Department for Environment Food and Rural Affairs

Environment Agency

delivering benefits through evidence

The risk of widespread flooding – Capturing spatial patterns in flood risk from rivers and coasts

SC060088/R1 Spatial Coherence of Flood Risk - Technical Methodology Report

Flood and Coastal Erosion Risk Management Research and Development Programme

The Environment Agency is the leading public body protecting and improving the environment in England and Wales.

It's our job to make sure that air, land and water are looked after by everyone in today's society, so that tomorrow's generations inherit a cleaner, healthier world.

Our work includes tackling flooding and pollution incidents, reducing industry's impacts on the environment, cleaning up rivers, coastal waters and contaminated land, and improving wildlife habitats.

This report is the result of research commissioned by the Environment Agency's Evidence Directorate and funded by the joint Environment Agency/Defra Flood and Coastal Erosion Risk Management Research and Development Programme.

Published by:
Environment Agency,Horizon House, Deanery Road, Bristol, BS1 5AH
www.environment-agency.gov.uk

ISBN: 978-1-84911-247-5

Further copies of this report are available from:
The Environment Agency's National Customer Contact Centre by emailing:
enquiries@environment-agency.gov.uk
or by telephoning 08708 506506.

Author(s):
Caroline Keef, Rob Lamb, Jonathan Tawn, Paul Dunning, Crispian Batstone, Mark Lawless

Dissemination Status:
Publicly available

Keywords:
Spatial dependence, extremes, flood risk

Research Contractor:
Dr Rob Lamb
JBA Consulting
South Barn
SKIPTON
North Yorkshire
BD23 3AE
01756 799 919

Environment Agency's Project Manager:
Stefan Laeger, Exeter

Collaborator(s):
Professor Jonathan Tawn

Project Number:
SC060088

Product Code:
SCHO1011BUBR-E-E

Evidence at the Environment Agency

Evidence underpins the work of the Environment Agency. It provides an up-to-date understanding of the world about us, helps us to develop tools and techniques to monitor and manage our environment as efficiently and effectively as possible. It also helps us to understand how the environment is changing and to identify what the future pressures may be.

The work of the Environment Agency's Evidence Directorate is a key ingredient in the partnership between research, guidance and operations that enables the Environment Agency to protect and restore our environment.

This report was produced by the Research, Monitoring and Innovation team within Evidence. The team focuses on four main areas of activity:

- **Setting the agenda**, by providing the evidence for decisions;
- **Maintaining scientific credibility**, by ensuring that our programmes and projects are fit for purpose and executed according to international standards;
- **Carrying out research**, either by contracting it out to research organisations and consultancies or by doing it ourselves;
- **Delivering information, advice, tools and techniques**, by making appropriate products available.

Miranda Kavanagh

Director of Evidence

Executive summary

Overview

This is the technical methodology report for the Environment Agency R&D project SC060088 'Spatial coherence of flood risk', which is a scoping study to identify, develop and trial methods for determining the likelihood of spatially extensive floods from single or multiple sources.

The project is about how we can estimate the risk (likelihood, severity and consequences) of flooding in more than one location. We need to be able to understand this spatial aspect of risk in order to properly estimate the probabilities of catastrophic emergencies and economic losses at a regional and national scale.

The overall objective of this project is to develop and test methodologies to assess the risk of widespread flooding, incorporating both the analysis of sources and consequences of flooding at different spatial scales, up to regional or national level.

The first objective is to review and develop methods for analysing and modelling dependence between multiple variables that affect flood risk, in particular the spatial dependence in extremes of river flow or level and sea levels.

The second objective is to show how the spatial and between-variable dependence can be extended to include risk pathways (such as defence failure) and receptors of risk to add a spatial dimension to probabilistic flood risk assessment methods that are being used or developed elsewhere.

Project reports

This technical methodology report contains a detailed description of the methods investigated in the project.

This report should be read in conjunction with the final proof of concept report which contains results from three demonstration case studies and a discussion of how the findings of this project can be used to realise benefits for flood risk management.

Technical methodology

This report reviews the statistical concepts and theory that are needed to understand spatial aspects of flood risk. It also reviews how spatial dependence can be important in flood risk management and how it has previously been studied. Following the review sections, the report sets out a methodology for modelling the risk of flooding that includes the risk of widespread floods by modelling spatial dependence in river flows and storm surge. We show how it can be integrated with current procedures used in the National Flood Risk Assessment to obtain estimates of aggregated economic damage over a whole country. Techniques for quantifying uncertainty within modelling approaches are discussed.

The final section of the report sets out a programme of work for implementation of the technical methods to deliver benefits through three business applications.

Acknowledgements

We would like to thank Matt Crossman (Defra), Edward Evans (MAR theme adviser) and Katharine Evans (Environment Agency) for taking time to discuss the aims of the research project.

We are grateful for helpful comments received during the course of this research from Adam Baylis, Bill Donovan, Stefan Laeger, Ian Meadowcroft, Suresh Surendran (Environment Agency), Edward Evans and Jon Wicks (Halcrow). We are particularly grateful for the detailed review comments from Jim Hall (Newcastle University).

We are grateful to Gwyn Rees of the National River Flow Archive (CEH Wallingford) for assistance with river flow and sea level data previously supplied to JBA Consulting.

We are grateful to the Environment Agency, Ben Gouldby (HR Wallingford), Matt Horrit and Jon Wicks (Halcrow) for supplying documentation and contributions to telephone consultations regarding RASP and NaFRA.

Contents

Tables

Figures

1 Introduction

In recent years Defra (Department for Environment, Food and Rural Affairs) and the Environment Agency have led the way in promoting risk-based policies and practices for flood management through Making Space for Water, the Flood Foresight project and changes to policy and process. Many of these changes have been supported by projects in the R&D programme.

Historically, the methods used for flood risk management had limited capability to deal with the spatial dependence structure in the sources of risk since they focussed on single points or local systems rather than on a wider area. In other words, it has not been possible to answer questions such as ‘what is the chance that many different locations will be affected by severe flooding?’ or ‘what is the chance that widespread river flooding will coincide with high tides and storm surge?’

This limits our ability to manage flood risk. We have not previously had tools to identify and assess the range of possible flood events that we must manage, and we do not have reliable tools to assess the national risk ‘profile’ to describe the relative likelihoods of different scales of flood event.

1.1 Overall objectives

This project is a scoping study to identify, develop and trial a method for assessing flood risk when aggregated over large spatial scales.

The overall objective of this project is therefore to develop and test methodologies to assess the risk of widespread flooding, incorporating both the analysis of sources of flooding and the consequences at different spatial scales, up to regional or national level.

The first objective is to review and develop methods for analysing and modelling dependence between multiple variables that affect flood risk, in particular the spatial dependence in extremes of river flow or level and sea levels, so as to assess the likelihood of spatially extensive floods from single or multiple sources.

The second objective is to show how the spatial and between-variable dependence can be extended to include risk pathways (such as defence failure) and receptors of risk to add a spatial dimension to probabilistic flood risk assessment methods such as the Risk Assessment for System Planning (RASP) approach that underlies the Environment Agency’s National Flood Risk Assessment (NaFRA).

1.2 Approach

Meeting the objectives of this project requires understanding and quantification of the probability and consequences of flooding at different spatial scales. A suitable statistical approach is needed to meet these aims.

After a review of requirements and possible approaches, the project has built upon a statistical method developed by Heffernan and Tawn (2004) and first applied to river flow and rainfall data by Keef (2007). The method provides a very flexible model for the joint probability of large sets of inter-related variables, including the probability of extreme values being experienced in more than one of those variables. This makes it well-suited to model the physical sources of flooding, such as high river flows and sea levels.

The work in this project has shown how the existing statistical model can be linked conceptually into an integrated flood risk model, based on the source – pathway – receptor approach adopted by the Environment Agency. This provides a conceptual framework for adding a spatial dimension to assessments of flood risk that consider defence system performance (‘pathways’) and economic or other consequences of flooding (‘receptors’), which has the potential to enhance our assessment of flood risk regionally and nationally. The scope for realising these benefits has been demonstrated in three case studies, which are presented in the accompanying proof of concept report.

1.3 Scope

The scope of the project is to examine fluvial and coastal (tide and surge) flooding. Urban ‘surface water’ flooding and groundwater flooding are excluded.

Rainfall is not included in this study because the primary requirement for river flooding is to model the statistics of high river flows and observations from river gauges are therefore a much more direct measure. It is important to note that the modelling approach is not a time series model for river flows but instead captures the statistics of observed high flow events. Hence it has not been necessary to model antecedent conditions because river flow data in effect gather together temporal and spatial variations in rainfall.

However, the statistical methodology that has been used is based on a general approach to representing the joint probability of multiple variables. It should therefore be suitable, in principle, for extension to include rainfall information or other sources of flooding.

1.4 Contents of the technical methodology report

This report outlines a generic conceptual model for flood risk analysis and methods that can address the above objectives and technical requirements. The contents of the methodology report are as follows:

Section 1	Background	
Section 2	Conceptual model	Overview of the generic conceptual model adopted for the project.
Section 3	Flood risk methods review	Review of existing methods and tools for flood risk analysis.
Section 4	Statistical methods review	Review of the statistical features required for a spatial flood risk model and relevant statistical methods.
Section 5	Statistical model for observed locations	Outlines the statistical model adopted for this study and shows how it can be used to model the spatial distribution of flood risk at the source level.
Section 6	Implementation of the statistical model	Describes the steps needed to implement the statistical model so that it can be linked with pathway and receptor models.

Section 7	Integration of pathway and receptor models	The proposed method for linking the spatial model to existing tools to represent flood defence systems and consequences of flooding.
Section 8	Uncertainty and confidence	Commentary on sources of uncertainty and quantification of confidence in the results.
Section 9	Methodology conclusions	

1.5 Contents of the proof of concept report (accompanying report)

An accompanying report summarises the proof of concept results that were presented at a stakeholder workshop on 18 March 2009. It also discusses the aims of the project further and how it supports the Environment Agency's flood risk management business. The report is structured as follows:

Section 1	Background	
Section 2	Proof of concept approach	Summary of how proof of concept has been established.
Section 3	Fluvial flood risk	Proof of concept results for river flooding in North East Region of the Environment Agency.
Section 4	Coastal flood risk	Proof of concept results for five coastal locations on the north east coast.
Section 5	Joint inland and coastal flood risk	Initial findings from joint analysis of regional river flood risk and coastal flooding.
Section 6	Integration with NaFRA	Commentary on proposed approach to integrate spatial analysis into NaFRA.
Section 7	Stakeholder workshop	Report and discussion of a workshop help to seek feedback from interested stakeholders
Section 8	Proof of concept conclusions	

1.6 Project management

1.6.1 R&D programme

The project was managed through the 'Modelling and Risk' (MAR) theme of the Defra/Environment Agency joint Flood and Coastal Erosion Risk Management (FCERM) research programme. The MAR theme leader is Suresh Surendan.

The Environment Agency's project manager for this contract was Stefan Laeger. Ian Meadowcroft was the Environment Agency's project executive.

1.6.2 Research contractors

The research was carried out by a consortium of JBA Consulting (Dr Rob Lamb, Dr Caroline Keef, Paul Dunning, Dr Crispian Batstone) and Professor Jonathan Tawn. The contractor's project manager was Rob Lamb of JBA Consulting.

1.7 Scoping study programme and outputs

The project started in January 2008 and completed in Autumn 2009. The main outputs of the project are listed below.

- Environment Agency R&D summary.
- Final summary and proof of concept report.
- Final methodology report (this document).
- Paper at FCM>08 conference, Manchester (July 2008).
- Paper at 'UK Extremes' statistics conference, Lancaster (September 2008).
- Paper at FloodRisk2008 conference, Oxford (October 2008).

1.8 Strategic programme

This scoping study is phase one of a potential longer term initiative to support a more integrated approach to flood risk management and will be particularly relevant to national policy and investment, catchment, coastal and estuary strategies, asset management, flood incident management, development control and mapping process.

Depending on the findings of this study and the prevailing development priorities, a second phase would see the proposed techniques integrated with Flood Risk Management tools such as the National Flood Risk Assessment (NaFRA), and the Modelling and Decision Support Framework (MDSF2).

2 Conceptual model overview

In this section of the report we introduce a generic conceptual model for spatially aggregated flood risk. Starting with the simplest outline structure, we describe an approach using the source – pathway – receptor (S-P-R) concepts for flood risk and set the scene for linking with the RASP methods for flood risk modelling, and therefore with tools that implement RASP.

2.1 The source – pathway – receptor framework

This source – pathway – receptor concept allows the analysis of flood risk to be approached by characterising three components separately:

Sources: The conditions that create a load on flood defence systems.

Pathways: The physical mechanisms by which flood waters and hence risk are propagated from a source to a receptor.

Receptors: The consequent impacts of the hazard.

In source – pathway – receptor terms, the loading on the flood defence system is the ‘source’, the response of the system is part of the ‘pathway’, and the impacts are felt by the ‘receptors’, which may include population, property, business, infrastructure, insurance and others.

2.2 Generic conceptual model

2.2.1 Single (localised) source of flooding

A good starting point to describe the proposed generic approach is a statistical model for the ‘source’ component of the S-P-R flood risk model. In the simplest case, where we consider only one source of flooding, the model must describe the probability distribution of the physical variable that represents the source term. Two obvious examples are the probability distributions of water levels at a given location in a river or at the coast. For flood risk analysis, we are particularly interested in finding a good model for the large extremes of the source variable.

The overall aim of the risk model is to estimate the consequences of flooding, that is, the ‘receptor’ term in the S-P-R concept. Consequences are represented by one or more variables that we will refer to generically as a ‘cost function’. Examples of the cost function could be financial loss from property damage, economic loss, an index of social impacts or simply maximum flood water depth. The desired outcome here is the probability distribution of impacts on the receptor, especially at the extremes.

If the statistical distribution representing the source variable is simple and the cost function is also simple, then it might be possible to derive the distribution of the cost function (that is, whatever we chose to represent receptor impacts) directly by algebraic manipulation. This might be possible, for example, if the source component was represented by a simple univariate distribution such as the Gumbel distribution and the cost function was a direct and simple function of the source variable such as a fixed proportion. An integration of this type is the simplest and most generic expression of our conceptual model, illustrated in Figure 2-1.

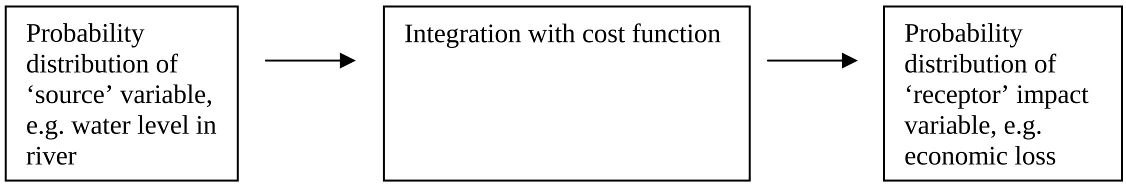

Figure 2-1: Simple generic risk model structure.

In general, the source distribution and cost function are much too complicated for the integration to be done analytically. In particular, to capture the joint occurrence of flooding over large areas (or to capture the joint occurrence of multiple sources) the source distribution has to be a multivariate distribution that can capture the dependence between the multiple locations or variables (the consequences of not doing so correctly can be serious bias in risk estimates, as shown in Section 4.5). There are few multivariate distributions that could be integrated analytically and so it is necessary to use some form of numerical integration to calculate the distribution of the cost function.

The simplest approach is a 'brute force' Monte Carlo simulation. The basic generic conceptual model can then be represented as in Figure 2-2.

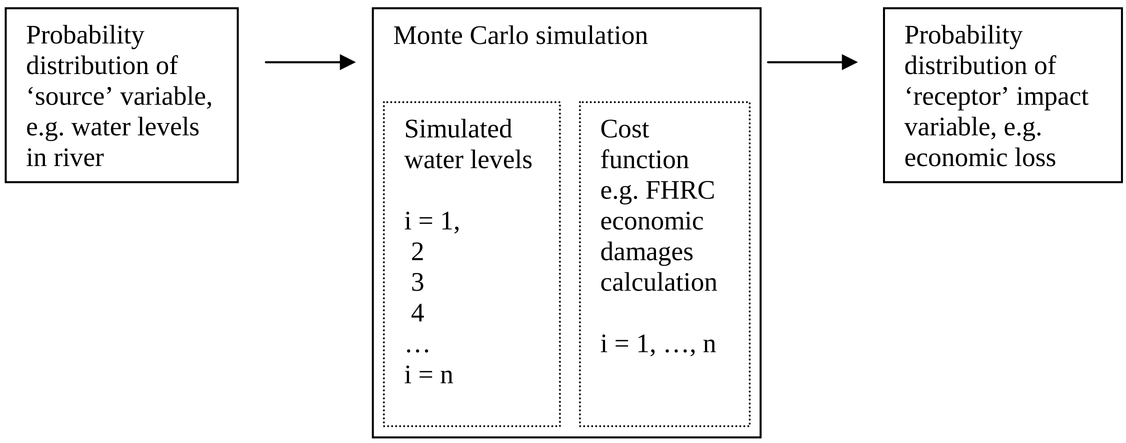

Figure 2-2: Simple generic risk model structure with a numerical approximation of the integration.

In the generic structure given above we have continued to make a very simple assumption, for the sake of argument, that the cost function can be calculated directly for any given value of the physical source variable. However, this is unlikely to be valid in most flood risk analysis, where the consequences of the flooding depend on the physical source variable via complex processes of flow routing and flood defences. Together, these modifying processes are the 'pathway' for the flood risk, and their complexity is a further reason why some type of numerical integration is needed.

To incorporate the risk pathway, we must add components to the generic model to represent the modification of the physical source variable on the cost function. For flow routing, this may be done by using a hydraulic model (or several linked models) to 'translate' the source water level values into flood depths that can be used in the cost function. Now, the generic risk model may look like Figure 2-3.

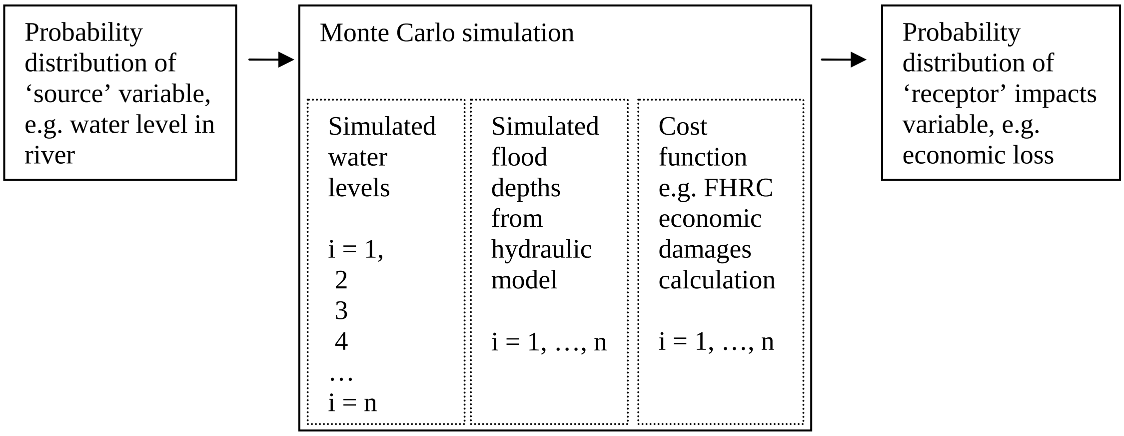

Figure 2-3: Generic risk model structure incorporating flow routing within the risk 'pathway'.

Here, hydraulic modelling is being used to represent a deterministic flood risk pathway, where a given water level loading from river or sea would always produce the same flood depths and hence cost function. However, not all aspects of the flood risk pathway can be considered in this way. In particular, there is generally insufficient local information (or physical understanding) to represent the performance of flood defences in a purely deterministic manner.

Instead, a stochastic approach is commonly taken where the defence performance is represented by a probability distribution. In the RASP methodology, the defence performance is represented by a fragility curve, which is a distribution of failure probability conditional on the water level loading and asset condition. We can generalise this to express the probability of the defence being in a certain state (failed or not failed, but potentially others as well) conditional on the 'intensity' of the loading (which may be measured physically by the water level but could potentially also include other source variables such as velocity and wave direction).

Because the modifying effect of flood defences on flow routing and depths is also complex, the integration of the conditional defence performance distribution into the basic scheme involves another Monte Carlo simulation step, where the defence state for each Monte Carlo sample is generated such that the simulated states correspond to the conditional failure distribution given by the fragility curve. The basic S-P-R conceptual model can then be represented as in Figure 2-4.

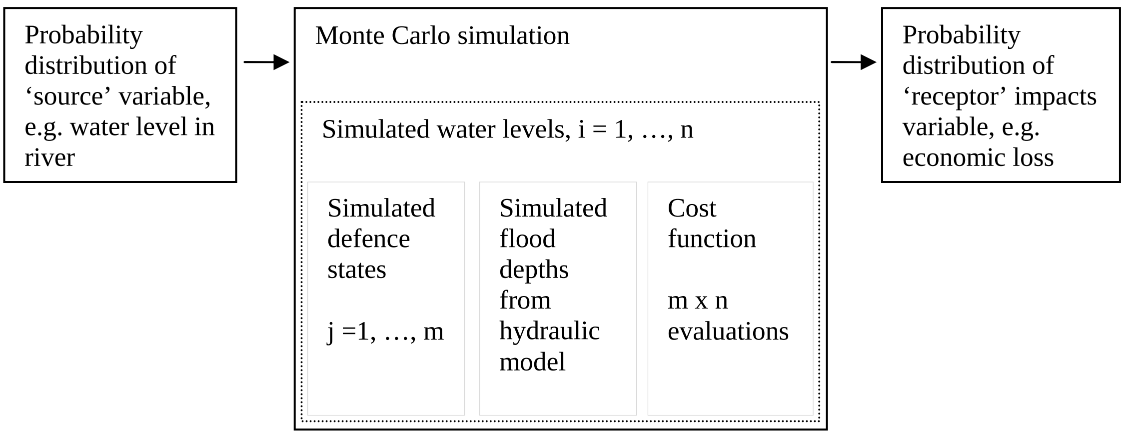

Figure 2-4: Generic risk model structure incorporating flood defence performance and routing in the pathway component.

The Monte Carlo simulation outlined in Figure 2-4 is a numerical approximation of the integration of the probability distributions of the source variable and defence failure state. The approximation introduces the potential for errors in the integration. The simulation can be carried out until some convergence has been achieved (see, for example, Dawson and Hall, 2006; Gouldby *et al.*, 2008).

It may be necessary to strike a balance between the density of the Monte Carlo simulation (the number of simulations) and the level of physical detail in the engineering models of the risk pathway. In our conceptual model for spatially aggregated risk we propose an approximation that allows for high spatial resolution (physical detail) at points within the Monte Carlo simulation space but relies on interpolation between those samples. It is important to bear in mind that this is not a fundamental decision about the conceptual basis of the model, but rather a practical choice based on available computing power and data. Different choices are possible within the same conceptual framework, for example use of a cruder but fast rapid flood spreading model (as proposed in MDSF2) or even greater use of more detailed, but more costly 2D modelling.

2.2.2 Multiple (distributed) sources of flooding

In Figure 2-4 above the source of flooding is considered as a univariate distribution corresponding to one physical variable. This may be defined at one location or at a set of locations that are deterministically dependent on each other (which effectively means treating those locations as if they were one point).

Conceptually it is straightforward to extend the generic model above to include a source variable that is spatially distributed. This extension is necessary if we are to calculate risk over arbitrarily large spatial scales of aggregation. For a spatial source component, we need to extend the basic model so that the integration takes place over a vector-valued source variable; that is, a source of flooding that can exist at multiple locations in any one event. In the Monte Carlo approximation, it will then be necessary to simulate a vector representing the physical source variable for each Monte Carlo sample. The structure is illustrated in Figure 2-5.

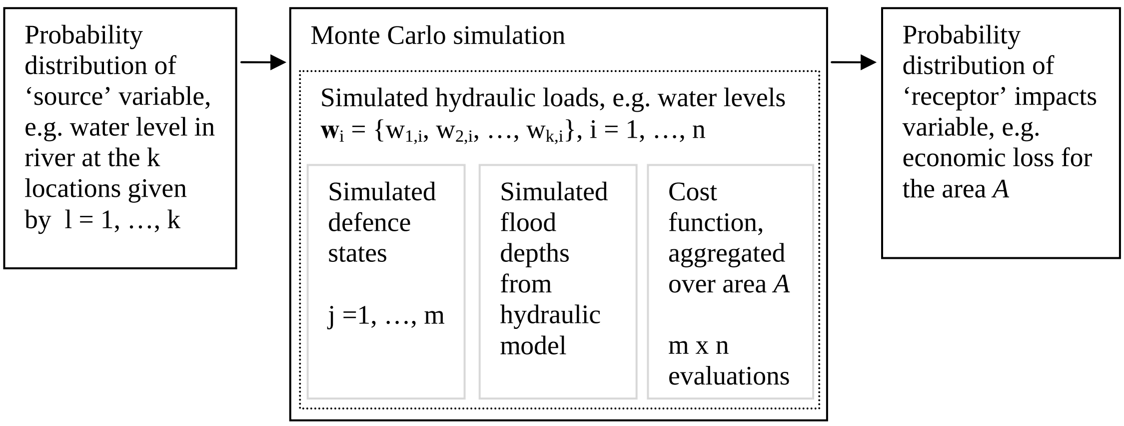

Figure 2-5: Spatially aggregated generic risk model structure with pathway and multiple (spatial) source variables. $\mathbf{w}_i = \{w_{1,I}, w_{2,I}, \ldots, w_{k,i}\}$ **is a spatial vector of hydraulic loads (such as water levels) at each of k locations (such as river network nodes) for the** i^{th} **of n Monte Carlo samples.**

The description of the joint probability distribution of the flood source variable is critical in this calculation. There are three important cases that need to be considered in view of their impact on both the realism of the analysis and the simulation approach. These three cases describe the degree of dependence between the physical source variable at different locations. The possibilities are:

- Complete independence, in which case values of the source variable can be simulated independently from the marginal distribution at each of the k locations.
- Complete dependence, in which case the value at the first location is simulated from its marginal distribution (giving the q-th quantile) and then the value at all other locations are set equal to their respective q-marginal quantile. If all locations have the same marginal distribution this simulation gives equal values at all locations.
- Partial dependence, in which case the water level vectors must be simulated from the joint distribution in a way that preserves the dependence between locations.

An assumption of independence is straightforward but unrealistic for flood risk. An assumption of complete dependence may be valid for localised analysis. For regional or national scale analysis, the joint distributions of important source variables such as river flows, sea levels or rainfall exhibit dependence. The focus of this project is on effective and efficient methods to represent the joint distribution function for rivers and sea levels, and generating data from these distributions that can be used within a simulation framework as outlined above.

In this scheme, the statistical model is effectively 'interfaced' with deterministic hydraulic models for the pathway and receptor. The statistical model can be used to generate sample vectors from the joint distribution of boundary conditions for a broad scale hydraulic model (for example at gauged tributaries to a major river or the tidal boundary plus river(s) inflowing to an estuary). These can then be used to drive a corresponding number of hydraulic model runs to obtain water levels throughout the system, with consequential damage and hence risk. The modelling for TE2100 was a bivariate version of this. A direct link between the statistical model and hydraulic models may be conceptually straightforward but not necessarily efficient or flexible. We discuss in Section 2.3.2 and Section 7 how the integration can be done in a modular, efficient and flexible way.

It is likely that if the load on the flood management system is defined over a number of different locations then the defence system may also be split into sections. In this case, there may be a very large number of combinations of defence states to consider and in the simple generic model this would have to be reflected in the size of the Monte Carlo sample (the value chosen for m in Figure 2-5). It may be a reasonable approximation to consider only those scenarios up to a specified number of defence failures. It is also possible to calculate probabilities of failure scenarios directly using a reliability analysis approach. Both points are discussed by Hall *et al*. (2003).

For the calculation of spatially aggregated risk, there is also an additional step required within the Monte Carlo procedure where the cost function is aggregated over the area of interest. This may add to the computational load required for modelling the risk pathway, for example if economic damage calculations are needed for every property in a national database. Again, there may be scope for making approximations to the aggregated cost function for large spatial scales.

2.3 Conceptual model development

2.3.1 Joint distribution of the source variables

The scope of this project is to concentrate on flooding from rivers and the sea (including storm surge but not wave action). The data and analysis methods are discussed in detail in later sections of this report. A brief overview is given here.

For river flows we will work with daily mean flow data. This is because quality checked daily mean flows are readily available via the National Water Archive maintained by the Centre for Ecology and Hydrology and daily flows are suitable for characterising the spatial and temporal dependence in most catchments. (Note that the statistical methods used here are not restricted to daily flows and is possible to incorporate sub daily data for the dependence analysis or instantaneous peak flows for marginal distributions).

For sea levels the variable of most interest is the skew surge, defined on a daily maximum scale from hourly sea level data. Using skew surge allows for statistical separation of the surge and astronomical tide components of the sea level.

The statistical method to be used to model the joint (spatial) distribution of the variables is the conditional exceedance model of Heffernan and Tawn (2004), which has been chosen for its flexibility in representing a range of features observed in the data.

2.3.2 Integration approach

The simplest implementation of the conceptual model would be a brute force Monte Carlo simulation evaluating randomly generated combinations of source variables (spatially distributed water levels, or related variables) and defence states with a suitable ‘hazard’ model, such as a gridded 2D hydraulic model and cost function. However, this would require an enormous computing effort to work for large scales and is therefore considered impractical.

It is not desirable to build a ‘black-box’ system where the entire simulation engine is closed because it is then difficult to take a modular approach, where choices can be made about how to implement different components of the model depending on priorities and available resources.

We therefore suggest an approach that breaks the generic model down into separate parts and introduces useful approximations wherever possible. This is already a feature of the statistical model for the joint distribution of river or sea ‘loading’ in that the marginal distributions of the physical variables are separated from the dependence structure. It is also an approach that allows for future enhancements of the methods.

Source, receptor and simple hydraulic pathway

To begin with, assume for simplicity that the flood defence ‘pathway’ can be ignored. The most time-consuming part of the simulation process is then likely to be the hydraulic modelling of floodplain depths (assuming that a relatively high resolution, dynamic model is used). In practice, small changes in the boundary conditions for the hydraulic model may often result in small changes in the depth and extent. In this case, it is reasonable to evaluate the hydraulic model for set points within the Monte Carlo simulation space and to approximate its behaviour for points in between. A quick and simple approximation is to interpolate between the depth data produced by detailed evaluations of the hydraulic model. This is, in effect, a simple form of statistical emulator for the full hydraulic model.

A straightforward approach is to run the hydraulic model ‘off line’ for a number of prescribed load conditions corresponding to set probabilities of the load, and then to use the resulting depth grids as a lookup table for events simulated in the Monte Carlo procedure. There could be thresholds where the hydraulic model response does not vary smoothly with the boundary conditions and the Monte Carlo approximation could be improved by carrying out further evaluations between points where there is a high rate of change in its response.

The approach is illustrated in Figure 2-6, where the shaded boxes replace the original direct simulation using the hydraulic model. The structure is modular, allowing different choices to be made if refinements are needed. The process starts with a large set of vectors simulated from the multivariate distribution of the source variable. Physically, this corresponds to vectors containing water levels for all locations within the aggregation area. Each simulated set of water levels can be thought of as representing a plausible event that could occur. We therefore refer to the set of load vectors as an ‘event set’.

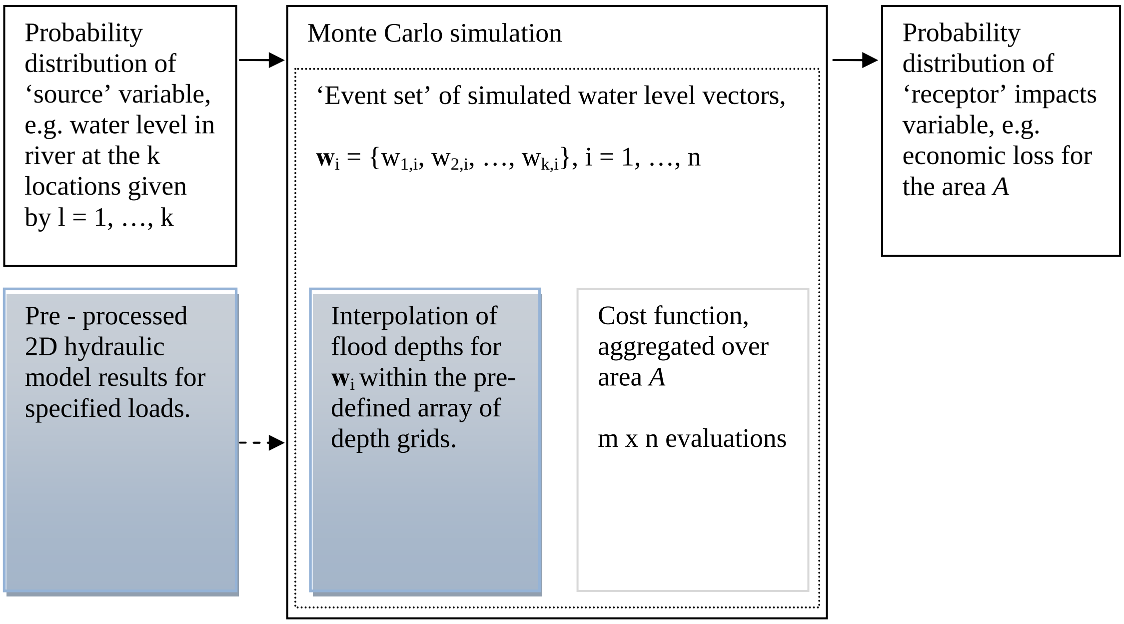

Figure 2-6: Generic spatial risk model structure with exhaustive evaluation of the pathway model replaced by emulation (shaded boxes). $w_i = \{w_{1,I}, w_{2,I}, \ldots, w_{k,i}\}$ **is a spatial vector of hydraulic loads, such as water levels, at each of k locations, for the i^{th} of n Monte Carlo samples.**

The interpolation method used within this procedure is also open to choice. One option may be to interpolate based on volumes flowing onto the floodplain. This is the method that has been used for the Environment Agency’s north east region to provide a regionally-consistent flood risk assessment for 65 flood risk management and catchment change scenarios for 10,000km of watercourse.

Although the most obvious physical variable to represent the source of flood risk is the water level, it is difficult to construct useful water level estimates for risk analysis over large areas because of the lack of comprehensive data to describe river channel hydraulics or rating curves. Working with levels is difficult when a model may have to combine uncertain representation of the river channel with floodplain topographic (DTM) data of varying accuracy and precision. These problems are also faced by

NaFRA, where volumetric approaches have been used to avoid having to specify water levels when there is too much uncertainty.

Another option, that may be efficient for large-scale risk modelling, is to define the event set not in terms of direct physical loading (such as water level) but instead in terms of the 'intensity' of the load. The most convenient measure of intensity is the probability of the load. In this case, depth data per event can be interpolated using probability as the indexing variable. This provides a convenient opportunity to make use of existing flood depth grids produced, for example, for flood mapping projects. The model structure is illustrated in Figure 2-7, where the shaded boxes again represent the module that stands in for the full evaluation of a hydraulic model.

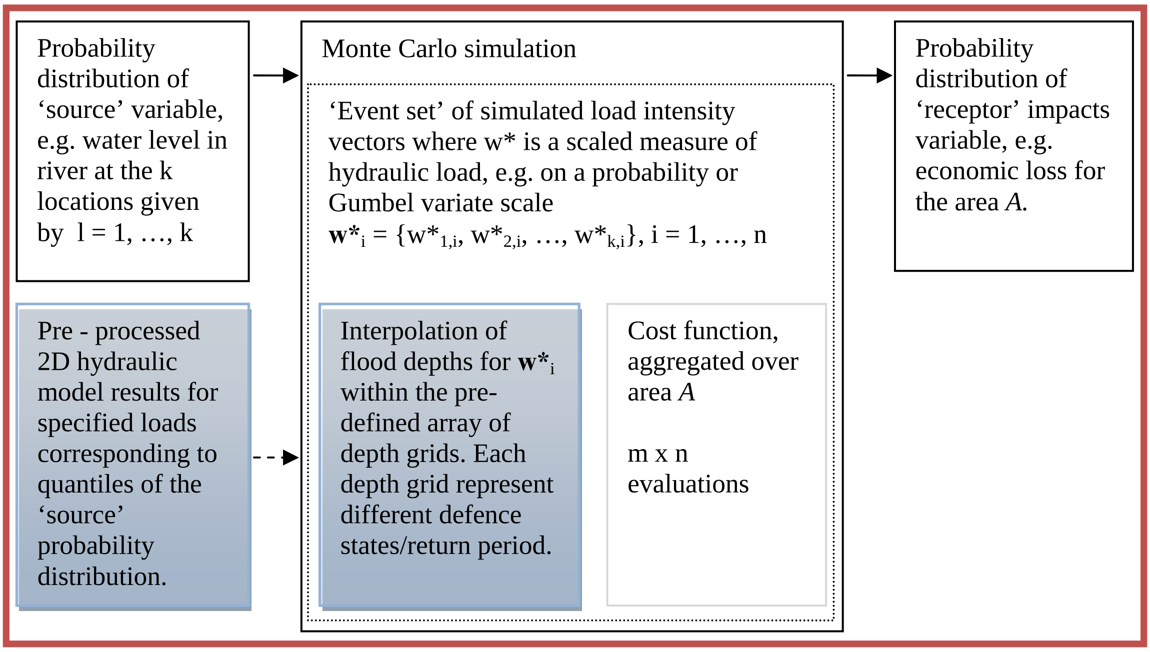

Figure 2-7: Generic spatial risk model structure with emulation (shaded boxes) of pathway model and hydraulic load expressed using an 'intensity' variable, such as probability, rather than water level. $\mathbf{w}^*_i = \{w^*_{1,i}, w^*_{2,i}, \ldots, w^*_{k,i}\}$ is a spatial vector of the 'intensity' variable at each of k locations, for the i^{th} of n Monte Carlo samples.

We have used this approach in the demonstration cases for this project. However, by exposing the event set as a complete entity, the generic approach above does allow for other choices to be made about how the simulation method is carried out.

Flood defence system 'pathway'

The addition of flood defence systems to the risk 'pathway' adds greater complexity to the above model. In the generic model, it would be straightforward, if expensive, to generate defence states and compute flood depths for each Monte Carlo step. To reduce the computational expense, an approximation could be made, similar to the approach described above, based on simulating defence states and then interpolating the resulting flood depths between evaluations of a hydraulic model for pre-defined combinations of defence failure scenarios and load 'intensity'.

Depending on the number of defences within a defence system, this could greatly increase the required number of hydraulic model evaluations, but some constraint could reasonably be put on the number of defence failure scenarios (as in Hall *et al*., 2003) because of the very low probability of the most 'extreme' failure cases. For rivers, this is further justifiable because there is limited volume available to flow to the floodplain, regardless of how many defences fail.

The Environment Agency has adopted the RASP approach to represent the performance of flood defence systems in modelling flood risk. It is therefore important to explore how the generic model for spatially aggregated flood risk can be aligned with RASP. In this project we have established that a link is possible in principle. This is discussed further in Section 7.

3 Flood risk methods review

In Section 2 we outlined a generic conceptual flood risk model based on the source – pathway – receptor concept. In this section we review previous research projects to identify which aspects of flood risk they relate to and how they are relevant to analysis of widespread flooding.

3.1 Methods reviewed

Table 3-1 summarises the projects reviewed. We have also included a number of more academic studies of flood risk variables, which are referred to in the text.

FD2308 Joint probability: dependence mapping and best practice	Review of existing (as of 2002) methods of joint probability analysis and extremal dependence estimation. Analysis and mapping of pairwise dependence between variables that relate to river, tidal and precipitation.
FD2105 Improved methods for national spatial-temporal rainfall and evaporation modelling for broad scale modelling	Development of a complex statistical model that can be used to simulate rainfall at differing temporal and spatial scales. A method of modelling potential evaporation at a single site was also developed.
FD2113 Spatial-temporal rainfall modelling with climate change scenarios	Applied the method developed in FD2105 under a wide range of climate change scenarios.
Environment Agency rainfall and weather impacts generator (EARWIG) and UKCP09 weather generator	Weather generators that, although intended to be used at a single site, can be used to generate weather data for small catchments. Incorporated into the UKCP09 climate change projections.
FD2020 Regionalised impacts of climate change on flood flows	An examination of the ways in which changes in monthly weather may affect changes in flood flows.
Risk assessment for flood and coastal defence for strategic planning (RASP)	Family of methods based on simulation of flood risk source and pathway using the fragility curve concept to represent dependence between the two.
Modelling and decision support system 2 (MDSF2)	Software application being developed on behalf of the Environment Agency to provide a flood risk analysis tool that includes probabilistic treatment of defence systems (using RASP).
Social vulnerability mapping	Mapped high concentrations of people with high vulnerability to flooding and critical infrastructure.

