

**Flood Events:  
Extreme Value Problems  
and  
Efficient Estimation of Loss**

Anna Maria Barlow, M.Math.(Hons.), M.Res



Submitted for the degree of Doctor of Philosophy at  
Lancaster University.

March 2021



# Abstract

Widespread flood events have heavy consequences on society and the environment. Gaining insight into the occurrence and impact of these rare flood events is thus of interest to many parties such as governments, environmental organisations and insurance companies. To assess flood risk, past events are studied and used to fit statistical models from which plausible flood events are simulated over large areas and large periods of time. These simulated extreme events then drive other models, such as models of loss for insurance purposes, to provide insight into the possible impact of future flood events.

This thesis addresses problems in the analysis of extreme river flows which cause flooding, and the inefficiency of simulation of yearly loss due to flooding.

Firstly, many extreme value analyses are conducted in reaction to the occurrence of a large flooding event. This timing of the analysis introduces bias and poor coverage probabilities into the associated risk assessments subsequently leading to over-designed flood protection schemes. These problems are explored through studying stochastic stopping criteria and new likelihood-based inferences are proposed that mitigate against these difficulties.

Simulated extreme events are used along with geographical knowledge and property information to simulate losses at each property for each flood event over many years. These simulations are then aggregated to obtain total yearly losses and to estimate return levels of yearly loss. The large number of simulations needed makes this process computationally expensive. A new method is proposed, using novel concentration inequalities, which reduces the number of years that need to be simulated.

Finally, modelling extreme flood events is complicated due to temporal dependence and the spatial dependencies of river flows between multiple locations with the presence of time lags between locations. The theory of multivariate temporally dependent extremes is explored, with focus on measures of dependence, and areas of further research are highlighted.

# Acknowledgements

Firstly, I would like to thank my supervisors, Jon Tawn and Chris Sherlock, for their guidance. Jon's continual support and encouragement has been a huge help in motivating me and calming me down when I was stressed. I am grateful to Chris for his enthusiasm and our long meetings problem solving and diving into details which led to lots of discoveries and things to pursue – I'm thankful for his time and patience with this.

I am very grateful for the financial support of EPSRC and JBA. I would also like to thank JBA and Emma Raven for hosting me on my visits and for our discussions - this greatly aided my understanding and gave me a taste of the industry. I would like to specifically thank Carl Fischer for his help in fixing my implementation of the loss simulation procedure.

My time at Lancaster has been greatly enhanced by my fellow year group at STOR-i – in this bubble I felt a strong sense of friendship and support. I would especially like to thank Harjit, Luke, Emily, David, Jake and Sam for their support and listening ears - enduring my monologues in the corridor (Jake), at lunchtimes (Luke) and so on! Harjit in particular was great at calming me down and cheering me up as a house mate and at the office.

I am extremely grateful to have found early on my PhD journey my future husband, Stefan, who has been a constant support and stood by me when things have been tough. I know I can be very difficult in times of high stress or low self esteem and I am so thankful to Stefan for enduring this and helping me get out of negative spirals.

Finally, I thank my family for their love and support throughout my life and particularly the last 5 years; my mum and grandma for caring for me and meine kleine Schwester for giving me such joy and much needed breaks from PhD work. Most of all I thank my grandad who I'd like to now speak to directly.

You were always a strong voice of encouragement, reassuring me when I doubted myself and reminding me 'you can only do your best'. I will try to always hold on your words and your positive attitude. There have been many times over my life and particularly the PhD years

in which you gave me strength. I always felt you were a father to me – though I never used the word, to me grandad is definitely right but the feeling is of a father. You were so proud of me doing this PhD – I wish you could have seen me finish it – but you are still with me in my heart always. Thank you for everything.

# Declaration

I declare that the work in this thesis has been done by myself and has not been submitted elsewhere for the award of any other degree.

Chapter 3 has been published in *J. R.Stat. Soc. Series C*, 69:765-789 as Barlow, A. M., Sherlock, C. and Tawn, J. (2020). Inference for extreme values under threshold-based stopping rules.

There are approximately 55944 words in this thesis.

Anna Maria Barlow

# Contents

<b>Abstract</b>	<b>I</b>
<b>Acknowledgements</b>	<b>II</b>
<b>Declaration</b>	<b>IV</b>
<b>Contents</b>	<b>X</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.1.1 Problems in statistical inference for extreme river flows . . . . .	1
1.1.2 Loss simulation and return level estimation . . . . .	3
1.2 Thesis overview . . . . .	4
<b>2 Literature Review</b>	<b>6</b>
2.1 Extreme Value Theory . . . . .	6
2.1.1 Block maxima approach . . . . .	8
2.1.2 Threshold approach and the Generalised Pareto distribution . . . . .	9
2.1.3 Point process representation . . . . .	11
2.2 Sampling from rare events . . . . .	12
2.2.1 Concentration inequalities . . . . .	12
2.2.2 Variance Reduction methods . . . . .	14
2.2.3 Splitting . . . . .	19
<b>I Investigating Extreme Values under Stopping Rules</b>	<b>21</b>
<b>3 Inference for extreme values under threshold-based stopping rules</b>	<b>22</b>
3.1 Introduction . . . . .	22
3.2 Inference under stopping rules . . . . .	28
3.2.1 Introduction . . . . .	28

3.2.2	Likelihood in presence of a stopping rule . . . . .	29
3.2.3	Fixed-threshold stopping rule with exponential observations . . . . .	30
3.2.4	Variable-threshold stopping rule with gamma observations . . . . .	32
3.3	Alternative Methods for Parameter Inference . . . . .	34
3.3.1	New conditional likelihoods . . . . .	34
3.3.2	Application to exponential observations . . . . .	36
3.4	Simulation results . . . . .	36
3.4.1	Simulation design . . . . .	37
3.4.2	Fixed-threshold stopping rule . . . . .	39
3.4.3	Variable-threshold stopping rule . . . . .	40
3.4.4	Use in Practice . . . . .	41
3.5	Case Study - Lune at Caton . . . . .	42
3.5.1	Fixed-threshold stopping rule . . . . .	42
3.5.2	Variable-threshold stopping rule . . . . .	44
3.5.3	Non-stationarity . . . . .	46
3.6	Discussion . . . . .	47
<b>4</b>	<b>Investigating confidence intervals for return-level estimators on data generated by threshold-based stopping rules</b>	<b>49</b>
4.1	Introduction . . . . .	49
4.2	Profile likelihood in Lune setting . . . . .	50
4.2.1	Relationship between confidence interval widths and sample size . . . . .	51
4.2.2	Similar data sets to the Lune data set . . . . .	52
4.3	Bootstrap methods with fixed-threshold stopping rule . . . . .	57
4.3.1	Methods . . . . .	57
4.3.2	Comparison of confidence interval methods . . . . .	65
4.4	Bootstrap CIs with variable-threshold stopping rule . . . . .	69
4.5	Influence of $c$ on conditional CIs . . . . .	75
4.6	Summary . . . . .	82
<b>II</b>	<b>Efficient Loss Estimation</b>	<b>84</b>
<b>5</b>	<b>Loss Estimation and Concentration Inequalities</b>	<b>85</b>

5.1	Notation and standard procedure . . . . .	86
5.1.1	Simulating flood events . . . . .	87
5.1.2	Portfolio information . . . . .	87
5.1.3	Water depths and distribution . . . . .	88
5.1.4	Vulnerability function and distribution . . . . .	89
5.1.5	Effective damage distribution . . . . .	89
5.1.6	Sampling and aggregation . . . . .	90
5.1.7	Test data . . . . .	93
5.2	Concentration inequalities . . . . .	93
5.2.1	Basic inequalities . . . . .	94
5.3	Concentration inequalities for sums of bounded random variables . . . . .	95
5.3.1	Assumptions, definitions and notation . . . . .	96
5.3.2	Preliminaries . . . . .	96
5.3.3	Chernoff's inequality for the sum of independent variables . . . . .	99
5.3.4	Hoeffding's inequality for sums of bounded random variables . . . . .	99
5.3.5	Chernoff-Hoeffding inequality . . . . .	99
5.3.6	Bernstein's inequality . . . . .	103
5.3.7	Bennett's inequality . . . . .	105
5.3.8	Jebara's Bennett refinement . . . . .	105
5.3.9	Comparing existing concentration inequalities on simulated data . . . . .	106
5.4	Tighter concentration inequalities . . . . .	109
5.4.1	Convexity tricks with the Chernoff-Hoeffding bound . . . . .	109
5.4.2	Improving Bennett's inequality . . . . .	114
5.4.3	Comparing our concentration inequalities on simulated binary data . . . . .	116
5.5	Connecting concentration inequality notation with the loss estimation setting . . . . .	119
5.6	Summary . . . . .	126
<b>6</b>	<b>Improving Loss Estimation</b> . . . . .	<b>128</b>
6.1	Efficient loss simulation via concentration inequalities . . . . .	128
6.1.1	Best Case from consideration of Monte Carlo samples . . . . .	129
6.1.2	Results . . . . .	130
6.1.3	Comparison of standard and exclude procedure times . . . . .	133
6.2	Non-equivalence probability . . . . .	135

6.3	Estimating return levels via the Berry-Esseen inequality . . . . .	141
6.3.1	Normal approximation . . . . .	141
6.3.2	Berry-Esseen procedure . . . . .	145
6.3.3	Combined procedure . . . . .	153
6.3.4	Discussion . . . . .	155
<b>III Extreme Values of Temporally Dependent and Multivariate Sequences</b>		<b>157</b>
<b>7</b>	<b>Extremes of dependent sequences</b>	<b>158</b>
7.1	Stationary dependent sequences . . . . .	158
7.1.1	Cluster size distribution and point process formulation . . . . .	160
7.1.2	Measures of extremal serial dependence . . . . .	161
7.1.3	Declustering . . . . .	163
7.2	The ARMAX process . . . . .	166
7.2.1	Extremal Index . . . . .	167
7.2.2	Cluster size distribution . . . . .	168
7.2.3	Measures of serial dependence . . . . .	169
7.2.4	Simulation and declustering . . . . .	169
7.3	Multivariate Extreme Value Theory . . . . .	171
7.3.1	Componentwise maxima approach . . . . .	172
7.3.2	Dependence Structure . . . . .	174
7.3.3	Copulas . . . . .	176
7.3.4	Measures of dependence . . . . .	178
7.3.5	Point process representation . . . . .	181
<b>8</b>	<b>Extremes of multivariate temporally dependent sequences</b>	<b>185</b>
8.1	Multivariate extremes of dependent sequences . . . . .	185
8.1.1	Limit distributions . . . . .	185
8.1.2	Multivariate Extremal Index . . . . .	188
8.1.3	Cluster size distribution(s) and point processes . . . . .	190
8.1.4	Coefficient of asymptotic dependence for dependent multivariate sequences . . . . .	193

8.2	The multivariate ARMAX process . . . . .	194
8.2.1	The MARMAX process . . . . .	194
8.2.2	Multivariate Extremal Index . . . . .	196
8.2.3	Examples . . . . .	196
8.2.4	$(\chi, \bar{\chi})$ for the MARMAX process . . . . .	198
8.2.5	MARMAX simulations and estimation of $\theta(\mathbf{x})$ . . . . .	199
8.3	The M4 process . . . . .	202
8.3.1	Theory . . . . .	202
8.3.2	Measures of dependence . . . . .	203
8.3.3	Cluster size distribution . . . . .	206
8.3.4	M4 simulations and estimation of $\theta(\mathbf{x})$ . . . . .	211
8.4	Summary . . . . .	212
<b>9</b>	<b>Conclusion</b> . . . . .	<b>213</b>
9.1	Summary . . . . .	213
9.1.1	Stopping bias . . . . .	213
9.1.2	Loss Estimation . . . . .	214
9.1.3	Extremes of dependent and multivariate sequences . . . . .	216
9.2	Possible further work . . . . .	216
9.2.1	Stopping Bias . . . . .	216
9.2.2	Loss simulation and return level estimation . . . . .	217
9.2.3	Extremes of dependent and multivariate sequences . . . . .	222
	<b>Appendices</b> . . . . .	<b>227</b>
<b>A</b>	<b>Appendix to Chapter 3</b> . . . . .	<b>227</b>
A.1	Proof of results from Chapter 3 §3.2 . . . . .	227
A.1.1	Proof of Proposition 3.2.3.1 . . . . .	227
A.1.2	Proof of Theorem 3.2.4.1 in Chapter 3 . . . . .	228
A.2	Properties of the GEV shape parameter . . . . .	230
A.2.1	Fixed-threshold stopping rule . . . . .	230
A.2.2	Variable-threshold stopping rule . . . . .	232

<b>B Appendix to Chapter 4</b>	<b>233</b>
B.1 Profile-likelihood based confidence intervals . . . . .	233
B.2 Comparing bootstrap-based confidence intervals . . . . .	236
B.3 Optimisation issues near the boundary . . . . .	242
<b>C Appendix to Chapter 5</b>	<b>248</b>
C.1 Proof of Hoeffding's Inequality (Theorem 5.3.4.1) . . . . .	248
C.2 Proof and exploration of Proposition 5.3.9.1 . . . . .	249
C.3 Issues with Jebara's Bennett refinement . . . . .	253
<b>D Appendix to Chapter 6</b>	<b>256</b>
D.1 Investigating the effect of event set size . . . . .	256
D.2 Exclude procedure results . . . . .	262
<b>E Appendix to Chapter 7</b>	<b>268</b>
E.1 Heffernan and Tawn's conditional multivariate extremes model . . . . .	268
E.1.1 Inference in the bivariate setting . . . . .	270
E.2 ARMAX simulation . . . . .	271
<b>Bibliography</b>	<b>273</b>

# Chapter 1

## Introduction

### 1.1 Motivation

Flooding can have a severe impact on society causing huge disruptions to life and great loss to homes and businesses. The December 2015 floods across Cumbria, Lancashire and Yorkshire caused widespread damage and tens of thousands of properties were left without power. Governments, environmental agencies and insurance companies are keen to know more about the causes and the probabilities of the re-occurrence of such events to prepare for future events. Therefore we wish to better understand the flood risk and the magnitude of losses that can be incurred. This PhD project with JBA Risk Management focuses on modelling such extreme events and estimating the total impact (in terms of financial losses).

JBA is a group of companies concerned with environment, engineering and risk. This project is with part of this group, JBA Risk Management, which specialises in consulting for flood risk and other natural extreme events. Their clients range from reinsurance companies to utility companies to local authorities and they provide a range of services such as flood risk assessments and portfolio analysis. In order to assess the risk from flooding one needs to simulate extreme flood events, identifying and tackling problems in extreme value simulation of river flow is one main focus of the project. The second topic of the project is concerned with studying the tail of the loss distribution and improving the efficiency of estimation of the quantiles of this distribution.

#### 1.1.1 Problems in statistical inference for extreme river flows

The simulation of flood events is important in understanding the flood risk and determining the loss distribution. Typically there is little loss history available and so extrapolation purely from this data would be unreliable. Thus, we instead consider the mechanism that leads to these losses, *i.e.* the extreme weather events. This part of the project is based on extreme

value theory since we are interested in the events that create the greatest losses of which there may be little or no past data to extrapolate from. Extreme value theory is the development of statistical models and techniques for describing rare events. In other words, concern lies in fitting the tails of a distribution correctly unlike much statistical theory which is generally more focussed on the body of the distribution of interest. An introduction to and overview of classical univariate extreme value theory is provided in Chapter 2.

JBA currently use the Heffernan and Tawn model (a conditional approach to modelling extremes developed by Heffernan and Tawn (2004), see Appendix E) to obtain river flows at a set of points, corresponding to gauge sites, for events simulated over a long period, typically 10000 years. As part of the PhD project we aim to address some of the current issues with the extreme value analysis underlying this model.

Firstly, the point in time at which we decide to analyse the data can have an effect on inference. Renewed analyses of river flow data are usually performed as a consequence of the occurrence of a major flood event, *i.e.*, the timing of the analysis depends on observing a large river flow and so the size of the data set of river flows is random. Such analysis is generally performed in order to assess the efficiency of existing and proposed defence schemes. The UK currently spends £400-500M per year on flood defence infrastructure. It is important that the flood risk analysis is as accurate as possible to make decisions on future investments; underestimation of flood risk may lead to inadequate flood defences whereas overestimation of risk may lead to money spent unnecessarily on flood defences which could be put to better use elsewhere. Performing an analysis after a major flood event introduces a positive bias in the estimated flood risk using the standard inference methods. In this thesis we focus on inference in the univariate setting (for river flows at one gauge site) with two threshold-based rules to decide on when to perform a statistical analysis.

Secondly, modelling flood events can be quite complicated since we need to model in both space and time due to the presence of different locations and lags between events at different locations. On one hand the occurrence and intensity of extreme values are likely to be similar at nearby locations. For example, rainfall will generally be similar at neighbouring locations as it is likely to be part of the same weather system and also the geography of the area may affect the rainfall-runoff process. On the other hand, there is some dependence within the time series as river flows on consecutive days are likely to be similar. So when extreme values are observed they often occur in clusters with similar clusters at nearby locations and/or

further along the same river network. This complex dependence structure complicates the classical extreme value methods. For loss calculations JBA need to aggregate losses, caused by extreme river flows, across a region and over some period of time so having techniques and understanding of multivariate temporally dependent extremes is important. In this thesis we explore the theory of multivariate temporally dependent extremes with focus on measures of dependence.

### 1.1.2 Loss simulation and return level estimation

The second aim of the project is to improve the efficiency of the estimation of total loss incurred from the modelled flooding information. In the following we consider clients to be insurance or reinsurance companies. JBA's clients have *portfolios* containing a number of insured properties/locations called *risks*. Each risk can be associated with multiple insurance coverage types: contents, buildings and business interruption. A standard procedure to estimate the loss distribution is to simulate extreme events, use these to model the water depths (essentially) everywhere, and use these water depths combined with portfolio information to create a loss distribution for each event, risk and coverage type. Losses are then simulated for all the risks in a portfolio and all the simulated events in a year then summed to give a simulation of the total loss for one year. JBA have 10000 years of simulated extreme events, referred to as the event set, and so can simulate the total loss for each of these hypothetical years.

JBA's clients are interested in the distribution of total loss per year from flood events, in particular the mean, variance and  $t$ -year *return levels*, for a range of *return periods* from  $t = 2$  to  $t = 5000$ . The  $t$ -year return level is the value which is exceeded in any given year with probability  $\frac{1}{t}$ . JBA find estimates of the return levels by simulating the yearly losses over a portfolio for each of 10000 years multiple times (typically 100 times). For each simulation they use all the ordered losses over all years to estimate the  $\frac{1}{t}$ th quantile. The quantile estimates over all simulations can then be used to obtain both a  $t$ -year return level estimate by taking the mean or median, and rough confidence intervals by taking quantiles of these estimates. Chapter 5 provides more detail of this standard procedure. JBA's end product is a 'curve' of return-level estimates and their approximate 95% confidence intervals plotted for a range of return periods. This is referred to in the insurance industry as a loss estimation curve and is used by clients to compare against their own historical data. The

200-year return level (the loss exceeded with probability  $0.005 = 1/200$ ) is of special interest since it is specifically required by the UK Government's 2015 solvency regulation (Swain and Swallow, 2015).

The method currently used to estimate the return levels is computationally expensive since it involves simulating from all risk and coverage combinations in all 10000 years of the event set multiple times. Portfolios can be extremely large (for example, covering Western Europe) and so can contain up to  $10^7$  risks. Clearly the large number of risks and events involved in this process has a huge burden on computation. In 2016 JBA's software took 20 hours to analyse approximately 2 million risks. We aim to find a more efficient method to estimate the return levels of yearly loss especially for high return periods.

This problem of efficient estimation of losses is faced across the flood/windstorm insurance sector where generally losses over large portfolios are determined via computationally expensive simulation.

## 1.2 Thesis overview

We begin in Chapter 2 with a review of classical univariate extreme value theory and rare event simulation. The thesis is then split into three parts covering our topics of interest: I – extreme values under stopping rules; II – efficient loss estimation; and III – extremes of multivariate temporally dependent sequences.

In Part I Chapter 3 we explore the effect on inference of extreme values when the timing of an analysis is dependent on the observation of an extreme event. In particular we find that this timing of the analysis introduces bias and poor coverage probabilities into the associated risk assessments and leads subsequently to inefficient flood protection schemes. We explore these problems through studying stochastic stopping criteria and propose new likelihood-based inferences that mitigate against these difficulties. Our methods are illustrated through the analysis of the river Lune, following it experiencing the UK's largest ever measured flow event in 2015. We show that without accounting for this stopping feature there would be substantial over-design in response to the event.

In Chapter 4 we continue our study of Chapter 3 by exploring in depth the features of profile-likelihood based confidence intervals and a variety of bootstrap-based confidence intervals for estimators under the stopping criteria of Chapter 3. We concentrate on the

Lune river flow data and create data sets of the same size to compare confidence interval methods and discuss the advantages and disadvantages of each method particularly when we have a small data set. We find that the profile-likelihood based confidence intervals are the best choice in terms of coverage, however, there is promise that the bootstrap methods could be improved to provide similarly well performing confidence intervals. We also discuss the problems which can arise using the likelihoods developed in Chapter 3 when the final observation is just large enough to triggered the analysis.

In Part II we discuss the estimation of the return levels of the loss distribution and our approaches to increase the computational efficiency of this estimation process. In Chapter 5 we describe in detail JBA's standard procedure to estimate quantiles of the loss distribution for a given portfolio from simulated events. A review of classic concentration inequalities is given and novel, tighter, bounds are developed for sums of bounded random variables with emphasis on random variables that are 0 with large probability but have large upper bounds.

In Chapter 6 we introduce a novel approach to reduce simulations using concentration inequalities and evaluate this procedure with a test portfolio and event set provided by JBA. We also discuss a possible method to estimate the return levels with low return periods.

In Part III we discuss the extension of classical extreme value theory to sequences with serial dependence and multiple dimensions. In Chapter 7 we focus on these two extensions separately then in Chapter 8 we bring both extensions together. In particular we investigate the multivariate extremal index, a measure of average cluster size of extreme events over locations and times, and we extend a measure of extremal dependence to describe structures with two sets of components over different time lags. We derive these measures for two multivariate stationary processes, the MARMAX process and the M4 process, and discuss estimation of the multivariate extremal index through simulations of these processes.

Finally, in Chapter 9 we conclude the thesis with a summary of the results of each part and discuss possible future research directions in each topic. In the appendices we provide details of various investigations in the main thesis text including proofs, figures and tables.

# Chapter 2

## Literature Review

This literature review focuses on two main areas of interest for our work: extreme value theory (§2.1) and rare event sampling (§2.2).

### 2.1 Extreme Value Theory

Often it is of interest to estimate the probability of events rarer (more extreme) than those observed and so some kind of extrapolation is required. Extreme value theory is concerned with modelling the tails of a distribution (the rare events) and is based on an asymptotic argument in a similar vein to the central limit theorem. In this literature review we focus on univariate extreme value theory for independent sequences – extensions to dependent and multivariate sequences are considered later in Chapters 7 and 8. We describe the block maxima and threshold approaches to modelling extremes for sequences of independent random variables and the associated distributions. The point process representation is briefly described in §2.1.3. A good introduction to the subject can be found in Coles (2001).

We are interested in the extreme values of a particular process represented by the sequence of random variables,  $\{X_t\}_{t \geq 1}$ . For example, this process could be daily rainfall at a site, in which case  $X_t$  would be the rainfall on day  $t$ . In the simplest case we consider each  $X_t$  to be independent and identically distributed over an observation period of length  $n$ . We consider the sequences of random variables with serial dependence later in §7.1. Then, if the common distribution,  $F$ , were known we could easily find the distribution of the maximum over the observation period, since:

$$\mathbb{P}(M_n \leq z) = \mathbb{P}(X_1 \leq z) \dots \mathbb{P}(X_n \leq z) = F^n(z),$$

where  $M_n = \max\{X_1, \dots, X_n\}$ .

Of course in reality the distribution  $F$  is unknown and simply estimating  $F$  can lead to

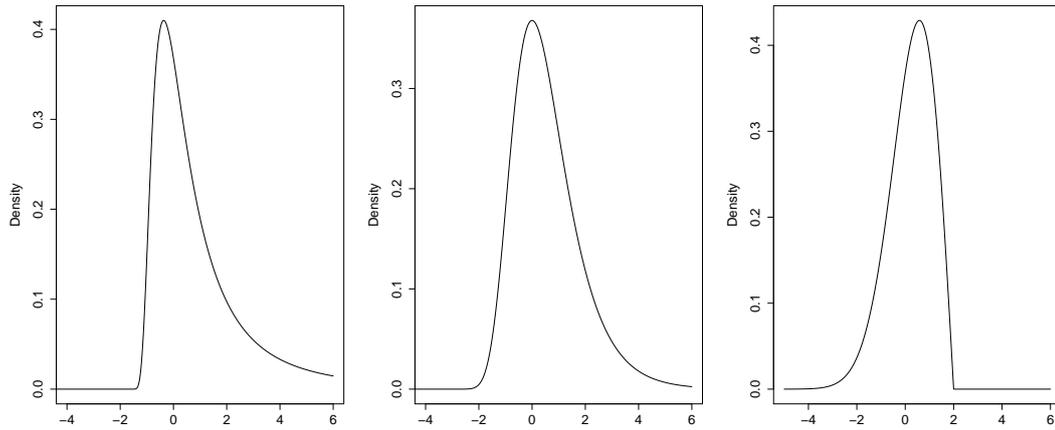


Figure 2.1.1: The Fréchet (left), Gumbel (middle) and Weibull (right) probability densities.

large error in the estimate of  $F^n$ , so instead we adopt a limiting distribution for the maxima as  $n \rightarrow \infty$ . However, the distribution of  $F^n$  reduces to a point mass at the upper end point of  $F$  with zero mass elsewhere as  $n \rightarrow \infty$ . To circumvent this we normalise using a sequence of constants ( $a_n > 0, b_n$ ) and consider the limiting distribution of

$$Z_n = \frac{M_n - b_n}{a_n}. \quad (2.1.1)$$

If such a sequence of normalising constants,  $b_n$  and  $a_n > 0$ , exists such that the distribution of  $Z_n$  in the limit as  $n \rightarrow \infty$  is non-degenerate then the limit distribution of  $Z_n$  is a member of the family of *Generalized Extreme Value* (GEV) distributions with cdf:

$$G(x) = e^{-\tau(x)} \quad \text{with} \quad \tau(x) = \begin{cases} [1 + \xi \left(\frac{x-\mu}{\sigma}\right)]_+^{-\frac{1}{\xi}} & \xi \neq 0 \\ e^{-\frac{(x-\mu)}{\sigma}} & \xi = 0 \end{cases} \quad (2.1.2)$$

where  $\mu$ ,  $\sigma$  and  $\xi$  are the location, scale and shape parameters respectively and  $[y]_+$  is 0 for  $y < 0$ . This distribution consists of three classes of distribution: *Fréchet* (when  $\xi > 0$ ), *Gumbel* (when  $\xi = 0$ ) and *Weibull* (when  $\xi < 0$ ). The shape parameter controls the rate of decay of the tails of  $F$  as illustrated in Figure 2.1.1. When  $\xi = 0$  the rate of decay in the upper tail is exponential whereas when  $\xi > 0$  the upper tail is heavier and for  $\xi < 0$  the distribution has a finite upper point so the maximum value possible is constrained by this upper bound.

The key to the proof of the limit distribution lies in the concept of *max stability*. For  $Z_n$

to have a non-degenerate limit the distribution  $G$  must be max-stable, that is there exists constants  $A_k > 0$  and  $B_k$  such that

$$G^k(A_k z + B_k) = G(z) \quad \forall k \in \mathbb{N}. \quad (2.1.3)$$

It then arises that a distribution is max-stable if and only if it belongs to the family of generalised extreme value distributions.

A distribution function,  $F$ , is said to belong to the *domain of attraction* of its limiting distribution,  $G$ , if and only if there exists sequences  $a_n > 0$  and  $b_n$  such that  $F^n(a_n x + b_n) \rightarrow G(x)$  as  $n \rightarrow \infty$ . For example, the distribution functions in the domain of attraction of the Gumbel distribution are the distribution functions for which

$$\lim_{n \rightarrow \infty} \mathbb{P}(Z_n \leq z) = G(z) = e^{-e^{-\frac{(z-\mu)}{\sigma}}},$$

*i.e.*,  $G$  is the Gumbel distribution function.

### 2.1.1 Block maxima approach

In practice to use the GEV distribution the data are split into blocks of equal length and the maxima of each of these blocks is modelled by the GEV distribution. This method is natural for some types of data, for example when only the annual maxima are recorded. The limiting distribution, (2.1.2), applies with block size tending to infinity, therefore taking block sizes too small will result in bias due to poor approximation in the limit. On the other hand, if blocks are taken to be too large there will be fewer data points available to fit the GEV distribution and hence large variance in the parameter estimates.

One can easily obtain estimates for the parameters and combinations thereof by maximizing the likelihood. Of particular interest is the return level - the return level corresponding to the  $t$ th return period is the value which is exceeded on average once every  $t$  periods. In most settings the relevant period is a year. It is informative to plot the estimated return levels along with confidence intervals against  $\log(-\log(1 - \frac{1}{t}))$  or  $\log(\frac{1}{t})$  in a *return level plot*; the two choices for the  $x$ -axis are approximately equal for large  $t$ . With this choice of  $x$ -axis we obtain a linear return level plot when the shape parameter,  $\xi$ , is 0, convexity when  $\xi > 0$  and concavity with a finite bound on the return level when  $\xi < 0$ . Confidence intervals can then be found using the delta method or, the more accurate, profile likelihood method.

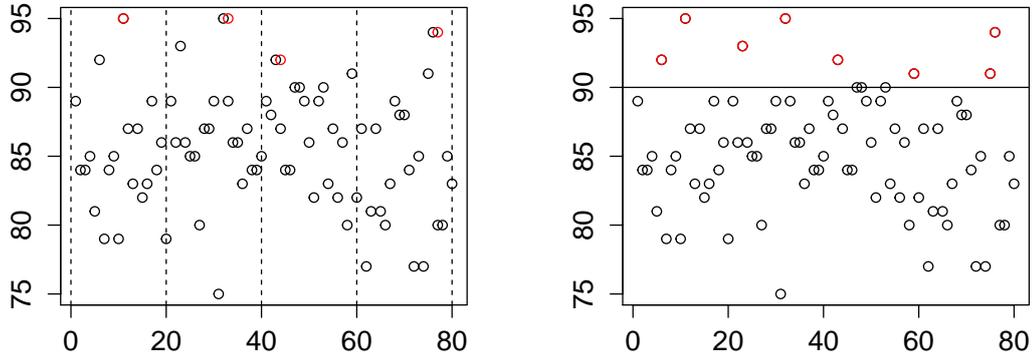


Figure 2.1.2: The block maxima approach with block size 20 (left) and the threshold approach with threshold,  $u = 90$  (right) applied to a test data set.

### 2.1.2 Threshold approach and the Generalised Pareto distribution

The block maxima approach is wasteful if we have more data available on the extreme values. This extra information can be included and further analysis improved by adopting a threshold approach. For suitably large  $u$  the exceedances of this threshold are typically assumed to be exactly modelled by their limiting distribution as the threshold tends to the upper end point of the distribution. This limiting distribution is the *Generalised Pareto* distribution (GPD) which has distribution function for  $y > 0$ :

$$H(y) = \mathbb{P}(X \leq u + y | X > u) = \begin{cases} 1 - \left(1 + \frac{\xi y}{\sigma_u}\right)_+^{-\frac{1}{\xi}} & \xi \neq 0 \\ 1 - \exp\left(-\frac{y}{\sigma_u}\right) & \xi = 0, \end{cases} \quad (2.1.4)$$

where  $\xi$  and  $\sigma_u > 0$  are the shape and scale parameter respectively. Note that if  $\xi$  is zero then the GPD is equivalent to the exponential distribution with rate parameter  $\sigma_u^{-1}$ . The shape parameter  $\xi$  is the same as that under the GEV distribution whereas the scale parameter changes with threshold with  $\sigma_u = \sigma + \xi(u - \mu)$  where  $(\mu, \sigma, \xi)$  are the associated GEV parameters. For modelling using the GPD we also need to model the rate at which the threshold  $u$  is exceeded.

There are similar issues with threshold choice as with block length choice for the first approach. The limit distribution of Equation (2.1.4) will only hold if the threshold,  $u$ , is large enough whereas, taking a threshold too high reduces the amount of data we can use to

fit the GPD. One way to check whether the threshold is large enough is to inspect the mean residual life plot. This is the plot of the threshold  $u$  against

$$\frac{1}{n_u} \sum_{i=1}^{n_u} (x_{(i)} - u) \quad \text{for } u < \max_t X_t,$$

where  $x_{(1)}, \dots, x_{(n_u)}$  are the  $n_u$  observations that exceed  $u$ . The expected exceedance for a given threshold, *i.e.*, the mean residual, is a linear function of the threshold where the GPD is valid. Therefore, above a particular threshold,  $\tilde{u}$ , the return-level plot will be approximately linear suggesting that the approximation in the limit is valid for higher thresholds,  $u > \tilde{u}$ . The estimates for both the shape parameter,  $\xi$ , and  $\sigma_u - \xi u$  are constant for high enough thresholds so plotting the estimates of these for increasing threshold, along with confidence intervals, can give another indication as to which threshold the GPD is valid above.

Additionally one can assess the model fit in the ‘usual way’ by checking histograms, probability plots and quantile plots and also by plotting the empirical estimates for the return level on the return level plot to see if they are in agreement with those predicted by the model.

The methods of modelling checking and threshold selection described above are quite subjective and can be time consuming. Scarrott and Macdonald (2012) provide a review of threshold selection methods including more recent approaches. More recently Wadsworth (2016) presented an automated procedure to select the threshold using a likelihood ratio test. The issue of threshold selection remains an area of considerable focus still.

All of the above theory and methods were presented after assuming we have a sequence of independent and identically distributed random variables. However, in many cases we may have data which are non-stationary (*i.e.*, the underlying distribution of the data is changing through time) or dependent. The case of dependent data is explored in §7.1. Non-stationarity of a series could be seasonal (*e.g.*, temperature) and/or be due to some other, possibly latent, processes. Such factors can be incorporated into the model by writing the parameters as a function of time and/or some covariate(s) (see, for example, Eastoe and Tawn 2009; Eastoe 2019; Turkman et al. 2010).

### 2.1.3 Point process representation

The point process representation unifies the GP and GEV models. With  $\{X_t\}_{t \geq 1}$ ,  $M_n$  and  $G(z)$  defined as in §2.1, consider the collection of points:

$$P_n = \left\{ \frac{X_{t+i} - b_n}{a_n} \right\}_{i=1, \dots, n},$$

for sequences  $b_n, a_n > 0$  such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \frac{M_n - b_n}{a_n} < z \right) = G(z).$$

Let  $z_-$  and  $z_+$  be the upper and lower points of  $G(z)$  respectively and define the point process,  $N_n$ , counting the points  $P_n$  above  $z$  as

$$N_n([z, z_+)) = \sum_{i=1}^n \mathbb{1} \left\{ \frac{X_{t+i} - b_n}{a_n} \in [z, z_+) \right\} \quad \text{where } z > z_-.$$

Then  $N_n \rightarrow N$  as  $n \rightarrow \infty$  where  $N$  has a Poisson distribution with mean

$$\Lambda[z, z_+) = \left[ 1 + \xi \left( \frac{z - \mu}{\sigma} \right) \right]_+^{-\frac{1}{\xi}}. \quad (2.1.5)$$

This is instantly recognisable as the exponent part of the GEV distribution function (2.1.2). We say that the process that distributes the points,  $P_n$ , converges on the set  $[z, z_+)$  with  $z > z_-$  to a *non-homogeneous Poisson point process*,  $P$ , with *intensity*, on  $[z, z_+)$ , given by  $\Lambda[z, z_+)$ .

The probability of the normalised maximum being less than a certain value is equivalent to there being no points above this value in the point process, thus the limit distribution of  $Z$  is the GEV distribution as before. Mathematically,

$$\begin{aligned} \mathbb{P}(Z \leq z) &= \lim_{n \rightarrow \infty} \mathbb{P}(\text{No points of } P_n \text{ in } [z, \infty)) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}(N_n([z, \infty)) = 0) \\ &= \mathbb{P}(N([z, \infty)) = 0) \\ &= \exp(-\Lambda[z, \infty)). \end{aligned}$$

For regions  $(u, \infty)$ , with  $u$  suitably large, we assume  $N_n = N$  and we absorb the norming

sequences,  $a_n$  and  $b_n$  into  $\mu$  and  $\sigma$ . Then the Poisson point process likelihood is

$$L(\mu, \sigma, \xi) = \exp \{-\Lambda(u, \infty)\} \prod_{i=1}^{n_u} \lambda(x_i), \quad (2.1.6)$$

where  $\lambda(x) = \frac{d}{dz} \Lambda[z, z_+]|_{z=x}$ ,  $\{x_i\}$  are the points above  $u$  and  $n_u$  is the number of points in the region. This likelihood can be maximised to find estimates for the parameters  $\mu, \sigma, \xi$ .

Alternatively, the threshold excess likelihood can be found by including information on the probability,  $p_u$ , of a point  $x_i$  exceeding  $u$ . Let  $f_{X|X>u}$  denote the density of  $X$  given  $X > u$ . Then we arrive at the likelihood:

$$\begin{aligned} L(p_u, \tilde{\sigma}, \xi) &= \mathbb{P}(X_i \leq u)^{n-n_u} \prod_{i=1}^{n_u} f_{X|X>u}(x_i) \mathbb{P}(X_i > u) \\ &= (1 - p_u)^{n-n_u} \prod_{i=1}^{n_u} p_u h(x_i - u) \\ &= (1 - p_u)^{n-n_u} p_u^{n_u} \prod_{i=1}^{n_u} h(x_i - u), \end{aligned}$$

where  $h(y)$  is the derivative of the generalised Pareto distribution function (2.1.4).

## 2.2 Sampling from rare events

To improve the efficiency of the estimation of the return levels of yearly loss due to flooding we draw on ideas from a range of sampling techniques. We make extensive use of concentration inequalities in particular; we provide a brief description and literature review of these in §2.2.1 but leave the mathematical details to §5.2 of Chapter 5. We also consider variance reduction methods (§2.2.2) and splitting methods (§2.2.3).

### 2.2.1 Concentration inequalities

Concentration inequalities provide bounds on the probability of a random variable deviating from a particular value, such as its expectation, by at least some margin, and so are especially helpful in finding bounds for tail probabilities. An advantage to finding bounds on probabilities using concentration inequalities is that these bounds are absolute, unlike the approximate bounds obtained from the central limit theorem.

A wide range of concentration inequalities has been developed, requiring varying amounts of information about the random variable of interest. The most basic concentration inequality,

the Markov inequality, requires knowledge of the expectation of the random variable only. The tighter Chebyshev and Cantelli inequalities are acquired by also using the variance of the random variable.

When the random variable of interest is a sum of independent random variables then tighter concentration inequalities can be developed; these inequalities are generally derived from the Chernoff inequality (Boucheron, 2013) which requires knowledge of the moment generating function of the random variables. However, the moment generating function may be unknown or difficult to compute so looser but tractable concentration inequalities have been developed for specific cases such as the sum of independent bounded random variables (*e.g.*, Bernstein's inequality (Bernstein, 1946)). Bennett's inequality (Bennett, 1962), which uses information on the maximum deviation of the random variables in the sum from their expectations, is one of the tightest inequalities known for sums of *bounded* random variables.

Refinements to the classic concentration inequalities including Bennett's inequality have been considered in the literature (*e.g.*, Jebara (2018); Zheng (2017); From and Swift (2013)) but these variations tend to be intractable and/or difficult to compute so the classic inequalities are more commonly used in practice. Hertz (2020) presents an improved version of Hoeffding's inequality (§5.3.4) based on an improvement of Hoeffding's Lemma when the distribution is skewed to the left and Kutin (2002) proves an extension of Bernstein's inequality with an upper bound on the probability of each independent random variable exceeding some value.

A range of concentration inequalities including those mentioned above is presented in detail in Chapter 5. In addition, examples and more details of the inequalities outlined above can be found in Ross (1996) and a detailed overview of concentration inequalities is given by Boucheron et al. (2004).

Gollini and Rougier (2015) apply the Markov, Cantelli and Chernoff inequalities in the insurance setting to estimate the tail probabilities of the total loss. They assume a general form for the loss distribution corresponding to a particular event and assume events arrive as a Poisson process. This results in a compound Poisson distribution for the total loss over a period of time. Gollini and Rougier (2015) also apply the generalised Markov inequality to the  $k$ th power of the random variable of interest,  $S^k$ , to obtain the so-called moment inequality, which requires knowledge of  $\mathbb{E}[S^k]$ . Gollini and Rougier (2015)'s overall conclusion was that the moment bound was the best bound to use however  $\mathbb{E}[S^k]$  is not easily calculated in

general and so calculation of the bound can become challenging.

## 2.2.2 Variance Reduction methods

A standard procedure to estimate the quantities of interest from the yearly loss distribution is to use simulation and Monte Carlo methods. As we are especially interested the tails of this distribution this method is inefficient since many of the simulations do not contribute to these tails, which by definition consist of rare values. It would be more efficient to find an estimator for the quantiles of interest which has a lower variance than the Monte Carlo estimator, so we obtain more accurate estimates. Methods that seek to accomplish this are called *variance reduction methods*.

Before exploring variance reduction methods we first consider the Monte Carlo estimator for comparison. In Monte Carlo estimation the aim is to estimate an expectation (*i.e.* an integral), which we write as  $I = \mathbb{E}[g(X)]$  where  $X$  has a distribution  $F$  and a density  $f$ . The simple (or naive) Monte Carlo estimate is then:

$$\hat{I}_{MC} = \frac{1}{n} \sum_{i=1}^n g(x_i),$$

where  $(x_1, \dots, x_n)$  are independently simulated from the target distribution. Note that estimates of probabilities,  $\mathbb{P}(X \in A)$ , can be found this way by replacing  $g(x)$  by the indicator function,  $\mathbb{1}_{\{X \in A\}}$ . The naive Monte Carlo estimate is unbiased ( $\mathbb{E}[\hat{I}_{MC}] = I$ ) and its variance is:

$$\text{Var}(\hat{I}_{MC}) = \frac{1}{n^2} \sum_{i=1}^n \text{Var}(g(X_i)) = \frac{1}{n} \text{Var}(g(X)) = \frac{1}{n} \left( \mathbb{E}_f [g(X)^2] - \mathbb{E}_f [g(X)]^2 \right).$$

JBA's main goal is to estimate the  $t$ -year return level (or  $(1 - \frac{1}{t})$ -quantile) of total yearly loss for multiple return periods  $t \geq 2$ . Quantile estimation via simulation is usually done by first estimating the cumulative distribution function of the random variable of interest and then inverting this estimate to get a quantile estimate. Using the simple Monte Carlo method we have the following estimate of the true distribution function,  $F$ , of the random variable  $X$ :

$$\hat{F}_{MC} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{X_i \leq x\}}, \quad (2.2.1)$$

which has variance,  $\text{Var}(\hat{F}_{MC}) = p(1-p)/n$  where  $p = \mathbb{P}(X \leq x)$ . Then the  $(1 - \frac{1}{t})$ -quantile

estimate is  $\hat{q}_{MC} = \hat{F}_{MC}^{-1}(\frac{1}{t})$ . The empirical distribution function,  $\hat{F}_{MC}$ , is a step-function so the inverse  $\hat{F}_{MC}^{-1}$  is not well defined. One convention is to take  $\hat{F}_{MC}^{-1}(p)$  to be the  $[np]$ th smallest simulated value, we detail and use another convention in §5.1.6.

### Conditional Monte Carlo and Latin Hypercube Sampling

The conditional Monte Carlo (CMC) method (Hammersley, 1956) reduces the variance of the Monte Carlo estimator of  $F$  by replacing the terms in the sum of (2.2.1) by the probability of  $X \leq x$  conditional on some auxiliary random variable,  $Y$ , which is easily observed and provides information on  $X$ :

$$\hat{F}_{CMC} = \frac{1}{n} \sum_{i=1}^n g(x, Y_i),$$

where  $g(x, Y) = \mathbb{P}(X \leq x|Y)$  and  $\{Y_i\}_{i=1}^n$  are  $n$  independent and identically distributed (iid) random variables. By variance decomposition we have

$$\begin{aligned} \text{Var}(\mathbb{1}_{\{X \leq x\}}) &= \text{Var}(\mathbb{E}[\mathbb{1}_{\{X \leq x\}}|Y]) + \mathbb{E}[\text{Var}(\mathbb{1}_{\{X \leq x\}}|Y)] \\ &\geq \text{Var}(\mathbb{E}[\mathbb{1}_{\{X \leq x\}}|Y]) = \text{Var}(\mathbb{P}(X \leq x|Y)), \end{aligned}$$

so  $\text{Var}(\hat{F}_{CMC}) = \frac{1}{n} \text{Var}(g(x, Y)) \leq \frac{1}{n} \text{Var}(\mathbb{1}_{\{X \leq x\}}) = \text{Var}(\hat{F}_{MC})$ . Asmussen (2018) discuss CMC in insurance setting for the distribution of the sum of iid random variables with conditioning on the random variables forming the partial sums. Nakayama (2007) discusses the conditional Monte Carlo method for quantile estimation of the sum of iid random variables with conditioning on some  $Y$  where the joint distribution of  $(X, Y)$  is bivariate normal.

Dong and Nakayama (2017) improve upon the CMC method by combining it with *Latin Hypercube Sampling* (of the auxiliary random variable) rather than simple random sampling. Latin Hypercube sampling is a variance reduction method which extends stratified sampling to high dimensions so the sample space is better explored in some sense. This is easiest to visualise if we assume that the random variable of interest,  $X$  ( $Y$  for CMC), can be written as a function of standard uniform random variables,  $\mathbf{U}$ . Then in Latin Hypercube sampling we essentially split the sample space of  $\mathbf{U}$  into a finite grid and the realisations of  $\mathbf{U}$  are spread such that (in two dimensions for ease of explanation) there is one realisation in each row and column of the grid. Avramidis and Wilson (1998) show that Latin Hypercube Sampling reduces the variance of the Monte Carlo estimator compared to simple random sampling and

in a similar manner Dong and Nakayama (2017) show the combination of Latin Hypercube sampling and CMC reduces the variance of the CMC estimator.

### Importance sampling

An important example of a variance reduction method is *Importance Sampling*. This involves sampling from a different distribution,  $q(x)$ , to the distribution of interest (the target distribution),  $f(x)$ , specifically from one which generates samples more frequently in the region of interest (the importance region). That is, rather than simply using the brute force of Monte Carlo simulation, the simulation is designed such that there is more concentration on the values we are more interested in. Therefore, less computational time is needed to reach the same level of accuracy of the estimate as compared to the Monte Carlo estimate.

In the importance sampling method the distribution from which we sample is termed the *proposal* distribution and denoted by  $q(x)$ . Observe that:

$$\mathbb{E}_f [g(X)] = \int g(x)f(x)dx = \int \frac{g(x)f(x)}{q(x)}q(x)dx = \mathbb{E}_q \left[ g(X)\frac{f(X)}{q(X)} \right].$$

This gives rise to the following unbiased estimator:

$$\hat{I}_{IS} = \frac{1}{n} \sum_{i=1}^n g(x_i)w(x_i), \quad (2.2.2)$$

where  $(x_1, \dots, x_n)$  are independent samples from  $q(x)$  and we define  $w(x) = \frac{f(x)}{q(x)}$  to be the importance weight. A good choice of proposal,  $q(x)$ , would be one which is easy to sample from and preferably have some nice density form so any subsequent calculations are not too difficult. More desirable properties of  $q(x)$  are discussed later.

Here we also note that sometimes it is more convenient to use an alternative estimator using normalised weights:

$$\tilde{I}_{IS} = \frac{\sum_{i=1}^n g(x_i)w(x_i)}{\sum_{i=1}^n w(x_i)}.$$

This estimator is biased; however, this bias decreases with increasing sample size.

The variance of the unbiased estimator (2.2.2) is:

$$\text{Var} \left( \hat{I}_{IS} \right) = \text{Var} \left( \frac{1}{n} \sum_{i=1}^n g(X_i)w(X_i) \right)$$

$$\begin{aligned}
&= \frac{1}{n^2} \sum_{i=1}^n \text{Var} (g(X)w(X)) \\
&= \frac{1}{n} \left( \mathbb{E}_q [g(X)^2 w(X)^2] - \mathbb{E}_q [g(X)w(X)]^2 \right) \\
&= \frac{1}{n} \left( \mathbb{E}_f [g(X)^2 w(X)] - \mathbb{E}_f [g(X)]^2 \right) \\
&= \frac{1}{n} \left( \int \frac{g(x)^2 f(x)}{q(x)} f(x) - I^2 q(x) dx \right) \\
&= \frac{1}{n} \int \frac{g(x)^2 f(x)^2 - q^2(x) I^2}{q(x)} dx. \tag{2.2.3}
\end{aligned}$$

So the variance of the estimator compared to the Monte Carlo estimate is reduced if:

$$\begin{aligned}
\text{Var} \left( \hat{I}_{MC} \right) - \text{Var} \left( \hat{I}_{IS} \right) &= \frac{1}{n} \left( \mathbb{E}_f [g(X)^2] - \mathbb{E}_f [g(X)^2 w(X)] \right) \\
&= \frac{1}{n} \int g(x)^2 (1 - w(x)) f(x) dx > 0.
\end{aligned}$$

Logically, the proposal must ‘cover’ the target distribution, that is, any point which can be sampled from the target can be sampled from the proposal. This necessity is also clear since the proposal distribution appears on the denominator of the estimator (and the variance of the estimator) thus, for all  $x$ ,  $q(x) > 0$  if  $f(x) \neq 0$ . By considering the right hand side of equation (2.2.3) we see that the variance of the estimate can become large when  $q(x)$  is close to zero and cause problems in the tails. To handle this we need to ensure the tails of the proposal are heavier than the tails of the target distribution,  $f(x)$ . We also note that the variance of  $\hat{I}_{IS}$  is small when the numerator is close to 0. This occurs when  $q(x) \propto g(x)f(x)$  and, moreover, if

$$q(x) = \frac{g(x)f(x)}{I},$$

then  $q(x)$  is the optimal proposal since then the variance of the estimate would be zero. Of course, this optimal proposal cannot be used in practice since when inserted into Equation 2.2.2 the importance sampling estimate,  $\hat{I}_{IS}$ , becomes the true unknown value,  $I$ . Nonetheless this knowledge is useful in determining ‘good’ proposals.

Finally we remark that the importance weights can provide us with information about the efficiency of the sampling procedure. For example, if one weight is much larger than the others then the sample is essentially equivalent to just one independent sample. A measure

of the efficiency of the scheme is the effective sample size (Robert and Casella, 2004):

$$ESS = \frac{(\sum_{i=1}^n w_i)^2}{\sum_{i=1}^n w_i^2}.$$

Beck and Zuev (2005) apply and compare Monte Carlo simulation and importance sampling with a range of proposals. Much like ourselves they were concerned with rare events, in particular estimating the probability of rare events. They consider discrete-time models of dynamic systems which have some stochastic input and define rare events as the set of inputs for which a function of the outputs exceeds some quantity. They find that the choice of proposal distribution can greatly affect the estimate and that the variance of the estimator becomes worse as the dimension of the problem increases. They also describe and apply a splitting method, which is the focus of the next section.

One strategy for choosing the proposal distribution is *exponential tilting* or *exponential twisting* – this was first used in the importance sampling context by Siegmund (1976). For exponential tilting, proposals are of the form:

$$q_{\theta}(\mathbf{x}) = \frac{\exp(\theta^T \mathbf{x})}{\mathbb{E}[\exp(\theta^T \mathbf{X})]} f(\mathbf{x}), \quad (2.2.4)$$

for some  $\theta \in \mathbb{R}_d$  where  $f(\mathbf{x})$  is the true (target) distribution. The ‘best’ proposal is (2.2.4) with  $\theta$  chosen such that the variance of the estimator,  $\hat{I}_{IS}^{\theta}$ , is reduced. Exponential tilting/twisting has been extensively explored in the area of rare-event probability estimation (*e.g.*, Sadowsky (1993); Ridder and Rubinstein (2007); Dieker and Mandjes (2005)). This method is particular ‘nice’ when  $f(\mathbf{x})$  is a member of the exponential family since then the proposal is also a member of the exponential family and the weights and resulting importance sampling estimate have a simple form. Asmussen et al. (2016) develop an exponential-tilting importance-sampling estimator for the left tail probability of the sum of iid lognormals through consideration of the exponential family. However in general the proposal  $q_{\theta}(\mathbf{x})$  may not be straight-forward and  $\theta$  often needs to be numerically optimised. These issues were raised by Ben Rached et al. (2021) in the context of estimation of the left tail of the sum of iid random variables. They propose an alternative estimator, using the Gamma distribution with suitably chosen parameters as the proposal distribution, which was shown to have similarly ‘good’ performance compared to Exponential twisting while circumventing the issues above. McLeish and Men (2015) argue and demonstrate that a proposal distribution

from the family of generalised extreme value distributions is advantageous (at least in the 1-dimensional setting) over exponential tilting.

### 2.2.3 Splitting

The notion of exploring the region of interest more in order to improve efficiency is also the basis of so-called splitting methods. The idea is that if one has multiple simulations of a stochastic process over time the trajectories of some simulations are more likely to enter the region of interest than others. These more promising simulations are replicated at a particular time and this process is repeated many times. An estimate of the rare event probability is similar to that for importance sampling in that it is a weighted version of the Monte Carlo estimate.

The occurrence of the rare event of interest is equivalent to the process entering some set  $A$  before a more likely set,  $B$ . We consider a nested sequence of subsets corresponding to increasingly likely events:

$$A_1 \supset A_2 \supset \dots \supset A_N = A.$$

The probability of the rare event,  $A$ , then becomes:

$$p^A := \mathbb{P}(A) = \mathbb{P}(A_N|A_{N-1}) \dots \mathbb{P}(A_2|A_1) \mathbb{P}(A_1),$$

which should be easier to calculate since larger probabilities are easier to estimate.

In the simplest one-dimensional case we consider the rare event occurring when the process exceeds some threshold,  $M$ . We define a range of increasing thresholds,  $M_1 < M_2 < \dots < M_N = M$ , where exceedance of  $M_i$  corresponds to event  $A_i$ . Cérou and Guyader (2007) work in this setting and give a description of classical multilevel splitting. Consider  $n$  simulations of the process,  $X(t)$ , for  $t = 0, \dots, T$ , starting at some initial value  $x_0$ . All those simulations that reach the first threshold before entering the more likely set  $B$  *split* - that is the process is replicated some chosen  $r$  times up to the point the threshold was reached and from that point onwards is simulated with the threshold as it's initial value. This method is then repeated for a higher threshold and so on until the threshold of interest,  $M$ , is reached (after  $N$  iterations). In this way we are creating more processes that are likely to reach the rare threshold we are interested in. The resulting estimate of  $p^A$  is:

$$\hat{p}^A = \frac{n_M}{n} r^{-N}, \quad (2.2.5)$$

where  $r$  is the number of replicas at each iteration,  $n_M$  is the number of simulations which have exceeded the final threshold  $M$  and  $N$  is the number of iterations. This estimate is simply the usual Monte Carlo estimate ( $n_M/n$ ) with a weight to account for the way the processes are sampled to be more likely to exceed  $M$  before entering set  $B$ . For example, if  $N = 2$  and  $r = 2$ , the processes with high trajectories are duplicated at the first iteration and so we would expect the probability of processes exceeding  $M_2 = M$  to be doubled – this is accounted for in the estimate by dividing by 2.

As noted by Glasserman et al. (1999) the issues with the classical approach are the choice of thresholds and the number of replicas at each iteration. Glasserman et al. (1999) focus on the second issue whereas Cérou and Guyader (2007) present an algorithm which chooses thresholds adaptively. Such algorithms are called *Adaptive Multilevel Splitting* (AMS) algorithms. Cérou and Guyader’s method is to sort the sample by the largest value attained by each simulation before entering set  $B$  and keep the  $k$  largest of these. The  $k$ th largest value attained is then taken to be the initial value for  $n - k$  new simulations. This process is then repeated with the kept and new simulations ( $k + (n - k) = n$  simulations) until the  $k$ th largest value is greater than the threshold of interest,  $M$ . The estimate of  $p^A$  is then:

$$\hat{p}^A = \frac{n_M}{n} \left( \frac{k}{n} \right)^N. \quad (2.2.6)$$

Similar to the classical splitting estimator (2.2.5) this estimator is a weighted version of the Monte Carlo estimator; the weights here are determined by the proportion of processes kept at each iteration.

AMS was also used by Beck and Zuev (2005) in a similar manner to Cérou and Guyader (2007), specifying the probability of reaching the next subset,  $p_0$ , rather than the number of simulations kept at each iteration,  $k$ . They confirm that the splitting strategy is much more efficient at estimating the probabilities of rare events than the standard Monte Carlo method and note that, unlike an importance sampling estimator, the splitting estimator does not deteriorate with increasing dimension.

## Part I

# Investigating Extreme Values under Stopping Rules

## Chapter 3

# Inference for extreme values under threshold-based stopping rules

### 3.1 Introduction

The UK currently spends £400-500M per year on coastal and river flood defence infrastructure, with 2 million properties exposed to the risk of flooding (Environmental Agency, 2020). The agencies responsible for this spend monitor the effectiveness of their investment at giving the level of protection expected. After major flooding events renewed analysis is performed to assess both existing flood defences and the cost benefit of potential new schemes, proposed in response to the flooding.

Statistical extreme value methods, with likelihood-based inference, have proved a core component of the required analysis in terms of minimising the costs without jeopardising the level of accepted risk, and hence have financial and societal benefits. However, there is a problem with using these methods when the statistical analysis has been prompted by the occurrence of a recent large event, since in this case the data-set size itself is also random. This can lead to substantially biased inference and poor coverage properties and so result in inefficient flood-defence designs. Omitting the new extreme data value from the data set also seems unsuitable, as intuition suggests that flood risk will then be underestimated; moreover it would appear perverse to flood management agencies to ignore events of the type most relevant to the design specification.

In this chapter we aim to identify the extent of the inference problems when an analysis has been triggered by a large event and to develop new conditional-likelihood methods which appear to overcome these problems. We do not *suggest* when the timing of the data analysis should take place but study the analysis given that its timing has been determined by a data-dependent decision making process.

We consider modelling the extreme events of a time series of independent and identically distributed (iid) random variables  $X_1, X_2, \dots$ . The classical approach to do this is to split the time series into blocks of equal size (often a year) and to model the maxima of these blocks. Normalisation is necessary since as the block size tends to infinity the distribution of the maxima degenerates to a point mass at the upper end point of the distribution of  $X$ . The generalised extreme value (GEV) distribution (Coles, 2001) is the only non-degenerate limiting distribution of the normalised maxima as the block size tends to infinity. The GEV has distribution function:

$$G(x) = \exp \left( - \left[ 1 + \xi \left( \frac{x - \mu}{\sigma} \right) \right]_+^{-\frac{1}{\xi}} \right), \quad (3.1.1)$$

where  $\mu, \sigma > 0$  and  $\xi$  are the location, scale and shape parameters respectively and  $[z]_+ = \max(z, 0)$ . The shape parameter determines the behaviour of the upper tail of the distribution: for  $\xi < 0$  the distribution has an upper end point, for  $\xi = 0$  the tail is exponential and for  $\xi > 0$  the distribution has a power-decaying tail. There is particular interest in the occurrence of extreme events and so an important part of the analysis is the estimation of return levels (quantiles). Under stationarity, the  $y$ -year return-level,  $x_y$ , is the value which is exceeded on average once every  $y$  years. For the GEV distribution this can be calculated as:

$$x_y = G^{-1} \left( 1 - \frac{1}{y} \right) = \begin{cases} \mu - \frac{\sigma}{\xi} \{ 1 - [-\log(1 - 1/y)]^{-\xi} \} & \xi \neq 0 \\ \mu - \sigma \log[-\log(1 - 1/y)] & \xi = 0. \end{cases} \quad (3.1.2)$$

One can also consider modelling daily observations above some high threshold (rather than just modelling the block maxima) by the asymptotically justified generalised Pareto distribution (GPD) (Davison and Smith, 1990). Threshold methods typically benefit from using more extreme-value data and hence are more efficient in their inferences than block maxima methods (Coles, 2001). We focus most of our analysis and developments on the GEV case, as similar benefits are found for both GEV and GPD inference, but with the GPD also sensitive to threshold choice. We illustrate some GPD results in the supplementary material to Barlow et al. (2020).

We consider the analysis of annual maxima of daily peak river flow data obtained from UK CEH (2018) for the Lune at Caton, just outside Lancaster, from 1968 to 2015 (Figure 3.1.1, left panel) and illustrate the inference issues due to the timing of analysis being determined by the occurrence of a flood event. Under the assumption that the annual maxima

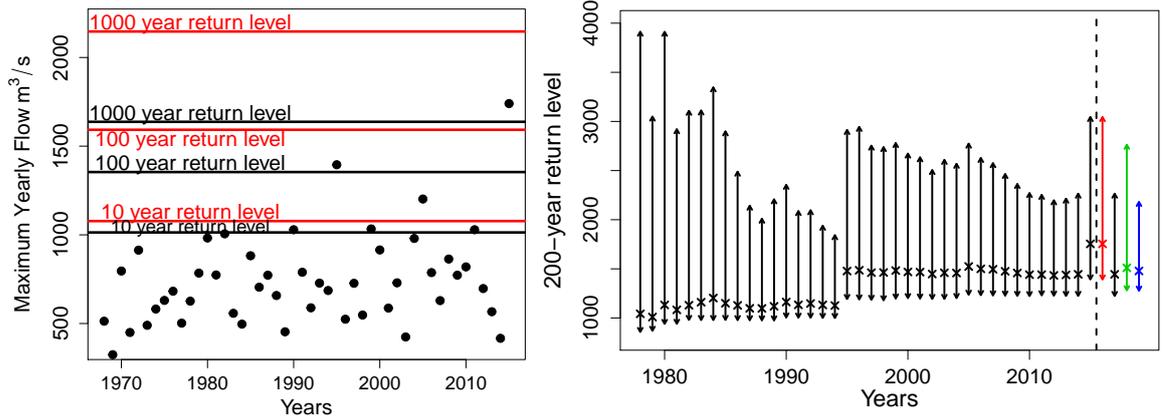


Figure 3.1.1: *Left: The annual maxima of daily peak river flow data for the Lune at Caton with return level estimates before (black) and after (red) the 2015 flood. Right: 200-year return-level estimates based on all the data up to and including the current year for the Lune at Caton with 95% profile likelihood-based confidence intervals. The four 200-year return-level estimates and associated 95% confidence intervals to the right of the vertical dotted line are our new estimates that aim to address a fixed-threshold stopping rule of  $c_k = 1568$  based on the all the data up to and including 2015: standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue).*

are independent and identically distributed (i.i.d.) we can fit the GEV distribution to the annual maxima using likelihood-based inference (the likelihood is  $\prod_{i=1}^n g(x_i; (\mu, \sigma, \xi))$  where  $n$  is the sample size and  $g$  is the density,  $g = dG/dx$ ), and estimate return levels using (3.1.2). The estimated 10, 100 and 1000-year return levels are shown in Figure 3.1.1 (left panel) for the data up to 2014 (black) and including 2015 (red). The December 2015 floods resulted in the river Lune recording the highest peak river flow ( $1740 \text{ m}^3/\text{s}$ ) of all UK rivers over all years of records. This value is higher than the 1000-year return-level estimate based on the observations up to 2014. However, once the 2015 event is included in the analysis the return-level estimates become much higher. If we were to take these 2015 point estimates as the truth we would expect to observe an event as extreme as that in 2015 approximately once every 200 years. For design purposes this level of sensitivity is highly undesirable, as the costs for flood protection would change dramatically.

Figure 3.1.1 (right panel) shows a reanalysis of all data available at each year between 1978 and 2015. It provides the point estimate and profile likelihood-based 95% confidence interval of the 200-year return-level, as it would have been produced in that year. The four additional point estimates and confidence intervals to the right of the vertical dotted line correspond to estimators introduced in §3.3 and their corresponding profile likelihood-based confidence intervals. At the beginning of the data collection the return-level estimates vary

considerably, but they become more stable as the number of years increases, with the width of the confidence intervals generally decreasing over time. However, even after many years of data collection, the largest events can be seen to cause sharp increases in the estimates and their associated uncertainty. For example, the return-level estimate following the January 1995 floods and the 2015 floods are larger than those of previous years.

This illustrative example is typical of when an analysis is performed immediately after a large event. Unless further analysis is undertaken it is unclear whether by analysing the data with the final extreme event we are introducing a positive bias into the inference. For example, the lower bound of the 95% confidence interval of the 200 year return-level after the 1995 event is larger than almost all previous point estimates - directly after the event (without the knowledge of later years) this could have been seen as an indication of positive bias in the standard estimator. However, after the 1995 event the return-level estimates were fairly stable and larger than those before 1995, so it would seem the standard estimator for 1995 may not have been overestimating and before the event the shape parameter was estimated too low.

An alternative approach is to simply ignore the most recent year of data when an analysis has been requested because we have large observations in that year, in which case the return-level estimate is lower and the confidence interval is narrower - in particular the upper bound is lower. However, we speculate (see also §3.2.3) that this estimator is now negatively biased due to the loss of information about the extreme event. Moreover the estimator is inefficient since the larger data values are the most informative about the upper tail (Davison and Smith, 1990). Finally, it would be hard to convince practitioners to exclude the largest events; for example, an event may be observed which is larger than the upper end point estimated from previous data, in which case it would be perverse not to make some update to the previously estimated return levels.

The key issue that the Lune example illustrates is that when meeting the flood management agencies' needs, the time to undertake the extreme value analysis is stochastic and triggered by a large event. Thus, there is effectively some form of unwritten stopping rule, determined by the flood management agencies, which determines the timing of the analysis. In contrast with a standard iid sample of fixed size, when we use a stopping rule the time at which we stop (the sample size) is variable, we denote this by  $N$ .

One can attempt to mathematically formulate the characteristics of the stopping rule,

though in reality a precise mathematical rule does not exist. The stopping decision may depend on (i) some absolute threshold, such as the height of existing flood defences or a critical level which when exceeded leads to severe flooding, or (ii) an assessment, based upon all observations to date, of what might constitute an ‘exceptional’ event. We consider two simple stopping rules based on a series of iid random variables,  $X_1, X_2, \dots$  which, in a sense, bracket this range of possibilities and we discuss other possibilities in §3.6.

### 1. Fixed-threshold stopping rule

Stop when an observation exceeds a specified value,  $c_k$ , *i.e.*:

$$N = \inf\{n \in \mathbb{N} : X_n > c_k\}, \quad (3.1.3)$$

where  $k$  is the true (but unknown) return period of  $c_k$ .

### 2. Variable-threshold stopping rule

Stop when an observation exceeds the return-level estimate,  $\hat{x}_k$ , corresponding to a chosen fixed return period of  $k$  years, calculated using previous observed values, *i.e.*, when:

$$N = \inf\{n \in \mathbb{N} : X_n > \hat{x}_k(X_1, \dots, X_{n-1})\}. \quad (3.1.4)$$

We do not suggest the stopping rule to use but study the analysis given that its timing has been determined by a stopping rule. As far as we are aware, there has been no study of stopping rules and their effects on likelihood estimation in the extreme-value setting.

Using a stopping rule to determine the sample size can lead to estimators, such as the maximum likelihood estimator (MLE), having different sampling properties to the fixed-sample case. To illustrate this feature consider an iid sequence of Bernoulli random variables,  $Y_1, Y_2, \dots$ , each with probability of success of  $\theta$ . If one fixes the number of trials,  $n$ , the number of successes,  $R$ , in these trials is binomially distributed and  $\hat{\theta} = R/n =: \hat{\theta}_1$ ; whereas if the number of successes is fixed as  $r$ , the number of trials,  $N$ , is negative-binomially distributed and  $\hat{\theta} = r/N =: \hat{\theta}_2$ . In both cases the MLE of the probability of success is the proportion of successes, however,  $\mathbb{E}[\hat{\theta}_1] = \theta$  whereas  $\mathbb{E}[\hat{\theta}_2] = r\mathbb{E}[1/N] \geq r/\mathbb{E}[N] = \theta$  by Jensen’s inequality. The presence of a stopping rule affects the performance of the estimator which motivates an investigation into the performance of return-level estimation under stopping

rules.

Testing the data against some ‘stopping criterion’ at regular intervals falls into the setting of sequential analysis, which has a rich literature covering applications from quality control (Wald, 2004), to clinical trials (Todd et al., 1996) and abundance modelling (Barry and Coggan, 2010). Many studies have considered the influence of such stopping rules on likelihood inference, *e.g.*, Barndorff-Nielsen and Cox (1984) consider the distribution of the likelihood-ratio statistic under different stopping rules for systems with Brownian motion and Poisson processes, and in the clinical trial setting Whitehead (1986) derives an expression for the bias of the MLE of the treatment effect tested under a sequential probability ratio test. Some papers compare the bias under different experimental designs or stopping criteria (*e.g.*, Bauer et al. (2010)). Cox (1952), Whitehead (1986) and Stallard and Todd (2005) propose bias-reduced estimators by approximating the bias and subtracting this from the usual estimate. One such approach uses an iterative method corresponding to a bootstrap bias correction (Efron, 1990).

Kenward and Molenberghs (1998) consider iid sampling from a Normal distribution using a deterministic stopping rule and study the estimation of the mean parameter of this Normal distribution. Molenberghs et al. (2014) extend this setting to the use of a probabilistic stopping rule. They note that an unbiased estimator of the mean parameter can be obtained from the conditional likelihood (we derive such estimators in §3.3) however, at the cost of an increased mean squared error (MSE) in comparison to the MSE of the sample average (the standard estimator if the sample size was fixed). The increased variance of a bias-reduced estimator appears to be an issue for many of the proposed bias-reduction methods. For example, bias reduction using Rao-Blackwellisation (Bowden and Glimm, 2008) and shrinkage estimators (Carreras and Brannath, 2013) often have a worse MSE than the standard MLEs.

In §3.2 we introduce the notation used throughout the chapter and discuss likelihood inference under stopping rules. In §3.2.3 and §3.2.4 we discuss the bias under the fixed-threshold and variable-threshold stopping rules respectively and derive expressions for the bias when sampling from some simple distributions. We introduce two conditioning-based likelihood estimators in §3.3. In §3.4 we perform a simulation study for sampling from the GEV distribution using the two stopping rules and discuss the properties of the estimators in this setting. We apply our estimators to the Lune river flow data in §3.5 and discuss our conclusions, the practical usage of the methods and extensions in §3.6.

## 3.2 Inference under stopping rules

### 3.2.1 Introduction

Throughout this chapter we restrict our attention to sequences of iid observations arising from some distribution with a density of  $f(x; \theta)$ , where  $\theta$  is the parameter vector for the distribution. When a fixed number,  $n$ , of observations,  $x_1, \dots, x_n$ , is analysed, the likelihood and log-likelihood for the data are

$$L_{fixed}(\theta; \mathbf{x}_{1:n}) = \prod_{i=1}^n f(x_i; \theta) \quad \text{and} \quad \ell_{fixed}(\theta; \mathbf{x}_{1:n}) = \sum_{i=1}^n \log f(x_i; \theta),$$

where  $\mathbf{x}_{1:n}$  denotes the vector of observations  $(x_1, \dots, x_n)$ . In practice it is usual to assume the sample size,  $n$ , is fixed however, for us,  $n$ , is not fixed; we sample consecutively until some stopping criterion is met and denote the (random) time at which it is met by  $N$ . In this chapter we consider both the full data and excluding the last data point, for which the log-likelihoods are, up to an additive constant:

$$\ell_{std}(\theta; n, \mathbf{x}_{1:n}) = \ell_{fixed}(\theta; \mathbf{x}_{1:n}) \tag{3.2.1}$$

$$\ell_{ex}(\theta; n, \mathbf{x}_{1:n}) = \ell_{std}(\theta; n, \mathbf{x}_{1:n}) - \log f(x_n; \theta). \tag{3.2.2}$$

In §3.2.2 we reproduce the proof that the likelihood for the data  $(n, \mathbf{x}_{1:n})$  is  $L_{fixed}(\theta; n, \mathbf{x}_{1:n}) \propto L_{std}(\theta; \mathbf{x}_{1:n})$ . Given data  $(n, \mathbf{x}_{1:n})$  an estimate of the parameter vector is obtained by maximising the log likelihood:  $\hat{\theta}(n, \mathbf{x}_{1:n}) = \arg \max_{\theta} \ell(\theta)$ . When the nature of the data is clear we abbreviate this to  $\hat{\theta}$ , and depending on the likelihood used we have estimators  $\hat{\theta}_{std}$  or  $\hat{\theta}_{ex}$ .

In practice we would not consider estimating return levels (particularly for large return periods) from a sample of only a very small number of observations. However, the fixed-threshold stopping rules can result in samples of size 1, and this can lead to parameter identifiability issues for data sets simulated from the hypothesised data-generating mechanism. In reality, if an analysis has been requested then sufficient information would be available to derive a meaningful estimate. This information could be historical information, hydrological knowledge, data from other sites, or data at the current site collected before the instigation of a stopping rule. We call this the *historical data* and, for simplicity in this article, code the historical data as some number,  $n_0$  of data values collected before the stopping rule could be

invoked. Real decisions will incorporate this information, and our analysis should allow for this.

We are interested in a set of  $y$ -year return-levels,  $x_y(\theta)$  ( $y \in \mathcal{Y}$ ), for some set  $\mathcal{Y}$ , such as  $\{50, 200, 1000\}$ . In particular, we wish to understand the behaviour of the estimators  $x_y(\hat{\theta}(N, \mathbf{X}_{1:N}))$  (with  $x_y$  given by expression (3.1.2) for GEV sampling) when the dataset arises from a stopping rule. In this section we focus on the relative bias, and in §3.4 we look at other properties including the relative root-mean-squared error, given respectively by:

$$\text{RelBias}(\hat{x}_y) = \frac{1}{x_y(\theta)} \mathbb{E} \left[ x_y(\hat{\theta}(N, \mathbf{X}_{1:N})) \right] - 1 \quad (3.2.3)$$

$$\text{RRMSE}(\hat{x}_y) = \frac{1}{x_y(\theta)} \sqrt{\mathbb{E} \left[ \{x_y(\hat{\theta}(N, \mathbf{X}_{1:N})) - x_y(\theta)\}^2 \right]}, \quad (3.2.4)$$

where  $x_y(\theta)$  is the true  $y$ -year return-level.

In §3.2.2 we detail a well-known result that the likelihood for the data  $(n, \mathbf{x}_{1:n})$  with a random stopping time is the same as for data  $\mathbf{x}_{1:n}$  with  $n$  fixed. However, the properties of the estimator, such as its bias and variance as well as the coverage of any confidence interval, may be influenced by the different data-generating mechanism.

The properties of likelihood-based estimators of tail quantiles under our stopping rules are intractable for data arising from the GEV or GPD distributions. However, for a particular special case of the GPD, the exponential distribution, certain properties are tractable and this provides insight into the behaviour observed in the simulation studies of §3.4 for the GEV. Specifically, in §3.2.3 we derive the bias in quantile estimates for exponential data under the fixed-threshold stopping rule, and in §3.2.4 show that, under the variable-threshold stopping rule, quantile estimates for gamma data (including the exponential as a special case) with a known shape parameter are unbiased.

### 3.2.2 Likelihood in presence of a stopping rule

Now, following Pawitan (2013), we derive the true likelihood for the data sampled using a general stopping rule which is a function of the data and not the unknown parameter vector. We define a stopping region  $\mathcal{S}_n = \mathcal{S}_n(\mathbf{x}_{1:n-1})$  such that we stop sampling if  $X_n \in \mathcal{S}_n$  and continue to sample otherwise. We abbreviate  $\mathbb{P}(X_i \in \mathcal{S}_i)$  by  $p_i$  and we let  $f_{X_i|\mathcal{S}_i}$  and  $f_{X_i|\mathcal{S}_i^c}$  denote the densities of  $X_i$  conditional on  $X_i \in \mathcal{S}_i$  and  $X_i \in \mathcal{S}_i^c$ . The likelihood for the full data is

$$\begin{aligned}
L_{std}(\theta; n, \mathbf{x}_{1:n}) &= \mathbb{P}(N = n, \mathbf{X}_{1:n} = \mathbf{x}_{1:n} | \theta) \\
&= \mathbb{P}(N = n) \mathbb{P}(\mathbf{X}_{1:n} = \mathbf{x}_{1:n} | N = n, \theta) \\
&= p_n \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=1}^{n-1} (1 - p_i) \mathbb{1}_{\mathcal{S}_i^c}(x_i) \times f_{X_n | \mathcal{S}_n}(x_n | \mathcal{S}_n) \prod_{i=1}^{n-1} f_{X_i | \mathcal{S}_i^c}(x_i | \mathcal{S}_i) \quad (3.2.5) \\
&= L_{fixed}(\theta; \mathbf{x}_{1:n}) \times \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=1}^{n-1} \mathbb{1}_{\mathcal{S}_i^c}(x_i) \\
&\propto L_{fixed}(\theta; \mathbf{x}_{1:n}).
\end{aligned}$$

The logic here is that to have a sample of size  $n$  the final observation must be in the stopping region and all other observations outside their respective stopping regions, hence  $\mathbb{P}(N = n)$  includes indicator functions of the observations being in the correct sets. The last step follows since the indicator functions do not depend on the unknown parameter,  $\theta$ , and so are absorbed into the proportionality constant. Thus inference purely from the likelihood leads to the same conclusions whether we have a random sample size according to some stopping rule or a fixed sample size. In particular, the MLE,  $\hat{\theta}$ , and the observed Fisher information are the same in both cases. However, the properties of the estimators are different since the distribution of  $\{N, X_1, \dots, X_N\}$  is different to the distribution of  $\{X_1, \dots, X_n\}$  for some fixed  $n$ . In particular, estimators obtained from  $L_{std}$  can be biased even when estimators from  $L_{fixed}$  are unbiased, as seen in §3.1 for Bernoulli sampling.

### 3.2.3 Fixed-threshold stopping rule with exponential observations

Let  $X_i$  have an exponential distribution with an unknown rate parameter of  $\beta$ , which is a special case of the GPD used to model the tails of a distribution and is given by expression (2.1.4) with  $\xi = 0$  and  $\sigma = \beta^{-1}$ . The  $y$ -observation return level is  $x_y = (\log y)/\beta$  and, since this is proportional to  $1/\beta$ , the relative bias is  $\beta/\hat{\beta} - 1$  whatever the value of  $y$ . The MLE of  $\beta^{-1}$  for a sample of size  $n$ , whether fixed or random is simply  $\bar{x}$ , where  $\bar{x}$  is the sample mean. When  $n$  is fixed, the MLE,  $(\hat{\beta}_{fixed})^{-1} = \bar{X}_n$ , is an unbiased estimator of  $1/\beta$ ; however with the fixed-threshold stopping rule  $N$  follows a geometric distribution where  $1/k$  is the probability of a ‘success’ *i.e.*, an exceedance. The geometric distribution is a special case of the negative-binomial distribution and so we know the estimator of the probability of exceedance of a fixed threshold is positively biased (§3.1 under (3.1.4)). Now  $(\hat{\beta}_{std})^{-1} = \bar{X}_N$  and, similarly,

the MLE when excluding the final observation is  $(\widehat{\beta}_{ex})^{-1} = \bar{X}_{N-1} = \sum_{i=1}^{N-1} X_i / (N-1)$ . It is straightforward to show (see Appendix A.1) the following.

**Proposition 3.2.3.1.** *Let  $X_1, X_2, \dots$  be a sequence of iid random variables with  $X_i \sim \text{Exp}(\beta)$ . Let  $N$  arise from the fixed-threshold stopping rule (3.1.3) giving data  $(N; \mathbf{X}_{1:N})$ . Let  $\hat{x}_y^{std} = x_y(\widehat{\beta}_{std}(N; \mathbf{X}_{1:N}))$  be the estimator of the  $(1 - \frac{1}{y})$ th quantile (equivalently the  $y$ -observation return-level) obtained from the MLE for  $\beta$  from the full likelihood and let  $\hat{x}_y^{ex} = x_y(\widehat{\beta}_{ex}(N; \mathbf{X}_{1:N}))$  be the estimator from the likelihood excluding the final observation.*

*Then the relative biases of the return-level estimators are*

$$\begin{aligned} \frac{1}{x_y} \mathbb{E} [\hat{x}_y^{std}] - 1 &= \frac{\beta c_k}{e^{\beta c_k} - 1} \left( \frac{\beta c_k}{1 - e^{-\beta c_k}} - 1 \right) \\ \frac{1}{x_y} \mathbb{E} [\hat{x}_y^{ex} | N > 1] - 1 &= -\frac{\beta c_k}{e^{\beta c_k} - 1}. \end{aligned}$$

In Proposition 3.2.3.1, when excluding both the final observation (and the fact that it *is* the final observation), when  $N = 1$  the MLE is undefined since there are no data;  $x_1$  is unknown and the fact that  $N$  would be greater than zero was known before the data-collection process began; we therefore condition on  $N > 1$ .

Proposition 3.2.3.1 shows that the estimator of any return level using the full likelihood is always positively biased, whereas if the final observation is omitted the estimator of any return level is always negatively biased. The final data observation is the largest and has been shown by Davison and Smith (1990) to be the most influential on the MLE fit so when this value, together with the information that it exceeded the threshold, is omitted from the dataset this changes the bias and, potentially, also the variance of the return-level estimator and risks being inefficient. Nevertheless, for thresholds with only a small chance of exceedance, *i.e.*, large values of  $\beta c_k$ ,  $\text{RelBias}(\hat{x}_y^{std}) \sim (\beta c_k)^2 \exp(-\beta c_k)$ , whereas  $\text{RelBias}(\hat{x}_y^{ex}) \sim -\beta c_k \exp(-\beta c_k)$ , that is, the bias is a factor  $(\beta c_k)^{-1}$  smaller for estimates where the final observation is ignored. The higher the threshold, the larger the typical data set that is generated before the stopping criterion is met and the less biased the estimate of any return level.

In Figure 3.2.1 we compare the relative bias of the estimates of  $\beta^{-1}$  (and hence also for the return-level estimates) both when including and excluding the final observation when varying  $k$ , the true return period of the stopping threshold,  $c_k$ . The two additional curves correspond to estimators that will be introduced in §3.3. The maximum relative bias in the standard

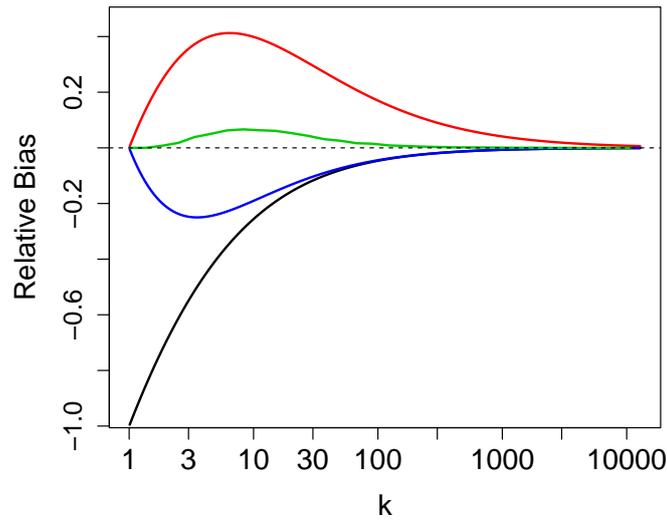


Figure 3.2.1: *Relative bias of the return-level estimates against the return period,  $k$ , of the fixed-threshold stopping rule  $c_k$  when sampling from the Exponential distribution with the fixed-threshold stopping rule using: standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue). The latter two methods are introduced in §3.3.1*

return-level estimator is 0.4; *i.e.*, the estimator is around 1.4 times the true value. This occurs for a threshold corresponding to  $k \approx 7$ , *i.e.*, when we stop sampling if an observation exceeds the 7-observation return level. Clearly this will generally result in a very small sample so we would expect return-level estimates to also be highly variable in this case.

### 3.2.4 Variable-threshold stopping rule with gamma observations

The positive bias in return-level estimates that arises from the fixed stopping rule is partly a result of the geometric distribution of  $N$  (see §3.2.3). For the variable-threshold rule  $N$  no longer has a geometric distribution and we find empirically for the GEV (see §3.4.3) that the bias is typically reduced; as we now show, at least for one parametric family of distributions, the bias disappears entirely.

Let  $X_i \sim \text{Exp}(\beta)$ , with unknown rate parameter,  $\beta > 0$ . In §3.2.1 we noted the need for a historical sample in practice; here, to reflect this, we suppose that the stopping rule is only implemented after an initial sample of independent  $\text{Exp}(\beta)$  variables,  $X_{-n_0}, \dots, X_{-1}$ , whose mean is denoted by  $\bar{X}_0$ , with  $\bar{X}_0 \sim \text{Gamma}(n_0, n_0\beta)$ .

As noted earlier, the return level is proportional to  $\beta^{-1}$  and the MLE for  $\beta^{-1}$ , from the full likelihood, is  $\bar{x}$ . Thus, for some constant of proportionality  $\gamma$  (depending on the return

period,  $y$ ), the variable-threshold stopping rule is equivalent to

$$N = \inf\{n \geq 1 : X_n > \gamma \bar{X}_{n-1}\}, \quad (3.2.6)$$

where

$$\bar{X}_k = \frac{n_0 \bar{X}_0 + X_1 + \dots + X_k}{n_0 + k} \quad k \geq 1. \quad (3.2.7)$$

**Theorem 3.2.4.1.** *With  $N$ ,  $\mathbf{X}_{1:N}$  and  $\bar{X}_k$  as defined in (3.2.6) and (3.2.7), for all  $n \in \mathbb{N}$ :*

$$\bar{X}_N | N = n \stackrel{d}{=} \bar{X}_n \sim \text{Gamma}((n + n_0), \beta(n + n_0)).$$

From Theorem 3.2.4.1 we see that  $\mathbb{E}[\bar{X}_n | N = n] = 1/\beta$ , and hence:

**Corollary 3.2.4.2.** *For a sample obtained as in Theorem 3.2.4.1, the sample mean and  $y$ -year return-level estimate are unbiased:*

$$\begin{aligned} \mathbb{E}[\bar{X}_N] &= \frac{1}{\beta}. \\ \mathbb{E}[x_y(\hat{\beta}^{std}(N, \mathbf{X}_{1:N}))] &= x_y(\beta). \end{aligned}$$

Contrasting Corollary 3.2.4.2 with Proposition 3.2.3.1, both of which apply to the exponential distribution, we see that the standard estimator can be unbiased for the variable-threshold stopping rule even though it is strongly positively biased for the fixed-threshold stopping rule.

Theorem 3.2.4.1 and Corollary 3.2.4.2 can be extended to  $X_i \sim \text{Gamma}(\alpha, \beta)$  random variables, where the shape parameter,  $\alpha > 0$ , is known and the the rate parameter,  $\beta > 0$ , is unknown (and must be estimated). As with the exponential distribution, the return levels of the gamma distribution are proportional to  $\beta^{-1}$ , with the constant of proportionality depending on  $\alpha$ . Furthermore the MLE from the full likelihood satisfies  $\hat{\beta}^{-1} = \bar{x}/\alpha$ . So, the variable-threshold stopping rule is (3.2.6) with the constant of proportionality  $\gamma$  depending on  $\alpha$  as well as the return period,  $y$ . Theorem 3.2.4.1 and Corollary 3.2.4.2 are retrieved by setting  $\alpha = 1$ .

A proof for Theorem 3.2.4.1 in this more general case is provided in Appendix A.1.2.

