University of Bath

PHD

Attribution of large-scale drivers for environmental change

Brady, Aoibheann

Award date:
2020

Awarding institution:
University of Bath

Link to publication

Alternative formats
If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

Attribution of large-scale drivers for environmental change

submitted by

Aoibheann Brady

for the degree of Doctor of Philosophy

of the

University of Bath

Department of Mathematical Sciences

September 2019

COPYRIGHT NOTICE

DECLARATION OF ANY PREVIOUS SUBMISSION OF THE WORK

The material presented here for examination for the award of a higher degree by research has not been incorporated into a submission for another degree.

Signature of Author . . . .

Aoibheann Brady

DECLARATION OF AUTHORSHIP

I am the author of this thesis, and the work described therein was carried out by myself personally in collaboration with my supervisors.

Chapter 2 is reproduced from a Hydrological Sciences Journal article (Brady et al., 2019).

Chapter 3 is reproduced from a prepared and revised manuscript.

Signature of Author . . . .

Aoibheann Brady

Abstract

In this thesis, spatial multilevel models are developed within a causal framework for the attribution of long-term changes in environmental studies to large-scale drivers of interest. In particular, these methods are applied to river flows in this thesis. Two common themes are apparent throughout the thesis. The first involves the accurate detection of long-term changes in river flows in Great Britain. The development of spatial multilevel methods for the accurate detection of trends is a major focus of the thesis. We firstly focus on the detection of countrywide trends in which the correlation between stations is modelled through a Gaussian process. Using a multilevel approach allows for the pooling of gauging station data to accurately estimate trends in peak river flows that may not be possible using traditional at-site analyses. This approach resulted in a first detection of trends in peak river flows at a countrywide level in Great Britain. We then switch in focus towards the analysis of a single river network, exploiting this network structure in order to understand how flows in a given region will evolve. This method considers the river gauging station measurements on the network as a graph, which is comprised of nodes (the river gauging stations themselves) and edges between nodes (the direction of flow between each station). The network structure of this river is encoded using a first-order conditional autoregressive (CAR) model. This method allows for the use of fast inference methods through the construction of a sparse precision matrix, and also respects the physical structure of the network.

The second theme focuses on the development of a causal framework for the attribution of long-term, large-scale changes in environmental studies to some climate drivers of interest. We first perform a preliminary attempt at attribution, where a clear association is seen between the East Atlantic (EA) index and peak river flows, even when a multivariate approach is used to account for temporal confounding. We then provide a more rigorous approach towards the attribution of long-term, large-scale drivers of environmental change. A systematic checklist is developed to provide a thorough causal assessment of drivers of environmental changes, demonstrating through this checklist that changes in peak river flows can be attributed to the EA index.

ACKNOWLEDGEMENTS

Contents

List of Figures

14

Glossary

annual maximum The annual maximum river flow data series contains the largest observed flow (in $\mathrm{m}^3\mathrm{s}^{-1}$) in a water year. Also referred to as "peak river flows". 9, 22, 23, 25–27, 37, 39, 42, 44–46, 48, 50–56, 58, 59, 62, 65, 87–89, 92, 95, 96, 108, 145, 149–157, 159, 160, 177

BFIHOST The baseflow index, derived using the hydrology of soil types (HOST) classification, provides a measure of catchment responsiveness. 11, 51, 62, 97, 146

daily mean flow The daily mean flow data series is the mean river flow in cubic metres per second in a water day. 14, 22, 25, 31, 33, 37, 108, 109, 113, 128, 136–140, 143, 144, 146, 155, 157, 158, 160

discharge The discharge is the total volume of water flowing through a channel at any given point, measured in cubic metres per second. 114, 139

peaks over threshold The peaks over threshold (POT) data consists of all peak flows which exceed a given threshold flow. This threshold is typically set to include an average of 5 events per year. 22, 27, 51

Acronyms

CAR Conditional Autoregessive model. 8, 14, 25, 29, 31, 122, 123, 127–131, 134–136, 138, 141, 142, 145, 146, 150, 155, 156, 158, 159

EA East Atlantic index. 11–13, 41, 53–56, 58–64, 87–99, 101, 149, 151, 154, 160

GAM Generalised Additive Model. 8, 92, 115, 121

GMRF Gaussian Markov Random Field. 31, 34, 35, 126, 130, 131, 145, 156

INLA Integrated Nested Laplace Approximation. 8, 27, 33–35, 129, 133, 138, 140, 141, 145

LGM Latent Gaussian Models. 34

MCMC Markov Chain Monte Carlo. 32–35, 49, 129, 181

NAO North Atlantic Oscillation index. 11, 12, 41, 52–56, 58–61, 63, 64, 67, 87–92, 95, 96, 99, 101, 102, 143, 151, 152, 154

NRFA National River Flow Archive. 22, 51, 88, 97, 113, 146, 159

SAAR Standard-period Average Annual Rainfall. 13, 97, 98, 113

URBEXT An index of urban and suburban land cover in a given period expressed as a fraction. 51

Chapter 1

Introduction

Climate-driven changes in river flow regimes worldwide can have considerable impacts on society. There are expectations that changes in climate will result in an increase in flood risk, fluctuations in water availability and a reduction in water quality. Understanding the patterns underlying these changes is vital for ensuring adequate flood defences and water resource management plans are in place for the future.

There are growing concerns that climate change may result in an increased risk of flooding, for example, with several large flood events hitting the UK in recent years. The resulting destruction caused by these floods led to the question of whether current flood defences were fit for purpose. The Environment Agency stated that a complete overhaul of the UK's flood defence infrastructure was needed, as the UK was moving from a period of "known extremes" of weather to one of "unknown extremes" (Climate Home News, 2015). Cologna et al. (2017) urged a change in coverage of major flooding incidents towards being prepared for an overall increase in flood risk due to climate change, instead of portraying them as once in a lifetime events. There are also fears regarding the impact of climate change on the availability of water resources in the UK in the future. While climate-driven increases in winter precipitation should result in an improvement in the reliability of water resources in northern England, it is also expected that there will be a reduction in summer precipitation leading to increased vulnerability to drought (Fowler and Wilby, 2010).

The current perception is that there has been an increase in the frequency and severity of flooding in recent years (Taylor et al., 2014), yet conflicting claims in the literature mean that no clear conclusions can be drawn. Climate change projections suggest an increase in rainfall in northern Europe, with a possible associated flood risk, over time

(Bates, 2009). However, this has not yet been verified by the observed river flow data, from which no compelling evidence of increasing trends has been inferred (Hannaford, 2015). Trend tests on river flow data, which are obtained from modelling each individual gauging station separately, often lead to unclear signals. These tests are not very powerful in a statistical sense, and may not be able to fully differentiate between confounders. The relatively short data records of annual maximum flows in the UK mean it is difficult to provide compelling statistical evidence of any long term trends.

In this thesis, we aim to reconcile the differences in results between climate model projections and observational river flow data. Through modelling of trends in annual maximum flow data, we will provide compelling evidence that flooding is likely to increase, and obtain an understanding of the factors, such as climate change, driving this increase. We will also demonstrate that this increase is paralleled with an increase in daily mean flows, with an expectation that this will impact future water resources management.

1.1 Data availability and limitations

The primary source of river flow data in the United Kingdom is the National River Flow Archive (National River Flow Archive (NRFA)) (Dixon, 2010; Dixon et al., 2013), which has data from a dense river monitoring network of approximately 1400 gauging stations. Through the NRFA, it is possible to access gauged daily mean flow data and annual maxima, both of which are derived from 15 minute recordings of the flow. Daily flow data are available for the majority of these gauging stations, and approximately 1000 stations also have annual maximum flow data (i.e. annual maximum flow or peaks over threshold data). Annual maximum flow data are used for flood frequency estimation and will be used in this thesis to investigate trend in extreme river flows. Interrogating the annual maximum flow data will help to determine whether changes in flooding are consistent with climate change projections. Daily data are useful for water resource management assessment, and will be used in this thesis to assess changes in overall long-term means in a single river network.

While this would suggest that the UK has a rich source of observational data for trend estimation, many records only started after the 1970s, and those stations with longer records are often compromised due to changes in gauging practice (Hannaford, 2015). Data quality is also a concern, particularly in the monitoring of extreme flows, meaning that it may be difficult to obtain accurate estimates of long-term trends in such

flows (Dixon et al., 2013). Moreover, there may be interference due to anthropogenic influences, which can result in the detection of trends unrelated to climate activity. This was overcome in part, however, by the development of a reference network of "benchmark catchments", a series of near-natural river gauging stations with good quality data and relatively unimpacted by anthropogenic changes (Bradford and Marsh, 2003). The latest version of this is the UK Benchmark Network V2 (UKBN2), which consists of 146 benchmark catchments as a representative reference network (Harrigan et al., 2017).

1.2   Statistical challenges

In an at-site analysis, as each individual hypothesis test is insufficiently powered to detect trends, we may not be able to correctly quantify changes in environmental risks such as flooding. One issue lies in the length of individual gauging station records. Prosdocimi et al. (2014) observed that a station-by-station approach to annual maximum river flow analysis may require a sample size of hundreds of years, while reliable records of river flow are typically much shorter than this. Renard et al. (2008) noted at-site studies which were limited in their statistical power to detect any changes, suggesting instead that a regional analysis could increase this power. In addition, as noted by Prosdocimi et al. (2014), current models largely assume stationarity of the river flow process, which in practice means assuming the probability of an extreme event is constant. This can result in a failure to accurately estimate flood risk, which can be extremely costly to the government and the public alike. Additionally, consistently underestimating any trends in annual maximum river flows may mean that the UK's flood infrastructure is not fit for purpose for future extreme events.

Current methods for detection of changes in river flow data do not take into account the spatial nature of rivers. These approaches often involve modelling each station separately, meaning the result at each station is based entirely on data available at the station itself. Monitoring stations are often geographically close to each other (sometimes gauging the same river) and exposed to similar weather and climate, meaning that these stations may display similar trends. Information from nearby stations, which may help to enhance the signal at the station in question, is not used in such approaches. However, once it is assumed that the trend is similar across all stations, it is possible instead to model all stations together in a mixed effect model. We do so using a Bayesian multilevel framework, by modelling the spatial relationships between gauging stations through a spatial random effect.

A further issue with many current river flow analyses is the sole focus on the detection of time trends, instead of a joint approach of trend detection and attribution of trends to a variable of interest. Merz et al. (2012) noted the need for a switch in focus towards a more holistic approach, which would incorporate change detection into the more challenging problem of attribution. Understanding how river flows are changing is beneficial when it comes to choosing effective interventions, and ensuring adequate water resources and flood defences are in place in the long term. However, the causal nature of relationships between variables in a long-term, large-scale environmental setting such as this are not well understood. The issues of separating anthropogenic change from natural variability, the inability to randomise exposures and the presence of many confounding variables mean that typical causal tools cannot always be directly applied to environmental problems. For example, while commonly-used tools such as propensity score matching allow for the construction of counterfactual outcomes in epidemiological observational studies, these are not always fit for the purpose of assessing for causality in environmental studies. In such situations, it can be difficult to have well-defined counterfactual outcomes. One could use geographical characteristics in spatial studies to construct such counterfactual outcomes, however in many environmental studies, such as river flow studies, such characteristics do not provide much, if any, explanatory power. Additionally, it can be difficult to consider all possible sources of confounding with additional unknown variables.

In recent years, there have been a number of studies aiming to attribute changes in the frequency of flooding to some driver of interest (see, for example, Hannaford and Marsh (2008); Lins and Slack (2005); Villarini et al. (2011)). Merz et al. (2012) provided a review of the state of the art of such attribution studies in flooding, dividing them into categories of "soft" and "hard" attribution. **"Soft" attribution** studies make use of hypotheses and references to past literature to corroborate any statements of attribution to some cause of change. On the other hand, **"hard" attribution** consists of the following criteria

1. Evidence of changes must be consistent with the proposed cause.

2. Evidence of changes must be inconsistent with any alternative causes proposed.

3. Some confidence level must be provided for the attribution statement.

Merz et al. (2012) had noted that no such attribution studies had satisfied the requirements of hard attribution. Fulfilling these criteria is challenging in this context, as natural variability and confounding play a large part in these studies. To see the true

causal nature of climate drivers on flooding, it is necessary to develop approaches which both rigorously test for associations between variables, but also overcome the issues of confounding and natural variability in these studies.

1.3 Key contributions and themes of the thesis

The goal of this thesis is the development of spatial multilevel methods for detecting and attributing long-term changes in environmental observational studies to large-scale drivers of interest. This will bridge the gap between current climate change projections (for example Bates (2009)) and observational data, from which evidence of increasing trends has not yet been produced.

This goal can be split into two common themes. The first involves the accurate detection of long-term changes in environmental observational studies, with a focus on river flows in Great Britain. In particular, the development of spatial multilevel methods for the accurate detection of trends is the main focus of Chapters 2 and 4. In Chapter 2, the focus is on the detection of countrywide trends in annual maximum river flow data, where the spatial correlation between stations is modelled through a Gaussian process. Using a multilevel approach allows for the pooling of gauging station data to accurately detect trends in annual maximum river flows that may not be possible using traditional at-site approaches. In this approach, we focus on a set of near-natural "benchmark" gauging stations in order to detect those trends solely driven by natural climate variability. In Chapter 4, there is a switch towards the analysis of a single river network, exploiting this network structure in order to understand long-term trends in flows for a single river. There is also a change in focus towards daily mean flows, which are of interest for water resource management. This method considers the river gauging station measurements on the network as a graph, which is comprised of nodes (the river gauging stations themselves) and edges between nodes (the directed river stretches between each station). The network structure of this river is encoded using a first-order conditional autoregressive (CAR) model. CAR models are primarily used for assessing disease incidences in areal data, and have not often been applied in an environmental setting such as this. The proposed method allows for the use of fast inference methods through the construction of a sparse precision matrix, and also respects the physical structure of the network.

In this thesis, methods are also developed for the accurate attribution of such trends to a large-scale driver of interest, in this case climate indices such as the North Atlantic

Oscillation and East Atlantic index. This attribution is the main focus of Chapters 2 and 3. We focus on such indices instead of precipitation (which one might expect to impact flooding in a more direct way than climate indices), as the indices are impacted by climate change in a more direct manner and thus represent a proxy for climate change (Guimarães Nobre et al., 2017). They can be modelled in a more simple manner in comparison with precipitation, which exhibits more complex behaviour of low variation but with intermittent peaks. Additionally, Sharma et al. (2018) highlight the lack of evidence for claims that an increase in precipitation will result in increased flooding. Indeed, Sharma et al. (2018) note that, in some cases, there are reductions in flooding magnitudes with increasing precipitation. Finally, the focus of this thesis is the exploration and attribution of the long-term, large-scale drivers of flooding, rather than short-term, local-scale drivers such as precipitation. The attribution method developed are illustrated throughout using the example of river flows in the UK, and take into account concerns over confounding with non-anthropogenic changes.

A more holistic approach to trend analysis is proposed in Chapter 2, incorporating detection of such trends and the "soft" attribution of these trends to large-scale climate drivers. However, further evidence is required to demonstrate a true causal relationship. In environmental observational studies such as the one seen in Chapter 2, where natural variability is a major factor, it is not trivial to apply standard methods to ensure that Merz et al. (2012)'s criteria are met in full. In Chapter 3, we propose a more rigorous approach to attribution of changes in these studies, developing a systematic checklist for assessing causality in environmental observational studies. This approach both fulfils and goes beyond the scope of Merz et al. (2012)'s "hard" attribution criteria. Changes in annual maximum river flows in Great Britain are observed to be consistent with changes in the East Atlantic index, and inconsistent with potential alternative causes discussed within the case study. A credible interval is provided for the estimate of the effect of the climate index upon flows. Beyond these requirements, the method also considers the interpretability of the model to the end-user, controls for unmeasured confounding and frames the problem in the context of causal diagrams which are revisited throughout the analysis to ensure diligence in considering possibilities. As a result, we obtain evidence that changes in annual maximum river flows in Great Britain can, in part, be attributed to the East Atlantic climate index.

1.4 Statistical methods

A wide variety of general statistical tools are used throughout this thesis in order to tackle the challenges posed by environmental observational studies. In this section, we summarise those methods common to the three main chapters of the thesis, and the novelty of their application to the field. We present the use of multilevel models as a useful framework for incorporating the spatial structure of river flow data. Multilevel models are used to estimate trends in river flows across multiple river gauging stations in Great Britain in a single model, instead of an at-site analysis. Potential spatial structures are discussed for representing the spatial correlation between river gauging stations, both on a countrywide and single river basis. The method of Bayesian inference on parameters is briefly summarised, including potential implementation methods such as the integrated nested Laplace approximation (INLA) approach. Generalised additive models are discussed for removing the seasonal effect from time series data of daily mean river flows. Finally, we briefly discuss causal methods for the attribution of trends in environmental observational studies such as these, however, such methods which are used on a single-chapter basis will be thoroughly discussed in the context of the particular chapter.

Note that throughout, the log of the river flows is taken and normality of the resulting data assumed, both in the annual maximum flows investigated in Chapters 2 and 3, and in the daily mean flows in Chapter 4. Taking the logarithm and assuming normality of this data has been found to be fit UK river flow data well (Prosdocimi et al., 2014; Vogel et al., 2011). One could fit generalised extreme value (GEV) or peaks over threshold (peaks over threshold) models to investigate high flows, however GEV models are complex to fit due to difficulties computing the maximum likelihood associated with the skewness (Coles, 2001), an issue that would be compounded using a more complex spatial multilevel model. Annual POT data are not as widely available as annual maxima flow data. Vogel et al. (2011) notes that using a log-normal pdf when combined with a log-linear trend model produces a simple yet appropriate model of flood frequency. As we are interested in understanding and attributing changes in annual maximum flows due to climate change on a UK-wide basis, it is beneficial to use this simpler approach for understanding changes in flows both in terms of model complexity and interpretation. These log-transformed flows will be used in the methods that follow.

1.4.1 Multilevel models

The common problem of detecting and attributing trends across multiple river gauging stations in Great Britain may be considered as a multilevel model, i.e. one that can vary at more than one level. Formally, a **multilevel model** is a regression model where model parameters have their own probability models (Gelman and Hill (2006)). This second-level model itself has parameters, known as *hyperparameters*, which are also estimated from the observed data. Thus, the multilevel model has varying coefficients and a model for these coefficients. They are particularly useful when the data for observational units are organised at multiple levels, and allow for one to take into account that there is variation between groups, while assuming that within-group observations are similar. Starting from a simple linear model where we have measurements from one source only,

$$Y_i = \beta_0 + \beta_1 X_i + \epsilon_i,$$

where Y is the outcome (for example, river flow data), X the exposure (a climate index), and ϵ represents the variation in the outcome that cannot be explained by the linear relationship with the exposure. To extend the model beyond a single measurement source, we need to allow for variation amongst different groups (river gauging stations throughout the thesis). To do this, we can allow the intercept in the model above to vary from group to group:

$$Y_{ij} = \beta_{0j} + \beta_1 X_{ij} + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma_\epsilon^2).$$

where Y_{ij} is the i^{th} measurement on the j^{th} group, and we now have β_{0j} which represents this group variation. This is the first-level model which accounts for the variation in the individual measurements on a single group. We can write β_{0j} as

$$\beta_{0j} = \beta_0 + u_{0j}, \quad u_{0j} \sim \mathcal{N}(0, \sigma_u^2).$$

This is the second-level model which accounts for the variation from one group to another. We can also allow the slopes of groups to vary. Incorporating a *random slope* within a model allows the explanatory variable to have a differing effect for each group. This enables us to determine whether a trend is consistent across groups or not. In a Bayesian setting, one can check this by inspection of the posterior of the variance of the random slope. By adding a random term β_{1j} to the coefficient of X_{ij} so that it can

vary for each group, the model then becomes

$$Y_{ij} = \beta_{0j} + \beta_1 X_{ij} + \beta_{1j} X_{ij} + \epsilon_{ij},$$

This model framework can take into account group-level predictors, for example region-based covariates. It can also incorporate interaction terms in the same way as for a single level model (Goldstein et al., 2014), and may also include smooth terms (Pedersen et al., 2019). Multilevel modelling allows one to more easily account for the variation between groups through the inclusion of random effects. Traditional multiple regression techniques treat the subjects under analysis as independent observations, rather than clustered or grouped, meaning that the standard errors of regression coefficients may be underestimated. When we consider multilevel models within a Bayesian framework and set priors to our parameters, this becomes a Bayesian hierarchical model. Bayesian models will be discussed further in Section 1.4.3.

1.4.2 Spatial random effects

To accurately capture the structure of river flow data, one random effect that we wish to incorporate in a multilevel model is a spatial random effect. This will allow us to model all stations together in a single model and capture the spatial correlation between stations. At first, we are interested in modelling a spatial effect at a countrywide level, where the spatial domain of river gauging station locations is not continuous. In such cases, as seen in Chapter 2, we propose a simple distance-based random effect, making use of spatial correlation functions implemented in the model through a Gaussian process. In Chapter 4, where the focus is across a single river network, we instead take advantage of the inherent graph structure of the gauging network, and implement the spatial random effect through a conditional autoregressive (CAR) model.

1.4.2.1 Spatial correlation structures and Gaussian processes

Spatial correlation structures are designed to model dependence in data, which is indexed by continuous two-dimensional position vectors, such as geostatistical data as presented in this thesis. The spatial correlation structures considered here are continuous functions of distance between these position vectors, and so may be generalised to any finite number of position dimensions - thus, they may be used with time series data

(see Pinheiro and Bates (2000)) such as river flow data as seen throughout this thesis. In Chapter 2, we propose the use of a simple exponential correlation structure with distance based on this Euclidean distance between river stations. Euclidean distances are chosen over stream distances in this case, as Ver Hoef et al. (2006) notes that stream distances can give rise to invalid covariance matrices. This will be further discussed in Chapter 4, along with potential alternative approaches, however for simplicity the Euclidean distance will initially be used in Chapter 2. The purely spatial correlation function depends on the monitoring locations through the Euclidean spatial distance only. Thus, we must assume that the spatial process is stationary and isotropic. Such a correlation structure can be included as a spatial random effect by modelling it through a Gaussian process.

A **Gaussian process** (GP) $f(x)$ is a collection of random variables, any finite number of which have a joint Gaussian distribution. A Gaussian process is completely specified by its mean function $\mu(x)$ and its covariance Σ. For $n \in \mathbb{N}$ and $x_1, ..., x_n$:

$$(f(x_1), ..., f(x_n))^T \sim \mathcal{N}((\mu(x_1), ..., \mu(x_n))^T, \Sigma)$$

A benefit of using Gaussian processes is that they can be completely defined by their second-order statistics (Bishop, 2006). This means that, if we assume a Gaussian process has mean zero, defining the covariance function will completely define the process' behaviour. Diggle et al. (1994) note that, while Gaussian processes are often used as models for spatial data such as the river flow data seen in this thesis, such models may lack physical justification. However, they are a convenient choice of empirical model which may capture a wide range of spatial behaviour, depending on how their correlation structure is specified. If the data of interest are not spatially dependent, then one can assume a diagonal covariance for the Gaussian process. However, if the data are spatially structured, as is the case in river monitoring networks, there is a need to include some terms for the spatial correlation. An example of such a covariance function is one with exponential correlation structure, in which nearby stations are highly dependent and far away stations are modelled to be less dependent:

$$\Sigma_{ij} = \eta^2 \exp\left(\frac{-d_{ij}}{\rho}\right),$$

where d_{ij} is the Euclidean distance between river stations. Here, the hyper-parameter ρ is the characteristic length-scale (or range) of the process. In practice, this represents how proximal two stations i and j must be in order to significantly influence one another, while η is the marginal standard deviation which controls the magnitude of

the range. In a Bayesian setting, priors can then be set on these parameters when specifying the statistical model used. Incorporating the spatial effect in this way in Chapter 2 will capture the spatial variability within the model when all stations are modelled together.

1.4.2.2 CAR models

Another means of modelling spatial structure between monitoring locations presented in this thesis involves defining neighbourhoods across the locations of interest. Two areas may be considered neighbours if they share a common boundary. This is also applicable to the example of representing the spatial effect of river gauging stations, which are immediately connected through a directed river stretch, as seen in Chapter 4. One can consider the structure of this network as a graph, where the nodes are the river gauging stations and the edges are these directed river stretches. Making use of this graph structure, autocorrelation is then modelled via a Gaussian distribution which has zero mean and a precision matrix Q (the inverse of the covariance matrix Σ) which will model correlation between neighbours. We assume that the latent effects are a Gaussian Markov Random Field (GMRF), i.e. a Gaussian random field which satisfies $p(x_i|\{x_j : j \neq i\}) = p(x_i|\{x_j : j \in \mathcal{N}_i\})$, where $\mathcal{N}_i$ are neighbours of locations s_i. this is to say that the distribution of x_i given all other nodes/stations is only influenced by the neighbouring nodes and not by the other nodes in network. Then the precision matrix Q will be such that elements $Q_{ij} = 0$ if i and j are not neighbours. Thus, Q is often a very sparse matrix, as there will only be entries in the matrix when i and j are neighbours. A benefit of this sparsity is that it will give rise to faster inference. This is particularly useful when dealing with large-scale data, such as the daily mean flow data over a long time period, as seen in Chapter 4.

One such structure for the spatial dependence is the conditional autoregressive (CAR) model as described in Besag et al. (1991), which averages over direct neighbours. CAR models display the spatial Markov property, i.e. they are spatially memoryless. For example, in the case of a river monitoring network, flows at a given station are influenced only by flows from direct neighbouring gauging stations, and not by neighbours of neighbours. Further details of CAR models and their use in representing the spatial random effect in a single river network will be discussed in Chapter 4.

1.4.3 Bayesian models

To make use of the inherent structure of spatial and spatio-temporal data for river flows, it is helpful to write models in a hierarchical manner. A hierarchical model is a particular type of multilevel model where parameters are nested within one another. We propose the use of a Bayesian hierarchical model as a flexible framework for statistical modelling, allowing us to perform inference to quantify levels in the models such as the underlying latent process (Gelman et al., 2013). Bayesian hierarchical models provide a flexible framework for the modelling of spatial data such as river flow data. While Bayesian methods have been used in river flow trend estimation on a very small scale (see Renard et al. (2006)), such models have not yet been used in the estimation of trends in a relatively large network of river gauging stations as seen in this thesis. We will make use of these Bayesian multilevel models throughout.

In general, the Bayesian approach allows us to make inference about model parameters through analysis of the posterior distribution via Bayes' rule. As we wish to use all individual river gauging stations to make inferences about the entire population of stations, we make use of the hierarchical form of Bayes rule as follows

$$\underbrace{p(\boldsymbol{\alpha}, \boldsymbol{\theta} \,|\, \mathbf{y})}_{posterior} \propto \underbrace{p(\mathbf{y} \,|\, \boldsymbol{\theta}, \boldsymbol{\alpha})}_{data} \; \underbrace{p(\boldsymbol{\theta} \,|\, \boldsymbol{\alpha})}_{process} \; \underbrace{p(\boldsymbol{\alpha})}_{prior}. \tag{1.4.1}$$

Again, $\mathbf{y}$ represents the variable under study (river flow data in this thesis), $\boldsymbol{\alpha}$ represent the population-level parameters and $\boldsymbol{\theta}$ the individual-level model parameters. The prior, $p(\boldsymbol{\alpha})$, represents the uncertainty in a (hyper) parameter prior to the data being observed. The data level, $p(\mathbf{y} \,|\, \boldsymbol{\theta}, \boldsymbol{\alpha})$ represents the likelihood for the observed river flow data, while $p(\boldsymbol{\theta} \,|\, \boldsymbol{\alpha})$ is the model for the parameters $\boldsymbol{\theta}$ which define the latent data generating process. The product of these three quantities is proportional to the posterior density, $p(\boldsymbol{\alpha}, \boldsymbol{\theta} \,|\, \mathbf{y})$. This quantity is the distribution of the parameters $\boldsymbol{\theta}$ having observed the data $\mathbf{y}$.

The exact form of this posterior distribution can rarely be computed analytically and thus must be derived via numerical methods. Often, a traditional Markov Chain Monte Carlo (Markov Chain Monte Carlo (MCMC)) method (Metropolis et al., 1953; Smith and Roberts, 1993) is employed to derive samples from the target posterior distribution. In most practical applications (including those presented in this thesis) a large number of parameters need to be estimated. As a consequence, generating predictions may be computationally difficult. This is due to the need to manipulate large matrices. We will discuss the use of a Hamiltonian Monte Carlo approach via Stan for inference

in Chapter 2, which will speed up inference compared to typical MCMC approaches in high dimensional settings. Stan is a general-purpose C++ program which can be used to derive samples from Bayesian models to obtain posterior simulations given a user-defined model and data (Carpenter et al., 2016). Stan uses a modified Markov chain Monte Carlo approach, known as Hamiltonian Monte Carlo (HMC), for sampling from these models. This approach was first developed as hybrid Monte Carlo in 1987 by Duane et al. (1987) to deal with difficult computations in lattice quantum chromodynamics. The usefulness of this approach in Bayesian problems was noted by Betancourt (2017); Neal (2012). The key difference between Hamiltonian Monte Carlo and the Metropolis Hastings algorithm typically used in MCMC schemes is that HMC reduces correlation between successive sampled states by using a Hamiltonian evolution between states, and additionally by targeting states with a higher acceptance criteria than the observed probability distribution. It adopts physical system dynamics rather than a probability distribution to propose future states in the Markov chain, allowing the chain to explore the target distribution much more efficiently. Further details on the Hamiltonian Monte Carlo scheme can be seen in Betancourt (2017) and are summarised in Appendix B.

We make use of an alternative to MCMC when the problem increases in complexity, such as in Chapter 4. This approach is briefly described in the section below.

1.4.4 Alternative solutions for Bayesian inference — approximation

While fitting models in Stan allows for user-defined model specifications and provides a flexible framework for modelling, it suffers from the typical efficiency problems of MCMC-based inference. As models grow in complexity, for example when large datasets such as the one in Chapter 4 are analysed, it is advantageous to move to a more efficient approximation-based approach. We employ the integrated nested Laplace approximation (INLA) technique proposed in Rue et al. (2009), which performs Bayesian inference based upon making a series of Laplace approximations and numerical integrations. This technique is a computationally attractive alternative to MCMC, as it does not require full MCMC sampling to be performed. The drawback of using an approximation method is that it is only easily applicable for a particular class of models (known as latent Gaussian models), and is not a general-purpose tool. However, the model discussed in Chapter 4 can be cast as one such model, and thus it is possible to make use of this efficient technique. In Chapter 4, the size of the data is much larger than in Chapter 2 (as daily mean flow data are used instead of

annual maxima). Using the INLA method instead of sampling-based approaches have a number of computational benefits to large-scale data such as this, which we briefly summarise here.

INLA provides a faster alternative to MCMC for a general class of models known as Latent Gaussian Models (LGM). LGMs are hierarchical models which include unobserved normally distributed random variables. LGMs have three levels; a likelihood model of observational data $\boldsymbol{y}$, a latent Gaussian field $\boldsymbol{z}$ and hyperparameters $\boldsymbol{\theta}$. It is possible to reframe many typical models (such as time series and spatial models as we have) as LGMs. We can formulate these models as follows:

$$y = \beta_0 + x^T \beta + \sum_k f_k(c_k),$$

where y represents the observational river flow data, β_0 is the intercept, β the regression coefficients of a linear predictor x, such as time or climate indices, and $f_k(\cdot)$ are non-linear smooth functions of other covariates. If we gather all model parameters in the linear predictor in a latent field $\boldsymbol{z} = \{\beta_0, \beta, f_k(\cdot)\}$, then a latent Gaussian model is obtained by setting Gaussian priors with zero mean and precision matrix $Q(\boldsymbol{\theta})$ to all elements of $\boldsymbol{z}$. The observations y are assumed to be conditionally independent given the latent field $\boldsymbol{z}$ and parameters $\boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is the vector of parameters describing the precision structure. The conditional independence assumptions are necessary for modelling through the INLA approach as discussed in Chapter 4, as approximate inference via INLA assumes that the latent field z is Gaussian, and fulfils the conditional independence property that any two latent effects z_i and z_j are conditionally independent given the remaining latent effects z_{-ij}. Note that, while the assumption of Gaussianity is not necessarily appropriate for the raw daily mean river flow data, taking the logarithm and assuming normality of this data has been found to be appropriate (Bowers et al., 2012; Prosdocimi et al., 2014) for such data.

The posterior distribution of this model is then given by

$$p(\boldsymbol{\theta}, \boldsymbol{z}|\boldsymbol{y}) = p(\boldsymbol{\theta})p(\boldsymbol{z}|\boldsymbol{\theta}) \prod_{t=1}^{T} p(\boldsymbol{y}_t|\boldsymbol{z}, \boldsymbol{\theta}) \qquad (1.4.2)$$

To preserve the underlying conditional independence structure of a GMRF (as defined in Section 1.4.2.2), the parameterisation must be constrained. This is difficult to achieve using a covariance matrix Σ, as typically this will be dense, i.e. have very few null

entries. Instead, we use the precision matrix $Q = \Sigma^{-1}$. It can be shown that

$$z_i \perp z_j | \mathbf{z}_{-ij} \iff Q_{ij} = 0.$$

It follows that, under conditional independence assumptions, Q will tend to be sparse, leading to faster inference and an ability to perform calculations on larger problems. Letting the dimension of a sparse precision Q for a typical two-dimensional model be $n \times n$, any calculations such as doing inference or computing normalising constants requires computations of order $\mathcal{O}(n^{3/2})$ compared to storage costs of $\mathcal{O}(n^2)$ and computation costs of $\mathcal{O}(n^3)$ for the corresponding dense Σ (Bakka et al., 2018). In Chapter 4, exploiting the network structure of a single river allows for the construction of a sparse precision matrix, meaning that inference can be performed on a large number of observations of daily mean data across the network for a 40 year period (a total of 155,398 observations) in a very short period of time (minutes). This would not be possible to achieve using standard MCMC methods, where the sparsity structure of precision matrices is not exploited.

Further details on the INLA approach to approximation can be seen in Rue et al. (2009) and are summarised in Appendix B. This approach to Bayesian inference is a much more efficient computational technique than MCMC methods due to the favourable computational properties of GMRFs and sparse precision matrices, and hence is ideal for use on large-scale problems as seen in Chapter 4.

1.4.5   Generalised additive models

A generalised additive model (GAM) is a generalized linear model in which the linear predictor depends linearly on unknown smooth functions of some predictor variables (Hastie and Tibshirani, 2017). In general, this has the form:

$$g(\mu_i) = \beta_0 + f_1(x_{1i}) + f_2(x_{2i}) + f_3(x_{3i}, x_{4i}) + ... \tag{1.4.3}$$

where x_i represents covariates, $\mu_i \equiv \mathbb{E}(Y_i)$ and Y_i is a member of some exponential family. The smooth term f is the sum of some number of basis functions,

$$f(x_1) = \sum_{j=1}^{k} b_j(x_1)\beta_j$$

35

GAMs allows for a flexible, relatively simple specification of a model with non-linear smooth functions. One can investigate whether any covariates have non-linear relationships with the response, in order to see whether an approach using GAMs is necessary. Additionally, GAMs allow for the modelling of seasonal effects using a cyclical smooth term, as will be demonstrated in Chapter 4. In this case, basis functions used are similar to those above, however the end points of the spline are constrained to be equal, which is necessary for modelling a cyclical variable such as days, months or years.

In R, models will be implemented using the `mgcv` package (Wood, 2017), with smooth functions specified through the `s()` function. Such smooth functions are specified as `s(x,bs="cc",k="")`, where x is the covariate of interest, `bs` represents the choice of basis, which is in Chapter 4 given by a cyclic cubic regression spline. For a normal cubic regression spline, `mgcv` estimates a coefficient for each basis function, and the resulting spline is a weighted sum of basis functions whose weights are the estimated coefficients. However this gives rise to a large discontinuity in the value that the spline takes at the end. If the covariate represents the day of year, as in Chapter 4, this is not ideal. Instead, we use the cyclic cubic spline, which has a further constraint that there may not be a discontinuity at the end point. Finally, `k` represents the basis dimension size, which must be specified by the user. The value `k-1` sets the upper limit on the degrees of freedom allowed, as one degree is taken up by the identifiability constraint on the smooth (Wood, 2017). This value should be chosen to be large enough so that the truth is represented well and so one does not oversmooth, but small enough so that computations are efficient (Wood, 2017). This use of GAMs for removing seasonality from time series data will be demonstrated in Chapter 4.

1.4.6   Causal approaches

In this thesis, we propose a systematic causal checklist for attribution in these studies, making use of relevant ideas in prior causal approaches and incorporating both data and non-data assumptions into our method. By taking inspiration from methods such as the Bradford Hill criteria (Hill, 1965), causal directed acyclic graphs (DAGs) (Greenland et al., 1999), the method of multiple working hypotheses (Chamberlin, 1890) and weighting of evidence methods (Suter et al., 2017), along with proposing a number of additional steps, it is instead possible to develop a systematic approach for the attribution of long-term, large-scale changes in environmental processes. These methods are discussed in detail in Chapter 3 and thus will not be further discussed here.

1.5   Thesis structure

The structure of this thesis is as follows:

- In Chapter 2, we present a multilevel Bayesian model for the detection and an initial attribution of trends in annual maximum river flows in Great Britain.

- In Chapter 3, we propose a systematic checklist for the causal assessment of environmental observational studies. The checklist is illustrated through the example of annual maximum river flows in Great Britain.

- In Chapter 4, we investigate a directional approach to detecting long-term trends in daily mean flows in a single river network. This is illustrated by the example of the river Eden network, which has been hit with a large number of flooding events in the past couple of decades.

- In Chapter 5, we provide a summary of the results and contributions of the thesis, as presented in Chapters 2-4. Conclusions and potential future opportunities in the areas of attribution and network-based methods for river flow time series are discussed. Finally, the key impacts of this thesis are summarised.

Chapter 2

Attribution of long-term changes in peak river flows in Great Britain

This chapter is reproduced from the authors' accepted manuscript of an article published as the version of record in Hydrological Sciences Journal (Brady et al., 2019). Related work on attributing long-term changes in peak river flows for the island of Ireland can be seen in Brady et al. (2018).

In this chapter, we demonstrate how the use of a Bayesian multilevel approach allows for the detection of trends in annual maximum river flows which would be missed in an at-site approach due to poorly powered statistical tests. The use of a Bayesian approach here provides a more intuitive inference over the frequentist approach of p-values (O'Hagan, 2004). The p-value does not inform us of how likely the null hypothesis given the data, while the Bayesian approach provides a statement on the probability that the hypothesis is true given provided evidence. Spatial correlation between stations is incorporated using a Gaussian process, with covariance specified using an exponential correlation structure. Time trends in annual maximum flows are estimated on a countrywide scale, with the aim of bridging the gap between current climate projections of increased flooding and observational data, from which no concrete evidence of trends have been found. In addition, a preliminary approach to attributing such trends to a large-scale cause (climate indices) is investigated. Limitations to this approach are discussed at the end of this chapter.

The major contribution of this chapter is the use of Bayesian spatial multilevel meth-

ods to accurately assess trends in peak river flows in Great Britain. This is a novel application of such methods, as typically trend analyses for river flow studies occur on an at-site basis. It is demonstrated in this chapter that assuming an overall country-wide trend with some spatial variability (with relationships between stations modelled through a Gaussian process) leads to a better ability to detect any signal present in the data. This chapter also provides some evidence towards attributing such changes in peak river flows to a climate index of interest. The switch towards a combined detection and attribution approach in this paper differs from past approaches, which tend to put a large emphasis on detection only.

