$ of random variables with a joint probability density function (pdf), $f(x_1, ..., x_d)$. Sklar's theorem (Sklar, 1959) states that any multivariate continuous distribution function $F(x_1, ..., x_d)$ with marginals $F_1(x_1), ..., F_d(x_d)$ can be written as:

$$F(x_1, ..., x_d) = C(F_1(x_1), ..., F_d(x_d)) \tag{6}$$

for some appropriate d-dimensional copula C, which is uniquely determined on [0,1]^d.

The density of X, $f(x_1, ..., x_d)$, can be found by taking the partial derivatives with respect to X:

$$f(x_1, ..., x_d) = c(u_1, ..., u_d) \prod_{i=1}^{d} f_i(x_i) \tag{7}$$

where $c(u_1, ..., u_d)$ is the copula density, given by

$$c(u_1, ..., u_d) = \frac{\vartheta^d C(u_1, ..., u_d)}{\vartheta u_1 ... u_d} \tag{8}$$

Expression 7 is important in terms of modelling because it permits to define a multivariate density as the product of marginal densities and a copula function that captures the dependence between the random variables (Abbara and Zevallos, 2014). For a theoretical introduction to copulas, see Nelsen (2006); Meucci (2011); Joe (2014); Durante and Sempi (2015); for a practical/engineering approach and guidelines, see Genest and Favre (2007); Salvadori and Michele (2007); Salvadori et al. (2014, 2015)

To quantify the dependence between variables, different measures have been defined, addressing different aspects of dependence. A common measure of overall dependence is the Kendall rank correlation coefficient, commonly referred to as Kendall's τ coefficient (Genest and Favre, 2007). However, dependence of rare events cannot be measured by overall correlations: even if two variables are completely uncorrelated, there can be a significant probability of a concurrent extreme event in the two, i.e., they can still be tail dependent. Tail dependence describes the amount of dependence in the lower tail or upper tail of a bivariate distribution. For its mathematical definition see Haff et al. (2015).

One important complication is that identifying the appropriate d-dimensional copula is not an easy task. In high dimensions, the choice of adequate families is rather limited (Brechmann and Schepsmeier, 2013). Standard multivariate copulas, either do not allow for tail dependence (i.e. multivariate Gaussian) or have only a single parameter to control tail dependence of all pairs of variables (Student-t and archimedean multivariate copulas). This is particularly problematic for catastrophe modeling applications, where a flexible modeling of tails is vital to assess reliably the extreme behaviour of natural events.

Vine copulas provide a flexible solution to this problem based on a pairwise decomposition of a multivariate model into bivariate (conditional and unconditional) copulas, where each pair-copula can be choosen independently from the others. In particular, asymmetries and tail dependence can be taken into account as well as (conditional) independence to build more parsimonious models. Vines thus combine the advantages of multivariate copula modeling, that is separation of marginal and dependence modeling, and the flexibility of bivariate copulas (Brechmann and Schepsmeier, 2013).

As an example, in a 4-dimensional case, the joint pdf can be decomposed as a product of 6 pair-copulas (3 uncoditional and 3 conditional) and 4 marginal densities as shown in Eq. 9:

$$\begin{aligned} f(x_1, x_2, x_3, x_4) = & f(x_1) f(x_2) f(x_3) f(x_4) \\ & \times c_{12}(F_1(x_1), F_2(x_2)) \\ & \times c_{23}(F_2(x_2), F_3(x_3)) \\ & \times c_{34}(F_3(x_3), F_4(x_4)) \\ & \times c_{13|2}(F_{1|2}(x_1 \mid x_2)), F_{3|2}(x_3 \mid x_2))) \\ & \times c_{24|3}(F_{2|3}(x_2 \mid x_3)), F_{4|3}(x_4 \mid x_3))) \\ & \times c_{14|23}(F_{1|23}(x_1 \mid x_2, x_3)), F_{4|23}(x_4 \mid x_2, x_3))) \end{aligned} \tag{9}$$

The above decomposition is not unique and Bedford and Cooke (2002) introduced a graphical structure called regular vine (R-Vine) structure to represent this decomposition with a set of nested trees. More details on vine copulas can be found in Aas et al. (2006); D. and Schirmacher (2008); Czado (2010); Schepsmeier (2013).

3.3.1 Selection of the Regular Vine Model (RVM)

In this study, the joint multivariate hazard distribution of AFI across all the model domain (170 cells) is decomposed as a product of marginal and pair-copula densities (in a similar way as shown for the 4-d case above). F(x) and f(x) represent the margnial GEV distributions here as defined by Eq. 2 and 10. The pair-copulas are fitted using the R (https://www.r-project.org/) package VineCopula (Schepsmeier et al., 2017; Brechmann and Schepsmeier, 2013). The method follows an automatic strategy of jointly searching for an appropriate R-Vine tree structure, its pair-copula families, and estimating their parameters developed by Dißmann et al. (2013). This algorithm selects the tree structure by maximizing the empirical Kendall's τ values, based on the intuition that variable pairs with high dependence should contribute significantly to the model fit and should be included in the first trees.

The copula family types for each selected pair in the first tree are determined by using the Akaike information criterion (see (Brechmann and Schepsmeier, 2013)). For computational reasons, the two-parameter Archimedean copulas are excluded from this analysis, which however has only a negligible impact in the results as discussed in Sect. 4.2.3. The copula parameters are estimated sequentially (using maximum likelihood estimation) starting from the top tree until the last tree, as described in Czado et al. (2013). This approach only involves estimation of bivariate copulas and has been chosen since it is computationally much less demanding than joint maximum likelihood estimation of all parameters at once.

The percentage of family types used for the first few trees of the selected RVM is shown in Table 2. The large majority of the pairs in all levels are estimated to be independent (64%), but these pairs occur mainly at the higher levels, since the most important dependencies are captured in the first trees (Brechmann and Schepsmeier, 2013; Dißmann et al., 2013). Large dependencies, with Kendall's tau coefficients greater than 0.90, are found as expected between neighbouring cells, but remain important across the whole model domain due to the nature of the hazard: AFI assess the freezing temperatures during the entire winter and, thus, is less associated with small scale local phenomena that can cause important spatial variation.

At the first tree, 48.5% of the selected bivariate copulas are found to belong to the Gumbel family, which has a greater dependence in the positive tail than in the negative and thus implies greater dependence at larger AFI values than at lower ones. This is also the case for the 180° Clayton copula which has the largest contribution at the second level (9.5%), but with a much more uniform split between families. A sensitivity analysis on the influence of the selected copula families in the resulting return periods is presented in Sect. 4.2.3. From the third tree and onwards, the percentage of independent families is always larger than 40%.

The small sample size used (51 years of data) in conjunction with the high dimensions of the modelled pdf (170) is of concern in this study since this can lead to large uncertainties in the resulting pdf, which can also propagate in the estimated return periods. However, as discussed previously, the large majority of the pairs are estimated to be independent and the most important dependencies are captured at the first levels. Both together result in a virtual reduction in the dimensions of the pdf.

Table 2. Percentage of family types used for the first five levels of the R-Vine Model. Results for all levels can be found in the supplementary material.