### 3.3 Alternative Methods for Parameter Inference

#### 3.3.1 New conditional likelihoods

Motivated by the lack of bias in the stopping rule of Molenberghs et al. (2014), we propose a similar estimator for our scenarios by conditioning on the fact that only the final observation met the stopping criterion. The likelihood, therefore, consists of the conditional densities of the data values given that each of the first  $n - 1$  is outside its stopping region and the  $n$ th is inside its stopping region. Using the same notation as §3.2.2 the full-conditioning likelihood is, from (3.2.5),

$$\begin{aligned} L_{fc}(\theta; n, \mathbf{x}_{1:n}) &= \frac{\mathbb{P}(N = n, \mathbf{X}_{1:n} = \mathbf{x}_{1:n} | \theta)}{\mathbb{P}(X_n \in \mathcal{S}_n) \prod_{i=1}^{n-1} \mathbb{P}(X_i \in \mathcal{S}_i^c)} \\ &= f_{X_n | \mathcal{S}_n}(x_n | \mathcal{S}_n) \prod_{i=1}^{n-1} f_{X_i | \mathcal{S}_i^c}(x_i | \mathcal{S}_i) \times \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=1}^{n-1} \mathbb{1}_{\mathcal{S}_i^c}(x_i). \end{aligned}$$

The log likelihood,  $\ell_{fc}$ , is as given in (3.3.1) and the corresponding estimate is denoted  $\hat{\theta}_{fc}$ . For the fixed-threshold stopping rule this effectively conditions out the geometric distribution for  $N$  (§3.2.3); it might be hoped, therefore, that it might remove that part of the positive bias that is due to the randomness of  $N$ .

By conditioning on the final observation exceeding its stopping threshold and all other observations not exceeding theirs we are effectively losing all of this information which will lead to larger uncertainty in the estimates, *e.g.*, giving wider confidence intervals. Hence, we consider a further likelihood which conditions *only* on the fact that the final observation exceeds its threshold:

$$\frac{\mathbb{P}(N = n, \mathbf{X}_{1:n} = \mathbf{x}_{1:n} | N = n, \theta)}{\mathbb{P}(X_n \in \mathcal{S}_n)} = f_{X_n | \mathcal{S}_n}(x_n | \mathcal{S}_n) \prod_{i=1}^{n-1} f_{X_i}(x_i) \times \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=1}^{n-1} \mathbb{1}_{\mathcal{S}_i^c}(x_i).$$

As with full conditioning, this results in the stochasticity of  $N$  being less influential. We refer to this method as partial conditioning with log likelihood, denoted by  $\ell_{pc}$ , given in (3.3.2). The corresponding estimate is denoted by  $\hat{\theta}_{pc}$ .

In summary, the two new log likelihoods we consider are:

$$\ell_{fc}(\theta; n, \mathbf{x}_{1:n}) = \ell_{std}(\theta; n, \mathbf{x}_{1:n}) - \log \bar{F}(s_{k,n}; \theta) - \sum_{i=1}^{n-1} \log F(s_{k,i}; \theta) \quad (3.3.1)$$

$$\ell_{pc}(\theta; n, \mathbf{x}_{1:n}) = \ell_{std}(\theta; n, \mathbf{x}_{1:n}) - \log \bar{F}(s_{k,n}; \theta) \quad (3.3.2)$$

where  $s_{k,i}$  is the lower boundary of the stopping set for the  $i$ th observation; then for the variable-threshold stopping rule  $s_{k,i} = \hat{x}_k^{std}(\mathbf{x}_{1:i-1})$ , *i.e.*, it is the standard estimate of the  $k$ -year return-level using all the data up to and including the previous observation, and for the fixed-threshold stopping rule  $s_{k,i} = c_k$  for all  $i$ .

The examples of §3.2 have exponential tails with an unknown scale parameter. When data values are modelled using the GEV, uncertainty in the shape parameter,  $\xi$ , has a much larger impact on estimates of high quantiles than the uncertainty in the other two parameters,  $\mu$  and  $\sigma$  (Coles, 2001). So we now consider estimation of the shape parameter and high quantiles using the standard and partial-conditioning likelihoods. For simplicity we focus on an idealised scenario where we take  $\mu = 0$  and  $\sigma = 1$  as known, so  $X$  has a distribution function of  $F(x; \xi) = \exp\left(-[1 + \xi x]_+^{-1/\xi}\right)$  and a survival function of  $\bar{F} = 1 - F$ .

For quantile estimation the standard likelihood estimator of  $\xi$ , *i.e.*,  $\hat{\xi}_{std}$ , leads to a positive bias for high quantiles. This can be seen as follows. The  $y$ -year return level can be written as  $\bar{F}^{-1}(1/y; \xi) = [\exp(a_y \xi) - 1]/\xi$ , with  $a_y = -\log[-\log(1 - 1/y)]$ , where  $a_y \geq 0$  provided  $y \geq e/(e - 1) \approx 1.6$ . Return levels as low as 1.6 years are of no practical interest in our setting. When  $a_y > 0$ ,  $\bar{F}^{-1}$  is an increasing, convex function of  $\xi \in \mathbb{R}$ , so, whatever the likelihood, Jensen's inequality gives  $\mathbb{E}_\xi [\bar{F}^{-1}(1/y; \xi)] \geq \bar{F}^{-1}(1/y; \mathbb{E}_\xi [\xi])$ . The monotonicity of  $\bar{F}^{-1}$  implies that even if  $\hat{\xi}_{std}$  is unbiased, we should expect a positive bias in all quantile estimates, and this will only be exaggerated if (as we find in our stopping-rule simulations)  $\xi$  is positively biased.

This bias in the estimator for high quantiles is guaranteed to be less positive when using the partial-conditioning likelihood rather than the standard likelihood. To see this first note that  $\ell_{pc}(\xi) = \ell_{std}(\xi) - \log[\bar{F}(c; \xi)]$  where  $c$ , the stopping threshold, has been standardised. The resulting MLE,  $\hat{\xi}_{pc}$ , satisfies  $\ell(\hat{\xi}_{pc}) - \log[\bar{F}(c; \hat{\xi}_{pc})] > \ell(\xi) - \log[\bar{F}(c; \xi)] \forall \xi$ . Also,  $\bar{F}$  is an increasing function of  $\xi$  since

$$\frac{\partial}{\partial \xi} \log \{-\log F(x; \xi)\} = \frac{1}{\xi^2} \left\{ \log[1 + \xi x] - \frac{\xi x}{1 + \xi x} \right\} \geq 0, \quad \forall \xi x > -1.$$

So, as  $\ell(\hat{\xi}_{std}) > \ell(\xi) \forall \xi$ , it follows that

$$\begin{aligned} \ell(\hat{\xi}_{std}) - \log[\bar{F}(c; \hat{\xi}_{pc})] &> \ell(\hat{\xi}_{pc}) - \log[\bar{F}(c; \hat{\xi}_{pc})] > \ell(\hat{\xi}_{std}) - \log[\bar{F}(c; \hat{\xi}_{std})] \\ \Rightarrow -\log[\bar{F}(c; \hat{\xi}_{pc})] &> -\log[\bar{F}(c; \hat{\xi}_{std})] \\ \Rightarrow \hat{\xi}_{pc} &< \hat{\xi}_{std}. \end{aligned}$$

Thus, if the standard estimator is positively biased the partial-conditioning method will be less positively biased. Given that this effect is magnified for return levels, as shown above, we should expect improvements in return-level estimates using the partial-conditioning likelihood. An analogous argument to the above also applies to data modelled using the GPD, except that there is no restriction on  $y$ , and  $a_y = \log y$ .

### 3.3.2 Application to exponential observations

Consider iid sampling from the exponential distribution with rate parameter,  $\beta$ , using the fixed-threshold stopping rule. The relative bias for return-level estimators using the log-likelihoods (3.2.1) and (3.2.2), are detailed in Proposition 3.2.3.1 and plotted in Figure 3.2.1. Figure 3.2.1 also plots the relative bias for the likelihoods in (3.3.1) and (3.3.2), the latter has the form

$$\text{RelBias}(\hat{x}_y^{pc}) = \text{RelBias}(\hat{x}_y^{std}) - \frac{\beta^2 c_k^2}{e^{\beta c_k} - 1} = \frac{\beta c_k}{e^{\beta c_k} - 1} \left[ \frac{\beta c_k}{1 - e^{-\beta c_k}} - 1 - \beta c_k \right]. \quad (3.3.3)$$

Estimator  $\hat{x}_y^{pc}$  is negatively biased but the bias is smaller than that for  $\hat{x}_y^{ex}$ . The bias of  $\hat{x}_y^{pc}$  tends to 0 as  $c_k$  tends to infinity at the same fast rate as for  $\hat{x}_y^{ex}$  (§3.2.3).

We were unable to obtain a tractable expression for the bias of the full-conditional estimator. In Figure 3.2.1 this bias was found using Monte Carlo methods. The bias is very low and tends towards 0 much faster than any of the other estimators considered. This finding is similar to that of Molenberghs et al. (2014) for the mean of normally distributed observations with a probabilistic stopping rule; however, the MSE of the unbiased estimator was found to be poor compared to that of the standard estimator. In §3.4 we show that in our ‘extremes’ setting, the full-conditional MSE for a return level is often lower relative to the MSE of the standard estimator since the high variance of return-level estimators using the standard likelihood is in part due to the final observation being large. Furthermore in §3.4 we show that the partial-conditioning approach results in estimators with much reduced variance and that this leads to lower MSE compared to the standard likelihood approach.

## 3.4 Simulation results

In this section we focus on the return-level inference when sampling from the GEV distribution with the two stopping rules of §3.1. In §3.4.2 we calculate the fixed stopping-threshold,  $c_k$ ,

for a range of return periods,  $k$ , using (3.1.2) and our knowledge of the true parameters  $\mu, \sigma$  and  $\xi$ . In §3.4.3 we consider the variable-threshold stopping rule over a range of  $k$ . Similar simulation results are given in the supplementary material to Barlow et al. (2020) for the GPD.

### 3.4.1 Simulation design

We investigate true return-periods,  $k$ , between 20 and 2000. When generating the data, for each  $k$ , for the fixed-threshold stopping rule, we set  $c_k$  to be the true  $(1 - 1/k)$ th quantile of the data-generating distribution (*i.e.*, the  $k$ -yr return level) whereas for the variable-threshold rule the threshold is the estimated  $(1 - 1/k)$ th quantile; with both rules we stop at the first exceedance. In the simulation study, for each combination of  $\theta$ , stopping rule and  $k$ , a large number of data sets were simulated to evaluate the RMSE, bias and variance of the estimators. Given the likelihood  $\ell_M$  for  $M \in \{\text{std}, \text{ex}, \text{fc}, \text{pc}\}$ , detailed in equations (3.2.1), (3.2.2), (3.3.1) and (3.3.2), profile-likelihood confidence intervals for a return level are studied in terms of their coverage and width.

One major issue with simulating data sets with stopping rules is parameter identifiability. For observation  $i$ , the stopping decision of the variable-threshold rule is based on the parameter MLEs using observations  $1, \dots, i-1$ . However, with  $N \leq 2$  observations contributing to a likelihood the GEV parameters are strictly not identifiable, and for larger but low values of  $N$  the parameters are still not practically estimable. As discussed in §3.2.1, in practice there is typically additional information which is incorporated into decisions, and our analysis should allow for this also. Such historical information is treated as *fixed* and introduces a fixed extra penalty term,  $P_{hist}(\theta)$ , into the log-likelihood; in a Bayesian analysis it would constitute prior information about the parameter vector. As our simulation studies are conducted without such evidence we treat the first  $n_0$  simulated values as providing historical information on  $\theta$ ; we call  $\tilde{x} := (x_1, \dots, x_{n_0})$  the *historical data*. Thus each simulated data set has the penalty contribution to the likelihood:

$$P_{hist}(\theta) = \sum_{j=1}^{n_0} \ell(\theta; x_j), \quad (3.4.1)$$

a contribution that does not depend on the stopping rule since we imagine that these data were available *before* decisions to stop and analyse the data were being made.

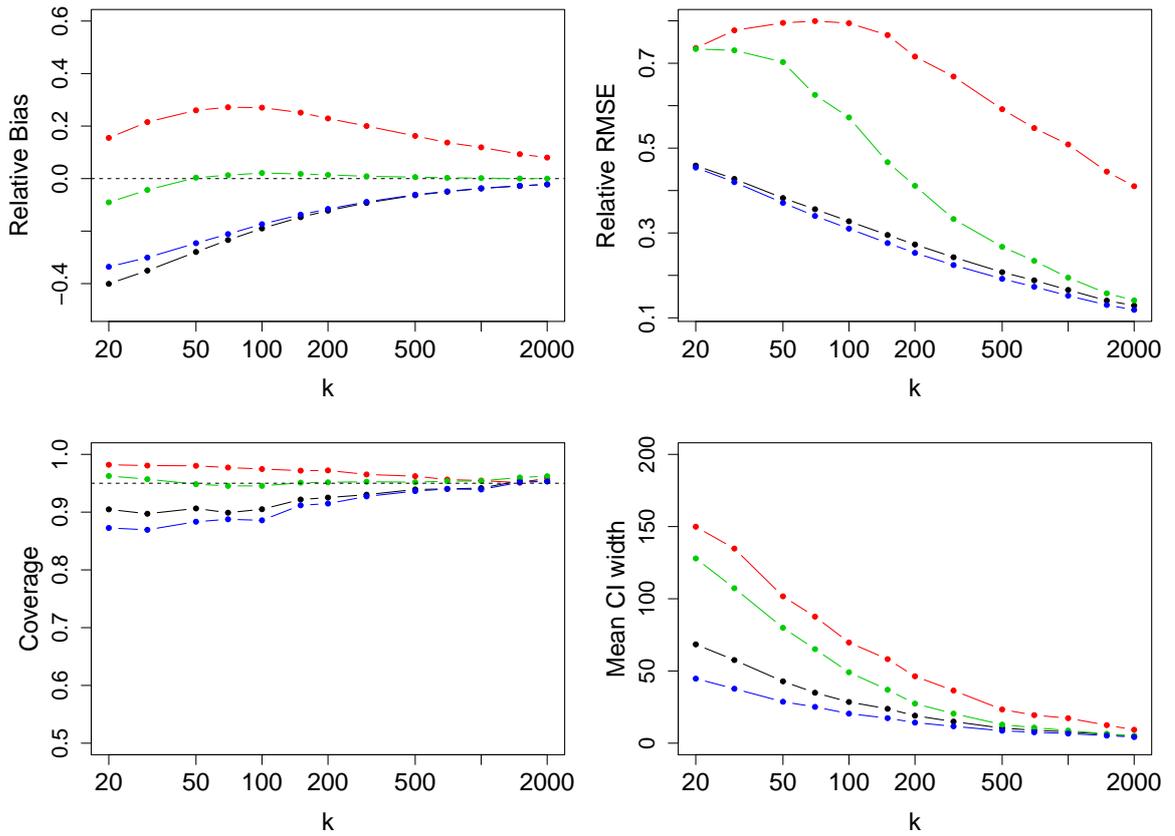


Figure 3.4.1: *200 year return-level estimates when sampling from the GEV distribution with  $(\mu, \sigma, \xi) = (0, 1, 0.2)$  using the fixed-threshold stopping rule over a range of thresholds. From left to right. Top: relative bias and relative RMSE. Bottom: coverage and average CI width. Colour scheme is the same as in Figure 3.2.1. Based on  $10^5$  replicated samples with the historical data created using approach (3.4.2). Coverage is based on 5000 replicated samples.*

We fix the historical data,  $\tilde{x}$ , using an even spread of values:

$$x_j = G^{-1}(j/(n_0 + 1); \theta) \quad \text{for } j = 1, \dots, n_0 \quad (3.4.2)$$

where  $G$  is the distribution function of the data-generating GEV distribution. In addition to providing a natural spread of values and stabilising the likelihood, for the fixed-threshold rule, provided  $c_k$  is greater than the  $1/(n_0 + 1)$  return level, no historical value exceeds the stopping threshold. The stopping threshold,  $\hat{x}_k(\mathbf{x}_{1:i-1})$  is now, implicitly, also a function of  $\tilde{x}$ . We take  $n_0 = 10$ , the smallest value that gave reliable numerical estimates for  $\hat{x}_k(\mathbf{x}_{1:i})$  with  $i \geq n_0$ , across the set of different true values for  $\theta$  that were used in the simulation study.

### 3.4.2 Fixed-threshold stopping rule

In the appendix §A.2 we describe in detail the behaviour of the shape parameter estimates in our simulations. In particular we found the shape parameter estimator using  $\ell_{std}$  has both large positive bias and large variance. The formulae in §3.3.1 show that return-level estimates are exponential in the shape parameter. For high return periods moderately large  $\xi$  estimates can lead to unrealistically high return-level estimates which exert unwarranted influence on statistics based on empirical averages, such as estimated bias. Hence, we use trimmed averages here.

Figure 3.4.1 shows the relative bias and RMSE of the 200 year return-level estimators when sampling using the fixed-threshold stopping rule from the GEV distribution with  $\xi = 0.2$ . Similar sets of plots for the 50 and 1000 year return-level estimator and  $\xi = -0.2$  can be found in the supplementary material to Barlow et al. (2020). The main driver of RRMSE in all cases is found to be the variance of the estimators, so changes in bias are not too important in this regard. Overall, the return-level estimator which results in the lowest RRMSE most consistently is  $\hat{x}_y^{pc}$ , mostly due to the low variance of these estimates whereas  $\hat{x}_y^{fc}$  has the lowest bias. Both conditioning estimators,  $\hat{x}_y^{pc}$  and  $\hat{x}_y^{fc}$ , improve upon the  $\hat{x}_y^{std}$  especially when we are estimating very high return levels (*i.e.*, for larger  $y$ ) and/or the underlying distribution is heavy tailed. Although  $\hat{x}_{200}^{ex}$  has somewhat similar properties to  $\hat{x}_{200}^{pc}$  for  $\xi = 0.2$  it has larger RRMSE for  $\xi = -0.2$ . The fitted distribution using  $\ell_{ex}$  typically has a lighter tail and can even have an upper end point which is less than the excluded observation.

The coverage for all likelihoods gets closer to the correct value (here 95%) as  $k$  increases for any return period,  $y$ . For  $\ell_{std}$  we have overcoverage and the widest confidence intervals on average and using  $\ell_{fc}$  we have good coverage, particularly when the distribution is heavy tailed. For the other likelihoods there is mostly undercoverage (coverage ranging from 80-95%) due to upper bounds being too low. The exclusion of upper tail information results in relatively narrow confidence intervals from  $\ell_{ex}$ . In contrast,  $\ell_{fc}$  produces a higher upper confidence limit and hence a wider confidence interval than  $\ell_{pc}$  and  $\ell_{ex}$  because the likelihood essentially neglects the distribution of  $N$ , *i.e.*, the threshold exceedance counts, which contain some information about the upper tail of the distribution. The confidence intervals produced using  $\ell_{fc}$  vary greatly in width across our simulations with a larger median width than those using  $\ell_{std}$ .

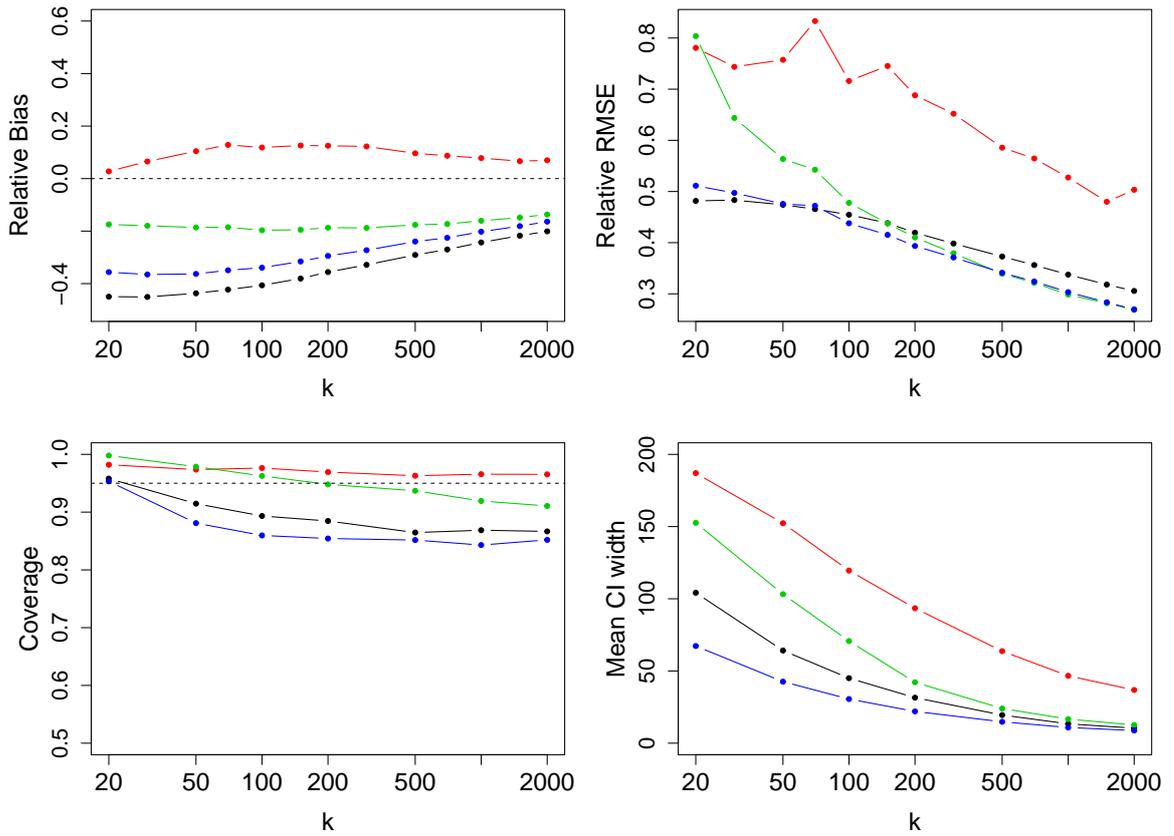


Figure 3.4.2: 200 year return-level estimates when sampling from the GEV distribution with  $(\mu, \sigma, \xi) = (0, 1, 0.2)$  using the variable-threshold stopping rule over a range of  $k$ . See Figure 3.4.1 for associated detail. Based on 10000 replicated samples with the historical data created using approach (3.4.2). Coverage is based on 3000 replicated samples.

### 3.4.3 Variable-threshold stopping rule

Within the samples simulated we find the stopping thresholds,  $\hat{x}_k(x_{1:m})$ , over  $m < N$  are generally less than the true  $k$ -year return level,  $x_k$ . As a result the samples are both smaller in size and consist of smaller values than when using the fixed-threshold stopping rule. So return-level estimates calculated using  $\ell_{std}$  have a small positive bias and those calculated using the other three likelihoods have a larger negative bias than observed for the fixed-threshold stopping rule.

The properties of the 200-year return-level estimators for  $\xi = 0.2$  are shown in Figure 3.4.2, the 50 and 1000-year return levels and  $\xi = -0.2$  are considered in the supplementary material to Barlow et al. (2020). For  $\xi = 0.2$  the conditioning methods provide the best return-level estimators in terms of RMSE despite the estimators having a larger squared bias than  $\hat{x}_y^{std}$ . The reason for this is that for heavy tailed distributions the variances of return-level estimators are generally larger than the bias. However, for lighter tailed distributions the

bias plays a larger role as the relative variances of the different estimators are much closer together. As a result  $\hat{x}_{200}^{pc}$  and  $\hat{x}_{200}^{fc}$  can perform marginally worse than  $\hat{x}_{200}^{std}$  in terms of RRMSE when the distribution has a light tail.

The coverage for  $\ell_{std}$  is high (96-98%), decreasing only slightly as  $k$  increases but it has the widest confidence intervals generally. Using either  $\ell_{pc}$  or  $\ell_{ex}$  leads to undercoverage, as  $k$  increases ranging from approximately 95% to 83-87% for  $\ell_{ex}$  and from 93% to 78-85% for  $\ell_{pc}$  with coverage higher when the distribution is heavy tailed. The coverage for  $\ell_{fc}$  also reduces with increasing  $k$  from 99 – 100% for  $k = 20$  to approximately 90% for larger  $k$ . On average the confidence intervals using  $\ell_{fc}$  are narrower than using  $\ell_{std}$  but generally wider than those using  $\ell_{ex}$  or  $\ell_{pc}$ .

Overall,  $\ell_{fc}$  provides the ‘best’ results when using the variable-threshold stopping rule. The RRMSE of  $\hat{x}_{200}^{fc}$  is generally lower than that of  $\hat{x}_{200}^{std}$ , coverage is above 90% and the confidence intervals are narrower on average than those using the  $\ell_{std}$ . Although  $\ell_{pc}$  provides estimators with a lower RRMSE than  $\ell_{fc}$ , particularly when the distribution is heavy tailed, it has more severe undercoverage.

### 3.4.4 Use in Practice

In practice, for the analysis of data that we believe has been obtained by the flood management agencies using the fixed-threshold rule we must set a threshold,  $c$ , and if they use a variable-threshold rule we must set a return period,  $k$ , neither of which may be known. This is important since the behaviour of the estimators can vary depending on the return period,  $k$ , associated with the stopping threshold (as we have seen in §3.4.2 and 3.4.3). For the fixed-threshold rule,  $c$  should lie between  $\max_{i < n} x_i$  and  $x_n$ . For the variable-threshold rule  $k$  should be such that  $x_i \leq \hat{x}_k(x_{1:i-1})$  for all  $i < n$ , but  $x_n > \hat{x}_k(x_{1:n-1})$ . To use the simulation study results to understand the properties of the estimators it is useful to narrow down a range of feasible  $k$ . For the variable rule, a range of possible  $k$  can be determined from the data. However, for the fixed-threshold stopping rule  $k$  is unknown. Nevertheless, we are likely to have some idea of the range of  $k$  which corresponds to  $c$ , *i.e.*, we have a prior belief for  $k$ .

The ‘historical data’ also needs to be determined, maybe incorporating prior knowledge in some way. The simplest approach is to start using the stopping rule after the first  $n_0$  observations of the data set and use these  $n_0$  values as the historical data. The choice of

$n_0$  only affects the point estimates and confidence intervals using  $\ell_{fc}$ . However  $n_0$  and the historical data itself can have a large impact on the properties of the estimators, particularly when the sample size is small. In the simulation study, out of necessity, we have restricted ourselves to a particular fixed historical sample, so for low  $k$  the properties of the estimators will differ slightly in practice.

### 3.5 Case Study - Lune at Caton

We now consider the analysis of the 48 annual maximum river flow observations from the Lune at Caton introduced in §3.1. Figure 3.1.1, right panel, shows the inference for the 200-year return level of the data, at yearly intervals as new data are observed, with the analysis not accounting for any stopping rule. We now estimate this return level using the four inference methods (standard, exclude, and our full- and partial-conditional) for both fixed- and variable-threshold stopping rules for a range of levels ( $c$  and  $k$  respectively), where we drop the subscript of  $c$  as the return period of the stopping threshold is unknown. The following discussion assumes that the sampling procedure is well approximated by these respective stopping rules for the selected  $c$  and  $k$ . In all cases we take the historical data to be the first  $n_0 = 10$  observations as in practice no estimates of long period return levels would be attempted from smaller samples. We also consider the implications if a trend in the annual maxima is also simultaneously estimated.

#### 3.5.1 Fixed-threshold stopping rule

First we discuss the inference using the fixed-threshold stopping rule with  $c = 1568m^3/s$ , where, for illustration purposes,  $c$  is taken to be the mid-point between the 1995 and 2015 levels and the realised value of  $N$  is 38, *i.e.*, we stop after 2015. Figure 3.1.1, right panel, to the right of the vertical dotted line, shows the estimates and the associated 95% confidence intervals for the four inference methods. The estimates  $\hat{x}_{200}^{std}$  and  $\hat{x}_{200}^{ex}$  are identical to the estimates in the right panel of the figure for years 2015 and 2014 respectively. Both  $\hat{x}_{200}^{fc}$  and  $\hat{x}_{200}^{pc}$  (evaluated at 2015) are only slightly larger than the  $\hat{x}_{200}^{std}$  estimates for the years before 2015 and  $\hat{x}_{200}^{ex}$ , despite the inclusion of the 2015 value. From §3.4.2 we know that, when employing the fixed-threshold stopping rule,  $\hat{x}_{200}^{std}$  is positively biased,  $\hat{x}_{200}^{fc}$  is close to being unbiased and both  $\hat{x}_{200}^{ex}$  and  $\hat{x}_{200}^{pc}$  have some negative bias, therefore it is reassuring to

see that  $\hat{x}_{200}^{std} \gg \hat{x}_{200}^{fc} > \hat{x}_{200}^{pc} > \hat{x}_{200}^{ex}$ .

The confidence interval for 2015 using  $\ell_{std}$  is wider than the intervals of the previous 15 years, especially the 2014 interval (*i.e.*, using  $\ell_{ex}$ ), and both the lower and upper bounds are much larger. In this case study, the confidence interval of  $x_{200}$  using  $\ell_{pc}$  is similar but slightly narrower than when using  $\ell_{ex}$ . However using  $\ell_{fc}$  the interval is wider (since the upper bound increases) than if we just ignored the 2015 event (using  $\ell_{ex}$ ) so we are capturing some of the increased uncertainty in the heaviness of the tail that this event has caused. Nevertheless, the upper confidence bound of  $x_{200}$  is lower than that using  $\ell_{std}$ .

The behaviour of the confidence intervals of these methods appears to be in line with our coverage and width results in §3.4.2. Indeed, here the shape parameter estimates,  $(\hat{\xi}_{std}, \hat{\xi}_{ex}, \hat{\xi}_{fc}, \hat{\xi}_{pc})$ , are  $(0.04, -0.07, -0.04, -0.05)$  so we expect coverage to be between the coverage values found in the simulation study for  $\xi = 0.2$  and  $\xi = -0.2$ . In the study we found that using  $\ell_{std}$  with the fixed-threshold stopping rule leads to overcoverage (95-98% for  $\xi = 0.2$ , 97.5-99.5% for  $\xi = -0.2$ ) and the upper bound of the confidence interval found using  $\ell_{std}$  is lower than  $x_{200}$  only 1-2% of the time, so it is likely that for the Lune data the upper bound of the confidence interval using  $\ell_{std}$  is too high. This is further emphasised for the Lune estimates by the upper bound for 2015 exceeding the associated values for the previous 30 years (Figure 3.1.1). In §3.4.2 we found that  $\hat{x}_{200}^{ex}$  and  $\hat{x}_{200}^{pc}$  exhibited narrow confidence intervals which together with their negative bias led to undercoverage, with the upper bounds being too low, especially when  $\xi = -0.2$  and  $k$  is low. For our chosen  $c = 1568$  we can obtain estimates of the corresponding return period,  $k$ , of  $c$ ; in particular  $\hat{k}_{std} = 90$  and  $\hat{k}_{ex} = 550$  and we expect  $k$  to lie between these two values. Thus, using the simulation study results, we expect that the coverage of the  $\ell_{pc}$  and  $\ell_{ex}$  confidence intervals to lie between 85 and 95%. However the lower bounds of these confidence intervals were found to be less than  $x_{200}$  for almost 100% of simulated samples so it is highly likely that the true 200-year return level for the Lune data is above the lower bounds given by the  $\ell_{pc}$  and  $\ell_{ex}$  confidence intervals. For  $\ell_{fc}$  and  $90 < k < 550$ , the coverage is 94-95% with the percentage of upper bounds too low being 3-6% suggesting that with the Lune data the upper bound of the  $\ell_{fc}$  confidence interval is likely to be higher than the true 200-year return level,  $x_{200}$ .

The above discussion assumed that  $c$  was known. In some cases this may be true as  $c$  could represent a known physical limit linked to flooding. This is not the case for the Lune at Caton, with our value chosen subjectively for illustrative purposes although it could be

argued that lower  $c$  values in this range would be more reasonable since the 1995 river flow observation was considered high as it led to flooding. To assess the impact of  $c$  we consider a range of values for  $c$  between the 1995 and 2015 observations, with the inference for the four methods presented in Figure 3.5.1, left panel (when  $c = 1568$  the estimates are those shown in Figure 3.1.1 right panel). Now,  $\hat{x}_{200}^{std}$  and  $\hat{x}_{200}^{ex}$  and the corresponding confidence intervals are invariant to  $c$  but as  $c$  increases  $\hat{x}_{200}^{fc}$  and  $\hat{x}_{200}^{pc}$  both decrease. As noted earlier,  $\hat{x}_{200}^{fc} > \hat{x}_{200}^{pc}$  but they become closer as  $c$  tends to the 2015 event level because the information that  $\ell_{fc}$  discards, *i.e.*, the probability of the event that  $c$  was not exceeded on the first  $n - 1$  observations, becomes less informative. The confidence intervals using the conditioning likelihoods notably narrow with increasing  $c$ ; the lower bounds slightly decrease but the largest reduction is in the upper bounds. For lower  $c$  values the  $\ell_{fc}$  intervals are wider than for  $\ell_{std}$ , in contrast for the largest possible  $c$  values the interval is very narrow (a reduction in size of factor 14 over the range of  $c$  possible). For  $\ell_{pc}$  the upper bounds are smaller than those using the  $\ell_{std}$  for all values of  $c$  and are slightly larger than those for  $\ell_{ex}$  for low  $c$ . However, for large  $c$  the upper bounds of both conditioning confidence intervals are much lower than that using  $\ell_{ex}$  since the information that  $c$  was exceeded on this observation becomes more informative about the tail of the distribution as  $c$  approaches the 2015 observation. Thus if we stop after the first minor exceedance of  $c$  we can be reasonably sure the tail is short. This is an unexpected but helpful finding. Further investigation into the confidence intervals can be found in Chapter 4.

### 3.5.2 Variable-threshold stopping rule

Now we consider the variable-threshold stopping rule and first determine a range of  $k$  from the data. In the Lune data the maximum river level in 2015 corresponds to  $\hat{k} = 2561$  given the data up to 2015 and to  $\hat{k} = 188$  using all the data. However, the river level in 1995 corresponds to  $\hat{k} = \infty$  (*i.e.*, it is larger than the point estimate of the upper end point of the GEV fitted to the data up to 1995) so the variable-threshold rule as given in (3.1.4) cannot have been applied for any  $k < \infty$ . Furthermore, the river level in 1980 corresponds to  $\hat{k} = 111$ . If the variable-threshold stopping rule had motivated a request for an analysis of the data up to and including 2015, the request must have been triggered by the second such exceedance. In our analysis we explore values of  $k$  between 200 and 2500 and simply amend  $\ell_{fc}$  slightly by replacing the  $\ell_{fc}$  contribution of the 1995 observation ( $i = 28$ ),  $g(x_{28}; \theta)/G(\hat{x}_k^{std}(\mathbf{x}_{1:27}; \theta))$ ,

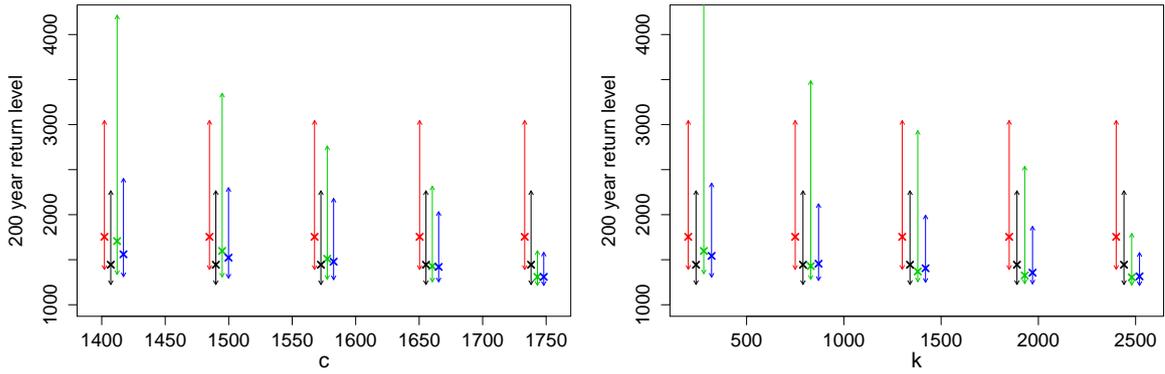


Figure 3.5.1: *200-year return-level estimates based on all the data up to and including 2015 for the Lune at Caton with 95% profile likelihood-based confidence intervals: left with the fixed-threshold stopping rule over a range of  $c$  and right with the variable-threshold stopping rule over a range of  $k$ : standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue). Each group of 4 estimates applies for the same  $c/k$  as for the standard estimate in each group and have been horizontally shifted for clarity.*

by  $g(x_{28}; \theta)$ .

Figure 3.5.1, right panel, shows the same inferences as the left panel, but for the variable-threshold over a range of return periods  $k \in [200, 2500]$ . Given the rarity of all events in this range we would expect a ‘true’  $k$  to be towards the lower end of this range. The estimates  $\hat{x}_{200}^{std}$  and  $\hat{x}_{200}^{ex}$  and the corresponding confidence intervals are invariant to  $k$  (and independent of the stopping rule used) but as  $k$  increases  $\hat{x}_{200}^{fc}$  and  $\hat{x}_{200}^{pc}$  both decrease. For small  $k$ ,  $\hat{x}_{200}^{fc} > \hat{x}_{200}^{pc}$ , as we would expect from our bias results in the simulation study. However, the inequality reverses for large  $k$  perhaps as a result of there being more than one exceedance of the threshold. This is hinted at by the bias results and also since if one omits the 1995 observation from the data set then  $\hat{x}_{200}^{fc} > \hat{x}_{200}^{pc}$  for all  $k$ . More investigation into the estimators when there are multiple exceedances would be useful.

The intervals using the conditioning likelihoods and variable-threshold stopping rule behave similarly to those using the fixed-threshold stopping rule. Again the  $\ell_{fc}$  intervals are highly influenced by the ‘extremeness’ of the stopping threshold. With the lowest possible  $k$  for this data set (ignoring the 1995 exceedance) the  $\ell_{fc}$  interval is more than double the width of the confidence interval using  $\ell_{std}$  whereas for a large  $k$  value it is less than half the width. The  $\ell_{pc}$  confidence intervals also reduce in width with increasing  $k$  but not as dramatically.

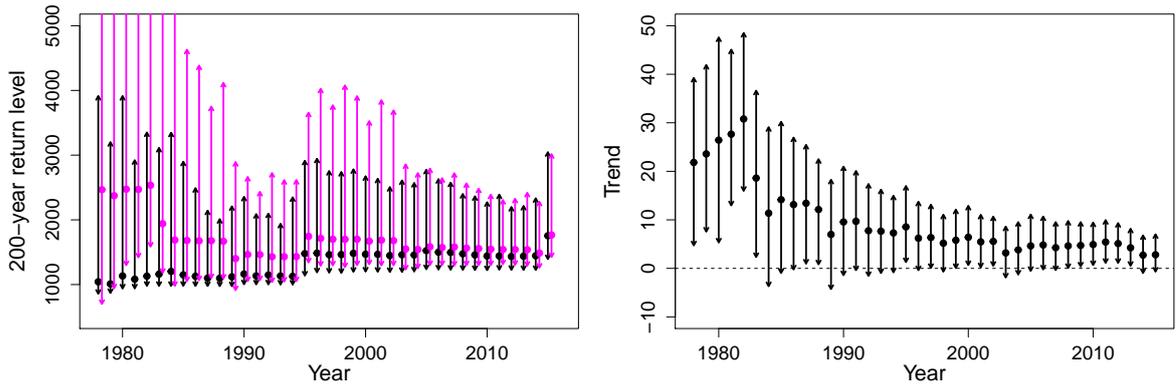


Figure 3.5.2: *Fitting a GEV to all the data up to and including the current year for the Lune at Caton. Left: 200yr return-level estimates for 2015 using progressively more data over the years with and without a trend in the location parameter (pink and black respectively). Each group of 2 estimates applies for the same year and have been horizontally shifted for clarity. Right: Slope parameter,  $\hat{\beta}$ , and its 95% confidence interval.*

### 3.5.3 Non-stationarity

The implications of using stopping rules on the estimation of trends in extreme levels is also a concern, as stopping with the final observation being large is likely to have a similar biasing effect as found in §3.2 and §3.4 for return levels. This is particularly important given the interest in whether trends in extreme values differ from trends in mean levels (Eastoe and Tawn, 2009; Hannaford and Marsh, 2008). In Figure 3.5.2 we illustrate the analysis of the Lune data with a GEV distribution including a linear trend  $\mu_t = \alpha_0 + \beta t$ , showing both the resulting estimates of the 200-year return level for 2015, *i.e.*, the estimates of the 0.995 quantile of the annual maximum in 2015, and the associated trend estimate  $\hat{\beta}$  using progressively more data over time. With few data used the trend is estimated to be unrealistically large, with huge uncertainty, and this results in very different point estimates of return levels relative to the analysis with no trend. As more data are observed we can see that the trend estimates generally decrease, with reduced uncertainty, with positive jumps in  $\hat{\beta}$  estimates after the large 1995 and 2015 events. Although the 2015 river flow is more extreme than that of 1995 its impact on  $\hat{\beta}$  is much less. Furthermore, we see that  $\hat{\beta}$  is not larger than  $\beta = 0$  at the 2.5% significance level. Thus here the effect of including the estimated trend is small on the 200-year return-level estimate and the stopping rule seems to have almost no effect on the trend estimate.

### 3.6 Discussion

In this chapter and the supplementary material associated with Barlow et al. (2020) we have shown that return-level estimators based on the standard likelihood are positively biased when sampling from the GEV or GP distributions using certain stopping rules. The extent of the stopping bias is lower for lighter tailed distributions and when estimating low return levels. We have proposed conditioning upon the stopping threshold in the likelihood. In most cases we have found that conditioning on the final observation exceeding the stopping threshold (partial conditioning) results in return-level estimates with the lowest RMSE despite the estimator being negatively biased.

A balance must be struck between low RMSE and good coverage, however. Partial conditioning results in undercoverage despite the low RMSE of  $\hat{x}_y^{pc}$ . The full-conditional likelihood, which also conditions on the non-exceedance of all previous observations, gives the closest to 95% coverage and though the intervals are wide, they are typically narrower than the confidence intervals obtained from the standard likelihood. The interval widths using the full and partial conditional likelihoods are smaller the closer the stopping threshold is to the final observation as the occurrence of the final exceedance becomes more informative on the tail of the distribution (see §3.5.1).

Overall, the conditioning estimators presented here outperform the standard estimator when the decision to analyse data at a particular time was triggered by what was perceived to be a large observation. For the fixed-threshold stopping rule, partial conditioning has the best combination of RMSE and coverage for a range of  $\xi$  with moderate  $k$  and particularly when the distribution is heavy tailed, as is the case for most UK rivers (CEH, 1999). For the variable-threshold stopping rule, full conditioning provides the best balance of coverage and low RMSE. To apply the conditioning estimators in practice if the rule of the flood management agency is unknown the statistician needs to choose a suitable stopping threshold,  $c$ , for the fixed-threshold stopping rule and a suitable stopping ‘period’,  $k$ , for the variable-threshold stopping rule if the values are unknown. A range of  $c$  and  $k$  can be considered provided that the observed data are below the resulting stopping threshold(s) up to the final observation.

The decision to analyse data will likely be based on a confluence of many factors. Our work attempts to simplify the true decision making procedure by using stopping rules based on the

occurrence of a *single* large observation exceeding some threshold. An analysis may instead be prompted by a prolonged period of quite large (but not necessarily ‘extreme’) observations or the observation of large values at many locations simultaneously, requiring more complex multivariate analysis since the observations at nearby locations will be dependent in some way (Keef et al., 2009; Asadi et al., 2015).

In practice if the stopping rule is unknown and the analysis is triggered by a large event, we suggest using the full conditional return-level estimator. However if  $k$  is thought to be less than 50, or the full-conditional estimate and/or confidence interval are clearly too large then partial conditioning should be used instead. We argue that the decision to ‘stop’ and analyse data would in part be based on both past return-level estimates and thresholds set due to current infrastructure and so the ‘true’ stopping rule is a mixture of the two rules considered here. Hence the ‘true’ bias, RMSE and coverage of the estimators can be expected to lie between those which we found under the two stopping rules. It should be noted that this work does not address the question of *when* the data should be analysed, but rather how we can reduce the bias given the use of a particular stopping rule. Nevertheless, if we are at a point in time where a stopping criterion has been met and triggered an analysis, this study can give guidance on the behaviour of return-level estimators calculated at the current time whether based on the full likelihood, partial or full conditioning, or even excluding the most recent, ‘triggering’ event.

In our theoretical and simulation studies we have not accounted for the possibility of a trend in the data, such as river flows gradually increasing over the years. We saw in §3.5.3 that the Lune data has a slight positive trend in the location parameter and fitting such a model at an earlier point in time resulted in a very large positive trend. This could cause problems for the fixed-threshold stopping rule, in particular it might become necessary to change  $c$  after a certain number of years. Nonetheless, doing this is probably not too unrealistic since, for example, the height of a flood defence might be increased if there has been evidence of higher flow in recent years. On the other hand the variable-threshold stopping rule is more robust to data with an underlying trend as it is directly a function of the observed data.

## Chapter 4

# Investigating confidence intervals for return-level estimators on data generated by threshold-based stopping rules

### 4.1 Introduction

In Chapter 3 we discussed two stopping rules based on the exceedance of some large value and proposed two estimators based on conditioning on the occurrence of such an exceedance. We compared the coverage of the different estimators using profile-likelihood based confidence intervals. However there was some concern that in many cases the profile-likelihood based confidence intervals appeared to be very wide. Here we investigate these intervals in more depth and compare to a variety of bootstrap-based confidence intervals, concentrating on data sets similar to the Lune data set.

First, in §4.2, we consider the profile-likelihood based confidence intervals when sampling using the fixed-threshold stopping rule from the GEV with the parameters set to the standard MLEs for the Lune data,  $\hat{\theta}_{std}^{Lune}$ , and explore the relationship between sample size and confidence interval width for the different estimators. Then, we create samples with the same procedure but such that all samples are of the same size as the Lune data (48 observations) in order to investigate the properties of the confidence intervals we would expect in the Lune setting and with similar data (in general, peak river flow data are unlikely to consist of more than 50 years of observations).

In §4.3 we consider the standard bootstrap confidence intervals and bias reducing vari-

ations thereof. We compare the coverage and confidence intervals of profile-likelihood and bootstrap-based methods when the primary data sets are of size 48 and are sampled from the GEV with parameters set to  $\hat{\theta}_{std}^{Lune}$  using the fixed-threshold stopping rule. We find that there is a trade-off between reducing the confidence interval widths and increasing coverage with the standard bootstrap method resulting in narrow intervals and poor coverage. The bootstrap variations generally increase coverage but the confidence interval widths are comparable to the profile-likelihood based interval widths; the latter is much faster computationally so remains the preferred confidence interval method. Nevertheless it appears reasonable that some reduction in width should be possible without drastically reducing the coverage; more investigation into confidence intervals in the stopping rule setting could be useful.

When using the variable-threshold stopping rule the creation of bootstrap confidence intervals requires more thought. In §4.4 we describe an importance-weighted bootstrap to create confidence intervals when using the variable-threshold stopping rule and also when there are multiple exceedances. We find these importance sampling confidence intervals are narrower than the profile-likelihood based intervals (like the standard bootstrap method for the fixed-threshold stopping rule) and are highly negatively biased due to the negative bias in the return-level estimators and the resulting bootstrap samples.

Finally, in §4.5 we step back to the profile-likelihood based confidence intervals for a more detailed investigation into the effect of the choice of the fixed stopping threshold,  $c_k$ , on the confidence intervals of the conditioning estimators (this was briefly discussed in §3.5 for the Lune case study) and discuss the issues that arise for very low and high  $c_k$ .

## 4.2 Profile likelihood in Lune setting

Consider using the fixed-threshold based stopping rule to sample from the GEV with the parameters set to  $\hat{\theta}_{std}^{Lune}$  (the MLE for the Lune data when using the standard likelihood,  $\ell_{std}$ ). In §4.2.1, we explore the relationship between sample size and confidence interval width in such a setting for the four return level estimators (based on the standard, exclude, full-conditioning and partial-conditioning likelihood as defined in §3.2 and §3.3 of Chapter 3). To further explore the behaviour of the estimators for data sets similar to the Lune data set, in §4.2.2 we consider the properties of the confidence intervals when the sample created using the fixed-threshold stopping rule is of size 48, the same size as the Lune data. Indeed,

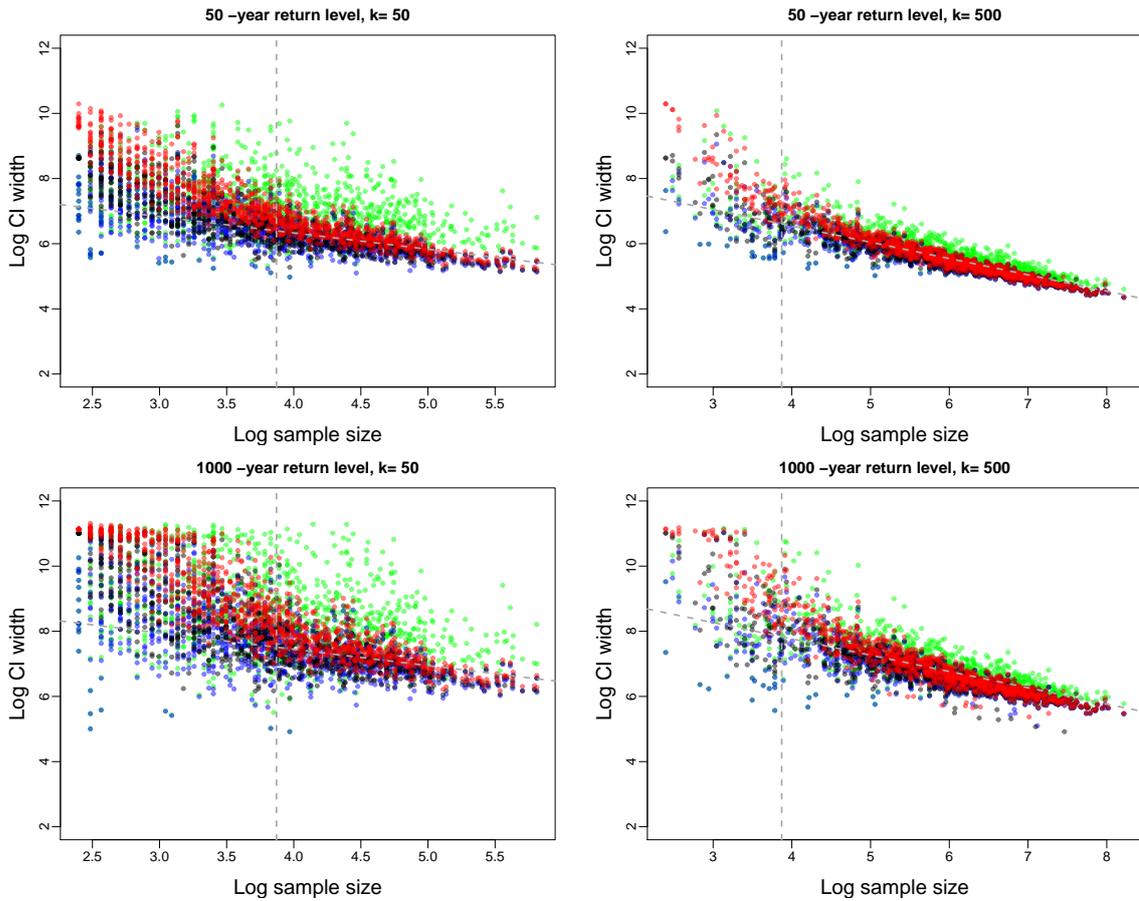


Figure 4.2.1: *Log CI width vs log sample size for the 50-year (top) and 1000-year (bottom) return level estimates. Profile-likelihood confidence intervals found using the standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue) based on 5000 samples from the GEV distribution with parameters equal to the standard MLEs for the Lune data and sample size determined by the fixed-threshold stopping rule with stopping threshold  $c_{50}$  (left) and  $c_{500}$  (right).*

maximum river flow data are unlikely to have been collected for more than 50 years so such a sample size is of particular interest.

In the following we concentrate on the 95% confidence interval and denote the true  $y$ -year return level by  $x_y$  and estimates  $\hat{x}_y^L$  with return period  $y \in \{50, 200, 1000\}$  and  $L \in \{std, ex, fc, pc\}$  indicating the likelihood used.

#### 4.2.1 Relationship between confidence interval widths and sample size

Figures 4.2.1 and B.1.1 show the log confidence interval widths against log sample size,  $\log(n)$ , for each of the four 50-year and 1000-yr return-level estimators with the random sample size determined by the fixed-threshold stopping rule where the return period of the stopping threshold is  $k \in \{50, 100, 500\}$ . The confidence interval width is proportional to  $\frac{1}{\sqrt{n}}$  *i.e.*,

we expect the points to be concentrated around the line with gradient  $-\frac{1}{2}$  (as a guide the dashed lines in Figures 4.2.1 and B.1.1 have such a gradient with  $y$ -intercept equal to the median of the standard confidence interval widths for mid sample sizes.). For medium to large samples the points are clustered around a line with such a gradient whereas for small samples confidence intervals are larger than expected on average, particularly when using  $\ell_{std}$ . Full conditioning generally leads to wider confidence intervals than the other likelihoods considered but also has a wider spread of confidence interval widths. The confidence interval widths using the other likelihoods are more similar but, for small sample sizes in particular, the standard confidence intervals are wider than for those excluding the final observation which in turn are wider than the partial conditioning intervals.

Reassuringly, the coverage results (Figure 4.2.2 left) lie between those for the simulation study of §3.4.2 with  $\xi = -0.2$  and  $\xi = 0.2$  (recall we are sampling from the GEV shape parameter  $\hat{\xi}_{std}^{Lune} = 0.04$  as found in Chapter 3 §3.5). Using  $\ell_{std}$  leads to overcoverage whereas for all other likelihoods the coverage is less than 95% (but increases with  $k$ ), largely due to the upper bounds of the confidence interval for  $\hat{x}_y$  being lower than  $x_y$  for several simulations. The small upper bounds are a particular problem when using  $\ell_{ex}$  or  $\ell_{pc}$ , with coverage ranging between 85% and 93%, whereas  $\ell_{fc}$  fares better with coverage ( $\sim 94\%$ ) close to 95%. For all likelihoods the percentage of lower bounds below  $x_y$  was often much higher than the desired 97.5% and reduced slightly with increasing  $k$ .

Overall full conditioning appears to do quite well in terms of coverage, however, this can be at the cost of wide confidence intervals. In Figure 4.2.1 the vertical dashed line is at sample size 48, the same size as the Lune data set. For all the proposed methods there is quite some spread in confidence interval widths for samples of this size, especially when using  $\ell_{fc}$ . As we expect annual maxima data sets to be of such size in practice we now concentrate on the confidence intervals for samples size 48.

#### 4.2.2 Similar data sets to the Lune data set

We continue our analysis sampling from the GEV with  $\hat{\theta}_{std}^{Lune}$  using the fixed-threshold stopping rule but in such a way that the sample is of the same size as the Lune data, *i.e.*,  $n = 48$ ; this is achieved by sampling  $n - n_0 - 1$  times (recall, §3.4,  $n_0$  is the size of the ‘historical data’) from the GEV truncated above by  $c_k$  and once from the GEV truncated below by  $c_k$  to obtain the final observation. From our knowledge of the bias of the different likelihood estimators

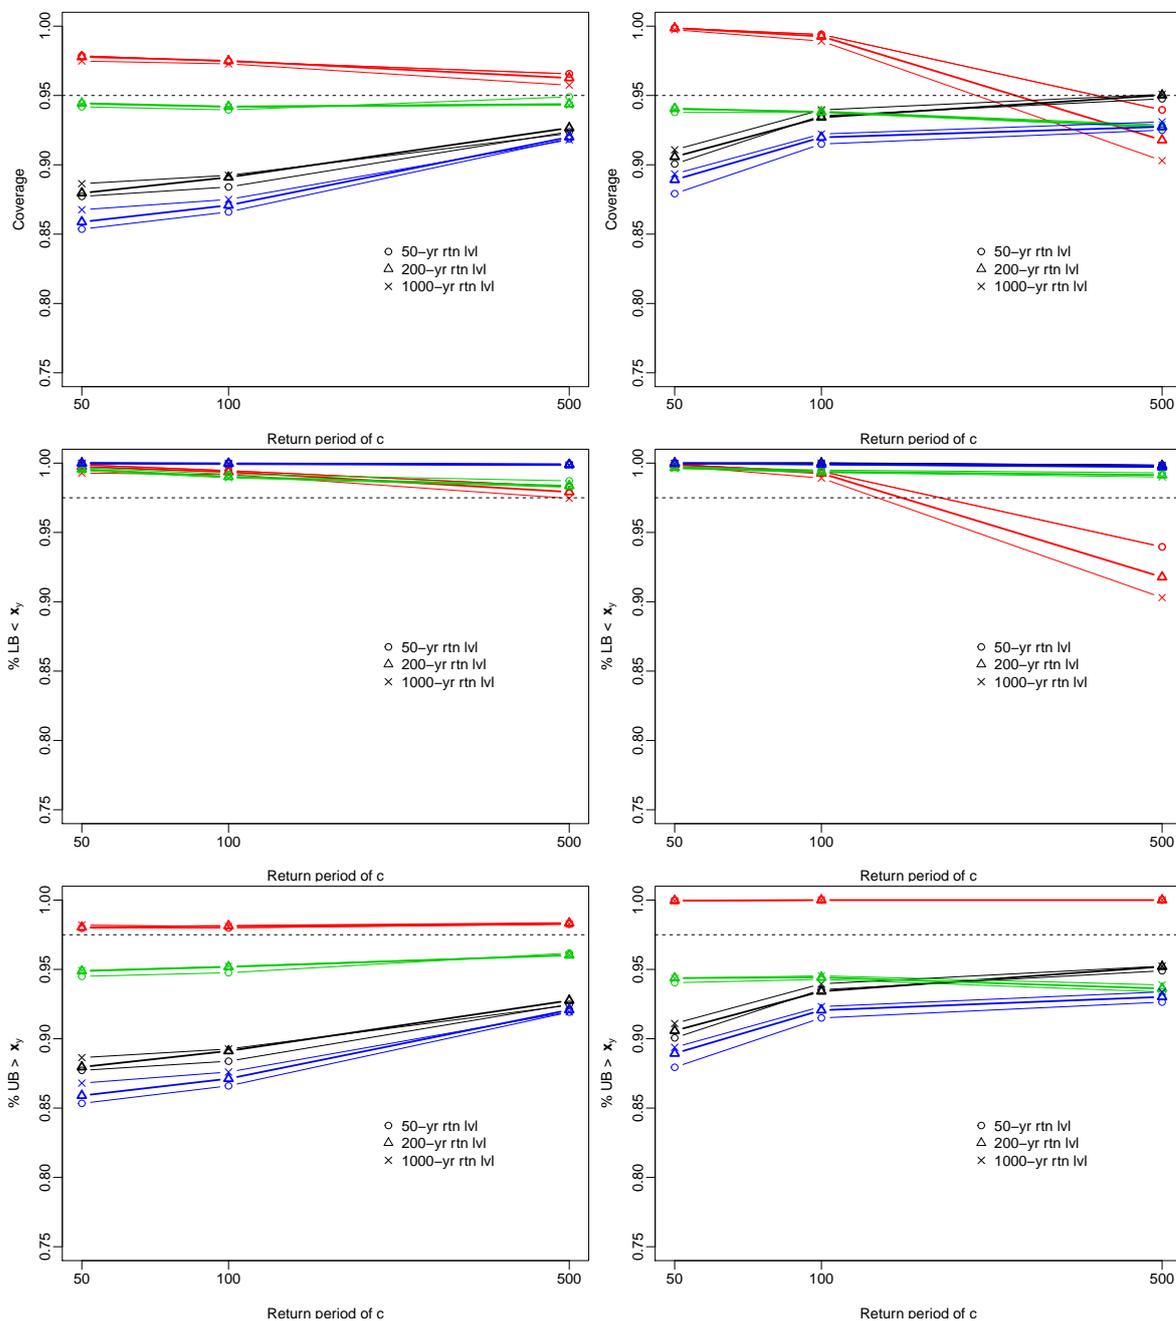


Figure 4.2.2: Coverage, % of lower bounds below  $x_y$ , % of upper bounds  $x_y$  of the profile-likelihood confidence intervals for the  $y$ -year return level,  $x_y$ , found using  $\ell_{std}$  (red),  $\ell_{ex}$  (black),  $\ell_{fc}$  (green) and  $\ell_{pc}$  (blue). There are 5000 samples from the GEV distribution with parameters equal to the standard MLEs for the Lune data and sample size determined by the fixed-threshold stopping rule (left) and equal to 48 (right) with stopping threshold  $c_k$ .

(§3.5) we expect  $\hat{k}_{std} < k = (1 - G(c_k; \hat{\theta}_{std}^{Lune}))^{-1}$ , in other words, a value of  $c_k$  or higher is estimated to be more likely using the standard estimator than it actually is according to the true distribution ( $\hat{k}_{std}^{-1} > k^{-1}$ ). Thus samples created from the GEV with  $\hat{\theta}_{std}^{Lune}$  will generally consist of larger values than the Lune data set.

We will consider the coverage, percentage of lower bounds below the truth and percentage of upper bounds greater than the truth as in Figure 4.2.2; compare the confidence interval widths using the different likelihoods (Figures B.1.2-B.1.3); and compare the confidence intervals themselves (Figures 4.2.4 and B.1.4) by plotting, for each integer in the range of the confidence intervals, the percentage of the 5000 confidence intervals that contain the integer. We also consider in Figures 4.2.3, B.1.2-B.1.3 the ‘CI width to MLE ratio’ which is the confidence interval width divided by the return-level estimate for a given sample. This ratio gives an indication of the usefulness of the confidence interval; we take a value of less than 1 to indicate reasonably sized confidence intervals.

The coverage results when the sample size is constrained to be 48 are similar to that when the sample size is determined by the stopping rule (see Figure 4.2.2 for comparison). The largest difference is in the confidence intervals using  $\ell_{std}$ ; they perform worse when the random sample size is  $n = 48$  and have undesirable properties. Firstly, the interval widths are very large (worse the larger the return period of interest is); the widths are rarely less than  $\hat{x}_y^{std}/2$  and can be more than  $6\hat{x}_y^{std}$  when the return period of interest,  $k$ , is large (e.g., Figure 4.2.3). Secondly, it can be seen in Figures 4.2.4 and B.1.4 that the intervals using  $\ell_{std}$  are skewed towards larger values than the true return level,  $x_y$ , (given by the vertical dashed line). This is particularly the case when  $k$  is large; with  $k = 500$  all 5000 data sets generated resulted in such confidence intervals covering values above  $x_y$  (the 100% peak occurs at values higher than  $x_y$ ). The upper confidence interval bounds using  $\ell_{std}$  were greater than  $x_y$  for more than 99.94% of data sets and when  $k = 500$  the lower confidence interval bounds are too high in 6-10% of simulated data sets (see Figure 4.2.2) thus leading to poorer coverage than the random sample size case.

More desirable confidence interval properties are obtained when using  $\ell_{ex}$  or  $\ell_{pc}$ . The confidence intervals have much smaller widths than those using  $\ell_{std}$  (the mean and median widths are 40-70% of those using  $\ell_{std}$ ) and are mostly centred around or just below the truth,  $x_y$ , with this value being close to the most ‘covered’ value when  $k$  is large and/or the return period of interest is high (the peaks in Figure 4.2.4 are close to  $x_y$ , the dashed line). The

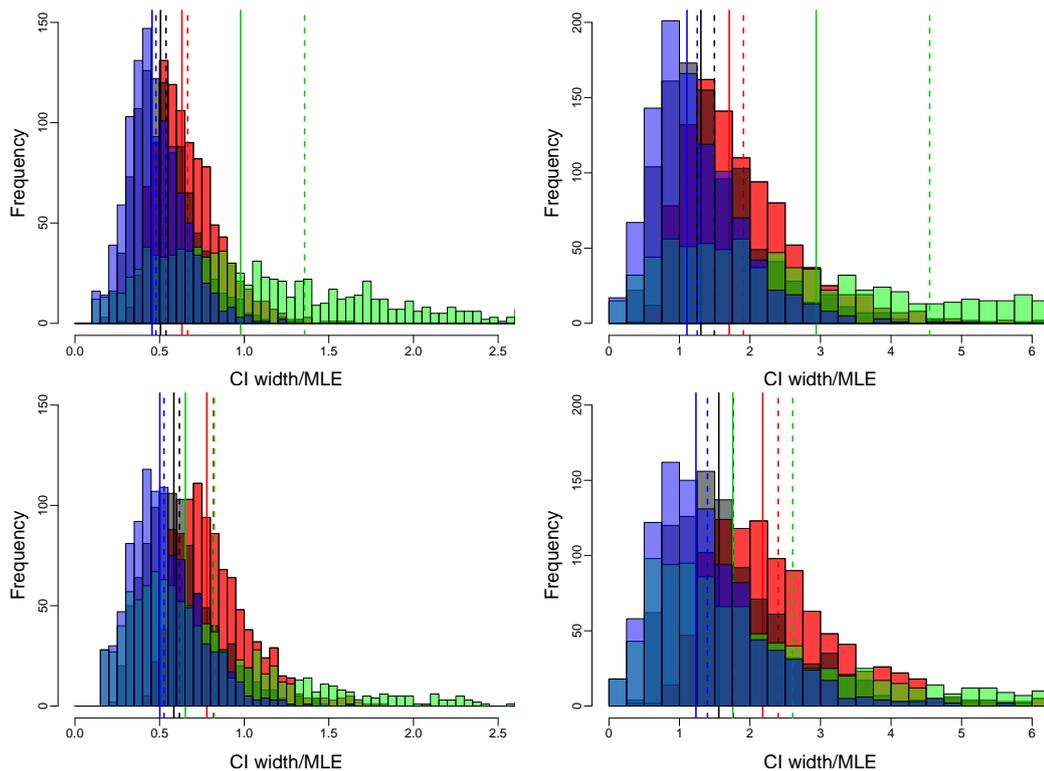


Figure 4.2.3: Profile-likelihood confidence interval widths over MLE for the 50 (top) and 1000 (bottom) -year return level found using  $\ell_{std}$  (red),  $\ell_{ex}$  (black),  $\ell_{fc}$  (green) and  $\ell_{pc}$  (blue). There are 5000 samples of size 48 from the GEV distribution with parameters equal to the standard MLEs for the Lune data and using stopping threshold  $c_k$  with  $k = 50/500$  (left/right).

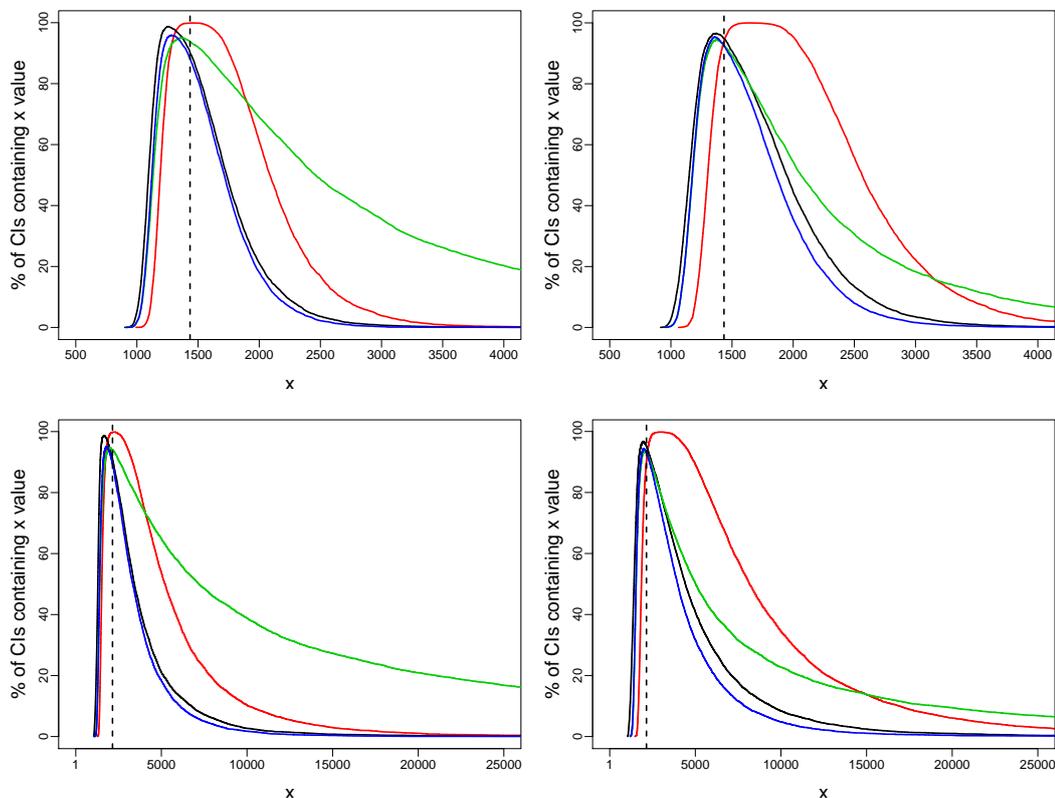


Figure 4.2.4: Percentage of confidence intervals covering each integer  $x$ . Further details in Figure 4.2.3 caption.

interval widths are reasonable in relation to the MLE, however, this comes at the price of an upper bound which is lower than  $x_y$  5-12% of the time giving coverage from 88-93% for  $\ell_{pc}$  and 90-95% for  $\ell_{ex}$ . Nevertheless, this coverage is an improvement on that found with random sample size,  $N$ .

For all three likelihoods discussed above, the confidence intervals generally become wider with increasing  $k$ . In fact, their upper bounds are larger (*e.g.*, Figure 4.2.4), since for larger  $k$  the values in the sample are larger, and so there is higher coverage with larger  $k$ . In contrast, the upper bounds of the confidence intervals using  $\ell_{fc}$  are smaller. This leads to the opposite coverage pattern in Figure 4.2.2. Most of the intervals using  $\ell_{fc}$  cover similar values to the intervals using  $\ell_{ex}$  or  $\ell_{pc}$ , however, a large percentage include extremely large values, especially when  $k$  is low. For example, with  $k = 50$  approximately 50% of the confidence intervals for  $\hat{x}_{50}^{fc}$  include the value 2500, whereas less than 20/10% of confidence intervals for  $\hat{x}_{50}^{std}/\hat{x}_{50}^{ex}$  or  $\hat{x}_{50}^{pc}$  contain 2500 (Figure 4.2.4). The confidence interval widths are much wider spread than those using the aforementioned likelihoods and they can be extremely large (much worse than using  $\ell_{std}$ ). Nevertheless, for large  $k$  the median of CI width to MLE is smaller than that using  $\ell_{std}$ .

Without knowledge of the use of a stopping rule, the small sample size (particularly compared to  $k$ ) would suggest that the observations in the sample, particularly the last observation, are more likely than they actually are. In practice there *will* be a relatively low number of annual river flow observations and so return-level estimators and, as investigated here, confidence intervals based on the standard likelihood will be positively biased if applied after a large flood event. We have found that partial conditioning provides a better profile-likelihood based confidence interval for such a small data set with only slight undercoverage. In many cases full conditioning also provides ‘good’ confidence intervals with better coverage but it is sensitive to the particular features of a sample and can result in unreasonably large intervals. As the full-conditioning return-level estimator was found to have low bias (Chapter 3) we recommend using this estimator and its corresponding confidence intervals unless those intervals are unreasonably large (for example, larger than the standard intervals).

### 4.3 Bootstrap methods with fixed-threshold stopping rule

We now consider, when using the fixed-threshold stopping rule, a variety of bootstrap-based confidence intervals and compare with intervals using the profile-likelihood method. We describe the standard parametric bootstrap in our setting, the bias-correction method of Efron (1981) and two further bias correcting amendments based on this. First we briefly discuss the profile-likelihood confidence intervals which were used in Chapter 3 and §4.2, and the Wald confidence intervals.

#### 4.3.1 Methods

##### Profile-likelihood (deviance-based) confidence intervals

The  $y$ -year return level,  $x_y$ , can be written in terms of  $\boldsymbol{\theta} = (\mu, \sigma, \xi)$  so we can instead work with the reparametrised likelihood with parameters  $\tilde{\boldsymbol{\theta}} = (x_y, \sigma, \xi)$ . The classic deviance-based method we have used is to find the likelihood-ratio test statistic and use its asymptotic properties to create a confidence interval. In our setting the likelihood-ratio test statistic (*i.e.*, deviance) is

$$D(x_y) = 2 \left\{ \ell(\hat{x}_y, \hat{\sigma}, \hat{\xi}) - \max_{\sigma, \xi} \ell(x_y, \sigma, \xi) \right\} \quad (4.3.1)$$

$$= 2 \{Pl(\hat{x}_y) - Pl(x_y)\}, \quad (4.3.2)$$

where  $Pl(x_y)$  denotes the profile likelihood of  $x_y$ . The likelihood-ratio test statistic (4.3.1) is asymptotically  $\chi_1^2$  distributed under some consistency and regularity conditions (see, *e.g.*, Pawitan (2013) §9.4, 9.5). One of these regularity conditions is that the MLE,  $\hat{\boldsymbol{\theta}}$ , is an interior point of the parameter space; this can cause problems in some extreme cases and is explored in §4.5. An approximate  $100(1 - \alpha)\%$  profile-likelihood (deviance-based) confidence interval for  $x_y$  is formed by finding the two points for which the deviance,  $D(x_y)$ , is equal to the  $(1 - \alpha)$  percentile of  $\chi_1^2$ :

$$[x_y : D(x_y) = \chi_{1, (1-\alpha)}^2].$$

The resulting interval can be asymmetric since it captures the possible asymmetry of the profile likelihood.

### Wald confidence intervals

The commonly used Wald confidence intervals are based on the asymptotic normality of the MLE,  $\hat{\boldsymbol{\theta}}$ . In practice one assumes

$$\sqrt{I(\hat{\boldsymbol{\theta}})}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \sim N(0, 1),$$

where  $I(\boldsymbol{\theta})$  is the observed Fisher information calculated as  $-\nabla\nabla^T\ell(\boldsymbol{\theta})$ .

To obtain a Wald confidence interval for one parameter,  $x_y$ , one assumes

$$W = \frac{\hat{x}_y - x_y}{se(\hat{x}_y)} \lesssim N(0, 1),$$

where the standard error of  $\hat{x}_y$ ,  $se(\hat{x}_y) = ([I(\hat{\boldsymbol{\theta}})^{-1}]_{11})^{\frac{1}{2}}$ , the squareroot of the 1st diagonal term of the inverse matrix of the observed Fisher information (Pawitan (2013) §9.7, 9.9). The resulting approximate  $100(1-\alpha)\%$  confidence interval is  $[\hat{x}_y \pm z_{\alpha/2} se(\hat{x}_y)]$ , where  $z_\alpha = \Phi^{-1}(\alpha)$  and  $\hat{x}_y$  is the ML estimate.

The Wald method uses similar distribution approximations to the deviance-based method;  $W^2$  is a quadratic approximation to the deviance  $D(x_y)$  so if  $D(x_y) \sim \chi_1^2$  then  $W^2 \sim \chi_1^2$  approximately and it follows that  $W \sim N(0, 1)$ . To obtain the approximation use the Taylor approximation around  $\hat{x}_y$  (assuming we can differentiate the profile likelihood,  $Pl(x_y)$ ):

$$\begin{aligned} D(x_y) &= 2 \{Pl(\hat{x}_y) - Pl(x_y)\} \\ &= 2 \left\{ Pl(\hat{x}_y) - \left[ Pl(\hat{x}_y) + \frac{d}{dx_y} Pl(x_y) \Big|_{x_y=\hat{x}_y} (x_y - \hat{x}_y) + \frac{1}{2} \frac{d^2}{dx_y^2} Pl(x_y) \Big|_{x_y=\hat{x}_y} (x_y - \hat{x}_y)^2 + \dots \right] \right\} \\ &\approx -\frac{d^2}{dx_y^2} Pl(x_y) \Big|_{x_y=\hat{x}_y} (x_y - \hat{x}_y)^2. \end{aligned} \quad (4.3.3)$$

It can be shown (*e.g.*, Pawitan (2013) §9.11) that the curvature of the profile likelihood is  $([I(\hat{\boldsymbol{\theta}})^{-1}]_{11})^{-1}$ , *i.e.*, it is equal to  $se(\hat{x}_y)^{-2}$ , so (4.3.3) =  $W^2$ .

The Wald confidence intervals are symmetric about the MLE; this is often undesirable since if the (profile) likelihood is asymmetrical the Wald confidence intervals will cover values with a lower likelihood than some not in the interval. Furthermore, if the MLE is not normally distributed the Wald confidence interval will not be correct, however, there may be some transformation of the MLE which is normally distributed. The profile-likelihood method essentially finds this transformation (if it exists) automatically and provides an in-

terval without needing to know the transformation. For the bootstrap methods described in the next sections we also use a normal assumption on the MLE transformed in some way.

### Standard (parametric) bootstrap

The data set generated from the true distribution and from which the MLEs are calculated is referred to as the primary data set. For the bootstrap methods multiple new datasets must be simulated, these new datasets should arise from the same sampling distribution as the original (primary) data set. Here the primary data set is simulated as in §4.2.2 with historical data given by (3.4.2) in Chapter 3 with  $n_0 = 10$ . The same historical data are used for all bootstrap samples associated with that primary data set since it represents the historical knowledge for that particular data set. The remainder of the bootstrap sample is generated using the fixed-threshold stopping rule from the GEV distribution with parameters estimated from the primary data set,  $\hat{\theta}^{\text{prim}}$ . Let the number of bootstrap samples be  $n_B$ ; then we have bootstrap samples  $\mathbf{x}_j$   $j = 1, \dots, n_B$  (each  $\mathbf{x}_j$  being a vector of length  $n_j^*$  determined by the stopping rule) and for each sample we find the maximum likelihood estimate,  $\hat{\theta}_j^*$ .

We assume there is some transformation,  $g$ , such that:

$$g(\hat{x}_y) - g(x_y) \sim \mathcal{N}(0, \sigma^2), \quad (4.3.4)$$

for some fixed  $\sigma^2$ , where  $x_y$  is the true  $y$ -year return level and  $\hat{x}_y := x_y(\hat{\theta})$  is the  $y$ -year return-level *estimator* based on the primary data set. Then an exact  $100(1 - \alpha)\%$  confidence interval for  $x_y$  is  $[g^{-1}(g(\hat{x}_y) \pm z_{\alpha/2}\sigma)]$  - this is the *percentile interval lemma* in Efron and Tibshirani (1993) §13.3. However,  $\sigma$ , is unknown so we use the bootstrap samples to find an approximate  $100(1 - \alpha)\%$  confidence interval.

Since the bootstrap samples are sampled in the same way as the primary sample we assume that  $g(\hat{x}_y^*)$ , where  $\hat{x}_y^* := x_y(\hat{\theta}^*)$  is the bootstrap  $y$ -year return-level estimator of  $g(x_y)$ , is similarly normally distributed but centred around the estimate  $g(\hat{x}_y^{\text{prim}})$  rather than the true  $g(x_y)$ . That is, we assume

$$g(\hat{x}_y^*) - g(\hat{x}_y^{\text{prim}}) \sim \mathcal{N}(0, \sigma^2), \quad (4.3.5)$$

where  $\hat{x}_y^{\text{prim}} = x_y(\hat{\theta}^{\text{prim}})$  is the fixed  $y$ -year return-level *estimate* based on the primary data set. We can also use the bootstrap samples to estimate the distribution of  $g(\hat{x}_y^*)$  by its

empirical distribution:

$$F_g^*(x) = \frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}\{g(\hat{x}_y^{*(j)}) \leq x\};$$

the larger  $n_B$  the better this approximation will be. So by assumption (4.3.5) we have that  $F_g^*(x) \approx \Phi\left(\frac{x - g(\hat{x}_y^{\text{prim}})}{\sigma}\right)$  as a function of  $x$ .

Consider  $F_g^*(x) = \frac{\alpha}{2}$  then

$$\begin{aligned} \frac{\alpha}{2} &\approx \Phi\left(\frac{F_g^{*-1}\left(\frac{\alpha}{2}\right) - g(\hat{x}_y^{\text{prim}})}{\sigma}\right) \\ \Rightarrow g(\hat{x}_y^{\text{prim}}) + \Phi^{-1}\left(\frac{\alpha}{2}\right)\sigma &\approx F_g^{*-1}\left(\frac{\alpha}{2}\right) \\ \Rightarrow g(\hat{x}_y^{\text{prim}}) + z_{\frac{\alpha}{2}}\sigma &\approx g(\hat{x}_y^*)_{\frac{\alpha}{2}}, \end{aligned}$$

where  $g(\hat{x}_y^*)_{\alpha/2}$  is the 100 $\frac{\alpha}{2}$  percentile of  $g(\hat{x}_y^*) = \{g(\hat{x}_y^{*(1)}), \dots, g(\hat{x}_y^{*(n_B)})\}$ . In other words the lower bound of the confidence interval for  $g(x_y)$ ,  $g(\hat{x}_y^{\text{prim}}) + z_{\alpha/2}\sigma$ , is approximately the 100 $\frac{\alpha}{2}$  percentile of the bootstrap estimates for  $g(x_y)$ . A similar argument follows for the upper bound and so we arrive at the following approximate 100(1 -  $\alpha$ )% confidence interval for  $g(x_y)$ :  $[g(\hat{x}_y^*)_{\frac{\alpha}{2}}, g(\hat{x}_y^*)_{1-\frac{\alpha}{2}}]$ .

For a monotone function,  $g$ , a percentile-based confidence interval is transform respecting and so the standard parametric bootstrap approximate 100(1 -  $\alpha$ )% confidence interval for  $x_y$  is:

$$[\hat{x}_{y(\frac{\alpha}{2})}^*, \hat{x}_{y(1-\frac{\alpha}{2})}^*], \quad (4.3.6)$$

where  $\hat{x}_{y(a)}^*$  is the 100 $a$  percentile of the bootstrap estimates  $\hat{x}_y^* = \{\hat{x}_y^{*(1)}, \dots, \hat{x}_y^{*(n_B)}\}$ . Therefore, just like the profile likelihood, we do not need to know the transform,  $g$ , in order to obtain a confidence interval for  $x_y$ , rather we just assume that there exists a function  $g$  for the which  $g(\hat{x}_y)$  is normally distributed.

### Efron's bias correction

Above we assumed that not only is the transformed MLE,  $g(\hat{x}_y)$ , normally distributed but also that it is an unbiased estimator of  $g(x_y)$ . In many cases it may be that such a function  $g$  that satisfies (4.3.4) and (4.3.5) does not exist and this false assumption will be reflected in large coverage error. Efron (1981) considers an extension of the standard bootstrap confidence

intervals, namely Efron's bias-corrected (BC) intervals, which allows bias in the normal assumption. In our setting, to use Efron's BC method, we assume there is some transformation,  $g$ , such that:

$$g(\hat{x}_y) - g(x_y) \sim \mathcal{N}(-b, \sigma^2), \quad (4.3.7)$$

for some fixed  $\sigma$ , where  $b$  is the bias term and  $x_y$  and  $\hat{x}_y$  are defined as in (4.3.4).

Similarly, we assume:

$$g(\hat{x}_y^*) - g(\hat{x}_y^{\text{prim}}) \sim \mathcal{N}(-b, \sigma^2), \quad (4.3.8)$$

with  $\hat{x}_y^{\text{prim}}$  and  $\hat{x}_y^*$  defined as in (4.3.5). Then an exact  $100(1 - \alpha)\%$  confidence interval for  $x_y$  is simply:

$$g^{-1}(g(\hat{x}_y) + b \pm \sigma z_{\frac{\alpha}{2}}). \quad (4.3.9)$$

A transformation which satisfies (4.3.8) approximately is  $g = \Phi^{-1}\hat{F}^*$ , where  $\hat{F}^*$  is the empirical cdf of  $\hat{x}_y^*$ :  $\hat{F}^*(x) = \frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}\{\hat{x}_y^{*(j)} \leq x\}$ . For the remainder of this section we continue with  $g = \Phi^{-1}\hat{F}^*$ . We have  $g(\hat{x}_y) = \Phi^{-1}\hat{F}^*(\hat{x}_y) \sim \mathcal{N}(0, 1)$  (with the approximation being better for larger  $n_B$ , the number of bootstrap samples). The expectation and variance of  $g(\hat{x}_y^*) - g(\hat{x}_y^{\text{prim}})$  is equal to  $-g(\hat{x}_y^{\text{prim}}) = -\Phi^{-1}\hat{F}^*(\hat{x}_y^{\text{prim}})$  and 1 respectively since  $g(\hat{x}_y^{\text{prim}})$  is fixed; so under assumption (4.3.8) we obtain the following bias and variance:

$$b = \Phi^{-1}\hat{F}^*(\hat{x}_y^{\text{prim}}) \quad \text{and} \quad \sigma^2 = 1.$$

Substituting this into (4.3.9) we obtain Efron's BC approximate  $100(1 - \alpha)\%$  confidence interval for  $x_y$ :

$$\hat{F}^{*-1}\Phi(2b \pm z_{\frac{\alpha}{2}}). \quad (4.3.10)$$

Note that  $\hat{F}^{*-1}$  is not well defined since  $\hat{F}^*$  is a step function. One convention is to take  $\hat{F}^{*-1}(x)$  to be the  $[n_B x]$ th smallest of the bootstrap estimates  $\hat{\mathbf{x}}_y^* = \{\hat{x}_y^{*(1)}, \dots, \hat{x}_y^{*(n_B)}\}$ . Notice that interval (4.3.10) is formed from percentiles of the bootstrap  $x_y$  estimates like the standard bootstrap CI; in particular, when the bias is 0 (corresponding to the MLE,  $\hat{x}_y^{\text{prim}}$ , being the median of the bootstrap MLEs) we recover the standard CI (4.3.6).

Similar to the standard bootstrap there may be cases for which a function  $g$  that satisfies (4.3.7) and (4.3.8) does not exist. The function  $g = \Phi^{-1}\hat{F}^*$  satisfies (4.3.8) approximately

but then the assumption (4.3.7) becomes

$$\Phi^{-1}\hat{F}^*(\hat{x}_y) \sim \mathcal{N}(\Phi^{-1}\hat{F}^*(x_y) - b, 1), \quad (4.3.11)$$

and this may not be satisfied. We now considered generalising the assumption further by allowing the standard error of  $g(\hat{x}_y)$  to vary with  $x_y$ .

### Variations on Efron's BC interval

Efron (1987) introduce the more general BCa confidence interval which corrects for the standard error of  $g(\hat{\theta})$  varying with  $\theta$  by introducing an acceleration quantity,  $a$ , which is the fixed rate of change in the standard error of  $g(\hat{\theta})$  with respect to  $g(\theta)$ . Efron describes ways to estimate this quantity using derivatives of the likelihood for one parameter models and also by using 'jackknife values' of  $g(\hat{x}_y)$ <sup>1</sup>. However, the estimation of  $a$  can be difficult, particularly for multi-parameter models. We attempt instead to use our knowledge of the stopping rule to estimate the standard error and obtain simpler confidence intervals. For clarity and simplicity from here onwards we refer to the ML parameter estimate based on the primary data set,  $\hat{\theta}^{\text{prim}}$ , as  $\hat{\theta}$  and denote the ML estimator by  $\hat{\Theta}$  so  $\hat{x}_y = x_y(\hat{\Theta})$ .

