

Modelling the effect of the El Niño-Southern Oscillation on extreme spatial temperature events over Australia

Hugo C. Winter¹, Jonathan A. Tawn² and Simon J. Brown³
Email: hugo.winter@edfenergy.com, j.tawn@lancaster.ac.uk,
simon.brown@metoffice.gov.uk

¹*EDF Energy R&D UK Centre, Cardinal Place, 80 Victoria Street, London, SW1E 5JL, UK.,*

²*Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4YF, U.K.,*

³*Met Office Hadley Centre, FitzRoy Road, Exeter, EX1 3PB, UK.*

Abstract

When assessing the risk posed by high temperatures, it is necessary to consider not only the temperature at separate sites but also how many sites are expected to be hot at the same time. Hot events that cover a large area have the potential to put a great strain on health services and cause devastation to agriculture, leading to high death tolls and much economic damage. South-eastern Australia experienced a severe heatwave in early 2009; 374 people died in the state of Victoria and Melbourne recorded its highest temperature since records began in 1859 (Nairn and Fawcett, 2013). One area of particular interest in climate science is the effect of large scale climatic phenomena, such as the El Niño-Southern Oscillation (ENSO), on extreme temperatures. Here, we develop a framework based upon extreme value theory to estimate the effect of ENSO on extreme temperatures across Australia. This approach permits us to estimate the change in temperatures with ENSO at important sites, such as Melbourne, and also whether we are more likely to observe hot temperatures over a larger spatial extent during a particular phase of ENSO. To this end, we design a set of measures that can be used to effectively summarise many important spatial aspects of an extreme temperature event. These measures are estimated using our extreme value framework and we validate whether we can accurately replicate the 2009 Australian heatwave, before using the model to estimate the probability of having a more severe event than has been observed.

Keywords: conditional extremes, covariates, El Niño-Southern Oscillation, extremal dependence, extreme temperature, severity-area-frequency curves, spatial extremes

1 Introduction

The early 2009 heatwave event was one of the most extreme to hit south-eastern Australia. Melbourne recorded its highest temperature since records began in 1859, at 46.4°C, and Adelaide its third highest temperature over the same observational period at 45.7°C. In total there were 374 heat related deaths in Victoria with over 2,000 people treated for heat related illness (Nairn and Fawcett, 2013). A particular challenge when modelling any environmental process across Australia

is the spatial distribution of the population and agricultural activity across the country. Four of the five largest cities are located on the coast in the south-eastern region and most agriculture occurs in the south-eastern region. A hot event occurring over this region will lead to increased mortality and economic losses. As such, for mitigation purposes, it is necessary to be able to give accurate estimates of the risk posed by high temperatures over specific regions of interest. Extreme value theory provides a statistical framework for modelling rare events. To model this problem sufficiently using extreme value statistics we require not only a univariate extreme value model that focuses on very high temperatures, but also a flexible model that accurately captures the spatial dependence between high temperatures at different sites.

There is much interest in how certain large-scale climatic phenomena will affect extreme events; both currently and under future climate change. One particular phenomenon known to affect the climate of Australia is the El Niño-Southern Oscillation (ENSO). It is a large-scale naturally occurring fluctuation in sea surface temperatures (SSTs) in the equatorial Pacific. Two limiting cases, corresponding to higher and lower SSTs in the equatorial Pacific Ocean, are called El Niño and La Niña respectively. During El Niño conditions, weaker easterly trade winds blowing across the Pacific can cause warm surface water to flow eastwards. This leads to increased convection in the central Pacific and reduces the amount of precipitation over Australia and other countries in southern Asia. In contrast, during La Niña conditions stronger trade winds blow warmer surface water to the west Pacific and cooler SSTs are observed in central and eastern Pacific regions (Wang and Picaut, 2004). The effect of ENSO on mean global temperatures has been well studied but the impact on extreme temperature is less well understood.

Looking at Australia specifically, Kenyon and Hegerl (2008) showed that El Niño conditions lead to increased temperatures over eastern and northern regions whereas during La Niña conditions temperatures will be lower over eastern and northern regions. Strong El Niño conditions do not guarantee higher temperatures and patterns are not uniform across space. The early 2009 heatwave event over south-eastern Australia occurred during a moderate La Niña event. The event covered much of southern and south-eastern Australia and as such had a great impact leading to record temperatures in certain places; this was not a uniform pattern across the whole of Australia with some regions affected by moderate heat only.

The aim of this study is to develop a better understanding of how ENSO has an effect on extreme temperatures over Australia. Perkins and Alexander (2013), Min et al. (2013) and Alexander and Arblaster (2009) have explored the effect of ENSO on the distribution of annual and seasonal maxima temperatures in Australia. They fit the generalised extreme value distribution with covariates in the location and scale parameters and map return level estimates over sites to produce spatial plots. However, none of these papers explicitly model spatial dependence and therefore cannot be used to estimate the probability of heatwave events occurring at multiple sites over space. We analyse the effect of ENSO on not only the marginal distribution of extreme temperatures, using more efficient threshold models (Coles, 2001), but also their spatial dependence structure.

Two approaches that have been used for the analysis of spatial processes, when extreme values are of interest, are geostatistics and max-stable processes. The broad area of geostatistics provides the most widely used approaches for spatial modelling and is based on an assumption that the process being modelled is Gaussian (Cressie, 1993). These approaches tend to focus on the main

body of data and as such can lead to misleading results when analysing rare events such as extreme temperatures. The most popular approach to spatial extreme value modelling is to fit a max-stable process to componentwise maxima, such as the site-by-site annual maxima. Key max-stable papers include Smith (1990), Coles (1993), Schlather (2002), Davison et al. (2012) and Dombry et al. (2013). A max-stable process arises as the limiting process derived by taking an affine normalisation of pointwise maxima over a sequence of n independent and identically distributed replicates of a random spatial process as $n \rightarrow \infty$. Max-stable models are often computationally intensive to fit and difficult to conditionally simulate from. More critically, both Gaussian and max-stable processes have restrictive features to their extremal dependence structure as explained below.

To help us consider what mathematical properties we require for our model for spatial extreme events we first introduce an important limiting pairwise measure of extremal dependence between random variables Y_1 and Y_2 with continuous distribution functions F_1 and F_2 respectively. The tail coefficient χ is given by

$$\chi = \lim_{p \rightarrow 1} \text{P}(F_2(Y_2) > p \mid F_1(Y_1) > p). \quad (1)$$

When $\chi > 0$, Y_1 and Y_2 are said to be asymptotically dependent, i.e., the conditional probability of concurrent extremes in Y_1 and Y_2 has some non-zero probability in the limit. The variables Y_1 and Y_2 are asymptotically independent when $\chi = 0$. All dependent max-stable processes are asymptotically dependent for all pairs of sites. In contrast all non-perfectly dependent Gaussian processes, are asymptotically independent for all pairs of sites (Ledford and Tawn, 1996). If a spatial process was asymptotically dependent for nearby sites and dependent but asymptotically independent for more separated sites neither a max-stable process or a Gaussian process could model its extreme events without leading to biased inferences.

To accurately model extremal dependence we build a flexible multivariate model based upon the conditional extremes approach of Heffernan and Tawn (2004) that fully takes into account spatial dependence on a spatial lattice within the framework of extreme value theory. The conditional extremes model leads to a class of multivariate distributions that allow for both asymptotic dependence and asymptotic independence between pairs of sites. As such, this model is suitable if the process being modelled is either max-stable or Gaussian whilst also permitting the analysis of more general processes. Thus this approach embeds both the two standard approaches within a general framework. A major benefit of the conditional extremes approach is that inference for extreme events does not require the choice of asymptotic dependence or asymptotic independence for each different pair of sites in advance. Uncertainty in estimates of extreme events derived from this method accounts for the evidence for each type of extremal dependence. Furthermore conditional simulation of extreme events is straightforward under this approach. The conditional extremes framework permits the estimation of not only joint extremes at different sites, but also how ENSO affects the spatial extent of a hot event.

To analyse the marginal effect of ENSO on extreme temperatures we will estimate the change in return levels at each site. Other existing measures of univariate temperature series quantify the effect of heat on mortality and other factors; see Alexander and Arblaster (2009) and Winter and Tawn (2016). A key contribution here is that we are the first to introduce spatial risk measures. Let the set of all sites be denoted by S , and the values of the daily maximum temperatures at these sites be $\mathbf{Y} = (Y_1, \dots, Y_l)$, where $|S| = l$. The most widely used existing measures are based on

equation (1), but they only describe the dependence between pairs of sites, they do not measure dependence when the variables are asymptotically independent, and they condition on a specific site being extreme. Our measures overcome these restrictions. We find the most informative spatial risk measure to be the severity-area-frequency (SAF) curve, which we adapt from drought analysis (Henriques and Santos, 1999). The SAF curve ($\gamma_j, j \geq 1$) gives the average marginal return period of an event at the j worst affected sites.

By analysing our suite of spatial risk measures we are able to explore the spatial extent of temperature extremes across Australia and see how the measures alter with ENSO conditions. We also test the validity of our approach by comparing risk measure values from observations from the heatwave event in early 2009 to simulations of hot days generated by our model, thus demonstrating that our model can capture such events accurately. We then illustrate how our approach can be used to estimate extremal features for rarer events than have been observed previously. Thus for the first time it is possible to answer questions such as *what is the probability of observing a spatial event more extreme than the early 2009 event?*

In Section 2 we introduce the gridded daily maximum temperature data along with the ENSO covariate. Section 3 presents the models for the margins and dependence structure of spatial extreme temperatures. The measures for assessing spatial risk are developed in Section 4. In Section 5 an approach for simulating spatial fields using the conditional extremes model is given. Results for the marginal and dependence parameters are provided in Section 6 along with estimates of important extremal measures. Finally, discussion and conclusions are given in Section 7.