Tree	Indep	Gaussian	Student t	Clayton	Gumbel	Frank	Joe	180°Clayton	180°Gumbel	180°Joe	90°Clayton	90°Gumbel	90°Joe	270°Clayton	270°Gumbel	270°Joe
1	0	9.5	13	0	48.5	11.8	7.7	5.3	2.4	1.8	0	0	0	0	0	0
2	35.1	8.3	6.5	4.2	7.7	5.4	4.2	9.5	1.8	3	3	0.6	1.2	4.8	1.8	3
3	43.7	3	3.6	6	4.2	16.2	3	5.4	1.8	1.2	2.4	0.6	2.4	3.6	1.2	1.8
4	40.4	4.8	8.4	4.2	5.4	6.6	2.4	4.2	3	4.2	1.8	2.4	4.2	2.4	1.8	3.6
5	48.5	7.3	3	1.2	1.8	10.3	6.1	4.8	3	1.8	0.6	0	3.6	3.6	1.2	3
All	64.2	3.7	2.2	2.6	1.7	5.4	2.6	2.8	1.2	2.4	2.5	0.9	2.1	2.3	1.0	2.2

Table 3. Goodness-of-fit values for the Cramer von Mises statistic based on different methods implemented in the VineCopula R package.

Method	CvM	p.val
ECP2	0.019	1
ECP	2.3	0.73
Breymann	1.8	0.185
Berg2	0.19	1
Berg	1.11	1

Sensitivity analysis indicates that the first few trees of the RVM drive the majority of the dependencies, as further discussed in Sect. 4.2.3. Moreover, the impact of the short sample size on the uncertainties in the results is quantified using a bootstrap technique, as described in the following section.

Goodness-of-fit (GOF) is calculated using the Cramer von Mises test, which compares the final selected RVM with the empirical copula. The RVineGofTest algorithm of the same R package implements different methods to compute the test, which however perform usually poorly in cases of small sample sizes and at higher dimensions as is the case for this work (Schepsmeier, 2013). Nevertheless, table 3 shows the GOF results for each of these methods. In all cases, the p.value is larger than 0.05, which is an indication that the fitted RVM cannot be rejected at a 5% significance level. However, given also the quite large p.values, a Type II error cannot be excluded. Nevertheless, the suitability of the model, in comparison to the empirical data, is further discussed in the the results section as well.

3.3.2 Stochastic simulation

The RVM is used to simulate 10K years of winter-seasons in the UK. For each year, the simulated AFI values at each grid cell depend on the other cells based on the fitted RVM. As an example, the AFI maps for the first six simulated winter-seasons are shown in Fig. 4.

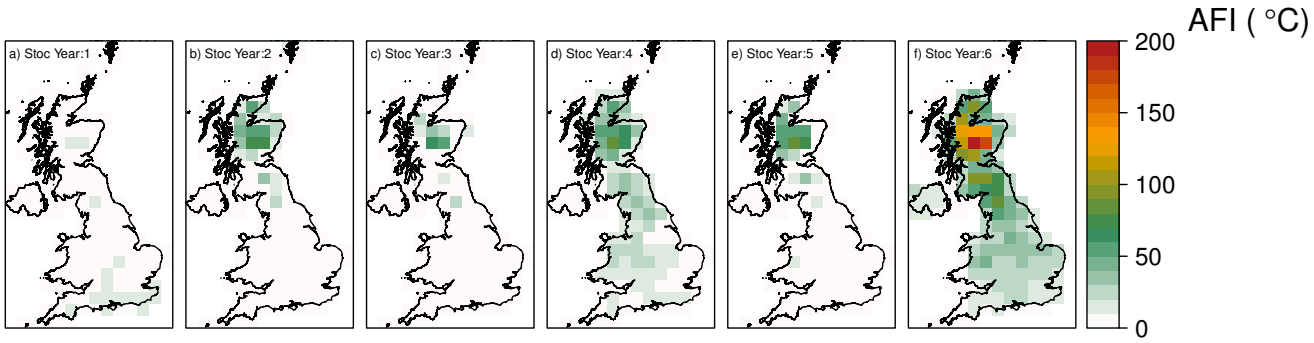

Figure 4. AFI maps (in °C) for the first six (out of 10K) years of the stochastic set.

Following Bevacqua et al. (2017), the model uncertainty is assessed using a parametric bootstrap approach, where a large number of models are created using as basis, instead of observations, randomly simulated data from the selected RVM. The resulting uncertainty accounts for both the uncertainty in the selected RVM and the uncertainty associated with the 10K years of the Monte Carlo simulation. In particular, confidence intervals are constructed as follows:

- A simulation with the same length as the observed data (i.e. 51 years) is repeated for 1,000 times.
- For each of these 1,000 simulations, a new RVM is fitted, whose structure is the same as the selected RVM (i.e. the one fitted with the observed data), while the pair-copula families and parameters are re-selected.
- For each of the resulting 1,000 RVMs, a simulation of 10K years of winter-seasons is performed and the corresponding return period levels are estimated.
- The uncertainty in the return levels is estimated by identifying the 95% confidence interval (i.e. the range 2.5–97.5 %) from these 1,000 return level curves.

The uncertainty associated solely due to the Monte Carlo sampling, is calculated as follows:

- Using the selected RVM, the simulation of 10K years of winter seasons is repeated for 1,000 times.
- For each of these 1,000 simulations, the corresponding return period levels are calculated.

– The uncertainty in the return periods is estimated by identifying the 95% confidence interval (i.e. the range 2.5–97.5 %) from these 1,000 return level curves.

4 Results and discussion

4.1 Return period maps

The obtained GEV fits for each cell (see Sect. 3.2) are used to create return period maps, as shown in Fig. 5. The top panels represent the AFI values that occur once every 10, 25, and 50 years. The largest AFI values follow, as expected, the orography of UK peaking in the northern region of the highlands where the elevation reaches 1000 meters. In this region, values of AFI greater than 200°C are reached often, approximately once every 5 years. The southern part of UK shows much lower values, not exceeding 200°C even at the 50 year RP. Slightly milder temperatures are also evident around the London area denoting the urban micro-climate effect. Other urban regions (e.g. Manchester or Midlands area) do not stand out as much as a result of the low grid resolution.

The empirical return periods are also plotted for comparison (bottom panels, Fig. 5). These are calculated for each cell as 1/(1-P), where P represents the cumulative probabilities of the ranked values and is calculated based on the Weibull formula P=i/(n+1) (Makkonen, 2006). The AFI values from the GEV fits correspond well with the empirical estimates, apart from the southern part of UK where the empirical values are approximately 20-30% larger at 50 years RP. This difference is driven by the 1962/63 event which empirically is estimated at 1 in 52 years while it is estimated to be less frequent according to the GEV fits. The probability of such an event happening today and its influence in the inhabited areas is discussed in detail in Sect. 4.2.

At higher return periods (100, 200, and 500 years, top panels of Fig. 6), AFI values exceeding 300°C are predicted to be able to occur not only in the north but in the southern part of UK, as well. The extreme AFI values in the south are again driven by the 1962/63 winter; exluding this winter from the analysis results in almost two times lower AFI values in most of the region (bottom panels in Fig. 6).

4.2 Regional return period AFI curves

The vine copula methodology permits the estimation of the hazard return periods over aggregated regions in the UK, which is particularly useful for insurance portfolio loss analysis. Results here are shown, apart from the entire UK, for three latitudinal regions: South England, North England & Northern Ireland, and Scotland.