Abstract

We investigate the evidence for changes in the magnitude of peak river flows in Great Britain. We focus on a set of 117 near-natural "benchmark" catchments to detect trends not driven by land use and other human impacts, and aim to attribute trends in peak river flows to some climate indices such as the North Atlantic Oscillation (NAO) and the Eastern Atlantic (EA) Index. We propose modelling all stations together in a Bayesian multilevel framework to be better able to detect any signal which is present in the data by pooling information across several stations. This approach leads to the detection of a clear countrywide time trend, which is consistent with climate change projections (Bates, 2009). Additionally, in a univariate approach, both the EA and NAO indices appear to have a considerable association with peak river flows. When a multivariate approach is taken to unmask the collinearity between climate indices and time, the association between NAO and peak flows disappears, while the association with EA remains clear. This demonstrates the usefulness of a multivariate and multilevel approach when it comes to accurately attributing trends in peak river flows.

2.1 Introduction

The 2013-14 winter floods saw the destruction of a railway line, thousands of homes without electricity and costs in excess of £100 million (Chatterton et al., 2016). A further series of heavy rainfall events hit Great Britain in the 2015-16 winter, with Storm Desmond breaking the UK 24-hour rainfall record and Christmas floods resulting in the highest water levels of every river in Lancashire (Barker et al., 2016). Insurance payouts for that flood period were estimated to be around £1.3 billion. As a consequence, there is a fear and widespread suspicion that there has been an increase in the frequency and severity of flooding in Great Britain. Climate change projections suggest an increase in rainfall (Bates, 2009) over the coming decades, with the belief that this may contribute to an increase in peak river flows.

Much research has been devoted to the identification of trends in river flow records, yet current methods do not appear to be fit for purpose. These have mostly involved performing some tests at each gauging station separately — see for example Hannaford and Marsh (2006); Mediero et al. (2014); Villarini et al. (2009). However, these at-site tests using the relatively short observed river flow data records do not display compelling evidence of increasing trends. Such an approach tends to involve fitting a model for some yearly summary value (e.g. annual maximum flow, annual number of events, etc) at each individual station, and evidence for monotonic trends is often derived using specific statistical tests, for example the Mann-Kendall test (Kendall, 1948; Mann, 1945). These tests are not very powerful in a statistical sense (that is, there is a non-negligible probability of not detecting a trend) — Prosdocimi et al. (2014) noted that a sample size of hundreds of years may be needed when using a station-by-station approach, while reliable records of river flow are typically much shorter than a hundred years. As a consequence, we may not be detecting some trends in the annual maximum flow, as each individual hypothesis test is not sufficiently powerful to detect them — meaning we may be unable to correctly quantify changes in flood risk. Additionally, as this approach involves modelling each station separately, the result at each station is based solely on the data available at the station itself. Many gauging stations are geographically close and are exposed to similar weather conditions, yet this approach does not use information from nearby stations, which should enhance the signal at the station in question, and improve the ability to detect trends provided this trend is the same across all stations.

Another key issue with current models (as noted by Merz et al. (2012)) is the focus on the detection of time trends, rather than the attribution of these trends to meaningful

flood-generating processes. Gaining a clearer understanding of how these flows change over time is certainly of interest, however time itself cannot explain any of the variability in peak river flows, instead it is a surrogate for variables which vary with flows in the same way as time. Thus, this focus on detection does not give us any insight into the drivers of change in peak river flows. Gaining an understanding of these drivers of change may help to better inform future flood defences, thus we instead propose to investigate some potential candidates for attribution of changes in these flows alongside an analysis of these time trends.

Merz et al. (2012) give a (non-exhaustive) list of the several studies which have investigated trends in peak river flows and associated flood risk over the past twenty years, typically with a focus on detection of changes, i.e. on the analysis of time series data, and attribution has largely followed as "an appendix" of a hypothesis test for significance of such changes. The authors noted the need for a switch in focus towards a more holistic approach, which would incorporate change detection into the more challenging problem of attribution. The authors discuss attribution in the context of "soft" and "hard" attribution. Here, "soft" attribution refers to studies which use hypotheses and references to previous studies to back up any attributions to drivers of change. On the other hand, "hard" attribution studies must provide evidence that detected changes are both consistent with the proposed driver of change and inconsistent with potential alternative drivers. Additionally, Merz et al. (2012) require that such "hard" attribution studies provide some confidence level in the attribution statement. The approach proposed in this paper will fulfil the criteria of "soft" attribution and the majority of the "hard" attribution requirements. A small number of studies have investigated whether climate indices can describe the observed variability in the frequency of flooding in river networks (see (Lins and Slack, 2005; Mallakpour and Villarini, 2015; Villarini et al., 2011) amongst others). Tootle et al. (2005) considered a network of 1009 "unimpaired" catchments in the United States with data from 1948-1988. By applying a non-parametric rank-sum test to test for significance, they demonstrated that a number of climate indices influence stream flow variability in the US. Hodgkins et al. (2017) investigated trends in floods for a set of over 1200 catchments across North America and Europe over time periods from 1961 to 2010 and 1931 to 2010, noting a much larger link between these occurrences and the Atlantic Multidecadal Oscillation (AMO) when compared to long-term time trends. However, these investigations tend to be on an at-site basis, leading to small sample sizes and low power of hypothesis tests. Additionally, Mallakpour and Villarini (2016) noted that the choice of optimal large-scale drivers of climate is particularly challenging so care must be taken when identifying appropriate climate indices.

There is a need for a new approach to the modelling of trends in peak river flow, in order to overcome these issues and improve the ability of tests to detect signals. Instead of focusing solely on detection of time trends, we instead propose focusing on the combined approach of the detection and attribution of trends. We relate trends detected in peak river flows to large-scale climate indices (which are proxies for climate variability) such as the North Atlantic Oscillation or the East Atlantic Index. Additionally, it is often difficult to separate out anthropogenic changes from natural climate variability, making it difficult to accurately attribute any such trends. To avoid the presence of additional potential confounding variables such as urbanisation levels, we focus on a set of near-natural "benchmark" catchments as defined by Harrigan et al. (2017) so that those predominantly climate-driven trends can be detected. In order to improve the statistical power of our approach compared with at-site testing, we propose modelling all stations together in a multi-level model framework. Specifically, Bayesian multilevel models are employed, which have widely been used for the modelling of spatial and spatio-temporal environmental data (Diggle et al., 2010; Pirani et al., 2014; Renard et al., 2006). Such models provide a framework which allows the pooling of information between stations, improving the possibility of detecting of signals which may be missed in an at-site approach. A Bayesian approach also allows for a clear uncertainty statement for attribution. Finally, the necessity for a multivariate approach to modelling trends in peak river flows is demonstrated, in order to accurately separate out the net effect of individual covariates. This combined multilevel multivariate approach will help to give a clearer picture of the drivers of peak river flows in Great Britain.

In section 2.2, multilevel models are introduced in the context of the attribution of such trends. A framework is developed for modelling of spatial dependence in these peak river flows. In section 2.3, the annual maximum river flow data series, climate indices and the reference network of near-natural catchments used in the model are presented. In section 2.4, this model is implemented for various climate indices to peak river flows in Great Britain and findings in both a univariate and multivariate setting are presented. The work is summarised, and future possibilities discussed in section 2.5.

2.2   Methods

In this section, an extension to linear models used to ascertain trends in peak river flows is proposed. Multilevel mixed effect models are introduced, in order to incorporate all stations together into one Bayesian modelling framework. Bayesian multilevel mixed

effect models allow for all stations to be modelled together in a unique model which pools the information across stations, which may help to better detect signals that cannot be found using an at-site analysis. The addition of a spatial structure to these models is also introduced, in order to account for similarities between stations which have some proximity to each other.

2.2.1 Spatial dependence

Many river gauging stations in Great Britain are geographically close to each other, and one might expect the peak river flows of these nearby stations to follow a similar pattern. Spatial correlation structures are designed to model dependence in data, such as times series obtained at fixed gauging stations. The work of Kjeldsen and Jones (2010) on spatial correlation in British annual maximum river flow data suggests that one can expect data from nearby gauging stations to be correlated with each other. This correlation structure is exploited in the approach discussed here, by including a spatial random effect to account for this similarity between nearby stations in the multilevel models proposed in Section 2.2. This is a correlation-based approach which can be used generally in any spatial setting, particularly for scenarios where measurement locations are fixed and the spatial domain is not continuous (such as river gauging station data). This pooling of information should help enhance any signal, helping us to obtain evidence of any trends in river flows that would have been too weak to detect otherwise.

Diggle et al. (1994) noted that the most common form of empirical behaviour for stationary correlation structure is that the correlation between sites i and j decreases as the distance between them increases. This shape of correlation appears to be valid for British rivers as shown in Kjeldsen and Jones (2010). Thus, we seek models whose theoretical correlation structure behaves in this way. We propose the use of an exponential correlation structure which satisfies this requirement; this simple structure is often used in environmental studies (see for example, Reich et al. (2011)). This purely spatial correlation function depends on the locations through the Euclidean spatial distance d_{ij} only, so that for gauging stations i and j, the covariance matrix is described as:

$$\Sigma_{ij} = \eta^2 \exp\left(\frac{-d_{ij}}{\rho}\right).$$

The parameter ρ describes the range over which sites i and j influence each other i.e. how close two points must be to influence each other significantly, while η is the marginal standard deviation controlling the magnitude of this range. Note that, as in

Kjeldsen and Jones (2010), the distance between two stations is taken to be the distance between the centroid of the catchments upstream of each station rather than the distance between the stations themselves. It has been observed that correlation between sampling errors can at least partially be described as a function of the geographical distance between river catchments (Kjeldsen and Jones, 2009; Stedinger et al., 1993).

2.2.2 Multilevel models

Given that a sample size of hundreds of years may be needed to construct powerful tests for trends when using an at-site approach, we propose extending the approach of Prosdocimi et al. (2014) by modelling all stations together within a multilevel framework. This has the benefits of improving the power of any models used by incorporating all data together within one model, making use of the natural structure of the spatial data provided. Including all stations in a model together with a spatial random effect may help to obtain evidence of any trends in river flows that would have been too weak to detect otherwise.

In particular, a Bayesian perspective to multilevel models (see section 2.2.3) is adopted, thus considering the model parameters as random variables with their own probability models (Gelman and Hill, 2006). These probability models themselves have parameters (known as hyper-parameters) which are also estimated from the observed data. Looking first to the original at-site model which states that the log of the standardised annual maximum flows, denoted Y, are affected by covariates $\boldsymbol{X}$ in the following way

$$Y_t = \beta_0 + \boldsymbol{X}_t \beta_1 + \epsilon_t$$

for a given matrix of covariates $\boldsymbol{X}$ which vary across years t, with regression coefficients β_1, error term $\epsilon_t \sim N(0, \sigma^2)$, where σ^2 represents the variation in flows after controlling for covariates X, and intercept β_0. . Here β_0 and β_1 represent the intercept and slope, which do not vary in time. Note that the log of the standardised annual maximum flows is used, i.e. flows divided by the median of the annual maximum series (QMED). The log of the annual maximum flow data is assumed to be normally distributed. Using a log normal distribution has been found to fit UK annual maximum flow data reasonably well (Prosdocimi et al., 2014). Using the median is considered to be more robust to outliers than the mean (Robson and Reed, 1999). A simple version of this model in which the water year is the only explanatory variable, i.e. $\mathbf{X}_t = [\text{Water Year}]_t$ (i.e. the value for the explanatory variable at time t), was used by Vogel et al. (2011) and Prosdocimi et al. (2014).

At-site investigations often fail to identify trends (Prosdocimi et al., 2014). It is difficult for authorities to make decisions regarding flood defences based upon this approach, which can prove unreliable. Instead, all stations are modelled together to better detect any signals, assuming that the peak flows are affected by some countrywide trend. A multilevel approach is used to allow for different station-specific effects which are expressed as random effects. The peak flows can then be modelled as

$$Y_{it} = \beta_0 + \mathbf{X}_t\beta_1 + r_i + \epsilon_{it}$$

for a gauging station i at time t. We assume that $r \sim N(0, \sigma_i^2)$ and the remaining error is now $\epsilon_{it} \sim N(0, \sigma^2)$, where σ_i^2 represents the variation in flows due to differences between gauging stations after controlling for the covariates X. Now β_0 and β_1 represent the overall countrywide intercept and slope respectively. The assumption of a countrywide trend is a strong assumption, but necessary to ensure that even with little data, it should be possible to detect trends. To balance this assumption to some extent, station-specific effects are included to allow for some variability between stations.

Nearby stations can be expected to be impacted in a similar way by external variables, thus a spatial correlation structure s is included within the multilevel model. For station i at time t, this can be expressed as:

$$Y_{it} = \beta_0 + \mathbf{X}_t\beta_1 + r_i + s_i + \epsilon_{it}, \tag{2.2.1}$$

where $\mathbf{X}$ is the matrix of explanatory variables we are investigating, r_i is a random effect to allow for variation between stations with $r \sim N(0, \sigma_i^2)$, s_i is a spatial random effect distributed as a multivariate normal ($s \sim MVN(0, \Sigma)$) to allow for correlation between nearby stations and $\epsilon_{it} \sim N(0, \sigma^2)$ represents the error term. The exponential correlation structure discussed in Section 2.2.1 is used here. A further modification of the model in equation (2.2.1) could include station-specific properties such as catchment size or altitude for each station i as explanatory variables. However, through some investigations it was found that such modifications do not improve the model performance in terms of explanatory power and are therefore not discussed further.

2.2.3 Bayesian inference

When a model such as the one shown in equation (2.2.1) is considered within a Bayesian framework and prior distributions are set on model parameters (i.e. before any data is observed), this becomes a Bayesian hierarchical model (Gelman et al., 2013). Using

such an analysis can allow for the incorporation of further data, by pooling information across gauging stations in one model. It also allows us to incorporate prior knowledge about parameters before observing the data, and provides a straightforward framework to assess the uncertainty in the estimation of parameters and functions of the parameters. It can provide a more intuitive and meaningful inference over the frequentist approach of p-values, through the inclusion of prior information about model parameters (O'Hagan, 2004). Additionally, as discussed in detail in Section 2.4.1, a Bayesian analysis avoids the issue of the multiple testing problem (Gelman et al., 2012). This approach allows us to make inference about model parameters through analysis of the posterior distribution via Bayes' rule:

$$\underbrace{p(\theta\,|\,y)}_{posterior} \propto \underbrace{p(y\,|\,\theta)}_{likelihood}\ \underbrace{p(\theta)}_{prior} \tag{2.2.2}$$

for parameters $\theta = \{\beta, r, s, \sigma^2\}$ and observations $y = \log(\text{Flow})$. The prior represents uncertainty about a parameter(s) before the data is observed, while the likelihood is the conditional density of the data given the parameters. The product of this is proportional to the posterior density, which describes the uncertainty about the unknown parameter(s) having observed the data. This posterior density is the output of a Bayesian inference. This quantity is of interest and in particular one often looks at the 95% credible interval, i.e. given the data observed, this is the interval in which the parameter is contained with 95% probability. Bayesian hierarchical models provide a flexible framework for statistical modelling of spatial data such as this, allowing one to perform inference to quantify levels in the models such as the underlying latent process. For further information on Bayesian methods in trend estimation, see Renard et al. (2006).

As we wish to use all individual gauging stations to make inferences about the entire population, the hierarchical form of (2.2.2) will be used as follows:

$$\underbrace{p(\alpha,\,\theta\,|\,y)}_{posterior} \propto \underbrace{p(y\,|\,\theta,\,\alpha)}_{data}\ \underbrace{p(\theta\,|\,\alpha)}_{process}\ \underbrace{p(\alpha)}_{prior}$$

for population-level parameters α_i, individual-level parameters $\theta_i = (r_i, s_i)$ and observations of flow data y_i. For the proposed model (Equation 2.2.1), the data level is modelled via a Gaussian likelihood where the log of the annual maximum flow observations is taken:

$$p(y|\beta, \mathbf{r}, \mathbf{s}, \sigma^2) = \prod_{i=1}^{n} N(y_i | x_i^T \beta + r_i + s_i, \sigma_\epsilon^2)$$

48

The process level (i.e. the physical drivers of peak river flows) is determined by the priors over β, r and the Gaussian process s given α, which corresponds to the parameters for the exponential covariance structure for the spatial random effect. Priors are then specified on those parameters $\alpha = (\rho, \eta^2)$ which will be estimated. Standard procedure is followed for prior specification of regression coefficient parameters β_i, which are given independent Gaussian priors

$$\beta_i \sim N(\mu_\beta, \Sigma_\beta),$$

where μ_β is a q-dimensional mean vector, and Σ_β is a $q \times q$-dimensional covariance matrix. We suggest a mean of 0, and a relatively large precision in this case. The variance σ^2 is given a half-Cauchy prior with scale 2.5 as suggested by Gelman et al. (2008)

$$\sigma^2 \sim \text{Cauchy}(0, 2.5 * b),$$

where b is the standard deviation of the residuals of a linear regression of covariates X against flows. The covariance parameters ρ and η^2 are given weakly informative half-normal priors following the recommendations of Gelman et al. (2017),

$$\rho \sim \text{half-normal}(0, \sigma_\rho), \quad \eta^2 \sim \text{half-normal}(0, \sigma_{\eta^2}),$$

where half-normal means that values are constrained to lie above zero. In most cases, the posterior is not mathematically easily tractable except in the cases of small numbers of dimensions. However, samples from this distribution may be generated using Markov Chain Monte Carlo (MCMC) methods (Smith and Roberts, 1993). Models proposed here are implemented through Stan (see Section 2.4), a general purpose tool which makes use of Monte Carlo techniques to carry out Bayesian inference (Carpenter et al., 2016). Stan is a C++ program which draws samples from Bayesian models to obtain posterior simulations given a user-defined model and data (Carpenter et al., 2016). This approach uses a modified Markov chain Monte Carlo approach for sampling from these models (see Appendix B). Diagnostics are carried out during modelling to check for convergence of the modelling procedure and to ensure the effective sample size is sufficient. The potential scale reduction factor, $\hat{R}$, provides an estimate of convergence, which can be interpreted as the factor by which the variance of an estimate can be reduced with longer chains. We seek values close to 1 (and at most $\hat{R} < 1.1$), which will happen as the number of simulations approaches infinity. If samples obtained by the sampler are independent, then the effective sample size N_{eff} is equal to the actual sample size. On the other hand, if the correlation between samples decreases so slowly that the sum in the denominator diverges, the effective sample size is zero (Ripley,

2009). Markov chains tend to explore the parameter space very slowly, leading to low effective sample size numbers, and parameters may not be accurately estimated.

2.3 Peak river flow data for Great Britain

2.3.1 Benchmark catchments and peak river flow data

It has been noted (Hannaford and Marsh, 2006, 2008) that there can be considerable difficulties in accurately attributing climate-driven trends to peak river flows in Great Britain, largely due to the impact of humans and changes in hydrometric performance in gauging stations over time. This has led to a focus on developing a series of dedicated networks of natural catchments, in order to study trends over time. In the UK, an initial benchmark network (UK Benchmark Network V1 (UKBN1)) consisting of 122 catchments was developed by Bradford and Marsh (2003). This aimed to use catchments which had long records of good hydrometric quality, relatively near-natural and representative of UK hydrology. Such natural gauged catchments tend to be small and rural, located predominantly in Wales, Scotland and the south west of England. An updated version of the benchmark network, obtained as a compromise between geographical coverage and a lack of external interference on river flows, was therefore developed for the detection and attribution of climate trends. This known as the UK Benchmark Network V2 (UKBN2) (Harrigan et al., 2017) and consists of 146 benchmark catchments as a representative reference network. Annual maximums of instantaneous annual maximum flow data for the V2 benchmark catchments will be used to investigate the effect of non-anthropogenic (i.e. non-human driven) changes on annual maximum river flows in Great Britain.

This analysis focuses on annual maximum river flow data from Great Britain across the series of reference benchmark catchments introduced by Harrigan et al. (2017). This data set contains the largest observed instantaneous annual maximum flows in each water year (which runs from October to September), measured in m^3/s. After removing stations for Northern Ireland in order to include a spatial effect, and including all stations for which there are available observations, in total, there are 5475 observations from 117 benchmark gauging stations in Great Britain, ranging from 1851-2015 — these station locations can be seen in Figure 2-1. The average record length of stations within this network is 46 years, with a minimum of 21 and maximum of 86 years. There is an average of 1.4% of records missing, with only 5 stations having records of 10–30%

missing. The average catchment size of the network is 210km^2, with a minimum size of 3.07km^2 and a maximum of 1500km^2. A total of 92% of stations may be considered to be "essentially rural" under the FEH An index of urban and suburban land cover in a given period expressed as a fraction (URBEXT) criteria (i.e. with less than 2.5% of the catchment area covered by urban landmass, see the Flood Estimation Handbook (Robson and Reed, 1999)). Only two stations exceed 10% urbanisation; both of these stations are in areas with spatial gaps in the network, and are included in UKBN2 as a compromise to ensure full spatial coverage. Plots of some of the FEH catchment descriptors can be seen in Figure 2-1. These plots show the catchment area, the baseflow index (BFIHOST) and average annual rainfall for each catchment.
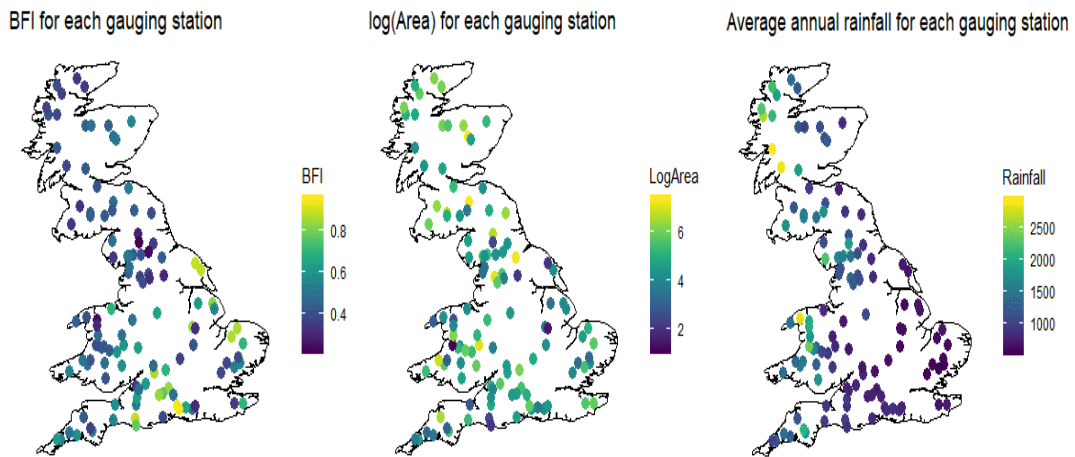
Figure 2-1: Plots of FEH catchment descriptors for each gauging station – the baseflow index (BFIHOST, left), log of the catchment area (middle) and average annual rainfall (right).

The data for the UKBN2 catchments, annual maximum series, benchmark catchments, peaks over threshold, and catchment descriptors can be obtained from the UK National River Flow Archive (NRFA) (Dixon, 2010), which is the primary UK source of hydrometric data.

2.3.2 Climate indices

Climate is defined as the average state of the atmosphere over long time periods, thus changes in climate are considerably slower than the weather. A climate index is defined as some calculated value that describes the state of the climate system, and any changes, including weather, occurring in the system (Integrated Climate Data Center,

2011). These indices are impacted by climate change in a more direct manner than precipitation (Guimarães Nobre et al., 2017), so they are proxies for climate which is changing (but are also variable to begin with). Moreover, these indices are constant across the whole country for a given time. In section 2.4, we investigate whether changes in annual maximum river flow can be attributed to some of these indices. Note that the values used are the average of the monthly values for December, January and February in a given year. The two key climates indices used in this study are the North Atlantic Oscillation and the East Atlantic index. These indices have been indicated in previous studies as potential drivers of variability in peak flow records (see for example Guimarães Nobre et al. (2017)), and are introduced briefly below.

2.3.2.1 North Atlantic Oscillation

The North Atlantic Oscillation (NAO) is a mode of natural climate variability, which impacts the weather and climate of the North Atlantic region and surrounding continents, particularly Europe (National Climatic Data Center, 2017). Usually, the North Atlantic surface pressure is relatively high in the subtropics at latitudes 20°N to 40°N ("the Azores High"), and lower further north at latitudes 50-70°N (known as the "Icelandic Low"). This state extends through higher levels in the atmosphere, and affects the north-south pressure difference, which determines the strength of the westerly winds directed from North America towards Europe. The NAO describes these fluctuations in north-south pressure differences.

When the NAO index is well above normal, the chances of above average seasonal temperatures in northern Europe increase. Precipitation patterns are more localised, with an increased chance of higher rainfall in northwest Europe and lower rainfall in southern Europe associated with a higher than usual NAO. When the NAO index is well below normal, the opposite tends to occur. The fluctuations in the NAO occur on a wide range of time-scales. There are day-to-day changes associated with weather systems, and slower changes associated with seasonal and longer term variability in other climate system components such as ocean temperature. Previous studies have shown that the NAO has been related to the variability in floods (Hannaford, 2015; Hannaford and Marsh, 2006; Kingston et al., 2006), while Macdonald and Sangster (2017) found statistically significant relationships between the British flood index and the North Atlantic Oscillation Index with historical data records. This provides strong motivation for inclusion as a possible driver of changes in peak river flows in Great Britain.

A correlation plot between NAO and time can be seen in Figure 2-2, which suggests a correlation of 0.49 between the two. This indicates towards possible confounding issues in the analysis, i.e. the inability to separate the effect of time from that of the NAO index. A confounding variable is one which influences both the dependent variable (the log of the standardised annual maximum flows in this analysis) and the independent variables, leading to spurious associations (see for example Faraway (2014)). Confounding will need to be taken into account in the analysis.

2.3.2.2   East Atlantic Index

The East Atlantic (EA) index (National Climatic Data Center, 2017) is a mode of low-frequency variability over the North Atlantic, similar in structure to the North Atlantic Oscillation (NAO) index. It consists of a north-south dipole of anomaly centres spanning the North Atlantic from east to west. Positive phases of the EAindex are associated with above-average surface temperatures in Europe in all months, above average precipitation over northern Europe and Scandinavia, and below average precipitation across southern Europe. The EAindex exhibits strong multi-decal variability across records from 1950-2004. Guimarães Nobre et al. (2017) noted that positive (negative) phases of both the NAO and EA are associated with more (less) frequent and intense seasonal extreme rainfall over large areas of Europe. A correlation plot between EA and time can be seen in Figure 2-2. Again, there appears to be a positive correlation ($r = 0.53$) between EA and Water Year, indicating a possible confounding between the two variables.

Figure 2-2: Correlation plots for water year with NAO (left) and EA (right) where R is the correlation coefficient between variables, and the red line represents a linear fit.

2.4 Results

Models investigating the relationship between the log of the standardised annual maximum flows and time, NAO and EA are fitted in this section. We look to fit models of the form in Equation 2.2.1, changing the covariates X_t in each case (where the subscript t indicates that the covariate is indexed by time).

1. **Model A:** $X_t = [\text{Water Year}]_t$ (Figure 2-4)

2. **Model B:** $X_t = [\text{EA}]_t$ (Figure 2-5)

3. **Model C:** $X_t = [\text{NAO}]_t$ (Figure 2-5)

4. **Model D:** $X_t = [\text{Water Year, EA}]_t$ (Figure 2-6)

5. **Model E:** $X_t = [\text{Water Year, NAO}]_t$ (Figure 2-7)

6. **Model F:** $X_t = [\text{Water Year, EA, NAO}]_t$ (Figure 2-8)

These models are implemented in Stan (Carpenter et al., 2016). In terms of diagnostics, all values for the $\hat{R}$ for models run fall between 0.999 and 1.001. For a model run with 4 chains, each of 1000 iterations (and a burn-in period of 500 iterations), the N_{eff} is at least 1500 for each parameter in the model, which is more than sufficient. Diagnostic plots are included as an appendix.

54

2.4.1 At-site approach for comparison

At first, models are fitted using the at-site approach described in Prosdocimi et al. (2014) for comparison. A linear model is fitted to observations at each individual station, then the p-value is extracted to determine whether there is a significant trend. The null hypothesis is that there is no significant trend (for time, EA or NAO) in the annual maximum flow data. If the p-value is less than 0.05, this null hypothesis is rejected and the conclusion is that a significant difference does exist.

First, time is considered as the explanatory variable, and the annual maximum river flow as the response. A map showing stations for which the null hypothesis of no trend is rejected can be seen in Figure 2-3. This map shows negative and positive trends in red and blue respectively, with significant trends denoted by a circle and non-significant trends by a cross. A total of 30 out of 117 stations exhibit significant trends (seen par-
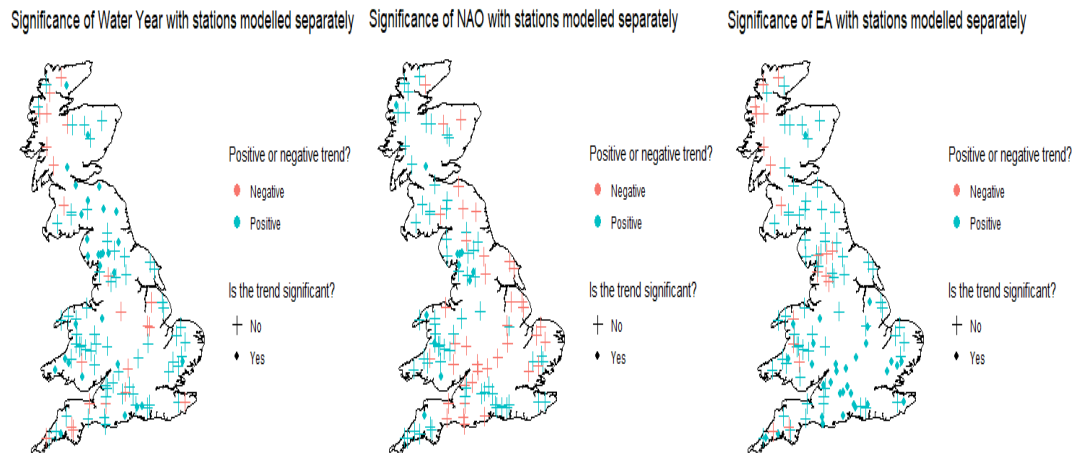
Figure 2-3: Significance of time and climate index trends for river flows in Great Britain.

ticularly in the north-west of England) - this represents of 25% of all stations. However, if all stations were independent, 5% of these stations could show a trend by chance. This large proportion of false positives produced when running multiple hypothesis tests is known as a multiple testing problem. Often, false discovery rate (FDR) controlling approaches are used to limit the number of these false positives. However, the need for using such methods to overcome this problem can disappear almost entirely when using a Bayesian multilevel approach. Gelman et al. (2012) proposes such an approach in scenarios when multiple comparisons occur. The authors note that classical inference techniques only use information from a particular site to get effect estimates at that site (thus ignoring key information from other sites), and tend to keep point

estimates fixed. Multiple comparisons are adjusted for by increasing interval width. On the other hand, the authors point towards multilevel models, which employ partial pooling of information, ensuring that each site's estimate will get pulled towards the overall estimate. This has the consequence of making multilevel model estimates more conservative, as is appropriate as the resulting intervals are more likely to include zero, and are more likely to be valid. Thus, the use of Bayesian multilevel models in the sections that follow ensures that there is no need to be concerned about the multiple testing problem unlike for the at-site case.

We repeat the at-site approach with NAO as the explanatory variable instead of time. In this case, very few stations seem to have a significant NAO trend - only 12 out of 117 stations (10%) - however, these are clustered in a similar manner to those stations with a time trend, further suggesting that confounding may be an issue with these two variables. Finally, for the EA index, a total of 32 stations display significant trends, accounting for only 27% of all stations in the dataset. Given the proximity of significant and non-significant gauging stations it seems this approach may not be fit for the purpose of detecting long-term trends in peak river flows. In comparison, the approach discussed in Section 2.2 and implemented below demonstrates a clear ability to detect trends on a countrywide level.

2.4.2   Univariate models

Posterior distribution plots for the fixed effect parameters of models A, B and C can be seen in Figures 2-4 and 2-5. These plots show the posterior density of the parameter of interest conditional on the data observed, along with a credible interval representing the uncertainty about the given parameter. Here, what one is interested in is where the large proportion of the density lies – the size of the y axis itself is not of interest. The median of this distribution is indicated by the thick vertical line on each plot, and the 95% credible interval (i.e. given the data and the model, there is a 95% chance the true values of the parameters lie in that interval) is represented by the shaded region of the plots. On the x-axis is the range of values that the posterior distribution can take. This represents the change in the log of the annual maximum flows given the covariate modelled.

2.4.2.1 Time trends in peak flow data

Firstly, a model investigating the relationship between peak river flows and time is fitted. A clear time trend is observed in Figure 2-4 upon inspection of the posterior distribution. This plot shows the range of the 95% most credible values that the regression coefficient for water year can take based on the given data. Using this posterior distribution plot, it can be seen that the fixed effect of water year is highly likely to be greater than 0.6 because the median value (indicated by the thick vertical line) and most of the mass of the distribution lies to the right of this value. As this is on a log scale, this corresponds to time adding at least 6% to peak flows each year suggesting that peak flows have been increasing considerably over time. This is of particular interest as it demonstrates the enhanced ability of the multilevel approach to detect signals that have previously been missed using an at-site approach.

However, time itself cannot cause changes in peak river flows. Instead it is acting as a proxy for some other unknown variables which vary with peak river flows in the same way as time, for example global warming. In itself, it does not provide information on potential causes of changes in peak river flows, and thus is not the only area of interest when it comes to the attribution of such changes. Instead, we focus on relating these changes to large-scale climate indices, which themselves represent changes in climate.
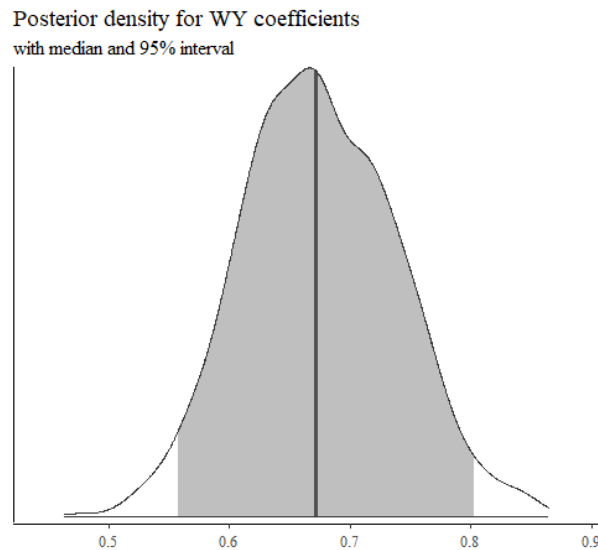
Figure 2-4: The fixed effect posterior for Model A. This shows the posterior density of the fixed effect of water year conditional on the data observed, with a 95% credible interval representing uncertainty about this parameter (shaded region). The median is indicated by the thick vertical line. The x-axis shows the range of values that the density can take.

2.4.2.2 The relationship between climate indices and peak river flows

We now focus on the attribution of changes in peak river flows to climate indices. In Model B, the effect of the Eastern Atlantic (EA) index alone is investigated. From looking at the posterior plot shown in Figure 2-5, this appears to have a strong association with peak river flows. The posterior density of the regression parameter is clearly centred away from zero as over half of its mass lies above 0.10, and does not contain zero in its credible interval. This corresponds to a 10% increase in the median peak river flows when going from a null EA value to a positive anomaly of size 1. This suggests that the EA index has some positive association with peak river flows in Great Britain.
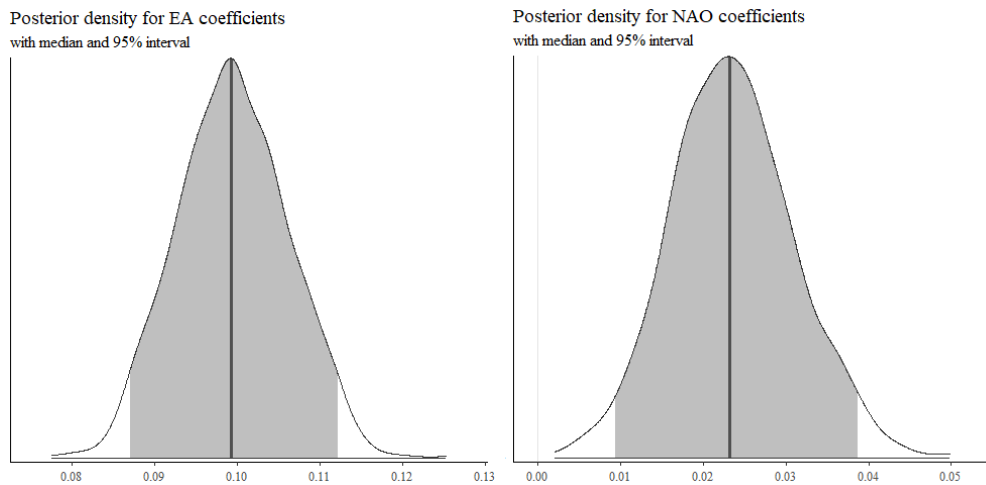
Figure 2-5: The fixed effect posteriors for Models B (left) and C (right).

When fitting Model C (see Figure 2-5), it can be seen that the NAO seems to have some association with peak river flows. The posterior distribution also has slight overlaps with zero, although it does not contain zero in its 95% credible interval. One might expect the NAO to have some association with these annual maximum flows, roughly a 2% increase in the median when going from a null NAO value to a positive anomaly of size 1.

2.4.3 A multivariate approach

In Section 2.4.2 we showed that both time and climate indices are related to peak river flows in Great Britain. It is clear from the plots in Figure 2-2, however, that the

climate indices are both correlated with time, suggesting the possibility of confounding between these variables. There is a need to separate out the net effect of these indices when time has been taken into account. It is necessary to use a multivariate approach to accurately determine the scale of the association between these indices and peak river flows after time has been taken into account.

It is not unreasonable to believe that the effect of time may be preventing us from seeing the true effect of climate indices such as EA and NAO on the peak annual river flows or that, vice-versa, the detected effect of time is actually the result of the effect of climate indices which happen to also change in time. Note that there may be other unobserved variables that change with time, that also drive change, as time itself cannot drive change — it is however a good proxy for these variables. We will demonstrate that this effect is masking the true extent of the associations between climate indices by utilising a multivariate approach.

To overcome this potential confounding effect of time, the peak river flows are modelled as a function of both time and climate indices together. Modelling more than one covariate at a time is key to identifying collinearity between variables and will provide a clearer picture of what is driving peak river flows in Great Britain. We noted that time trends have primarily been the focus in past approaches, but isolating the effect of climate indices to attribute these trends has largely been overlooked thus far (Merz et al., 2012). Combinations of time and these climate indices are now considered to observe whether those associations seen in Models B and C remain when time has been included in the model.