2 Data and Exploratory Analysis

Daily maximum near-surface air temperatures for Australia are taken from HadGHCND, a global gridded dataset (<http://hadobs.metoffice.com/hadghcnd/>) of quality-controlled station observations compiled by the U.S. National Climatic Data Center (Caesar et al., 2006). An angular distance weighting technique is used to interpolate observed anomalies onto a 2.5° latitude by 3.75° longitude grid which results in 72 boxes covering Australia over the period 1957-2011. Whilst this is a relatively coarse resolution heatwaves are large meteorological phenomena and surface air temperatures have long correlation length scales, for which Caesar et al. (2006) found values of between 700km and 1400km for the 0°S to 30°S latitude band. For Australian surface air temperatures Avila et al. (2015) found that their extremal characteristics and correlations with ENSO were preserved across a range of grid resolutions from 0.25° to 2.5° . Hot days are most likely to occur in summer months, here defined as the 90 day period from December to February (91 day period for a leap year); these three month periods are extracted from each year. No missing data values exist within the summer months of the years for which the data are provided.

To measure the effect of ENSO the Niño3.4 index is used. This is a measurement of the monthly SST anomaly, with respect to the average SST for 1981-2010, in a region bounded by 5°N to 5°S and 170°W to 120°W . Other ways of measuring ENSO variability are available; for example the Southern Oscillation Index which is based on atmospheric changes as opposed to changes in SSTs (Jones and Trewin, 2000). However, Niño3.4 is commonly used to characterise ENSO (Kenyon and Hegerl, 2008). Large positive values of this index indicate El Niño events, whereas large negative values correspond to La Niña events. In this paper values of $+1^\circ\text{C}$ and -1°C are used to define

El Niño and La Niña events respectively. Our framework permits estimates for any value of Niño3.4.

To help determine our modelling strategy we first explore the spatial-temporal dependences between the daily maximum temperatures at the grid box corresponding to Melbourne. We denote this temperature at time t by $Y_{s_M,t}$, and at site s at time $t - h(s)$ by $Y_{s,t-h(s)}$ for all $s \in S$. Figure 1 shows the spatial cross-correlation function $\text{corr}(Y_{s_M,t}, Y_{s,t-h(s)})$, $s \in S$ and for two choices of $h(s)$. The left plot corresponds to time lag 0, i.e., $h(s) = 0$, for all $s \in S$, and the right to

$$\max_{h(s)} \text{corr}(Y_{s_M,t}, Y_{s,t-h(s)}) - \text{corr}(Y_{s_M,t}, Y_{s,t}),$$

with $h(s)$ the value that gives this maximum. These spatial-temporal dependence summaries show that Melbourne temperatures tend to be strongly positively correlated with concurrent temperatures at other most sites in Australia except for western and north-east regions. For the sites with positive correlation at lag $h(s) = 0$, then this is typically the temporal lag with the maximum dependence. However, for the other sites, the dependence is maximised for $h(s) = 1$ to 3 for s in the west and typically $h(s) = -1$ to -4 for s in the north-east. The correlations highlight that although spatial dependence does decay broadly with distance it is anisotropic and different processes appear to be active in coastal regions. Similar plots for data in high and low ENSO phases give very similar correlations as ENSO effects are small relative to natural spatial variation.

We also explored the temporal behaviour at different sites, finding stationarity to be a reasonable description within each summer, that ENSO accounts for a small proportion of temporal variation, and that the process decorrelates over 5 days in the south and 20 days in the north.

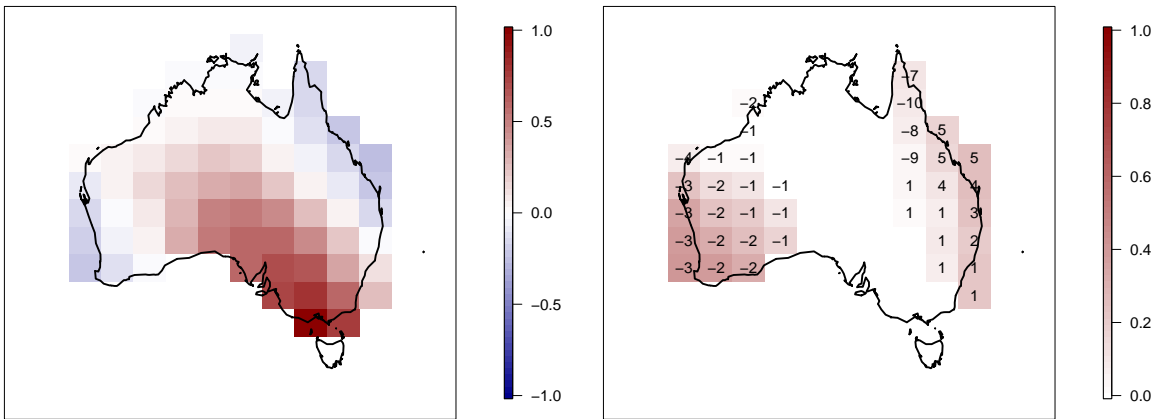


Figure 1: Estimated spatial cross-correlation function for Melbourne daily maximum temperatures and other sites: (left) the lag 0 and (right) the difference between the maximum value of the cross-correlation function and value of the lag 0 cross-correlation function. Data are for the years 1957-2011. Numbers in squares represent the lag value at which the maximum cross-correlation occurs; a blank square represents lag 0.

3 Modelling and inference for extreme values

3.1 Strategy

We are interested in modelling the spatial extent of heatwaves over Australia, in particular the south-east region given its population density and economic importance. Figure 1 shows that there is little difference between the spatial structure over south-eastern Australia when considering lags other than lag 0. Therefore, for simplicity, we will model concurrent temperatures, i.e., $\{Y_{s,t}, s \in S\}$. We could have studied the field $Y_{s,t-h(s)}$ where $h(s)$ are the lags shown in Figure 1 (right). We do not feel this extra complication is justified here as extreme values may have different temporal lags than typical values and given that the ENSO covariate changes monthly.

Based on the exploratory data analysis, during our modelling and inference we will ignore the effect of temporal dependence to focus on the impact of ENSO on spatial dependence. Thus we follow a similar strategy to Chavez-Demoulin and Davison (2005) and derive estimates by making a false assumption of temporal independence. If we derived the sampling distribution under this false assumption we would underestimate the variability of this distribution but our estimates will be unbiased (Self and Liang, 1987). So, we recognise the effect of temporal dependence in our inference through the use of a model-based block bootstrap approach to derive the sampling distributions of our estimates. Specifically, we take temporal blocks of 20 days.

In summary we will model the daily maximum temperatures process $\{Y_{s,t}, s \in S\}$ as independent over time with a covariate g_t which varies with time but not space that effects both marginal and dependence structure of $\{Y_{s,t}, s \in S\}$. Our strategy for modelling the probabilistic behaviour of extremes $\{Y_{s,t}, s \in S\}$ is two-fold. Firstly, we model the marginal structure using a threshold based approach at each site $s \in S$ separately. Once the marginal structure has been modelled, we transform the data from each site onto common margins and model the extremal dependence structure using the conditional extremes approach. The sampling distribution of estimates accounts for the temporal structure through the use of a block bootstrap.

3.2 Marginal modelling

Daily maximum temperatures at a site vary with the ENSO covariate g_t . As we are interested in the behaviour of extreme temperatures we need to be able to model the effect of a covariate on tail behaviour at a site. Davison and Smith (1990) and Northrop and Jonathan (2011) propose different modelling approaches to achieve this by focusing exclusively on the effect of the covariate on the tail. Here we adopt the pre-processing approach of Eastoe and Tawn (2009) where a pre-processing step removes covariate effects from the body of the distribution and then residual influences of the covariates on the tails are accounted for using the methods of Davison and Smith (1990). As noted in Section 2, the ENSO signal has a small effect on the series relative to natural spatio-temporal variations and so modelling of pre-processed extreme residuals corresponds to modelling the extremes of the original series. The pre-processing approach has close parallels with Northrop and Jonathan (2011) since the threshold for the extreme value modelling is derived to be covariate dependent. However, the pre-processing approach has major benefits in efficiently estimating covariate effects if the effect of covariates is somewhat similar in the body and tail of the distribution.

For the pre-processing we fit a location-scale model in the margins, i.e., for daily maximum tem-

perature $Y_{s,t}$ at location s and time t we have

$$Y_{s,t} = \psi_s(g_t) + \tau_s(g_t)W_{s,t}, \quad (2)$$

for $t = 1, \dots, n$ and $s \in S$, where $(\psi_s(g_t), \tau_s(g_t))$ are the location-scale parameters, g_t is a time-varying covariate and $W_{s,t}$ is the zero mean residual. In this paper all covariates are included linearly with an appropriate link function such that

$$\psi_s(g_t) = \psi_s^{(0)} + \psi_s^{(1)}g_t \quad \log \tau_s(g_t) = \tau_s^{(0)} + \tau_s^{(1)}g_t,$$

with parameters $\psi_s^{(0)}, \psi_s^{(1)}, \tau_s^{(0)}$ and $\tau_s^{(1)}$ each in \mathbb{R} . It is assumed that covariate effects in the body of the distribution are accounted for by the location-scale transform and as such the distribution of $W_{s,t}$ is independent of t in its body. In contrast this transformation may not completely capture all the covariate effect in the extremes of $W_{s,t}$, defined by $W_{s,t} > u_s$, where u_s is a high threshold. The distribution of $W_{s,t} \mid W_{s,t} > u_s$ is therefore modelled as a generalised Pareto distribution (GPD) with scale and shape parameters that depend the covariates s and g_t . Exploratory data analysis revealed that the shape parameter was independent of g_t and the scale parameter was log-linear in g_t . Although both parameters were dependent on s , there was no simple relationship on latitude and longitude and so s is treated as categorical. Hence we denote the shape parameter by ξ_s and the scale parameter by $\sigma_s(g_t)$, with $\log \sigma_s(g_t) = \sigma_s^{(0)} + \sigma_s^{(1)}g_t$, with parameters $\sigma_s^{(0)}, \sigma_s^{(1)}$ and ξ_s each in \mathbb{R} . As a result, our model for the distribution function of the residual variable is $W_{s,t}$ is such that

$$F_W(w; s, t) = \begin{cases} 1 - \tilde{F}_W(u_s; s) [1 + \xi_s (w - u_s) / \sigma_s(g_t)]_+^{-1/\xi_s} & \text{if } w > u_s \\ \tilde{F}_W(w; s) & \text{if } w \leq u_s, \end{cases} \quad (3)$$

where $\tilde{F}_W(\cdot; s)$ is the empirical cumulative distribution function of $\{W_{s,t}\}_{t=1}^n$ at site s .