Assuming that the standard error of  $g(\hat{x}_y)$  depends on  $\theta$  we now have the following normal assumption:

$$g(\hat{x}_y) - g(x_y) \sim \mathcal{N}(-b, \sigma^2(\theta)), \quad (4.3.12)$$

where  $b$  is the bias term and  $x_y$  is defined as in (4.3.4). The variance of  $g(\hat{x}_y)$ ,  $\sigma^2(\theta)$ , is an unknown function of  $\theta$ . The bootstrap assumption follows with the standard error of  $g(\hat{x}_y^*)$  depending on  $\hat{\theta}$ :

$$g(\hat{x}_y^*) - g(\hat{x}_y^{\text{prim}}) \sim \mathcal{N}(-b, \sigma^2(\hat{\theta})), \quad (4.3.13)$$

with  $\hat{x}_y^{\text{prim}}$  and  $\hat{x}_y^*$  defined as in (4.3.5).

As for Efron's BC method we use  $g = \Phi^{-1}\hat{F}^*$ , so  $g(\hat{x}_y^*) = \Phi^{-1}\hat{F}^*(\hat{x}_y^*) \sim \mathcal{N}(0, 1)$ . So the expectation and variance of  $g(\hat{x}_y^*) - g(\hat{x}_y^{\text{prim}})$  is approximately equal to  $-g(\hat{x}_y^{\text{prim}}) = -\Phi^{-1}\hat{F}^*(\hat{x}_y^{\text{prim}})$  and 1 respectively since  $g(\hat{x}_y^{\text{prim}})$  is fixed. Then under assumption (4.3.13)

<sup>1</sup>The jackknife value  $\hat{\theta}_{(i)}$  is the MLE based on deleting the  $i$ th observation. To calculate  $\hat{a}$  one needs to calculate many jackknife values:  $\hat{\theta}_{(i)}$   $i = 1, \dots, n$ .

we obtain the following approximate bias and variance:

$$b = \Phi^{-1} \hat{F}^*(\hat{x}_y^{\text{prim}}) \quad \text{and} \quad \sigma^2(\hat{\theta}) = 1.$$

From (4.3.12) an approximate confidence interval for  $g(x_y)$  is

$$g^{-1}(2b \pm \sigma(\theta)z_\alpha), \quad (4.3.14)$$

where  $b = \Phi^{-1} \hat{F}^*(\hat{x}_y^{\text{prim}})$  and  $\sigma(\theta)$  is unknown.

Let  $N(\theta)$  and  $N(\hat{\theta})$  be the sample sizes when sampling using the fixed-threshold stopping rule from the GEV with parameters  $\theta$  and  $\hat{\theta}$  respectively. The variance of a function of the MLE is inversely proportional to the *known* sample size:

$$\text{Var}(g(\hat{x}_y)|N(\theta) = n) = \frac{C(\theta)}{n} \quad \text{and} \quad \text{Var}(g(\hat{x}_y^*)|N(\hat{\theta}) = n^*) = \frac{C(\hat{\theta})}{n^*}, \quad (4.3.15)$$

where  $n$  is the primary sample size,  $n^*$  is the sample size of a *particular* bootstrap sample and  $C(\theta)$  and  $C(\hat{\theta})$  are constants depending on  $\theta$  and  $\hat{\theta}$  respectively. Here we ignore the effect of the historical sample as that is fixed over all bootstrap samples and is equal to the historical sample of the primary data set. Recall that the sample size,  $N$ , under the fixed-threshold stopping rule is geometric with expectation equal to the inverse probability of exceeding  $c$  and let  $k(\theta) := \mathbb{E}[N(\theta)] = (\mathbb{P}(X > c|\theta))^{-1}$  and  $k(\hat{\theta}) := \mathbb{E}[N(\hat{\theta})] = (\mathbb{P}(X > c|\hat{\theta}))^{-1}$ . Using the law of total variance we have for *known*  $\hat{\theta}$  (and so fixed known  $\hat{x}_y^{\text{prim}}$ ):

$$\begin{aligned} \sigma^2(\hat{\theta}) &= \text{Var}(g(\hat{x}_y^*)) = \mathbb{E} \left[ \text{Var} \left( g(\hat{x}_y^*) | N(\hat{\theta}) \right) \right] + \text{Var} \left( \mathbb{E} \left[ g(\hat{x}_y^*) | N(\hat{\theta}) \right] \right) \\ &= C(\hat{\theta}) \mathbb{E} \left[ \frac{1}{N(\hat{\theta})} \right] + \text{Var} \left( g(\hat{x}_y^{\text{prim}}) - b \right) \quad (\text{by (4.3.13)}) \\ &= C(\hat{\theta}) \sum_{n=1}^{\infty} \frac{1}{k(\hat{\theta})} \frac{(1 - 1/k(\hat{\theta}))^{n-1}}{n} \quad (\text{from Geometric distribution}) \\ &= C(\hat{\theta}) \frac{\log k(\hat{\theta})}{k(\hat{\theta}) - 1}. \end{aligned}$$

In the second step  $\text{Var}(g(\hat{x}_y^{\text{prim}}) - b) = 0$  since  $g(\hat{x}_y^{\text{prim}})$  is fixed given  $\hat{\theta}$  is known. Similarly, for fixed  $\theta$  and *unknown random variable*  $\hat{\Theta}$ :

$$\sigma^2(\theta) = \text{Var}(g(\hat{x}_y)) = C(\theta) \frac{\log k(\theta)}{k(\theta) - 1}.$$

We assume that the variance of  $g(\hat{x}_y)$  based on a sample of size  $n$  is approximately the same as the variance of  $g(\hat{x}_y^*)$  based on a bootstrap sample also of size  $n$ . Then  $C(\boldsymbol{\theta}) \approx C(\hat{\boldsymbol{\theta}})$  and, since  $\sigma^2(\hat{\boldsymbol{\theta}}) = 1$ ,

$$\sigma^2(\boldsymbol{\theta}) = \frac{\log k(\boldsymbol{\theta})}{k(\boldsymbol{\theta}) - 1} \times \frac{k(\hat{\boldsymbol{\theta}}) - 1}{\log k(\hat{\boldsymbol{\theta}})} \approx \frac{k(\hat{\boldsymbol{\theta}})}{k(\boldsymbol{\theta})},$$

since  $k(\boldsymbol{\theta}) \gg 1$  and  $k(\hat{\boldsymbol{\theta}}) \gg 1$  in practice and  $\log k(\boldsymbol{\theta})/\log k(\hat{\boldsymbol{\theta}}) \approx 1$  as  $\hat{\boldsymbol{\theta}}$  is close to  $\boldsymbol{\theta}$ .

Now  $k(\boldsymbol{\theta})$  is unknown so we use the bootstrap assumption that the bias of the MLE is approximately the same as the bias in the bootstrap estimates:

$$\mathbb{E} \left[ \log k(\hat{\boldsymbol{\Theta}}) \right] - \log k(\boldsymbol{\theta}) \approx \mathbb{E} \left[ \log k(\hat{\boldsymbol{\theta}}^*) \right] - \log k(\hat{\boldsymbol{\theta}}), \quad (4.3.16)$$

where  $\hat{\boldsymbol{\Theta}}$  is the ML estimator and  $\hat{\boldsymbol{\theta}}$  is the fixed known ML estimate. We further approximate (4.3.16) by approximating the expectations:

$$\log k(\hat{\boldsymbol{\theta}}) - \log k(\boldsymbol{\theta}) \approx \log \hat{k}^* - \log k(\hat{\boldsymbol{\theta}}), \quad (4.3.17)$$

where  $\hat{\boldsymbol{\theta}}$  is the ML estimate and  $\hat{k}^*$  is the median of the bootstrap estimates of the inverse probability of exceedance. The usual Monte Carlo approximation of  $\mathbb{E} \left[ \log k(\hat{\boldsymbol{\theta}}^*) \right]$  would be the empirical mean:  $\frac{1}{n_B} \sum_{j=1}^{n_B} \log k(\hat{\boldsymbol{\theta}}_j^*)$ . We instead use the median of the estimates as it is robust and invariant to monotonic transformations, however we note in our setting that  $\frac{1}{n_B} \sum_{j=1}^{n_B} \log k(\hat{\boldsymbol{\theta}}_j^*) > \text{med}(\log k(\hat{\boldsymbol{\theta}}_j^*))$ .

Then (4.3.17) leads to the following approximation:

$$\sigma^2(\boldsymbol{\theta}) \approx \frac{k(\hat{\boldsymbol{\theta}})}{k(\boldsymbol{\theta})} \approx \frac{\hat{k}^*}{k(\hat{\boldsymbol{\theta}})}. \quad (4.3.18)$$

This approximation is expected to be an underestimation of the true standard error of  $\hat{\boldsymbol{\theta}}$  due to the median approximation.

Substituting (4.3.18) into (4.3.14) we obtain the following confidence interval which we refer to as the bias and variance-corrected (BVC) interval:

$$g^{-1} \left( 2b \pm \sqrt{\frac{\hat{k}^*}{k(\hat{\boldsymbol{\theta}})} z_\alpha} \right). \quad (4.3.19)$$

We also consider the case where  $k(\hat{\boldsymbol{\theta}})$  is replaced by the size of the primary sample,  $n$ , which

is a direct estimate of  $k(\hat{\theta})$  resulting in the following confidence interval:

$$g^{-1} \left( 2b \pm \sqrt{\frac{\hat{k}^*}{n}} z_\alpha \right). \quad (4.3.20)$$

We call this the bias and variance-corrected interval with the estimate  $n$  (BVCn).

### 4.3.2 Comparison of confidence interval methods

In the following we concentrate on primary data sets of size  $n = 48$  simulated from the  $\text{GEV}(\hat{\theta}_{std}^{Lune})$  with the fixed-threshold stopping rule. In a similar manner to §4.2 we compare coverages (Figure 4.3.1), confidence interval widths using boxplots (Figures 4.3.2, B.2.3-B.2.6) and the interval values (Figures 4.3.3, B.2.7-B.2.10). The latter figures are obtained by counting the percentage of times each  $x$ -value is contained within the confidence interval, essentially providing the coverage for each value on the  $x$ -axis. The dashed line is the true return level,  $x_y$ , for the return period,  $y$ , stated in the plot captions so reading the percentages on the  $y$  axis where the dashed lines cross the curves gives the coverage as in Figure 4.3.1.

Before discussing the results we highlight the problem with naively ignoring the final observation. The fact that this value has been observed needs to be included in the analysis in some way to ensure one does not arrive at impossible estimates since when using the exclude likelihood  $k(\hat{\theta})$  can be estimated to be infinite (*i.e.*,  $c$  is larger than the estimated upper bound). In our simulations  $k(\hat{\theta})$  was infinite in 3(11)% of cases on average when  $k = 50(500)$ ; these cases were omitted when using the BVC (4.3.19) and BVCn (4.3.20) confidence interval method. Care must also be taken when maximising the likelihood to find the bootstrap MLEs,  $\hat{\theta}^*$ , when the final observation is close to  $c$ . This occurs particularly when  $\hat{\xi}$  is very small, *e.g.*,  $-0.2$  and  $k(\hat{\theta})$  is large; such  $\hat{\xi}$  and  $k(\hat{\theta})$  values are more likely when using the exclude likelihood to find  $\hat{\theta}$  since  $\hat{\xi}_{ex}$  is negatively biased (Chapter 3). Bootstrap samples with parameters based on the exclude likelihood are then mostly large in size and have a very light tail with upper bound close to  $c$ . Optimising the likelihood with such bootstrap samples can be difficult since the MLE is near the boundary of the parameter space; this issue is explored further in §4.5 with the partial and full conditioning likelihoods.

In terms of coverage the profile-likelihood confidence intervals outperform the bootstrap confidence intervals considered, especially when  $k$  is large compared to  $n$ . The standard bootstrap performs badly since it is highly influenced by the bias of the MLE of the primary

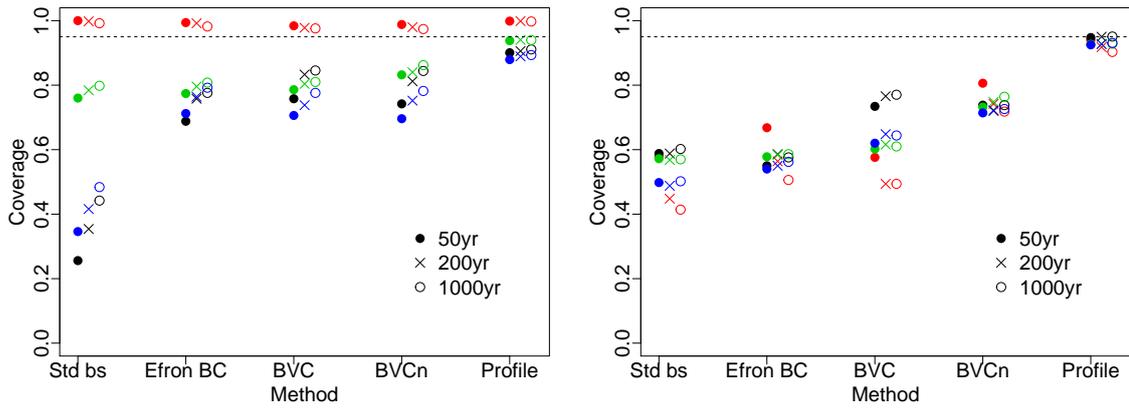


Figure 4.3.1: Coverage of the  $y$ -year return level,  $x_y$ , found using different confidence interval methods with  $\ell_{std}$  (red),  $\ell_{ex}$  (black),  $\ell_{fc}$  (green) and  $\ell_{pc}$  (blue). Primary data sets simulated from  $GEV(\hat{\theta}_{std}^{Lune})$  with sample size,  $n$ , determined by the fixed-threshold stopping rule with stopping threshold  $c_{50}$  (left) and  $c_{500}$  (right) such that  $n = 48$ . Bootstrap samples are created from the same sampling process as the primary data set with no restriction on sample size.

data set. As expected the bias-reducing bootstrap variations improve upon the standard bootstrap, with BVCn having the best coverage properties of these, however the coverage is still very poor when  $k$  is large (e.g., Figure 4.3.1 right panel).

The properties of the confidence intervals can be quite different when using the standard likelihood ( $\ell_{std}$ ) compared to using one of the other three likelihoods ( $\ell_{ex}, \ell_{fc}, \ell_{pc}$ ) so we now consider these cases separately.

Firstly, using the standard likelihood results in undesirable confidence interval properties with all methods considered. The upper confidence interval bounds for  $x_y^{std}$  are almost always much larger than  $x_y$  (Figure B.2.2 bottom row) whereas coverage is low when  $k \gg n$  due to the lower bounds being too large (Figure B.2.2 middle); these behaviours indicate positive bias in the confidence intervals using  $\ell_{std}$ . This bias can also be clearly seen in the interval plots (Figures 4.3.3, B.2.7-B.2.10 top left); using  $\ell_{std}$  all of the confidence intervals have a similar lower bound and cover approximately the same values which are around and larger than  $x_y$  (resulting in the ‘square-shaped’ curves centred around or above the dashed vertical line). When  $k$  is large compared to  $n$ , the primary data set used includes an observation which is expected to be seen in a much larger sample. This leads to a negative bias in  $k(\hat{\theta})$  and positive bias in  $\hat{\xi}_{std}$  which is then used to generate bootstrap samples. Hence the bootstrap samples will be smaller and contain larger values generally than there would be using the true  $\theta$  and  $k$ . As a result many bootstrap confidence intervals do not cover  $x_y$  but values

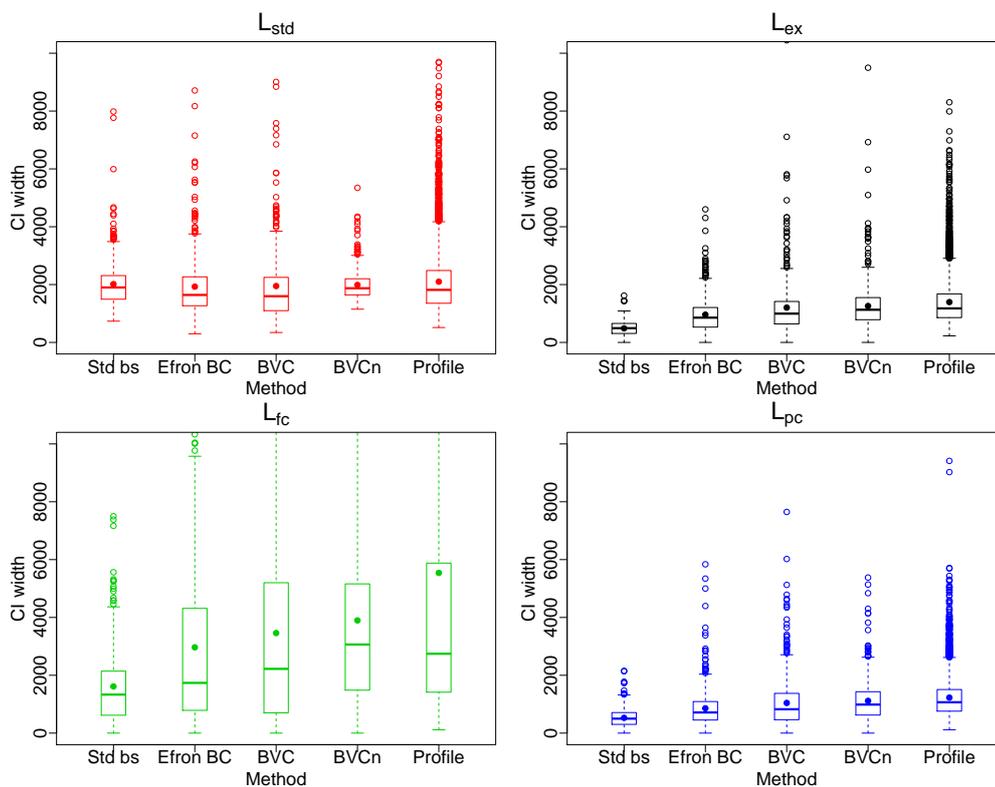


Figure 4.3.2: Box plots of confidence interval widths for  $x_{200}$ , the 200-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ . Crosses indicate the mean confidence interval width for each method.

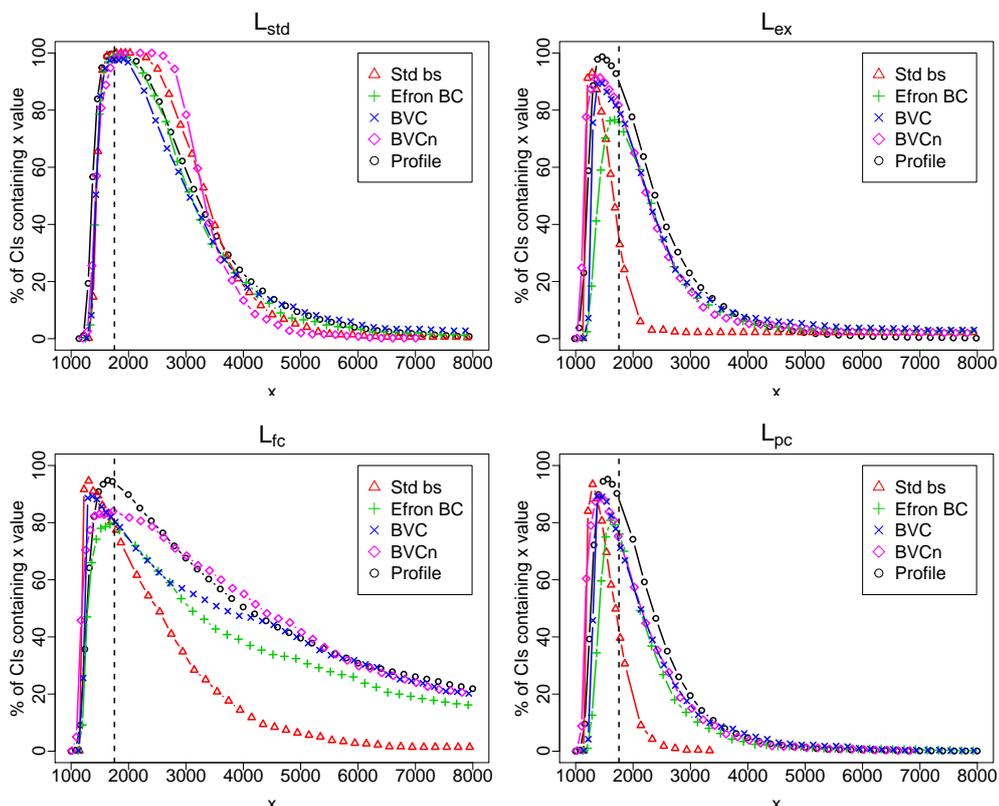


Figure 4.3.3: Confidence intervals for  $x_{200}$ , the 200-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ .

larger than  $x_y$  as seen by the peaks for the bootstrap methods in Figures B.2.9-B.2.10 being shifted to the right of the true  $y$ -year return level.

From the top left panels of Figures 4.3.2, B.2.3-B.2.4 we can see that the confidence interval widths when using  $\ell^{std}$  and  $k = 50$  are mostly similar across all methods (the lower and upper quartiles are around similar values), however, the profile likelihood method results in a higher percentage of large widths than the bootstrap methods. The BVCn method generally results in the smallest confidence intervals with  $k = 50$  whereas for larger  $k$  the method produces the largest confidence intervals (Figures B.2.5-B.2.6) since  $n = 48$  is a poor estimate of  $k(\hat{\theta})$  when the true  $k$  is 500. When  $k$  is large the Efron BC and BVC confidence intervals are generally narrower than the standard bootstrap intervals, particularly when estimating  $x_y$  with large return period,  $y$ .

In contrast using  $\ell_{ex}$ ,  $\ell_{fc}$  or  $\ell_{pc}$  the standard bootstrap confidence interval widths have the smallest mean and variation; Figures 4.3.3, B.2.7-B.2.10 show that the standard bootstrap intervals cover a much smaller set of values than the other intervals, in particular less high values. The intervals of the other bootstrap methods are more variable and larger in width, however, they also mostly cover values lower than the truth (the corresponding peaks in the figures are left of the dashed line). This results in undercoverage as the upper bounds are lower than the truth more often than in the desired 2.5% of cases.

Overall out of all the methods considered here the profile-likelihood based confidence intervals appear to have the best performance. There is a clear trade off between coverage and confidence interval width; the bootstrap-based methods often result in narrower intervals but poorer coverage compared to the profile-likelihood method. These intervals tend to be too high when using the standard likelihood and too low when using the exclude, full or partial conditioning likelihoods due to the positive/negative bias of the parameter estimates using these likelihoods. However it may be possible to improve upon these bootstrap methods to obtain narrower intervals than the profile-likelihood based intervals while retaining the good coverage properties of the latter. For example, it would be desirable to have a confidence interval for the  $y$ -year return level with the widths and tails similar to the standard bootstrap intervals but shifted such that  $x_y$  is more frequently covered.

If one instead considers a random sample size,  $N$ , for the primary data set as in §4.2.1 the bootstrap methods perform much better for large  $k$  (Figure B.2.1). In particular the lower bounds when using  $\ell_{std}$  are reduced and the upper bounds for the other likelihoods are

higher so that  $x_y$  is within the confidence intervals more often. In this case, the coverage of BVCn confidence intervals are similar to that of the profile method. However, the BVCn confidence intervals can be very wide (and are wider than the profile likelihood intervals when  $k = 500$ ) and they also take more computational power to produce compared to the profile likelihood method. Therefore, the profile-likelihood method is the preferred method to obtain confidence intervals for  $x_y$ , especially when the sample size of the primary data set is small.

#### 4.4 Bootstrap CIs with variable-threshold stopping rule

In the previous section we concentrated on bootstrap samples of random size as determined by the fixed-threshold stopping rule. Generating bootstrap samples with different sizes seems sensible as the sample size is variable for the data-generating mechanism of the primary data sets. The confidence intervals produced in this way are answering the question ‘*Given that we’ve stopped at a random time and estimated the parameters, what is the uncertainty in the  $\hat{x}_y$ ?*’. Creating bootstrap samples in this setting is simple as it just requires sequential sampling from GEV with the parameters estimated from the primary data set until the threshold (fixed  $c$  or the  $k$ -year return level estimated using the bootstrap sample up to that point) is exceeded.

Alternatively we could consider generating bootstrap samples of the same size as the primary data set. Then the question of interest would be ‘*What is the uncertainty in  $\hat{x}_y$  given that the random stopping time according to the stopping rule is  $n$ ?*’. By conditioning on  $n$  we are removing some of the uncertainty and the resulting confidence intervals are smaller than when we allow the bootstrap sample size to be random. With the fixed-threshold stopping rule, this results in smaller coverage for the standard likelihood as the upper confidence interval bounds are lower than those for random bootstrap sample sizes and the opposite effect occurs for the other likelihoods. For the variable-threshold stopping rule the bootstrap sampling procedure is more complex; the rest of this section is dedicated to developing an efficient way to do this bootstrap sampling.

Generating bootstrap samples with the same size,  $n$ , as the primary data set is more complex than allowing for variable sample sizes; one cannot simply sample sequentially from the GEV with parameters  $\hat{\theta}$  truncated at  $\hat{x}_k(\mathbf{x}_{1:i-1})$  for  $n_0 + 1 < i < n$ . A correct, but inefficient, procedure would be to use rejection sampling; sampling  $n - n_0$  times from the

GEV distribution with  $\hat{\boldsymbol{\theta}}$  and accepting the bootstrap data set if the only exceedance occurs on the final sample, *i.e.*,  $\mathbb{1}\{x_n > \hat{x}_k(\mathbf{x}_{1:n-1})\} = 1$  and  $\sum_{i=n_0+1}^n \mathbb{1}\{x_i > \hat{x}_k(\mathbf{x}_{1:i-1})\} = 1$ . There is likely to be a very large number of rejected samples using such a procedure since it requires exactly the  $n$ th sampled point to be an exceedance and also that all previous sampled points are non-exceedances.

We now suggest an importance-weighted (recall §2.2.2) bootstrap sampling procedure using our knowledge of the stopping rule. As in §3.2.2, we define a stopping region  $\mathcal{S}_n = \mathcal{S}_n(\mathbf{x}_{1:n-1})$  such that we stop sampling if  $X_n \in \mathcal{S}_n$  and continue to sample otherwise. We refer to the true data-generating distribution, where  $f(\cdot; \boldsymbol{\theta})$  is the density of the GEV distribution with parameters  $\boldsymbol{\theta}$ , as

$$\begin{aligned} h(\mathbf{x}; \boldsymbol{\theta}) &= \left( \prod_{i=n_0+1}^n f(x_i; \boldsymbol{\theta}) \right) \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=n_0+1}^{n-1} \mathbb{1}_{\mathcal{S}_i^c}(x_i) \\ &= \left( \prod_{i=n_0+1}^n f(x_i; \boldsymbol{\theta}) \right) \mathbb{1}\{x_n > \hat{x}_k(\mathbf{x}_{1:n-1})\} \prod_{i=n_0+1}^{n-1} \mathbb{1}\{x_i \leq \hat{x}_k(\mathbf{x}_{1:i-1})\}. \end{aligned}$$

The inefficient rejection sampling method above simulates from the distribution  $h(\mathbf{x}; \hat{\boldsymbol{\theta}})$  to obtain bootstrap samples.

Instead of sampling from the true data-generating distribution,  $h(\mathbf{x}; \boldsymbol{\theta})$ , to simulate bootstrap samples we sample from some *proposal* distribution,  $q(\mathbf{x})$  and obtain a weighted sample of bootstrap return-level estimates. Consider sampling sequentially from the GEV, fitted with the primary data set, right-truncated at  $\hat{x}_k(\mathbf{x}_{1:i-1})$  for  $i < n$  and left-truncated for  $i = n$ . The density of this proposal would be

$$q(\mathbf{x}; \hat{\boldsymbol{\theta}}) = \left( \prod_{i=n_0+1}^n \frac{f(x_i; \hat{\boldsymbol{\theta}})}{F(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i \neq n\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i=n\}}} \right) \mathbb{1}_{\mathcal{S}_n}(x_n) \prod_{i=n_0+1}^{n-1} \mathbb{1}_{\mathcal{S}_i^c}(x_i). \quad (4.4.1)$$

The corresponding weights are

$$\omega(\mathbf{x}) = \frac{h(\mathbf{x}; \hat{\boldsymbol{\theta}})}{q(\mathbf{x}; \hat{\boldsymbol{\theta}})} = \prod_{i=n_0+1}^n F(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i \neq n\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i=n\}}. \quad (4.4.2)$$

For each bootstrap sample,  $\mathbf{x}^{(j)}$   $j = 1, \dots, n_B$ , generated from  $q(\mathbf{x}; \hat{\boldsymbol{\theta}})$  we calculate the  $y$ -year return-level estimate  $\hat{x}_y^{\star(j)} = x_y(\hat{\boldsymbol{\theta}}^{\star})$ . Then we can take  $\{\hat{x}_y^{\star(j)}, \tilde{\omega}^{(j)}\}_{j=1}^m$ , with normalised weights  $\tilde{\omega}^{(j)} = \omega(\mathbf{x}^{(j)}) / \sum_{l=1}^m \omega(\mathbf{x}^{(l)})$ , as a weighted sample of  $y$ -year return-level estimates

from the true data-generating mechanism.

Recall the approximate  $100(1 - \alpha)\%$  standard bootstrap confidence interval for  $\hat{x}_y$  is  $[\hat{F}^{\star-1}(\alpha/2), \hat{F}^{\star-1}(1 - \alpha/2)]$  where  $\hat{F}^{\star}$  is the empirical CDF of the bootstrap  $y$ -year return-level estimator:

$$\hat{F}^{\star}(x) = \frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}\{\hat{x}_y^{\star(j)} \leq x\}.$$

In practice we find the interval  $[\hat{x}_y^L, \hat{x}_y^U]$  by solving

$$\frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}\{\hat{x}_y^{(j)} \leq \hat{x}_y^L\} = \frac{\alpha}{2} \quad \text{and} \quad \frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}\{\hat{x}_y^{(j)} \leq \hat{x}_y^U\} = 1 - \frac{\alpha}{2}.$$

The bootstrap confidence interval with our importance sampling method outlined above is  $[\hat{x}_y^L, \hat{x}_y^U]$  with

$$\sum_{j=1}^{n_B} \tilde{\omega}^{(j)} \mathbb{1}\{\hat{x}_y^{(j)} \leq \hat{x}_y^L\} = \frac{\alpha}{2} \quad \text{and} \quad \sum_{j=1}^{n_B} \tilde{\omega}^{(j)} \mathbb{1}\{\hat{x}_y^{(j)} \leq \hat{x}_y^U\} = 1 - \frac{\alpha}{2}.$$

Using the variable-threshold stopping rule with the Lune data there were at least two exceedances (depending on the chosen  $k$ ) and multiple exceedances could be common in practice. The procedure when the bootstrap sample has random sample size is easily extended to multiple exceedances but the fixed bootstrap sample procedure requires more thought. We outline the importance-weighted bootstrap sampling procedure when there are exactly  $r$  exceedances with the final exceedance being the final sampled point. The true data-generating distribution in this setting is

$$h(\mathbf{x}; \boldsymbol{\theta}) = \left( \prod_{i=n_0+1}^n f(x_i; \boldsymbol{\theta}) \right) \mathbb{1}\{x_n > \hat{x}_k(\mathbf{x}_{1:n-1})\} \mathbb{1}\left\{ \sum_{i=n_0+1}^{n-1} \mathbb{1}\{x_i > \hat{x}_k(\mathbf{x}_{1:i-1})\} = r - 1 \right\}.$$

Let  $\mathcal{I}^*$  be the set of all subsets  $\mathbf{i}^*$  of size  $r - 1$  such that  $\mathbf{i}^* \subset \{1, \dots, n - 1\}$ . Every  $\mathbf{i}^* \in \mathcal{I}^*$  is a possible set of the  $r - 1$  exceedance ‘times’ before the final exceedance at ‘time  $n$ ’. We denote the region corresponding to these exceedances and non-exceedances by  $\mathcal{R}_{\mathbf{i}^*}$ , *i.e.*, we let

$$\mathcal{R}_{\mathbf{i}^*} = \{x_n > \hat{x}_k(\mathbf{x}_{1:n-1})\} \bigcup_{\mathbf{i} \in \mathbf{i}^*} \{x_i > \hat{x}_k(\mathbf{x}_{1:i-1})\} \bigcup_{\mathbf{i} \in \{1, \dots, n-1\} \setminus \mathbf{i}^*} \{x_i < \hat{x}_k(\mathbf{x}_{1:i-1})\}.$$

Then we can rewrite  $h(\mathbf{x}; \boldsymbol{\theta})$  as a sum over the possible exceedance times:

$$h(\mathbf{x}; \boldsymbol{\theta}) = \sum_{\mathbf{i}^* \in \mathcal{I}^*} h^*(\mathbf{x}, \mathbf{i}^*; \boldsymbol{\theta}) = \sum_{\mathbf{i}^* \in \mathcal{I}^*} \left( \prod_{i=1}^n f(x_i; \boldsymbol{\theta}) \mathbb{1}\{\mathcal{R}_{\mathbf{i}^*}\} \right).$$

Consider the following proposal: Let the times of exceedance,  $\mathbf{i}^*$ , be randomly sampled from  $\{1, 2, \dots, n-1\}$  without replacement; then sequentially simulate from the GEV( $\hat{\boldsymbol{\theta}}$ ) (with parameters estimated from the primary data set) left-truncated at the  $k$ -year return-level estimate,  $\hat{x}_k(\mathbf{x}_{1:i-1})$ , for the exceedances  $i \in \mathbf{i}^* \cup \{n\}$  and right-truncated otherwise. This proposal has the joint density function,

$$q(\mathbf{x}, \mathbf{i}^*; \hat{\boldsymbol{\theta}}) = q(\mathbf{x}|\mathbf{i}^*) \cdot q(\mathbf{i}^*), \quad (4.4.3)$$

where  $q(\mathbf{i}^*) = (r-1)!(n-r)!/(n-1)!$  is the distribution of possible  $r-1$  exceedance times and

$$q(\mathbf{x}|\mathbf{i}^*) = \left( \prod_{i=n_0+1}^n \frac{f(x_i; \hat{\boldsymbol{\theta}})}{F(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{1-\mathbb{1}\{i \in \mathbf{i}^*\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i \in \mathbf{i}^*\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i=n\}}} \right) \mathbb{1}\{\mathcal{R}_{\mathbf{i}^*}\}. \quad (4.4.4)$$

Proposal distribution (4.4.3) is an extension of (4.4.1) to multiple exceedances. The product in (4.4.4) is the likelihood conditioning on exceedances at times  $\mathbf{i}^*$  and  $n$  and non-exceedance at other times. The factor  $\mathbb{1}\{\mathcal{R}_{\mathbf{i}^*}\}$  is needed such that  $q(\mathbf{x}|\mathbf{i}^*) = 0$  if  $\mathbf{x}$  is outside the stopping region formed by exceedance times  $\mathbf{i}^*$  and  $n$ .

For each bootstrap sample,  $\mathbf{x}^{(j)}$   $j = 1, \dots, n_B$ , we sample a new set of  $r-1$  exceedances:  $\mathbf{i}^{*(j)}$ . Then we generate bootstrap samples from  $q(\mathbf{x}, \mathbf{i}^{*(j)}; \hat{\boldsymbol{\theta}})$ , calculate the  $y$ -year return level estimates  $\hat{x}_y^{*(j)} = x_y(\hat{\boldsymbol{\theta}}^{*(j)})$  and calculate corresponding weights;

$$\begin{aligned} \omega(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)}) &= \frac{h(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)}; \hat{\boldsymbol{\theta}})}{q(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)}; \hat{\boldsymbol{\theta}})} \\ &= C \prod_{i=n_0+1}^n F(\hat{x}_k(\mathbf{x}_{1:i-1}^{(j)}); \hat{\boldsymbol{\theta}})^{1-\mathbb{1}\{i \in \mathbf{i}^{*(j)}\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}^{(j)}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i \in \mathbf{i}^{*(j)}\}} \bar{F}(\hat{x}_k(\mathbf{x}_{1:i-1}^{(j)}); \hat{\boldsymbol{\theta}})^{\mathbb{1}\{i=n\}}, \end{aligned} \quad (4.4.5)$$

where  $C = (n-1)!/((n-r)!(r-1)!)$ .

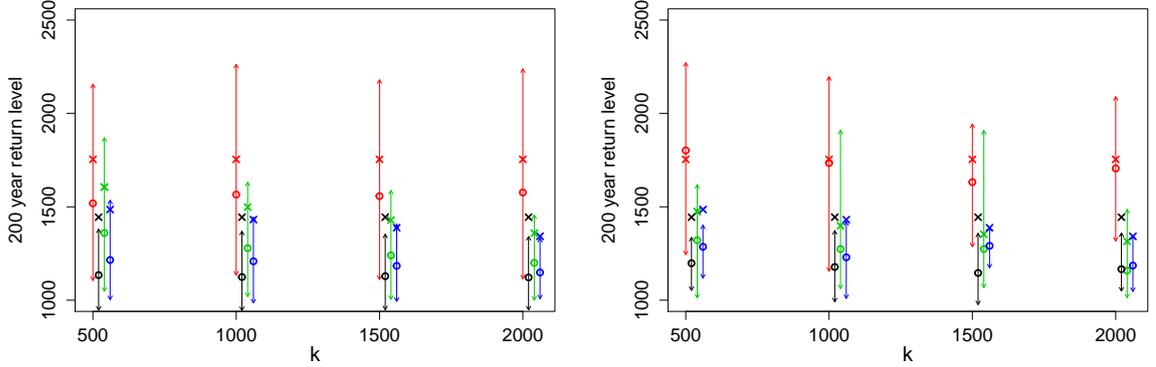


Figure 4.4.1: 200-year return-level estimates (crosses) based on all the data up to and including 2015 for the Lune at Caton with 95% bootstrap confidence intervals with the variable-threshold stopping rule over a range of  $k$ . Left: assuming one exceedance, right: with two exceedances. The means of the 500 bootstrap return-level estimates for each  $k$  and likelihood are shown as circles. Each group of 4 estimates applies for the same  $k$  as for the standard estimate in each group and have been horizontally shifted for clarity.

As above we normalise the weights,

$$\tilde{\omega}^{(j)} := \frac{\omega(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)})}{\sum_{l=1}^{n_B} \omega(\mathbf{x}^{(l)}, \mathbf{i}^{*(l)})},$$

to obtain a weighted sample of  $y$ -year return level estimates,  $\{\hat{x}_y^{(j)}, \tilde{\omega}^{(j)}\}_{j=1}^{n_B}$ .

A useful measure of the benefit of the importance-weighted bootstrap method is the effective sample size (ESS):

$$ESS = \frac{(\sum_{j=1}^{n_B} \omega(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)}))^2}{\sum_{j=1}^{n_B} (\omega(\mathbf{x}^{(j)}, \mathbf{i}^{*(j)}))^2} = \frac{1}{\sum_{j=1}^{n_B} (\tilde{\omega}^{(j)})^2}.$$

An efficient sampling scheme would be one with an effective sample size close to  $n_B$  and so occurs when we have approximately equal importance weights. In contrast, if one weight is much larger than the others then the sample is essentially equivalent to just one independent sample.

We tested the above importance sampling methods with the Lune data set with  $500 < k < 2000$  for estimating the 200-year return level. Using the variable-threshold stopping rule with this range of  $k$  results in two exceedances, one in 1995 and the other final exceedance in 2015. First we ignore the knowledge that there was an exceedance in 1995 and use the

proposal (4.4.1) and weights (4.4.2) as if the only exceedance is the final exceedance. Then we consider the importance sampling method with two exceedances, *i.e.*, proposal (4.4.3) and weights (4.4.5) with  $r = 2$ . We compare the MLEs based on the primary data set and the mean bootstrap  $y$ -year return-level estimates and confidence intervals using the four likelihoods  $\ell_{std}$ ,  $\ell_{ex}$ ,  $\ell_{fc}$  and  $\ell_{pc}$  in Figure 4.4.1. For the two exceedance method we amend  $\ell_{fc}$  slightly by replacing the  $\ell_{fc}$  contribution of the  $i^*$ th observation (the first exceedance) by  $f(x_{i^*}; \boldsymbol{\theta}) / \bar{F}(\hat{x}_k(\mathbf{x}_{1:i^*-1}); \boldsymbol{\theta})$ .

The most striking observation is the strong negative bias in the bootstrap return-level estimator,  $\hat{x}_{200}^{*L}$ , when using the likelihoods  $L = \ell_{ex}$ ,  $\ell_{fc}$  or  $\ell_{pc}$  or using  $\ell_{std}$  assuming only one exceedance. We saw in Chapter 3 that for variable-threshold stopping rule the return-level estimators are negatively biased, except the standard estimator which has slight positive bias for low return periods and large  $\xi$ , so we expect negative bias when using  $\ell_{ex}$ ,  $\ell_{fc}$  or  $\ell_{pc}$ . Bootstrapping amplifies this negative bias since the bootstrap samples are from the GEV with a shape parameter which is likely to be smaller than the true shape parameter. This bias is also apparent in how far the MLEs are to the confidence interval upper bounds. For the standard likelihood the MLE lies comfortably within the bound, with  $\ell_{fc}$  it is further into the upper tail as the full-conditional estimator has negative bias,  $\hat{x}_y^{*pc}$  is more negatively biased and the MLE  $\hat{x}_y^{pc}$  is outside the confidence interval in some cases, and  $\hat{x}_y^{*ex}$  is the most negatively biased with corresponding intervals far below the MLE  $\hat{x}_y^{ex}$ .

The bias in the bootstrap estimates is reduced when the two exceedance method is employed since the resulting bootstrap samples correctly have two exceedances and so are more similar to the primary data set than the bootstrap samples under the one exceedance method. Using the one exceedance method the bootstrap samples are likely to have less ‘large’ values than the primary data set and so, for any likelihood, there is a negative bias in the bootstrap shape parameter estimator due to this false sampling. Thus, in particular, assuming one exceedance leads to a negative bias in the standard bootstrap return-level estimator whereas the bias is small when bootstrapping two exceedances (the crosses and circles are closer in Figure 4.4.1 right panel). Moreover, for  $\ell_{std}$ ,  $\ell_{ex}$  and  $\ell_{pc}$  the confidence intervals are narrower when including a second exceedance as the extra exceedance information reduces the uncertainty. However, when assuming one exceedance we found the ESS as a percentage of total bootstraps,  $n_B$ , ranged from 40-54% whereas using two exceedances it was only 16.8-34.8%. For both the one and two exceedance methods the effective sample sizes are smaller when  $k$

is larger.

Finally we note that, similar to the findings of §4.3.2, the confidence intervals for both methods are narrower than those obtained using the profile-likelihood deviance method (cf. Chapter 3 Figure 3.5.1 right panel).

In conclusion, the bootstrap intervals obtained using the correct (multiple exceedance) procedure are reasonable when using  $\ell_{std}$  and are narrower than the profile-likelihood based intervals. However, for other likelihoods the intervals are too heavily influenced by the negative bias in the estimators and so some sort of bias correction is needed, perhaps incorporating Efron's BC confidence intervals. More exploration into these intervals in the variable-threshold stopping rule setting would be useful, including a simulation study to look at coverage properties etc.

## 4.5 Influence of $c$ on conditional CIs

In Chapter 3 we saw that the profile-likelihood based confidence intervals from the two conditioning methods decrease in width with increasing stopping threshold  $c$ /return period of stopping threshold  $k$ . Here we concentrate on the fixed-threshold stopping rule and investigate the impact of the threshold,  $c$ , in more depth. Recall that  $c$  must lie between the largest and second largest observed value; for  $c$  close to the largest value,  $x_n$ , maximising the likelihood can be difficult and so care must be taken in the calculation of confidence intervals.

The full-conditional and partial-conditional likelihoods (recall §3.3.1) are:

$$L_{fc}(\boldsymbol{\theta}) = \left( \prod_{i=1}^{n-1} \frac{f(x_i)}{F(c)} \right) \frac{f(x_n)}{\overline{F}(c)} \quad L_{pc}(\boldsymbol{\theta}) = \left( \prod_{i=1}^{n-1} f(x_i) \right) \frac{f(x_n)}{\overline{F}(c)},$$

where  $f$  and  $F$  are the pdf and cdf respectively of the GEV distribution with parameters  $\boldsymbol{\theta} = (\mu, \sigma, \xi)$ . Now, for  $c \gg \mu$ , using the Taylor expansion on the denominator we have

$$\frac{f(x_n)}{\overline{F}(c)} = \frac{F(x_n) [1 + \xi (\frac{x_n - \mu}{\sigma})]^{-(1/\xi+1)}}{\sigma \left( 1 - \exp \left\{ - [1 + \xi (\frac{c - \mu}{\sigma})]^{-1/\xi} \right\} \right)} \approx \frac{F(x_n) [1 + \xi (\frac{x_n - \mu}{\sigma})]^{-(1/\xi+1)}}{\sigma [1 + \xi (\frac{c - \mu}{\sigma})]^{-1/\xi}}. \quad (4.5.1)$$

When  $c$  is very close to  $x_n$ ,

$$\frac{F(x_n) [1 + \xi (\frac{x_n - \mu}{\sigma})]^{-(1/\xi+1)}}{\sigma [1 + \xi (\frac{c - \mu}{\sigma})]^{-1/\xi}} \approx \frac{F(x_n)}{\sigma [1 + \xi (\frac{x_n - \mu}{\sigma})]} = \frac{\exp \left\{ - [1 + \xi (\frac{x_n - \mu}{\sigma})]^{-1/\xi} \right\}}{\sigma + \xi(x_n - \mu)} \quad (4.5.2)$$

which is monotonically decreasing in  $\xi$  for all possible  $\xi$  values, thus (4.5.2) is maximised by setting  $\xi = -\frac{\sigma}{x_n - \mu}$ . The upper end point of the GEV distribution when  $\xi < 0$  is  $x^U = \mu - \frac{\sigma}{\xi}$  so setting  $\xi = -\frac{\sigma}{x_n - \mu}$  is the same as making the upper end point,  $x^U$ , equal to  $x_n$ . The parameter space is constrained by the upper end point being larger than  $x_n$ , so when  $c$  is close to  $x_n$  the MLE,  $\hat{\theta}$ , is close to this boundary. This can cause problems both with numerical maximisation of the likelihood (discussed below) and makes the theoretical basis of the confidence intervals more complex.

Recall (§4.3.1) to create confidence intervals for the  $y$ -year return level we use the profile likelihood,  $Pl(x_y)$ , and the asymptotic distribution of the deviance under certain regularity conditions. One of these regularity conditions is that the MLE is an interior point of the parameter space. Here, when  $c$  is close to the largest observation,  $x_n$ ,  $(\hat{x}_y, \hat{\sigma}, \hat{\xi})$  lies close to/on a boundary of the parameter space. There is much literature on estimation and asymptotic properties when the parameter is on the boundary of the parameter space, *e.g.*, Andrews (1999). We do not explore this theoretical issue further but concentrate on the practical numerical optimisation of the likelihood for such parameters.

The second issue raised is the numerical maximisation of the likelihood,  $\ell(x_y, \sigma, \xi)$ , to obtain the profile likelihood,  $Pl(x_y)$ . Consider fixing the  $y$ -year return level,  $x_y$ , and searching over a grid of possible  $\xi$  and  $\sigma$  values for the largest likelihood. This is a rather slow but effective way to approximate  $Pl(x_y)$  (the finer the grid chosen the better the approximation but also slower computationally) and the likelihoods at each grid point can be plotted to show the parameter space for a particular  $x_y$  value. Since the full and partial-conditioning likelihoods are almost identical for large  $c$  we present results for full conditioning only. In Figures 4.5.1/B.3.1 we fixed the 200-year return level at values between 1100 and 1800 with  $c = 1735/1739$  respectively (recall  $x_n = 1740$ ) and searched over a grid of  $\xi$  values from  $-0.3$  to  $0.3$  and  $\sigma$  up to 400. Figure B.3.1 uses the standard likelihood with  $c = 1739$  to illustrate the case where the MLE *is* in the interior of the parameter space. The colour scale indicates the magnitude of the likelihood at each of these grid points, blue being low likelihood and yellow high likelihood. Grey signifies combinations of  $\xi$  and  $\sigma$  which are (very close to) impossible for the chosen  $x_{200}$  and  $c$ .

From the shape of the contours on the parameter space plots we can see that  $\sigma$  and  $\xi$  are negatively ‘correlated’. This behaviour is expected since increasing the scale of the GEV distribution will have a similar effect on the model to increasing the shape parameter.

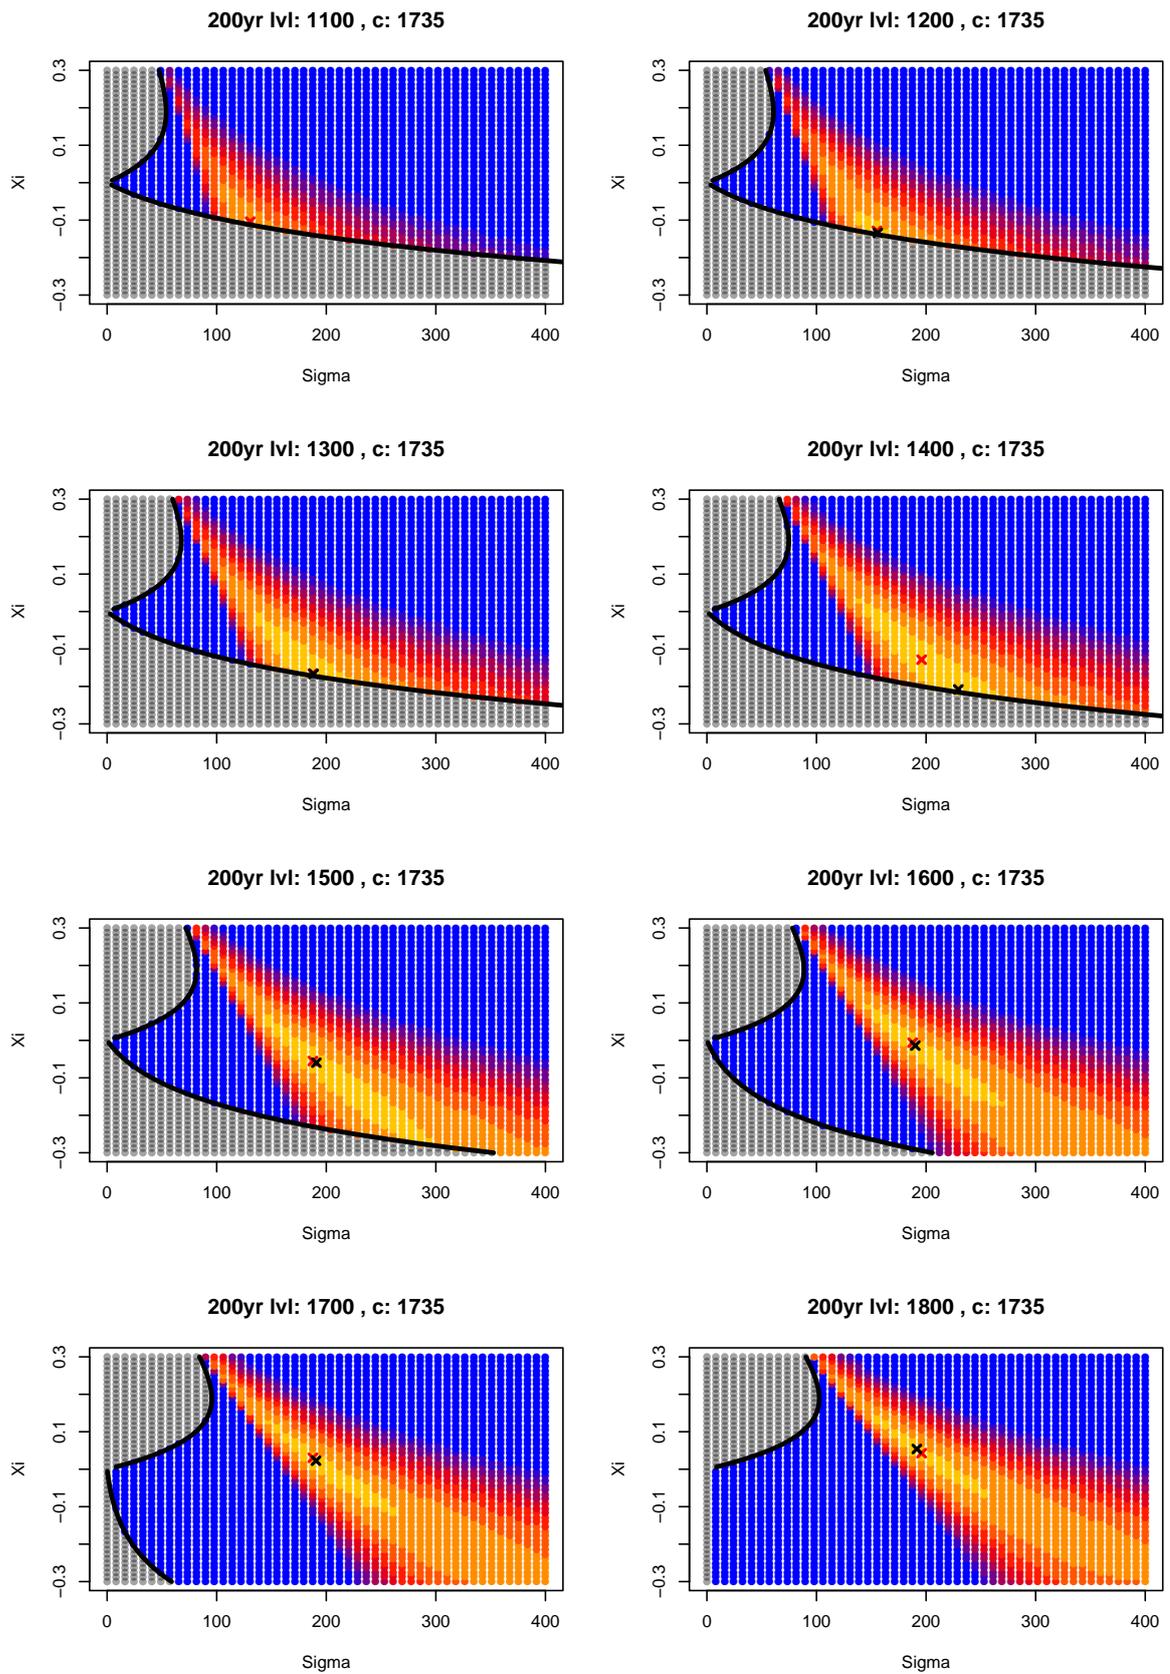


Figure 4.5.1: Full conditioning likelihood for each combination of  $\sigma$  and  $\xi$  given fixed  $x_{200}$  and  $c = 1735$  over a range of  $x_{200}$  values. Low/high likelihood regions are coloured in blue/yellow with grey being outside the parameter space. The black/red crosses are the MLEs using the optimisation/grid method respectively.

However the lighter the upper tail is, the wider the range of  $\sigma$  which will result in the same likelihood. We also see the larger we fix the value of the 200-yr return level the more elliptical and tighter the high likelihood contours become.

The boundaries of the parameter space are shown in black. The upper boundary curve is the consequence of the exponent in the likelihood being positive, this ensures that the smallest observation is larger than the lower end point when the shape parameter is positive. The lower boundary curve signifies when the largest observation is less than the upper end point when  $\xi$  is negative. Rearranging  $\bar{F}(x_y) = 1/y$  we have

$$\mu = x_y - \frac{\sigma}{\xi} \left[ (-\log(1 - 1/y))^{-\xi} - 1 \right].$$

So when  $x_y$  is fixed and  $\xi < 0$  the upper end point constraint is:

$$\begin{aligned} \mu - \frac{\sigma}{\xi} > x_n &\Rightarrow x_y - x_n - \frac{\sigma}{\xi} (-\log(1 - 1/y))^{-\xi} > 0 \\ &\Rightarrow \sigma > |\xi|(x_n - x_y)(-\log(1 - 1/y))^\xi, \end{aligned} \quad (4.5.3)$$

which is smaller the larger  $x_y$  is and negative when  $x_y > x_n$ , hence the lower bound on the plots reduces and disappears as  $x_y$  increases. Similarly, when  $\xi > 0$ , the lower end point constraint is:

$$\sigma > \xi(x_y - \min(x))(-\log(1 - 1/y))^\xi.$$

When  $x_{200}$  is fixed to lower values the upper end point boundary (4.5.3) moves towards the high likelihood region and in particular the MLE  $(\hat{\sigma}, \hat{\xi})$  given  $x_{200}$ . For the full and partial conditioning methods when  $c > 1734$  these MLEs suddenly become very close to the boundary when  $x_{200}$  is less than some critical value which increases with increasing  $c$ . For example, see Figure B.3.2, when  $c = 1735$  this jump to the boundary occurs when  $x_{200}$  is less than 1389 which is close to the second largest observation in the data set (1395.22). It would seem the jump here is due to the change in the restriction on the probability of exceedance of the second largest observation, that is the exceedance probability is less than  $1/200$  whereas for  $x_{200} > 1395.22$  it is greater than  $1/200$ . For very high  $c$  (in our investigation for  $c > 1738$ ) there is a small range of  $x_{200}$  for which the MLE jumped between to and away from the boundary (for  $c = 1738$  this was from 1415 to 1426) indicating instability in the estimates near the boundary.

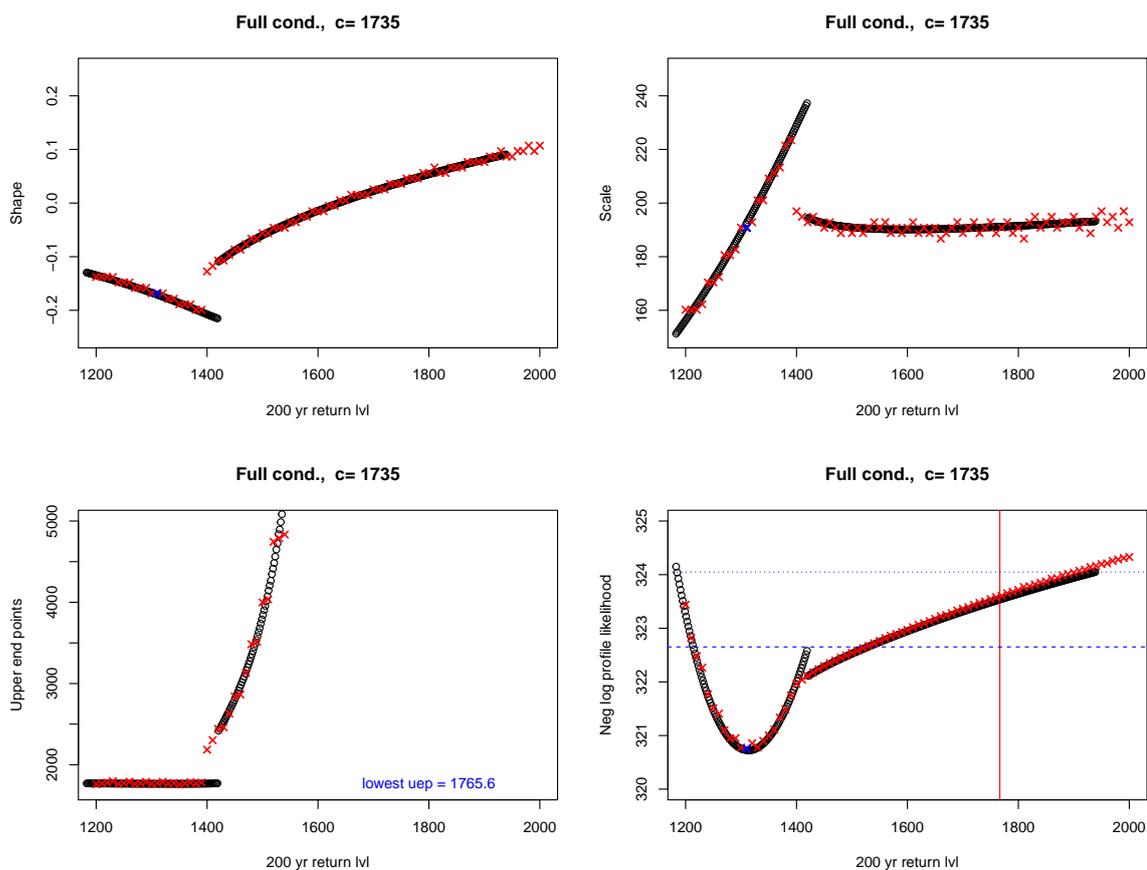


Figure 4.5.2: Top: MLE of shape(left) and scale(right) parameters given the  $x_{200}$  value on the  $x$ -axis. Bottom: Estimated upper end point (left) and profile likelihood (right) over different  $x_{200}$  values. The black/red points refer to the optimisation/grid method respectively with the full conditional likelihood with  $c = 1735$ . The blue crosses are at the MLE using the grid method and the red vertical line on the profile likelihood plot is the estimated upper end point at the MLE. The lowest upper end point (uep) estimate for the range of 200-yr return level considered is given in the corner of the bottom left panel. The blue dotted and dashed lines are the thresholds to obtain the 99% and 95% confidence intervals respectively.

Figures 4.5.2, B.3.3-B.3.5 show the maximum likelihood shape and scale parameter estimates (top), upper end point estimates (bottom left) and profile likelihood (bottom right) over a range of fixed 200 year return level values for  $c = 1730, 1735$  and  $1739$ . The black circles are the results of what we will refer to as the optimisation method. For this method we consider a fine grid of  $x_{200}$  values and for each fixed  $x_{200}$  obtain  $Pl(x_{200})$  by using the optim function in R to maximise the resulting likelihood with respect to  $\xi$  and  $\sigma$ . For the optim input we use initial parameters set to the MLE found for the previous  $x_{200}$  value. The red crosses are the results when searching over a grid of  $\xi$  and  $\sigma$  values as described for the parameter space plots. This latter method is much more computationally demanding as the likelihood is calculated at every point in the  $(\sigma, \xi)$  grid and for every  $x_{200}$  value. The

blue crosses are at the MLE using the grid method and the red vertical line on some of the profile-likelihood plots is the estimated upper end point at the MLE. The blue dotted and dashed lines are the thresholds to obtain the 99% and 95% confidence intervals respectively.

There is instability in both the optimisation and grid methods when  $c$  is extremely close to  $x_n$  (e.g.,  $c = 1739$ , Figure B.3.3). As expected for the conditioning methods the closer  $c$  is to  $x_n$  the closer the maximum likelihood upper end point estimate is to  $x_n = 1740$ , i.e., the closer  $(\hat{x}_{200}, \hat{\sigma}, \hat{\xi})$  are to the boundary of the parameter space (Figures 4.5.2, B.3.3-B.3.5 bottom left). The most striking feature of the optimisation method is the discontinuity when  $c$  is large (Figures B.3.3 and 4.5.2). On the parameter space plots this is where the MLE suddenly shifts towards the boundary. For lower  $c$  (e.g., Figure B.3.5) there is no discontinuity, however, a sudden change in the shape of the profile likelihood can still be seen as the estimated upper end point estimate moves away from  $x_n$ . For comparison the corresponding plots with the standard likelihood and  $c = 1735$  are shown in Figure B.3.6, the profile likelihood is much smoother with the standard likelihood and has no discontinuities.

In contrast to the optimisation method, the grid method does not result in a discontinuity in the profile likelihood (this is seen clearly for  $c = 1735$  in Figure B.3.7 top left). Nevertheless, for large  $c$ , there is still a discontinuity in MLEs given  $x_{200}$  as they jump away from the upper end point boundary, this discontinuity occurs at a lower  $x_{200}$  value than for the optimisation method (red crosses jump before the black circles). When  $x_{200}$  is fixed to some low value and is gradually increased the shape parameter becomes more negative and the scale parameter more positive to ensure the upper end point remains approximately equal to  $x_n$  (since, as discussed after (4.5.2), for  $x_n \approx c$  the likelihood is large when  $x^U \approx x_n$ ). However, if the upper bound of the fitted distribution is restricted to be approximately the largest observation,  $x_n$ , then  $\bar{F}(c) \approx \bar{F}(x_n) \approx 0$ . There will be some critical value of  $x_{200}$  above which it will not be possible to keep the probability of exceeding  $c$  so small and also have  $\bar{F}(x_{200}) = \frac{1}{200}$ . (This critical value will be greater the closer  $c$  is to  $x_n$  as the probability of exceeding  $c$  is then lower.) At this critical point the upper end point has to increase, thus the shape and scale parameter are no longer restricted to be close to the boundary - this is reflected in the sudden increase of the shape parameter and decrease in the scale parameter. This behaviour can be seen in the maximum likelihood parameter plots in the top row of Figures 4.5.2, B.3.3-B.3.5.

In Figure B.3.2 we show the parameter space as in Figure 4.5.1 but concentrating on 200-year return levels around the discontinuity and including the optimisation-based MLEs

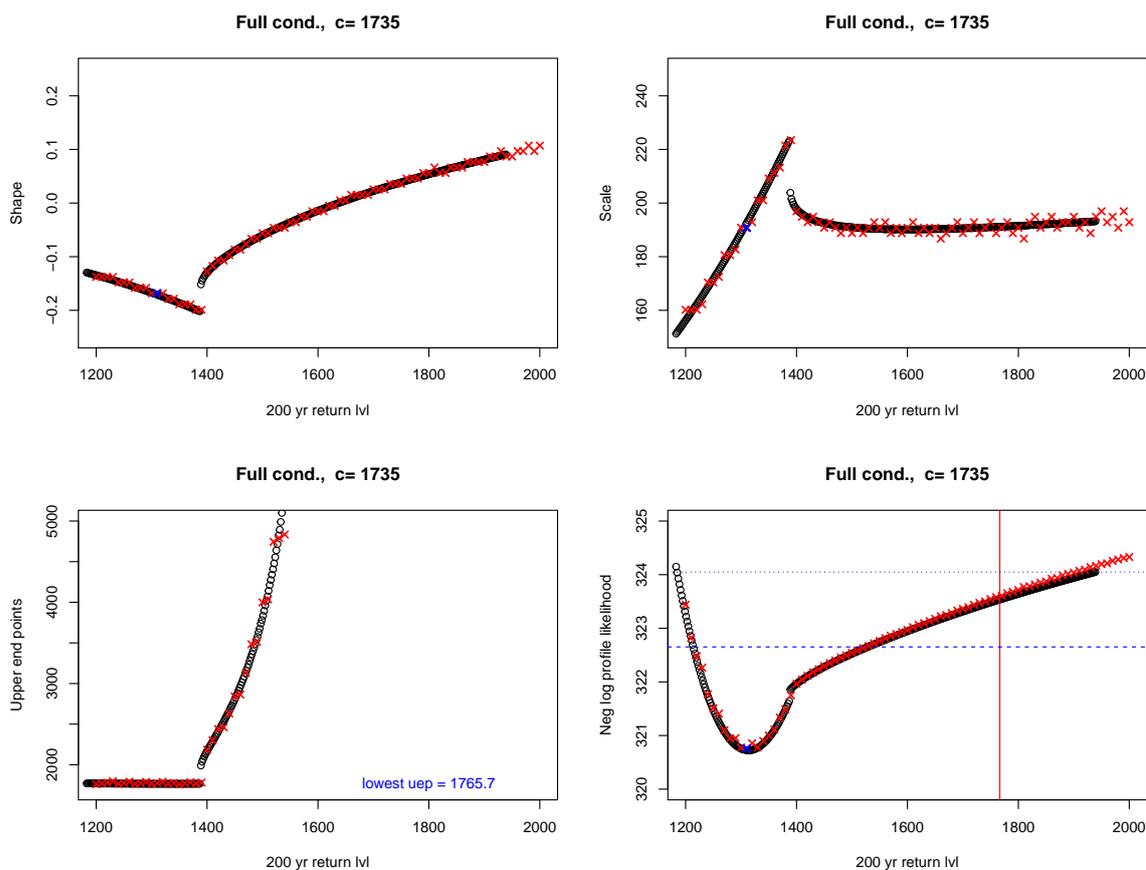


Figure 4.5.3: As Figure 4.5.2 but with with initial parameters shifted away from the boundary in the optimisation method.

(black crosses). It appears that when using the optimisation method the MLE ‘gets stuck’ on the boundary (since the initial parameter is on or close to the boundary) then moves inside to agree with the grid based MLE (red crosses) for larger  $x_{200}$  values - this sudden jump is causing the discontinuity. For example, in Figure B.3.7 the profile likelihood for  $1400 < x_{200} < 1420$  is estimated to be smaller than it truly is as the MLE is stuck on the boundary. This can cause the calculated deviance (4.3.1) to be much larger than it actually is, thus leading to too narrow confidence intervals.

One attempt to improve the optimisation near the critical  $x_{200}$  value is to shift the initial parameters  $(\xi, \sigma)$  (entered into `optim`) away from the boundary by increasing  $\xi$  and decreasing  $\sigma$  slightly. With such a shift the optimisation moves away from the boundary slightly too early (depending on the extent of the shift in initial parameters), so the discontinuity is shifted to earlier  $x_{200}$  values but is also reduced (see Figures 4.5.3, B.3.7, B.3.8).

Other attempts were made to improve upon the optimisation of the profile likelihood when  $c$  is close to  $x_n$ . We discuss one such improvement and its (in)effectiveness now. Consider

(4.5.1), when the estimated upper end point is close to  $x_n$ , then  $F(x_n) \approx 1$  and so

$$\begin{aligned}
 (4.5.2) &\approx \frac{1}{\sigma(1 + \xi(c - \mu)/\sigma)} \left[ \frac{\sigma + \xi(x_n - \mu)}{\sigma + \xi(c - \mu)} \right]^{-\left(1 + \frac{1}{\xi}\right)} \\
 &= \frac{1}{\sigma_c} \left[ \frac{\sigma_c - \xi(c - \mu) + \xi(x_n - \mu)}{\sigma_c - \xi(c - \mu) + \xi(c - \mu)} \right]^{-\left(1 + \frac{1}{\xi}\right)} \quad \text{where } \sigma_c = \sigma + \xi(c - \mu) \\
 &= \frac{1}{\sigma_c} \left[ 1 + \xi \left( \frac{x_n - c}{\sigma_c} \right) \right]^{-\left(\frac{1}{\xi} + 1\right)}.
 \end{aligned}$$

Therefore, when  $F(c)$  is close to 1 we can approximate the term  $\frac{f(x_n)}{F(c)}$  in the likelihood by  $g(x_n|X > c)$ , the GPD density of  $x_n$  given that it is greater than  $c$  with parameters  $(\xi, \sigma_c)$ . In practice the approximation makes little practical difference for the Lune data. It does improve the stability of the estimates for low 200 year return levels with extremely large  $c$  (Figure B.3.10) but doesn't change the resulting confidence intervals for the 200 year return level.

## 4.6 Summary

Overall, the conditioning estimators we presented in Chapter 3 outperform the standard estimator when the decision to analyse data at a particular time was triggered by what was perceived to be a large observation. However, we have seen that full-conditioning can lead to unrealistically large deviance-based confidence intervals, particularly when the sample size is small, and there are potential numerical issues which occur when the largest observation (which triggered the analysis) is close to the stopping threshold. In practice the second issue will be rare since an analysis is likely undertaken when there is an observation substantially greater than a large threshold rather than just surpassing it.

We also compared the deviance-based confidence intervals to the standard bootstrap confidence interval and bias-reducing variations thereof including our own version based on Efron's bias correction. The standard bootstrap method results in narrow intervals and poor coverage whereas the bootstrap variations generally increase coverage but with confidence interval widths comparable to the profile likelihood based interval widths; the latter is much faster computationally so remains the preferred confidence interval method. For the variable-threshold stopping rule we developed an importance-weighted bootstrap to create confidence intervals including an extension to multiple exceedances of the stopping threshold, however,

these intervals also suffer from negative bias similar to their fixed-threshold counterparts.

## Part II

# Efficient Loss Estimation

## Chapter 5

# Loss Estimation and Concentration Inequalities

JBA are interested in calculating the total loss incurred over time due to flooding events. Their clients are interested in the distribution of this total loss from flood events per year for a portfolio, in particular the mean, variance and  $t$ -year return levels,  $q_t$ . The 200-year return level (the loss expected to be exceeded with probability  $0.005 = 1/200$ ) is of special interest since it is specifically required by the government's 2015 solvency regulation (Swain and Swallow, 2015). JBA's end product is a set of return-level estimates for a range of return periods,  $t \in \{2, 5, 10, 20, 50, 75, 100, 150, 200, 250, 500, 1000, 1500\dots\}$  (referred to in the insurance industry as a loss curve) and their corresponding 95% confidence intervals. We are mostly interested in calculating high return periods; however, clients will compare the low quantiles of this curve with their historical data and so it is important that low quantiles are well estimated also.

In this chapter we discuss the estimation of the return levels of the loss distribution and explore and develop concentration inequalities for our approach to increase the computational efficiency of this estimation process (Chapter 6). In §5.1 we describe what we refer to as JBA's standard procedure to estimate quantiles of the loss distribution for a given portfolio from simulated events; this is the procedure used by JBA when this line of research was started. Since we are interested in estimating the tail probabilities of the total loss and thereby summing over events and risks (and subrisks), concentration inequalities for sums of independent random variables are useful; our approach to improving the computational efficiency (§6.1) uses concentration inequalities. In §5.2 we first review known concentration inequalities and then develop some novel, tighter, bounds in §5.4 and connect notation with the loss simulation setting in §5.5.

Term	Notation	Definition
Risk	$r$	Depending on the detail of a portfolio, a risk is either an insured property/location or a collection of insured properties in a certain postcode.
Value of risk $r$	$v_r$	Total insured value of risk.
Subrisk	$s$	Insured property/location.
Portfolio	$\mathfrak{R}$	A set of risks and their locations.
Coverage type	$c$	Insurance coverage types (B = building, C = contents, BI = business interruption).
Peril type	$\mathbf{p}$	Type of flooding (river, surface water or coastal).
Set of subrisks of risk $r$	$\mathcal{S}_r$	The collection of subrisks forming risk $r$ .
Event set	-	Set of simulated flood events and the year in which they occurred.
Events in year $y$	$\mathfrak{E}_y$	Set of events in the year $y$ simulation.
Hazard map	-	A fine grid of simulated water depths for an event and peril type covering all locations in the portfolio, for example, the whole of the UK.
Hazard distribution	$f_H(h; e, r, \mathbf{p})$	Distribution of water depths greater than 0 at risk $r$ for event $e$ and peril $\mathbf{p}$ .
Damage ratio	-	Loss as a fraction of total insured value.
Vulnerability function	-	Relates water depth to the mean and standard deviation of the damage ratio. Depends on particular aspects of a property such as building type.
Vulnerability distribution	$f_{V H}(x h; r, c)$	Distribution of damage ratio for risk $r$ and coverage type $c$ given water depth $h$ at the property. Assumed to be a Beta distribution.
Wet distribution	$f_X(x; e, r, \mathbf{p}, c)$	Distribution of damage ratio for risk $r$ with coverage $c$ given that there is flooding of type $\mathbf{p}$ during event $e$ .
Effective damage distribution	-	Distribution of damage ratio for risk $r$ with coverage $c$ due to flooding of type $\mathbf{p}$ during event $e$ .
Loss at risk $r$ in event $e$	$L_{e,r}$	
Total loss in year $y$	$S_y$	
Proportion of area affected	$p_{e,r,\mathbf{p}}$	Probability of flood of type $\mathbf{p}$ at risk $r$ during event $e$ .

Table 5.1.1: *Loss simulation definitions and notation.*

## 5.1 Notation and standard procedure

JBA's standard estimation procedure is detailed in this section; it involves many components combining flood and property data. Table 5.1.1 provides a list of definitions and notations for various terms used in the flood loss simulation setting. First, we give a brief overview of the procedure, with details in the following subsections.

Flood events are simulated for  $n_y = 10000$  years and the water level across the landscape is modelled (see §5.1.1) for each event for different types of flooding (river, coastal and surface water) known as *perils*. The distribution of water depths greater than 0 at a certain risk is

derived from this model and is referred to as the *hazard distribution*<sup>1</sup>. From the property information a *vulnerability distribution* is ‘found’ (see §5.1.4); this provides a distribution on the percentage of insured value loss for each risk, coverage and peril type (*e.g.*, contents loss at a particular property due to river flooding) given a certain water depth at the property. The vulnerability distribution and hazard distribution are then combined to form an *effective damage distribution* on the percentage loss of the insured value for each risk and event combination (§5.1.5). Finally, yearly losses are simulated by sampling from the effective damage distribution many times for each event and risk and accumulating the losses over a year, for each of the  $10^4$  years (§5.1.6). We aim to improve the efficiency of this final sampling and accumulation stage.

### 5.1.1 Simulating flood events

To determine the tails of the loss distribution we need to model the events which could contribute to the greatest losses. JBA incorporate the spatial structure of, *e.g.*, rainfall, by using the methodology in Heffernan and Tawn (2004) to model the extremes in an area conditional on the largest observation in the area. This neighbourhood is determined by the tail dependence, *i.e.* using  $\bar{\chi}$ , or equivalently  $\eta$  (§7.3.4). Keef et al. (2013) present a practical implementation of this method using all gauges rather than the localised version currently used.

This spatial extreme value model is then used to simulate extreme events over a 10000 year period on a network of rainfall, river and tidal gauges. Each event is given a unique ID and the set of event IDs and the year in which they occurred is called the *event set*. We denote the events in year  $y$  by  $\mathfrak{E}_y$ .

### 5.1.2 Portfolio information

The portfolio provides an ID and location for each risk in terms of postcode and sometimes also latitude and longitude. Some risks may in fact be a collection of multiple insured properties, referred to as *subrisks*, all assigned to the same location and ID. There is no information on individual subrisks only the whole collection, such as the total value of the subrisks, and the number of subrisks. We let each subrisk have a unique identifier,  $s \in \mathcal{S}_r$

---

<sup>1</sup>Not to be confused with the statistical definition of hazard in survival analysis. In the insurance industry a *hazard* is something that can cause loss, such as flooding.

where  $\mathcal{S}_r$  denotes the set of subrisks of risk  $r$ . For each subrisk,  $s$ , there is a mapping to the risk it is part of:  $r = r(s)$ . The number of subrisks of risk  $r$  is  $|\mathcal{S}_r|$ .

### 5.1.3 Water depths and distribution

The simulated water levels are inputted into JBA's software, JFLOW (Crossley et al., 2010), to simulate where the water will flow accounting for features of the landscape which cause blockages or water build up. From this they obtain a fine grid of water depths,  $h$ , for each event and peril type called the *hazard map*. This grid is so fine (each pixel is 5 square metres) that they essentially have the water depths resulting from the event everywhere in the UK. In the 2015 UK model these pixels are grouped into cells of 30 x 30m for river and coastal perils and 120 x 120m for surface water peril.

Using the location information in the portfolio each risk can be assigned to a model cell (or collection of cells if the postcode covers an area larger than one cell). Within this assigned area there will be multiple pixels, so multiple water depths given by the hazard map, which are used to obtain the hazard distribution. The hazard distribution is typically a parametric distribution *e.g.*, the Weibull distribution for which the parameters can be calculated from the 5th and 95th percentile of the (non-zero) water depths by solving an equation for these quantiles.

If a risk is flooded (non-zero water depth), the 'true' water depth is taken to be drawn from the hazard distribution corresponding to that risk. We denote the density of the hazard distribution by  $f_H(h; e, r, \mathbf{p})$  for event  $e$ , risk  $r$  and peril type  $\mathbf{p}$ .

The proportion of area affected,  $p_{e,r,\mathbf{p}}$ , is also deduced for each risk, event and peril type, as the percentage of pixels with non-zero water depth in the area corresponding to the risk.

For example, if a postcode is on a steep slope with one end near a river then  $p$  will be small (since flood water does not reach most of the hill) but the water depths near the river will be high, whereas for a postcode in a flat area near a river  $p$  will be close to 1 and the water depths will be similar across the postcode area.

The true distribution of water depth is 0 w.p.  $1 - p_{e,r,\mathbf{p}}$  and  $H$  w.p.  $p_{e,r,\mathbf{p}}$ , where  $H$  has density  $f_H(h; e, r, \mathbf{p})$ .

### 5.1.4 Vulnerability function and distribution

In order to translate the water depths into a loss some measure of how different water depths affect the resulting loss is needed. The *vulnerability function* gives a mean and standard deviation of the relative loss (fraction of the value of a risk that will be lost, termed the *damage ratio*) for each of a discrete number of depths. Each risk and coverage type (buildings, contents or business interruption) will have a different set of expectations and variances of the relative loss, this is due to particular aspects of the property such as the height of the doorstep or the type of building. In all cases, of course, the expected loss will be zero when the water depth is zero.

The *vulnerability distribution* of the damage caused relative to the value of the risk given a certain water depth,  $h$ , for risk  $r$  and coverage type  $c$  is assumed to follow a Beta distribution with parameters  $\alpha_{r,h,c}$  and  $\beta_{r,h,c}$  such that the expectation and standard deviation matching those given by the vulnerability function:  $\mathcal{V}_{r,h,c} \sim \text{Beta}(\alpha_{r,h,c}, \beta_{r,h,c})$ . We denote the density of the vulnerability distribution as  $f_{\mathcal{V}|H}(x|h; r, c)$ . The uncertainty here is due, in part, to variations in the state of the property, for example, if the owner is in the property at the time of flooding there may be less damage. There is no peril in the above formula since it is captured in the water depth.

### 5.1.5 Effective damage distribution

For each risk, coverage, peril and event combination, an *effective damage distribution* is built by combining hazard distributions with damage distributions.

If we knew the relative loss distribution for every possible water depth for a particular risk we could obtain the distribution for the damage ratio,  $x$ , by integrating the product of the vulnerability function and the depth distribution over the water depths:

$$f_X(x; e, r, \mathbf{p}, c) = \int_h f_{\mathcal{V}}(x|h; r, c) f_H(h; e, r, \mathbf{p}) dh$$

However, we only have hazard and vulnerability information for a discrete number of water depths. We partition the depth space into bins around the depths we have information for, so the marginal distribution becomes the sum of the product of the fitted beta distributions (vulnerability) and the probability of the water depth falling in a certain bin (from the hazard distributions):

$$f_X(x; e, r, \mathbf{p}, \mathbf{c}) = \sum_{i=1}^{n_B} \mathbb{P} \left( H \in [h_i^{lo}, h_i^{hi}); e, r, \mathbf{p} \right) f_V(x|h_i; r, \mathbf{c}),$$

where  $n_B$  is the number of bins,  $h_i^{lo}$  and  $h_i^{hi}$ ,  $i = 1, \dots, n_B$  with  $h_i^{hi} = h_{i+1}^{lo}$ ,  $i = 1, \dots, n_B - 1$  are the boundaries of the bin for the  $i$ th water depth and  $h_i = (h_i^{lo} + h_i^{hi})/2$ . Hence, the relative loss distribution for a given risk and event is a mixture of Beta distributions where the weights are determined by the distribution of water depths. The distribution,  $f_X(x; e, r, \mathbf{p}, \mathbf{c})$ , is called the *wet distribution* and is itself approximated by a Beta distribution with parameters determined by moment matching.

The final effective damage distribution is a mixture distribution such that the damage ratio is 0 with probability  $1 - p$  and follows the wet distribution,  $X$ , with probability  $p$ , where  $p$  is the proportion of area affected. Since there is no information on individual subrisks the hazard and vulnerability distribution for each subrisk is taken to be the same as that for the entire risk and the damage ratios for the subrisks are considered to be independent of each other. The effective damage distribution for every event,  $e$ , peril,  $\mathbf{p}$ , and coverage type,  $\mathbf{c}$  and subrisk,  $s$ , of risk  $r = r(s)$  is:

$$DR_{e,s,\mathbf{p},\mathbf{c}} = Z_{e,s,\mathbf{p}} X_{e,s,\mathbf{p},\mathbf{c}},$$

$$\text{where } Z_{e,s,\mathbf{p}} \sim \text{Bernoulli}(p_{e,r,\mathbf{p}})$$

$$\text{and } X_{e,s,\mathbf{p},\mathbf{c}} \sim \text{Beta}(\alpha_{e,r,\mathbf{p},\mathbf{c}}, \beta_{e,r,\mathbf{p},\mathbf{c}}).$$

The loss,  $L_{e,r,\mathbf{p},\mathbf{c}}$ , for a particular event, peril, risk and coverage type is the sum over the subrisks of the relative loss multiplied by the average subrisk value:

$$L_{e,r,\mathbf{p},\mathbf{c}} = \frac{v_{r,\mathbf{c}}}{|\mathcal{S}_r|} \sum_{s \in \mathcal{S}_r} DR_{e,s,\mathbf{p},\mathbf{c}}, \quad (5.1.1)$$

where  $\mathcal{S}_r$  is the set of subrisks of risk  $r$  and  $v_{r,\mathbf{c}}$  is the total insured value of the risk with coverage  $\mathbf{c}$ .

### 5.1.6 Sampling and aggregation

The final stage of the process is to simulate from the effective damage distributions many times for each event, (sub)risk, peril and coverage combination. Then we aggregate the losses over the portfolio in each of the 10000 years. However, portfolios can be extremely large (for example, Western Europe) and so can contain up to  $10^7$  risks. Clearly the large number of

risks and events involved in this process has a huge burden on computation. In 2016 JBA's software took 20 hours to analyse approximately 2 million risks. Hence we wish to find ways to improve the efficiency of this process by reducing the number of simulations necessary.

*Important: In what we do we are not trying to get to as close to the 'truth' as possible, rather produce results very close to those which would have been obtained using the standard procedure, whilst reducing computational cost.*

Three data sets are used within the simulation procedure: the event set, the portfolio data and the damage distribution data. The event set lists the unique IDs of each simulated event and the year  $(1, \dots, 10^4)$  in which it occurred. A portfolio consists of *risks*, their location and their insured values for each coverage type; each *risk* has a unique ID and consists of a number of subrisks. For each event and each of the 3 coverage types (building, contents and business interruption) and 3 peril types (river, coastal and surface water), the damage distribution table provides the mean and standard deviation of the wet distribution, the proportion of area affected and corresponding risk and event IDs.

For the rest of the chapter we consider the simplified setting where we have one coverage and peril type. We denote the set of events in year  $y$  by  $\mathfrak{E}_y$ , and the portfolio by  $\mathfrak{R}$ . The loss in year  $y$  is:

$$S_y = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} L_{e,r} = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \frac{v_r}{|\mathcal{S}_r|} \sum_{s \in \mathcal{S}_r} DR_{e,s} = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \frac{v_r}{|\mathcal{S}_r|} \sum_{s \in \mathcal{S}_r} Z_{e,s} X_{e,s},$$

where  $L_{e,r}$ ,  $DR_{e,s}$  and  $p_{e,r}$  are the loss, damage ratio and proportion of area affected for event  $e$  and risk  $r$ , subrisk  $s$  and  $v_r$  is the value of risk  $r$ . We now give the standard procedure, simulating the total loss in each of  $n_y$  years, to find return level estimates and confidence intervals of the yearly loss.

### Standard procedure

**Input:**  $m$  (the number of times the loss of each event and subrisk combination is simulated, typically  $m = 100$ )

**Step 1:** For each event and subrisk combination simulate  $DR_{e,s}$   $m$  times from the corresponding damage ratio distribution, and apply (5.1.1) to obtain realisations of the loss for

that event and risk,  $l_{e,r}^{(i)}$  for  $i = 1, \dots, m$ :

$$l_{e,r}^{(i)} = \frac{v_r}{|\mathcal{S}_r|} \sum_{s \in \mathcal{S}_r} Z_{e,s}^{(i)} X_{e,s}^{(i)}.$$

**Step 2:** For each simulation,  $i$ :

- a) For each year,  $y$ , sum over all the risks in the portfolio,  $\mathfrak{R}$ , and events in year  $y$ ,  $\mathfrak{E}_y$ , to produce a simulation of the loss for year  $y$ :

$$s_y^{(i)} = \sum_{r \in \mathfrak{R}} \sum_{e \in \mathfrak{E}_y} l_{e,r}^{(i)} \quad (5.1.2)$$

- b) To obtain  $q_t^{(i)}$ , the  $t$ -year return level estimate for simulation  $i$ , estimate the  $(1 - \frac{1}{t})$  quantile using the simulations  $\{s_y^{(i)}\}_{y=1, \dots, n_y}$ . We use the following estimate (see below for details):

$$q_t^{(i)} = (1 - \omega) s_{(n_y+1-\lceil k \rceil)}^{(i)} + \omega s_{(n_y+1-(\lceil k \rceil-1))}^{(i)},$$

where  $k = \frac{n_y+1}{t}$ ,  $\omega = \lceil k \rceil - k$ , and  $s_{(j)}^{(i)}$  is the  $j$ th order statistic ( $j$ th smallest loss simulated) of simulation  $i$ .