Model D investigates the link between both time and EA, and the log-transformed peak river flows. Even with time taken into account, EA is clearly associated with peak river flows. This can be seen in the plot on the right of Figure 2-6 — the posterior distribution still lies away from zero, with a median value of approximately 0.09. This plot represents the effect of the change in EA for a fixed point in time. We see that, even when time has been taken into account, EA still shows a 9% increase in the median peak flows when going from neutral to positive EA. This is crucial as it suggests that there is a clear association between the East Atlantic index and the annual maximum river flows in Great Britain. Note also that the size of the association of time is reduced when EA is added to the model. This again suggests the presence of confounding between these variables. On the other hand, the plot for Model E (Figure 2-7) suggests that the NAO no longer has any relation to the annual maximum flows when the time effect is taken into account — it can be seen that the posterior now has considerable overlap with zero, in contrast with Figure 2-5, and in fact the value becomes negative, suggesting there is
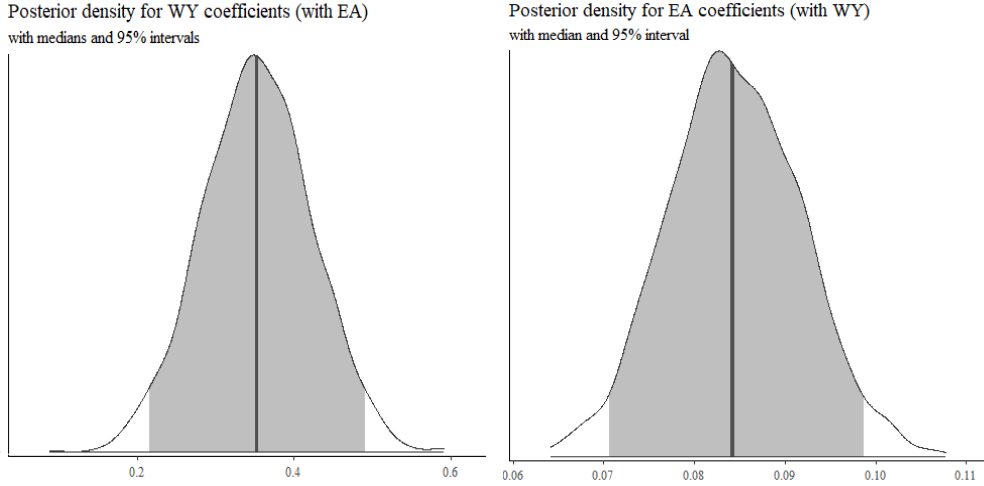
Figure 2-6: The fixed effect posteriors for Model D - Water Year (left) and EA (right).

some collinearity between NAO and Water Year. This lack of a clear association may also suggest that the NAO is not a key driver of change. However, we do not rule out this possibility, as if the NAO changes linearly with time in a very close manner it is possible that it is still a driver of peak flows in Great Britain. Finally, modelling the
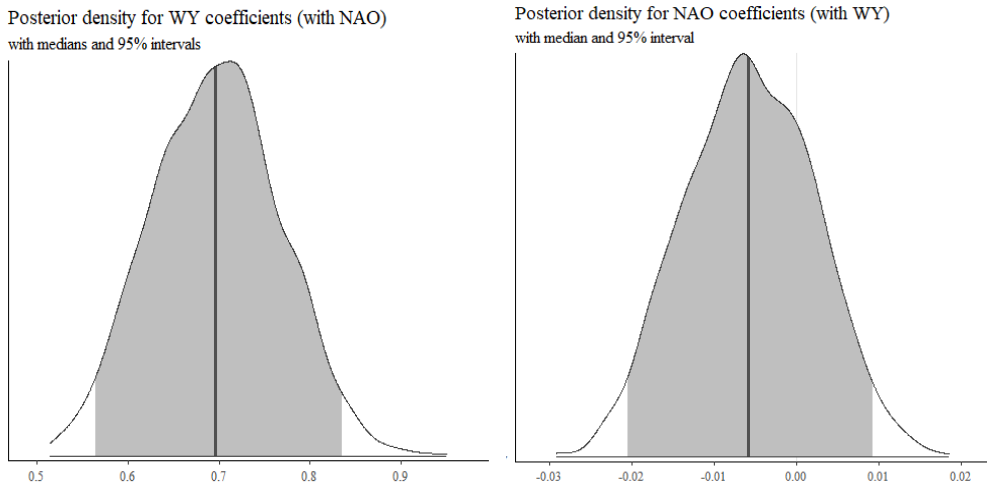
Figure 2-7: The fixed effect posteriors for Model E - Water Year (left) and NAO (right).

combined effect of time, EA and NAO on peak river flows (Figure 2-8) shows the same results – both time and EA appear to have an association with peak flows, while NAO again seems to have little to no association. These results, in particular the change in apparent relationship between NAO and peak river flows when accounting for time, suggest that it is necessary to use a multivariate approach to accurately estimate the

size of associations between climate indices and peak river flows in Great Britain.
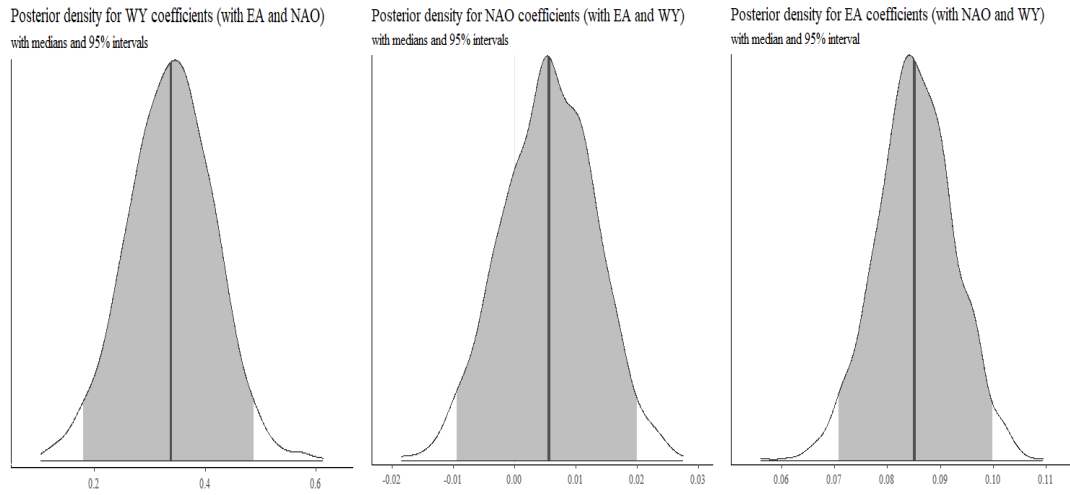
Figure 2-8: Fixed effect posteriors for Model F - Water Year (left), NAO (middle) and EA (right).

2.4.4 Spatial trends

We investigate whether any of these explanatory variables have some regional trends, in order to see whether there is any unexplained variance remaining that displays a spatial pattern. We plot the spatial random effect (adding the i.i.d. station effect r_i to the station-level spatial effect s_i) to check for regional behaviour: ideally, we would see positive and negative values randomly scattered across the country. As an example, we show the residual effects for model D (Figure 2-9): the residuals for the other models do not differ significantly. The scale of these effects can be considered approximately
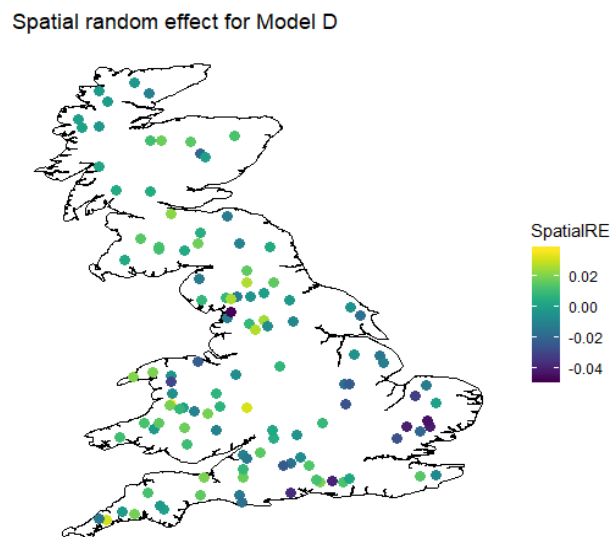
Figure 2-9: The mean of the posterior spatial random effect for Model D.

as percentage differences as we have taken the log of the annual maximum flows in the model. This means that an effect of 0.05 corresponds to a 5% difference in the median peak annual flow. There appears to be some difference in the size of the spatial random effect in the east and south-east of England compared with the rest of the country. This may be due to the fact that catchments in this region have high BFIHOST (soil permeability) values (see Figure 2-1), and were included in the reference network as compromise catchments to ensure full spatial coverage. However, adding this variable to the model did not make the spatial residual structure disappear. Further investigation suggests that the size of the EA index association is at its highest in this part of Great Britain.

2.5   Discussion and conclusions

In this paper we have presented a study investigating the dependence of extreme river flows upon climate indices for benchmark catchments of Great Britain. We have demonstrated that it is possible to use multilevel models to detect a countrywide trend at sites with short records by pooling information from nearby catchments, and that by using these more complex models, clear associations between these trends and some climate indices of interest are found. The use of near-natural "benchmark" catchments in our approach ensured that any signal found in the data cannot be related to anthropogenic changes other than climate. A clear countrywide time trend was detected, in contrast with the scattered signal seen in the at-site analysis (see Section 2.4.1). This demonstrates the value of a pooled approach for future analyses of trends, to get a more accurate picture. Note that this approach relies on the assumption of constant countrywide effects of the climate indices and on the linearity of these effects. These assumptions nevertheless do not appear to be too stringent for a first investigation and have been verified by checking the model residuals. The use of non-parametric regression models, as done for example in Villarini et al. (2009), to describe the relationship between peak flow and the explanatory variable might be a viable option to relax the strong linearity assumption.

We also investigated the effect of climate indices on peak river flows in our approach, and noted a clear signal for both the EA index in the univariate case, with an increase of 10% in the median when going from a neutral to positive EA value. The signal was still strong even when time was added to the model, suggesting that the EA index may have an impact on peak river flows even when confounding is accounted for. This strong association is a key step towards the accurate attribution of trends in peak river flows motivated by Merz et al. (2012). In particular, it fulfils the "soft" attribution criteria, along with two of the three criteria for "hard" attribution - detected changes in peak river flows appear to be consistent with the East Atlantic index, and an uncertainty interval has been provided with the Bayesian approach. It remains to check whether detected changes are inconsistent with potential drivers other than NAO, which was also investigated in this paper. However, this is not yet evidence of a causal link between the two, and given that it is impossible to control the value of climate indices, ensuring a full attribution is not trivial. It would be of interest to explore a causal framework for this model, to ensure that this association between the EA index and peak flows is indeed a causal relationship.

There can be concerns over the use of non-stationary models in hydrological studies (see

for example Koutsoyiannis and Montanari (2015)), specifically that the model structure in these cases can lead to an additional form of uncertainty when making predictions. However, if one can attribute changes in flows to the EA index, then predictions of how flows will change with this index in the future can be made. Additionally, aside from prediction, a key contribution of this paper is the method itself, which demonstrates the value of using a multilevel and multivariate approach to give a clearer insight into the countrywide drivers of peak river flows. To gain a full understanding of how these flows will change in the future, however, the approach should not only be based upon time series data, but requires additional information. Some element of this can be captured within prior information provided in a Bayesian analysis such as this, but additional non-data assumptions must be made in order to infer causal relationships, as suggested by Merz et al. (2014). The authors note the importance of statistical methods in hydrological studies, but that they must be complemented with investigations of causal relationships and key drivers of changes within this system. We believe that such an approach would be of benefit when exploring a causal framework for the method discussed in this paper.

Finally, in the naive univariate approach, there appeared to be a clear link between NAO and peak river flows in Great Britain, with an increase of 2% in the median of peak flows when going from neutral to positive NAO. However, including time in the model to address possible confounding between variables leads to this association going to zero or even becoming slightly negative, suggesting collinearity between variables. This demonstrates the necessity of a multivariate approach for accurately quantifying the strength of association between climate indices and peak flows. A combined multilevel, multivariate approach towards attribution helps to provide a clearer insight into changes in peak river flows in Great Britain, and the strength of the relationship between these flows and climate indices.

2.6   Chapter conclusions

As Brady et al. (2019) has shown, it is of benefit to use a Bayesian multilevel framework when it comes to the accurate detection and attribution of trends in peak river flows in Great Britain. By increasing the sample size available to the model in comparison with at-site methods typically used, it was possible to detect clear countrywide time trends that had been missed with the at-site approach. This result reconciles previous discrepancies between climate change projections of increased flooding and observational data, where no clear trends had previously been found. Note that, while it was found that

the assumption of a linear relationship between annual maximum river flows and both climate indices appeared reasonable, it may be of interest to further explore non-linear relationships between these variables to rule out any different potential structure. This linearity assumption is explored further in Chapter 3.

This paper explored a model which varied in space, but not in time. This was not seen as necessary, as using annual maximum flows means there is little to no temporal correlation. However, a space-time interaction term could be included to capture any residual spatio-temporal variation that may not be accounted for by the spatial random effect. This would be a simple extension of the model that may produce more reliable model estimates.

The approach in this chapter satisfied Merz's "soft" attribution criteria (Merz et al., 2012), along with two of the three "hard" attribution criteria. However, it was noted that further evidence must be provided for a causal link between climate indices such as the East Atlantic index and peak river flows in Great Britain. In environmental observational studies, where natural variability is a major factor, it is not trivial to apply standard methods to ensure that Merz's criteria are met in full. In Chapter 3, we propose a systematic checklist for assessing causality in environmental observational studies such as these, which fulfils and goes beyond the scope of these criteria.

As the set of locations at which rivers are measured is fixed (and the spatial domain is not continuous), we encoded a distance-based spatial effect through a Gaussian process. However, this does not exploit the network structure of an individual river network nor account for the fact that Euclidean distance may not be the most appropriate measure for describing river distance. Additionally, it is also of interest to look at a more regional level of modelling for local flood defence planning. The use of Bayesian methods described in this chapter may be useful for exploring long-term trends at a regional level, which again may be missed out in an at-site approach. In the setting of a single river, the spatial structure may be encoded through a conditional autoregressive (CAR) model, where one makes use of the graph structure by dividing the area into regions and defining neighbours to this region. The use of such structures gives rise to faster inference techniques, and will be discussed in Chapter 4.

Chapter 3

A systematic checklist for causal assessment of environmental observational studies

This chapter consists of a prepared and revised manuscript. In Chapter 2, we demonstrated a preliminary approach towards attributing long-term changes in peak river flows to a large-scale driver of interest. However, this approach was not a formal attribution, as it only demonstrated associations between climate indices and peak river flows in Great Britain. Further, it also did not satisfy all of the required criteria as set out by Merz et al. (2012) for "hard" attribution. In particular, while it was shown that detected changes in peak river flows were associated with the East Atlantic index with a provided uncertainty interval, it must be demonstrated that detected changes are consistent with further potential drivers than NAO . In this chapter, not only is it shown how this criterion can be fulfilled, but a thorough systematic checklist for assessing causality in environmental observational studies such as these is proposed, which goes further beyond the scope of Merz's attribution criteria.

The major contribution of this chapter is the development of a general checklist which can be used for the causal assessment of long-term, large-scale environmental studies.

Abstract

Natural environmental systems are highly variable, and this can have significant impact on safety, wealth and well-being. To make accurate predictions and informed decisions under different conditions, it is vital to understand and accurately attribute the mechanisms driving long-term, large-scale environmental change. However, many standard methods for establishing causality are usually either not appropriate or easily applicable to environmental observational studies. We review potential approaches for assessing evidence of a causal relationship in such problems, including instrumental variable methods, causal diagrams, methods of multiple working hypotheses, and the Bradford Hill criteria. We assess the fitness for purpose of these methods for this type of study and select a number of criteria appropriate to these problems. We also propose further steps to be taken within the analysis to address environment-specific issues, incorporating them into a framework which assesses for causal relationships through an audit of strength of evidence for a broad class of problems. We demonstrate this method through a case study of peak river flows in Great Britain.

3.1 Introduction

Understanding the causal mechanisms driving long-term, large-scale environmental problems such as air pollution, flooding, a rise in sea surface temperature, and tree defoliation is vital for choosing effective interventions and ensuring adequate defences are in place in the long-term. Norris et al. (2011) notes the challenges in determining causal relationships in such natural systems, given the presence of confounding, natural variability, and an inability to replicate studies. In addition, data records are often limited in size and scope, or suffer from further complexities which can be avoided in the gold-standard randomised control trial. For example, it is not usually possible to apply exposures randomly to observational units in an environmental setting (Norton and Suter, 2014). Randomisation of exposures removes confounding by other causes of outcome, and thus enables identification of the effect of exposures. In environmental observational studies, many variables may covary, often making it near-impossible to identify the causal role any of these will have upon the response. These variables can be one of several causes, a product of the cause, a precursor to the cause or in fact have no causal role whatsoever (and just happen to covary with the true cause in space or in time, not actually causing a change in the outcome). The same exposure in different sites in environmental observational data may not lead to the same or similar response — this is due to population heterogeneity. Though the exposure level appears to be the same, the activities that can impact specific ecosystems can be highly variable and modify its effect, thus an exposure's effect may differ considerably in practice, particularly given differing geographical characteristics. Such variables which impact the size of the exposure's effect, known in causal inference as modifiers, are not affected themselves by either the exposure or the outcome, but are associated with a different effect of the exposure on the outcome (Höfler, 2005).

Temporal confounding is a common difficulty encountered in observational environmental studies. It is possible that the apparent effect of time may be preventing us from seeing the true effect of the variables of interest on a given outcome. Time itself does not cause changes in an outcome, but acts as a surrogate for other unknown variables which vary with the exposure in the same way as time. It is a catch-all for complex processes such as global warming, which has been increasing monotonically in recent decades. Auffhammer and Vincent (2012) noted that unobserved effects of time may confound identification of climate changes, and that controlling for such effects appropriately allows for the true effect of an exposure to be seen. When time was not accounted for in their model considering the effect of weather-driven crop yields on

emigration, the effect of these yields was significant. However, with time included in the approach, no significant effect could be seen. This exemplifies that it is crucial to account for time when assessing causal relationships in environmental problems. Spatial confounding is another common issue in an environmental setting; Augustin et al. (2007) noted that a common problem of spatial survey data (in that case, forest survey data) is that there is often confounding between the spatial effect and covariates, particularly those covariates which are spatially correlated.

It is clear that environmental problems cannot always be assessed for causal relationships using standard causal tools — instead these tools should be adapted for this application. Though there are techniques such as propensity score matching which allow for the construction of counterfactual outcomes in observational studies, these are not always fit for the purpose of assessing for causality in environmental studies. In such situations, it can be difficult to have well-defined counterfactual outcomes. Instead, Cartwright (1989) suggested through the "No causes in, no causes out" principle that there must be some additional non-data assumptions made when it comes to a thorough understanding of causes. Pearl (2001) notes that "behind every causal conclusion there must lie some causal assumption that is not testable in observational studies". Additionally, Greenland et al. (1999) points out a common flaw of typical statistical models is that they are unable to capture all assumptions necessary for an accurate analysis. Pearce and Lawlor (2016) further highlight this point, noting that, while it is rarely acknowledged in practice, the process of attribution to a particular cause in epidemiology usually involves the consideration of a wide variety of evidence.

We believe that it is necessary to include non-data assumptions when assessing for a causal relationship in observational environmental studies, as this will add further weight to the strength of evidence provided from the data. We will discuss potential approaches that may be of benefit for assessing the strength of evidence for a causal relationship in environmental problems. We propose a systematic causal checklist for these studies, making use of relevant ideas in these approaches and incorporating both data and non-data assumptions into our method. We require that not only must the strength of association between the exposure and response be clear, but that the model is as free from sources of confounding as possible, is interpretable to decision-making bodies and explains as much variability as possible.

This approach will be of benefit to studies investigating long-term, large-scale environmental change problems such as causes of drought, the impact of climate change on sea surface temperatures, river flows, air pollution, water quality and forest tree defoliation amongst several others, where causality has rarely been explicitly discussed. We believe

70

that this will provide a contribution to a field that could develop in a similar way to epidemiology, where approaches to causality are constantly evolving and up for debate — see, for example, the special issue on causality in epidemiology (Davey-Smith and Ebrahim, 2016).

First, a selection of potential approaches to establishing evidence of a causal relationship in environmental statistics are discussed. These include the Bradford Hill criteria, causal diagrams, instrumental variable methods and an approach based on multiple working hypotheses. In the following section, a causal checklist for the assessment of environmental observational studies is proposed. We then illustrate our proposed checklist by means of a case study assessing the causal relationship between climate indices and peak river flows in Great Britain using the criteria proposed. Finally, these findings are summarised and the applicability of this systematic checklist to the wider field of environmental statistics is discussed.

3.2 A review of selected potential causal approaches

In a large majority of observational environmental studies, it is not possible to undertake a controlled experiment to estimate the effect of an exposure. Randomisation is infeasible in these studies, and thus methods from the gold-standard randomised control trial are not applicable. Instead, other approaches must be considered in order to gain an understanding of causal mechanisms driving environmental change. Here we discuss some methods used across econometrics, epidemiology and ecology which may be of benefit to environmental problems. Note that this is not an exhaustive review, but a set of methods that may have merit in an environmental setting and which help to form a basis for the checklist presented in the next section. While not all of these are causal methods in themselves, they comprise a selection of useful tools and approaches that may be beneficial to the causal analysis of environmental problems.

3.2.1 Instrumental variable methods

Instrumental variables (IVs, (Reiersøl, 1945)) allow for the controlling of confounding variables and measurement error in observational studies. They allow for the possibility of making causal inferences with observational data. They can adjust for both observed and unobserved confounding effects, thus providing a natural experiment. The instrumental variables approach is to find some variable that influences which exposure is

received, but that is independent of unmeasured confounders. These variables should only indirectly affect the outcome through the exposure. This variable can then be used to extract variation in exposure that is not subject to unmeasured confounders. This confounder-free variation will subsequently be used to estimate the causal effect of exposure.

Instrumental variable methods have been used to assess the impact of climate change, for example through station-level rainfall and temperature on agriculture as seen in Di Falco et al. (2012), which uses village-specific characteristics as instrumental variables. This is a useful approach, but cannot always be used effectively. The choice of instrumental variable is difficult in environmental observational problems. exposure randomisation here is not controlled by the experimenter, but rather induced by natural variability. The search for a suitable instrument may be difficult or even impossible depending on the problem at hand. As a consequence, this approach is difficult to apply directly to environmental observational studies, however the idea of controlling for an unmeasured variable will be used within our approach through a "partialling out" approach to regression.

3.2.2 Causal diagrams

A causal diagram, as described by Greenland et al. (1999), enables the visualisation of causal relationships between variables in a causal model. Greenland et al. (1999) notes that causal diagrams provide a useful aid to assessing causal relationships between variables. The authors discussed how these methods were formalised and could be used in an epidemiological setting, to complement the models used for analysis. In addition, one must start with a set of assumptions in order to make any inductions of causal relationships. Constructing this set of assumptions and the corresponding diagram will form the initial part of assessment within our approach in the next section.

Causal diagrams typically consist of a set of variables (called *nodes*) defined as being within the scope of the model being represented. A line or arrow connecting two variables is called an *edge*. When there are directed edges representing causal relations between variables and no cycles connecting the other edges, it becomes a directed acyclic graph (DAG). A variable A is said to be an *ancestor* (cause) of another variable B if there exists a directed path of arrows leading from A to B, and B is considered a *descendant* of A.

In observational studies, one variable is defined as the primary exposure, and another as the primary outcome. Expert knowledge then defines the directed edges from which

the status of other variables can be determined — these variables could be confounders, mediators or colliders, amongst others. The presence of any of these may impact any unadjusted association between the exposure and the outcome. Care must be taken to correctly identify variables, as including them within a regression analysis will change estimates of effects, either correctly or incorrectly. There are other variables such as modifiers and instrumental variables, but these will not be considered within the examples that follow.

1. **Confounder:** A confounder is a variable that is an ancestor of both the exposure and the outcome (along a path which doesn't include the exposure). For example, A is a confounder of the effect of B on C in the DAG to the left of Figure 3-1.

2. **Mediator:** A variable is known as a mediator when it is a descendant of the exposure and an ancestor of the outcome. One can choose to control for the mediator or not depending on the objective. A is a mediator in the DAG in the centre of Figure 3-1.

3. **Collider:** A variable is known as a collider when it is causally influenced by two or more variables — one that is (or is associated with) the exposure, and one that is (or is associated with) the outcome. Variables influencing this collider may not necessarily be associated. It is not appropriate to control for colliders as this can induce an association between variables B and C. A is a collider in the DAG to the right of Figure 3-1.

Figure 3-1: Examples of a DAG where variable A is a confounder (left), a mediator (centre) and a collider (right).

These causal diagrams are of benefit for environmental observational studies, when often there is clear confounding (for example between time and the exposure of interest, where time is a proxy representing other unknown variables which change with the exposure in the same way), along with mediating effects and presence of collider variables. In addition, there is often unmeasured confounding which may have some potential causal effect upon variables in this system. This approach can help to identify whether potential unmeasured variables are confounders, colliders or mediators, and how best to account for these within the modelling approach. Causal diagrams will be

used to formalise causal dependencies within an environmental observational setting, and help to determine which variables to adjust for or otherwise within the checklist introduced in the next section.

3.2.3 Method of multiple working hypotheses

The method of multiple working hypotheses (MMWH) was proposed in Chamberlin (1890). It is a method which involves the development of several hypotheses that might explain a phenomenon being observed, while recognising that more than one hypothesis may be simultaneously valid (thus also increasing the likelihood of discovering interactions between hypotheses). In essence, this method amounts to being rigorous and diligent when it comes to considering alternatives to a proposed cause of change in outcome.

Harrigan et al. (2013) demonstrated the use of this approach in hydrological studies, suggesting that it may be appropriate for wider use in environmental problems. The authors note that working with multiple simultaneous hypotheses in mind provides a more systematic approach to attribution to those which consider single causes of change in outcome only, and which may result in confirmation bias. Using this approach, the authors formulated several possible hypotheses for a set of detected changepoints. These were based upon several potential causes of change both in isolation and as combinations. Additionally, one such working hypothesis involved the acceptance that there may be unknown or unmeasured factors which could be contributing to or counteracting change. These were all assessed against past literature and expert knowledge before testing. Hypotheses were compared against criteria such as consistency prior to carrying out hypothesis tests. The authors noted that the application of this method can result in valuable insights, allowing for a systematic approach to considering multiple causes of climate change.

The method of multiple working hypotheses allows for the development of new hypotheses and highlighting weaknesses in current understanding, while allowing for the possibility of unknown causes. This approach was also used by Clark et al. (2011), where the authors noted that hypothesis testing in catchment-scale hydrology required the isolation and linking of several model hypotheses. They noted the practical concerns, namely trade-offs between model flexibility, complexity, comprehensiveness, and computational cost, which must be taken into account should one use this approach. The method's key features of diligence and rigour when considering alternative causes will be incorporated within the proposed checklist for assessing causal relationships in

an environmental observational setting.

Elliott and Brook (2007) provide an updated perspective of this method, noting that the intentions behind the approach can often be misinterpreted and incorrectly used in practice. The authors note that the true aim of Chamberlin's method was to avoid becoming attached to any single governing hypothesis, and to avoid arriving at premature explanations of observed phenomena. In addition, they discuss how one can consider hypotheses in parallel or in series as appropriate to the situation at hand. Each hypothesis is considered in turn in the case study for the checklist that follows, however this can easily be adapted to consider hypotheses in parallel should it be appropriate. Finally, Elliott and Brook (2007) discuss how various modelling and model selection techniques fit with the MMWH approach. For example, in contrast with the p-value, the authors note that Bayesian methods have the benefit of removing a reliance upon falsification of competing hypotheses, and allow for the incorporation of uncertainty in modelling and prior knowledge. While the checklist does not prescribe a particular choice of modelling, a Bayesian approach is utilised in the case study to illustrate this checklist.

3.2.4 Bradford Hill Criteria

The Bradford Hill criteria are a set of 9 principles developed in 1965 by Sir Austin Bradford Hill (Hill, 1965). They are most often used when attempting to establish epidemiological evidence of a causal relationship between some posited cause and an observed effect. They consist of the following (as paraphrased from the original):

1. **Strength (effect size):** While a small association does not imply that there is no causal effect, the larger an association is the more likely that it is causal.

2. **Consistency (reproducibility):** If results are found to be consistent across different observers, locations or samples, this strengthens the likelihood of an effect.

3. **Specificity:** If there is a very specific population at a specific site and disease without another likely explanation, then causation is said to be likely. The more specific the association, the greater the probability of a causal relationship.

4. **Temporality:** The cause must precede the effect.

5. **Biological gradient:** Greater exposure to a variable should, in general, result in

a greater incidence of the effect. In some cases, the presence of this variable alone may lead to the effect. In other cases, an inverse relationship may be observed.

6. **Plausibility:** There should be some theoretical basis for proposing an association between cause and effect.

7. **Coherence:** The proposed association should be compatible with existing knowledge, i.e. claims must be evaluated within the context of the current state of knowledge.

8. **Experiment:** Sometimes it may be possible to appeal to experimental evidence.

9. **Analogy:** The effect of similar factors may be considered.

Hill noted that while alone, none of these criteria provide indisputable evidence in favour of, or against, a cause and effect hypothesis, rather they instead help one to address if the exposure of interest provides the most likely way of explaining any evidence present, or if there is "any other answer equally, or more, likely than cause and effect?" (Hill, 1965). Thygesen et al. (2005) argued that these criteria should not be seen as rules, but weighted differently on a per-scenario basis, an approach we will make use of in the proposed checklist. Fedak et al. (2015) noted that these criteria were not meant as a set of rules, but as a flexible framework to "guide epidemiologic investigations and aid in causal inference". Additionally, they also noted that how each criterion is applied, interpreted and weighted must be measured against the types of data available in the given situation. It is the spirit of this approach that we believe must be maintained when assessing environmental observational studies for cause and effect. These criteria are all useful for a variety of problems, but some are more likely to be important (and evidence may be found) than others for environmental problems.

We believe that a systematic approach which assesses the strength of evidence in favour of, or against, a causal relationship between two variables is the most appropriate for such environmental problems. The Bradford Hill criteria in their current form are not entirely applicable to environmental observational studies, however those relevant criteria will be taken into consideration within the causal checklist introduced in the next section to weigh up this strength of evidence for a broad class of environmental problems.

3.2.5 Additional approaches in ecology — the weight of evidence framework

One key component of the proposed approach is the ability to systematically consider and rule out causes with diligence and rigour. Such a systematic framework has been proposed for ecological data by Norton and Suter (2014), which involves a number of steps. The causal pathway is broken in segments for analysis and alternative hypotheses are compared. The effects of individual variables are isolated, then multiple variables are modelled together. This is then supplemented with further evidence, and one must finally consider how to combine this evidence to draw appropriate conclusions.

As discussed in Suter et al. (2017), the U.S. Environmental Protection Agency developed a general framework for ecological studies. In this framework, firstly the evidence for a particular hypothesis is assembled. This evidence is weighted according to three properties — relevance, strength and reliability. Relevance can include biological, physical/chemical, and environmental relevance, the latter of which looks for "correspondence between test conditions and conditions at the assessed site" (Suter et al., 2017). Though these are specific to ecological studies, similar comparisons between mathematical/statistical model output and observational study data will be explored in the checklist proposed later. Strength suggests that a stronger signal should be given more weight than a weaker signal as it shows a higher degree of association (often using correlation coefficients) between a potential cause and the outcome of interest. This is similar to the strength criterion of Bradford Hill, and one which we include as an item of the causal checklist later. Finally, reliability comprises a number of properties that make the evidence more compelling, such as standardisation of methods, specificity, corroboration across multiple studies, minimised confounding and consilience (i.e. coherence with scientific theory). A number of these factors are considered within the checklist for environmental studies.

The weight of the body of evidence as a whole is then determined based upon the weights of these parts (using a scoring table, see Suter et al. (2017)) and interpreted to see whether any of the hypotheses is supported by the weight of evidence. Considering the evidence across a number of key criteria and ruling out causes in turn as a consequence is a vital component of the proposed method.

Finally, this approach further suggests sensible consideration of the quality of the data and the information used to choose the model. This is common to all analyses, rather than specific to causal analyses, and is assumed to be undertaken alongside the proposed checklist.

3.2.6 Potential alternative approaches which did not influence the systematic checklist

Below we briefly outline some approaches which are often used within a variety of observational studies, and justify why we do not consider them further within the systematic checklist proposed.

One such method is the Granger causality test, a statistical hypothesis test which determines whether one time series is useful in forecasting another (Granger, 1969). Granger argued that causality in economics could be tested for by measuring the ability to predict the future values of a time series by the use of prior values of another time series. Granger defined the causality relationship based on two principles — that the cause happens prior to its effect, and that the cause has unique information about the future values of its effect. One time series X at time t is said to Granger-cause another time series Y at time $t + 1$ if, with $\mathcal{I}(t)$ being all available knowledge up until time t:

$$\mathbb{P}[Y(t+1) \in A \mid \mathcal{I}(t)] \neq \mathbb{P}[Y(t+1) \in A \mid \mathcal{I}_{-X}(t)],$$

where A is an arbitrary non-empty set, $\mathcal{I}_{-X}(t)$ denotes all information up to time t, excluding X. That is, X_t contains some information about Y_{t+1} that is not part of the rest of the set $\mathcal{I}(t)$. This is usually tested with regressions to determine how informative the lagged values of X are about Y. Granger causality, despite its name, relates to linear prediction rather than true causation. It only captures whether a variable X precedes another variable Y, and if X can help to predict Y (Diebold, 2001). In addition, the Granger-causality tests are designed to handle pairs of variables, and may produce misleading results when the true relationship involves three or more variables. This lack of reliability when it comes to the multivariate case further ensures that it is unsuitable for environmental problems, where we may be trying to assess many combinations of covariates.

Structural equation models (SEM) are a collection of mathematical techniques, such as factor analysis, regression and path analysis, which evaluate complex hypotheses, often causal in nature, against data (Bollen, 1989; Grace, 2008). These techniques have become more popular in ecological studies in recent years (for example Belkhiri and Narany (2015); Grace (2008); Riseng et al. (2011)). There are two key model components to structural equation models — the structural model, which encodes potential causal relationships between variables, and the measurement model, which shows the relationships between latent variables and those variables believed to indicate them. There can be "free" pathways, i.e. those which hypothesise causal relationships to be

tested and "fixed" pathways where the relationship between variables is already known. A particular advantage of this method is that any latent variable can be a dependent variable in one defined set of relationships, while it can be an independent variable in another (Belkhiri and Narany, 2015). SEM is a confirmatory model technique — the proposed model is evaluated to confirm whether the data supports the proposed structure or not. This structure may be rejected if it does not match the data structure. A potential disadvantage of this method is interpretability (Jeon, 2015) — as a variety of techniques underpin this method and may be used in a single model, results can often be misinterpreted due to a poor understanding of the techniques. However, the predominant reason for not considering the SEM approach is that it is designed for a more complex causal DAG than the one considered in this approach, i.e. one which does not fit into the linear modelling framework.

In environmental studies, we are also interested in assessing the causal effect of exposures as a function of variables that may impact this exposure. These variables may vary with time. Structural nested models were proposed in Robins (1986) to model and estimate the joint effect of a (time) sequence of exposures. One may wish to estimate the joint effect of this sequence of exposures in the presence of some confounding variable C, which has 3 characteristics (Vansteelandt et al., 2014); that it is independently associated with outcome Y, either as it is a direct cause of the outcome, or it shares the same unmeasured cause, that it predicts subsequent exposure levels A_1, and it is affected by earlier exposure A_0. When one is interested in estimating the joint effect of a sequence of exposures, standard approaches such as regressing Y on A_0 and A_1, or a function of both, may not be appropriate, regardless of whether conditions on the confounder are adjusted for or not. If a time-varying confounder is affected by the prior exposure, then standard approaches to control for confounding may not be appropriate, as the covariate can act as both a confounder and a mediator of the effect of the exposure on the outcome (Williamson and Ravani, 2017).

Instead of using a regression-based approach such as structural nested models (that is, a model for the outcome conditional on the covariates), one could use an approach involving marginal structural models (a model for the counterfactual outcomes). Marginal structural models are a class of causal models for the estimation of the causal effect of a time-dependent exposure in the presence of time-dependent covariates that may be both confounders and mediators (Robins et al., 2000). This paper notes that parameters may be consistently estimated using a special class of estimators, known as the inverse probability of treatment weighted estimators. The marginal structural model is for the counterfactual outcomes, not the observed outcomes, and the expectation is

marginalized over the entire population (not conditional on the covariates). However, it is not always practical to build counterfactual outcomes for use in environmental observational studies. Krieger and Davey-Smith (2016) note that a pillar of such approaches is that the ability to quantify causal effects relies on the ability to pose counterfactuals involving exposures that could be randomised, though this has been refuted for epidemiological studies (Daniel et al., 2016). The potential outcomes approach is particularly useful when one wishes to consider the treatment assignment mechanism, which is useful in such studies where one can impact this mechanism. This is generally not possible however in environmental studies, where natural variability plays a large role. Randomisation in environmental observation studies is induced by this natural variability, and the studies cannot be easily manipulated to construct counterfactuals. In a spatial setting, in theory one could make use of location characteristics to construct counterfactuals, however in many cases (for example, in peak river flows, see later discussion) these characteristics do not provide much, if any, explanatory power. It is also not necessarily simple in this framework to consider potential associations with additional variables and hence confounding. On the other hand, the use of a graphical approach such as DAGs helps to visualise the overall causal structure, including such potential confounders and additional complex relationships between variables, that might not be considered in the potential outcomes approach. As a consequence, we believe it is of most benefit to frame the problem in terms of DAGs instead of the potential outcomes approach in the method presented, where we propose a systematic workflow for assessing causal relationships in environmental studies.

3.3 Causal assessment of environmental observational studies

There are considerable challenges in determining causal relationships in natural systems, given the presence of confounding, interference issues, natural variability, an inability to replicate studies and data records of limited length. In this section, we propose a checklist for assessing causal relationships in environmental studies, which takes into account these issues by providing a means for incorporating both data and non-data assumptions into the overall assessment. We demonstrate the necessity for these non-data assumptions as motivated by Pearl (2001), as they add further weight to the strength of evidence provided from the data. These assumptions must be clearly stated and backed up (for example, using expert opinion). We provide a framework for assessing these relationships within a practical checklist that is applicable to a broad

class of environmental problems. This is framed in the context of causal DAGs for simplicity of interpretation and the ease of visualising complex relationships between variables in the system. Within this approach, we incorporate those Bradford Hill criteria (within items 1, 4 and 5) that we believe are appropriate for environmental problems and justify these choices. In addition, we propose a number of additional checks (items 0, 2, 3, 6 and 7) that are necessary for accurately establishing causality in these problems.

3.3.1 A systematic checklist for assessing strength of evidence

0. **Write down a set of potential causal diagrams.** The directed acyclic graph (DAG) approach is chosen as a suitable framework for this class of problems, as it provides a useful framework for visualising and interpreting the often complex relationship between the exposure and outcome in the presence of confounding variables. The use of DAGs will be of benefit when it comes to considering additional variables which may influence the environmental system at hand.