3.3 Dependence modelling

The conditional extremes method of Heffernan and Tawn (2004) is used here to model extremal dependence. Using the methods outlined in Section 3.2, data are transformed onto common margins. The transformation onto common margins simplifies the estimation of extremal dependence quantities. This is especially important in the spatial problems encountered here since we are interested whether different sites have rare values simultaneously irrespective of the value of these rare values on the original temperature scale. Modelling using the conditional extremes approach is simplified if the choice of common margin is assumed to be the Laplace distribution, as the margins have exponential upper and lower tails which ensures models for positive and negative dependence are symmetric (Keef et al., 2013). As a consequence we make the transformation

$$X_{s,t} = \begin{cases} \log \{2F_W(W_{s,t}; s, t)\} & \text{if } F_W(W_{s,t}; s, t) < 1/2 \\ -\log \{2[1 - F_W(W_{s,t}; s, t)]\} & \text{if } F_W(W_{s,t}; s, t) \geq 1/2, \end{cases}$$

where $F_W(\cdot; s, t)$ is given by equation (3), and $X_{s,t}$ is now identically distributed over s and t .

Let $\mathbf{X}_t = (X_{1,t}, \dots, X_{l,t})$, where l is the number of sites in the region S , so the marginal distributions of \mathbf{X}_t are all Laplace. Furthermore, define $\mathbf{X}_{-s,t}$ as all the components of the vector \mathbf{X}_t without $X_{s,t}$. In what follows all vector calculations are to be interpreted as componentwise.

The aim is to model the distribution of $\mathbf{X}_{-s,t}$ given that $X_{s,t}$ exceeds some high threshold u . It is necessary that the conditional distribution $P\{\mathbf{X}_{-s,t} \leq \mathbf{x}_{-s,t} \mid X_{s,t} = x_{s,t}\}$ is non-degenerate as $x_{s,t} \rightarrow \infty$ and hence normalising sequences are required to ensure $\mathbf{x}_{-s,t}$ changes appropriately with $x_{s,t}$. Heffernan and Tawn (2004), Heffernan and Resnick (2007) and Keef et al. (2013) show that under broad conditions there exists vectors $\boldsymbol{\alpha}_{-s,t} = (\alpha_{1|s,t}, \dots, \alpha_{s-1|s,t}, \alpha_{s+1|s,t}, \dots, \alpha_{l|s,t}) \in [-1, 1]^{l-1}$ and $\boldsymbol{\beta}_{-s,t} = (\beta_{1|s,t}, \dots, \beta_{s-1|s,t}, \beta_{s+1|s,t}, \dots, \beta_{l|s,t}) \in (-\infty, 1)^{l-1}$ such that for $\mathbf{z} \in \mathbb{R}^{l-1}$ and $x > 0$

$$P\left(\frac{\mathbf{X}_{-s,t} - \boldsymbol{\alpha}_{-s,t}X_{s,t}}{X_{s,t}^{\boldsymbol{\beta}_{-s,t}}} \leq \mathbf{z}, X_{s,t} - u > x \mid X_{s,t} > u\right) \rightarrow G_{-s,t}(\mathbf{z}) \exp(-x), \quad (4)$$

as $u \rightarrow \infty$ where $G_{-s,t}$ is a time-varying $(l-1)$ -dimensional distribution function, non-degenerate in each margin, i.e., for $j \in S \setminus \{s\}$ the j th margin $G_{-s,t}^{(j)}$ of $G_{-s,t}$ is non-degenerate. Different values of the dependence parameters $\boldsymbol{\alpha}_{-s,t}$ and $\boldsymbol{\beta}_{-s,t}$ arise for different types of tail dependence. If $\alpha_{j|s,t} = \beta_{j|s,t} = 0$ and $G_{-s,t}^{(j)}$ is the Laplace distribution function, for $j \in S \setminus \{s\}$, the variables $(X_{s,t}, X_{j,t})$ are independent. On the other hand for $(X_{s,t}, X_{j,t})$, $\alpha_{j|s,t} = 1$ and $\beta_{j|s,t} = 0$, for $j \in S \setminus \{s\}$, corresponds to the situation of asymptotic dependence, $-1 \leq \alpha_{j|s,t} \leq 0$ is negative extremal dependence and $0 < \alpha_{j|s,t} < 1$ or $\alpha_{j|s,t} = 0$ and $\beta_{j|s,t} > 0$ corresponds to asymptotic independence with positive extremal dependence. Here, a time-varying covariate g_t is introduced into the dependence parameters such that

$$\tanh^{-1}[\boldsymbol{\alpha}_{-s,t}] = \boldsymbol{\alpha}_{-s}^{(0)} + \boldsymbol{\alpha}_{-s}^{(1)}g_t \quad \tanh^{-1}[\boldsymbol{\beta}_{-s,t}] = \boldsymbol{\beta}_{-s}^{(0)} + \boldsymbol{\beta}_{-s}^{(1)}g_t, \quad (5)$$

with parameters $\boldsymbol{\alpha}_{-s}^{(0)}$, $\boldsymbol{\alpha}_{-s}^{(1)}$, $\boldsymbol{\beta}_{-s}^{(0)}$ and $\boldsymbol{\beta}_{-s}^{(1)}$ are each in \mathbb{R}^{l-1} . The inverse tanh link function is used to ensure the parameters $\boldsymbol{\alpha}_{-s,t}$ and $\boldsymbol{\beta}_{-s,t}$ are restricted to the range $[-1, 1]^{l-1}$. The restriction on $\boldsymbol{\beta}_{-s,t}$ is satisfactory since in practice it is very unlikely that $\beta_{j|s,t} < -1$, for $j \in S \setminus \{s\}$, as this corresponds to $\mathbf{X}_{-s,t} - \boldsymbol{\alpha}_{-s,t}X_{s,t}$ rapidly tending to zero as $u \rightarrow \infty$ i.e., $\mathbf{X}_{-s,t}$ is essentially deterministic given large $X_{s,t}$.

Modelling using the conditional extremes approach requires the assumption that the limiting form of equation (4) holds exactly for all values of $X_{s,t} > u$ given that u is a sufficiently high threshold, from now on called the modelling threshold. From equation (4) we have our model for $X_{s,t} > u$ that

$$\mathbf{X}_{-s,t} = \boldsymbol{\alpha}_{-s,t}X_{s,t} + X_{s,t}^{\boldsymbol{\beta}_{-s,t}}\mathbf{Z}_{-s,t},$$

where $\mathbf{Z}_{-s,t} = (Z_{1|s,t}, \dots, Z_{s-1|s,t}, Z_{s+1|s,t}, \dots, Z_{l|s,t})$ is a random variable with distribution function $G_{-s,t}$ that is independent of $X_{s,t}$.

The multivariate distribution $G_{-s,t}$ does not take any simple parametric form, which motivates the inclusion of a false working assumption of Gaussianity as in Keef et al. (2013) solely for the estimation of $\alpha_{j|s,t}$ and $\beta_{j|s,t}$ with $j \neq s$. That is $Z_{j|s,t} \sim N(\mu_{j|s,t}, \theta_{j|s,t}^2)$ and as such for each $j \in S \setminus \{s\}$

$$X_{j,t} \mid \{X_{s,t} = x\} \sim N\left(\alpha_{j|s,t}x + \mu_{j|s,t}x^{\beta_{j|s,t}}, \theta_{j|s,t}^2x^{2\beta_{j|s,t}}\right) \quad \text{for } x > u.$$

The working assumption permits the estimation of the set of parameters $(\alpha_{j|s,t}, \beta_{j|s,t}, \mu_{j|s,t}, \theta_{j|s,t})$ by standard likelihood approaches. Each element of $\boldsymbol{\alpha}_{-s,t}$ and $\boldsymbol{\beta}_{-s,t}$ is estimated pairwise for a

particular $s \in S$. Covariates are included in the nuisance parameters such that

$$\boldsymbol{\mu}_{-s,t} = \boldsymbol{\mu}_{-s}^{(0)} + \boldsymbol{\mu}_{-s}^{(1)} g_t \quad \log \boldsymbol{\theta}_{-s,t} = \boldsymbol{\theta}_{-s}^{(0)} + \boldsymbol{\theta}_{-s}^{(1)} g_t, \quad (6)$$

where $\boldsymbol{\mu}_{-s,t} = (\mu_{1|s,t}, \dots, \mu_{s-1|s,t}, \mu_{s+1|s,t}, \dots, \mu_{l|s,t})$ and $\boldsymbol{\theta}_{-s,t} = (\theta_{1|s,t}, \dots, \theta_{s-1|s,t}, \theta_{s+1|s,t}, \dots, \theta_{l|s,t})$ and the parameters $\boldsymbol{\mu}_{-s}^{(0)}$, $\boldsymbol{\mu}_{-s}^{(1)}$, $\boldsymbol{\theta}_{-s}^{(0)}$ and $\boldsymbol{\theta}_{-s}^{(1)}$ are each in \mathbb{R}^{l-1} .

At this stage the Gaussian assumption is discarded and a non-parametric estimate of the distribution for $\mathbf{Z}_{-s,t}$ is formed. We assume that the effect of the time varying covariate on $\mathbf{Z}_{-s,t}$ is through its mean and variance only and so the distribution of $(\mathbf{Z}_{-s,t} - \boldsymbol{\mu}_{-s,t})/\boldsymbol{\theta}_{-s,t}$ is independent of t . Defining a new multivariate distribution G_{-s} by

$$G_{-s}(\mathbf{z}) = G_{-s,t} \left(\frac{\mathbf{z} - \boldsymbol{\mu}_{-s,t}}{\boldsymbol{\theta}_{-s,t}} \right),$$

gives that G_{-s} is independent of t . We estimate the distribution G_{-s} non-parametrically using replicates of $(\mathbf{Z}_{-s,t} - \boldsymbol{\mu}_{-s,t})/\boldsymbol{\theta}_{-s,t}$ over t . Specifically, where n_u is the number data points exceeding the threshold u , let t_1, \dots, t_{n_u} be the indices of $t = 1, \dots, n$ where $x_{s,t} > u$ then let

$$\hat{\mathbf{z}}_{-s,i} = \frac{\mathbf{x}_{-s,t_i} - \hat{\boldsymbol{\alpha}}_{-s,t_i} x_{s,t_i} - \hat{\boldsymbol{\mu}}_{-s,t_i} (x_{s,t_i})^{\hat{\boldsymbol{\beta}}_{-s,t_i}}}{\hat{\boldsymbol{\theta}}_{-s,t_i} (x_{s,t_i})^{\hat{\boldsymbol{\beta}}_{-s,t_i}}}, \quad (7)$$

for $i = 1, \dots, n_u$. In this way the empirical distribution of sample $\hat{\mathbf{z}}_{-s,i}$ provides a non-parametric estimate, \tilde{G}_{-s} , to the distribution function G_{-s} for conditioning site s .