Table 3-1: Summary of flood risk management projects reviewed.

3.2 Dependence in the source - pathway - receptor model

Although it is convenient to break the analysis of risk down into source, pathway and receptor, it is also important to take account of the interactions, or dependence, between the components. To date, attention in the flood risk management literature has concentrated only on certain aspects of dependence.

One area that has been considered is the relationship between source and pathway, which has been represented using the fragility curve concept where the failure probability of defences is made conditional on the magnitude of the loading. Another important aspect of dependence that has been studied in detail is the relationship between tide levels, wave and storm surge.

In both cases, established approaches only provide a limited description of the full spatial and temporal structure of the dependence between variables. Two examples are the RASP approach for modelling flood risk in the presence of defences and the POL112 method for sea level extremes. The RASP approach includes the fragility curve concept and also builds in some spatial structure by considering systems of defences, though the structure is a fixed one, imposed a priori by the analyst. More sophisticated approaches to represent spatial dependence in defence failure have been discussed but are not routinely available. The POL112 joint probability approach for extreme sea levels is essentially based on a point in space, though there is a deterministic interpolation of the marginal probability parameters over space to account for the spatial coherence of the surge process.

3.3 Projects investigating the flood risk ‘source’ term

The projects that relate mainly to characterising the source component of flood risk are FD2308, FD2105, FD2113, EARWIG and FD2020. The projects FD2105, FD2113 and EARWIG all simulate spatially and temporally coherent rainfall which can then be used in conjunction with a rainfall-runoff model to simulate river flows. The extremes of these river flows can then be analysed to assess the likelihood of multiple floods. Project FD2308 looked at existing (as of 2002) methods of joint probability analysis and extremal dependence estimation and mapped the pairwise dependence between extreme river flows, precipitation and sea surge. The ongoing project FD2020 is examining the ways in which changes in climate may affect flood flows.

3.3.1 Rainfall

Defra project FD2105 (Improved methods for national spatial-temporal rainfall and evaporation modelling for broad scale modelling, Wheater *et al.*, 2006) developed a complex suite of statistical models that can be used to simulate rainfall at differing temporal and spatial scales. A method of modelling potential evaporation at a single site was also developed. The rainfall simulation method can be used to simulate rainfall at time intervals of greater than an hour and at spatial scale up to a medium sized catchment.

The project FD2113 (Spatial-temporal rainfall modelling with climate change scenarios) uses the outputs from different climate models to change the parameters of the rainfall simulation model. The differences in simulated rainfall for different climate models were examined and a method of combining these outputs was also developed. One of the findings of the FD2113 project was that predicted changes in rainfall vary greatly depending on the chosen climate model.

The main output from the EARWIG project is a combined stochastic rainfall model and weather generator. The rainfall model is a non-spatial point process stochastic approach based on the Neyman-Scott rectangular pulses model. A 'weather generator' approach is used to generate data for temperature, vapour pressure, sunshine and wind speed conditional on the simulated rainfall. The rainfall model is essentially a point model, where the 'points' are defined on a 5km grid. However, it also incorporates a facility to work with precipitation averaged over an area to generate data for small catchments (areas less than ~1000km^2 are suggested in the EARWIG documentation).

A similar model is included with the UKCP09 climate change projections. It can also be used to simulate weather data for a large number of defined future climate change scenarios. This is done by applying change factors, derived from the climate model outputs, to perturb the rainfall statistics that are built into the model. The change factors are defined at the climate scenario grid scale (25km). Spatial patterns in the change factors are fixed. The spatial grids of UKCP09 projections are not spatially coherent, and therefore do not support generation of large scale rainfall fields. This and other aspects of the use of UKCP09 are being considered by the Environment Agency under R&D project SC080004.

The method used to simulate rainfall developed in project FD2105 has two stages. The first stage is to generate a wet/dry state at a site, that is, whether or not that site has any rain, and then to generate a rainfall amount. Different generalised linear models (GLMs) are fitted for the two stages. The spatial structures of the two linear models are also modelled separately.

For each simulated rainfall event the number of wet sites is modelled using a beta-binomial distribution. This distribution can then be used to generate number of wet sites for each rainfall event. The positions of these wet sites are then allocated in such a way as to preserve the marginal probabilities of rain at each site. An algorithm for achieving this is described in Chandler (2002).

The amount of rainfall on each wet day is modelled using logistic regression at each site. The daily spatial dependence structure is then modelled by assuming a multivariate normal distribution structure for the Anscombe residuals. At each site these Anscombe residuals have a univariate normal distribution and are equal to Y_i/μ, where Y_i is the amount of rainfall on the i^{th} wet day and μ is the mean of the fitted model.

Buishand *et al*. (2008) estimated the total areal rainfall that is exceeded once in 100 years. The approach taken was to simulate a large number of synthetic daily rainfall fields that can be used to estimate the total areal rainfall that can be expected to be exceeded once in 100 years. The data used in this study were from 32 rainfall stations records located in North Holland. The length of record used was 30 years and the data were recorded at daily intervals. Only the autumn months were examined. In simulating the data three assumptions were made. The first was that sites that are located close together are asymptotically dependent. The second was that the dependence between sites that are far apart can be wholly explained by the intervening sites. Taken together these imply asymptotic dependence over the whole set. The third assumption was that the only factor affecting dependence is distance, this would not be appropriate for river flow modelling or for precipitation over substantial hills.

3.3.2 River flow

Donner (in press) used the standard multivariate statistical technique of principal component analysis to assess the overall level of dependence of river flows in a spatial region. In this work the proportion of principal components that are needed to explain a specified proportion of the total variance was used as a dependence measure. It was found that fewer principal components were needed to explain the variance within a flood event than are needed when the flows are at a more usual level. This suggests that the dependence structure becomes simpler as the river levels become more extreme. It also shows that any method used to describe the spatial dependence of extreme river flows must be able to handle this change in dependence at extreme levels.

The aim of the research in Troutman and Karlinger (2003) was to estimate the probability of any site in a region having a T-year flood in any given year. They called this probability the regional flood probability (RFP). The data used in this study were annual maxima data from Washington State. A variety of different methods were used to estimate the RFP. The first was to transform the data so that it had standard normal marginal distributions. The empirical distribution function was used to estimate the marginal distributions of the data. Then a multivariate normal copula was assumed, from which artificial data could be simulated. Another method was to use the average between-station correlation to directly influence the RFP.

3.3.3 Sea levels

A study that used the Heffernan and Tawn (2004) conditional method to examine spatial dependence is that of Latham and Tawn (2007). This paper is an application of the method to sea level data in the North Sea. One of the findings of the paper was that there was lower dependence between sites at the far south eastern corner of Britain with other sites on the east coast of Britain than between sites at the far north eastern corner of Britain with other sites on the east coast of Britain.

This is a feature of the spatial dependence of sea surges that was not clear from the work of Svensson and Jones (2002, 2004); the reason for this is probably because of the fact that the Heffernan and Tawn method is capable of providing more information on the extremal dependence. Another finding in Latham and Tawn (2007) is that as the sea surge events become more extreme, the effects of dependence were seen to become more localised.

3.3.4 Multiple variables

The aim of the project FD2308 (joint probability: dependence mapping and best practice) was to review current methods of joint probability analysis and extremal dependence estimation and to encourage the wider industry to use these methods. Another part of the FD2308 project was to analyse and map pairwise dependence between variables that relate to river, tidal and precipitation. This dependence mapping looked at same day and lagged extremal dependence.

It is preferable to look at dependence between variables that arise from the same meteorological event because it is this dependence that affects the probability of widespread floods. The work of Svensson and Jones (2002, 2004) is an investigation of the pairwise dependence between extreme sea surge, river flow and precipitation around the coast of Britain. Svensson and Jones (2002) focuses on the east coast. Their 2004 work, which forms part of FD2308, focuses on south and west coasts. The dependence measure used in these papers is χ, defined in equation (4.6.1). This measure has been described in detail by Buishand (1984) and Coles *et al*. (1999). Buishand used it previously to describe the dependence in precipitation data in the Netherlands. By using the dependence measure χ, one of the main outcomes of these studies was a set of maps describing pairs of variables that exhibit asymptotic dependence. The variable pairs that were examined were all possible same-variable pairs (surge-surge, flow-flow and precipitation-precipitation) and also all different variable pairs (surge-flow, and so on). Same day dependence was examined along with dependence at a lag of plus or minus one day.

The main findings of these studies of Svensson and Jones are as follows. Dependence in extreme sea surges is higher on the west and east coasts than on the south coast. The dependence in extreme river flows from different catchments was found to be higher in areas draining to the west coast than the east coast, whereas the dependence of extreme precipitation is higher on the east coast than the west. It was also found that there was evidence of asymptotic dependence of sea surges over longer distances than for river flow or precipitation. It is expected that this should be the case due to the physical nature of the processes involved. The highest dependence for flow-surge and precipitation-surge pairs was found on the east coast of Scotland and the western half of the south coast of England. The seasonal analysis revealed that there was no consistent change in dependence between winter and summer for sea surges. However, both precipitation and river flows exhibited higher levels of dependence for the winter months.

Another approach to examining the dependence used in the Svensson and Jones papers was to look at how the differing storm tracks of mid-latitude cyclones cause extreme observations in single, or multiple variables. It was found that differing routes and speeds of storm tracks have an effect on which variables have simultaneous extreme observations.

Ancona-Naverrete and Tawn (2002) use the coefficient of tail dependence, η, developed by Ledford and Tawn (1996) to describe how extremal dependence changes with distance. The two variables studied were rainfall in south-west England and sea surges on the east coast of Britain. They found that rainfall was asymptotically independent but strongly correlated; that is, many sites are likely to experience an extreme event at the same time, but unlikely to experience their very largest event at the same time. The sea surge variable used in this analysis was the surge at high tide. They found evidence of asymptotic dependence of surges at distances up to 250km. At distances between 250km and 400km they found evidence of decreasing extremal dependence. At distances of greater than 400km they found that the level of extremal dependence was stable.

The work of Keef (2007) and Keef *et al.* (2009b) uses the Heffernan and Tawn model (Heffernan and Tawn, 2004) to estimate the risk measure $N(p)$, defined in equation (4.7.4) to examine the strength of spatial dependence in extreme daily river flows and extreme daily rainfall over Britain. The dependence of these two types of hydrological variable was examined in two separate studies. We define the set of variables of interest as $(X, \boldsymbol{Y})$ and we condition on variable X. The risk measure $N(p)$ can be thought of as the expected proportion of variables in $\boldsymbol{Y}$ that are extreme within a certain time window of t given that X is extreme at time t. For the river flow analysis the time window was taken to be plus or minus three days and for the rainfall analysis the time window was taken to be zero days (in other words, same day dependence). The site X was taken to be each measuring station in turn and the set of variables $\boldsymbol{Y}$ was taken to be all other measuring stations of the same type within a certain radius of the site X. These radii were taken to vary between 30km and 120km. A number of different probabilities, p, were also examined, these ranged from the probabilities corresponding to the 0.1 year to 55 year return periods for the river flow analysis and 0.1 year to 548 year return periods for the rainfall analysis. These probabilities are the probability of a certain daily mean flow or daily precipitation being exceeded on a particular day.

For river flows the spatial dependence was highest for areas that had homogeneous catchments, such as the steep valleys of South Wales, for example. Areas such as the upper reaches of the Thames catchment, which has both permeable and impermeable sub catchments, and the Lake District, where some gauges are located upstream of lakes and some downstream, have much lower levels of spatial dependence. Other findings from the river flow analysis were that dependence decreased as the probability threshold increased and also that the dependence decreased with distance. The effect of seasonality on the dependence structure was not examined. For rainfall, the overall level of spatial dependence was lower than that for river flow. The main factor that seemed to affect the spatial dependence of rainfall was the hilliness of the region. The flatter areas of south-east England exhibited a higher level of dependence than the more upland areas of the north and west. Similar to the river flow analysis, the level of dependence decreased with increasing threshold and increasing distance. In addition to the whole data series analysis, the data from the summer months and the data from the winter months were analysed separately. In general, the level of spatial dependence in summer was less than that in winter.

3.3.5 Climate change

The project FD2020 (Regionalised impacts of climate change on flood flows) takes an alternative approach to assessing the impacts of climate change. Instead of examining the outputs of existing climate models and seeing how they vary from each other and the present day, the project FD2020 will look at ways in which changes in monthly weather may affect changes in flood flows. By taking this approach it is hoped that the results can be used with both current and future climate models. The FD2020 project will look at both small and large catchments over the whole of Britain (England, Scotland and Wales). It will examine the effect of climate change on different types of catchment where catchments are classified according to their catchment descriptors.

The project FD2308 also looks at the effect of climate change on dependence. To do this it looks at the outputs from two climate models and how the dependence associated with these outputs differs from the current records.

3.4 Projects investigating the flood risk ‘pathway’ term

3.4.1 The RASP family

The Environment Agency and Defra initiated development of the risk assessment for flood and coastal defence for strategic planning (RASP) methodology in 2001, with the aim of providing a hierarchical risk-based analysis framework to help assess flood risk and in particular how flood defences, and investment in flood management, influence flood risk. There have been numerous reports and other outputs produced since then and the RASP approach has come to be regarded as a ‘family of methods’ for flood risk analysis. Key reports are Environment Agency (2004, 2007).

In principle, the RASP methods are quite simple. The idea is to represent flood risk by integrating the probability distributions of the source (which is usually interpreted as water level loading), the pathway (flood defences) and, where relevant, the receptor (for example by including economic damages as a function of flood depth).

Conceptually, the approach follows on closely from the methods proposed by the US Army Corps of Engineers in the 1990s (National Research Council, 2000) and thereby represents a significant shift in thinking away from deterministic 'standards of protection' towards a risk-based methodology for flood management.

The integration of the various probability distributions is most easily accomplished using a stochastic simulation approach, and the 'RASP family' is in effect a set of related simulation tools that allow the underlying concept to be applied at different levels of detail and with different types of data. The stochastic simulation consists, in the most basic form, of a large number of possible 'flood events' being generated from the probability distribution of the source variable (such as water level), which are then taken as inputs to a model of the defence system behaviour. The complexity of the published reports and methods stems largely from the need to accommodate a range of different circumstances within the simulation approach.

Whilst the aims of the RASP method were quite general, one central aspect is the influence of flood defences on risk. There can be a spatial dimension here, because the risk at a particular location is generally influenced not just by defences at a single point, but by a whole defence system. There are also important dependencies to consider, which could include:

- Local dependence between loading (the flood 'source term') and probability of defence failure.
- System wide dependence in the relationship between loading and probability of defence failure.
- Spatial dependence of loading, in the presence or absence of defence failure.

The RASP methods use the fragility curve concept to represent the first bullet point. A fragility curve specifies how likely a defence failure is given specified water level loading. The last two items are not fully included, as we will discuss below.

3.4.2 Dependence assumptions in RASP

Current RASP implementations make assumptions of 'Dependence of load' and 'Independence of defence strength'. More sophisticated versions of RASP may progressively relax these two assumptions. Pointers to future development were given in an Environment Agency report 'Scoping the development and implementation of flood and coastal risk models', which is the main output from R&D project SC050065.

Dependence of load

This assumption means that the load applied to a defence system has been taken to be equal over the system. It is a fixed, deterministic dependence. A 'defence system' is taken here to mean the combination of flood defence assets that together provide some protection from flooding along a river or coastline.

Future RASP tools are likely to allow for spatially varying loads. An example of how this could be important is where local features in bathymetry and the direction of storm travel can lead to significantly different conditions at adjacent positions on a coastline. Incorporating this kind of spatial variation is relatively straightforward as forecast loads are generally known (from a numerical model output) and hence can be included within the RASP analysis without changing its basic form.

It should be noted that this is a fixed form of dependence rather than the association between random variables that is implied by the term 'dependence' elsewhere in this report. Hence although allowing spatial variation in loading would relax the original assumption of equal loading on a defence system, the relaxation envisaged by SC050065 remains purely deterministic. Over a relatively small scale this is reasonable. However over larger scales, particularly for fluvial flood risk, it would be implausible to assume a deterministic variation in loading within a probabilistic method.

Independence of defence strength

This assumption means that the probability of failure in a given defence system component (which is in effect a point-wise estimate) is taken as not having any influence on other components. The two main sources of dependence in defence systems are the effect of breaching on the loads (water levels) that neighbouring defences are subject to and the spatial dependence in statistical variation in strength.

Relaxing the assumption of independence of strength increases the computational complexity of the risk simulation, requiring secondary influences on asset strength to be included. For example, it could be necessary to model the probability of failure in several different locations given dependency in the physical condition of the defence asset and also in the loading over the system. The scoping study for future RASP methods concluded that it is likely to be sometime before such a capability exists.

A physical example of this type of problem could be a defence system where failure of a flood gate allows scour and progressive failure of an earth bank, simultaneously changing water levels experienced on another part of the defence system a short distance downstream. Representing this type of scenario in a deterministic simulation is difficult enough. A stochastic simulation method would then have to represent numerous possible scenarios of this type, each with a realistic frequency of occurrence within the set of simulated events.

3.4.3 Spatial scale and dependence

The discussion of dependence in RASP is concentrated at the spatial scale of a discrete defence system (that is, a collection of defence assets, gates, embankments and so on that can be treated as a single system protecting a given flood area). Large scale spatial variation of loading, in the stochastic sense that we understand it for this project, is not included. This is unlikely to be a limitation for a small scale analysis at the asset system level.

However, for larger scale analysis, in particular the development of national or regional ‘risk profiles’, it is not realistic to assume a deterministic spatial pattern (including constant loading) for the stochastic simulation of the source of flood risk. It is at the larger geographical scale that the present project can make the most significant contribution to risk assessment that may currently use RASP-based methods. There is no conflict between the approaches; rather the development of models for spatial dependence in sources of risk enhances the capability of RASP or other risk assessment methods.

This issue of dependence over different scales is clearly very relevant to the continuing development of NaFRA. Without accounting for the dependence, there would be bias in aggregated risk profiles, for the reasons discussed in Section 2 of this report.

3.5 Projects investigating the flood risk 'receptor' term

3.5.1 Modelling and decision support framework (MDSF)

The modelling and decision support framework (MDSF) tool was developed initially in 2001 to provide a tool for quantifying economic and social impacts of flooding at a catchment scale for present day conditions, future scenarios and with flood management options. It draws on water levels generated from external hydraulic models and includes simple methods for linearly spreading these water levels and evaluating flow depths. It has been applied widely for flood and erosion risk assessment as part of catchment flood management plans (CFMP) and shoreline management plan (SMP) programmes and has also been used on strategy studies, pre-feasibility studies and other similar scheme appraisals.

The present version of MDSF, however, uses a simplified representation of the role of defences and does not properly take account of defence performance in the analysis of risks and their management. This is a particularly crucial point in the context of understanding and managing actual risk. The Environment Agency has therefore commissioned the development of a second version, MDSF2, to incorporate RASP methods to take into account the performance of flood defences. (The project also addresses a number of software issues, such as GIS platform, which have been obstacles to widespread uptake within the Environment Agency). The MDSF2 will therefore include probabilistic flood modelling algorithms, based on rapid, physically approximate flood spreading (though it is also planned to have the capability to add in more detailed, external model data where this is useful).

MDSF2 will incorporate the National Property Database (NPD2) dataset (amongst others such as the Valuation Office Database and Multi-coloured Manual) for estimating property damages, however, the user will have the flexibility to edit or change these values. It will allow estimation of impacts such as property and agriculture damage, indirect infrastructure costs, risk to people and a range of indicators such as numbers of properties at risk, length of rail or road at risk (using simplified methods).

The MDSF2 has been referred to in its design report (Environment Agency R&D report SC05001/SR 'Task 1g - MDSF2 System Design', undated) as a tool that can implement the RASP methods 'at any scale'. From the perspective of this report, it should be apparent that there is a gap in current probabilistic risk methodologies at scales large enough for dependence in the source of flood risk to become important. The current project will identify methods to account for the spatial dependence of sources or risk that will be capable of linking with the algorithms in MDSF2. This would provide the capability to allow a scale-independent analysis of risk within software applications such as MDSF.

3.5.2 Broad scale modelling data for the North East

In 2007 the North East Region of the Environment Agency commissioned a broad scale modelling study to produce consistent flood risk data for the region using a 2D modelling approach. The completed data set assimilates 1,250,000 separate 2D gridded flood depth models into risk information suited to CFMP or strategic studies ranging from regional spatial strategies to strategic flood risk assessments. The data produced for the study were built into a database containing 19 new flood depth grids and 19 indicative hazard index grids over the North East Region, covering 190,000km of river length. These depth grids were completed in four months using a parallelised 2D diffusion wave flood model run on a 10x10m spatial grid. The flow data used in the modelling was derived from automated Flood Estimation Handbook (FEH) estimates licensed by the Environment Agency (but interpolated to intermediate return periods) with some local adjustments based on more detailed comparisons with existing modelling studies. Defences were represented using simple standard-of-protection assumptions.

The data were supported by GIS-based software to allow scenario exploration, including climate change scenarios, vulnerability maps, flood risk to people, and probability maps. Reporting was per asset system area, but other spatial aggregations could also be adopted. The spatial (grid or outline) and frequency distribution outputs include:

- Flood depths at properties.
- Damage at properties (based on the Flood Hazard Research Centre depth-damage curves giving the same output as MDSF NPD2).
- Hazard ratios at properties.
- Vulnerability indices based on spatial analysis of Census data.
- Consequence scores that combine hazard with vulnerability.
- Annual average damages expressed as a function of defence standard for property and for agriculture.
- Agricultural damages aggregated over asset system areas (calculated as per MDSF1).
- Depth grid statistics.
- Probability maps.

Although the database produced for this study was designed to assist with scenario exploration, it is also well suited to demonstrating this link with receptor impacts within an efficient simulation framework, which, as we will discuss in Section 7, offers a potential route to deliver a spatially aggregated model of flood risk.

3.6 Flood risk analysis in the insurance sector

3.6.1 Catastrophe models

The large scale spatial aspects of flood risk are of direct concern to the insurance and reinsurance industries. The insurance companies take reinsurance cover (or use equivalent financial instruments) to ensure sufficient capital is available in the event of a catastrophic loss. Reinsurers need to know the probabilities of extensive flood scenarios in order to guide their pricing. Reinsurance pricing is generally guided by a combination of actuarial analysis of historical events and the outputs from catastrophe ('cat') models. Here we provide an outline of the general principles of cat modelling and give three examples. For a detailed description of cat modelling and more examples see Sanders *et al.* (2002).

The details of each cat model differ, however there are certain broad principles that emerge. In general, cat models are made up of four different modules. The first is a stochastic model which randomly generates a synthetic sample of catastrophic events, corresponding to the 'source' of risk in S-P-R terms. The second component, the hazard module, takes each random catastrophic event and derives its physical impact (for example, flood depth). The third module is a vulnerability function that calculates the damage to buildings and other structures, likely business interruption and other financial implications of each event. The final component is a financial analysis module that, in the simplest case, adds up the total financial loss for each event (though in practice the calculations are more complicated). The financial losses from each randomly generated event can then be ranked in size order and used to produce a risk profile, which can then be used to asses total exposure and hence inform pricing decisions.

3.6.2 Examples

Detailed information about the workings of individual cat models is generally commercially privileged. However, the principles of models produced by major risk analysis firms are published or presented at technical seminars. Here we summarise information known about three models.

One example is the river flood model developed by Risk Management Solutions (RMS). The probabilistic component of the RMS model is produced by simulating spatial and temporal weather data (precipitation, potential evaporation and temperature). The simulation is carried out by stochastic perturbation of meteorological data from historical patterns, with variations in season and altitude taken into account. We have not been able to find sufficient published details of the mechanism for modelling dependence and for simulation of spatial rainfall fields to comment at this time on how well the simulation would account for the statistical properties discussed in Section 4.2 of this report. In general, it is worth noting that simulation based on perturbing observed data leads to a dependence structure that will not weaken as larger events are produced, and so implicitly assumes asymptotic dependence.

The hazard component of the RMS river flood model incorporates a soil moisture component, a rainfall-runoff component and a flood depth model. The soil moisture must be treated as a random variable and can be expected in reality to have a complex dependency on the rainfall history - it is not known whether this is represented by the model. The flood depth model takes into account the probability of defence failure using the fragility curve concept. The vulnerability model includes information on the vulnerability of different types of buildings and the location of these buildings.

The ABI East Coast Flood Study, carried out by Entec, RMS and Risk & Policy Analysts, examined the spatial impact of storm surge events on the east coast of England using the RMS storm-surge model (Thurston et al., 2007, Muir-Wood, 2005). Three storm surge scenarios were considered, focussing in turn on the Humber Estuary, the East Anglian coast and the Thames Estuary. Each had a return period of 200-250 years. The model used by RMS generates storm surges for the whole of the east coast of Britain based on a stochastic wind storm model and correlation analysis of surge heights. It also incorporates aspects of defence failure. This study also examined the effect that sea level rise and how possible changes in the distribution of ages in the population may affect the likely impact of storm surge events. Another aspect of the study was to assess the economic benefit of improving coastal defences.

A third cat model example is the Swiss Re exposure calculation tool for the Czech Republic. Unlike the RMS model, the Swiss Re approach is to produce spatially realistic events using river flow data. The Czech Republic is split into Administrative Units (AUs). Each of the simulated flood events consists of a synthetic flow 'observation' at each AU. The data used to derive the model are monthly maxima records from gauging stations over the Czech Republic. These records were then transformed to have standard normal marginal distributions. The spatial dependence is modelled using a multivariate normal copula. Monthly maxima flows at AUs that do not contain a gauging station are obtained by interpolating the gauged data using kriging. So, to generate stochastic flow events, the covariance matrix of the observed flow data is calculated and events are then generated from the multivariate normal distribution with this covariance matrix. These generated, normally distributed, events can then be used to work out a corresponding water depth and return period of flooding at each AU.

3.7 Discussion

Both recent advances in flood risk management and insurance industry cat models adopt stochastic simulation to perform the integration of the complex joint distributions of sources of flood risk and the defence pathways and 'receptors' (vulnerability functions). The basic structure of the simulation approach is also relevant for the conceptual model development in this project.

Within this common basic framework, the various cat models and flood risk management models can differ greatly in the assumptions made about the distributions of the variables, whether explicitly or implicitly. In particular, assumptions about the dependence structure can be quite subtle, but ,as we show in Section 4.5 of this report, have a conspicuous effect on the results of the risk analysis. This is why we propose to adopt an approach that is designed to be flexible enough both to test and represent the dependence structure found in the data.

4 Statistical methods review

In this section we discuss the statistical science relevant to understanding and modelling the spatial structure of data for flood risk analysis, in particular river flows (or levels) and sea level.

In the statistical review, we start by identifying features of the data that need to be considered and then assess how well the available and developing statistical techniques can incorporate these features.

We then summarise the theory behind fitting marginal distributions to extremes of variables and the principles behind the techniques of separating marginal and dependence characteristics (the theory of copulas).

This section also describes relevant groups of statistical techniques. These groups are, broadly speaking:

- Methods that are only suitable for asymptotically dependent data (where the largest values of each variable tend to occur together).
- Methods than can only be used for pairs of variables.
- Multivariate methods that do not account for changes in dependence at extreme levels.
- Multivariate extreme value methods that do account for changes in dependence at extreme levels.

4.1 'Dependence' and 'coherence'

Two variables are statistically dependent if the value of one of the variables affects the likely values of the other variable. Dependence between different locations is arguably the most important feature in a statistical model of flood risk if it is to work over a wide range of spatial scales. The spatial dependence structure is a description of the strength and type of dependence between one or more variables over a range of distance scales.

We use the term 'spatial coherence' to describe the dependence between the parameters of the marginal distributions of variables at sites that are geographically close together. In other words, the coherence tells us about how smoothly the probabilities of an individual 'flood risk variable' would vary between a given location and its neighbouring region.

The coherence can therefore be seen as a 'by-product' of a model for how the full joint distribution varies over space. Although the project is called 'Spatial coherence of flood risk', it is in fact the dependence structure that has to be studied rather than the coherence alone.

4.2 Key statistical issues

A suitable model for use in this project must be able to give estimates of the likelihood of spatially extensive floods and of simultaneous floods from different sources. In order to do this it must take account of the following features.

4.2.1 Joint distribution function

The chosen model should describe the full tail structure of the joint distribution. This joint distribution is required so that the probability of all possible combinations of the source variables, either simultaneously or at various time lags of interest, can be estimated. Consequently we require a model for both the marginal distributions and dependence structure (both simultaneously and at different lags) between of a set of variables.

4.2.2 Multidimensionality and consistency in marginal distributions

The model should be able to handle multivariate datasets of high dimensionality. Another desirable feature is that the absence (or presence) of any variable should not affect the marginal characteristics of, or dependence between, the remaining variables. It must also be able to handle multivariate data sets of low dimensionality, such as, sets with only two or three variables.

4.2.3 Change in dependence at extreme levels

As a pair, or set, of variables become extreme the dependence between them may change. If a pair of variables is highly dependent at lower levels it does not automatically follow that they will be highly dependent at extreme levels. Two differing examples are flood flows in different areas of the country and extreme sea levels and river flows. On a normal day-to-day basis river flows in different areas of the country are likely to be highly correlated. This correlation is largely due to the fact that flows are dominated by seasonality and average previous rainfall, both of which are largely consistent over the whole country. However, extreme flows are usually precipitated by a single heavy rainstorm on wet ground. The spatial extent of rain that is heavy enough to cause extreme flows is much less likely to cover large areas, and so we would expect a decrease in dependence in extreme flows.

The opposite change in dependence is true for river flow and total sea level. There is evidence for extremal dependence between river flow and sea surge. However, total sea level is usually dominated by the variation of the tidal component so when the sea level is lower, sea levels and river flows may appear to be independent as the tide is independent of the river flow. But the highest sea levels are caused by large surges occurring with the highest tides. As the largest tidal levels are similar to each other the dependence between large river flows and large sea surges is likely to become apparent through a stronger dependence between sea levels and river flows in the largest events. A similar feature can also be deduced from the increases in dependence observed between wave height and sea level for increasing threshold levels in the FD2308 R&D project (Defra/Environment Agency, 2005a).

4.2.4 Incorporating both extreme and non-extreme values

The chosen model should be able to model events where some variables are extreme, other variables are moderately large, and some are non-extreme. Most established approaches for modelling the joint tail distribution, or associated summary measures, make an inappropriate assumption (at least for river flows) that the very biggest events in all the variables have a possibility of occurring together. Hence, they can only handle occurrences when all the variables have an extreme observation and are restricted in practice to small numbers of sites that are close together in space. There is further discussion of the different types of dependence, and how dependence assumptions can be checked, in Section 4.6 of this report.

4.2.5 Differing temporal and spatial scales

The chosen method should be able to model dependence of different temporal and spatial scales. It must be able to handle temporal data observed at a temporal resolution capable of resolving flood peaks up to the duration of large flow events. It must also be able to handle temporal dependence in order to identify which physical occurrences qualify as single events, for example when we ask whether flooding within the same hour, day, week or month counts as a single 'event'. In terms of spatial scale it must be able to handle dependencies over distances ranging from a few kilometres up to hundreds of kilometres. It must also be possible to express these dependencies in an understandable way at each of these temporal and spatial scales.

4.2.6 Covariates

A model suitable for the objectives of this project should ideally be able to incorporate covariates to bring additional information about spatial patterns. These covariates could include seasonality, meteorological indices, spatial location and geological information. This is a desirable feature although perhaps less important than those discussed above.

4.2.7 Impacts of flooding on 'receptors'

Methods of describing these impacts at the receptor are discussed in Section 3.5. One way to handle impacts of flooding on receptors spatially is to use a statistical model for the spatial source variables that can be integrated with existing methods for modelling pathway and receptor.

4.3 Marginal models

When fitting a statistical model to a single source variable it is tempting to fit a standard statistical distribution to the whole data. However, this approach is very unlikely to lead to a good fit in the tails of the distribution, which is the region of most concern when extrapolating. Instead, it is common to focus on the extreme values. When working with univariate extreme values there are two main approaches: block maxima (or minima) and threshold exceedances (often termed the 'peaks over thresholds' method). It is widely known that the block maxima method is less efficient in estimating return levels than the threshold exceedance approach.

The standard approach for modelling threshold exceedances is to specify a high threshold, for example the 95th or 99th per centile of the marginal variable, and to estimate the rate of exceedance of that threshold and the distribution of the excesses of that threshold. The model used for the excesses of the threshold is the Generalised Pareto Distribution (GPD) (see Davison and Smith, 1990), as this model has the following properties:

- Following asymptotic probability theory it is the only distribution with which it is possible to describe the excesses as the threshold tends to the maximum possible value of the threshold.

- If excesses of some threshold follow a GPD then excesses for any higher threshold will also follow a GPD. For marginal variable W the generalised Pareto has cumulative distribution function G_u with

$$G_u(W) = P(W < w|W > u) = 1 - \left[1 + \xi \frac{w-u}{\sigma_u}\right]^{-1/\xi} \quad \text{for} \quad u < w < \infty,$$

where $\sigma_u > 0$ and ξ are scale and shape parameters respectively.

No attempt is made to fit a parametric model to the distribution of values below the threshold in this case. Davison and Smith (1990) show that the threshold exceedance rate and GPD parameters can be extended to include covariate effects.

Often in spatial studies the key covariate information is in the formulation of knowledge that the process being modelled is spatially coherent. The standard methods used to exploit this information are to impose that the rate and GPD parameters change smoothly over space. Two examples of the implementation of this type of strategy are given by Coles and Tawn (1990) and Butler *et al*. (2007), with the former using explicit parametric models based on distance between sites to impose smoothness and the latter imposing it through non-parametric smoothing. In both cases the key is that the shape parameter of the GPD is constrained to change very slowly over distance. Incorporating this information leads to improvements in estimates of return levels at different sites, particularly when some sites have more data than others and that data from longer records is implicitly transferred to the neighbouring sites.

There are a number of different approaches used in the flood risk management industry to estimate the return period of different river or sea levels. The approach recommended in the Flood Estimation Handbook (FEH) to estimate flood flows is to fit the generalised logistic distribution to the annual maxima. These fitted distributions can then be used to estimate the median annual flood (QMED) and the growth curve, which defines how flows at larger return periods relate to QMED. The FEH also contains methods for estimating QMED from catchment characteristics and for pooling data from different sites to estimate the shape of the growth curve. Another approach used to estimate the return period of different flows is to fit the Generalised Extreme Value (GEV) distribution to the annual maxima data.

Various approaches are used to estimate the return period of different sea levels. The simplest of these is to fit the GEV distribution to the annual maxima or a variant of this (the asymptotic distribution of the r-largest order statistics) to the r-largest total sea level data (Tawn, 1988). For sea level data this approach has the drawback that variation in the annual maxima sea level is often dominated by year to year deterministic variations in tide. So, in using this approach the fitted distribution is dominated by the tidal cycle, masking the contribution of the surge component to the total sea level. Hence extrapolation of the fitted distribution to long period return levels is biased and tends to underestimate return levels (see Dixon and Tawn, 1999).

Related approaches for estimating the return periods of total sea level that overcome this disadvantage are the Joint Probabilities Method (JPM) of Pugh and Vassie (1979, 1980) and the Revised Joint Probabilities Method (RJPM) of Tawn and Vassie (1989). Both of these methods separate the sea level into the tide and surge components and analyse the distribution of each separately. In the JPM the empirical distribution of the hourly surges is used, and in the RJPM a combination of the empirical distribution function for the non-extreme surges and a modelled distribution function for the extreme surges is used. The RJPM also takes account of the serial dependence of hourly tidal and surge measurements and tide-surge interaction. Dixon and Tawn (1994, 1997) illustrate the r-largest method, JPM and RJPM for a range of UK A class gauge sites. Dixon *et al*. (1998) apply a spatial extension of the RJPM to exploit similarity in the surge characteristics over sites. A version of the RJPM is currently being updated and systematically implemented for UK sea levels as part of Environment Agency project SC060064.

4.3.1 Transforming margins

Most methods for estimating the dependence structure of extreme variables rely on an assumption that the variables have a standard marginal form. This assumption is, in general, inappropriate for most data sets. However, it is possible to change the distribution of any univariate random variable. The transformation is performed by use of the probability integral transform.

The probability integral transformation states that if any continuous random variable X has distribution function F_X, then the random variable Y, defined by $Y = F_X(X)$, has a Uniform(0,1) distribution, so $Y \sim U(0, 1)$, and the converse also holds.

If we start with random variable X, with distribution function F_X, and we wish to transform this variable to a random variable Z which has distribution function H_X we apply the probability integral transformation in two steps. The first step is to transform the variable X to the random variable Y, with $Y = F_X(X)$, where Y has a uniform marginal distribution. The distribution function F_X that is used can be estimated by the methods described above for the marginal models or by use of the empirical distribution function. To transform the uniformly distributed variable Y to a random variable Z, which has distribution function H_X, we take $Z = H_X^{-1}(Y)$, that is, the inverse of the distribution function H_X.

4.4 Copulas

It is possible to separate the features of any multivariate distribution into marginal characteristics and dependence structure. This can be useful when examining the extremes of a multivariate process. For instance, if we are interested in the joint tail of a multivariate process where the variables have very different scales, or some of the variables have light tails and some heavy tails, then the extremes will be dominated by these variables with heavy tails, and/or larger scales. By examining the dependence characteristics separately from the marginal characteristics it is possible to obtain a clearer picture of the dependence structure.

This separation of marginal and dependence characteristics can be achieved using the copula function (Nelsen, 1999) which can be described as follows. Let $F(\boldsymbol{x}) = F(x_1, \ldots, x_d)$ be a d-dimensional multivariate distribution function with marginal distributions $F_i(x_i)$, $i = 1, \ldots, d$. Then we can write $F(\boldsymbol{x})$ as

$$F(\boldsymbol{x}) = C\{F_1(x_1), \ldots, F_d(x_d)\}.$$

The function C is called the copula. It has domain $[0, 1]^d$ and contains all the information about the dependence structure of $F(\boldsymbol{x})$. Each multivariate distribution has a unique copula so it is possible to construct any multivariate distribution using just its copula and marginal distributions.

A commonly used copula is the multivariate normal copula. Using this copula is equivalent to assuming the data have a standard multivariate normal distribution when the marginal variables are each transformed to following a normal distribution. One of the major attractions of the multivariate normal distribution is its flexibility; each marginal variable has a normal distribution, each linear combination of marginal variables has a normal distribution, all conditional distributions are multivariate normal and each subset of variables has a multivariate normal distribution.

Figure 4-1 shows scatter plots of data from bivariate variables with different dependence structures after transformation to uniform margins (which is, a simple standardisation of the data). Figure 4-2 shows scatter plots of the same data, but after transformation of the marginal distribution to Gumbel margins. When Figures 4-1 and 4-2 are compared we can see that the type of extremal dependence between pairs of variables is much easier to see when the data are plotted after transformation to Gumbel margins. For example, if we examine the plot of the data in Figure 4-2 (top right panel) simulated from a bivariate extreme value distribution we can see that the largest values of both variables occur together. This is a feature of asymptotic dependence, discussed in more detail in Section 4.6.1 below. However, if we examine the plot of data simulated from a bivariate normal distribution (Figure 4-2, top left) we can see that the largest values in both variables do not occur together (this is a feature of asymptotic independence, again see Section 4.6.1).