   All variables which will be used in the analysis should be set out from the beginning, where possible. It is assumed that there is sensible consideration regarding the quality of the data and the information behind the choice of model as seen in the weight of evidence method seen earlier, as the DAG will only be useful if the variables used are reliably measured. A set of suitable causal diagrams, which can be assessed using available data, must be written down for the data generating process from the outset in order to be clear on the mechanism being proposed. Some idea of the direction of the relationship (positive or negative) should be informed from past results or physical models — if the effect found is inconsistent upon analysing the data, this may suggest the possibility of confounding variables to be considered. The causal diagram should be kept in mind throughout the process of conducting a statistical analysis for these environmental studies, which often assess the impact of a large-scale cause of change upon some outcome over some (long-term) period of time. An example of this is the three-variable diagram seen in Figure 3-2, where E represents the large-scale exposure, O the outcome and T time which usually acts as a confounder. In practice, as E is usually a time-varying exposure, the DAG will have separate nodes for E at each time point, but we represent it here as a three-node diagram for simplicity. Note also that this diagram may represent a more complex causal chain, however the most immediate cause is the one of interest in this approach. One cause will be

examined at a time, while other possible causes will be treated as confounders.

It is important to note that time itself is not the cause of any change, but represents other unknown variables which vary in the same way (i.e. which lead to the same observed change in response). In other studies, time can be replaced by another confounder, for example, space. It is vital to consider whether the

Figure 3-2: A potential DAG for large-scale environmental studies, where E is the (time-varying) exposure, O the outcome and T time.

proposed direction of the edges exist in practice. If it can be shown that if there is an edge from T to E and from T to O (for example, by inspection of time series plots for these variables), then an association between E and O should exist when one controls for T, assuming all of the explanatory power is not captured by T. Controlling for time will be discussed in item 3 of this checklist. Controlling for space will be explored in the discussion.

1. **Is this mechanism plausible and backed up by expert opinion?** This is an adaptation of the Bradford Hill criterion of plausibility. There must be good reason to consider that a variable may have an effect on the outcome under consideration, and this association should be compatible with physics and existing knowledge. This should be backed up by a thorough review of the relevant literature.

2. **Have best efforts been made to ensure the approach is interpretable to the end-user?** There is a need to balance model complexity with model interpretability. Many environmental models inform decision-making processes and must be transparent for this purpose. Simpler models are, in general, preferred, where simple means that the model is both parsimonious and interpretable. Checks must be carried out to ensure that one has chosen an appropriate model for assessing causal relationships by trying many combinations of variables believed to have an effect on the outcome of interest, but also that this model can be sensibly interpreted by decision-making bodies. In addition, model outputs should be intepretable to the end user.

3. **Has known (time-varying) confounding been accounted for?** In envi-

ronmental problems, natural systems are usually slow moving across time with some fast moving variations on top of those (for example climate and weather, sea levels and tide), and variables driving changes in an outcome tend to be highly correlated with time. In addition, the monitoring length for the outcome may not happen on the same time scale as the changes in the exposure variable. For example, climate indices tend to show seasonal and decadal behaviour, while peak river flows are an annual measurement built from 15-minute measurements. It is necessary to account for potential confounding between time and exposure variables in a manner appropriate for the problem at hand. Time itself does not cause a change in an outcome of interest, but can be used as a catch-all in environmental studies, when often these processes are not necessarily known (or are a combination of processes) but are long-term and slowly varying. In itself, time does not provide information on potential causes of changes in an outcome of interest but can help control for unknown variables which do. A "partialling out" approach to removing the effect of time from the exposure and the outcome will be used in the case study. The result of this is that the estimate now represents the true relationship between the exposure and outcome, without the presence of confounders represented by time. Care must be taken with such an approach. Regressing out time is meant to be a catch-all for which time is a proxy (e.g. climate change). It is a good, if not perfect, catch-all for a number of unknown variables in environmental studies, such as global warming. Including time instead will mask anything collinear with time, while excluding time will stop us from seeing the true effect of the variable of interest. There is a trade-off between wanting to unmask the effect of a variable that is unrelated to time, but the risk is that we may be removing some portion of this that can be attributed to climate change.

There are potential causal variables which are collinear with time, and without regressing time out, they will not be distinguishable from time. They may be causal, but we cannot reliably detect them in these situations. For those variables that do vary from time, we will be able to draw causal conclusions.

Spatial confounding is also of concern in these problems. Augustin et al. (2009) noted that ignoring either spatial or temporal correlations in tree defoliation data can lead to biased estimates in modelling. Spatial confounding will be explored in further detail in the discussion.

4. **Is there clear evidence of association between exposure and outcome?**
   This is inspired by the Bradford Hill criterion on strength of evidence. Strong

associations are more likely to be causal, though weak associations in practice are common and cannot rule out a causal relationship by themselves. A clear signal between the exposure and outcomes is sought, after known confounding has been controlled for. However, caution should be taken when making inferences about small area effects — weakness of evidence should not be taken as non-causality. In practice, the strength of the effect can be observed with uncertainty expressed in a frequentist setting through a confidence interval, or in a Bayesian setting by inspection of the posterior distribution. A clear statement on the size and sign of the effect and the associated uncertainty is required. If the sign of the effect does not match with intuition or physical models, the possibility of confounding variables will need to be considered.

5. **Do the model results match with physical model outputs?** If the results are consistent with experiments or large-scale mathematical/physical simulations, then there is a higher chance of a causal relationship, as noted in the Bradford Hill criteria and the weight of evidence approach seen in Suter et al. (2017). It is not always practical to run lab-based experiments under more controlled conditions, so climate or physical model outputs can be taken as the proxy for experiment in these problems. The sign of the effect from the previous item should be consistent with these results.

6. **Has the approach been diligent in considering alternative causes?** Thorough checks must be carried out to ensure that there is no likely alternative explanation for the observed data, and that there has been a rigorous search to ensure the approach has considered the impact of unmeasured confounders. This is partly an adaptation of the multiple working hypotheses approach as proposed by Chamberlin (1890), and the minimisation of confounding as discussed in Suter et al. (2017). In practice, this means that every effort must be made to consider all potential scenarios which may cause some change in outcome — though noting that it is impossible to consider every possible scenario. For example, in assessing changes in river flows, one might want to consider river management or other human factors. Artificial scenarios should be constructed to ensure that there is nothing critical remaining that resembles some structure present in the data. This can include a residual analysis, which checks whether any structure remaining can be seen in the residuals, or even multiple structures. One can also explore including covariates, which impact the outcome only (i.e. they do not impact the effect like a covariate would), in order to further reduce model variability. This is of benefit to environmental studies, where explanatory power is often poor. It

is important to identify whether the variable is a confounder, mediator or covariate, and decide whether to include them in the model as appropriate. Additional variables should be considered in collaboration with an expert in the field.

7. **Revisit the original causal diagram — is this still a sensible mechanism?**
One must return to the causal diagram in item 0 to consider potential scenarios for further unmeasured hypothetical variables not considered in item 6.

Take for example the specific case of time as a confounding variable and suppose we introduce a fourth variable Q into the system, which we have not measured but may have some possible causal effect on the other variables. We consider two possibilities; that Q is (almost) collinear with T or that Q is (almost) orthogonal to T. A mixture of the two is possible but the individual components can be treated separately. There may be many unmeasured variables, but these can be similarly decomposed and reassembled.

When Q is collinear with time, as time itself cannot be a cause, then Q takes the place of T in the DAG in item 0, but causal directions are not necessarily the same. In this scenario, four possibilities exist:

- Q is a confounder. The DAG is the same as that in item 0. Thus, if a directed edge exists from Q to E, and from Q to O, then once Q is controlled for, a directed edge (and thus a causal relationship) exists from the exposure to the outcome.

- Q is a mediator. In this case, one can control for Q depending on their objective. Either way, the exposure causes the outcome, either directly or indirectly.

- Q is a collider. One should not control for Q as this will induce an association between the exposure and the outcome. As Q is not measured in practice, this is not an issue. Note that this relationship still has a directed edge between the exposure and outcome without Q.

- Q is a mediator but in reverse, i.e. the outcome variable might cause the mediator in this case. However, in this case, we assume the outcome cannot cause the exposure (for example, flooding cannot cause a change in the North Atlantic Oscillation climate index), so this is not possible.

These possibilities can be seen in Figure 3-3. Note that no matter the DAG, when Q is collinear with T, the same conclusion is reached — we do not need to know Q explicitly to control for it or otherwise. As a consequence, once Q is appropriately controlled for, we can conclude that there is a causal relationship

between the exposure and outcome.

The other option is that Q is orthogonal to T. One possible DAG can be seen in Figure 3-3. Here Q enters as a confounder. T does not cause Q as they are

Figure 3-3: Examples of a DAG where **(a)** Q is a confounder, **(b)** Q is a mediator, **(c)** Q is a collider, **(d)** Q is a reverse mediator, or **(e)** Q is orthogonal to T.

orthogonal. Q can potentially confound the causal effect of E on O, whether we adjust for T or not. However, for Q to exist, a variable must be found that has components orthogonal to time and has causal effects on both the exposure and the outcome. Possibilities for Q should be considered and ruled out on a case-by-case basis to finally conclude on whether the exposure of interest can be considered a cause of the outcome or not. There are other four-node DAGs which are possible, but if Q is a mediator or collider rather than a confounder,

we do not need to control for them in a causal analysis such as this (as discussed previously).

3.4 Case study: attribution of trends in peak river flows in Great Britain

In this section, we present a qualitative case for a causal relationship between climate indices and the annual maximum of peak river flows in Great Britain, following the causal workflow proposed earlier. In this case study, we consider the East Atlantic (EA) or North Atlantic Oscillation (NAO) index as a potential cause of the process of interest, and annual maximum river flows at a large number of river gauging stations as the outcome.

3.4.1 Data and model used

The North Atlantic Oscillation (NAO) index is a mode of natural climate variability, representing surface pressure anomalies in the North Atlantic region and surrounding continents, particularly Europe (National Climatic Data Center, 2017). Positive phases of the NAO are associated with above average temperatures and precipitation across northern Europe. The East Atlantic (EA) index (National Climatic Data Center, 2017) is a mode of low-frequency variability over the North Atlantic, consisting of a north-south dipole of anomaly centres spanning the North Atlantic from east to west. It is similar in structure to the North Atlantic Oscillation (NAO) index. Positive phases of the EA index are associated with above-average surface temperatures in Europe and above average precipitation over northern Europe. Climate data was obtained from the National Oceanic and Atmospheric Administration website (National Climatic Data Center, 2017).

These indices are impacted by climate change in a more direct manner than precipitation (which we might consider as a more obvious influence of peak river flows), so they are proxies for climate which is changing but are also quite variable to begin with (Guimarães Nobre et al., 2017). Additionally, they can be modelled more accurately than precipitation, a complex process which mostly exhibits low variation but with intermittent peaks. There is also a lack of evidence for the widespread claims that increased precipitation leads to an increase in flooding, as noted by Sharma et al. (2018), who point out that in some cases there are even reduced flooding magnitudes with

increased precipitation. Further, it is of interest to explore the large-scale, longer term drivers of flooding, not the shorter term and local scale choice of precipitation. River flows themselves also exhibit complex behaviour, however the annual maximum values are themselves relatively simple to model by taking the log and assuming these values are normally distributed. This has been found to be a reasonable assumption for UK annual maximum flow data (Prosdocimi et al., 2014) and is further discussed in Appendix A. We investigate whether changes in peak river flow can be attributed to some of these indices. A key problem when assessing for causal relationships in this setting is that values for climate indices are constant across all river gauging stations in Great Britain for a given year, i.e. the exposures vary only in time. As each station receives the same exposure in one year and this exposure is not manipulable, the counterfactual outcomes framework is not easily applicable. Matching river gauging stations based upon catchment characteristics is a potential approach, however it will be shown later that these characteristics do not provide much explanatory power. As a consequence, such counterfactual-based methods for causality are not considered in this setting.

We focus on annual maximum river flow data from Great Britain across the series of reference benchmark catchments introduced by Harrigan et al. (2017). The catchments were chosen as they have relatively long records of good hydrometric quality, are representative of Great Britain's hydrology and are relatively near-natural. The benefit of using this reference network is that catchments were chosen specifically to overcome the difficulties in accurately attributing climate-driven trends. They ensure that any signal found in the data cannot be related to any other anthropogenic changes, thus partially controlling for unmeasured confounding. This data set contains the largest observed instantaneous peak flows in each water year (which runs from October to September), measured in m^3/s. In total, there are 5475 observations from 117 benchmark gauging stations in the UK, with an average record length of 46 years. This data can be obtained from the UK National River Flow Archive (NRFA) (Dixon, 2010), which is the primary UK source of hydrometric data. The NRFA gathers and quality controls such hydrometric data from all UK gauging stations. Records for EA and NAO are available only from 1950, thus we only use records with these values available.

The DAG proposed (as per item 0 of the checklist) can be seen in Figure 3-4. The model used is the Bayesian multilevel model approach for attributing peak river flows as presented in Brady et al. (2019). This assumes that a country-wide trend is present in the flow data and employs a multilevel approach to allow different station-specific effects which are expressed as random effects r_i. Nearby stations can be expected to be impacted in a similar way by external climatological variables, thus we include a

spatial correlation structure s within the multilevel model. For station i at time t, we have:

$$\log(\text{Flow})_{it} = \alpha + \boldsymbol{X}_t \boldsymbol{\beta} + r_i + s_i + \varepsilon_{it}, \qquad (3.4.1)$$

where $\boldsymbol{X}$ is a matrix of explanatory variables, $\epsilon_{it} \sim N(0, \sigma_i^2)$ is the measurement error, r_i is a random effect to allow for variation between stations with $r \sim N(0, \tau_0)$ and s_i is a spatial random effect ($s \sim \text{MVN}(0, \Sigma)$) to allow for correlation between nearby stations (where Σ is taken to be of an exponential correlation structure).

3.4.2 Write down a set of potential causal diagrams.

We propose the simple causal diagram in Figure 3-4, where C represents a large-scale climate index (EA or NAO, where we will treat each cause in turn as noted previously), F the outcome (peak flows or flooding) and T time (which acts as a surrogate for other unknown variables, such as global warming or ocean temperature). We know that neither C nor F can cause T, and that F cannot cause C. Thus, the only potential DAG for this choice of variables is the one shown in Figure 3-4. In this case C and F represent the climate index values and annual maximum flows in a given year respectively. In practice, the DAG will have separate nodes for C and F at each time point, but we represent it as a three-node diagram for simplicity.

Figure 3-4: A DAG for attribution of changes in peak river flows where C represents a climate index, F is flooding/peak river flows and T time.

3.4.3 Is this mechanism plausible and backed up by expert opinion?

Firstly, motivation must be provided for considering climate indices as exposure variables. Reasoning must be provided for their inclusion, which is compatible with existing theory and for which evidence can be found in the existing literature. In this section, we consider the evidence in the literature linking these indices to peak river flows in Great Britain.

Previous studies have shown that the NAO explains much of the variability in river

runoff, motivating its selection for use in our approach (Hannaford, 2015; Hannaford and Marsh, 2006; Kingston et al., 2006). Guimarães Nobre et al. (2017) noted that positive (negative) phases of both the NAO and EA are associated with more (less) frequent and intense seasonal extreme rainfall over large areas of Europe. We might expect an increase in precipitation to result in increased flooding. In the UK, there have been limited studies investigating the impact of climate indices on precipitation and flooding. Huntingford et al. (2014) proposed a number of potential drivers for the record precipitation that caused flooding in southern England in the winter of 2013/14 including NAO, while Van Oldenborgh et al. (2015) noted that the strongest relationship with precipitation in southern England was with NAO (though this had only a small correlation of 0.23).

3.4.4 Have best efforts been made to ensure the approach is interpretable to the end-user?

Models attributing trends in peak river flows should be transparent as these models can inform long-term flood defence strategies. We seek a simple relationship between flows and some key large-scale covariates that explains a suitable proportion of the variability in the high flow data and is still interpretable to these stakeholders, instead of a "perfect" and detailed model such as a rainfall runoff model which is complex both in model structure and computational power. Trenberth (1999) notes that highly variable rates of rain and spatial variability not only make estimating mean precipitation difficult, but also near impossible to predict how flows will change as the climate itself changes. As we hope to quantify the impact of climate change, we instead make use of climate indices, which are impacted by climate change in a more direct manner than precipitation (Guimarães Nobre et al., 2017), so they can be used as proxies for climate which is changing. Mahlstein et al. (2015) note that using climate indices may be more appropriate than simple climate means, as they are generally closer related to climate and potential impacts — and thus of more benefit to users. The use of these indices ensure that we obtain a more direct quantification of environmental impacts (Della-Marta et al., 2009).

Model outputs should also be intepretable to the end user. In this case, this requires a statement regarding how the size of the change in peak river flows varies with the climate index being investigated. Brady et al. (2018) and Brady et al. (2019) noted a clear positive association between EA and peak river flows in Ireland and Great Britain. This suggested that there should be a 10% increase in the median peak river flows when

going from a null EA value to a positive anomaly of size 1.

3.4.5 Has known (time-varying) confounding been accounted for, and is there clear evidence of association?

These two parts of the checklist are considered together, by accounting for confounding within the model and estimating the relationship to peak river flows after this has happened. The chosen exposures (i.e. the climate indices) are confounded with time (as can be seen in Figure 3-5), which itself represents the effect of some unknown covariates. The effect of these unmeasured confounders may be masking the true effect of climate indices on the peak annual river flows. It is of interest to know if positive or negative anomalies of these indices have a different impact upon peak flows, and, whether it would be possible to observe any such differences given the variable is likely to be confounded with time in the observation period available in the analysis. We demonstrate how controlling for time-varying confounders by regressing time out of the model impacts the size of the association between climate indices and peak annual river flows in Great Britain.

We first seek a relationship between peak river flows y_{it} (for location i and time t), and time (as a proxy for unknown confounders) and climate indices of the form (for example, for EA):

$$y_{it} = f(\text{Water Year}_t) + g(\text{EA}_t) + \varepsilon_i$$

where EA is also changing with time. $f(\cdot)$ and $g(\cdot)$ represent some smooth functions of time and EA respectively. Here we require that EA varies in a way that is not a simple smooth function of time. However, if the variables represented by time can vary rapidly, one may not easily be distinguished from the other. If the effect is somewhat smooth, it should be possible to separate out the effects of the climate index from the confounder(s). In order to proceed, it is necessary to assume that the effect of time is smooth — otherwise there may be some unknown covariate that varies with these climate indices in the same way. In practice, it also does not seem unreasonable to assume that the effect of time (as a proxy for unknown variables) is smooth. Time itself does not change anything alone, but is a proxy for other changes. These changes of interest are long-term and slowly varying. Such processes can be expected to be roughly monotone or at most include a hockey stick type change point (Mann et al., 1998), so assuming a linear relationship should be appropriate. Time series plots of the data can be examined to identify any possible trends. A time series plot of NAO can be seen in Figure 3-5, showing that there appears to be some slowly varying time

trend, and that NAO has shifted direction towards positive values. A time series plot for EA can also be seen in Figure 3-5. The linear trend over time is not so clear here, though there does seem to be an overall slight upward trend.

In order to control for this confounding in the model, we aim to find some parametric relationship between climate indices and time. Such relationships can be investigated using a generalised additive model (GAM). Using the approach of Wood (2012), p-values of the fixed effect smooth terms can be calculated to test whether these smooth terms differ significantly from 0. These smooths (along with their associated confidence intervals) can be seen as blue lines in Figure 3-5, while the linear regression line can be seen in red. Note that for EA, the relationship is essentially linear, and these two lines overlap entirely. For the NAO, these two do not overlap, though both fall within the confidence interval,. Following the approach of Wood (2012), is clear that these smooth
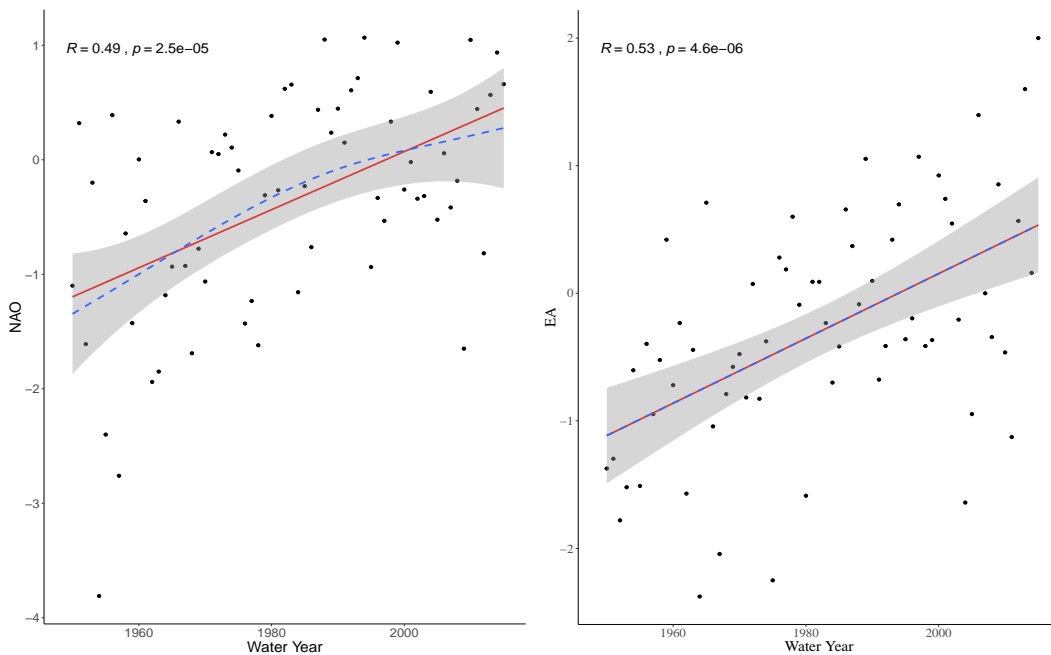
Figure 3-5: Time series plots for water year with NAO (left) and EA (right) with linear regression line in red & smooth line in blue, correlation coefficients R above and with a p-value for a significance test on the correlation being null.

terms are not significantly different from zero, with a p-value of 0.127 for the smooth term for NAO and 0.607 for EA. This suggests that a linear relationship between time and each of the covariates can be assumed. To remove this time-varying component, the effect of time will be regressed out both from the exposure (EA or NAO) and the response (the log of the annual maximum flows). What remains will represent the relationship between climate indices and peak river flows, once time is removed from

the equation.

3.4.5.1 Regressing out the effect of time — EA

First, a multi-level model for the log of standardised peak flows against EA is fitted, following the approach of Equation 3.4.1. The output of fitting a Bayesian model as per the approach in Brady et al. (2019) seen in the results that follow is a posterior density, which describes the uncertainty about the unknown parameter(s) having observed the data. This posterior density is shown in all results along with a credible interval representing the uncertainty about the fixed effect parameter of interest. The size of the y axis itself is not of interest, what is relevant is where the large proportion of the density lies. The median of this density is indicated by the thick vertical line on each plot, and the 95% credible interval (i.e. given the data and the model, there is a 95% chance that the true values of the parameters lie in that interval) is represented by the shaded region of the plots.

The posterior plot shown on the left in Figure 3-6 suggests that the magnitude of the fixed effect of EA seems to be considerable — the posterior does not contain zero in its 95% credible interval, and most of the mass lies above 0.09, suggesting (as we are working on the log scale) a change of 9% in the median peak annual flows going from a null value to a positive anomaly of size 1.

It is of interest to see if this effect size of EA is changed at all when the effect of time is removed. Under the assumption of the smooth time effect, it should be possible to "regress out" the effects of time. In this case, time is regressed out of the exposure (EA) in a simple linear model

$$\mathrm{EA}_t = \alpha + \beta \mathrm{Water\ Year}_t.$$

The residuals from this fit are then extracted and stored as EA_i^*. This is known as a "partialling out" approach as described in Wooldridge (2015), which removes the time-varying portion of the EA effect. The residuals EA_i^* represent the part of EA that is uncorrelated with time. The effect of time is also regressed out of the response (the log of peak annual flows) in the same way. Finally, the residuals of both of these models are taken, and one is regressed against the other through a Bayesian model as follows:

$$y_{it}^* = \alpha^* + \beta^* \mathrm{EA}_t^* + \varepsilon, \tag{3.4.2}$$

93

where y_i^* refers to the time-regressed out version of the flows, and EA_i^* the time-regressed out version of EA. The β^* measures the relationship between peak river flows y and EA once the time component has been removed from both. The posterior density for the parameter β^* can be seen on the right in Figure 3-6. It appears that EA has a

Figure 3-6: The fixed effect posteriors for EA before time is regressed out as in Equation 3.4.1 (β, left) and after time is regressed out as in Equation 3.4.2 (β^*, right). The area under the curve represents the posterior density of the fixed effect of EA conditional on the data observed, with a the shaded region showing the 95% credible interval representing uncertainty about this parameter. The median is indicated by the thick vertical line. The x-axis shows the range of values that the density can take.

clear and positive association with the log of the peak annual flows in Great Britain, even with the time-varying component regressed out. The size of the effect is reasonably large and positive. Looking at the posterior distribution plot shown in Figure 3-6, the relationship is clearly centred away from zero as over half of the mass lies above 0.08, and does not contain zero in its credible interval. This corresponds to an 8% increase in the median peak river flows when going from a null EA value to a positive anomaly of size 1 in any given year. Thus, it appears that the EA is a driver of changes in peak river flows in Great Britain, even when time has been taken into account.

3.4.5.2 Regressing out the effect of time - NAO

The same process is repeated for NAO, using the multilevel model of the log of the annual maximum flows against NAO as seen in Equation 3.4.1. Looking at the plot on the left of Figure 3-7, NAO appears to have some association with peak annual flows. The time-varying portion of the NAO and of the log of the standardised annual maximum flow is then removed. The multilevel model in Equation 3.4.2 is then run, in order to assess the fixed effect of NAO on annual maximum flows once time is removed. It appears that the fixed effect of NAO has disappeared or even become slightly negative when the time effect is regressed out. A plot of the fixed effect posterior for the β^* parameter in the time-regressed model can be seen on the right in Figure 3-7. This
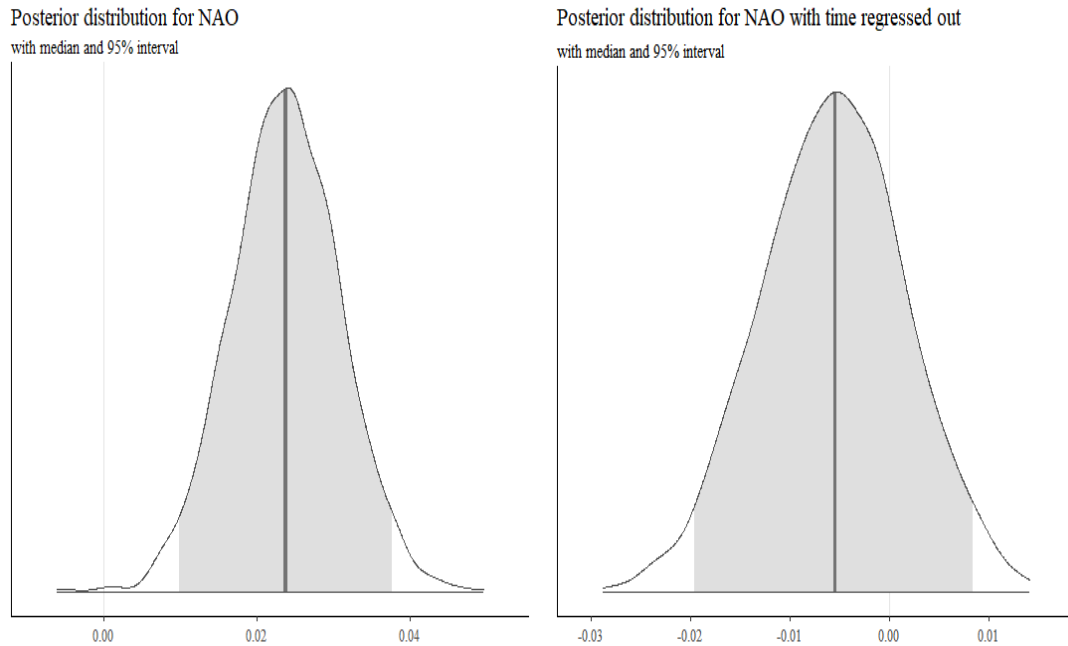
Figure 3-7: The fixed effect posterior for NAO before time is regressed out (β, left) and after time is regressed out (β^*, right). Again, the area under the curve represents the posterior density of the fixed effect of EA conditional on the data observed, with a the shaded region showing the 95% credible interval representing uncertainty about this parameter. The median is indicated by the thick vertical line. The x-axis shows the range of values that the density can take.

suggests that NAO may have little association with the annual maximum flows when the time effect has been removed. The posterior now has considerable overlap with zero in contrast with the left plot of Figure 3-7. The value has become negative, suggesting there is some collinearity between NAO and time. As a consequence, is not possible to

show that there is a causal effect of NAO upon annual maximum flows. It may suggest that the NAO may not be a key cause of change in peak river flows across Great Britain once time is taken into account, but there is also a possibility that the effect of NAO is masked by the effect of some unknown variable represented by time.

3.4.6 Do the model results match with physical model outputs?

In this case study, "experiment" refers both to mathematical climate models which can reproduce large-scale atmospheric patterns and hydrological model outputs. These mathematical models are a proxy for experiments where such an experiment is impractical. Global climate models simulate global and regional climate systems. Their outputs are downscaled to regional level, which then become inputs for hydrological models (Chiew et al., 2010). If model results are consistent with these physical model outputs, then there is a higher chance of a causal relationship.

Climate change projections suggest an increase in precipitation (Bates, 2009), with an expectation that this will lead to an increase in flows in the UK. Climate change is expected to further change UK catchment hydrological responses in the future (Dawson et al., 2016; Watts et al., 2015). Observed changes in climate in recent years have been linked with changing precipitation patterns, which in turn can be linked to changing patterns of the EA and NAO . It is expected that there will be an increase in the magnitude of floods under these climate change projections (Fowler and Wilby, 2010).

Results from the Bayesian multilevel model discussed previously noted a positive relationship between peak river flows and both time and EA (which is a mode of climate variability, and thus expected to change in line with climate change projections), suggesting consistency with these physical model outputs.

3.4.7 Has the approach been diligent in considering alternative causes?

In this case study, this item involves consideration of further possible scenarios that might either cause flooding, or further explain some of the variability missed in the climate index model — these can be confounding variables, mediators or covariates. Note that while the variables considered in the following section were selected as the most likely to contribute to the variability peak river flows in Great Britain, there may be many further options which are not considered.

3.4.7.1 Additional known variables — catchment descriptors

Additional hypotheses must also be considered in order to discover whether any important variables have been left out — these may be confounders, mediators or covariates, and it is vital to treat these carefully and only include such variables within the model where appropriate.

A number of catchment properties are known to have an impact on river flows, and an additional analysis is carried out to investigate whether catchments of different types are found to be impacted differently. Within the NRFA dataset, a number of catchments descriptors are given to describe properties of the catchment; we focus on the variables which Kjeldsen et al. (2008) found to be associated to median flow, as these are often found to be key elements in explaining river flow variability. In particular, these are the catchment drainage area in km^2, the catchment average annual rainfall in the standard period (1961-1990) in millimetres (SAAR) and the base flow index (i.e. a measure of catchment permeability and responsiveness) derived from the catchment soil types (BFIHOST). Transformations for each of these are recommended in Kjeldsen et al. (2008) — the square of BFIHOST, 1000/SAAR and the natural log of area. Similar transformations of these descriptors are included as additional covariates in the time-regressed model of peak river flows against EA as in Equation 3.4.1. Catchment rainfall may be a mediator, as bigger catchments can respond differently to catchment rainfall levels. We would not include a mediator in the model, however a covariate which notably improves the explained variability would be worth including. The other variables, which only impact the outcome, can be considered as covariates, and could be included in the model should they increase explanatory power.

We look first at the relationship between flows and catchment drainage area using the time-regressed model, i.e. the model given by

$$y_{it}^* = \log(\text{Flow})_{it}^* = \alpha^* + \beta^* \boldsymbol{EA}_t^* + \gamma \text{Area}_i + r_i + s_i + \varepsilon_{it} \tag{3.4.3}$$

The posterior distribution for this can be seen in Figure 3-8. This credible interval has slight overlap with zero, and does not appear to explain much of the variability in the model. We repeat this process for the catchment's baseflow index, taking the square of these values, and replacing area with BFIHOST in model 3.4.3. The posterior distribution for the fixed effect of BFIHOST on peak river flows can be seen in Figure 3-8. Here, the credible interval again has overlap with zero, and in fact appears to be slightly negative. It does not seem to explain much of the variability in the model through the fixed effect.
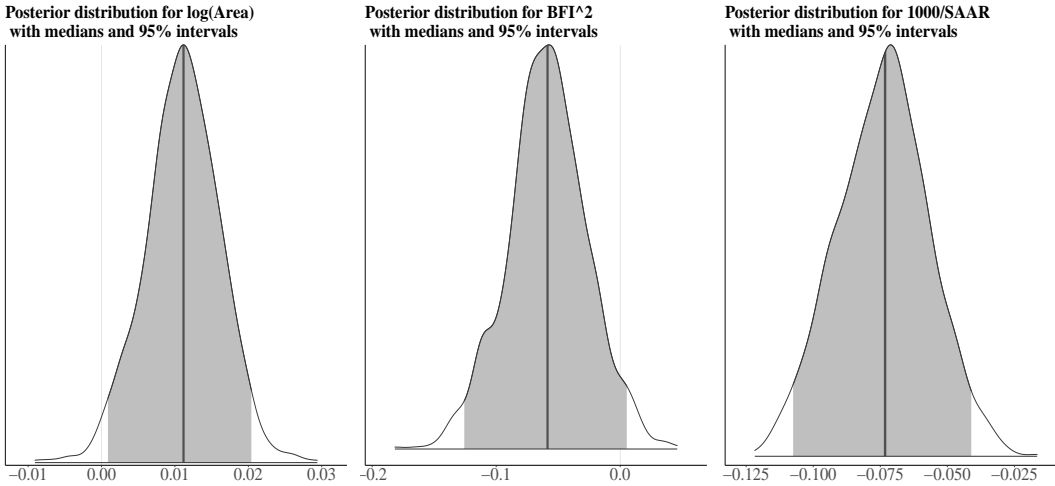
Figure 3-8: The fixed effect for log(Area), BFIHOST^2 and log(SAAR) (left, middle and right) when modelled with EA (time regressed out). Again, the area under the curve represents the posterior density of the fixed effect of EA conditional on the data observed, with a the shaded region showing the 95% credible interval representing uncertainty about this parameter. The median is indicated by the thick vertical line. The x-axis shows the range of values that the density can take.

We also investigate whether the catchment average annual rainfall has any relationship with peak river flows. The posterior distribution for this can be seen in Figure 3-8. Note that the fixed effect of this covariate is highly likely to be greater than 0.05 because the median value (indicated by the thick vertical line) and most of the mass of the distribution lies to the right of this value. Thus, the catchment average annual rainfall appears to have some relationship with peak river flows in Great Britain, i.e. wetter catchments tend to experience larger peak flow values, and could be included in a model for peak river flows. However, as it is likely to be a mediator, it is not appropriate to include the catchment rainfall in this situation. These results are all comparable to the coefficients for these variables estimated in Kjeldsen et al. (2008). Note, however, that the inclusion of these additional covariates does not impact upon the estimates of the magnitude and direction of the EA index's effect upon peak river flows. This is due to the fact that they are covariates as opposed to confounders, however it can be useful to include them in a model to further explain the variability in the outcome.

3.4.7.2 Residual analysis

Another common way to check for unexplained variability and potential unmeasured confounding is to look at the residuals of a regression — i.e. to inspect what is left

98

over after explaining the variation in the dependent variable using independent variable(s). The aim of this is to see whether an additional predictor could be constructed that resembles the structure of these residuals, and whether there are any other model assumption violations (such as heteroscedasticity). The plot on the left of Figure 3-9 shows the posterior yearly average residuals against time, for the model of peak river flows against EA with time regressed out. These residuals appear to vary randomly

Figure 3-9: A time series plot of model posterior mean annual residuals with standard error (left) and a time series plot of number of sunspots per year (right)

.

around the mean, in particular from the 1970s onwards, where data records are available for the majority of sites and thus the residuals are much less variable. This appears to be a smooth curve, which can be checked using the method of Marra and Wood (2012) discussed in in the case study. This results in a p-value for the smooth term of 0.271. It is again clear that the smooth terms are not significant, suggesting a linear relationship can be assumed.

It is possible, however, that a predictor whose behaviour resembles that of the residuals once time is regressed out could be constructed. A climate index other than NAO or EA is one possible option, given the apparent multi-decadal behaviour of the residuals. However, this approach aims to quantify the impact of climate variability upon peak river flows. Having this predictor be an additional climate index, by definition a proxy for climate which is changing, would add further to the strength of evidence

that climate change is indeed causing changes to these flows. As an illustration, one alternative predictor that might resemble this structure is the number of sunspots per year. Macdonald and Sangster (2017) noted strong correlations between flood-rich phases and solar magnetic activity. It is possible that solar activity impacts rainfall and thus we use it to showcase the exploration of further alternatives. Such potential confounders should, in practice, be identified with an expert, however we explore this possibility here to highlight the approach used. A plot of a time series of numbers of annual sunspots may be seen on the right of Figure 3-9. This is not entirely dissimilar in structure to the residual plot to the left of Figure 3-9 and expert opinion should be sought to verify this further. However, this is likely to be unrelated to anthropogenic climate change (i.e. it cannot be impacted by humans) and regressing the number of sunspots against the model mean annual residuals does not result in a significant effect, thus this potential confounder does not require further consideration within our model. It seems unlikely that other predictors could be constructed that both display the behaviour seen in the residual plot and would interfere with the ability to detect the true impact of climate change upon peak river flows in Great Britain. This will be further discussed in the context of the causal diagram later in this section.

3.4.8 Revisit the original causal diagram — is this still a sensible mechanism?

We follow item 7 of the systematic checklist and discuss further possibilities for Q as a potential unmeasured confounding variable in the case of attributing changes in peak river flows to climate indices. One such possibility is global warming as seen in Figure

Figure 3-10: A causal diagram for peak river flows in Great Britain including unmeasured variable G (representing global warming) which is orthogonal to T.

3-10. This will cause changes in climate indices and perhaps peak river flows; however it is increasing with time (see Pachauri et al. (2014)) and occurs at an earlier point in the causal diagram — i.e. the effect of global warming on flooding is mediated

by climate indices. Regressing time out of the model ensures that any confounder such as this which is collinear with time is accounted for. Another possibility for a confounder is changes in land use. This can lead to changes in peak river flows but is unlikely to cause changes in climate index values. Additionally, the use of benchmark catchments in this case study controls for land use changes, and thus this possibility can be ruled out. Ocean temperatures may have some causal effect on climate indices and potentially also on changes in peak river flows. This is a potential candidate for an unknown confounder, but it is again increasing with time (Pachauri et al., 2014) and thus can be accounted for in the same way as global warming. Rainfall is a mediator, so it is not necessary to control for this covariate.

Other than these examples, it is difficult to identify further potential candidates which may have a confounding effect upon this system, suggesting that the original causal diagram is appropriate for the problem at hand and that it is possible to say that EA has some causal effect upon increasing peak river flows in Great Britain.