Since our focus is mainly on inhabited areas, for each simulation year and for each region, I compute the "weighted AFI" (wAFI), where the AFI value at each cell j is weighted by the corresponding number of residential properties (n_j), as shown in Eq. 10. The weighted AFI thus places more weight on the hazard over large populated urban areas than agricultural or mountaineous areas. The number of residential properties in the UK is taken from the PERILS Industry Exposure Database (https://www.perils.org/), which contains up-to-date high quality insurance market data at Cresta level ("Catastrophe Risk

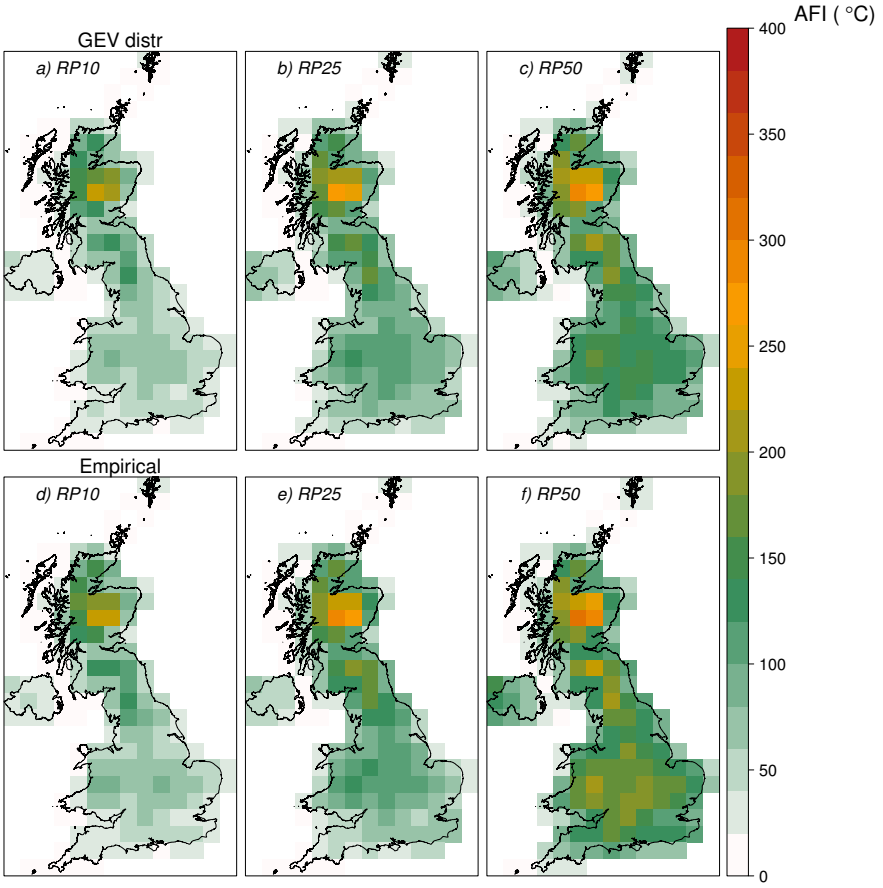

Figure 5. (Top panels) Maps of AFI values (in °C) for return periods of a) 1 in 10, b) 1 in 25, and c) 1 in 50 years calculated assuming a GEV distribution. (Bottom panels) Maps of the corrsponding empirical AFI values.

Evaluation and Standardizing Target Accumulations", https://www.cresta.org/) based on data directly collected from insurance companies writing property business in the UK.

$$wAFI_{year} = \frac{\sum AFI_{j,year} \cdot n_j}{\sum n_j} \tag{10}$$

Return period wAFI curves for both the empirical and the stochastic data is shown in Fig. 7. The empirical return periods calculation is described in Sect. 4.1, while their uncertainty intervals are computed via the 2.5^{th} and 97.5^{th} quantile of the beta probability distribution function (Folland and Anderson, 2002). The stochastic curve and confidence intervals are computed as described in Sect. 3.3.2. Analogous return period plots based on the UK average AFI (mAFI), i.e. without weighting, can be found in the Appendix (Fig. A1).

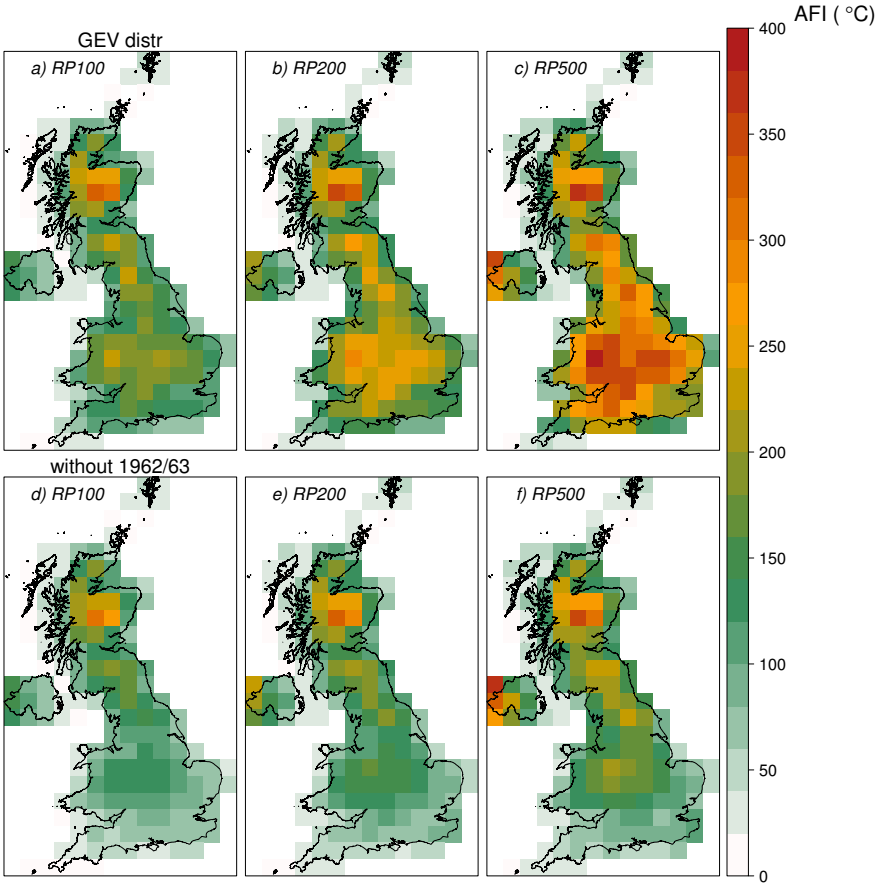

Figure 6. (Top panels) Maps of AFI values (in °C) for return periods of a) 1 in 100, b) 1 in 200, and c) 1 in 500 years. (Bottom panels) Same as top panels but without taking into account the extreme winter of 1962/63.