**Step 3:** We then have  $m$   $t$ -year return level estimates,  $(q_t^{(1)}, \dots, q_t^{(m)})$ , from which we calculate the mean, median and variance and take the 2.5th and 97.5th quantiles for an approximate 95% confidence interval.

### Quantile estimation in Step 2b of standard procedure

We estimate the CDF of the yearly loss by  $\tilde{F}$ , the empirical CDF slightly adapted so that  $\tilde{F}(\max_y(s_y)) < 1$ :

$$\tilde{F}(q) = \frac{1}{n_y + 1} \sum_{y=1}^{n_y} \mathbb{1}_{\{s_y^{(i)} \leq q\}}. \quad (5.1.3)$$

Setting (5.1.3) equal to  $1 - \frac{1}{t}$  and solving for  $q$  gives us an estimate of the  $(1 - \frac{1}{t})$  quantile which we denote by  $q_t$ . The equation to solve for  $q_t$  can be written as:

$$\sum_{y=1}^{n_y} \mathbb{1}_{\{s_y > q_t\}} = \frac{n_y + 1}{t} - 1 \quad (5.1.4)$$

When  $\frac{n_y+1}{t} \in \mathbb{N}$ , (5.1.4) is solved by any  $q_t \in \left[ s_{(n_y+1-\frac{n_y+1}{t})}, s_{(n_y+1-(\frac{n_y+1}{t}-1))} \right)$ , *i.e.*,  $q_t$

can be any value between the  $\left(\frac{n_y+1}{t} - 1\right)$ th and  $\left(\frac{n_y+1}{t}\right)$ th largest simulated loss, including the lower endpoint but not the upper endpoint. We take  $\hat{q}_t^{(i)} = s_{\left(n_y+1-\frac{n_y+1}{t}\right)}^{(i)}$  to be the return-level estimate for the  $i$ th simulation when  $\frac{n_y+1}{t} \in \mathbb{N}$ .

When  $\frac{n_y+1}{t} \notin \mathbb{N}$  (5.1.4) cannot be solved, however, we know that  $q_t$  must lie between the solutions to (5.1.4) with the right-hand side is rounded to the nearest integer up and down. So, using our estimate for  $\frac{n_y+1}{t} \in \mathbb{N}$ , we have

$$s_{\left(n_y+1-\lceil\frac{n_y+1}{t}\rceil\right)} < q_t < s_{\left(n_y+1-\left(\lceil\frac{n_y+1}{t}\rceil-1\right)\right)}.$$

We make the quantile estimate a continuous function of  $t$  by taking  $q_t$  to be a weighted sum of these bounds with weights,  $\omega$  given by the ‘distance’ from  $\frac{n_y+1}{t}$  to the next integer:

$$q_t^{(i)} = (1 - \omega)s_{\left(n_y+1-\lceil k \rceil\right)}^{(i)} + \omega s_{\left(n_y+1-\left(\lceil k \rceil-1\right)\right)}^{(i)},$$

where  $k = \frac{n_y+1}{t}$ ,  $\omega = \lceil k \rceil - k$ , and  $s_{(j)}^{(i)}$  is the  $j$ th order statistic ( $j$ th smallest loss simulated) of simulation  $i$ . This quantile estimate is equivalent to definition 6 in the paper of Hyndman and Fan (1996) which compares many quantile estimates; it is the default for some programming languages (Python, for example) and is highly recommended by Makkonen and Pajari (2014).

### 5.1.7 Test data

For this work JBA have provided a small portfolio, 1000-year event set, and damage distribution information for event and risk combinations in this portfolio and event set which have  $p_{e,r} > 0$ . Unless stated otherwise, all the examples presented in this chapter use these data or subsets thereof, particularly focusing on buildings insurance coverage and river flooding.

## 5.2 Concentration inequalities

Concentration inequalities provide bounds on the probability of a random variable deviating from a particular value, such as its expectation, by at least some margin, and so are especially helpful in finding bounds for tail probabilities. A wide range of concentration inequalities has been developed, requiring varying amounts of information about the random variable of interest. In this section we outline some basic inequalities and in §5.3 we discuss inequalities for the sum of independent random variables. More details on the concentration inequalities

presented here can be found in Boucheron et al. (2004). In §5.4 we present some new, tighter bounds building on the more complex inequalities of §5.3.

An advantage to finding bounds on probabilities using concentration inequalities is that these bounds are absolute – unlike the approximate bounds obtained from the central limit theorem. However, whether or not these bounds are useful will depend on the trade off between computational efficiency and the tightness of the bound. A loss estimation procedure using concentration inequalities is discussed in §6.1. In what follows we denote the random variable of interest by  $S$ .

### 5.2.1 Basic inequalities

The most basic concentration inequality, the *Markov inequality*, requires knowledge of the expectation of  $S$  only. For  $t > 0$  and  $S$  taking only non-negative values, the Markov inequality arises by taking expectations of both sides of the inequality  $t\mathbb{1}_{\{S \geq t\}} \leq S$ :

$$\mathbb{P}(S \geq t) \leq \frac{\mathbb{E}[S]}{t}. \quad (5.2.1)$$

From this we can derive the *Chebyshev inequality* by replacing  $S$  in (5.2.1) by  $(S - \mathbb{E}[S])^2$ ; for  $a > 0$ :

$$\mathbb{P}\left((S - \mathbb{E}[S])^2 \geq a\right) \leq \frac{\mathbb{E}\left[(S - \mathbb{E}[S])^2\right]}{a} = \frac{\text{Var}(S)}{a}.$$

Thus, letting  $t = \sqrt{a}$  we obtain the Chebyshev inequality:

$$\mathbb{P}(|S - \mathbb{E}[S]| \geq t) \leq \frac{\text{Var}(S)}{t^2}.$$

The *Cantelli inequality* is a generalisation of the Chebyshev inequality in that the bound is on the probability in a single tail rather than both tails, but the bound on that tail is slightly tighter than the bound obtained from Chebyshev's inequality and is particularly tighter in the body where  $t^2 < \text{Var}(S)$ :

$$\mathbb{P}(S - \mathbb{E}[S] \geq t) \leq \frac{\text{Var}(S)}{t^2 + \text{Var}(S)}. \quad (5.2.2)$$

To obtain this inequality, note the following is true for all  $x \in \mathbb{R}$ :

$$\begin{aligned} \mathbb{P}(S - \mathbb{E}[S] + x \geq t + x) &\leq \mathbb{P}(|S - \mathbb{E}[S] + x| \geq t + x) = \mathbb{P}((S - \mathbb{E}[S] + x)^2 \geq (t + x)^2) \\ &\leq \frac{\mathbb{E}[(S - \mathbb{E}[S] + x)^2]}{(t + x)^2} = \frac{\text{Var}(S) + x^2}{(t + x)^2}, \end{aligned}$$

where the final inequality arises from Markov's inequality. Since this is true for all  $x \geq 0$  we can choose  $x$  such that the right-hand-side bound is as small as possible. Differentiating we find this value of  $x$  to be  $\frac{\text{Var}(S)}{t}$ , which – substituted into the above – gives the desired result (5.2.2).

We can achieve stronger bounds by including more information about the random variable,  $S$ . For example, the Markov inequality can be generalised by replacing  $S$  in (5.2.1) by  $g(S)$  where  $g$  is a non-negative monotonically increasing function; the Chebyshev inequality is a special case of this. Another important special case is when the Markov inequality is applied with  $g(S) = e^{\lambda S}$  for some  $\lambda > 0$ :

$$\mathbb{P}(S \geq t) \leq \inf_{\lambda \geq 0} \frac{\mathbb{E}[e^{\lambda S}]}{e^{\lambda t}} = \inf_{\lambda \geq 0} \frac{M_S(\lambda)}{e^{\lambda t}}, \quad (5.2.3)$$

where  $M_S(\lambda)$  is the moment generating function of  $S$ . This bound is known as the *Chernoff inequality* and plays an important role in the proofs of many more complicated concentration inequalities.

Examples and more details of the inequalities outlined above can be found in Ross (1996) and a detailed overview of concentration inequalities is given by Boucheron et al. (2004). In the next section we will detail some of these inequalities, specifically for sums of independent, bounded random variables.

### 5.3 Concentration inequalities for sums of bounded random variables

In this section we consider inequalities that apply specifically to the tail probability of a sum of random variables. First, we state our assumptions and introduce useful results for the derivation of the concentration inequalities in later subsections.

### 5.3.1 Assumptions, definitions and notation

We are interested in concentration inequalities of the sum of independent, bounded random variables,  $X_i$  ( $i = 1, \dots, n$ ), *i.e.*,  $X_i$  satisfying the following assumption:

**Assumption 1.**  $X_i$  ( $i = 1, \dots, n$ ) are independent random variables with  $b_i \leq X_i \leq c_i$  where  $b_i$  and  $c_i$  are constants.

We present many of the concentration inequalities with the following further assumption on the lower bounds of the random variables:

**Assumption 2.** The random variables,  $X_i$ , satisfy Assumption 1 with  $b_i = 0 \forall i$ .

We denote the sum  $S_n = \sum_{i=1}^n X_i$ , define  $p_i = (\mathbb{E}[X_i] - b_i)/(c_i - b_i)$  and let  $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$ . We also denote  $Y_i = X_i - \mathbb{E}[X_i]$ ,  $m_i = b_i - \mathbb{E}[X_i]$  and  $a_i = c_i - \mathbb{E}[X_i]$ . Then under Assumption 1,  $Y_i$  are independent random variables with  $m_i \leq Y_i \leq a_i$  for each  $i$ . Under Assumption 2,  $p_i = \mathbb{E}[X_i]/c_i$  so  $\mathbb{E}[S_n] = \sum_{i=1}^n \mathbb{E}[X_i] = \sum_{i=1}^n c_i p_i = n\bar{c}\bar{p}$  where we have defined  $\bar{c}\bar{p} = \frac{1}{n} \sum_{i=1}^n c_i p_i$ . Furthermore,  $m_i = -c_i p_i$  and  $a_i = c_i(1 - p_i)$ .

Further we define:

$$\begin{aligned} c_{\min} &:= \min_{i=1, \dots, n} c_i & b_{\min} &:= \min_{i=1, \dots, n} b_i & a_{\min} &:= \min_{i=1, \dots, n} a_i & m_{\min} &:= \min_{i=1, \dots, n} m_i \\ c_{\max} &:= \max_{i=1, \dots, n} c_i & b_{\max} &:= \max_{i=1, \dots, n} b_i & a_{\max} &:= \max_{i=1, \dots, n} a_i & m_{\max} &:= \max_{i=1, \dots, n} m_i \end{aligned}$$

### 5.3.2 Preliminaries

All of the inequalities we will examine derive from Chernoff's inequality (5.2.3) applied to  $S_n$  (see 5.3.3). Many of the derivations also involve the following bound for a convex function,  $f(x)$ :

$$f(x) \leq \frac{c-x}{c-b} f(b) + \frac{x-b}{c-b} f(c), \quad b \leq x \leq c \quad (5.3.1)$$

In particular, setting  $f(x) = e^{\lambda x}$  and taking expectations, this leads to

$$\mathbb{E} \left[ e^{\lambda X} \right] \leq \frac{c - \mathbb{E}[X]}{c - b} e^{\lambda b} + \frac{\mathbb{E}[X] - b}{c - b} e^{\lambda c}. \quad (5.3.2)$$

In the following sections we often make use of the *arithmetic-geometric mean inequality* which states for non-negative real numbers  $\nu_1, \dots, \nu_n$ :

$$\left( \prod_{i=1}^n \nu_i \right)^{\frac{1}{n}} \leq \frac{1}{n} \sum_{i=1}^n \nu_i \quad (5.3.3)$$

It will also be useful to introduce here the *Kullback-Leibler divergence*, a measure of the ‘difference’ between two probability distributions:

**Definition 5.3.2.1.** *For two discrete distributions,  $Q$  and  $P$ , defined on the same countable probability set  $\mathcal{X}$ , the Kullback-Leibler divergence is*

$$D_{KL}(Q||P) = \sum_{x \in \mathcal{X}} Q(x) \log \left( \frac{Q(x)}{P(x)} \right).$$

For  $Q \sim \text{Bernoulli}(q)$  and  $P \sim \text{Bernoulli}(p)$  the Kullback-Leibler divergence is

$$D_{KL}(Q||P) = q \log \left( \frac{q}{p} \right) + (1-q) \log \left( \frac{1-q}{1-p} \right), \quad (5.3.4)$$

which we will refer to as  $D_{KL}(q||p)$  for simplicity.

**Proposition 5.3.2.2.** *An expression in the form below can be written in terms of K-L divergence:*

$$q \log \left( \frac{(1-p)q}{p(1-q)} \right) - \log \left( 1-p + \frac{(1-p)q}{1-q} \right) = D_{KL}(q||p).$$

**Proof** Let  $A$  denote the quantity on the left hand side of the above equation. Then

$$\begin{aligned} A &= q \log \left( \frac{q}{p} \right) + q \log \left( \frac{1-p}{1-q} \right) - \log \left( \frac{(1-p)(1-q) + (1-p)q}{1-q} \right) \\ &= q \log \left( \frac{q}{p} \right) - q \log \left( \frac{1-q}{1-p} \right) - \log \left( \frac{(1-p)}{1-q} \right), \end{aligned}$$

$$\begin{aligned} \text{and} \quad D_{KL}(q||p) &= q \log \left( \frac{q}{p} \right) + (1-q) \log \left( \frac{1-q}{1-p} \right) \\ &= q \log \left( \frac{q}{p} \right) - q \log \left( \frac{1-q}{1-p} \right) - \log \left( \frac{(1-p)}{1-q} \right). \end{aligned}$$

■

Finally, for proving some of the more complex concentration inequalities the Lambert W

function will be useful. The Lambert W function,  $W(x)$ , is the solution to  $W(x)e^{W(x)} = x$ . The Lambert W function is only defined on the real line for  $x \geq -e^{-1}$ , and for  $-e^{-1} < x < 0$  the function is double-valued (on the reals). So for  $x < 0$  the Lambert W function is split into two branches; the lower branch,  $W_{-1}(x) = W(x) : W(x) \leq -1$ , ranges from  $-1$  at  $x = -e^{-1}$  to  $-\infty$  at  $x = 0$  whereas the upper branch,  $W_0(x) = W(x) : W(x) > -1$ , ranges from  $-1$  at  $x = -e^{-1}$  to  $0$  at  $x = 0$ . It is important to be aware of which branch (if any) is the one of interest when using this function.

General expressions of the form,  $e^{\lambda a} = \lambda b + d$ , for some fixed constants  $a, b, d, \lambda$  can be rearranged into the form  $we^w = x$  and, hence, solved for  $\lambda$  in terms of  $a, b, d$  by use of the Lambert W function:

$$\begin{aligned} e^{\lambda a} = \lambda b + d &\Leftrightarrow e^{-\lambda a}(\lambda b + d) = 1 \\ &\Leftrightarrow e^{-\lambda a} \left( -\lambda a - \frac{ad}{b} \right) = -\frac{a}{b} \\ &\Leftrightarrow e^{-(\lambda a + \frac{ad}{b})} \left( -\lambda a - \frac{ad}{b} \right) = -\frac{a}{b} e^{-\frac{ad}{b}} \quad \Rightarrow \quad we^w = x, \end{aligned} \quad (5.3.5)$$

where  $x = -\frac{a}{b} e^{-\frac{ad}{b}}$  and  $w = -(\lambda a + \frac{ad}{b})$ . Since the Lambert W function is defined as the solution,  $w$ , to the equation (5.3.5) we can say  $W(-\frac{a}{b} e^{-\frac{ad}{b}}) = -(\lambda a + \frac{ad}{b})$ . Thus,

$$\lambda = -\frac{1}{a} \left[ W \left( -\frac{a}{b} e^{-\frac{ad}{b}} \right) + \frac{ad}{b} \right]. \quad (5.3.6)$$

In practice it is useful to be able to approximate or bound Lambert's W. The Taylor series of the upper branch is:

$$W_0(x) = \sum_{n=1}^{\infty} \frac{(-n)^{n-1}}{n!} x^n$$

For large values of  $x$ ,  $W_0 = \log(x) - \log(\log(x)) + o(1)$  asymptotically. Similarly, when  $x$  approaches 0 the lower branch,  $W_{-1} = \log(-x) - \log(-\log(-x)) + o(1)$ . Hoorfar and Hassani (2008) derive the following bounds on  $W_0$  for  $x \geq e$ :

$$\log(x) - \log(\log(x)) + \frac{\log(\log(x))}{2 \log(x)} \leq W_0(x) \leq \log(x) - \log(\log(x)) + \frac{e}{e-1} \frac{\log(\log(x))}{\log(x)}.$$

In the remainder of this section and §5.4 we present the proofs for the upper bounds on  $\mathbb{P}(S_n - \mathbb{E}[S_n] \geq nt)$  only. To obtain the upper bounds on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  we notice that  $\mathbb{P}(S_n - \mathbb{E}[S_n] \leq -nt) = \mathbb{P}(-S_n - \mathbb{E}[-S_n] \geq nt)$ . Thus, unless otherwise stated, the

same logic as used in the proofs for  $\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt)$  can be applied but with  $X_i$  replaced by  $-X_i$ .

### 5.3.3 Chernoff's inequality for the sum of independent variables

Applying the Chernoff inequality (5.2.3) to  $S_n$  but with  $t$  replaced by  $\mathbb{E}[S_n] + nt$ , we have the following bounds:

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \inf_{\lambda \geq 0} e^{-\lambda(\mathbb{E}[S_n] + nt)} \prod_{i=1}^n \mathbb{E} \left[ e^{\lambda X_i} \right] \quad (5.3.7)$$

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \inf_{\lambda \geq 0} e^{\lambda(\mathbb{E}[S_n] - nt)} \prod_{i=1}^n \mathbb{E} \left[ e^{-\lambda X_i} \right] \quad (5.3.8)$$

If the moment generating function,  $\mathbb{E}[e^{\lambda X}]$  is tractable the bound can be optimised numerically. However, for every iteration of the optimisation scheme there is an  $\mathcal{O}(n)$  operation so numerical optimisation can be computationally expensive. All of the concentration inequalities presented henceforth derive from (5.3.7) and are therefore looser than numerically optimising (5.3.7) but avoid many repeated  $\mathcal{O}(n)$  calculations.

### 5.3.4 Hoeffding's inequality for sums of bounded random variables

**Theorem 5.3.4.1** (Hoeffding (1963), Theorem 2). *Under assumption 1, for any  $t > 0$ ,*

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \exp \left( -\frac{2(nt)^2}{\sum_{i=1}^n (c_i - b_i)^2} \right) \\ \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &\leq \exp \left( -\frac{2(nt)^2}{\sum_{i=1}^n (c_i - b_i)^2} \right). \end{aligned} \quad (5.3.9)$$

The proof of Theorem 5.3.4.1 can be found in Appendix C.1.

The Hoeffding bound has two desirable properties: it has Gaussian tails, that is it behaves as  $e^{-nt^2}$ ; it also takes into account the different ranges of each random variable,  $X_i$ , in the sum. However it does not make use of any information on the variance of the individual  $X_i$ . Tighter bounds can be found by including this information, in particular we obtain *Bennett's inequality* (§5.3.7).

### 5.3.5 Chernoff-Hoeffding inequality

First we present another bound from Hoeffding (1963) which uses more information on the individual  $p_i$  (in the Hoeffding bound this information is only in  $\mathbb{E}[S_n] = \sum_{i=1}^n (b_i + (c_i -$

$b_i)p_i)$ ) and so provides tighter bounds than Hoeffding's inequality in some cases. Theorem 1 of Hoeffding (1963) presents a bound on the probability of deviations from the mean of independent random variables,  $X_i$ , with  $0 \leq X_i \leq 1$ . This is often referred to as the additive form of the Chernoff-Hoeffding Theorem and in literature it is simply referred to as a Chernoff bound referencing one or both Hoeffding and Chernoff (e.g., Impagliazzo and Kabanets (2010), Mulzer (2018)).

Zheng (2017) and From and Swift (2013) build on Hoeffding's Theorem 1 presenting bounds on the sum of independent random variables,  $X_i, i = 1, \dots, n$ , with  $\mathbb{P}(0 \leq X_i \leq 1) = 1$ ; the former by using a refined arithmetic-geometric mean bound, the latter splitting the expectations of the  $X_i$  into two groups.

Hoeffding (1963) remarks how his Theorem 1 extends to  $b \leq X_i \leq c$  via an affine transformation. Letting  $\mu = \frac{1}{n} \mathbb{E}[S_n]$ :

$$\begin{aligned} \mathbb{P}(S_n \geq nt + \mathbb{E}[S_n]) &\leq \exp\left(n \left[ \frac{\mu - b + t}{c - b} \log\left(\frac{\mu - b}{\mu + t - b}\right) + \frac{c - \mu - t}{c - b} \log\left(\frac{c - \mu}{c - \mu - t}\right) \right]\right) \\ &= \exp\left(-n D_{KL}\left(\frac{\mu - b + t}{c - b} \middle\| \frac{\mu - b}{c - b}\right)\right) \\ &= \exp\left(-n D_{KL}\left(\bar{p} + \frac{t}{c - b} \middle\| \bar{p}\right)\right), \end{aligned}$$

with  $p_i = (\mathbb{E}[X_i] - b_i)/(c_i - b_i)$  as defined in §5.3.1. Bounds of this form are often referred to as Chernoff-Hoeffding bounds.

Here we present (and later use) a slightly more general, novel, Chernoff-Hoeffding bound allowing the  $X_i$  to have different upper bounds.

**Theorem 5.3.5.1** (Chernoff-Hoeffding inequality). *Under Assumption 1 and 2 and with  $p_i, c_{\max}$  and  $c_{\min}$  defined as in §5.3.1 the Chernoff-Hoeffding inequalities are,*

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp\left(-n D_{KL}\left(\frac{\bar{c}\bar{p} + t}{c_{\max}} \middle\| \bar{p}\right)\right), \quad \bar{p}c_{\max} - \bar{c}\bar{p} < t < c_{\max} - \bar{c}\bar{p}, \quad (5.3.10)$$

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \exp\left(-n D_{KL}\left(\frac{\bar{c}\bar{p} - t}{c_{\min}} \middle\| \bar{p}\right)\right), \quad \bar{c}\bar{p} - c_{\min}\bar{p} < t < \bar{c}\bar{p}; \quad (5.3.11)$$

alternatively,

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \exp\left(-nD_{KL}\left(\frac{\bar{c}p' + t}{c_{\max}} \middle| \middle| \bar{p}'\right)\right), \quad \bar{p}'c_{\max} - \bar{c}p' < t < c_{\max} - \bar{c}p', \quad (5.3.12)$$

where  $p' = 1 - p$  and  $D_{KL}(q||p)$ , is the Kullback-Leibler divergence, as defined in (5.3.4), between Bernoulli distributed random variables with parameters  $p$  and  $q$  respectively.

The alternative bound (5.3.12) can also be written as:

$$\begin{aligned} \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &\leq \exp\left(-nD_{KL}\left(\frac{\bar{c} - \bar{c}p + t}{c_{\max}} \middle| \middle| 1 - p\right)\right) \\ &= \exp\left(-nD_{KL}\left(\frac{c_{\max} - \bar{c} + \bar{c}p - t}{c_{\max}} \middle| \middle| \bar{p}\right)\right), \end{aligned}$$

for  $\bar{c}p - \bar{p}c_{\max} + c_{\max} - \bar{c} < t < \bar{c}p + c_{\max} - \bar{c}$ . Notice that the two upper bounds, (5.3.11) and (5.3.12), on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  are the same when  $c_i = c \forall i$ . Below we give the proof for the bound (5.3.10) and outline the proofs for (5.3.11) and (5.3.12).

**Proof** For some fixed  $\lambda > 0$  we have by Chernoff's inequality (5.3.7):

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp(-\lambda(\mathbb{E}[S_n] + nt)) \prod_{i=1}^n \mathbb{E}[\exp(\lambda X_i)]. \quad (5.3.13)$$

Since  $e^{\lambda X_i}$  is a convex function we have (using (5.3.2))  $\mathbb{E}[e^{\lambda X_i}] \leq \frac{c_i - \mathbb{E}[X_i]}{c_i} + \frac{\mathbb{E}[X_i]}{c_i} e^{\lambda c_i} = 1 - p_i + p_i e^{\lambda c_i}$ .

Substituting this into (5.3.13) and using the inequality (5.3.3) relating the arithmetic and geometric means:

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \prod_{i=1}^n (1 - p_i + p_i e^{\lambda c_i}) e^{-\lambda(\sum_{i=1}^n c_i p_i + nt)} \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n (1 - p_i + p_i e^{\lambda c_i})\right)^n e^{-\lambda n(\bar{c}p + t)} \end{aligned} \quad (5.3.14)$$

$$\leq (1 + \bar{p}(e^{\lambda c_{\max}} - 1))^n e^{-\lambda n(\bar{c}p + t)}, \quad (5.3.15)$$

where  $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$ ,  $\bar{c}p = \frac{1}{n} \sum_{i=1}^n c_i p_i$  and  $c_{\max} = \max_i c_i$ , the maximum value of all the  $X_i$ .

The expression (5.3.15) is minimised by choosing

$$\lambda = \frac{1}{c_{\max}} \log \left( \frac{(1 - \bar{p})(\bar{c}\bar{p} + t)}{\bar{p}(c_{\max} - \bar{c}\bar{p} - t)} \right), \quad (5.3.16)$$

when  $\bar{p}c_{\max} - \bar{c}\bar{p} < t < c_{\max} - \bar{c}\bar{p}$ . The bound is 1 when  $t < \bar{p}c_{\max} - \bar{c}\bar{p}$  and 0 when  $t > c_{\max} - \bar{c}\bar{p}$ .

Finally substituting (5.3.16) into (5.3.15) we arrive at our Chernoff-Hoeffding inequality:

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \\ \exp \left\{ -n \left[ \left( \frac{\bar{c}\bar{p} + t}{c_{\max}} \right) \log \left( \frac{\bar{c}\bar{p} + t}{c_{\max}\bar{p}} \right) + \left( \frac{c_{\max} - \bar{c}\bar{p} - t}{c_{\max}} \right) \log \left( \frac{c_{\max} - \bar{c}\bar{p} - t}{c_{\max}(1 - \bar{p})} \right) \right] \right\}. \end{aligned}$$

Using (5.3.4) we can write this more concisely in terms of the Kullback-Leibler divergence,  $D_{KL}(q||p)$ , between Bernoulli distributed random variables with parameters  $p$  and  $q$  respectively so obtain bound (5.3.10).

**Bound 5.3.11** For some fixed  $\lambda > 0$  we have by Chernoff's inequality (5.3.8):

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \exp(-n\lambda(t - \bar{c}\bar{p})) \prod_{i=1}^n \mathbb{E}[\exp(-\lambda(X_i))].$$

Since  $e^{-\lambda X_i}$  is a convex function we have (using (5.3.2))  $\mathbb{E}[e^{-\lambda X_i}] \leq 1 - p_i + p_i e^{-\lambda c_i}$ .

Substituting this and using the inequality (5.3.3) relating the arithmetic and geometric means:

$$\begin{aligned} \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &\leq \left( \frac{1}{n} \sum_{i=1}^n (1 - p_i + p_i e^{-\lambda c_i}) \right)^n e^{-\lambda n(t - \bar{c}\bar{p})} \\ &\leq (1 + \bar{p}(e^{-\lambda c_{\min}} - 1))^n e^{-\lambda n(t - \bar{c}\bar{p})}, \end{aligned}$$

where  $\bar{p} = \frac{1}{n} \sum_{i=1}^n p_i$ ,  $\bar{c}\bar{p} = \frac{1}{n} \sum_{i=1}^n c_i p_i$  and  $c_{\min} = \min_i c_i$ , the minimum value of all the  $X_i$ .

The expression is minimised by choosing

$$\lambda = -\frac{1}{c_{\min}} \log \left( \frac{(1 - \bar{p})(-\bar{c}\bar{p} + t)}{\bar{p}(-c_{\min} + \bar{c}\bar{p} - t)} \right),$$

when  $\bar{c}\bar{p} - c_{\min} < \bar{c}\bar{p} - c_{\min}\bar{p} < t < \bar{c}\bar{p}$ . The bound is 1 when  $t < \bar{c}\bar{p} - c_{\min}\bar{p}$  and 0 when  $t > \bar{c}\bar{p}$ . Substituting this  $\lambda$  value we obtain (5.3.11). *To see this use Proposition 5.3.9.1 with  $q = (\bar{c}\bar{p} - t)/c_{\min}$  and  $p = \bar{p}$ .*

**Bound 5.3.12** The bound (5.3.12) is obtained by letting  $Z_i = c_i - X_i$  and considering

$$\begin{aligned} \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &= \mathbb{P}\left(\sum_{i=1}^n c_i - \sum_{i=1}^n Z_i \leq n\bar{c}p - nt\right) = \mathbb{P}\left(\sum_{i=1}^n Z_i \geq n(t + \bar{c}p')\right) \\ &\leq \frac{\prod_{i=1}^n \mathbb{E}[\exp(\lambda Z_i)]}{\exp(\lambda n(t + \bar{c}p'))}, \end{aligned}$$

by Chernoff's inequality. Then since

$$\mathbb{E}\left[e^{\lambda Z_i}\right] \leq \frac{c_i - \mathbb{E}[Z_i]}{c_i} + \frac{\mathbb{E}[Z_i]}{c_i} e^{\lambda c_i} = 1 - p'_i + p'_i e^{\lambda c_i},$$

we have

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \prod_{i=1}^n (1 - p'_i + p'_i e^{\lambda c_i}) e^{-n\lambda(\bar{c}p' + t)}.$$

The proof is then the same as that for (5.3.10) with  $p$  replaced by  $p' = 1 - p$  throughout. ■

It is not immediately clear which of (5.3.11) and (5.3.12) is the tighter bound. In fact it depends largely on the value of  $t$  and whether the bounds on  $t$  given in (5.3.11) and (5.3.12) overlap. When  $\bar{p}' > \bar{c}/c_{\max}$  the lower bound on  $t$  in (5.3.12),  $\bar{p}'c_{\max} - \bar{c}p' = \bar{p}'c_{\max} - \bar{c} + \bar{c}p'$ , is larger than the upper bound in (5.3.11),  $\bar{c}p'$ . So for  $t < \bar{c}p'$  bound (5.3.12) is 1 (the case  $t > \bar{c}p'$  is trivial since  $\mathbb{P}(S_n < E[S_n] - n\bar{c}p) = \mathbb{P}(S_n < 0) = 0$ ) whereas bound (5.3.11) is less than 1 for  $\bar{c}p' - \bar{p}c_{\min} < t < \bar{c}p'$ . However, when  $c_{\min}$  is small the interval is very narrow and so bound (5.3.11) is only useful for a small range of  $t$  values.

The scenario of interest has small values of  $p$ . When  $p$  is small the condition  $\bar{p}' > \bar{c}/c_{\max}$  is likely to be satisfied (unless  $\bar{c}$  is very close to  $c_{\max}$ ) and so bound (5.3.11) is tighter than (5.3.12). Thus henceforth we only consider (5.3.11). The proofs of Chernoff-Hoeffding+ and Chernoff-Hoeffding++ bounds in §5.4.1 use the same initial steps as bound (5.3.12).

### 5.3.6 Bernstein's inequality

An inequality for the sum of independent variables which makes use of the variance of each component of the sum is *Bernstein's inequality*, however this inequality assumes uniform bounds on the random variables.

**Theorem 5.3.6.1** (Bernstein (1946)). *Under assumption 1 with  $a_{\max}$  and  $m_{\min}$  defined as*

in §5.3.1 the Bernstein inequalities are, for  $t > 0$ ,

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp\left(-\frac{\frac{1}{2}(nt)^2}{\sum_{i=1}^n \text{Var}(X_i) + \frac{nt}{3}a_{\max}}\right) \quad (5.3.17)$$

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \exp\left(-\frac{\frac{1}{2}(nt)^2}{\sum_{i=1}^n \text{Var}(X_i) - \frac{nt}{3}m_{\min}}\right), \quad (5.3.18)$$

**Remark 5.3.6.2.** *These bounds have different tail properties depending on the deviation  $t$  in relation to the variance of the sum. For small deviations,  $t \ll \frac{3\text{Var}(S_n)}{na_{\max}}$ , the upper bound has Gaussian tails, i.e., behaves like  $e^{-nKt^2}$  for some constant  $K$ . Whereas for large deviations the bound has exponential tails  $e^{-nKt}$ .*

**Proof** Let  $f(x) := \frac{e^x - x - 1}{x^2} = \sum_{k=2}^{\infty} \frac{x^{k-2}}{k!}$ . Using the fact that  $k! \geq 2(3^{k-2}) \forall k \geq 1$ , we have for  $x < 3$ :

$$f(x) \leq \sum_{k=2}^{\infty} \frac{x^{k-2}}{2(3^{k-2})} = \frac{1}{2} \sum_{k=0}^{\infty} \left(\frac{x}{3}\right)^k = \frac{1}{2(1 - \frac{x}{3})}. \quad (5.3.19)$$

Defining  $\sigma_i^2 = \mathbb{E}[Y_i^2] = \text{Var}(X_i)$  we have from the definition of  $f$ :

$$\mathbb{E}\left[e^{\lambda Y_i}\right] = 1 + \lambda^2 \mathbb{E}\left[Y_i^2 f(\lambda Y_i)\right] \leq 1 + \sigma_i^2 \lambda^2 f(\lambda a_i) \leq \exp(\sigma_i^2 \lambda^2 f(\lambda a_i)), \quad (5.3.20)$$

where the first inequality follows because  $f(x)$  is monotone increasing and the second follows as  $1 + x \leq e^x$ . So, for some fixed  $0 < \lambda < 3/a_{\max}$ ,

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp(-\lambda nt) \prod_{i=1}^n \mathbb{E}\left[\exp(\lambda Y_i)\right] \quad (\text{by Chernoff's inequality (5.3.7)})$$

$$\leq \exp(-\lambda nt) \exp\left(\lambda^2 \sum_{i=1}^n \sigma_i^2 f(\lambda a_i)\right) \quad (\text{substituting (5.3.20)}) \quad (5.3.21)$$

$$\leq \exp(-\lambda nt) \exp\left(\lambda^2 f(\lambda a_{\max}) \sum_{i=1}^n \sigma_i^2\right). \quad (5.3.22)$$

Then, using (5.3.19),

$$(5.3.22) \leq \exp\left(-\lambda nt + \frac{\lambda^2}{2} \frac{\text{Var}(S_n)}{1 - \frac{a_{\max}\lambda}{3}}\right) \quad (\text{using (5.3.19)}).$$

The right-hand side is minimised by

$$\lambda = \frac{nt}{\text{Var}(S_n) + a_{\max}nt/3},$$

and leads to (5.3.17). The proof for (5.3.18) follows by replacing  $Y_i$  by  $-Y_i$ ,  $a_i$  by  $-m_i$  and  $a_{\max}$  by  $-m_{\min}$  in the above. ■

### 5.3.7 Bennett's inequality

**Theorem 5.3.7.1** (Bennett (1962)). *Under assumption 1 with  $a_{\max}$  and  $m_{\min}$  defined as in §5.3.1 the Bennett inequalities are, for  $t > 0$ ,*

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \exp\left(-\frac{\text{Var}(S_n)}{a_{\max}^2} h\left(\frac{na_{\max}t}{\text{Var}(S_n)}\right)\right) \\ \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &\leq \exp\left(-\frac{\text{Var}(S_n)}{m_{\min}^2} h\left(\frac{-nm_{\min}t}{\text{Var}(S_n)}\right)\right), \end{aligned}$$

where  $h(x) = (1+x)\log(1+x) - x$ .

**Proof** Following the first part of the proof of the Bernstein inequality we have

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp(-\lambda nt + \lambda^2 f(\lambda a_{\max}) \text{Var}(S_n)) \quad (\text{from (5.3.22)}). \quad (5.3.23)$$

This is minimised at  $\lambda^* = \frac{1}{a_{\max}} \log\left(1 + \frac{nta_{\max}}{\text{Var}(S_n)}\right)$ , giving the Bennett bound above. ■

Notice that Bennett's inequality is strictly tighter than Bernstein's inequality since it does not use the loosening step bounding  $f(x)$ .

**Remark 5.3.7.2.** *Similar to Bernstein we have different tail behaviour for different  $t$  values. For  $t \ll \frac{\text{Var}(S_n)}{na_{\max}}$ ,  $h(na_{\max}t) \approx \left(\frac{na_{\max}t}{\text{Var}(S_n)}\right)^2$  so the Bennett bound is approximately  $\exp\left(-\frac{n^2t^2}{\text{Var}(S_n)}\right)$ , i.e., it has Gaussian tails. On the other hand, when  $t \geq \frac{\text{Var}(S_n)}{na_{\max}}$ ,  $h(na_{\max}t) \geq \frac{na_{\max}t}{\text{Var}(S_n)} \log\left(\frac{na_{\max}t}{\text{Var}(S_n)}\right)$  so the Bennett bound is less than or equal to  $\exp\left(-\frac{nt}{\text{Var}(S_n)} \log\left(\frac{na_{\max}t}{\text{Var}(S_n)}\right)\right) = \left(\frac{\text{Var}(S_n)}{na_{\max}t}\right)^{\frac{nt}{a_{\max}}}$ , which is a Poisson-like tail.*

### 5.3.8 Jebara's Bennett refinement

Jebara (2018) presents a more complex refinement of Bennett's concentration bound which which can be, but is not always, slightly tighter than the Bennett bound. First omitting the

loosening step (5.3.20) in the proof of Bennett's inequality, (5.3.21) becomes:

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq e^{-\lambda nt} \prod_{i=1}^n (1 + \lambda^2 \sigma_i^2 f(\lambda a_i)) \quad (5.3.24)$$

$$\leq \exp\left(\sum_{i=1}^n \left[\log(1 + \lambda^2 \sigma_i^2 f(\lambda a_i)) - \lambda t \frac{a_i}{\bar{a}}\right]\right). \quad (5.3.25)$$

Jebara (2018) then finds the  $\lambda_i$  which minimises each term,  $b_i(\lambda)$ , in the summation and uses this  $\lambda_i$  as part of a quadratic bound on each term:

$$\log(1 + \lambda^2 \sigma_i^2 f(\lambda a_i)) - \lambda t \frac{a_i}{\bar{a}} \leq \frac{a_i^2}{1 - e^{-a_i^2/\sigma_i^2}} \frac{(\lambda - \lambda_i)^2}{2} + \log(1 + \lambda_i^2 \sigma_i^2 f(\lambda_i a_i)) - \lambda_i t \frac{a_i}{\bar{a}}. \quad (5.3.26)$$

The bound (5.3.26) is then inserted into (5.3.25) and the optimal  $\lambda$  is found, which we denote by  $\lambda_Q^*$ . The resulting bound for the sum,  $S_n$ , of  $n$  independent variables,  $X_i$ , with  $X_i - \mathbb{E}[X_i] \leq a_i$  and  $t \in (0, \bar{a})$  is bound (5.3.24) with  $\lambda$  replaced by  $\lambda_Q^*$  where

$$\lambda_Q^* = \left(\sum_{i=1}^n \frac{a_i^2}{1 - e^{-a_i^2/\sigma_i^2}}\right)^{-1} \sum_{i=1}^n \frac{a_i^2 \lambda_i}{1 - e^{-a_i^2/\sigma_i^2}}, \quad (5.3.27)$$

$$\lambda_i = \frac{\bar{a}}{t a_i} + \frac{a_i}{\sigma_i^2} - \frac{1}{a_i} - \frac{1}{a_i} W\left(\exp\left(\frac{\bar{a}}{t} + \frac{a_i^2}{\sigma_i^2} - 1 + \log\left(\frac{\bar{a} - t}{t}\right)\right)\right).$$

So to calculate this bound one needs to work out the Lambert W function for each  $X_i$ , that is  $n$  times. The individual variances,  $\sigma_i^2$ , and  $a_i$  are also needed unlike the other bounds which only require some summary statistics such as the mean of the variances.

### 5.3.9 Comparing existing concentration inequalities on simulated data

The upper tail bounds presented in the previous sections are now compared for the sum of  $n$  independent binary variables,  $S_n = \sum_{i=1}^n X_i$  with  $X_i \sim c_i \text{Bernoulli}(p_i)$ . We consider two cases for each of  $p_i$  and  $c_i$ :  $p_i = 0.013 \forall i$  and  $P_i \sim \text{Beta}(0.3, 22)$ , and  $c_i = 1 \forall i$  and  $C_i \sim 0.1 + \text{Exp}(1)$ . The parameters are chosen in this way to mimic the loss simulation setting where the expected loss is small since probabilities are very small but the variance is high with large potential loss. Figure 5.3.1 shows the various concentration bounds on the probability of the sum of 1000 independent binary variables exceeding a certain number,  $n_{sd}$ , of standard deviations,  $\sigma = \sqrt{\sum p_i(1-p_i)c_i^2}$ , above the mean,  $\mathbb{P}(S_n \geq \mathbb{E}[S_n] + n_{sd}\sigma)$ . The concentration inequalities presented are bounds on  $\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt)$  so throughout this

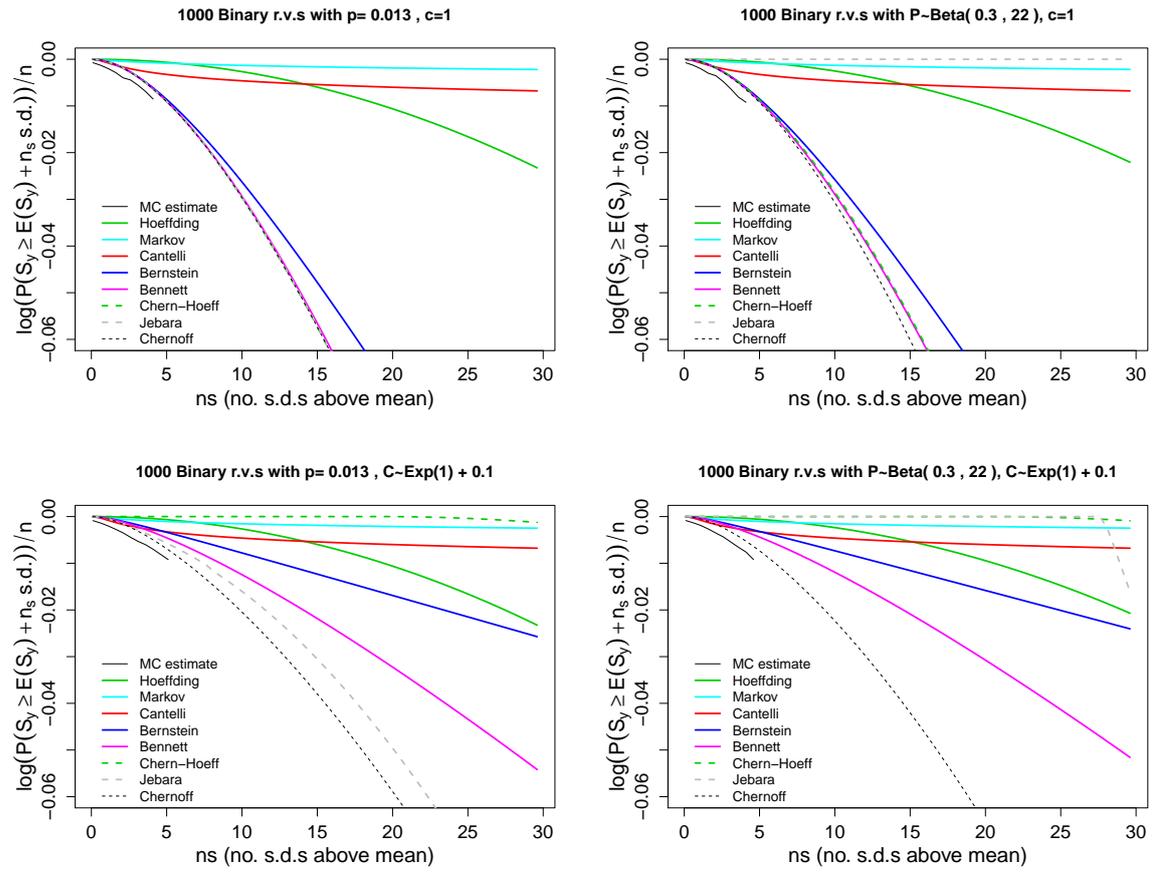


Figure 5.3.1: Bound on the exceedance probability,  $\mathbb{P}\left(S_n \geq \mathbb{E}[S_n] + n_{sd}\sqrt{\text{Var}(S_n)}\right)$ , against  $n_{sd}$  for the sum of 1000 binary random variables,  $X_i \sim c_i \text{Bernoulli}(p_i)$ . Top left:  $p = 0.013$ ,  $c_i = 1 \forall i$ , Top right:  $p_i \sim \text{Beta}(0.3, 22)$ ,  $c_i = 1$ , Bottom left:  $p = 0.013$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ , Bottom right:  $p_i \sim \text{Beta}(0.3, 22)$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ .

Bounds used are: Markov (5.2.1) (cyan), Cantelli (5.2.2) (Red), Chernoff (5.3.7) numerically optimised (dashed black), Hoeffding (5.3.4) (solid green), Chernoff-Hoeffding (5.3.10) (green dashed), Bernstein (5.3.17) (blue), Bennett (5.4.12) (pink), Jebara (C.3.1) (grey dashed). The solid black line indicates the Monte Carlo estimate based on 10000 simulations of  $S_n$ .

subsection we let  $t = n_{sd}\sigma/n$ . The Chernoff bound (5.3.7) is shown by a dashed black line, all of our bounds presented (aside from Markov and Cantelli) derive from this bound and so this provides a best possible bound for us to compare against. When both  $c$  and  $p$  are fixed the Chernoff bound can be found analytically (since the MGF is tractable), otherwise the bound is found by numerical optimisation. The Monte Carlo estimate based on 10000 simulations of  $S_n$  is also shown for comparison.

Of the classical concentration inequalities reviewed in §5.2 and §5.3 Bennett's inequality (solid pink) performs best for exceedances more than 5 standard deviations from the mean. The Cantelli inequality (solid red) is usually the tightest bound for small  $t$ , *i.e.*,  $n_{sd}$  small. When  $p$  is large Bernstein's inequality is very similar to Bennett's bound however for small  $p$

it is considerably looser and performs worse when the maximum value of each  $X_i$  is different (Figure 5.3.1 bottom, blue). When  $c$  is not fixed the Chernoff-Hoeffding bound is very loose. However, when both  $c$  and  $p$  are fixed the Chernoff-Hoeffding bound,  $\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp(-nD_{KL}(t+p||p))$ , is equal to the Chernoff bound, *i.e.*, it is the best possible bound of this form. Hoeffding's inequality (solid green) performs particularly poorly in all four cases since it does not incorporate any information about the small  $p_i$  values, *i.e.*, there is no information on the variance of the random variables, only their upper and lower bounds. On the other hand, it is one of the tightest bounds when  $\bar{p} \approx 0.5$ , particularly when  $c$  is fixed, since this leads to the largest possible variance and so the other bounds (which use the variance information) do not perform as well.

For small  $p$  and when all  $c_i = c$  the Bennett bound is almost equal to but slightly larger than the Chernoff-Hoeffding bound. The following proposition formulates this observation.

**Proposition 5.3.9.1.** *Let  $p_{CH}(t)$  and  $p_B(t)$  be the Chernoff-Hoeffding and Bennett upper tail bound (the right-hand sides of (5.3.10) and (5.4.12)) respectively. Define  $\rho := (1 - p_{\min})/(1 - \bar{p}^2/\bar{p})$ . If  $c_i = c \forall i$  then*

$$-\frac{1}{n} \log p_{CH}(t) = \bar{p}h\left(\frac{t}{c\bar{p}}\right) + \frac{(t/c)^2}{2(1-\bar{p})} + \mathcal{O}\left(\left[\frac{t}{c}\right]^3\right).$$

for  $t \leq \epsilon \bar{a}$  for some  $0 < \epsilon < 1$ . Furthermore, if  $t/[c\bar{p}] > b$  for any fixed  $b > 1$  then

$$-\frac{1}{n} \log p_B(t) = \bar{p}h\left(\frac{t}{c\bar{p}}\right) \times \frac{1}{1-p_{\min}} \rho^{\{t/[c\bar{p}] - \log(1+t/[c\bar{p}])\}/h(t/[c\bar{p}])} \times \{1 + \mathcal{O}([\log \rho]^2)\}. \quad (5.3.28)$$

Proposition 5.3.9.1 tells us if  $c_i = c \forall i$  and  $c\bar{p} < t < \bar{a} = c(1-\bar{p})$  then our Chernoff-Hoeffding bound and the Bennett bound are close since each log bound is  $-n\bar{p}h(t/[c\bar{p}])$  modulo a small correction. For the correction term in  $p_B$ , first notice that  $p_{\min}\bar{p} \leq \bar{p}^2 \leq p_{\max}\bar{p}$ , so  $1 \leq \rho \leq \frac{1-p_{\min}}{1-p_{\max}}$  and  $\rho \approx 1$  when  $p$  are small. Also the power of  $\rho$  in (5.3.28) with fixed  $t$  and  $c$  is a monotonically increasing function in  $\bar{p}$  which is bounded above by 1 and tends to 0 as  $\bar{p} \rightarrow 0$ . So when  $\bar{p}$  is small the correction term in  $p_B$  is very small and, unless  $t$  is very small, it is larger than the correction term in  $p_{CH}$ . The proof of Proposition 5.3.9.1 along with examples and more details are given in Appendix C.2.

For *Jebara's* bound (grey dashed) one must take care evaluating the Lambert W function; often approximations for  $W(x)$  at large or small  $x$  are needed. When  $p$  is fixed Jebara's bound

is tighter than the Bennett bound whereas for  $P \sim \text{Beta}(0.3, 22)$  the bound is very loose; for small  $t$  it is 1 and only starts to improve upon other bounds when  $t$  is large. This poor behaviour seems to be driven by very small values of  $p$ . We also found that Jebara's bound is much looser than Bennett when  $p$  is fixed but very small (*e.g.*,  $p = 0.001$ ) so it appears the larger  $\bar{p}$  is the tighter the bound becomes. This suggests that the minimiser (5.3.27) which Jebara uses is not optimal when  $c$  and/or  $p$  are not fixed and  $p$  is small unless  $t$  is very large. We now explore this issue in more detail in Appendix C.3.

## 5.4 Tighter concentration inequalities

In this section we present some novel, tighter concentration bounds based on the Chernoff-Hoeffding (§5.3.5) and Bennett (§5.3.7) bounds in §5.4.1 and §5.4.2 respectively.

### 5.4.1 Convexity tricks with the Chernoff-Hoeffding bound

We first present two refinements to the Chernoff-Hoeffding bound found by bounding convex functions. These bounds are compared to other concentration inequalities (in particular the Chernoff-Hoeffding bound) in §5.4.3. When  $c_i$  is the same for all  $i$  these bounds both reduce to the standard Chernoff-Hoeffding bound (5.3.10).

#### Chernoff-Hoeffding+

In the derivation of the Chernoff-Hoeffding bound we bound  $c_i$  by  $c_{\max}$  after applying the arithmetic-geometric mean inequality (going from (5.3.14) to (5.3.15)). In the following we omit this loosening step and instead use a tighter bound, using the convexity of  $e^{\lambda c}$ , before applying the arithmetic-geometric mean inequality.

**Theorem 5.4.1.1** (Chernoff-Hoeffding+). *Under Assumption 1 and 2, and with  $p_i$  and  $c_{\max}$  defined as in §5.3.1 the Chernoff-Hoeffding+ inequalities are*

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp\left(-nD_{KL}\left(\frac{\bar{c}p + t}{c_{\max}} \middle\| \frac{\bar{c}p}{c_{\max}}\right)\right) \quad 0 < t < c_{\max} - \bar{c}p \quad (5.4.1)$$

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \exp\left(-nD_{KL}\left(\frac{\bar{c}p - t}{c_{\max}} \middle\| \frac{\bar{c}p}{c_{\max}}\right)\right) \quad 0 < t < \bar{c}p. \quad (5.4.2)$$

**Proof** Notice that, since  $e^{\lambda c}$  is convex in  $c$ , if  $0 \leq c \leq c_{\max}$  then  $e^{\lambda c} \leq \left(1 - \frac{c}{c_{\max}}\right)e^0 + \frac{c}{c_{\max}}e^{\lambda c_{\max}}$ , so

$$1 - p_i + p_i e^{\lambda c_i} \leq 1 - p_i + p_i \left( 1 - \frac{c_i}{c_{\max}} \right) + \frac{p_i c_i}{c_{\max}} e^{\lambda c_{\max}}. \quad (5.4.3)$$

So we can replace the sum in (5.3.14) by the sum of (5.4.3) over  $i$ :

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq e^{-\lambda n(\overline{c\bar{p}}+t)} \left[ \frac{1}{n} \sum_{i=1}^n \left( 1 - \frac{c_i p_i}{c_{\max}} + \frac{c_i p_i}{c_{\max}} e^{\lambda c_{\max}} \right) \right]^n \\ &= e^{-\lambda n(\overline{c\bar{p}}+t)} \left( 1 + \frac{\overline{c\bar{p}}}{c_{\max}} (e^{\lambda c_{\max}} - 1) \right)^n. \end{aligned}$$

Since  $\overline{c\bar{p}}/c_{\max} \leq \bar{p}$  this bound is no larger than (5.3.15) and is strictly smaller if the  $c_i$ s differ.

The right hand side is minimised by choosing

$$\lambda = \frac{1}{c_{\max}} \log \left( \frac{(c_{\max} - \overline{c\bar{p}})(\overline{c\bar{p}} + t)}{\overline{c\bar{p}}(c_{\max} - (\overline{c\bar{p}} + t))} \right), \quad (5.4.4)$$

where  $0 < t < c_{\max} - \overline{c\bar{p}}$ . When  $t > c_{\max} - \overline{c\bar{p}}$ ,  $\mathbb{P}(S_n - \mathbb{E}[S_n] \geq nt) = 0$ . Substituting into (5.4.4) we obtain the following bound:

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \\ &\exp \left\{ -n \left[ \left( \frac{\overline{c\bar{p}} + t}{c_{\max}} \right) \log \left( \frac{\overline{c\bar{p}} + t}{\overline{c\bar{p}}} \right) + \left( \frac{c_{\max} - \overline{c\bar{p}} - t}{c_{\max}} \right) \log \left( \frac{c_{\max} - \overline{c\bar{p}} - t}{c_{\max} - \overline{c\bar{p}}} \right) \right] \right\}. \end{aligned}$$

Following Proposition 5.3.2.2, this can be written in terms of Kullback-Leibler divergence to obtain the forms given in the theorem.

**Bound 5.4.2** For the bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  note that since  $e^{-\lambda c}$  is convex in  $c$ ,  $e^{-\lambda c} \leq (1 - c/c_{\max}) + (c/c_{\max})e^{-\lambda c_{\max}}$ . So  $1 - p_i + p_i e^{-\lambda c_i} \leq 1 + (e^{-\lambda c_{\max}} - 1)c_i p_i / c_{\max}$  and we can replace the sum in (4.23) by the sum of these and proceed as before arriving at:

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq e^{-\lambda n(t - \overline{c\bar{p}})} \left( 1 + \frac{\overline{c\bar{p}}}{c_{\max}} (e^{-\lambda c_{\max}} - 1) \right)^n. \quad (5.4.5)$$

Unlike the upper tail bound, we cannot say at this step that (5.4.2) will be tighter than (5.3.11);  $e^{-\lambda c_{\max}} < 1$  so

$$1 + \frac{\overline{c\bar{p}}}{c_{\max}} (e^{-\lambda c_{\max}} - 1) > 1 + \bar{p}(e^{-\lambda c_{\max}} - 1).$$

However,  $1 + \frac{\overline{c\bar{p}}}{c_{\max}} (e^{-\lambda c_{\max}} - 1)$  may be smaller than  $1 + \bar{p}(e^{-\lambda c_{\min}} - 1)$  since  $1 + \bar{p}(e^{-\lambda c_{\max}} - 1)$  is smaller than  $1 + \bar{p}(e^{-\lambda c_{\min}} - 1)$ .

The right-hand side of (5.4.5) is minimised at

$$\lambda = -\frac{1}{c_{\max}} \log \left[ \frac{(t - \bar{c}p)(c_{\max} - \bar{c}p)}{\bar{c}p(\bar{c}p - t - c_{\max})} \right],$$

where  $0 < t < \bar{c}p$ . Otherwise  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) = 0$ . Substituting in we arrive at the correct bound.  $\blacksquare$

Again, an alternative bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  can be found which is the same as (5.4.1) with  $p' = 1 - p$ :

$$\begin{aligned} \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &\leq \exp \left( -n D_{KL} \left( \frac{\bar{c}p' + t}{c_{\max}} \left\| \frac{\bar{c}p'}{c_{\max}} \right\| \right) \right) \quad 0 < t < c_{\max} - \bar{c}p' \\ &\leq \exp \left( -n D_{KL} \left( \frac{\bar{c}}{c_{\max}} - \frac{\bar{c}p - t}{c_{\max}} \left\| \frac{\bar{c}}{c_{\max}} - \frac{\bar{c}p}{c_{\max}} \right\| \right) \right), \end{aligned} \quad (5.4.6)$$

for  $0 < t < c_{\max} - (\bar{c} - \bar{c}p)$ . When  $\bar{c} \approx c_{\max}$  the bound is approximately the same as (5.4.2) (Again, when  $c_i = c \forall i$  the bound is the same as (5.3.11)). As we noted for the Chernoff-Hoeffding alternative bound, this bound is only tighter than bound (5.4.2) when  $p$  is large so we only consider (5.4.2).

### Chernoff-Hoeffding++

Now we derive an even tighter bound on the upper tail, which we call Chernoff-Hoeffding++, by using the convexity in  $c$  of the function  $h(c; \lambda) := (e^{\lambda c} - 1)/c = \lambda + \lambda^2 c/2 + \lambda^3 c^2/6 + \dots$

**Theorem 5.4.1.2** (Chernoff-Hoeffding++). *Under assumption 1 and 2 and with  $p_i$  and  $c_{\max}$  defined as in §5.3.1 the Chernoff-Hoeffding++ inequality is,*

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp \left\{ -n \left[ \lambda^*(\bar{c}p + t) - \log \left( 1 + \bar{c}p\lambda^* + \overline{pc^2} \left( \frac{e^{\lambda^* c_{\max}} - 1 - \lambda^* c_{\max}}{c_{\max}^2} \right) \right) \right] \right\},$$

for  $0 < t < c_{\max} - \bar{c}p$  where  $\overline{pc^2} = \frac{1}{n} \sum_{i=1}^n p_i c_i^2$  and  $\lambda^* = -\frac{1}{c_{\max}} [W(-ke^{-l}) + l]$  with

$$k = \frac{\overline{pc^2}(c_{\max} - \bar{c}p - t)}{(c_{\max}\bar{c}p - \overline{pc^2})(\bar{c}p + t)} \quad \text{and} \quad l = \frac{\overline{pc^2}(c_{\max} - \bar{c}p - t) + c_{\max}^2 t}{(c_{\max}\bar{c}p - \overline{pc^2})(\bar{c}p + t)}$$

The upper bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  for  $0 < t < c_{\max} - \bar{c} + \bar{c}p$  is as above but with  $\bar{c}p$  and  $\overline{pc^2}$  replaced by  $\bar{c}p' = \bar{c} - \bar{c}p$  and  $\overline{p'c^2} = \frac{1}{n} \sum_{i=1}^n (1 - p_i) c_i^2$  respectively.

Before proving Theorem 5.4.1.2 we introduce and prove the following useful lemma.

**Lemma 5.4.1.3.**

For any integer  $k \geq 0$  and  $x > -(k+4)/3$ ,  $h_k(x) = \frac{1}{x^{k+1}} \left( e^x - \sum_{j=0}^k \frac{x^j}{j!} \right)$  is convex.

**Proof** [Lemma 5.4.1.3] We have

$$h_k(x) = \frac{1}{x^{k+1}} \sum_{j=k+1}^{\infty} \frac{x^j}{j!} = \sum_{j=k+1}^{\infty} \frac{x^{j-k-1}}{j!}.$$

Therefore,

$$\begin{aligned} h_k''(x) &= \sum_{j=k+3}^{\infty} (j-k-1)(j-k-2) \frac{x^{j-k-3}}{j!} \\ &= \sum_{j=0}^{\infty} (j+1)(j+2) \frac{x^j}{(j+k+3)!} \\ &= \sum_{j=0}^{\infty} (2j+1)(2j+2) \frac{x^{2j}}{(2j+k+3)!} + (2j+2)(2j+3) \frac{x^{2j+1}}{(2j+k+4)!} \\ &= \sum_{j=0}^{\infty} (2j+2) \frac{x^{2j}}{(2j+k+3)!} \left[ 2j+1 + (2j+3) \frac{x}{(2j+k+4)} \right]. \end{aligned}$$

The factor outside of the bracket is non-negative for all  $x$  so we have  $h_k''(x) > 0$  if the square bracket is also non-negative, this happens when

$$x > \frac{-(2j+k+4)(2j+1)}{2j+3} := f(j)$$

for all  $j$  i.e.,  $x \geq \max_{j=0,1,2,\dots}(f(j))$ . Now  $f'(j) < 0$  so  $f(j)$  is an decreasing function in  $j$ , thus the largest value is at  $j = 0$  and we arrive at the following condition on  $x$  for which  $h_k(x)$  is convex:  $x > \frac{-(k+4)}{3}$ . ■

**Remark 5.4.1.4.** We believe the result of Lemma 5.4.1.3 is true for all  $x$ , but have been unable to prove it.

**Proof** [Theorem 5.4.1.2]

Since, by Lemma 5.4.1.3, the function  $h(c; \lambda) := (e^{\lambda c} - 1)/c = \lambda + \lambda^2 c/2 + \lambda^3 c^2/6 + \dots$  is convex in  $c$  using (5.3.1) we have

$$\frac{e^{\lambda c} - 1}{c} \leq \left( 1 - \frac{c}{c_{\max}} \right) h(0; \lambda) + \frac{c}{c_{\max}} h(c_{\max}; \lambda) \leq \left( 1 - \frac{c}{c_{\max}} \right) \lambda + \frac{c}{c_{\max}} \frac{e^{\lambda c_{\max}} - 1}{c_{\max}}.$$

Since  $1 - p + pe^{\lambda c} = 1 + pc h(c; \lambda)$  we can bound the sum in (5.3.14) as follows:

$$\begin{aligned} \sum_{i=1}^n (1 - p_i + p_i e^{\lambda c_i}) &\leq \sum_{i=1}^n \left\{ 1 + p_i c_i \left[ \left( 1 - \frac{c_i}{c_{\max}} \right) \lambda + \frac{c_i}{c_{\max}} \frac{e^{\lambda c_{\max}} - 1}{c_{\max}} \right] \right\} \\ &\leq n \left[ 1 + \overline{cp} \lambda + \left( \frac{1}{n} \sum_{i=1}^n p_i c_i^2 \right) \times \left( \frac{e^{\lambda c_{\max}} - 1 - \lambda c_{\max}}{c_{\max}^2} \right) \right]. \end{aligned}$$

Setting  $\overline{pc^2} = \frac{1}{n} \sum_{i=1}^n p_i c_i^2$  and substituting into (5.3.14) we obtain the bound for  $t > 0$ :

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp \left\{ -n \left[ \lambda(\overline{cp} + t) - \log \left( 1 + \overline{cp} \lambda + \overline{pc^2} \left( \frac{e^{\lambda c_{\max}} - 1 - \lambda c_{\max}}{c_{\max}^2} \right) \right) \right] \right\}. \quad (5.4.7)$$

Note that (5.4.7) is tighter than the Chernoff-Hoeffding+ bound since:

$$\begin{aligned} \log \left( 1 + \overline{cp} \lambda + \overline{pc^2} \left( \frac{e^{\lambda c_{\max}} - 1 - \lambda c_{\max}}{c_{\max}^2} \right) \right) &\leq 1 + \overline{cp} \lambda + \frac{\overline{pc^2}}{c_{\max}^2} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) \\ &\leq 1 + \overline{cp} \lambda + \frac{\overline{cp} c_{\max}}{c_{\max}^2} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) \\ &= 1 + \frac{\overline{cp}}{c_{\max}} (e^{\lambda c_{\max}} - 1). \end{aligned}$$

If  $t > c_{\max} - \overline{cp}$  (5.4.7) is a decreasing function in  $\lambda$  and so is minimised at  $\lambda = \infty$ . This is a trivial case since  $\mathbb{P}(S_n > \mathbb{E}[S_n] + nc_{\max} - n\overline{cp}) = \mathbb{P}(S_n > nc_{\max}) = 0$ . Otherwise the (local) minimum occurs where the gradient is 0 so the optimal  $\lambda$  is the solution to

$$\begin{aligned} \overline{cp} + t - \frac{\overline{cp} + (e^{\lambda c_{\max}} - 1) \frac{\overline{pc^2}}{c_{\max}}}{1 + \overline{cp} \lambda + \frac{\overline{pc^2}}{c_{\max}^2} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max})} &= 0 \\ \Rightarrow \overline{cp} + t + \overline{cp}(\overline{cp} + t)\lambda + \frac{\overline{pc^2}}{c_{\max}^2} (\overline{cp} + t)(e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) & \\ &= \overline{cp} + \frac{\overline{pc^2}}{c_{\max}} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) + \lambda \overline{pc^2} \\ \Rightarrow \frac{\overline{pc^2}}{c_{\max}^2} (c_{\max} - \overline{cp} - t)(e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) &= t - (\overline{pc^2} - \overline{cp}(\overline{cp} + t))\lambda \\ \Rightarrow \frac{\overline{pc^2}}{c_{\max}^2} \frac{c_{\max} - \overline{cp} - t}{\overline{pc^2} - \overline{cp}(\overline{cp} + t)} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max}) &= \frac{t}{\overline{pc^2} - \overline{cp}(\overline{cp} + t)} - \lambda \end{aligned} \quad (5.4.8)$$

The right-hand side is simply a linear function of  $\lambda$  with negative gradient and we know  $e^{\lambda c_{\max}} - 1 - \lambda c_{\max}$  is an increasing convex function of  $\lambda$ , in particular its behaviour is quadratic around 0. So there is a solution of (5.4.8) for positive  $\lambda$  when either  $\overline{pc^2} - \overline{cp}(\overline{cp} + t) > 0$  and

$c_{\max} - \overline{cp} - t > 0$  or  $\overline{pc^2} - \overline{cp}(\overline{cp} + t) < 0$  and  $c_{\max} - \overline{cp} - t > 0$ . These conditions reduce to  $t < c_{\max} - \overline{cp}$  which is always satisfied so there is a positive solution for  $\lambda$ .

Rearranging (5.4.8) further we obtain

$$e^{\lambda c_{\max}} = \frac{(c_{\max}\overline{cp} - \overline{pc^2})(\overline{cp} + t)}{\frac{\overline{pc^2}}{c_{\max}}(c_{\max} - \overline{cp} - t)}\lambda + 1 + \frac{c_{\max}^2 t}{pc^2(c_{\max} - \overline{cp} - t)}.$$

This equation is of the form  $e^{\lambda a} = \lambda b + d$  so the optimal  $\lambda$  is (5.3.6) with  $a = c_{\max}$ ,  $b = (c_{\max}\overline{cp} - \overline{pc^2})(\overline{cp} + t) \left(\frac{\overline{pc^2}}{c_{\max}}(c_{\max} - \overline{cp} - t)\right)^{-1}$  and  $d = 1 + c_{\max}^2 t \left(\overline{pc^2}(c_{\max} - \overline{cp} - t)\right)^{-1}$ . Note that  $b > 0$  (since  $t < c_{\max} - \overline{cp}$  and  $c_{\max}\overline{cp} > \overline{pc^2}$ ) so the argument of the Lambert W function is negative and the function is either undefined (argument less than  $-1/e$ ) or doubled valued. Here the value of the Lambert W function is  $-(\lambda a + \frac{ad}{b}) < -\lambda a - 1 < -1$  so the lower branch is the correct branch. Then setting  $l = \frac{ad}{b}$  and  $k = a/b$  we obtain the bound as given in Theorem 5.4.1.2.

A bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  cannot be found by replacing  $X_i$  by  $-X_i$  in the above since  $1 - p + pe^{-\lambda c}$  cannot be written in the form  $1 + pcg(c; \lambda)$  with  $g(c; \lambda)$  a convex function of  $c$ . Instead we note that with  $Z_i = c_i - X_i$  we have by Chernoff's inequality:

$$\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) \leq \frac{\prod_{i=1}^n \mathbb{E}[\exp(\lambda Z_i)]}{\exp(\lambda n(t + cp'))}$$

and we have the following bound:  $\mathbb{E}[e^{\lambda Z_i}] \leq 1 - p'_i + p'_i e^{\lambda c_i} = 1 + p'_i c_i h(c_i; \lambda)$ . Then the upper bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  follows by replacing  $p_i$  by  $p'_i$  in the above proof. ■

The lower tail bound here behaves as the upper tail bound with  $p$  replaced by  $1 - p$ , so if (5.4.7) is not tight for large  $p$ , the corresponding lower tail bound is not tight for small  $p$ . Therefore, in practice if the upper bound on  $\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt)$  is tight it may be better to use a different concentration inequality to obtain an upper bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$ . Also, while the lower tail bound is strictly tighter than the Chernoff-Hoeffding(+) alternative lower bounds ((5.3.12) and (5.4.6)), it can be looser than the more useful lower tail bounds ((5.3.11) and (5.4.2)).

## 5.4.2 Improving Bennett's inequality

Here we present the Bennett+ inequality, a tighter concentration bound found by tightening various steps in the proof for Bennett's inequality.

**Bennett+**

**Theorem 5.4.2.1** (Bennett+). *Under Assumption 1 and 2 and with  $a_{\max}$  defined as in §5.3.1 the Bennett+ inequality is, for  $0 < t < a_{\max}$ ,*

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \exp \left\{ -n \left[ \lambda_{\star} t - \frac{\lambda_{\star}^2}{2} \overline{\sigma^2} - \left( \frac{e^{\lambda_{\star} a_{\max}} - \frac{1}{2} \lambda_{\star}^2 a_{\max}^2 - \lambda_{\star} a_{\max} - 1}{a_{\max}^3} \right) \overline{a\sigma^2} \right] \right\}, \quad (5.4.9)$$

where  $\overline{\sigma^2} = \text{Var}(S_n)/n$ ,  $\overline{a\sigma^2} = \frac{1}{n} \sum_{i=1}^n a_i \sigma_i^2$  and

$$\lambda_{\star} = -\frac{1}{a_{\max}} \left[ W \left( -\frac{\overline{a\sigma^2}}{a_{\max} \overline{\sigma^2}} \exp \left( -\frac{\overline{a\sigma^2} + a_{\max}^2 t}{a_{\max} \overline{\sigma^2}} \right) \right) + \frac{\overline{a\sigma^2} + a_{\max}^2 t}{a_{\max} \overline{\sigma^2}} \right].$$

The upper bound on  $\mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt)$  for  $0 < t < -m_{\min}$  is as above but with  $a_{\max}$  and  $\overline{a\sigma^2}$  replaced by  $-m_{\min}$  and  $-\overline{m\sigma^2} = -\frac{1}{n} \sum_{i=1}^n m_i \sigma_i^2$  respectively.

**Proof** In the proof for Bennett's inequality we defined the function  $f(x) := (e^x - x - 1)/x^2 = 1/2 + x/6 + \dots$ . By Lemma 5.4.1.3  $f(x)$  is convex so we can bound it using (5.3.1):

$$f(\lambda a_i) \leq \left( 1 - \frac{a_i}{a_{\max}} \right) f(0) + \frac{a_i}{a_{\max}} f(\lambda a_{\max}) = \frac{1}{2} + \frac{a_i}{a_{\max}} \left( \frac{e^{\lambda a_{\max}} - \lambda a_{\max} - 1}{\lambda^2 a_{\max}^2} - \frac{1}{2} \right). \quad (5.4.10)$$

Inserting the bound (5.4.10) into (5.3.21) we arrive at the following:

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \exp \left\{ -\lambda nt + \frac{\lambda^2}{2} \sum_{i=1}^n \sigma_i^2 + \left( \frac{e^{\lambda a_{\max}} - \frac{1}{2} \lambda^2 a_{\max}^2 - \lambda a_{\max} - 1}{a_{\max}^3} \right) \sum_{i=1}^n a_i \sigma_i^2 \right\} \\ &\leq \exp \left\{ -n \left[ \lambda t - \frac{\lambda^2}{2} \overline{\sigma^2} - \left( \frac{e^{\lambda a_{\max}} - \frac{1}{2} \lambda^2 a_{\max}^2 - \lambda a_{\max} - 1}{a_{\max}^3} \right) \overline{a\sigma^2} \right] \right\}. \end{aligned} \quad (5.4.11)$$

The  $\lambda$  which minimises this bound is the solution to

$$\begin{aligned} t - \lambda \overline{\sigma^2} - \frac{\overline{a\sigma^2}}{a_{\max}^3} (a_{\max} e^{\lambda a_{\max}} - \lambda a_{\max}^2 - a_{\max}) &= 0 \\ \Rightarrow e^{\lambda a_{\max}} &= \lambda \left( a_{\max} - \frac{a_{\max}^2 \overline{\sigma^2}}{a_{\max}^3} \right) + 1 + \frac{a_{\max}^2 t}{a_{\max}^3}. \end{aligned} \quad (5.4.12)$$

The right hand side of the equation is a linear function in  $\lambda$  with negative gradient and intercept greater than 1, hence there will be one, positive, solution for  $\lambda$ . The equation

(5.4.12) is of the form  $e^{\lambda a} = \lambda b + d$  so the optimal  $\lambda$  is (5.3.6) with  $a = a_{\max}$ ,  $b = a - \frac{a^2 \overline{\sigma^2}}{a \sigma^2}$  and  $d = 1 + \frac{a^2 t}{a \sigma^2}$  i.e.,

$$\lambda = -\frac{1}{a_{\max}} \left[ W \left( -\frac{\overline{\sigma^2}}{a \sigma^2 - a_{\max} \overline{\sigma^2}} \exp \left( -\frac{\overline{\sigma^2} + a_{\max}^2 t}{a \sigma^2 - a_{\max} \overline{\sigma^2}} \right) \right) + \frac{\overline{\sigma^2} + a_{\max}^2 t}{a \sigma^2 - a_{\max} \overline{\sigma^2}} \right].$$

Since  $b < 0$  and  $a > 0$  the argument of the Lambert W function is positive and so the function is defined and single-valued.

Similarly, an upper bound on  $\mathbb{P}(S_n - \mathbb{E}[S_n] \leq -nt)$  is obtained by replacing  $a_{\max}$  by  $-m_{\min}$  and  $a_i$  by  $-m_i$  in the above proof.  $\blacksquare$

**Remark 5.4.2.2.** *Since  $0 \leq X_i \leq c_i \forall i$  we have  $0 \leq S_n \leq \sum_{i=1}^n c_i = n\bar{c}$  and thus  $-n\bar{c}p \leq S_n - \mathbb{E}[S_n] \leq n(\bar{c} - \bar{c}p)$  i.e., it is impossible for the sum  $S_n$ , to deviate from it's expectation by more than  $n\bar{c}p$  below or more than  $n\bar{a}$  above. Therefore we have the following trivial probabilities:*

$$\begin{aligned} \mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &= 0 & \text{for } t > \bar{a} = \bar{c} - \bar{c}p \\ \mathbb{P}(S_n \leq \mathbb{E}[S_n] - nt) &= 0 & \text{for } t > \bar{c}p. \end{aligned}$$

**Remark 5.4.2.3.** *When  $c_i = c \forall i$  the bound, (5.4.11), to be minimised becomes*

$$\exp \left( -\lambda nt + n \left[ \lambda^2 f(\lambda a_{\max}) \frac{\overline{(1-p)\sigma^2}}{1-p_{\min}} + \frac{\lambda}{2} \left( \overline{\sigma^2} - \frac{\overline{(1-p)\sigma^2}}{1-p_{\min}} \right) \right] \right),$$

*which is close to (5.3.23) of the Bennett proof when  $p$  is small and so the Bennett+ bound will be close to Bennett when  $c_i = c \forall i$  and  $\overline{p\sigma^2} \approx p_{\min} \overline{\sigma^2}$  is small. Furthermore, when  $p_i = p \forall i$  the Bennett+ and Bennett bounds are identical.*

### 5.4.3 Comparing our concentration inequalities on simulated binary data

We now compare our new bounds to the best performing existing bounds when  $p$  is small. We consider the sum of  $n$  independent binary variables with the same set up as in §5.3.9 but only consider the case  $C_i \sim \text{Exp}(1) + 0.1$  since when  $c_i = c \forall i$  the Chernoff-Hoeffding+ and ++ bounds are equal to Chernoff-Hoeffding and Bennett is close to Bennett+ when  $\overline{p\sigma^2} \approx p_{\min} \overline{\sigma^2}$ . Figure 5.4.1 is the same as Figure 5.3.1 (§5.3.9) except with our new concentration bounds and dropping the loser existing bounds. Here we also consider the performance of the lower tail bounds for  $C_i \sim \text{Exp}(1) + 0.1$  (Figure 5.4.2).

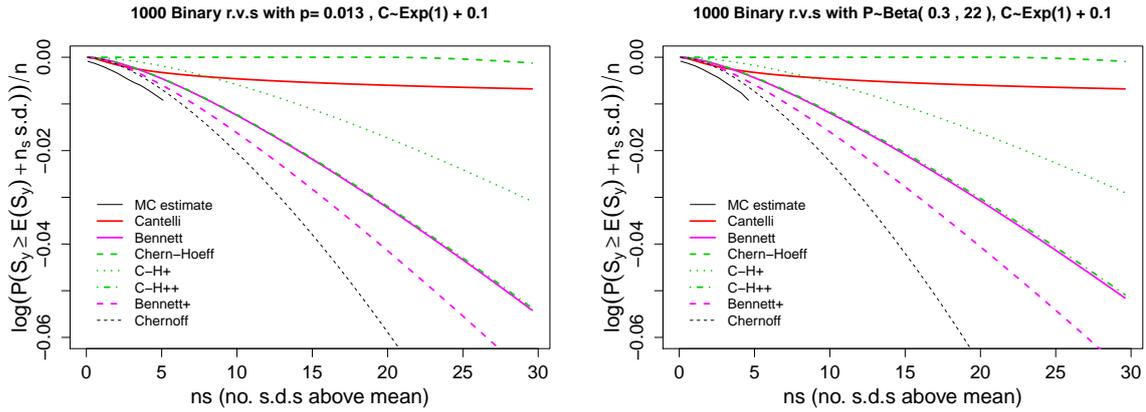


Figure 5.4.1: Bound on the exceedance probability,  $\mathbb{P}\left(S_n \geq \mathbb{E}[S_n] + n_{sd}\sqrt{\text{Var}(S_n)}\right)$ , against  $n_{sd}$  for the sum of 1000 binary random variables,  $X_i \sim c_i \text{Bernoulli}(p_i)$ . Left:  $p = 0.013$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ , Right:  $p_i \sim \text{Beta}(0.3, 22)$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ . Bounds used are: Cantelli (5.2.2) (Red), Chernoff (5.3.7) numerically optimised (dashed black), Chernoff-Hoeffding (5.3.10) (green dashed), Chernoff-Hoeffding+ (5.4.1) (green dotted), Chernoff-Hoeffding++ (5.4.7) (green dot-dash), Bennett (5.4.12) (pink), Bennett+ (5.4.9) (pink dashed). The solid black line indicates the Monte Carlo estimate based on 10000 simulations of  $S_n$ .

Of all the concentration inequalities considered in this chapter Bennett+ is the most consistent in its good performance, particularly when  $p$  is small and we have differing  $c$  values. When all  $c_i = c$  the Bennett+ bound is very close to (or when  $p_i = p \forall i$  equal to) the Bennett bound and the upper tail bound is very similar but slightly looser than the Chernoff-Hoeffding bounds (see §5.3.9 Proposition 5.3.9.1). However, when the  $c_i$ 's are different Bennett+ (5.4.9) provides a much tighter bound for both tails. In all cases where  $p$  is small the Bennett+ bound (5.4.9) is consistently close to the optimal Chernoff bound.

As noted in the proof, the Chernoff-Hoeffding+ upper tail bound (Figure 5.4.1 green dots) is tighter than the Chernoff-Hoeffding bound since  $\bar{c}p/c_{\max} \leq \bar{p}$  so the bound uses more information of the different  $c_i$ 's. When  $p$  is fixed the bound will be tighter the smaller  $\bar{c}$  is compared to  $c_{\max}$ . The Chernoff-Hoeffding++ upper tail bound (green dash-dot) is even tighter (as shown in the proof) and is close to Bennett's bound when  $p$  is small.

In our setting with Bernoulli random variables and small  $p$ , we have  $a_{\max} \approx c_{\max}$  and  $\text{Var}(S_n) \approx n\bar{p}c^2$ .

The right-hand side of (5.3.23) of the Bennett proof is:

$$\exp(-\lambda nt + \lambda^2 f(\lambda a_{\max}) \text{Var}(S_n)) \approx \exp\left(-\lambda nt + \frac{n\bar{p}c^2}{c_{\max}^2} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max})\right). \quad (5.4.13)$$

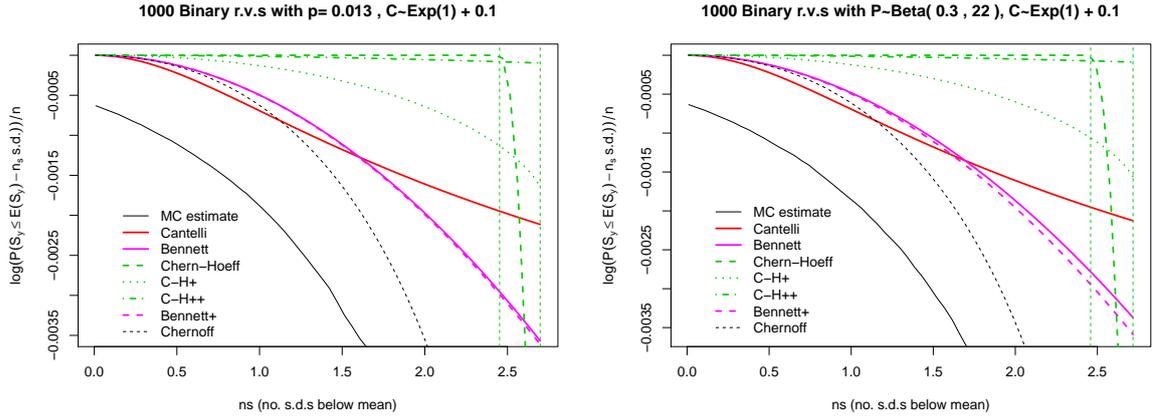


Figure 5.4.2: Bound on the exceedance probability,  $\mathbb{P}\left(S_n \leq \mathbb{E}[S_n] - n_{sd} \sqrt{\text{Var}(S_n)}\right)$ , against  $n_{sd}$  for the sum of 1000 binary random variables,  $X_i \sim c_i \text{Bernoulli}(p_i)$ . Left:  $p = 0.013$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ , Right:  $p_i \sim \text{Beta}(0.3, 22)$ ,  $c_i \sim 0.1 + \text{Exp}(1)$ . Colour and line scheme as in Figure 5.4.1. The vertical dashed green lines indicate the range of  $t$  for which the Chernoff-Hoeffding bound can be used, i.e.,  $n_s$  such that  $\bar{c}\bar{p} - c_{\min}\bar{p} < t = \frac{n_s \sigma}{n} < \bar{c}\bar{p}$ .

The right-hand side of (5.4.7) of the Chernoff-Hoeffding++ proof:

$$\begin{aligned} & \exp \left\{ -n\lambda t - n\lambda\bar{c}\bar{p} + n \log \left( 1 + \bar{c}\bar{p}\lambda + \frac{\bar{p}c^2}{c_{\max}^2} \left( e^{\lambda c_{\max}} - 1 - \lambda c_{\max} \right) \right) \right\} \\ & \approx \exp \left( -\lambda n t - n\lambda\bar{c}\bar{p} + n\lambda\bar{c}\bar{p} + \frac{n\bar{p}c^2}{c_{\max}^2} \left( e^{\lambda c_{\max}} - 1 - \lambda c_{\max} \right) \right) = (5.4.13), \end{aligned}$$

if  $\lambda\bar{c}\bar{p} + \frac{\bar{p}c^2}{c_{\max}^2} (e^{\lambda c_{\max}} - 1 - \lambda c_{\max})$  is small since  $\log(1+x) \approx x$  near  $x=0$ . This happens when  $\lambda$  is sufficiently small (in our example we need  $\lambda < \sim 1.8$  and  $\lambda_{Benn}^* < 1$  and  $\lambda_{CH++}^* < 1$ ). So for small  $p$  the Bennett and Chernoff-Hoeffding++ upper tail bound are approximately the same. Similarly when  $p$  is large the Bennett and Chernoff-Hoeffding++ lower tail bounds are similar.

In contrast the Chernoff-Hoeffding++ bound for the lower tail performs poorly, even looser than the Hoeffding bound, since  $p' = 1 - p$  is large (Figure 5.4.2). The range of values of  $t$  for which the classic Chernoff-Hoeffding lower tail bound holds is quite small in this setting so the bound is 1 for most  $t$  and sharply decreases towards 0 for  $t > \bar{c}\bar{p} - c_{\min}\bar{p}$ . Recall that in the Chernoff-Hoeffding+ lower tail proof (Bound 5.4.2) we could not show that the bound is tighter than the Chernoff-Hoeffding bound. Here, in Figure 5.4.2, we see that Chernoff-Hoeffding+ is tighter for small  $t$  but as  $t$  approaches  $\bar{a}$  it is looser than the classic bound.

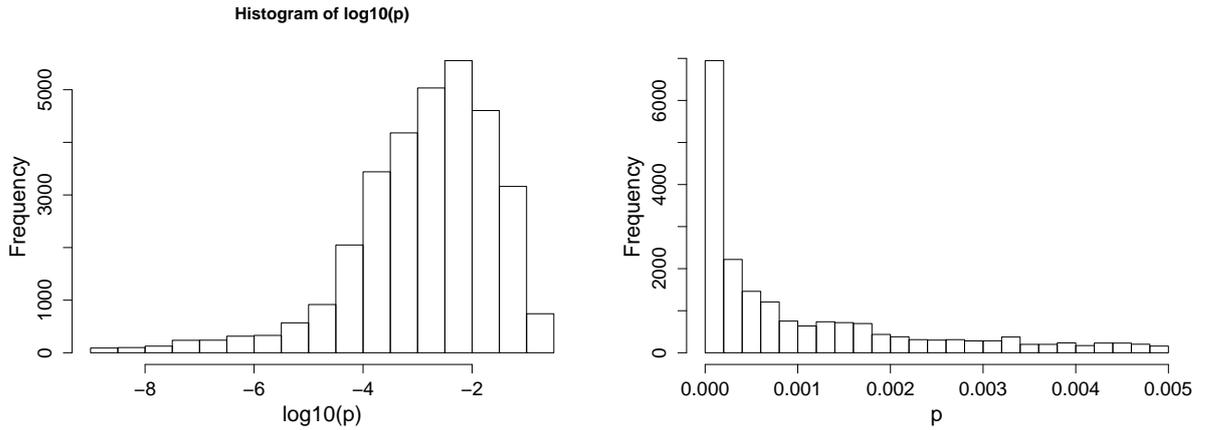


Figure 5.5.1: Histograms of  $\log_{10}(p)$  and  $p$  for all 1000 years of the JBA test data where  $p$  is the expectation of the loss as a fraction of maximum possible loss for each event and subrisk combination.