3.4.9   Summary of case study

This case study demonstrates the use of the systematic checklist proposed to attribute long-term changes in peak river flows in Great Britain to some large-scale driver of interest. This approach started with a causal diagram between large-scale climate indices (EA and NAO) and peak flows, with time as a proxy for unmeasured confounding. The choice of climate indices as a plausible, coherent and interpretable choice for potential causal variables was justified. The case study demonstrated the usefulness of regressing time out of the model in order to gain a better understanding of the residual effects of climate indices on peak river flows once time has been taken into account.

In particular, before time is regressed out of the model, the NAO index appears to have some small impact upon peak river flows. However, once time has been taken into account, this association disappears completely. The collinearity between NAO and time means that it is not possible to show that there is a causal link between NAO and peak flows — it may be causal, however we cannot arrive at this conclusion due to its collinearity with time. On the other hand, the EA has a clear association with peak flows both before and after time has been regressed out of the model, suggesting it has a causal impact upon peak flows. This is clear to observe as the EA varies with time. The case study also exemplifies the complexity of finding potential additional candidate confounding variables and that it is likely that the EA has a causal relationship with peak river flows in Great Britain.

3.5 Discussion

There are many difficulties when assessing for causal relationships in environmental studies including confounding, limited data records and natural variability. Most standard methods for establishing causality are not easily applicable. There is often no control with which to compare exposed units, meaning that typical counterfactual-based methods for causality cannot be directly applied. In addition, there is a need to consider non-data assumptions which help to build up the strength of evidence for a particular relationship — the data alone is not sufficient for such studies. This is best summarised by Cartwright (1989)'s "No causes in, no causes out" approach, which in its essence states that it is impossible to learn anything of a causal structure using only statistical correlations. Instead, additional knowledge must be incorporated in order to make causal statements.

We have proposed a systematic method for assessing causality in environmental observational studies, which incorporates both data and non-data assumptions. This approach provides guidelines on how to conduct a thorough analysis to ensure all possible avenues are explored through a methodical checklist, which brings together multiple workflows by bridging the gap from climate scientists to the relevant authorities. The method helps to improve accuracy and interpretability for decision-making bodies, as well as being thorough in ruling out additional causes and controlling for confounding where possible. We also claim that it is necessary to sense check the proposed models with causal diagrams throughout the process, in order to consider as wide a range of possibilities as is feasible.

We applied this method to a case study of peak river flows in Great Britain, where the aim was to attribute changes in these flows to climate change (represented by the East Atlantic and North Atlantic Oscillation indices). A thorough review was carried out, which addressed each item within the checklist. The result from this approach suggests that changes in these flows can, in part, be attributed to the East Atlantic index and hence climate change. The case study, through assessing the impact of the NAO index upon peak river flows, also demonstrated the need to account for time-varying confounding and hence gain a further insight into the true associations between the exposure and the outcome of interest.

Though the case study focused specifically on peak river flows in Great Britain, this approach is applicable to the attribution of changes in such flows in any location. Another spatial correlation structure may be appropriate or indeed the climate index governing such changes may differ from that considered here, however the core model

in Equation 3.4.1 and checklist for ruling out additional possibilities would hold in any location. As discussed in Chapter 1, Merz et al. (2012) divides attribution studies in flood trend studies into "soft" and "hard" attribution. "Soft" attribution uses hypotheses and references to previous studies to back up attributions to causes of change. "Hard" attribution requires evidence that detected changes are both consistent with the proposed cause of change and inconsistent with potential alternative causes, and further that some confidence level in the attribution statement is provided. The authors of Merz et al. (2012) note that, by that particular point in time, no such studies had fulfilled the requirements of hard attribution. In contrast, this systematic checklist both fulfils these three requirements and goes beyond their scope. Changes in peak river flows in Great Britain are consistent with changes in the East Atlantic index, and inconsistent with potential alternative causes discussed within the case study. A credible interval is provided for the estimate of the effect of the climate index upon flows. Beyond these three requirements, the method also considers the interpretability of the model to the end-user, controls for unmeasured confounding and frames the problem in the context of causal diagrams which are revisited throughout the analysis to ensure diligence in considering possibilities.

The approach used in the case study on river flows can also be directly applied to many other environmental observational studies, such as for example the assessment of the impacts of climate change on water quality. Xia et al. (2015) discusses the importance of improving methods in order to separate the effects of climate change from confounders such as land use changes and pollution. Land use changes will vary in time and can be controlled for in a similar approach to the "partialling out" method used in the case study, and following the checklist discussed would ensure that any relationship uncovered in the analysis can be deemed to be causal. One example where the approach does not apply, however, is in the case of feedback loop environmental problems, which cannot be summarised by a DAG (Pearce and Lawlor, 2016).

The river flow case study presented primarily focused on time-varying confounding, however spatial confounding can also be an issue in environmental observational studies. Typically, models for spatial observational data will include both a fixed effect for the covariate(s) of interest along with random effects representing the spatial correlation in the data. The aim of adding these spatially correlated random effects is to reduce bias in variable estimates resulting from unobserved covariates. However, if these variables are also spatially correlated (as well as the response), then the introduction of such spatial effects may result in the confounding of the covariate effect (Reich et al., 2006; Thaden and Kneib, 2018) demonstrated that the effect of spatial

confounding may be so strong that a coefficient that is significant using a non-spatial model may become non-significant with the inclusion of a spatial random effect.

Such a scenario arises, for example, when estimating the impact of site-specific characteristics on tree defoliation, as noted by Augustin et al. (2007). This confounding, which makes it impossible to distinguish between spatial effects modelled through location co-ordinates and site-specific effects (such as soil type and local weather conditions), means that there has been a focus on prediction of trends in defoliation rather than establishing a causal relationship between the exposure and the response (Augustin et al., 2009). This spatial confounding has primarily been investigated in the discrete case, i.e. areal spatial data, where the main approach involves the constraining of spatial random effects to be orthogonal to the fixed effects (Hodges and Reich, 2010), while Hanks et al. (2015) discussed the application of this approach to the continuous case, noting that while there are computational benefits to this method, care must be taken as it assumes that all variation in the same direction as the fixed effects can be attributed to them. Thaden and Kneib (2018) propose the use of structural equation modelling methods to separate the direct effect of a covariate from the indirect effect which arises from correlations with other variables. The authors suggest a different interpretation of the role of space, allowing space to have an effect on both the exposure and the response simultaneously, while the other approaches discussed allow for information within the response only. This method provides reliable estimates, even when it is unclear which variables are impacted by unknown spatially varying variables, and thus may be incorporated within a checklist for environmental observational studies where it is anticipated that spatial confounding may occur.

3.6 Conclusions

The accurate attribution of the mechanisms driving long-term, large-scale changes in environmental problems is crucial for informing future interventions, yet standard methods for establishing causality are often not directly applicable to such studies. In this paper, we provided a selected review of potential approaches for assessing the evidence for causal relationship in these studies. Such methods as causal diagrams, the Bradford Hill criteria, methods of multiple working hypotheses, and the weight of evidence approach as seen in Suter et al. (2017) were assessed for suitability to environmental problems. We proposed a systematic checklist for the causal assessment of these issues, incorporating a number of these methods, while also including additional steps to address environment-specific issues. This framework can be used to assess

for causal relationships in a broad class of environmental problems by laying out the strength of evidence according to this checklist. This method was illustrated through a case study of peak river flows in Great Britain, in which a strong effect of the East Atlantic (EA) index on high flows was identified.

3.7   Chapter conclusions

The checklist proposed in this chapter provides a simple, usable tool that provides a conclusion to the attribution analysis conducted in Chapter 2. In particular, the "hard" attribution criteria set out by Merz et al. (2012) has been satisfied, and it is possible to state that changes in peak river flows in Great Britain can, in part, be attributed to the East Atlantic index and hence to climate change. Another important contribution of this chapter is the general applicability of this causal approach to environmental observational studies, for which standard causal methods are often not appropriate.

However, it is clear from this case study that collinearity, particularly with time, is an issue that is not easily overcome in environmental observational studies such as this. As a consequence, it is not always possible to attribute changes in some outcome to a variable of interest. For example, it was not possible to attribute changes in peak river flows in Great Britain to the North Atlantic Oscillation index due to its collinearity with time — it is possible that there is a causal relationship, but using these methods it is not possible to conclude this. The difficulty in establishing causality is discussed again in Chapter 4, when a more regional assessment of trends in river flows is carried out.

Chapter 4

Assessing trends on a single river network

4.1 Introduction

Water-related natural hazards can have tremendous impacts on the well-being of communities; water levels severely below average can bring periods of drought and water scarcity, while those severely above average can be connected to floods. These types of hazards are typically spatial and temporal in nature. In the case of rivers, in which flows tend to move only downstream (tidal systems being an exception), the measurements for a single river with multiple stations form a network. Learning how river flows evolve throughout a network in space and in time is beneficial when it comes to accurately estimating the availability of water in the network and the probability at each location of exceeding some threshold for flooding.

The key contribution of this chapter is the development of spatial methods for trend detection on a single river network across a 40 year period, by exploiting the inherent network structure of such rivers. As noted in Chapter 2, at-site methods are often poorly powered and unable to detect trends when a trend is present. The use of a countrywide trend assumption ensured a better ability to detect trends present in the data. However, this assumption of a countrywide trend might be too strong, although it was necessary for any trend to be detected with the limited available data. The assumption was balanced in part using station-specific effects, which allowed for some variability between stations. In this chapter, we instead propose a river-wide assumption, which is

a more relaxed assumption. There is a switch to focus on a single river network, resulting in a more accurate estimate of trends present in the data. The structure of a river monitoring network of stations is encoded using a first-order conditional autoregressive (CAR) model. CAR models, while often seen in assessing disease incidences in areal data, have not often been applied in an environmental setting such as this. We believe that this contribution is beneficial for accurately encoding spatial relationships between gauging stations in a network. A benefit of this approach is the ability to learn about the regional behaviour of river flows, and leads to reduced variability when compared to an at-site approach. It may also allow for the prediction of intermediate flows along a network, which would not be possible with at-site methods. The full record of the flow data is used (i.e. daily mean flows), not just the annual maximum, in comparison with Chapter 2. The use of daily mean flows, while also giving information surrounding water availability in the network, ensures there is a larger amount of data entering the model — this is a benefit as we are better able to detect trends, however as we are using daily mean flows we might expect to see smaller trends. The amount of data may also lead to computational issues. Another potential drawback is that the focus is on an entire river basin rather than a single river, leading to a trade-off between accuracy and generality of the estimation.

The structure of this chapter is as follows. We provide a brief review of previous approaches to incorporating river flow structure and direction into the estimation of trends in flows. We then investigate a new approach to modelling trends in daily mean flows on a river station network which exploits the network structure of rivers. This is then used to model the spatial covariance between stations, making use of conditional independence and directed graphs to map out these relationships through a sparse precision (inverse covariance) matrix. This spatial structure will be included within a model to estimate trends in daily mean flows across the river network. It is shown that the addition of this spatial structure reduces the variability in trend estimates when compared to a simpler (i.i.d.) model.

We also discuss in brief potential alternative methods which incorporate the network structure in the mean, allowing for the attribution of trends detected to some local drivers of interest. The difficulty of accurately representing the physics of the river network within a simple statistical model is discussed. The methods discussed is showcased using the network of station in the river Eden catchments in the northwest of England, which has experienced a series of devastating floods in the last 20 years.

4.2 Proposed approaches for modelling flows on a river network

In this chapter, we introduce novel approaches for explaining the flows at each given location in a network, both of which exploit the network structure of rivers. One method, investigated in detail this chapter, exploits the network structure to model the spatial covariance (precision) between stations, making use of conditional independence and directed graphs to map out these relationships. Modelling the spatial covariance between stations given the directional nature of river flows has previously been developed by Ver Hoef et al. (2006). The novel approach presented in this chapter instead makes use of conditional independence relationships between stations to represent the network structure. In this way, we obtain a sparse precision matrix, giving rise to fast inference for a large number of observations. The aim of this approach is to accurately detect long-term trends in daily mean flows in a single river network. In particular the methods proposed in Section 4.7 and implemented in Section 4.8 exploit the network structure of a single river. The approach makes use of direct connections between river gauging stations to model flows at stations as a function of their nearest neighbours. This is illustrated by the river Eden network in north England. The method will be compared to modelling the stations as independent from each other, in order to demonstrate the benefit of utilising the spatial structure of the network. A further approach to utilising the network structure, by modelling the structure in the mean, will be discussed in brief later in this chapter.

Before these approaches are explored in detail, we first investigate prior approaches to modelling the spatial structure between monitoring stations and how these might be of benefit to the method devised in Section 4.7.

4.3 A review of past approaches

Spatial patterns in a river gauging network will depend both upon distance between gauging station locations along the network, but also the direction of the river flow. As a consequence, as noted by Gallacher (2016), it may not be appropriate to use a purely Euclidean distance in this setting. Instead, O'Donnell et al. (2014) discussed the possibility of using a different approach to spatial distance, noting that models for spatial river flow data may need to be constructed in a different way to using standard distances. Instead of using the Euclidean distance between sites, it was proposed

one can instead use the stream distance — see Ver Hoef et al. (2006), which defines this quantity as "the shortest distance between two locations, where distance is only computed along the stream network". Curriero (2006) demonstrated that substituting this stream distance in place of Euclidean distance does not produce a valid spatial covariance model, except for when an exponential covariance structure is used. For example, Gardner et al. (2003) made use of stream distances with a spherical covariance function, without ensuring that this leads to a valid covariance matrix. Using these stream distances to determine a covariance matrix can give rise to negative eigenvalues. This means the matrix is no longer positive definite and thus cannot be considered to be a covariance matrix any longer.

Instead, Ver Hoef et al. (2006) made use of moving average constructs to define a broad class of valid spatial covariance models which use stream distance. These methods, based on a moving average approach, incorporated stream distances, river flow and appropriate weights at confluence points in the network. It had been shown by Barry and Ver Hoef (1996) that a large class of autocovariances may be developed by producing random variables as the integration of a moving average function over a white noise process. In order to have a stationary process, the function must have finite volume. The moving-average construction allowed for a valid autocovariance function. These moving average functions allowed Ver Hoef et al. (2006) to build valid stream models which account for directional flow, specifically in that these where moving average functions will only be positive upstream from a monitoring station. At a confluence point, the moving average function will continue upstream but will split according to some defined spatial weights that have been assigned. The weights may be determined according to flow volume, catchment area or some other appropriate descriptor. These models, known as Tail-up models, assign a correlation of zero to pairs of locations which are not connected. This approach was formalised by Cressie et al. (2006) by defining the process according to some stochastic integral with respect to a Brownian motion. Such a model only allows for spatial dependence between two points which are upstream of each other. While such models will not be directly applied in this thesis, the idea of only allowing for those stations directly upstream to influence values taken at a monitoring station of interest will be used in the approach discussed in section 4.9.

Monestiez et al. (2005) used conditional probabilities to define spatial dependence in a river flow network using directed trees. This defined two types of tree models - one based upon graph theory, characterising topological relationships between parts of the river, and one consisting of segments of the river (some set of points with a metric which are linked together to construct the topology of the first kind of tree). However,

Cressie et al. (2006) noted that as such approaches are not based on covariances, they are not as easily adapted for use in kriging as tail-up methods. In the approach proposed, we make use of conditional relationships to define spatial dependence between river gauging stations in this chapter, however, in order to respect river flow direction, we cannot construct a valid covariance matrix. Instead, we use these conditional relationships to produce a valid sparse precision matrix leading to a fast inference approach which respects the structure of the network.

O'Donnell et al. (2014) proposed an alternative approach of adapting kernel methods for stream distance and penalized splines for a finite discrete approximation to river networks. These were then used to estimate spatial trends instead of a covariance function. The authors noted that, while covariance functions provide a simple framework for statistical modelling of spatial data, in cases where trends may be non-parametric in nature (as can often be the case in environmental studies) it may be more appropriate to instead directly model these trends using some form of flexible regression. This method would then incorporate some form of spatial error. Such methods were generally described in the additive modelling framework by Wood (2006). O'Donnell et al. (2014) proposed the use of a set of basis functions $\phi_j(x), j = 1, ..., p$ as components in a regression model. These would represent the estimate in the form $\hat{f}(x) = \sum_j \beta_j \phi_j(x)$. Using basis functions means that estimates can be expressed in proper functional form through specification of the coefficients β_j.

A penalised approach which was useful for controlling the smoothness of the estimate was proposed by O'Donnell et al. (2014). The smoothness of the β at adjacent stream locations j and k was measured by $(\beta_j - \beta_k)^2$, once no confluence point existed between the two. Where such a confluence point existed, then this measure of smoothness must reflect the relative levels of flow in the contributing streams, which were denoted f_a and f_b. At the confluence point c, it was expected that $f_c = f_a + f_b$. Then, smoothness was achieved at the confluence point c using the penalty

$$\lambda\{\omega_a^2(\beta_a - \beta_c)^2 + \omega_b^2(\beta_b - \beta_c)^2\},$$

where the ω_i represent weights which can be determined by the relevant flows of inputting streams. Here λ controls smoothness through the weight that is attached to the penalty. A P-spline model was then formulated, which did not directly depend on measures of stream distance or spatial location, but through a mixing process. The first-order penalty used in O'Donnell et al. (2014) described this mixing by imposing smoothness within segments of a stream, but also across confluences. For the latter of these, a conditional independence assumption was made that, given immediately

upstream catchments, flows at a station of interest are independent of all those further upstream. A conditional independence assumption will be used in the approach proposed in this chapter.

The method proposed in O'Donnell et al. (2014) allowed for the borrowing of strength from nearby locations (the benefits of which are discussed for annul maximum flows in Chapter 2), while preserving the directionality and structure of a river flow network. While the approach proposed in this chapter differs in its implementation by instead making use of conditional independence relationships between stations, the same motivation of borrowing strength from nearby location (thus improving statistical power) and respecting the river network structure is followed throughout this chapter.

4.3.1   Summary

As discussed above, a number of methods have been used to model river flows on a single network, incorporating both distance and river flow direction. The tail-up method proposed by Ver Hoef et al. (2006) overcame the issues of stream distances giving rise to invalid covariance matrices. This was achieved by utilising moving average constructs to define a broad class of valid spatial covariance models which use stream distance. These methods respect the direction of flows by specifying that moving average functions will only be positive upstream from a monitoring station. The approach of Monestiez et al. (2005) made use of conditional relationships o define spatial dependence in a river flow network using directed trees. Finally, O'Donnell et al. (2014) adapted kernel methods for stream distance and penalized splines for a finite discrete approximation to river networks, thus preserving the structure of a river network while borrowing strength from nearby locations.

Aspects of all of these methods will be incorporated within the approach proposed in this chapter. Conditional independence relationships between stations will be used to produce a sparse precision matrix which encode the network structure, while a multilevel modelling approach is used to pool data from multiple sites together, allowing for the borrowing of information across monitoring stations.

4.4   Data

The method proposed in Section 4.7 will be exemplified using daily mean river flow data across a series of river gauging stations spanning the river Eden. These catchments

112

are located in the north-west of England. This data set contains the daily mean flows in each day of the water year (which runs from October to September), measured in m^3/s. In total, there are 247,807 observations from 18 gauging stations on the river Eden network, ranging from 1951-2017 — these station locations can be seen in Figure 4-1 and some basic information for each station on the River Eden can be seen in Table 4.1. This includes active years, length of river flow records, catchment area and SAAR, the catchment's average annual rainfall in the standard period (1961-1990) in millimetres. The average record length of stations within this network is 40.77

Station ID	Year opened	Year closed	Record length	Area (km^2)	SAAR (mm)
76001	1951	-	66	32.34	2438
76002	1966	1998	32	1374.85	1272
76003	1960	-	57	406.99	1768
76004	1962	-	55	156.10	1828
76005	1970	-	57	618.21	1142
76007	1967	-	60	2276.00	1182
76008	1967	-	60	333.42	1073
76009	1968	2000	32	147.2	NA
76010	1969	2014	45	157.58	940
76011	1967	-	60	1.63	1096
76014	1971	-	56	66.82	1492
76015	1970	-	57	149.36	2149
76017	2004	-	13	1371.70	1273
76019	1999	-	18	63.09	983
76020	2006	-	11	NA	NA
76806	2000	-	17	223.03	1270
76809	1997	-	20	249.70	1211
76811	1999	-	18	33.97	1428

Table 4.1: A summary of record lengths and catchment descriptors at each station

years, with a minimum of 11 and maximum of 66 years. There is an average of 2% of records missing across the network. The average size of catchments within the network is 578km^2, with a minimum size of 1.63km^2 and a maximum of 2276km^2. All data related to rivers and river flows can be obtained from the UK National River Flow Archive (NRFA) (Dixon, 2010), which is the primary UK source of hydrometric data.
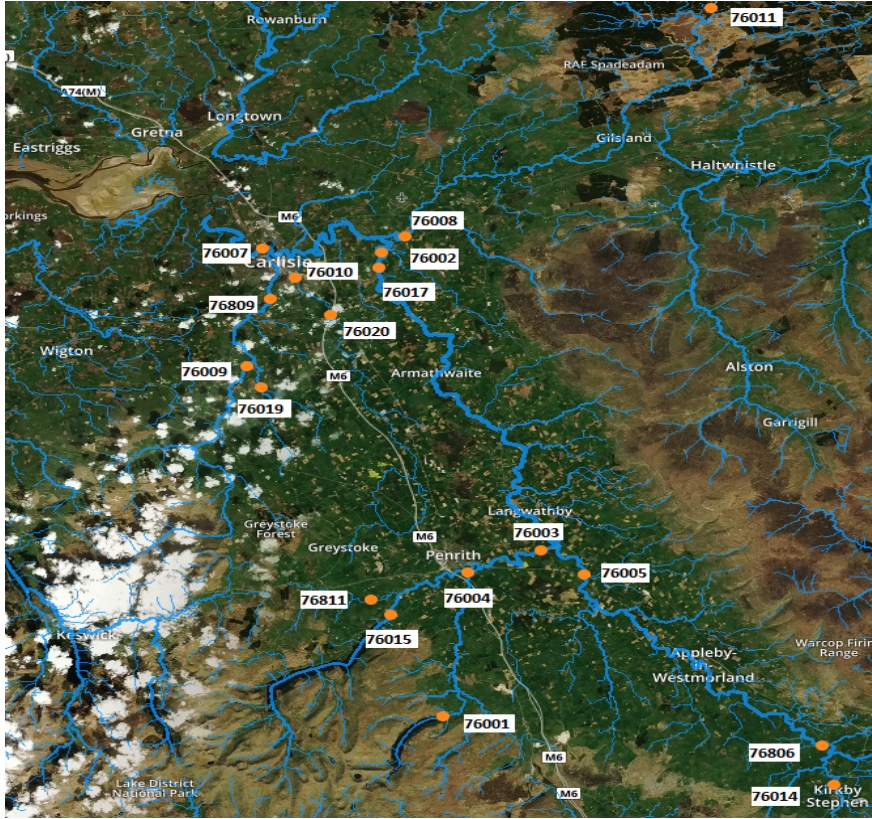
Figure 4-1: Gauging station locations for the River Eden catchments in the UK, from National River Flow Archive (2017).

4.5 Exploratory analysis and data standardisations

Because of the different sizes of catchment area drained at each location, river discharge measurements at the different stations tend to have different means and variances, potentially along with differing skewness and distribution. To make these measurements more comparable across time periods, transformations are applied. As we wish to model all stations together, if stations exhibit considerably different behaviour to the others, they may be excluded from the analysis or treated differently.

4.5.1 Data standardisations

In order for measurements on any given day across any given year to be comparable in the network model, the data requires some transformation and removal of seasonality. This will be achieved by taking the log of the flows, then adjusting for seasonality

114

through fitting a generalised additive model (GAM — see Section 1.4.5) with the day of year as a smooth function. Using a log normal distribution has been found to fit UK flow data reasonably well (Prosdocimi et al., 2014), while adjusting for seasonality will ensure that longer-term trends are more clearly detected.

We start with removing the seasonality of the log-adjusted flows obtained by fitting a GAM with a day of year (DOY) smooth. We then adjust the predicted values using the fitted value at that station at the start of the water year (October 1st). That is, for station i at time t, the log and seasonal adjusted flow will be

$$\log Y_{it} = X_{1it} - X_{2it} + X_{3it} + \varepsilon_i, \ \varepsilon_i \sim \mathcal{N}(0, \sigma^2). \tag{4.5.1}$$

where Y_{it} is the adjusted daily mean flow, X_{1it} is the original river flow, X_{2it} the predicted flow and X_{3it} the river flow on October 1st (i.e. the first day of the water year) at station i and time t. The predicted flow in Equation 4.5.1 is calculated from the GAM described in the following paragraphs.

The log of the daily mean flows is taken, and a generalised additive model is fitted for each station for this quantity against a smooth function of the day of year and a linear term for water year as follows

$$\log(\bar{Y}_i) = \alpha + f(Z_{1i}) + \beta Z_{2i} + \varepsilon_i, \ \ \varepsilon_i \sim \mathcal{N}(0, \sigma^2) \tag{4.5.2}$$

where $\bar{Y}$ is the original daily mean river flow, Z_1 represents the day of year, $f(\cdot)$ represents a smooth function of the day of year and Z_2 the water year. A cyclic cubic spline is used as the basis function, which is necessary for the seasonal term as there should be no discontinuity between September and October (as discussed in Section 1.4.5). The addition of the linear water year term ensures that values across years are also comparable.

This GAM is fitted, with the choice of k (the basis dimension size) chosen following the discussion in Section 1.4.5. There are 365 separate values for day of year, so one can choose higher than the default of $k = 10$, however no more than $k = 12$ was seen as necessary to observe the structure of the smooth and this choice cut down on computational time compared to a larger k. The fitted smooth terms for the seasonal (day of year) component, i.e. $f(Z_1)$, are extracted for each station. The adjusted flow according to Equation 4.5.1 is then calculated using these seasonal component values as predicted flows. Plots for the individual station-level smooths can be seen in Figures 4-2 and 4-3, where the line in each case shows the fitted smooth effect, while the shaded region displays the confidence interval. These all follow a similar structure with peaks

115

in winter months, with peaks roughly around the 100 day mark (December/January) in the water year. However, there appears to be some variation at different timepoints in the year — it can be seen that stations 76001 and 76017 displays somewhat differing behaviour to the remaining stations, which may need to be taken into account when it comes to the selection of data for analysis. These differences may be down to the level of smoothing (which differs slightly in each case, note that the smoothness parameter could be fixed to ensure the level of smoothness is consistent across stations.) versus variation, however further investigation may be required. We will look at station 76001 as a particular case study of candidate stations for removal from the analysis.
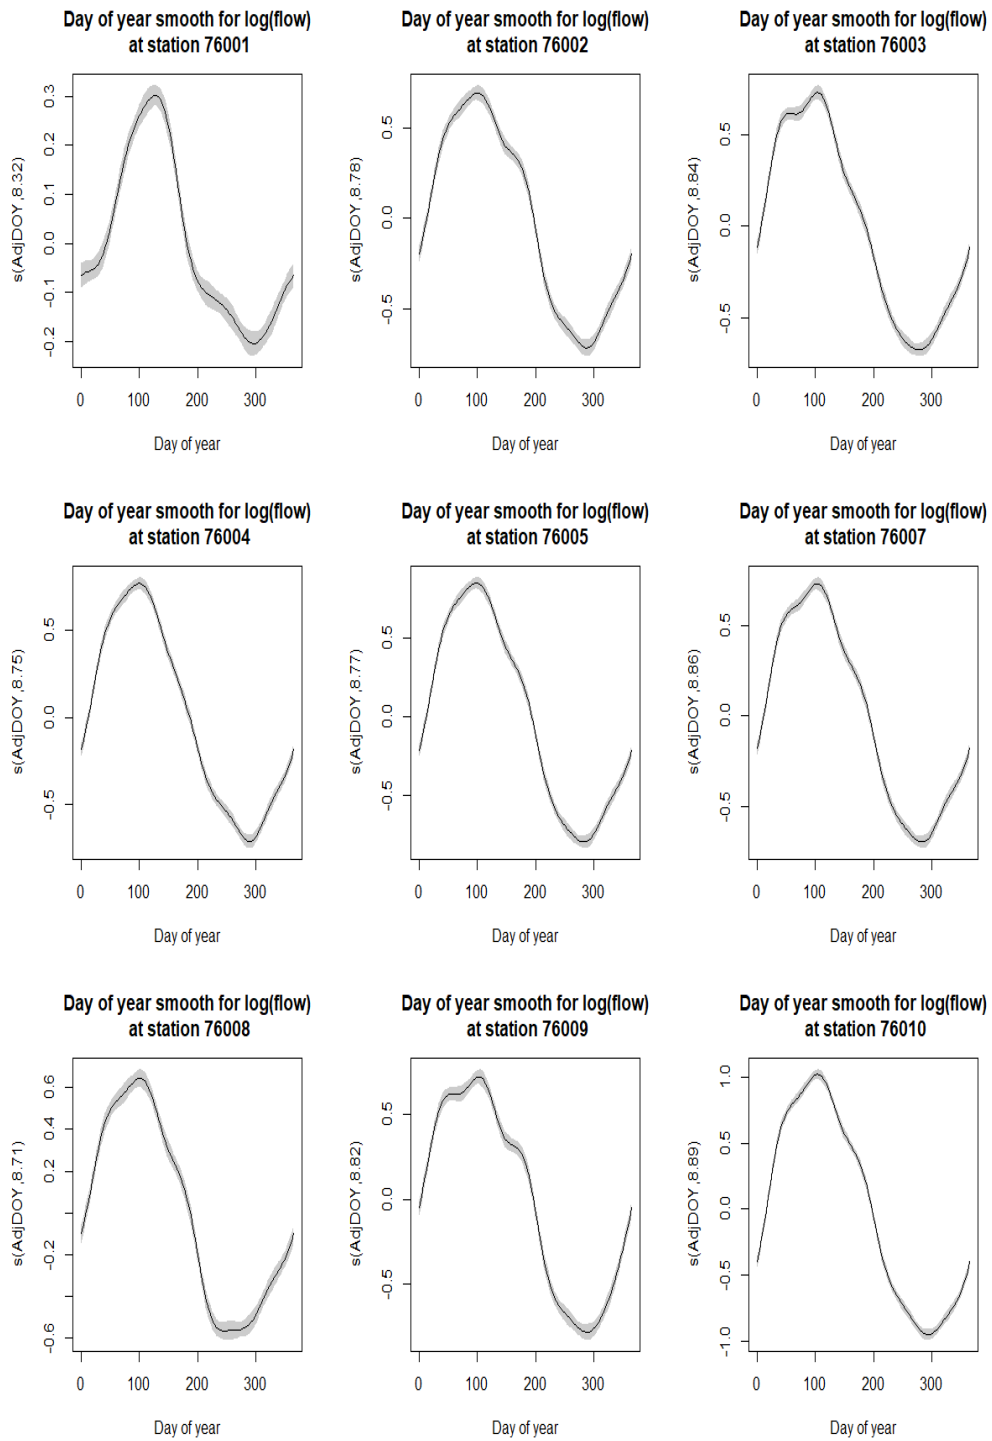
Figure 4-2: A plot of the daily smooth function (with uncertainty) for stations 76001-76010. The line represents the fitted smooth effect, while the shaded region represents the confidence interval.

117
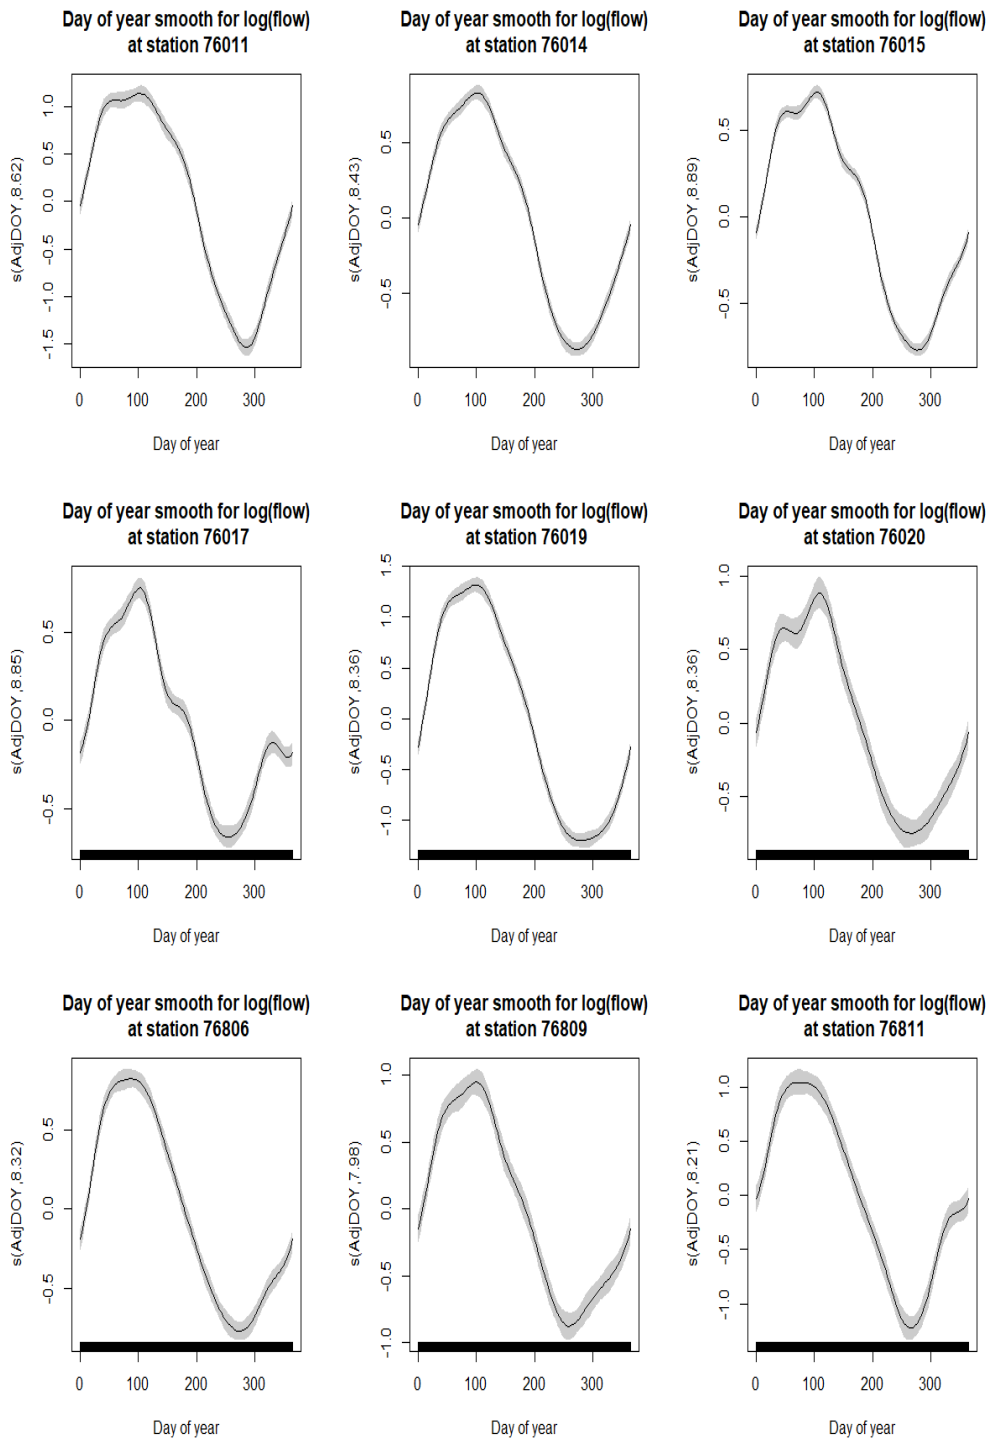
Figure 4-3: A plot of the daily smooth function (with uncertainty) for stations 76011-76811. The line represents the fitted smooth effect, while the shaded region represents the confidence interval.

4.5.2   Functional boxplots

One simple way to identify potential stations that behave differently is to plot all the time series curves for each station through a functional boxplot (Sun and Genton, 2011). A functional boxplot involves ordering of functional data from the centre outwards, and consequently can provide a measure of the "outlyingness" of a particular gauging station's time series.

One such example of a functional boxplot can be seen in Figure 4-4, where each curve represents one year (2003) of data from each of the gauging stations on the network. Here, the red curves represent potential outlier candidates, the black curve represents the median curve, blue curves are the "envelopes" while magenta represents the 50% central region. In this case, the median curve is for station 76007, which is the most downstream station in the network. Despite this plot showing a number of potential outlier candidates, for the most part these curves all display a similar shape. The one counterexample for this is for station 76001, which we will investigate in more detail.
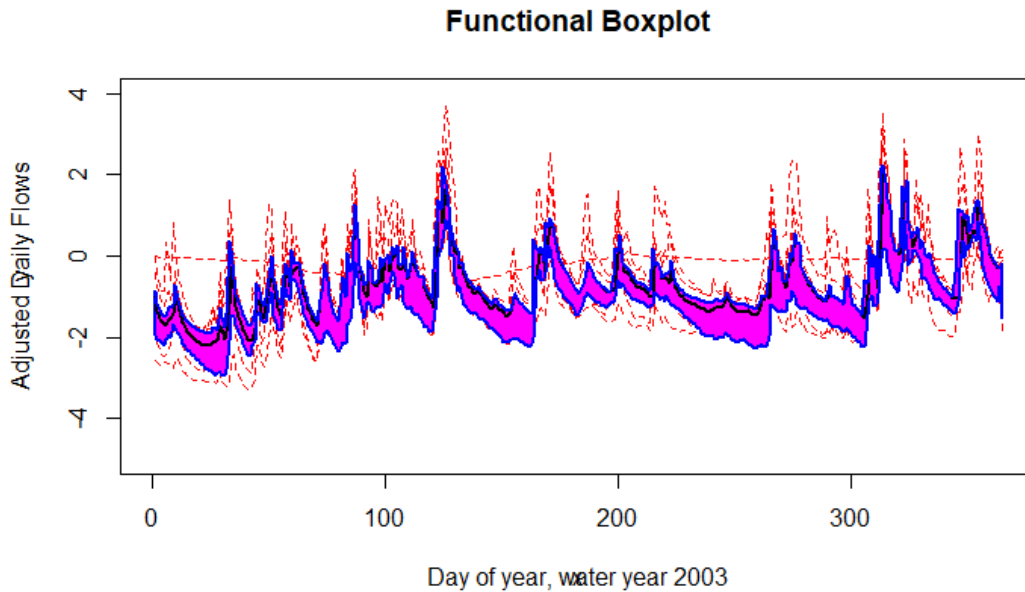
Figure 4-4: A functional boxplot for station-level data in 2003. Red curves represent potential outlier candidates, the black curve represents the median curve, blue curves are the "envelopes" (the border of the 50% central region), and magenta represents the 50% central region.

4.5.3   Station 76001

Station 76001, which gauges the Haweswater Beck at Burnbanks, is one of the most upstream gauging stations in the river Eden network. It is a very wet catchment, just 500m downstream of the Haweswater reservoir (National River Flow Archive, 2017). In dry years, all flood flows are absorbed by the reservoir. The resulting flow pattern observed at station 76001 reflects the heavy abstraction caused by this, meaning the natural flow pattern of this station may not always be observed. The fact it is affected by a reservoir is in itself a good justification for excluding this station from the analysis as it is expected that this will result in differing behaviour of the flow measurements. Additionally, as noted in Chapters 2 and 3, human-altered flow is not informative of natural changes which are what we are interested in.