4.2.1 The 1962/63 winter return period

The stochastic curve (black line, Fig. 7a) follows closely the empirical one (grey line and circles) apart from the last historical point (with a wAFI of 141°C), which corresponds to the 1962/63 severe cold winter. This event is estimated empirically as 1 in 52 years event, but with a large uncertainty around this estimate due to the small size of the historical record, as shown by the uncertainy lines in Fig. 7a. In the stochastic set, such an extreme winter represents a larger return period of 89 years, with 95% confidence intervals of 81 to 120 years (see table 4). Especially for Southern England (Fig. 7b), this winter has been particularily rare; the model suggests a return period of 1 in 96 (91-138) years, which corresponds well with other independent point measurements. For example, according to the Central England Temperature (CET) record, the oldest continuously running temperature dataset in the world (Manley, 1974), only two other winters (1683/84 and 1739/40) have been colder than 1962/63 in the last 350 years, suggesting a return period in the range of 110-120 years.

Table 4. Empirical and model return periods in years, including 95% confidence interval ranges in parenthesis, for the 1962/63 winter in the UK.

Region	Empirical RP (95% CI)	Model RP (95% CI)
UK	52 (14-2015)	89 (81-120)
England S.		96 (91-138)
England N. & N. Ireland		79 (70-101)
Scotland		55 (49-65)

However, recent studies suggest that cold weather in the UK is likely to be less severe, to occur less frequently, and to last for a shorter period of time than was historically the case due to anthropogenic induced climate change (on Climate Change, 2017), although this is still under debate. Massey et al. (2012) used climate model simulations to demonstrate that cold December temperatures in the UK are now half as likely as they were in the 1960s. Christidis and Stott (2012) also indicate that human influence has reduced the probability of such a severe winter in UK by at least 20% and possibly by as much as 4 times, with a best estimate that the probability has been halved. On the other hand, recent studies have argued that warming in the Arctic could favor the occurrence of cold winter extremes, and might have been also responsible for the unusually cold winters in the UK of 2009/10 and 2010/11 (Francis and Vavrus, 2012; Tang et al., 2013). This hypothesis though is still largely under debate, see for example Barnes and Screen (2015) and Wallace et al. (2014).

As shown in Fig. 7b, South England is in general warmer than the North England and Northern Ireland region, partially driven by the urban micro-climate effect of the London area. The 1962/63 winter was less extreme in this region (wAFI of 139 °C) with an estimated return period of 1 in 79 years. On the other hand, Scotland is usually significantly colder than the rest of UK, reaching for example AFI values of 100 °C almost 2 times more often. However, the curve flattens out shortly after that and more extreme winters are estimated to be less probable than in the South. The winter of 1962/63 was less extreme in comparison to the south and therefore shows a lower return period of approximately once every 55 years.

4.2.2 NAO influence

The winters 1962/63, 1985/86, 2009/10, 2010/11 were associated with a negative NAO phase (Murray, 1966; Osborn, 2011). The NAO has a profound effect on winter climate variability around the Atlantic basin, accounting more than half of the year-to-year variability in winter surface temperature over UK (Scaife et al., 2005; Scaife and Knight, 2008). Not surprising, the average AFI over the entire UK is found to be significantly anti-correlated (ρ = -0.59, pval=4.810^{-6}) with the winter (December through March) station-based NAO index (NAOI) (Hurrell, 2017), as shown in Fig. 2.

In order to investigate this further, a generalized linear model (GLM) has been introduced into the location parameter of the GEV distributions. More precisely, the location parameter for each cell has been defined as a function of the NAOI: $\mu = \beta_0 + \beta_1 NAOI$ (see Eq. 1 in Sect. 3.3). However, the non-stationary fits were statistically similar to the default model, with

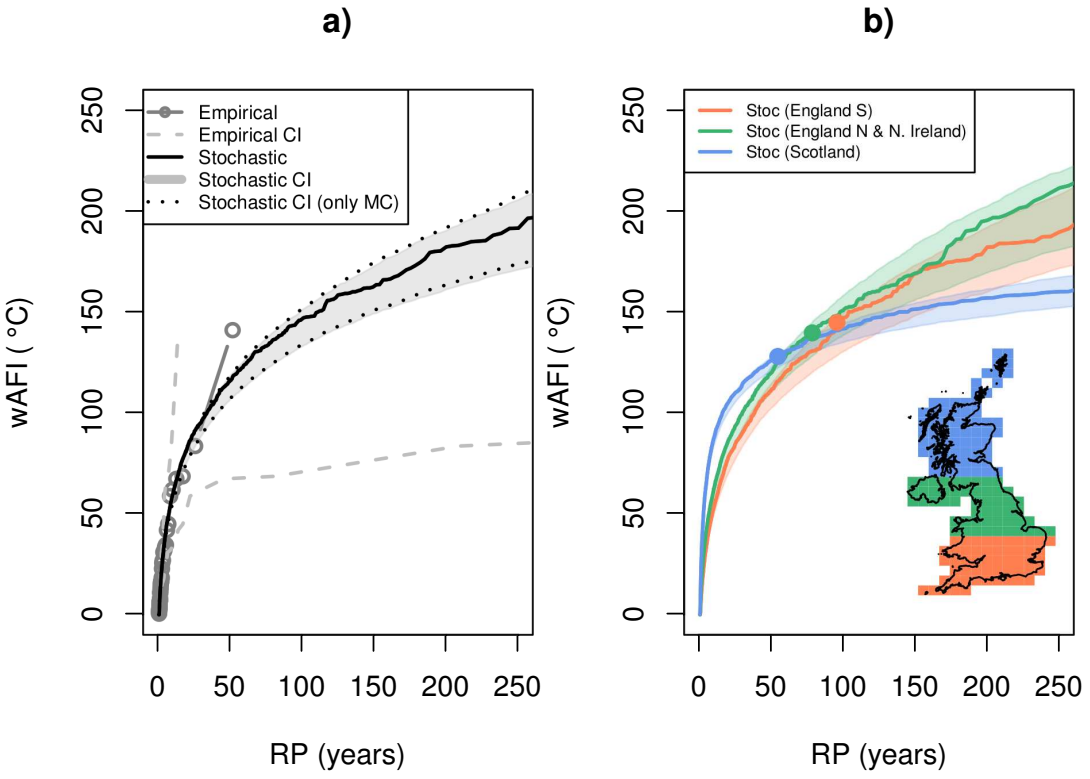

Figure 7. a) Return period curves of wAFI (in °C) based on the historical data (grey) and the stochastic model (black). The 95% confidence intervals are shown as dashed grey lines for the historical data and as a shaded grey area for the stochastic model. The dotted black lines represent the 95% confidence intervals of the stochastic model resulting from the length of the Monte Carlo simulation alone.
b) Modelled return period curves of wAFI (in °C), for South England (red), North England & Northern Ireland (green), and Scotland (blue). The shaded areas denote the respective 95% confidence intervals for each region. The circles specify the winter 1962/63 values on the curves.

β_1 parameters not significantly different from zero. This is probably related to the quite noisy character of the phenomenon and the relatively short historical record used in this study, which makes it difficult to discern the statistical differences in the extreme temperatures between positive and negative NAO winters. The effect of NAO in the hazard dependency structure has not been taken into account here. Recently, Bevacqua et al. (2017); Bevacqua (2017) developed a methodology that offers the possibility to include such meteorological predictors in a vine copula model and is something to be addressed in a future study. Because of its intrinsic chaotic behaviour, NAO is difficult (if even possible) to be predicted (Kushnir et al., 2006). Nevertheless, numerical seasonal forecast systems are currently rapidly improving and have even shown some success in the past (Graham et al., 2006; Folland et al., 2006). Incorporating such information in models could be very useful from the catastrophe risk management perspective.