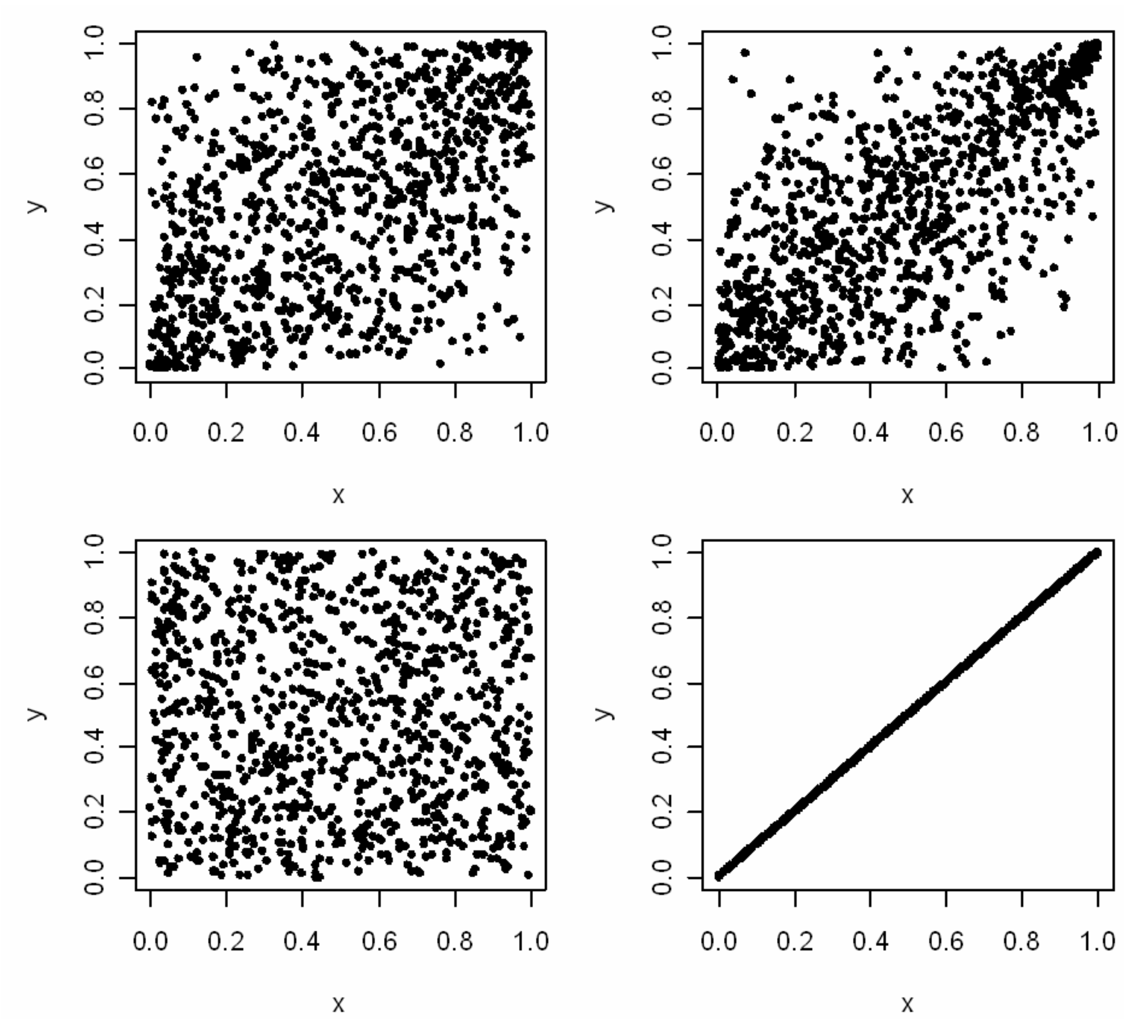

Figure 4-1: Data from different distributions on uniform margins. Clockwise from top left, Bivariate normal with ρ = 0.5, bivariate extreme value distribution with logistic dependence structure, α = 0.5, complete dependence and complete independence.

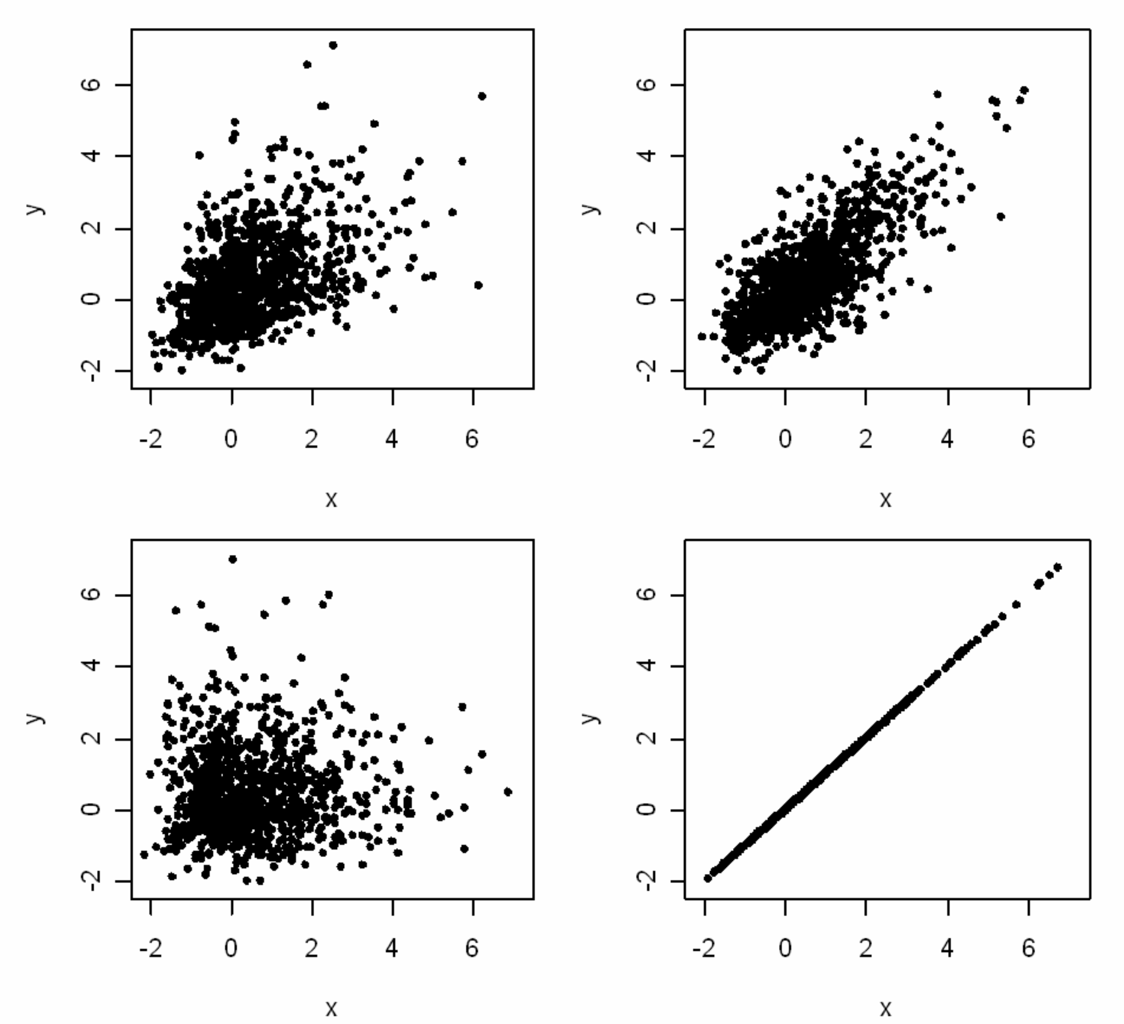

Figure 4-2: Data from different distributions on Gumbel margins. Clockwise from top left, Bivariate normal with ρ = 0.5, bivariate extreme value distribution with logistic dependence structure, α = 0.5, complete dependence and complete independence.

4.5 Consequences of incorrect estimation of dependence

The importance of correctly estimating the dependence structure can be seen in Figure 4-3. Here we show combinations of the bivariate variables discussed above plotted against return period (in relation to flood risk management, the bivariate data we have generated could, for example, be water levels at different locations).

If the minimum of the two variables is of interest (for example in assessing the risk of crossing a threshold) and we assume that they are independent when they are not, then we can see that we will underestimate the chance of large values. The same is true when the sum of two variables is of interest, for example to assess the risk of an aggregated cost function. Another feature it is possible to see from Figure 4-3 is that if we overestimate the dependence by assuming perfect dependence we will overestimate the maximum possible value of the minimum or sum of two random variables. The same is true with models for partial dependence, with different dependence models leading to different return level curves for the impacts at the receptor.

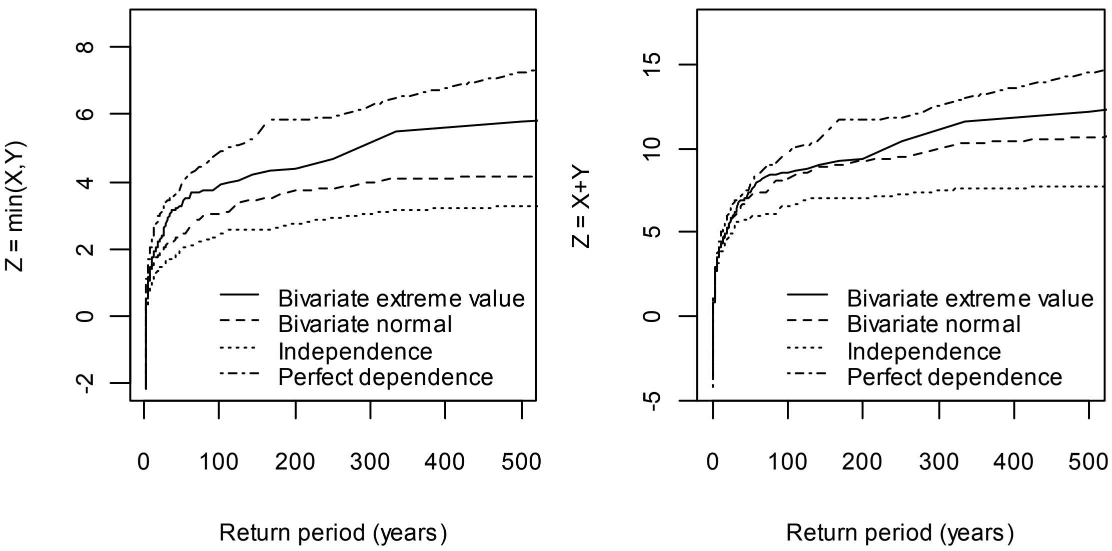

Figure 4-3: Distribution function of functionals of simulated data shown in Figure 4-2. The variables X and Y have been standardised onto the Gumbel scale.

Figure 4-3 illustrates a number of key aspects of the analyses to be undertaken in this project. Each subplot corresponds to the impact on a receptor, with different combinations of the source variables being important depending on the nature of the receptor. In Figure 4-3, it is large combinations as determined by the minimum and the sum of the source variables that is important. In applications the function of the source variables that is important will be determined by the form of the receptor and the pathway distributions. By controlling the design, the combinations of source variables that determine impacts on the receptor can be affected.

Here we have illustrated the evaluation of the distribution of the impacts on the receptor using a very large sample size, 1000 years of synthetic data. Of course in practice we typically have only between 5 and 30 years of data in any application so we cannot estimate these distributions beyond the five-year level with any reliability unless we adopt a statistical model. If a statistical model is adopted then we can simulate a large sample from that model and derive empirically any characteristics of the joint distribution of the source variables we are interested in, such as the distribution of impacts on the receptor.

Simulation from the fitted joint distribution of the source variables is therefore key to our evaluation of properties of interest. To be able to simulate in this way, we require a statistical model for the joint distribution of the source variables. This joint distribution of source variables will be fixed, but unknown, for a given problem. As discussed above, the joint distribution for the source variables requires models to be adopted for both marginal and dependence characteristics. Here we considered the margins to be known to be Gumbel distributed so that we can focus on illustrating the dependence characteristics.

Figure 4-3 illustrates that there can be considerable sensitivity on the distribution of impacts of receptor for different forms of dependence between the source variables. In practice a statistical model for the dependence structure has to be selected, so Figure 4-3 can be considered an illustration of the importance of this aspect of the statistical modelling; in other words, **to avoid bias, the dependence model that is used must possess sufficient flexibility to model a full range of different dependence structures between the source variables**.

From any model for the joint distribution of the source variables (however complex its marginal distributions or copula function) it is possible to simulate repeated realisations and hence integrate this with pathway information to derive the distribution of impacts at the receptor. As the marginal distributions of the source variables are likely to be similarly modelled by most analysts, it is tempting to think that the simulated impacts are realistic as they produce an impact distribution and have every element of the process modelled required in their derivation of this impact distribution.

However, even in the simple bivariate example, the results in Figure 4-3 show that the choice of statistical dependence model (copula) is fundamental in this process with careful selection necessary. The importance of this selection is magnified as the dimension of the problem is increased. In many applications, simple copulas have been used because of limited exploration of the importance of this step, of other possible copulas or of how they can be simulated from. For example, Buishand *et al*. (2008) simulate rainfall over the whole of the Netherlands using a dependence model that only allows a particular form of spatial dependence (asymptotic dependence, defined in Section 4.6.1) but no check on the adequacy of this model is made and if, as is likely, the process does not exhibit this particular dependence structure over the whole country then spatial dependence would be overestimated in this case.

To reduce the risk of making biased estimates of the impact distribution we propose an approach to statistical modelling of dependence that has two benefits:

- Specialist statistical tools are used to ensure that the selected models for dependence between variables captures adequately the observed dependence in data and any knowledge there exists about the physical structure of this dependence.
- We use a class of statistical dependence models that have proven mathematical properties giving them greater statistical flexibility than any other model for extreme values so that they can capture a very wide range of dependence forms between the source variables.

Details of the models we propose and some of the measures of dependence that we plan to exploit as tools for assessing the adequacy of the selected models for the data are described in the following section.

4.6 Multivariate extreme value methods

As in the univariate case, modelling the whole dependence structure via copula methods is likely to lead to substantial bias as the extreme events may possess a different dependence structure than the typical day to day values. Therefore simply taking a multivariate normal copula is not likely to be a sufficiently good model of dependence for the extreme events.

For example, extreme river flow events at different sites often show weaker dependence than non extreme events and so fitting a multivariate normal copula to all the river flow data in this case would lead to an over estimate of the joint probability of large river flows occurring simultaneously at the different sites. Therefore specific methods are required for describing the dependence of the extreme values.

4.6.1 Classes of extremal dependence – asymptotic dependence and asymptotic independence

Asymptotic dependence

There are two main classes of extremal dependence. The first is asymptotic dependence. If two variables are asymptotically dependent then the largest observations in both variables would occur together with a positive probability. Note that 'largest observations' in this context, and below, means the values approached asymptotically and not the largest real observations in a given sample.

Asymptotic independence

If two variables are asymptotically independent then the probability of the largest observations of each variable occurring together is zero. There are three sub classes of asymptotic independence: positive extremal association, near independence, and negative extremal association. These three classes correspond respectively to joint extremes of two variables occurring more often than, approximately as often as, or less often than joint extremes if all components of the variable were independent. Variables that have a multivariate normal dependence structure with correlation function, ρ, greater than zero are examples of asymptotic independence with positive extremal association.

Implications

It is important to know which class of extremal dependence a pair of variables falls into. For example, if river flows on neighbouring catchments are asymptotically dependent then there is a chance that they will both experience severe flooding at the same time and so this possibility should be taken into account in modelling the flood risk. If two rivers are asymptotically independent, but have positive extremal association, then the risk of joint flooding can be considerably greater than if the river flows at the different sites were independent, but severe flooding at both sites in the same flood event is less likely (particularly for long period return level events) than if the variables are asymptotically dependent. How close asymptotically independent variables are to being independent and to asymptotically dependent is important to quantify as any form of positive extremal association can significantly raise the receptor risk relative to an approximation of that risk when assessed under an independence assumption.

4.6.2 Pairwise dependence measures

The first collection of methods that are based on multivariate extreme value theory are those that describe asymptotic dependence and asymptotic independence but are only defined in a bivariate context. The simplest of these methods are the dependence measures χ and $\bar{\chi}$, these are both described fully by Coles *et al.* (1999) although the measure χ had been used previously (Buishand, 1984). These measures can be described as follows. Let (X, Y) be a bivariate pair of random variables, not necessarily identically distributed. We can transform (X, Y) to uniform margins using the probability integral transform. Let $U = F_X(X)$ and $V = F_Y(Y)$ denote the variables after transformation to uniform margins and let them have copula function C. The dependence measure χ is defined as:

$$\chi = \lim_{u \to 1} \Pr(V > u | U > v). \quad (4.6.1)$$

So, χ is equal to the limiting probability that the variable X is above a high probability threshold conditional on the variable Y being above the same probability threshold. It is more convenient to obtain χ using the following, asymptotically equivalent, function:

$$\begin{aligned} \Pr(V > u | U > u) &= \frac{1 - 2u + C(u, u)}{1 - u} \\ &\sim 2 - \frac{\log C(u, u)}{\log u} \quad \text{as} \quad u \to 1. \end{aligned}$$

We define

$$\chi(u) = 2 - \frac{\log C(u, u)}{\log u} \quad \text{for} \quad u \in [0, 1],$$

and it follows that $\chi = \lim_{u \to 1} \chi(u)$. So, if $\chi = c$, $c \in (0, 1]$ then (X, Y) are asymptotically dependent and if $\chi = 0$ then (X, Y) are asymptotically independent. Also if (X, Y) are independent then $\chi(u) = 0$ for all u.

Figure 4-4 contains plots of estimates of $\chi(u)$ for the bivariate data shown in Figure 4-2 and Figure 4-1. The estimates of the copula function C are obtained using the empirical joint distribution of the data in Figure 4-2. Estimates of $\chi(u)$ for each of the four data sets show quite different behaviour:

- For the multivariate normal data the value $\chi(u)$ decreases to zero as $u \to 1$ (the maximum possible value), suggesting that $\chi = 0$ but as $\chi(u) > 0$ for all $u < 1$ then the variables are asymptotically independent but exhibit positive extremal association.
- For the bivariate extreme value data as $u \to 1$ then $\chi(u) \to 0.5$ suggesting that $\chi = 0.5 > 0$ so the variables are asymptotically dependent.
- For the complete dependence data $\chi(u) = 1$ for all u so $\chi = 1$ and the variables are asymptotically dependent.
- Finally for completely independent data $\chi(u) \approx 0$ for all u and so $\chi = 0$ and the variables are asymptotically independent.

In all cases the sampling variability increases as $u \to 1$ making it difficult to estimate χ precisely from this approach. As $\chi = 0$ for all pairs of asymptotically independent variables (for example the multivariate normal and the completely independent data) it is necessary to define a second dependence measure to provide information on the relative strength of dependence under asymptotic independence. We denote the joint survivor function of (X, Y) as $\bar{F}(x, y)$. The copula survivor function $\bar{\chi}$ is defined as follows:

$$\begin{aligned} \bar{F}(x, y) &= 1 - F_X(x) - F_Y(y) + F(x, y) \\ &= \bar{C}[F_X(x), F_Y(y)]. \end{aligned}$$

So $\bar{C}(u, v) = 1 - u - v - C(u, v)$. By analogy with $\chi(u)$, $\bar{\chi}(u)$ is defined (Coles *et al.*, 1999) as:

$$
\begin{aligned}
\bar{\chi}(u) &= \frac{2\log \Pr(U > u)}{\log \Pr(U > u, V > v)} - 1 \\
&= \frac{2\log(1-u)}{\log \bar{C}(u,u)} - 1 \quad \text{for} \quad u \in [0,1].
\end{aligned}
$$

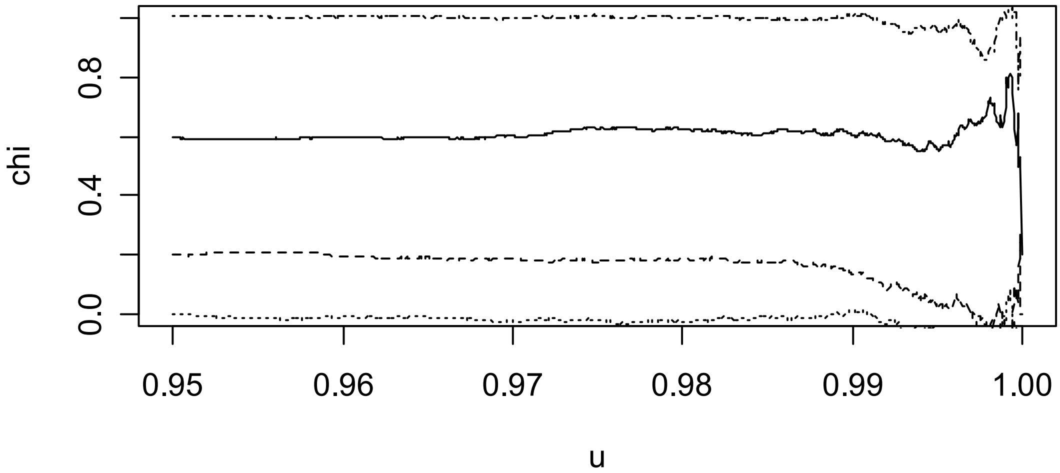

Figure 4-4: Estimates of χ for simulated data. Solid line bivariate extreme value distribution, α = 0.5, dashed line bivariate normal distribution, ρ = 0.5, dotted complete independence, dot-dash perfect dependence.

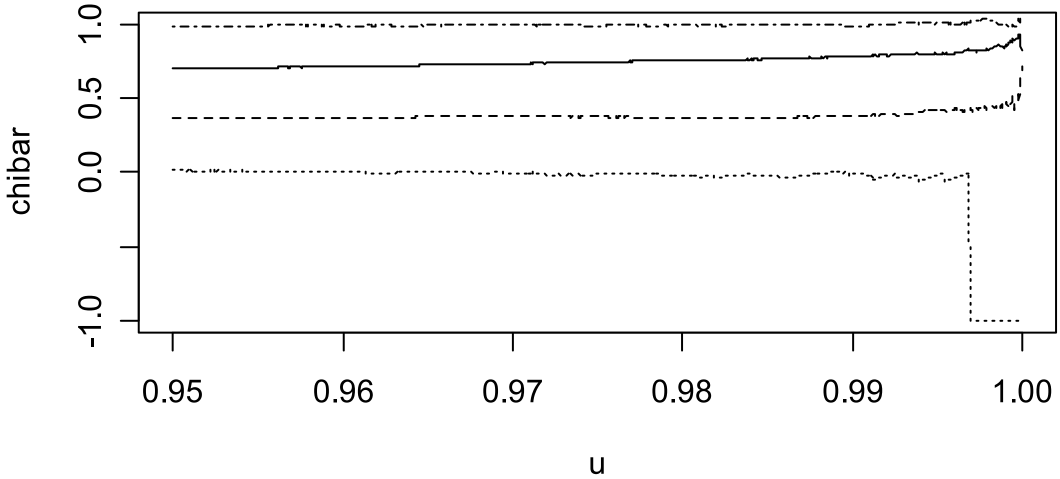

Figure 4-5: Estimates of χ bar for simulated data. Solid line bivariate extreme value distribution, α = 0.5, dashed line bivariate normal distribution, ρ = 0.5, dotted complete independence, dot-dash perfect dependence.

We have that $\bar{\chi}(u) \in (-1, 1]$ and define $\bar{\chi}$ as $\lim_{u \to 1} \bar{\chi}(u)$. If (X, Y) are asymptotically dependent then $\bar{\chi} = 1$, if (X, Y) are asymptotically independent then $\bar{\chi} < 1$. So we can see that both χ and $\bar{\chi}$ are needed to describe the extremal dependence of bivariate random variables.

Figure 4-5 contains plots of $\bar{\chi}$ for the bivariate data shown in Figure 4-2 and Figure 4-1. Again the estimates of $\bar{\chi}$ show different behaviour for the different dependence assumptions:

- For the completely dependent data set the value of $\bar{\chi} = 1$.
- For the bivariate extreme value data the estimate of $\bar{\chi}(u)$ suggests that $\bar{\chi} = 1$ asymptotic dependence, as $\bar{\chi}(u)$ increases to 1 as $u \to 1$.
- For the completely independent data the $\bar{\chi} = 0$ suggesting near independence in the asymptotic independence class.
- For the multivariate normal data then $\bar{\chi}(u) \to 0.4$ as $u \to 1$, so the variables are asymptotically independent but with positive extremal association.

Again the sampling variability of the estimates is poor as $u \to 1$. The sudden jump to a value of -1 for the complete independence case is due to the fact that above a certain threshold the estimated value of $\bar{C}(u, u)$ is equal to zero.

The dependence measure χ is used in Svensson and Jones (2002, 2004), however in only using χ without also estimating $\bar{\chi}$ it is not possible to use their results to completely classify the extremal dependence of the flood risk variables around the coast of Britain. Estimating χ when the variables are asymptotically independent frequently leads to the false conclusion that $\chi > 0$ (that is, that the variables are asymptotically dependent) as the alternative used is that the variables are completely independent (instead of the correct alternative that the variables are asymptotically independent).

A theoretical foundation for both χ and $\bar{\chi}$, and other pairwise dependence measures, and an improved framework for their estimation is given by Ledford and Tawn (1996). This theory can be stated as follows. Let (S, T) be bivariate random variables with unit Fréchet marginal distributions so that $\Pr(S > s) = \Pr(T > s) \approx s^{-1}$ for large s. It is possible to assume unit Fréchet margins without loss of generality as it is possible to change the marginal distributions of any random variables using the probability integral transform (see Section 4.3.1).

Also let $\bar{F}(s, s) = \Pr(S > s, T > s)$ be the joint survivor function of (S, T). Under certain modelling assumptions about the distribution of $\bar{F}(s, s)$ as s gets large then following Ledford and Tawn (1996), we may model $\bar{F}(s, s)$ as:

$$\bar{F}(s, s) = \mathcal{L}(s) s^{-1/\eta} \tag{4.6.2}$$

where $\eta \in (0, 1]$ is called the coefficient of tail dependence and $\mathcal{L}(\bullet)$ is a function that can be treated as if it is approximately constant as s gets large (a slowly varying function).

From asymptotic property (4.6.2) we have that:

$$\Pr(T > s | S > s) \approx \mathcal{L}(s) s^{1-1/\eta} \tag{4.6.3}$$

Using expression (4.6.3), we can classify the dependence pattern of (S,T) as $s \to \infty$ as follows

$$\begin{array}{lll} \eta = 1 & \mathcal{L} \not\to 0 & \text{asymptotic dependence} \\ \eta = 1 & \mathcal{L} \to 0 & \text{asymptotic independence} \\ \eta < 1 & & \text{asymptotic independence.} \end{array}$$

It is also possible to identify subclasses of asymptotic independence:

$$\begin{array}{rl} 1 > \eta > 1/2 & \text{positive association} \\ \eta = 1/2 & \text{near independence} \\ \eta < 1/2 & \text{negative association.} \end{array}$$

Ancona-Naverrete and Tawn (2002) show how it is possible to use η to show how dependence changes with distance. They applied this method to rainfall and sea surges and we discuss these findings in Section 3.3.4.

It is easy to see links between χ and $\bar{\chi}$ and the measures η and $\mathcal{L}$. If $\eta = 1$ and $\mathcal{L} \not\to 0$ then $\bar{\chi} = 1$ and asymptotic dependence between the variables is measured by $\chi = \lim_{s\to\infty} \mathcal{L}(s) > 0$. If $\eta < 1$ then $\chi = 0$ and the asymptotic independence between the variables is measured by $\bar{\chi} = 2\eta - 1$.

Ledford and Tawn (1996) show how η and $\mathcal{L}$ can be estimated using data from the joint tail in an efficient way. As a consequence of the connections between η and $\mathcal{L}$ with χ and $\bar{\chi}$ that are identified above these methods can be used to estimate χ and $\bar{\chi}$.

4.6.3 Models based on an asymptotic dependence assumption

A large proportion of the previous work of multivariate extreme value theory is based upon the assumption that the data exhibit asymptotic dependence. This is equivalent to assuming that the probability of the largest values of each variable occurring simultaneously is greater than zero. Over short distances this assumption may be appropriate for rainfall and sea surge (Svensson and Jones, 2002, 2004), however over longer distances and for river flows this assumption is inappropriate. We have discussed in Section 3.3.1 the findings of Buishand *et al*. (2008), who used the theory of continuous stochastic processes to describe the spatial extremes of rainfall under the assumption of asymptotic dependence.

4.6.4 Models based on an asymptotic independence assumption

Two variables are asymptotically independent if the probability of the largest observations on each variable occurring simultaneously is zero. The multivariate Gaussian tail model of Bortot *et al*. (2000) is a dependence model that makes the assumption that all variables are from a multivariate normal copula in the joint tail region. This corresponds to assuming that the variables are asymptotically independent of each other. This is the methodology on which the JOINSEA software was based. However, this is a highly restrictive assumption that limits the form of dependence, for example the dependence is symmetrical in variables and computational problems restrict application to 3-4 dimensional problems. For JOINSEA this model is only available in two variables (with dependence on a third variable produced through a different method).

JOINSEA also has a more sophisticated dependence model, termed the 'mixture of normals' model. This model comes from using a copula corresponding to a mixture of two biviarate normal distributions.The mixture of normals model allows changing association with the level of the variables (unlike the Gaussian tail model) but is restricted to being asymptotically independent.

4.7 Conditional model

4.7.1 Introduction

Heffernan and Tawn (2004) present asymptotic theory which provides the basis for a model for the behaviour of the conditional distribution of $\boldsymbol{Y}|X$, when X is large, under the assumption that observations of $\boldsymbol{Y}$ and X are independent, identically distributed, and without missing values. From this theory they developed a method for multivariate extreme values that can be used for both asymptotically dependent and asymptotically independent data. It is a conditional approach that uses all the observations such that the conditioning variable is above a certain threshold. This approach separates the marginal and dependence characteristics of the data and models them separately. The dependence characteristics of the data are accounted for by use of a specialised regression model. This model has the property that its residuals are independent of the size of the conditioning variable so that it can be used to extrapolate beyond the range of the data.

The method is based on a model of the distribution of $\boldsymbol{Y}|X$ when X is large. Here $(X, \boldsymbol{Y})$ is a vector variable, X, of dimension one and $\boldsymbol{Y}$ of dimension d, with known, identical, Gumbel margins. So X is a single variable and $\boldsymbol{Y}$ is a set of variables.

4.7.2 General method

The method involves the relatively weak assumption (compared with distributional assumptions made in some other extreme value methods) that there exist vector-valued normalising functions, $\boldsymbol{a}(x)$ and $\boldsymbol{b}(x)$, such that

$$\Pr\left(\frac{\boldsymbol{Y}-\boldsymbol{a}(x)}{\boldsymbol{b}(x)}|X=x\right) \to G(z) \quad \text{as} \quad x \to \infty \tag{4.7.1}$$

where the j^{th} marginal distribution G_j of G is a non-degenerate distribution function for all $j \in \Delta$ where Δ is the set $\{1, \ldots, d\}$. Here and throughout the vector algebra is to be interpreted as componentwise. To ensure that G is well-defined the following additional condition is required:

$$\lim_{z_j \to \infty} G_j(z_j) = 1 \quad \text{for all} \quad j,$$

so there is no mass at $+\infty$ in any margin.

The model is based on the approximation that limiting relationship (4.7.1) holds exactly for all values of $X > v_p$ for a suitably high threshold v_p which has probability p of being exceeded. A consequence of this assumption is that when $X = x$, with $x > v_p$, the random variable $\boldsymbol{Z}$, defined by

$$\boldsymbol{Z} = \frac{\boldsymbol{Y}-\boldsymbol{a}(x)}{\boldsymbol{b}(x)}$$

is independent of X and has distribution function G. It is this assumption of independence that allows us to extrapolate the model beyond the range of the data.

The assumption of the existence of normalising functions can be expressed simply as the assumption that the distribution function of the normalised variables, $G(\boldsymbol{z})$, does not tend to a constant as the conditioning variable X tends to infinity. This is equivalent to the assumptions made in using the GPD distribution to model the marginal distributions of the data. The assumption of independence of the random variables $\boldsymbol{Z}$ and the conditioning variable X is much more important. If this assumption is invalid then any results inferred above the range of the observations of X will be invalid. Therefore a key step in the method is the assessment of the validity of this assumption. This simply requires a test of independence between $\boldsymbol{Z}$ and X when $X > v_p$.

The theory suggests that there should always be some level v_p above which independence of X and $\boldsymbol{Z}$ is an appropriate assumption. Therefore v_p needs to be selected large enough to achieve this independence.

4.7.3 Choice of normalising functions

In Heffernan and Tawn (2004) the normalising functions $\boldsymbol{a}(x)$ and $\boldsymbol{b}(x)$ were derived for a number of different distributions. They found that the functions were all special cases of the parametric family

$$\boldsymbol{a}(x) = \boldsymbol{a}x + I_{\{\boldsymbol{a}=\boldsymbol{0},\, \boldsymbol{b}<\boldsymbol{0}\}}(\boldsymbol{c} - \boldsymbol{d}\log(x)),$$
$$\boldsymbol{b}(x) = x^{\boldsymbol{b}}$$

where, on the right hand side, $\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{c}$ and $\boldsymbol{d}$ are vector constants and I is an indicator function. The vectors of constants have components such that $0 \le a_j \le 1$, $-\infty < b_j < 1$, $-\infty < c_j < \infty$ and $0 \le d_j \le 1$ for all $j \in \Delta$.

The dependence class into which pairs of variables Y_j, X fall can be determined by the vector parameters in the following way: if $a_j = 1$ and $b_j = 0$ then Y_j and X are asymptotically dependent, otherwise the variables are asymptotically independent. If at least one of $0 < a_j < 1$ or $b_j > 0$ holds then the variables exhibit positive extremal dependence; if $a_j = d_j = 0$ and $b_j \le 0$ then the variables exhibit extremal near independence; and if $a_j = 0$, $d_j < 0$ and $b_j < 0$ then the variables exhibit negative extremal dependence. In practice, for most flood risk applications it is possible to assume that the variables exhibit either asymptotic dependence or asymptotic independence with positive extremal association. So we can set $\boldsymbol{c} = \boldsymbol{d} = \boldsymbol{0}$ and so the normalising functions are simply

$$\boldsymbol{a}(x) = \boldsymbol{a}x, \quad \boldsymbol{b}(x) = x^{\boldsymbol{b}}.$$

4.7.4 Bivariate features

For a pair of variables (X, Y) it is possible to think of the Heffernan and Tawn model as a regression model of Y upon X conditional on X being above a certain threshold. The constant parameters a and b describe the strength of dependence between X and Y. Two examples are given in Figure 4-5 showing samples of (X, Y) when X > 4.5. In the left hand plot both large and small values of Y can occur with large X, with large values of Y more likely than if X and Y were independent. In the right hand plot, only large values of Y occur with large values of X.

The parameter a describes the overall strength of dependence between the two variables, as a increases the overall strength of dependence between X and Y increases. The main purpose of parameter b is to describe how the dependence changes with threshold; for positive values of b the variance of $Y|X = x$ increases as x increases. Consequently large values of Y can occur with large values of X if either a is large or if b is large. However, small values of Y occur with large X only when b is large.

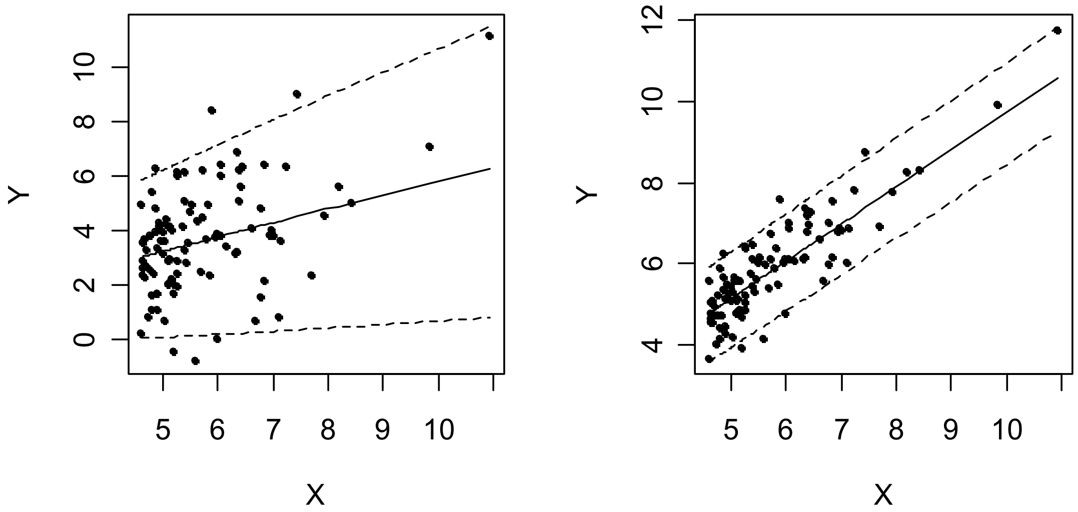

Figure 4-6: Plots of data simulated for Heffernan and Tawn model, solid lines median value of Y|X, dashed lines 2.5th and 97.5th percentiles of the distribution of Y|X. Left plot a = 0.3, b = 0.7, right plot a = 0.9, b = 0.1. In both Z follows a normal distribution with mean 0.5, variance 0.25.

4.7.5 Multivariate features

Figure 4-7 shows how the dependence model is structured for more than two variables. The dependence of X and each variable in $\boldsymbol{Y}$ is modelled parametrically using the Heffernan and Tawn model. So, in this case, we estimate the distribution of

$$Z_j = \frac{Y_j - a_j x}{x^{b_j}}, \quad \text{for all} \quad j = 1, \ldots, 4. \tag{4.7.2}$$

By modelling the dependencies of the individual variables $Y_j \in \boldsymbol{Y}$ with X we remove some of the dependence between the Y_j variables. This is because, if all variables in $\boldsymbol{Y}$ are positively associated with X, if X is big all variables in $\boldsymbol{Y}$ are also likely to be big. This tendency is captured by the a_j and b_j parameter constants. However, not all the dependence between the Y_j variables is captured in the a_j and b_j parameters. This additional dependence is captured by the Z_j variables, which are modelled as being correlated with each other. The dependencies between the Z_j parameters are modelled non-parametrically.

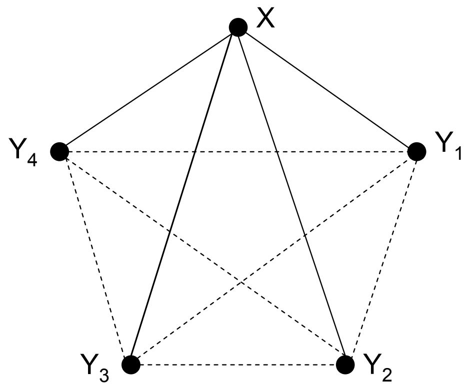

Figure 4-7: Diagram of modelled dependencies, solid lines indicate dependencies modelled parametrically (through a and b), dashed lines indicate dependences modelled non-parametrically (through the random variable Z).

4.7.6 Inference

The method of fitting the Heffernan and Tawn model to a set of variables $(X, \boldsymbol{Y})$ is as follows.

Stage 1 Obtain Gumbel margins.

The data are transformed onto Gumbel margins by repeated use of the probability integral transform separately for each marginal variable.

Stage 2 Threshold selection.

Select threshold, v_p, to be high enough for the asymptotic results to hold and low enough to retain as much data as possible. One way to select the threshold is to use the lowest threshold so that the estimates of the dependence parameters are unchanged for higher thresholds.

Stage 3 Estimate parameters and residuals.

To achieve this estimation we assume that each residual random variable Z_j, defined in equation (4.7.2), has mean μ_j and standard deviation σ_j. From this assumption, and that of expression (4.7.1), it follows that the variables $Y_j|X = x,\ x > v_p$, have mean

$$\mu_j(y_j) = a_j y_j + \mu_j x^{b_j}$$

and standard deviation

$$\sigma_j(y_j) = \sigma_j x^{b_j}.$$

We then make the false assumption that Z_j are independent random variables from a normal distribution. This allows us to use the technique of maximum likelihood to estimate the parameters. This is equivalent to minimising the following expression with respect to μ_j, σ_j, a_j and b_j.

$$\sum_{t=1}^{T_j}\left[\log\{\sigma_j x^{b_j}\} + \frac{1}{2}\left\{\frac{y_j - a_j x - \mu_j x_j^b}{\sigma_j x^{b_j}}\right\}^2\right], \qquad (4.7.3)$$

where T_j is the number of observations of y_j when $x > v_p$. The estimated residuals at time t are defined as

$$\hat{Z}_{j,t} = \frac{y_{j,t} - \hat{a}_j x_t}{x_t^{\hat{b}_j}},$$

where $\hat{a}_j$ and $\hat{b}_j$ are the values of a_j and b_j that minimise function (4.7.3). By defining the residuals in this way we do not need to take μ_j and σ_j into account in the rest of the analysis.