4 Measures for summarising spatial dependence

To analyse the spatial behaviour of hot events, we require measures that can adequately capture spatial characteristics. As noted in Section 1, the limiting measure χ , defined in expression (1), has a number of limitations for spatial risk assessment. Our solution is to propose a number of measures to address these weaknesses. By using a selection of different measures, we aim to characterise extremal dependence well and identify any changes in spatial structure that may occur due to a change in ENSO. These measures are valuable for model checking and enable comparisons between empirical and modelled values. For notational simplicity we drop the time index on the variables in this section. Here we define the marginal distribution function of Y_s , incorporating all steps of the pre-processing outlined in Section 3.2, by F_s .

First we present a pairwise sub-asymptotic extension of χ , proposed by Coles et al. (1999), namely

$$\chi_{s_2|s_1}(p) = \text{P} (F_{s_2}(Y_{s_2}) > p \mid F_{s_1}(Y_{s_1}) > p), \quad (8)$$

where $s_1, s_2 \in S$ and p is some high level. Often p is taken to be the non-exceedance probability associated with a critical return level. The benefit of this measure over χ is that is able to discriminate between different levels of extremal dependence irrespective of whether the variables are asymptotically dependent or asymptotically independent, particularly when studied over a range of large p . In the spatial context, $\chi_{s_2|s_1}(p)$ is most usefully applied by fixing s_1 and p as the site and level of interest respectively and estimating $\chi_{s_2|s_1}(p)$ for all other sites $s_2 \in S$. If $s_2 = s_1$ then $\chi_{s_2|s_1}(p) = 1$, but for $s_2 \neq s_1$ values of $\chi_{s_2|s_1}(p)$ typically decrease as sites are further apart. By

assessing this measure for a range of large p , we can discriminate between pairs of asymptotically dependent and asymptotically independent pairs of sites as for the former (latter) $\hat{\chi}_{s_2|s_1}(p)$ is constant (decreasing), other than for sampling variability, as p is increased.

A restriction of $\chi_{s_2|s_1}(p)$ is that it provides only pairwise dependence information, so tells us nothing about the occurrence of concurrent extremes at more than two sites at a time. Many extensions are possible, but we propose a useful and practically informative measure that evaluates the expected number of sites in the set of interest, R say with $R \subseteq S$, that exceed a critical level given that Y_s exceeds the same critical level. Specifically,

$$\phi_{R|s}(p) = \mathbb{E}(N_R(p) \mid F_s(Y_s) > p), \quad (9)$$

where $N_R(p) = \#\{j \in R : F_j(Y_j) > p\}$ gives the number of variables that concurrently exceed the probability level p in R . Therefore larger values of this measure suggest that there is a greater spatial risk from the event. Again this measure is studied over all $s \in S$.

Measure (9) requires a particular conditioning site to be defined prior to estimation. In general assuming that a hot event must strike a particular site is restrictive. We propose a measure corresponding to the probability of an exceedance of a critical level in a region R' given that there is an exceedance somewhere within a region R , i.e.,

$$\omega_{R'|R}(p) = \mathbb{P}(N_{R'}(p) \geq 1 \mid N_R(p) \geq 1),$$

for some regions $R, R' \subseteq S$. Mostly we are interested in sets of the form $R' \subset R$; but other sets, such as $R' \cap R = \emptyset$, can be considered. A special case of this measure occurs where $R' = \{s\}$ which gives the probability of an exceedance at site s given that there is an exceedance in region R .

A weakness of the $\chi_{s_2|s_1}(p)$, $\phi_{R|s}(p)$ and $\omega_{R'|R}(p)$ measures of spatial risk is the requirement to select a critical level p . Our final risk measure overcomes this weakness as well as the limitations of the other measures. This measure is an adaption of the severity-area-frequency (SAF) curves used in hydrology (Henriques and Santos, 1999). The SAF curve $(\gamma_j, j \geq 1)$ gives the average marginal return period of an event at the j worst affected sites, where the sites affected need not be contiguous. Specifically

$$\gamma_j = \frac{1}{j} \sum_{i=1}^j [1 - F_{(i)}(Y_{(i)})]^{-1}, j = 1, \dots, l, \quad (10)$$

where $F_{(1)}(Y_{(1)}) \geq \dots \geq F_{(l)}(Y_{(l)})$ are the ordered values of $(F_1(Y_1), \dots, F_l(Y_l))$ for the event.

The SAF measure permits spatial information to be compressed into a single curve that is easily interpretable by statisticians and climate scientists. The SAF curve is a monotone non-increasing function; the larger the value of γ_j the more severe the event is at that scale. In cases where the SAF curve for one event exceeds the SAF curve for another event we can talk about an ordering between the size of these events. Otherwise we need to reference against the spatial scale j . SAF curves can have distinctly different behaviours for different types of spatial process. If the process is strongly asymptotically dependent at all sites, the SAF curve will decrease very slowly with increasing j . If it is asymptotically dependent only very locally and independent otherwise there

will be a sudden drop off in values for some j with $j \ll l$. For asymptotically independent processes the SAF curve will decay more rapidly for events with larger γ_1 . For asymptotically independent processes the SAF curve decays faster as the level of dependence becomes weaker.

We can use the SAF curve to assess the probability of an event being more severe than a previously observed event, such as the early 2009 event. Let γ_j^{obs} be the SAF curve of an observed event. Here, we evaluate

$$\rho_j(R) = \text{P}\left(\gamma_j > \gamma_j^{\text{obs}}\right), j = 1, \dots, |R|, \quad (11)$$

where the distribution of $(\gamma_j, j \geq 1)$ is derived using the methods of Section 5 for the fitted model.

5 Simulating spatial fields

To estimate the measures of spatial dependence introduced in Section 4 we need to be able to simulate spatial gridded fields from the model fitted in Section 3. Heffernan and Tawn (2004) and Keef et al. (2013) give simulation schemes for the conditional extremes approach conditional upon an exceedance at a specified site. These schemes are adequate to obtain estimates of $\chi_{s'|s}(p)$ and $\phi_{R|s}(p)$ and they form the basis of the simulation schemes outlined here. Estimation of measures that condition upon an exceedance within a region require a more involved algorithm for generating simulated spatial gridded fields. The use of SAF curves for model validation also requires conditions on the value taken by the maximum spatial event. To distinguish between observed and simulated fields we use X^* to denote the simulated variable X , and $X^{*|s}$ when the simulated variable is conditional on the field being extreme at site s .

Throughout this section we present simulations of the spatial field on the Laplace margin scale, i.e., for $\{X_{s,t}^*; s \in S\}$. If interest is in fields on the temperature scale, each simulated field can simply be transformed back to the original temperature scale as follows. For a simulated field $\{X_{s,t}^*; s \in S\}$ the simulated field on the temperature scale is $\{Y_{s,t}^*; s \in S\}$ where

$$Y_{s,t}^* = \psi_s(g_t) + \tau_s(g_t)F_W^{-1}\{F_X(X_{s,t}); s, t\},$$

for all $s \in S$ and any t , where $(\psi_s(g_t), \tau_s(g_t))$ and F_W are given in expressions (2) and (3) respectively, and F_X is the distribution function of a Laplace random variable. As we evaluate these fields only at fixed time/covariate values we drop reference to t in our notation, but in Section 6 we will simulate fields for a range of covariate values.

First we present the simplest algorithm for simulating spatial fields conditional on the field being extreme at a specified site s , i.e., simulating from $\mathbf{X}_{-s}|X_s > v_p$, where $v_p = -\log\{2(1-p)\}$ is the critical level on Laplace scale associated to the non-exceedance probability p .

Algorithm 1: Generates fields with exceedances at site s

1. Sample $\tilde{\mathbf{z}}_{-s}^{*|s}$ from \tilde{G}_{-s} , i.e., the empirical distribution of the sample in equation (7).
2. Obtain $\mathbf{z}_{-s}^{*|s} = \boldsymbol{\mu}_{-s} + \boldsymbol{\theta}_{-s}\tilde{\mathbf{z}}_{-s}^{*|s}$ where $\boldsymbol{\mu}_{-s}$ and $\boldsymbol{\theta}_{-s}$ are defined by equation (6).
3. Simulate an exceedance $X_s^{*|s} > v$ as the sum of v and a unit Exponential random variable.

4. Spatial field $\mathbf{X}_{-s}^{*|s} = \boldsymbol{\alpha}_{-s} X_s^{*|s} + (X_s^{*|s})^{\beta_{-s}} \mathbf{z}_{-s}^{*|s}$, where $\boldsymbol{\alpha}_{-s}$ and β_{-s} are from equation (5).

The final simulated spatial field generated using Algorithm 1 is

$$\mathbf{X}^{*|s} = (X_s^{*|s}, \mathbf{X}_{-s}^{*|s}) = (X_1^{*|s}, \dots, X_l^{*|s}),$$

where $X_s^{*|s} > v$. To estimate extremal measures $\chi_{s'|s}(p)$ and $\phi_{R|s}(p)$ for $s' \in S$ and $R \subseteq S$, Algorithm 1, with $v = v_p$, is repeated m times to obtain $\mathbf{X}_1^{*|s}, \dots, \mathbf{X}_m^{*|s}$ and then

$$\hat{\chi}_{s'|s}(p) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}(X_{s',i}^{*|s} > v_p) \quad \text{and} \quad \hat{\phi}_{R|s}(p) = \frac{1}{m} \sum_{i=1}^m \sum_{j \in R} \mathbb{I}(X_{j,i}^{*|s} > v_p), \quad (12)$$

where $\mathbb{I}(\cdot)$ is the indicator function, and the second subscript denotes the i th replicated field. For later use, note that the probability that site s is the maximum of the field over R , given that $X_s > v_p$, is denoted by

$$q_s^{*|s}(p; R) = \mathbb{P}\left(X_s = \max_{k \in R}(X_k) \mid X_s > v_p\right), \quad (13)$$

and can be estimated as

$$\hat{q}_s^{*|s}(p; R) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}\left(X_{s,i}^{*|s} > \max_{k \in R \setminus \{s\}}(X_{k,i}^{*|s})\right). \quad (14)$$

To estimate $\omega_{R'|R}(p)$ and the SAF curve, extensions of Algorithm 1 are required as we are interested in events that are hot for at least one site over a region R , i.e., $N_R(p) \geq 1$ for $R \subseteq S$ as opposed to an extreme temperature at site s . Our strategy for this simulation is as follows. We select a site that exceeds v_p by picking it to be the site with the maximum value of the spatial field over R . The probability that site j is largest over R , given that $N_R(p) > 1$, varies with j due to the changing dependence structure over space; we denote this probability by q_j , with

$$q_j(p; R) = \mathbb{P}\left(X_j = \max_{k \in R}(X_k) \mid N_R(p) \geq 1\right), \quad j \in R. \quad (15)$$

We then use Algorithm 1 with s set at the site selected to be the maximum over R . As Algorithm 1 can generate fields with the maximum over R larger than at site s we reject these fields. This is captured by Algorithm 2.