4.2.3 Model uncertainty and sensitivity analysis

Any modelling results need to be interpreted being aware of their uncertainties. In this study, the main uncertainty in the model and the subsequent return periods stems from four sources:

(a) the uncertainty due to the short historical record length

(b) the uncertainty resulting from the length of the Monte Carlo simulation (i.e. the number of simulated winters)

(c) the uncertainty in the joint pdf (i.e. in the RVM) due to the limited historical record

(d) the uncertainty due to the model assumption of a stationary climate

The uncertainty in the historical data is sunbstantial as shown by the dashed grey lines in Fig. 7a (notice that measurement and interpolation errors in the historical dataset are assumed to be negligible). As an example, the estimated RP 95% confidence interval infor the 1962/63 winter ranges from 14 to 2015 years. Extreme-value theory is considered as a state-of-the-art procedure to find values for return periods that amply exceed the record length and has been utilized in this study. However, a common difficulty with extremes is that, by definition, data is rare and as a result, the shorter the record length, the more inaccurate is the estimation of the GEV parameters. As a (rather extreme) example, excluding the 1962/63 winter from the analysis would result in a significant reduction of the modelled probability of the occurrence of such an extreme winter to 1 in 831 years. A longer therefore record is needed to reduce this uncertainty. Meteorological reanalysis datasets could provide a comprehensive and consistent gridded temperature dataset over a very long period (e.g. Poli et al. (2016); Compo et al. (2011)), but higher spatial and temporal resolution is required in order to accurately calculate the air freezing index.

The dotted black lines in Fig. 7a show the estimated confidence intervals due to number of years used for the Monte Carlo simulation alone. The obtained uncertainty is important but smaller in comparison to the uncertainty in the empirical curve, i.e. directly from the observed data. Furthermore, the accuracy can be improved by increasing the number of simulated years, but at a computation cost.

The confidence intervals resulting from the Monte Carlo simulation almost entirely account for the model uncertainty estimated using the parametric bootstrap approach (i.e. which encompasses both (a) and (b) types of uncertainty, see Sect. 3.3.2), suggesting that the uncertainty in the RVM is negligible. As mentioned in Sect. 3.3.1, the reason of the small uncertainy in the RVM is twofold: the large majority of the pairs are estimated to be independent and also the most important dependencies are captured at the first trees. Both reasons lead to a virtual reduction in the dimensions of the pdf. Figure 8a shows the return period plots for the same RVM but truncated above the first seven levels (i.e. using independent copulas above level 1, 2, 3, up to 7). The same seed as for the default RVM is used in the simulation of these truncated models in order to avoid differences associated with the Monte Carlo sampling. The return period curves are quite similar for the RVMs with truncation above level 2, indicating that the first two levels capture most of the dependency structure.

Another source of uncertainty stems from the fact that the model has been developed under the assumption of a stationary climate, i.e. that the climate has not changed significantly during the last 51 years. Despite the observed winter warming in

the UK during the last decades (Jenkins et al., 2009), its effects in the frequency and magnitude of extreme cold spells is still unclear, as also discussed in Sect. 4.2. To test the non-stationarity assumption, a linear covariate is incorporated in the location parameter of the GEV distributions ($\mu = \beta_0 + \beta_1 year$) in order to account for an annual trend in AFI. The resulting β_1 parameters were not significantly different from zero, indicating an unsubstantial linear trend in AFI during the last five decades. Due to its high year-to-year variability (see Fig. 2), longer monitoring records are needed to identify statistically significant trends.

Additional sensitivity tests have been performed in order to investigate the influence of selected RVM copula families and parameters to the estimated return periods. Figure 8b shows the return period curves based on a range of vine copula models that are fitted using subset or single copula families. In comparison to the selected RVM (black line), the choice of a single copula family in the fitting process reduces significantly the co-occurrence of extreme values in the case of Gaussian, Clayton, and Frank copulas. This is to be expected since all those copulas do not show upper tail dependence in the limits. In fact, away from the extremes, the Gaussian copula shows greater upper tail concentration than Frank or Clayton copulas, as also found in the RP results. On the other hand, the copulas that show upper tail dependence (Joe, Gumbel, and Student t), lead to results that are comparable to the default model, which indicates that a selection of a more parsimonious model might be possible. Finally, including the two-parameter Archimedean copulas in the model fitting process (red line in Fig. 8b) also has minor impact in the estimated return periods in comparison to the default model.

5 Conclusions

This paper presents a probabilistic model of extreme cold winters in the United Kingdom. The hazard is modeled using the Air Freezing Index, an index which takes account both the magnitude and the duration of air temperature below freezing and is calculated from temperature data from the last 51 years. Extreme value theory has been applied in order to estimate the probability of extreme cold winters spatially across the UK. More importantly, the spatial dependence between regions in the UK has been assessed through a novel approach which takes advantage of the vine copula methodology. This approach allows the modeling of concurrent high AFI values across the country which is necessary in order to assess reliably the extreme behaviour of such events.

A stochastic set of 10K years is generated which is used to estimate the return period of extreme cold winters in UK, such as the "Big Freeze of 1962/63". According to the model, such a cold winter is estimated to occur once every 89 years in UK, with 95% confidence intervals ranging from 81 to 120 years. Especially for South England, this winter has been particularily rare with a return period equal to 1 in 96 years. It is important to note, though, that considerable uncertainty exists in these estimates. First and foremost, the 52-year historical record used in this study is short in order to estimate with enough confidence the frequencies of such extreme events. A longer record of temperature data would be necessary in order to reduce the uncertainty and high quality long-term reanalysis products could help towards this direction. Additional uncertainty in the model stems from the impacts of our changing climate due to anthropogenic forcing, but further research is necessary in order to discern how exactly extreme winter temperatures are affected in the UK. Singificant improvements are expected to come