## 5.5 Connecting concentration inequality notation with the loss estimation setting

Since we are interested in estimating the tail probabilities of the total loss and thereby summing over events and risks (and subrisks) we can use the bounds described in §5.3 and §5.4. In the loss estimation setting, each subrisk within a risk is sampled independently so we can connect the concentration inequality notation with the loss simulation notation of §5.1 as follows:

Concentration inequality notation	Loss estimation notation
$X_i$	$L_{e,r,s} := \frac{v_r}{ \mathcal{S}_r } Z_{e,s} X_{e,s}$
$c_i$	$v_r /  \mathcal{S}_r $
$p_i = \mathbb{E}[X_i] / c_i$	$p_{e,r} \mathbb{E}[X_{e,r}]$

For loss simulation we are interested in the sum  $S_y = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \sum_{s \in \mathcal{S}_r} L_{e,r,s}$  which is the same as  $\sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \sum_{s \in \mathcal{S}_r} L_{e,r,s} \mathbb{1}\{p_{e,r} > 0\}$  since  $L_{e,r,s} = 0$  when  $p_{e,r} = 0$ . Therefore our ‘ $n$ ’ is the number of flooded subrisk and event combinations:  $\sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} |\mathcal{S}_r| \mathbb{1}\{p_{e,r} > 0\}$ .

A range of summary statistics is needed to calculate the various concentration inequality bounds, we have:

$$c_{\max} = \max \left( \frac{v_r}{|\mathcal{S}_r|} \right)$$

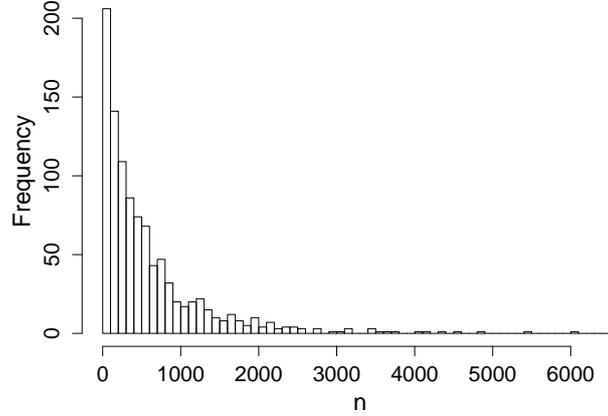


Figure 5.5.2: Histogram of  $n$ , the number of independent random variables (i.e., the number of event and subrisk combinations with possible flooding), for each of the 1000 years of the JBA test data.

$$a_{\max} = \max(c_i(1 - p_i)) = \max_{r \in \mathfrak{R}} \left( \frac{v_r}{|\mathcal{S}_r|} \left( 1 - \min_{e \in \mathfrak{E}_y} p_{e,r} \mathbb{E}[X_{e,r}] \right) \right) \quad (5.5.1)$$

$$m_{\min} = \min(-c_i p_i) = - \max_{r \in \mathfrak{R}} \left( \frac{v_r}{|\mathcal{S}_r|} \max_{e \in \mathfrak{E}_y} p_{e,r} \mathbb{E}[X_{e,r}] \right) \quad (5.5.2)$$

$$\bar{p} = \frac{1}{n} \left( \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} |\mathcal{S}_r| p_{e,r} \mathbb{E}[X_{e,r}] \right)$$

$$\overline{c\bar{p}} = \frac{1}{n} \left( \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} v_r p_{e,r} \mathbb{E}[X_{e,r}] \right)$$

$$\overline{c^2\bar{p}} = \frac{1}{n} \left( \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \frac{v_r^2}{|\mathcal{S}_r|} p_{e,r} \mathbb{E}[X_{e,r}] \right)$$

$$\overline{\sigma^2} = \frac{1}{n} \left( \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \frac{v_r^2}{|\mathcal{S}_r|} p_{e,r} \left[ \text{Var}(X_{e,r}) + (1 - p_{e,r}) \mathbb{E}[X_{e,r}]^2 \right] \right)$$

$$\overline{a\sigma^2} = \frac{1}{n} \left( \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \frac{v_r^3}{|\mathcal{S}_r|^2} p_{e,r} \left[ \mathbb{E}[(X_{e,r})^2] - p_{e,r} \mathbb{E}[X_{e,r}] (\mathbb{E}[X_{e,r}] + \text{Var}(X_{e,r})) + (p_{e,r})^2 \mathbb{E}[X_{e,r}]^3 \right] \right)$$

For our JBA test data  $p_i$  ranges from  $1 \times 10^{-9}$  to 0.24 with a mean of 0.0132, the distribution of  $p$  is highly skewed towards low values as can be seen in the histogram of  $p$  between 0 and 0.005 (Figure 5.5.1). The  $c$  values are spread between 108900 and 155200 with only 11 unique

Year	Summary statistics							
	$n$	$\mathbb{E}[S_y]$	$\overline{cp}$	$\overline{p}$	$p_{\min}$	$p_{\max}$	$p$ median	Subrisks hit per flood event
18	0	0	0	0	0	0	0	No flood events
680	227	72420	319	$2.22 \cdot 10^{-3}$	$8.33 \cdot 10^{-5}$	$8.84 \cdot 10^{-2}$	$5.24 \cdot 10^{-4}$	200, 26, 1
22	450	393600	875	$6.27 \cdot 10^{-3}$	$1.87 \cdot 10^{-6}$	$6.23 \cdot 10^{-2}$	$6.90 \cdot 10^{-4}$	3, 390, 40, 17
909	674	1362241	2021	$1.57 \cdot 10^{-2}$	$4.00 \cdot 10^{-8}$	$1.42 \cdot 10^{-1}$	$1.58 \cdot 10^{-3}$	28, 7, 48, 4, 350, 237
699	6597	22480000	3407	$2.52 \cdot 10^{-2}$	$2.00 \cdot 10^{-9}$	$2.34 \cdot 10^{-1}$	$7.87 \cdot 10^{-3}$	2, 30, 6565
All	587	1139816	1423	$1.32 \cdot 10^{-2}$	$3.49 \cdot 10^{-4}$	$8.32 \cdot 10^{-2}$	$4.09 \cdot 10^{-3}$	155

Table 5.5.1: *Summary statistics for selected years of the event set. The first row is one of the years with smallest expected loss (i.e., zero loss); the second to fourth row are the years with expected losses close to the 1st, 2nd and 3rd quartile respectively; the fifth row is the year with the highest expected loss; and the last row gives the average over all years of the summary statistic given by the column title. The final column gives the number of subrisks which may experience flooding for each of the flood events in the year (events which cannot cause flooding for this portfolio are removed).*

values.<sup>2</sup> The number of terms in the sum  $S_y$  varies greatly from year to year, Figure 5.5.2 shows the histogram of  $n$  over all years, the average is 587 with almost a fifth of years having  $n < 100$ . Some summary statistics are given in Table 5.5.1 for a selection of 5 years (the years with the highest and lowest expected loss and the years with expected losses close to the 1st, 2nd and 3rd quartile respectively) covering the entire range of expected yearly losses. The number of events causing flooding in a year ranges from 0 to 10 with an average of 3.783, whereas the number of subrisks possibly hit per flood event ranges from 1 to 6565 with high skew towards lower numbers (for example, the median is 47). The final column of Table 5.5.1 gives the number of subrisks potentially flooded for all events,  $e \in \mathfrak{E}_y$ , in the given year for which  $\exists r \in \mathfrak{R}$  such that  $p_{e,r} > 0$ . For example, in year 680 there are three flood events with possible damage to 200, 26 and 1 subrisk(s) respectively. Out of the 1000 years there are 23 years (e.g., year 18 of Table 5.5.1) with no events which can lead to flooding of risks in the portfolio, so for these years  $n = 0$  and  $S_y = 0$ . However with a larger, i.e., more realistic, portfolio we are unlikely to have years with no possible flood damage (the more risks the more likely that at least one  $p_{e,r} > 0$  in each year) and  $n$  would also be much larger.

If we are interested in the probability of being a certain number,  $n_s$ , of standard deviations above the mean then  $nt = n_s \sqrt{\text{Var}(S_y)}$  in the inequalities of the previous sections.

<sup>2</sup>There is some doubt in JBA and our understanding whether these values of  $c$  are realistic. The total value of a risk has a wider spread (Figure 5.5.3 left) but when one divides by subrisks there is little variability. This phenomenon could perhaps be down to a deliberate choice of the number of subrisks per risk so as to ensure similar values for all subrisks; however, informal communication with JBA suggested that this was not the case. Since the JBA employee who supplied us with the data subset has since left the company it was not possible to dig down to find the cause. This is why, in our simulations of §5.3.9 and §5.4.3, we cover a large range of possibilities for the variability of  $c$  by investigating sums of random variables with the same fixed maximum size,  $c$ , and sums of random variables with maximum sizes,  $C \sim \text{Exp}(1) + 0.1$ .

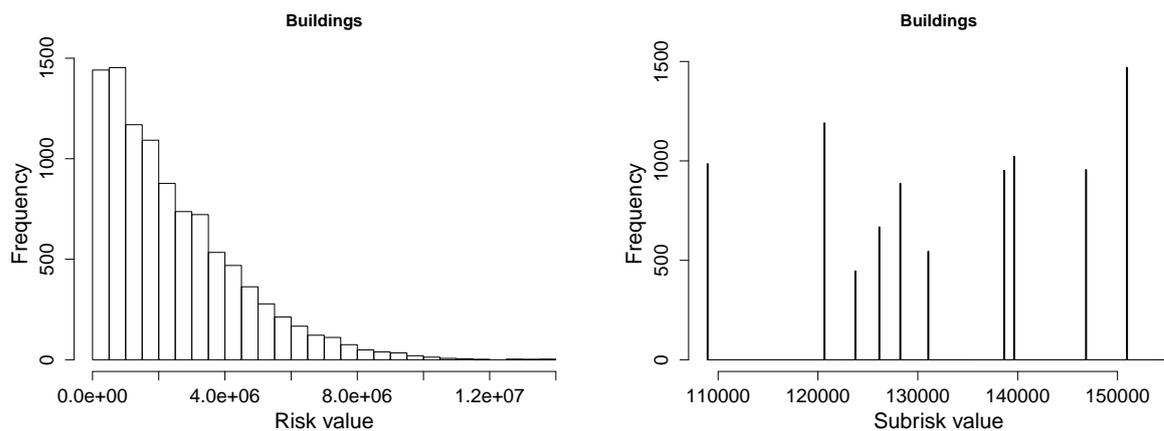


Figure 5.5.3: Histograms of  $v_r$  and  $v_r/|S_r|$  for all 1000 years of the JBA test data.

Figure 5.5.4 show the log probability bounds on the probability of the total loss (for coverage type B) of each year exceeding the expected loss for that year by some number of standard deviations. The plots of Figure 5.5.4 are very similar to those for the sum of 1000 Bernoulli random variables in §5.3.9 and §5.4.3; the largest difference is in the behaviour of the Chernoff-Hoeffding bounds. We noted in §5.4.3 when  $c_i = c \forall i$  all the Chernoff-Hoeffding bounds are equal (and perform similarly to Bennett for small  $p$ ) whereas when  $C_i \sim 0.1 + \text{Exp}(1)$  the bounds were very different with the classic Chernoff-Hoeffding being 1 or close to 1 and Chernoff-Hoeffding++ being close to the much tighter Bennett bound. For the JBA data the  $c$ 's have a small spread (Figure 5.5.3 right panel) and so we expect the bounds to be a mixture of these two behaviours - this is the case since the bounds are close but not the same, all being larger than Bennett but much smaller than 1. The Chernoff-Hoeffding++ bound is substantially larger than the Bennett bound despite the small  $\bar{p}$  since the variance approximation,  $\text{Var}(S_y) \approx n\bar{p}c^2$ , used in §5.4.3 does not hold due to the large  $c$  values.

Jebara's Bennett refinement was also considered but was almost always worse than the Bennett bound, being particularly poor for years with large  $n$ . This reinforces our decision not to use Jebara's Bennett refinement in further work. As expected all concentration bounds except Cantelli and Jebara perform better when  $n$  is larger.

The horizontal lines indicate where the bounds are equal to 0.001 and  $10^{-6}$  and so the intersections with these lines tell us the number of standard deviations above the expected loss we need to achieve an upper concentration bound of  $0.001/10^{-6}$ . For example, for year 909 the Bennett and Bennett+ bounds are 0.001 for deviations of approximately  $6\sigma$  from  $\mathbb{E}[S_y]$  whereas Cantelli requires  $31\sigma$  and Hoeffding  $37\sigma$ . It is likely when we apply our

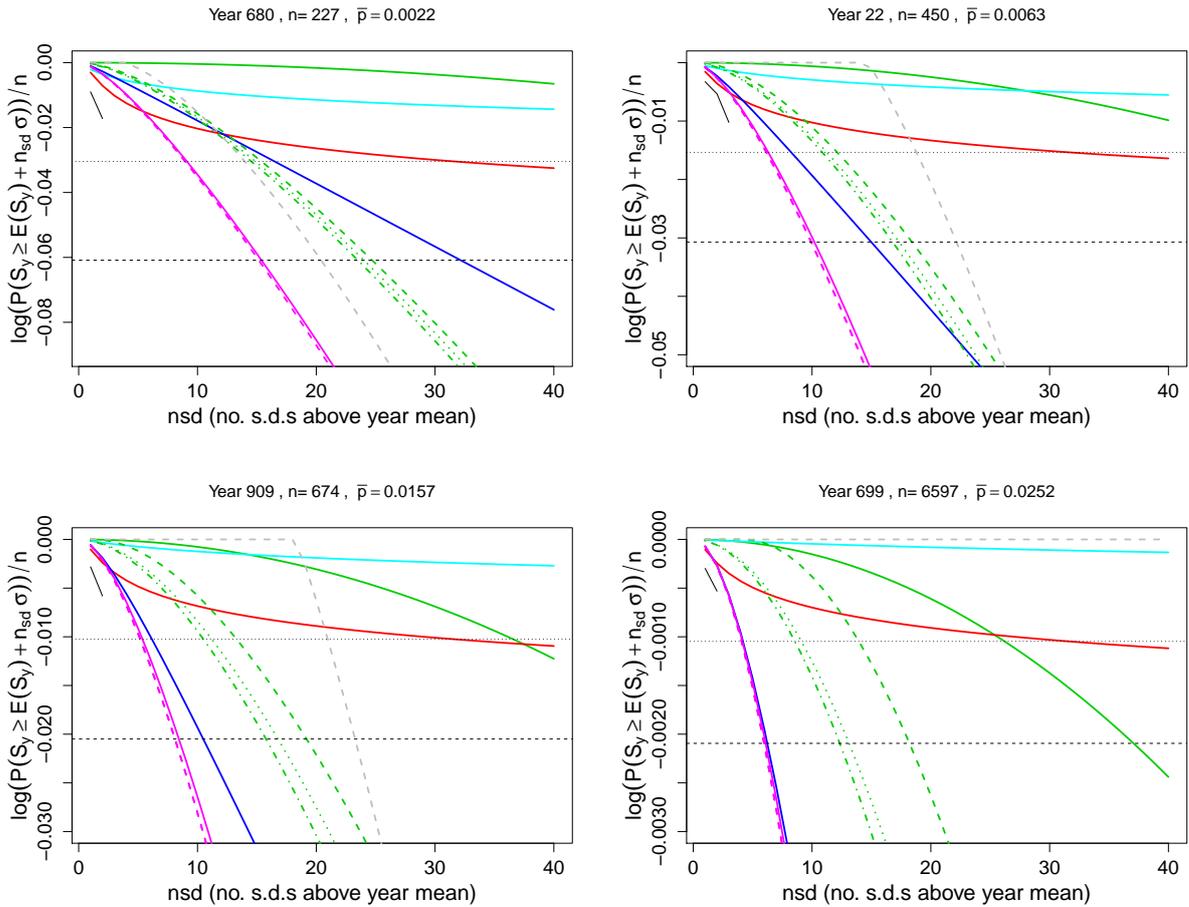


Figure 5.5.4: Log probability bounds on the probability of the total loss (for coverage type B) exceeding the expected loss for that year by some number of standard deviations for years 680, 22, 909 and 699. Colour and line scheme as in Figures 5.3.1 and 5.4.1. The horizontal grey dotted lines indicate the level corresponding to  $\varepsilon = 0.001$  and  $\varepsilon = 10^{-6}$ .

loss estimation procedure outline in §6.1 that Bennett and Bennett+ will be most useful concentration bounds.

### Examples of loss bounds

We now present, for a selection of concentration inequalities, the upper and lower bounds,  $UB_y$  and  $LB_y$  respectively, on  $S_y$  such that  $\mathbb{P}(S_y > UB_y) < \varepsilon$  and  $\mathbb{P}(S_y < LB_y) < \varepsilon$ , where  $\varepsilon$  is some small probability. We do this by setting the concentration bounds presented in §5.2-5.4 equal to  $\varepsilon$  and solving for  $t$ .

Using the Cantelli inequality the upper and lower bounds for the loss in year  $y$  are:

$$UB_y^{Cantelli} = \mathbb{E}[S_y] + \sqrt{\frac{1-\varepsilon}{\varepsilon} \text{Var}(S_y)} \quad LB_y^{Cantelli} = \mathbb{E}[S_y] - \sqrt{\frac{1-\varepsilon}{\varepsilon} \text{Var}(S_y)}. \quad (5.5.3)$$

For Hoeffding's inequality we have

$$\begin{aligned} UB_y^{Hoeffding} &= \mathbb{E}[S_y] + \sqrt{-\frac{1}{2} \log(\varepsilon) \sum_{e \in \mathcal{E}_y} \sum_{r \in \mathcal{R}} \frac{v_r^2}{|S_r|}}, \\ LB_y^{Hoeffding} &= \mathbb{E}[S_y] - \sqrt{-\frac{1}{2} \log(\varepsilon) \sum_{e \in \mathcal{E}_y} \sum_{r \in \mathcal{R}} \frac{v_r^2}{|S_r|}}. \end{aligned} \quad (5.5.4)$$

Using the Bernstein inequality we obtain:

$$\begin{aligned} UB_y &= \mathbb{E}[S_y] - \frac{a_{\max}}{3} \log(\varepsilon) + \sqrt{\left(\frac{a_{\max}}{3} \log(\varepsilon)\right)^2 - 2 \log(\varepsilon) \text{Var}(S_y)} \\ LB_y &= \mathbb{E}[S_y] - \frac{m_{\min}}{3} \log(\varepsilon) - \sqrt{\left(\frac{m_{\min}}{3} \log(\varepsilon)\right)^2 - 2 \log(\varepsilon) \text{Var}(S_y)}, \end{aligned}$$

where  $m_{\min} \leq L_{e,r} - \mathbb{E}[L_{e,r}] \leq a_{\max}$  for all  $r$  and  $e$  in year  $y$  with  $a_{\max}$  and  $m_{\min}$  as defined in (5.5.1) and (5.5.2).

Using the Bennett's inequality is more complicated since it requires finding the root of an equation. The bounds are:

$$UB_y = \mathbb{E}[S_y] + \frac{\text{Var}(S_y)}{a_{\max}} \cdot u_{UP}^* \quad LB_y = \mathbb{E}[S_y] - \frac{\text{Var}(S_y)}{-m_{\min}} \cdot u_{LOW}^*,$$

where  $u_{UP}^* = h^{-1}(-a_{\max}^2 \log \varepsilon / \text{Var}(S_y))$  *i.e.*, it is the root of the following equation:

$$(1 + u) \log(1 + u) - u + \frac{a_{\max}^2 \log(\varepsilon)}{\text{Var}(S_y)} = 0,$$

and  $u_{LOW}^*$  is the equivalent with  $a_{\max}$  replaced by  $-m_{\min}$ .

Figure 5.5.5 shows the simulated losses (top) and the base-ten logarithm of the simulated losses (bottom) for years  $10i, i = 1, \dots, 100$ , where years have been reordered according to their expected total loss. For each of these years upper and lower bounds derived from various concentration inequalities with  $\varepsilon = 1 \times 10^{-6}$  are shown. The Bernstein, Bennett and Bennett+ bounds (pink) are very tight with Bennett+ having a slightly tighter lower bound. In agreement with Figure 5.5.4, the Chernoff-Hoeffding, Chernoff-Hoeffding+ and Chernoff-Hoeffding++ upper bounds are similar, with Chernoff-Hoeffding++ providing the tightest bounds. For the bottom 500 years in terms of expected yearly loss these bounds are tighter than the Bernstein bounds (*e.g.*, Figure 5.5.4 top left) but are otherwise looser.

All of the concentration inequalities except Cantelli, Hoeffding and Chernoff-Hoeffding

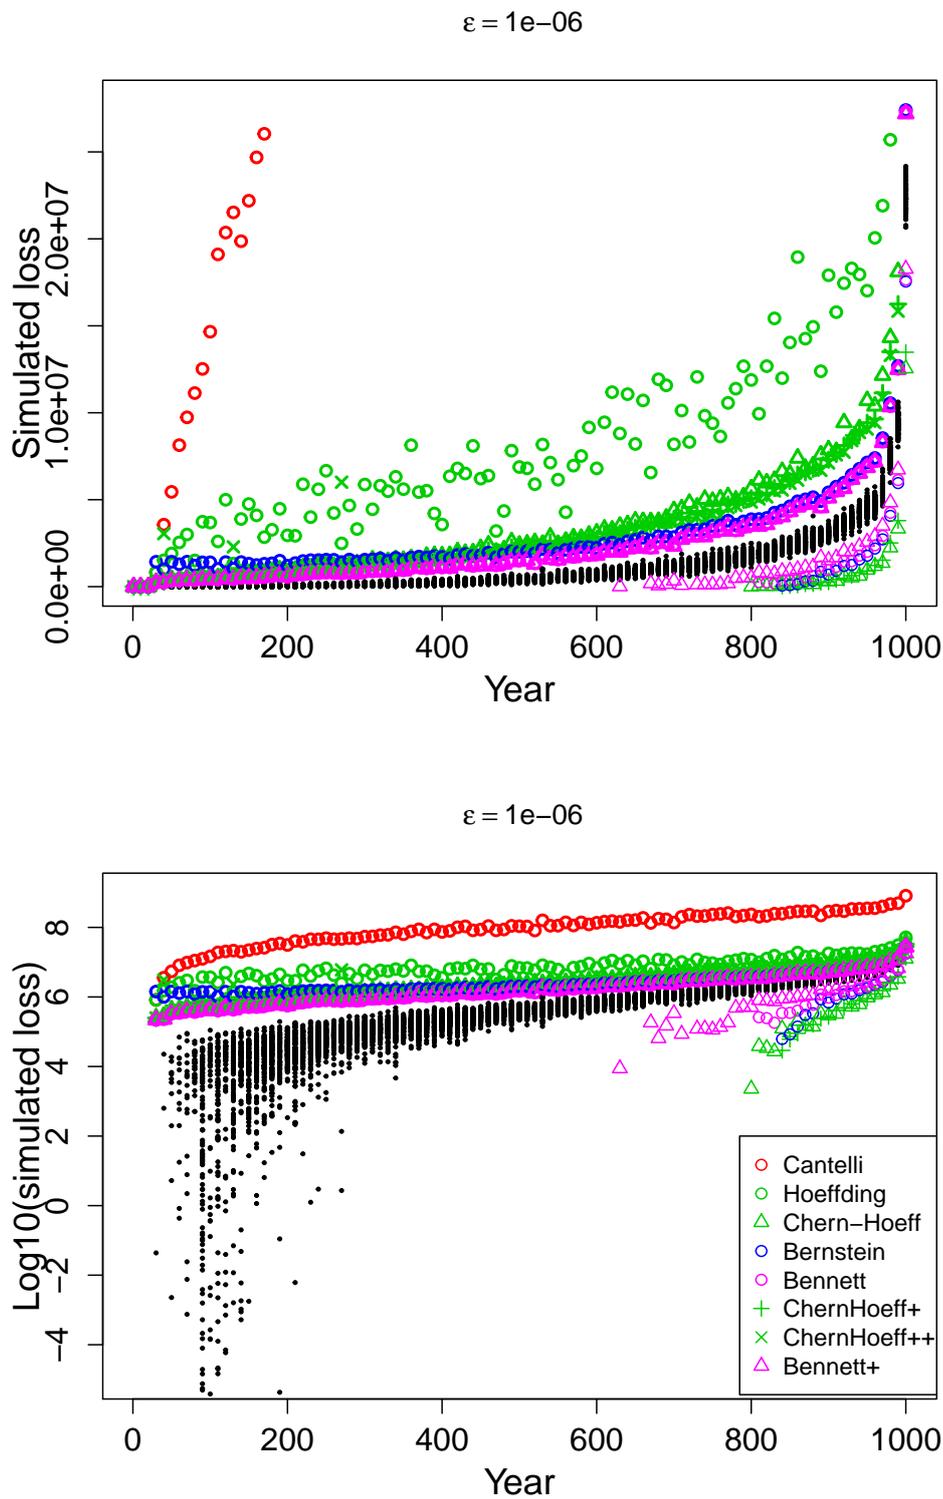


Figure 5.5.5: 100 simulated yearly losses (black) for years  $10i$ ,  $i = 1, \dots, 100$  of 1000 years ordered by yearly expected loss with upper and lower bounds on the total loss in each year derived by setting the various concentration inequalities to  $\varepsilon = 1 \times 10^{-6}$ . The lower bounds that do not appear on the plots, which includes all lower bounds for Cantelli, Hoeffding and Chernoff-Hoeffding, are zero. The top and bottom plots show the same quantities, but differ in terms of the scale of the y-axis: the top plot uses loss whereas the bottom plot shows its base 10 logarithm.

provide lower bounds greater than 0 for years with large  $n$  when  $\varepsilon = 1 \times 10^{-6}$ , in particular the top 338 years (in order of expected loss) have positive lower bounds when using the Bennett+ inequality. We know that  $\mathbb{P}(S_y \leq 0) = 0$ ; however, for a given year and concentration bound, as  $\mathbb{E}[S_y] + nt$  approaches zero (from above)  $\mathbb{P}(S_y \leq \mathbb{E}[S_y] + nt)$  does not approach zero, tending, instead to some small positive value,  $\epsilon^*$ , say. Inverting this, for  $\varepsilon \leq \epsilon^*$  the concentration inequality gives the unhelpful value for  $LB_y = 0$ . In practice, because  $p$  is small, the expected loss for each year is small compared to total loss possible, *i.e.*,  $\mathbb{E}[S_y] \ll \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} v_r$  and  $\epsilon^*$  is large for the Hoeffding, Cantelli and Chernoff-Hoeffding inequalities.

## 5.6 Summary

In this chapter we have detailed a standard procedure used by insurers to describe the yearly loss distribution due to flood events; this procedure is computationally expensive so we aim to obtain a more efficient procedure. Our approach to reducing computational cost uses concentration inequalities so in §5.2.1 we reviewed known concentration inequalities, particularly for sums of independent bounded random variables; and, in §5.4, we developed novel bounds, Chernoff-Hoeffding+ and ++ and Bennett+, improving on the Chernoff-Hoeffding and Bennett bounds respectively. Our Bennett+ concentration bound was the most consistent in its good performance out of all the concentration inequalities considered, particularly when the independent random variables had differing upper bounds and a very small expectation.

Finally, we applied the concentration bounds in the loss simulation setting, applying the various concentration bounds to the losses in each year of a test data set. Figure 5.5.5 is useful to visualise the loss simulations, in particular showing how the simulated years differ, and gives an impression of the tightness of the concentration inequalities in the loss simulation setting. However, in further work we will not simply be considering a fixed  $\varepsilon$  for each year. Instead we consider an overall  $\varepsilon_0$  bounding the total probability of both a particular set of years all exceeding a particular threshold and a particular set of years being below the same threshold. So the actual performance of each concentration inequality depends on their tightness at a range of  $t$  values. Details of the probabilities and connection to  $\varepsilon_0$  are given in §6.2. First, in Chapter 6, we introduce our method to reduce the number of years of losses simulated. For most years the losses are relatively low whereas the top years can have extremely high losses. Intuitively, our method will perform best when there is high

between-year but low within-year variance leading to a greater separation between years.

## Chapter 6

# Improving Loss Estimation

In this chapter we discuss a novel approach to improving the computational efficiency (§6.1) of the loss estimation procedure using concentration inequalities. There is some small probability that the return-level estimates of the yearly loss using our approach will differ from those using the standard loss estimation procedure; this probability is discussed in §6.2. Finally, in §6.3 we discuss a method to estimate the return levels with low return periods.

### 6.1 Efficient loss simulation via concentration inequalities

We propose a simple method, the ‘exclude method’, to reduce the number of simulations by not simulating from (*i.e.*, excluding) years which are very unlikely to change the estimate of the quantile. Recall that usually the loss for each of 10000 years is simulated 100 times and each of the 100 simulations gives an estimate of one or more return levels of interest. The method we propose here will require simulating 100 times (and obtaining 100 return-level estimates) but from a reduced number of years, thus reducing the total simulations necessary. This method is for estimating high quantiles in particular, for lower quantiles we consider other methods.

The key idea is as follows. Suppose that for some threshold,  $u$ , the losses in at least  $a$  years are all extremely likely to exceed  $u$ , and the losses in at least  $b$  years are all extremely unlikely to exceed  $u$ , then the  $\frac{n_y+1}{a}$ -year return level estimated by completely ignoring the  $b$  years is very likely to be the same as the  $\frac{n_y+1}{a}$ -year return level that would be estimated if the  $b$  ‘low’ years were also simulated. Using concentration inequalities with the quantities defined in §5.5 we can obtain an upper bound,  $\varepsilon_+^y$ , on the probability,  $\mathbb{P}(S_y > u)$ , of the total loss in year  $y$  being greater than some value,  $u$ . Similarly we can obtain an upper bound,  $\varepsilon_-^y$ , on  $\mathbb{P}(S_y < u)$ . These bounds are used to obtain a bound on the probability that the  $t$ -year return-level estimate using the ‘exclude method’ differs from the  $t$ -year return-level estimate obtained when simulating from all the years. We refer to this probability as the

*non-equivalence probability.* This probability is a function of the threshold,  $u$ , and a set of years to discard (see §6.2 for more details). The aim is to maximise the number of years discarded while keeping the non-equivalence probability low.

### Exclude procedure

**Step 1:** Find the optimal combination of threshold,  $u$ , and set of years to discard such that the non-equivalence probability of one simulation is bounded by some small, chosen,  $\varepsilon_0$  and the number of years discarded is as large as possible (see §6.2). We call the resulting set of years to discard the *maximum discard set* and denote by  $\mathcal{D}$ .

**Step 2:** Simulate total losses from the years not in the maximum discard set (Steps 1-2a of §5.1.6 standard procedure).

**Step 3:** The  $t$ -year return-level estimate for simulation  $i$ ,  $q_t^{(i)}$ , is the solution to

$$\sum_{y \in \mathcal{D}^c} \mathbb{1}_{\{s_y > q_t\}} = \frac{n_y + 1}{t} - 1.$$

As in the standard procedure we take the estimate to be a linear interpolation between the  $\left(\frac{n_y+1}{t}\right)$ th and  $\left(\frac{n_y+1}{t} - 1\right)$ th largest simulated loss:

$$\hat{q}_t^{(i)} = (1 - \omega) s_{(|\mathcal{D}^c|+1-[\kappa])}^{*(i)} + \omega s_{(|\mathcal{D}^c|+1-([\kappa]-1))}^{*(i)}, \quad (6.1.1)$$

where  $k = \frac{n_y+1}{t}$ ,  $\omega = [\kappa] - k$ , and  $s_{(j)}^{*(i)}$  is the  $j$ th order statistic ( $j$ th smallest loss) of the non-discarded years in simulation  $i$ .<sup>1</sup>

**Step 4:** Find the  $t$ -year return-level estimates and approximate 95% confidence interval as in Step 3 of the standard procedure (§5.1.6).

#### 6.1.1 Best Case from consideration of Monte Carlo samples

In Step 1 of the exclude procedure we sort the yearly losses into order of descending expectation and denote the ordered set of random variables by  $\{\mathbb{S}_k\}_{k=1}^{n_y}$  (so the loss from the year with the  $k$ th largest expectation is  $\mathbb{S}_k$  etc.). Discard sets are formed as sets of years corresponding to the random variables  $\{\mathbb{S}_k\}_{k=j}^{n_y}$  for some index  $j \geq 1$ , that is, if we discard the year with the  $j$ th largest expected loss then we also discard all years with smaller expected loss. Therefore to obtain the same  $t$ -year return-level estimate with the exclude procedure

<sup>1</sup>Note that  $n_y$  is the total number of years, not just those kept. We can imagine we have simulated from other years but they were all too low to contribute to the return level of interest.

as the standard procedure we cannot discard any years which have an higher expected loss than the year for which the  $\lceil \frac{n_y+1}{t} \rceil$ th largest simulated loss was obtained. This motivates the ‘BC-MC’ column of Table 6.1.1 which provides a ‘best case scenario’ for the number of years which we can discard while ensuring the quantile estimates,  $q_t^{(i)}$ , for each simulation,  $i = 1, \dots, 100$ , are the same and hence  $\hat{X}_t^{excl} = \hat{X}_t^{std}$  for this particular set of 100 simulations. We obtain this number in the following way:

1. For each simulation,  $i$  in  $1, \dots, 100$ :

Find the maximum index,  $k_{max}$ , of the  $\mathfrak{s}_k$  corresponding to the  $\lceil \frac{n_y+1}{t} \rceil$  top years in terms of simulated loss:

$$k_{max}^{(i)} = \max \left( k : \mathfrak{s}_k^{(i)} \geq s_{\left(n_y+1 - \lceil \frac{n_y+1}{t} \rceil\right)}^{(i)} \right).$$

To ensure that  $q_t^{excl,(i)} = q_t^{std,(i)}$  we cannot discard any years with expected loss more than  $\mathbb{E} \left[ \mathfrak{S}_{k_{max}^{(i)}} \right]$ , the  $k_{max}^{(i)}$ th largest expected loss. So for simulation  $i$  we need to keep at least  $k_{max}^{(i)}$  years, *i.e.*, we can discard at most  $n_y - k_{max}^{(i)}$  years.

2. The lower bound on the number of years that losses must be simulated from needs to be acceptable for all 100 simulations so we take our lower bound,  $k^*$ , to be  $\max_i \left( k_{max}^{(i)} \right)$ .

This  $k^*$  is only valid for the particular set of simulations it was calculated with; if only the largest  $k^*$  years in terms of expected loss had been simulated then we could not have known that none of the other years would exceed the standard return-level estimate. So BC-MC represents a sample from the distribution of a quantity that itself represents an unachievable goal but is a useful lower bound on the percentage of years it is necessary to simulate from.

### 6.1.2 Results

Table 6.1.1 shows the number of years ‘kept’, *i.e.*, not discarded, out of the 1000 years of events in the JBA test data set when using the exclude method, for a range of concentration inequalities and return periods, with the non-equivalence probability bounded by  $\varepsilon_o = 0.001$ . Both Cantelli and Hoeffding are omitted since for all return periods, all 1000 years were kept; these inequalities were expected to perform poorly since the highest potential loss for each event and risk combination varies greatly and the probabilities of losses are often very small. In Appendix D.2 there are also results for  $\varepsilon_o = 1 \times 10^{-4}$ ; these show a similar pattern

Return period	No. of kept years						
	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	1000	1000	1000	1000	1000	1000	687
5	975	977	762	869	977	1000	303
10	862	517	415	398	835	1000	158
20	690	274	242	228	618	1000	99
50	376	81	78	75	299	621	30
75	310	47	43	40	206	364	21
100	266	33	32	32	169	318	14
150	147	20	20	20	85	176	7
200	139	20	20	19	81	143	7
250	128	17	15	14	71	131	7
500	125	14	14	14	62	107	7

Table 6.1.1: *Number of years kept out of 1000 years when using the exclude method for each concentration inequality with  $\varepsilon_o = 0.001$ . The meaning behind the final column is described in §6.1.1.*

of results but with slightly more years kept since the non-equivalence probability bound is stricter.

In the setting of Table 6.1.1, we find that the BC-MC lower bound on the percentage of years kept is less than 30.5% when we are estimating return levels with return period of at least 5, with this percentage rapidly decreasing with increasing return period. For the 2-year return level, however, the number of years simulated cannot be reduced by more than 30.1% for this particular set of simulations. For most years the simulated losses will be relatively low whereas years with the highest expected loss can have extremely high simulated losses; for example, in Figure 5.5.5 there is a large jump in simulated losses from the year with the 10th highest expected loss to the year with the highest expected loss. Since our method works better the more separation there is between years, the behaviour of this data leads to a large reduction in years simulated for high return levels using the exclude procedure but a high percentage of years needed to estimate the 2-year return level.

Using the Bernstein, Bennett or Bennett+ concentration inequalities leads to the largest reduction in the number of years simulated. In particular, for high return periods, the number of kept years is close to the best case scenario given by BC-MC (for example, with  $\varepsilon_o = 0.001$ , for the 100-year return period only losses from 3.2-3.3% of years need to be simulated). The discard sets using these three concentration inequalities in the exclude procedure are (essentially) the same for high return periods, but for return periods of at least 10 years Bennett+ generally leads to the largest reduction in years simulated followed closely by Bennett.

The Chernoff-Hoeffding+ inequality reduces the number of years simulated compared to

the classic Chernoff-Hoeffding for return periods larger than 5 years, however, the number of kept years is up to 5 times that using Bernstein, Bennett or Bennett+.

The Chernoff-Hoeffding++ inequality is generally poor; it performs similarly to the classic Chernoff-Hoeffding inequality for high return periods but results in a much smaller, or even empty, discard set for mid to low return periods. The non-equivalence probability is lowest (and so the exclude method performs best) when the lower/upper concentration bounds are tight for years with high/low expected loss. Recall that when  $p$  is small, and so  $p' = 1 - p$  is large, the Chernoff-Hoeffding++ probability bound on the lower tail is quite poor (as we saw for simulated binary data in Figure 5.4.2). For the JBA test data we have small  $p$ 's, even for the years with high expected loss, and so the C-H++ lower bound is poor; for example, an upper bound on the lower tail exceedance probability,  $\mathbb{P}(S_y \leq LB_y)$ , of at most  $10^{-6}$  can only be achieved using the Chernoff-Hoeffding++ inequality by  $LB_y = 0$  (crosses do not appear below the simulated values of Figure 5.5.5). This poor lower tail performance restricts the number of years which can be discarded while keeping the non-equivalence probability below  $\varepsilon_o = 0.001$ . Thus it may be more useful to use C-H+ for the lower bound and C-H++ for the upper bound to obtain a larger discard set.

Unfortunately, when estimating the 2-year return level using the exclude method with these concentration inequalities, there is no reduction in the number of years from which losses must be simulated, even when  $\varepsilon_o$  is large. This is not surprising as our procedure concentrates on bounding tail probabilities so its strength is in calculating return levels with high return periods. Recall that the maximum discard set is found such that our return-level estimates using the exclude procedure will be the same as the those from the standard procedure with probability greater than  $1 - \varepsilon_o$ . Suppose that we use the Bennett inequality and choose to discard 23.8% of the years, then the non-equivalence probability will be around 0.001 for the 5-year return level (Table 6.1.1) and much lower for higher return levels, whereas the 2-year return-level estimate will almost certainly be different to that obtained from the standard procedure. Nevertheless, given the large number of years, the 2-year return-level estimate found by simulating losses from 76.2% of the years is likely to be relatively close to the 2-year return-level estimate based on simulating all the years and also, crucially, close to the 'true' 2-year return level.

Return period	Percentage of simulations performed						
	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	100.0	100.0	100.0	100.0	100.0	100.0	95.5
5	100.0	100.0	97.2	99.1	100.0	100.0	69.6
10	99.0	86.9	79.8	78.5	98.7	100.0	50.5
20	95.1	66.4	62.4	60.5	91.9	100.0	37.2
50	76.5	32.9	31.9	31.1	69.2	92.1	15.2
75	70.2	22.8	21.5	20.6	57.4	75.4	14.9
100	65.3	18.2	17.9	17.9	51.8	71.1	11.4
150	48.0	13.0	13.0	13.0	33.9	52.8	7.3
200	46.4	13.0	13.0	12.4	32.9	47.2	6.1
250	44.1	11.4	10.5	10.0	29.9	44.6	6.1
500	43.5	10.0	10.0	10.0	27.7	39.5	6.1

Table 6.1.2: *Simulated subrisk and event combinations for the exclude procedure as a percentage of the standard procedure time with 1000 years for each concentration inequality with  $\varepsilon_o = 0.001$ .*

### 6.1.3 Comparison of standard and exclude procedure times

Overall the results in Table 6.1.1 are promising, however, the percentage of years kept is not the same as the percentage of total standard simulation time taken by the new procedure. On one hand, many of the years discarded are those with a lower number of flood events and risks which experience loss therefrom, *i.e.*, the sum (5.1.2) has a smaller number of terms, and so these years are quicker to simulate relative to the kept years. Therefore, the percentage simulation time saved is less than the percentage of years discarded. In fact, the computational time is proportional to the total number of subrisk losses that must be simulated summed over all events in all the years simulated. Table 6.1.2 gives the number of subrisk and event combinations simulated when using the exclude procedure as a percentage of the total number of event and subrisk combinations simulated with the standard procedure. Furthermore, the exclude method introduces a new step to the loss estimation procedure, namely calculation of the discard set,  $\mathcal{D}$ , which adds to the total procedure time. We can write the procedure time for the exclude procedure as follows:

$$T \text{ (Procedure time)} = T_{sim} \text{ (Simulation)} + T_{disc} \text{ (Finding discard set)} + T_{other}$$

The third term,  $T_{other}$ , is negligible compared to  $T_{sim}$  and  $T_{disc}$ . Figure 6.1.1 shows for each concentration inequality the % of years kept (right) and the time  $T$  (left) as a percentage of  $T^{std}$ , the total simulation time for the standard procedure.

For return periods of 75 years and above the loss simulation stage with the Bernstein,

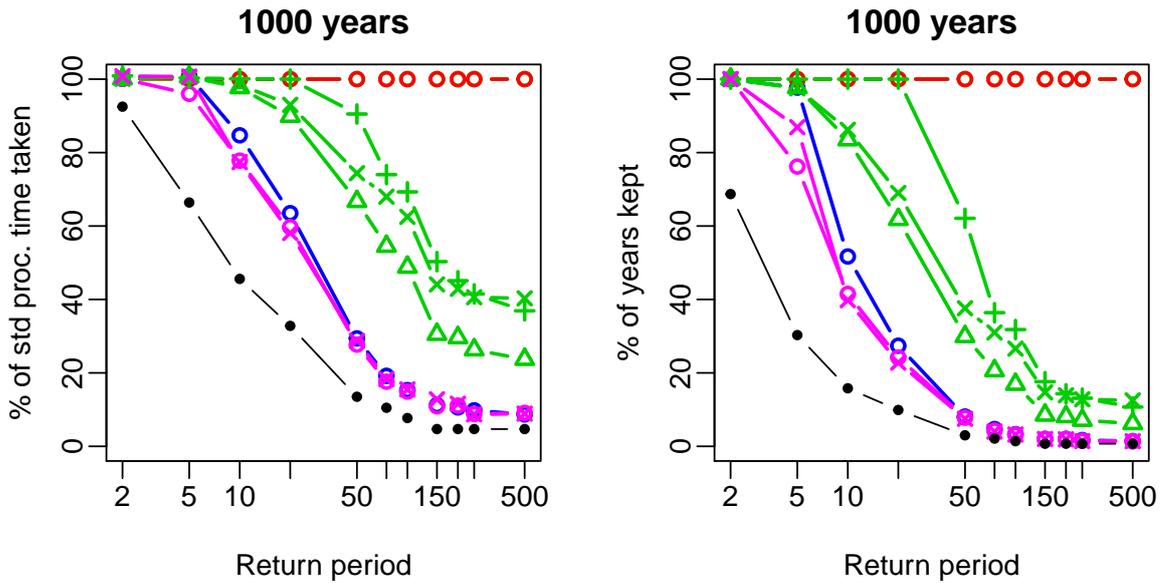


Figure 6.1.1: Comparing performance of our procedure with various concentration inequalities against the standard procedure with a 1000 year event set and  $\varepsilon_o = 0.001$ . Left:  $T$ , all as a percentage of  $T^{std}$  and right: % of years kept. Usual color scheme for concentration inequalities (Red: Cantelli, Green + : C-H, Green  $\Delta$  : C-H+, Green  $\times$  : C-H++, Blue: Bernstein, Pink  $\circ$ : Bennett, Pink  $\times$  : Bennett+). The black filled circles is the lower bound BC-MC.

Bennett and Bennett+ inequalities takes less than 20% of the standard procedure time,  $T^{std}$ , to simulate losses from up to 4.7% of the years. For lower return periods the time taken is much larger, *e.g.*, the time taken for the exclude procedure with Bennett’s inequality and  $\varepsilon_o = 0.001$  to obtain the 10-year return-level is 79% of  $T^{std}$  despite only simulating 42% of years. Overall  $T_{disc}$  is small compared to  $T_{sim}$  but for Bennett+ finding the discard set can take up to 8% of the standard procedure time (due to the cost of evaluating the Lambert W function) making it perform slightly worse overall (in terms of total procedure time) than Bennett, which tends to have a very similar maximum discard set.

In conclusion, for estimating return levels with return periods of at least 5 the exclude method with Bennett’s inequality reduces computational time by the largest amount, with time savings over 70/80% for return periods of 50/75 or higher.

Most of our work uses  $n_y = 1000$  years of events, but JBA uses  $n_y = 10000$  years of events. Multiplying the number of years by 10 whilst attempting to discard the same percentage of years might be expected to increase the non-equivalence probability by a factor of around 10. Alternatively, if the non-equivalence probability is to be maintained then fewer years may be discarded. We investigate the effect of event set size on the non-equivalence probability in Appendix D.1. Experiments comparing  $n_y = 500$  with  $n_y = 1000$  showed a general reduction

in the percentage of discarded years when increasing  $n_y$  and a small increase in the computing time as a percentage of the standard computing time. (Figures and tables for the exclude procedure applied to a random selection of 500 of the years in the portfolio are given in Appendix D.2.)

On the other hand, the portfolio sizes considered by JBA are typically much larger than the 10000-risk portfolio that we considered. Since expected losses are proportional to portfolio size and standard deviations are proportional to the square root of the size, a larger portfolio size is expected to lead to a much clearer delineation between years and a much larger percentage of discards. For example, if we increase the number of risks in a portfolio by a factor of 10 while retaining the distribution of values,  $v_r$ , and  $p_{e,r}$  etc., then most of the summary statistics ( $\overline{c\bar{p}}$ ,  $\bar{p}$ ,  $p_{max}$ , ...) would remain the same but we would have  $10n$  and  $\frac{t}{\sqrt{10}}$  instead of  $n$  and  $t$  respectively. It is straightforward to see that, when the deviation from the expectation is a fixed number of standard deviations, the Bernstein bound will be tighter for larger portfolio size. For Bennett and the Chernoff-Hoeffding inequalities the behaviour is less clear but in simulations these bounds appears to be tighter for larger portfolio sizes.

Overall we advocate our proposed procedure for the estimation of  $t$ -year return levels with  $t \geq 20$  since it requires substantially less computational effort compared to the standard procedure. For smaller return periods ( $2 \leq t < 20$ ), in the worst case scenario no years are discarded with a little additional effort to work out this empty discard set, whereas in the best case scenario substantial computational effort and time is saved by simulating from less years.

There is some probability  $< \varepsilon_0$  that the simulated loss from one or more of the discarded years would be larger than that from one or more of the kept years. This would cause the return-level estimate of the exclude method to be smaller than that using the standard procedure. In the example(s) shown here the return-level estimates using both procedures were the same. We discuss, in §6.2, an upper bound on this non-equivalence probability.

## 6.2 Non-equivalence probability

We now discuss the *non-equivalence probability*, the probability of return-level estimates using the procedure of §6.1 differing from those using the standard procedure, and how we use this probability to decide which years to discard in the simulation procedure.

Let  $\hat{X}_t^{std}$  and  $\hat{X}_t^{excl}$  be the  $t$ -year return-level estimators using the standard procedure and the exclude procedure respectively. Recall that for simulation,  $i$  of  $m$ , the  $t$ -th year return-level estimate,  $\hat{q}_t^{(i)}$ , is a weighted sum of the the  $\lceil \frac{ny+1}{t} \rceil$ th and  $(\lceil \frac{ny+1}{t} \rceil - 1)$ th largest  $s_y^{(i)}$  out of the years simulated (6.1.1). The exclude estimator,  $\hat{X}_t^{excl} = \frac{1}{m} \sum_{i=1}^m \hat{q}_t^{(i),excl}$ , will never be greater than the standard, Monte Carlo estimator; either all years discarded would have simulated losses less than  $s_{(ny+1-\lceil \frac{ny+1}{t} \rceil)}$  and so  $\hat{q}_t^{(i),std} = \hat{q}_t^{(i),excl}$ , or at least one discarded year would have a simulated loss greater than  $s_{(ny+1-\lceil \frac{ny+1}{t} \rceil)}$  so the  $\lceil \frac{ny+1}{t} \rceil$ th largest loss simulated of the kept years is smaller than  $s_{(ny+1-\lceil \frac{ny+1}{t} \rceil)}$  and  $\hat{q}_t^{(i),excl} < \hat{q}_t^{(i),std}$ .

Therefore,  $\hat{X}_t^{std}$  will always be greater than or equal to  $\hat{X}_t^{excl}$  so we are interested in the probability  $\mathbb{P}(\hat{X}_t^{std} \neq \hat{X}_t^{excl}) = \mathbb{P}(\hat{X}_t^{std} > \hat{X}_t^{excl})$ . This probability can be written in terms of the probabilities of the standard quantile estimate of one simulation exceeding the exclude quantile estimate:

$$\begin{aligned}
\mathbb{P}(\hat{X}_t^{std} \neq \hat{X}_t^{excl}) &= \mathbb{P}\left(\sum_{i=1}^m \hat{q}_t^{(i),std} \neq \sum_{i=1}^m \hat{q}_t^{(i),excl}\right) \\
&= 1 - \mathbb{P}\left(\sum_{i=1}^m \hat{q}_t^{(i),std} = \sum_{i=1}^m \hat{q}_t^{(i),excl}\right) \\
&= 1 - \mathbb{P}\left(\bigcup_{i=1}^m \{\hat{q}_t^{(i),std} = \hat{q}_t^{(i),excl}\}\right) \quad (\text{since } \hat{q}_t^{(i),std} \geq \hat{q}_t^{(i),excl} \forall i) \\
&= 1 - \prod_{i=1}^m \mathbb{P}(\hat{q}_t^{(i),std} = \hat{q}_t^{(i),excl}) \quad (\text{simulations independent}) \\
&= 1 - \left[1 - \mathbb{P}(\hat{q}_t^{(1),std} > \hat{q}_t^{(1),excl})\right]^m \quad (\text{identically distributed}) \quad (6.2.1)
\end{aligned}$$

We drop the superscript (1) in the following for simplicity.

The probability  $\mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl})$  can be decomposed by conditioning on the exceedance or non-exceedance of some arbitrary threshold,  $u$ :

$$\begin{aligned}
\mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl}) &= \mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl} | \hat{q}_t^{std} \leq u) \mathbb{P}(\hat{q}_t^{std} \leq u) + \mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl}, \hat{q}_t^{std} > u) \\
&\leq \mathbb{P}(\hat{q}_t^{std} \leq u) + \mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl}, \hat{q}_t^{std} > u). \quad (6.2.2)
\end{aligned}$$

This is true for all  $u$  so the right-hand side of (6.2.2) minimised over  $u$  gives an upper bound on the non-equivalence probability. Next we find upper bounds on  $\mathbb{P}(\hat{q}_t^{std} \leq u)$  and  $\mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl}, \hat{q}_t^{std} > u)$  in terms of  $u$ .

The probability,  $\mathbb{P}(\hat{q}_t^{std} \leq u)$ , is smaller than the probability that the  $\lceil \frac{ny+1}{t} \rceil$ th largest

simulated loss is less than  $u$ . We sort the yearly losses into order of descending expectation and denote this ordered set of random variables by  $\{\mathbb{S}\}_{k=1,\dots,n_y}$  (so the loss from the year with the  $k$ th largest expectation is  $\mathbb{S}_k$  etc.). We define the events

$$E_1(u) = \text{“at least } \left\lceil \frac{n_y+1}{t} \right\rceil \text{ of the simulated losses are } > u\text{”}$$

$$E_2(u) = \text{“the } \left\lceil \frac{n_y+1}{t} \right\rceil \text{ years with the largest expectations all have simulated losses } > u\text{”}.$$

Then,

$$\begin{aligned} \mathbb{P}\left(\hat{q}_t^{std} \leq u\right) &\leq \mathbb{P}\left(S_{(n_y+1-\lceil \frac{n_y+1}{t} \rceil)} \leq u\right) \\ &= 1 - \mathbb{P}(E_1(u)) \\ &\leq 1 - \mathbb{P}(E_2(u)) \quad \text{since } E_2(u) \Rightarrow E_1(u) \\ &= 1 - \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} \mathbb{P}(\mathbb{S}_k > u) \quad \text{(Independence)} \\ &= 1 - \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} [1 - \mathbb{P}(\mathbb{S}_k \leq u)] \\ &\leq 1 - \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} \left(1 - \varepsilon_{-, (k)}^u\right) =: b_1(u, t), \end{aligned} \tag{6.2.3}$$

where  $\varepsilon_{-, (k)}^u$  is the upper bound on  $\mathbb{P}(\mathbb{S}_k \leq u)$  found using a concentration inequality. For  $E_2$  we could use we could use a different ordering of the years, for example ordering by some chosen quantile. The tightest bound will be achieved with an ordering that reflects the ordering of the simulated losses.

To bound the probability  $\mathbb{P}(\hat{q}_t^{std} > \hat{q}_t^{excl}, \hat{q}_t^{std} > u)$  we notice that to have both  $\hat{q}_t^{std} > \hat{q}_t^{excl}$  and  $\hat{q}_t^{std} > u$  there must be one or more of the simulated losses from the discarded years greater than  $u$ . So to derive an upper bound on this probability we need to know which years are discarded; we will refer to this set of years as  $\mathcal{D}$ . We denote the event  $E_3$  “ $\exists y \in \mathcal{D}$  such that  $S_y > u$ ” then

$$\mathbb{P}\left(\hat{q}_t^{std} > \hat{q}_t^{excl}, \hat{q}_t^{std} > u\right) \leq \mathbb{P}(E_3) \leq \sum_{y \in \mathcal{D}} \mathbb{P}(S_y > u) \leq \sum_{y \in \mathcal{D}} \varepsilon_{+, y}^u =: b_2(u, t), \tag{6.2.4}$$

where  $\varepsilon_{+, y}^u$  is the upper bound on  $\mathbb{P}(S_y \geq u)$  found using a concentration inequality. In §5.3.9, §5.4.3 and §5.5 we explored how  $\varepsilon_{\pm}^u$  behaves as  $u$  changes for different concentra-

tion inequalities and, in particular, in Figure 5.5.5 of §5.5 we plotted the bounds such that  $\varepsilon_{+,y}^{\text{upperbound}} = 10^{-6}$  and  $\varepsilon_{-,y}^{\text{lowerbound}} = 10^{-6}$ . Now, each of these individual probabilities contributes to an overall probability bound,  $\varepsilon_o$ .

Inserting (6.2.3) and (6.2.4) into (6.2.2) we obtain the following upper bound on the *non-equivalence probability* for a particular simulation:

$$\mathbb{P}\left(\hat{q}_t^{std} > \hat{q}_t^{excl}\right) \leq 1 - \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} \left(1 - \varepsilon_{-, (k)}^u\right) + \sum_{y \in \mathcal{D}} \varepsilon_{+,y}^u = b_1(u; t) + b_2(u; t). \quad (6.2.5)$$

This is valid for any chosen threshold so, if the set  $\mathcal{D}$  is known, we can find the optimal threshold,  $u^*$ , which minimises this bound.

Thus, using (6.2.1), an upper bound on the probability of the return-level estimates of the exclude method differing to the standard return-level estimate with  $m$  simulations is

$$\mathbb{P}\left(\hat{X}_t^{std} \neq \hat{X}_t^{excl}\right) \leq 1 - [1 - b_1(u; t) + b_2(u; t)]^m. \quad (6.2.6)$$

Finally, using the binomial expansion, probability (6.2.6) is bound above by  $m(b_1(u; t) + b_2(u; t))$  and is approximately equal to  $m(b_1(u; t) + b_2(u; t))$  when the non-equivalence probability bound,  $b_1(u; t) + b_2(u; t)$ , is small. Using the exclude method we wish to simulate from as few years as possible while keeping  $\mathbb{P}\left(\hat{X}_t^{std} \neq \hat{X}_t^{excl}\right)$  less than  $m\varepsilon_o$  for some small, chosen,  $\varepsilon_o$ , *i.e.*, we want to keep  $b_1(u; t) + b_2(u; t) < \varepsilon_o$ .

We employ a two step procedure to find the maximum number of years,  $d^* = |\mathcal{D}|$ , that can be discarded. Firstly, to simplify and speed up the procedure, we only consider possible discard sets of the form  $\{k, \dots, n_y\}$ , with  $(n_y + 1)(1 - 1/t) < k \leq n_y$ , where the years are ordered in terms of their expected loss from highest to lowest. Thus, if a year with some particular expected loss is discarded then all years with a smaller expected loss are also discarded. Next, given a threshold,  $u$ , we perform a binary search over possible discard sets for the largest  $d = |\mathcal{D}|$  such that the non-equivalence probability bound (6.2.5) is close to, but below, our chosen  $\varepsilon_o$ . Finally, we find the maximum of this function,  $d^*$ , by optimising over possible threshold,  $u$ , using the ‘R’ function ‘optimise()’.

The above algorithm finds a threshold,  $u$ , that maximises the size of the discard set,  $|\mathcal{D}|$ . However, this threshold is not unique: all values of  $u$  over some contiguous interval would lead to the same discard set. In order to automatically find the threshold that minimises the

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	0
5	0	0	1	1	1	1	1	0
10	0	0	0.04	0.2	0.3	0.4	0.05	0
20	0	0	0.001	0.0009	0.0009	0.0009	0.0009	0
50	0	0	$2 \cdot 10^{-8}$	$1 \cdot 10^{-9}$	$8 \cdot 10^{-11}$	$2 \cdot 10^{-11}$	$4 \cdot 10^{-8}$	0
75	0	0	$2 \cdot 10^{-10}$	$4 \cdot 10^{-12}$	$3 \cdot 10^{-14}$	$5 \cdot 10^{-15}$	$4 \cdot 10^{-10}$	0
100	0	0	$1 \cdot 10^{-11}$	$1 \cdot 10^{-13}$	$4 \cdot 10^{-16}$	$7 \cdot 10^{-18}$	$1 \cdot 10^{-11}$	0
150	0	0	$1 \cdot 10^{-15}$	$1 \cdot 10^{-19}$	$1 \cdot 10^{-29}$	$3 \cdot 10^{-33}$	$8 \cdot 10^{-15}$	0
200	0	0	$1 \cdot 10^{-15}$	$9 \cdot 10^{-22}$	$1 \cdot 10^{-33}$	$1 \cdot 10^{-37}$	$9 \cdot 10^{-16}$	0
250	0	0	$1 \cdot 10^{-16}$	$9 \cdot 10^{-22}$	$1 \cdot 10^{-33}$	$1 \cdot 10^{-37}$	$8 \cdot 10^{-16}$	0
500	0	0	$1 \cdot 10^{-51}$	$9 \cdot 10^{-40}$	$1 \cdot 10^{-71}$	$7 \cdot 10^{-79}$	$4 \cdot 10^{-40}$	0
No. yrs kept	1000	1000	690	274	242	228	618	1000
% simulations	100.0	100.0	95.1	66.4	62.4	60.5	91.9	100.0

Table 6.2.1: *Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 20-year return level non-equivalence probability to be less than  $\varepsilon = 0.001$  (4th row of Table 6.1.1).*

bound on the non-equivalence probability,  $b_1(u; t) + b_2(u; t)$ , instead of maximising  $|\mathcal{D}|$ , we maximise  $|\mathcal{D}| - (b_1(u; t) + b_2(u; t))$ .

Tables 6.2.1 and D.2.8 in the appendix show the non-equivalence probability when simulating the number of years necessary for the  $t^*$ -year return level non-equivalence probability to be less than  $\varepsilon_o = 0.001$  and  $\varepsilon_o = 1 \times 10^{-4}$  respectively with  $t^* = 20$ , *i.e.*, the number of years discarded correspond to the 4th rows of Tables 6.1.1 and D.2.1 respectively. Similar tables (D.2.7, D.2.9, D.2.10 and D.2.12) are also given in the appendix for simulating the years necessary for the 10-year and 50-year return level non-equivalence probability to be less than  $\varepsilon_o = 0.001$  and  $\varepsilon_o = 1 \times 10^{-4}$ . Even when  $\varepsilon_o = 0.001$  the non-equivalence probability of the 20-year return levels being less than 0.001, the Cantelli, Hoeffding and Chernoff-Hoeffding++ inequalities lead to no years being discarded. For the remaining inequalities, the upper bound on the non-equivalence probability dramatically decreases the larger the return period of interest. The Bennett and Bennett+ inequalities lead to the smallest upper bounds on the non-equivalence probability for all return periods higher than  $t^* = 20$  with non-equivalence probability less than  $\varepsilon_o$ , despite these inequalities resulting in the largest discard sets.

It should be highlighted here that we do not know how tight the bound on the non-equivalence probability is; the derivation of the bound involved quite a few loosening steps and so the non-equivalence probability may well be much smaller than the bounds. It is also important to remember that the non-equivalence probability is the probability of the return-level estimate obtained from the exclude method differing from the standard estimate and is not a measure of the size of the discrepancy of our estimate from the true return level.

So a large non-equivalence probability does not imply that the return-level estimate from the exclude method is far from the true return level; on the contrary, for the estimation of low return levels to a desired ‘accuracy’ the number of years which need to be simulated is much less than 1000 (albeit larger than  $\left\lceil \frac{ny+1}{t} \right\rceil$ ). In the centre of the loss distribution the discrepancy between ordered losses is small so the standard estimate based on 1000 years will be very ‘accurate’ (close to the true return level) and the exclude estimate from simulating a much lower number of years will likely be close to this. So, despite the large non-equivalence probability bound, the lower  $t$ -year return-level estimates are likely to be close to the standard estimates if at least  $\left\lceil \frac{ny+1}{t} \right\rceil$  years are simulated - indeed in all simulations we have performed to date the estimates have been the same - however, we cannot determine how close to the true return level they will be. This motivates the need for a different method to estimate the lower return levels which can be used in conjunction with the fast exclude procedure for higher return levels. In §6.3 we attempt to develop such a method based on the Central Limit Theorem and the Berry-Esseen inequality.

In practice one needs to decide, for each return period, how important computational saving is compared to ensuring the return-level estimate is almost always the same as that using the standard procedure. For example, suppose we use our exclude method with Bennett’s inequality and  $\varepsilon_o = 0.001$ . If we simulate 415 of the 1000 years, we ensure that the 10-year return-level estimate is the same as the standard estimate with probability greater than 0.999 (Table 6.1.1) and for higher return levels this probability is much higher. In this case the exclude procedure takes 79% of the time of the standard procedure (Figure 6.1.1 left) as 79.8% of the event and subrisk combinations need to be simulated (Table D.2.10 3rd row). On the other hand, if we are willing to allow the 10-year return-level estimate under the exclude procedure to differ from the standard estimate we only need to simulate 242 years to ensure the non-equivalence probability of return levels with return periods of 20 years or higher is less than  $\varepsilon_o = 0.001$ , saving more than 19% of computational effort (simulating 62.4% of event and risk combinations). Moreover, we can save more than 67/80% of computation time if we are only concerned about ensuring that the non-equivalence probability of the 50-yr/100-yr return level and above is less than  $\varepsilon_o = 0.001$ . However, with  $\varepsilon_o = 0.001$ , the probability that the quantile estimates,  $\hat{X}_t^{std}$  and  $\hat{X}_t^{excl}$ , obtained by averaging over the 100 simulations will be the same is bounded above by  $100 \times 0.001 = 0.1$ .

## 6.3 Estimating return levels via the Berry-Esseen inequality

### 6.3.1 Normal approximation

The standard procedure estimates the  $t$ -year return level of yearly loss, and confidence intervals thereof, by simulating  $m$  (typically 100) times the loss from each year (1000 years in our JBA test data, 10000 years in practice). Each of these  $mn_y$  yearly loss simulations is a sum of realisations of the losses from each event and subrisk in the relevant year. Since the yearly loss,  $S_y$ , is a sum of many independent random variables we can approximate (under some mild conditions) the distribution of loss in each year by the normal distribution with the same mean and variance as the true distribution of  $S_y$ . This approximation is justified by a variant of the central limit theorem, the *Lyapunov CLT*, which requires the summands to be independent but not necessarily identically distributed random variables.

**Theorem 6.3.1.1** (Lyapunov CLT, Billingsley (1995)). *Let  $\{X_i\}_{i=1}^{\infty}$  be a sequence of independent random variables and let  $S_n = \sum_{i=1}^n X_i$ . Then if  $S_n$  has finite expectation,  $\mu_n = \sum_{i=1}^n \mathbb{E}[X_i]$ , and variance,  $\sigma_n^2 = \sum_{i=1}^n \text{Var}(X_i)$ , and the Lyapunov condition:*

$$\lim_{n \rightarrow \infty} \frac{1}{\sigma_n^{2+l}} \sum_{i=1}^n \mathbb{E} \left[ |X_i - \mathbb{E}[X_i]|^{2+l} \right] = 0$$

*holds for some  $l > 0$ , then*

$$\frac{1}{\sigma_n} (S_n - \mu_n) \xrightarrow{d} N(0, 1).$$

In the loss estimation setting we have  $S_y = \sum_{r \in \mathfrak{R}} \sum_{e \in \mathfrak{E}_y} \sum_{s \in \mathfrak{S}_r} L_{e,r,s}$  (as defined in §5.5) with the loss from each event and subrisk combination,  $L_{e,r,s}$ , being independent random variables. Thus, in our context  $n$  in Theorem 6.3.1.1 is the number of event and subrisk combinations (over all risks). There are two ways these combinations could be considered to be part of an infinite sequence: (i) more and more risks, and consequently more subrisks, could be added to the portfolio, and (ii) more and more events could occur in year  $y$ . We do not believe that there will be many more events in any given year, but we know that our portfolio set is a subsample of all the risks that could be in the portfolio, so we consider case (i). Provided that the values of all subrisks in the potentially infinite sequence are upper bounded by some constant,  $c$ , we have bounded finite expectations,  $\mathbb{E}[L_{e,r,s}] < c$ , so  $\mathbb{E}[S_y] < nc$ . We assume

that the underlying distributions of the loss modelling process<sup>2</sup> are invariant to permutations in  $r$  (that is, the distributions do not depend on the order in which risks are added to the portfolio) and that the variance,  $\text{Var}(S_y)$ , is proportional to  $n$ . Then the Lyapunov condition on the sequence of  $L_{e,r,s}$  with increasing portfolio size,  $|\mathfrak{R}| \rightarrow \infty$ , holds for any  $l > 0$ .

The normal approximation to the distribution of  $S_y$  is

$$\mathbb{P}(S_y \leq s) \approx \Phi\left(\frac{s - \mu_y}{\sigma_y}\right),$$

where  $\mu_y$  is  $\mathbb{E}[S_y]$  and  $\sigma_y$  is  $\sqrt{\text{Var}(S_y)}$  (see §5.5). For each year,  $y$ , of the  $n_y$  years, we simulate  $\tilde{m}$  times from the normal approximation for  $S_y$ . Then a  $t$ -year return-level estimate,  $\hat{X}_t^{approx}$ , can be obtained by following Steps 2b and 3 of the standard procedure (§5.1.6).

The error in this approximation is known to be larger in the tails (see the Berry-Esseen Theorem, §6.3.2 below). Similarly, the tail behaviour of a large number of simulations from the normal approximation to the yearly loss will not accurately reflect the true tail behaviour of the losses for that year. The accuracy with which simulations from the CLT will reflect the true distributions depends on how the year distributions are themselves distributed; here we focus on discussing the simulations from the CLT for the year distributions in the JBA test data.

Figure 6.3.1 shows the empirical cdf of the total loss for a selection of years (the years with the 90th, 334th, 667th and 920th largest expected losses) using losses simulated from the standard procedure (red). The cdf of the corresponding normal approximation is shown in black with grey indicating absolute bounds on the error in the approximation (the Berry-Esseen error provided in Theorem 6.3.2.1 of §6.3.2). The normal approximation is closest to the actual loss distribution for years with the largest expected losses (which are the years with the largest  $n$  and  $p$  generally). For the years with small  $n$ , (*e.g.*, Figure 6.3.1 top left) the loss distribution is skewed towards smaller values (with a large positive probability of having zero loss) but there is still a small probability of very large losses; for such years the body of the normal approximation is around larger loss values and the upper tail is lighter than that of the actual loss distribution.

In the left panel of Figure 6.3.2 the return-level estimates,  $\hat{X}_t$ , and 95% confidence intervals are plotted against return period,  $t$ , for the JBA test data using the standard procedure (red) with  $m = 100$  simulations and using the normal approximation (black) as described above

<sup>2</sup>*e.g.*, the distributions of risk value  $v_r$ , probability of flooding  $p_{e,r}$  and damage ratio  $X_{e,r}$ .

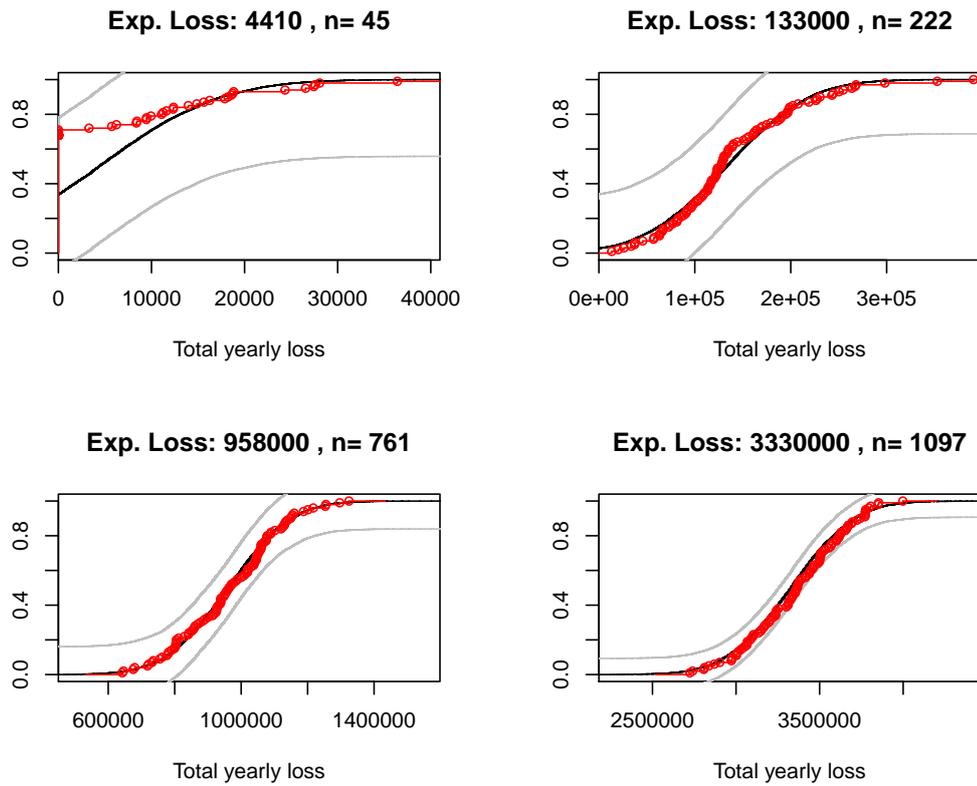


Figure 6.3.1: Comparing the normal approximation to the empirical distribution for a selection of years. The empirical cdf of the total loss using losses simulated from the standard procedure with  $m = 100$  simulations is in red, the cdf of the corresponding normal approximation is in black. Grey indicates the Berry-Esseen error i.e., the absolute bounds on the error in the normal approximation.

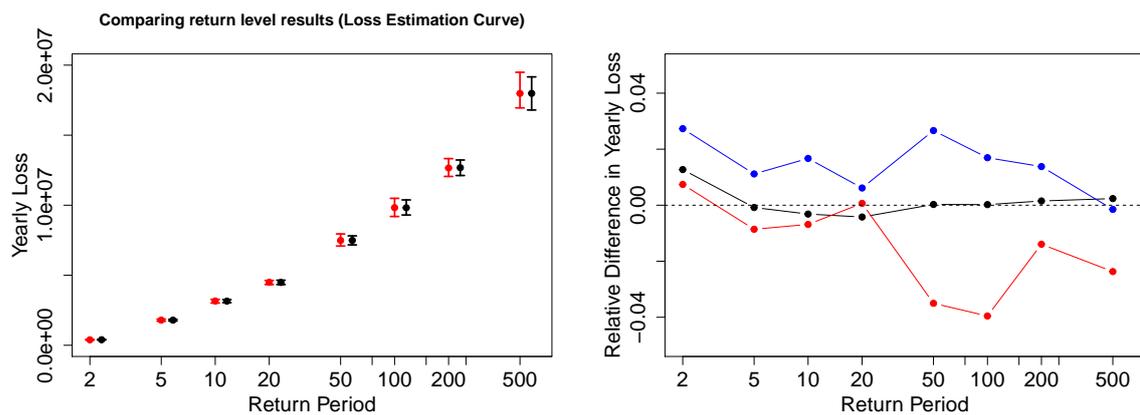


Figure 6.3.2: Comparing return-level estimates and 95% confidence intervals for the JBA test data using the standard procedure with  $m = 100$  and using the normal approximation method with  $\tilde{m} = 10000$ . Left: The Loss estimation curve using the standard procedure (red) and using the normal approximation (black). Right: The relative difference of the return level estimates (black), lower (blue) and upper (red) confidence interval bounds of the normal approximation compared to the standard procedure.

with  $\tilde{m} = 10000$  simulations. We show the differences of the return-level estimates (black) and lower (blue) and upper (red) confidence interval bounds from the two methods relative to the standard procedure in the right panel of Figure 6.3.2. The return-level estimates,  $\hat{X}_t^{std}$  and  $\hat{X}_t^{approx}$  are close for all return periods, with a relative difference less than 1.5%. The confidence intervals using the normal approximation are almost the same as the standard confidence intervals but slightly narrower with less skew towards larger values, particularly for large return periods.

The error in the normal approximation is greater in the tails (see the Berry-Esseen Theorem below) so for high return periods the return-level estimates based on the normal approximation will be less ‘accurate’ than for low return periods. Given a simulation from each of 1000 years, for example, the 100-year return level is taken to be the 10th highest simulation. The distribution of this might depend on the body of the distribution for the handful of years with the very highest expected loss, but also on the upper tail of the next 10 or so years. By contrast, the 2-year return level is estimated as the 500th highest simulation. The distribution of this might depend on the upper tail of some years with a loss that is typically relatively low and the lower tail of some years which have a loss that is typically relatively high; however, it will mainly depend on the bodies of the distributions for the many years with typical losses that are around the median. Thus, we expect low return levels to be estimated well via the CLT, and high return levels to be estimated less well.

For each year the total number of random variables simulated under the normal approximation is  $\tilde{m}n_y$  where  $n_y$  is the number of years and  $\tilde{m}$  is the number of times the approximate loss for each year is simulated. In contrast, the number of simulations under the standard procedure is  $mN$ , where  $m$  is the number of times the total loss for each year is simulated and  $N$  is the number of event and subrisk combinations which may lead to flooding. Denoting the number of event and subrisk combinations in year  $y$  by  $n(y) = \sum_{e \in \mathcal{E}_y} \sum_{r \in \mathcal{R}} |\mathcal{S}_r| \mathbb{1}\{p_{e,r} > 0\}$  we have  $N = \sum_{y=1}^{n_y} n(y)$ . For our test data we have  $N = 586459$  so almost  $5.9 \times 10^7$  simulations are needed under the standard procedure. Using the normal approximation the number of times the loss in each year is simulated can be increased while keeping computation cost low; since  $N$  is large and typically much larger than  $n_y$  we can have  $\tilde{m} > m$  and still have a low number of simulations overall, *i.e.*,  $\tilde{m}n_y \ll mN$ . For example, for our data we can have  $\tilde{m} = 10000$  and still have fewer simulations:  $10000 \times 1000 = 10^7 < 5.9 \times 10^7$ .

Overall the normal approximation method provides estimates and confidence intervals

close to those found using the standard procedure with the advantage of much reduced computational effort. Furthermore, intuitively, when  $t$  is small, using the normal approximation with  $\tilde{m} > m$  should result in  $t$ -year return-level estimates closer to the ‘true’ return level than the standard procedure. However, we do not have a way of ‘measuring’ how the return-level estimates,  $\hat{X}_t^{approx}$ , from the normal approximation will compare to those from the standard procedure (such as *e.g.*, the non-equivalence probability for the exclude procedure (§6.2)). Also, unfortunately, for 39.5% of the 1000 years the normal approximation to the distribution of yearly losses is not justified since  $\sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} |\mathcal{S}_r| p_{AA} < 20$ , that is, these years have a small number of event and subrisk combinations (recall Figure 5.5.2 in §5.5) and/or the distribution of losses in the year is skewed towards 0. Nevertheless, the high return-level estimates are driven mostly by the years with higher expected loss rather than the years with large approximation error and the low return levels are less affected by the error in the tails (as described above) so  $\hat{X}_t^{approx} \approx \hat{X}_t^{std}$ .

### 6.3.2 Berry-Esseen procedure

Under some stronger conditions, the Berry-Esseen theorem provides bounds on the error of the normal approximation to the distribution of the sum of independent (not necessarily identically distributed) random variables.

We now propose a method extending the normal approximation method to obtain an estimated range for return levels which is certain to contain the standard return-level estimate. The Berry-Esseen inequality provides an *absolute* bound on the error of the normal approximation to the distribution of the sum of  $n$  independent (not necessarily identically distributed) random variables for all  $n \in \mathbb{N}$  without the need to satisfy the Lyapunov condition. The ‘error’ is the Kolmogorov-Smirnov distance - comparing how ‘close’ two distributions are. It is the supremum of the absolute difference between the cdfs over all possible values.

**Theorem 6.3.2.1** (Berry-Esseen inequality (Berry, 1941; Esseen, 1942)).

*For the sum  $S_n = \sum_{i=1}^n X_i$  the Berry-Esseen inequality bounding the error in the normal approximation to the distribution of standardised sum is*

$$\left| \mathbb{P} \left( \frac{S_n - \mathbb{E}[S_n]}{\sqrt{\text{Var}(S_n)}} \leq t \right) - \Phi(t) \right| \leq C_0 \frac{\sum_{i=1}^n \mathbb{E}[|X_i - \mathbb{E}[X_i]|^3]}{(\text{Var}(S_n))^{3/2}}, \quad (6.3.1)$$

where  $C_0$  is a constant bounded below by  $\frac{\sqrt{10+3}}{6\sqrt{2\pi}} \approx 0.4097$  (Esseen, 1956) and above by 0.5600

(Shevtsova, 2010).

Notice that the right-hand side of (6.3.1) is the left-hand side of the Lyapunov condition with  $l = 1$  multiplied by the constant  $C_0$ . Applying the Berry-Esseen inequality to the yearly loss,  $S_y$  (as defined in §5.5) and letting  $\mu_y = \mathbb{E}[S_y]$  and  $\sigma_y = \sqrt{\text{Var}(S_y)}$ , we have:

$$\left| \mathbb{P} \left( \frac{S_y - \mu_y}{\sigma_y} \leq t \right) - \Phi(t) \right| \leq C_0 \frac{\sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} \sum_{s \in \mathfrak{S}_r} \mathbb{E} [|L_{e,r,s} - \mathbb{E}[L_{e,r,s}]|^3]}{\sigma_y^3}. \quad (6.3.2)$$

We denote the Berry-Esseen error for year  $y$  (*i.e.*, the right-hand side of (6.3.2)) by  $C_0\psi_y$ . Let  $F_y(s)$  be the true cumulative distribution function of  $S_y$  then the Berry-Esseen inequality tells us

$$\Phi \left( \frac{s - \mu_y}{\sigma_y} \right) - C_0\psi_y \leq F_y(s) \leq \Phi \left( \frac{s - \mu_y}{\sigma_y} \right) + C_0\psi_y.$$

Setting  $u = F_y(s)$  and rearranging the two inequalities gives

$$\mu_y + \sigma_y \Phi^{-1}(u - C_0\psi_y) \leq s \leq \mu_y + \sigma_y \Phi^{-1}(u + C_0\psi_y).$$

Using this information we can obtain an absolute upper and lower bound on quantiles of  $S_y$ ,  $s_{ub}(u; y)$  and  $s_{lb}(u; y)$  respectively:

$$s_{lb}(u; y) = \begin{cases} \max(\mu_y + \Phi^{-1}(u - C_0\psi_y)\sigma_y, 0) & \text{if } u > C_0\psi_y \\ 0 & \text{if } u \leq C_0\psi_y \end{cases},$$

$$s_{ub}(u; y) = \begin{cases} \min(\mu_y + \Phi^{-1}(u + C_0\psi_y)\sigma_y, s_y^{max}) & \text{if } u < 1 - C_0\psi_y \\ s_y^{max} & \text{if } u \geq 1 - C_0\psi_y \end{cases},$$

where  $s_y^{max}$  is the maximum possible loss in year  $y$ , *i.e.*,  $s_y^{max} = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} v_r \mathbb{1}\{p_{e,r} > 0\}$ , with  $v_r$  and  $p_{e,r}$  defined as in §5.5. We use these upper and lower bounds to obtain an upper and lower bound on the simulated loss for each year. We treat these losses in the same manner as the simulated yearly loss values,  $s_y^{(i)}$ ,  $i = 1, \dots, m$ , in the standard procedure to obtain an estimated range for the return level and a conservative 95% confidence interval.

### Berry-Esseen procedure

**Input:**  $\tilde{m}$

**Step 1:** For each year,  $y = 1, \dots, n_y$ :

- a) Calculate the Berry-Esseen error,  $C_0\psi_y$ , of the normal approximation as defined in

(6.3.2). See below, directly following the procedure, for details.

b) Calculate  $s_{y,max} = \sum_{e \in \mathfrak{E}_y} \sum_{r \in \mathfrak{R}} v_r \mathbb{1}\{p_{e,r} > 0\}$ , with  $v_r$  and  $p_{e,r}$  defined as in §5.5.

**Step 2:** For each simulation,  $j$ :

a) For each year,  $y$ , simulate  $u^{(j)}$  from the uniform distribution on  $[0, 1]$  and calculate:

$$s_{y,lb}^{(j)} = s_{lb}(u^{(j)}; y) \quad s_{y,ub}^{(j)} = s_{ub}(u^{(j)}; y)$$

Note that simulating a loss value from year  $y$  using the inverse probability integral transformation of the true distribution of the loss in year  $y$  and  $u^{(j)}$  would result in a value between  $s_{y,lb}^{(j)}$  and  $s_{y,ub}^{(j)}$ . Therefore  $s_{y,lb}^{(j)}$  and  $s_{y,ub}^{(j)}$  are giving a range on the possible loss simulated in year  $y$  given the simulated  $u^{(j)}$ .

b) Calculate the lower and upper  $t$ -year return-level estimate,  $q_t^{lb,(j)}$  and  $q_t^{ub,(j)}$  respectively as we described in the standard procedure in §5.1.6:

$$\begin{aligned} q_t^{lb,(j)} &= (1 - \omega) s_{lb,(n_y+1-[k])}^{(j)} + \omega s_{lb,(n_y+1-([k]-1))}^{(j)} \\ q_t^{ub,(j)} &= (1 - \omega) s_{ub,(n_y+1-[k])}^{(j)} + \omega s_{ub,(n_y+1-([k]-1))}^{(j)} \end{aligned}$$

where  $k = \frac{n_y+1}{t}$ ,  $\omega = [k] - k$  and  $s_{ub,(l)}^{(j)}$  is the  $l$ th order statistic of  $s_{ub}^{(j)}$ .

**Step 3:** We have  $\tilde{m}$  upper and lower  $t$ -year return-level estimates,  $(q_t^{ub,(1)}, \dots, q_t^{ub,(\tilde{m})})$  and  $(q_t^{lb,(1)}, \dots, q_t^{lb,(\tilde{m})})$ .

a) Take the median of the lower return-level estimates to obtain a lower  $t$ -year return-level estimate,  $\hat{X}_t^{lb}$ . Similarly obtain an upper  $t$ -year return-level estimate,  $\hat{X}_t^{ub}$ .

b) Obtain a conservative 95% confidence interval for the  $t$ -year return-level estimate by taking the 2.5th quantile of  $(q_t^{lb,(1)}, \dots, q_t^{lb,(\tilde{m})})$  and the 97.5th quantile of  $(q_t^{ub,(1)}, \dots, q_t^{ub,(\tilde{m})})$ .

We now derive the Berry-Esseen error used in the above procedure. The variance of each year is known so we only need to derive the numerator of (6.3.2). Recall (§5.5)  $L_{e,r,s} = \frac{v_r}{|\mathfrak{S}_r|} Z_{e,s} X_{e,s}$  where  $v_r$  is the value of risk  $r$ ,  $Z_{e,s} \sim \text{Bernoulli}(p_{e,r})$  and  $X_{e,s} \sim \text{Beta}(\alpha_{e,r}, \beta_{e,r})$  for event  $e$

and subrisk  $s \in \mathcal{S}_r$ . Then:

$$\sum_{r \in \mathfrak{R}} \sum_{e \in \mathfrak{E}_y} \sum_{s \in \mathcal{S}_r} \mathbb{E} [|L_{e,r,s} - \mathbb{E} [L_{e,r,s}]|^3] = \sum_{r \in \mathfrak{R}} \left( \frac{v_r}{|\mathcal{S}_r|} \right)^3 \sum_{e \in \mathfrak{E}_y} \sum_{s \in \mathcal{S}_r} \mathbb{E} [|Z_{e,s} X_{e,s} - \mathbb{E} [Z_{e,s} X_{e,s}]|^3]. \quad (6.3.3)$$

Let  $Y_{e,s} = Z_{e,s} X_{e,s}$  and denote  $\mathbb{E} [X_{e,r}]$  by  $\mu_{e,r}$ . We have  $\mathbb{E} [Y_{e,s}] = p_{e,r} \mu_{e,r}$  since  $Z_{e,s}$  and  $X_{e,s}$  are independent and for all  $s \in \mathcal{S}_r$ ,  $\mathbb{E} [X_{e,s}] = \mu_{e,r}$ . The random variable  $Y_{e,s}$  is 0 with probability  $1 - p_{e,r}$  and is  $\text{Beta}(\alpha_{e,r}, \beta_{e,r})$  distributed with probability  $p_{e,r}$  hence

$$\begin{aligned} \mathbb{E} [|Y_{e,s} - \mathbb{E} [Y_{e,s}]|^3] &= (1 - p_{e,r}) \mathbb{E} [|0 - \mathbb{E} [Y_{e,s}]|^3] + p_{e,r} \mathbb{E} [|X_{e,r} - \mathbb{E} [Y_{e,s}]|^3] \\ &= (1 - p_{e,r}) p_{e,r}^3 \mu_{e,r}^3 + p_{e,r} \mathbb{E} [|X_{e,r} - p_{e,r} \mu_{e,r}|^3]. \end{aligned}$$

Inserting into (6.3.3) we have:

$$\sum_{r \in \mathfrak{R}} \sum_{e \in \mathfrak{E}_y} \sum_{s \in \mathcal{S}_r} \mathbb{E} [|L_{e,r,s} - \mathbb{E} [L_{e,r,s}]|^3] = \sum_{r \in \mathfrak{R}} \frac{v_r^3}{|\mathcal{S}_r|^2} \sum_{e \in \mathfrak{E}_y} (1 - p_{e,r}) p_{e,r}^3 \mu_{e,r}^3 + p_{e,r} \mathbb{E} [|X_{e,r} - p_{e,r} \mu_{e,r}|^3]. \quad (6.3.4)$$

To calculate  $\mathbb{E} [|X_{e,r} - p_{e,r} \mu_{e,r}|^3]$  we need to consider the distribution of  $X_{e,r}$ , using Beta functions to obtain an expression for this expectation. The probability distribution function of the Beta distribution with parameters  $\alpha$  and  $\beta$  is

$$f(x; \alpha, \beta) = \frac{x^{\alpha-1} (1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad (6.3.5)$$

where  $B(\alpha, \beta) := \int_0^1 t^{\alpha-1} (1-t)^{\beta-1} dt$  is the Beta function.