Algorithm 2: Generates fields with at least one exceedance in R

1. Sample J with probability $\mathbb{P}(J = j) = q_j(p; R)$ with $j \in R$.
2. Set $s = J$ and apply Algorithm 1.
3. If $\max_{k \in R}(X_k^{*|s}) > X_s^{*|s}$ reject this spatial field and repeat Algorithm 1 for the selected s until the simulated field is not rejected. The rejection probability is $1 - q_s^{*|s}(p; R)$ given by expression (13).

Algorithm 2, with $v = v_p$, is repeated m times, giving values j_1, \dots, j_m for J and the resulting fields $\mathbf{X}_1^{*|j_1}, \dots, \mathbf{X}_m^{*|j_m}$ and then

$$\hat{\omega}_{R'|R}(p) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}\left(\max_{k \in R'} X_{k,i}^{*|j_i} > v_p\right).$$

Similarly we can derive SAF curves using the fields generated in Algorithm 2 with $R = S$, giving for large p

$$\hat{\gamma}_j = \frac{2}{m} \sum_{k=1}^m \sum_{i=1}^j \exp\left(X_{(i)}^{*|j_k}\right),$$

where $X_{(1)}^{*|j_k} > \dots > X_{(l)}^{*|j_k}$ are the ordered values of the simulated field on Laplace marginals.

To use SAF curves for validation, we need to simulate replicate events that have similar characteristics to a particular event, e.g., the early 2009 heatwave. This necessitates fixing the maximum at the observed peak and corresponding site s , (η, s) say, with η being the maximum value of the field after transformation to the Laplace marginals. This restriction is achieved using Algorithm 2 with steps 1 and 2 removed and Algorithm 1 step 3 changed to $X_s^{*|s} = \eta$. See Winter (2016) for details of simulating these fields under the additional constraints that the maximum of the field is either equal to, or greater than, η when the site that maximum occurs at is not specified.

It remains to provide an estimate for $q_j(p; R)$ defined in expression (15). If v_p is small enough an empirical estimate of this probability may be sufficient but for large v_p we need to use our fitted conditional model. First note that

$$\mathbb{P}(N_R(p) \geq 1) = \sum_{k \in R} \mathbb{P}\left(X_k = \max_{i \in R}(X_i), X_k > v_p\right) = \sum_{k \in R} pq_k^{|k}(p; R).$$

It follows that

$$\begin{aligned} q_j(p; R) &= \mathbb{P}\left(X_j = \max_{i \in R}(X_i) \mid N_R(p) \geq 1\right) \\ &= \frac{\mathbb{P}(X_j = \max_{i \in R}(X_i), N_R(p) \geq 1)}{\mathbb{P}(N_R(p) \geq 1)} \\ &= \frac{\mathbb{P}(X_{j,t} = \max_{i \in R}(X_i), X_j > v_p)}{\mathbb{P}(X_j > v_p)} \frac{\mathbb{P}(X_j > v_p)}{\mathbb{P}(N_R(p) \geq 1)} \\ &= \mathbb{P}\left(X_j = \max_{i \in R}(X_i) \mid X_j > v_p\right) \frac{p}{\sum_{k \in R} pq_k^{|k}(p; R)} \\ &= \frac{q_j^{|j}(p; R)}{\sum_{k \in R} q_k^{|k}(p; R)}. \end{aligned} \tag{16}$$

Thus we can estimate $q_j(p; R)$ using expression (16) and estimate (14).

6 Analysis of Australia temperature data

The extreme value framework built in Section 3 is now combined with the summary measures defined in Section 4 to evaluate the characteristics of hot days over Australia for the gridded observations introduced in Section 2. Firstly, pre-processing is applied to the original data to model the marginal structure and transform values onto identical marginals. The choice of the conditional extremes approach is validated by comparing against other extreme value approaches that do not account for asymptotic independence. Finally, the variability of the spatial extent of hot events under El Niño and La Niña conditions is estimated. This culminates in estimating whether the framework can replicate similar events to the early 2009 heatwave event over Australia and how this event varies with the phase of ENSO.

6.1 Marginal structure

The pre-processing approach of Section 3.2 is now used to estimate the effect of ENSO on marginal quantities such as return levels. The covariate used to summarise the effect of SST on temperatures is Niño3.4 as introduced in Section 2. Figure 2 gives plots of the pre-processing parameters, with hashed boxes indicating sites where the parameter does not exhibit a significant change with the ENSO covariate (at the 5% significance level) based on a likelihood ratio test. In the right hand plots of Figure 2 darker shaded boxes show an increase in parameter values with an increase in Niño3.4 from -1°C to $+1^{\circ}\text{C}$; i.e., the difference between estimates under El Niño and La Niña conditions. The top row gives estimates of the location parameters $(\psi_s^{(0)}, \psi_s^{(1)})$ with the estimates of $\psi_s^{(0)}$ showing that warmer extreme temperatures are observed in northern and central regions of Australia but with cooler temperatures in coastal areas. The estimate of $\psi_s^{(1)}$ shows that an increase in Niño3.4 causes an increase in the location parameter over the most of Australia, with the largest increases being in eastern and western regions. For the scale parameter $\tau_s(g_t)$, the largest changes seem to be over western regions where El Niño conditions reduce temperature variability.

For each parameter we investigate for how many grid boxes the covariate is significant using likelihood ratio tests for each site at the 5% significance level. A decision is then made as to whether the covariate effect is included in the final model. The right hand plot of Figure 2 shows that out of a total of 72 grid boxes, 64 show a significant change in the location parameter with the ENSO covariate. This clear signal is not fully repeated by the scale parameter $\tau_s(g_t)$ which shows a significant change in 29 grid boxes out of 72. Although the result for the scale parameter is less significant we keep both covariate effects for all grid squares as we desire to have the same covariate structure incorporated in each parameter for all grid boxes. As such, we use the most general form of pre-processing outlined in Section 3.2.

Estimates of the GPD scale and shape parameters are given in Figure 2 (bottom row) and in Figure 3 respectively. Standard diagnostics (Coles, 2001) suggest the 90% quantile at each site is an appropriate threshold choice. As outlined in Section 3.2, the aim of this step is to take the approximately stationary time-series and ensure that the extremes are identically distributed over time. The estimates for $\tau_s^{(1)}$ and $\sigma_s^{(1)}$ possibly offset one another in the south-east corner of Australia. To check this we fixed the value of $\tau_s(g_t) = \tau_s^{(0)}$ at all sites and re-estimated $\sigma_s^{(1)}$ and found that the significant changes in the south-east are still present and therefore the changes are real. The shape parameter of the GPD is found to be negative at all sites over Australia, indicating a finite upper bound to the distribution at each site.

The clearest picture of the effect of the covariate can be seen when examining return levels after transforming onto the original scale. Figure 4 gives the 1- and 50-year return levels on the original margins during an El Niño event (i.e., the value of Niño3.4 is $+1^{\circ}\text{C}$) along with the change relative to a La Niña event (i.e., the value of Niño3.4 is -1°C). It is observed that the central regions of Australia are hotter than coastal regions as expected. There is an increase of up to 1°C in the 1-year return level between an El Niño event and a La Niña event. From a spatial perspective, the largest increases in the temperature occur in western and mid-eastern regions. The change in the 50-year return level is broadly similar, however southern and some northern areas show a larger decrease in temperatures with an increase in Niño3.4 due to the covariate effect on the GPD scale parameter shown in Figure 2.