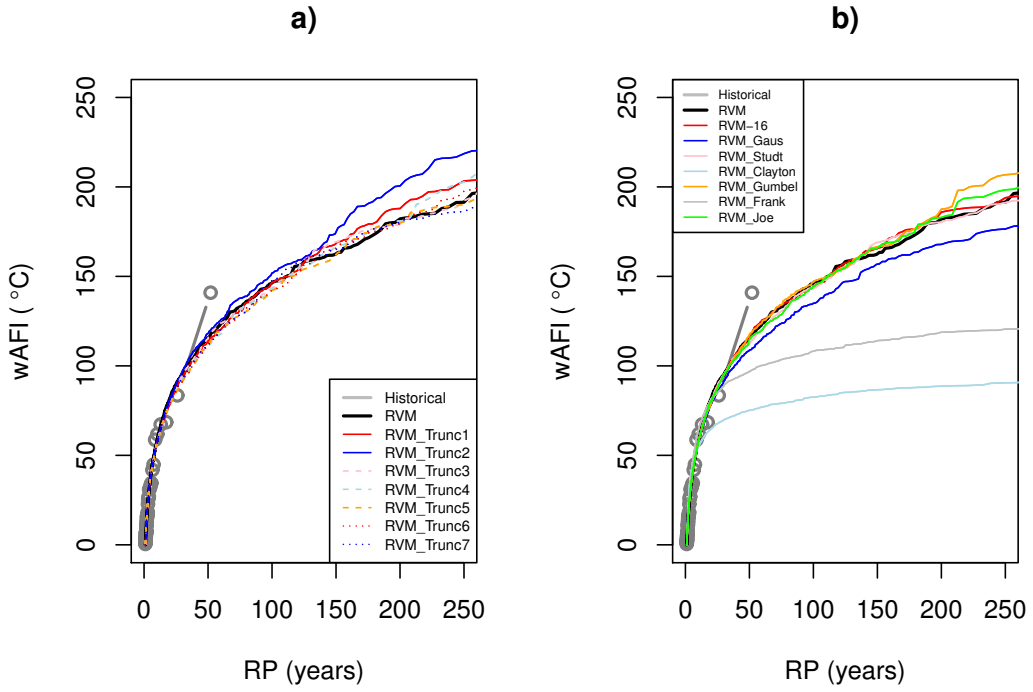

Figure 8. a) wAFI RP curves for the whole UK. The empirical curve is shown in grey. Stochastic results are shown for trucated RVMs above levels 1 to 7 (colored lines) or with no truncation (default, black line).
b) Sensitivity tests for the RP curves of wAFI based on RVM fitted using: all available copula families (black line), all but the two-parameter Archimedean copulas (i.e. default RVM, red line), and only one Copula family each time, i.e. Gaussian (blue line line), Student's t (dashed grey), Clayton (dotted grey line), Gumbel (dotdash grey line), Frank (red line), and Joe (grey line).

with increasing availability of data, increasing understanding of the science, and with advancements in computing capability and technology. This model is part of a probabilistic catastrophe model for insurance losses due to burst pipes resulting from freezing temperatures.

Appendix A

A1

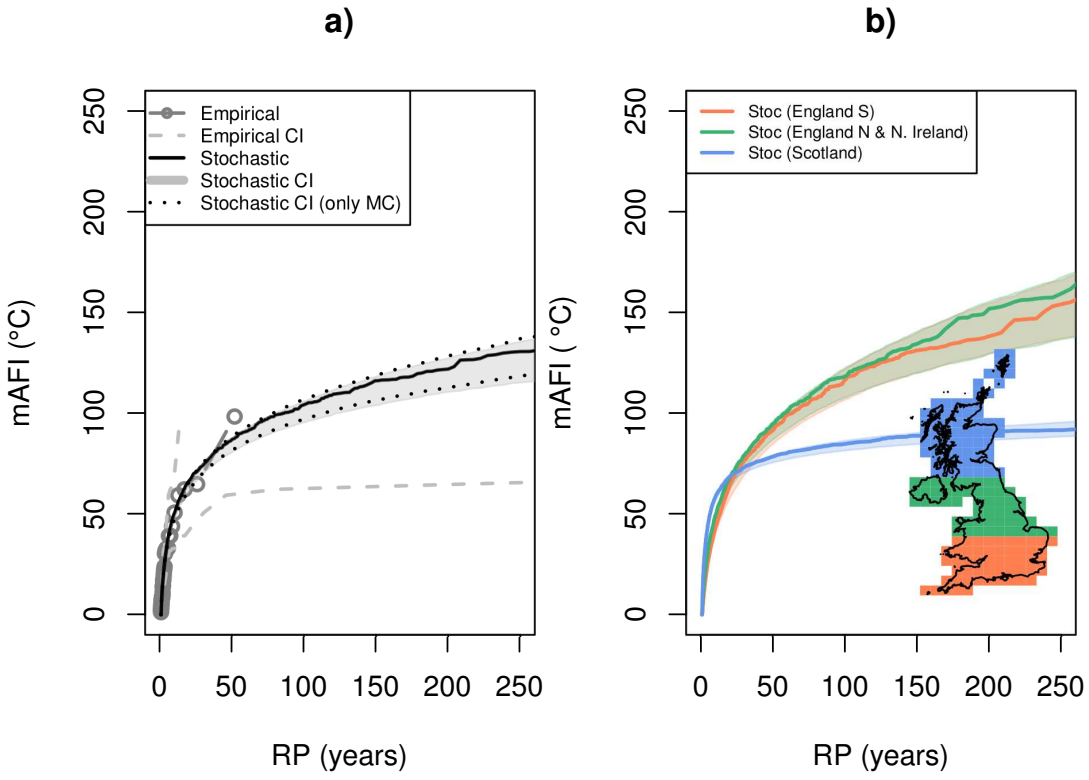

Figure A1. Similar to Fig. 7 but for mAFI (without weighting, in °C).

Disclaimer. TEXT

Acknowledgements. TEXT

References

Aas, K., Czado, C., Frigessi, A., and Bakken, H.: Pair-copula constructions of multiple dependence,, Tech. rep., Munich University, Institute for Statistics, http://epub.ub.uni-muenchen.de/, 2006.

Abbara, O. and Zevallos, M.: Assessing stock market dependence and contagion, Quantitative Finance, 14, 1627–1641, doi:10.1080/14697688.2013.859390, https://doi.org/10.1080/14697688.2013.859390, 2014.

ABI: Industry Data Downloads, Tech. rep., Association of British Insurers, https://www.abi.org.uk/data-and-resources/industry-data/free-industry-data-downloads/, 2017.

AIR: About Catastrophe Models, Tech. rep., AIR Worldwide, https://www.air-worldwide.com/Publications/Brochures/documents/About-Catastrophe-Models/, 2012.

Barnes, E. A. and Screen, J. A.: The impact of Arctic warming on the midlatitude jet-stream: Can it? Has it? Will it?, Wiley Interdisciplinary Reviews: Climate Change, 6, 277–286, doi:10.1002/wcc.337, http://dx.doi.org/10.1002/wcc.337, 2015.

Bedford, T. and Cooke, R. M.: Vines–a new graphical model for dependent random variables, Ann. Statist., 30, 1031–1068, doi:10.1214/aos/1031689016, https://doi.org/10.1214/aos/1031689016, 2002.

Bevacqua, E.: CDVineCopulaConditional: Sampling from Conditional C- and D-Vine Copulas. R package version 0.1.0, R package version 0.1.0, https://cran.r-project.org/web/packages/CDVineCopulaConditional/index.html, 2017.

Bevacqua, E., Maraun, D., Hobæk Haff, I., Widmann, M., and Vrac, M.: Multivariate statistical modelling of compound events via pair-copula constructions: analysis of floods in Ravenna (Italy), Hydrology and Earth System Sciences, 21, 2701–2723, doi:10.5194/hess-21-2701-2017, https://www.hydrol-earth-syst-sci.net/21/2701/2017/, 2017.

Bilotta, R., Bell, J. E., Shepherd, E., and Arguez, A.: Calculation and Evaluation of an Air-Freezing Index for the 1981–2010 Climate Normals Period in the Coterminous United States, Journal of Applied Meteorology and Climatology, 54, 69–76, doi:10.1175/JAMC-D-14-0119.1, http://dx.doi.org/10.1175/JAMC-D-14-0119.1, 2015.