4.7.7 Estimation of summary measures

Because the Heffernan and Tawn model covers the whole joint distribution of the variables $(X, \boldsymbol{Y})$ conditional on X being above the chosen threshold it is possible to use it to estimate a wide variety of summary measures. Two such summary measures which have been used in practice are:

$$P_j(p) = \Pr\left(Y_j > v_p | X > v_p\right),$$

$$N(p) = \frac{E\left(\#\{j \in \Delta : Y_j > v_p\} | X > v_p\right)}{\#\{j \in \Delta\}}.$$

Here $P_j(p)$ is equal to the probability that variable $Y_j, j \in \Delta$ is extreme and $N(p)$ is the expected proportion of variables in Δ that are extreme both conditionally given that X is extreme.

4.7.8 Missing data

One of the main limitations of the original Heffernan and Tawn method is that it does not deal with missing data. This has the implication that for estimating the joint distribution of a set of variables it is only capable of using data when all the variables are observed, all other data must be discarded. For flow data this often results in a data set that has too few observations to be of any use. To overcome this problem we use the extension to the Heffernan and Tawn method of Keef *et al*. (2009a). The principle behind this extension is that the observed data tell us something about the unobserved data.

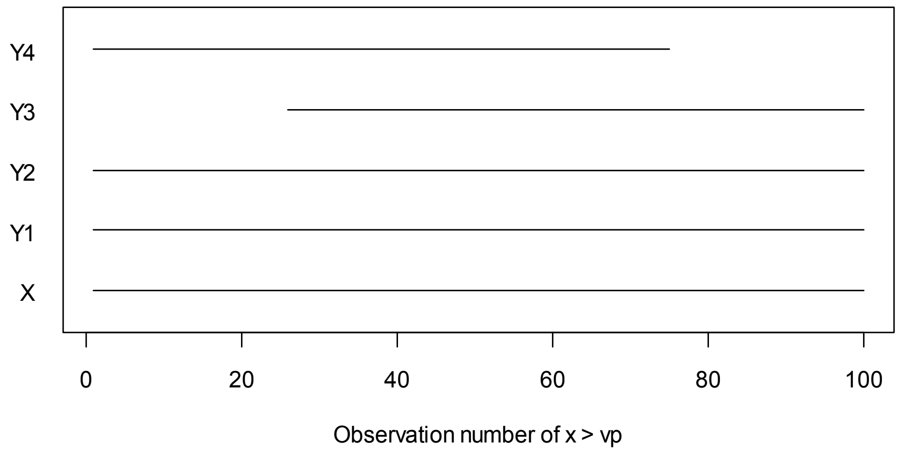

Figure 4-8: Illustration of missing data pattern

To illustrate the method we use the set of variables shown in Figure 4-7, where we have a conditioning variable X and four non-conditioning variables $Y_1, \dots, Y_4$. Figure 4-8 shows an illustration of a typical missing data pattern. The first 25 per cent of variable Y_3 is missing, along with the last 25 per cent of variable Y_4, and so if the original Heffernan and Tawn approach is used only the central 50 per cent of data can be used.

The main strategy behind the method of accounting for missing data is to estimate the distribution function of the missing observations of Y_3 conditional on Y_1, Y_2, Y_4 and X, when $X > v_p$ and similarly for Y_4 conditional on Y_1, Y_2, Y_3 and X, when $X > v_p$. These conditional distributions are estimated from the conditional distributions of the residual random variables, $Z_3|(Z_1, Z_2, Z_4)$ and $Z_4|(Z_1, Z_2, Z_3)$. This is done by assuming a multivariate normal copula for the Z_j variables. The main reason this copula is used is convenience, if a set of variables $\boldsymbol{Z}$ have a multivariate normal distribution then it is very simple to compute the conditional distribution of any subset Z_M, where $Z_M \in \boldsymbol{Z}$, of these variables, conditional on the other variables in $\boldsymbol{Z}$.

This treatment of missing data means that we can make the most use of the available data. There is less uncertainty than if we only used the data for days where all sites were observed, but more uncertainty than if we had complete records for all sites. For the case study example for North East Region reported in the accompanying proof of concept summary report, the range of record lengths was from 19 to 47 years, with a median of 9 years of non-overlapping data.

4.7.9 Temporal dependence

The method of handling temporal dependence is to use the Heffernan and Tawn model to model the joint distribution of $Y_{j,t+\tau}|X_t > v_p$ where $\tau \in A_j$, $A_j = \{L_j, \dots, L^j\}$ with L_j and L^j being the upper and lower limits that define a set of lags of interest between the variables X and Y_j for all $Y_j \in \boldsymbol{Y}$. It is then possible to extend the definitions of $P_j(p)$ and $N(p)$ to

$$P_j(p) = \Pr\left(\max_{\tau \in A_j} Y_{j,t+\tau} > v_p | X_t > v_p\right),$$

$$N(p) = \frac{E\left(\#\{j \in \Delta : \max_{\tau \in A_j} Y_{j,t+\tau} > v_p\}|X_t > v_p\right)}{\#\{j \in \Delta\}}. \quad (4.7.4)$$

The value $P_j(p)$ is the probability that Y_j has at least one threshold exceedance within a specified window of time around a threshold exceedance of X. The quantity $N(p)$ in equation (4.7.4) is the expected proportion of variables in Δ that have at least one threshold exceedance within a window of time around a threshold exceedance of X.

To simulate a single event we need to simulate at different locations, and at different time points. We therefore model the dependence between different variables at the same time point (X_t and $Y_{j,t}$), and at different time points (X_t and $Y_{j,t+\tau}$ where $\tau \in A_j$). Figure 4-9 illustrates this approach; the variables Y_j at lags $\tau \in A_j$ are simply treated as additional variables. We model the distribution of $Y_{j,t+\tau}|X_t > v_p$ for all $Y_j \in \boldsymbol{Y}$ and all $\tau \in A_j$ using the same method. Because we wish to simulate 'events' it is sensible to take A_j to be the same for all j.

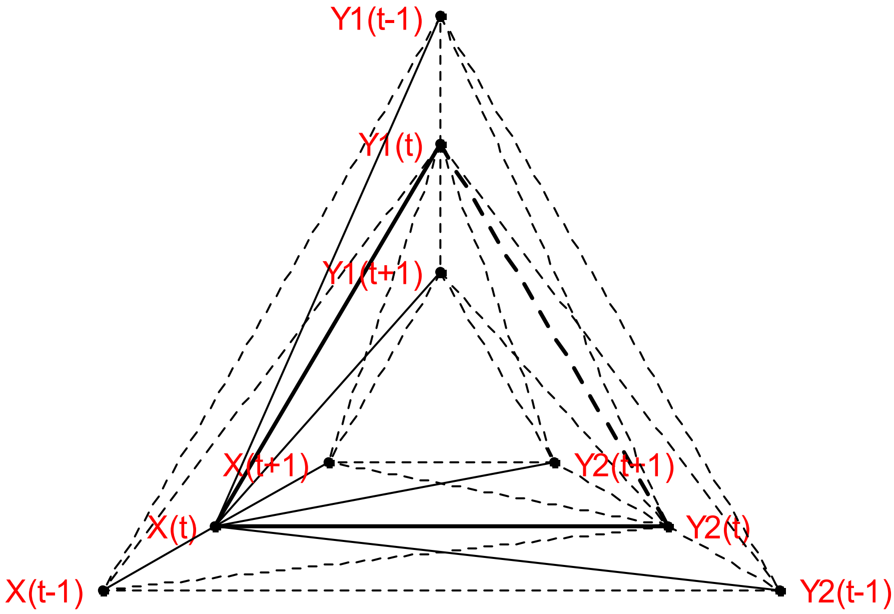

Figure 4-9: Extended version of Figure 4-7 to show modelled temporal dependencies. Solid lines show dependencies modelled parametrically (through a and b), dashed lines indicate dependencies modelled non-parametrically (through the random variable Z). Bold lines indicate dependencies modelled for no time lag, so directly equivalent to those in Figure 4-7.

It is also possible to define the univariate dependence measure, $P^{(\tau)}(p)$ where

$$P^{(\tau)}(p) = \Pr(X_{t+\tau} > v_p | X_t > v_p) \tag{4.7.5}$$

is the probability that X is big at time $t+\tau$ given that X is big at time t.

The main limitations of using the Heffernan and Tawn method are as follows. The first is the amount of data that is needed to produce reliable estimates, however, this limitation is common to all statistical extreme value methods. Typically periods of overlap in record length should be at least 20 years. In theory the Heffernan and Tawn model extends to any number of variables, and any number of time lags. However, due to computational issues in fitting the model (the process is memory intensive), in practice the number of variables/time lags it is possible to model is limited. These limits are in the region of around 50 variables/site with time lags of +/- 5 time steps ($L_j = L^j = 5$), or more variables with fewer time steps.

4.8 Summary of statistical methods to model dependence

Table 4-1 contains a summary of which techniques meet the original criteria set out earlier in this section. The method that looks most promising is the conditional model of Heffernan and Tawn. The only criterion that it does not meet is that of incorporating dependence covariates into the model. However, this is also true for all other method groups examined in this study.

	Method groups				
	Copula methods	Asymptotic dependence	Pairwise dependence measures	Gaussian tail model	Conditional model
Joint distribution function	✓	✓	-	✓	✓
High dimensionality	✓	✓	-	-	✓
Change in dependence at extreme levels	-	✓	✓	✓	✓
Different types of extremal dependence	-	-	✓	-	✓
Extremes and non-extremes	-	-	✓	✓	✓
Differing spatial and temporal scales	✓	✓	✓	✓	✓
Covariates	-	-	-	-	-

Table 4-1: Summary of statistical methods. Ticks indicate established capability to handle the requirements shown in the left most column of the table.

Based on ideas in Davison and Smith (1990), the addition of covariates into extreme value models has been accomplished in many different univariate applications. Due to the semi-parametric nature of the Heffernan and Tawn model, including covariates in this analysis would require additional work, however this would not preclude a method to include covariates from being developed.

5 Statistical model for sources of flooding – analysis of gauged data

In this section of the report we describe a statistical model for the joint distribution of river flows or sea levels that will be used to supply the source (hydraulic load) component for a source-pathway-receptor flood risk model. This section therefore describes in detail the proposed form of the probability distribution for the 'source' variables (the first box in the generic model illustrated in Figure 2-1 to Figure 2-7).

5.1 Introduction

In Section 4 we outlined the statistical requirements of a model to estimate the spatial dependence of flooding. We also identified a method that is capable of meeting these requirements. This is the conditional method of Heffernan and Tawn. This has two components: models for the marginal distribution of the variables, and a separate model for the dependence structure.

For any vector variable $\boldsymbol{X} = \{X_i, i \in \Delta\}$ where Δ is the set $\{1, \ldots, d+1\}$ with continuous marginal distribution functions F_i, for $i \in \Delta$ i.e. $F_i(x) = \Pr(X_i < x)$ the joint distribution can be written as

$$\Pr(X_1 \leq x_1, \ldots, X_{d+1} \leq x_{d+1}) = C\{F_1(x_1), \ldots, F_{d+1}(x_{d+1})\} \quad (5.1.1)$$

where C is a unique function, known as the copula, which determines the dependence structure of $\boldsymbol{X}$ (see Joe, 1997 and Nelsen, 1999). The copula formulation separates the joint distribution into the $d+1$ marginal distribution functions and a joint distribution function C for the variables on a common marginal distribution. In expression (5.1.1), C is the joint distribution function for uniform $[0,1]$ variables, however the choice of common marginal distribution does not matter and so in different studies different common marginal distribution forms are selected for the convenience of the problem.

To model the joint distribution of $\boldsymbol{X}$ we need models for $F_1, \ldots F_{d+1}$ and C. We present our marginal models in Section 5.2. These are standard models which describe the behaviour each variable alone.

To study the dependence structure through a copula we initially transform our data, for example river flow or sea levels, to a common marginal distribution. For the purposes of modelling extreme values it is best to use standard Gumbel margins as this scale induces the most linearity in the dependence (Heffernan and Tawn, 2004). Suppose we are most interested in the extremes at site $i = 1$, that is, the values of X_1. This is not restrictive as later we will consider each site in turn. We term X_1 after it has been transformed to follow a Gumbel distribution by X, and collectively term $\{X_2, \ldots, X_{d+1}\}$, after it has been transformed to follow Gumbel margins by $\boldsymbol{Y}$, where $\boldsymbol{Y}$ is a vector of dimension d. Details of how these transformations are made are given in Section 4.3.1.

The aim is to be able to estimate features of the distribution of the $\boldsymbol{Y}$ variables when the X variable is large. The strategy is first to model the distribution of X, which is given due to our choice of Gumbel marginal distribution. Then we model the conditional distribution of $\boldsymbol{Y}|X = x$ for large values of x. Here the conditional distribution is the distribution of all the elements of $\boldsymbol{Y}$ given that X is fixed equal to x.

We propose a model for $\boldsymbol{Y}|X = x$ motivated by asymptotic probabilistic theory for this conditional distribution as $x \to \infty$. We assume that this model will be appropriate for all values of x above a high threshold v_p, where v_p has probability p of being exceeded, and use the observations of $\boldsymbol{Y}$ with $X > v_p$ to fit the model. The key step in all our dependence modelling is the modelling of the conditional distribution of $\boldsymbol{Y}|X = x$ for large x as all subsequent aspects of the inference hinge on this step.

Having fitted the conditional distribution the next stage is the estimation of features of the distribution of $\boldsymbol{Y}$ when $X > v$ where $v \geq v_p$. To illustrate this consider the feature

$$\Pr(Y_j > v|X > v) \tag{5.1.2}$$

which describes the dependence between two variables, that is, the probability that a variable exceeds a high level provided that the other variable has exceeded the same high level. This corresponds to the probability that variable X_{j+1} exceeds its T year return level given that X_1 exceeds its T year return level, with T being determined by the value of v.

The simplest way to evaluate probability (5.1.2) numerically using the fitted model is to generate a large sample from the fitted distribution: first simulate $X > v$ giving a value x^* say, second simulate $\boldsymbol{Y}|X = x^*$ (for probability (5.1.2) only $Y_j|X = x^*$ is needed), repeat these first two steps to simulate replicate samples, and then estimate the feature of interest empirically from this simulated sample. For estimating probability (5.1.2) the Monte Carlo estimate is simply the proportion of simulated points with $Y_j > v$ out of the subset of the simulated sample with $X > v$. The sample size of the simulated sample is taken to be sufficiently large that the Monte Carlo estimate has very small uncertainty. The threshold v_p will necessarily be inside the range of the data sample to facilitate inference about the conditional distribution of $\boldsymbol{Y}|X = x$. However, the level v can be arbitrarily large and so the proposed strategy provides estimates of features about the distribution of $\boldsymbol{Y}$ within the observed tail of X through to extrapolation beyond the maximum X observation.

For more general functions of the simulated sample, which are required for converting a loading into a damage cost, we need to transform back from $(X, \boldsymbol{Y})$ to $\boldsymbol{X}$. This can simply be achieved by using the inverse of the transformation we used to convert the data into Gumbel margins (see Section 4.3.1 for details).

5.2 Marginal models

The approach that we use to model the marginal distributions of river flow data is to fit a parametric distribution (GPD) to the exceedances of a threshold, and to simply use the empirical distribution below the threshold.

To model the marginal distributions of the sea level data we use the version of the revised joint probability method (RJPM) of Tawn and Vassie (1989) that is being developed as part of the Environment Agency project SC060064 (Development and dissemination of information on coastal and estuary extremes). The RJPM models the deterministic tide component of sea level data separately from the stochastic surge component. The surge residuals are modelled using a distribution for threshold exceedances.

It is worth noting that because of the separation of the marginals from the dependence structure, other choices of marginal distribution could be used within the proposed model. For example, it would be possible to replace the marginal distributions we use with distributions taken from an FEH analysis. There is a chance that this could introduce some inconsistency between the marginal and dependence models, although it is unlikely to change the results of the analysis significantly. This is a point that would be suitable for investigation outside of this project but is not considered essential.

5.2.1 River flow marginal model

There are two reasons we fit a parametric distribution to the threshold exceedances rather than just use the models contained in the Flood Estimation Handbook (FEH). Both are related to the use of daily mean flow data.

The FEH includes distributions suitable for representing annual (block) maxima and threshold exceedances (peaks over threshold, or PoT, data). In general the annual maximum procedures are preferred by FEH users as they are simpler and considered good enough for most flood estimation studies in practice. The distribution function of the exceedances of a threshold is different to the distribution function of block maxima. Our model is fitted to exceedances on the daily mean flow scale, not to annual maxima (which cannot capture the dependence on an event-by-event basis).

The second reason is that the data we have are daily mean flows, rather than instantaneous peak flow rates. It is likely that the distribution function of daily mean flow will be slightly different to the distribution function of (for instance) daily maximum flow, especially for small or flashy catchments. The daily maximum must always be at least as big as the daily mean flow.

5.2.2 Still water sea level marginal model

For coastal flooding we require a method to estimate the probability of flooding from extreme still sea water levels and to derive the probability of joint occurrence of extreme still water levels at different coastal sites. Still water sea levels are made up of two components, tide and surge, with the predicted astronomical tidal component being deterministic and the surge component being a random process.

For evaluating the distribution of extreme still water levels at a site it is best to separate the still water level into its tide and surge components, analyse each of these individually, and recombine them to give the distribution of still water level. This approach allows for the extrapolation of still water levels to be influenced by the extrapolation of surges, all the possible combinations of tide and surge, and the nodal cycle structure of the tidal series. Similarly, the analysis of still water levels at more than one site is best achieved by studying separately the dependence between surges at the sites and the dependence between tide and surge at the sites. Precise details of how the different aspects of these analyses will be developed in this project are discussed below.

Skew surge

The surge residual is equal to the actual, observed, still water sea level, minus the predicted astronomical tidal level at a particular point in time, as illustrated in Note:
TSL – total sea level
AST – astronomical tide

Figure 5-1. The surge residual has traditionally been used in the statistical analysis of extreme sea levels. This however leads to a complication in that the complex tide-surge interaction needs to be accounted for. Whilst tides and surges are driven by completely different physical processes, tidal currents and surge currents interact in shallow water, affecting the timing of each other (Horsburgh and Wilson, 2007). The result of this interaction is that the maximum surge residuals tend to occur at low to mid tidal levels on the rising tide, see Note:
TSL – total sea level
AST – astronomical tide

Figure 5-1. Consequently, if the full probability distribution of surge residuals is used in a joint probability analysis, then careful statistical modelling is required to ensure that large surge residual values recorded at low to mid tidal levels are not coupled with high tidal levels, as to do so would lead to an overestimation of the probability of very extreme events.

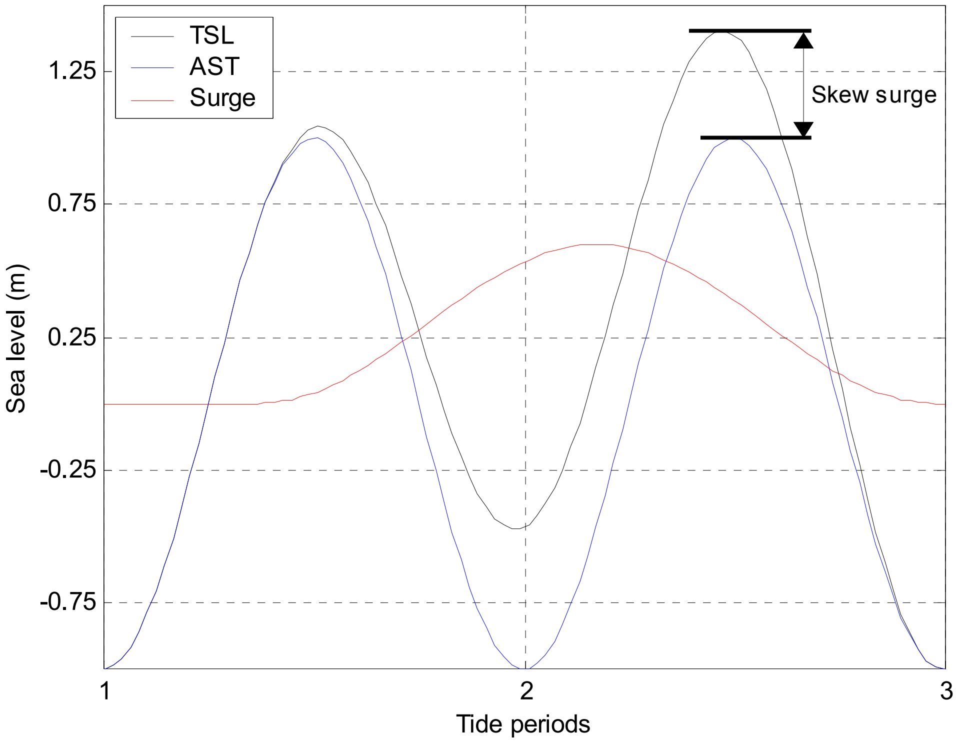

Note:
TSL – total sea level
AST – astronomical tide

Figure 5-1: Definition of skew surge.

What is important in terms of flooding is how much a storm event raises the sea level above the predicted tidal level and whether this would lead to flooding. If a very large surge residual is recorded at low tide, it is irrelevant given that flooding would be unlikely.

Increasingly, there has been a move to use the skew surge in preference to surge residual, particularly for operational forecasts (for example, see Verlaan *et al.*, 2005). This term refers to the difference between the maximum recorded sea level during a tidal cycle and the predicted maximum tidal level for that cycle, irrespective of their timing, as illustrated on Note:
TSL – total sea level
AST – astronomical tide

Figure 5-1. Modelling the distribution of skew surge, rather than surge residual, has the advantage that we can avoid having to account for the complex issue of tide-surge interaction. Skew surge also has the advantage that its magnitude is largely independent of the associated tidal height. This is indicated by the scatter plots in Figure 5-2, which plot skew surge magnitude against tidal height. There is a clear lack of relationship shown in these plots which suggests that it is reasonable to model the skew surge as being independent of the high tide level in the associated tidal cycle.

In the Environment Agency project SC060064 (Development and dissemination of information on coastal and estuary extremes) the use of statistical models for the distribution of skew surge is being compared with the modelling of the more complex surge residuals for the evaluation of return levels for still water level. Preliminary results suggest that the usage of skew surge gives similar results to the established surge residual methods but requires much less statistical modelling and expertise, and hence results are more robust to statistical assumptions.

Therefore in this study we will focus on using the skew surge. It is of practical relevance for flooding, less complicated than the surge residual to use, and it captures better the element of a storm surge that is most important in terms of flooding.

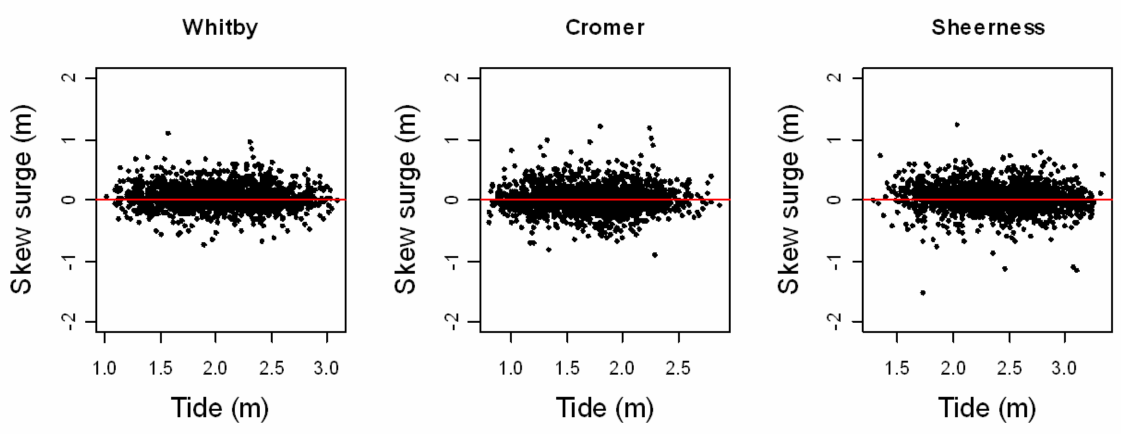

Figure 5-2: Plots of skew surge against tide for three sites on the east coast of England.

5.3 Event definitions – fluvial

An event is defined here as a period of time in which floods at different sites can be classed as ‘simultaneous’ (that is, belonging to the same event). This period of time could either be defined hydrologically or by the user. Examples of users who might have differing definitions of which events would be classified as ‘simultaneous’ are transport planners and the reinsurance industry. For planning diversions around flooded areas it may only be important to know if two possible transport routes are flooded on the same day (or even shorter time spans). In the reinsurance industry there are fixed periods of time in which flood events are classified as the same event. Any events that last longer than this period of time are classified as separate events for the purpose of financial analysis.

The hydrological definition of an ‘event’ may seem obvious but is also open to choice. To define events in the hydrological sense we first need to analyse the data to define a time interval at which events at the same or different sites can be classed as independent. A fixed time period could easily be chosen if a user wished to concentrate on a particular type of event, for example flooding over the timescale of one or two months, as in Summer 2007.

In this section we illustrate methods that can be used to assess duration and travel time. When conducting a study into spatial dependence for a given region it will be necessary to perform an analysis to examine these temporal aspects of flooding in order to be able to assess the true level of spatial dependence in that region.

5.3.1 Duration

The flows or water levels taken to be a single 'event' in our model may in fact occur at different points in time because the physical event may have a different duration in different places. Temporal dependence in river flows is largely a function of the catchment characteristics of the river. In particular, the permeability and size of the catchment have large effects Figure 5-3 shows empirical estimates of $P(X_{t+\tau} > v_p | X_t > v_p)$ for v_p equal to the 0.99 probability threshold of gauging station X against BFI and catchment area for gauging station$X$. The relationship between $P(X_{t+\tau} > v_p | X_t > v_p)$ and BFI is clear, the relationship between $P(X_{t+\tau} > v_p | X_t > v_p)$ and catchment area is less clear, however for the largest catchment the temporal dependence in high flows does appear to be slightly higher.

Figure 5-4 shows plots of the correlation between X_t and $X_{t+\tau}$ using Spearman's ρ (a standard non-parametric dependence measure) and $P(X_{t+\tau} > v_p | X_t > v_p)$ for three gauging stations X_t. The three stations chosen were 39020 which has a BFI value of 0.94 (high) and area of 106.7 km^2; 27025, which has a BFI value of 0.53 (medium) and area of 352.2 km^2; and 72002, which has a BFI value of 0.32 (low) and area of 275.0 km^2. We chose these three stations as they had neither very large nor very small catchments and they had very different BFI values.

We can clearly see that the dependence in the extremes is much lower than the dependence in the main body of the data, and falls to zero much more rapidly. The relationship between temporal dependence and BFI is also clear. We can also see that for the medium and low BFI stations the dependence in high flows drops to near zero (independence) for very small lags. However, for the station with high BFI the temporal dependence is much higher, and drops to zero at a much slower rate. This suggests that for some sites it will be difficult to define a short window of time at which threshold exceedances can be classed as belonging to different events. However, for most sites a time window of around plus or minus five days would be sufficient.

The sequencing of flood events is also of importance, especially if defence system or economic damage calculations need to differentiate between a single event and multiple events occurring in succession (for example to avoid double counting damages). If a given catchment experiences a number of floods within a short time of each other, then other catchments that have some dependence are also more likely to experience floods around the same times. To a certain extent this will be dealt with by that fact that we use the flow data directly to estimate the spatial dependence. Occurrences of sequences of events will be present within the data, and so will be included in the modelling process. However, this feature has not been examined in further detail for this project.

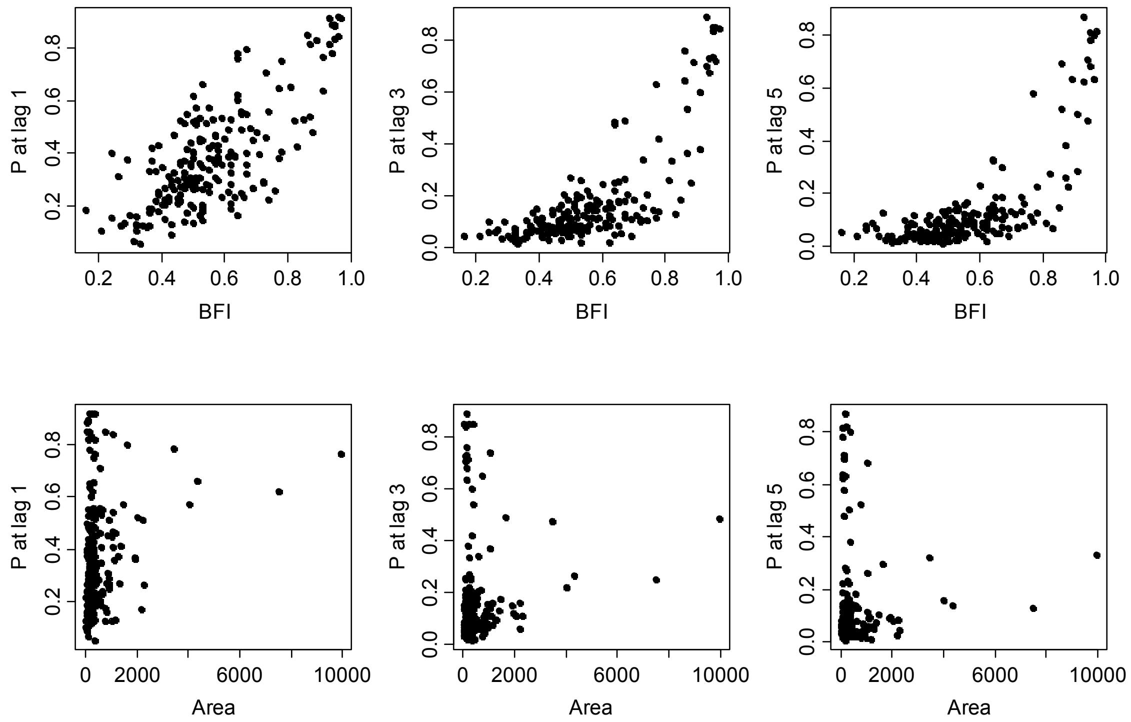

Figure 5-3: Top row $P(X_{t+\tau} > v_p | X_t > v_p)$ against BFI, bottom row against catchment area. Left $\tau = 1$, middle $\tau = 3$, right $\tau = 5$ days.

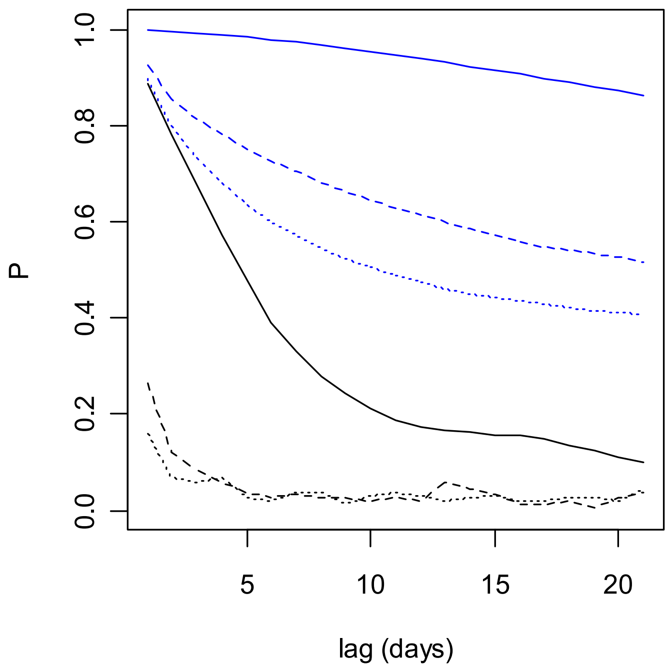

Figure 5-4: $P(X_{t+\tau} > v_p | X_t > v_p)$, and acf shown in black and blue respectively. Solid lines high BFI, dashed lines medium BFI, dotted lines low BFI. Threshold v_p equal to 99% on the daily scale.

5.3.2 Travel time

In addition to the effect of duration, an 'event' may also occur physically at different times in different places because of the pure time lag as the flood travels downstream. The definition of travel time that we use is the lag at which dependence between two stations on the same river peaks. It is important to note that this is not the same as the physical travel time for a single flood wave to travel down a river. We use the lag at which dependence peaks because it is more relevant to modelling the joint probability distribution and implicitly captures influences of catchment characteristics and rainfall patterns.

The data that we use are daily data, so we are unlikely to see any travel time effects for shorter rivers. The river that we use to illustrate travel time is the Severn. The reasons we chose the Severn are that it is well gauged, is one of the longest rivers in the country and that it has a relatively homogeneous catchment. The Thames is longer and is also well gauged but is has a very diverse catchment. The three gauging stations on the Severn that we use are 54012 (most upstream of the three), 54095 (located in the middle of the three), and 54001 (most downstream of the three), these stations are illustrated in the Appendix in Figure A-3.

Figure 5-5 shows plots of the lagged correlation between 54012 and 54095; 54012 and 54001; and 54095 and 54001. The correlation measures that we use are Spearman's ρ correlation measure and $P^{(\tau)}(p) = P(Y_{t+\tau} > v_p | X_t > v_p)$ where v_p is the 0.99 probability threshold for X and Y. The first thing to note is that temporal dependence in the extremes of the flow data is much lower than that in the main body of the data. We can also see that the dependence between 54012 and the other stations appears to peak when 54095 and 54001 are at a lag of plus one day to station 54012. This suggests that there is a travel time effect between 54012 and the other stations. Additionally, the dependence between 54095 and 54001 is almost the same for no lag and a lag of plus one day. One way to interpret this is that there is a travel time effect between 54095 and 54001, but that this lag is between 12 and 24 hours and so cannot be fully identified from daily data.

From this analysis we can see that although there does seem to be a travel time effect on dependence, for these stations at least the dependence falls off within plus or minus three to five days. We adopt a time window of plus or minus three days in the analysis for this project (this is in addition to working with the five-day maxima to account for differing event durations, as described in the previous section).

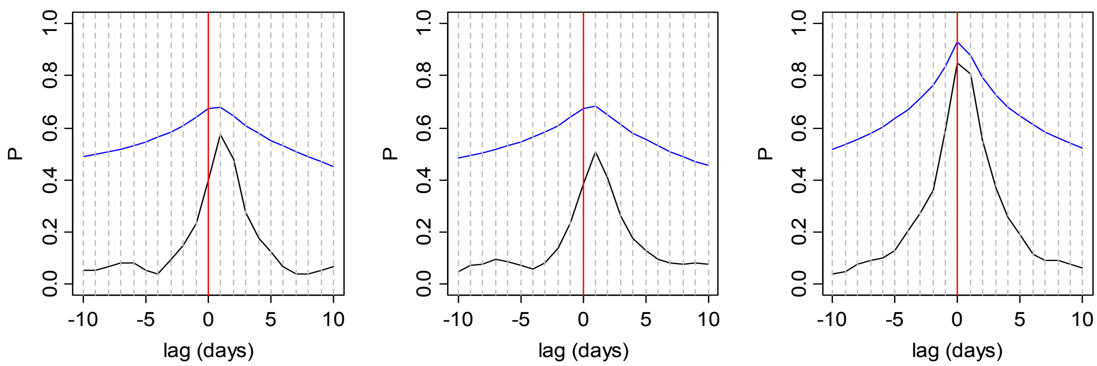

Figure 5-5: Left plot X = 54012, Y = 54095, middle plot X = 54012, Y = 54001, right plot X = 54095, Y = 54001. Blue lines Spearman's ρ, black lines $P(Y_{t+\tau} > v_p | X_t > v_p)$.

5.4 Event definitions – coastal

The issues surrounding event definition for coastal events are the same as those for fluvial events. Surges last for a period of time, and, at least on the east coast, they reach different sites at different times as surge waves travel down the coast. The methods used to take this temporal dependence into account when simulating events are the same as those used for fluvial events. So we need to be able to define a window of time within which extreme surges can be classed as dependent.

Figure 5-6 shows plots of the highest ten surge events at North Shields. All events are clearly defined and the duration of each is around 30 hours. Although not all coastal sites have surge events that are as well defined as here, the duration of surge events are generally shorter than the duration of most fluvial events. When we repeated the analysis shown in Sections 5.3.1 and 5.3.2 for sea level data we found that a time window of plus or minus one day was sufficient to capture full events.

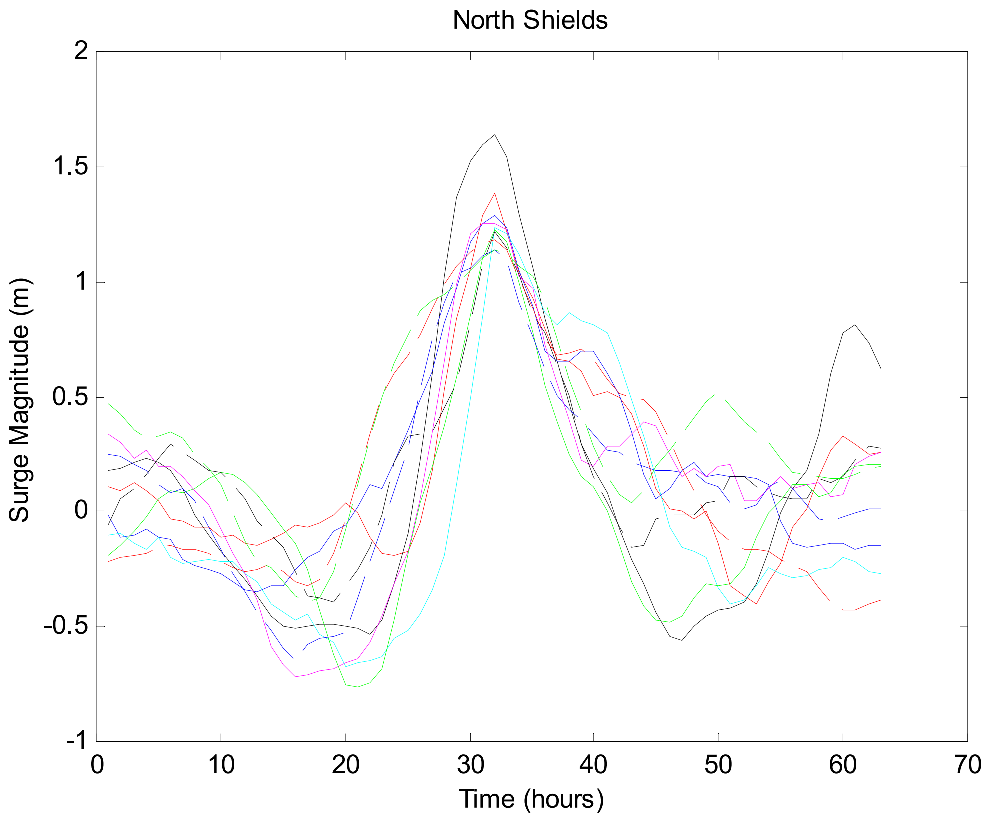

Figure 5-6: Top ten surges for North Shields plotted so that their peak value occurs concurrently.

5.5 Seasonality

A useful way to investigate seasonality is to examine the number of high flow days observed in each month. We define high flow days as days on which the daily mean flow exceeds the 99 per cent probability threshold for the whole year. If there were no seasonal effect we would expect the proportion to be constant for each month, and equal to 0.01.

Figure 5-7 shows plots of the proportion of high flow days for selected months. The gauging stations used in this analysis are a subset of the whole data set that have long records. The reason we carried out the analysis on this subset was to ensure that there was enough data at each site to obtain accurate results. These plots show clearly two aspects of flooding that are well known: first that there are more high flow days in the winter months; and second that the actual number of high flow days in a particular month varies over the country. In general the south east of the country experiences more high flow days later in the water year (running Autumn to Autumn) than the north and west.