It was observed in Section 4.5.2 that the shape of the data from station 76001 appears to be an outlier when compared to the remaining stations; this is not surprising due to the abstraction caused by the reservoir noted above. it is clear that station 76001 does not display similar behaviour to the remaining stations from the boxplot in Figure 4-4 and the daily smooths in Figure 4-2. As a consequence, this station will not be considered within the analysis that follows. Further criteria for the selection of candidate stations are discussed in the following section.

4.6   Criteria for selection of stations for analysis

A long period of records across all stations is needed to accurately estimate temporal trends within the network, with the aim to estimate trends over a 40 year period from 1978–2017. However, this is not possible given the dropout and addition of stations throughout the time period of the network. As a consequence, criteria are proposed for the selection of stations to ensure that each station in the analysis has a long record. This is a trade-off — ideally, we would want to use as much data as possible to get more precise estimates of trends in the network, however for the detection of long-term trends we need those stations with longer records.

Only stations meeting the following criteria are included in the study:

1. Station data must begin by 1978, in order to ensure a sufficient record length (40 years) for the station. This rules out stations **76017**, **76019**, **76020**, **76806**, **76809** and **76811**.

2. Stations with a full year of data missing during this period will be dropped from the analysis. This rules out station **76001**, which is missing data for 10 separate years across this period. In addition, this station was observed to exhibit unusual behaviour compared to the remaining stations, and was deemed inappropriate for inclusion for that reason also. **76014** has complete missingness of the 1978 and 1979 water years, and will not be included in the analysis.

3. Stations with more than two months of data missing for a given year will have that year dropped entirely from the analysis. Station **76003** has 153 days missing in water year 1979 and 92 days in water year 1980, and as a consequence these two years are removed from the analysis. **76009** has 117 days missing from water year 1997, so this year will be removed from the analysis. Station **76011** has 253 days missing from water year 1972 and 76 missing from water year 1984.

This leaves stations 76002, 76003, 76004, 76005, 76007, 76008, 76009, 76010, 76011 and 76015 available for the analysis which follows.

This subset of the data consists of 155,398 observations from 10 river gauging stations for the time period 1978–2017, with flow records missing from 1298 entries, a total of 0.83% of records.

4.6.1 At-site trend analysis using GAMs

We now wish to obtain a preliminary estimate of the yearly effect at a station-by-station basis, having transformed and deseasonalised the data. We investigate a simple at-site approach using GAMs and `mgcv` first (Wood, 2017). Using this method will allow us to see whether it is reasonable to assume that flows from different gauging stations behave similarly, i.e. if the assumption of a network-wide trend is acceptable within our proposed approach.

A simple generalised additive model (GAM) of the adjusted flows against water year is fitted using `mgcv` to determine whether any evidence can be seen of a common trend across stations, and if the relationship can be reasonably approximated by a linear term.

$$Y_{it} = \alpha_i + f_i(X_{it}) + \varepsilon_i, \ \ \varepsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$$

where $f(\cdot)$ represents a smooth function of time, Y are the adjusted flows discussed in Section 4.5.1, i represents the gauging station location and X an explanatory variable measured at time t, in this application X is taken to be the year of record/time. The

values of α and f are estimated separately for each station. Figure 4-5 shows the

Figure 4-5: A plot of station-by-station time trends in daily flows using mgcv

relationship between water year and the adjusted log of daily flows on a station-level basis. This plot suggests that, on the whole, the stations in the river Eden network follow a similar time trend, meaning that it is reasonable to assume a network-wide trend. As a consequence, instead of using an at-site approach, the hierarchical and spatial structure of the data may be better exploited by using a network-based approach to model the spatial relationships between stations. Modelling all stations together should also allow for the better detection of trends that might be missed in an at-site approach (see Chapter 2, for example). We discuss such methods in the following section.

4.7 Modelling the covariance (precision) between stations

In this section, we propose incorporating the inherent spatial structure of the network within a model, in order to more accurately estimate trends of the network as a whole. This will be achieved by encoding the conditional independence relationships between stations (only those that are connected to each other on an undirected graph) through a sparse precision matrix, and modelling this spatial structure using a first-order CAR

model. Such models have not previously been applied to environmental observational studies such as this. The use of CAR models to represent the spatial relationships between these river stations is a novel approach in the area, which will allow for a good representation of changes in measuring networks, such as river flows.

4.7.1 Conditional independence and directed acyclic graphs

We propose that the river gauging stations on a single river can be considered as a network, which is comprised of nodes (the river gauging stations themselves) and edges between nodes (the stretches of river between each station). We will demonstrate how conditional independence between nodes can be inferred upon inspection of the graph structure. It is important to note that a node's parents separate it from further ancestors. These conditional independences are related to the graph itself, and graph separation. In this section, we discuss the relationship between conditional independence and graph separation, illustrating this relationship via three simple examples. Further detail on these examples can be seen, for example, in Koller and Friedman (2009); Murphy (2012); Pearl (2009).

4.7.1.1 Indirect connections

In Figure 4-6, the lack of an edge between nodes A and C can be interpreted to be conditional independence. This can be inferred from the graph structure (and, in practice, by considering how a river flows):

$$P(A, B, C) = P(A)P(B|A)P(C|B),$$

which suggests that

$$P(C|A, B) = \frac{P(A, B, C)}{P(A, B)} = \frac{P(A)P(B|A)P(C|B)}{P(A)P(B|A)} = P(C|B).$$

This example is one of the most basic conditional independencies. Here, B is the parent of C, while A is a non-parent ancestor of C and thus is conditionally independent of C given B (i.e. a Markov chain). In terms of graph separation, one can clearly observe that B separates A and C.

$$A \longrightarrow B \longrightarrow C$$

Figure 4-6: An example of an indirect connection between nodes.

4.7.1.2 Common cause

In Figure 4-7, the lack of an edge between nodes A and C can again be interpreted to be conditional independence. This can be inferred from the graph structure, which suggests that

$$P(A|C, B) = \frac{P(A, B, C)}{P(B, C)} = \frac{P(B)P(A|B)P(C|B)}{P(B)P(C|B)} = P(A|B).$$

Again, graph separation clearly demonstrates independence here. If B is unknown, however, then A and C may appear to be dependent (a hidden variable scenario). Note however that this type of structure rarely occurs in river networks.

Figure 4-7: An example of a node which is a common "cause" for two others.

4.7.1.3 V-structure

Finally, in Figure 4-8, the lack of an edge between nodes A and C cannot be interpreted to be conditional independence, but instead *marginal independence*. Here, once one knows the value of B, then A and C may depend on one another. In the case of river networks, this means that tributaries that combine together are not necessarily conditionally independent of one another.

Figure 4-8: An example of a v-structure of nodes.

4.7.2 Network spatial structure for 1978–2017 period

River flows across the 10 stations in the river network for the period 1978-2017 are included in the analysis. We begin by re-ordering the network into a directed acyclic graph (DAG) where one can more clearly see the parent, child and ancestor nodes, and consequently identify the conditional independence between nodes. The network diagram for this time period can be seen in Figure 4-9.

Figure 4-9: A graph (right) of the stream network structure in the 1978–2017 time period.

Now, for station $X_i = 7600i$, the factorisation theorem allows us to represent the

probability of $P(X_2, X_3, X_4, X_5, X_7, X_8, X_9, X_{10}, X_{11}, X_{15})$ as

$$
\begin{aligned}
P(X_2, X_3, X_4, X_5, X_7, X_8, X_9, X_{10}, X_{11}, X_{15}) = {} & P(X_4)P(X_5)P(X_9)P(X_{10})P(X_{11}) \\
& \cdot P(X_{15})P(X_3|X_4, X_{15})P(X_2|X_3, X_5) \\
& \cdot P(X_8|X_{11})P(X_7|X_2, X_8, X_9, X_{10})
\end{aligned}
$$

We then proceed by "moralising" the DAG. A moral graph is used here to find the equivalent undirected form of a DAG. The moral graph may be formed by adding edges between the parents joining the parents of common children, and subsequently making all edges in this graph undirected — this is equivalent to saying that a moral graph is an undirected graph in which every node of the original DAG is now connected to its Markov blanket (a set of nodes composed of the node's parents, its children, and its children's other parents). The moralised graph version of the previous figure may be seen in Figure 4-10. This moralised version of the graph allows one to construct a

Figure 4-10: A moralised version of the DAG as seen in Figure 4-9.

precision matrix of a Gaussian Markov random field (GMRF), where entries are zero if X_i and X_j are conditionally independent of each other and only non-zero otherwise. The use of a precision matrix over the covariance is useful, as Rue and Held (2005) demonstrated that a number of common models with complex and dense covariance structures have sparse precision matrices. Such sparse matrices are computationally

efficient given they are mostly comprised of zeroes which are not stored in a computer. This sparsity structure arises from the conditional independence structure of the network. Here, conditional dependence can be seen by identifying cliques (i.e. a subset of nodes of the undirected graph such that every two distinct nodes in the clique are adjacent) in the graph. Then, the general structure for the precision matrix Q will be as follows (where subscript i represents node $7600i$ and jk represents node $760jk$):

$$Q_{ij} = \begin{bmatrix} q_{22} & q_{23} & 0 & q_{25} & q_{27} & q_{28} & q_{29} & q_{210} & 0 & 0 \\ q_{32} & q_{33} & q_{34} & q_{35} & 0 & 0 & 0 & 0 & 0 & q_{315} \\ 0 & q_{43} & q_{44} & 0 & 0 & 0 & 0 & 0 & 0 & q_{415} \\ q_{52} & q_{53} & 0 & q_{55} & 0 & 0 & 0 & 0 & 0 & 0 \\ q_{72} & 0 & 0 & 0 & q_{77} & q_{78} & q_{79} & q_{710} & 0 & 0 \\ q_{82} & 0 & 0 & 0 & q_{87} & q_{88} & q_{89} & q_{810} & q_{811} & 0 \\ q_{92} & 0 & 0 & 0 & q_{97} & q_{98} & q_{99} & q_{910} & 0 & 0 \\ q_{102} & 0 & 0 & 0 & q_{107} & q_{108} & q_{109} & q_{1010} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & q_{118} & 0 & q_{1111} & 0 & 0 \\ 0 & q_{153} & q_{154} & 0 & 0 & 0 & 0 & 0 & 0 & q_{1515} \end{bmatrix}$$

This forms a sparse precision matrix. We make use of this sparse matrix for fast inference, through the specification of a first-order CAR model to represent this spatial structure of the network. A simple alternative to using this CAR structure would be to use a general covariance matrix which would not enforce the network structure, or to have an i.i.d. model. These approaches will be compared with the CAR model in section 4.8, to see how including a spatial structure can improve the variability of model estimates and lead to faster inference when compared to a dense matrix.

4.7.3 CAR models

Conditional autoregressive (CAR) models are frequently used to represent spatial random effects for areal spatial data (Banerjee et al., 2014). Here, the river gauging data provided is point-level, however stream gauging stations are each directly connected to the nearest neighbouring gauging station through the river stretches, so the CAR specification remains appropriate. If there is a random quantity $\phi = (\phi_1, \phi_2, ..., \phi_n)'$ at n locations representing the spatial effect, the CAR model may be expressed via full

conditional distributions:

$$\phi_i | \phi_j, j \neq i \sim N \left(\alpha \sum_{j=1}^{n} b_{ij} \phi_j, \tau_i^{-1} \right),$$

where τ_i is a spatially varying precision parameter, and $b_{ii} = 0$. Then, the joint distribution of ϕ is a multivariate normal with mean 0 and variance Q^{-1}, where Q is the precision matrix discussed in the previous section, i.e. $\phi \sim N\left(0, Q^{-1}\right)$. Q can be decomposed into $D_\tau (I - \alpha B)$ so that we can write

$$\phi \sim N \left(0, [D_\tau (I - \alpha B)]^{-1} \right).$$

Here $D_\tau = \tau D$, $D = \text{diag}(m_i)$ a diagonal matrix where m_i is the number of neighbours for a given location i, I the identity, α a parameter which controls spatial dependence ($0 \leq \alpha \leq 1$, where 0 implies spatial dependence and $\alpha = 1$ results in an intrinsic CAR specification), $B = D^{-1}W$ the scaled adjacency and W the adjacency matrix ($w_{ii} = 0, w_{ij} = 1$ if i is a neighbour of j and 0 otherwise). Then the CAR model specification reduces to (see Banerjee et al. (2014) for further details):

$$\phi \sim N \left(\mathbf{0}, [\tau(D - \alpha W)]^{-1} \right). \tag{4.7.1}$$

Here, $\tau(D - \alpha W)$ is the precision matrix. The precision parameter τ and spatial dependence parameter α will be estimated during inference. This representation is useful as it is entirely specified through the precision matrix. This is easily obtained from the graph of the network by calculating the matrix W of adjacent nodes (river stations) in the network.

Given the adjusted river flow data $y_1, y_2, ..., y_n$ at n locations, and under the assumption that we expect that neighbouring locations will have correlated flow levels with likelihood (for location i at time t):

$$y_{it} \sim N \left(\alpha + X_{it}\beta + \phi_i + u_{0i} + \epsilon_{it}, \sigma^2 \right),$$

where X is the matrix of covariates (in this case, the water year), β is a vector of coefficients, ϕ represents the spatial random effect and $u_{0i} \sim N(0, \sigma_u^2)$ represents the random intercept to allow differences in scale between locations, $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ represents the remaining error, where σ_ϵ^2 represents the variation in daily mean flows after controlling for covariates X and intercept α. If a proper CAR prior is specified for ϕ, then ϕ is distributed according to Equation 4.7.1. The complete Bayesian specification

will include priors for the remaining parameters α, τ and β such that the posterior distribution is

$$P\left(\phi, \beta, \alpha, \tau | y\right) \propto P\left(y | \beta, \phi\right) P\left(\phi | \alpha, \tau\right) P\left(\alpha\right) P\left(\beta\right) P\left(\tau\right).$$

When constructing the adjacency matrix in practice, nodes can be defined to be "neighbours" if they are connected in the moralised graph in Figure 4-10. Then, D is the diagonal matrix whose entries consist of the number of these neighbours, and W is given by

$$W = \begin{bmatrix} 0 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

4.7.4   R-INLA implementation of CAR model

The Integrated Nested Laplace Approximation (INLA) is an approximate Bayesian method which performs numerical calculations of posterior densities using a series of Laplace approximations of latent Gaussian models (Rue et al., 2009). If a model can be formulated as a latent Gaussian process, it can be implemented in INLA. R-INLA provides an interface to allow INLA to be performed in R (Rue et al., 2012). By using these approximations instead of performing MCMC, INLA allows for faster inference (as detailed in Chapter 1 and Appendix B) and thus is computationally attractive for larger datasets such as this. In addition, R-INLA has a number of inbuilt models for different specifications of CAR models.

Before we proceed with implementing CAR models for the spatial effect using R-INLA, we must first introduce the concept of a Gaussian Markov random field, and the properties of the various models used to fit this CAR model.

4.7.4.1 Graphs and GMRFs

As discussed in section 4.7.2, we use a (undirected) graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to represent the conditional independence properties between nodes, where

- $\mathcal{V}$ represent the vertices, $1, 2, ..., n$

- $\mathcal{E}$ represents the edges, $\{i, j\}$

with no edge between i and j if $x_i \perp x_j | \mathbf{x}_{-ij}$ and an edge i and j if $x_i \not\perp x_j | \mathbf{x}_{-ij}$. The resulting graph represents the sparsity structure of the precision matrix Q.

A random vector $x = (x_1, ..., x_n)^T$ is called a *Gaussian Markov random field* with respect to the graph $\mathcal{G} = (\mathcal{V} = \{1, ..., n\}, \mathcal{E})$ with mean μ and precision matrix $Q > 0$ if and only if its density has the form

$$x \sim \mathcal{N}(\mu, Q^{-1})$$

and $Q_{ij} \neq 0 \iff \{i, j\} \in \mathcal{E} \; \forall i \neq j$.

In summary, if $x_i \perp x_j | \mathbf{x}_{-ij}$ for a set of $\{i, j\}$, then we need to constrain the parameterisation of the GMRF. This is difficult to achieve with the covariance matrix but easy with the precision matrix (as seen in Chapter 1), and use of the GMRF results in a much faster inference. An alternative to specifying a GMRF by its mean and precision matrix is to specify it implicitly through the full conditionals, leading to the first-order CAR model discussed in the previous section.

4.7.4.2 Besag model

The simplest form of the CAR model for spatial correlation is the intrinsic model proposed in Besag et al. (1991). This is referred to as the *Besag or ICAR (intrinsic conditional autoregressive) model*, which assumes spatial dependence among neighbouring geographical regions. This model states that

$$y_i = \alpha + X_i^T \beta + \phi_i, \tag{4.7.2}$$

where

$$x_i - x_j \sim N(0, \sigma^2)$$

130

if i and j are neighbours. Besag et al. (1991) introduced the use of a first-order intrinsic GMRF (Rue and Held, 2005) to model the spatial effect ϕ. If we define $\mathcal{G}$ as the conditional independence graph (obtained from the adjacency matrix), with ∂i the set of neighbours to a node/location i and $n_{\partial i}$ the number of neighbours. The conditional distribution for ϕ_i is given by

$$\phi_i | \phi_{-\mathbf{i}}, \tau_\phi \sim \mathcal{N} \left(\frac{1}{\mathbf{n}_{\partial \mathbf{i}}} \sum_{\mathbf{j} \in \partial \mathbf{i}} \phi_{\mathbf{j}}, \frac{1}{\mathbf{n}_{\partial \mathbf{i} \tau_\phi}} \right). \tag{4.7.3}$$

Here τ_ϕ is the precision parameter for the spatial random effect.

For a Besag model with an intercept in the model, there must be some constraint ($\sum_i x_i = 0$) preventing the Besag term ϕ from taking on the effect of the intercept. One issue with the Besag model is that it only accounts for similarities between regions, but doesn't account for the fact that each region may vary slightly. This means that it is necessary to add an i.i.d. random effect per region (station in this case). This is equivalent to the BYM model discussed next, which will be the primary model investigated in Section 4.8.

4.7.4.3   BYM model

The Besag model only models within-region over-dispersion as spatially correlated noise, and so this approach will tend to overstate spatial dependence. As a result, the estimate of the random effects will be too spatially smoothed. Instead, we propose using the *Besag, York and Mollie model* (Besag et al., 1991), which is based on a conditional autoregressive (CAR) model for spatial random effects. As noted, in such a model, spatial dependence is expressed conditionally, requiring that the random effect in a given area depends only on some neighbouring values. The BYM model is a model where the random effect associated with any given area is the sum of two components $\phi_i + \theta_i$. ϕ_i represents a spatially structured random effect which is assigned an improper CAR prior, while θ_i represents i.i.d zero-mean unstructured random effects (a random intercept).

In the BYM model we define the adjusted log flows as

$$y_i = \alpha + x_i^T \beta + \phi_i + \theta_i \tag{4.7.4}$$

where α is the overall intercept, β measures the effect of covariates x_i (which may be location-specific), θ is a zero-mean Gaussian with precision matrix $\tau_v I$ representing

some unstructured random effect which accounts for regional variability, and again ϕ is a spatial component which specifies that nearby regions/locations are similar. The conditional distribution of this spatially structured effect ϕ is given as in Equation 4.7.3.

A key part of the BYM model is the specification of the neighbourhood structure, which is typically based upon adjacency relationships (Rodrigues and Assunção, 2012). In this case, it will be given based on the adjacency matrix W defined in section 4.7.3. An issue with this is that the neighbourhood structure determines the degree of spatial smoothing (Duncan et al., 2017), and may over-smooth the response when a typical adjacency structure is used. It may be more practical to have an adaptive neighbourhood structure, particularly if we observe that the spatial structure evolves over time. Adapting the network structure will be discussed in the conclusions.

Another problem with the BYM model lies in its choice of priors. The unstructured component θ cannot be seen independently from the spatially structured component ϕ. θ is partially included within ϕ in the case of no spatial dependence, leading to an identifiability issue (Eberly and Carlin, 2000). Thus, as noted by Simpson et al. (2017), care should be taken to ensure that priors for τ_θ and τ_ϕ are dependent. Penalised complexity priors may be used in this case. These are discussed in section 4.7.4.4.

There is also a scaling issue with this method — it is important that hyperpriors used in one application will have the same interpretation when used in another (Sørbye and Rue, 2014). A modification of this, as noted by Riebler et al. (2016), states that the structured effect $\mathbf{b}$, where $\mathbf{b} = \phi + \theta$ must be scaled such that

$$\sigma_{GV}^2(b) = \frac{1}{\tau_b}$$

so that the precision τ_b represents the deviance from a constant level, and is independent of the underlying graph provided. Here GV represents the generalized variance, which is calculated by taking the geometric mean of the marginal variance. Then, Simpson et al. (2017) propose the use of a scaled structured component $\mathbf{u}_*$, with Q_* the scaled precision matrix, as scaled according to the above. According to Riebler et al. (2016), this gives rise to a modified $\mathbf{b}$,

$$\mathbf{b} = \frac{1}{\sqrt{\tau_{\mathbf{b}}}} \left(\sqrt{1 - \psi}\theta + \sqrt{\psi}\phi_* \right).,$$

for some mixing parameter $\psi \in [0,1]$. When $\psi = 1$, the model reduces to the Besag, while when $\psi = 1$ it reduces to pure overdispersion. This has covariance matrix

$$\text{Var}(b|\tau_b, \psi) = \tau_b^{-1} \left((1 - \psi)\mathbf{I} + \psi \mathbf{Q}_*^{-1} \right).$$

This representation leads to the scaling of $\mathbf{b}$ as a result, and the prior on τ_b will always have the same interpretation across graphs. The hyperparameters will also be identifiable under this representation. This representation of the BYM model can be implemented in R-INLA through the `bym2` choice of model.

4.7.4.4 Choice of priors & the penalised complexity (PC) prior

By inspection of R-INLA inputs, we observe that the default prior on the fixed effect is a normal with zero mean and a precision (1/variance) of 0.001, which is often not particularly informative. R-INLA places a gamma prior on the precisions by default. In particular, for the unstructured effect precision τ_v in the BYM model, a minimally informative prior is specified on $\log \tau_\nu$, i.e. $\log \tau_\nu \sim \log \text{Gamma}(1, 0.0005)$. Another weakly informative prior is set on the log of the spatially structured effect (τ_u), i.e. $\log \tau_u \sim \log \text{Gamma}(1, 0.0005^2)$. Whether this is sufficient or not depends on how much prior information is available to incorporate into the data.

R-INLA allows for the specification of different priors however. One possibility is to make use of the *penalised complexity* (PC) prior as seen in Simpson et al. (2017). This proposes a widely applicable method for the specification of priors on parameters which may be difficult or impossible to elicit from expert information. The approach makes use of the inherent nested structure for many model components. It defines the model component to be a flexible extension of a simple base model. It ensures that simpler models are given more importance, by defining a proper prior to put a penalty any complexity induced by deviating from this base model. Such priors are then formulated after inputting a user-defined scaling parameter for that model component. Depending on how the user tunes a single scaling parameter, this PC prior may be vague, weakly or strongly informative.

As pointed out in Simpson et al. (2017), the procedure is invariant to reparameterisation, since the prior is defined on the distance d, which is then transformed to the corresponding prior for ξ. This means the prior may be constructed without taking the specific parameterisation into account. The use of PC priors will be investigated in section 4.8.3.

4.7.5   Model selection

We can compare the models via the Deviance Information Criterion (DIC), which is a hierarchical modeling generalization of the Akaike information criterion (AIC) and Bayesian information criterion (BIC). We define the deviance as

$$D(\theta) = -2\log(p(y|\theta)) + C\,,$$

where y, are the data, θ, are the unknown parameters of the model and $p(y|\theta)$ is the likelihood function. C is a constant that cancels out. The DIC is calculated as or equivalently as

$$DIC = D(\bar{\theta}) + 2p_D.$$

where

$$p_D = \bar{D} - D(\bar{\theta}),$$

where $\bar{\theta}$ is the expectation of θ.

The comparison of DIC values between models is common method to test effects of a model. Complex models incur a penalty for the number of parameters, so under DIC these models might not be preferred. The model with the smaller DIC is usually seen as the better fit for the data. We use this to compare the goodness-of-fit of the models implemented in Section 4.8. The DIC requires approximate normality, but INLA overcomes this via the evaluation of the posterior mode of θ instead of the posterior mean. Note that, as the effective number of parameters is expected to be similar, we do not expect the DIC to vary greatly. As a result, we choose the more complex BYM model over the Besag model, as the Besag model only accounts for similarities between regions, but doesn't account for the fact that each region may vary slightly.

4.8   Results

In this section, we verify that the conditional independence assumption used in the first-order CAR model is an appropriate choice. Results of the BYM model shown in Equation 4.7.4 are presented, to see if there is a clear time trend to be found in the River Eden network using this representation. This is compared at first to the simple i.i.d. case, where no spatial structure is included in the data, to assess whether including a spatial component to the model provides any noticeable difference in the estimation.

We also provide a comparison to a dense precision matrix, in order to demonstrate the speed-up in computation provided by exploiting the network structure of the river. The spatial residual structure is examined and discussed.

4.8.1 Verifying the CAR assumption

The key modelling assumption made in the first-order CAR model in this setting is that river flows measured at a given station only depends on its upstream station and is independent of other stations given the direct upstream station. This is because the first-order CAR model obeys the spatial version of the Markov property. This first order assumption is somewhat restrictive, and its use in this approach requires justification. We can illustrate the validity of this assumption by checking the conditional independence relationship holds for some stations within the network.

4.8.1.1 Example 1: Stations 76002, 76003 and 76004

Station 76002 should be conditionally independent of station 76004 given station 76003 — see Figure 4-9. To check this, one can regress flows at 76002 against the flows at 76003 only, and the flows at 76002 against the flows at both 76003 and 76004 together. Including 76003 as a predictor alone should provide as much information as including both 76003 and 76004 — this is because all the information in flows from 76004 should be contained in the flows from 76003.

Using a linear regression between flows at 76003 and flows at 76002 results in an R^2 of 0.848, suggesting that the majority of the variability in flows at 76002 are explained by flows at 76003. We then also estimate the relationship of flows from both 76003 and 76004 on flows at 76002. If the conditional independence assumption holds perfectly, this will not provide any additional information. The fixed effect of flows from both stations appear to be significant, however the R^2 value increases only slightly to 0.8519, suggesting that the added station does not appear to improve the explanatory power of the model.

4.8.1.2 Example 2: Stations 76007, 76008 and 76011

We repeat this analysis for another set of stations on the network. Here, station 76007 should be conditionally independent of station 76011 given station 76008. Again, we

regress flows at 76007 against the flows at 76008 only, and the flows at 76007 against the flows at both 76008 and 76011 together.

The relationship between flows at 76008 and flows at 76007 again appears to be significant, with an R^2 value of 0.6649. We then also estimate the fixed effect of flows from both 76008 and 76011 on flows at 76007. Again, if the conditional independence assumption holds perfectly, this will not provide any additional information. The association between flows at both 76008 and 76011 with 76007 appear to be significant, but the R^2 value again only increases very slightly to 0.6656, suggesting that the addition of further information again does not contribute further in terms of explanatory power.

This suggests that while the first-order CAR model may not be a perfect assumption (as station flows further upstream appear to be significant), it seems that the flows are reasonably well explained by the flows at the most immediately upstream station. As a consequence, and as the most simple of these spatial structures, we will explore the first-order CAR model in this chapter. Relaxing of this assumption to a second-order could be of interest when considering networks of differing resolutions, and will be discussed briefly in future work.

4.8.2 Posterior population-wide effects for the river Eden network

All models in this section were run over the common period of 1978–2017.

4.8.2.1 Base case: trends estimated using the i.i.d. model

We wish to demonstrate the usefulness of including a spatial structure for the estimation of trends in daily mean flows on a river network. For comparison, we start by fitting a model of the form

$$y_{it} = \alpha + x_{it}^T \beta + \theta_i + \epsilon_i \tag{4.8.1}$$

where α is the overall intercept, β measures the effect of covariates z_i (which may be location-specific), $\theta \sim N(0, \sigma_\theta)$ is an i.i.d. (non-spatial) effect which accounts for between-station variability and $\epsilon \sim N(0, \sigma_\epsilon)$ the residual error. This is the simplest multilevel model that one could fit ignoring the spatial variability between stations. The resulting plot of the posterior fixed effect of time can be seen in Figure 4-11. From inspection of this plot, the fixed effect of time is highly likely to be greater than 0.001, because the median value and most of the mass of the distribution lies to the right of this value. As this is on a log scale, this corresponds to time adding at least 0.1%

136

Figure 4-11: The fixed effect of time on daily mean flows on the river Eden network using an i.i.d. model. Dashed lines correspond to the 2.5 and 97.5 percentiles of the posterior distribution.

to the daily mean flows each year (or a 1% increase in the median every 10 years), suggesting that daily mean flows are clearly increasing over time. The model DIC here is 288790.82.

A plot of the posterior standard deviations for the hyperparameters can be seen in Figure 4-12. The standard deviations correspond to the precisions on the fixed effect and the i.i.d. station random effect. As can be seen in Figure 4-11, the use of an i.i.d. effect only leads to a reasonably high level of variability in the estimate of the posterior standard deviations for both the Gaussian observation and i.i.d. station-specific effects. We will demonstrate next that the use of a spatial effect which respects the river network structure will provide an improvement in the variability of these estimates.

Posterior sd for Gaussian observations   **Posterior sd for station iid effect**
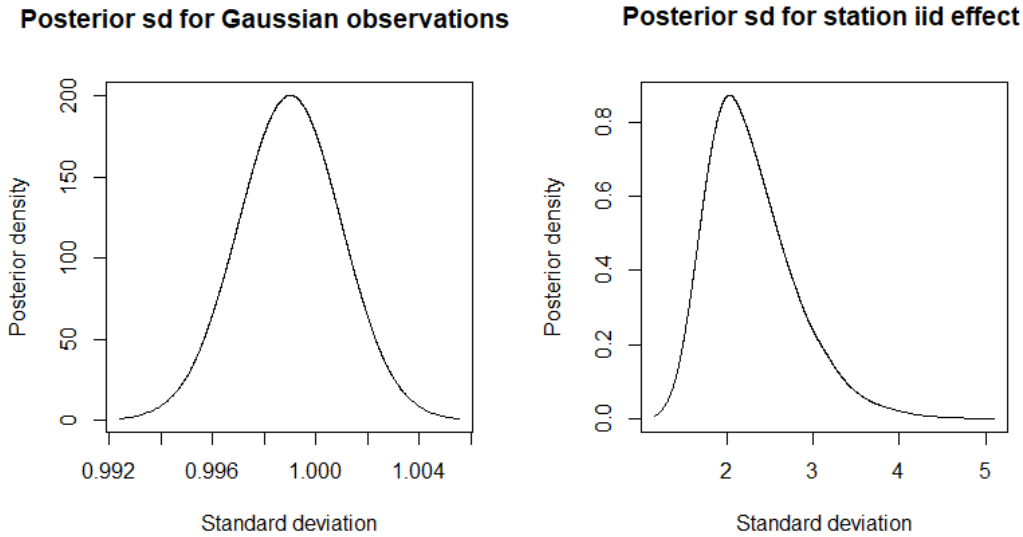
Figure 4-12: A plot of the posterior standard deviations for the Gaussian observations σ_ϵ (left) and the station i.i.d. effect σ_θ (right).

4.8.2.2   Trends estimated using a CAR model as a spatial random effect

The BYM model seen in Equation 4.7.4 and given by

$$y_i = \alpha + x_i^T \beta + \phi_i + \theta_i \tag{4.8.2}$$

is now used to incorporate a spatial effect (through a first-order CAR model) when estimating time trends present in the river Eden network.

Using a default prior in R-INLA, with a model fitted as in Equation 4.7.4, the posterior fixed effect of time on daily mean flows can be seen in Figure 4-13. This posterior fixed effect estimate is consistent with that seen in Figure 4-11 (i.e. a clear time trend is observed again). This is not surprising; however, we might expect to see some difference in the posterior standard deviations of the hyperparameters.

A plot of the posterior standard deviations for the hyperparameters can be seen in Figure 4-14. These are achieved using default priors as discussed previously. Now, the standard deviations correspond to the precisions on the fixed effect, the i.i.d. component of the BYM model, θ_i and the spatial random slope effect ϕ_i. The addition of a structured spatial effect along with the i.i.d. effect leads to a reduction in the variability in the estimates of the hyperparameters. This model, along with the enhanced ability to detect a trend compared to standard methods (see discussion of at-site approaches

138

Posterior fixed effect of Water Year

Figure 4-13: The fixed effect of time on daily mean flows on the river Eden network using a BYM model.

in Chapter 2), allows for a better explanation of the variability in the data compared to the i.i.d. approach and thus a more precise inference.

Note that it is clear that these default priors result in skewed posterior standard deviations for each of these quantities — while this is not a problem as one would expect to see such skew, it is apparent that these distributions appear to display some bimodality. As a consequence, the use of PC priors will instead be investigated. Note here that the model DIC is 288791.01, which is very similar to what was observed for the i.i.d. case, suggesting that the BYM model is more appropriate for the data at hand. We have observed here that the increase in daily mean river flow is consistent with the predicted increase in rainfall totals from climate projections (Bates, 2009). More work is necessary to disentangle components of this increase in flows, but it appears that more precipitation will correspond broadly to more river discharge.

Figure 4-14: A plot of the posterior standard deviations for the Gaussian observations σ_ϵ (left), the station i.i.d. effect σ_θ (centre) and the station spatial effect σ_ϕ (right) using default priors.

4.8.3 PC priors

Ideally, the PC prior is best used when one has prior knowledge of the underlying problem. Instead, we use the standard deviation of the response (the adjusted daily mean flows) to help set the scale of this prior. These priors are specified and the model run again in R-INLA. The results of this can be seen in Figure 4-15. This leads to the
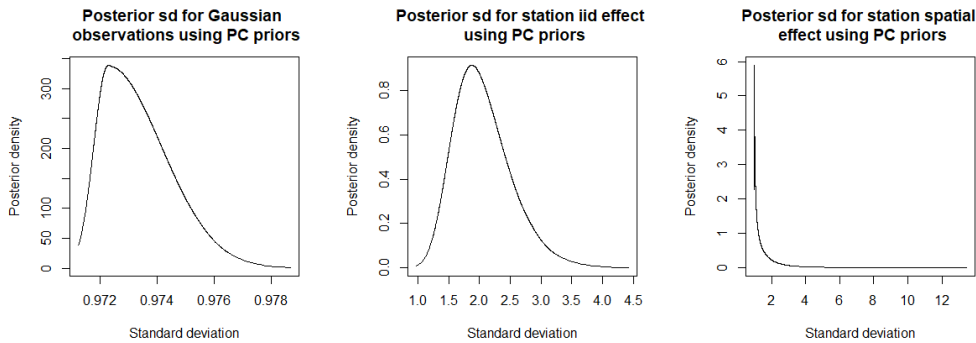
Figure 4-15: A plot of the posterior standard deviations for the Gaussian observations σ_ϵ (left), the station i.i.d. effect σ_θ (centre) and the station spatial effect σ_ϕ (right) using PC priors.

removal of the previously observed bimodality for both the posterior standard deviation of the Gaussian observations and the station i.i.d. effect, and the results appear to be more appropriate.

4.8.4 Model efficiency — comparison to full precision matrix

Suppose instead that we did not want to make use of the sparse precision matrix, but wanted to specify a general covariance. In this case, we will have a dense precision matrix, where every node in the network is connected to every other node. The resulting fixed effect estimates will be similar; however it is expected that the speed of calculations will be slower.

As the precision matrix is only 10×10, calculation times are small. However, we still note a considerable speed-up in calculation times. We obtain a model run time of 252 seconds using the sparse precision matrix method proposed in this chapter. On the other hand, using the dense precision matrix here takes 466 seconds — almost twice as slow as the CAR model approach. This will increase very quickly with any increases in number of river gauging stations, due to the $\mathcal{O}(n^2)$ order of storage and $\mathcal{O}(n^3)$ of computations required when using the dense matrix compared to $\mathcal{O}(n^{3/2})$ of a sparse precision matrix (Bakka et al., 2018). This suggests that not only does the CAR model approach discussed in this chapter lead to improved estimates through inclusion of a spatial effect, but also improved efficiency of computations when the network structure can be encoded in this way. The choice of a CAR model is advantageous both in respecting the river network structure and efficiency.

4.8.5 Posterior spatial structure from R-INLA models

First, we observe the mean of the adjusted log daily flows by station. This can be seen to the left of Figure 4-16. On inspection, it may be reasonable to think that there is
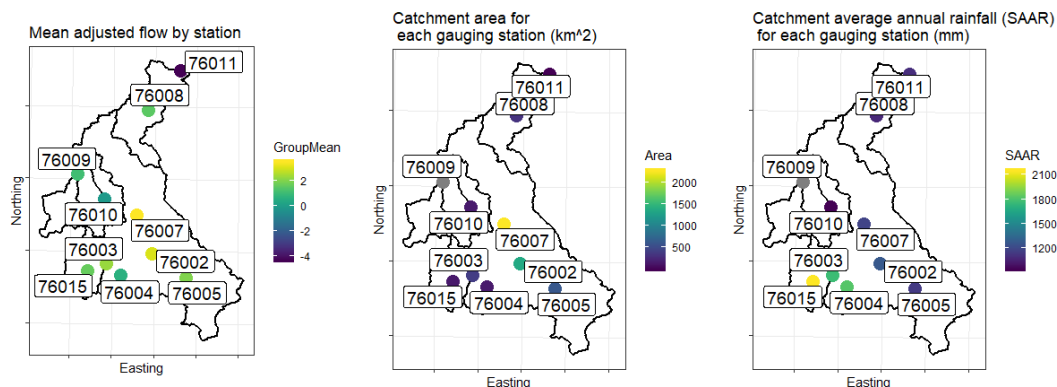
Figure 4-16: A plot of the mean adjusted flow by station (left), the catchment area (centre) and the average annualised rainfall (right).

not necessarily a clear spatial structure in the data, given that nearby stations do not necessarily have similar sized flows. This is due to the size of the catchment area (see the plot to the centre of Figure 4-16) and the river structure — smaller catchments such as 76011 will display smaller flows, and thus it would not be clear from this that, for example, stations 76008 and 76011 are related. However, we know from the river structure in Figure 4-9 that station 76011 drains into 76008 and consequently there must be some relationship between the two. As a result, use of the conditional independence-based models appear reasonable.

The posterior spatial effect $\phi + \theta$ using the BYM model from Equation 4.7.4 can be seen in Figure 4-17. It can be seen that there has been some smoothing due to the first-order CAR model, in that, as anticipated, this results in a spatial structure in which neighbours on the graph are more closely related than those which are not considered neighbours. Again, station 76011 appears to be an outlier compared to the other stations, however it is the smallest station by catchment area and the furthest away from the majority of the remaining gauging stations by distance, so this appears reasonable in practice.