Bonazzi, A., Cusack, S., Mitas, C., and Jewson, S.: The spatial structure of European wind storms as characterized by bivariate extreme-value Copulas, Natural Hazards and Earth System Sciences, 12, 1769–1782, doi:10.5194/nhess-12-1769-2012, https://www.nat-hazards-earth-syst-sci.net/12/1769/2012/, 2012.

Bowman, G., Coburn, A., and Ruffle, S.: Freeze - Profile of a Macro-Catastrophe Threat Type, Cambridge centre for risk studies working paper series, Cambridge Risk Framework, https://hazdoc.colorado.edu/handle/10590/6787?show=full, 2012.

Brechmann, E. C. and Schepsmeier, U.: Modeling Dependence with C- and D-Vine Copulas: The R Package CDVine, Journal of Statistical Software, doi:10.18637/jss.v052.i03, https://www.jstatsoft.org/article/view/v052i03, 2013.

Cawthorne, R. A. and Marchant, J. H.: The effects of the 1978/79 winter on British bird populations, Bird Study, 27, 163–172, doi:10.1080/00063658009476675, http://dx.doi.org/10.1080/00063658009476675, 1980.

Christidis, N. and Stott, P. A.: Lengthened odds of the cold UK winter of 2010/2011 attributable to human influence [in "Explaining Extreme Events of 2011 from a Climate Perspective"], Bulletin of the American Meteorological Society, 93, 1060–1062, doi:10.1175/BAMS-D-12-00021.1, https://doi.org/10.1175/BAMS-D-12-00021.1, 2012.

Compo, G. P., Whitaker, J. S., Sardeshmukh, P. D., Matsui, N., Allan, R. J., Yin, X., Gleason, B. E., Vose, R. S., Rutledge, G., Bessemoulin, P., Brönnimann, S., Brunet, M., Crouthamel, R. I., Grant, A. N., Groisman, P. Y., Jones, P. D., Kruk, M. C., Kruger, A. C., Marshall, G. J., Maugeri, M., Mok, H. Y., Nordli, O., Ross, T. F., Trigo, R. M., Wang, X. L., Woodruff, S. D., and Worley, S. J.: The Twentieth Century

Reanalysis Project, Quarterly Journal of the Royal Meteorological Society, 137, 1–28, doi:10.1002/qj.776, http://dx.doi.org/10.1002/qj.776, 2011.

Czado, C.: Pair-Copula Constructions of Multivariate Copulas. In: Jaworski P., Durante F., Härdle W., Rychlik T. (eds) Copula Theory and Its Applications. Lecture Notes in Statistics, Springer, doi:https://doi.org/10.1007/978-3-642-12465-5_4, 2010.

Czado, C., E.C.Brechmann, and Gruber, L.: Selection of Vine Copulas. In: Copulae in Mathematical and Quantitative Finance. Lecture Notes in Statistics, Springer, 2013.

D., S. and Schirmacher, E.: Multivariate Dependence Modeling Using Pair-Copulas, Tech. rep., Society of Actuaries., 2008.

Dißmann, J., Brechmann, E., Czado, C., and Kurowicka, D.: Selecting and estimating regular vine copulae and application to financial returns, Computational Statistics & Data Analysis, 59, 52 – 69, doi:https://doi.org/10.1016/j.csda.2012.08.010, http://www.sciencedirect.com/science/article/pii/S0167947312003131, 2013.

Durante, F. and Sempi, C.: Principles of copula theory, CRC/Chapman & Hall, 2015.

Folland, C. and Anderson, C.: Estimating Changing Extremes Using Empirical Ranking Methods, Journal of Climate, 15, 2954–2960, doi:10.1175/1520-0442(2002)015<2954:ECEUER>2.0.CO;2, 2002.

Folland, C. K., Parker, D. E., Scaife, A. A., Kennedy, J. J., Colman, A. W., Brookshaw, A., Cusack, S., and Huddleston, M. R.: The 2005/06 winter in Europe and the United Kingdom: Part 2 –Prediction techniques and their assessment against observations, Weather, 61, 337–346, doi:10.1256/wea.182.06, http://dx.doi.org/10.1256/wea.182.06, 2006.

Francis, J. A. and Vavrus, S. J.: Evidence linking Arctic amplification to extreme weather in mid-latitudes, Geophysical Research Letters, 39, n/a–n/a, doi:10.1029/2012GL051000, http://dx.doi.org/10.1029/2012GL051000, l06801, 2012.

Frauenfeld, O. W., Zhang, T., and Mccreight, J. L.: Northern Hemisphere freezing/thawing index variations over the twentieth century, International Journal of Climatology, 27, 47–63, doi:10.1002/joc.1372, http://dx.doi.org/10.1002/joc.1372, 2007.

Genest, C. and Favre, A.-C.: Everything You Always Wanted to Know about Copula Modeling but Were Afraid to Ask, Journal of Hydrologic Engineering, 12, 347–368, doi:10.1061/(ASCE)1084-0699(2007)12:4(347), https://ascelibrary.org/doi/abs/10.1061/%28ASCE%291084-0699%282007%2912%3A4%28347%29, 2007.

Gordon, J. R.: An Investigation into Freezing and Bursting Water Pipes in Residential Construction, Research report 96-1, Building Research Council. School of Architecture. College of Fine and Applied Arts. University of Illinois at Urbana-Champaign, http://hdl.handle.net/2142/54757, 1996.

Graham, R. J., Gordon, C., Huddleston, M. R., Davey, M., Norton, W., Colman, A., Scaife, A. A., Brookshaw, A., Ingleby, B., McLean, P., Cusack, S., McCallum, E., Elliott, W., Groves, K., Cotgrove, D., and Robinson, D.: The 2005/06 winter in Europe and the United Kingdom:Part 1 –How the Met Office forecast was produced and communicated, Weather, 61, 327–336, doi:10.1256/wea.181.06, http://dx.doi.org/10.1256/wea.181.06, 2006.

Guirguis, K., Gershunov, A., Schwartz, R., and Bennett, S.: Recent warm and cold daily winter temperature extremes in the Northern Hemisphere, Geophysical Research Letters, 38, n/a–n/a, doi:10.1029/2011GL048762, http://dx.doi.org/10.1029/2011GL048762, l17701, 2011.

Haff, I. H., Frigessi, A., and Maraun, D.: How well do regional climate models simulate the spatial dependence of precipitation? An application of pair-copula constructions, Journal of Geophysical Research: Atmospheres, 120, 2624–2646, doi:10.1002/2014JD022748, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JD022748, 2015.

Hurrell, J.: The Climate Data Guide: Hurrell North Atlantic Oscillation (NAO) Index (station-based)., Tech. rep., National Center for Atmospheric Research, https://climatedataguide.ucar.edu/climate-data/hurrell-north-atlantic-oscillation-nao-index-station-based, 2017.

Jenkins, G., Perry, M., and Prior, J.: The climate of the UK and recent trends, Tech. rep., Hadley Centre, Met Office, Exeter, http://ukclimateprojections.metoffice.gov.uk/media.jsp?mediaid=87932&filetype=pdf, 2009.

Joe, H.: Dependence Modeling with Copulas. CRC Monographs on Statistics & Applied Probability., Chapman & Hall, London, 2014.

Kemp, M.: Tail Weighted Probability Distribution Parameter Estimation, Tech. rep., Nematrian Limited, http://www.nematrian.com/Docs/TailWeightedParameterEstimation.pdf, 2016.