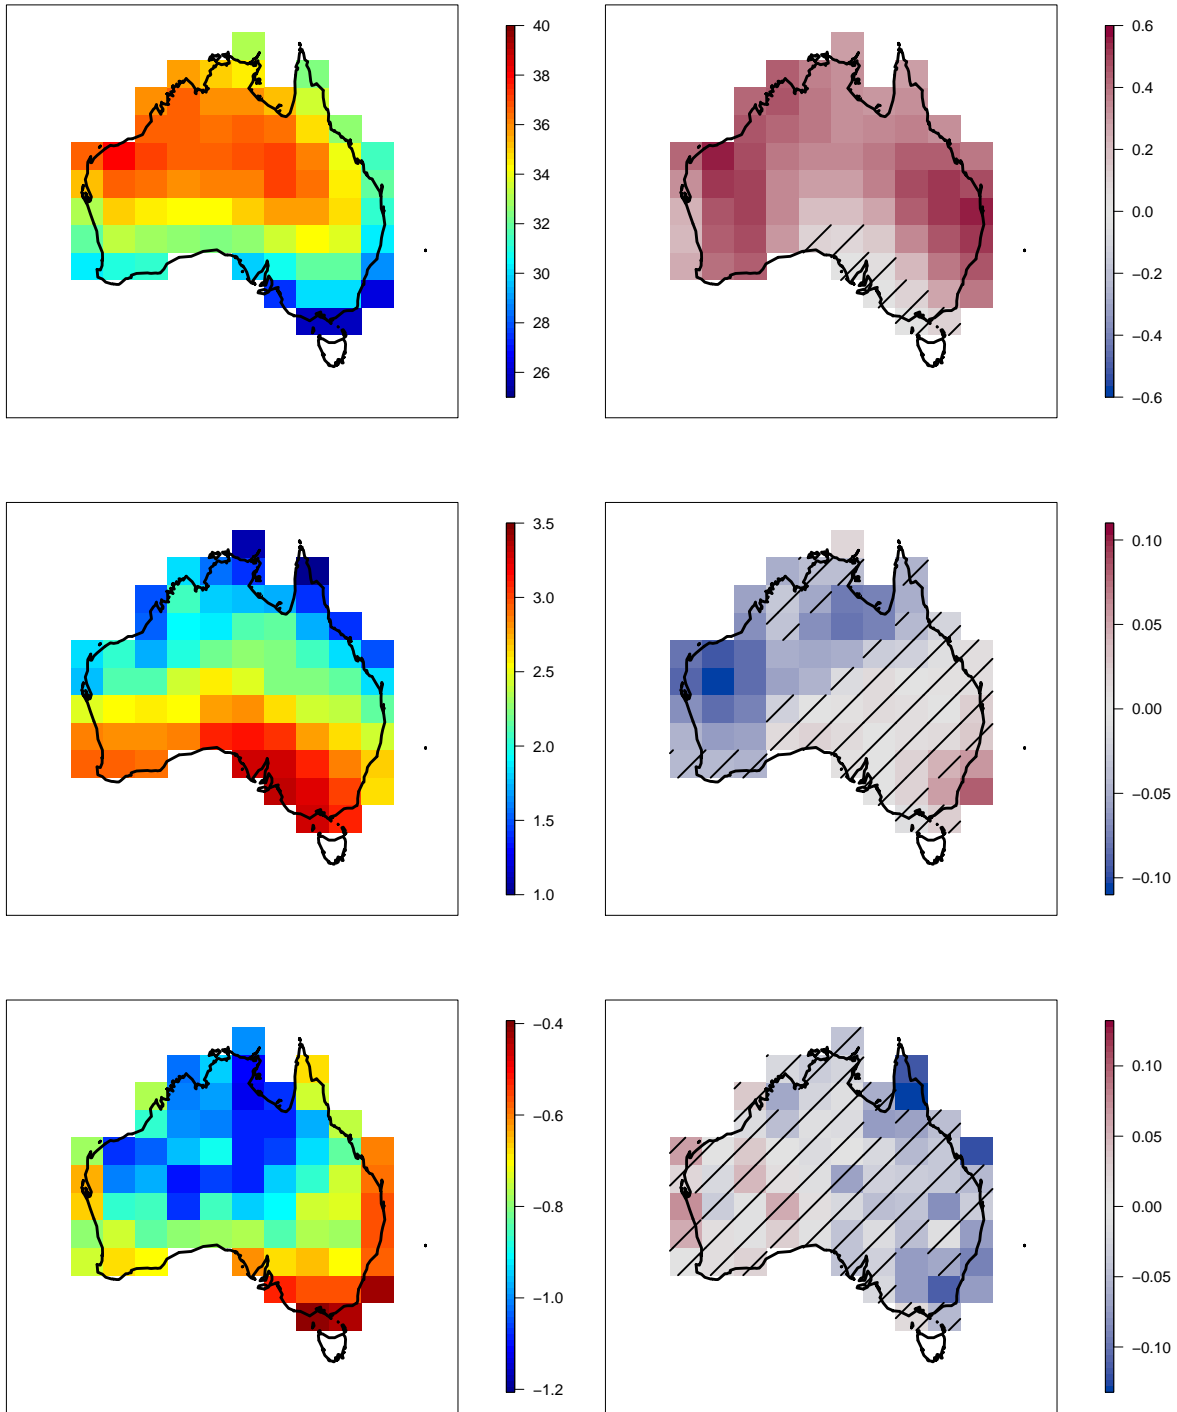


Figure 2: Estimates of pre-processing location and scale parameters $(\psi_s^{(0)}, \psi_s^{(1)})$ (top) and $(\tau_s^{(0)}, \tau_s^{(1)})$ (middle) and GPD scale parameters $(\sigma_s^{(0)}, \sigma_s^{(1)})$ (bottom). Hashed squares correspond to boxes where the change with covariate is not significant at a 5% significance level, tested using a likelihood ratio test.




Figure 3: Estimates of the GPD shape parameter ξ_s .

6.2 Pairwise spatial dependence

We now model the spatial pairwise dependence of the transformed data. Here we focus our presentation on pairs with Melbourne, s_M , being one of the sites. First we show that our conditional approach, that covers both asymptotic independence and asymptotic dependence, fits the data substantially better than the methods that can account for asymptotic dependence only. This initial analysis is with ENSO ignored, we then examine the effect of ENSO on our conditional model fit. This is appropriate as our exploratory analysis reveals that the covariate effects are small relative to spatial variation more generally.

Figure 5 shows estimates of the extremal dependence measure $\chi_{s|s_M}(v)$ with $v = v_{0.9}$ and the one-year return level $v = v_1^{RL}$. There are three estimates: empirical using the observed data and model-based using the stationary conditional extremes approach with both $\alpha_{s|s_M}$ and $\beta_{s|s_M}$ unconstrained and with $\alpha_{s|s_M} = 1$ and $\beta_{s|s_M} = 0$ for all $s \in S$ (i.e., asymptotic dependence). The empirical estimates of $\chi_{s|s_M}(v)$ show that for fixed v the decay of extremal dependence is not directly proportional to distance or invariant to direction, and that as v increases extremal dependence weakens. Critically this means that when an extreme event at a site is defined as the temperature exceeding the marginal T year return level, then the most extreme events at Melbourne become more localised as T increases. Both model-based estimates appear to be capturing the spatial dependence well for level $v_{0.9}$. However at higher levels the asymptotically dependent approach is substantially overestimating the amount of dependence across the field, especially at sites further from Melbourne. This is because the modelled dependence is independent of level once the levels are sufficiently extreme. In contrast, the general conditional extremes approach captures the observed weakening of dependence, with increasing levels, very well at all spatial separations.

Now we include the ENSO covariate in the analysis. Figure 6 gives estimates of the extremal dependence parameters $\alpha_{-s_M,t}$ and $\beta_{-s_M,t}$. We observe that the value of $\hat{\alpha}_{s|s_M}^{(0)}$ is broadly higher for sites s that are located closer to site s_M . The change in $\alpha_{-s_M,t}$ with the covariate is shown by the estimate $\hat{\alpha}_{s|s_M}^{(1)}$ which demonstrates an increase in extremal dependence, as g_t increases, over northern




Figure 4: 1-year (left) and 50-year (right) return levels plotted on original margins during El Niño conditions with SST temperature anomaly of $+1^\circ\text{C}$ (top) and change between return levels for El Niño and La Niña conditions under temperature anomaly of $+1^\circ\text{C}$ and -1°C respectively (bottom).

regions with a slight decrease in the east. The estimates of $\hat{\beta}_{s|s_M}^{(1)}$ seem to be consistently negative across northern region. These parameter estimates suggest that extreme temperature events that are occurring over Melbourne are more likely to extend over northern regions of Australia during El Niño conditions. It is easier to understand how the covariates effect the extremal dependence measures as the dependence parameters are not orthogonal. In Figure 7 a map of $\hat{\chi}_{s|s_M}(v_1^{RL})$ is given for an El Niño event along with a map of the difference in $\hat{\chi}_{s|s_M}(v_1^{RL})$ between El Niño and La Niña conditions. This inference suggests that if a hot day occurs at Melbourne, then if it was an El Niño year, the spatial extent of the event is likely to increase over southern coastal regions, including Adelaide, but it will not cover as much of the south-eastern region.

6.3 Spatial dependence measures

We now look at the characteristics of extreme spatial events by estimating the new quantities defined in Section 4. Firstly, we estimate the expected number of sites across Australia affected by extreme hot temperatures, given site s is extreme; i.e., $\phi_{R|s}(p)$ for $s \in R = S$ and p corresponding