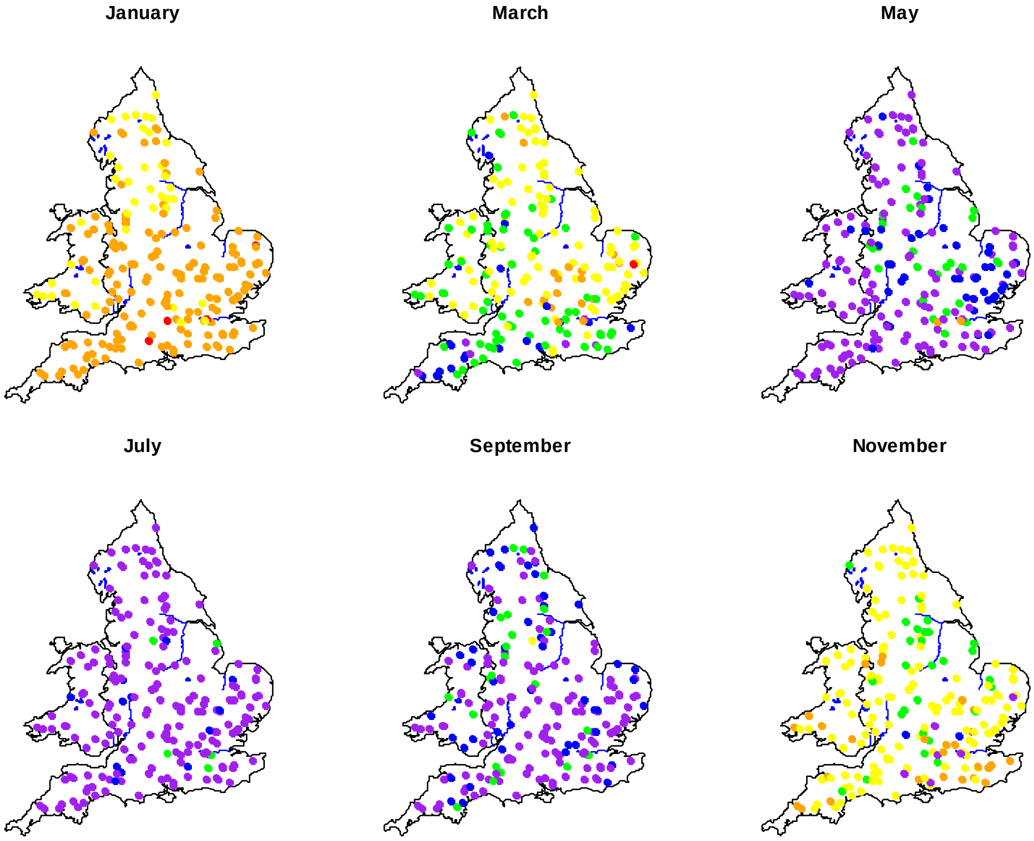

Figure 5-7: Proportion of days each month exceeding 99 per cent probability threshold for selected months for each gauging station. Colouring of sites gives the value of the observed proportion in relation to what that proportion would be if there were no seasonality. Purple: more than four times less likely than no seasonality; blue: between two and four times less likely; green: between two and one times less likely; yellow: between one and two times more likely; orange: between two and four times more likely; red: more then four times more likely.

The second stage of the analysis of seasonality is to see if there are any common factors that appear to affect whether or not a station has a high proportion of high flow days in the summer months. From the first stage of the analysis we split the year into two seasons, 'winter' running from 1 October to 30 April and 'summer' running from 1 May to 30 September. In choosing these seasons we looked at the number of high flow days in each month, rather than trying to get seasons of equal length. Table 5-1 shows all stations where the proportion of high flow days that occur in the summer months is greater than 15 per cent. If high flow days occurred without any seasonality all stations would have around 42 per cent of high flow days in the summer months. From Table 5-1 we can see that although some of these stations are highly urbanised, many are not. However, when the gauging station descriptions are examined (Marsh and

Hannaford, 2008) many of these stations (27030, 28003, 28080, 30001, 38003, 39005, 39012, 39014, 54006, 71004) have large artificial influences on the flow regime and one is tidally influenced at spring tide (twice a month).

Station number	Station name	AREA	BFI	URBEXT1990	Proportion of high flow days in summer months
25005	Leven at Leven Bridge	196.3	0.44	0.01	0.153
27030	Dearne at Adwick	310.8	0.61	0.097	0.155
28003	Tame at Water Orton	408	0.62	0.403	0.365
28080	Tame at Lea Marston Lakes	799	0.69	0.268	0.202
30001	Witham at Claypole Mill	297.9	0.67	0.19	0.157
38003	Mimram at Panshanger Park	133.9	0.94	0.043	0.265
38007	Canons Brook at Elizabeth Way	21.4	0.41	0.173	0.174
39005	Beverley Brook at Wimbledon Common	43.5	0.64	0.379	0.299
39012	Hogsmill at Kingston upon Thames	69.1	0.74	0.207	0.211
39014	Ver at Hansteads	132	0.86	0.079	0.185
39023	Wye at Hedsor	137.3	0.93	0.07	0.294
52009	Sheppey at Fenny Castle	59.6	0.68	0.029	0.155
54006	Stour at Callows Lane, Kidderminster	324	0.72	0.165	0.192
65001	Glaslyn at Beddgelert	68.6	0.31	0	0.158
68004	Wistaston Brook at Marshfield Bridge	92.7	0.62	0.083	0.168
71004	Calder at Whalley Weir	316	0.43	0.098	0.151
72002	Wyre at St Michaels	275	0.32	0.006	0.152

Table 5-1: Gauging stations with high proportion of high flow days in summer months.

One of the ways of accounting for seasonality in our analysis is to split the data up into two parts, summer and winter. However, due to the small number of high flow days in summer, for most stations there are not enough data points to accurately estimate the spatial dependence for the summer months. So, to see if ignoring seasonality had any effect we estimated dependence measure $N(p)$ (defined in equation (4.7.4)) empirically from the data for all observations, and for observations from the winter months, keeping the threshold, v_p, to be the 99 per cent probability threshold for the whole dataset for both analyses.

Figure 5-8 shows the correlation between the estimates of $N(p)$ obtained for the winter months and for all months. There is clearly a strong correlation between the two, with only four sites having a large difference (greater than 0.05) between the winter months and all months. These were 28003, 39005, 39023 and 54006, which all have a relatively high proportion of summer high flows. However, for almost all sites the estimates obtained for all months were slightly lower than those obtained for the winter months. This suggests that the level of spatial dependence is lower in summer than in winter, but that ignoring seasonality does not have a huge influence on the results of the spatial dependence analysis. We do not, therefore, include seasonality in the rest of the analysis.

This analysis of seasonality is likely to be influenced by the use of daily flow data and might not be fully representative of 'flash flood' events. However, these events are also not necessarily captured by river flow gauges. Accounting for localised, rapid onset flood events is one of the motivations to examine extending the modelling to include rainfall data, which is recommended as a further research topic.

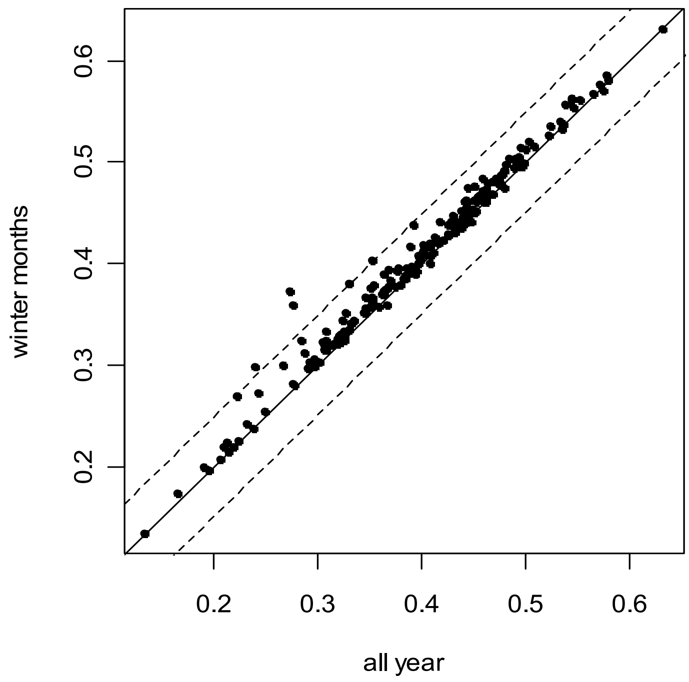

Figure 5-8: Plot of dependence measure $N(p)$ obtained for winter months against that obtained for all months. Solid line perfect correlation, dashed lines +/- 0.05.

5.6 Large scale dependence

One of the requirements of the conceptual model is that it can estimate dependence over large spatial and temporal scales. For England and Wales, following quality checking, we have data from 432 flow gauging stations, so to estimate the dependence at the highest level of detail for a whole year would involve estimating 432 * 431 * 365 a and b parameters with associated residuals. The method to handle missing data increases these numbers.

Although in theory there is no limit to the number of dimensions in which the Heffernan and Tawn model can be applied, in practice most computers would struggle with allocating memory to data arrays of the size needed for the highest level of detail, even with careful programming. The approach that we adopt to overcome this problem is to thin the data and so reduce the level of detail of the model. Up to a point it is possible to do this without any loss of useful information. However, there are situations in which a thinned dataset will not be able to capture the details required. These situations are generally smaller scale. For instance if we wish to set inputs to a hydraulic model, either at an estuary or for a large scale river model; or if we wish to inform emergency planners about how likely two roads are to be blocked by flooding at the same time; then it may be important to use as much detail as possible. In these situations a separate, detailed analysis using the full dataset for the region of interest would be appropriate. Note that the more detailed analysis would, in this situation, be consistent with the larger scale model, that is it would add to, rather than replace, the large scale data.

The first way in which it is possible to thin the data is to reduce the number of stations included in the fitted model. A sensible way to do this is to remove stations of shorter, lower quality records and which have very high dependence with other stations. In this way it is possible to reduce the number of gauging stations used in the analysis. For example, Keef *et al*. (2009b) used a set of 271 gauging stations over Great Britain (including Scotland) with a separation of around 25-30km between them.

The second way in which it is possible to thin the data is to look at time intervals of greater than one day. For instance it would be possible to fit models to five-day maxima. The main drawback with this approach is that it is inevitable that we will split some events. An alternative to setting fixed time windows is to fit the model to local maxima of the individual time series, where the window of time corresponds to particular events. This approach means that the average length of an event can vary considerably over the country. For permeable catchments the time window over which to look for local maxima should be around 20 days, whereas for impermeable catchments it should be around five days. It therefore requires some skill to match up the maxima from different parts of the country and to work out which values should be taken for non-extreme time periods.

Figure 5-9 shows an example based on five day maxima. The plot shows time series from summer 2007 for two flow stations in the Yorkshire Ouse catchment (27002, Wharfe at Flint Mill Weir and 27041, Derwent at Buttercrambe) and two in the Thames catchment (39001, Thames at Kingston and 39016, Kennet at Theale). Along with the daily time series we have plotted the five-day maxima from this series. If we look at the bottom three plots we can pick out instances when flood peaks have been split into two time windows – for instance the highest peaks for all three. However, for the first time series plots the five-day separation appears to capture all the major peaks in the flow series.

In summary, we estimate large scale dependence by applying the fitted model to a carefully selected subset of the gauging stations to maxima of blocks of a certain number of days. The actual length of the blocks should be as short as possible (to obtain independent events) but can be set to a specified interval if required to ensure that certain types of events are included, such as the flooding of summer 2007. The number of gauging stations chosen in the analysis should be as large as is feasible computationally.

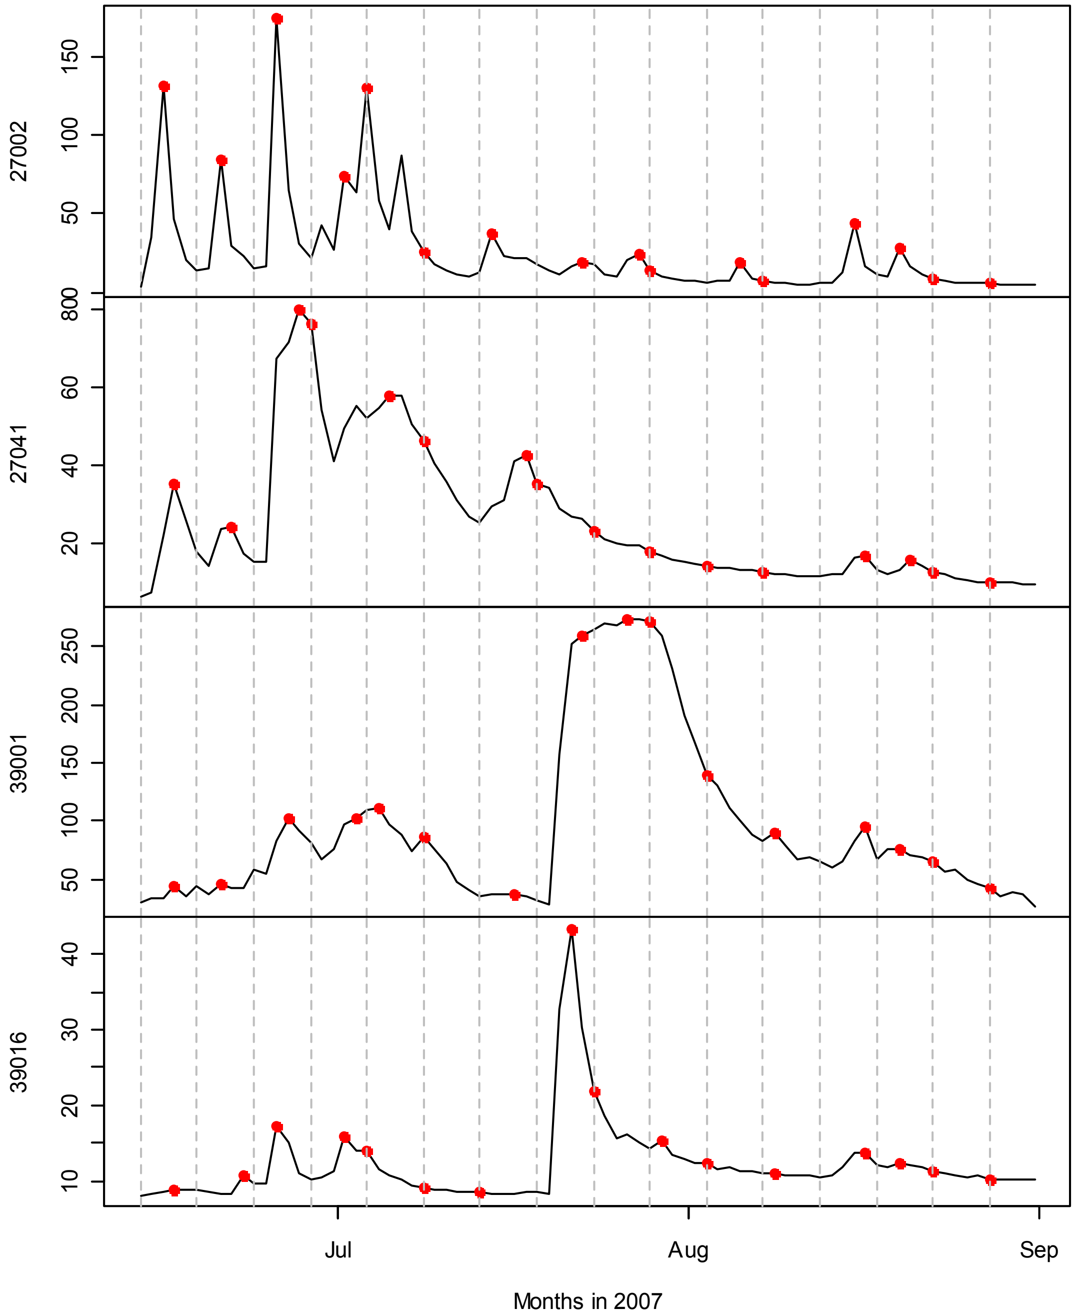

Figure 5-9: Lines show daily time series, red dots show five-day maxima. Grey dashed lines indicate the first day of the five-day window.

5.7 Model checks

An important feature of the Heffernan and Tawn model is that due to its asymptotic derivation it requires relatively few assumptions about the underlying form of the marginal distributions or the dependence structure. Specifically, the variables can follow essentially any distributional form as for large enough values of these variables all possible marginal and dependence structures must follow the Heffernan and Tawn model. The key decision is therefore deciding what 'large enough' corresponds to - this is determined by the threshold level that we will use (so far we have used the 99 per cent probability threshold).

To be able to apply the Heffernan and Tawn model in any application, the validity of the assumptions needs checking for the required application data. This corresponds to determining if the selected threshold is large enough for the asymptotic derivation to be justified. Clearly the assumptions are more likely to be valid the higher the threshold chosen. However, the higher the threshold the less data we have to make inferences about the statistical model, so for model fitting we require the threshold to be a low as possible. Therefore we are looking for the threshold to be sufficiently high for the asymptotic derivation to be justified yet sufficiently low for there to be enough data to reliably fit the statistical model.

If for a selected value of the threshold the assumptions appear appropriate that gives us confidence in using the associated statistical model, whereas if some assumptions are found to be inappropriate then further investigation is usually required, and typically a higher threshold needs to be used. Here we list the assumptions that are made and outline how these modelling assumptions can be assessed by use of data.

5.7.1 Assumption 1 – Tails of the marginal distributions are generalised Pareto distributions

Our marginal model assumes that a generalised Pareto distribution is an appropriate form of statistical distribution to describe the exceedances of a threshold u. This is consistent with the assumption that the annual maxima follow a generalised logistic distribution or the generalised extreme value distribution. For marginal variable W the generalised Pareto has cumulative distribution function G_u with

$$G_u(W) = P(W < w | W > u) = 1 - \left[1 + \xi \frac{w-u}{\sigma_u}\right]^{-1/\xi} \quad \text{for} \quad u < w < \infty,$$

where $\sigma_u > 0$ and ξ are scale and shape parameters respectively. If the generalised Pareto distribution is appropriate above the threshold u then for a higher threshold v $(v > u)$ the exceedances of v have cumulative distribution function G_v where

$$G_v(W) = P(W < w | W > v) = 1 - \left[1 + \xi \frac{w-v}{\sigma_v}\right]^{-1/\xi} \quad \text{for} \quad v < w < \infty,$$

where $\sigma_v = \sigma_u + \xi(v-u)$. Thus the exceedances of v also follow a generalised Pareto distribution with shape parameter also being ξ but scale parameter σ_v. Thus both ξ and $\sigma_v^* = \sigma_v - \xi(v-u)$ are invariant to the choice of threshold v.

A detailed way to assess if a selected threshold u is sufficiently large enough is to check if the estimates of both ξ and σ_v^* are stable over different choices of threshold v, $v > u$. Lack of stability in these parameter estimates suggests a higher threshold is required. For example, Figure 5-10 shows plots of estimated σ_v^* and ξ against the non-exceedance probability of different thresholds v for the tide gauge data at Whitby. Two features are clear; there is a slight increase in σ_v^* and corresponding slight decrease in ξ around a probability threshold of 0.97, which suggests a lower probability than this would be too low; and that as we get to much higher thresholds (above a probability threshold of 0.99) the parameter estimates change rapidly. This is due to the small numbers of exceedances above this threshold. From this plot we can see that taking a threshold equal to the 0.97 probability threshold is likely to be sufficient.

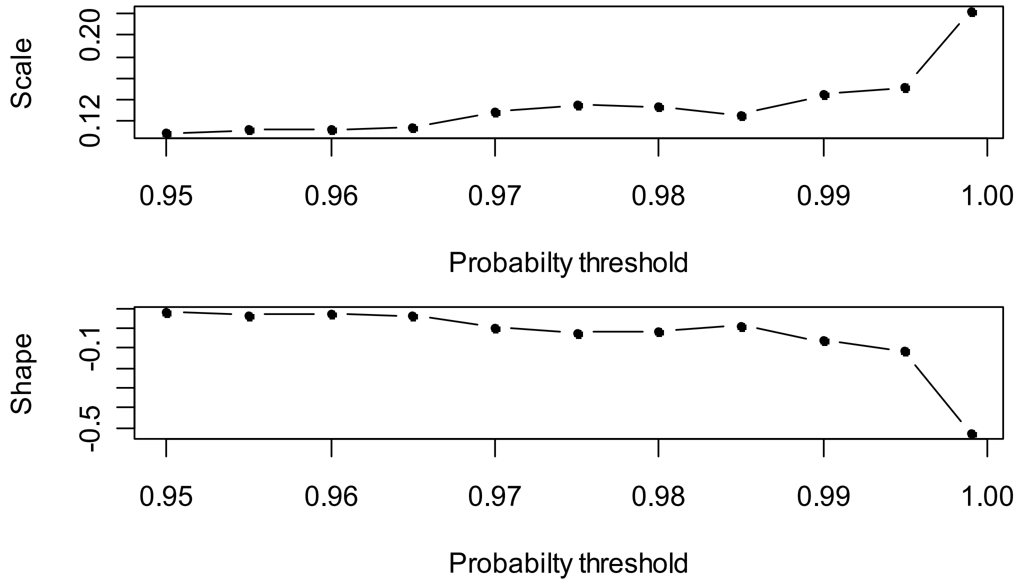

Figure 5-10: Plots indicating stability of parameter estimates over different thresholds. Top plot indicates scale parameter estimates σ_v^* for fitting a GPD model to skew surges at Whitby. Lower plot indicates shape parameter estimates ξ for the same fitted models.

Once a threshold has been selected that satisfies the stability property above, the second stage of the diagnostic method is to check if the distribution of the exceedances follows a generalised Pareto distribution. Q-Q plots, which compare the observed quantiles of the sample data of exceedances, with the predicted quantiles from the fitted generalised Pareto distribution, are the standard tool for making this assessment. Figure 5-11 shows a Q-Q plot for a fitted GPD to skew surge data, again from Whitby, above the 0.97 probability threshold. We can see that the fitted points lie close to the line of perfect fit. We have visually inspected Q-Q plots of this type for all of the (river and coastal) stations used in the analysis.

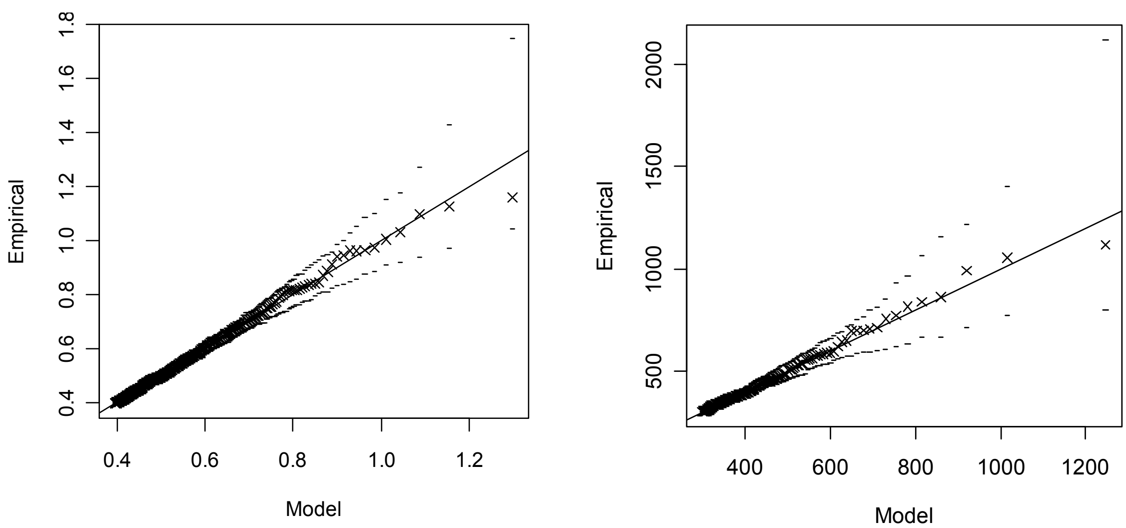

Figure 5-11: Q-Q plots showing quantiles on empirical scale versus quantile on modelled scale (fitted GPD) left hand plot skew surge at Whitby above the 0.97 probability threshold, right hand plot daily mean flow data for the Tyne about the 0.99 probability threshold. Solid line perfect fit, crosses data points, horizontal lines modelled 95 per cent confidence intervals for the data points.

5.7.2 Assumption 2 – The correct threshold for fitting the Heffernan and Tawn model has been selected

There are two features of the Heffernan and Tawn method that it is possible to test to ensure that the asymptotic assumptions the model is based on are correct. The first is the assumption that the standardised residuals $\boldsymbol{Z}$ are independent of the conditioning variable X. The second is that the parameter estimates are constant above the fitted threshold. A method to test this second feature is related to the parameter stability plots used in threshold selection for the GPD. You fit the Heffernan and Tawn model at a series of different thresholds and select the lowest threshold above which the parameter estimates are stable.

Independence of the conditioning variable and the standardised residuals

The dependence structure component of the Heffernan and Tawn model requires the selection of a threshold u such that the joint distribution of standardised residuals $\boldsymbol{Z}$, defined by

$$\boldsymbol{Z} = \frac{\boldsymbol{Y} - \boldsymbol{a}(x)}{\boldsymbol{b}(x)}$$

is independent of the conditioning variable X, where $\boldsymbol{a}(x) = \boldsymbol{a}x$ and $\boldsymbol{b}(x) = x^{\boldsymbol{b}}$.

To make an assessment of this independence assumption we examine the distribution of $\boldsymbol{Z}$ for different ranges of X. Typically it is sufficient to split the range of X for $X > u$ into three equally likely, non-overlapping ranges and then to plot the associated samples of $\boldsymbol{Z}$ for each range. Monte Carlo tests of independence are then used to formally test for differences between these three distributions.

If the $\boldsymbol{Z}$ and X data are found to exhibit dependence, and hence violate the independence assumption, then this may indicate that the selected forms $\boldsymbol{a}(x) = \boldsymbol{a}x$ and $\boldsymbol{b}(x) = x^{\boldsymbol{b}}$ may not be suitable or a higher threshold is required than u. In most cases we have identified a failure of the independence assumption; the best solution has been to raise the threshold to be used.

5.7.3 Assumption 3 – The marginal distributions of the standardised residuals are non-degenerate with no mass at infinity

This assumption of non-degeneracy and no mass at infinity are of theoretical relevance and included for completeness. Constraints are applied in fitting the model to ensure no values of any component of $\boldsymbol{Z}$ can be infinity for any given sample. Provided more than one different value of $\boldsymbol{Z}$ is obtained for each different marginal variable of $\boldsymbol{Z}$ then the assumption is satisfied. In practical terms, the assumption will be met for any plausible input data.

5.7.4 Assumption 4 – The fitted joint distribution is consistent with the data

Key to all aspects of model checking is that the fitted distribution is consistent with the observed data. We check this by simulating from the fitted model (see Section 6) a large sample and then comparing empirically quantities from the data sample and simulated sample. The features that we compare include: marginal distributions; the copula; distributions of combinations of the variables, such as $T(\boldsymbol{Y})$, for some function of interest T.

As the simulated sample can be of any length, we typically generate samples which are much larger than the data sample. This allows us to derive rare return period events. An important part of the assessment of the assumptions of the model we use is that these large return level events are consistent with broader knowledge (not contained in the data) of the variables being simulated, this may correspond to knowledge of historical flooding events.

6 Implementation of statistical model for rivers and coasts – event simulation

6.1 Introduction

Given a suitable model for the joint probability distribution of the 'source' variables (that is, hydraulic loads), our generic conceptual risk model requires a procedure to simulate random samples from this distribution in order to integrate with 'pathway' and 'receptor' models. In the previous section of this report we have described a suitable joint distribution model. Now, we discuss the simulation of a set of 'events' from that distribution, corresponding to the Monte Carlo process illustrated in the generic model in Figure 2-6.

The events will consist of simulations of data representing the severity of the source variables (flows and/or sea levels) for selected study areas. There are two reasons that we use a simulation approach. The first is the semi-parametric nature of the statistical model that we propose to use. The second is that this method will fit in with complex pathway and receptor models such as in RASP and NaFRA.

Each simulated event will consist of standardised flows or sea levels for all points of interest. Each of these standardised flows/sea levels will have standard Gumbel distributions, so can be transformed to any scale of interest, for example the original daily scale or annual exceedance probability scale.

Note that the simulation is of a sample of independent flood events (with temporal dependence having already been accounted for in the distributional model). It is important to realise that the event set is not a time series simulation.

6.1.1 Pointwise simulation at gauging sites

The choice of method used to simulate a large event set of artificial data for a set of gauging stations is independent of the choices of time window and subset of gauging stations. The method is as follows.

Before simulating data we fit the Heffernan and Tawn model using each station as the conditioning station in turn. This means estimating sets of $\boldsymbol{a}$ and $\boldsymbol{b}$ parameters and residuals, $\boldsymbol{Z}$, taking X to be each gauge in turn and $\boldsymbol{Y}$ to be all other gauges. These parameter sets and residuals will cover same-day dependence and also lagged dependence in the way demonstrated in Section 4.7.9. Each of these sets of parameters will contain estimates of $a_{j,\tau}$ and $b_{j,\tau}$ for all $j \in \Delta$ and for all $\tau \in A$ where Δ is the set of all non-conditioning gauging stations and A is the set of lags of interest. Because we are simulating an 'event' we set the lags of interest to be the same for the whole dataset, based on the findings of the temporal dependence analysis (see Section 5.3) or user-supplied definitions. Each of the sets of residuals will consist of n subsets, $Z_{j,\tau}$, for all $j \in \Delta$, $\tau \in A$ and $t \in \{1, \ldots, n\}$ where n is the number of exceedances of the threshold by X.

The procedure for simulating a single event is to select a conditioning site, simulate a value at that site and then generate values at other sites preserving the dependence structure captured by the model through the parameters $\boldsymbol{a}$ and $\boldsymbol{b}$, and the residuals, $\boldsymbol{Z}$.

These steps are repeated many times to get a large event set. In order to generate the correct number of threshold exceedances at each site each conditional gauging station (model) must be simulated from the correct number of times. In simulating each single event, the value at the conditional gauging station is the maximum value over all sites. So in generating the event set the proportion of events for which a particular site is conditioned upon should be the same as the true proportion of events where that site is the maximum. For any particular site this proportion can be estimated by first generating a large number of events conditioning on that site; these events will be representative of all events where this site is above the threshold. The proportion of these events where the conditioning site is the maximum can be used to estimate the true proportion of events where this site is the maximum. Because each event has at least one threshold exceedance, a site cannot be the maximum if it is below the conditioning threshold.

The distribution function of spatially aggregated damages can then be obtained by combining the distribution function of aggregated damage for the simulated events with a distribution for occurrences of events, as discussed further in Section 7.

6.2 Fluvial events

6.2.1 Simulated events at gauging stations

The simulation procedure outlined above results in an event set defined at gauging station locations. Figure 6-1 shows three events simulated in this way. Each event consists of ‘flows’ (on the chosen transformed scale) at each gauging station. Of course to model flood risk we need to have information about the event severity at any location on a river or coast, not just the gauged points. Before these events can be used to assess the associated impacts on receptors we must therefore first interpolate the event data between the gauging stations and into headwater catchments.

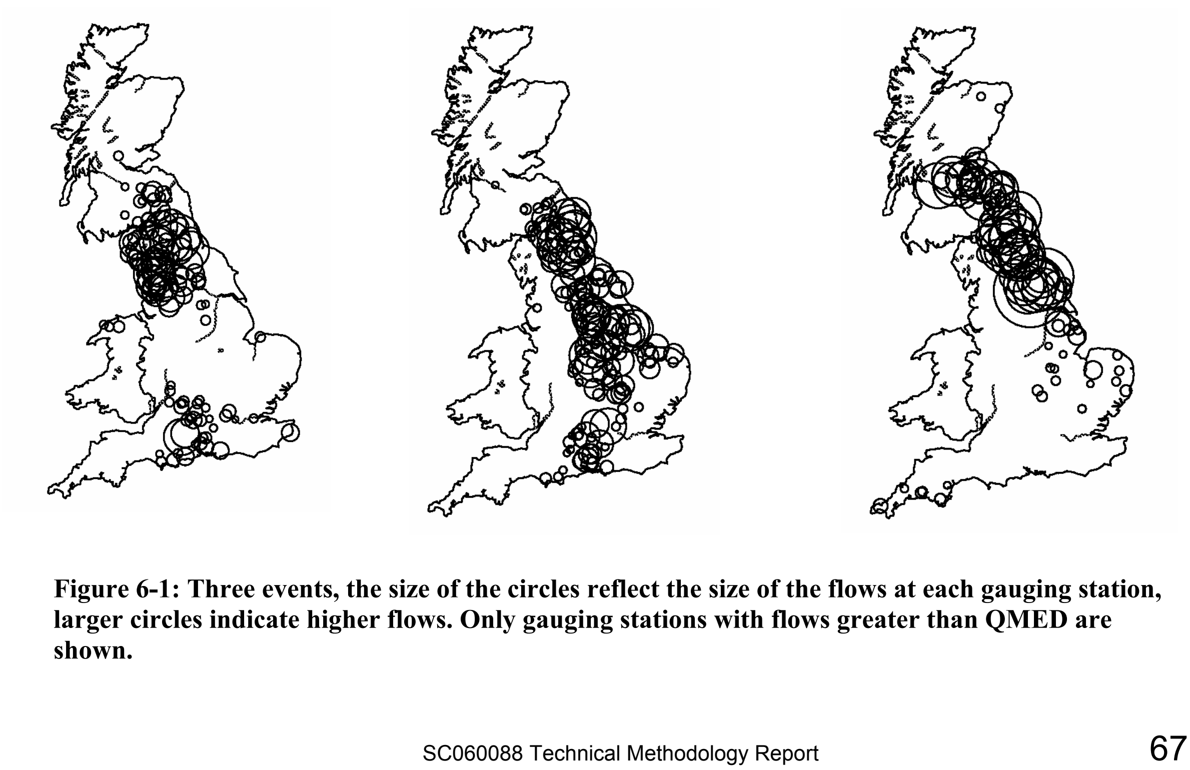

Figure 6-1: Three events, the size of the circles reflect the size of the flows at each gauging station, larger circles indicate higher flows. Only gauging stations with flows greater than QMED are shown.

6.2.2 Flows between gauging stations

The method used to estimate dependence is applied to gauging station flow records; it is not capable of directly estimating the probability of multiple floods at ungauged sites. But for a large scale risk analysis we need to include locations in between gauging stations and on ungauged rivers. There are two possible ways of overcoming this difficulty. The first is to build covariates into the model; this would create a model that could be used to predict the dependence at a set of ungauged flow sites. The second is to use the flows simulated at the gauged sites to predict the flows at ungauged sites. This prediction should be carried out using information about the catchment.

There are advantages and disadvantages to both approaches. The main advantage of including covariates in the dependence model is that the resulting regression model could be used in catchments that were completely ungauged. The main disadvantages of including covariates in the dependence model are that it is difficult to define quantitative relationships between catchment descriptors and the parameters of the dependence model, and the numerical difficulty of including covariates in the Heffernan and Tawn model.

Two examples of investigations into obtaining quantitative relationships between pairwise dependence of extreme river flows and catchment descriptors are Keef *et al*. (2009b) and Office of Public Works (OPW, 2008). In both investigations although correlations with differences in catchment descriptors were found, the predictive capability of all differences in catchment descriptors was very low. The strongest relationships found were whether or not two sites were ‘connected’ (that is, had the same water flowing through them) and the geographical separation between the catchments of the sites. Instead of building ungauged sites directly into the joint distribution model, we therefore consider two interpolation methods to use the flows simulated at gauging stations to predict flows at ungauged sites.

Interpolation methods

The two interpolation methods we consider are:

- Interpolating the flows along the flow network using shared catchment area.
- Estimating the flows at ungauged sites, based on the distance between the catchment centroid of the ungauged site and the catchment centroids of the gauged sites.

An obvious method of interpolating through the river network is to use distance along the river, for example defined as stream length from any point to the sea. These changes in river length are always smooth. However, we would expect changes in relative flood severity to change more suddenly in some places, for example at confluences, where upstream catchment area can also change abruptly. This could be particularly visible when moving from a main river to a tributary, especially if that tributary is small compared to the size of the main river. One of the chosen methods for predicting flows at ungauged sites is therefore to use differences in catchment area, which can be calculated from a DTM for any ungauged river location as well as at the gauged sites.

Alternatively, we can use geographical distance between catchment centroids as a distance measure to represent the position of ungauged locations relative to gauged sites in order to estimate the scaled ‘flow’ variable. One reason for using catchment centroids is that there could be two flow sites located close to each other but with catchments that are relatively far apart, an example could be tributaries descending from opposite sides of a valley that join the main river close to each other. We would expect that the flows at these points are more different to each other than for nearby flow sites located on the same river. Conversely, headwater sites that are close together but in different catchments would appear closer using the centroids distance measure. Catchment centroids have recently been adopted in the revised FEH procedure for data transfers from gauged to ungauged sites, partly for the above reasons (Environment Agency project SC050050/SR; Kjeldsen and Jones, 2007).

Both methods that we use to predict flows at ungauged sites can be viewed as interpolation. The first simply along the river network by viewing catchment area at points along the network as a 1D surface, the second by viewing the geographical area of interest as a 2D surface with distances on the surface represented by distances between catchment centroids.

When interpolating river flows it is necessary to ensure that the interpolated flows at each site along the river are sensible given the range of possible flows at that site. The interpolated flows must take into account the probability distribution of flows at each site. One way to do this is to interpolate on a standardised scale. Possible standardised scales are the return period scale, the probability scale and the Gumbel scale.

In this study we have used the Gumbel scale because, for most flow sites, the tail of the Gumbel distribution is a similar shape to the tail of the true distribution of the gauged sites. The relative changes in flow with return period are similar on both the Gumbel scale and for the original flow data. The method used to transform data onto the Gumbel scale is given in Section 4.3.1.

For both distance measures used to interpolate (difference in catchment area and distance between catchment centroids) we use simple linear interpolation to obtain the expected flow at intermediate points. For distance between catchment centroids this equates to a simple weighted average, given in equation (6.2.1):

$$X = \frac{\sum_{i=1}^{N} X_i d_i^{-1}}{\sum_{i=1}^{N} d_i^{-1}} \tag{6.2.1}$$

where d_i is equal to the geographical distance between the catchment centroids for sites i and the site of interest (X), and X_i is the value on the Gumbel scale at site i.

For the difference in catchment area the interpolation formula is more complicated. The reason for this is that we have to take account of the connectivity of the river system. We cannot use this method to interpolate points that are not directly connected to each other. We define connected in the following way; two flow sites are connected if they have the same water flowing though them. In the left plot of Figure 6-2 there are two connected gauging stations, one upstream and one downstream. We denote the flow at the upstream gauging station X_{UP} and the flow at the downstream gauging station X_{DOWN}. Similarly we denote the catchment area of the upstream gauging station A_{UP} and of the downstream station A_{DOWN}. The formula we use to interpolate flow at point X using catchment area is given in equation (6.2.2):

$$X = X_{DOWN} + (X_{UP} - X_{DOWN}) \left(\frac{A_{DOWN} - A}{A_{DOWN} - A_{UP}} \right). \tag{6.2.2}$$

The interpolation formula given in equation (6.2.2) is insufficient if, for a given interpolation point, there is more than one upstream gauging station. In this case we define A_{UP} as the sum of all upstream catchment areas, so $A_{UP} = \sum_{k=1}^{n_{UP}} A_k$, where

n_{UP} is the number of upstream gauging stations, and X_{UP} as a weighted average of all upstream flows, so $X_{UP} = \sum_{k=1}^{n_{UP}} X_k A_k / A_{UP}$.

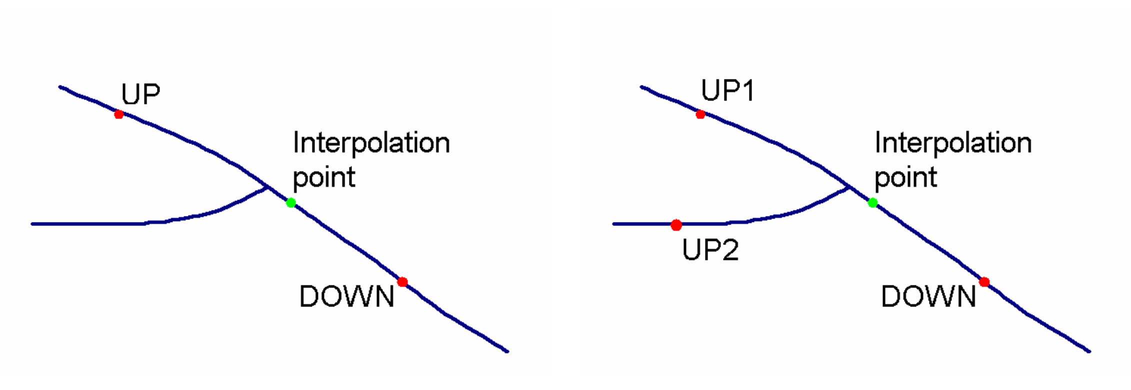

Figure 6-2: Illustration of interpolation method. Left plot one upstream gauge, right plot two upstream gauges.