**Definition 6.3.2.2** (The regularised incomplete Beta function). *The regularised incomplete Beta function is*

$$I_x(\alpha, \beta) := \int_0^x \frac{t^{\alpha-1} (1-t)^{\beta-1}}{B(\alpha, \beta)} dt = \int_0^x f(t; \alpha, \beta) dt = F(x; \alpha, \beta),$$

where  $f(t)$  is defined as in (6.3.5) and  $F(x; \alpha, \beta)$  is the cumulative distribution function of the  $\text{Beta}(\alpha, \beta)$  distribution.

The following equality is useful for later calculations: for any  $c \in [0, 1]$  and integer  $k \geq 0$ ,

$$\begin{aligned}
& \int_c^1 x^{\alpha+k-1}(1-x)^{\beta-1} dx - \int_0^c x^{\alpha+k-1}(1-x)^{\beta-1} dx \\
&= B(\alpha+k, \beta) \left( \int_c^1 \frac{x^{\alpha+k-1}(1-x)^{\beta-1}}{B(\alpha+k, \beta)} dx - \int_0^c \frac{x^{\alpha+k-1}(1-x)^{\beta-1}}{B(\alpha+k, \beta)} dx \right) \\
&= B(\alpha+k, \beta) (1 - 2I_c(\alpha+k, \beta)), \tag{6.3.6}
\end{aligned}$$

where  $B(\alpha, \beta)$  is the Beta function and  $I_c(\alpha, \beta)$  is the regularised incomplete Beta function evaluated at  $c$ .

Dropping the subscripts of  $\mu_{e,r}$ ,  $p_{e,r}$ ,  $\alpha_{e,r}$  and  $\beta_{e,r}$  for clarity we have

$$\begin{aligned}
\mathbb{E} [|X_{e,r} - p_{e,r}\mu_{e,r}|^3] &= \int_{p\mu}^1 (x - p\mu)^3 f(x; \alpha, \beta) dx + \int_0^{p\mu} (p\mu - x)^3 f(x; \alpha, \beta) dx \\
&= \int_{p\mu}^1 (x - p\mu)^3 \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)} dx + \int_0^{p\mu} (p\mu - x)^3 \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)} dx \\
&= (p\mu)^3 (2I_{p\mu}(\alpha, \beta) - 1) + 3(p\mu)^2 \frac{B(\alpha+1, \beta)}{B(\alpha, \beta)} (1 - 2I_{p\mu}(\alpha+1, \beta)) + \\
&\quad 3p\mu \frac{B(\alpha+2, \beta)}{B(\alpha, \beta)} (2I_{p\mu}(\alpha+2, \beta) - 1) + \frac{B(\alpha+3, \beta)}{B(\alpha, \beta)} (1 - 2I_{p\mu}(\alpha+3, \beta)) \\
&= \sum_{k=0}^3 (-1)^k \binom{3}{k} (p\mu)^{3-k} \frac{B(\alpha+k, \beta)}{B(\alpha, \beta)} [2I_{p\mu}(\alpha+k, \beta) - 1]. \tag{6.3.7}
\end{aligned}$$

Finally inserting (6.3.7) into (6.3.4) we have the Berry-Esseen error for year  $y$ .

We followed the Berry-Esseen procedure with the JBA test data and  $\tilde{m} = 10000$  to obtain lower and upper  $t$ -year return-level estimates ( $\hat{X}_t^{lb}$  and  $\hat{X}_t^{ub}$ ) and conservative 95% confidence intervals for return periods  $t = (2, 5, 10, 20, 50, 100, 150, 200, 250, 500)$ . These estimates and confidence intervals are shown in blue in Figure 6.3.3. The return-level estimates,  $\hat{X}_t^{approx}$ , based on simulating from the normal approximation for the loss in each year as described in §6.3.1 are shown as solid blue dots. The return-level estimate,  $\hat{X}_t^{std}$ , and confidence interval when using the standard procedure are in red. Recall that our main aim is to have a procedure which is more computationally efficient than the standard procedure while providing essentially the same return-level estimates. The plot of these estimates as in Figure 6.3.3 is referred to as a loss estimation curve and is used by JBAs clients to check against historical loss data.

The lower Berry-Esseen estimates,  $\hat{X}_t^{lb}$ , are relatively close to  $\hat{X}_t^{approx}$  except when  $t = 2$ , in which case  $\hat{X}_t^{lb} \approx 0.4\hat{X}_t^{approx}$ . The upper Berry-Esseen return-level estimates on the other

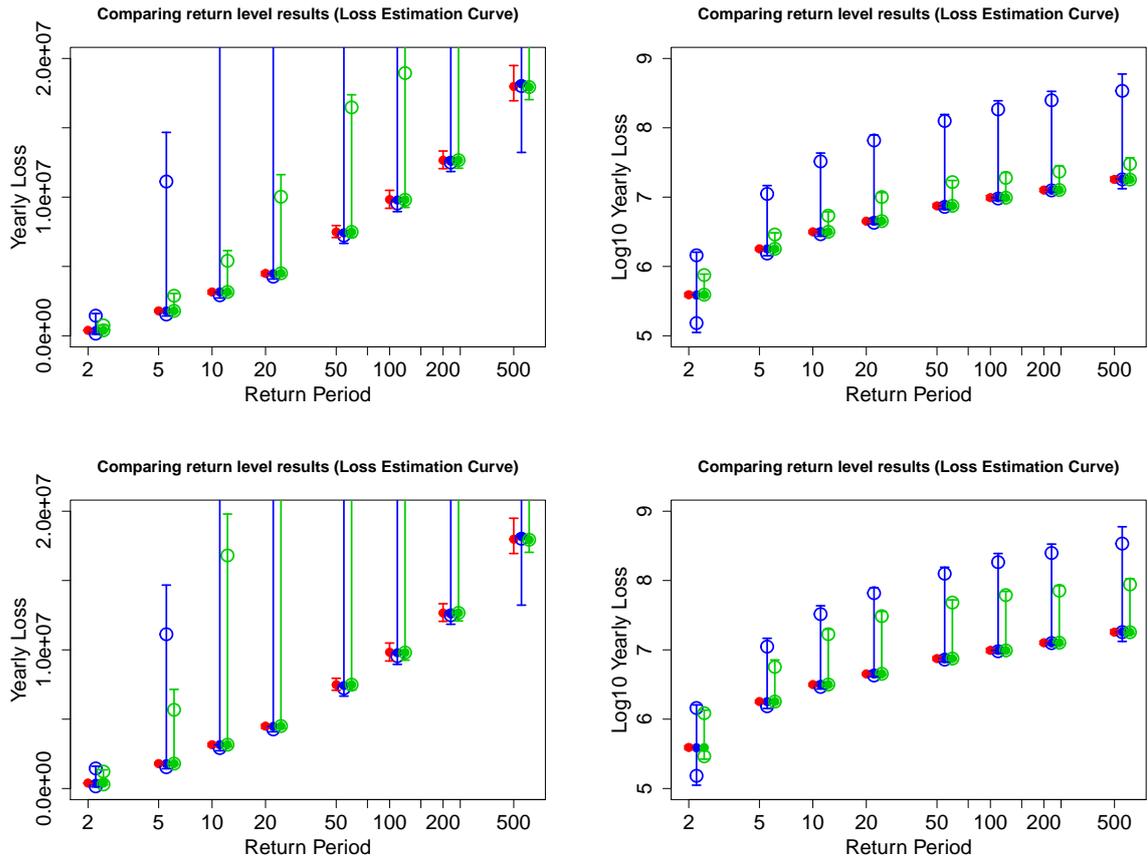


Figure 6.3.3: Loss (left) and  $\log_{10}$  Loss (right) estimation curves (return-level estimates and 95% confidence intervals) for the JBA test data using the standard procedure (red), the Berry-Esseen method (blue) and the combined method (green) introduced in §6.3.3. Solid dots are the return-level estimates (standard estimate or using the normal approximation) and circles are the upper and lower return-level estimates using the Berry-Esseen/combined method. The only difference in the top and bottom plots is in the combined method which in the top/bottom row uses simulations from years kept using Bennett’s inequality in the exclude procedure when ensuring the 5-yr/10-yr return level is the same as the standard estimate with probability greater than  $1 - \varepsilon = 0.999$ . For combined method there are 762 and 415 kept years for top and bottom respectively.

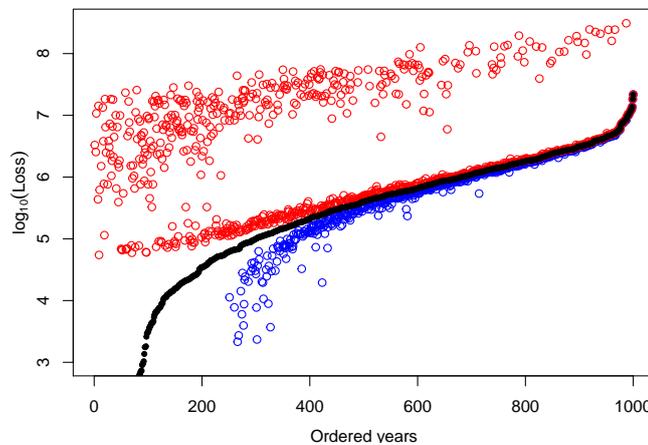


Figure 6.3.4: Upper (red) and lower (blue) bounds on the simulated loss ( $s_{y,ub}^{(j)}$  and  $s_{y,lb}^{(j)}$  respectively) and the loss simulated based on the normal approximation,  $s_y^{approx,(j)}$  (black), for each year of a given simulation.

hand are much larger than  $\hat{X}_t^{approx}$  with factor increasing with increasing return period; for the 2-year return level  $\hat{X}_2^{ub} \approx 5\hat{X}_2^{approx}$  and for the 500-year return level  $\hat{X}_{500}^{ub} \approx 20\hat{X}_{500}^{approx}$ . The Berry-Esseen confidence intervals are similarly skewed towards large values. *Note that comparing confidence intervals here is somewhat unfair as the Berry-Esseen confidence intervals are very conservative since the Berry-Esseen error bound is an absolute bound whereas the standard confidence intervals are only based on 100 simulations from the loss distribution in each year.* These poor upper return-level estimates and confidence bounds are due to a combination of large Berry-Esseen error and the asymmetry of the loss distribution.

Figure 6.3.4 shows for a particular simulation (*i.e.*, a certain  $j$ ) the losses simulated using the normal approximation,  $s_y^{approx}$ , and the upper and lower bounds on the simulated loss ( $s_{y,ub}$  and  $s_{y,lb}$  respectively) for each year ordered by expected loss. For each simulation,  $j = 1, \dots, \tilde{m}$ , 37-42% of the lower bounds are 0; those which are not zero are reasonably close to the loss simulated from the normal approximation. Using the Berry-Esseen error calculated we find the expected percentage of years for which  $s_{y,lb} = 0$  is approximately 39.3%. Similarly the expected percentage of years with the upper bound,  $s_{y,ub}$ , set to the maximum possible loss for the year is 33.0% and in our simulations the percentage ranged from 31 to 36%. The maximum possible loss is in many cases more than 100 times the simulated loss; these losses are seen clearly in a separate ‘cluster’ in the plot ranging from  $10^{5.5}$  to  $10^{8.5}$ . All the upper Berry-Esseen return-level estimates,  $\hat{X}_t^{ub}$ , with return period greater than 3 ( $\approx 1/0.33$ ) will be in this higher cluster of values.

Unfortunately, the Berry-Esseen error is greater than 1 for approximately 7.7% of all years; in these years the lower bound on the simulated loss,  $s_{y,lb}^{(j)}$ , will always be 0 and the upper bound on the simulated loss,  $s_{y,ub}^{(j)}$ , will be the maximum possible loss,  $s_{y,max}$ . Thus for a return period  $t \geq 13 \approx 1/0.077$  the upper Berry-Esseen quantile estimate,  $q_t^{ub,(j)}$ , will be a linear interpolation of the maximum possible loss in two years. Under the ‘true’ data generating mechanism described in §5.1, the probability of simulating a loss between  $s_{y,max}/2$  (say) and  $s_{y,max}$  can be bounded above using the concentration inequalities in §5.4, and is extremely low. Out of the years with Berry-Esseen error of at least 1 this probability using Bennett+ is at most  $3 \times 10^{-23}$ , and for the year with the highest loss it is at most  $5.3 \times 10^{-13}$ . Thus the upper bound of  $s_{y,max}$  occurs far too frequently.

The top right panel of Figure 6.3.5 shows the Berry-Esseen error against  $n$ ; as expected the error is smaller the larger the number of independent terms in the sum,  $S_y$ . In particular,

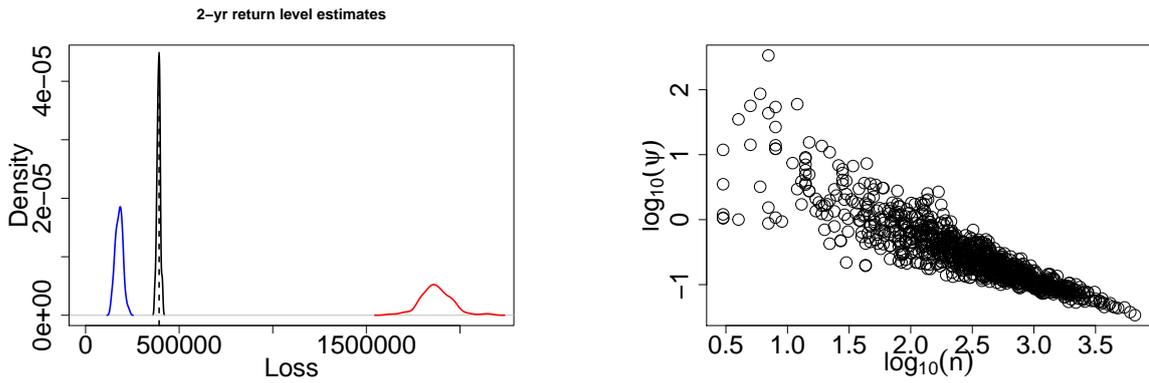


Figure 6.3.5: Applying the Berry-Esseen method to the JBA data set. Left: Densities of  $\tilde{m} = 100$  lower and upper Berry-Esseen 2-year return-level estimates (blue and red respectively) and the 2-year return-level estimate using the normal approximation (black). The dashed vertical line indicated the 2-year return-level estimate using the standard procedure. Right: Log base 10 of the Berry-Esseen error versus  $\log_{10}(n)$  for each of the 1000 years in the event set.

for all years with Berry-Esseen error more than 1 we have  $n < 200$ . Note that for a larger (and so more realistic) portfolio  $n$  would be larger and so we would expect the Berry-Esseen error to be smaller; for example a portfolio 100 times larger with the same characteristics (average number of subrisks affected for each event in the year, average variance of  $L_{e,s}$  over all event and subrisk combinations) will have a tenth of the Berry-Essen error of the smaller portfolio, leading to only approximately 1.2% of the years having Berry-Esseen errors larger than 1.

One could consider simulating directly from the years with large Berry-Esseen error; this would increase the estimation procedure time but only slightly since the worst years have small  $n$ . The smallest Berry-Esseen error over all the years is 0.0339; thus for all years there is more than 1 in 30 chance of simulating  $s_{y,ub} = s_{y,max}$ . Moreover, for a particular simulation we expect 39.3% of  $s_{y,ub} = s_{y,max}$  (as seen, for example, in Figure 6.3.4) resulting in a very large  $q_t^{ub,(j)}$ . Similarly, the expected number of years with  $s_{y,lb} = 0$  is  $\approx 330$ . However, even though this is large, the estimates of the lower bounds on the return levels are less affected by these bounds than the upper bounds were of  $s_{y,ub} = s_{y,max}$ . The lowest return period of interest is the two year return period, which for a particular replicate of simulations from the  $n_y$  years is estimated by the median, so  $q_t^{lb,(j)}$  with  $t = 2$  will typically not be influenced by any of the zero values.

### 6.3.3 Combined procedure

The Berry-Esseen procedure can be improved by combining it with the exclude procedure of §6.1. Recall that the exclude procedure simulates from a subset of the years, with this set chosen such that for  $t$  greater than the desired minimum return period,  $t^*$ , the  $t$ -year return-level estimate differs from the standard  $t$ -year return-level estimate with probability less than some chosen small  $\varepsilon$ . The smaller the minimum return period,  $t^*$ , the lower the reduction in years simulated; with the JBA test data the exclude procedure with  $t^* = 2$  resulted in no simulation reduction whereas  $t^* = 10$  reduced the number of years simulated by more than 60%. With  $t^* = 10$ , however, the non-equivalence probability bound for return periods less than 10 is large - even as large as 1 - so to estimate return levels with small return periods we need more simulations. Moreover, it may even occur that for the chosen  $t^*$  there are fewer years simulated than are necessary to calculate small return levels; for example, using the exclude procedure with Bennett+,  $\varepsilon_o = 0.001$  and  $t^* = 10$  leads to only 39.7% years simulated so the 2-year return level cannot be estimated.

We need to be able to estimate return levels with both low and high return periods. For low return periods we need simulations from the body of the yearly loss distribution whereas for high return periods we need simulations from the tail. For high return periods we can obtain good estimates and also reduce the number of years substantially using the exclude procedure as we only ‘keep’ years with high losses. This approach cannot be used for low return periods; as noted in §6.3.1 the 2-year return level depends mainly on the bodies of distributions of many years with typical losses around the median. So we need some way of estimating losses from the bodies of loss distributions of years likely to be discarded using the exclude procedure with  $t^* > 10$  (say). In the standard procedure this is achieved by simulating  $m$  times *each risk and event combination* over all years ( $N$  combinations altogether) - we propose to do this only for the ‘kept’ years according to the exclude procedure with some  $t^*$  and to simulate  $m$  times the approximate *total year loss* (not each risk and event combination) for each of the years discarded. Thus, using the same notation as in §6.3.1, we have  $m(N - N_{\mathcal{D}} + |\mathcal{D}|) < mN$  simulations overall where  $N_{\mathcal{D}} = \sum_{y \in \mathcal{D}} n(y)$ ,  $\mathcal{D}$  is the discard set and  $n(y)$  is the number of event and subrisk combinations in year  $y$ . Then we have good estimates and confidence intervals for return levels with  $t \geq t^*$  via the exclude procedure and good return-level estimates for low return periods via the normal approximation.

**Combined estimation procedure****Input:**  $\varepsilon_o$ ,  $m$ ,  $t^*$ 

1. Follow the exclude procedure with the maximum discard set  $\mathcal{D}$  found such that the non-equivalence probability for the  $t^*$ -year return level,  $\mathbb{P}\left(\hat{X}_{t^*}^{std} > \hat{X}_{t^*}^{excl}\right)$ , is less than  $\varepsilon_o$ .
2. Set  $\tilde{m} = m$  (the number of simulations used for the exclude procedure).
3. Follow steps 1 and 2a of the Berry-Esseen procedure for years in the discard set,  $y \in \mathcal{D}$ .
4. Set  $s_{y,lb}^{(k)} = s_{y,ub}^{(k)} = s_y^{(k)}$  for  $y \in \mathcal{D}^c$  and  $k = 1, \dots, m$ .
5. Follow steps 2b and 3 of the Berry-Esseen procedure taking the order statistics over the upper and lower bounds of all years.

We now discuss the combined estimation procedure applied of the JBA test data and discuss areas for improvement.

We applied the combined estimation procedure with  $t^* = 5$  and  $t^* = 10$  using Bennett's inequality and  $\varepsilon = 0.001$ ; the resulting return-level estimates and confidence intervals are shown in Figure 6.3.3 in green. The estimates and confidence interval bounds lie between those for the standard procedure and the Berry-Esseen procedure since the simulations used to obtain the estimates are a mixture of those from both procedures. The estimates and confidence intervals are closer to those of the standard procedure when  $t^*$  is smaller since then more years are simulated. Note that plots illustrate the return-level estimates and confidence intervals from the Berry-Esseen method for all return periods but in practice we would use the estimates and confidence intervals in red for return periods  $t > t^*$ . For lower return periods the non-equivalence probability is not bounded so we use the normal approximation and Berry-Esseen method.

When discussing using the Berry-Esseen procedure alone we noted it could be improved by simulating directly from the years with large Berry-Esseen error, however, since these years are generally those with low expected losses doing this in conjunction with the exclude procedure could lead to simulating almost all of the years. For  $t^* = 5$  approximately 76% of years are simulated in the exclude procedure and this takes up to 97% of the standard procedure time so additionally using the Berry-Esseen procedure could lead to a procedure time larger than the standard. For  $t^* = 10$  approximately 41% of years are simulated in

the exclude procedure and this takes up 79% of the standard procedure time. If the Berry-Esseen procedure and simulating losses from years with low expected loss is fast then the combined method with  $t^* = 10$  could provide time savings; unfortunately, the upper Berry-Esseen return-level estimate,  $\hat{X}_t^{ub}$ , and Berry-Esseen confidence bound is very poor for  $t = 2$  and  $t = 5$  (Figure 6.3.3 bottom row). So we return to the problem of needing to have a method which calculates both return levels with high return periods and those with low return periods effectively.

### 6.3.4 Discussion

To conclude, return-level estimates based on simulations from the normal approximation to the yearly loss distributions perform well, and, empirically, using the normal approximation saves much computational effort, but we cannot show to a degree of certainty how close the return-level estimates will be to the standard return-level estimates (or to the true return level). For many years the error in the normal approximation is too large for the Berry-Esseen procedure to work well, however, this should improve with larger portfolios. The combined procedure improves upon the Berry-Esseen results but still results in upper Berry-Esseen return-level estimates,  $\hat{X}_t^{ub}$ , much larger than the standard estimates along with the cost of extra computation. Overall, it seems reasonable to use the exclude procedure with Bennett's inequality for  $t \geq 10$  and the normal approximation for the 2 and 5-year return periods since we have seen that return-level estimates based on the normal approximation are very close to the standard estimates. The upper and lower Berry-Esseen return-level estimates,  $\hat{X}_t^{ub}$  and  $\hat{X}_t^{lb}$ , give a rough guide to the possible error in the 2 and 5-yr return-level estimates,  $\hat{X}_t^{approx}$ , from the normal approximation method.

The main issue with the Berry-Esseen procedure is the large error bound on the normal approximation for the loss in each year. This could be improved if it is possible in this setting to develop a tighter error bound than the Berry-Esseen bound for the normal approximation. In particular it would be useful to have a tighter bound on  $\mathbb{P}\left(\frac{S_y - \mu_y}{\sigma_y} \leq t\right) - \Phi(t)$ ; this seems plausible for our set of bounded random variables<sup>3</sup> since, as we saw using concentration inequalities, for many years  $s_{y,ub} = s_y^{max}$  is much too large. Currently the upper Berry-Esseen return-level estimates,  $\hat{X}_t^{ub}$ , and confidence interval bounds are created from essentially the upper tails of the loss in each year - this is extremely unlikely to occur so the high estimates

<sup>3</sup>In the general setting, the Berry-Esseen theorem tells us that it is not possible to get a much tighter bound using the expected sum of cubes of absolute values.

and confidence bounds are too conservative. Some sort of cross-year distribution would be needed to create estimates and confidence bounds which are not based on taking the ‘worst case’ in every year rather have a few ‘bad’ years. Another consideration would be whether there is another distribution which better approximates the true loss distribution and if so, whether some sort of error bound could be developed for this. This could be of particular use for the years with small  $n$ , (*e.g.*, Figure 6.3.1 top left) for which the loss distribution is skewed towards 0.

Finally, one could obtain a range for the return-level estimate using concentration bounds instead of the Berry-Esseen bound. This could be done by following the Berry-Esseen procedure but replacing  $s_{y,ub}^{(j)}$  in Step 2a by  $\mathbb{E}[S_y] + nt$  where  $t$  solves  $C(t) = u^{(j)}$  and  $C(t)$  is the concentration bound on the upper tail (*e.g.*,  $\sigma_y/(\sigma_y + t^2)$  for Cantelli). Similarly  $s_{y,lb}^{(j)}$  can be replaced using the concentration bound on the lower tail.

## Part III

# Extreme Values of Temporally Dependent and Multivariate Sequences

# Chapter 7

## Extremes of dependent sequences

In this chapter we discuss two extensions to the classical univariate extreme value theory. First, in §7.1, we focus on extreme values when the series exhibits temporal dependence and discuss so called declustering methods. In §7.2 we introduce a stationary process, the ARMAX process, for which we derive the extremal index and coefficient of asymptotic dependence and apply declustering methods. Second, in §7.3, we address the extension of univariate extreme value theory to many dimensions. This theory is substantially more complex than the univariate case since now the dependence structure between variables needs to be incorporated into the model. We build up through the literature from the classical approach of componentwise maxima to the point process representation of multivariate extremes with particular interest in measures of dependence. This material is mostly well-known and is presented here primarily as a stepping stone in the developments for Chapter 8.

### 7.1 Stationary dependent sequences

In many processes there is a degree of temporal dependence between the consecutive observations. For example, if an area experiences intense flooding one day it is likely that the next day there will also be flooding. This violates the independence assumption which the extreme value analysis has been based upon in the previous sections. To handle this violation, when the process is stationary, we assume that extreme events are near-independent if there are far enough apart in time, which is a reasonable assumption for many physical processes. Under such an assumption the limit distribution of the block maxima is GEV and has a nice connection to the limit for the corresponding independent sequence of random variables. We now formalise this finding mathematically.

Let  $\{X_t\}_{t \geq 1}$  be a stationary sequence of dependent variables with common distribution  $F$  and define  $M_n := \max(X_{t+1}, \dots, X_{t+n})$ . Let  $\{\tilde{X}_t\}_{t \geq 1}$  be the corresponding sequence of independent variables with common distribution  $F$ . We define  $\tilde{M}_n := \max(\tilde{X}_{t+1}, \dots, \tilde{X}_{t+n})$  for

some arbitrary  $t$  so  $\mathbb{P}(\tilde{M}_n \leq x) = \prod_{i=1}^n \mathbb{P}(\tilde{X}_{t+i} \leq x) = F^n(x)$  and examine the connection between  $M_n$  and  $\tilde{M}_n$ . We normalise using a sequence of constants  $(a_n > 0, b_n)$  and consider the limiting distribution of  $\tilde{Z}_n = (\tilde{M}_n - b_n)/a_n$ . If such a sequence of normalising constants,  $a_n > 0$  and  $b_n$ , exists such that the distribution of  $\tilde{Z}$  in the limit is non-degenerate then

$$\mathbb{P}(\tilde{M}_n \leq a_n z + b_n) = F^n(a_n z + b_n) \rightarrow \tilde{G}(z)$$

and the limit distribution of  $\tilde{Z}$ ,  $\tilde{G}(z)$ , is a member of the family of Generalized Extreme Value (GEV) distributions. We set  $u_n = u_n(z) = a_n z + b_n$  for the remainder of §7.1.

To obtain a similar result for the dependent sequence  $\{X_t\}_{t \geq 1}$  we restrict the dependence structure by assuming that extreme events are near-independent if they are far enough apart in time, which is a reasonable assumption for many physical processes. Leadbetter (1983a) presents this condition of asymptotic independence at long ranges as the  $D(u_n)$  condition. Consider two sets of indices  $I_1 = \{i_1, \dots, i_p\}$  and  $I_2 = \{j_1, \dots, j_q\}$  such that  $1 \leq i_1 < \dots < i_p < j_1 \dots < j_q \leq n$  and  $j_1 \geq i_p + l$ . Let  $I = I_1 \cup I_2$  and  $M_{\{I\}} = \max(X_i : i \in I)$  then

$$|\mathbb{P}(M_{\{I\}} \leq u_n) - \mathbb{P}(M_{\{I_1\}} \leq u_n) \mathbb{P}(M_{\{I_2\}} \leq u_n)| < \alpha_{n,l} \quad (D(u_n) \text{ condition}) \quad (7.1.1)$$

and  $\alpha_{n,l_n} \rightarrow 0$  as  $n \rightarrow \infty$  for some  $\{l_n\}$ ,  $l_n = o(n)$ , *i.e.*,  $l_n/n \rightarrow 0$  as  $n \rightarrow \infty$ . This condition essentially ensures that the maxima over  $I_1$  and  $I_2$  become independent as the  $n$  increases.

If condition (7.1.1) holds for all  $z$  such that  $\tilde{G}(z) > 0$  and the limit,  $G(z)$ , exists then

$$\mathbb{P}(M_n \leq a_n z + b_n) \rightarrow G(z) = \tilde{G}^\theta(z), \quad (7.1.2)$$

where the constant  $\theta$  ( $0 \leq \theta \leq 1$ ) is the extremal index and  $G(z)$  is a GEV distribution with the same shape parameter as  $\tilde{G}(z)$ . The extremal index lies between 0 and 1 with dependence at extreme levels increasing as  $\theta$  decreases. An extremal index of 1 is indicative of independence at asymptotically high levels (since then the limit is the same as in the independent case), however, there may still be dependence at extreme but non-limit levels, *e.g.*,  $\theta = 1$  for all Gaussian processes but such processes can exhibit any correlation,  $\rho$ , with  $|\rho| < 1$ , at lag 1.

The extremal index is also shown to have the following definition from O'Brien (1987) for  $t \geq 1$ :

$$\theta = \lim_{n \rightarrow \infty} \mathbb{P}(M_{r_n} \leq u_n | X_t > u_n), \quad (7.1.3)$$

where  $r_n = o(n)$ , that is  $r_n/n \rightarrow 0$  as  $n \rightarrow \infty$  and  $M_{r_n} = \max(X_{t+1}, \dots, X_{t+r_n})$ .

### 7.1.1 Cluster size distribution and point process formulation

The more dependence there is between consecutive values, *i.e.*, the smaller  $\theta$  is, the more the process tends to cluster so that extremes occur at similar times. In fact, the extremal index is asymptotically equal to the reciprocal of the mean cluster size as the threshold tends to the upper end point (Leadbetter, 1983b). The observations within a cluster are thought of as being part of one extreme event.

Let  $K$  be the cluster size. The cluster size distribution is defined for  $k \in \mathbb{Z}^+$  as

$$\pi(k) := \mathbb{P}(K = k | K > 0) = \lim_{n \rightarrow \infty} \mathbb{P} \left( \sum_{i=1}^{r_n} \mathbb{1}\{X_{t+i} > u_n\} = k \mid \sum_{i=1}^{r_n} \mathbb{1}\{X_{t+i} > u_n\} > 0 \right). \quad (7.1.4)$$

It is assumed that as  $n \rightarrow \infty$  the exceedances in a block of length  $r_n$  belong to the same cluster so (7.1.4) gives the probability of  $k$  exceedances in a block given that there is at least one exceedance, *i.e.*, that there is a cluster.

We can also define a point process count,  $N_n$ , on a scaled time axis  $[0, 1]$ , which counts the exceedances of the threshold,  $u_n(z)$ :

$$\begin{aligned} N_n([0, 1] \times [z, z^U]) &= \sum_{i=1}^n \mathbb{1} \left\{ \frac{i}{n} \in [0, 1] \right\} \mathbb{1}\{X_{t+i} > u_n(z)\} \\ &= \sum_{i=1}^n \mathbb{1} \left\{ \frac{i}{n} \in [0, 1] \right\} \mathbb{1} \left\{ \frac{X_{t+i} - b_n}{a_n} \in (z, \infty) \right\}, \end{aligned} \quad (7.1.5)$$

where  $z^U$  is the upper end point of  $G(z)$ . We can rewrite the cluster size distribution, (7.1.4), in terms of this point process:

$$\pi(k) = \lim_{n \rightarrow \infty} \mathbb{P} (N_n([0, r_n/n] \times [z, z^U]) = k \mid N_n([0, r_n/n] \times [z, z^U]) > 0). \quad (7.1.6)$$

**Definition 7.1.1.1.** *A marked point process is a sequence of pairs of random variables, one following a point process and the other a randomly assigned ‘mark’ or ‘size’ associated to each point.*

A *compound Poisson process* arises as the sum of the ‘marks’ over the counting point process in a marked point process:

**Definition 7.1.1.2.** *A compound Poisson process is a stochastic process where the time of events are determined by a Poisson process and the ‘size’ of each event is independent and identically distributed; we call this distribution of sizes the ‘mark’ distribution or the ‘multiplicity’.*

Assume that there exists  $a_n > 0$  and  $b_n$  such that (7.1.2) holds with  $G$  non-degenerate. Then Hsing et al. (1988) show under mild mixing conditions that  $N_n \rightarrow N$ , where  $N$  is a compound Poisson process with intensity measure

$$\Lambda((t_1, t_2) \times [z, \infty)) = -(t_2 - t_1) \log(G(z)) = -\theta(t_2 - t_1) \log(\tilde{G}(z)), \quad (7.1.7)$$

for  $(t_1, t_2) \in [0, 1]$  and ‘mark distribution’ or ‘multiplicity’  $\pi$  as described in Definition 7.1.1.2.

Here, the ‘events’ are the independent clusters and the ‘multiplicity’ is the cluster size distribution. Notice that, in the limit, events occur at single points in time (corresponding to a point in the point process) whereas in practice we will have clusters of exceedances spread over time which belong to the same independent event. The intensity measure is the expected number of independent events in the limit. The expected number of exceedances in the time frame  $[0, 1]$  is the product of the expected number of clusters/events in  $[0, 1]$  and the expected cluster size:

$$\mathbb{E}[N([0, 1] \times [z, \infty))] = \Lambda([0, 1] \times [z, \infty)) \mathbb{E}[K] = -\theta \log(\tilde{G}(z)) \sum_{k=1}^{\infty} k\pi(k).$$

Under very mild conditions, (see, for example, Hsing et al. 1988,)  $\theta = (\sum_{k=1}^{\infty} k\pi(k))^{-1}$ . This relation is assumed true throughout the thesis. An estimate for the extremal index can therefore be constructed as the inverse empirical mean of cluster sizes.

### 7.1.2 Measures of extremal serial dependence

As we wish to measure the serial dependence of extreme values it is informative to consider the probability of an observation being extreme given that an observation earlier in the sequence

was extreme. A measure of extremal serial dependence for the process lag  $\tau$  apart is

$$\chi^{(\tau)} = \lim_{x \rightarrow \infty} \mathbb{P}(X_{t+\tau} > x | X_t > x),$$

when  $X_{t+\tau}$  and  $X_t$  have the same margins (*i.e.*, the process is stationary). The notation here follows from the *coefficient of asymptotic dependence*,  $\chi$ , which is a measure of dependence between multiple variables and is explored more later (§7.3.4). Here  $\chi^{(\tau)}$  ranges between 0 and 1 with 0 corresponding to asymptotic independence. For  $\chi^{(\tau)} > 0$  we have the case of asymptotic dependence and  $\chi^{(\tau)}$  gives a measure of the different degrees of dependence for each lag.

Now:

$$\begin{aligned} \mathbb{P}(X_{t+\tau} > x | X_t > x) &= \frac{\mathbb{P}(X_{t+\tau} > x, X_t > x)}{\mathbb{P}(X_t > x)} \\ &= \frac{1 - \mathbb{P}(X_{t+\tau} \leq x) - \mathbb{P}(X_t \leq x) + \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x)}{1 - F(x)} \\ &= \frac{1 - 2F(x) + \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x)}{1 - F(x)} \\ &= 2 - \frac{1 - \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x)}{1 - F(x)} \\ &\sim 2 - \frac{\log \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x)}{\log F(x)} \quad \text{as } x \rightarrow \infty, \end{aligned} \quad (7.1.8)$$

where  $F$  is the common distribution of the stationary sequence  $\{X_t\}_{t \geq 1}$ , so  $\chi^{(\tau)}$  can be written as the limit as  $x \rightarrow \infty$  of a function  $\chi^{(\tau)}(x)$ :

$$\chi^{(\tau)} = \lim_{x \rightarrow \infty} \chi^{(\tau)}(x) = 2 - \frac{\log \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x)}{\log F(x)}. \quad (7.1.9)$$

When we are in the case of asymptotic independence  $\chi^{(\tau)}$  gives no measure of the different degrees of dependence at finite levels. The need for such a measure when we have asymptotic independence led to the development of a dual measure,  $\bar{\chi}^{(\tau)}$ , in the bivariate setting (Coles et al., 1999). This measure ranges between -1 and 1 with  $\bar{\chi}^{(\tau)} = 1$  corresponding to the case of asymptotic dependence and  $\bar{\chi}^{(\tau)} < 1$  for asymptotic independence, with the value  $\bar{\chi}^{(\tau)} < 1$  of informative about the degree of asymptotic independence.

The dual  $\bar{\chi}^{(\tau)}$  can be defined in a similar manner to  $\chi^{(\tau)}$ :

$$\bar{\chi}^{(\tau)} = \lim_{x \rightarrow \infty} \bar{\chi}^{(\tau)}(x),$$

where

$$\begin{aligned}\bar{\chi}^{(\tau)}(x) &= \frac{2 \log \mathbb{P}(X_t > x)}{\log \mathbb{P}(X_t > x, X_{t+\tau} > x)} - 1 \\ &= \frac{2 \log(1 - F(x))}{\log(1 - 2F(x) + \mathbb{P}(X_{t+\tau} \leq x, X_t \leq x))} - 1.\end{aligned}\tag{7.1.10}$$

Together the two measures give a complete description of the extremal dependence: whether we have asymptotic dependence at lag  $\tau$ , ( $\chi^{(\tau)} > 0, \bar{\chi}^{(\tau)} = 1$ ), or asymptotic independence, ( $\chi^{(\tau)} = 0, \bar{\chi}^{(\tau)} < 1$ ), and a measure of the strength of the dependence at finite levels of the process lag  $\tau$  apart.

### 7.1.3 Declustering

We now consider inference for extremes of stationary dependent data. As we saw earlier, the limit distribution of the ‘normalised’ maxima is GEV (7.1.2) and is connected to the independence case through the extremal index,  $\theta$ . If one is only interested in inference for the block maxima of data with short-term dependence, such as meteorological data, then it is sufficient to use the GEV model, estimating parameters using the likelihood with the block maxima observed. The model fit will be less accurate, however, the stronger and longer term the dependence between subsequent observations is.

Inference is more complicated for threshold exceedances and cluster characteristics due to the temporal dependence within clusters. Leadbetter (1991) shows that the limiting distribution of the maximum threshold excesses in each independent cluster is a GPD. This leads to the idea of declustering in order to obtain sets of extreme observations (clusters) which we assume to be independent from one cluster to another and are part of one extreme event. The maxima of each cluster can then be considered independent and the GP distribution can be fitted to these maxima.

A simple declustering technique is so called *runs declustering* (Smith and Weissman, 1994). Once the threshold has been exceeded, subsequent observations above the threshold are taken to belong to the same event/cluster until there are  $r$  consecutive observations which fall below the threshold thus implying that the end of that cluster has been reached. However, the choice of  $r$  can greatly affect the clusters obtained - the familiar problem of bias versus variance is again relevant. If  $r$  is chosen to be too small we can obtain clusters split by only a few observations when in fact they may have truly belonged to the same event, *i.e.*, we are assuming there are too many independent clusters thus leading to bias in the parameter

estimators due to the false independence assumption. On the other hand, taking  $r$  too large can result in the number of clusters being underestimated reducing the number of cluster maxima and increasing the variance of the estimates. An illustration of this method with  $r = 4$  is shown in Figure 7.1.1 leading to the identification of 2 clusters. Note that if we instead chose  $r < 4$  or  $r > 4$  we would obtain 4 and 1 clusters respectively.

Clearly there is much room for improvement on this method, in particular note that for the consecutive observations below the threshold there is no measure of their distance below the threshold. One could imagine a situation where following some exceedances there are a set of observations only just below the threshold before crossing the threshold again, in which case these observations could arguably be part of the same event. Furthermore, we could have the situation where we have a small number of observations between exceedances which are far below the threshold and so are conceivably not part of the same event. This incorporation of the trajectory of the process is considered in Laurini and Tawn (2003). The idea commonly used by hydrologists is to include a second lower threshold; in the river setting this could be the base river level. If observations are still above this level after dropping below the higher threshold they may contribute to further flooding and so can be considered as part of the same event. Laurini and Tawn (2003) introduce a declustering method which combines runs declustering and this lower chosen threshold,  $v$ . This improvement on the runs declustering method reduces the sensitivity to the choice of run length although it is also sensitive to the choice of the lower threshold. Laurini and Tawn's method is applied with  $r = 4$  and  $v = 10$  in Figure 7.1.1 right panel; three clusters are now identified since the low observation at index 15 leads to a reasonable split in the large cluster identified using the runs method.

Ferro and Segers (2003) introduced a declustering method which doesn't rely on such arbitrary cluster identification parameters. The idea is to consider the limiting distribution as the threshold increases (towards the upper end point of the common distribution,  $F$ ) of the time between threshold exceedances. The time between events, or inter-cluster times, are exponentially distributed asymptotically whereas within events the times between exceedances will not be exponentially distributed. Hence the limiting distribution of the inter-exceedance times is a mixture distribution of an exponential distribution and a point mass at zero if time is scaled to  $[0, 1]$  in the limit. Ferro and Segers (2003) show that the proportion of inter-cluster (non-zero interexceedance) times is given by the extremal index,  $\theta$  and also that the mean of the exponential distribution of inter-cluster times is the inverse of this ( $\theta^{-1}$ ).

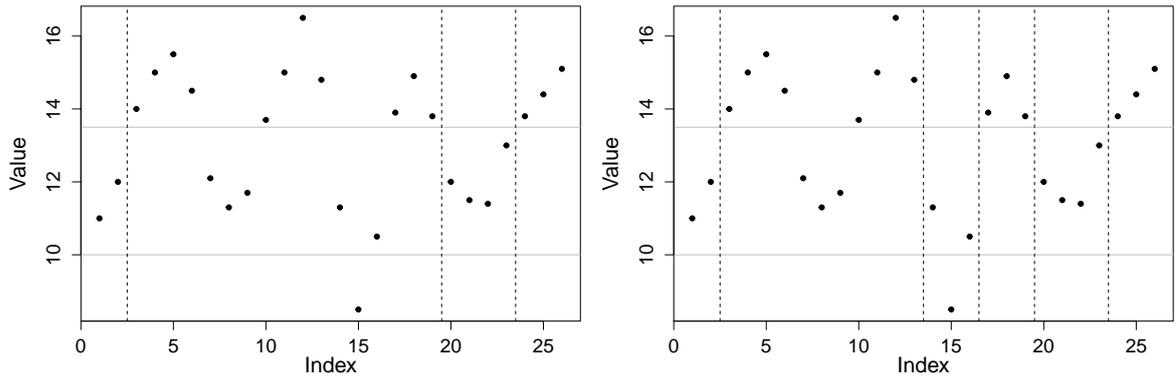


Figure 7.1.1: Illustration of declustering methods with threshold 13.5. Runs declustering with  $r = 4$ , resulting in 2 clusters (left) and Laurini and Tawn’s method with  $v = 10$  (lower threshold) and  $r = 4$ , resulting in 3 clusters (right).

That is, the probability density function of the inter-exceedance times is:

$$(1 - \theta)\delta(t) + \theta^2 \exp(-\theta t) \quad t \geq 0, \tag{7.1.11}$$

where  $\delta(t)$  is the Dirac delta function<sup>1</sup>. So if we plot the exponential quantiles against the inter-exceedance times we should find that it is non linear for small inter-exceedance times (corresponding to intra-cluster, *i.e.*, within cluster, times) and then becomes linear with a gradient equal to  $\theta^{-1}$  (corresponding to inter-cluster times).

The extremal index can be estimated based on the moments of the limiting distribution (7.1.11) and Ferro and Segers (2003) refer to this estimator as the *intervals estimator*:

$$\hat{\theta}(u) = \begin{cases} \min \left( 1, \frac{2(\sum_{i=1}^{n_u-1} (t_i-1))^2}{(n_u-1) \sum_{i=1}^{n_u-1} (t_i-1)(t_i-2)} \right) & \text{if } \max\{t_i : 1 \leq i \leq n_u - 1\} > 2 \\ \min \left( 1, \frac{2(\sum_{i=1}^{n_u-1} t_i)^2}{(n_u-1) \sum_{i=1}^{n_u-1} t_i^2} \right) & \text{if } \max\{t_i : 1 \leq i \leq n_u - 1\} \leq 2, \end{cases} \tag{7.1.12}$$

where  $n_u$  is the number of exceedances of the threshold  $u$  and  $t_i$  are the interexceedance times. Ferro and Segers (2003) show in a simulation study that this estimator is robust to both the threshold and the ‘true’ extremal index. The estimate,  $\hat{\theta}$ , can then be used as part of a declustering scheme since the extremal index is the proportion of inter-cluster times; it can be assumed that the largest  $\lfloor \theta n_u \rfloor + 1$  interexceedance times, where  $n_u$  are the

<sup>1</sup>The Dirac delta function can be properly defined as a measure. Informally,  $\delta(t) = \infty$  if  $t = 0$  and is 0 otherwise and  $\int_{-\infty}^{\infty} \delta(t)dt = 1$ .

number of exceedances observed, are approximately independent inter-cluster times. Note other declustering schemes estimate  $\theta$  after declustering whereas the intervals estimate for  $\theta$  requires no declustering, rather it is used to decluster. Ferro and Segers (2003) also describe a bootstrap method for estimating the uncertainty of  $\hat{\theta}$  and functionals of the clusters.

## 7.2 The ARMAX process

We now introduce a stationary process, the ARMAX process, for which we can decide the common distribution and the level of serial dependence. We can derive the true extremal index and cluster size distribution for the ARMAX process and use it to test the various declustering methods.

The ARMAX process (Alpuim, 1989) is for  $t \geq 1$ :

$$X_t = \max(cX_{t-1}, \epsilon_t) \quad 0 < c < 1, \quad (7.2.1)$$

where  $X_0 \sim F_0$  for some distribution  $F_0$  independent of  $\epsilon_t$ ,  $\{\epsilon_t\}_{t \geq 1}$  are independent and  $\epsilon_t \sim G$  for some distribution  $G$ . We set  $G(x) = \exp(-(1-c)/x)$  as this ensures unit Fréchet margins for the common distribution  $F$  of  $X$ . Then, for  $t \geq 0$  and  $n \geq 1$ ,

$$X_{t+n} = \max(c^n X_t, \max_{j=1:n} c^{n-j} \epsilon_{t+j}). \quad (7.2.2)$$

So for  $X_n$ ,

$$F_n(x) := \mathbb{P}(X_n \leq x) = \mathbb{P}\left(X_0 \leq \frac{x}{c^n} \prod_{j=1}^n \mathbb{P}\left(\epsilon_j \leq \frac{x}{c^{n-j}}\right)\right) = F_0\left(\frac{x}{c^n}\right) \prod_{j=0}^{n-1} G\left(\frac{x}{c^j}\right)$$

and the distribution of the stationary univariate ARMAX process is

$$F(x) = \lim_{n \rightarrow \infty} F_n(x) = \prod_{j=0}^{\infty} G\left(\frac{x}{c^j}\right) \quad (7.2.3)$$

$$\begin{aligned} &= \prod_{j=0}^{\infty} \exp\left(-\frac{c^j(1-c)}{x}\right) \\ &= \exp\left(-\frac{1}{x}\right) \quad \text{for } x > 0, \end{aligned} \quad (7.2.4)$$

that is,  $F$  is unit Fréchet. To see the stationarity, note that from (7.2.1),

$$F_t(x) = F_{t-1}\left(\frac{x}{c}\right)G(x) \quad \text{and} \quad F(x) = F\left(\frac{x}{c}\right)G(x). \quad (7.2.5)$$

The above relations are useful in later calculations.

The limit distribution  $F$  is *max-stable*:  $F^n(a_nx + b_n) = F(x)$ , when we choose  $a_n = n$ ,  $b_n = 0$ . Therefore  $F$  is also the limiting distribution of  $\tilde{M}_n := \max(\tilde{X}_{t+1}, \dots, \tilde{X}_{t+n})$  where  $\{\tilde{X}_t\}_{t \geq 1}$  are i.i.d. with the same marginals as  $\{X_t\}_{t \geq 1}$ .

### 7.2.1 Extremal Index

Now let  $M_n := \max(X_{t+1}, \dots, X_{t+n}) = \max(cX_t, \epsilon_{t+1}, \dots, \epsilon_{t+n})$  by (7.2.2). We have

$$\mathbb{P}(M_n \leq nx) = F_t\left(\frac{nx}{c}\right)G^n(nx) \xrightarrow{n \rightarrow \infty} G(x),$$

since  $G$  is max-stable. The  $D(u_n)$  condition holds since  $\{X_t\}_{t \geq 1}$  is a Markov chain, *i.e.*, each state depends only on the previous state (see Asmussen 1987), so we can use (7.1.2) to deduce that  $G(x) = F^\theta(x)$ . So, since we set  $G(x) = \exp(-(1-c)/x)$  and  $F$  is unit Fréchet we have  $\theta = 1 - c$ .

Alternatively, we can find  $\theta$  using O'Brien's formulation (7.1.3):

$$\begin{aligned} \theta &= \lim_{n \rightarrow \infty} \mathbb{P}(X_{t+1} \leq u_n, \dots, X_{t+r_n} \leq u_n | X_t > u_n) \\ &= \lim_{n \rightarrow \infty} \frac{\mathbb{P}(\epsilon_{t+1} < u_n, \dots, \epsilon_{t+r_n} < u_n, u_n < X_t < \frac{u_n}{c})}{\mathbb{P}(X_t > u_n)} \quad (\text{using (7.2.2)}) \\ &= \lim_{n \rightarrow \infty} G^{r_n}(u_n) \frac{F\left(\frac{u_n}{c}\right) - F(u_n)}{1 - F(u_n)} \\ &= \lim_{n \rightarrow \infty} G^{r_n}(u_n) F\left(\frac{u_n}{c}\right) \frac{1 - G(u_n)}{1 - F(u_n)} \quad (\text{using (7.2.5)}). \end{aligned} \quad (7.2.6)$$

Now  $G^{r_n}(u_n) = \exp(-r_n(1-c)/nx) \rightarrow 1$  as  $n \rightarrow \infty$  since  $r_n = o(n)$ . Also since  $F$  and  $G$  are both max stable,  $n(1 - F(u_n)) \xrightarrow{n \rightarrow \infty} -\log(F(x))$  and  $n(1 - G(u_n)) \xrightarrow{n \rightarrow \infty} -\log(G(x))$ . So (7.2.6) becomes

$$\theta = \frac{-\log(G(x))}{-\log(F(x))} = \frac{1/x}{1/(x(1-c))} = 1 - c.$$

## 7.2.2 Cluster size distribution

For the stationary ARMAX process a set of consecutive exceedances begins with an independent innovation, *i.e.*, an  $\epsilon$  exceeding  $u_n$ . In fact there can only be one such exceedance by  $\{\epsilon_t\}$  of  $u_n$  in a cluster as the probability of an independent innovation is  $1 - G(u_n)$  which is  $O(\frac{1}{n})$  and in a run of length  $r_n$  there are  $r_n$  places this could occur. Thus the probability of an independent innovation somewhere in the run is  $r_n O(\frac{1}{n})$  which tends to 0 as  $n \rightarrow \infty$ . Therefore we can just consider the exceedances in the run to be consecutive with no independent innovation after  $X_t = \epsilon_t$ , *i.e.*, for each  $j \in 1 : r_n$   $X_{t+j} = cX_{t+j-1}$ , so (7.1.4) becomes

$$\begin{aligned}
 \pi(k) &= \lim_{n \rightarrow \infty} \mathbb{P}(X_{t+1} > u_n, \dots, X_{t+k-1} > u_n, X_{t+k} \leq u_n | X_t > u_n) \\
 &= \lim_{n \rightarrow \infty} \mathbb{P}\left(cX_t > u_n, \dots, c^{k-1}X_t > u_n, c^k X_t \leq u_n | X_t > u_n\right) \\
 &= \lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{u_n}{c^{k-1}} < X_t \leq \frac{u_n}{c^k} | X_t > u_n\right) \\
 &= \lim_{n \rightarrow \infty} \frac{\mathbb{P}\left(\frac{u_n}{c^{k-1}} < X_t \leq \frac{u_n}{c^k}\right)}{\mathbb{P}(X_t > u_n)} \\
 &= \lim_{n \rightarrow \infty} \frac{F\left(\frac{u_n}{c^k}\right) - F\left(\frac{u_n}{c^{k-1}}\right)}{1 - F(u_n)} \\
 &= \lim_{n \rightarrow \infty} \frac{e^{-\frac{c^k}{nx}} - e^{-\frac{c^{k-1}}{nx}}}{1 - e^{-\frac{1}{nx}}} \\
 &= c^{k-1}(1 - c) \\
 &= \theta(1 - \theta)^{k-1}.
 \end{aligned}$$

Note as  $c$  tends to 0 we get closer to independence and so the probability of a cluster of size 1 tends to 1. The cluster size distribution is Geometric with probability of ‘success’,  $\theta = \lim_{n \rightarrow \infty} \mathbb{P}(X_{t+1} \leq u_n | X_t > u_n)$ , being the probability that given an exceedance the next ‘observation’ is a non-exceedance. This probability is independent of previous values of  $\{X_i\}_{i=1}^{t-1}$  given the current exceedance and so the distribution is memoryless.

As noted earlier, we can confirm that the inverse of the limiting mean cluster size is the extremal index (Leadbetter, 1983b):

$$\left(\sum_{k=1}^{\infty} k\pi(k)\right)^{-1} = \left((1 - c) \sum_{k=1}^{\infty} kc^{k-1}\right)^{-1} = \left(\frac{1 - c}{(1 - c)^2}\right)^{-1} = 1 - c = \theta. \quad (7.2.7)$$

### 7.2.3 Measures of serial dependence

We now derive the joint distribution of  $X_t$  and  $X_{t+\tau}$ , *i.e.*, the process lag  $\tau$  apart, for an ARMAX sequence with unit Fréchet common distribution: for  $x_0 > 0$  and  $x_\tau > 0$ ,

$$\begin{aligned}
 F_{t,t+\tau}(x_0, x_\tau) &= \mathbb{P}\left(X_t \leq x_0, X_{t+\tau} \leq \frac{x_\tau}{c^\tau}, \epsilon_{t+1} \leq \frac{x_\tau}{c^{\tau-1}}, \dots, \epsilon_{t+\tau} \leq x_\tau\right) \\
 &= \mathbb{P}\left(X_t \leq \min(x_0, \frac{x_\tau}{c^\tau})\right) \prod_{j=1}^{\tau} \mathbb{P}\left(\epsilon_{t+j} \leq \frac{x_\tau}{c^{\tau-j}}\right) \\
 &= \exp\left(-\frac{1}{\min(x_0, \frac{x_\tau}{c^\tau})} - \frac{1-c}{x_\tau} \sum_{j=1}^{\tau} c^{\tau-j}\right) \\
 &= \exp\left(-\max\left(\frac{1}{x_0}, \frac{c^\tau}{x_\tau}\right) - \frac{c^\tau - 1}{x_\tau}\right) \\
 &= \exp\left(-\max\left(\frac{1}{x_0} - \frac{c^\tau - 1}{x_\tau}, \frac{1}{x_\tau}\right)\right). \tag{7.2.8}
 \end{aligned}$$

So, using (7.1.9),

$$\begin{aligned}
 \chi^{(\tau)} &= \lim_{x \rightarrow \infty} \left(2 - \frac{\log(F_{t,t+\tau}(x, x))}{\log(\exp(-\frac{1}{x}))}\right) \\
 &= \lim_{x \rightarrow \infty} \left(2 - \frac{\max\left(\frac{1}{x} - \frac{c^\tau - 1}{x}, \frac{1}{x}\right)}{\frac{1}{x}}\right) \\
 &= \lim_{x \rightarrow \infty} (2 - \max(2 - c^\tau, 1)) = c^\tau, \tag{7.2.9}
 \end{aligned}$$

where  $0 < c^\tau < 1$ , so  $\chi^{(\tau)}$  decays geometrically.

Similarly using (7.1.10) we find  $\bar{\chi}_\tau = 1$ , that is we have asymptotic dependence with the strength of dependence at lag  $\tau$  being  $c^\tau$ . Moreover, the joint distribution (7.2.8) is a bivariate extreme value distribution (see 7.3.4) and, as stated in Coles et al. (1999), all bivariate extreme value distributions have  $\bar{\chi} = 1$  as they are asymptotically dependent.

### 7.2.4 Simulation and declustering

We now consider the performance of the declustering methods of §7.1.3 for a range of ARMAX sequences with unit Fréchet common distribution,  $F$ . First, Figure 7.2.1 gives an example ARMAX sequence with  $c = 0.8$  - for this example the extremal index is 0.2 and thus the limiting mean cluster size is 5. According to the limiting cluster size distribution there is less than 0.33 probability of a cluster size greater than 5. Figure 7.2.1 shows the result of runs declustering with  $r = 10$  (top) and intervals declustering (bottom) for exceedances of

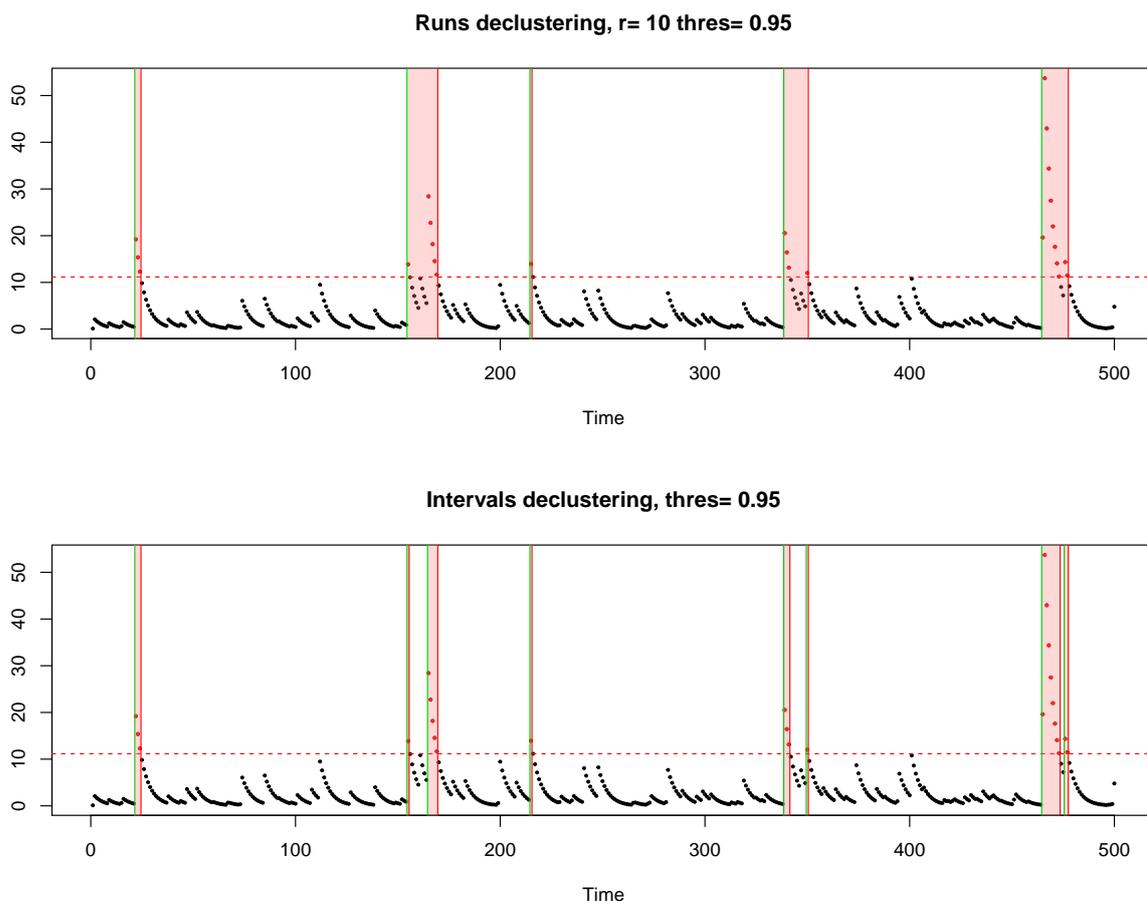


Figure 7.2.1: Simulated ARMAX sequence with  $c = 0.8$  and unit Fréchet common distribution declustered using runs declustering with  $r = 10$  (top) and Ferro and Seger's intervals declustering (bottom).

the 95% quantile of the data. When applied to an ARMAX sequence the intervals method almost always chooses clusters to end when the series dips below the threshold - even when this is only one value just below the threshold. Sequences for which this is not the case are rare for all  $c$  but are more common the smaller  $c$  is, so for the ARMAX process the intervals method is generally equivalent to the runs method with run length 1. In contrast the runs declustering method with  $r > 1$  would group two such sets of exceedances into one cluster. Of course in this setting it makes sense that the next upward jump in the process is the start of a new cluster.

We simulated 100 ARMAX sequences with unit Fréchet common distribution,  $F$ , for each  $c \in (0, 0.2, 0.4, 0.6, 0.8)$  and applied the the runs, intervals and Laurini and Tawn's declustering methods to the simulated sequences. From the counts of the cluster sizes one can obtain an empirical estimate of the cluster size distribution; an estimate for the extremal

index,  $\theta$ , can then be obtained using this empirical distribution in place of  $\pi(k)$  in (7.2.7). An equivalent estimate,  $\hat{\theta}$ , is the number of clusters divided by the the total number of exceedances.

The relative bias of the empirical  $\theta$  estimators for all three declustering methods and a range of run lengths as well as the intervals estimator (7.1.12) are shown in Figure E.2.1. The intervals estimator has close to zero bias - in fact the estimator does have zero first-order bias (Ferro and Segers, 2003) - whereas the empirical estimators from all three declustering methods have positive bias; this bias is smaller the larger  $c$  is, *i.e.*, the more serial dependence there is. The runs estimator for  $\theta$  has the largest bias followed by Laurini and Tawn's method with a lower threshold set to the 80% quantile of the simulated values; the bias is reduced for both of these methods the smaller the run length,  $r$ , is, particularly when  $c$  is small.

The relative RMSE is slightly harder to interpret (Figure E.2.2). In most cases the relative RMSE increases as  $c$  increases with the intervals/runs estimator having the smallest RMSE for small/large  $c$ . The most notable exception is for the runs declustering with  $r > 1$ ; for large  $c$  the RMSE is similar and increasing for all run lengths whereas for small  $c$  the relative RMSE is larger the larger the run length is and is greatest for  $c = 0$ , *i.e.*, when we have an independent sequence of random variables.

Recall that the intervals estimator is based on the limiting distribution of (normalised) interexceedance times which is a mixture distribution of a point mass at 0 and a standard exponential distribution (7.1.11). In Appendix E Figure E.2.3 shows example of a diagnostic plot comparing standard exponential quantiles and normalised interexceedance times; the vertical line corresponds to the  $(1 - \hat{\theta})$  quantile, where  $\hat{\theta}$  is the intervals estimator, and the diagonal line has gradient  $1/\hat{\theta}$ .

### 7.3 Multivariate Extreme Value Theory

Multivariate extremes occur naturally in many environmental settings as often multiple physical processes are linked to each other (for example, rainfall and wind speed) or there is interest in a physical process at multiple locations, such as river flows on a network of river gauges. In the latter case there will be some interaction between observations at different sites; the occurrences of extreme values at multiple locations simultaneously is of interest as well as determining which occurrences 'belong' to the same 'extreme event'. In this section we con-

concentrate on independent multivariate sequences of random vectors and return to sequences with serial dependence in §8.1.

In §7.3.1 we describe the componentwise maxima approach to modelling multivariate extremes and in §7.3.3 we discuss how the copula function encompasses the dependence structure of a multivariate distribution. Measures of dependence are then described in §7.3.4. Finally, the multivariate extension of the point process representation is discussed in §7.3.5.

In what follows  $\{\mathbf{X}_t\}_{t \geq 1}$  is a sequence of independent  $d$ -dimensional random vectors. Each dimension corresponds to a physical process and, as in previous sections, the index  $t$  corresponds to time. It is convenient for much of the following theory to transform the components such that they follow a unit Fréchet distribution (so  $\mathbb{P}(X_j \leq x) = \exp(-1/x)$  for  $x > 0$ ) - this can be easily done using the probability integral transform.

### 7.3.1 Componentwise maxima approach

Maxima are less straightforward to define in the multivariate setting than in the univariate setting as there is no natural ordering. The classical approach discussed here is to consider the componentwise maxima:

$$\mathbf{M}_n = (M_{n,1}, \dots, M_{n,d}) = \left( \max_{i=1, \dots, n} X_{t+i,1}, \dots, \max_{i=1, \dots, n} X_{t+i,d} \right). \quad (7.3.1)$$

However, with this approach the maxima of one component does not necessarily occur at the same time as the maxima of another so the vector of componentwise maxima is not necessarily an observed vector of values.

We define the marginal distributions of the normalised maxima for each component  $j \in (1, \dots, d)$  as

$$Z_j = \lim_{n \rightarrow \infty} Z_{n,j} = \lim_{n \rightarrow \infty} \frac{M_{n,j} - b_{n,j}}{a_{n,j}},$$

for some sequences  $a_{n,j} > 0$  and  $b_{n,j}$  for  $j = 1, \dots, d$ . We know from the univariate theory that, if they exist, these marginal distributions for the maxima of each component must be GEV. When the components are unit Fréchet distributed we can choose the normalising constants  $a_{n,j} = n$  and  $b_{n,j} = 0$  then the normalised maxima are also unit Fréchet:

$$\mathbb{P}(Z_{n,j} \leq z) = \mathbb{P}(M_{n,j} \leq nz) = \left[ \exp\left(-\frac{1}{nz}\right) \right]^n = \exp\left(-\frac{1}{z}\right).$$

We now introduce the family of *multivariate extreme value distributions* as the limit distribution of the normalised maxima with unit Fréchet marginal distributions.

**Theorem 7.3.1.1** (Pickands (1981)). *Let  $\{(\mathbf{X}_{t+i})_{i=1,\dots,n}$  be a sequence of independent  $d$ -dimensional random variables with standard Fréchet marginal distributions and define  $Z_{t,j} = M_{n,j}/n$  where  $M_{n,j}$  are the component-wise maxima (7.3.1). Let  $\mathbf{z} = (z_1, \dots, z_d)$ . If*

$$\lim_{n \rightarrow \infty} \mathbb{P}(Z_{n,1} \leq z_1, \dots, Z_{n,d} \leq z_d) = G(\mathbf{z}),$$

where  $G$  is a non-degenerate distribution function in each margin, then  $G$  is of the form

$$G(\mathbf{z}) = \exp(-V(\mathbf{z})), \quad z_1 > 0, \dots, z_d > 0, \quad (7.3.2)$$

with

$$V(\mathbf{z}) = \int_{S_d} \max_{j=1,\dots,d} \left( \frac{w_j}{z_j} \right) dH(\mathbf{w}), \quad (7.3.3)$$

where  $S_d$  is the unit simplex and  $H(\mathbf{w})$  is the spectral measure of dimension  $d-1$  satisfying

$$\int_{S_d} dH(\mathbf{w}) = d \quad \text{and} \quad \int_{S_d} w_j dH(\mathbf{w}) = 1 \quad \text{for } j = 1, \dots, d-1. \quad (7.3.4)$$

A unit simplex is basically a generalisation of a triangle (the unit simplex in 2 dimensions) to other dimensions, that is, the  $d$ -dimensional unit simplex is the convex hull of  $d+1$  vertices. In the bivariate case,  $(Z_1, Z_2)$ , the simplex is simply the line segment  $[0, 1]$  so the integrals above are then 1-dimensional from 0 to 1.

The *spectral measure*,  $H$ , can be thought of as a distribution function on  $S_d$  with mean  $\frac{1}{d}$  if normed to give  $\frac{H}{d}$ . This concept is easier to visualise in the bivariate case since then  $\frac{H}{2}$  is the set of distribution functions on  $[0, 1]$  satisfying the moment constraint (7.3.4) with  $d = 2$ . This set includes differentiable distribution functions for which  $dH(w)$  simply becomes  $h(w)dw$  and also non-differentiable functions, for example  $H$  could place a point mass of a half at 0 and at 1. Insight into what  $w$  and  $H(w)$  represent is gained when we consider the point process representation for multivariate extremes in §7.3.5.

Note that we obtain the unit Fréchet marginal distributions if we set all but one component

equal to infinity in the *exponent measure*,  $V$ , for example:

$$\begin{aligned} G(z_1, \infty, \dots) &= \exp(-V(z_1, \infty, \dots)) = \exp\left(-\int_{S_d} \max\left(\frac{w_1}{z_1}, \frac{w_2}{\infty}, \dots\right) dH(\mathbf{w})\right) \\ &= \exp\left(-\frac{1}{z_1} \int_{S_d} w_1 dH(\mathbf{w})\right) \quad (\text{for } z_1 > 0) \\ &= \exp\left(-\frac{1}{z_1}\right). \end{aligned}$$

The max-stability property (2.1.3) is also satisfied here since  $V$  is homogeneous of order  $-1$ :

$$G^n(n\mathbf{z}) = \exp(-nV(n\mathbf{z})) = \exp(-V(\mathbf{z})) = G(\mathbf{z}).$$

### 7.3.2 Dependence Structure

We can obtain various dependence structures from different choices of the spectral measure  $H$ . Independence is obtained if the spectral measure is chosen such that  $V(\mathbf{z}) = \frac{1}{z_1} + \dots + \frac{1}{z_d}$ , since then  $G(\mathbf{z})$  can be factorised. This occurs when  $H$  places point masses at the vertices of the simplex and nowhere else. Complete dependence is obtained when the spectral measure only places a point mass in the centre of the simplex, then  $V(\mathbf{z}) = \max(\frac{1}{z_1}, \dots, \frac{1}{z_d})$ .

There are (infinitely) many possibilities for the spectral measure,  $H$ , despite being subject to the moment constraint. The drawback of this is that there is no finite parametrisation for the limit distributions. One approach is to find parametric models/subfamilies which can approximate the entire family of distributions for  $G$ . An example of such a family is the logistic model (Gumbel, 1960) for which

$$V(\mathbf{z}) = \left(\sum_{j=1}^d z_j^{-\frac{1}{\alpha}}\right)^\alpha \quad \text{for } 0 < \alpha \leq 1. \quad (7.3.5)$$

For this model the dependence structure is symmetric (the variables are exchangeable), with dependence determined by the parameter  $\alpha$ . When  $\alpha \rightarrow 0$  we have complete dependence and when  $\alpha \rightarrow 1$  independence. There are also models, such as the bilogistic model (Joe et al., 1992), which allow asymmetry in the dependence structure. A range of other parametric models can be found in Kotz and Nadarajah (2000).

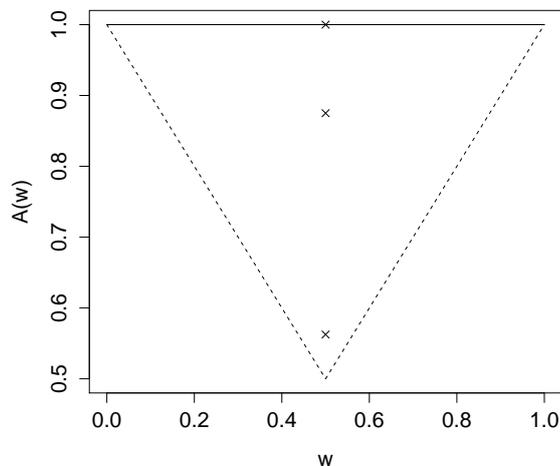


Figure 7.3.1: The space of possible Pickands' dependence functions. The crosses correspond to the points the function must pass through when  $\theta = 2$ ,  $7/8$  and  $9/16$ .

### Pickands' dependence function

In the bivariate setting we can derive models by choosing a suitable function,  $A(w)$   $w \in [0, 1]$ , related to the spectral measure. The function  $A(w)$  is termed the *Pickands' dependence function* (Pickands, 1981) and is linked to the exponent measure as

$$V(z_1, z_2) = \left( \frac{1}{z_1} + \frac{1}{z_2} \right) A \left( \frac{z_1}{z_1 + z_2} \right).$$

To satisfy the marginal conditions, *i.e.*, to obtain the standard Fréchet distribution when we set  $z_1$  or  $z_2$  to  $\infty$  we must have  $A(0) = A(1) = 1$ . The function  $A(w)$  must also be convex with  $\max(w, 1 - w) \leq A(w) \leq 1$  and so the function is restricted to the dashed triangular region shown in Figure 7.3.1. Nonetheless there are infinite possibilities for this function within this space. The upper bound on  $A(w)$  corresponds to independence between the variables whereas the lower bound corresponds to complete dependence.

Pickands (1981) showed that the functions  $H$  and  $A$  are connected in the following way:

$$A'(w) = H(w) - 1 \quad A''(w) = H'(w) \equiv h(w),$$

at points where these functions are differentiable. Coles and Tawn (1991) showed that given the exponent measure we can find the spectral measure through consideration of the deriva-

tives of  $V$ :

$$V_{1,\dots,d}(z) = -\frac{1}{(\sum_i z_i)^{d+1}} h\left(\frac{z}{\sum_i z_i}\right),$$

when no mass exists on the boundaries of  $H$ .

We can estimate  $A(w)$  for a range of  $w$  by considering the empirical distribution of  $T_w = \max\{(1-w)Z_1, wZ_2\}$  since

$$\mathbb{P}(T_w < z) = \mathbb{P}\left(Z_1 < \frac{z}{1-w}, Z_2 < \frac{z}{w}\right) = \exp\left(-\frac{A(w)}{z}\right),$$

thus

$$\mathbb{P}(T_w^{-1} > z) = \mathbb{P}\left(T_w < \frac{1}{z}\right) = \exp(-zA(w)).$$

This is the survivor function of the exponential distribution with mean  $\frac{1}{A(w)}$  hence  $A(w)$  can be estimated as the reciprocal of the mean of the  $1/T_w$  observations. However, these estimates for a range of  $w$  may go outside the bounds on  $A(w)$  and/or together they may result in a non-convex estimate of  $A(w)$ . Alternative methods can be used to improve on this estimation, for example, Hall and Tajvidi (2000) present a modified version which ensures the end points are correct ( $A(0) = A(1) = 1$ ) and that the lower bound on  $A(w)$  is satisfied.

### 7.3.3 Copulas

To fully define a multivariate distribution we require both the marginal distributions and the dependence structure of the variables. Obtaining the marginals is simple since we can easily transform between margins using the probability integral transform. In §7.3.1 we transformed the data to Fréchet margins and based subsequent analysis of the model for the componentwise maxima on the assumption of these margins. Fréchet margins are convenient in the componentwise maxima approach since they lead to a nice form for the limiting distribution of the maxima, however, different margins can be more convenient for other approaches *e.g.*, Gumbel margins with the Heffernan and Tawn model (§E.1). Though all these marginals are equivalent, different margins can give different indications of the structures present in the data. For example, Fréchet margins emphasise the extremes whereas in Gumbel margins the relationship *between* variables can be seen more clearly.

The dependence structure of a multivariate distribution can be completely explained by

its *copula*. Sklar's theorem (Sklar, 1959) states that the joint distribution function of the continuous random variables  $X_1, \dots, X_d$  can be written as a unique function of its corresponding marginal distributions,  $F_1, \dots, F_d$ :

$$F(\mathbf{x}) = \mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d) = C(F_1(x_1), \dots, F_d(x_d)).$$

This function,  $C$ , is the copula and it is a distribution function with uniform margins:

$$C(\mathbf{u}) = \mathbb{P}(X_1 \leq F_1^{-1}(u_1), \dots, X_d \leq F_d^{-1}(u_d)) = F(F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)).$$

Copulas can also be written in different margins<sup>2</sup>, in which case each element inside the copula must be transformed by the distribution function of the desired margin. This is in order to ensure the joint distribution function  $F$  retains the correct margins. Let  $F_*$  be the distribution function of the desired common margins then the copula function  $C_*$  in these margins is, for  $\mathbf{x} = (x_1, \dots, x_d)$ ,

$$C_*(\mathbf{x}) = F(F_1^{-1}(F_*(x_1)), \dots, F_d^{-1}(F_*(x_d))),$$

then

$$\begin{aligned} C(F_1(x_1), \dots, F_d(x_d)) &= F(\mathbf{x}) \\ &= \mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d) \\ &= \mathbb{P}(X_1 \leq F_1^{-1}(F_*(F_1^{-1}(F_1(x_1))))), \dots, X_d \leq F_d^{-1}(F_*(F_d^{-1}(F_d(x_d)))))) \\ &= C_*(F_1^{-1}(F_1(x_1)), \dots, F_d^{-1}(F_d(x_d))). \end{aligned}$$

For example, in Fréchet margins we obtain the copula,  $C_F$ , for which

$$F(\mathbf{x}) = C_F \left( -\frac{1}{\log F_1(x_1)}, \dots, -\frac{1}{\log F_d(x_d)} \right),$$

since the inverse distribution function for the Fréchet distribution,  $F_*^{-1}(u)$ , is  $-(\log u)^{-1}$ . In the componentwise maxima approach of §7.3.1 the multivariate (MVE) copula (7.3.2) was

---

<sup>2</sup>Technically copulas have uniform margins but we can create similar functions in different common margins and for convenience we also call these copulas.

presented in Fréchet margins:

$$G(\mathbf{z}) = C(F_1(z_1), \dots, F_d(z_d)) = C_F^{MVE}(\mathbf{z}) = \exp(-V(\mathbf{z})).$$

It is also useful to consider copulas with Gumbel margins:

$$F(\mathbf{x}) = C_G(-\log(-\log(F_1(x_1))), \dots, -\log(-\log(F_d(x_d)))).$$

Heffernan and Tawn (2004) work in Gumbel margins in their conditional approach to modelling extremes (§E.1).

### 7.3.4 Measures of dependence

In order to gauge the dependence between extremes a number of dependence measures have been developed some of which we discuss in this section.

#### Coefficient of extremal dependence

From the exponent measure we can obtain a measure of the extremal dependence between a set of variables. If we denote the indices of the set of variables by  $C$  then by considering equal margins for the variables in this set and marginalising over the other variables (*i.e.*, setting  $z_j = \infty$  for  $j \notin C$  and  $z_j = z$  for  $j \in C$ ), we find:

$$\begin{aligned} G(\mathbf{z}) &= \exp\left(-z^{-1} \int_{S_d} \max_{i \in C}(w_i) dH(\mathbf{w})\right) \\ &= \exp(-z^{-1})^{\theta_C} \quad \text{with} \quad \theta_C = \int_{S_d} \max_{i \in C}(w_i) dH(\mathbf{w}), \end{aligned}$$

where we have exploited the homogeneity of the exponent function. Then  $\theta_C$  is termed the *coefficient of extremal dependence* for these variables. This measure can be interpreted as the effective number of independent variables in this set. This interpretation is clear if we consider the full set of variables,  $D = \{1, \dots, d\}$ , in which case  $1 \leq \theta_D \leq d$ . When the extremal coefficient equals the number of variables,  $\theta_D = d$ , the joint distribution,  $G(\mathbf{z})$ , clearly factorises into  $d$  unit Fréchet distributions corresponding to independence between the variables. On the other hand,  $\theta_D = 1$  corresponds to complete dependence between the variables. Schlather and Tawn (2003) derived bounds on the extremal coefficient for different sets of variables, with the bounds being a function of extremal coefficients of lower

dimensional sets of variables. This is intuitive since we would expect information about the dependence structure of lower dimensions to provide information about the dependence at higher dimensions.

In the bivariate setting the extremal coefficient is related to Pickands' dependence function (§7.3.2) as  $\theta = V(1, 1) = 2A(\frac{1}{2})$ . Thus the function  $A(w)$  is restricted as it must pass through  $\frac{\theta}{2}$  at  $w = \frac{1}{2}$  while maintaining convexity and this restriction is stronger the larger  $\theta$  is - this is illustrated in Figure 7.3.1.

### Coefficient of asymptotic dependence

To measure the extent to which values of one variable occur with large values of another variable it is informative to consider the probability of variables being extreme given that another variable is extreme. This probability is presented as the limit as the values get more and more extreme and is termed the *coefficient of asymptotic dependence* (Coles et al., 1999):

$$\chi_C = \lim_{u \rightarrow 1} \frac{\mathbb{P}\left(\bigcap_{j \in C} \{X_j > F_j^{-1}(u)\}\right)}{1 - u},$$

where  $C \subset (1, \dots, d)$ . If the joint distribution has common margins,  $F$ , this is simply:

$$\chi_C = \lim_{x \rightarrow \infty} \frac{\mathbb{P}\left(\bigcap_{j \in C} \{X_j > x\}\right)}{1 - F(x)}.$$

This measure is similar to  $\chi^{(\tau)}$  for measuring the dependence in a univariate sequence at lag  $\tau$  (§7.1.2) and has the same properties: for  $\chi_C = 0$  we have asymptotic independence whereas  $\chi_C > 0$  gives a measure of degree of dependence when there is asymptotic dependence but now for all variables in  $C$ . Likewise there is a dual measure  $\bar{\chi}_C$  for measuring the degree of dependence at finite levels in the set  $C$  when there is asymptotic independence.

Restricting to the bivariate setting and omitting subscripts for simplicity and following the same arguments used to derive (7.1.8) we can write  $\chi = \lim_{u \rightarrow 1} \chi(u)$  with

$$\chi(u) = 2 - \frac{\log C(u, u)}{\log u}, \quad (7.3.6)$$

where  $C(u_1, u_2) = \mathbb{P}(X_1 \leq F_1^{-1}(u_1), X_2 \leq F_2^{-1}(u_2))$  is the copula describing the dependence structure of the random variables  $(X_1, X_2)$ .

Similarly the dual  $\bar{\chi}$  is:

$$\bar{\chi} = \lim_{u \rightarrow 1} \bar{\chi}(u) = \lim_{u \rightarrow 1} \frac{2 \log(1-u)}{\log(1-2u+C(u,u))} - 1.$$

A graphical method for determining  $\chi$  and  $\bar{\chi}$  is to plot empirical estimates of  $\chi(u)$  and  $\bar{\chi}(u)$  for increasing threshold  $u$ . For all distributions falling in the class of bivariate extreme values distributions  $\chi(u)$  is constant for all  $u$ . So the graphical method can be used as a diagnostic test to check the suitability of the bivariate extreme value model as in this case the plot should be linear. However, we can only estimate  $\chi(u)$  and  $\bar{\chi}(u)$  for  $u < 1$ , so for distributions where  $\chi(u)$  varies with  $u$  there can be issues of convergence. Coles et al. (1999) give an example of this behaviour for random variables from the bivariate normal distribution, which exhibits asymptotic independence. However, when the correlation coefficient is positive the convergence of  $\chi(u)$  to zero is slow and so in practice will be positive near  $u = 1$ , falsely suggesting asymptotic *dependence*.

The consideration of both of the measures,  $(\chi, \bar{\chi})$ , is important since then we can decide whether or not the data exhibit asymptotic independence with more certainty. Assuming asymptotic dependence when in fact we have asymptotic independence can lead to overestimation of the extreme values. This overestimation can occur when we employ the bivariate extreme value distribution since for all cases of positive dependence these distributions are all asymptotically dependent ( $\chi \neq 0$ ) (Ledford and Tawn, 1996). Thus the bivariate extreme value distribution is not a good model for the extremes when  $\bar{\chi} \ll 1$ .

### Coefficient of tail dependence

A better way of estimating  $\bar{\chi}$  is based on the following asymptotically justified parametric approach for unit Fréchet distributed  $X_j$  (Eastoe and Tawn 2012 extension of Ledford and Tawn 1996)

$$\mathbb{P} \left( \min_{j \in A} (X_j) > x \right) \sim \mathcal{L}_A(x) x^{-\frac{1}{\eta_A}}, \quad (7.3.7)$$

where  $A$  is the set of indices corresponding to the set of variables of interest and  $\mathcal{L}_A(z)$  is a slowly varying function as  $z \rightarrow \infty$ . (A slowly varying function is a function,  $f$ , for which:  $f(xt)/f(t) \rightarrow 1$  as  $t \rightarrow \infty$  for fixed  $x > 0$ ). The *coefficient of tail dependence*,  $\eta_A$ , with  $0 < \eta_A \leq 1$ , is another useful measure of extremal dependence. The value of  $\eta_A$  covers

the four classes of extremal dependence as described by Ledford and Tawn (1997). When  $\eta_A = 1$  the variables corresponding to the set  $A$  are asymptotically dependent. All other cases correspond to asymptotic independence with the value of  $\eta_A$  giving an indication of the nature of the extremal dependence: *positive extremal dependence* when  $(\frac{1}{|A|} < \eta_A < 1)$  indicating that the joint extremes occur more than one would expect if the variables were independent; *negative extremal dependence* when  $(0 < \eta_A < \frac{1}{|A|})$  for extremes occurring less often than one would expect; and *near extremal independence* when  $(\eta_A = \frac{1}{|A|})$  indicating that extremes occur as often as one would expect.

It arises that in the bivariate case  $\bar{\chi} = 2\eta - 1$  (Coles et al., 1999) and so inferences on  $\eta$  automatically lead to inferences for  $\bar{\chi}$ . By considering the variable  $T = \min(X_1, X_2)$  (7.3.7) becomes

$$\mathbb{P}(T > t) = \mathbb{P}(X_1 > t, X_2 > t) \sim \mathcal{L}(t)t^{-\frac{1}{\eta}}. \quad (7.3.8)$$

Then  $\eta$  can be estimated as the shape parameter of variable  $T$  using threshold-based likelihood inference. For  $t > u$  where  $u$  is some large threshold we can approximate (7.3.8) as:

$$\mathbb{P}(T > t) = \frac{c}{t^{\frac{1}{\eta}}} \quad t > u,$$

for some unknown  $c$  and  $\eta$ . The observations which lie below this threshold contribute information to the likelihood as  $\mathbb{P}(T < u)$ . Therefore the likelihood is:

$$L(c, \eta) = \mathbb{P}(T < u)^{n-n_u} \prod_{i=1}^{n_u} f_T(t_i) = \left(1 - \frac{c}{u^{\frac{1}{\eta}}}\right)^{n-n_u} \prod_{i=1}^{n_u} \frac{c}{n} t_i^{1-\frac{1}{\eta}},$$

where  $f_T$  is the density of  $T$  and  $n_u$  are the number of observations of  $T$  above the threshold  $u$ . The maximum likelihood estimate,  $\hat{\eta}$ , is found to be the Hill's estimator (Hill, 1975). Then  $\hat{c}$  is estimated from the proportion of points above the threshold:

$$\frac{n_u}{n} = \frac{\hat{c}}{u^{\frac{1}{\hat{\eta}}}}.$$

### 7.3.5 Point process representation

The point process representation of §2.1.3 extends to the multivariate setting and likewise all multivariate models are unified in this representation. A particularly 'nice' feature of this representation is that it gives insight into the interpretation of the spectral measure,  $H$ ,

which was introduced in §7.3.1.