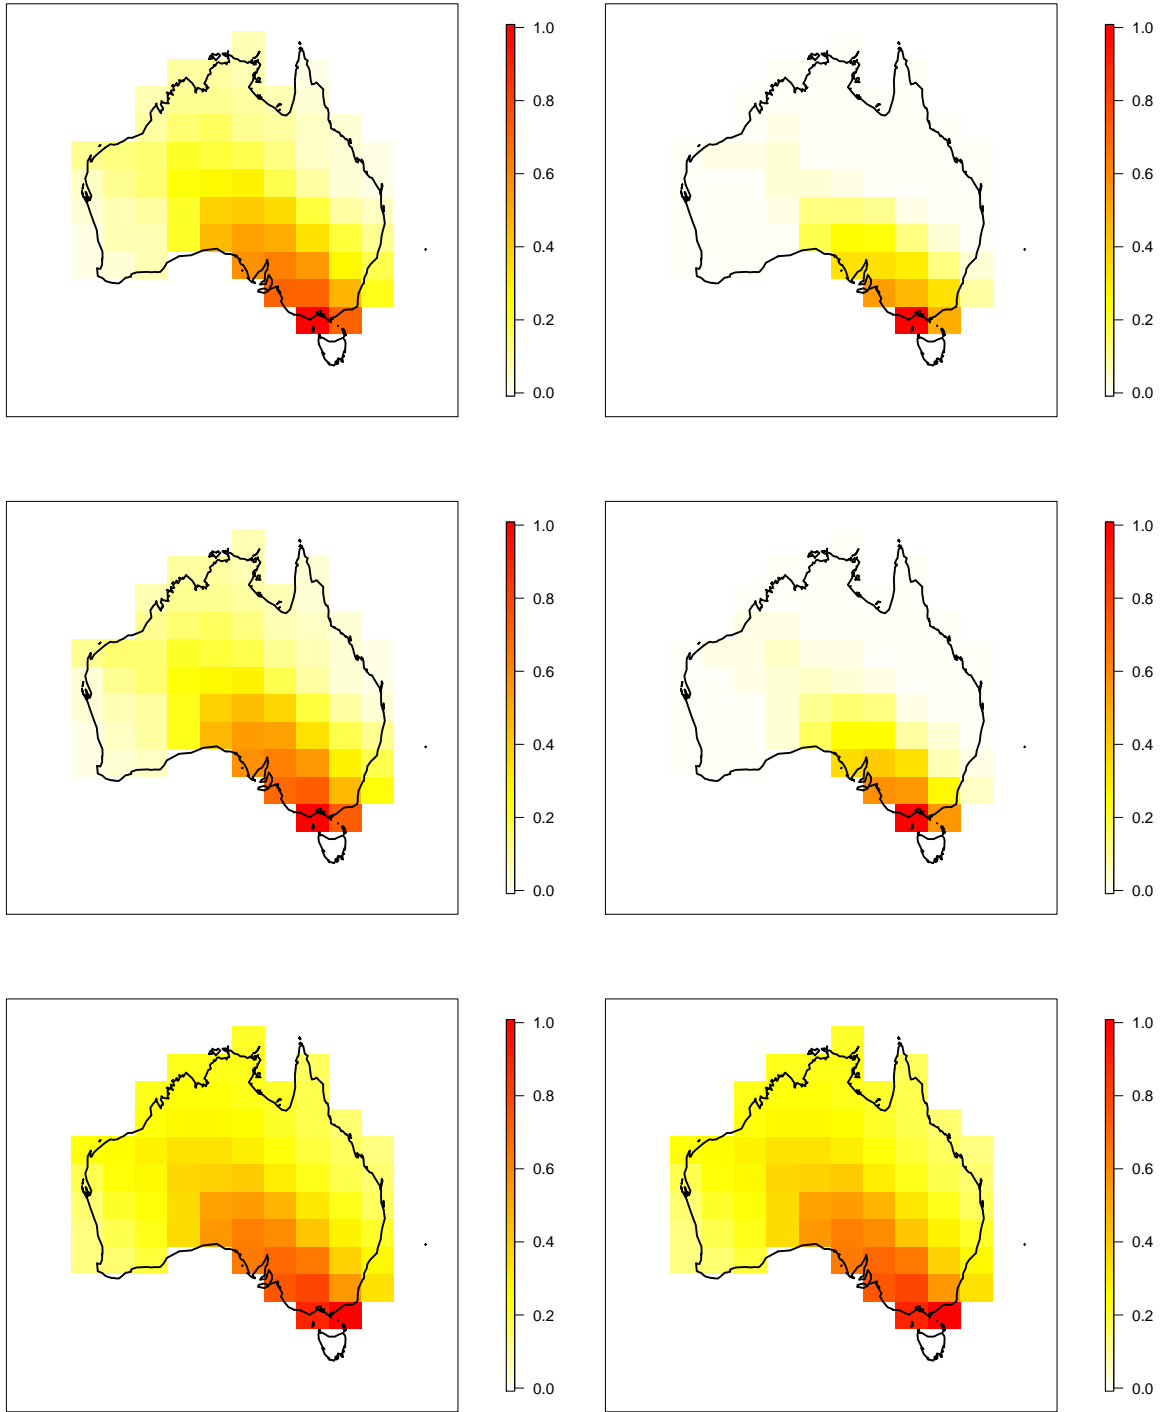


Figure 5: Values of $\chi_{s|s_M}(v)$ with $v = v_{0.9}$ (left) and one $v = v_1^{RL}$ the one year return level (right) for empirical (top), conditional extremes (centre) and asymptotic dependence (bottom) approaches. Here, the conditioning site s_M is the Melbourne grid-box. No ENSO covariate effects are included.

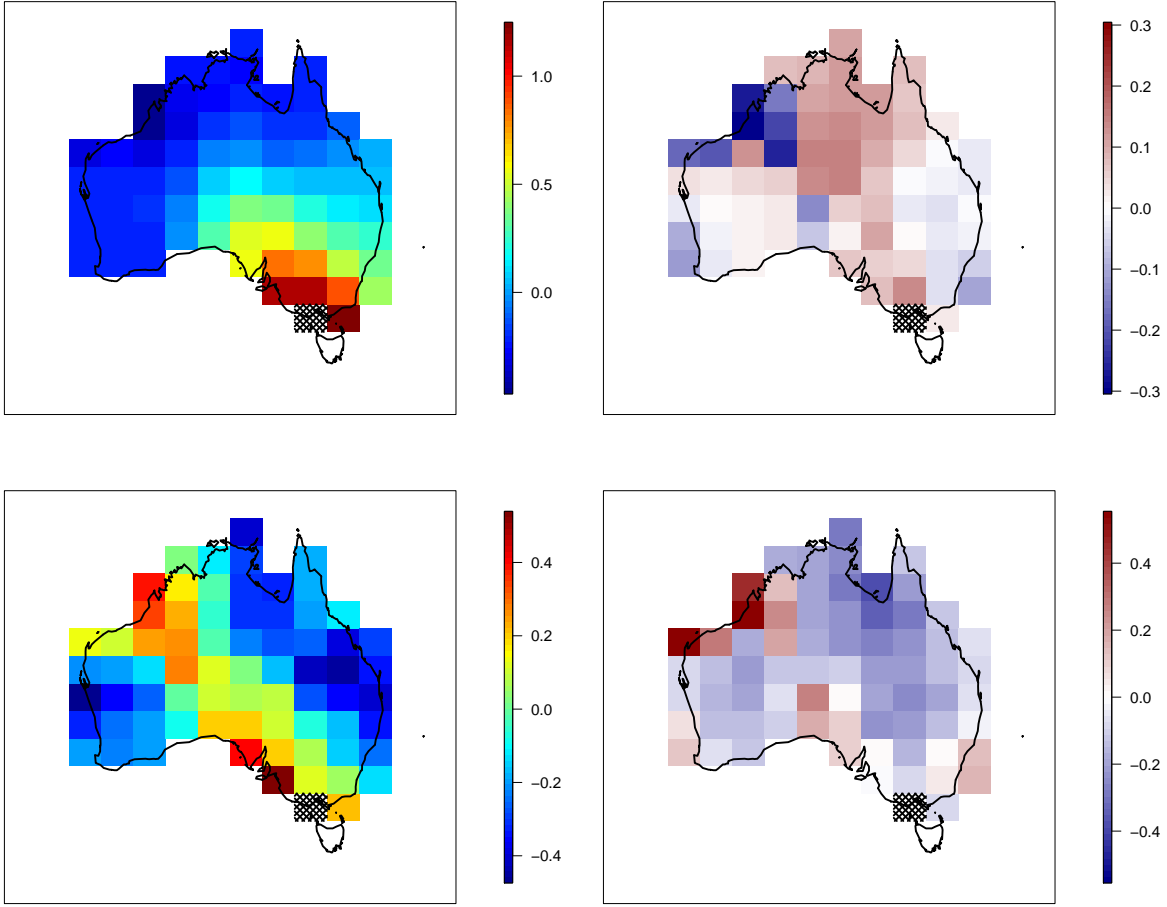


Figure 6: Conditional extremes dependence parameters $\hat{\alpha}_{s|s_M}^{(0)}$ (top left), $\hat{\alpha}_{s|s_M}^{(1)}$ (top right), $\hat{\beta}_{s|s_M}^{(0)}$ (bottom left) and $\hat{\beta}_{s|s_M}^{(1)}$ (bottom right), conditioning on site s_M (black hashed).

to the 1 year level, defined in equation (9). Estimates of $\phi_{R|s}(p)$ are given in Figure 8 under El Niño conditions and showing the change in estimates between El Niño and La Niña. It is observed that events occurring in the middle and east of Australia seem to have a greater spatial extent than for the west side during an El Niño event. The change in $\phi_{R|s}(p)$ between an El Niño event and a La Niña event suggests that El Niño conditions lead to a reduction in the spatial extent of hot days across most of Australia. Figure 8 suggests that during La Niña conditions the difference in the spatial extent of hot days between the east and west will become more pronounced. We also observe that results obtained conditioning upon Melbourne are typical of coastal grid boxes in the south-eastern region.

We are also interested in the probability of a hot event occurring over Melbourne given that an extreme temperature is observed somewhere in south-east Australia, i.e., $\omega_{\{s_M\}|R}(p)$ defined in Section 4 with R the set of 14 sites in south east Australia and p corresponding to the 1 year level. We have that $\hat{\omega}_{s_M|R}(p) = 0.20$ (0.19) under El Niño (La Niña) conditions respectively. This feature of the extreme events appears insensitive to the ENSO signal.




Figure 7: Extremal dependence measure $\chi_{s|s_M}(v)$ for control site over Melbourne under El Niño conditions $g_t = +1$ (left) and difference between extremal dependence measures during El Niño and La Niña ($g_t = -1$) years (right).

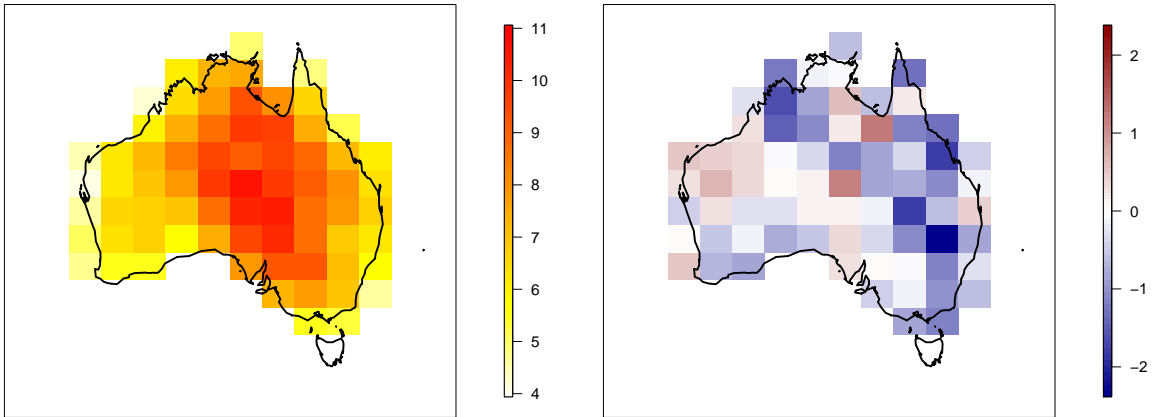


Figure 8: Estimates of $\phi_{R|s}(p)$ across Australia under El Niño conditions (left) and the change in estimates of $\phi_{R|s}(p)$ between an El Niño and La Niña year.