Interpolation method checks

To check the performance of our method of interpolation we have used the method of 'leave one out' cross-validation on five river catchments from the UK. These are the Bristol Avon, the Yorkshire Ouse, the Severn, the Thames, and the Wye. These were selected to provide a wide range of differing types of catchment in terms of geographical location, size, shape, relief and geology. Maps of these catchments are given in the Appendix to show the distribution and number of gauging stations in each case. Figure 6-3 shows an event on the Wear catchment in north east England before and after interpolation.Two possible methods of interpolating these flows are given in Section 6.2.2.

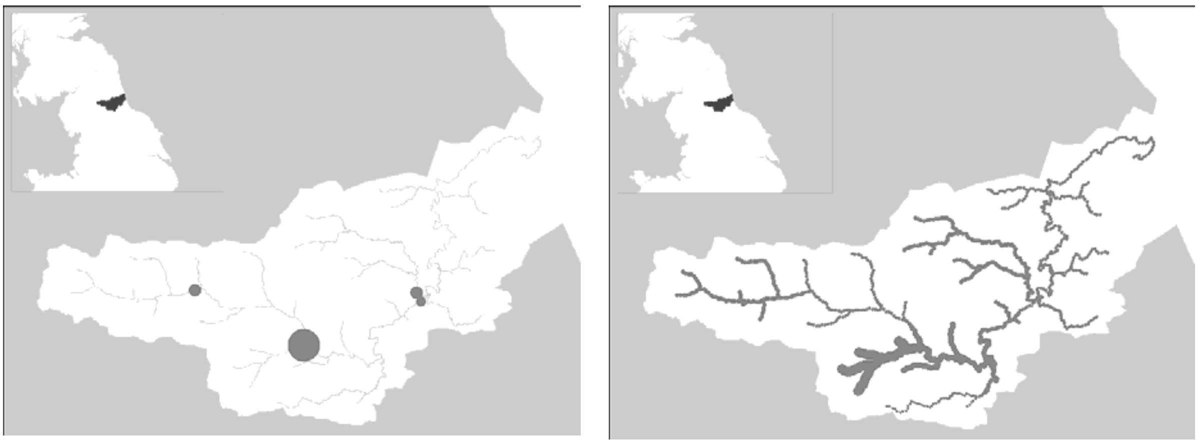

Figure 6-3: An event on the Wear catchment. Left plot shows flows simulated at gauging stations (larger discs indicates higher flow), right plot shows the interpolated event (thicker lines indicate high flows).

The two interpolation methods were assessed separately and also for each catchment in turn. We carried out two cross validation studies for each catchment. The aim of these studies was to evaluate how good each interpolation method is at predicting the flows at ungauged flow sites.Obviously the only data to make this comparison is at gauged sites. Hence, the principle of the cross validation technique is to remove each gauging station in turn, predict the flows at the removed gauging station (by interpolating from the other stations) and then compare the predicted flows with the real flows.

The method used in these studies is as follows. For each catchment we denote the set of gauging stations to be Δ where $\Delta = \{1, \ldots, n\}$, n being the number of gauging stations in the whole catchment. For each site $j \in \Delta$ we then used equations (6.2.1) and (6.2.2) to estimate the flows at site j.

Figure 6-4 shows summaries of the correlation between the observed 'flows' (on the Gumbel scale) and the 'flows' obtained by interpolating from the other gauging stations. To obtain these summaries we first used both methods of interpolation to predict 'flows' at each gauging station. We then calculated the correlations (Spearman's ρ) between the observed 'flows' and the interpolated 'flows' for each station in turn. We present the results as boxplots, grouped for each catchment. We can see that, for most sites, both interpolation methods perform well. With the exception of the Thames and the Severn, using area interpolation the 25 per cent quantile correlation (lower edge of the box) between the real and interpolated flows is well above 0.9 for all catchments.

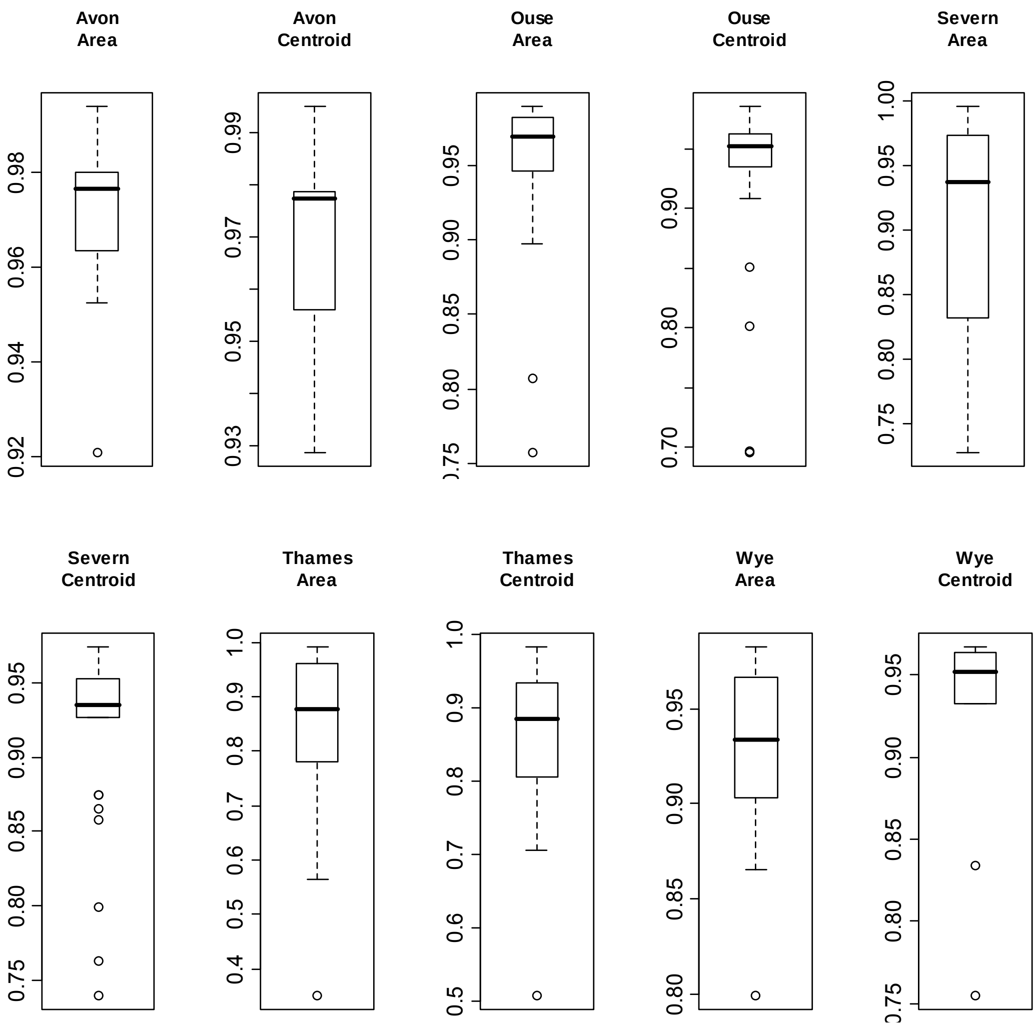

Figure 6-4: Boxplots showing range of estimated correlation (Spearman's ρ) coefficients between observed flows on Gumbel scale and interpolated flows on Gumbel scale for each method, for each river. Results are grouped into each catchment. Box indicates upper and lower quartiles of the data (estimated correlation coefficients), thick lines median of the data, whiskers indicate the largest and smallest data points that are no more than 1.5 times the length of the box away from the box. Circles indicate outliers.

The gauging station on the Thames that is not well represented by either of the interpolation methods is station 39065, Ewelme at Ewelme Brook, which has an area of 13.4km^2 and BFI 0.98. This gauging station has no upstream stations and the nearest downstream gauging station is 39072, the Thames at Windsor Park, which has an area of 7046.0km^2 and BFI 0.72. The station with the nearest catchment centroid is 39002, the Thames at Days Weir, which has an area of 3444.7 km^2 and BFI 0.64. When this information is examined it is not that surprising that station 39065 is poorly interpolated. It is likely to respond at a very different rate to the nearest flow station using catchment centroids, and it is very small compared to the nearest flow station using catchment area. So the correlation between flows at this site and flows at the other sites is likely to be very small. For sites like these, it is probably physically realistic to assume that they act effectively independently of the other flow sites. In our further investigation of the performance of the interpolation methods we also found that the stations for which interpolation was particularly poor were those where there was a large amount of significant intervening catchment between the gauge at which the flows are to be interpreted, and the nearest site in terms of distance between catchment centroids. The implications of this are that there are some locations where the interpolation method will not work as well as in other areas.

We can also see that for most sites the interpolation using catchment area seems to perform better, but for some the interpolation using catchment centroids seems to perform better. In particular for the neighbouring catchments of the Wye and the Severn the catchment centroids method performs best.

The results presented here will be biased towards suggesting that interpolation using catchment area will work better than interpolation using distance between centroids. The reason for this is that in each analysis we only use flows contained in the same catchment. For interpolation using catchment area this will not change when we examine a whole region. However, when using interpolation using distance between catchment centroids we will be able to use flows in neighbouring catchments to help interpolate the flows. So, assuming that flows in neighbouring catchments will give us more useful information we should obtain better results. For downstream gauges the difference is likely to be very small. However, for small upstream gauges the improvement may be quite large.

During our work on the proof of concept demonstrations for the north east region we found that, in general, the method of interpolation based on catchment centroids outperformed the method of interpolation based on catchment area. There were many gauges for which the method based on catchment centroids was best, and few for which the catchment area method was best. Additionally, the method based on catchment centroids can be used for all locations.

6.3 Coastal events

To estimate the joint probability of flooding from extreme still water levels at different coastal sites we model dependence between the skew surges at the sites using the Heffernan and Tawn model and the dependence between the associated tidal levels empirically from the predicted tidal series.

A complication with such an approach is the event definition to account for the duration of events and the travel times of surges and tides. For example, high tides do not occur at the same time along the whole coastline, also there are usually two high tides each day but some days there is only one. We must be able to match a single skew surge to a single high tide at site, and for a spatial event also do this between different sites along the coast. Additionally, we also want to be able to associate a still water level event with the daily mean river flow. Therefore we reduce the still water level data into a daily format, with days defined to run from 9am to 9am so that the river and coastal data are concurrent.

For each day, we select the highest skew surge and the highest tide. Together these are used to represent the still water level on that day at a site. They also provide a clearly defined set of values from different sites along the coast on that day. By defining daily maximum sea level in this way we will never underestimate the maximum sea level in the day, but we may overestimate it. Therefore he levels obtained will have a positive bias. Given the diurnal cycle in the tides leads to a differential in the two tidal levels in a day, we anticipate that the level of bias will be slight.

Interpolation between the sea level gauge sites is simpler than between the flow gauge sites because we only need to interpolate along one line, rather than a series of connected lines. One method of using our technique to simulate sea levels at a high enough resolution to use in pathway and receptor models is as follows:

- Using the model data, determine how far along the coast two sites can be considered to be highly dependent. The technique we have presented to assess at which time lag two flows on the same river can be considered to be dependent can be used here.
- Use the method detailed in Section 6.1.1 to simulate an event at each of the points in the network.
- Use simple linear interpolation between the network points to obtain sea levels at any point on the coastline.

6.4 Joint river and coastal

If we define daily maximum sea level and use daily mean river flows, then use of the Heffernan and Tawn model to estimate the probability of joint fluvial and coastal flooding follows automatically. We can model the dependence between flows and surge, in exactly the same way as between flows and flows, or surge and surge. So we will be able to obtain simulated data sets of spatially coherent river flows and skew surges. The skew surges can be combined with high tides in the way described in Section 6.3 to obtain an event set of spatially coherent river flows and skew surges. The interpolation between the gauge sites can be carried out independently for flows and sea levels using the methods described in the previous sections. For sea levels the interpolation is in any case easier because the variation is smoother than for river flows.

6.5 Model checks

The model is checked during fitting via test of the assumptions described in Section 5.7. In addition, there are some checks made of the event data simulated from the model as described in this section. Simulated values at the gauge sites are checked to ensure that they have the same distribution as the real values. This is done by inspecting plots of the distribution function and also by carrying out a parametric test (the Kolmogorov-Smirnov test is suitable). The second aspect of the simulation method to check is that the pairwise dependences within the simulated data reflect the pairwise dependences within the real data set. This can be checked by estimating pairwise dependence measures for both the simulated data and the real data and more simply by inspection of pairwise scatter plots, as shown in Figure 6-5 for an arbitrarily chosen set of river gauging stations.

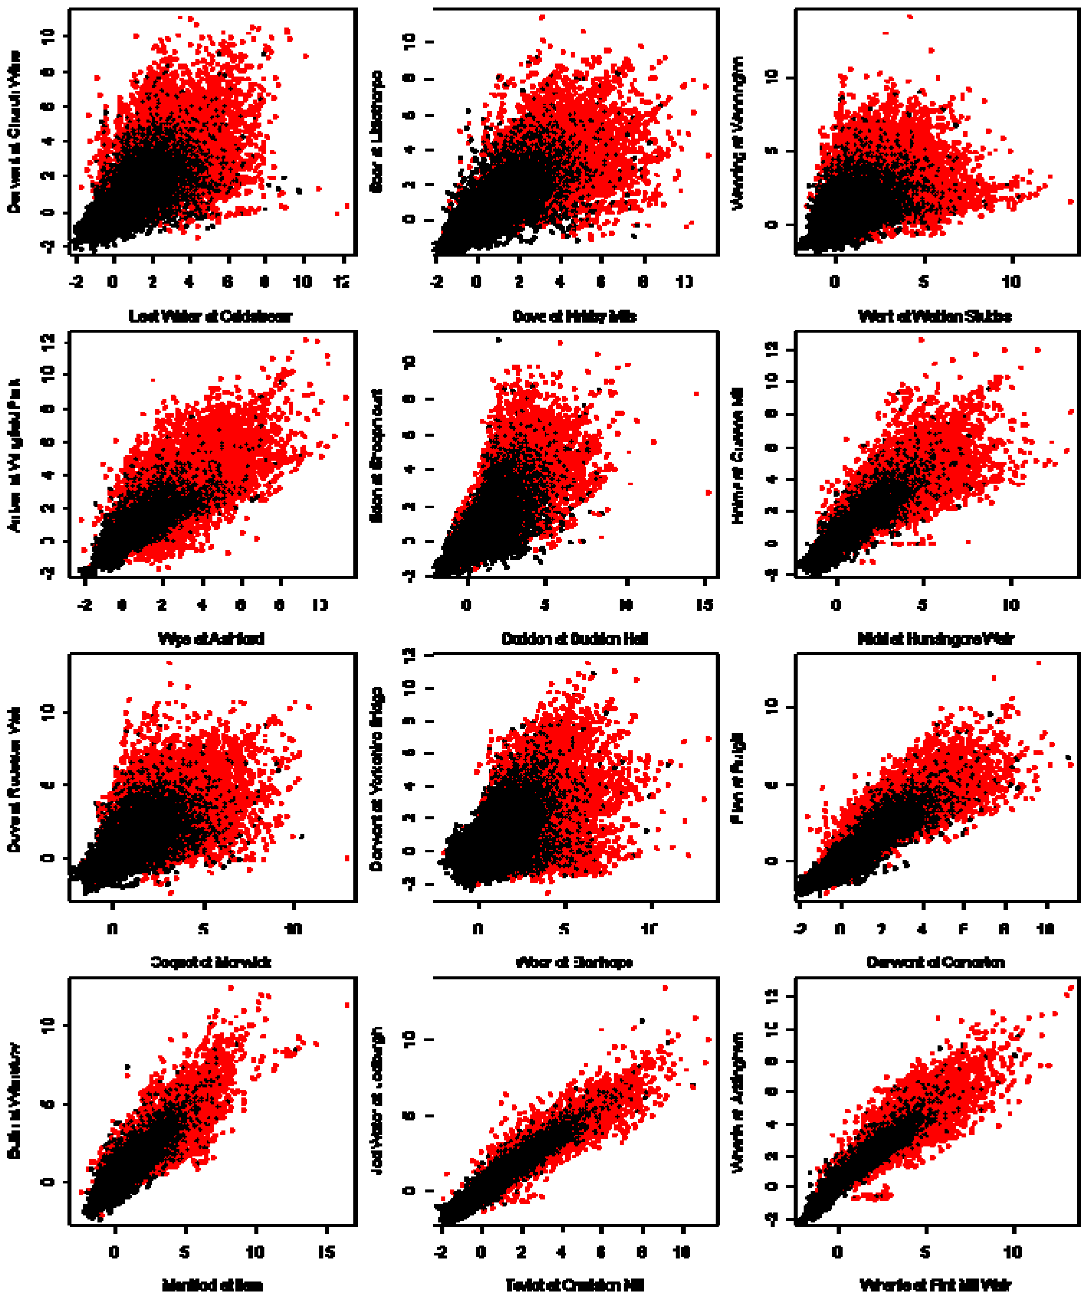

Figure 6-5: Scatter plots showing daily river flows on a standardised dimensionless scale for an arbitrarily selected set of gauging stations. Black dots are gauged data, red dots simulated data.

These checks are of course possible up to limits implied by the available gauged data - it is only possible to check the simulation method up to a threshold for which there is enough observed data to make the comparison. Even so, plots such as Figure 6-5 provide a useful check that the model is generating a plausible extrapolation from the gauged data.

7 Integration of pathway and receptor models

7.1 RASP (Risk assessment for system planning)

RASP is a probabilistic method for modelling flood risk that includes the performance of flood defences, attempting to take account of location, type, condition and failure modes. The RASP method seeks, in principle, to integrate the probability distribution for hydraulic loads on a flood defence system with the probability distribution of defences failing (conditional on the load) and the translation of the resulting flood volumes into depths and damages. It does this for systems of defences, exploring the many possible failure scenarios using Monte Carlo simulation as a numerical approximation to the probability integration.

As we will discuss below in more detail, RASP is essentially a univariate model in terms of the way the source of flooding (hydraulic load) is represented, although joint probabilities of fluvial and tidal load have also been considered. However, it does provide a method for representing the ‘pathway’ and ‘receptor’ components of the source-pathway-receptor concept for flood risk.

This section therefore looks at how we may be able to take advantage of the established RASP methods to integrate the defence system pathway with the spatially aggregated, multivariate risk model outlined in the previous sections.

The RASP methods are currently implemented in a nation-wide analysis within the National Flood Risk Assessment (NaFRA), but can also be applied to smaller study areas, such as the Thames Estuary in TE2100. Spatial aspects of risk become more relevant when working at large scales of aggregation and we have therefore concentrated on how a NaFRA implementation of RASP could be linked with a spatial model. However, the principles are generic and could be relevant for smaller scales, for example in modelling risk for a future Catchment Flood Management Plan.

Integration of RASP with the multivariate, spatial model has been considered within this project. Section 7.2 describes an approach that we have concluded would make the link possible in principle. We have also verified this conclusion in discussions with the Environment Agency and consultants working with the organisation on development and implementation of RASP and NaFRA (HR Wallingford and Halcrow).

7.1.1 NaFRA

The NaFRA methodology is based on assessing the flood risk for individual impact zones over the whole country. The description here is based on documentation supplied by the Environment Agency relating to recent implementations (to 2007). It is recognised that some details may change in the near future or already have changed.

NaFRA is a broad scale assessment of the likelihood and impact of flooding at a national scale. These impact zones are 50m grid squares that are at risk of flooding from rivers or the sea.

The receptors in each impact zone are accounted for by inclusion of three separate, national spatial datasets. The first is land use, the second what buildings are on the land (national property database), and the third dataset describes who lives in the buildings (social flood vulnerability index, derived from census data). For each impact zone the potential economic damage can be calculated from these datasets.

The flooding pathways to each impact zone are accounted for by inclusion of flood defence data from NFCDD. Each impact zone is related to a set of defences that protect it. Each impact zone may have multiple pathways by which it can flood. The probability of failure for each particular defence, given a certain loading condition, is calculated using fragility curves.

The loadings from the source variables on each defence are calculated differently for sea and rivers. For rivers water levels at the 100- and 1000-year return periods are obtained from existing hydraulic models or the Flood Zones projects where more detailed information is not available. Using the ground level as the two-year level, these levels are then interpolated to obtain water levels for other return periods. For coastal loadings sea levels and wave heights are obtained from the JOIN-SEA analysis method, using data from a variety of different research projects.

For the coastal analysis the coastline is split up into joint probability regions. These are regions in which the distributions of sea levels, wave heights, and dependence structure between sea levels and wave heights within each region are similar enough to be treated as the same. The regions currently in use are understood to be updates of joint probability regions produced originally in Defra R&D project FD2308.

The average annual damage for each impact zone is calculated by evaluating the probability and damage for each flood-causing scenario. The method used in the latest version of NaFRA is HR Wallingford's rapid flood spreading mechanism with a fixed volume from breaching formula to create a dataset of flood depths. The corresponding damage from these depths are then calculated, and combined with the probability of hydraulic load, to obtain the average annual economic damage.

7.1.2 MDSF (1 and 2)

The Modelling and Decision Support Framework (MDSF) is a tool originally designed to assist in catchment flood management plans and shoreline management plans.

In principle the datasets contained within MDSF perform the same function as the datasets considered within NaFRA. They allow the user to determine land use, building type, and social vulnerability of receptors. Within MDSF2 there are planned to be two ways of identifying flood depth. The first will be to use the rapid flood spreading mechanism embedded within the software, the second will be to import pre-determined depth grids produced by an external model.

The important distinction between MDSF1 and MDSF2 is that MDSF1 is a deterministic tool, designed to analyse scenarios, whilst MDSF2 is planned to include a probabilistic treatment of defence systems by implementing RASP. The RASP implementation appears very similar to that in NaFRA, but should also provide a user-friendly option to import water levels to represent the hydraulic load from external modelling.

7.2 Integrating RASP with a multivariate, spatial flood risk model

7.2.1 The RASP scheme for expected annual damages

'Dependence of load' is one of the assumptions made in the RASP methodology. However, the RASP concepts for defence failure analysis do not, in principle, preclude a spatially distributed source variable, or multiple sources. For example, in the Thames Estuary analysis described by Gouldby *et al*. (2008), the tidal nature of the river required joint probability analysis of tidal and fluvial boundary conditions, combined with a hydraulic model, to provide distributions of the water level loading along the river. Here, a joint probability method already existed for the limited case of fluvial/tidal boundary conditions in one estuary. But for larger scales of spatial aggregation, no such method has previously been available.

In RASP the assumption of 'dependence of load' generally refers to an assumption that the process of defence failure somewhere within the system being modelled does not have a hydraulic influence on the loading elsewhere. It is easy to think of situations where this may not hold true, but more difficult to build such hydraulic interactions into the risk model.

A similar assumption is made in the spatial risk model that we are developing in this project. Hydraulic interactions, for example in estuaries or for floodplains with large washlands, are implicitly included in the dependence analysis to the extent that they are reflected in the observed record, but will not necessarily be included for scenarios that have never been observed. This is a challenge for any flood risk model and will need further work in future. With the techniques currently available, some progress could be made by using hydraulic models within a simulation framework to represent the interactions explicitly, although this might only be practical in specific cases like TE2100 where there are the resources to allow it.

The spatial structure of the RASP method is shown in Figure 7-1. In the terminology of RASP, the floodplain is divided into 'impact zones' which may lie behind a system of defences, each of which can have an independent resistance to flood loading. Current implementations of RASP further divide the floodplain into 'impact cells', which could correspond to a hydraulic model or flood spreading algorithm grid, or some aggregation thereof. A series of nodes are located along the river or coastline representing the locations where the hydraulic load is defined.

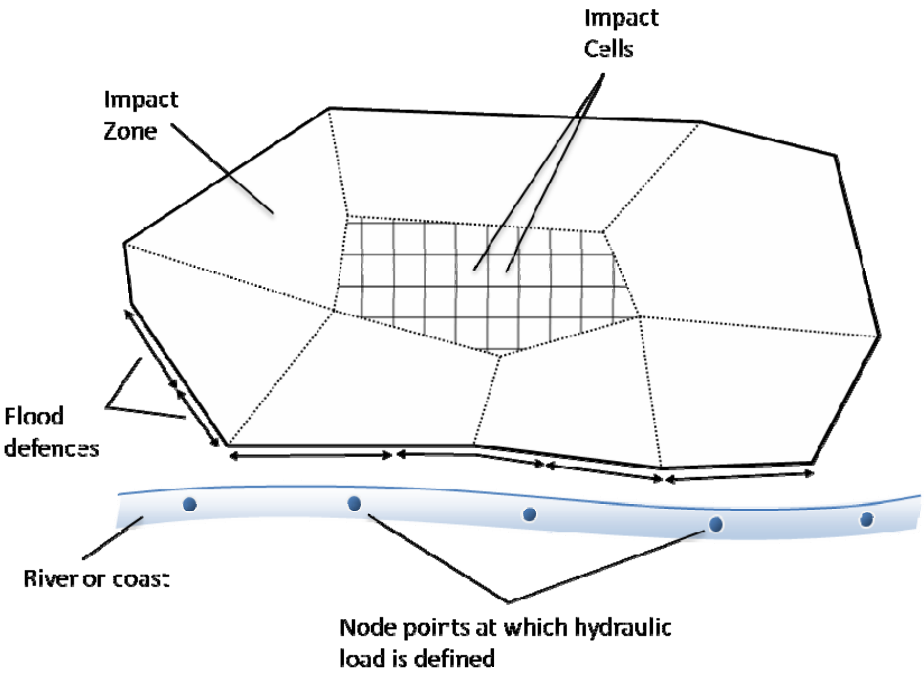

Figure 7-1: Schematic illustration of the RASP division of the floodplain.

If translated into the simple generic model structure of Section 2, the RASP approach corresponds in principle to Figure 2-4. If RASP methods were actually applied using a brute force Monte Carlo simulation of this type, then the spatial hydraulic load terms could simply be supplied as an additional stochastic input. In practice, even for a univariate hydraulic load term, this brute force approach would be inefficient because it would require repeated creation of a very large Monte Carlo sample. Instead, RASP replaces the exhaustive Monte Carlo procedure with a structured simulation in which a set of discrete loading conditions (corresponding to 39 return levels for the analysis in NaFRA) are combined with a Monte Carlo simulation of possible defence states and an approximate flood spreading algorithm.

This procedure generates an estimate of the probability distribution of flood depth, conditional on loading, incorporating what is known or assumed about the defence system performance. The estimate is defined spatially for each cell within a grid covering the area associated with a given defence system. In RASP, the primary result of interest is the Expected Annual Damage (EAD). Conceptually this is obtained (Hall *et al*., 2003) by integrating the derived probability distribution of flood depths $f(\boldsymbol{y})$ with the cost function, $C(\boldsymbol{y})$, which is also dependent on depth, where $\boldsymbol{y}$ is a vector of depths generated from the Monte Carlo simulation of defence system states. EAD is therefore defined by Hall *et al*. as:

$$EAD = \int_0^{\max(y)} f(\boldsymbol{y})C(\boldsymbol{y})dy.$$

Gouldby *et al*. (2008) give the alternative approximation:

$$EAD \approx \sum_{i=2}^{q-1} \left\{ \left[P\left(L \geq \frac{l_i + l_{i+1}}{2}\right) - P\left(L \geq \frac{l_i + l_{i-1}}{2}\right)\right] \bar{c}_i \right\}.$$

Here the load distribution has been discretised into q levels l_i, $i = 1, \ldots, q$,, where $P(L \geq l)$ is the annual probability of experiencing a load of level l or greater, and $\bar{c}_i$ is the mean economic damage, for a given hydraulic load, for a given flood impact cell.

7.2.2 Calculation of expected economic damages conditional on hydraulic load

The calculation of $\bar{c}_i$ is conditional on the load and is based on a Monte Carlo simulation of defence states, combined with flood spreading, for a given load. Hence there is a value of $\bar{c}_i$ for each of the q load conditions. The 'hydraulic load' is not defined precisely in the generic description given by Gouldby *et al*., but for their case study of the Thames Estuary it was the water level derived from a hydraulic river model based on detailed physical data.

We cannot always assume such detailed information about the river channel and so the load may have to be expressed more generally, for example in factorial terms relative to defence standard as proposed by Hall *et al*. These authors also set out a method for deriving volumes of water flowing onto the floodplain. Although highly generalised and approximate, in the absence of better information similar formulae could be used to estimate a volume to floodplain based on the return level for locations in our spatial model. Alternatively, given a spatial statistical model, river flows and sea levels could be computed directly from their marginal distribution at any location. For river flows it would then be necessary to define volumes to floodplain via assumptions about hydrograph shape and channel capacity. The assumptions needed would be similar to those discussed by Hall *et al*. or currently made in NaFRA.

Whatever method is chosen to derive a volume to floodplain from the distribution of the hydraulic loads, the RASP methodology currently requires the distribution function for the load to be specified (both to compute the expected damage per floodplain cell for a given load level and also to compute the EAD). In NaFRA we understand that this is done using flow frequency distributions based on automated application of the FEH methods (CEH flow grids) where more detailed flow data are not available. In Gouldby *et al*. (2008) a hydraulic model for the main channel of the Thames was used. In MDSF2 it is proposed that users in effect supply their own distribution via an import of water levels for specified return levels from external models, combined with an interpolation procedure within MDSF2. The exact choices are not always easy to summarise from the literature because the RASP methods have evolved through different applications. Details of the current specifications will require clarification should the methods in this project be taken forward to link with NaFRA or MDSF2.

If the joint distribution modelled with the Heffernan and Tawn method were to be used, then the required data would be the marginal distributions of the 'source' variable, that is, hydraulic load, at specified locations along the river or coast. These data are readily available from the spatial model. But because the spatial model and event set simulation can generate data on a probability scale, or a related scale such as a Gumbel reduced variate, it is possible instead to use the marginal distributions that are currently applied in NaFRA (that is, derived using FEH methods, if the CEH flow grids are applied as the input to RASP). What this means is that it is possible to make direct use of the expected damage data produced by NaFRA in combination with a spatial event set produced by the models developed in this project.

7.2.3 Integration of RASP conditional expected damages with a spatial event set

Given the expected damage per load level ($\bar{c}_i$), it should be possible to estimate the spatially-aggregated distribution of damage by using these conditional expected damages in combination with a spatial event set that specifies vectors of load levels per event. This requires the values of $\bar{c}_i$ for flood cells to be associated with the simulated loads for a specified location. Given an assumption that the spatial variation in load is at a larger scale than the flood zones used for calculating $\bar{c}_i$, then this should be a reasonable approximation. Such an assumption is in effect that of 'dependence of load' at the defence system scale, which is in any case made in RASP.

Calculation of the damage distribution would be made more accurate by considering not just the expected damager per load level, but also the distribution of values around the expectation. This distribution could be captured from the Monte Carlo simulation of defence states already performed in recent versions of RASP, and summarised as an empirical distribution function or, to an approximation, by the variance of that distribution.

For the sake of efficiency, it may be worth exploring the possibility of defining the expected conditional damages as the average for an entire flood impact zone, and then associating each impact zone with the hydraulic loads at one location, or perhaps the mean of the loads at the locations that RASP shows to be contributing to the risk for the particular impact zone. A possible configuration is illustrated in sketch form in Figure 7-2. Here, the choice is left open about mapping between locations defined in the spatial event set and RASP impact cells or impact zones.

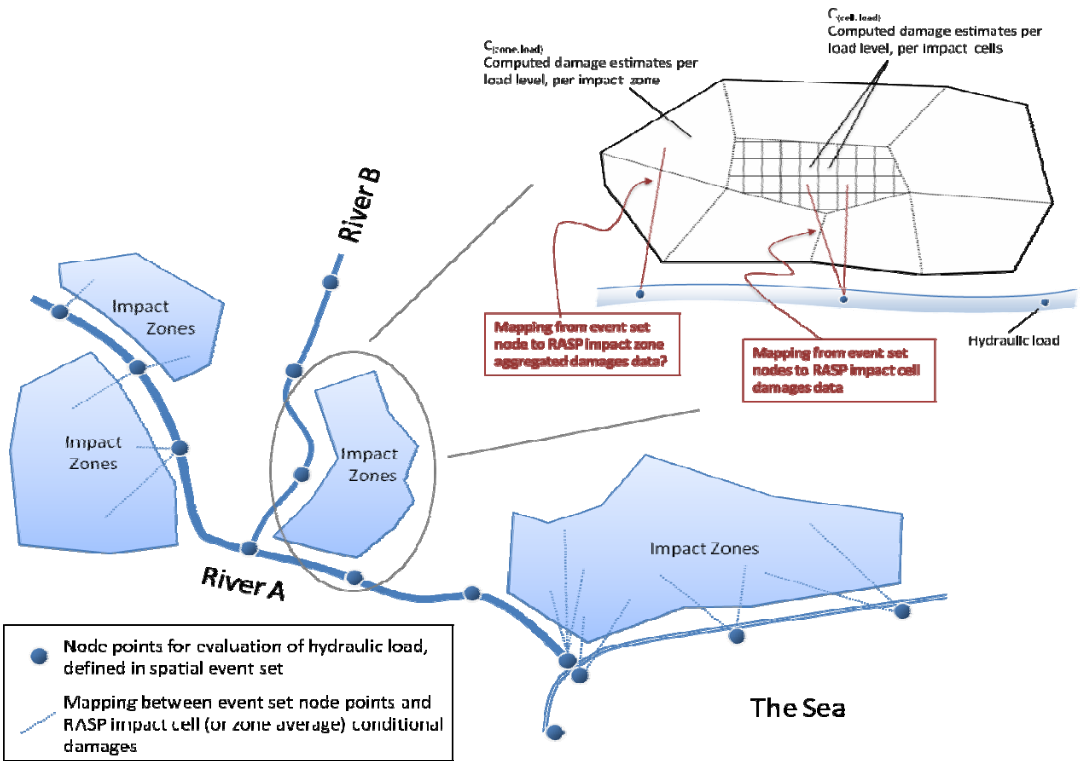

Figure 7-2: Schematic illustration of the an integration of RASP conditional expected damages with spatial event set.

In summary, an outline algorithm for linking our multivariate risk model to the RASP methods is set out as follows:

Integration with RASP

1. Use an implementation of RASP, for example in NaFRA, to calculate the expected economic damage conditional on hydraulic load for every flood impact cell or impact zone. This is a parameter that approximately captures the flood risk ‘pathway’ through the performance of the defence system and the routing of flood water, along with the ‘receptor’ or risk through the economic damages calculation.
2. Establish a suitable spatial connection between RASP impact cells (or at a higher level of spatial approximation, impact zones) and locations at which our spatial model can simulate the intensity of the hydraulic load. In effect this means tying the RASP conditional economic damages data to points on the river centre line or coast. This process could be assisted by the analysis done using RASP to associate risk with defence assets.
3. Generate a spatial event set from our multivariate model. Each event in the set consists of the hydraulic load intensity concurrently for each river or coastal location (node point).
4. For each event, calculate the expected damage corresponding to each location by combining steps 1, 2 and 3. If the RASP analysis is also used to provide information from the Monte Carlo simulation about the variance surrounding the conditional expectation of damage, this could be used to enhance the method efficiently by sampling values from the distribution of damages, rather than always taking the expectation.
5. For each event, accumulate the expected damages over the required spatial units, for example catchment or country.
6. The ordered event set of accumulated damages provides an empirical estimate for the probability distribution of the spatially aggregated damage.
7. Evaluate the distribution of the annual damage and its expectation and variance by accounting for the distribution of damage per event and the distribution of the number of events per year.

From steps 6 and 7 we can easily calculate a curve relating the damage distribution to exceedance probability or return period, in other words a risk profile. Because we use the empirical distribution of damages from each simulated event set to build up the risk profile we need a large number of events to reduce sampling noise in this risk profile.

The RASP approach uses the marginal distributions of hydraulic load to calculate EAD. Since EAD is the mean damage it is possible to aggregate it over spatial scales to estimate, for example, catchment or national expected annual damages. However, it does not say anything else about the probability distribution of damages at the aggregated scale. This is because calculation of the full distribution of the aggregated cost function has to take account of dependence to avoid potentially serious bias, as shown for a simple example in Section 4.

In contrast, our conceptual model aims to derive the full distribution of event damages, including the probabilities of extreme events, for any spatial scale of aggregation. In this sense, it should be seen as having the potential to augment RASP-based tools such as NaFRA and MDSF2 so that we can understand not just the expected, 'average' risk but also the risk of suffering a much more severe event.

To estimate expected annual damages, we can compute $EAD = E(D) * E(N)$, where $E(D)$ is the expected damage in a flood event and $E(N)$ is the expected number of events per year. To evaluate the full distribution of annualised damages, the distribution of the event damages is combined with a model for the distribution of the number of events per year as described in step 7 above. This is similar to the procedures used in single site peaks over threshold models.

8 Uncertainty

In estimating the return period of spatially aggregated loss there are a number of sources of uncertainty. In this section we investigate these sources.

8.1 Natural variability in observed source data

Because we only have flow observations from the past 40 years or so we only have a sample of all possible observations. Like any sample this is subject to sampling error which leads to estimation uncertainty of the fitted model. In previous work (Keef *et al.* 2009a, b) this sampling error has been accounted for by use of a block bootstrap, in Heffernan and Tawn (2004) a standard bootstrapping technique was used. The principle behind bootstrapping is to use the variability contained within the observed data to inform us about the overall variability in the full data. Block bootstrapping is an extension of the basic bootstrapping approach that is suitable for time series data. We first define the basic bootstrap method (for more details see Davison and Hinckley, 1997). Let X be a univariate sample of independent identically distributed (i.i.d.) data of size n. If we are estimating the value of a parameter θ of the distribution from which the sample X is taken and we wish to estimate the uncertainty in the estimate $\hat{\theta}$ of θ then the basic bootstrap method is as follows.

- Re-sample X with replacement to obtain a bootstrapped sample X^* of size n.
- Calculate $\hat{\theta}$ for X^*.

These two steps are repeated B times, where B is a large number, to obtain a sample $\hat{\theta}^B$, of estimates of $\hat{\theta}$ of size B. The variation in this sample can then be assessed and used as the estimate of uncertainty in the parameter estimation. In particular, if we wish to obtain a 95 per cent confidence interval then we can take the end points of this interval to be the 0.025 and 0.975 quantiles of $\hat{\theta}^B$.

The basic bootstrap can be extended to multivariate data in the following way. Let $\boldsymbol{X} = (X_1, \ldots, X_d)$ be a sample of i.i.d. multivariate data with n multivariate observations. In the multivariate case, in order to maintain the between-variable dependence, it is possible to sample the vector observations $\boldsymbol{X}_i$, $i = 1, \ldots, n$ with replacement. The observations of all variables $(X_1, \ldots, X_d)$ at these re-sampled observation points can then be used to create the bootstrapped samples $\boldsymbol{X}^*$.

If we used the basic bootstrap to estimate the uncertainty of a multivariate time series parameter then we would lose any temporal dependence present in the original series. This is because we re-sample the observation point independently of each other.

The assumption that all data is i.i.d. is invalid for river flow data, and for sea level data, which generally display seasonality and short-term dependence. In the block bootstrap the original multivariate time series $\boldsymbol{X}_t$ is divided into blocks. These blocks are then re-sampled with replacement to create the bootstrapped sample $\boldsymbol{X}_t^*$. The blocks are chosen to be large enough to preserve the temporal dependence in the time series but small enough to allow a large number of possible combinations in each re-sample. Due to the seasonality of the data it is sensible to choose blocks that correspond to a whole year. In making this choice we make the assumption that floods in one year are independent of floods in the previous year. If we chose blocks that corresponded to calendar years then this assumption would be invalid. This is because the start date of the calendar year falls in the middle of the flood season and what happens in the second half of a flood season is not independent of what happens in the first half. The start day of the block year should be chosen so that it is in the middle of the lower seasonal flow periods. In Keef (2007) it was found that a suitable date to choose as the start date of a year for the purposes of block bootstrapping river flows was 1 August, as most UK rivers are in the middle of seasonal low flows on this date. We have not investigated a suitable year start date for sea levels, but as highest surges generally occur in late winter/early spring this should also be a suitable date for sea levels.