Figure 4-17: The spatial random effect $\phi + \theta$ from fitting a CAR(1) model to the river Eden network

4.9 Attribution on a river network: possible avenues

We consider two potential approaches for the attribution of changes in a specific river flow network. One involves directly applying the causal checklist developed in Chapter 3. In this way, we can identify which large-scale variables are influencing daily mean river flows on the network in a long-term pattern. Alternatively, if we wish to understand how daily mean flows will evolve in the future, we can consider modelling the river network structure through the mean, by incorporating variables which represent the physical and directional nature of flows. This should help us to gain an understanding of what processes have a considerable influence on daily mean flows on a network.

4.9.1 Applying the causal checklist developed in Chapter 3

Given the parallels of the method discussed in this chapter to that of Chapter 2, it would be possible to use the causal checklist developed in Chapter 3 for the attribution of these time trends found in the river Eden network to some cause of interest. Attribution at a local level such as this is of interest as it may be beneficial when informing future flood defences in the area. One potential driver of flows in the north west of England could be the NAO index. This was found to not have a countrywide influence in Chapters 2 and 3, however it is associated with higher rainfall in the north west of England (O'Hare et al., 2014) and is potentially a local driver of changes in flows. Further potential local drivers of change include water management and land use management, which could be investigated using the checklist devised in Chapter 3. However, at such a local level, it is even more vital that such investigations are conducted in conjunction with an expert on the terrain of the area, and thus these drivers will not be investigated further in this chapter.

4.9.2 Modelling the network structure through the mean

While the covariance-based approach presented in this chapter is of benefit for estimating trends in daily mean flows and the attribution of trends at a regional level could be achieved by applying the checklist in Chapter 3, modelling the relationships between stations through the mean is more useful for understanding how river flows behave. It allows one to make predictions about future flows and thus is beneficial for water resources management. Additionally, O'Donnell et al. (2014) noted that it is often of

143

benefit to directly model trends in environmental observational studies through the mean, as this provides a more flexible means for modelling non-parametric trends. Ideally, one would simply model these relationships through the mean, selecting a number of appropriate explanatory variables for inclusion.

In order to be useful to practitioners, the model should respect the physical structure of a network (through the inclusion of direction and/or appropriate covariates), it should be interpretable by hydrologists and government agencies and explain a reasonable proportion of the variability in these flows. Many methods will be able to fulfil one or two of these three criteria with relative ease, however in order to be of benefit for the future management of water resources, we believe that is it necessary to fulfil all three of these criteria. By including appropriate covariates such as rainfall that are believed to influence flows, we represent the rainfall runoff process within the model. The physical structure of the network can be incorporated through modelling flows at a given station as a function of flows from stations directly upstream. Finally, an estimate of the remaining water entering the system can be modelled to be able to better represent this unknown quantity in future models.

The ultimate end goal of this approach is to obtain a more accurate estimation and attribution of flows in a given river network. However, some initial analysis shows that this is not so easily achieved. While upstream station flows and rainfall certainly appear to contribute towards flows at a station of interest when modelled separately, there is clear confounding when the two are modelled together. In addition, based on a preliminary analysis rainfall is not always a good predictor for daily mean flows, at least in a linear form — though considerably lagged versions of rainfall may prove to be a better predictor. While it is often expected that an increase in precipitation will result in an increase in flows, it is not so straightforward a relationship. Sharma et al. (2018) note the lack of evidence between increasing precipitation and floods, noting that there are even examples of reduced flooding magnitudes with increased precipitation. This perception of a direct relationship between the two assumes invariance of catchment conditions and that stream flows are generated solely from precipitation, which is not the case. Understanding the nature of this relationship between rainfall and daily mean flows is critical to this method, and is still under investigation (Blöschl et al., 2019).

Further issues with this method lie in accurately expressing the remaining water volume entering and exiting the system. Sometimes there is a greater volume of flow in stations directly upstream than the station of interest itself, suggesting some loss of flow to the external environment. There are also small streams which may not be

mapped or have a measurement which flow into a gauging station, increasing the flow volume. Obtaining an understanding of the process underlying this loss and gain of flow would require incorporating more physics into the system, which is not a simple procedure. It may be possible to have a hybrid statistical-PDE approach, as has been used in blood flow modelling (Azzimonti et al., 2015). This would involve using a PDE governing the flow behaviour as a prior to a statistical model such as the one described here. However, this will add considerable complexity to the method and may not be suitable for the purpose outlined here.

4.10 Chapter conclusions

In this chapter, we developed precision-based methods which encode conditional independence relationships between stations on a single river network. These conditional independent relationships are encoded through a sparse precision matrix resulting in fast inference methods. We have demonstrated that such methods lead to a reduction of the variability in posterior standard deviations for the station random effects compared to an i.i.d. structure. We also demonstrated the utility of PC priors for parameters which are difficult to elicit from expert information.

The use of a first-order CAR model for representing spatial relationships between these river stations is a novel approach which proved useful for assessing long-term trends in daily mean river flows on a single river network. The use of Bayesian methods in such a scenario have previously been proposed for annual maximum river flows in Great Britain (see Chapter 2 and Brady et al. (2019)) and Ireland (see Brady et al. (2018)). The benefit of using the CAR model on a network, however, is that using the sparse GMRF representation allows for fast inference using the INLA approach. Using a CAR structure is not possible on the countrywide data seen in Chapter 2, as the entire river network for Great Britain is not connected. However, in the case of a single river, it helps to encode relationships between neighbouring stations and speeds up inference compared to the use of the GP in a countrywide approach through the use of a sparse precision matrix. We demonstrated that inference using a dense precision matrix is almost twice as slow as this approach. It may be possible to combine the approaches seen in this chapter and Chapter 2, to have a countrywide approach which also respects the structure of individual river networks throughout. However, it was demonstrated in section 4.8.1 that a first-order CAR model may not be sufficient to fully explain flows at a particular station. It may be necessary to extend this to a second order CAR model, to ensure flows from stations further away are included within the approach.

A drawback of this is that it will result in a less sparse precision matrix, leading to slower computation times. On a small scale such as the one used in this study, however, it is not expected that this will dramatically increase computation time. The upper limit would be the full precision matrix, which was shown to take twice as long as the first-order CAR model.

It may also be possible to have a time-varying precision matrix, which encodes the spatial relationships between stations (which may drop in or out, thus changing the spatial structure of the network) at time t, however for the sake of simplicity and in order to assess trends across a longer time period in this chapter, we focused on one precision matrix only. This extension would be of interest for future developments of this work. It would also be of interest to explore the performance of the precision-based method developed in this chapter with the approaches of Ver Hoef et al. (2006) or O'Donnell et al. (2014), to see whether this approach will lead to a more accurate estimation of trends on a river network. Alternatively, a simulation study could be developed to demonstrate that the benefits of the CAR approach, as demonstrated on the river Eden network, scale to larger networks or can be applied to alternative applications.

This chapter has provided discussion on both Euclidean distances and stream distances for river networks. While the approach of Ver Hoef et al. (2006) develops a set of covariance functions in order to be able to incorporate stream distances, these distances between stations are not readily available from the NRFA. One must obtain the shapefile of a given river, then calculate the length between points on this river. If this were feasible, it may instead be possible to use the stream distance within this method. Alternatively, one could combine Euclidean and stream distance to benefit from the advantages of both. The stream distance represents the shortest distance between two river gauging stations along a stream network (Ver Hoef et al., 2006), and hence may be more appropriate for a true spatial representation of the network. The Euclidean distance can be easily extended (and appropriately weighted) to include catchment descriptors such as area and BFIHOST, thus providing a hydrological similarity measure. Combining both could further improve methods to accurately detect trends on a single river network.

While this method was useful for detecting trends on a network level, it only tells us whether changes have been occurring on the network. It is also of interest to learn how these changes are arising and what factors are driving this change. We noted in Section 4.9 how attribution at a local level might be explored. The various components that might contribute towards daily mean flows in stations on a network were briefly discussed, with the aim of obtaining a better understanding of how and why these flows

146

are changing over time. However, it was noted that constructing a statistical model that was both simple and respected the physical behaviour of river flows does not easily lead to sensible results for attribution of changes in these flows. Further exploration of these methods is needed for accurate attribution of changes in flows in the network.

Chapter 5

Conclusions and further work

The overall aim of this thesis was to develop spatial multilevel models for the detection and attribution of long-term changes in environmental studies to large-scale drivers of interest. Two common themes were explored throughout the thesis. The first of these was the use of spatial methods for the accurate detection of long-term changes in environmental observational studies, focusing on river flows in Great Britain. Typical at-site methods used in such studies are often poorly powered, meaning trends present in the data may be missed. The approaches presented here aim to overcome this problem, both at a countrywide and river-network basis. The second theme was the attribution of long-term, large-scale changes in environmental studies to some climate drivers of interest.

There are three major contributions of this thesis. The first of these is the first countrywide detection of long-term trends in annual maximum river flows in Great Britain, as seen in Chapter 2 and in Brady et al. (2019). This was achieved by using Bayesian spatial multilevel methods, which had not been previously applied in this setting. A major consequence of this result is that it successfully bridges the gap in estimated trends between climate change projections and observational river flow data. The second key contribution is the first attribution of these trends to large-scale climate drivers. In particular, a preliminary attempt at attribution was seen in Chapter 2, with a clear association seen between the East Atlantic (EA) index and annual maximum river flows. Chapter 3 focused on developing a systematic checklist for a rigorous causal assessment of drivers of environmental changes, demonstrating through this checklist that changes in annual maximum river flows can be attributed to the EA index. Finally, methods that utilised the spatial structure of the data under study were used to enhance the

ability to detect trends in river flows. This involved modelling the spatial structure through a Gaussian process in Chapter 2, when a countrywide approach is taken to estimating trends in annual maximum river flows. Such an approach can be used more generally in spatiotemporal modelling when the set of locations is fixed, and the spatial domain is not continuous. On the other hand, Chapter 4 looks at a single river network. The spatial structure of this network is considered as a graph and encoded through a first-order conditional autoregressive (CAR) model. This is a beneficial approach, as it respects the physical structure of the network and gives rise to a sparse precision matrix, leading to fast inference on the network. While this method was again applied to river flows, the approach could be applied to any application where there is a directional spatial structure. The results, benefits and limitations of the approaches used in each chapter are summarised in the following sections.

5.1 A first detection and "soft" attribution of countrywide trends in annual maximum river flows in Great Britain

In Chapter 2, we presented a study investigating the use of Bayesian multilevel, multivariate methods both for the detection of time trends in annual maximum river flows, and the attribution of such flows upon climate indices for a reference network of river gauging stations in Great Britain. At first, a typical at-site approach was investigated, to assess the ability of such methods to detect trends in annual maximum river flows should any exist. While it is expected that flooding events are increasing with time in Great Britain, just 25% of these gauging stations exhibited significant time trends. However, if measurements at all stations are independent from each other, 5% of stations could show a trend by chance due to multiple testing. This issue, along with the poor statistical power of such methods, led to a change in approach towards the use of Bayesian spatial multilevel models. The use of these models allowed for the pooling of information across river gauging stations. The spatial nature of relationships between these stations was modelled through the inclusion of a Gaussian process representing the spatial random effect. By increasing the sample size available to the model in this way, it was possible to detect clear countrywide time trends that had been missed with the at-site approach.

The use of a multilevel approach resulted in a first countrywide detection of a clear time trend in annual maximum flows in the country, in contrast with the scattered

signal observed in the at-site analysis. It was observed that the fixed effect of time on annual maximum river flows was highly likely to be greater than 0.6. As the method involved taking the log of the annual maximum flows, this corresponds to time adding at least 6% to annual maximum flows each year, which suggests that annual maximum flows have been increasing considerably over time. This in turn suggests that changes in flows match climate change projections (Bates, 2009), which predicts an increase in flooding.

Time itself does not cause such changes in annual maximum river flows, however. It represents a proxy for other unknown quantities, which themselves vary with annual maximum river flows in the same way as time. Time itself will not provide any useful information on potential drivers of changes in these flows, and thus detection should not be the sole focus of trend analyses in such studies. Instead, a switch in focus towards the combined detection and attribution of these changes to large-scale climate indices was proposed in Chapter 2. These climate indices themselves represent changes in climate. The effect of the East Atlantic (EA) and North Atlantic Oscillation (NAO) indices on annual maximum river flows was explored in Chapter 2, and a clear signal was obtained for the EA index in the univariate case, with an increase of 10% in the median when going from a neutral to positive EA value. This signal was still clear and positive when time was accounted for, suggesting that the EA index may have an impact on annual maximum river flows even when temporal confounding is accounted for. On the other hand, in the naive univariate approach, there appeared to be a clear link between NAO and annual maximum river flows, with an increase of 2% in the median of annual maximum flows when going from neutral to positive NAO. However, when time was included as a variable in the model (to address possible confounding between variables) resulted in this association overlapping with zero and becoming slightly, suggesting collinearity between this climate index and time. The difference observed going from the univariate to the multivariate setting demonstrated the necessity of a multivariate approach for accurately quantifying the true nature of associations between climate indices and annual maximum flows. A combined multilevel, multivariate approach towards attribution was deemed necessary to provide a clearer insight into changes in annual maximum river flows in Great Britain, and the strength of the relationship between these flows and climate indices.

This strong association between climate indices and annual maximum river flows was a crucial first step towards the accurate attribution of trends in annual maximum river flows motivated by Merz et al. (2012). In particular, the approach in Chapter 2 fulfilled Merz et al. (2012)'s "soft" attribution criteria, along with two of the three "hard"

attribution criteria. It was observed that detected changes in annual maximum river flows appear to be consistent with the East Atlantic index, and an uncertainty interval over the estimates was provided. However, the approach in essence only demonstrates clear associations between climate indices and annual maximum river flows, once time has been taken into account. Further evidence was deemed necessary to demonstrate a causal link between these climate indices and annual maximum river flows in Great Britain. For example, it remains to check whether detected changes are inconsistent with potential alternative drivers other than NAO, which was also investigated in Chapter 2. In Chapter 3, we proposed a systematic checklist for assessing causality in environmental observational studies such as these, which fulfils and goes beyond the scope of these criteria. The results of this approach are discussed in the next section.

It should be noted that the method presented in Chapter 2 relies on the assumption of a countrywide trend of both time and climate indices, and that the effects of climate indices are linear. These assumptions do not appear to be too strong for a preliminary approach, however. The use of non-parametric regression models, as seen in Villarini et al. (2009), to describe the relationship between annual maximum flow and the explanatory variable may be useful in relaxing the strong linearity assumption.

In conclusion, in Chapter 2 it was demonstrated that one can use multilevel models to detect countrywide trend at monitoring stations with short records by pooling information from nearby catchments, and that by using these more complex, yet still interpretable models, clear associations between these trends and some climate indices of interest are found. It was necessary to use a multivariate approach to uncover the effect of confounding between climate indices and time. By including time as a variable within the model for annual maximum river flows and a climate index of interest, we were able to observe the true relationship between such indices and flows. Additionally, the use of near-natural "benchmark" catchments in our approach ensured that any signal found in the data cannot be related to anthropogenic changes other than climate. The results from this approach could be beneficial to environmental monitoring agencies in understanding how floods may be changing in the future and be better prepared for this period of unknown extremes.

5.2 Causal methods for environmental observational studies

In Chapter 3, we presented a systematic checklist for the causal assessment of environmental observational studies. The challenges of assessing causal relationships in environmental observational studies were discussed, including the problems of confounding, limited data records and natural variability. Many standard causal methods are not easily applicable to such studies as a result. The need to consider non-data assumptions which help to build up the strength of evidence for a particular relationship was highlighted, and it was deemed necessary to incorporate additional expert knowledge in order to make causal statements about observational environmental studies.

Chapter 3 provided a review of a variety of causal tools from differing fields that were of potential benefit to environmental observational studies such as the case study of annual maximum river flows. Inspiration was taken from methods such as the Bradford Hill criteria, causal directed acyclic graphs (DAGs), the method of multiple working hypotheses and weighting of evidence methods, in order to propose a systematic approach for the attribution of long-term, large-scale changes in environmental processes. Using such methods, along with proposing a number of additional steps (incorporating both data and non-data assumptions), it was possible to develop a framework which assesses variables for a causal link with some outcome of interest, while ruling out further alternative causes. The Bradford Hill criteria was beneficial in framing the problem in terms of key causal questions, DAGs provided a means of visualising the overall causal structure (including potential confounding variables), while the method of multiple working hypotheses helped to rule out potential alternative hypotheses. Further steps such as accounting for sources of confounding and interpretability of model results to the end user were proposed as necessary for ensuring that the method was of use for decision-making bodies.

The approach proposed in Chapter 3 provided guidelines on how to conduct a thorough analysis to ensure all possible avenues are explored through a methodical checklist, which brings together multiple workflows by bridging the gap from climate scientists to the relevant authorities. The method helps to improve accuracy and interpretability for decision-making bodies, as well as being thorough in ruling out additional causes and controlling for confounding where possible. We also claim that it is necessary to sense check the proposed models with causal diagrams throughout the process, in order to consider as wide a range of possibilities as is feasible.

We applied this method to a case study of annual maximum river flows in Great Britain, where the aim was to attribute changes in these flows to climate change (represented by the East Atlantic and North Atlantic Oscillation indices). A thorough review was carried out, which addressed each item within the checklist. The result from this approach suggests that changes in these flows can, in part, be attributed to the East Atlantic index and hence climate change. The case study, through assessing the impact of the NAO index upon annual maximum river flows, also demonstrated the need to account for time-varying confounding and hence gain a further insight into the true associations between the exposure and the outcome of interest. This approach provided a thorough attribution of changes in annual maximum river flows in Great Britain to the EA index. In contrast to the approach seen in Chapter 2, this systematic checklist fulfils the three requirements set out by Merz et al. (2012) for "hard" attribution and goes beyond the scope of these requirements. It was observed that changes in annual maximum flows are consistent with changes in the EA index, and inconsistent with several potential alternative causes. In addition, a statement of uncertainty is provided through a credible interval on estimates. Further to these requirements, the checklist also ensured that results are interpretable to decision making bodies, that the method controlled for additional unmeasured confounding, and that the problem was framed in the context of causal diagrams, which are revisited throughout the analysis, to ensure diligence in considering potential alternative causes.

Though the case study focused specifically on annual maximum river flows in Great Britain, this approach is applicable to the attribution of changes in annual maximum flows anywhere. While the spatial correlation structure or climate variable under investigation may differ from the case study presented in Chapter 3, the overall approach and checklist for ruling out additional possibilities would hold in any location. The checklist can also be directly applied to many other environmental observational studies investigating the causal link between long-term, large-scale exposures and some outcome of interest. For example, this could easily be used in estimating the impact of climate change on water quality. However, the checklist would not apply in the case of feedback loop environmental problems, which cannot be summarised by a DAG (Pearce and Lawlor, 2016). Alternative approaches to the DAG must be considered to frame the problem in such a setting.

In summary, the checklist developed in Chapter 3 provides a simple, usable tool that provides an extension to the preliminary attribution analysis conducted in Chapter 2. The "hard" attribution criteria set out by Merz et al. (2012) has been satisfied, and it is possible to state that changes in annual maximum river flows in Great Britain

can, in part at least, be attributed to the East Atlantic index and hence to climate change. Another important contribution of this chapter is the general applicability of this causal approach to environmental observational studies, for which standard causal methods are often not appropriate.

5.3 Precision-based methods for modelling trends on a river network

In order to inform local flood defence management, is also of interest to look at a more regional level of modelling. The use of Bayesian methods described in Chapter 2 may also be useful for exploring long-term trends at a regional level, which again may be missed out in an at-site approach. However, adaptations are required. As the spatial domain of river gauging stations is not continuous on a countrywide level, the distance-based spatial effect was modelled in this approach through a Gaussian process. However, in the setting of a single river, the spatial structure may instead be encoded by exploiting the inherent graph structure of a single river network.

Chapter 4 focused on the detection of such changes in daily mean flows on a single river, by exploiting the inherent graph structure of this single river network. Prior approaches to respecting the directional structure of rivers were discussed, however these primarily made use of covariance-based methods. In this setting, one could use the Gaussian process approach seen in Chapter 2, however the covariance matrix encoding the spatial relationships between stations is dense, and using the Euclidean distance in this setting may not be appropriate Ver Hoef et al. (2006). Instead, we proposed the use of precision-based methods which encode conditional independence relationships between stations on a single river network. These conditional independent relationships are encoded through a sparse precision matrix resulting in faster inference methods than would be obtained with a dense covariance structure. We have demonstrated that such methods lead to an improvement of the variability in estimates of trends in these flows compared to an independent covariance structure. We also demonstrated the utility of PC priors for parameters which are difficult to elicit from expert information.

Modelling spatial relationships between river gauging stations on a network through a CAR model is a novel approach which proved beneficial for accurately assessing long-term trends in daily mean river flows on a single river network. As noted, the use of Bayesian methods in such a scenario have previously been proposed for annual maximum river flows in Great Britain (see Chapter 2 and Brady et al. (2019)). However,

155

the major advantage of using the CAR model on a network is that exploiting the sparse GMRF representation allows for faster inference than a dense matrix. It was not possible to use a CAR structure was not possible on the countrywide annual maximum river data seen in Chapter 2, as the entire river network for Great Britain is not connected. However, in the case of a single river, it helps to encode relationships between neighbouring gauging stations and speeds up inference compared to the use of the Gaussian process in a countrywide approach through the use of a sparse precision matrix. It was seen that, even for a small number of river gauging stations (10) and the corresponding 10×10 precision matrix, inference using a full precision matrix is approximately twice as slow as using the sparse precision approach from Chapter 4. It may be possible to combine the approaches seen in this chapter and Chapter 2, to have a countrywide approach which also respects the structure of individual river networks throughout and benefits from computational speed-up at a regional level.

5.4 Future work

There are a number of possible future avenues for this work that may be worth exploring. These are discussed in this section, and are grouped by future attribution approaches and network-based methods.

5.4.1 Additional attribution methods

The systematic checklist presented in Chapter 3 primarily discussed temporal confounding in attribution studies, but spatial confounding is another common issue in environmental observational studies, for example when estimating impact of climate change on tree defoliation (Augustin et al., 2007). Models for spatial observational data will typically include both a fixed effect for the covariate(s) of interest along with random effects representing the spatial correlation in the data, as is the case throughout this thesis. Adding these spatial random effects aims to reduce bias from unmeasured covariates. However, if the response and these covariates are also spatially correlated, including a spatial random effect may lead to confounding of the covariate effect (Thaden and Kneib, 2018). The impact of this spatial confounding may be such that, as seen in the case of temporal confounding in Chapter 2, a coefficient which is significant with no spatial random effect can become non-significant if a spatial random effect is included (Reich et al., 2006). This confounding ensures it is near-impossible

to distinguish between spatial effects modelled through location co-ordinates and site-specific effects. This spatial confounding has primarily been investigated in the case of areal spatial data, where the main approach involves constraining of spatial random effects to be orthogonal to the fixed effects (Hodges and Reich, 2010). A similar approach was discussed by (Hanks et al., 2015) for the case of continuous spatial data, however the method assumes that all variation in the same direction as the fixed effects can be attributed to these effects. Thaden and Kneib (2018) suggest using structural equation modelling in order to separate the direct effect of a covariate from the indirect effect which arises from correlations with other variables. This method allows space to have an effect on both the exposure and the response simultaneously, compared to with prior methods allowing for space to have an effect on the response only. The method proposed in Chapter 3 could be extended through an item accounting for spatial confounding with this method, and used for environmental observational studies where it is anticipated that spatial confounding may occur.

It was also clear from the case study in Chapter 3 that collinearity, in particular with time, is not easily overcome in environmental observational studies in which potential drivers are affected by climate change. As a consequence, it may not always be feasible to attribute changes in some outcome to a variable of interest using the systematic checklist proposed in Chapter 3. It was not possible to attribute changes in annual maximum river flows in Great Britain to the North Atlantic Oscillation index due to its collinearity with time — it is possible that there is a causal relationship, but using these methods it is not possible to conclude this. Further work is necessary to accurately uncover the true causal nature of relationships in the presence of collinear variables.

Chapter 4 provided a means of detecting long-term changes in daily mean flows on a single river network, however it did not provide any insight into the underlying causes of such changes. Attribution at a local level was discussed in Section 4.9, but care must be taken at such a fine spatial scale to consult with regional experts on potential drivers of change. This in turn may help to inform future water management practices at a regional level. Additionally, the various components that might contribute towards daily mean flows in stations on a network were noted in Chapter 4. Preliminary models were fitted in order to understand how and why these flows are changing over time. Variables investigated included upstream station flows, rainfall and remaining water entering the system. However, it was found that constructing a simple statistical model which respected the physical behaviour of river flows does not easily lead to sensible results for attribution of changes in these flows. Clear confounding was

observed between these variables, due to their highly collinear nature. Additionally, the relationship between rainfall and river flows is complex in nature. It is anticipated that increasing precipitation results in an increase in river flows, however Sharma et al. (2018) discuss the lack of evidence between increasing precipitation and floods, noting that there are even examples of reduced flooding magnitudes with increased precipitation. It is not possible to sensibly include rainfall in this approach without a thorough understanding of the true nature of this relationship between precipitation and daily mean flows, and further investigation must be carried out to accurately incorporate precipitation in future statistical models (Blöschl et al., 2019).

It was also noted in Chapter 4 that it is difficult to accurately express the remaining water volume entering and exiting the system. There is often loss of flows to the external environment, and small ungauged streams may exist which flow into a river gauging station, leading to an increase in river flow volume. Accurately representing these physical processes is complex to achieve, and may require the use of hybrid statistical-PDE approaches, as has been used in blood flow modelling (Azzimonti et al., 2015). This would require the inclusion of a PDE governing the flow behaviour as a prior to a statistical model. It is important to note however that this will further increase the complexity of any model, meaning it may no longer be interpretable to decision-making bodies, and thus not of use for future water management practices. Care must be taken to achieve a fine balance between model simplicity and explanatory power, in order for any such method to be of practical use.

5.4.2   Network-based methods

It was demonstrated in Chapter 4, that the use of a first-order CAR model may not always be enough to entirely explain flows on the river network. Instead, one could extend this to a second order CAR model, to ensure flows from stations further away are included within the approach. While using a second-order model may help to better explain flows on the monitoring network, a disadvantage is that it will result in a less sparse precision matrix, leading to slower computation times. However, as this is on a small scale, it is not expected that this will dramatically increase computation time for problems of the size of that presented in Chapter 4. The upper limit would be the full precision matrix, which was shown to take twice as long as the first-order CAR model. However, for a larger monitoring network, the trade-off between explanatory power and sparsity of the precision matrix will become more important and require careful thought.

It would also be of interest to compare the spatial structure of two differing time periods, to see whether the relationship between river gauging stations in the network has evolved over time. It may be possible to have a time-varying precision matrix, which encodes the spatial relationships between stations (which may drop in or out, thus changing the spatial structure of the network) at time t, however for the sake of simplicity and in order to assess trends across a longer time period in Chapter 4, the precision matrix was only spatial in nature. Knight et al. (2016) introduced a set of flexible models for network time series such as these known as the network autoregressive (integrated) moving average (NARIMA) processes. It was noted that such models could be used when the structure of the graph is evolving over time. Using this approach would provide an interesting extension to the work presented in Chapter 4.

In Chapter 4, the use of an unweighted CAR model was discussed, however the adjacency matrix describing the conditional relationships between river stations could be weighted according to the distance between sites, so that closer locations might have a greater influence over flows at a given river station. In this case, the entries of the adjacency matrix W may be constructed as follows:

$$
w_{ij} = \begin{cases} \omega_{ij} & \text{if an edge exists between nodes } x_i \text{ and } x_j \\ 0 & \text{if no edge exists between nodes } x_i \text{ and } x_j \end{cases}
$$

where ω_{ij} is the weight based upon the distance between nodes x_i and x_j. However, as noted in Chapter 4 and by Ver Hoef et al. (2006), the Euclidean distance may not be an appropriate choice for the network, so it is proposed that the stream distance be used. Stream distance between stations is not readily available from the NRFA. It would be necessary to obtain the shapefile of a given river, then calculate the length between points on this river. Should this be feasible, this would be a simple extension to the method which weights the spatial relationships between stations (as computed in Chapter 4) by stream distance.

5.5 Summary

A number of methods have been discussed and developed in this thesis which are of benefit to various research communities and authorities. The use of Bayesian methods for the more accurate detection of changes in river flows at a countrywide and local level are beneficial to environmental authorities. We have demonstrated that there is clear evidence of increasing trends of annual maximum river flows in Great Britain,

verifying climate change projections as seen in Bates (2009). As a consequence, care must be taken to ensure that appropriate flood defence systems are put in place to withstand future flooding events.

Through the use of a reference network of river gauging stations to avoid the interference of anthropogenic changes, we were also able to demonstrate clear associations between the East Atlantic index and annual maximum river flows. The causal checklist discussed in this thesis showed that this appears to be a causal link, suggesting that changes in the EA index are contributing towards this increase in annual maximum river flows in Great Britain. As this index is impacted by climate change, it is apparent that climate change is driving these changes in flooding events in Great Britain. The development of this causal checklist for the attribution of drivers of change in environmental observational studies is a useful tool for researchers in a number of areas in which climate change is the primary focus, such as air pollution, sea surface temperature and flooding.

While the evidence for change in flood risk at a national scale is of interest for large-scale planning, this might not translate directly into local changes in river flow. A different approach, which better exploits the known structure of the river measuring network, can then be employed for the study or regional changes. These changes in river flow patterns can also be seen at the regional level, where evidence of an increase in daily mean flows on the river Eden network over time was seen. Understanding these changes in the daily mean is important for water companies and future water resources management in general. In order to make this a more useful tool, steps should be taken to model the seasonal variation. It has been demonstrated for northern England (Fowler et al., 2003) that, while it is expected that increased winter rainfall will result in improved water resource reliability, there will also be an increased vulnerability to drought due to decreased summer rainfall. Modelling these daily trends across Great Britain with a renewed focus on seasonal changes will help to understand future risks to water resources at a countrywide level.

Bibliography

Auffhammer, M. and Vincent, J. R. (2012), 'Unobserved time effects confound the identification of climate change impacts', *Proceedings of the National Academy of Sciences* **109**(30), 11973–11974.

Augustin, N. H., Lang, S., Musio, M. and Von Wilpert, K. (2007), 'A spatial model for the needle losses of pine-trees in the forests of Baden-Württemberg: an application of Bayesian structured additive regression', *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **56**(1), 29–50.

Augustin, N. H., Musio, M., von Wilpert, K., Kublin, E., Wood, S. N. and Schumacher, M. (2009), 'Modeling spatiotemporal forest health monitoring data', *Journal of the American Statistical Association* **104**(487), 899–911.

Azzimonti, L., Sangalli, L. M., Secchi, P., Domanin, M. and Nobile, F. (2015), 'Blood flow velocity field estimation via spatial regression with PDE penalization', *Journal of the American Statistical Association* **110**(511), 1057–1071.

Bakka, H., Rue, H., Fuglstad, G.-A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D. and Lindgren, F. (2018), 'Spatial modeling with R-INLA: A review', *Wiley Interdisciplinary Reviews: Computational Statistics* **10**(6), e1443.

Banerjee, S., Carlin, B. P. and Gelfand, A. E. (2014), *Hierarchical Modeling and Analysis for Spatial Data*, Chapman and Hall/CRC.

Barker, L., Hannaford, J., Muchan, K., Turner, S. and Parry, S. (2016), 'The winter 2015/2016 floods in the UK: a hydrological appraisal', *Weather* **71**(12), 324–333.

Barry, R. R. and Ver Hoef, J. M. (1996), 'Blackbox kriging: spatial prediction without specifying variogram models', *Journal of Agricultural, Biological, and Environmental Statistics* pp. 297–322.

Bates, B. (2009), *Climate Change and Water: IPCC technical paper VI*, World Health Organization.

Belkhiri, L. and Narany, T. S. (2015), 'Using multivariate statistical analysis, geostatistical techniques and structural equation modeling to identify spatial variability of groundwater quality', *Water Resources Management* **29**(6), 2073–2089.

Besag, J., York, J. and Mollié, A. (1991), 'Bayesian image restoration, with two applications in spatial statistics', *Annals of the institute of statistical mathematics* **43**(1), 1–20.

Betancourt, M. (2017), 'A conceptual introduction to Hamiltonian Monte Carlo', *arXiv preprint arXiv:1701.02434* .

Betancourt, M. and Girolami, M. (2015), 'Hamiltonian Monte Carlo for hierarchical models', *Current trends in Bayesian methodology with applications* **79**, 30.

Bishop, C. M. (2006), *Pattern Recognition and Machine Learning*, Springer.

Blangiardo, M. and Cameletti, M. (2015), *Spatial and Spatio-temporal Bayesian Models with R-INLA*, John Wiley & Sons.

Blöschl, G., Bierkens, M. F. P., Chambel, A., Cudennec, C., Destouni, G., Fiori, A., Kirchner, J. W., McDonnell, J. J., Savenije, H. H. G., Sivapalan, M. et al. (2019), 'Twenty-three unsolved problems in hydrology (UPH)–a community perspective', *Hydrological Sciences Journal* **64**(10).

Bollen, K. A. (1989), *Structural Equations with Latent Variables*, John Wiley & Sons, New York.

Bowers, M. C., Tung, W. and Gao, J. (2012), 'On the distributions of seasonal river flows: lognormal or power law?', *Water Resources Research* **48**(5).

Bradford, R. B. and Marsh, T. J. (2003), Defining a network of benchmark catchments for the UK, *in* 'Proceedings of the Institution of Civil Engineers-Water and Maritime Engineering', Vol. 156, pp. 109–116.

Brady, A., Faraway, J. and Prosdocimi, I. (2018), 'Attribution of large-scale drivers of peak river flows in Ireland', *Proceedings of the 33rd International Workshop on Statistical Modelling 2018, IWSM 2018* pp. 54–58.
URL: *https://researchportal.bath.ac.uk/en/publications/attribution-of-large-scale-drivers-of-peak-river-flows-in-ireland*

Brady, A., Faraway, J. and Prosdocimi, I. (2019), 'Attribution of long-term changes in peak river flows in Great Britain', *Hydrological Sciences Journal* **64**(10), 1159–1170. **URL:** *http://www.tandfonline.com/doi/abs/10.1080/02626667.2019.1628964*

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, B. and Riddell, A. (2016), 'Stan: A probabilistic programming language', *Journal of Statistical Software* **20**, 1–37.

Cartwright, N. (1989), *Nature's Capacities and Their Measurement*, Clarendon Press.

Chamberlin, T. C. (1890), 'The method of multiple working hypotheses', *Science* **15**(366), 92–96.

Chatterton, J., Clarke, C., Daly, E., Dawks, S., Elding, C., Fenn, T., Hick, E., Miller, J., Morris, J., Ogunyoye, F. et al. (2016), 'The costs and impacts of the winter 2013 to 2014 floods', *Environment Agency, UK Government* pp. 5–6.

Chiew, F. H. S., Kirono, D. G. C., Kent, D. M., Frost, A. J., Charles, S. P., Timbal, B., Nguyen, K. and Fu, G. (2010), 'Comparison of runoff modelled using rainfall from different downscaling methods for historical and future climates', *Journal of Hydrology* **387**(1-2), 10–23.

Clark, M. P., Kavetski, D. and Fenicia, F. (2011), 'Pursuing the method of multiple working hypotheses for hydrological modeling', *Water Resources Research* **47**(9).

Climate Home News (2015), 'Global weather extremes underline climate threat, say experts', `https://www.climatechangenews.com/2015/12/29/global-weather-extremes-underline-climate-threat-say-experts/`. Accessed: 2018-02-20.

Coles, S. (2001), *An Introduction to Statistical Modeling of Extreme Values*, Vol. 208, Springer.

Cologna, V., Bark, R. H. and Paavola, J. (2017), 'Flood risk perceptions and the UK media: Moving beyond "once in a lifetime" to "be prepared" reporting', *Climate Risk Management* **17**, 1–10.

Cressie, N., Frey, J., Harch, B. and Smith, M. (2006), 'Spatial prediction on a river network', *Journal of Agricultural, Biological, and Environmental Statistics* **11**(2), 127.

Curriero, F. C. (2006), 'On the use of non-Euclidean distance measures in geostatistics', *Mathematical Geology* **38**(8), 907–926.

Daniel, R. M., De Stavola, B. L. and Vansteelandt, S. (2016), 'Commentary: The formal approach to quantitative causal inference in epidemiology: misguided or misrepresented?', *International Journal of Epidemiology* **45**(6), 1817–1829.

Davey-Smith, G. and Ebrahim, S. E. (2016), *International Journal of Epidemiology. Special Issue: Causality in Epidemiology* **45(6)**.

Dawson, R. J., Thompson, D., Johns, D., Gosling, S., Chapman, L., Darch, G., Watson, G., Powrie, W., Bell, S., Paulson, K., Hughes, P. and Wood, R. (2016), UK Climate Change Risk Assessment Evidence Report: Chapter 4, Infrastructure, Technical report.

Della-Marta, P. M., Mathis, H., Frei, C., Liniger, M. A., Kleinn, J. and Appenzeller, C. (2009), 'The return period of wind storms over Europe', *International Journal of Climatology* **29**(3), 437–459.

Di Falco, S., Yesuf, M., Kohlin, G. and Ringler, C. (2012), 'Estimating the impact of climate change on agriculture in low-income countries: Household level evidence from the Nile Basin, Ethiopia', *Environmental and Resource Economics* **52**(4), 457–478.

Diebold, F. X. (2001), *Elements of Forecasting*, 2nd edn, Cincinnati: South Western.

Diggle, P. J., Menezes, R. and Su, T. (2010), 'Geostatistical inference under preferential sampling', *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **59**(2), 191–232.

Diggle, P., Liang, K. Y. and Zeger, S. L. (1994), *Analysis of Longitudinal Data*, Oxford Statistical Science Series, Clarendon Press.

Dixon, H. (2010), 'Managing national hydrometric data: from data to information', *Global Change: Facing Risks and Threats to Water Resources* **340**, 451–458.

Dixon, H., Hannaford, J. and Fry, M. J. (2013), 'The effective management of national hydrometric data: experiences from the United Kingdom', *Hydrological Sciences Journal* **58**(7), 1383–1399.

Duane, S., Kennedy, A. D., Pendleton, B. J. and Roweth, D. (1987), 'Hybrid Monte Carlo', *Physics Letters B* **195**(2), 216–222.

Duncan, E. W., White, N. M. and Mengersen, K. (2017), 'Spatial smoothing in Bayesian models: a comparison of weights matrix specifications and their impact on inference', *International Journal of Health Geographics* **16**(1), 47.

Eberly, L. E. and Carlin, B. P. (2000), 'Identifiability and convergence issues for Markov chain Monte Carlo fitting of spatial models', *Statistics in medicine* **19**(17-18), 2279–2294.

Elliott, L. P. and Brook, B. W. (2007), 'Revisiting chamberlin: multiple working hypotheses for the 21st century', *BioScience* **57**(7), 608–614.

Faraway, J. J. (2014), *Linear Models with R*, Chapman and Hall/CRC.

Fedak, K. M., Bernal, A., Capshaw, Z. A. and Gross, S. (2015), 'Applying the Bradford Hill criteria in the 21st century: How data integration has changed causal inference in molecular epidemiology', *Emerging themes in epidemiology* **12**(1), 14.

Fowler, H. J., Kilsby, C. G. and O'Connell, P. E. (2003), 'Modeling the impacts of climatic change and variability on the reliability, resilience, and vulnerability of a water resource system', *Water Resources Research* **39**(8).