6.4 SAF results

Figure 9 shows estimated SAF curves for the early 2009 heatwave in the form of an estimate and 95% confidence intervals for γ_j . We fix the maximum value of our simulated fields to agree with the observed maximum value. We compare with the observed SAF curve with estimates obtained under both the conditional extremes model and its restriction to asymptotic dependence, with and without accounting for ENSO. Although we had earlier identified asymptotic dependence as a poor model formulation, we retain that model here to show how we could make incorrect inferences for SAF curves under the assumption of asymptotic dependence.

As expected, the model that allows asymptotic independence provides a better fit to the observed

SAF curve than the asymptotically dependent estimate. Asymptotic dependence leads to a grouping of large return periods that is too strong, with estimates that decay too slowly, irrespective of whether knowledge of ENSO is incorporated. Notice that the observed SAF curve corresponds to the lower endpoint of the 95% confidence interval of the SAF curve under asymptotic dependence as under this modelling assumption extrapolated spatial events simply scale observed events in terms of size but keep relative values identical to each other. So when the maximum of the event is the observed maximum then one realisation of the sampling distribution will be our observed event. The asymptotically dependent model produces much wider confidence intervals for SAF curves as it is a poor model fit and so the confidence intervals try to match both the false model and the structure of observed heatwaves.

For the conditional extremes approach ignoring ENSO leads to an overestimation for SAF at scales of up to 30 grid squares. However the observed SAF fits well inside 95% confidence intervals when the phase of the observed ENSO is accounted for. These results highlight the need to account for both asymptotic independence and ENSO in the spatial dependence structure.




Figure 9: Severity-area-frequency (SAF) curves, on a log scale, for the early 2009 heatwave: observed curve (black), $\hat{\gamma}_j$ under conditional extremes model (blue solid) and under asymptotically dependence (red solid). Left: no ENSO effect. Right: ENSO fixed at level observed on the day. Dotted lines are 95% confidence intervals obtained from 10000 replicates.

We estimate $\rho_j(R)$ from equation (11) with R a region of 14 sites in south-eastern Australia including Melbourne, where the field maximum can occur at any site in R . Figure 10 shows estimates of $\rho_j(R)$ under the observed La Niña ($g_t = -0.7$) and typical El Niño ($g_t = 1$) conditions. In the left plot, the maximum value is taken to be greater than v_1^{RL} . The left plot shows for low j ($1 \leq j \leq 5$), i.e., events considered locally, the observed event is rarer under El Niño, than La Niña, conditions. As j is increased there seems to be little difference between the ENSO conditions. It is noted that irrespective of the ENSO conditions, the observed event was very rare. The right plot shows the rarity of the observed event given that the maximum is fixed at the peak of the observed event. In this situation, at all values of j there is a difference between ENSO phases, with the observed event much less rare if it was to occur under El Niño conditions than for La Niña conditions.




Figure 10: Estimates of $\rho_j(R)$ conditional upon the maximum field value in R being: (left) greater than v_1^{RL} and (right) equal to the observe value in the early 2009 heatwave. Estimates are given for observed La Niña conditions (black) and typical El Niño conditions (grey). Here R is a set of 14 sites in south east Australia.

7 Discussion and Conclusion

In this paper we have modelled the spatial extent of extreme temperature events over Australia and motivated an approach for modelling gridded spatial data using the conditional extremes approach. Within this framework we have included the ability to account for covariates within the margins and the dependence structure which has allowed us to understand the effect of ENSO on extreme temperatures. Our approach has confirmed that El Niño conditions lead to higher temperatures across most of Australia and that the increase in temperature might not be uniform at all return levels, i.e., the effect of ENSO does not just cause a shift in the distribution of temperatures.

Results regarding the change in the spatial extent of heatwaves with ENSO value are more subtle than the changes in marginal structure and vary for different sites. We have shown that a hot event over Melbourne is likely to cover more of the south-eastern region during La Niña conditions. We have also estimated quantities that are not dependent on the process being extreme at a particular site, which have greater practical importance. These measures have highlighted drawbacks in current pairwise measures and as such need to be considered in future spatial analyses. In particular our proposed SAF curves succinctly present complex space-time information in a highly informative and interpretable way. We have also used the observations from the early 2009 heatwave event to estimate whether the event would have been more likely under El Niño or La Niña conditions. The quantities presented here are just a subset of potential measures that could be estimated. We have outlined a general approach for simulating spatial extreme temperature events that could be used to generate any quantity of interest for decision makers.

The impact of climate change on the spatial distribution of extreme temperature events has not been dealt with in this paper. This is clearly an important issue that could be included into our framework as another covariate, see Winter et al. (2016) for an illustration of such an approach for a single site. One problem concerns the uncertainty regarding the effect of climate change on

ENSO which is currently not well known and would preclude a comprehensive study of the joint effects of ENSO and climate change on extreme temperatures.

Finally, it is also noted that from a mortality perspective we may be interested in different measures. For example, fires can be caused from the combination of hot temperatures, low rainfall, high winds and low humidity. In many situations, runs of hot temperatures are more important than particular hot days. Winter and Tawn (2016) showed how temporal heatwave events can be simulated using the conditional extremes framework for a single site. The next step will be to combine these temporal approaches with the spatial approaches outlined in this paper to generate full space-time model on a lattice which incorporate asymptotic independence as well as asymptotic dependence, hence expanding on max-stable spatio-temporal models of Davis et al. (2013) and Huser and Davison (2014).

Acknowledgements

We gratefully acknowledge the support of the EPSRC funded EP/H023151/1 STOR-i centre for doctoral training. We would like to acknowledge funding and data from the Met Office and the helpful input from referees and the editor. This work was completed while Hugo Winter was at Lancaster University.

References

- Alexander, L. V. and Arblaster, J. M. (2009). Assessing trends in observed and modelled climate extremes over Australia in relation to future projections. *International Journal of Climatology*, 29:417–435.
- Avila, F. B., Dong, S., Menang, K. P., Rajczak, J., Renom, M., Donat, M. G., and Alexander, L. V. (2015). Systematic investigation of gridding-related scaling effects on annual statistics of daily temperature and precipitation maxima: A case study for south-east Australia. *Weather and Climate Extremes*, 9:6–16.
- Caesar, J., Alexander, L., and Vose, R. (2006). Large-scale changes in observe daily maximum and minimum temperatures: Creation and analysis of a new gridded data set. *Journal of Geophysical Research: Atmospheres*, 111(5):1–10.
- Chavez-Demoulin, V. and Davison, A. C. (2005). Generalized additive modelling of sample extremes. *Applied Statistics*, 54:207–222.
- Coles, S. G. (1993). Regional modelling of extreme storms via max-stable processes. *Journal of the Royal Statistical Society: Series B*, 55(4):797–816.
- Coles, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer Verlag.
- Coles, S. G., Heffernan, J. E., and Tawn, J. A. (1999). Dependence measures for extreme value analyses. *Extremes*, 2(4):339–365.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data*. Wiley, New York.

- Davis, R. A., Kluppelberg, C., and Steinkohl, C. (2013). Statistical inference for max-stable processes in space and time. *Journal of the Royal Statistical Society: Series B*, 75(5):791–819.
- Davison, A. C., Padoan, S. A., and Ribatet, M. (2012). Statistical modeling of spatial extremes. *Statistical Science*, 27(2):161–186.
- Davison, A. C. and Smith, R. L. (1990). Models for exceedances over high thresholds (with discussion). *Journal of the Royal Statistical Society: Series B*, 52(3):393–442.
- Dombry, C., Éyi-Minko, F., and Ribatet, M. (2013). Conditional simulation of max-stable processes. *Biometrika*, 100(1):111–124.
- Eastoe, E. F. and Tawn, J. A. (2009). Modelling non-stationary extremes with application to surface level ozone. *Journal of the Royal Statistical Society: Series C*, 58(1):25–45.
- Heffernan, J. E. and Resnick, S. I. (2007). Limit laws for random vectors with an extreme component. *The Annals of Applied Probability*, 17(2):537–571.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B*, 66(3):497–546.
- Henriques, A. G. and Santos, M. J. J. (1999). Regional drought distribution model. *Physics and Chemistry of the Earth, Part B: Hydrology, Oceans and Atmosphere*, 24(1-2):19–22.
- Huser, R. and Davison, A. C. (2014). Space-time modelling of extreme events. *Journal of the Royal Statistical Society: Series B*, 76(2):439–461.
- Jones, D. A. and Trewin, B. C. (2000). On the relationships between the El Nino-Southern Oscillation and Australian land surface temperature. *International Journal of Climatology*, 20:697–719.
- Keef, C., Papastathopoulos, I., and Tawn, J. A. (2013). Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the Heffernan and Tawn model. *Journal of Multivariate Analysis*, 115:396–404.
- Kenyon, J. and Hegerl, G. C. (2008). Influence of modes of climate variability on global temperature extremes. *Journal of Climate*, 21(15):3872–3889.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187.
- Min, S., Cai, W., and Whetton, P. (2013). Influence of climate variability on seasonal extremes over Australia. *Journal of Geophysical Research: Atmospheres*, 118(2):643–654.
- Nairn, J. and Fawcett, R. (2013). Defining heatwaves: heatwave defined as a heat-impact event servicing all community and business sectors in Australia. *Centre for Australian Weather and Climate Research, Technical Report*, 060:1–96.
- Northrop, P. J. and Jonathan, P. (2011). Threshold modelling of spatially dependent nonstationary extremes with application to hurricane-induced wave heights. *Environmetrics*, 22:799–809.
- Perkins, S. E. and Alexander, L. V. (2013). On the measurement of heat waves. *Journal of Climate*,

26(13):4500–4517.

Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes*, 5(1):33–44.

Self, S. G. and Liang, K. (1987). Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association*, 82:605–610.

Smith, R. L. (1990). Max-stable processes and spatial extremes. *Unpublished manuscript, University of Surrey*, pages 1–32.

Wang, C. and Picaut, J. (2004). Understanding ENSO physics - a review. *Geophysical Monograph Series*, 147:21–48.

Winter, H. C. (2016). Extreme Value Modelling of Heatwaves. *Lancaster University PhD thesis*.

Winter, H. C. and Tawn, J. A. (2016). Modelling heatwaves in central France: a case study in extremal dependence. *Journal of the Royal Statistical Society: Series C*, 65(3):345–365.

Winter, H. C., Tawn, J. A., and Brown, S. J. (2016). Detecting changing behaviour of heatwaves with climate change. *Submitted*.