To illustrate the block bootstrap method we show how it can be used to obtain confidence intervals for the economic risk profile curves associated with a pair of flow gauges, each of which is located close to a city in the north east region. An economic risk profile was calculated using depth and damage data as described in Section 3.5.2 and the simulation methods described in Section 7, as follows:

Steps in obtaining economic risk profile

- Step 1: Fit a GPD distribution to the data from each flow gauge.
- Step 2: Fit the Heffernan and Tawn model, conditioning on each gauging station. This results in two separate fitted models.
- Step 3: Use the simulation procedure described in Section 6.1.1 to simulate events equivalent to n years of data (here we choose n to be 1000).
- Step 4: For each simulated event calculate the return period at each site.
- Step 5: For each simulated event calculate the associated expected loss at each site, conditional on the return period at that site. Figure 8-1 shows the conditional expected damage curves we use in this illustration.

We also compare our results with those that would be obtained if we assumed complete dependence or complete independence. When using the Heffernan and Tawn model we only simulate events such that at least one of the gauging stations has an observation above the modelling threshold. So to make the comparison, we also only simulate events for complete dependence and independence where at least one of the gauging stations has an observation above the threshold. The simulation procedure assuming complete dependence is to replace steps two and three above by simulating a single value above the modelling threshold and to use this value for both sites. The simulation procedure assuming complete independence is to replace steps two and three by simulating two values independently, with one of these values being above the threshold. All simulation is carried out so that the simulated values at each site have the correct marginal distributions.

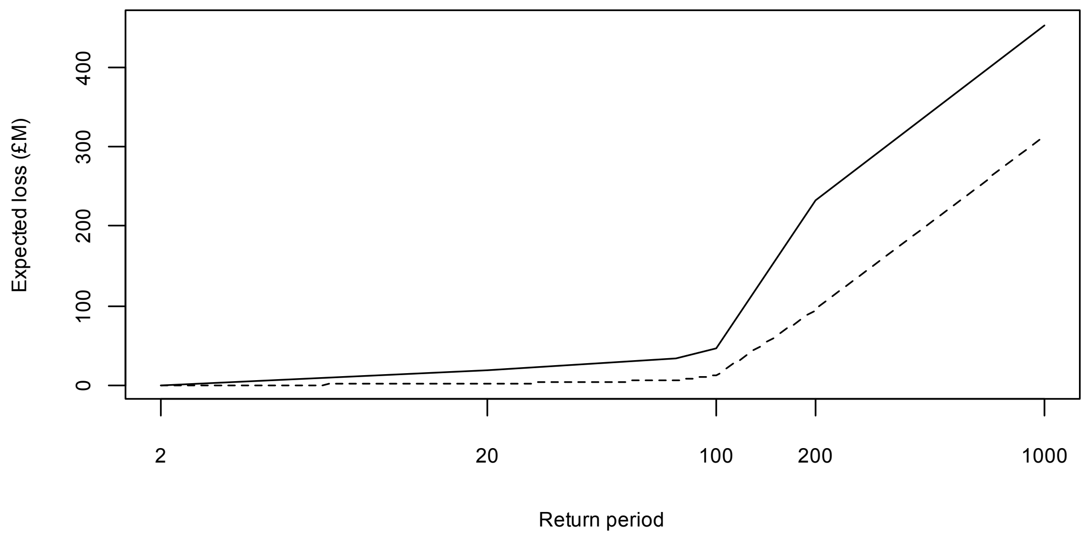

Figure 8-1: Assumed conditional expected damage curves for two cities in the north east region (undefended).

When using block bootstrapping the extension to the procedure defined above needed to obtain 95 per cent confidence intervals is as follows:

- Re-sample the blocks (years running from 1 August to 31 July) with replacement to obtain a bootstrapped sample of the flow data at both sites.
- For this bootstrapped sample carry out steps 1-5 for obtaining an economic risk profile as outlined above.
- Repeat these steps 1 and 2 to get a large number of bootstrapped loss risk profile curves.
- Calculate the 0.025 and 0.975 quantiles of the curves.

Figure 8-2 shows the resulting risk profile curves and associated confidence intervals. For clarity we have plotted these curves twice, first with confidence intervals surrounding the modelled risk profiles, and again with confidence intervals surrounding the curves obtained using complete dependence and complete independence.The first thing to note in Figure 8-2 is that the (hypothetical, undefended) risk profile for the two cities is somewhere between the complete dependence and complete independence cases. More importantly, for higher return periods the confidence intervals for the modelled risk profile curves do not contain the curves for either the complete dependence or complete independence cases. Additionally the confidence intervals for the complete dependence and complete independence cases do not contain the modelled risk profile.

The implications of this for flood risk management are that it is not appropriate to make either of the assumptions of complete dependence, or complete independence. In other words, it would not be sensible to assume that floods at the two sites were completely independent of each other as that would underestimate the likely losses, or that the two sites always experience the same return period at the same time as that would overestimate the likely losses. Even allowing for uncertainty, any model of the aggregated risk of flooding should include a model of the dependence structure of the flood risk variables.

Another point to note is that that the confidence intervals for the modelled curve are wider than those for the false assumptions and also get wider as the return period increases. This reflects that the Heffernan and Tawn model allows a range of dependence structures and the information in the data is not sufficient to discriminate completely between them. In contrast under the false assumptions of complete independence and complete dependence there is no statistical variation in the modelling of dependence which gives narrower confidence intervals but biased estimates.

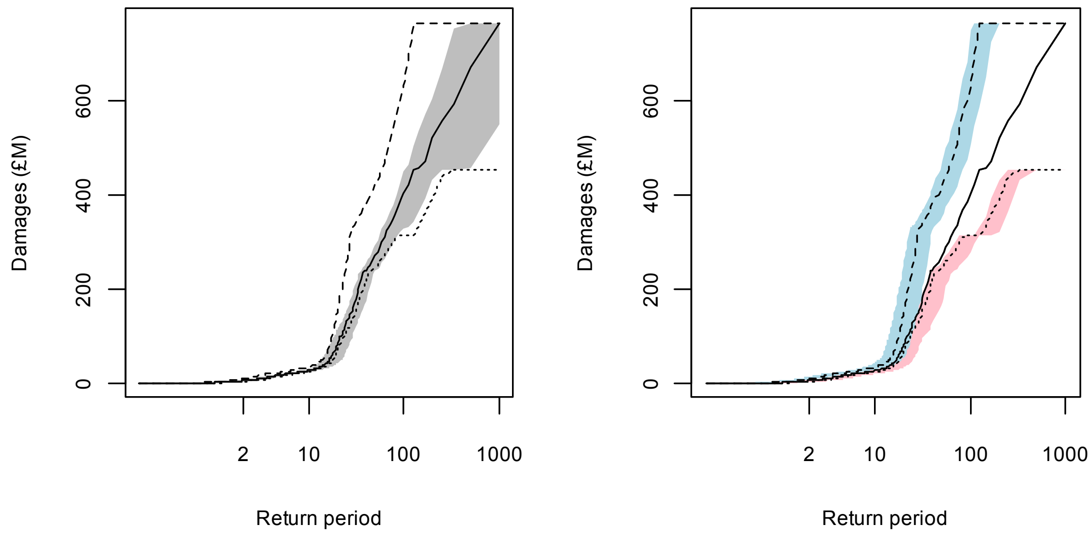

Figure 8-2: Combined risk profile plots for two cities in the north east region. In both plots the solid black line is the fitted model, the dashed line assumes complete dependence, the dotted line complete independence. In the left plot the bootstrapped confidence intervals for the modelled curve are shown in grey, in the right hand plot confidence intervals assuming complete dependence are shown in blue and confidence intervals assuming complete independence are shown in pink.

8.2 Simulation uncertainty

Because we simulate events at random from the fitted model there is some additional uncertainty involved. This uncertainty can be reduced by simulating a very large number of events. There are two ways to go about this. The first is to simply generate a very large number of events in one simulation. The second is to simulate a smaller number of events replicated many times and then average over these results. As methods to reduce simulation uncertainty there is very little to choose between these two approaches. However, the second approach is much more efficient computationally. This approach was used to produce the plots in Figure 8-2. For each bootstrapped sample we repeated the simulation procedure 100 times (steps 3-6 in steps in obtaining the risk profile curve), we then took the median of these curves to be the value for that bootstrapped sample.

8.3 Interpolation

The spatial interpolation procedure used to transfer information from gauged reaches to ungauged reaches is also a source of uncertainty. There are two factors to be considered here. The first is the potential for bias in the interpolated values, not least because in our method of interpolation no value will have a return period higher than the highest simulated value. The second is variation around the expected values at each interpolation point.

It is possible to build realistic variation into the interpolation procedure as a future development. In doing this the value obtained at each interpolation point would not be simply the expected value at that point, but rather a sample from a distribution around that expected value. However, there are a number of difficulties in this approach, particularly for river systems. The main difficulty would be ensuring that the implied event peak flow values would vary smoothly along the rivers, and also according to physical principles at confluences. Building in appropriate constraints would that take into account our knowledge of the physical river system is a refinement that should be considered in future. This is one of the main sources of uncertainty and one of the most difficult to resolve. Adding a more sophisticated method for interpolation is a recommendation for further research work, starting with analysis of how different choices about the interpolation contribute to uncertainty in the method, which could be carried out at a simple level in a phase two implementation project.

8.4 Measurement uncertainty

Any study that uses data is constrained by the quality of the data. Where the data are measurements of river flows or sea levels, any assessment of data quality must take into account measurement error. In this project we have sought to minimise this source of uncertainty by only using good records from quality controlled archives and by making basic checks on data consistency in preparation for the analysis. The block bootstrapping method outlined in Section 8.1 will account for observed random measurement error. However, systematic measurement error and bias are assumed to be minimal and are not accounted for separately.

8.5 Non-stationarity

Our work is based on the assumption that the data observed in the past 40 years is representative of what we can expect to observe in the time period over which the results are used to make planning decisions. There are several reasons why this assumption may be invalid to some degree. The first is natural morphological variations such as river channel erosion/deposition and coastal morphological change, including wave action and isostatic movement. A second possible source of non stationarity is climate change. A third is human intervention, such as urban development. Additionally, it is possible that the period of record is simply not representative of current conditions or the recent past owing to decadal scale natural variability.

For coastal flooding there are two features associated with climate change. The first is sea level rise, and the second is a possible change in the frequency distribution of mid-latitude cyclones affecting the UK, and so a possible change in the number of storm surge events. Additionally the distribution of future storm surge events may not be the same as the distribution of current storm surge events. Changes in extreme sea levels have been studied and the uncertainties around projected sea level rises are well documented. Changes in sea levels due to human intervention other than climate change are less likely.

For fluvial flooding estimating the influence of climate change is more complicated. The main reason for this is that river flows are affected by a number of factors including the intensity, frequency and seasonality of precipitation and total evapotranspiration, affecting soil moisture. Land use and vegetation changes (whether natural or anthropogenic) may also have an effect. To model the spatial structure of river flows the full joint distribution structure of these factors, along with the temporal structure, would have to be correctly accounted for. This may require an approach using process based models (such as climate models and rainfall runoff models) to generate proxy observations.

An investigation into the effect of non-stationarity on spatial dependence would be a suitable topic for further research. Some statistical methods for analysis of extremes can be extended to include non stationarity, although this has not yet been done for the methods applied here.

8.6 Pathway and receptor uncertainty

In this project we have focused on estimating the dependence between flood risk source variables at different locations. The reason for this is because the dependence of flood risk between receptors' different locations is caused by the dependence in the source variables.

Our approach to representing the impacts of dependence of flood risk on receptors is to use outputs from previous work. In particular we have shown how it would be possible to integrate our work with the RASP method used in NaFRA. In terms of handling uncertainty this has the advantage that any work done to quantify and reduce uncertainty or otherwise improve those methods would be available to link with the spatial model adopted here.

9 Methodology conclusions

This project is a scoping study to identify, develop and trial a method for assessing flood risk when aggregated over large spatial scales. The overall objective of the project was to develop and test methodologies to assess the risk of widespread flooding. There were also a number of specific tasks within the project; those that we have addressed within this report are as follows:

- To develop a sound theoretical and statistical understanding for assessing the spatial joint distributions and associated likelihood of flooding from single or multiple sources.
- To show how the spatial joint distribution can be extended to include risk pathways and receptors of risk to add a spatial dimension to probabilistic flood risk assessment methods such as RASP.

In developing the statistical method for assessing the spatial joint distribution of flooding we first reviewed a number of different reports of methods that have been or could be used to tackle this problem. In general, these reports fell into three categories. The first category is academic papers and Environment Agency/Defra R&D reports of dependence estimation of flood risk variables (rainfall, sea levels, river flows, and multiple variables). The second category is insurance and reinsurance industry documentation on catastrophe modelling for flood risk. The third category is statistical methods that could be used to estimate the joint distributions of flood risk variables.

In the review we considered each method against a list of requirements that any method to estimate this joint distribution must meet. There was one method that met almost all of these requirements - the Heffernan and Tawn model. In this report we have investigated this method and shown how it is possible to use it to estimate the joint distribution of a large number of flood risk variables. We have also developed this method and shown how it is possible to link it to current RASP procedures. In the accompanying proof of concept report we have demonstrated that it is possible to use this method to develop risk profiles of flooding over a large spatial area.

The main benefits of the method that we have chosen to use in this project are as follows:

- It correctly estimates the spatial structure of flood events at all gauged locations.
- It is easily verifiable against observed data.
- It is modular, so components can easily be improved without disturbing the rest of the model.

The method is best suited to capturing the large-scale spatial structure of flood events. It has the advantage that it is based on observed flow data, hence avoiding uncertainties associated with rainfall-runoff modelling, although this means that the spatial structure of floods between observed locations must be interpolated. The method that we have used here is a simple linear interpolation which is not intended to represent small scale variation in flooding. Another aspect of the flooding process that operates over smaller distances is the effect of defence performance; a defence failure in one location may affect loading conditions elsewhere, and so the spatial dependence structure.

An alternative method that is able to estimate spatial dependence over small scales is continuous simulation, in which simulated rainfall is used as a driver to rainfall-runoff models to generate long time series of river flows. Continuous simulation has the disadvantage that the modelling procedure is much more complicated, but has the advantage that it can be used on completely ungauged catchments, assuming that the runoff model can be parameterised well enough (which is not a trivial matter). In some situations, for instance ungauged headwater catchments, continuous simulation could be able to represent some of the spatial pattern in flood flows by using rain gauge data, which has not been incorporated into the method developed here. It would be more difficult to use continuous simulation to estimate a large scale (for example regional or national) risk profile for flooding due to the sheer size of the model that would need to be run. More fundamentally, such an application would depend of the representation of the spatial dependence and marginal statistics in rainfall at a wide range of space and time scales, as well as the parameterisation and structure of the chosen runoff model. These are all sources of uncertainty.

A beneficial future extension of our model would be to include covariates. It is not yet possible to use additional information to tell us about the spatial aspects of flooding. Possible sources of information that it would be useful to include in a model of the spatial flooding are rainfall data, soil moisture data (antecedent conditions) and large scale atmospheric circulation patterns such as the North Atlantic Oscillation. Inclusion of covariate information within the spatial model would also have the potential to add assessment of the impact of climate change on widespread flooding based on interpretation of precipitation patterns or large scale climatic indices from climate model outputs.

In the meantime, the method described here models flood risk coherently over a wide range scales, at both source and receptor level, to support national risk assessment. Case study demonstrations are given in the accompanying proof of concept summary report, which also describes potential benefits, as identified through stakeholder consultation, that could be realised by applying the technical approach. The final section of this methodology report outlines a programme of phase two work to implement the methods so as to deliver benefits for three such business applications.

10 Phase two recommendations

10.1 Main outputs of the phase one scoping study

Phase one of this R&D project has defined a method for flood risk assessment that accommodates the risk of widespread flooding and is designed to complement existing methods based on RASP already in use by the Environment Agency. The new method has the potential to augment current national scale risk assessments and provides a coherent framework for analysis over a wide range of scales.

The phase one scoping study has produced the following outputs:

- Review of existing work on spatial aspects of flood risk, particularly at regional or national scales.
- Review of relevant statistical methods.
- Detailed statement of a suitable statistical model.
- Three proof of concept demonstration case studies presented at a stakeholder workshop.

These outputs have been delivered by means of this technical methodology report, a proof of concept summary report, an Environment Agency R&D summary, a paper submitted to the *Journal of Flood Risk Management* and several conference papers.

The method that has been demonstrated is compatible with the 'source – pathway – receptor' concept for risk assessment. It is based on a spatially consistent approach to modelling risk at the source level, combined with existing models for the pathway and receptor components.

The phase one scoping study has considered river and coastal flooding and the method is intended to be capable of including other sources of flooding. For practical reasons, other sources of flooding, wave overtopping and climate change were not incorporated in the scoping study.

The proof of concept demonstration studies include analysis of risk at the receptor level and were carried out at a regional scale for practical reasons. A national scale analysis at the source level has also been prepared as an extension of the scoping study.

10.2 Users' requirements

The scoping study has engaged with potential users of the science through initial consultations, a stakeholder workshop and further consultations to define the 'extended' national demonstration. We identified distinct groups of business user needs that are summarised in Table 10-1. The methods demonstrated in this phase one scoping study can meet these needs through a programme of further development and implementation work set out in this section.

Table 10-1. Business user needs, decisions that may benefit from the new methods and relevant technical capacity

Business users	Relevant decision making processes	Technical capacity and access to data
Environment Agency investment planning	Economic risk assessment to support investment planning.	Strong – Good technical understanding of risk assessment methods. Possible customers are NaFRA and MDSF2 projects. Spatially aggregated risk assessment presents opportunities for a new view on risk but will require some new interpretation as well.
Strategic Emergency Response Planners for Flood Risk	Assessing realistic expectations for flood incident response. Planning for realistic flood emergency scenarios.	Capacity to realise benefits from analysis products in support of national scale assessment of resource needs and in preparation of realistic flood exercise scenarios. End product user - unlikely to have requirements or capacity to work with the methods or underlying detailed information without further support.
Defra emergency planning	Planning for realistic flood emergency scenarios.	End product user. Would gain benefits from better understanding of the quantified level of risk associated with planning scenarios.
Cabinet Office	Improved understanding of risk for ‘catastrophic’ type flood events. Analysis of risk to critical infrastructure at multiple locations.	End product user. Strong capacity to incorporate high level information about likelihood of high-consequence event scenarios.
Environment Agency strategic planning	Regional or catchment scale flood risk assessment.	Varying capacity to use methods and access to relevant data. Potential users for detailed information from method to support catchment or regional risk assessment and joint probability analysis for catchment scale river modelling.

10.3 Business applications

We indentified three business applications that can be enabled by applying the methods developed under this scoping study to meet user requirements:

- Business application 1: National and regional flood risk profiles incorporating the risk of widespread river and coastal flooding.
- Business application 2: Probabilistic assessment of emergency planning scenarios.
- Business application 3: Spatial joint probability tools for river basin modelling.

They were identified with the help of feedback gathered at the stakeholder workshop in March 2009 and are described in more detail below. Table 10-2 summarises the three applications and their envisaged benefits.

Table 10-2. Business applications identified following scoping and user consultation

Application	Key outputs	Benefits
National and regional risk profiles	Assessment of the return period of floods (including economic loss or other measures of consequence) occurring anywhere within a region, or anywhere in England and Wales. Analysis of the risk of any individual event or of damages over any given year.	Enhance current NaFRA products. Include quantification of aggregated risk for severe, widespread flood events. Assessment of the resilience of investment decisions and understanding exposure to risk of damaging widespread events.
Emergency planning scenario assessment	Assessment of the risk (probability and consequence) of set spatial patterns of flooding. Generation of plausible flood event scenarios with quantifiable return periods.	Improved understanding of the large scale exposure to flood risk in emergency response planning (for example, how likely are current arrangements to 'fail'?) Information to support strategic thinking about the appropriate deployment and overall level of resources for recovery from flooding. Better quantification of realistic scenarios at different levels of risk for emergency planning exercises.
Spatial joint probability tools for river basin modelling	Methods for setting inflows to catchment models to deal with the joint probability of multiple inflows.	Consistent, scientifically well-founded methodology for assessing joint catchment model inflows.

The applications and their benefits are described in more detail below.

10.3.1 Business application 1: National and regional flood risk profiles incorporating the risk of widespread river and coastal flooding

Purpose and background

Historically, assessment methods used for flood management have had only limited capability to deal with spatially aggregated measures of risk since they focus on single points or local systems rather than considering effects over a wider area.

But widespread floods do occur (such as those in Autumn 2000 and Summer 2007) and can be associated with severe economic and social costs.

Whilst current national flood risk assessment methods can provide information on the expected, long term average economic losses aggregated over the country (or any region), they cannot inform us about the risk of any one, large and damaging event occurring.

This scoping study has shown how recently developed statistical methods can now provide this type of analysis for flood risk, allowing us to assess our exposure to risk from widespread flood events in a way that has not previously been possible. The analysis can be summarised in the form of a 'risk profile', which gives the probability of exceeding a certain loss.

Outcome and business impact

The work would enhance the current NaFRA products to include quantification of aggregated risk for severe, widespread flood events.

This information has the potential to realise benefits for flood risk managers and policy makers in helping assess the resilience of investment decisions and understanding exposure to risk within high level flood management policies.

High level investment planning, policy and strategic management functions will have to adjust to using this new information

Overall objective

To provide a national and regional assessment of the risk of widespread, severe flooding from rivers and coasts.

The project will extend current Environment Agency tools to include more extreme and extensive flood events, and will be based on and hence consistent with NaFRA data.

Specific objectives

To implement the spatially aggregated flood risk model, harmonised with current Environment Agency risk assessment data incorporated in NaFRA, to develop risk profiles for defined administrative regions, including national analysis.

To develop a communication plan to explain the assumptions behind the results and their interpretation.

To explore scenarios for alternative policy options.

Outputs/results and key milestones

0. Inception and project plan.

1. Extraction of required intermediate results from contemporary NaFRA run.

2. Regional and national risk profiles, showing the expected economic losses or other consequences for return periods up to approximately 1000 years.

3. Communication plan to help assimilate the results in investment planning and risk management policy.

Benefits

The work would enhance the current NaFRA products to include quantification of aggregated risk for severe, widespread flood events. It will enable planners and policy makers to understand for the first time how likely we are to suffer the damaging consequences of severe and widespread events such as Autumn 2000 or 1953.

This information has the potential to realise benefits for flood risk managers and policy makers in helping assess the resilience of investment decisions and understanding exposure to risk not just at the level of annual average consequences but also for more extreme, widespread events within high level flood management policies.

10.3.2 Business application 2: Probabilistic assessment of emergency planning scenarios

Purpose and background

The phase one stakeholder workshop identified a number of pertinent flood risk management questions that the new method, developed as part of the Spatial Coherence project, can help to answer. These include:

- What is the probability of two or more critical infrastructure facilities being flooded at the same time within a region?
- What is the maximum number of emergency response resources such as pumps or rescue boats that are likely to be needed in any one flood event?
- What is the best strategy for locating emergency response resources so that they are most likely to be in the right place when floods happen?
- How many flood recovery resources are we likely to need in a 'worst case' flood event within a given planning time horizon?

It was also identified that the methods developed for the Spatial Coherence project could be used to provide extreme scenarios for emergency planning exercises.

Outcome and business impact

The work would improve the Environment Agency's ability to plan for future widespread flood events.

This information has the potential to realise benefits for high-level emergency planning in helping to better understand the implications of the decisions made in planning for widespread flood events.

It will also provide readily available and scientifically well founded reference information to help emergency responders communicate about the relative severity and rarity of future flood events. This may not require any substantial change in how stakeholders operate but would provide an opportunity to provide additional information when communicating about risk.

Overall objective

To provide a suite of data products and guidance for emergency planning and incident management consisting of:

- A library of plausible future widespread flood scenarios and guidance for their use in emergency planning.
- An evaluation of the probability of specific historical events to help guide emergency planners in understanding preparedness for flooding and to give a context for communicating about the severity of future flood events.

To provide an assessment of the resilience of emergency response resources to severe, low probability events using the data.

Specific objectives

To develop a methodology statement describing how the techniques applied for phase one of this project can be used to generate emergency planning scenarios.

To assess requirements for scenarios and derive a library of scenarios with associated probability and consequence measures.

To provide risk profiles in terms of populations at risk or emergency response resource demands (rather than economic damages).

To provide quantification of the severity and likelihood of a number of significant past flood events.

To provide guidance on how to assess the resilience of emergency response resources to severe, widespread river and coastal flooding.

Outputs/results and key milestones

0. Inception and project plan.

1. Consultation with end-users on scenarios to be considered.

2. Methodology statement.

3. Flood scenarios.

4. Guidance for analysing the resilience of emergency response resources to severe, low probability, widespread events.

Benefits

Improved understanding of the large scale exposure to flood risk in emergency response planning – how likely are current arrangements to ‘fail’?

Information to support strategic thinking about the appropriate deployment and overall level of resources for recovery from flooding.

Better quantification of realistic scenarios at different levels of risk for emergency planning exercises.

10.3.3 Business application 3: Spatial joint probability tools for river basin modelling

Purpose and background

In many applications of hydraulic river models it is necessary to model a flood event of specified probability at all points in the river system. For larger catchments or at significant confluences this can raise conceptual and practical difficulties because a single, unique 'T-year' event does not occur in reality at every location in the drainage network. For studies such as a CFMP it can be difficult to know what realistic combinations of inflows should be used, requiring multiple runs, each of which aims to simulate the design condition in part of the river system. One difficulty is that there is no guarantee that hydrographs scaled to match design flows at model inflow points will result in the preferred design flows being reproduced further downstream within the model.

There are other applications which require simulation of a realistic flood event throughout the river system. These include model calibration and simulation of events for flood warning studies. The problems boil down to setting combinations of inflows at confluences to result in the required design condition downstream.

The methods used in phase one of this project offer a flexible method for representing the joint probability of multiple river reaches that can help with this practical problem of setting inflows for catchment modelling. The same approach has already been adopted in the Flood Studies Update Guidance for River Basin Modelling (OPW, 2008).

Outcome and business impact

An improved method for catchment-wide modelling that will require less time to calibrate than current procedures. This has the potential benefit of improving the approach used for applying hydraulic river or routing models in catchment scale studies such as CFMPs. It would also allow results that are consistent with broader scale analysis (if the phase one methods are taken forward into broad scale assessments such as NaFRA, for example).

River modellers in the Environment Agency and consultants would have to adapt to revised guidance that incorporated the new methods. Tools or datasets would need to be published and maintained.

Overall objective

To provide a methodology for representing the joint probability of multiple river inflows in catchment modelling.

To extend current tools for catchment-scale river modelling.

Specific objectives

To assess requirements for joint probability analysis of catchment model inflows and where the methods used in phase one can help. This would draw on experiences and lessons learned from work for the OPW, Ireland. It would also draw upon developments applying the same principles as the phase one spatial joint probability method that are underway in FRMRC2, FRACAS and that have now been applied in SFRM2 modelling studies.

To set out a technical approach to meet these requirements.

To develop datasets, guidelines and optionally software tools to support catchment modelling.

Outputs/results and key milestones

0. Inception and project plan

1. Research report giving analysis and classification of catchment model inflow joint probability problems.

2. Methodology, data or prototype tool set to provide guidance on setting model inflows.

3. Report and guidance on how to use the new method.

4. Specification for implementation in a software tool

Benefits

Consistent, scientifically well founded methodology for assessing catchment model inflows.

Avoidance of cost of unnecessary or inappropriate model inflow scenarios in catchment modelling studies.

Consistency with national scale analysis (if taken forward).

10.4 Tasks required to deliver phase two business applications

Table 10-3 summarises the work required to implement the proof of concept methods from this scoping study in order to deliver the business applications identified above.

The work is presented in Table 10-3 as a modular set of tasks that are listed in three categories as follows:

Foundation tasks	Essential work necessary to underpin the business application and to deliver initial results. This includes the basic statistical model development, generation of results for the 'source' level; that is, river flows or storm surge, and consultation and training actions. These 'foundation' tasks alone are sufficient to deliver useful products, but not the full range of outputs to meet the user requirements identified in this scoping study.
Business analysis tasks	Work needed to deliver the full scope of each of the identified business applications, including analysis of economic risk, linkages with RASP and NaFRA data and supporting analysis tools, where relevant. These work items are not 'extras' but necessary elements to meet fully the user needs identified in this scoping study and hence realise the potential benefits of the applications.
Enhancements	Additional work that would significantly enhance the business applications or add to knowledge, involving some additional scoping and technical development work and also new sources of data. These work items would enhance the proposed applications beyond the scope initially identified, for example by including surface water flooding or research seeking to link extreme flood scenarios to larger scale meteorological indices.

Each task is described in detail in the following sections, followed by an indicative programme outlining how the tasks could be built into a 23-month work plan to deliver against each of the three identified applications.

Table 10-3. Summary of modular work items required to deliver the identified business applications.

Application	Foundation tasks	Business analysis tasks	Enhancements
National and regional risk profiles	F1. Consultation and communications. F2. Conditional exceedance model – rivers and coasts. F3. Event scenario library. F4. Probability analysis of river flows and sea level extremes. F5. Training and handover.	BA1. NaFRA conditional economic damages analysis. BA2. National/regional river and coastal economic risk profile analysis. BA3. Economic damages comparison checks and validation. BA4. Event database analysis tools.	E1. Scope and extend conditional exceedance model with additional source variables. E2. Surface water flooding economic damages analysis. E3. Covariate analysis. E4. Climate change modifications. E5. River network interpolation enhancements.
Emergency planning scenarios	F1. Consultation and communications. F2. Conditional exceedance model – rivers and coasts. F3. Event scenario library. F4. Probability analysis of river flows and sea level extremes. F5. Training and handover.	BA5. Emergency planning threshold levels analysis. BA6. Emergency planning scenario timing analysis.	E1. Scope and extend conditional exceedance model with additional source variables. E3. Covariate analysis. E4. Climate change modifications.
Spatial joint probability tools for river basin modelling	F1. Consultation and communications. F2. Conditional exceedance model – rivers and coasts. F3. Event scenario library. F5. Training and handover. F6. River model inflows joint probability database.	BA7. River model inflows joint probability tool.	E4. Climate change modifications.

10.4.1 Foundation tasks

F1. Consultation and communications	
What	Consultation with business users to ensure suitable definition and presentation of outputs and to plan any steps needed to prepare users to receive and apply the results. Specification of required event definitions.
Why	Ensures relevance and uptake of the products so as to realise benefits in full.
How	This task should take place with early consultation activities to confirm user needs. It should identify key business customer representatives who can work with the project to define and agree specific points of detail about issues such as choices of spatial data formats, geographical units for analysis, formats for presentation of results, choice of database file formats and delivery of software routines. Liaison with CIS if appropriate. The task should include early planning for and subsequent delivery of dissemination activities (such as workshops, seminars), publishing and coordination between business users,R&D programme and consultants. This task should identify and plan for any immediate and on-going training needs to complete hand over of products generated by the project. It should identify any relevant future updates or developments.
Depends on	-
Application	National and regional risk profiles. Emergency planning scenarios. Joint probability boundary conditions for river models.
Data required	-
Outcome	Detailed specification of formats for presentation of results, data/databases and software procedures. Communications and training plan. ***An essential outcome of this task is to specify the flood event definitions (in terms of durations and temporal resolution) that will be analysed during the construction of the conditional exceedances model (task F2) and the event simulation task (F3)***.

F2. Conditional exceedance model – rivers and coasts	
What	Fitting conditional exceedance model for the joint distribution of extremes in gauged river flows and tide gauge data.

Why	Provides the basic statistical model of spatial dependence in extremes.
How	Approach as set out and tested in the phase one scoping study technical methodology report and in this report.
Depends on	-
Application	National and regional risk profiles. Emergency planning scenarios. Joint probability boundary conditions for river models.
Data required	Daily or sub-daily river flow gauge data. Sub daily tide gauge data.
Outcome	Implemented statistical model and parameters.

F3. Event scenario library	
What	Simulation of flood events from conditional exceedance model.
Why	Primary output from statistical model for river flows and storm surge data.
How	Approach as set out and tested in the phase one scoping study technical methodology report and in this report.
Depends on	Conditional exceedance model – rivers and coasts.
Application	National and regional risk profiles. Emergency planning scenarios.
Data required	Conditional exceedance model.
Outcome	National database of simulated flood events expressed on a physical or probability scale for river and coastal model node points. Database documentation.

F4. Probability analysis for river flows and sea levels	
What	Statistical analysis of event scenario library to provide estimates of probabilities for events or summary measures of flood severity at the ‘source’ level (that is, river flows and sea levels, not accounting for flood defence system performance or economic damages).
Why	Primary analysis of outputs from the statistical model for river flows and storm surge data at the ‘source’ level.
How	Approach as set out and tested in the phase one scoping study technical methodology report, in this report and in subsequent extensions to phase one.
Depends on	Conditional exceedance model – rivers and coasts. Event scenario library.
Application	National and regional risk profiles.

	Emergency planning scenarios.
Data required	Conditional exceedance model event scenario library outputs.
Outcome	Analysis of event scenario library at ‘source’ level giving return period for different levels of flood severity.

F5. Training and handover	
What	Formal handover of final products, tools and data. Training, if necessary, for specialist users.
Why	To ensure full realisation of benefits from each of the phase two projects.
How	Publication (internal or external) of final results, preparation of briefing notes, guidance, presentations and so on, Preparation of training material and delivery of initial training.
Depends on	Project outputs.
Application	National and regional risk profiles. Emergency planning scenarios. Joint probability boundary conditions for river models.
Data required	Project outputs.
Outcome	Take-up and full utilisation of project outputs.

F6. River model inflows joint probability database	
What	Database of joint probability relationships for setting boundary conditions to catchment scale river models.
Why	Primary resource to assist with setting multiple inflows in catchment river models so as to provide a consistent basis for resolving joint probabilities in catchment scale modelling.
How	Catchment based analysis of events simulated from conditional exceedance model. This task should create a spatial dataset giving the dependence between tributary inflows for selected points on the river network and a library of suitable events contributing to a given downstream return period.
Depends on	Conditional exceedance model – rivers and coasts.
Application	Joint probability boundary conditions for river models.
Data required	Conditional exceedance model. River centre line and logical river network. Catchment boundaries.
Outcome	Database and map outputs showing joint probability of river flows corresponding to target return periods at specified locations.

10.4.2 Business analysis tasks

BA1. NaFRA conditional economic damages analysis	
What	Create database of conditional distributions of flood depths and consequence (economic damages) based on RASP calculations embedded in NaFRA.
Why	Provides pathway and receptor information for economic risk analysis, links with NaFRA.
How	Modification of NaFRA procedures to derive conditional damages data as defined in the phase one technical methodology report and in this report. This task needs to include a detailed specification for the precise format of the derived data, for example whether it is delivered in the form of raw RASP simulation outputs, histograms or summary statistics (mean and variance of economic damage per depth interval). Requires close coordination with NaFRA programme and technical teams.
Depends on	Close coordination with national/regional river and coastal economic risk profile analysis task so as to ensure consistent definition of data formats and agreement on methodology. Close co-ordination with NaFRA programme.
Application	National and regional risk profiles.
Data required	NaFRA inputs, spatial grid to define resolution at which outputs are produced.
Outcome	Database documentation. Technical report detailing algorithm and implementation details, validation checks and summary statistics. Documented code and validation checks.

BA2. National/regional river and coastal economic risk profile analysis	
What	Integrate event scenarios generated from the conditional exceedance model with conditional economic damages data to develop economic risk profiles at regional and national level.
Why	Key business user output.
How	Monte Carlo simulation approach as set out and tested in the phase one scoping study technical methodology report and in this report.
Depends on	Conditional exceedance model. NaFRA conditional economic damages analysis.
Application	National and regional risk profiles.
Data required	Outputs from conditional exceedance model and NaFRA conditional economic damages analysis.
Outcome	Results giving estimate of return period of differing levels of economic damage from flooding for rivers and coasts at a range of spatial

	scales.

BA3. Economic damages comparison checks and validation	
What	Comparison with NaFRA results and other economic damages data to provide higher-level validation of results.
Why	To help explain the relationship between the spatial risk profile results and other broad scale risk assessments. To provide high-level 'sense checks'.
How?	Comparison of annual scale results with NaFRA expected annual damages. There is a theoretical correspondence, for a given geographical area, between the annual average damage (EAD) calculated by the RASP methodology in NaFRA and the mean damage on the annual scale calculated using the methodology developed in this scoping study. This correspondence should be checked to ensure that the approach to extracting economic damages data from the NaFRA process and linking it to the conditional exceedance model outputs has been correctly implemented. Comparison of extreme event scenarios with existing estimates of economic consequences from historical events. There are estimates available for economic damages for some notable past floods. These should be compared with the outputs of the risk profile analysis and any discrepancies investigated. It is noted that estimates of economic damages from historical events are not necessarily detailed or complete.
Depends on	National/regional river and coastal economic risk profile analysis.
Application	National and regional risk profiles.
Data required	NaFRA economic analysis results. Economic damage analysis from past flood events.
Outcome	Comparison of expected annual damages with NaFRA results and recorded event damages. Comparisons of dependence in losses between different sub regions. Sense checks on risk profile outputs. Validation report.

BA4. Event database analysis tools	
What	Code and software tools to allow re-running of economic risk profile for different damages data, event scenarios or spatial regions.
Why	To provide greater flexibility and capacity to update results in future or explore alternative investment planning scenarios.
How	Documentation and technical user guidance on the algorithms and software procedures used in analysis of the event scenario library produced using the conditional exceedance model.
Depends on	National/regional river and coastal economic risk profile analysis.

Application	National and regional risk profiles.
Data required	NaFRA economic analysis results. Economic damage analysis from past flood events.
Outcome	Algorithm specification, software procedures and guidance. This may be in the form of software routines and scripts that are geared towards a technical specialist user, rather than a fully featured end-user application with a sophisticated front end graphical user interface.

BA5. Emergency planning threshold levels analysis	
What	Analysis of event scenario probability based on defining events in terms of river flows and sea levels exceeding set thresholds at specified locations.
Why	Provides a rational event definition for emergency planning.
How	Identify suitable locations for use in event definition, such as flood warning areas. Specify suitable thresholds to use in characterising the occurrence of flooding at those locations, such as river flow return period, water level. Perform return period analysis on scenarios in terms of numbers of threshold exceedances.
Depends on	Event scenario library and return period analysis.
Application	Emergency planning scenarios.
Data required	Spatial data to define flood locations. Threshold exceedance levels.
Outcome	Comparison of expected annual damages with NaFRA results and recorded event damages. Comparisons of dependence in losses between different sub regions. Sense checks on risk profile outputs. Validation report.

BA6. Emergency planning scenario timing analysis	
What	Additional analysis of flood event scenarios to consider the probability of widespread events for different scenarios of event timing, for example two regions experiencing extreme flooding within one month, or successive months.
Why	User requirement for emergency planning.
How	Application of conditional exceedance model for event definitions that include specified patterns in temporal lag.
Depends on	Event scenario library and return period analysis.

Application	Emergency planning scenarios.
Data required	Event definitions to be agreed with business users.
Outcome	Comparison of expected annual damages with NaFRA results and recorded event damages. Comparisons of dependence in losses between different sub regions. Sense checks on risk profile outputs. Validation report.

BA7. River model inflows joint probability tool	
What	Software and database tool to package results of joint probability analysis for ease of use by river modelling practitioners.
Why	Promote update and benefits realisation.
How	Define and develop a tool to assist river modellers in setting boundary conditions with joint probabilities. May be a combination of software application, database, web-based data, GIS layers, maps and guidance notes.
Depends on	River model inflows joint probability database.
Application	Joint probability boundary conditions for river models.
Data required	River model inflows joint probability database. Background mapping, logical river network.
Outcome	Comparison of expected annual damages with NaFRA results and recorded event damages. Comparisons of dependence in losses between different sub regions. Sense checks on risk profile outputs. Validation report.