Fowler, H. J. and Wilby, R. L. (2010), 'Detecting changes in seasonal precipitation extremes using regional climate model projections: Implications for managing fluvial flood risk', *Water Resources Research* **46**(3).

Gallacher, K. M. (2016), Using river network structure to improve estimation of common temporal patterns, PhD thesis, University of Glasgow.

Gardner, B., Sullivan, P. J. and Lembo, Jr, A. J. (2003), 'Predicting stream temperatures: geostatistical model comparison using alternative distance metrics', *Canadian Journal of Fisheries and Aquatic Sciences* **60**(3), 344–351.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A. and Rubin, D. B. (2013), *Bayesian Data Analysis*, Chapman and Hall/CRC.

Gelman, A. and Hill, J. (2006), *Data Analysis using Regression and Multilevel/Hierarchical Models*, Cambridge University Press.

Gelman, A., Hill, J. and Yajima, M. (2012), 'Why we (usually) don't have to worry about multiple comparisons', *Journal of Research on Educational Effectiveness* **5**(2), 189–211.

Gelman, A., Jakulin, A., Pittau, M. G., Su, Y.-S. et al. (2008), 'A weakly informative default prior distribution for logistic and other regression models', *The Annals of Applied Statistics* **2**(4), 1360–1383.

Gelman, A., Simpson, D. and Betancourt, M. (2017), 'The prior can often only be understood in the context of the likelihood', *Entropy* **19**(10), 555.

Goldstein, H., Carpenter, J. R. and Browne, W. J. (2014), 'Fitting multilevel multivariate models with missing data in responses and covariates that may include interactions and non-linear terms', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **177**(2), 553–564.

Grace, J. B. (2008), 'Structural equation modeling for observational studies', *The Journal of Wildlife Management* **72**(1), 14–22.

Granger, C. W. J. (1969), 'Investigating causal relations by econometric models and cross-spectral methods', *Econometrica: Journal of the Econometric Society* pp. 424–438.

Greenland, S., Pearl, J. and Robins, J. M. (1999), 'Causal diagrams for epidemiologic research', *Epidemiology* pp. 37–48.

Guimarães Nobre, G., Jongman, B., Aerts, J. and Ward, P. J. (2017), 'The role of climate variability in extreme floods in Europe', *Environmental Research Letters* **12**(8).

Hanks, E. M., Schliep, E. M., Hooten, M. B. and Hoeting, J. A. (2015), 'Restricted spatial regression in practice: Geostatistical models, confounding, and robustness under model misspecification', *Environmetrics* **26**(4), 243–254.

Hannaford, J. (2015), 'Climate-driven changes in UK river flows: A review of the evidence', *Progress in Physical Geography* **39**(1), 29–48.

Hannaford, J. and Marsh, T. (2006), 'An assessment of trends in UK runoff and low flows using a network of undisturbed catchments', *International Journal of Climatology* **26**(9), 1237–1253.

Hannaford, J. and Marsh, T. J. (2008), 'High-flow and flood trends in a network of undisturbed catchments in the UK', *International Journal of Climatology* **28**(10), 1325–1338.

Harrigan, S., Hannaford, J., Muchan, K. and Marsh, T. J. (2017), 'Designation and trend analysis of the updated UK Benchmark Network of climate sensitive river flow stations: The UKBN2 dataset', *Hydrology Research* **48**(5).

166

Harrigan, S., Murphy, C., Hall, J., Wilby, R. L. and Sweeney, J. (2013), 'Attribution of detected changes in streamflow using multiple working hypotheses', *Hydrology and Earth System Sciences* **10**, 12373–12416.

Hastie, T. and Tibshirani, R. (2017), *Generalized Additive Models*, CRC Press.

Hill, A. B. (1965), 'The environment and disease: Association or causation?', *Journal of the Royal Society of Medicine* **58**(5), 295–300.

Hodges, J. S. and Reich, B. J. (2010), 'Adding spatially-correlated errors can mess up the fixed effect you love', *The American Statistician* **64**(4), 325–334.

Hodgkins, G. A., Whitfield, R. H., Burn, D. H., Hannaford, J., Renard, B., Stahl, K., Fleig, A. K., Madsen, H., Mediero, L., Korhonen, J. et al. (2017), 'Climate-driven variability in the occurrence of major floods across North America and Europe', *Journal of Hydrology* **552**, 704–717.

Höfler, M. (2005), 'Causal inference based on counterfactuals', *BMC medical research methodology* **5**(1), 28.

Huntingford, C., Marsh, T., Scaife, A. A., Kendon, E. J., Hannaford, J., Kay, A. L., Lockwood, M., Prudhomme, C., Reynard, N. S., Parry, S. et al. (2014), 'Potential influences on the United Kingdom's floods of winter 2013/14', *Nature Climate Change* **4**(9), 769.

Integrated Climate Data Center (2011), 'Climate Indices', `https://icdc.cen.uni-hamburg.de/1/daten/climate-indices.html`. Accessed: 2018-02-20.

Jeon, J. (2015), 'The strengths and limitations of the statistical modeling of complex social phenomenon: Focusing on SEM, path analysis, or multiple regression models', *Int J Soc Behav Educ Econ Bus Ind Eng* **9**(5), 1594–1602.

Kendall, M. G. (1948), *Rank Correlation Methods*, Charles Griffin.

Kingston, D. G., Lawler, D. M. and McGregor, G. R. (2006), 'Linkages between atmospheric circulation, climate and streamflow in the northern North Atlantic: Research prospects', *Progress in Physical Geography* **30**(2), 143–174.

Kjeldsen, T. R. and Jones, D. A. (2009), 'An exploratory analysis of error components in hydrological regression modeling', *Water Resources Research* **45**(2).

Kjeldsen, T. R. and Jones, D. A. (2010), 'Predicting the index flood in ungauged UK catchments: On the link between data-transfer and spatial model error structure', *Journal of Hydrology* **387**(1-2), 1–9.

Kjeldsen, T. R., Jones, D. A. and Bayliss, A. C. (2008), *Improving the FEH Statistical Procedures for Flood Frequency Estimation*, Environment Agency, United Kingdom.

Knight, M. I., Nunes, M. A. and Nason, G. P. (2016), 'Modelling, detrending and decorrelation of network time series', *arXiv preprint arXiv:1603.03221* .

Koller, D. and Friedman, N. (2009), *Probabilistic Graphical Models: Principles and Techniques*, Adaptive Computation and Machine Learning series, MIT Press.

Koutsoyiannis, D. and Montanari, A. (2015), 'Negligent killing of scientific concepts: the stationarity case', *Hydrological Sciences Journal* **60**(7-8), 1174–1183.

Krieger, N. and Davey-Smith, G. (2016), 'The tale wagged by the DAG: broadening the scope of causal inference and explanation for epidemiology', *International Journal of Epidemiology* **45**(6), 1787–1808.

Landau, L., Lifshitz, E., Sykes, J. and Bell, J. (1976), *Mechanics*, Butterworth-Heinemann, Elsevier Science.

Lins, H. F. and Slack, J. R. (2005), 'Seasonal and regional characteristics of US streamflow trends in the United States from 1940 to 1999', *Physical Geography* **26**(6), 489–501.

Macdonald, N. and Sangster, H. (2017), 'High-magnitude flooding across Britain since AD 1750', *Hydrology and Earth System Sciences* **21**(3), 1631–1650.

Mahlstein, I., Spirig, C., Liniger, M. A. and Appenzeller, C. (2015), 'Estimating daily climatologies for climate indices derived from climate model data and observations', *Journal of Geophysical Research: Atmospheres* **120**(7), 2808–2818.

Mallakpour, I. and Villarini, G. (2015), 'The changing nature of flooding across the central United States', *Nature Climate Change* **5**(3), 250–254.

Mallakpour, I. and Villarini, G. (2016), 'Investigating the relationship between the frequency of flooding over the central United States and large-scale climate', *Advances in Water Resources* **92**, 159–171.

Mann, H. B. (1945), 'Nonparametric tests against trend', *Econometrica: Journal of the Econometric Society* pp. 245–259.

Mann, M. E., Bradley, R. S. and Hughes, M. K. (1998), 'Global-scale temperature patterns and climate forcing over the past six centuries', *Nature* **392**(6678), 779–787.

168

Marra, G. and Wood, S. N. (2012), 'Coverage properties of confidence intervals for generalized additive model components', *Scandinavian Journal of Statistics* **39**(1), 53–74.

Mediero, L., Santillán, D., Garrote, L. and Granados, A. (2014), 'Detection and attribution of trends in magnitude, frequency and timing of floods in Spain', *Journal of Hydrology* **517**, 1072–1088.

Merz, B., Aerts, J. C. J. H., Arnbjerg-Nielsen, K., Baldi, M., Becker, A., Bichet, A., Blöschl, G., Bouwer, L. M., Brauer, A., Cioffi, F. et al. (2014), 'Floods and climate: emerging perspectives for flood risk assessment and management', *Natural Hazards and Earth System Sciences* **14**(7), 1921–1942.

Merz, B., Vorogushyn, S., Uhlemann, S., Delgado, J. and Hundecha, Y. (2012), 'HESS opinions: "More efforts and scientific rigour are needed to attribute trends in flood time series"', *Hydrology and Earth System Sciences* **16**(5), 1379–1387.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. and Teller, E. (1953), 'Equation of State Calculations by Fast Computing Machines', *The Journal of Chemical Physics* **21**(6), 1087–1092.

Monestiez, P., Bailly, J.-S., Lagacherie, P. and Voltz, M. (2005), 'Geostatistical modelling of spatial processes on directed trees: Application to fluvisol extent', *Geoderma* **128**(3-4), 179–191.

Murphy, K. (2012), *Machine Learning: A Probabilistic Perspective*, Adaptive Computation and Machine Learning series, MIT Press.

National Climatic Data Center, N. (2017), 'Climate Indices: Monthly Atmospheric and Ocean Time Series', `https://www.esrl.noaa.gov/psd/data/climateindices/list/`. Accessed: 2017-01-22.

National River Flow Archive (2017), 'National River Flow Archive', `http://nrfa.ceh.ac.uk/`. Accessed: 2017-01-22.

Neal, R. M. (2012), *Bayesian Learning for Neural Networks*, Vol. 118, Springer Science & Business Media.

Norris, R. H., Webb, J. A., Nichols, S. J., Stewardson, M. J. and Harrison, E. T. (2011), 'Analyzing cause and effect in environmental assessments: using weighted evidence from the literature', *Freshwater Science* **31**(1), 5–21.

Norton, S. B. and Suter, G. W. (2014), 'Causal assessment in environmental studies', *Wiley StatsRef: Statistics Reference Online* .

O'Donnell, D., Rushworth, A., Bowman, A. W., Scott, E. M. and Hallard, M. (2014), 'Flexible regression models over river networks', *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **63**(1), 47–63.

O'Hare, G., Sweeney, J. and Wilby, R. (2014), *Weather, Climate and Climate Change: Human Perspectives*, Routledge.

O'Hagan, A. (2004), 'Bayesian statistics: principles and benefits', *Frontis* pp. 31–45.

Pachauri, R. K., Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., Church, J. A., Clarke, L., Dahe, Q., Dasgupta, P. et al. (2014), *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, IPCC.

Pearce, N. and Lawlor, D. A. (2016), 'Causal inference—so much more than statistics', *International Journal of Epidemiology* **45**(6), 1895–1903.
URL: *http://dx.doi.org/10.1093/ije/dyw328*

Pearl, J. (2001), Bayesianism and causality, or, why I am only a half-Bayesian, *in* 'Foundations of Bayesianism', Springer, pp. 19–36.

Pearl, J. (2009), *Causality*, Cambridge University Press.

Pedersen, E. J., Miller, D. L., Simpson, G. L. and Ross, N. (2019), 'Hierarchical generalized additive models: an introduction with mgcv', *PeerJ* **7**, e6876.

Pinheiro, J. C. and Bates, D. M. (2000), *Mixed-Effects Models in S and S-plus*, Statistics and Computing, Springer-Verlag New York.

Pirani, M., Gulliver, J., Fuller, G. W. and Blangiardo, M. (2014), 'Bayesian spatiotemporal modelling for the assessment of short-term exposure to particle pollution in urban areas', *Journal of Exposure Science and Environmental Epidemiology* **24**(3), 319.

Prosdocimi, I., Kjeldsen, T. R. and Svensson, C. (2014), 'Non-stationarity in annual and seasonal series of peak flow and precipitation in the UK', *Natural Hazards and Earth System Sciences* **14**, 1125–1144.

Reich, B. J., Eidsvik, J., Guindani, M., Nail, A. J. and Schmidt, A. M. (2011), 'A class of covariate-dependent spatiotemporal covariance functions', *The Annals of Applied Statistics* **5**(4), 2265.

Reich, B. J., Hodges, J. S. and Zadnik, V. (2006), 'Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models', *Biometrics* **62**(4), 1197–1206.

Reiersøl, O. (1945), Confluence analysis by means of instrumental sets of variables, PhD thesis, Almqvist & Wiksell.

Renard, B., Garreta, V. and Lang, M. (2006), 'An application of Bayesian analysis and Markov chain Monte Carlo methods to the estimation of a regional trend in annual maxima', *Water Resources Research* **42**(12).

Renard, B., Lang, M., Bois, P., Dupeyrat, A., Mestre, O., Niel, H., Sauquet, E., Prudhomme, C., Parey, S., Paquet, E. et al. (2008), 'Regional methods for trend detection: Assessing field significance and regional consistency', *Water Resources Research* **44**(8).

Riebler, A., Sørbye, S. H., Simpson, D. and Rue, H. (2016), 'An intuitive Bayesian spatial model for disease mapping that accounts for scaling', *Statistical Methods in Medical Research* **25**(4), 1145–1165.

Ripley, B. D. (2009), *Stochastic Simulation*, Vol. 316, John Wiley & Sons.

Riseng, C. M., Wiley, M. J., Black, R. W. and Munn, M. D. (2011), 'Impacts of agricultural land use on biological integrity: A causal analysis', *Ecological Applications* **21**(8), 3128–3146.

Robins, J. (1986), 'A new approach to causal inference in mortality studies with a sustained exposure period–application to control of the healthy worker survivor effect', *Mathematical Modelling* **7**(9-12), 1393–1512.

Robins, J. M., Hernan, M. A. and Brumback, B. (2000), 'Marginal structural models and causal inference in epidemiology', *Epidemiology* **11**(5), 550–560.

Robson, A. and Reed, D. (1999), *Flood Estimation Handbook (five volumes)*, Institute of Hydrology, Wallingford.

Rodrigues, E. C. and Assunção, R. (2012), 'Bayesian spatial models with a mixture neighborhood structure', *Journal of Multivariate Analysis* **109**, 88–102.

Rue, H. and Held, L. (2005), *Gaussian Markov Random Fields: Theory and Applications*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability, CRC Press.

Rue, H., Martino, S. and Chopin, N. (2009), 'Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **71**(2), 319–392.

Rue, H., Martino, S. and Lindgren, F. (2012), 'The R-INLA Project', *R-INLA* `http://www.r-inla.org` .

Sharma, A., Wasko, C. and Lettenmaier, D. P. (2018), 'If precipitation extremes are increasing, why aren't floods?', *Water Resources Research* **54**(11), 8545–8551.

Simpson, D., Rue, H., Riebler, A., Martins, T. G., Sørbye, S. H. et al. (2017), 'Penalising model component complexity: A principled, practical approach to constructing priors', *Statistical Science* **32**(1), 1–28.

Smith, A. F. M. and Roberts, G. O. (1993), 'Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* pp. 3–23.

Sørbye, S. H. and Rue, H. (2014), 'Scaling intrinsic Gaussian Markov random field priors in spatial modelling', *Spatial Statistics* **8**, 39–51.

Stedinger, J., Vogel, R. and Foufoula-Georgiou, E. (1993), 'Frequency analysis of extreme events', *Handbook of Hydrology* .

Sun, Y. and Genton, M. G. (2011), 'Functional boxplots', *Journal of Computational and Graphical Statistics* **20**(2), 316–334.

Suter, G., Cormier, S. and Barron, M. (2017), 'A weight of evidence framework for environmental assessments: Inferring qualities', *Integrated Environmental Assessment and Management* **13**(6), 1038–1044.

Taylor, A., de Bruin, W. B. and Dessai, S. (2014), 'Climate change beliefs and perceptions of weather-related changes in the United Kingdom', *Risk Analysis* **34**(11), 1995–2004.

Thaden, H. and Kneib, T. (2018), 'Structural equation models for dealing with spatial confounding', *The American Statistician* **72**(3), 239–252.

Thygesen, L. C., Andersen, G. S. and Andersen, H. (2005), 'A philosophical analysis of the Hill criteria', *Journal of Epidemiology & Community Health* **59**(6), 512–516. **URL:** *http://jech.bmj.com/content/59/6/512*

Tootle, G. A., Piechota, T. C. and Singh, A. (2005), 'Coupled oceanic-atmospheric variability and US streamflow', *Water Resources Research* **41**(12).

Trenberth, K. E. (1999), 'Conceptual framework for changes of extremes of the hydrological cycle with climate change', *Climatic Change* **42**(1), 327–339.

Van Oldenborgh, G. J., Stephenson, D. B., Sterl, A., Vautard, R., Yiou, P., Drijfhout, S. S., Von Storch, H. and Van Den Dool, H. (2015), 'Drivers of the 2013/14 winter floods in the UK', *Nature Climate Change* **5**(6), 490.

Vansteelandt, S., Joffe, M. et al. (2014), 'Structural nested models and g-estimation: The partially realized promise', *Statistical Science* **29**(4), 707–731.

Ver Hoef, J. M., Peterson, E. and Theobald, D. (2006), 'Spatial statistical models that use flow and stream distance', *Environmental and Ecological statistics* **13**(4), 449–464.

Villarini, G., Smith, J. A., Baeck, M. L. and Krajewski, W. F. (2011), 'Examining flood frequency distributions in the Midwest US', *Journal of the American Water Resources Association* **47**(3), 447–463.

Villarini, G., Smith, J. A., Serinaldi, F., Bales, J., Bates, R. D. and Krajewski, W. F. (2009), 'Flood frequency analysis for nonstationary annual peak records in an urban drainage basin', *Advances in Water Resources* **32**(8), 1255–1266.

Vogel, R. M., Yaindl, C. and Walter, M. (2011), 'Nonstationarity: Flood magnification and recurrence reduction factors in the United States', *Journal of the American Water Resources Association* **47**(3), 464–474.

Watts, G., Battarbee, R. W., Bloomfield, J. P., Crossman, J., Daccache, A., Durance, I., Elliott, J. A., Garner, G., Hannaford, J., Hannah, D. M. et al. (2015), 'Climate change and water in the UK–past changes and future prospects', *Progress in Physical Geography* **39**(1), 6–28.

Williamson, T. and Ravani, P. (2017), 'Marginal structural models in clinical research: when and how to use them?', *Nephrology Dialysis Transplantation* **32**(suppl_2), ii84–ii90.

Wood, S. N. (2006), 'On confidence intervals for generalized additive models based on penalized regression splines', *Australian & New Zealand Journal of Statistics* **48**(4), 445–464.

Wood, S. N. (2012), 'On p-values for smooth components of an extended generalized additive model', *Biometrika* **100**(1), 221–228.

Wood, S. N. (2017), *Generalized Additive Models: An Introduction with R*, 2nd edn, CRC press.

Wooldridge, J. M. (2015), *Introductory Econometrics: A Modern Approach*, Nelson Education.

Xia, X. H., Wu, Q., Mou, X. L. and Lai, Y. J. (2015), 'Potential impacts of climate change on the water quality of different water bodies', *J. Environ. Inform* **25**(2), 85–98.

Appendix A

Additional exploratory analysis and model diagnostics for Chapter 2

We present some additional exploratory data analysis and model diagnostics that were not included in the published version of Chapter 2.

A.1 Exploratory data analysis

A histogram and density plot of annual maximum flow data for the UKBN2 benchmark catchments in Great Britain, as presented in Chapter 2, can be seen in Figure A-1. This is very skewed in nature, as can be seen further by examining the QQplot in Figure A-2.
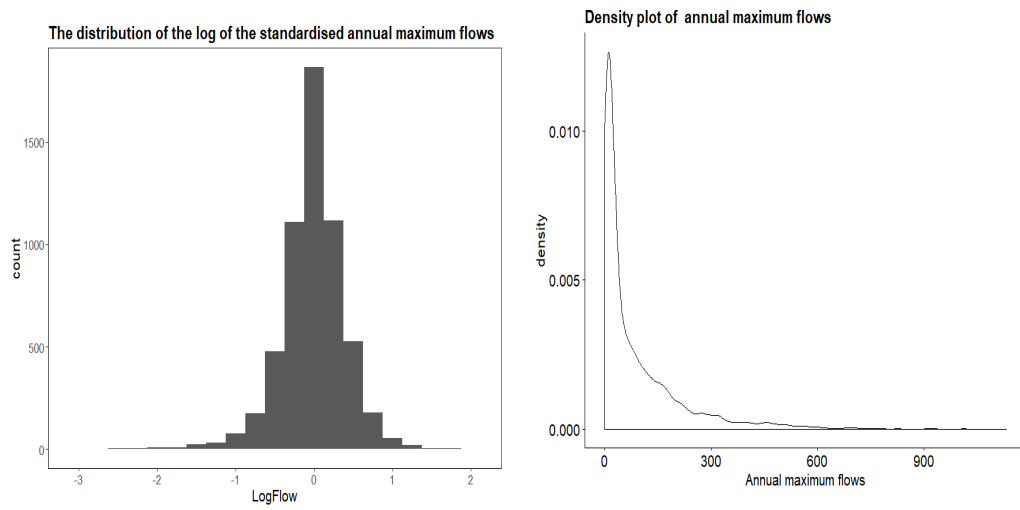
(a) A histogram of annual maximum flow data

(b) A density plot of annual maximum flow data

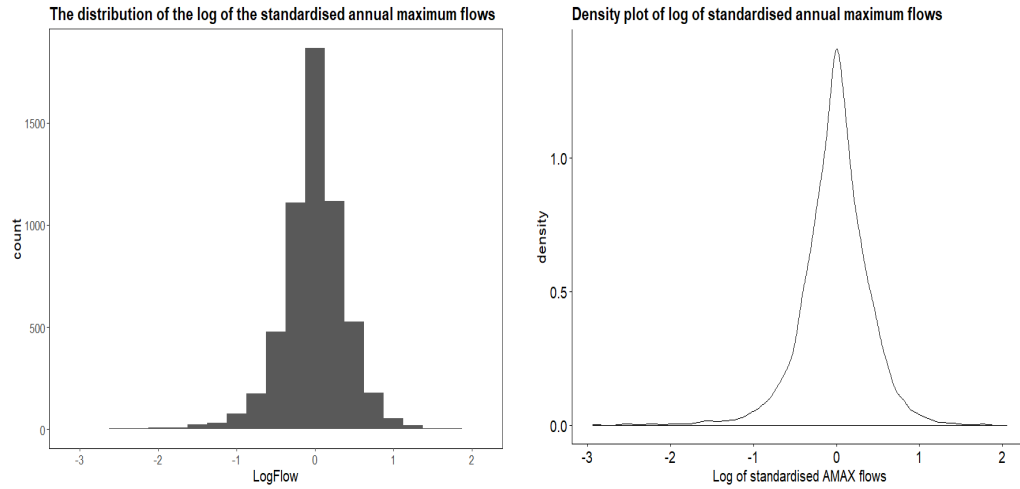Figure A-1: Distribution of annual maximum flow data for the UKBN2 benchmark catchments

Figure A-2: A qqplot for the annual maximum flow data

We try some transformations as taking the log of the annual maximum flows and assuming normality has been seen to be appropriate (Prosdocimi et al., 2014). First, the data is standardised by dividing the annual maximum flows by the median of the annual maximum series (QMED), then the log is taken. The histogram and density plots of this may be seen in Figure A-3. This is much improved. There is still some

(a) A histogram of the log of the standardised annual maximum flow data

(b) A density plot of the log of the standardised annual maximum flow data

Figure A-3: Distribution of the log of the standardised annual maximum flow data for the UKBN2 benchmark catchments

skewness in the data (see Figure A-4), however it appears somewhat more reasonable to approximate this density by a normal distribution. The skewness is due to behaviour of some individual stations (generally those with smaller flows), however it can be seen by examining some stations that the log-normal fits the data well. For example, a qqplot for the log of the standardised annual maximum flow data for station 42010 can be seen in Figure A-5. This is seen to fit the data very well.
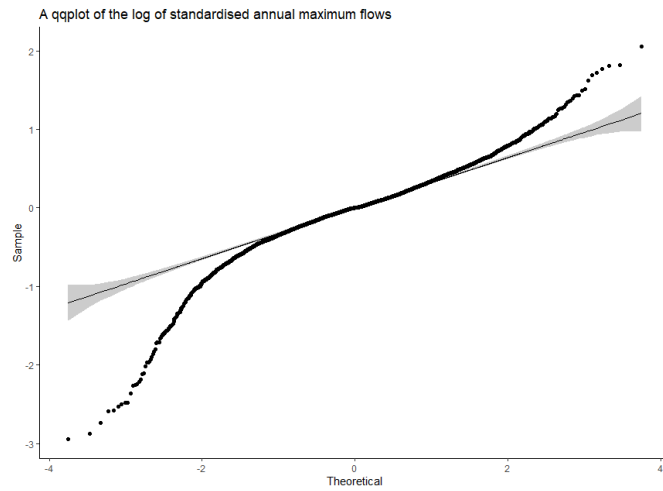
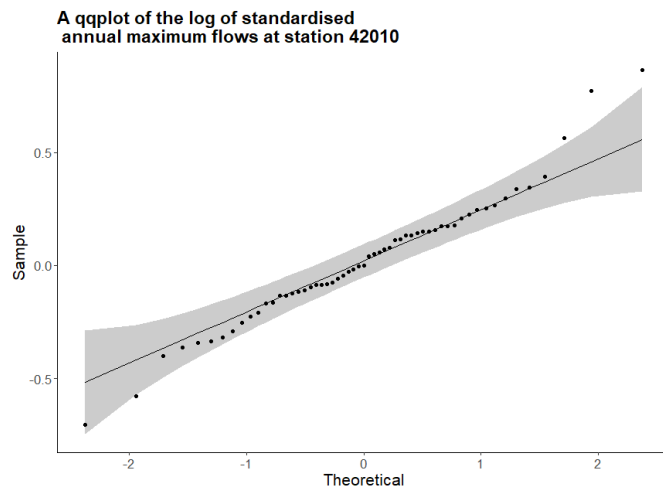Figure A-4: A qqplot for the log of the standardised annual maximum flow data

Figure A-5: A qqplot for the log of the standardised annual maximum flow data at station 42010

A semivariogram for the log of the standardised annual maximum flow data can be seen in Figure A-6, which demonstrates that the choice of an exponential correlation is reasonable. Note that there is a very steep dropoff, which suggests that it may be possible to set the correlation to zero after a chosen distance to speed up modelling, by using a more sparse correlation matrix.


Figure A-6: A variogram for the log of the standardised annual maximum flow data with distance in metres.

A.2   Model diagnostics

We can look at some diagnostics such as the traceplot for the β parameters from Model A (Figure A-7). The chains appear to be approximately stationary. We can also look at the traceplot for the random effect parameters, η^2 (Figure A-8a) and ρ (Figure A-8b). These also appear to be satisfactory.
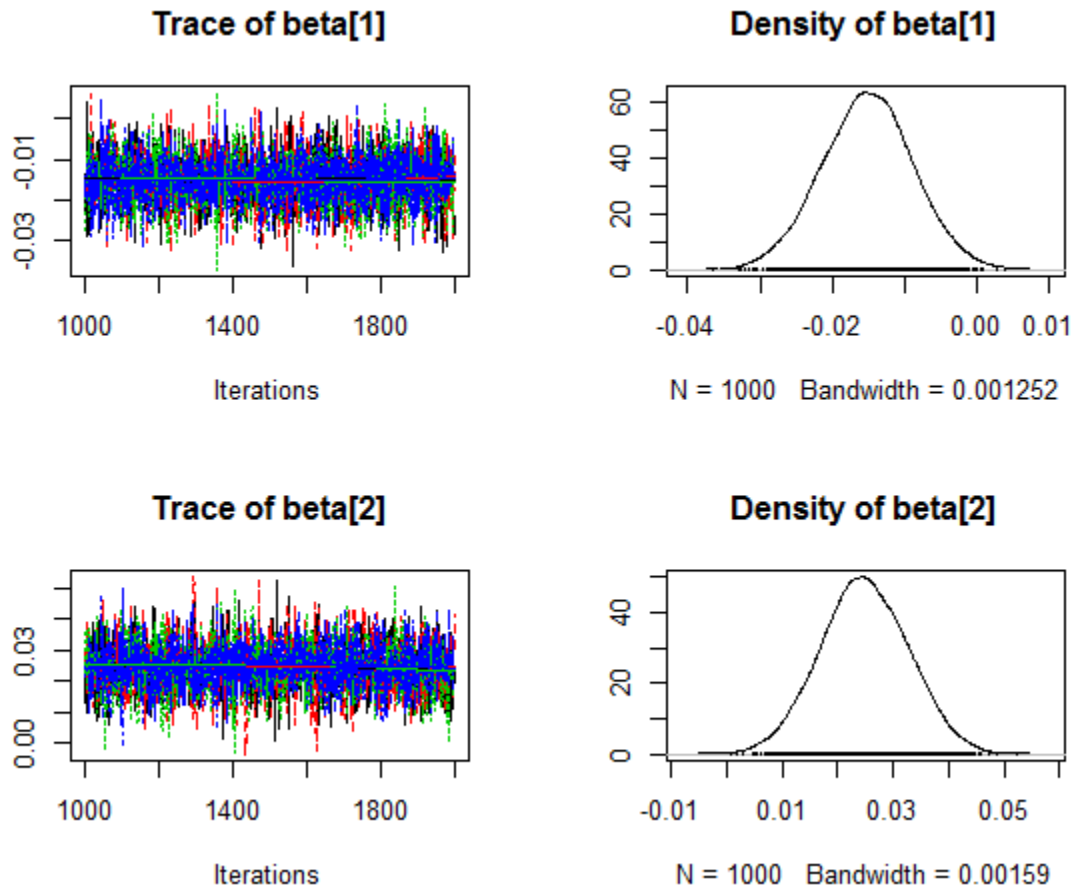
Figure A-7: A traceplot for the β_i

(a) A traceplot for parameter η^2
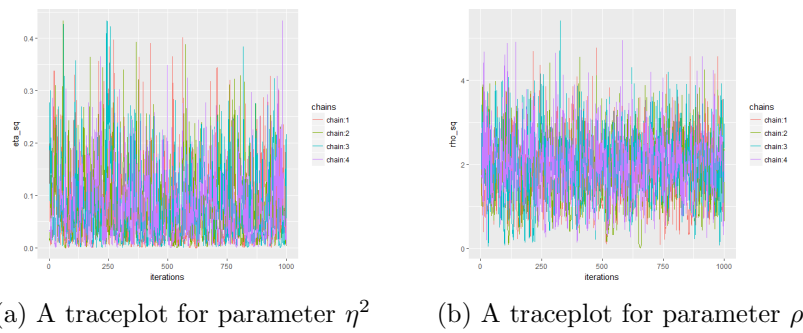
(b) A traceplot for parameter ρ

Figure A-8: Traceplots for the random effect parameters

Appendix B

Technical details for Stan and INLA

Here we present further details of the underlying methods behind Stan (used in Chapters 2 and 3) and INLA (used in Chapter 4).

B.1 Stan

Stan is a C++ based probabilistic programming language for statistical inference. It may be used to derive samples from Bayesian models to obtain posterior simulations given a user-defined model and data (Carpenter et al., 2016). It makes use of a modified Markov chain Monte Carlo (MCMC) approach, known as Hamiltonian Monte Carlo (HMC), for sampling from these models. This method uses the derivatives of the density function being sampled to generate efficient transitions spanning the posterior density (Betancourt and Girolami, 2015). The usefulness of this approach in Bayesian problems was noted by Betancourt (2017); Neal (2012). The difference between HMC and typical MCMC schemes such as Metropolis Hastings is that HMC reduces correlation between successive sampled states by using a Hamiltonian evolution between states, and additionally by targeting states with a higher acceptance criteria than the observed probability distribution. It adopts physical system dynamics by developing a Hamiltonian function $H(x, p)$, instead of a probability distribution to propose future states in the Markov chain. This allows the chain to explore the target distribution much more efficiently. Hamiltonian dynamics describe an object's motion in terms of

its location **x** and momentum **p**. For more details on Hamiltonian systems, see Landau et al. (1976).

We use the Hamiltonian dynamics as a proposal function for a Markov Chain in order to explore the target density $\boldsymbol{p(x)}$ more efficiently than using a proposal probability distribution. Starting at some initial state $(\boldsymbol{x_0}, \boldsymbol{p_0})$, we can simulate Hamiltonian dynamics for a short time using the Leap Frog method, described by:

$$p_i(t + \delta/2) = p_i(t) - (\delta/2)\frac{\partial U}{\partial x_i(t)}$$
$$x_i(t + \delta) = x_i(t) + \delta\frac{\partial K}{\partial p_i(t + \delta/2)}$$
$$p_i(t + \delta) = p_i(t + \delta/2)(\delta/2)\frac{\partial U}{\partial x_i(t + \delta)}$$

We then use the current state of the position and momentum variables at the end of the simulation as our proposed states variables $\boldsymbol{x^*}$ and $\boldsymbol{p^*}$. The proposed state is accepted or rejected via an analogous rule to the Metropolis acceptance criterion. That is, if the probability of the proposed state after Hamiltonian dynamics is greater than that before, i.e. if:

$$p(\boldsymbol{x^*}, \boldsymbol{p^*}) \propto \exp(-U(\boldsymbol{x^*}), K(\boldsymbol{p^*}) > p(\boldsymbol{x_0}, \boldsymbol{p_0}) \propto \exp(-U(x(t-1), K(p(t-1))) \quad \text{(B.1.1)}$$

then the proposed state is accepted, otherwise, the proposed state is accepted randomly. If the state is rejected, the next state of the Markov chain is set as the state at $t - 1$. For a given set of initial conditions, Hamiltonian dynamics will follow contours of constant energy in phase space. Therefore we must randomly perturb the dynamics so as to explore all of the target density. This is done by simply drawing a random momentum from the corresponding canonical distribution p(p) before running the dynamics prior to each sampling iteration. Combining these steps, we define the HMC algorithm for drawing M samples from a target distribution:

Algorithm 1: Hamiltonian Monte Carlo
---

1. Set $t = 0$.

2. Generate an initial position state $x^{(0)} \sim \pi^{(0)}$.

3. Repeat until $t = M$.

4. Set $t = t + 1$.

5. Sample a new initial momentum variable from the momentum canonical distribution $\boldsymbol{p}_0 \sim p(\boldsymbol{p})$.

6. Set $\boldsymbol{x}_0 = \boldsymbol{x^{t-1}}$.

7. Run Leap Frog algorithm starting at $(\boldsymbol{x}_0, \boldsymbol{p}_0)$ for L steps and stepsize δ to obtain proposed states $\boldsymbol{x^*}$ and $\boldsymbol{p^*}$.

8. Calculate the Metropolis acceptance probability:
   $\alpha = \min(1, \exp(U(\boldsymbol{x^*}) + U(\boldsymbol{x_0}) - K(\boldsymbol{p^*}) + K(\boldsymbol{p_0}))$.

9. Draw a random number u from $\mathrm{Unif}(0,1)$.

10. **if** $u \leq \alpha$ **then**

    accept the proposed state position $\boldsymbol{x^*}$ and set the next state in the Markov chain;

    else

    set $\boldsymbol{x}^{(t)} == \boldsymbol{x}^{(t-1)}$;

    end

B.2 INLA

The Integrated Nested Laplace Approximation (INLA) is an approximate Bayesian method performs numerical calculations of posterior densities using a series of Laplace approximations of hierarchical latent Gaussian models. We summarise the details of the INLA method following the approach of Rue et al. (2009).

Given the model

$$\log(y_i) = \alpha + \sum_{j=1}^{n_f} f^{(j)}(u_j) + \sum_{k=1}^{n_\beta} \beta_k z_{ki} + \varepsilon_i \tag{B.2.1}$$

with latent effects $x = \{\log(y_i), \alpha, \beta_k, ...\}$, where f_k may be used to model non-linear effects, spatial dependencies, etc., and x a GMRF as defined in Chapter 1. Then, the posterior distribution of model latent effects may be written as:

$$\pi(x, \theta|y) \propto \pi(\theta)\pi(x|\theta) \prod_{i \in \mathcal{I}} \pi(y_i, x_i, \theta)$$

$$\propto \pi(\theta)|Q(\theta)|^{-1/2} \exp\left(\frac{1}{2}x^T Q(\theta)x\right) + \sum_{i \in \mathcal{I}} \log(\pi(y_i|x_i, \theta))$$

The aim of the INLA approach is to find the marginal posterior densities for each of the latent variables x_i as well as for the hyperparameters θ given the observed data $\boldsymbol{y}$,

$$\pi(x_i|y) = \int \pi(x_i|\theta, y)\pi(\theta|y) \, d\theta$$

$$\pi(\theta_j|y) = \int \pi(\theta|y) \, d\theta_{-j}.$$

To estimate these marginals, we need either $\pi(\theta|y)$, or a good approximation to it:

$$\hat{\pi}(\theta|y) \propto \frac{\pi(x, \theta, y)}{\pi_G(x|\theta, y)},$$

where $\pi_G(x|\theta, y)$ is the Gaussian approximation to the full conditional of x. This was used to compute the marginal of x_i numerically (Rue et al., 2009):

$$\hat{\pi}(x_i|y) = \sum_k \hat{\pi}(x_i|\theta_k, y)\hat{\pi}(\theta_k|y)\Delta_k,$$

where the Δ_k are weights associated with the ensemble of values θ. For the first term, we need a good approximation. A Gaussian may be sufficient but this is not always

the case. Rue et al. (2009) introduced a Laplace approximation:

$$\hat{\pi_{LA}}(x_i|\theta, y) \propto \frac{\pi(x, \theta, y)}{\pi_{GG}(x_{-i}|x_i, \theta, y)}$$

where $\hat{\pi_{LA}}$ denotes the Laplace approximation, about the mode. Finally, Rue et al. (2009) developed a simplified version of this approximation to improve this, using a series approximation of the Laplace approximation about x_i.

This approach to Bayesian inference is a much more efficient computational technique than Markov Chain Monte Carlo (MCMC) methods and hence is ideal for use on large scale problems. If we can formulate a model as a latent Gaussian model (as introduced in Chapter 1), we can implement it using INLA.

We can implement such models in `R` via R-INLA (Rue et al., 2012). R-INLA provides a user interface similar to that of (generalised) linear models in `R`, and is available from `http://www.r-inla.org`. The interface can handle a multitude of different models including ones that have fixed effects, non-linear terms or random effects. Thsi flexible interface allows for the specification of different priors and model fitting options using a `formula()` argument. The model is fitted with a call to function `inla()`, which will return results from the fitted model, which include the marginal distributions of the latent process and model parameters, as well as summary statistics. R-INLA can produce other quantities up on specification to the call to INLA (Blangiardo and Cameletti, 2015). For more details, see Blangiardo and Cameletti (2015); Rue et al. (2012).