Let  $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,d})$  for  $i = 1, \dots, n$  be independent and identically distributed random vectors. As described in Coles and Tawn (1991) the marginals,  $Z_j$ , are taken to be unit Fréchet (we can simply transform to ensure this). Then the point process,

$$P_n = \left\{ \frac{\mathbf{X}_i}{n} \right\}_{i=1, \dots, n},$$

converges to the non-homogeneous Poisson point process,  $P$ , on sets bounded away from the origin, *i.e.*, on  $\mathbb{R}_+^d / \{\mathbf{0}\}$ . The origin is excluded since small points in the Poisson process  $P_n$  will tend to zero in the limit resulting in a point mass at zero. By considering pseudo polar co-ordinates we arrive at a useful form for the intensity of the limiting process and an interpretation for the spectral measure,  $H$ . The pseudo polar coordinates are defined as:

$$R_i = \sum_{j=1}^d X_{i,j} \quad W_{i,j} = \frac{X_{i,j}}{R_i}.$$

Then it emerges (Coles and Tawn, 1991) that the limiting process,  $P$ , has intensity on the set  $A \subset \mathbb{R}^2 \setminus \{0\}$  with

$$\Lambda(A) = \int_A \frac{dr}{r^2} dH(\mathbf{w}),$$

where  $H$  is a positive measure on the  $d - 1$  dimensional simplex,  $S_d$ . The radial component gives a sense of how extreme the values are whereas the angular component measures the relative size of each component. In 2 dimensions the angular component can be interpreted as a ray, *e.g.*,  $w = \frac{1}{2}$  is the ray  $X_1 = X_2$ , whereas  $w = 0$  and  $w = 1$  correspond to the rays  $X_1 = 0$  and  $X_2 = 0$  respectively. The spectral measure gives a sense of the ‘spread’ of the points *i.e.*,  $H$  places more weight on values of  $w$  corresponding to rays on which observations are more likely to occur. So if the extremes of two components are likely to occur at the same time more weight will be on values of  $w$  close to  $\frac{1}{2}$ . On the other hand if the extremes are near-independent the density  $h(w)$  of  $H(w)$  will be larger near  $w = 0$  and  $w = 1$ .

We obtain the multivariate extreme value distribution of the componentwise maxima approach (7.3.3) by considering the point process on the set  $A = \{\max_j (X_j - x_j) > 0\}$  with  $x_j > 0$  for all  $j$  (*i.e.*, the set with at least one large component being above its respective  $x_j$ ) since:

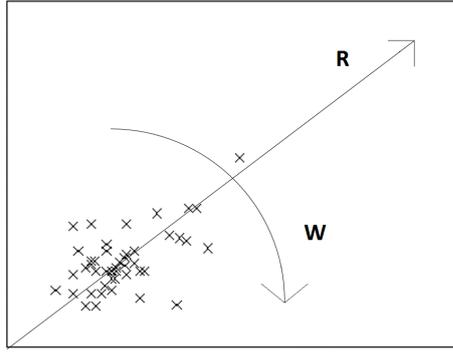


Figure 7.3.2: Visualisation of point process representation in 2-dimensions.

$$\begin{aligned}
 \max_j (X_j - x_j) > 0 &\Leftrightarrow \max_j (Rw_j - x_j) > 0 \\
 &\Leftrightarrow \max_j (R - x_j/w_j) > 0 \\
 &\Leftrightarrow R > \min_j \left( \frac{x_j}{w_j} \right) \quad \text{as all } w_j > 0,
 \end{aligned}$$

then

$$\begin{aligned}
 \Lambda(A) &= \int_{S_d} \int_{R > \min_j \left( \frac{x_j}{w_j} \right)} \frac{dr}{r^2} dH(\mathbf{w}) = \int_{S_d} \left[ -\frac{1}{r} \right]_{\min_j \left( \frac{x_j}{w_j} \right)}^{\infty} dH(\mathbf{w}) \\
 &= \int_{S_d} \frac{1}{\min_j \left( \frac{x_j}{w_j} \right)} dH(\mathbf{w}) \\
 &= \int_{S_d} \max_j \left( \frac{w_j}{x_j} \right) dH(\mathbf{w})
 \end{aligned}$$

and  $\mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d) = \mathbb{P}(N(A) = 0) = \exp(-\Lambda(A)) = \exp(-V(\mathbf{x}))$ .

If we instead consider the set  $A = \{R > u\}$ , then  $\Lambda(A) = \frac{d}{u}$ . Another choice for the set,  $A = \{R > u, \mathbf{W} \in C\}$ , results in  $\Lambda(A) = \frac{H(\{C\})}{u}$ . Combining the intensity measures for these two sets we can find an estimate for  $H$  based on the following:

$$H(\{C\}) = d \frac{H(\{C\})/u}{d/u} = d \frac{\log(\mathbb{P}(R > u, \mathbf{W} \in C))}{\log(\mathbb{P}(R > u))}.$$

A useful property of the point process in Fréchet margins is that as  $t \rightarrow \infty$

$$\mathbb{P}(X \in tA) \sim \frac{1}{t} \mathbb{P}(X \in A),$$

since  $\Lambda(tA) = \frac{1}{t} \Lambda(A)$ . So we can use the probability of being in set  $A$ , which we can estimate

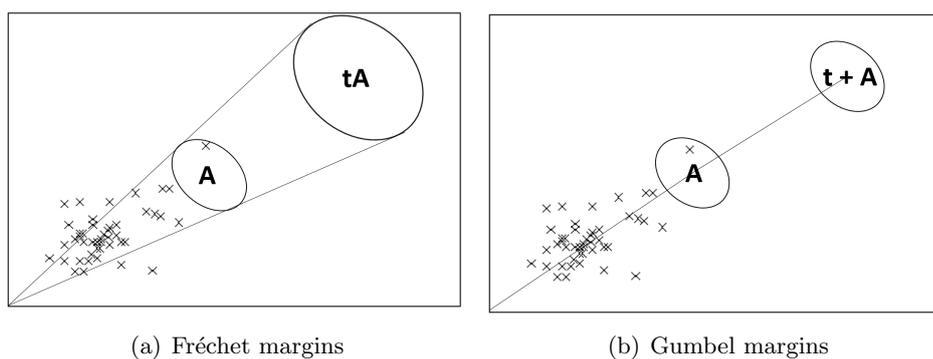


Figure 7.3.3: *Estimating probabilities of extreme sets.*

from the data empirically, to find the probability of lying in a more extreme set  $tA$  in which we have no observations.

It is perhaps more useful to work in Gumbel margins, in which case we can instead find the probability of lying in the set translated by some value  $t$ ,  $\mathbb{P}(X \in t + A)$ . However in both cases, as can be seen in Figure 7.3.3, this technique is only useful when extreme values occur together. When there is asymptotic independence the probability of points lying in these sets will be smaller than each of the associated marginal probabilities and we are more interested in the areas nearer to the axes, *i.e.*, extremes in a subset of the variables. This leads on to the conditional approach of Heffernan and Tawn (2004) to modelling extremes detailed in Appendix E.

## Chapter 8

# Extremes of multivariate temporally dependent sequences

In this chapter we extend the univariate results for dependent sequences to the multivariate setting and present multivariate versions of the extremal index, cluster size distribution and coefficient of asymptotic dependence. We explore two multivariate stationary processes, the MARMAX and M4 process, in §8.2 and §8.3 respectively, and derive the multivariate extremal index and coefficient of asymptotic dependence for both processes. For the M4 process we also derive the multivariate cluster size distribution; both this and  $(\chi^\tau, \bar{\chi}^\tau)$  for the M4 process are believed to be novel contributions. We complete each subsection with simulations of the processes focusing on estimation of the multivariate extremal index.

### 8.1 Multivariate extremes of dependent sequences

#### 8.1.1 Limit distributions

Let  $\{\mathbf{X}_t = (X_{t,1}, \dots, X_{t,d})\}_{t \geq 1}$  be the stationary sequence of  $d$ -dimensional random variables with stationary distribution  $F(\mathbf{x})$  and let  $\{\tilde{\mathbf{X}}_t = (\tilde{X}_{t,1}, \dots, \tilde{X}_{t,d})\}_{t \geq 1}$  be the corresponding i.i.d. sequence with stationary distribution  $F(\mathbf{x})$ . Let  $F_j(x) = \mathbb{P}(X_{t,j} \leq x)$ ,  $j = 1, \dots, d$ , be the marginal distributions. We define the pointwise maxima over a period of length  $n$  as  $\mathbf{M}_n = (M_{n1}, \dots, M_{nd})$  with  $M_{nj} = \max(X_{1,j}, \dots, X_{n,j})$  for  $j = 1, \dots, d$ . First we determine the distribution of the maxima of the independent sequences,  $\tilde{\mathbf{M}}_n$ , by generalising Theorem 7.3.1.1 to allow  $F$  to have different margins.

**Theorem 8.1.1.1** (Nandagopalan (1994)). *For each dimension,  $j \in 1, \dots, d$ , let  $u_{nj}(x) = a_{nj}x + b_{nj}$ , where  $a_{nj} > 0$  and  $b_{nj}$  are a sequence of constants such that the limit distribution of  $X_j$  is non-degenerate, then*

$$\mathbb{P}\left(\tilde{M}_{nj} \leq u_{nj}(x)\right) = F_j^n(u_{nj}(x)) \rightarrow \tilde{G}_j(x),$$

where  $\tilde{G}_j(x)$  is a GEV CDF. If these  $u_{nj}(x)$  also satisfy

$$n(1 - F(\mathbf{u}_n(\mathbf{x}))) \rightarrow -\log \tilde{G}(\mathbf{x}),$$

where  $\mathbf{u}_n(\mathbf{x}) = (u_{n1}(x_1), \dots, u_{nd}(x_d))$  and  $\tilde{G}$  is non-degenerate, then  $\tilde{G}(\mathbf{x})$  is a MEV distribution. The margins of  $\tilde{G}(\mathbf{x})$  are GEV with parameters depending on the margins of  $F$ .

Recall that the MEV distribution in unit Fréchet margins is written as  $G(\mathbf{z}) = \exp(-V(\mathbf{z}))$  with the exponent function  $V$  defined as in Theorem 7.3.1.1. The MEV,  $\tilde{G}(\mathbf{x})$ , in Theorem 8.1.1.1 does not (necessarily) have unit Fréchet margins but we can write  $\tilde{G}(\mathbf{x})$  in terms of  $V$  by using the probability integral transform in the same manner as with copulas in §7.3.3:

$$\begin{aligned} \tilde{G}(\mathbf{x}) &= \exp\left(-V\left(-\frac{1}{\log \tilde{G}_1(x_1)}, \dots, -\frac{1}{\log \tilde{G}_d(x_d)}\right)\right) \\ &= \tilde{G}^{\text{II}}\left(-\frac{1}{\log \tilde{G}_1(x_1)}, \dots, -\frac{1}{\log \tilde{G}_d(x_d)}\right), \end{aligned} \quad (8.1.1)$$

where  $\tilde{G}^{\text{II}}$  is the MEV with unit Fréchet margins (Fréchet distributions are also known as a type II extreme value distributions).

It is often convenient to work in terms of  $\tau(\mathbf{x}) = (-\log(\tilde{G}_1(x_1)), \dots, -\log(\tilde{G}_d(x_d)))$  then (8.1.1) is simply  $\tilde{G}^{\text{II}}(\tau(\mathbf{x})^{-1})$ . Working in  $\mathbf{x}$  we are comparing extremes on the data scale of each margin whereas with  $\tau^{-1}$  we are comparing extremes on a *unified* unit Fréchet scale. The multivariate theory can thus be developed purely in terms of  $\boldsymbol{\tau} = (\tau_1, \dots, \tau_d) \in (0, \infty)^d$  without ‘thought of’  $\mathbf{x}$ . We suppose we can find a sequence of thresholds,  $u_{nj}(\tau_j)$  which satisfy  $n(1 - F_j(u_{nj}(\tau_j))) \rightarrow \tau_j$ . As Robert (2008) notes, a natural choice for the threshold is  $u_{nj}(\tau_j) = \inf\{x \in \mathbb{R} : F_j(x) \geq 1 - \tau_j/n\}$ . If these  $u_{nj}(\tau_j)$  also satisfy  $n(1 - F(\mathbf{u}_n(\boldsymbol{\tau}))) \rightarrow -\log \tilde{H}(\boldsymbol{\tau})$  for some non-degenerate  $\tilde{H}$  then

$$\mathbb{P}\left(\tilde{M}_n \leq \mathbf{u}_n(\boldsymbol{\tau})\right) \rightarrow \tilde{H}(\boldsymbol{\tau}) = \tilde{G}^{\text{II}}(\boldsymbol{\tau}^{-1}), \quad (8.1.2)$$

where  $\tilde{G}^{\text{II}}(\mathbf{z})$  is a MEV distribution function with unit Fréchet margins. The margins in terms of  $\mathbf{x}$  can be recovered by simply replacing  $\boldsymbol{\tau}$  by  $\tau(\mathbf{x}) = (-\log(\tilde{G}_1(x_1)), \dots, -\log(\tilde{G}_d(x_d)))$ . Moreover we can formulate the theory for extremes on a *unified* standard uniform scale with  $\mathbf{t} = (\tilde{G}_1(x_1), \dots, \tilde{G}_d(x_d))$  - this is perhaps useful as it gives an indication as to how extreme

we are considering each component in comparison to the other components.

All three scales have their merits - for the remainder of the chapter we mostly use  $\mathbf{x}$  but in this section and §8.1.2 we also present results in terms of  $\boldsymbol{\tau}$  to provide further insight.

As we did in the univariate case we can obtain a similar limit result for the dependent sequence  $\{\mathbf{X}_t = (X_{t,1}, \dots, X_{t,d})\}_{t \geq 1}$  by restricting the dependence structure. Leadbetter's  $D(u_n)$  condition can be extended to the multivariate setting (Hsing (1989)) whereas Nandagopalan (1994) presents a slightly stronger mixing condition,  $\Delta(\mathbf{u}_n)$ . Like the univariate counterpart both conditions are said to hold if  $\alpha_{n,l_n} \rightarrow 0$  as  $n \rightarrow \infty$  for some  $l_n = o(n)$  where  $\alpha_{n,l_n}$  is in some sense giving a distance between the independent and dependent case. For the multivariate  $D(\mathbf{u}_n)$  condition:

$$\alpha_{n,l_n} = \max(|\mathbb{P}(A^c \cup B^c) - \mathbb{P}(A^c)\mathbb{P}(B^c)| : A \in \mathcal{B}_1^k(u_n), B \in \mathcal{B}_{k+l_n}^n(u_n), 1 \leq k < k + l_n \leq n)$$

and for  $\Delta(\mathbf{u}_n)$ :

$$\alpha_{n,l_n} = \sup(|\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)| : A \in \mathcal{B}_1^k(u_n), B \in \mathcal{B}_{k+l_n}^n(u_n), 1 \leq k < k + l_n \leq n), \quad (8.1.3)$$

where  $\mathcal{B}_m^r(u_n)$  is the  $\sigma$ -field of events  $\{X_{tj} > u_{nj}\}$ ,  $m \leq t \leq r$ ,  $1 \leq j \leq d$ .

**Theorem 8.1.1.2** (Hsing (1989)). *Let  $\mathbf{u}_n(\mathbf{x}) = \mathbf{a}_n \mathbf{x} + \mathbf{b}_n$  be such that  $\mathbb{P}(M_{nj} \leq u_{nj}(x))$  has a non-degenerate limit distribution,  $G_j(x)$ , for  $j = 1, \dots, d$ . If  $D(\mathbf{u}_n(\mathbf{x}))$  holds and the limit exists then*

$$\mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n(\mathbf{x})) \rightarrow G(\mathbf{x}),$$

where  $G(\mathbf{x})$  is a multivariate extreme value distribution.

From the univariate theory, the margins of  $G(\mathbf{x})$  are  $G_j(x_j) = \tilde{G}_j(x_j)^{\theta_j}$  for each  $j$ , where  $\theta_j$  is the extremal index for the sequence of the  $j$ th component and  $\tilde{G}_j$  is the limit distribution of the corresponding independent sequence of the  $j$ th component. Then we can rewrite  $G(\mathbf{x})$

in terms of the univariate extremal indices:

$$\begin{aligned} G(\mathbf{x}) &= \exp\left(-V\left(-\frac{1}{\log G_1(x_1)}, \dots, -\frac{1}{\log G_d(x_d)}\right)\right) \\ &= \exp\left(-V\left(-\frac{1}{\theta_1 \log \tilde{G}_1(x_1)}, \dots, -\frac{1}{\theta_d \log \tilde{G}_d(x_d)}\right)\right) \\ &= \exp\left(-V\left(\frac{1}{\theta_1 \tau_1}, \dots, \frac{1}{\theta_d \tau_d}\right)\right) \quad \text{using } \tau_j = -\log \tilde{G}_j(x_j). \end{aligned}$$

Equivalently, with some restrictions on  $\boldsymbol{\tau}$  (see Robert 2008) and  $\mathbf{u}_n(\boldsymbol{\tau})$  chosen such that the limit exists and  $D(\mathbf{u}_n)$  holds, we have

$$\mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n(\boldsymbol{\tau})) \rightarrow H(\boldsymbol{\tau}) = G^{\text{II}}(\boldsymbol{\tau}^{-1}),$$

where  $G^{\text{II}}(\mathbf{x})$  is a MEV distribution with (non-unit) Fréchet margins (Robert, 2008).

### 8.1.2 Multivariate Extremal Index

The multivariate extremal index is a function describing the relation between the independent multivariate process and the dependent process. The index,  $\theta(\mathbf{t})$ , was introduced by Nandagopalan (1994) in terms of,  $\mathbf{t} = (\tilde{G}_1(x_1), \dots, \tilde{G}_d(x_d))$  but is most often presented in terms of  $\boldsymbol{\tau}$  (Robert, 2008; Ferreira, 1994; Martins and Ferreira, 2005):

$$\theta(\boldsymbol{\tau}) = \frac{\log(H(\boldsymbol{\tau}))}{\log(\tilde{H}(\boldsymbol{\tau}))}, \quad (8.1.4)$$

where  $\tilde{H}(\boldsymbol{\tau})$  is as defined in (8.1.2). In terms of  $\mathbf{x}$ :

$$\theta(\mathbf{x}) = \frac{\log(G(\mathbf{x}))}{\log(\tilde{G}(\mathbf{x}))} = \frac{\log(G^{\text{II}}(\boldsymbol{\tau}(\mathbf{x})^{-1}))}{\log(\tilde{G}^{\text{II}}(\boldsymbol{\tau}(\mathbf{x})^{-1}))}, \quad (8.1.5)$$

where  $G^{\text{II}}$  and  $\tilde{G}^{\text{II}}$  have Fréchet margins with  $\boldsymbol{\tau}(\mathbf{x}) = -(\log(\tilde{G}_1(x_1)), \dots, \log(\tilde{G}_d(x_d)))$ . In particular, when the stationary distribution,  $F$ , has unit Fréchet margins  $\boldsymbol{\tau}(\mathbf{x}) = 1/\mathbf{x}$  and  $\tilde{G} = \tilde{G}^{\text{II}}$ . Note that the margins of  $G$  and  $\tilde{G}$  are different GEVs.

Like the univariate extremal index we have the following relation equivalent to (8.1.5):

$$\tilde{G}(\mathbf{x})^{\theta(\mathbf{x})} = \exp(\theta(\mathbf{x}) \log(\tilde{G}(\mathbf{x}))) = G(\mathbf{x})$$

and

$$\theta(\boldsymbol{\tau}) = \frac{\log(G^{\Pi}(\boldsymbol{\tau}^{-1}))}{\log(\tilde{G}^{\Pi}(\boldsymbol{\tau}^{-1}))} = \frac{V\left(\frac{1}{\theta_1\tau_1}, \dots, \frac{1}{\theta_d\tau_d}\right)}{V\left(\frac{1}{\tau_1}, \dots, \frac{1}{\tau_d}\right)}.$$

### Properties of the multivariate extremal index

1.  $0 \leq \theta(\boldsymbol{\tau}) \leq 1 \quad \forall \boldsymbol{\tau} \in (0, \infty)^d \Leftrightarrow 0 \leq \theta(\boldsymbol{\tau}(\mathbf{x})) \leq 1 \quad \forall \mathbf{x} \in (0, \infty)^d$ .
2.  $\theta(c\boldsymbol{\tau}) = \theta(\boldsymbol{\tau}) \Leftrightarrow \theta(\boldsymbol{\tau}(t^c)) = \theta(\boldsymbol{\tau}(t))$ ,  $\forall c > 0$ .
3. The extremal index for the sequence of the  $j$ th component,  $\theta_j$ , is recovered by setting  $\boldsymbol{\tau} = (0, \dots, 0, \tau_j, 0, \dots, 0)$  into (8.1.4) or  $\mathbf{x} = (\infty, \dots, \infty, x_j, \infty, \dots, \infty)$  into (8.1.5).
4. By property 3, if the components of  $G^{\Pi}$  and  $\tilde{G}^{\Pi}$  are independent:

$$\theta(\boldsymbol{\tau}) = \frac{\sum_{j=1}^d \log(G_j^{\Pi}(\tau_j^{-1}))}{\sum_{j=1}^d \log(\tilde{G}_j^{\Pi}(\tau_j^{-1}))} = \frac{\sum_{j=1}^d \theta_j \tau_j}{\sum_{j=1}^d \tau_j}.$$

5. The extremal index can be written as a function of  $d - 1$  rather than  $d$  variables (Nandagopalan, 1994). For example,  $\theta(\boldsymbol{\tau}) = \theta(\mathbf{a})$  where  $\mathbf{a} = \left(\frac{\tau_1}{\tau_d}, \dots, \frac{\tau_{d-1}}{\tau_d}\right) \in (0, \infty)^{d-1}$ .
6. When  $\tau_j = \tau \forall j$  we recover a relationship between the extremal index and the coefficient of extremal dependence of the  $D = \{1, \dots, d\}$  variables,  $\phi_D$ , defined in §7.3.4:

$$\theta(\boldsymbol{\tau}) = \frac{V\left(\frac{1}{\theta_1}, \dots, \frac{1}{\theta_d}\right)}{V(1, \dots, 1)} = \frac{V\left(\frac{1}{\theta_1}, \dots, \frac{1}{\theta_d}\right)}{\phi_D}.$$

The denominator describes the extremal dependence between variables when there is no temporal dependence.

**Remark 8.1.2.1.** *The intuition behind property 5 is clear when one considers the univariate setting. The univariate extremal index is a constant - it does not depend on  $\tau$  - as it is a limit in one component only. In higher dimensions the extremal index is a function since there is interaction between the components and so the limit depends on the value of the components compared to one another. For example, in two dimensions when  $\tau_1 = \tau_2$  the two components are being considered at the same extremal level and the extremal index is a measure of dependence when the two components are similarly extreme. On the other hand if  $\tau_1 \gg \tau_2$ , then the first component is more extreme than the 2nd component and the extremal index is capturing the dependence at a higher level for the first component than the second.*

The multivariate extremal index can also be found as the univariate extremal index of an associated sequence depending on  $\boldsymbol{\tau}$  (Smith and Weissman, 1996). Let  $\{\mathbf{X}_t\}_{t \geq 1}$  have unit Fréchet margins then the sequence  $\{Z_t(\boldsymbol{\tau})\}_{t \geq 1}$  where  $Z_t(\boldsymbol{\tau}) = \max_{j=1:d}(\tau_j X_{tj})$  has extremal index  $\theta(\boldsymbol{\tau})$ . Applying O'Brien's formulation of the extremal index, (7.1.3), to the sequence  $\{Z_t(\boldsymbol{\tau})\}_{t \geq 1}$  with  $\mathbf{u}_n = \{n/\tau_1, \dots, n/\tau_d\}$  and  $r_n = \lfloor n/k_n \rfloor$  where  $k_n$  satisfies the appropriate conditions<sup>1</sup> we have:

$$\theta(\boldsymbol{\tau}) = \lim_{n \rightarrow \infty} \mathbb{P} \left( X_{t+i,j} \leq \frac{n}{\tau_j}, 1 \leq i \leq r_n, 1 \leq j \leq d \mid \max_j \left( \frac{X_{t,j} \tau_j}{n} \right) > 1 \right).$$

One can also consider other definitions of 'extreme', for example, defining observations as extreme if all components exceed the threshold rather than at least one component exceeding. Nandagopalan (1994) gives a more general formulation of the extremal index on the set of interest:

$$\theta(A) = \frac{\log(\lim_{n \rightarrow \infty} \mathbb{P}(M_n \in u_n(A)))}{\log(\lim_{n \rightarrow \infty} \mathbb{P}(\tilde{M}_n \in u_n(A)))}$$

where  $A = \bigcup_j(x_j, x_j^U)$ , with  $x_j^U$  being the upper end point of  $\tilde{G}_j$ , for the usual component-wise maxima and  $A = \bigcap_j(x_j, x_j^U)$  for the case where all components must exceed to be considered extreme.

### 8.1.3 Cluster size distribution(s) and point processes

In the univariate setting there is a clear definition of a cluster as a set of exceedances of a high threshold,  $u_n$ , in a run of length  $r_n = o(n)$  as  $n \rightarrow \infty$ . In the multivariate setting there are multiple thresholds so the exceedances depend on the thresholds relative to one another. Therefore, as we have seen for the extremal index, the cluster size distribution and underlying point process representation of the exceedances depends on  $\boldsymbol{x}$ . Furthermore, there is not a clear definition of what is extreme in the multivariate setting - it may refer to joint exceedances or exceedance in at least one component. Additionally, one can consider many variants of cluster size distribution, for example conditioning on exceedances in certain components. We now present some of these point process representations and cluster size distributions.

<sup>1</sup>The mixing condition,  $\Delta(\mathbf{u}_n)$ , holds for some  $l_n$  and we can find a sequence  $\{k_n\}_{n \geq 1}$  such that  $k_n \rightarrow \infty$ ,  $k_n l_n / n \rightarrow 0$ ,  $k_n \alpha_{n, l_n} \rightarrow 0$  as  $n \rightarrow \infty$ .

Extending (7.1.5) we obtain the multivariate point process:

$$\mathbf{N}_n^{(\mathbf{x})}([a, b]) = \sum_{i=1}^n \mathbb{1} \left\{ \frac{i}{n} \in [a, b] \right\} (\mathbb{1}\{X_{t+i,1} > u_{n,1}(x_1)\}, \dots, \mathbb{1}\{X_{t+i,d} > u_{n,d}(x_d)\}),$$

where  $0 \leq a < b \leq 1$ . This is a marked point process with marks on  $[0, 1]^d$ .

Let  $\{m_n\}$  be a sequence of positive integers that satisfy

$$m_n \rightarrow \infty, \quad m_n l_n / n \rightarrow 0 \quad \text{and} \quad m_n \alpha_{n,l_n} \rightarrow 0 \quad \text{as} \quad n \rightarrow \infty, \quad (8.1.6)$$

with  $l_n$  satisfying the  $\Delta(u_n)$  condition (8.1.3), and set  $r_n = \lfloor n/m_n \rfloor$ . We then define the multivariate cluster size distribution given there is an exceedance in at least one component as:

$$\begin{aligned} \pi^{(\mathbf{x})}(\mathbf{k}) &= \lim_{n \rightarrow \infty} \pi_n^{(\mathbf{x})}(\mathbf{k}) \\ &= \lim_{n \rightarrow \infty} \mathbb{P} \left( \mathbf{N}_n^{(\mathbf{x})}([0, r_n/n]) = \mathbf{k} \mid \mathbf{N}_n^{(\mathbf{x})}([0, r_n/n]) \neq \mathbf{0} \right) \quad \mathbf{k} \in \mathbb{Z}_+^d / \{\mathbf{0}\}. \end{aligned} \quad (8.1.7)$$

A similar result to Theorem 7.1.7 is obtained:

**Theorem 8.1.3.1.** *Assume that we have a sequence of constants,  $\mathbf{a}_n$  and  $\mathbf{b}_n$  such that the conditions of Theorem 8.1.1.2 are satisfied,  $\{m_n\}$  and  $\{l_n\}$  satisfy (8.1.6) and the limit  $\pi^{(\mathbf{x})}(\mathbf{k})$  exists. Then it follows that  $\mathbf{N}_n^{(\mathbf{x})} \rightarrow \mathbf{N}$ , where  $\mathbf{N}$  is a compound Poisson process with mark distribution  $\pi$  on  $\mathbb{Z}_+^d / \{\mathbf{0}\}$  and intensity measure  $\Lambda((t_1, t_2) \times [x, \infty)^d) = -(t_2 - t_1) \log(G(\mathbf{x}))$ .*

Nandagopalan (1994) also described a one dimensional distribution,  $\pi^{(\mathbf{x})}$ , of the total cluster size over all components. Consider the one-dimensional point process counting when there is an exceedance in at least one component:

$$N_n^{(\mathbf{x})}([a, b]) = \sum_{i=1}^n \mathbb{1} \left\{ \frac{i}{n} \in [a, b] \right\} \mathbb{1} \left\{ \bigcup_{j=1}^d \{X_{t+i,j} > u_{n,j}(x_j)\} \right\}.$$

Using this point process we arrive at the following ‘cluster size’ distribution:

$$\pi^{(\mathbf{x})}(k) = \lim_{n \rightarrow \infty} \mathbb{P} \left( N_n^{(\mathbf{x})}([0, r_n/n]) = k \mid N_n^{(\mathbf{x})}([0, r_n/n]) > 0 \right) \quad k \in \mathbb{Z}_+ / \{0\}. \quad (8.1.8)$$

Under mild conditions, including that the mixing condition  $\Delta(\mathbf{u}_n)$  holds, Nandagopalan

(1994) showed that (8.1.8) has a nice relation with the extremal index:

$$\theta(\mathbf{x}) = \left( \sum_{k=1}^{\infty} k \pi^{(\mathbf{x})}(k) \right)^{-1}.$$

That is the extremal index is the reciprocal of the limiting mean number of exceedances as found in the univariate setting.

One can also consider the distribution of joint exceedances as (8.1.8) but with:

$$\tilde{N}_n^{(\mathbf{x})}([a, b]) = \sum_{i=1}^n \mathbb{1} \left\{ \frac{i}{n} \in [a, b] \right\} \mathbb{1} \left\{ \bigcap_{j=1}^d \{X_{t+i,j} > u_{n,j}(x_j)\} \right\}, \quad \text{with } 0 \leq a < b \leq 1.$$

Now we consider the connection between the distribution of cluster sizes in the multivariate setting with that of the marginals. Let  $N_{n,j}$  be the point process, as defined in (7.1.5), counting exceedances in component  $j$ . In (8.1.7) we condition on there being an exceedance in some component, we can instead define the cluster size distribution given an exceedance in a particular component,  $j$ :

$$\begin{aligned} \pi_j^{(\mathbf{x})}(\mathbf{k}) &= \lim_{n \rightarrow \infty} \mathbb{P} \left( \mathbf{N}_n^{(\mathbf{x})}([0, r_n/n]) = \mathbf{k} \mid N_{n,j}([0, r_n/n]) > 0 \right) & (8.1.9) \\ &= \lim_{n \rightarrow \infty} \mathbb{P} \left( \left( \sum_{i=1}^{r_n} \mathbb{1} \{X_{t+i,l} > u_{n,l}\} = k_l, l = 1, \dots, d \right) \mid \sum_{i=1}^{r_n} \mathbb{1} \{X_{t+i,j} > u_{n,j}\} > 0 \right) \\ &= \lim_{n \rightarrow \infty} \frac{\mathbb{P} \left( \left( \sum_{i=1}^{r_n} \mathbb{1} \{X_{t+i,1} > u_{n,1}\} = k_1, \dots, \sum_{i=1}^{r_n} \mathbb{1} \{X_{t+i,d} > u_{n,d}\} = k_d \right) \right)}{1 - \mathbb{P} \left( \bigcap_{i=1}^{r_n} \{X_{t+i,j} \leq u_{n,j}\} \right)} \\ &\qquad \qquad \qquad \mathbf{k} \in \mathbb{Z}_+^d, k_j > 0. \end{aligned}$$

This distribution is different to (8.1.7) only in the denominator; the ratio of the cluster size distributions (8.1.7) and (8.1.9) is the ratio of the probability of an exceedance in one component to the probability of at least one exceedance in component  $j$ . That is, we have  $\pi_j^{(\mathbf{x})}(\mathbf{k}) = q \pi^{(\mathbf{x})}(\mathbf{k})$  for  $\mathbf{k} \in \mathbb{Z}_+^d$ ,  $k_j > 0$  with

$$\begin{aligned} q &= \lim_{n \rightarrow \infty} \frac{1 - \mathbb{P}(\mathbf{N}_n([0, r_n/n]) = \mathbf{0})}{1 - \mathbb{P}(N_{n,j}([0, r_n/n]) = 0)} = \lim_{n \rightarrow \infty} \frac{1 - \mathbb{P} \left( \sum_{i=1}^{r_n} \sum_{j=1}^d \mathbb{1} \{X_{i,j} > u_{n,j}\} = 0 \right)}{1 - \mathbb{P} \left( \sum_{i=1}^{r_n} \mathbb{1} \{X_{i,j} > u_{n,j}\} = 0 \right)} \\ &= \lim_{n \rightarrow \infty} \frac{1 - \mathbb{P} \left( \bigcap_{i=1}^{r_n} \bigcap_{j=1}^d \{X_{i,j} \leq u_{n,j}\} \right)}{1 - \mathbb{P} \left( \bigcap_{i=1}^{r_n} \{X_{i,j} \leq u_{n,j}\} \right)}, \end{aligned}$$

and  $1 < q \leq 1/\pi^{(\mathbf{x})}(\mathbf{k})$ .

The distribution of cluster sizes given the size in one component can be obtained from (8.1.9) and the corresponding marginal cluster size distribution  $\pi_j(k_j)$ :

$$\begin{aligned} \pi^{(\mathbf{x})}(\mathbf{k}^{-j}|k_j) &= \lim_{n \rightarrow \infty} \mathbb{P} \left( \left( \sum_{i=1}^{r_n} \mathbb{1}\{X_{t+i,l} > u_{n,l}\} = k_l, l = 1, \dots, d \right) \middle| \sum_{i=1}^{r_n} \mathbb{1}\{X_{t+i,j} > u_{n,j}\} = k_j \right) \\ &= \frac{\pi_j^{(\mathbf{x})}(\mathbf{k})}{\pi(k_j)}. \end{aligned}$$

In summary variants of cluster size distribution include:

- $\pi_j(k_j)$  - the marginal cluster size distribution from the univariate setting (7.1.6);
- $\pi^{(\mathbf{x})}(\mathbf{k})$  - the multivariate cluster size distribution given an exceedance in some component (8.1.7);
- $\pi^{(\mathbf{x})}(k)$  - the one dimensional distribution of exceedances over all components (union of exceedances (8.1.8));
- $\tilde{\pi}^{(\mathbf{x})}(k)$  - the one dimensional distribution of joint exceedances (intersection of exceedances);
- $\pi_j^{(\mathbf{x})}(\mathbf{k})$  - the multivariate cluster size distribution given an exceedance in a particular component,  $j$  (8.1.9);
- and  $\pi(\mathbf{k}^{-j}|k_j)$  - the  $d - 1$  dimensional cluster size distribution given the exceedance in the  $j$ th component.

#### 8.1.4 Coefficient of asymptotic dependence for dependent multivariate sequences

We extend the extremal dependence measures  $(\chi, \bar{\chi})$  of §7.3.4 to the multivariate setting by considering the dependence between two different sets of components at some lag  $\tau$ . We define

$$\chi_{C,C^*}^{(\tau)} = \lim_{u \rightarrow 1} \chi_{C,C^*}^{(\tau)}(u) = \lim_{u \rightarrow 1} \mathbb{P} \left( \bigcap_{j^* \in C^*} \{F_{j^*}(X_{t+\tau,j^*}) > u\} \middle| \bigcap_{j \in C} \{F_j(X_{t,j}) > u\} \right),$$

where  $C \subseteq \{1, \dots, d\}$  and  $C^* \subseteq \{1, \dots, d\}$ .

It is simple, and perhaps most useful, to consider the measure for pairs of variables  $(j, j^*)$  split by a certain lag. Then the coefficient of asymptotic dependence,  $\chi_{j,j^*}^\tau$ , has a similar

form to (7.3.6) for bivariate random variables with temporal independence:

$$\chi_{j,j^*}^{(\tau)} = \lim_{u \rightarrow 1} \chi_{j,j^*}^{(\tau)}(u) \quad \text{where} \quad \chi_{j,j^*}^{(\tau)}(u) = 2 - \frac{\log(C_{j,j^*}^{(\tau)}(u, u))}{\log u}, \quad (8.1.10)$$

where

$$C_{j,j^*}^{(\tau)}(u, u) = \mathbb{P}(F_j(X_{t,j}) \leq u, F_{j^*}(X_{t+\tau,j^*}) \leq u).$$

Similarly an extension to the dual measure,  $\bar{\chi}$ , is

$$\bar{\chi}_{j,j^*}^{(\tau)} = \lim_{u \rightarrow 1} \bar{\chi}_{j,j^*}^{(\tau)}(u) = \lim_{u \rightarrow 1} \frac{2 \log(1 - u)}{\log(1 - 2u + C_{j,j^*}^{(\tau)}(u, u))} - 1. \quad (8.1.11)$$

## 8.2 The multivariate ARMAX process

### 8.2.1 The MARMAX process

We now present the MARMAX process, a multivariate stationary process introduced by Ferreira and Ferreira (2013) which extends the ARMAX process of §7.2. We derive the multivariate extremal index and coefficient of asymptotic dependence for this process and apply and discuss estimation of the multivariate extremal index.

Consider the sequence  $\{\mathbf{X}_t\}_{t \geq 1} = \{X_{t,1}, \dots, X_{t,d}\}_{t \geq 1}$  such that

$$X_{t,j} = \max(c_j X_{t-1,j}, \epsilon_{t,j}) \quad 0 < c_j < 1 \quad j = 1, \dots, d,$$

where  $\epsilon_t \sim G$ ,  $\{\epsilon_t\}_{t \geq 1}$  are independent and  $\mathbf{X}_0 \sim F_0$  independent of  $\epsilon_t$  for some  $d$ -dimensional distributions  $F_0$  and  $G$ . Following the same argument as in §7.2  $\{\mathbf{X}_t\}_{t \geq 1}$  is a stationary sequence with common distribution:

$$F(\mathbf{x}) = \lim_{n \rightarrow \infty} F_0 \left( \frac{x_1}{c_1^n}, \dots, \frac{x_d}{c_d^n} \right) \prod_{t=0}^{n-1} G \left( \frac{x_1}{c_1^t}, \dots, \frac{x_d}{c_d^t} \right) = \prod_{t=0}^{\infty} G \left( \frac{x_1}{c_1^t}, \dots, \frac{x_d}{c_d^t} \right) \quad (8.2.1)$$

by extension of (7.2.3) and  $F$  satisfies, extending (7.2.5),

$$F(\mathbf{x}) = G(\mathbf{x}) F \left( \frac{x_1}{c_1}, \dots, \frac{x_d}{c_d} \right). \quad (8.2.2)$$

If  $F$  is a MEV distribution with unit Fréchet margins then  $G$  is also a MEV distribution

albeit with different margins and vice versa. If we want the stationary distribution to be MEV with unit Fréchet margins, *i.e.*,  $F(\mathbf{x}) = \exp(-V_F(x_1, \dots, x_d))$ , then  $G$  must have marginal distributions

$$G_j(x_j) = \frac{F_j(x_j)}{F_j\left(\frac{x_j}{c_j}\right)} = \exp\left(-\frac{1-c_j}{x_j}\right) = F_j(x_j)^{1-c_j} \quad j = 1, \dots, d$$

and

$$G(\mathbf{x}) = \exp\left(-\left[V_F(x_1, \dots, x_d) - V_F\left(\frac{x_1}{c_1}, \dots, \frac{x_d}{c_d}\right)\right]\right) \quad \text{using (8.2.2).}$$

Note that the extremal index of the marginal ARMAX process is  $\theta_j = 1 - c_j$ . It follows that  $G$  is a MEV distribution since it has GEV margins and its copula is max-stable:

$$\begin{aligned} C_G^k(\mathbf{u}^{\frac{1}{k}}) &= \left[ G\left(G_1(X_1) \leq u_1^{\frac{1}{k}}, \dots, G_d(X_d) \leq u_d^{\frac{1}{k}}\right) \right]^k \\ &= \exp\left(-k \left[ V_F\left(G_1^{-1}(u_1^{\frac{1}{k}}), \dots, G_d^{-1}(u_d^{\frac{1}{k}})\right) - V_F\left(\frac{G_1^{-1}(u_1^{\frac{1}{k}})}{c_1}, \dots, \frac{G_d^{-1}(u_d^{\frac{1}{k}})}{c_d}\right) \right]\right) \\ &= \exp\left(-k \left[ V_F(kG_1^{-1}(u_1), \dots, kG_d^{-1}(u_d)) - V_F\left(k\frac{G_1^{-1}(u_1)}{c_1}, \dots, k\frac{G_d^{-1}(u_d)}{c_d}\right) \right]\right) \\ &= \exp\left(-\left[ V_F(G_1^{-1}(u_1), \dots, G_d^{-1}(u_d)) - V_F\left(\frac{G_1^{-1}(u_1)}{c_1}, \dots, \frac{G_d^{-1}(u_d)}{c_d}\right) \right]\right) \\ &= C_G(\mathbf{u}). \end{aligned}$$

Thus we can write  $G(\mathbf{x}) = \exp(-V_G(x_1^*, \dots, x_d^*))$  where  $x_j^* = F_j^{-1}G_j(x_j) = x_j/(1 - c_j)$ . Examples of  $G$  and the extremal index given different choices of  $V_F$  are shown later (§8.2.3).

A useful consequence if  $F(\mathbf{x})$  and  $G(\mathbf{x})$  are MEV is that both distributions are max-stable:  $C_F^k(\mathbf{u}^{\frac{1}{k}}) = C_F(\mathbf{u})$  and  $C_G^k(\mathbf{u}^{\frac{1}{k}}) = C_G(\mathbf{u})$ . For  $F$  with unit Fréchet margins and  $G$  determined from  $F$  as above we have  $F^n(n\mathbf{x}) = F(\mathbf{x})$  and  $G^n(n\mathbf{x}) = G(\mathbf{x})$ . Let  $\mathbf{u}_n(\mathbf{x}) = n\mathbf{x}$  and let  $\{\tilde{\mathbf{X}}_t\}_{t \geq 1}$  be the independent sequence of vectors corresponding to the MARMAX sequence  $\{\mathbf{X}_t\}_{t \geq 1}$ . The components of each vector,  $\tilde{\mathbf{X}}_t$ , are potentially dependent but each vector  $\tilde{\mathbf{X}}_t$  in the sequence  $\{\tilde{\mathbf{X}}_t\}_{t \geq 1}$  is independent. Then we have the following limits:

$$\mathbb{P}\left(\tilde{\mathbf{M}}_n \leq \mathbf{u}_n(\mathbf{x})\right) = \mathbb{P}\left(\tilde{\mathbf{X}}_{t+1} \leq n\mathbf{x}, \dots, \tilde{\mathbf{X}}_{t+n} \leq n\mathbf{x}\right) = F^n(n\mathbf{x}) \rightarrow F(\mathbf{x}),$$

and, since  $M_{n,j} = \max(c_j X_{t,j}, \epsilon_{t+1,j}, \dots, \epsilon_{t+n,j})$ ,

$$\begin{aligned} \mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n(\mathbf{x})) &= \mathbb{P}(\mathbf{X}_{t+1} \leq n\mathbf{x}, \dots, \mathbf{X}_{t+n} \leq n\mathbf{x}) \\ &= F\left(\frac{nx_1}{c_1}, \dots, \frac{nx_d}{c_d}\right) G^n(n\mathbf{x}) \rightarrow G(\mathbf{x}), \quad n \rightarrow \infty. \end{aligned}$$

Ferreira and Ferreira (2013) show that the mixing condition  $\Delta(\mathbf{u}_n)$  (8.1.3) holds for the MARMAX sequence with general  $F$  and  $G$ .

## 8.2.2 Multivariate Extremal Index

For the following we consider  $F$  as a MEV distribution with unit Fréchet margins. Since the necessary mixing condition holds, using (8.1.5) the extremal index is  $\theta(\mathbf{x}) = \frac{\log G(\mathbf{x})}{\log F(\mathbf{x})}$ , where  $F$  and  $G$  are both MEV distributions with Fréchet margins. We can write  $\theta(\mathbf{x})$  in terms of just one of the distributions by using the relations (8.2.2) and (8.2.1):

$$\theta(\mathbf{x}) = 1 - \frac{\log F\left(\frac{x_1}{c_1}, \dots, \frac{x_d}{c_d}\right)}{\log F(\mathbf{x})} \quad \text{and} \quad \theta(\mathbf{x}) = \frac{\log G(\mathbf{x})}{\sum_{t=0}^{\infty} \log G\left(\frac{x_1}{c_1^t}, \dots, \frac{x_d}{c_d^t}\right)}.$$

Moreover, since  $F$  and  $G$  are both MEV distributions we can write the multivariate extremal index in terms of the exponent functions,  $V_F(\mathbf{x}) = -\log F(\mathbf{x})$  and

$$V_G(\mathbf{x}) = -\log G(G_1^{-1}F_1(x_1), \dots, G_d^{-1}F_d(x_d)) = -\log G((1-c_1)x_1, \dots, (1-c_d)x_d):$$

$$\theta(\mathbf{x}) = \frac{V_G\left(\frac{x_1}{1-c_1}, \dots, \frac{x_d}{1-c_d}\right)}{V_F(x_1, \dots, x_d)} = 1 - \frac{V_F\left(\frac{x_1}{c_1}, \dots, \frac{x_d}{c_d}\right)}{V_F(x_1, \dots, x_d)}.$$

The extremal index for component  $j$ ,  $\theta_j = 1 - c_j$ , is recovered by setting all  $x_i, i \neq j$  to infinity. Since  $V_F$  and  $V_G$  are homogeneous functions of order  $-1$ ,  $\theta(a\mathbf{x}) = \theta(\mathbf{x})$  where  $a > 0$  is some constant. We can therefore reparameterise,  $\tilde{x}_j = x_j / \sum_{j=1}^d x_j$  (so that  $0 < \tilde{x}_j < 1$ , which is an easier input to interpret), without changing the extremal index, *i.e.*,  $\theta(x_1, \dots, x_d) = \theta(\tilde{x}_1, \dots, \tilde{x}_d)$ .

## 8.2.3 Examples

Here we consider various dependence structures for the stationary distribution  $F$ . In order to simulate a MARMAX process with such a stationary distribution one must calculate  $G$ . We let  $\mathbf{x}^* = (x_1^*, \dots, x_d^*)$  with  $x_j^* = x_j / (1 - c_j)$  throughout.

**Example 1a (General MEV with unit Fréchet margins and constant  $x$ )**

If  $x_j = x \forall j$  then

$$\begin{aligned} V_G\left(\frac{1}{\theta_1}, \dots, \frac{1}{\theta_d}\right) &= V_G\left(\frac{1}{1-c_1}, \dots, \frac{1}{1-c_d}\right) = V_F(1, \dots, 1) - V_F(c_1^{-1}, \dots, c_d^{-1}) \\ &= \phi_D - V_F(c_1^{-1}, \dots, c_d^{-1}), \end{aligned}$$

where  $\phi_D$  is the coefficient of extremal dependence of the  $D = \{1, \dots, d\}$  variables (§7.3.4)

for  $F$ . The multivariate extremal index is

$$\theta(\mathbf{x}) = 1 - \frac{V_F(c_1^{-1}, \dots, c_d^{-1})}{V_F(1, \dots, 1)} = \frac{V_G\left(\frac{1}{1-c_1}, \dots, \frac{1}{1-c_d}\right)}{V_F(1, \dots, 1)} = \frac{V_G\left(\frac{1}{\theta_1}, \dots, \frac{1}{\theta_d}\right)}{\phi_D}.$$

**Example 1b (General MEV with unit Fréchet margins and constant  $c$ )**

Clearly if  $c_j = c \forall j$  then  $V_G(\mathbf{x}^*) = (1-c)V_G(\mathbf{x}) = V_F(\mathbf{x}) - cV_F(\mathbf{x})$ , i.e.,  $V_G(\mathbf{x}) = V_F(\mathbf{x})$ , and  $\theta(\mathbf{x}) = 1 - c$ .

**Example 2 (Independence)**

Consider  $V_F(\mathbf{x}) = \sum_{j=1}^d x_j^{-1}$ , then

$$V_G(\mathbf{x}^*) = V_G\left(\frac{x_1}{1-c_1}, \dots, \frac{x_d}{1-c_d}\right) = V_F(\mathbf{x}) - V_F\left(\frac{x_1}{c_1}, \dots, \frac{x_d}{c_d}\right) = \sum_{j=1}^d \frac{1-c_j}{x_j}.$$

So  $V_F(\mathbf{x}) = V_G(\mathbf{x})$  and

$$\theta(\mathbf{x}) = \frac{\sum_{j=1}^d \frac{1-c_j}{x_j}}{\sum_{j=1}^d \frac{1}{x_j}} \quad (\text{c.f., Nandagopalan (1994) Prop. 3.4}).$$

In particular,  $\theta(x, \dots, x) = \frac{\sum_{j=1}^d 1-c_j}{d} = \frac{\sum_{j=1}^d \theta_j}{d}$ .

**Example 3 (Perfect dependence)**

Consider  $V_F(\mathbf{x}) = \max_j(x_j^{-1})$ , then

$$\begin{aligned} V_G(\mathbf{x}^*) &= V_G\left(\frac{x_1}{1-c_1}, \dots, \frac{x_d}{1-c_d}\right) = \max_j(x_j^{-1}) - \max_j\left(\frac{c_j}{x_j}\right) \\ \text{and} \quad \theta(\mathbf{x}) &= 1 - \frac{\max_j\left(\frac{c_j}{x_j}\right)}{\max_j(x_j^{-1})}, \end{aligned}$$

with  $V_G(x, \dots, x) = V_F(x, \dots, x)$ , and  $\theta(x, \dots, x) = 1 - \max_j(c_j)$ .

**Example 4 (Logistic)**

Consider  $F$  with logistic dependence function:  $V_F = \left( \sum_{j=1}^d x_j^{-\frac{1}{\alpha}} \right)^\alpha$  where  $0 < \alpha < 1$ . Examples 2 and 3 are the special case of the logistic dependence function where  $\alpha = 1$  and  $\alpha \rightarrow 0$  respectively.

$$V_G(\mathbf{x}^*) = \left( \sum_{j=1}^d x_j^{-\frac{1}{\alpha}} \right)^\alpha - \left( \sum_{j=1}^d \left( \frac{c_j}{x_j} \right)^{\frac{1}{\alpha}} \right)^\alpha \quad \text{and} \quad \theta(\mathbf{x}) = 1 - \frac{\left( \sum_{j=1}^d \left( \frac{c_j}{x_j} \right)^{\frac{1}{\alpha}} \right)^\alpha}{\left( \sum_{j=1}^d x_j^{-\frac{1}{\alpha}} \right)^\alpha}.$$

When  $\alpha$  is close to 0 the largest  $\frac{c_j}{x_j}(x_j)$  dominates the sum on the numerator(denominator) so  $\alpha \rightarrow 0$  gives same result as example 3.

**Example 4a (Logistic with constant  $x$ )**

If  $x_j = x \forall j$  then

$$V_G(\mathbf{x}^*) = \frac{1}{x} \left( d^\alpha - \left( \sum_{j=1}^d c_j^{\frac{1}{\alpha}} \right)^\alpha \right) \quad \text{and} \quad \theta(\mathbf{x}) = 1 - \frac{\left( \sum_{j=1}^d c_j^{\frac{1}{\alpha}} \right)^\alpha}{d^\alpha}.$$

**8.2.4  $(\chi, \bar{\chi})$  for the MARMAX process**

We now present the extremal dependence measures  $(\chi_{j,j^*}^{(\tau)}, \bar{\chi}_{j,j^*}^{(\tau)})$  for the dependence between component  $j$  and  $j^*$  at a lag  $\tau$  apart in the MARMAX sequence with common distribution  $F$  MEV with unit Fréchet margins. First we need to derive the copula of the joint distribution of these two components with lag  $\tau > 0$ :

$$\begin{aligned} C_{j,j^*}^{(\tau)}(u, u) &= \mathbb{P}(F_j(X_{t,j}) \leq u, F_{j^*}(X_{t+\tau,j^*}) \leq u) \\ &= \mathbb{P}\left(F_j(X_{t,j}) \leq u, F_{j^*}(X_{t,j^*}) \leq F_{j^*}^{-1}\left(\frac{F_{j^*}^{-1}(u)}{c_{j^*}^\tau}\right), \epsilon_{t+1,j^*} \leq \frac{F_{j^*}^{-1}(u)}{c_{j^*}^{\tau-1}}, \dots, \epsilon_{t+\tau,j^*} \leq F_{j^*}^{-1}(u)\right) \\ &= \mathbb{P}\left(F_j(X_{t,j}) \leq u, F_{j^*}(X_{t,j^*}) \leq u^{c_{j^*}^\tau}\right) \prod_{i=0}^{\tau-1} G_{j^*}\left(\frac{F_{j^*}^{-1}(u)}{c_{j^*}^i}\right) \\ &= C_{j,j^*}^{(0)}\left(u, u^{c_{j^*}^\tau}\right) \prod_{i=0}^{\tau-1} u^{(1-c_{j^*}^i)c_{j^*}^i} \\ &= C_{j,j^*}^{(0)}\left(u, u^{c_{j^*}^\tau}\right) u^{1-c_{j^*}^\tau} \quad \text{since } \sum_{i=0}^{\tau-1} c_{j^*}^i = \frac{1-c_{j^*}^\tau}{1-c_{j^*}}. \end{aligned}$$

The copula between 2 components,  $(j, j^*)$ , with no lag, evaluated at  $(u, u^{c_{j^*}^\tau})$  is equivalent to  $F(\mathbf{y}) = \exp(-V_F(\mathbf{y}))$  with  $y_j = F_j^{-1}(u) = (-\log u)^{-1}$ ,  $y_{j^*} = F_{j^*}^{-1}(u^{c_{j^*}^\tau}) = (-c_{j^*}^\tau \log u)^{-1}$  and  $y_k = \infty \forall k \in (1, \dots, d) \setminus \{j, j^*\}$ . This simplifies to  $C_{j,j^*}^{(0)}(u, u^{c_{j^*}^\tau}) = u^{V_F(\mathbf{z})}$  where  $z_j = 1$ ,  $z_{j^*} = \frac{\log u}{c_{j^*}^\tau \log u} = c_{j^*}^{-\tau}$  and  $z_k = \infty \forall k \in (1, \dots, d) \setminus \{j, j^*\}$ . Then, using (8.1.10) and (8.1.11),

$$\chi_{j,j^*}^{(\tau)}(u) = 1 + c_{j^*}^\tau - V_F(\mathbf{z}) \quad \bar{\chi}_{j,j^*}^{(\tau)}(u) = \frac{\log(1 - 2u + u^2)}{\log\left(1 - 2u + u^{2-\chi_{j,j^*}^{(\tau)}}\right)} - 1.$$

As we would expect  $(\chi_{j,j^*}^{(\tau)}, \bar{\chi}_{j,j^*}^{(\tau)}) = (0, 0)$  when  $c_{j^*} = 0$  or  $\tau \rightarrow \infty$  since then we have independence between the two components at lag  $\tau$ . In all other cases  $\bar{\chi}_{j,j^*}^{(\tau)} = 1$ , indicating asymptotic dependence.

For the special case where  $F$  is MEV with the logistic dependence function (7.3.5) we have

$$C_{j,j^*}^{(\tau)}(u, u) = u^{\left(1 + c_{j^*}^{\frac{\tau}{\alpha}}\right)^\alpha} u^{1 - c_{j^*}^\tau}.$$

Therefore we obtain

$$\chi_{j,j^*}^{(\tau)}(u) = 2 - \frac{\log(C_{j,j^*}^{(\tau)}(u, u))}{\log u} = 1 + c_{j^*}^\tau - \left(1 + c_{j^*}^{\frac{\tau}{\alpha}}\right)^\alpha.$$

The bivariate logistic results for independent sequences are recovered when  $\tau = 0$  or  $c_{j^*} = 1$ , *i.e.*,  $\chi_{j,j^*}^{(\tau)} = 2 - 2^\alpha$ . We have independence (component-wise) when  $\alpha = 1$  so  $(\chi_{j,j^*}^{(\tau)}, \bar{\chi}_{j,j^*}^{(\tau)}) = (0, 0)$  in this case for any value of  $c_{j^*}$  or  $\tau$ . For  $0 < c_{j^*} \leq 1$  and fixed  $\alpha$  and  $\tau$  the dependence level increases towards  $2 - 2^\alpha$  as  $c_{j^*} \rightarrow 1$ , *i.e.*, as serial dependence decreases. Similarly for  $0 \leq \alpha < 1$  and fixed  $c_{j^*}$  and  $\tau$  the dependence level increases towards  $c_{j^*}^\tau$ , the univariate ARMAX  $\chi^\tau$  (7.2.9), as  $\alpha \rightarrow 0$ , *i.e.*, as component-wise dependence increases.

### 8.2.5 MARMAX simulations and estimation of $\theta(\mathbf{x})$

Here we consider the behaviour of the extremal index *and dependence measures*  $(\chi, \bar{\chi})$  for simulated bivariate MARMAX sequences with common distribution MEV with unit Fréchet margins and a logistic dependence function. We explore a possible multivariate declustering method and apply it to the simulated sequences.

Recall that the extremal index,  $\theta(\mathbf{x})$  is equal to  $\theta(\tilde{\mathbf{x}})$  with  $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_d)$  and  $\tilde{x}_j = x_j / \sum_{j=1}^d x_j$ . Moreover,  $\sum_{j=1}^d \tilde{x}_j = 1$  so one can define  $\tilde{x}_d$  as a function of  $\tilde{x}_1, \dots, \tilde{x}_{d-1}$ . Since

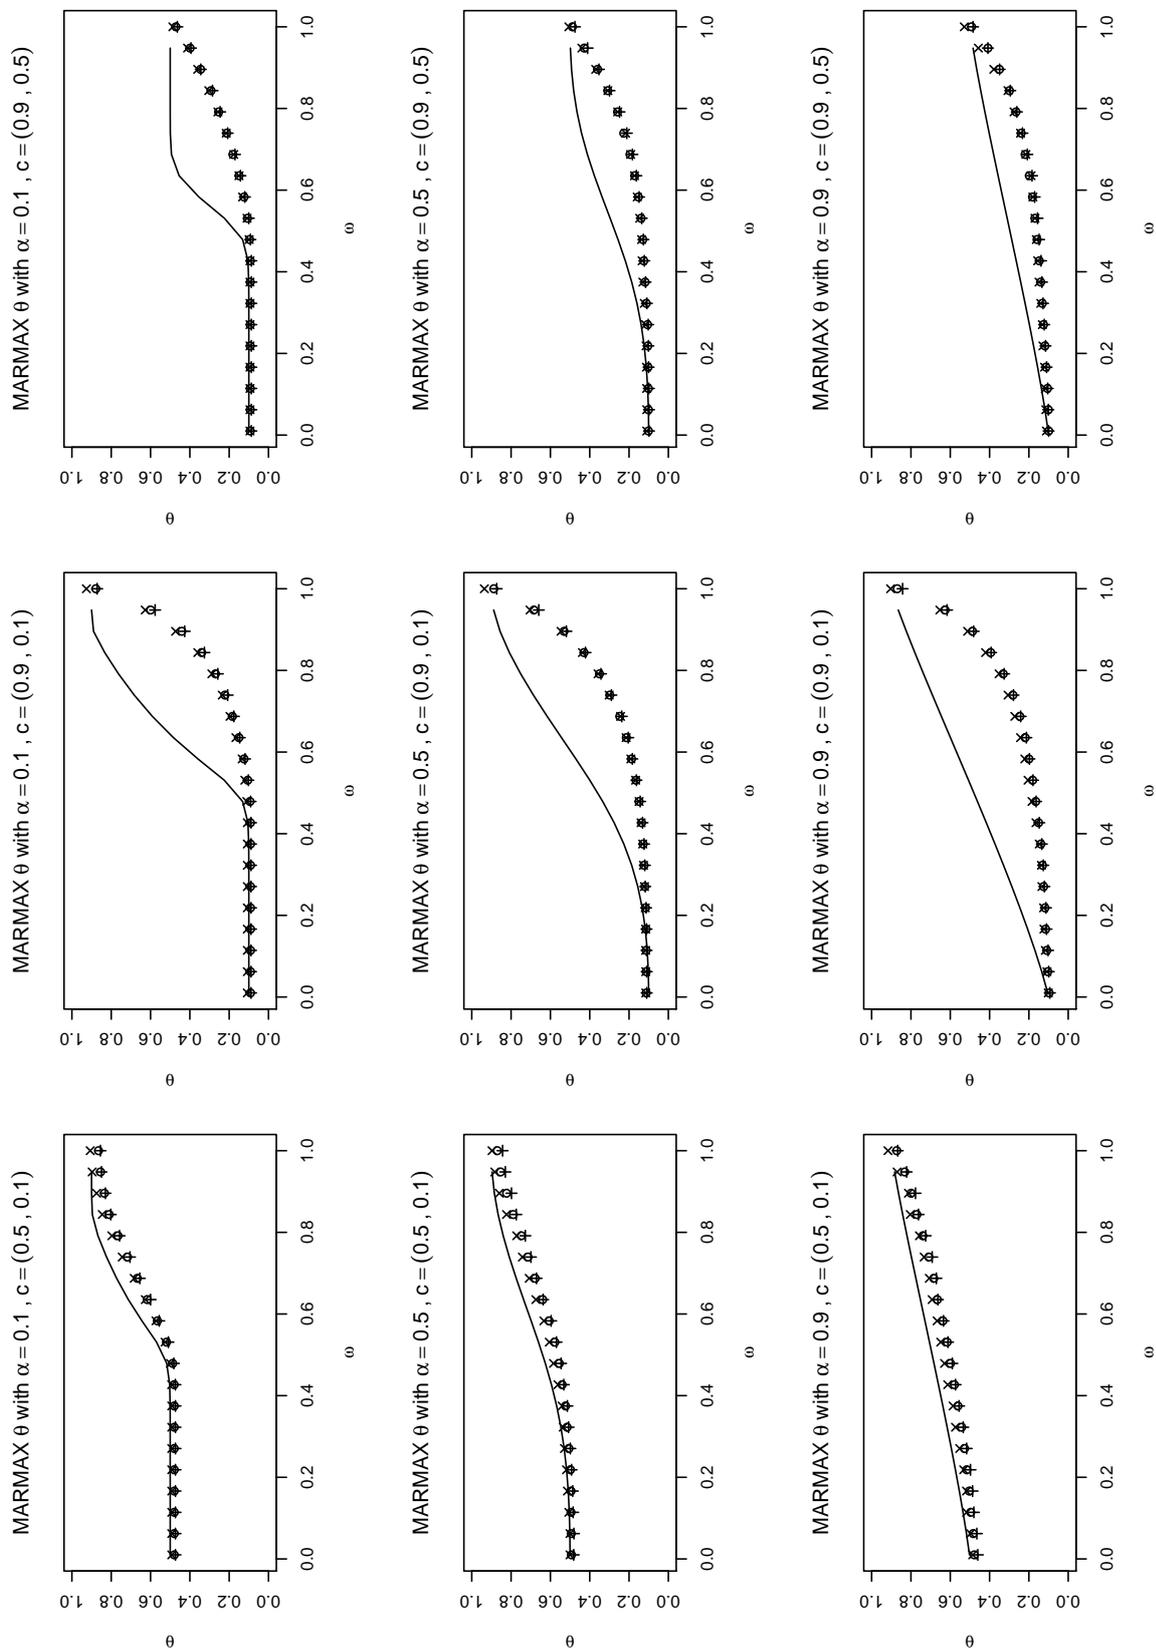


Figure 8.2.1: Extremal index for  $\tau = (\omega, 1 - \omega)$ ,  $0 < \omega < 1$ . True extremal index (solid line), runs estimator with run length 10 (o), empirical estimator from intervals declustering (+) and the intervals estimator (x). Each plot is based on a particular bi-variate ARMAX simulation where  $F$  has a logistic dependence structure with  $\alpha$  and  $c$  indicated in the title.

$0 < \tilde{x}_j < 1$  we can think of each  $\tilde{x}_j$  as a weight on component  $j$ . In the bivariate setting we can simply set  $\tilde{x}_1 = \omega$  and  $\tilde{x}_2 = 1 - \omega$  so we have the extremal index in terms of one parameter. For common distribution  $F = \exp\left(-\left(\sum_{j=1}^d x_j^{-\frac{1}{\alpha}}\right)^\alpha\right)$  where  $0 < \alpha \leq 1$  the bivariate ARMAX extremal index is

$$\theta(\omega, 1 - \omega) = 1 - \frac{\left(\left(\frac{\omega}{c_1}\right)^{-1/\alpha} + \left(\frac{1-\omega}{c_2}\right)^{-1/\alpha}\right)^\alpha}{(\omega^{-1/\alpha} + (1 - \omega)^{-1/\alpha})^\alpha},$$

for  $0 < \omega < 1$ . Figure 8.2.1 shows the extremal index,  $\theta(\omega, 1 - \omega)$ , and estimators thereof against  $\omega$  based on bi-variate ARMAX simulations where  $F$  has a logistic dependence structure with  $\alpha$  and  $\mathbf{c}$  indicated in the title. It is immediately clear, as must be the case, that  $\theta(0, 1) = \theta_1 = 1 - c_1$  and  $\theta(1, 0) = \theta_2 = 1 - c_2$ . When there is independence between components, *i.e.*, when  $\alpha = 1$ ,  $\theta(\omega, 1 - \omega)$  is a linear function of  $\omega$  whereas the stronger the between-component dependence is, the more the component with the stronger serial dependence, *i.e.*, the larger  $c_j$  dominates. Let  $c_1 > c_2$  (as in Figure 8.2.1), then for  $0 < \omega < \frac{1}{2}$  and small  $\alpha$ ,  $\theta(\omega, 1 - \omega) \approx \theta_1 = 1 - c_1$  and  $\theta(\frac{1}{2}, \frac{1}{2})$  is closer to  $\theta_1$  than  $\theta_2$  - how close and how fast  $\theta(\omega, 1 - \omega)$  moves towards  $\theta_2$  with increasing  $\omega$  is determined by  $\alpha$ .

One declustering method is to take the weighted maximum over all components at each time point and then decluster the resulting univariate sequence. We consider intervals declustering (Ferro and Segers, 2003) on the univariate sequence  $\{Z : Z_t = \max((1 - \omega)X_{t,1}, \omega X_{t,2})\}$  which, using the arguments of Smith and Weissman (1996), has the same extremal index as  $\theta(\omega, 1 - \omega)$ . The intervals estimator (7.1.12) for  $\theta(\omega, 1 - \omega)$  based on the sequence  $\{Z\}_t$ ,  $t = 1, \dots, 5000$  is shown in Figure 8.2.1 ( $\times$ ). We also show the runs estimator with run length 1 ( $\circ$ ) and the empirical estimator from intervals declustering ( $+$ ) for the  $Z$  sequence with both estimates calculated as the number of clusters identified over the total number of exceedances. All the estimators considered perform better when more emphasis is placed on the component with stronger serial dependence (low  $\omega$ ) and underestimate the truth when  $\omega$  is large. The bias in the extremal index estimators is smallest the smaller the difference between the components of  $\mathbf{c}$ , especially for small values of  $\mathbf{c}$ .

### 8.3 The M4 process

#### 8.3.1 Theory

The MARMAX process of §8.2 is restricted to a certain shape of exceedances, that is extremes appear as a sudden spike followed by decreasing points until the next large innovation. We now discuss the M4 process which is a multivariate stationary process for which extremes can follow different shapes known as signature patterns. The M4 (multivariate maxima of moving maxima) process was introduced by Smith and Weissman (1994) and is shown to closely approximate a max stable process. Smith and Weissman (1994) also showed that the limiting distribution of joint maxima is taken from a max stable process, *i.e.*, it is a multivariate extreme value distribution, hence motivating the use of M4 processes to study multivariate extremes.

Let  $\{Z_{si} : s \geq 1, -\infty < i < \infty\}$  be an array of independent unit Fréchet random variables. The M4 process is defined as  $\{\mathbf{X}_i\}_{-\infty < i < \infty}$  where

$$X_{ij} = \max_{s \geq 1} \max_{-\infty < k < \infty} a_{skj} Z_{s,i-k}, \quad j = 1, \dots, d, \quad (8.3.1)$$

for non-negative constants  $\{a_{skj} : s \geq 1, -\infty < k < \infty, 1 \leq j \leq d\}$  such that

$$\sum_{s=1}^{\infty} \sum_{k=-\infty}^{\infty} a_{skj} = 1 \quad \text{for } j = 1, \dots, d.$$

The M4 process has unit Fréchet margins:

$$\mathbb{P}(X_{ij} \leq x) = \mathbb{P}\left(Z_{s,i-k} \leq \frac{x}{a_{skj}}, s \geq 1, -\infty < k < \infty\right) = \prod_{s \geq 1} \prod_{k=-\infty}^{\infty} \exp\left(-\frac{a_{skj}}{x}\right) = e^{-\frac{1}{x}};$$

and is a stationary process with joint common distribution function

$$\begin{aligned} F(\mathbf{x}) &:= \mathbb{P}(\mathbf{X}_i \leq \mathbf{x}) = \mathbb{P}\left(Z_{s,i-k} \leq \frac{x_j}{a_{skj}}, s \geq 1, -\infty < k < \infty, 1 \leq j \leq d\right) \\ &= \mathbb{P}\left(Z_{s,i-k} \leq \min_{1 \leq j \leq d} \left(\frac{x_j}{a_{skj}}\right), s \geq 1, -\infty \leq k \leq \infty\right) \\ &= \prod_{s \geq 1} \prod_{k=-\infty}^{\infty} \exp\left(-\max_{1 \leq j \leq d} \frac{a_{skj}}{x_j}\right) = \exp\left(-\sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \max_{1 \leq j \leq d} \left(\frac{a_{skj}}{x_j}\right)\right). \end{aligned} \quad (8.3.2)$$

Notice that  $\mathbb{P}(\tilde{\mathbf{M}}_n \leq \mathbf{u}_n) = F^n(n\mathbf{x}) = F(\mathbf{x})$ , i.e., the max-stability property is satisfied. Similarly, setting  $l = i - k$  so  $1 - l \leq k \leq n - l$  and  $-\infty < l < \infty$  when  $1 \leq i \leq n$ ,

$$\begin{aligned} \mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n) &= \mathbb{P}(\mathbf{X}_i \leq n\mathbf{x}, 1 \leq i \leq n) \\ &= \mathbb{P}\left(Z_{s,l} \leq \frac{nx_j}{a_{skj}}, s \geq 1, 1-l \leq k \leq n-l, -\infty < l < \infty, 1 \leq j \leq d\right) \\ &= \mathbb{P}\left(Z_{s,l} \leq \min_{1 \leq j \leq d} \min_{1-l \leq k \leq n-l} \left(\frac{nx_j}{a_{skj}}\right), s \geq 1, -\infty \leq l \leq \infty\right) \\ &= \prod_{s \geq 1} \prod_{l=-\infty}^{\infty} \exp\left(-\max_{1 \leq j \leq d} \max_{1-l \leq k \leq n-l} \frac{a_{skj}}{nx_j}\right) \\ &= \exp\left(-\frac{1}{n} \sum_{s \geq 1} \sum_{l=-\infty}^{\infty} \max_{1 \leq j \leq d} \max_{1-l \leq k \leq n-l} \left(\frac{a_{skj}}{x_j}\right)\right). \end{aligned}$$

Smith and Weissman (1994) show that

$$\frac{1}{n} \sum_{l=-\infty}^{\infty} \max_{1 \leq j \leq d} \max_{1-l \leq k \leq n-l} \left(\frac{a_{skj}}{x_j}\right) \rightarrow \max_{-\infty \leq k \leq \infty} \max_{1 \leq j \leq d} \frac{a_{skj}}{x_j}$$

as  $n \rightarrow \infty$  and  $\lim_{n \rightarrow \infty} \mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n) = \exp(-\sum_{s \geq 1} \max_{-\infty \leq k \leq \infty} \max_{1 \leq j \leq d} \frac{a_{skj}}{x_j})$ . Thus the extremal index of M4 process is

$$\theta(\mathbf{x}) = \frac{\log(\lim_{n \rightarrow \infty} \mathbb{P}(\mathbf{M}_n \leq \mathbf{u}_n))}{\log(\lim_{n \rightarrow \infty} F^n(\mathbf{u}_n))} = \frac{\sum_{s \geq 1} \max_{-\infty \leq k \leq \infty} \max_{1 \leq j \leq d} \frac{a_{skj}}{x_j}}{\sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \max_{1 \leq j \leq d} \frac{a_{skj}}{x_j}}.$$

### 8.3.2 Measures of dependence

We now derive the coefficient of asymptotic dependence  $\chi_{\{j, j^*\}}^{(\tau)}$  (8.1.10) for the M4 process starting with the case of 0 lag. We first calculate the copula describing the dependence structure of all components at the same point in time. Let  $\mathbf{u} = (u, \dots, u) \in \mathbb{R}^d$  then

$$\begin{aligned} C(\mathbf{u}) &= \mathbb{P}(F_1(X_{i,1}) \leq u, \dots, F_d(X_{i,d}) \leq u) \\ &= \exp\left(-\sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \max_{1 \leq j \leq d} \left(\frac{a_{skj}}{F_j^{-1}(u)}\right)\right) \quad \text{using (8.3.2)} \\ &= \exp\left(-\sum_{s \geq 1} \sum_{k=-\infty}^{\infty} (-\log u) \max_{1 \leq j \leq d} a_{skj}\right) = u^{\sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \max_{1 \leq j \leq d} a_{skj}}. \end{aligned}$$

Thus, using (8.1.10) with  $C_{j,j^*}^{(0)}(u, u) = \mathbb{P}(F_j(X_{i,j}) \leq u, F_{j^*}(X_{i,j^*}) \leq u)$ , we have for the dependence of components  $j$  and  $j^*$

$$\chi_{j,j^*}^{(0)} = 2 - \frac{\log(C_{j,j^*}^{(0)}(u, u))}{\log u} = 2 - \sum_{s \geq 1} \sum_{-\infty \leq k \leq \infty} \max(a_{skj}, a_{skj^*})$$

and from (8.1.11) in §8.1.4 we obtain the dual measure

$$\bar{\chi}_{j,j^*}^{(0)} = \lim_{u \rightarrow 1} \frac{2 \log(1 - u)}{\log(1 - 2u + u^{\sum_{s \geq 1} \sum_k \max(a_{skj}, a_{skj^*})})} - 1.$$

For any set  $C \subset D$  with  $j^* \in C$  we have

$$1 = \sum_{s \geq 1} \sum_{k=-\infty}^{\infty} a_{skj^*} \leq \sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \max_{j \in C} a_{skj} \leq \sum_{s \geq 1} \sum_{k=-\infty}^{\infty} \sum_{j \in C} a_{skj} = |C|$$

so  $u^d \leq C_{j,j^*}^{(0)}(u, u) \leq u$  and thus  $0 \leq \chi_{j,j^*}^{(0)} \leq 1$  and  $0 \leq \bar{\chi}_{j,j^*}^{(0)} \leq 1$ .

**Example 1:** Let  $a_{skj} = a_{sk} \forall j$  then we have complete dependence,  $C_{j,j^*}^{(0)}(u, u) = u$  and  $(\chi_{j,j^*}^{(0)}, \bar{\chi}_{j,j^*}^{(0)}) = (1, 1)$ .

In practice we concentrate on  $1 \leq s \leq S$  and  $-K_1 \leq k \leq K_2$  where  $S$  is the number of so called *signature patterns* and  $(K_1 > 0, K_2 > 0)$  are ‘determining’ the range of serial dependence. So  $a_{skj} = 0$  for  $k > K_2, k < -K_1$  or  $s < S$ .

**Example 2:** Let  $a_{101} = 1, a_{112} = 1$  and all other  $a_{skj} = 0$  so there is only one signature pattern ( $S = 1$ ) and serial dependence is only between component 1 and component 2 a lag 1 behind component 1 ( $K_1 = 0, K_2 = 1$ ). So  $X_{i1} = \max(a_{101}Z_{1,i}, a_{111}Z_{1,i-1}) = Z_{1,i}$  and  $X_{i2} = Z_{1,i-1}$ . Then we have complete independence (at 0 lag),  $C_{1,2}^{(0)}(u, u) = u^2$  and  $(\chi_{1,2}^{(0)}, \bar{\chi}_{1,2}^{(0)}) = (0, 0)$ .

Now, back in the general case, we consider the joint distribution of *one component* separated

by some lag  $\tau$ :

$$\begin{aligned} \mathbb{P}(X_t \leq x_0, X_{t+\tau} \leq x_\tau) &= \mathbb{P}\left(Z_{s,t-k} \leq \frac{x_0}{a_{s,k}}, Z_{s,t+\tau-k} \leq \frac{x_\tau}{a_{s,k}}, s \geq 1, -K_1 \leq k \leq K_2\right) \\ &= \mathbb{P}\left(Z_{s,t-k} \leq \min\left(\frac{x_0}{a_{s,k}}, \frac{x_\tau}{a_{s,k+\tau}}\right), s \geq 1, -K_1 - \tau \leq k \leq K_2\right), \end{aligned}$$

with  $a_{s,-K_1-1} = \dots = a_{s,-K_1-\tau} = 0$  and  $a_{s,K_2-\tau} = \dots = a_{s,K_2}$ . So

$$\begin{aligned} C^{(\tau)}(u, u) &= \mathbb{P}(X_t \leq F^{-1}(u), X_{t+\tau} \leq F^{-1}(u)) \\ &= \mathbb{P}\left(Z_{s,t-k} \leq (-\log u)^{-1} \min\left(\frac{1}{a_{s,k}}, \frac{1}{a_{s,k+\tau}}\right), s \geq 1, -K_1 - \tau \leq k \leq K_2\right) \\ &= u^{\sum_{s \geq 1} \sum_{k=-K_1-\tau}^{K_2} \max(a_{s,k}, a_{s,k+\tau})}. \end{aligned}$$

Thus,

$$\chi^{(\tau)} = 2 - \sum_{s \geq 1} \sum_{k=-K_1-\tau}^{K_2} \max(a_{s,k}, a_{s,k+\tau}) \quad \bar{\chi}^{(\tau)} = \lim_{u \rightarrow 1} \frac{2 \log(1-u)}{\log(1-2u+u^{2-\chi^{(\tau)}})} - 1. \quad (8.3.3)$$

To evaluate  $(\chi^{(\tau)}, \bar{\chi}^{(\tau)})$  we need to consider various ranges of  $\tau$ . For simplicity consider  $S = 1$  (the following results are the same for  $S \geq 2$ ) and let  $r = K_1 + K_2 + 1 > 1$ . Then,  $\sum_{k=-K_1}^{K_2} a_k = 1$  by definition of the M4 process, and

$$\sum_{k=-K_1-\tau}^{K_2} \max(a_k, a_{k+\tau}) = \begin{cases} \sum_{k=-K_1}^{K_2} a_k = 1 & \text{if } \tau = 0 \\ \sum_{k=-K_1-\tau}^{-K_1-1} a_{k+\tau} + \sum_{k=-K_1}^{K_2-\tau} \max(a_k, a_{k+\tau}) + \sum_{k=K_2-\tau+1}^{K_2} a_k & \text{if } 1 \leq \tau < r \\ \sum_{k=-K_1-\tau}^{-K_1-\tau+r-1} a_{k+\tau} + \sum_{k=-K_1}^{K_2} a_k = 2 & \text{if } \tau \geq r. \end{cases}$$

So by inserting the above results into (8.3.3) we have  $0 \leq \chi^{(\tau)} \leq 1$  for general  $\tau$  with  $(\chi^{(\tau)}, \bar{\chi}^{(\tau)}) = (1, 1)$  when  $\tau = 0$  (complete dependence) and  $(\chi^{(\tau)}, \bar{\chi}^{(\tau)}) = (0, 0)$  when  $\tau \geq r$  (independence). For  $\tau < r$  we have asymptotic dependence as  $\bar{\chi}^{(\tau)} = 1$  with asymptotic dependence level  $\chi_{j^*}^{(\tau)}$ .

It is easy to extend the above ideas to obtain dependence measures over multiple components and lags. Let  $C$  be the set of components of interest at lag 0 and let  $j^*$  be the

component we are interested in at lag  $\tau$  from the components in  $C$ . Then the copula describing the dependence structure between the components of  $C$  and component  $j^*$  at lag  $\tau$  is

$$\begin{aligned} C_{C,j^*}^{(\tau)}(\mathbf{u}) &= \mathbb{P}\left(X_{t+\tau,j^*} \leq F_{j^*}^{-1}(u), X_{t,j} \leq F_j^{-1}(u) \quad \forall j \in C\right) \\ &= u^{\left[\sum_{s \geq 1} \sum_{k=-K_1-\tau}^{K_2} \max_{j \in C} (a_{skj}, a_{s,k+\tau,j^*})\right]}. \end{aligned}$$

So for the dependence structure between pairs of components  $j$  and  $j^*$  ( $j \neq j^*$ ) lag  $\tau$  apart:

$$\begin{aligned} \chi_{j,j^*}^{(\tau)} &= 2 - \sum_{s \geq 1} \sum_{k=-K_1-\tau}^{K_2} \max(a_{skj}, a_{s,k+\tau,j^*}) \\ \bar{\chi}_{j,j^*}^{(\tau)} &= \lim_{u \rightarrow 1} \frac{2 \log(1-u)}{\log(1-2u + C_{j,j^*}^{(\tau)}(u, u))} - 1 = \begin{cases} 1 & \text{if } \tau < r \\ 0 & \text{if } \tau \geq r. \end{cases} \end{aligned}$$

Following the same arguments as above we have in the binary setting  $(\chi_{j,j^*}^{(\tau)}, \bar{\chi}_{j,j^*}^{(\tau)}) = (0, 0)$  for  $\tau \geq r$ , and  $0 < \chi_{j,j^*}^{(\tau)} < 1$ ,  $\bar{\chi}_{j,j^*}^{(\tau)} = 1$  for  $0 < \tau \leq r$ . We do not (necessarily) have complete dependence ( $\chi_{j,j^*}^{(\tau)} = 1$ ) when  $\tau = 0$  since this depends on the dependence between components at 0 lag also.

**Example 2 cont.:** We have  $r = 0 + 1 + 1 = 2$  so for lag  $\tau \geq 2$  we have independence  $(\chi_{j,j^*}^{(\tau)}, \bar{\chi}_{j,j^*}^{(\tau)}) = (0, 0)$ . Recall  $a_{101} = a_{112} = 1$  and all other  $a_{skj} = 0$ . For  $\tau = 1$ ,  $\bar{\chi}_{j,j^*}^{(\tau)} = 1$  and

$$\begin{aligned} \chi_{j,j^*}^{(\tau)} &= 2 - \sum_{k=-1}^1 \max(a_{1kj}, a_{1,k+1,j^*}) \\ &= 2 - (a_{10j^*} + \max(a_{10j}, a_{11j^*}) + a_{11j}) = \begin{cases} 1 & \text{if } j = 1, j^* = 2 \\ 0 & \text{if } j = 2, j^* = 1. \end{cases} \end{aligned}$$

So there is complete dependence between component 1 and component 2 one time step ahead but independence between component 1 and component 2 one time step behind.

### 8.3.3 Cluster size distribution

A cluster of extreme values occurs when there is a large realisation of  $Z_{s^*,t^*}$  for some  $t^*$  and  $s^*$ . When this occurs  $Z_{s^*,t^*}$  dominates the maximum in (8.3.1) in the neighbourhood of  $Z_{s^*,t^*}$  resulting in a sequence of related realisations from time  $t^* - K_1$  to time  $t^* + K_2$  which we refer to as an event. The realisations of  $X_{i,j}$  for  $t^* - K_1 \leq i \leq t^* + K_2$  will be multiples of  $Z_{s^*,t^*}$ , i.e.,  $(a_{s^*,-K_1,j}, \dots, a_{s^*,K_2,j})Z_{s^*,t^*}$  for  $j = 1, \dots, d$ . Here we consider only one signature

pattern,  $S = 1$ , so (8.3.1) simplifies to

$$X_{i,j} = \max_{-K_1 \leq k \leq K_2} a_{kj} Z_{i-k} \quad j = 1, \dots, d. \quad (8.3.4)$$

We define the event  $A_{t,j}$  in component  $j$  starting at time  $t$  by

$$\begin{aligned} A_{t,j} &= \left\{ a_{ij} Z_{t+K_1} = \max_{-K_1 \leq k \leq K_2} a_{kj} Z_{t+K_1+i-k}, \quad i = -K_1, \dots, K_2 \right\} \\ &= \{X_{t+K_1+i,j} = a_{ij} Z_{t+K_1}, \quad i = -K_1, \dots, K_2\}. \end{aligned}$$

Zhang (2002) calculated the probability of this event and Zhang and Smith (2004) show that such an event occurs infinitely often, *i.e.*, there are an infinite number of times at which the process  $\{X_{t,j}\}_{t \leq 1}$  is determined by a single large jump and the signature pattern given by the constants  $a_{i,j}$ ,  $i = -K_1, \dots, K_2$ . The probability of two events within  $K_1 + K_2$  time points of each other is 0; by definition each event is a sequence of  $K_1 + K_2 + 1$  random variables so if the time between two events was less than  $K_1 + K_2 + 1$  then at least one of these random variables would overlap leading to contradiction as they have different definitions according to each event. For example in event  $A_{t,j}$  we have  $X_{t+K_1+i,j} = a_{ij} Z_{t+K_1}$  whereas in event  $A_{t+l,j}$  we have  $X_{t+K_1+i,j} = a_{i-l,j} Z_{t+l+K_1}$ . It is also shown that events are independent when there are further apart in time thus:

$$\mathbb{P}(A_{t,j}, A_{t+l,j}) = \begin{cases} [\mathbb{P}(A_{t,j})]^2 & \text{if } l > K_1 + K_2 \\ 0 & \text{if } 1 \leq l \leq K_1 + K_2. \end{cases}$$

Independent extreme events of the process  $\{\mathbf{X}_t\}_{t \geq 1}$  are formed from the exceedances of  $\mathbf{u}_n$  in the independent events  $\{A_{t_i^*}\}_{i \geq 1} = \bigcap_{j=1}^d \{A_{t_i^*,j}\}_{i \geq 1}$  where  $\{t_i^*\}_{i \geq 1}$  are the times of these events.