10.4.3 Enhancements

E1. Scope and extend conditional exceedance model with additional source variables	
What	Development work to extend the conditional exceedance model set out in this scoping study to include additional variables, which may include: • Rainfall (to allow inclusion of surface water flood risk). • Wave overtopping (to allow enhancement of coastal risk analysis). • Wind (to enhance emergency planning scenario analysis). This task should comprise initial scoping and feasibility analysis followed by implementation.
Why	Provide extension of risk assessment to include surface water flooding, wave overtopping and wind storm occurrence for emergency

	planning scenarios. These features have been identified by business users and technical reviewers as worthwhile extensions of the scoping study methods. The phase one methodology uses a statistical model of the joint probability distribution of storm surge or river flows at gauging stations. The approach would also be suitable for modelling the joint distribution of rainfall based on gauged data. This is required to represent the 'source' component for surface water flood risk. There are a number of rainfall models already available as research outputs or operational tools for applications such as continuous rainfall-runoff modelling and climate change impacts analysis. These models tend to be based on point process approaches, combined using a cellular structure that represents some aspects of meteorological patterns.
How	Application of the conditional exceedance model to suitable data sets describing rainfall and wind speed or storm occurrence. Extension of sea level simulation to include wave overtopping. Waves are a main determinant of coastal structural failure, though they may often be depth-limited at the coast. As they are driven by storm conditions, wave heights do show spatial dependence over scales of interest. Good archives of wave data are available on a 25km grid from the UKMO wave model. Future research should be directed at understanding how best to include rainfall in the risk modelling methodology proposed here, considering both the statistical modelling approach and the practical issues of working with rain gauge or other rainfall information. Relevant questions include: • What type of rainfall data should be used? For example, gauges, gridded averages, radar, accumulations. • How should the time variability of rainfall be represented? • Can the modelling approach adopted in phase one be combined usefully with existing point process stochastic rainfall models? At what scales?
Depends on	
Application	National and regional risk profiles. Emergency planning scenarios.
Data required	Analysis of surface water flood requires rain gauge or gridded rainfall data for the source model and access to information on receptor impacts including national surface water flood map data. An atlas of swell wave heights has been developed as part of current Defra R&D on coastal extremes. Characterising large scale patterns in wind waves is likely to involve analysis of Met Office hindcast deep water wave model outputs (which may incur licensing costs) or modelled and reanalysis wind field data, combined with near shore wave transformations. Met Office of ECMWF wind storm data.
Outcome	Tested method and model for extension of conditional exceedance model to additional data sets at source. Application of model, linked to river and costal data.

E2. Surface water flooding economic damages analysis	
What	Development work to provide estimates of economic consequences for surface water flooding to be combined with conditional exceedance model. This task should comprise initial scoping and feasibility analysis followed by implementation. It will link to current Environment Agency R&D projects developing improved data for surface water flood risk assessment.
Why	Provide extension of risk assessment to include economic impacts of surface water flooding.
How	Analysis of surface water flood mapping combined with property data and economic damages calculations to estimate the distribution of damages expected from surface water flood events. Integration with simulations of the occurrence of surface water flooding as part of the extended conditional exceedance model.
Depends on	Scope and extend conditional exceedance model with additional source variables.
Application	National and regional risk profiles.
Data required	National surface water flood mapping. National Property Database.
Outcome	Database of event scenario library for estimated economic consequences of surface water flooding.

E3. Covariate analysis	
What	Technical development of the statistical methodology set out in this scoping study to build covariates into the conditional exceedance model and event simulation approach.
Why	The current method for modelling the spatially aggregated risk of flooding relies on the idea that river flow gauges capture most of the important statistical variability in flood flows. However, there is other information about the occurrence and severity of flooding that could be helpful in refining the accuracy of risk estimates, for example large scale weather patterns. Saturated soils were cited by Pitt (Cabinet Office, 2008) as being one of the contributing factors to the severity and extent of the 2007 floods, and also in the Bye review of the 1998 floods (Bye and Horner, 1998). Antecedent coastal storms can draw down beaches making them more vulnerable to subsequent storms. This task should test the significance of antecedent conditions to see whether they merit inclusion in future work, without necessarily embarking down the road of continuous simulation. For fluvial flooding a start would be to consider soil moisture as a covariate in the model, with the soil moisture modelled as dependent between events.

How	Further research is needed to gain an understanding of how this kind of information could be used and what difference it would make to a flood risk model, especially in the light of other uncertainties. In principle it may be possible to incorporate data about covariates directly into the joint distribution within our model. Alternatively, it could be beneficial to use covariate data to construct some weighting scheme to 'balance' the sample of events generated within a Monte Carlo simulation to ensure the correct proportion of different types of event. This type of approach may also allow future climate risk to be handled efficiently.
Depends on	
Application	National and regional risk profiles. Emergency planning scenarios.
Data required	Meteorological indices and reanalysis data.
Outcome	Scientific basis for inclusion of covariate information in the statistical model for joint distribution of flood risk.

E4. Climate change modifications	
What	Technical scoping and development to incorporate climate change.
Why	To allow for climate scenario analysis and relax assumptions of stationarity made in the current modelling. The phase one methodology seeks to avoid the uncertainties associated with models of rainfall-runoff processes by working directly with river flow data. Similarly, sea levels have been modelled using tide gauge data. One constraint in this approach is the assumption that the gauged record is a sample from a stationary process. This has two important implications for the method. It means that: (1) any non stationarity in the data must either be corrected or modelled, or it will inflate the uncertainty about the model, and (2) the model does not represent 'scenarios' that could be inconsistent with the observed joint distribution of extremes in the gauged data.
How	The issue of non stationarity has not been addressed in detail within the phase one methodology, although data used to fit the model are checked for obvious trends or step changes. Further work is recommended to look at how to incorporate non stationarity in the conditional exceedance model. There should be an assessment of the scope to apply the current approach using perturbed data sets, such as river flows developed in Defra project FD2020. A second requirement is to deal with scenarios such as changes in the source variables or in risk pathways. At the source level, climate change is of particular importance for rivers and coasts. Other types of change that affect the river flow regime (such as urbanisation) should also be considered. At the pathway level, there is a need to understand how to incorporate scenarios about changes in defence systems within the methodology.
Depends on	

Application	National and regional risk profiles. Emergency planning scenarios. Joint probability boundary conditions for river models.
Data required	Climate change scenario/impact data for river flows, sea levels and storm surge.
Outcome	Modified analysis results to include climate change scenarios.

E5. River network interpolation enhancements	
What	Improvement to the methods used in the phase one scoping study for interpolation of events at gauging stations through the river network.
Why	To provide improved quality in results.
How	Develop improved procedure for river network interpolation to include greater variability in the interpolation. The interpolation procedures used to transfer information from gauged to ungauged reaches are currently deterministic, which places some constraints on the hydraulic load that can be simulated at any given point. An obvious improvement would be to allow for some stochastic variation in the interpolation. It may be possible to build this in to the current methods with the addition of some constraints based on physical considerations about flow routing. The interpolation approach is likely to be more uncertain for rivers that have no gauging or are geographically far away from any other gauged river. These sites are generally on smaller watercourses or in the headwaters of larger catchments. Here, it is also possible that including additional information based on a rainfall covariate could improve the model of dependence. This work is highly recommended as it could provide a relatively cost effective technical improvement in accuracy of the economic analysis for national and regional risk profiles.
Depends on	-
Application	National and regional risk profiles.
Data required	River flow data at gauging stations.
Outcome	Technical report detailing methodology improvements, tests and implementation. Modified source code to implement the improved interpolation method.

10.5 Options

The tasks set out above provide the building blocks for implementation of the scoping study methodology for each of the business applications. The individual tasks could be combined in several ways. Four principal options have been identified, as follows. The cost and programme implications are indicated in terms of the seniority of personnel required and a realistic programme for delivery (assuming staff inputs may be at less than 100 per cent full time equivalent).

10.5.1 Option 1 – 'Do nothing'

No further R&D or implementation work.

Cost implications

Missed opportunity to realise business benefits identified during the scoping study. No results linking to economic or other measures of risk at the 'receptor' level. No link to flood defence performance models and no specific event scenario analysis. No capability to apply or adapt the scoping study methodology for further analysis.

Benefits

Scoping study outputs adds to scientific knowledge, which could be taken forward into business applications in future.

National analysis of probability of widespread high river flows or storm surge carried out in a short case study extension to the scoping study

10.5.2 Option 2 – Foundation tasks

Cost/programme implications

Up to approximately 12-month programme to deliver the common foundation tasks conditional exceedance model (F2) and event scenario library database (F3). Senior scientist level staff inputs.

Approximately four-month programme to deliver analysis of probabilities at 'source' level (F4). Senior scientist level.

Concurrent consultation activities to define precise user needs (such as choice of event definition). Senior scientist and technical director inputs.

Training or roll out of results – programme dependent on precise definition of outputs and user groups. Senior scientist and technical director inputs. Combination of briefing notes and workshop presentations.

Lack of business analysis to deliver full economic risk assessment product, event emergency planning scenarios or tool for catchment model inflows.

Additional benefits

Greater knowledge gain through creation of a national model and event library/database for Environment Agency use.

Analysis outputs for 'source variables' (river flow and sea level) at national scale tailored to each specific business application, greatly extending the demonstration outputs from the scoping study.

10.5.3 Option 3 – Business analysis

Cost/programme implications

Up to approximately 12-month programme to deliver the economic damages data from RASP/NaFRA (BA1). Senior scientist and graduate level staff inputs.

Approximately six- to eight-month programme to deliver economic risk analysis, validation checks and economic analysis tools. Technical director, senior scientist and graduate level.

Approximately six- to eight-month programme to deliver analysis tool for setting catchment inflows. Senior scientist and graduate programmer.

Additional benefits

Outputs designed to fully meet user requirements identified by this scoping study, including economic risk assessment, detailed emergency planning scenarios and catchment modelling inflows tool.

Capability for Environment Agency to repeat and extend analysis in house or through consultants using databases and tools delivered under this option.

10.5.4 Option 4 – Enhancements

Cost/programme implications

Between 3- and 18-month programme to deliver enhancements identified above, depend on choice of enhancement tasks. Technical director, senior scientist and graduate level staff inputs.

Additional benefits

Range from 'quick win' improvements to some aspects of methodology to significant extensions to include climate change and surface water flood risk.

10.6 Indicative phase two work plan

Figure 10-1 shows an indicative programme for phase two projects to deliver all three business applications. Milestones are indicated. Should more than one application be developed then there will be significant overlap between tasks, particularly the foundation tasks and enhancements, but it should be noted that the emphasis and precise definition of those tasks might vary slightly, depending on the intended application (for example, the precise choice of users to be consulted in task F1: *Consultation and communication* and the flood event definitions specified for use in task F2: *Conditional exceedance model* may vary according to the intended application).

Figure 10-1. Business user needs, decisions that may benefit from the new methods and relevant technical capacity

	Month																					
National and regional risk profiles																						
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**
Foundation tasks																							
F1. Consultation and communications.	■	■	■			■	■				■	■					■	■					
F2. Conditional exceedance model - rivers and coasts.	■	■	■	■	■	■																	
F3. Event scenario library.						■	■	■	■	■	■												
F4. Return period analysis (river flows and sea level).											■	■	■	■									
F5. Training and handover.																		■	■	■	■	■	■
Business analysis tasks																							
BA1. NaFRA conditional economic damages analysis.	■	■	■	■	■	■	■	■	■	■	■	■											
BA2. National/regional river and coastal economic risk profile analysis.												■	■	■	■	■							
BA3. Economic damages comparison checks and validation.																■	■	■					
BA4. Event database analysis tools.									■	■	■	■	■	■									
BA5. BAiver network interpolation enhancements.	■	■	■	■	■	■																	
Enhancements																							
E1. Scope and extend conditional exceedance model with additional source variables:	■	■	■	■	■	■	■	■	■	■	■												
E2. Surface water flooding economic damages analysis.	■	■	■	■	■	■	■	■	■	■	■	■											
E3. Covariate analysis.	■	■	■	■	■	■	■	■	■	■	■	■											
E4. Climate change modifications	■	■	■	■	■	■	■	■	■	■	■												
Milestones																							
MS1 Inception			x																				
MS2 Conditional exceedance model						x																	
MS3 Event scenario library											x												
MS4 Extraction of economic damages data from NaFRA (a) progress, (b) completion						x						x											
MS5 Return period analysis (source level)														x									
MS6 Risk profiles (receptor level)																		x					
MS7 Training and hand over delivered																							x
Emergency planning scenarios																							
Foundation tasks	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**
F1. Fonsultation and communications.	■	■	■			■	■				■	■			■	■							
F2. Fonditional exceedance model - rivers and coasts.	■	■	■	■	■	■																	
F3. Event scenario library.						■	■	■	■	■	■												
F4. Return period analysis (river flows and sea level).											■	■	■	■	■								
F5. Training and handover.															■	■	■	■	■	■			
Business analysis tasks																							
BA6. Emergency planning threshold levels analysis.										■	■	■	■	■	■								
BA7. Emergency planning scenario timing analysis				■	■	■				■	■	■	■	■	■								
Enhancements																							
E1. Scope and extend conditional exceedance model with additional source variables:	■	■	■	■	■	■	■	■	■	■	■												
E3. Covariate analysis.	■	■	■	■	■	■	■	■	■	■	■	■											
E4. Climate change modifications	■	■	■	■	■	■	■	■	■	■	■												
Milestones																							
MS1 Inception			x																				
MS2 Conditional exceedance model						x																	
MS3 Event scenario library											x												
MS4 Emergency planning scenarios produced															x								
MS5 Training and hand over delivered																				x			
Spatial joint probability tools for river basin modelling																							
Foundation tasks	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**
F1. Fonsultation and communications.	■	■	■			■	■				■	■			■	■							
F2. Fonditional exceedance model - rivers and coasts.	■	■	■	■	■	■																	
F3. Event scenario library.						■	■	■	■	■	■												
F6. River model inflows joint probability database											■	■	■	■	■	■							
F5. Training and handover.																■	■	■	■	■	■		
Business analysis tasks																							
BA8. BAiver model inflows joint probability tool.									■	■	■	■	■	■	■	■							
Enhancements																							
E4. Climate change modifications	■	■	■	■	■	■	■	■	■	■	■												
Milestones																							
MS1 Inception			x																				
MS2 Conditional exceedance model						x																	
MS3 Event scenario library											x												
MS4 Joint probability database																x							
MS5 Joint probability boundary conditions tool																x							
MS6 Training and hand over delivered																					x		

The indicative programme is designed to allow for production of a new conditional exceedance model to the exact specification required by the Environment Agency applications, along with time to investigate optional extensions to the model required to handle rainfall, wave overtopping, wind and climate change. The programme allows a period of 11 months to build these extensions into the statistical model and simulated flood event library. The details of work undertaken within these options would need to be specified with a pragmatic view to what can reasonably be achieved within the given time frame (for example by consideration of how many rain gauges are proposed for inclusion in analysis of rainfall, or whether it is feasible to include wind data to augment emergency planning scenarios).

The potential programme for extraction of economic damages data from the NaFRA process is not known with certainty at this time. The indicative programme allows 12 months for this analysis, but the actual timing of a project would need to be set in conjunction with the NaFRA technical programme. There may be scope to stagger the delivery of conditional economic damages data, working on a regional basis, for example. The methods proposed and tested in this scoping study are capable of accepting updates to conditional damages data in a phased manner, although for a national analysis it is of course necessary to have a national economic damages data set in place. This is also a consideration for local adjustments to NaFRA or future use of MDSF2 in generating NaFRA data. Future local modifications to the economic damages data could be incorporated in the risk profile analysis without having to repeat a full national NaFRA run.

11 References

ANCONA-NAVARRETE, M. A. AND TAWN, J. A. (2002). Diagnostics for pairwise extremal dependence in spatial processes. Extremes, 5, 271–285.

BORTOT, P., COLES, S. G. AND TAWN, J. A. (2000). The multivariate Gaussian tail model: an application to oceanographic data, Appl. Statist., 49, 31-49.

BUISHAND, T.A., (1984). Bivariate extreme-value data and the station-year method. *Journal of Hydrology*, 69(1-4), 77-95.

BUISHAND, T.A., DE HAAN, L. AND ZHOU, C. (2008) On spatial extremes: with application to a rainfall problem, *Ann. Applied Statistics. 2(2), 624-642*

BUTLER, A. HEFFERNAN J. E., TAWN, J. A., FLATHER, R. A. AND HORSBURGH, K. (2007). Extreme value analysis of decadal variations in storm surge elevations. Journal of Marine Systems, 67, 189-200.

BYE, P. AND HORNER, M. (1998) Easter 1998 Floods. Volume I: Findings Sections. Report by the Independent Review Team to the Board of the Environment Agency, Bristol.

CABINET OFFICE (2008) The Pitt Review - Learning Lessons from the 2007 floods, Cabinet Office, 22 Whitehall, London SW1A 2WH.

CHANDLER, R. E. (2002). GLIMCLIM: Generalized linear modelling for daily climate time series (software and user guide). Technical report, no. 227, Department of Statistical Science, University College London, London, WC1E 6BT.

COLES, S. G. AND TAWN, J. A. (1990). Statistics of coastal flood prevention, Phil. Trans Roy. Soc. (Lond), A, 332, 457–476.

COLES, S. G., HEFFERNAN, J. E. AND TAWN, J. A. (1999). Dependence measures for extreme value analyses. Extremes, 2, 339–365.

COLES, S. G. AND TAWN, J. A. (1994). Statistical methods for multivariate extremes: an application to structural design (with discussion), Appl. Statist., (1994), 43, 1-48.

COLES, S. G. AND TAWN, J. A. (1996) Modelling Extremes of the Areal Rainfall Process, *J. R. Statist. Soc. B*, **58**, 329-347.

DAVISON, A.C. AND HINCKLEY, D.V. (1997) Bootstrap Methods and their Application. Cambridge University Press.

DAVISON, A. C., AND SMITH, R. L. (1990) Model for exceedances over High Threshold, *J. R. Statist. Soc. B*, **52**, 393-442.

DAWSON, R.J. and HALL, J.W. (2006) Adaptive importance sampling for risk analysis of complex infrastructure systems, Proc. Royal Soc. A: Mathematical, Physical and Engineering Sciences, 462(2075) 3343–3362.

DEPARTMENT FOR ENVIRONMENT FOOD AND RURAL AFFAIRS/ENVIRONMENT AGENCY (2005a) Joint Probability: Dependence Mapping and Best Practice: Technical report on dependence mapping R&D Technical Report FD2308/TR1.

DIXON, M. J. AND TAWN, J. A. (1994). Extreme sea-levels at the UK A-class sites: site-by-site analyses, Proudman Oceanographic Laboratory, Internal document No. 65.

DIXON, M. J. AND TAWN, J. A. (1997). Spatial analyses for the UK, Proudman Oceanographic Laboratory report, 112, 200.

DIXON, M. J. AND TAWN, J. A. (1999). The effect of non-stationarity on extreme sea-level estimation, Appl. Statist., 48, 135–151.

DIXON, J. M., TAWN, J. A. AND VASSIE, J. M. (1998). Spatial modelling of extreme sea-levels, Environmetrics, 9, 283–301.

DONNER, R. (In Press) Spatial correlations of river runoffs in a catchment. In KROPP, J. and SCHNELLNHUBER, H-J. (Eds.), *In Extremis: Extremes, Trends and Correlations in Hydrology and Climate,* Springer, Berlin. Pre-print available from http://www.agnld.uni-potsdam.de/~reik/Donner_ScalingBook.pdf [accessed August 2009].

ENVIRONMENT AGENCY (2004) Risk Assessment of Flood and Coastal Defence Systems for Strategic Planning (RASP). HR Wallingford Report for the Environment Agency - R&D Technical Report W5B-030/TR.

ENVIRONMENT AGENCY (2007) Scoping the development and implementation of flood and coastal RASP models, SC050065/SR1.

GOULDBY, B., SAYERS, P., MULET-MARTI, J., HASSAN, M.A.A.M and BENWELL, D. (2008) A methodology for regional-scale flood risk assessment, Proc. Inst. Civil Engrs. (Water Management), WM0, 1-14.

HALL, J.W., DAWSON, R.J., SAYERS, P.B., ROSU, C., CHATTERTON, J.B. and DEAKIN, R. (2003) A methodology for national-scale flood risk assessment, Proc. Inst. Civil Engrs. (Water & Maritime Engineering), 156(3), 235-247.

HEFFERNAN, J. E. AND TAWN, J. A. (2004). A conditional approach to modelling multivariate extreme values (with discussion) J. Roy. Statist. Soc., B, 66, 497–547.

HORSBURGH, k. J., and WILSON, C. (2007), Tide-surge interaction and its role in the distribution of surge residuals in the North Sea, *J. Geophys. Res.*, *112*, C08003, doi:10.1029/2006JC004033

JOE, H. (1997). Multivariate Models and Dependence Concepts. Chapman & Hall, London.

KEEF, C. (2007) Spatial dependence of river flooding and extreme rainfall PhD Thesis Lancaster University

KEEF, C., TAWN, J. A. AND SVENSSON, C., (2009a). Spatial risk assessment for extreme river flows. Appl, Statist. 58, 601-618

KEEF, C., SVENSSON, C. and TAWN, J.A., (2009b) Spatial dependence of extreme river flows and precipitation in Britain, J. Hydrology. 378, 240-252

KJELDSEN, T.R. and JONES, D.A. (2007) Estimation of an index flood using data transfer in the UK. Hydrological Sciences Journal, 52(1), 86-98.

LATHAM AND TAWN (2007) A dimension reduction model for spatial extreme sea-level events, Submitted to Applied Statistics.

LEDFORD, A. W. AND TAWN, J. A. (1996). Statistics for near independence in multivariate extreme values Biometrika, 83, 169–187.

LEDFORD, A. W. AND TAWN , J. A., 1998. Concomitant tail behaviour for extremes. Adv,. Appl. Prob., 30, 197-215.

MARSH, T. J., and HANNAFORD, J., (2008). UK Hydrometric Register. Hydrological Data UK Series. Centre for Ecology and Hydrology, Wallingford.

MUIR WOOD, R., DRAYTON, M., BERGER, A. & BURGESS, P. (2005) Catastrophe loss modelling of storm surge flood risk in eastern England. Phil. Trans. R. Soc. A. Volume 363, Number 1831, 1407-1422.

NATIONAL RESEARCH COUNCIL, (2000) Risk Analysis and Uncertainty in Flood Damage Reduction Studies, Committee on Risk-Based Analysis for Flood Damage Reduction, Water Science and Technology Board, National Academy Press, Washington D.C., ISBN: 0-309-56982-6, 216 pages.

NELSEN, R. B., (1999). An Introduction to Copulas. Springer-Verlag, New York.

OFFICE OF PUBLIC WORKS, (2008). Flood Studies Update Programme, Work package 3.4, Guidance for River Basin Modelling.

PUGH, D.T, AND VASSIE, J.M. (1979) Extreme sea-levels from tide and surge probability. Procedings 16th Coastal Engineering Conference, 1978, Hamburgh, American Society of Civil Engineers, New York, 1, 911-930

PUGH, D.T, AND VASSIE, J.M. (1980) Applications of the joint probability method for exteme sea-level computations, Proc. Instn. Civ. Engrs., Part 2 69, 959-975

BRIX, A., DUFFY, P.,FORSTER, W, HARTINGTON, T, JONES, G, LEVI, C, PADDAM, P, PAPACHRISTOU, D, PERRY, G, RIX, S, ROSS, F, SMITH, A.J., SETH, A, WESTCOTT, D AND MARK WILKINSON, M.2002. The management of losses arising from extreme events, GIRO 2002. Available from: http://www.actuaries.org.uk/ [accessed 27 January 2010].

SVENSSON, C. AND JONES, D. A. (2002) Dependence between extreme sea surge, river flow and precipitation in eastern Britain, Int. J. Climatol., 22, 1149-1168.

SVENSSON, C. AND JONES, D. A. (2004) Dependence between extreme sea surge, river flow and precipitation in south and west Britain, Hydrol and Earth System Science, 8, 973-992.

TAWN, J. A. (1988). An extreme value theory model for dependent observations. J. Hydrology, 101, 227–250.

TAWN, J. A. AND VASSIE, J. M. (1989). Extreme sea levels: the joint probabilities method, revisited and revised. Proc. Instn. Civ. Engrs. Part 2, 87, 429–442.

TROUTMAN, B. M. AND KARLINGER, M. R. (2003) Regional flood probabilities, Water Resour. Res., 39(4) 1095.

THURSTON N, CROSSMAN M, MUIR-WOOD R, DRAYTON M, ASH J. (2007) Paper 05A-2 Coastal Flood Risk Thinking for Tomorrow, Acting today. DEFRA 2007 Flood Risk and Coastal Erosion Conference York

VERLANN, M., ZIJDERVELD, A., DE VRIES, H., KROOS, J.. (2005), Operational storm surge forecasting in the Netherlands: developments in the last decade, Phil. Trans. R. Soc. A, doi:10.1098/rsta.2005.1578

WHEATER, H.S., ISHAM, V.S., CHANDLER, R.E., ONOF, C.J., STEWART, E.J., BELLONE, C., YANG, C., LEKKAS, D., LOURMAS, G., SEGOND, M-L., FROST, A.J., PRUDHOMME, C., CROOKS, S. (2006) Improved methods for national spatial temporal rainfall and evaporation modelling for BSM, Defra Technical Report FD2105/TR.

12 Appendix

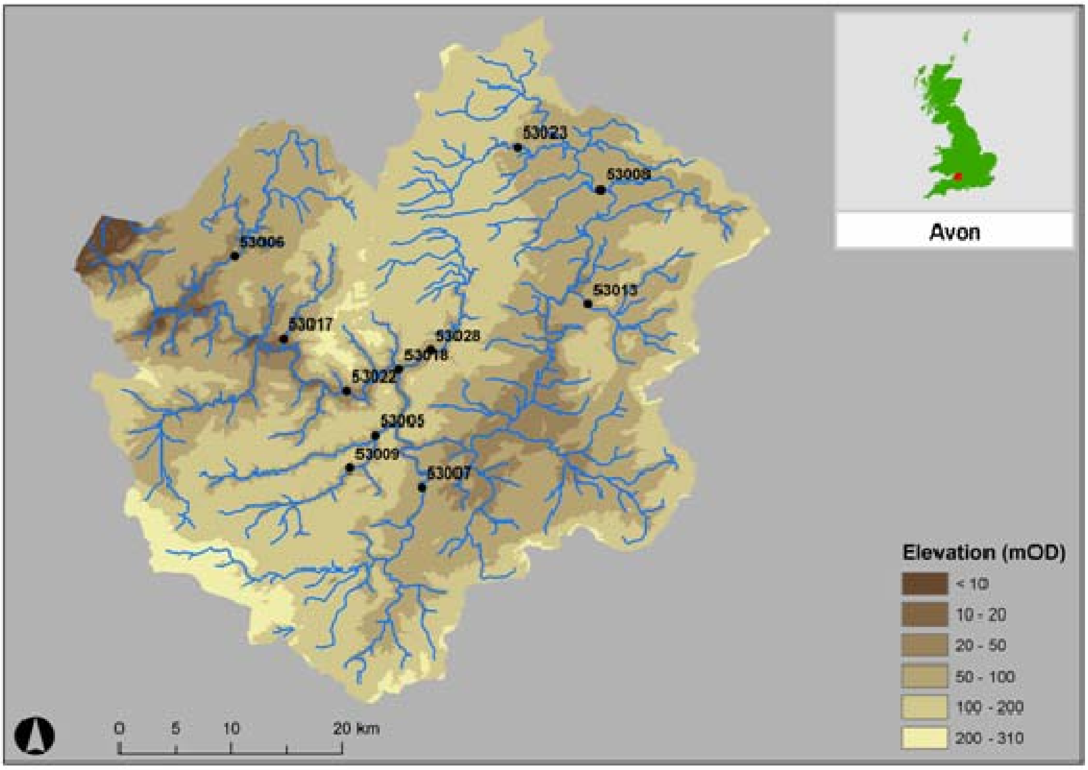

Figure A-1: Avon catchment

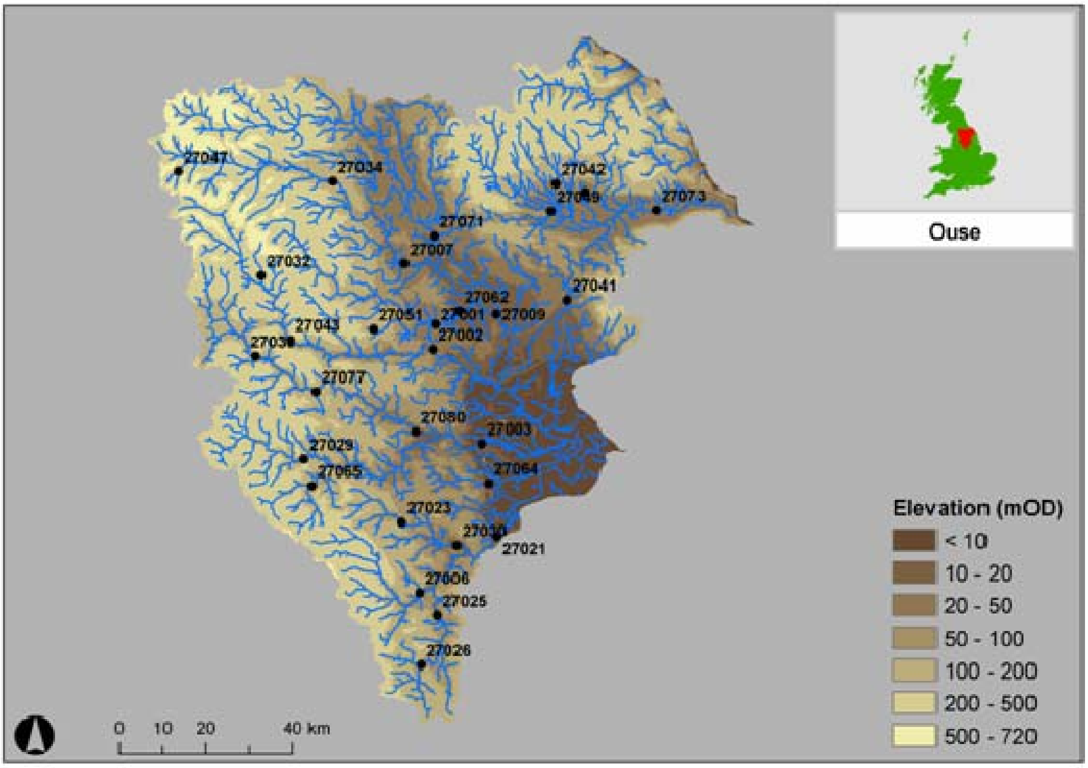

Figure A-2: Ouse catchment

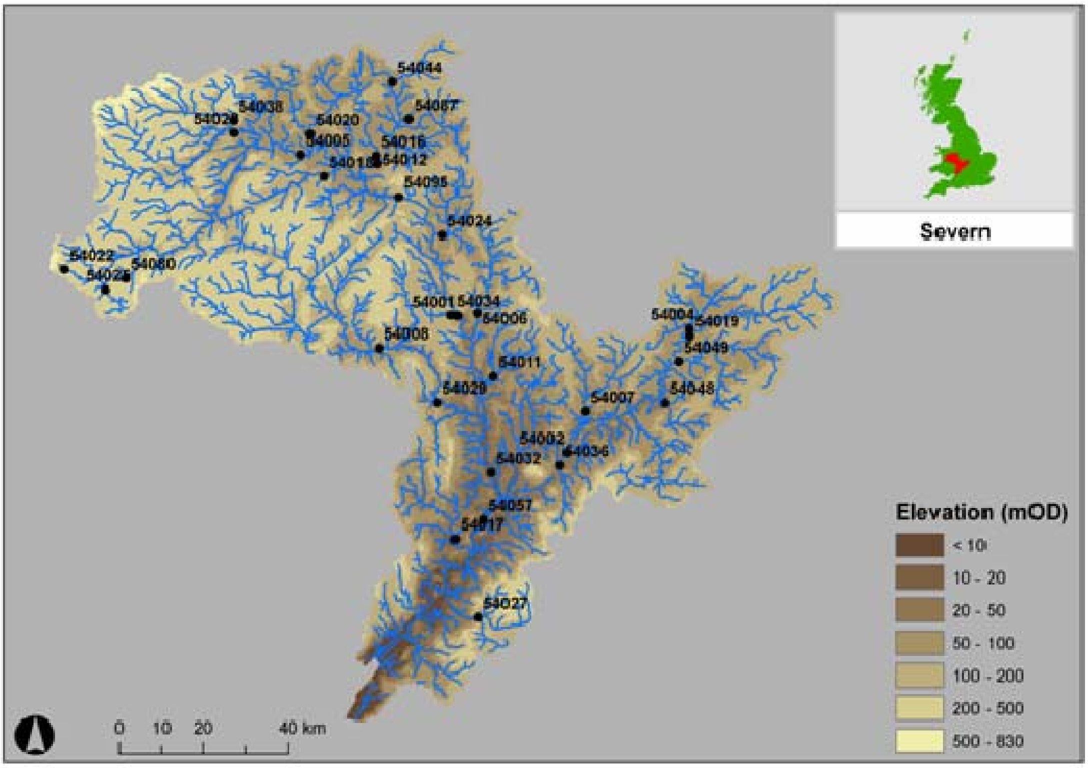

Figure A-3: Severn catchment

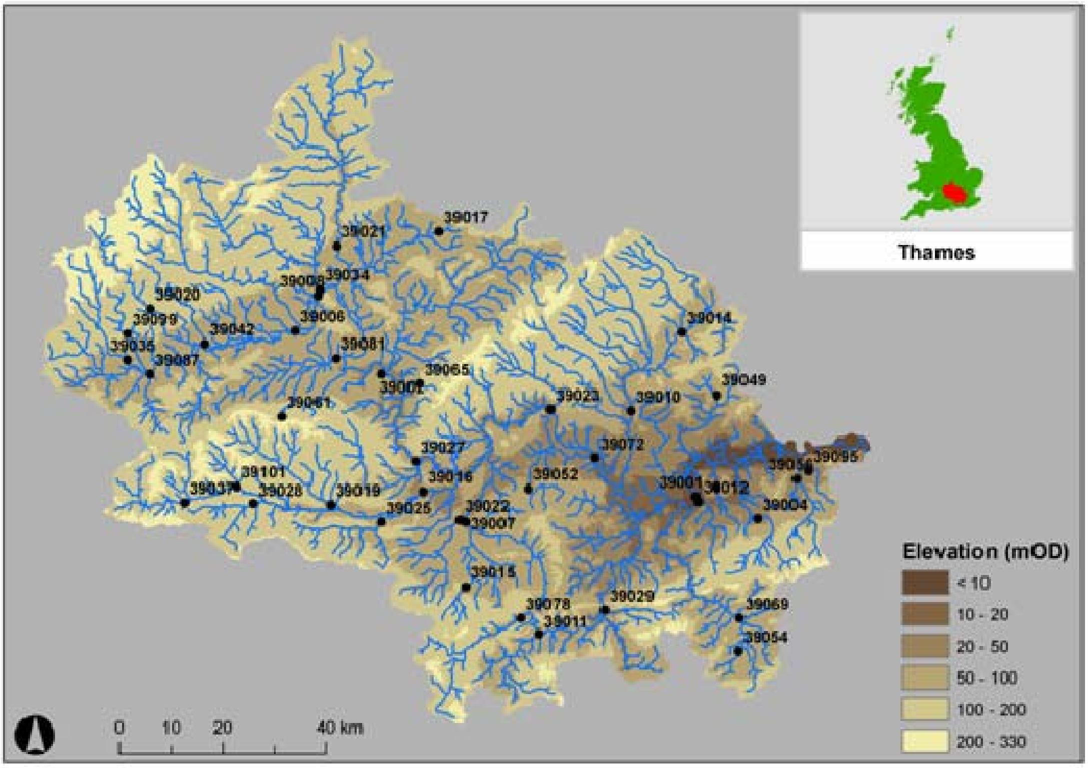

Figure A-4: Thames catchment

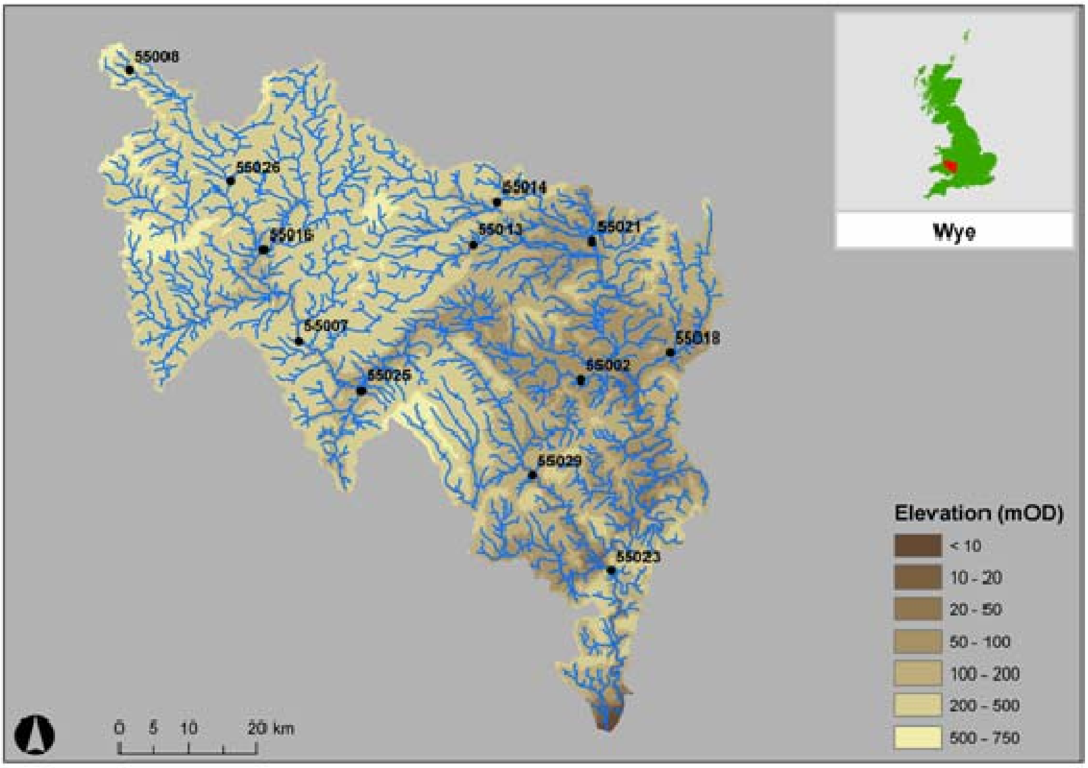

Figure A-5: Wye catchment

Glossary

Correlation	Two variables are correlated if large values in both variables are likely to occur together.
Dependence	Two variables are statistically dependent if the observed value of one of the variables affects the likely values of the other variable.
Dependence structure	This describes the dependencies between a set of variables.
Joint distribution function	The probability distribution function of a set of variables. This can be broken down to the marginal distributions and the dependence structure of the variables.
Marginal distribution	The probability distribution function of a single variable from a set of variables.
Pathway	The mechanism by which a source leads to flooding. Pathways could include defence breaching, defence overtopping, and failure of warning systems or emergency plans.
Receptor	Anything that is at risk from flooding. Receptors include; the general public and their houses; public and emergency services, critical infrastructure.
Source	Something that has the potential to cause flooding. Sources include rivers, sea, and rainfall. In an extreme state all these have the potential to cause flooding.
Spatial coherence	We use this term to describe the dependence between the parameters of the marginal distributions of variables at sites that are close together spatially.

We are The Environment Agency. It's our job to look after your environment and make it **a better place** – for you, and for future generations.

Your environment is the air you breathe, the water you drink and the ground you walk on. Working with business, Government and society as a whole, we are making your environment cleaner and healthier.

The Environment Agency. Out there, making your environment a better place.

Would you like to find out more about us, or about your environment?

Then call us on

08708 506 506* (Mon-Fri 8-6)

email

enquiries@environment-agency.gov.uk

or visit our website

www.environment-agency.gov.uk

incident hotline 0800 80 70 60 (24hrs)

floodline 0845 988 1188

* Approximate call costs: 8p plus 6p per minute (standard landline). Please note charges will vary across telephone providers