Kushnir, Y., Robinson, W. A., Chang, P., and Robertson, A. W.: The Physical Basis for Predicting Atlantic Sector Seasonal-to-Interannual Climate Variability, Journal of Climate, 19, 5949–5970, doi:10.1175/JCLI3943.1, http://dx.doi.org/10.1175/JCLI3943.1, 2006.

Makkonen, L.: Plotting Positions in Extreme Value Analysis, Journal of Applied Meteorology and Climatology, 45, 334–340, doi:10.1175/JAM2349.1, https://doi.org/10.1175/JAM2349.1, 2006.

Manley, G.: Central England temperatures: Monthly means 1659 to 1973, Quarterly Journal of the Royal Meteorological Society, 100, 389–405, doi:10.1002/qj.49710042511, http://dx.doi.org/10.1002/qj.49710042511, 1974.

Massey, N., Aina, T., Rye, C., Otto, F., Wilson, S., Jones, R., and Allen, M.: Have the odds of warm November temperatures and of cold December temperatures in Central England changed. Bulletin of the American Meteorological Society 93, 1057-1059, Bulletin of the American Meteorological Society, 93, 1057–1059, 2012.

McDonald, A., Bschaden, B., Sullivan, E., and Marsden, R.: Mathematical simulation of the freezing time of water in small diameter pipes, Applied Thermal Engineering, 73, 142 – 153, doi:https://doi.org/10.1016/j.applthermaleng.2014.07.046, http://www.sciencedirect.com/science/article/pii/S135943111400622X, 2014.

Meucci, A.: A Short, Comprehensive, Practical Guide to Copulas, GARP Risk Professional, pp. 22–27, https://ssrn.com/abstract=1847864orhttp://dx.doi.org/10.2139/ssrn.1847864, 2011.

Murray, R.: A note on the large scale features of the 1962/63 winter, Meteorol. Mag., 95, 339–348, 1966.

Nelsen, R. B.: An Introduction to Copulas, Springer-Verlag, New York., 2006.

on Climate Change, C.: UK Climate Change Risk Assessment 2017. Synthesis report: priorities for the next fi e years, Tech. rep., Committee on Climate Change, https://www.theccc.org.uk/wp-content/uploads/2016/07/UK-CCRA-2017-Launch-slidepack.pdf, 2017.

Osborn, T. J.: Winter 2009/2010 temperatures and a record-breaking North Atlantic Oscillation index, Weather, 66, 19–21, doi:10.1002/wea.660, http://dx.doi.org/10.1002/wea.660, 2011.

Perry, M. and Hollis, D.: The generation of monthly gridded datasets for a range of climatic variables over the UK, International Journal of Climatology, 25, 1041–1054, doi:10.1002/joc.1161, http://dx.doi.org/10.1002/joc.1161, 2005.

Perry, M., Hollis, D., and Elms, M.: The generation of daily gridded datasets of temperature and rainfall for the UK, Tech. rep., National Climate Information Centre, 2009.

Poli, P., Hersbach, H., Dee, D. P., Berrisford, P., Simmons, A. J., Vitart, F., Laloyaux, P., Tan, D. G. H., Peubey, C., Thépaut, J.-N., Trémolet, Y., Hólm, E. V., Bonavita, M., Isaksen, L., and Fisher, M.: ERA-20C: An Atmospheric Reanalysis of the Twentieth Century, Journal of Climate, 29, 4083–4097, doi:10.1175/JCLI-D-15-0556.1, https://doi.org/10.1175/JCLI-D-15-0556.1, 2016.

Salvadori, G. and Michele, C. D.: On the Use of Copulas in Hydrology: Theory and Practice, Journal of Hydrologic Engineering, 12, 369–380, doi:10.1061/(ASCE)1084-0699(2007)12:4(369), https://ascelibrary.org/doi/abs/10.1061/%28ASCE%291084-0699%282007%2912%3A4%28369%29, 2007.

Salvadori, G., Tomasicchio, G., and D'Alessandro, F.: Practical guidelines for multivariate analysis and design in coastal and off-shore engineering, Coastal Engineering, 88, 1 – 14, doi:https://doi.org/10.1016/j.coastaleng.2014.01.011, http://www.sciencedirect.com/science/article/pii/S0378383914000209, 2014.

Salvadori, G., Durante, F., Tomasicchio, G., and D'Alessandro, F.: Practical guidelines for the multivariate assessment of the structural risk in coastal and off-shore engineering, Coastal Engineering, 95, 77 – 83, doi:https://doi.org/10.1016/j.coastaleng.2014.09.007, http://www.sciencedirect.com/science/article/pii/S0378383914001811, 2015.

Scaife, A. A. and Knight, J. R.: Ensemble simulations of the cold European winter of 2005-2006, Quarterly Journal of the Royal Meteorological Society, 134, 1647–1659, doi:10.1002/qj.312, http://dx.doi.org/10.1002/qj.312, 2008.

Scaife, A. A., Knight, J. R., Vallis, G. K., and Folland, C. K.: A stratospheric influence on the winter NAO and North Atlantic surface climate, Geophysical Research Letters, 32, n/a–n/a, doi:10.1029/2005GL023226, http://dx.doi.org/10.1029/2005GL023226, l18715, 2005.

Schepsmeier, U.: Estimating standard errors and efficient goodness-of- t tests for regular vine copula models, Ph.D. thesis, Fakultat fur Mathematik Technische Universitat Munchen, http://mediatum.ub.tum.de/doc/1175739/document.pdf, 2013.

Schepsmeier, U., Stoeber, J., Brechmann, E. C., Graeler, B., Nagler, T., and Erhardt, T.: VineCopula-package: Statistical Inference of Vine Copulas, Tech. rep., R package, https://cran.r-project.org/web/packages/VineCopula/VineCopula.pdf, 2017.

Seager, R., Kushnir, Y., Nakamura, J., Ting, M., and Naik, N.: Northern Hemisphere winter snow anomalies: ENSO, NAO and the winter of 2009/10, Geophysical Research Letters, 37, n/a–n/a, doi:10.1029/2010GL043830, http://dx.doi.org/10.1029/2010GL043830, l14703, 2010.

Sklar, A.: Fonctions de répartition à n dimensions et leurs marges, Publications de l'Institut de Statistique de L'Université de Paris, 8, 229–231, 1959.

Tang, Q., Zhang, X., Yang, X., and Francis, J. A.: Cold winter extremes in northern continents linked to Arctic sea ice loss, Environmental Research Letters, 8, 014 036, http://stacks.iop.org/1748-9326/8/i=1/a=014036, 2013.

Wallace, J. M., Held, I. M., Thompson, D. W. J., Trenberth, K. E., and Walsh, J. E.: Global Warming and Winter Weather, Science, 343, 729–730, doi:10.1126/science.343.6172.729, http://science.sciencemag.org/content/343/6172/729, 2014.

Walsh, J. E., Phillips, A. S., Portis, D. H., and Chapman, W. L.: Extreme Cold Outbreaks in the United States and Europe, 1948–99, Journal of Climate, 14, 2642–2658, doi:10.1175/1520-0442(2001)014<2642:ECOITU>2.0.CO;2, 2001.