We now derive the cluster size distribution,  $\pi^{(\mathbf{x})}(k)$  given by (8.1.8), for the M4 process with  $S = 1$ . Roughly our calculation is as follows:

$$\begin{aligned} \pi^{(\mathbf{x})}(k) &= \lim_{n \rightarrow \infty} \frac{\mathbb{P}\left(N_n^{(\mathbf{x})}([0, r_n/n]) = k\right)}{\mathbb{P}\left(N_n^{(\mathbf{x})}([0, r_n/n]) > 0\right)} \\ &= \lim_{n \rightarrow \infty} \frac{\mathbb{P}(\text{'Exactly } k \text{ time points in an event } A \text{ have an exceedance'})}{\mathbb{P}(\text{'At least one time point in an event } A \text{ has an exceedance'})}. \end{aligned} \quad (8.3.5)$$

In the limit there is no cluster size larger than the range of temporal dependence therefore the

numerator of (8.3.5) is 0 for  $k > K_1 + K_2 + 1$ . For  $k \leq K_1 + K_2 + 1$  we consider the numerator and denominator separately before considering the whole limit. First, the denominator is

$$\begin{aligned}
& 1 - \mathbb{P}(X_{t+K_1+i,j} \leq u_{n,j}, j = 1, \dots, d, i = -K_1, \dots, K_2) \\
&= 1 - \mathbb{P}(a_{ij}Z_{t+K_1} \leq u_{n,j}, j = 1, \dots, d, i = -K_1, \dots, K_2) \\
&= 1 - \mathbb{P}\left(Z_{t+K_1} \leq \min_{j=1, \dots, d} \left( \frac{u_{n,j}}{\max_{i=-K_1, \dots, K_2} (a_{ij})} \right)\right) \\
&= \frac{1}{n} \max_{j=1, \dots, d} \left( \frac{\max_{i=-K_1, \dots, K_2} (a_{ij})}{x_j} \right) + O\left(\frac{1}{n^2}\right). \tag{8.3.6}
\end{aligned}$$

Let  $I^*$  be an arbitrary set of values of size  $k \leq K_1 + K_2 + 1$  from the set  $I = (-K_1, \dots, K_2)$ . For the numerator of (8.3.5) we note that in order to have exactly  $k$  exceedances in an event we need (i)  $X_{t+K_1+i,j} > u_{n,j}$  for at least one component  $j$  and for  $i \in I^* \subset I$  where  $|I^*| = k$  and (ii)  $X_{t+K_1+i,j} \leq u_{n,j} \forall j$  and for  $i \in I \setminus I^*$ . For (i) we have, for all  $i \in I^*$ ,

$$\begin{aligned}
X_{t+K_1+i,j} > u_{n,j} &\Leftrightarrow a_{ij}Z_{t+K_1} > u_{n,j} \quad \text{for at least one component } j \\
&\Leftrightarrow \max_{j=1, \dots, d} (a_{ij}Z_{t+K_1} - u_{n,j}) > 0 \\
&\Leftrightarrow Z_{t+K_1} > \min_{j=1, \dots, d} \left( \frac{u_{n,j}}{a_{ij}} \right) \\
&\Leftrightarrow Z_{t+K_1} > \frac{n}{b_i} \quad \Leftrightarrow Z_{t+K_1} > \frac{n}{b^{(k)}},
\end{aligned}$$

where in the final two steps we have defined  $b_i := \max_{j=1, \dots, d} \left( \frac{a_{ij}}{x_j} \right)$  and ordered these  $b_i$ ,  $i = -K_1, \dots, K_2$  such that  $b^{(m)}$  is the  $m$ th largest. For (ii) we have, for  $k < K_1 + K_2 + 1$  and  $\forall i \in I \setminus I^*$ ,

$$\begin{aligned}
X_{t+K_1+i} \leq u_{n,j} &\Leftrightarrow a_{ij}Z_{t+K_1} \leq u_{n,j} \quad \forall j \\
&\Leftrightarrow Z_{t+K_1} \leq \min_{j=1, \dots, d} \left( \frac{u_{n,j}}{a_{ij}} \right) \\
&\Leftrightarrow Z_{t+K_1} \leq \frac{n}{b_i} \quad \Leftrightarrow Z_{t+K_1} \leq \frac{n}{b^{(k+1)}}.
\end{aligned}$$

When  $k = K_1 + K_2 + 1$  all elements in the event are exceedances, so  $I^* = I$  and there is no upper bound on  $Z_{t+K_1}$  - the derivation of  $\pi^{(x)}(K_1 + K_2 + 1)$  follows below with  $b^{(k+1)}$

replaced by 0. Putting (i) and (ii) together the numerator becomes for  $k \leq K_1 + K_2$

$$\begin{aligned} \mathbb{P}\left(\frac{n}{b^{(k)}} < Z_{t+K_1} \leq \frac{n}{b^{(k+1)}}\right) &= \exp\left(-\frac{b^{(k+1)}}{n}\right) - \exp\left(-\frac{b^{(k)}}{n}\right) \\ &= \frac{1}{n}(b^{(k)} - b^{(k+1)}) + O\left(\frac{1}{n^2}\right). \end{aligned} \quad (8.3.7)$$

Thus, taking the limit as  $n \rightarrow \infty$  of (8.3.7) over (8.3.6) the cluster size distribution for  $k = 1, \dots, K_1 + K_2 + 1$  is

$$\pi^{(\mathbf{x})}(k) = \begin{cases} \lim_{n \rightarrow \infty} \frac{\frac{1}{n}(b^{(k)} - b^{(k+1)}) + O\left(\frac{1}{n^2}\right)}{\frac{1}{n}b^{(1)} + O\left(\frac{1}{n^2}\right)} = \frac{b^{(k)} - b^{(k+1)}}{b^{(1)}} & \text{if } k \leq K_1 + K_2 \\ \frac{b^{(k)}}{b^{(1)}} & \text{if } k = K_1 + K_2 + 1 \\ 0 & \text{otherwise.} \end{cases}$$

We recover the marginal cluster size distribution for the  $j^*$ th component,  $(a_{j^*}^{(k)} - a_{j^*}^{(k+1)})/a_{j^*}^{(1)}$  as stated by Robinson and Tawn (2000), by setting  $x_j = \infty$  for  $j \neq j^*$ .

**Example 3** Consider the 3-dimensional  $M_4$  process (8.3.4) with  $a_{01} = a_{12} = a_{23} = 1$  and  $a_{ij} = 0$  otherwise. Then  $X_{t1} = Z_t, X_{t2} = Z_{t-1}, X_{t3} = Z_{t-2}$  for all  $t$ , i.e., at a given time point the process is independent across components, however, there is dependence across components with a time delay (component 2 is the same as component 1 in the previous time step etc.). We have  $K_1 = 0$  and  $K_2 = 2$  so  $\pi^{(\mathbf{x})}(k) = 0$  for  $k > 3$ . Ordering  $x_1, x_2, x_3$  such that  $x^{(1)} \geq x^{(2)} \geq x^{(3)}$  we have  $b^{(i)} = \frac{1}{x^{(4-i)}}$  and

$$\begin{aligned} \pi^{(\mathbf{x})}(1) &= 1 - \frac{x^{(3)}}{x^{(2)}} \\ \pi^{(\mathbf{x})}(2) &= x^{(3)} \left( \frac{1}{x^{(2)}} - \frac{1}{x^{(1)}} \right) \\ \pi^{(\mathbf{x})}(3) &= \frac{x^{(3)}}{x^{(1)}}. \end{aligned}$$

(i) Consider  $x_1 = x_2 = x_3$ . Then  $\pi^{(\mathbf{x})}(1) = \pi^{(\mathbf{x})}(2) = 0$  and  $\pi^{(\mathbf{x})}(3) = 1$ , i.e., every cluster of extremes is of size 3 with one exceedance in each component at each time point.

(iii) Consider  $x_1 = 2y, x_2 = x_3 = y$ . Then  $\pi^{(\mathbf{x})}(1) = 0$  and  $\pi^{(\mathbf{x})}(2) = \pi^{(\mathbf{x})}(3) = \frac{1}{2}$ , i.e., half of the clusters are of size 3 with one exceedance in each component at each time point and the other half of clusters are size 2 consisting of exceedances in components 2 and 3 only.

(ii) Consider  $x_1 = 3y$ ,  $x_2 = 2y$ ,  $x_3 = y$ . Then  $\pi(\mathbf{x})(1) = \frac{1}{2}$ ,  $\pi(\mathbf{x})(2) = \frac{1}{6}$ ,  $\pi(\mathbf{x})(3) = \frac{1}{3}$  and the expected cluster size is  $\frac{11}{6}$ .

**Example 4** Consider the 2-dimensional M4 process (8.3.4) with  $a_{01} = a_{11} = a_{21} = \frac{1}{3}$ ,  $a_{32} = a_{42} = \frac{1}{2}$  and  $a_{ij} = 0$  otherwise. Then  $X_{t1} = \frac{1}{3} \max(Z_t, Z_{t-1}, Z_{t-2})$  and  $X_{t2} = \frac{1}{2} \max(Z_{t-3}, Z_{t-4})$ . We have  $K_1 = 0$  and  $K_2 = 4$  so  $\pi(\mathbf{x})(k) = 0$  for  $k > 5$ . We have  $b_0 = b_1 = b_2 = \frac{1}{3x_1}$  and  $b_3 = b_4 = \frac{1}{2x_2}$ .

(i) If  $\frac{1}{3x_1} < \frac{1}{2x_2}$ , then  $\pi(\mathbf{x})(1) = \pi(\mathbf{x})(3) = \pi(\mathbf{x})(4) = 0$ ,

$$\pi(\mathbf{x})(2) = 1 - \frac{2x_2}{3x_1} \quad \pi(\mathbf{x})(5) = \frac{2x_2}{3x_1}$$

and  $\mathbb{E}[K] = 2 + \frac{2x_2}{3x_1}$ . So clusters occur as exceedances in component 2 (so  $k = 2$ ) or as exceedances in both components ( $k = 2 + 3 = 5$ ).

(ii) If  $\frac{1}{3x_1} > \frac{1}{2x_2}$ , then  $\pi(\mathbf{x})(1) = \pi(\mathbf{x})(2) = \pi(\mathbf{x})(4) = 0$ ,

$$\pi(\mathbf{x})(3) = 1 - \frac{3x_1}{2x_2} \quad \pi(\mathbf{x})(5) = \frac{3x_1}{2x_2}$$

and, letting  $K$  denote the random cluster size,  $\mathbb{E}[K] = 3 + \frac{3x_1}{x_2}$ . So clusters occur as exceedances in component 1 (so  $k = 3$ ) or as exceedances in both components ( $k = 2 + 3 = 5$ ).

(iii) If  $\frac{1}{3x_1} = \frac{1}{2x_2}$ , we have  $\pi(\mathbf{x})(1) = \pi(\mathbf{x})(2) = \pi(\mathbf{x})(3) = \pi(\mathbf{x})(4) = 0$  and  $\pi(\mathbf{x})(5) = 1$ . So all clusters consist of exceedances in both components and are of size 5 in the limit.

All of the derivations in this subsection have been based on the M4 process with  $S = 1$ . In the case of multiple signature patterns there will be a particular  $s$  associated with a large  $Z_{s,t+K_1}$  realisation so each extreme event will be dominated by a particular signature pattern corresponding to this  $s$ . Different extreme events will be dominated by a different values of  $s$  and so events will follow a range of signature patterns. The cluster size distribution will reflect all the signature patterns. For example, if there are two signature patterns with one short term serial dependence and one long term then the clusters formed will be a mixture of these.

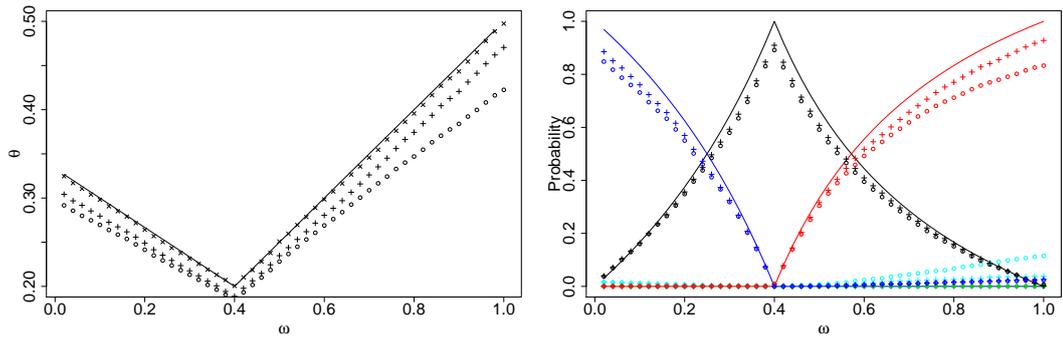


Figure 8.3.1: *Extremal index  $\theta(\omega, 1 - \omega)$ ,  $0 < \omega < 1$  and probabilities of cluster sizes (1-green, 2-red, 3-blue, 4-cyan, 5-black) for Example 2 of §8.3.3. True extremal index and cluster size probabilities (solid line), runs estimator with run length 1 (o), empirical estimator from intervals declustering (+) and the intervals estimator (x). Estimates are averages over 100 repeated simulations of the  $M_4$  process with 5000 time points.*

### 8.3.4 $M_4$ simulations and estimation of $\theta(x)$

We now simulate some bivariate  $M_4$  processes and discuss the corresponding multivariate extremal index for varying thresholds. We apply and evaluate declustering methods to these simulations and compare extremal index estimates and empirical cluster size distributions.

As in §8.2.5 we set  $\tilde{x}_1 = \omega$  and  $\tilde{x}_2 = 1 - \omega$  so we have the extremal index in terms of one parameter. The bivariate  $M_4$  multivariate extremal index is, for  $0 < \omega < 1$ ,

$$\theta(\omega, 1 - \omega) = \frac{\sum_{s \geq 1} \max_{-K_1 \leq k \leq K_2} \max(a_{sk1}(1 - \omega), a_{sk2}\omega)}{\sum_{s \geq 1} \sum_{k=-K_1}^{K_2} \max(a_{sk1}(1 - \omega), a_{sk2}\omega)}.$$

Figure 8.3.1 shows the extremal index,  $\theta(\omega, 1 - \omega)$ , and estimators thereof against  $\omega$  based on simulations of the bivariate  $M_4$  process, described in §8.3.3 Example 4. We also show the cluster size distribution and estimates thereof in Figure 8.3.1 right panel. Rewriting the cases in Example 4 in terms of  $\omega$  we have (i)  $\frac{1}{3x_1} < \frac{1}{2x_2} \Leftrightarrow \omega > \frac{2}{5}$ , (ii)  $\frac{1}{3x_1} > \frac{1}{2x_2} \Leftrightarrow \omega < \frac{2}{5}$  and (iii)  $\frac{1}{3x_1} = \frac{1}{2x_2} \Leftrightarrow \omega = \frac{2}{5}$ . The lowest extremal index value occurs when  $\omega = \frac{2}{5}$ , for this  $\omega$  all clusters are of size 5 in the limit and so the extremal index is 0.2. In case (ii) we found that in the limit clusters are of size 5 with probability  $\frac{3x_1}{2x_2} = \frac{3\omega}{2(1-\omega)}$  and are of size 3 (shown in blue) otherwise. The mean cluster size in case (ii) is  $\frac{3}{1-\omega}$  so the extremal index is  $\frac{1-\omega}{3}$ , a decreasing function in  $\omega$ , and when  $\omega = 0$  we recover the extremal index for component 1. In contrast, when in case (i), the extremal index is an increasing function of  $\omega$  and the extremal index for component 2 is recovered at  $\omega = 1$ .

As in §8.2.5 we consider intervals and runs declustering (Ferro and Segers, 2003) on the

univariate sequence  $\{Z : Z_t = \max((1 - \omega)X_{t,1}, \omega X_{t,2})\}$  with length 5000. We show the mean over 100 such simulations of the intervals estimate (7.1.12) ( $\times$ ), the runs estimate with run length 1 ( $\circ$ ) and the empirical estimate from intervals declustering ( $+$ ) for each  $\omega$ . The intervals estimates are close to the truth whereas the empirical estimates underestimate the extremal index. The underestimation of the empirical estimates is due to the too small/large proportion of small/large clusters as a result of the declustering methods – this can be seen particularly when  $\omega$  is large as the probability of a cluster of size 2 is underestimated using the runs method by almost 0.1 and there is a non-zero probability estimate of a cluster of size 4.

## 8.4 Summary

In Part III we have discussed the extension of classical extreme value theory to sequences with serial dependence and to multiple dimensions. Having techniques and understanding of multivariate temporally dependent extremes is important for companies such as JBA since losses, caused by extreme river flows, must be aggregated across a region (thus multiple locations/dimensions) over some period of time. The extreme events modelled need to capture both the dependence structure between different locations and at different time lags. Here we explored the theory of multivariate temporally dependent extremes with focus on measures of dependence; we investigated the multivariate extremal index and cluster size distribution and extended the coefficient of asymptotic dependence to describe the dependence between two sets of components some lag  $\tau$  apart, specifically for pairs of components. We explored two stationary processes, the MARMAX process (ARMAX in one dimension) and the M4 process and derived the multivariate extremal index and coefficient of asymptotic dependence for these processes. For the M4 process we also derived the cluster size distribution expanding on the univariate cluster size distribution in Robinson and Tawn (2000). In simulations of the MARMAX and M4 processes we focussed on estimation of the multivariate extremal index using univariate declustering methods. In the thesis discussion (§9.2.3) we discuss the difficulties of declustering in the multivariate setting and empirical estimation of the cluster size distribution illustrated with an example bi-variate M4 process.

# Chapter 9

## Conclusion

### 9.1 Summary

In this thesis we explored extreme value problems in the analysis of river flow data, particularly after flood events, and the efficient estimation of loss from such events. The thesis covered three main topics: Extreme values under stopping rules (Part I); efficient loss estimation (Part II) and extremes of multivariate dependent sequences (Part III). Part I was motivated by the stochastic nature of the time an analysis takes place due to increased interest after a flood event. We studied the extent of the inference problems due to the variable sample size under such ‘stopping rules’ and developed and evaluated new conditional-likelihood methods which appear to overcome these problems. Part II focused on the improving the efficiency of the standard (Monte Carlo) procedure used to simulate losses and estimate return levels from property and flood event information. We focussed on estimation of return levels with high return periods, using concentration inequalities as part of a novel simulation procedure which is faster than the standard procedure. We also developed our own tighter concentration inequalities. Finally, modelling flood events can be quite complicated due to the presence of extreme values occurring at different locations, with strong dependence between neighbouring locations, and short-term serial dependence (*e.g.*, similar river flow values on consecutive days). This motivated our exploration of the extreme value theory of multivariate and dependent sequences in Part III.

#### 9.1.1 Stopping bias

In Chapter 3 we showed that return-level estimators based on the standard likelihood are positively biased when sampling from the GEV distribution using a fixed-threshold stopping rule. We proposed two new likelihoods, full conditioning and partial conditioning, based on conditioning upon the stopping threshold. We found that full conditioning almost eliminates

the bias in the return-level estimators and gives close to the desired coverage but at the cost of large RMSE. In most cases we found that conditioning on the final observation exceeding the stopping threshold (partial conditioning) results in return-level estimates with the lowest RMSE despite the estimator being negatively biased. We noted that the interval widths using the full and partial conditional likelihoods are smaller the closer the stopping threshold is to the final observation since the occurrence of the final exceedance becomes more informative on the tail of the distribution. We investigated the confidence intervals for the return-level estimates in more depth in Chapter 4 concentrating on data sets similar to the Lune data set. In particular, in §4.5, we discussed the issues that arise with the profile-likelihood based intervals when the final observation is close to the stopping threshold.

In Chapter 4 we also considered the standard bootstrap confidence interval and bias-reducing variations thereof including our own version based on Efron's bias correction. We found that the standard bootstrap method results in narrow intervals and poor coverage whereas the bootstrap variations generally increase coverage but with confidence interval widths comparable to the profile likelihood based interval widths; the latter is much faster computationally so remains the preferred confidence interval method. For the variable-threshold stopping rule we developed an importance-weighted bootstrap to create confidence intervals including an extension to multiple exceedances of the stopping threshold. These importance sampling confidence intervals are narrower than the profile likelihood based intervals however they are highly negatively biased due to the negative bias in the return-level estimators and the resulting bootstrap samples.

Overall, the conditioning estimators we presented in Chapter 3 outperform the standard estimator when the decision to analyse data at a particular time was triggered by what was perceived to be a large observation. For the fixed-threshold stopping rule, partial conditioning has the best combination of RMSE and coverage for a range of  $\xi$  with moderate stopping threshold and particularly when the distribution is heavy tailed, as is the case for most UK rivers (CEH, 1999). For the variable-threshold stopping rule, full conditioning provides the best balance of coverage and low RMSE.

### 9.1.2 Loss Estimation

In Part II we discussed the estimation of the return levels of the loss distribution and our approaches to increase the computational efficiency of this estimation process. In Chapter 5

we reviewed known concentration inequalities, particularly for sums of independent bounded random variables, noting the issues with Jebara's Bennett refinement in our loss estimation setting. To better handle cases where the upper bounds on the magnitude of the individual terms in the sum are not uniform, we developed novel bounds, Chernoff-Hoeffding+ and ++ and Bennett+, improving on the Chernoff-Hoeffding and Bennett bounds respectively. In each case we exploited the convexity of a key function (see Lemma 5.4.1.3). Our Bennett+ concentration bound was the most consistent in its good performance out of all the concentration inequalities considered, particularly when the independent random variables had differing upper bounds and a very small expectation.

In Chapter 6 we used the concentration inequalities discussed in Chapter 5 as part of a novel approach, the 'exclude method', to improving the computational efficiency of the loss estimation procedure. Our method reduces the number of simulations by not simulating from (*i.e.*, excluding) years which are very unlikely to change the estimate of the quantile. The years to be discarded are determined using an upper bound on the non-equivalence probability; the probability that the  $t$ -year return-level estimate using the 'exclude method' differs from the  $t$ -year return-level estimate obtained when simulating from all the years. This upper bound is a function of concentration bounds on the losses in each year of the event set. For high return periods the results using the exclude procedure are very promising - substantial computational effort can be saved while ensuring the return-level estimate is almost always the same as that using the standard procedure.

The exclude method was developed with the estimation of high return levels in mind - for small return periods there is no or little computational saving. However, our criterion for deciding on the years to discard is very conservative in two ways: (i) we imply that only an estimate exactly equal to the standard estimate is acceptable rather than accepting being sufficiently close to the standard estimate or, even better, the true return level and (ii) the upper bound on the non-equivalence probability involved many loosening steps (*e.g.*, (6.2.4)) and is, at best, only as tight as the concentration inequalities. These are both topics for further work. The former problem requires some sort of measure of distance to the standard estimate or true return level. For (ii) a tighter upper bound could be found if we take the return-level estimate to be the median of the quantile estimates of each simulation rather than the mean. This would allow the  $m/2 - 1$  smallest quantile estimates to differ to those using the standard procedure.

In §6.3 we discussed a method to estimate the return levels with low return periods using the normal approximation and Berry-Esseen error. Using the normal approximation to the yearly loss distributions saved substantial computational effort and the return-level estimates based on such simulations performed well. The issue with just using the normal approximation is that we have no measure of how close the return-level estimates will be to the standard return-level estimates (or to the true return level). This motivated the incorporation of the Berry-Esseen bound into a loss estimation procedure using the normal approximation. Unfortunately, for many years in our test event set the error in the normal approximation was too large for the Berry-Esseen procedure to work well, however, this should improve with larger portfolios.

### 9.1.3 Extremes of dependent and multivariate sequences

Finally, in Part III we discussed the extension of classical extreme value theory to sequences with serial dependence and to multiple dimensions. As part of this work we investigated the multivariate extremal index and cluster size distribution and extended the coefficient of asymptotic dependence to describe the dependence between two sets of components some lag  $\tau$  apart, specifically for pairs of components. We explored two stationary processes, the MARMAX process (ARMAX in one dimension) and the M4 process and derived the multivariate extremal index and coefficient of asymptotic dependence for these processes. For the M4 process we also derived the cluster size distribution. In simulations of the MARMAX and M4 processes we focussed on estimation of the multivariate extremal index using univariate declustering methods. More exploration of declustering in the multivariate setting is the subject of further work.

## 9.2 Possible further work

### 9.2.1 Stopping Bias

In Chapter 3 we considered two stopping rules; one based on a fixed threshold and one based on past return-level estimates since the decision to ‘stop’ and analyse data would in part be based on both past experience and thresholds set due to current infrastructure. Our work attempted to simplify the true decision making procedure by using stopping rules based on the occurrence of a *single* large observation exceeding some threshold. An analysis may instead be

prompted by a prolonged period of quite large (but not necessarily ‘extreme’) observations or the observation of large values at many locations simultaneously (or within in a short period of time). Such stopping criteria requires more complex, multivariate, analysis to account for serial dependence and the dependence between observations at nearby locations building on the theory of Chapters 7 and 8.

Another area of further work is the investigation of stopping rules on data with a long-term trend, such as river flows gradually increasing over the years. The fixed-threshold stopping rule may be less appropriate in this setting, in particular it might become necessary to change the threshold after a certain number of years. Nonetheless, doing this is probably not too unrealistic since, for example, the height of a flood defence might be increased if there has been evidence of higher flow in recent years. The variable-threshold stopping rule is more robust to data with an underlying trend as it is directly a function of the observed data.

Finally, as we saw in Chapter 4, in many cases the full-conditioning method leads to wide profile-likelihood based confidence intervals. Bootstrap-based intervals were found to have smaller intervals but poorer coverage and the contest between interval width and coverage was seen in all the various bias-corrected methods compared. However, it appears reasonable that some reduction in width should be possible without drastically reducing the coverage; more investigation into confidence intervals in the stopping rule setting could be useful.

## 9.2.2 Loss simulation and return level estimation

### Improving concentration inequalities in the loss estimation setting

The concentration inequalities we have discussed and developed apply to *any* sum of independent, bounded random variables. More of our knowledge of the actual distribution (weighted sum of Betas) could also be employed to find a tighter bound for our particular loss estimation setting rather than the general setting considered so far. Also more research directions could be opened up since the bound itself does not need to be tractable to be useful in reducing computational efficiency.

A possible extension to the Chernoff-Hoeffding inequality is to split the random variables into groups with different upper bounds, using the information of these different upper bounds to obtain tighter concentration bounds than just using  $c_{max}$  as the upper bound for all  $X_i$ . The indices of the random variables can be split into  $J$  mutually disjoint sets such that for every  $j \in 1, \dots, J$  we have  $c_{M_{j-1}} < c_i \leq c_{M_j} \forall i \in I_j$  where  $0 = c_{M_0} < c_{M_1} < \dots < c_{M_J} = c_{max}$

and  $X_i \leq c_{M_j} \forall i \in I_j$ . In the Chernoff-Hoeffding proof the step from (5.3.14) to (5.3.15) is achieved by the bound  $e^{\lambda c_i} < e^{\lambda c_{max}}$ . Now using the bounds  $e^{\lambda c_i} < e^{\lambda c_{M_j}}$  for every  $i \in I_j, j \in 1, \dots, J$  we arrive at the bound:

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq \left(1 + \frac{1}{n} \sum_{j=1}^J n_j \bar{p}_j (e^{\lambda c_{M_j}} - 1)\right)^n e^{-\lambda n(\bar{c} + t)}, \quad (9.2.1)$$

where  $n_j = |I_j|$  and  $\bar{p}_j = \frac{1}{n_j} \sum_{i \in I_j} p_i$ . It is likely that numerical optimisation will be required to minimise (9.2.1) over  $\lambda > 0$ , however, this is only an issue if said optimisation is computationally expensive.

In our results of §6.1 we saw that in our setting the Chernoff-Hoeffding bounds and variations thereof are poorer (often substantially so) than the Bennett, Bennett+ and Bernstein bounds; it is unlikely that the improvements conjectured here will lead to bounds as ‘good’ as these, nevertheless, they could be useful in other settings and a variation of this idea could also be applied to other inequalities such as the Bennett inequality. On the other hand, the Bennett(+) concentration bounds are fairly tight and the resulting discard sets using the exclude procedure are almost as large as they can be for high return periods. This suggests that any improvement to the concentration inequalities is unlikely to make much difference to the performance of the exclude procedure. A better way forward would be to consider ways of improving the exclude procedure and/or completely different methods. We now consider an idea to improve the former and focus on the latter in subsequent subsections.

### Separating risk and event losses by variance

The exclude procedure of §6.1 reduced the computational cost of return-level estimation by reducing the number of years simulated. We now focus instead on reducing the computational effort needed to simulate the total loss in a given year by reducing the number of risk and event losses simulated.

There are a large number of risk and event combinations for which the simulated water depths are very small and so contribute little to the total loss of the event. Simulating these increases the computation time, however, they cannot simply be ignored since these small losses aggregated over a large number of such events may contribute substantially to the total loss. One consideration could be to circumvent this computational effort by bounding the losses or crudely replacing losses by their expected value from such event and risk combinations and

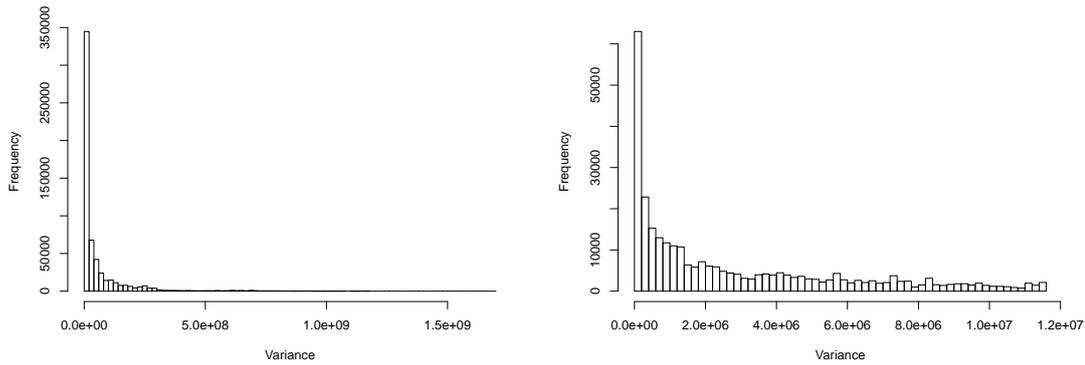


Figure 9.2.1: Histograms of variance of the loss from event and subrisk combinations. Left: All combinations. Right: Combinations with variance lower than 11465217, the median variance.

simulate losses only from more volatile event and risk combinations. Our idea is to separate the risks (for each event) into ‘low-variance risks’ and ‘high-variance risks’, simulating only losses from the ‘high-variance risks’. We denote the total loss in year  $y$  from high-variance risks by  $S_y^{hi} = \sum_{e \in \mathcal{E}_y} \sum_{r \in \mathcal{R}} L_{e,r} \mathbb{1}_{\{\text{Var}(L_{e,r}) \geq \gamma\}}$  and  $S_y^{lo} = \sum_{e \in \mathcal{E}_y} \sum_{r \in \mathcal{R}} L_{e,r} \mathbb{1}_{\{\text{Var}(L_{e,r}) < \gamma\}}$  for low-variance risks, where the threshold,  $\gamma$ , can be chosen, for example, by inspecting the histogram of variances of the loss from all event and risk combinations (which are given in the damage distribution table).

One way to include the contribution of ‘low-variance’ risks without simulation would be to use, for each year, the expected total loss of all such risks in the year,  $\mathbb{E}[S_y^{low}]$ , rather than simulating  $m$  times from the loss distributions of each ‘low-variance’ event and risk combination. This would reduce the simulations needed in Step 1 of the standard procedure (§5.1.6).

Another approach which uses more information from the ‘low variance risks’ whilst not resorting to computationally expensive simulation is to find bounds on the loss contribution of such risks in each year using concentration inequalities. The sum of the simulated losses (from ‘high variance’ risks) and the bound on loss from low variance risks in a particular year can be treated as a simulation of total loss in that year. Then one can find upper and lower  $t$ -year return-level estimates,  $q_t^{ub,(i)}$  and  $q_t^{lb,(i)}$ , as in Steps 2b and 3 of the standard procedure (in the same manner as in the Berry-Esseen procedure).

Both approaches described above could possibly be used in conjunction with the exclude procedure, however some thought would be needed to find the discard set; the return-level estimates *will* differ from the standard estimate and so one would need an alternative to

the non-equivalence probability, rather some sort of measure of ‘closeness’ to the standard estimate, to obtain a discard set.

If a large number of risks can be classified as ‘low variance’ then the methods in this subsection may reduce the computational cost substantially. For example, if 50% of subrisk and event combinations have low-variance loss then we have  $0.5Nm$  simulations rather than  $Nm$  (recall  $N$  is the total number of possible flood event and subrisk combinations over all years). A disadvantage of this method is that there is some arbitrariness in what we decide to be high or low variance. There would be more thought and testing needed to determine an appropriate threshold above which we consider the variance to be ‘high’. Some sensitivity analysis would be required since there would be a trade off between computational efficiency and the estimation performance depending on the threshold chosen.

### Importance Sampling

In Chapter 2 we presented importance sampling, a variance reduction method. Can an importance sampling scheme enable us to estimate return levels more efficiently by sampling more frequently from the high loss region of the yearly loss distribution? The standard procedure used for loss estimation (§5.1) is to estimate the cdf of the yearly loss distribution and invert it to find quantile estimates. Let  $f(s)$  be the target distribution, the (univariate) distribution of total loss in *any single* year. Then the unbiased importance sampling estimator of the cdf corresponding to that used in the standard procedure is, for simulation  $i$ :

$$\hat{I}_t^{IS(i)} = \frac{1}{n_y} \sum_{y=1}^{n_y} \frac{n_y}{n_y + 1} \mathbb{1}_{\{s_y^{(i)} \leq q_t\}} w(s_y^{(i)}),$$

where  $(s_1^{(i)}, \dots, s_{n_y}^{(i)})$  is an independent sample of size  $n_y$  from  $q(s)$ , the proposal distribution of yearly loss, and  $w(s) = \frac{f(s)}{q(s)}$  is the importance weight. The proposal distribution needs to ‘cover’ the target distribution, have heavier tails and preferably be easy to simulate from. The best proposal distributions will have high density around the quantiles of interest so there will be a high probability of sampling higher losses.

A major obstacle of such an importance sampling scheme shown here is the calculation of the importance weights. In our loss estimation setting we do not know the target distribution,  $f(s)$ , and so cannot calculate weights. Therefore in order to use importance sampling we would need to focus on proposals with a known relationship to the target distribution. One

possibility, for which the importance weight is tractable, is to focus on the part of the yearly loss distribution which determines whether there is a flood or not for each risk and event combination in a particular year. Recall that the true probability of flooding is  $p_{e,r}$  for event  $e$  and risk  $r$ . We can instead sample  $q_{e,r} = \text{logit}^{-1}(\log(p_{e,r}/(1 - p_{e,r})) + c)$  for some positive tuning constant  $c$  so more high probabilities of flooding are sampled. For each risk and event combination this skew towards flood events is then accounted for by the weight  $p_{e,r}/q_{e,r}$  if we simulate that there is a flood, and  $(1 - p_{e,r})/(1 - q_{e,r})$  if not. Then the product of these ratios over all events and risks in a year give us the importance weight.

### Latin Hypercube Sampling and Conditional Monte Carlo

In Chapter 2 we reviewed two other variance reduction methods: conditional Monte Carlo and Latin Hypercube sampling. The latter method should be easily applied to our loss estimation problem by considering the random elements contributing to the yearly loss *i.e.*, we would need to write  $S$  as some function of  $\mathbf{U}$  where the elements of  $\mathbf{U}$  a standard uniform random variables determining the year, the probability of each event and subrisk combination etc. Note that in some sense our procedure is already stratified over the element of  $\mathbf{U}$  describing the year since we sample once from each year in the 1000 year set.

For conditional Monte Carlo in our loss estimation setting it is not clear what auxiliary random variable,  $Y$ , we could condition on. The auxiliary random variable needs to be something that doesn't contain all the information about  $S$  (for example, a partial sum of losses in the year) but crucially we want  $\mathbb{P}(S \leq s|Y)$  to be easy to compute. This would require more careful thought for future work.

### Splitting

Another approach to make loss estimation more efficient is to consider ways to steer simulations to the region of interest (high loss). For a particular year, rather than simulating the loss for an entire portfolio this process could be split up over a number of sub-portfolios (*e.g.*, streets or towns as opposed to the whole country). In analogy to splitting methods (§2.2.3), the 'process',  $X_y(t)$ , is the cumulative total loss over the sub-portfolios in a particular year and at each 'time step',  $t$ , the loss due to a randomly selected sub-portfolio is realised and added onto the current total loss. Thus 'time' in this setting is the number of sub-portfolios added. Since sub-portfolios are randomly selected we expect the gradient to be approximately

constant so trajectories can be used to decide which years are most likely to achieve high total losses over the whole portfolio.

A simple method would be to discard the simulations from the half of the years with the lowest cumulative loss after a certain number of sub-portfolios have been added. Alternatively, since the less promising years may still contribute to the quantile of interest, instead of simply discarding all the simulations from the years with lower losses we could reduce the number of simulations from these years.

An extension to this idea, related to the approach of §9.2.2, could be to split the sub-portfolios into sets of low, medium and high variance portfolios and simulate from the high variance portfolios first. This would result in different gradients for the first, second and final section of the simulation but the gradients within each segment would be approximately constant if the sub-portfolios are randomly selected within these sets. With this procedure it may be feasible to use concentration inequalities to bound the contribution to the loss distribution of the low variance sub-portfolios thus reducing the number of simulations further.

### 9.2.3 Extremes of dependent and multivariate sequences

In our work on extremes of multivariate dependent sequences we focused on the properties of such sequences and estimation of measures such as the multivariate extremal index. In the univariate setting we described (and applied to an ARMAX process) declustering methods which allow us to identify independent events over a time series with the maxima of the events following a generalised Pareto distribution. Identification of events is much more complex in the multivariate setting. Firstly it is not clear how an event/cluster should be formed across margins – do they necessarily need to occur in the same time frame? How do we account for events happening independently in each margin? Or multiple events in one margin during one event in another margin? Secondly it is unclear what information we can/should extract from the identified events – are we interested in the componentwise maxima or the values of all components simultaneously where (at least) one component has its maximum value in the event?

For a brief discussion we illustrate in Figure 9.2.2 declustering on the bi-variate M4 process,  $\{(X_{t,1}, X_{t,2})\}$ , (8.3.4) with  $a_{01} = 0.7, a_{11} = 0.3$  and  $a_{i2}$  non-zero for  $i = 0, \dots, 11$  giving a signature pattern (shown in Figure 9.2.3) of a large peak followed by a smaller peak. The areas in blue are the clusters found using intervals declustering (§7.1.3) on the

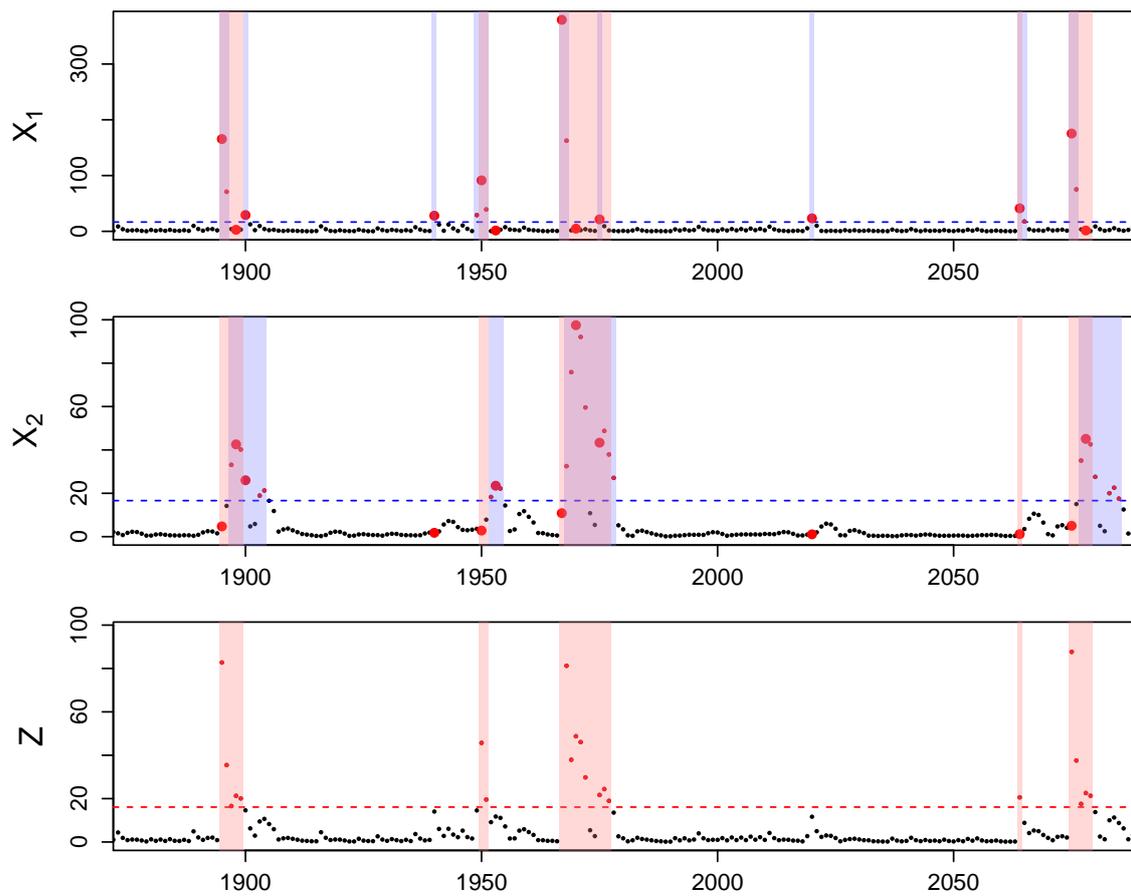


Figure 9.2.2: *Bi-variate  $M_4$  process declustered component-wise (blue areas) and using intervals declustering on  $Z$  with  $\omega = 0.5$  (red areas).*

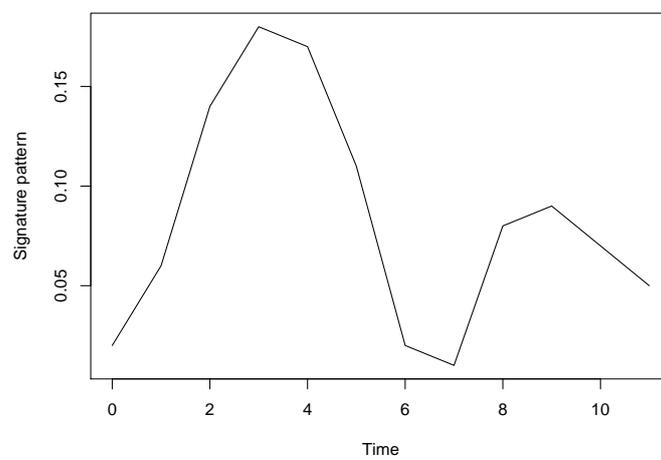


Figure 9.2.3: *Signature pattern for component 2 in declustering examples.*

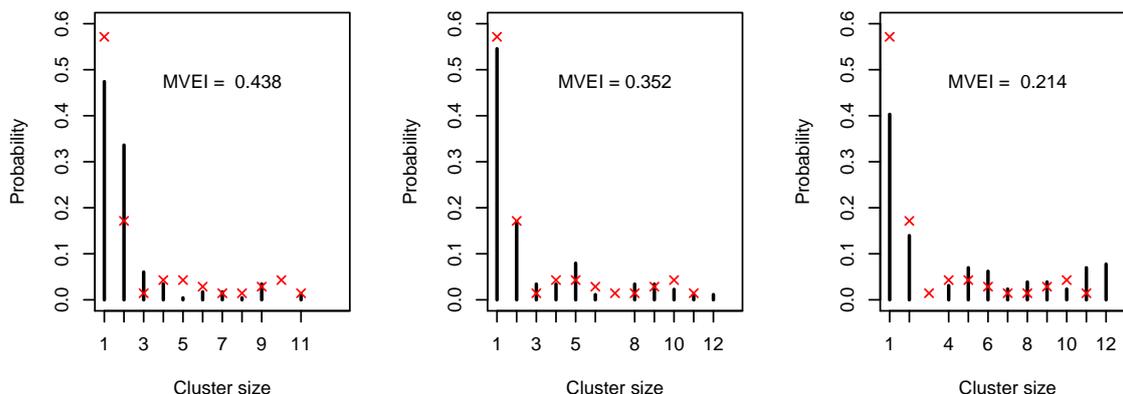


Figure 9.2.4: Empirical cluster size distributions based on clusters formed from a simulated  $M_4$  process. Left: Univariate intervals declustering in each component, middle: Intervals declustering on  $Z$ , right: Nadarajah's declustering method. True cluster size probabilities are shown by red crosses. The multivariate extremal index (MVEI) estimates based on these empirical distributions are also shown.

individual components whereas in red are the clusters found using intervals declustering on  $\{Z_t : Z_t = \max((1 - \omega)X_{t,1}, \omega X_{t,2})\}_{t=1}^{5000}$  where  $\omega = 0.5$ . The pairs shown in the top two panels of Figure 9.2.2 as thick red dots are independent cluster peaks in each margin with the corresponding value in the other margin. These pairs are not necessarily independent temporally despite the cluster maxima being independent in each margin, rather we may have multiple pairs from the same independent cluster, for example the pairs at 1950 and 1953.

The component with stronger serial dependence (component 2 here) dominates the declustering procedure; here the two independent clusters identified between 1965 and 1980 in component 1 are merged into one cluster when declustering the  $Z$  sequence and the clusters at 1900, 1940 and 2020 are no longer identified. So if one defines events over all components by the start and end times of the clusters identified using intervals declustering on the  $Z$  sequence, then information is lost about independent events in the margin with weaker serial dependence. Nadarajah (2001) presents two multivariate declustering methods which aim to capture the missing information of such events by (i) grouping events which occur too close together in time into a smaller set of independent events and (ii) by forming lower dimensional sets of independent events over components with lower temporal dependence which would be 'lost' when creating events over all components.

In Figure 9.2.4 we compare the empirical cluster size distributions and extremal index

estimates for some of the declustering methods discussed here. The true extremal index for our example is 0.3658 and the true probabilities of each cluster size are marked by red crosses. In the first panel we simply took the cluster sizes to be the union of the sizes in each component separately hence assuming both independence between components and that the events in each component did not overlap. Clearly this is a false assumption and this is reflected in a slight bias towards smaller cluster sizes and an overestimate of the extremal index. The second panel shows the cluster size distribution from intervals declustering on the sequence  $Z$ . This distribution is fairly close to the true cluster size distribution and the empirical extremal index estimate is only slightly too small. The final panel shows the distribution of cluster sizes determined using Nadarajah's declustering procedure. This procedure extracts independent pairs as events and so it is not clear how to define a cluster size, here the size was taken to be the difference of time from the first exceedance in one of the grouped events to the last exceedance. This leads to a bias towards larger clusters, due to the grouping of close events, and so a negatively biased extremal index estimate.

Finally, we saw in Figure 8.2.1 for MARMAX simulations that estimation of the multivariate extremal index is poorer when there is low component-wise dependence and a high difference in temporal dependence across margins. Low component-wise dependence could occur as a result of lag between the process in each margin so it seems sensible to reduce this lag or adjust for it some way in estimation of the extremal index and/or declustering methods. It may be useful to develop a scheme which uses the measure  $\chi^{(\tau)}$  (§8.1.4) since empirical estimates of  $\chi^{(\tau)}$  will be largest at the lag where dependence between components is strongest.

We have seen that there are many open problems in the area of multivariate and temporally dependent extremes – in particular declustering and the identification of events is unclear. Smith and Weissmann's technique of combining the series of each component into one univariate time series and then declustering is effective for the estimation of the multivariate extremal index, but it is not clear how to identify events in the individual components using such a method. In the univariate setting we saw that intervals declustering is an effective and theoretically justified method – the extension of this theory and declustering technique to multiple dimensions is one direction further work could pursue. Another idea, as mentioned above, is to incorporate other dependence measures such as  $\chi^{(\tau)}$  into a declustering method so that more knowledge of the dependence structure is being utilised.

# Appendices

# Appendix A

## Appendix to Chapter 3

### A.1 Proof of results from Chapter 3 §3.2

#### A.1.1 Proof of Proposition 3.2.3.1

For simplicity we denote  $c_k$  by  $c$ . Sampling from some general distribution with the first stopping rule, we have:

$$\mathbb{E} \left[ \frac{1}{N} \right] = \sum_{n=1}^{\infty} \bar{F}(c) \frac{F(c)^{n-1}}{n} = -\frac{\bar{F}(c)}{F(c)} \log(\bar{F}(c)),$$

where  $F(x)$  and  $\bar{F}(x) = 1 - F(x)$  are the CDF and survival function of the distribution of  $X_i$ ,  $i = 1, \dots, n$ . Thus,

$$\begin{aligned} \mathbb{E} [\bar{X}_N] &= \mathbb{E} \left[ \frac{1}{N} \mathbb{E} \left[ \sum_{i=1}^N X_i \mid N = n \right] \right] \\ &= \mathbb{E} \left[ \frac{1}{N} ((N-1) \mathbb{E} [X \mid X \leq c] + \mathbb{E} [X \mid X > c]) \right] \\ &= \mathbb{E} [X \mid X \leq c] + \mathbb{E} \left[ \frac{1}{N} (\mathbb{E} [X \mid X > c] - \mathbb{E} [X \mid X \leq c]) \right] \\ &= \mathbb{E} [X \mid X \leq c] + \mathbb{E} \left[ \frac{1}{N} \right] (\mathbb{E} [X \mid X > c] - \mathbb{E} [X \mid X \leq c]). \end{aligned}$$

Specifically for sampling from the exponential distribution:

$$\mathbb{E} \left[ \frac{1}{N} \right] = \frac{\beta c}{e^{\beta c} - 1}.$$

By the memoryless property of the exponential distribution:

$$\mathbb{E} [X \mid X > c] = \mathbb{E} [X] + c = \frac{1}{\beta} + c,$$

and rearranging  $\mathbb{E}[X] = F(c)\mathbb{E}[X|X \leq c] + \bar{F}(c)\mathbb{E}[X|X > c]$  gives

$$\mathbb{E}[X|X \leq c] = \frac{1}{F(c)} \left[ \frac{1}{\beta} - \bar{F}(c) \left( c + \frac{1}{\beta} \right) \right] = \frac{1}{\beta} - c \frac{\bar{F}(c)}{F(c)}. \quad (\text{A.1.1})$$

Therefore, for the exponential distribution,

$$\mathbb{E}[\bar{X}_N] = \frac{1}{\beta} + \frac{c}{e^{\beta c} - 1} \left( \frac{\beta c}{1 - e^{-\beta c}} - 1 \right). \quad (\text{A.1.2})$$

For the standard estimator based on the full sample we have  $1/\hat{\beta}_{std} = \bar{X}_N$ , the sample mean. The first part of Proposition 3.2.3.1 then follows from (A.1.2).

If the final data point is excluded from the sample then all included samples are from the distribution truncated at  $c$ , so, from (A.1.1),

$$\mathbb{E} \left[ \frac{1}{\hat{\beta}_{ex}} \mid N > 1 \right] = \mathbb{E}[X|X \leq c] = \frac{1}{\beta} - \frac{c}{e^{\beta c} - 1},$$

leading to the expression in the second part of Proposition 3.2.3.1.

### A.1.2 Proof of Theorem 3.2.4.1 in Chapter 3

We start by defining the following key quantities for each  $k \geq 1$ ,

$$S_k := (n_0 + k)\bar{X}_k = n_0\bar{X}_0 + \sum_{j=1}^k X_j \quad V_k := \frac{X_k}{S_k}.$$

Marginally  $S_k \sim \text{Gamma}((n_0 + k)\alpha, \beta)$  and  $V_k \sim \text{Beta}(\alpha, (n_0 + k)\alpha)$ ; we denote their marginal densities as:

$$\begin{aligned} f_{S_k}(s_k) &\propto s_k^{(n_0+k-1)\alpha-1} e^{-\beta s_k} \\ f_{V_k}(v_k) &\propto v_k^{\alpha-1} (1-v_k)^{(n_0+k-1)\alpha-1}. \end{aligned}$$

The stopping time,  $N$ , is  $n$  if  $X_n > \gamma\bar{X}_{n-1}$  and  $X_i < \gamma\bar{X}_{i-1}$  for  $1 \leq i < n$ . However,

$$\begin{aligned} X_n > \gamma\bar{X}_{n-1} &\Leftrightarrow X_n > \frac{\gamma}{n+n_0-1}(S_n - X_n) \\ &\Leftrightarrow \left( 1 + \frac{\gamma}{n+n_0-1} \right) X_n > \gamma \frac{1}{n+n_0-1} S_n \\ &\Leftrightarrow \left( 1 + \frac{\gamma}{n+n_0-1} \right) V_n > \frac{\gamma}{n+n_0-1} \end{aligned}$$

$$\Leftrightarrow V_n > \frac{\gamma}{n + \gamma + n_0 - 1}.$$

So the stopping rule can be written purely as function of the  $V$ s. Explicitly, we stop at time  $n$  if  $V_n > \frac{\gamma}{n + \gamma + n_0 - 1}$  and  $V_i < \frac{\gamma}{i + \gamma + n_0 - 1}$  for  $1 \leq i < n$ .

We define the statement  $\mathcal{A}_n := "V_1, \dots, V_n, S_n \text{ are mutually independent}"$ . Below, we will show by induction that  $\mathcal{A}_n$  holds for all  $n \geq 1$ . Thus  $\bar{X}_n \perp V_i \quad \forall i \leq n$ ; the distribution of  $\bar{X}_n$  is independent of whether or not the stopping rule has been triggered. Therefore,  $\bar{X}_N$  conditioned on  $N = n$  is equivalent to the mean of  $n$  i.i.d.  $\text{Gamma}(\alpha, \beta)$  random variables, as stated in the theorem.

$\mathcal{A}_{n-1} \Rightarrow \mathcal{A}_n$ : If  $\mathcal{A}_{n-1}$  holds then the joint pdf of  $V_1, \dots, V_{n-1}, S_{n-1}$  can be factorised:

$$f_{n-1}(v_1, \dots, v_{n-1}, s_{n-1}) = f_{S_{n-1}}(s_{n-1}) \prod_{i=1}^{n-1} f_{V_i}(v_i),$$

Consider the change of variables  $(V_1, \dots, V_{n-1}, S_{n-1}, X_n) \rightarrow (V_1, \dots, V_{n-1}, V_n, S_n)$ , where  $X_n = S_n V_n$  and  $S_{n-1} = S_n(1 - V_n)$ . The Jacobian for this transformation is:

$$|J| = \left| \frac{\partial(v_{1:n-1}, s_{n-1}, x_n)}{\partial(v_{1:n}, s_n)} \right| = \begin{vmatrix} I_{n-1} & 0 \\ 0 & A \end{vmatrix} = s_n.$$

where  $I_{n-1}$  is the  $(n-1) \times (n-1)$  identity matrix and

$$A = \frac{\partial(s_{n-1}, x_n)}{\partial(v_n, s_n)} = \begin{bmatrix} -s_n & 1 - v_n \\ s_n & v_n \end{bmatrix}.$$

So, since  $S_{n-1}$  and  $V_1, \dots, V_{n-1}$  are independent of  $X_n$ ,

$$\begin{aligned} f_n(v_{1:n}, s_n) &= f_{n-1}(v_{1:n-1}, s_{n-1}(s_n, v_n)) f_X(x(s_n, v_n)) |J| \\ &\propto \left( \prod_{i=1}^{n-1} f_{V_i}(v_i) \right) (s_n(1 - v_n))^{(n_0+n-1)\alpha-1} e^{-\beta s_n(1-v_n)} (s_n v_n)^{\alpha-1} e^{-\beta s_n v_n} s_n \\ &= \left( \prod_{i=1}^{n-1} f_{V_i}(v_i) \right) s_n^{(n+n_0)\alpha-1} e^{-\beta s_n} v_n^{\alpha-1} (1 - v_n)^{(n_0+n-1)\alpha-1} \\ &\propto \prod_{i=1}^n f_{V_i}(v_i) f_{S_n}(s_n). \end{aligned}$$

So  $\mathcal{A}_n$  holds.

$\mathcal{A}_1$  holds: We must show that  $V_1$  and  $S_1$  are independent. We do this by using the change of variables  $(\bar{X}_0, X_1) \rightarrow (V_1, S_1)$  to show that the joint pdf of  $V_1$  and  $S_1$  factorises.

We have

$$f_{\bar{X}_0, X_1}(\bar{x}_0, x_1) \propto \bar{x}_0^{n_0\alpha-1} e^{-n_0\beta\bar{x}_0} x_1^{\alpha-1} e^{-\beta x_1}$$

and  $X_1 = S_1 V_1$  and  $\bar{X}_0 = \frac{1}{n_0} S_1 (1 - V_1)$ . So Jacobian for the transformation is:

$$\left| \frac{\partial(\bar{X}_0, X_1)}{\partial(V_1, S_1)} \right| = \begin{vmatrix} -\frac{s_1}{n_0} & s_1 \\ \frac{1}{n_0}(1-v_1) & v_1 \end{vmatrix} = \frac{s_1}{n_0}.$$

Thus the joint pdf of  $V_1, S_1$  is:

$$\begin{aligned} f_{V_1, S_1}(v_1, s_1) &\propto s_1 (s_1 (1 - v_1))^{n_0\alpha-1} e^{-n_0\beta(s_1(1-v_1)/n_0)} \times (s_1 v_1)^{\alpha-1} e^{-\beta s_1 v_1} \\ &= s_1^{(n_0+1)\alpha-1} e^{-\beta s_1} v_1^{\alpha-1} (1 - v_1)^{n_0\alpha-1} \\ &\propto f_{S_1}(s_1) f_{V_1}(v_1), \quad \text{as required.} \end{aligned}$$

■

## A.2 Properties of the GEV shape parameter

### A.2.1 Fixed-threshold stopping rule

The shape parameter,  $\xi$ , is important in determining the tail behaviour. Figure A.2.1 shows the relative bias, variance and RMSE of each of the estimators when sampling using the fixed-threshold stopping rule for  $\xi = 0.2$  and  $-0.2$  (top and bottom rows respectively). Judged by RRMSE, we find that  $\ell_{pc}$  is generally best for moderate to large  $k$ , with clear benefits for  $\xi = -0.2$ ; however  $\ell_{fc}$  has generally quite similar RRMSE and low bias. As one would expect the lighter the tail of the distribution, the smaller both the relative variance and, in most cases, the relative bias of the shape parameter estimators resulting in smaller RRMSE. To help understand why these RRMSE results arise we now look at more detail at the performance of the four estimators.

The standard MLE for the shape parameter,  $\hat{\xi}_{std}$ , is almost always positively biased while  $\hat{\xi}_{ex}$  leads to quite large negative bias (with  $E(\hat{\xi}_{ex}) < 0.1$  when  $\xi = 0.2$  and  $k < 50$  (Figure A.2.1)) since we lose information about the upper tail of the underlying distribution. In particular, the fitted distribution typically has a lighter tail and can even have an upper

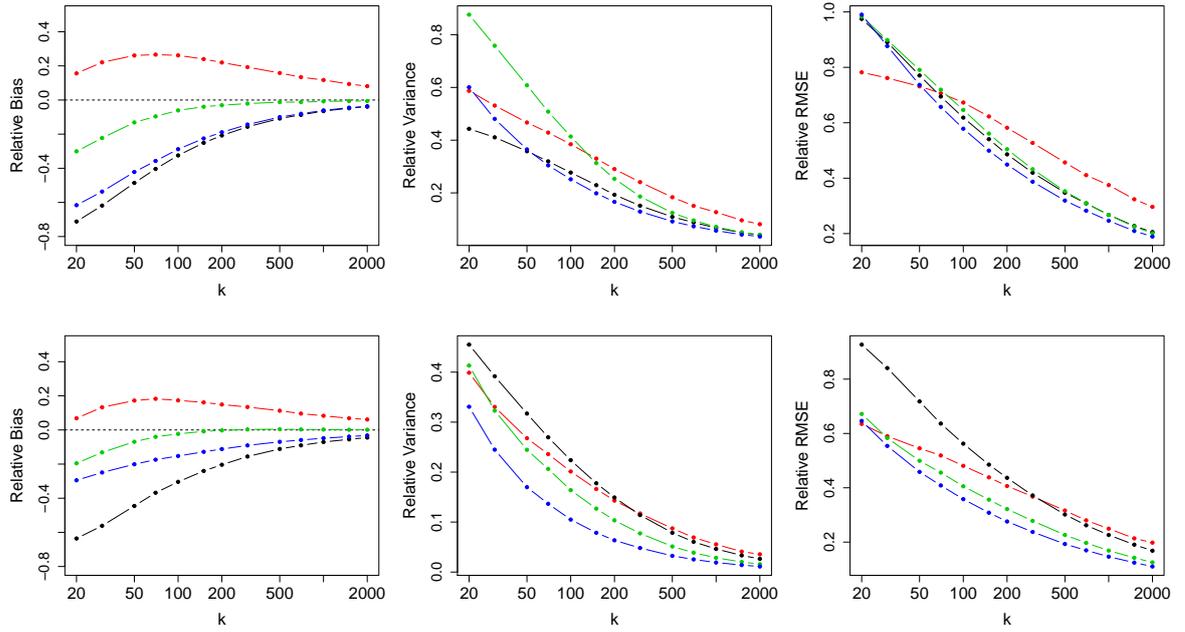


Figure A.2.1: Shape parameter estimates when sampling from the GEV distribution with  $(\mu, \sigma) = (0, 1)$  using the fixed-threshold stopping rule with threshold  $c_k$  and  $\xi = 0.2$  (top) and  $\xi = -0.2$  (bottom) both plotted against  $k$ . Left: relative bias, centre: relative variance, right: relative RMSE, using: standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue). Based on  $10^5$  replicated samples with the historical data created using approach (3.4.2).

end point which could be less than the excluded observation. Unlike all other estimators considered, the variance of  $\hat{\xi}_{ex}$  is not substantially lower when the tail is lighter and so has quite large RRMSE when  $\xi = -0.2$ .

The partial conditioning method generally has  $\hat{\xi}_{pc}$  lower than the truth however, for moderate  $k$ , they consistently have low variance relative to the other methods over a range of  $\xi$ . Therefore, partial conditioning provides  $\xi$  estimators with the lowest RRMSE for  $k > 100$ . In contrast,  $\ell_{fc}$  leads to very little bias in  $\xi$  estimates for  $k > 100$  but the variance can be large, particularly when  $\xi = 0.2$  with  $k < 100$ . This is in agreement with Molenberghs et al. (2014) findings that the full-conditional estimator has poor precision despite it's unbiasedness. However, unlike in Molenberghs et al. (2014), we find that, in our context, full conditioning can improve upon the standard estimator especially when the stopping threshold is high (*i.e.*, for large  $k$ ).

### A.2.2 Variable-threshold stopping rule

Properties of the shape parameter estimators under the variable stopping rule are shown in the supplementary material of Barlow et al. (2020). We find that in the variable threshold setting  $\hat{\xi}_{std}$  has very low bias (similarly recall in Chapter 3 §3.2.4 when sampling from the gamma distribution with this stopping rule we found the standard return-level estimator was unbiased) whereas all other  $\xi$  estimators are negatively biased, with  $\hat{\xi}_{ex}$  having the largest negative bias out of all the estimators for both values of  $\xi$  considered. We find that  $\hat{\xi}_{std}$  also has the lowest RRMSE of the estimators. Despite  $\hat{\xi}_{std}$  performing well under the variable threshold stopping rule, this is not always the case for the  $\ell_{std}$  return-level estimators.

# Appendix B

## Appendix to Chapter 4

### B.1 Profile-likelihood based confidence intervals

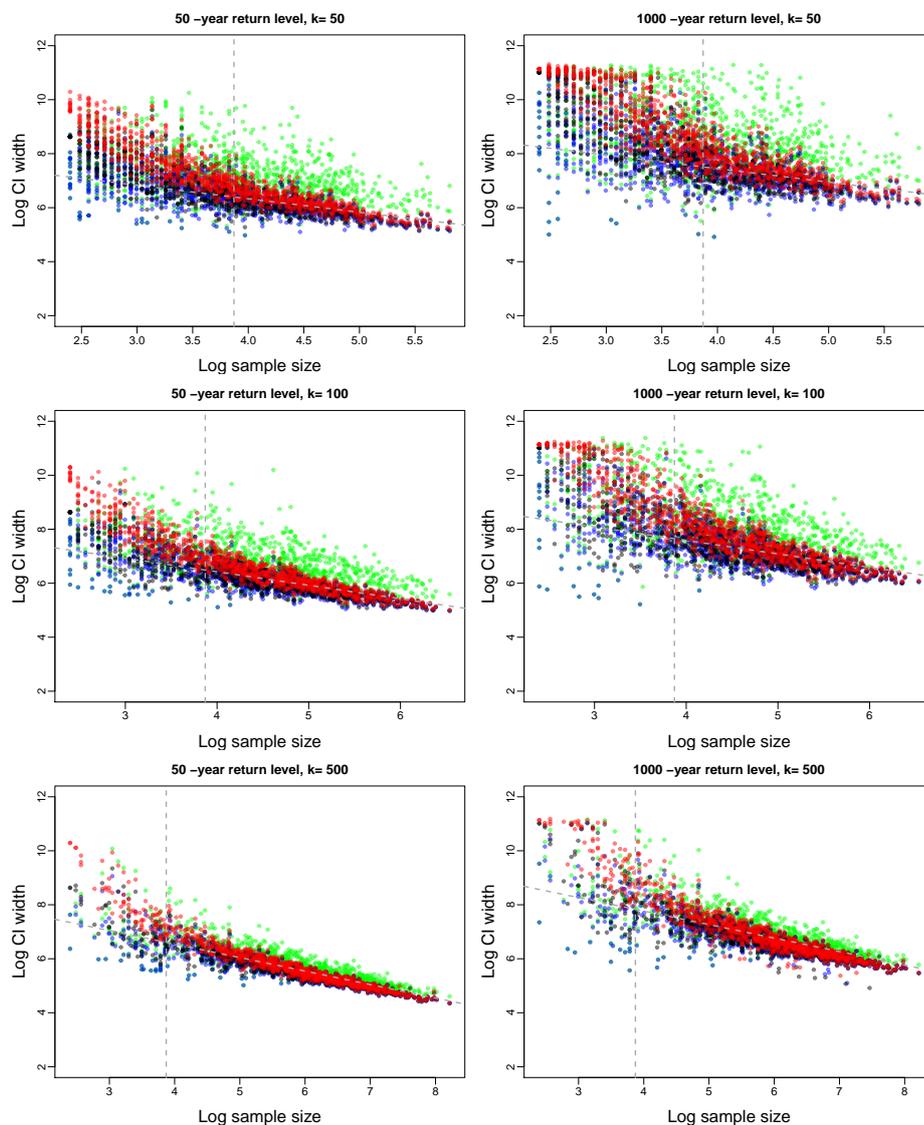


Figure B.1.1: *Log CI width vs log sample size for the 50-year (left) and 1000-year (right) return level estimates. Profile likelihood confidence intervals found using the standard likelihood (red), excluding the final observation (black), full conditioning (green) and partial conditioning (blue) based on 5000 samples from the GEV distribution with parameters equal to the standard MLEs for the Lune data and sample size determined by the fixed-threshold stopping rule*

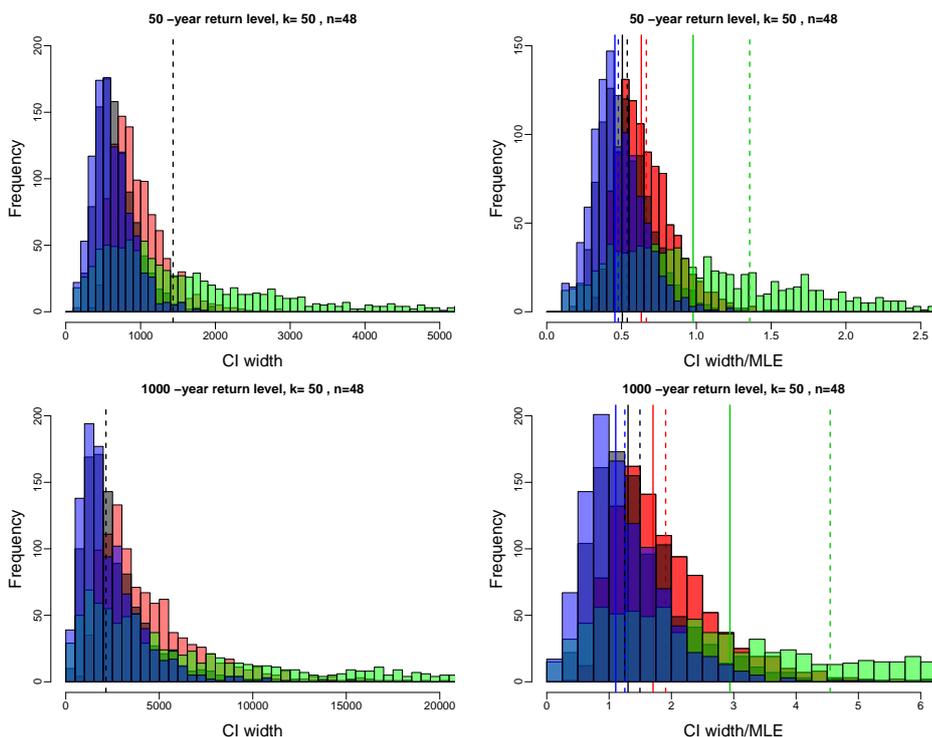


Figure B.1.2: *CI widths (left) and CI widths/MLE (right) for the 50-yr (top) and 1000-yr (bottom) return level using  $k = 50$  and Lune MLE to create samples of size 48.*

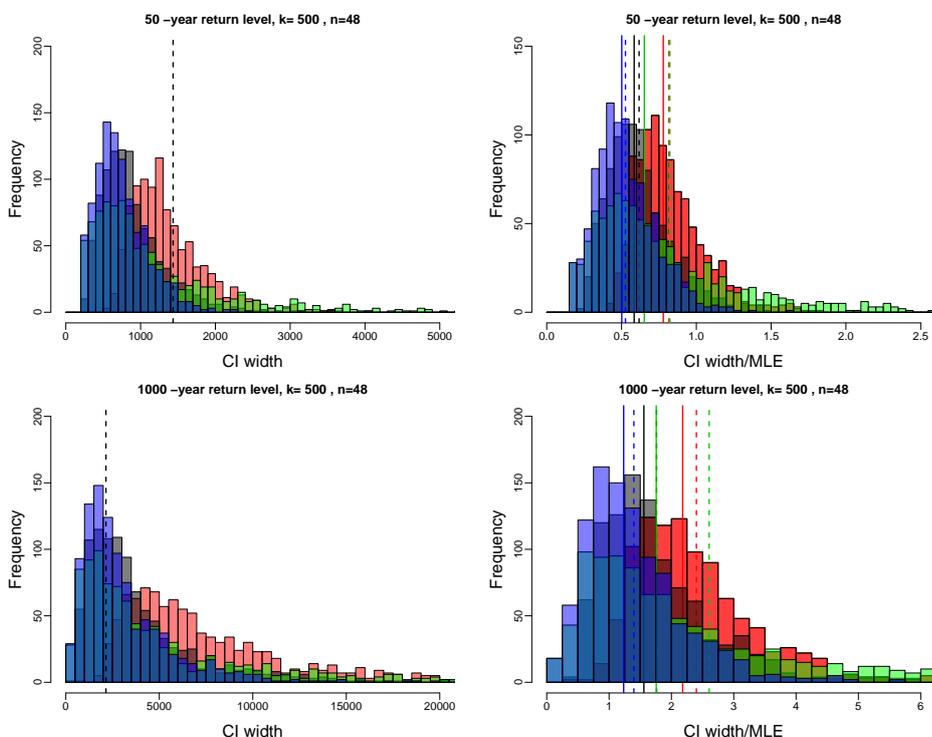


Figure B.1.3: *CI widths (left) and CI widths/MLE (right) for the 50-yr (top) and 1000-yr (bottom) return level using  $k = 500$  and Lune MLE to create samples of size 48.*

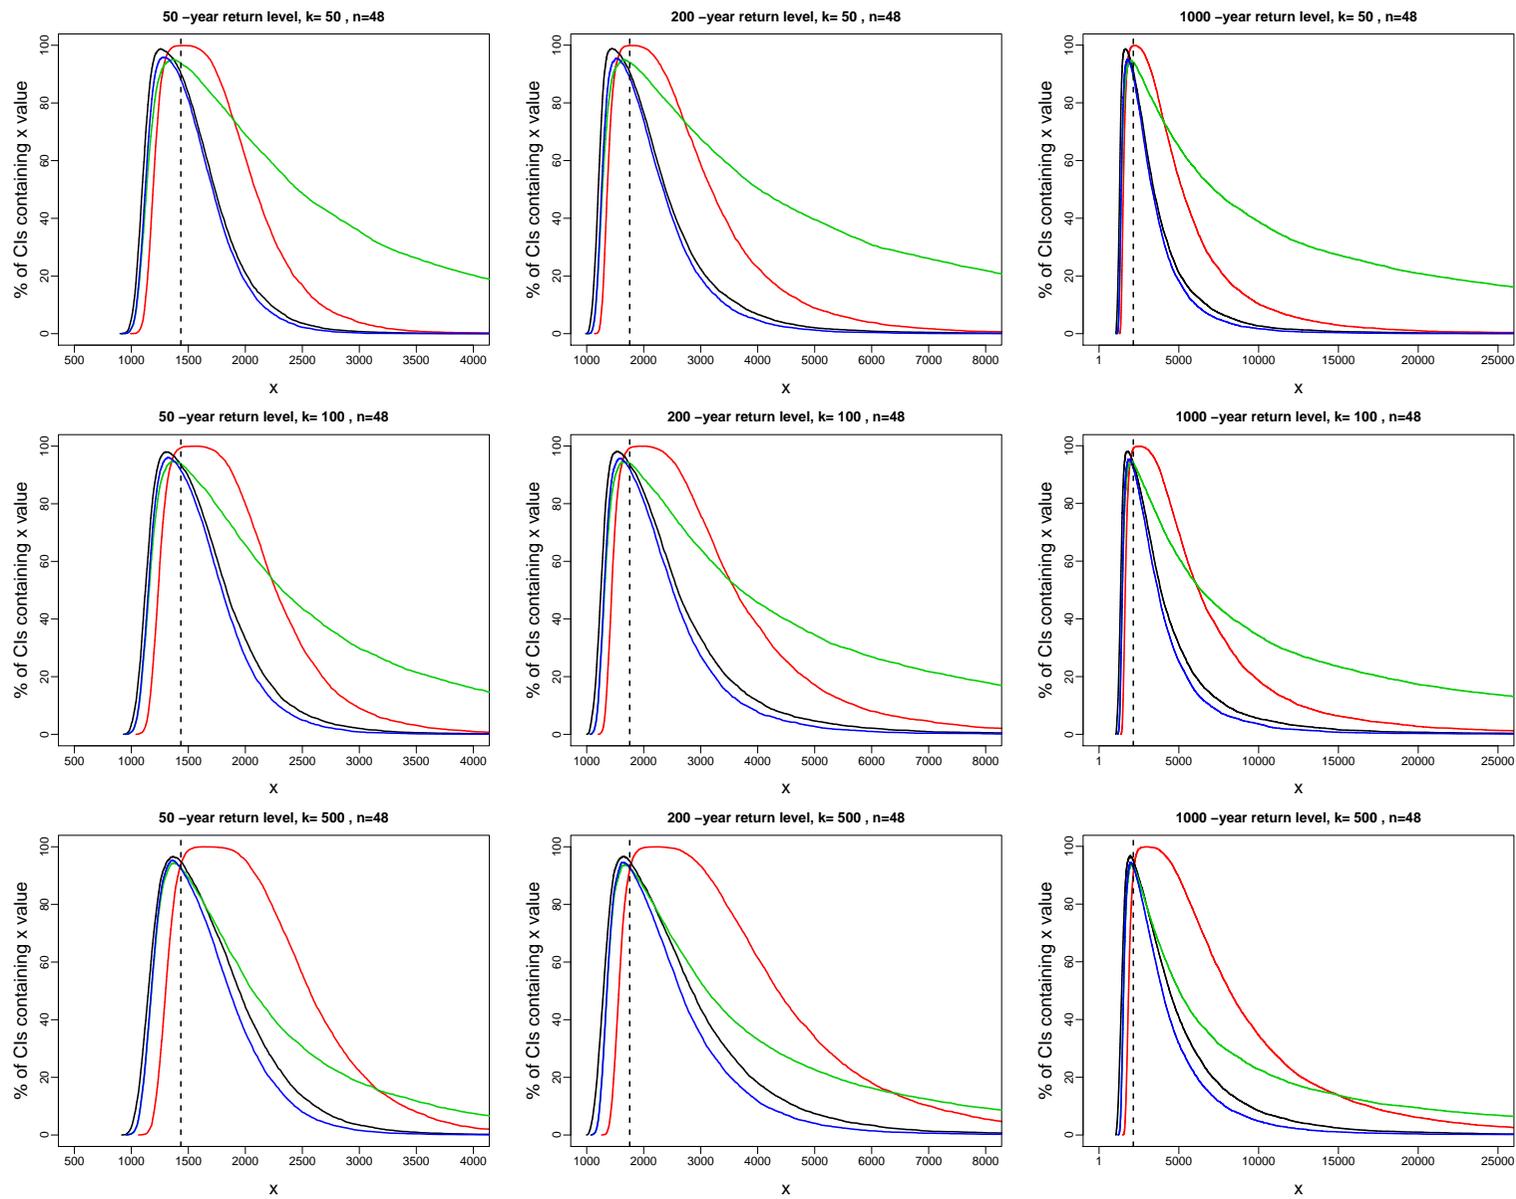


Figure B.1.4: CIs for the 50/200/1000-year return levels (from left to right) using  $k = 50, 100, 500$  (top to bottom) and the Lune MLE to create samples of size 48 using the fixed-threshold stopping rule with threshold  $c_k$ .

## B.2 Comparing bootstrap-based confidence intervals

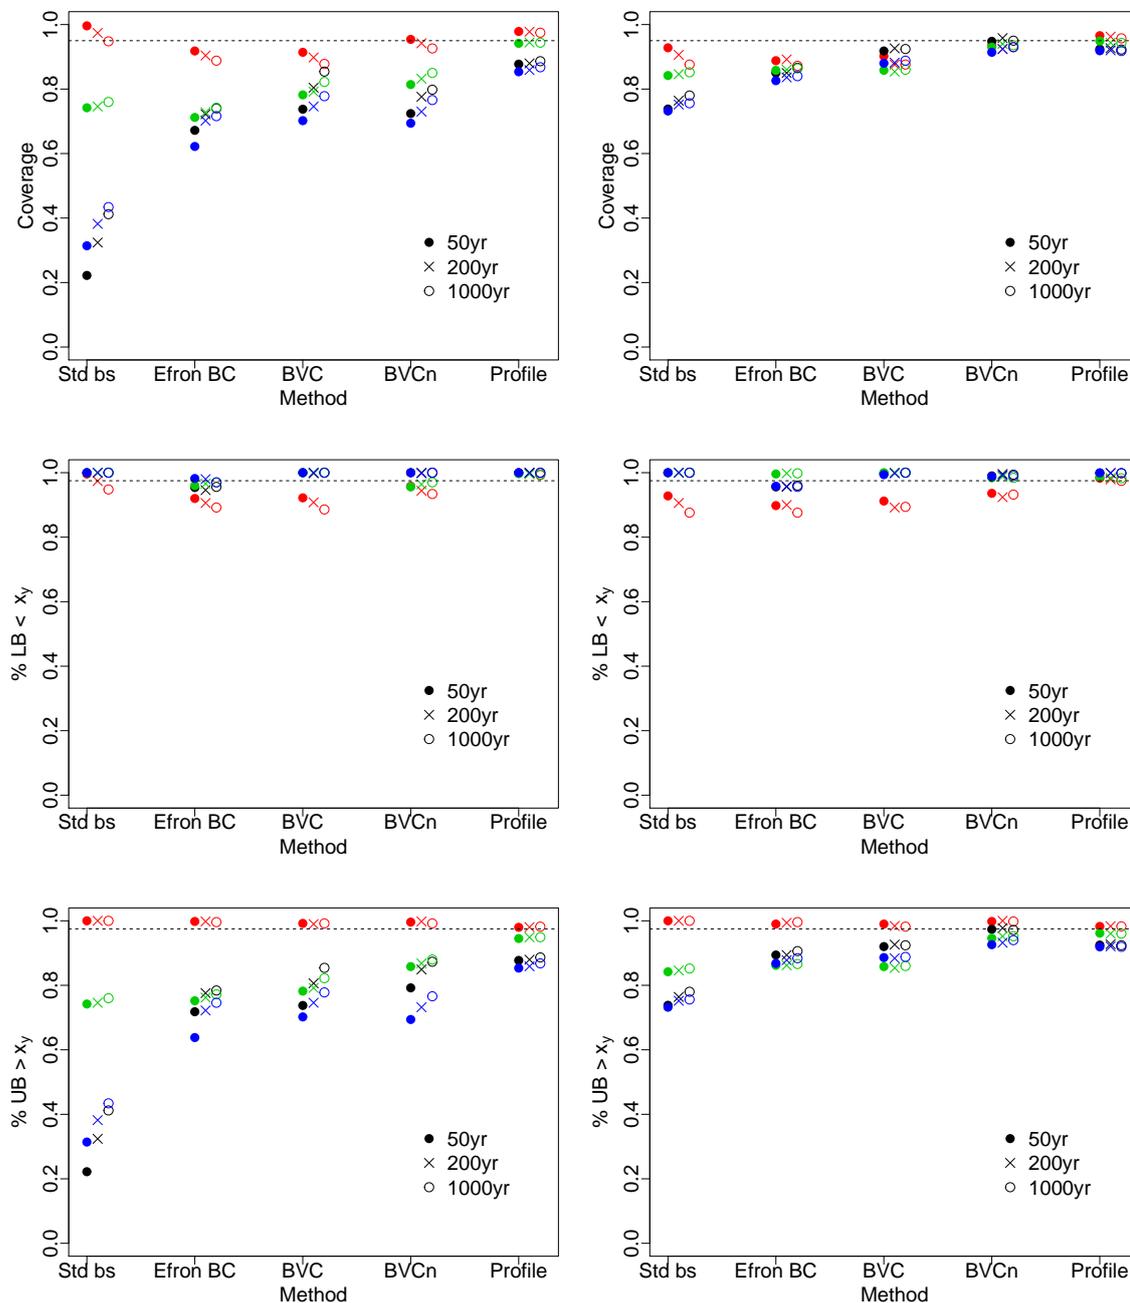


Figure B.2.1: Coverage, % of lower bounds below  $x_y$ , % of upper bounds  $x_y$  of confidence intervals for the  $y$ -year return level,  $x_y$ , found using different confidence interval methods with  $\ell_{std}$  (red),  $\ell_{ex}$  (black),  $\ell_{fc}$  (green) and  $\ell_{pc}$  (blue). Primary data sets simulated from  $GEV(\hat{\theta}_{std}^{Lune})$  with sample size determined by the fixed-threshold stopping rule with stopping threshold  $c_{50}$  (left) and  $c_{500}$  (right). Bootstrap samples are created from the same sampling process as the primary data set with no restriction on sample size.

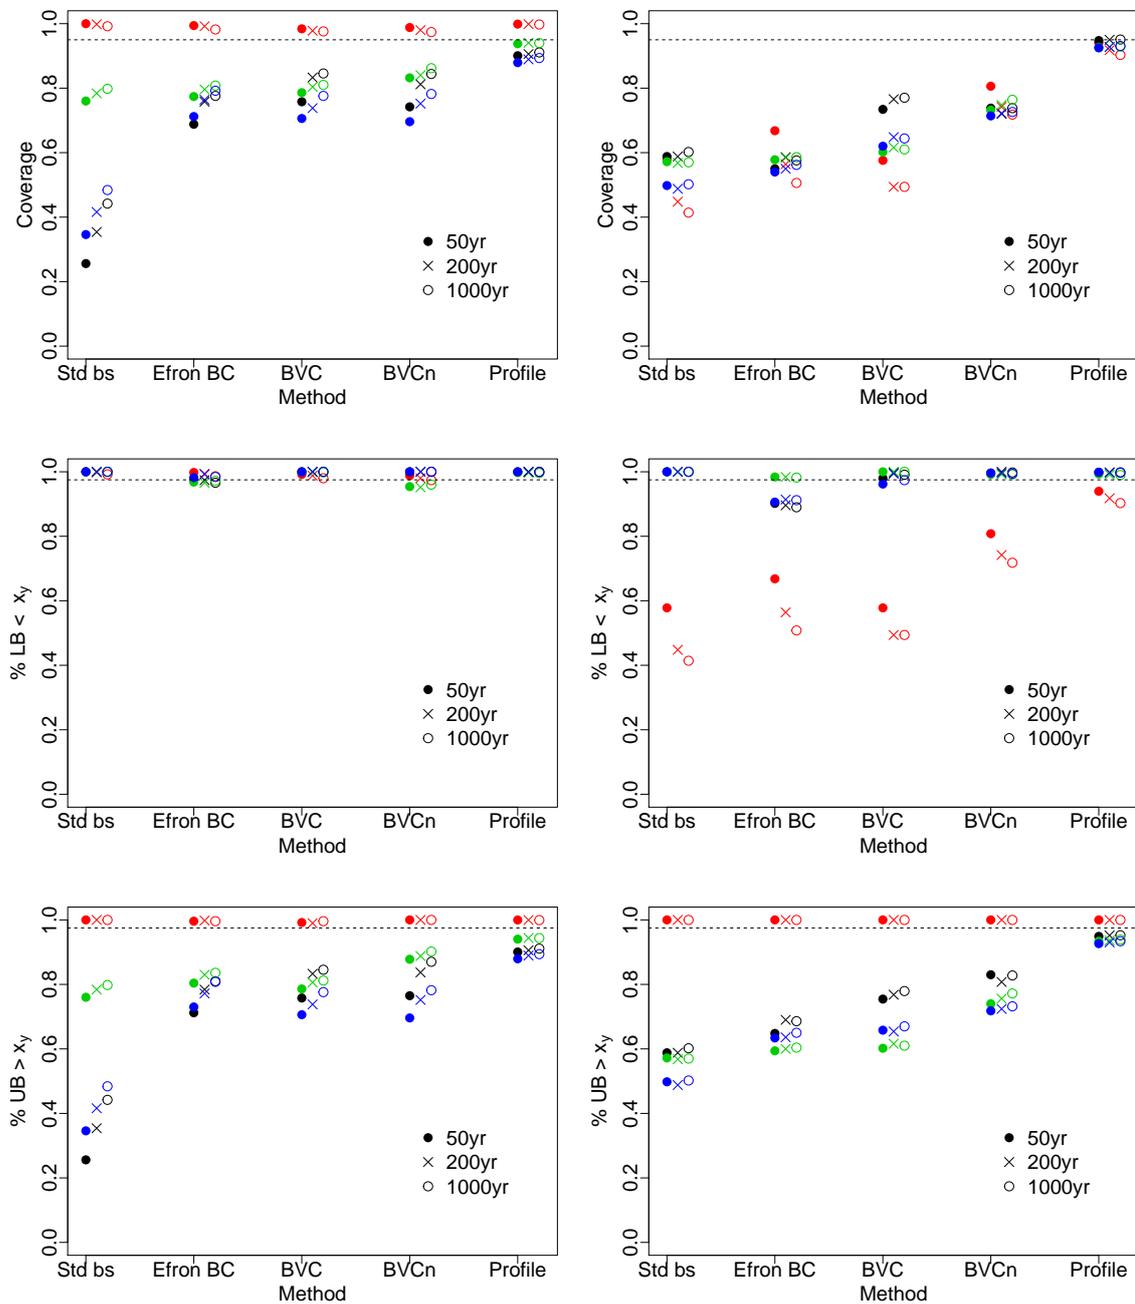


Figure B.2.2: Coverage, % of lower bounds below  $x_y$ , % of upper bounds  $x_y$  of confidence intervals for the  $y$ -year return level,  $x_y$ , found using different confidence interval methods with  $\ell_{std}$  (red),  $\ell_{ex}$  (black),  $\ell_{fc}$  (green) and  $\ell_{pc}$  (blue). Primary data sets simulated from  $GEV(\hat{\theta}_{std}^{Lune})$  with sample size,  $n$ , determined by the fixed-threshold stopping rule with stopping threshold  $c_{50}$  (left) and  $c_{500}$  (right) such that  $n = 48$ . Bootstrap samples are created from the same sampling process as the primary data set with no restriction on sample size.

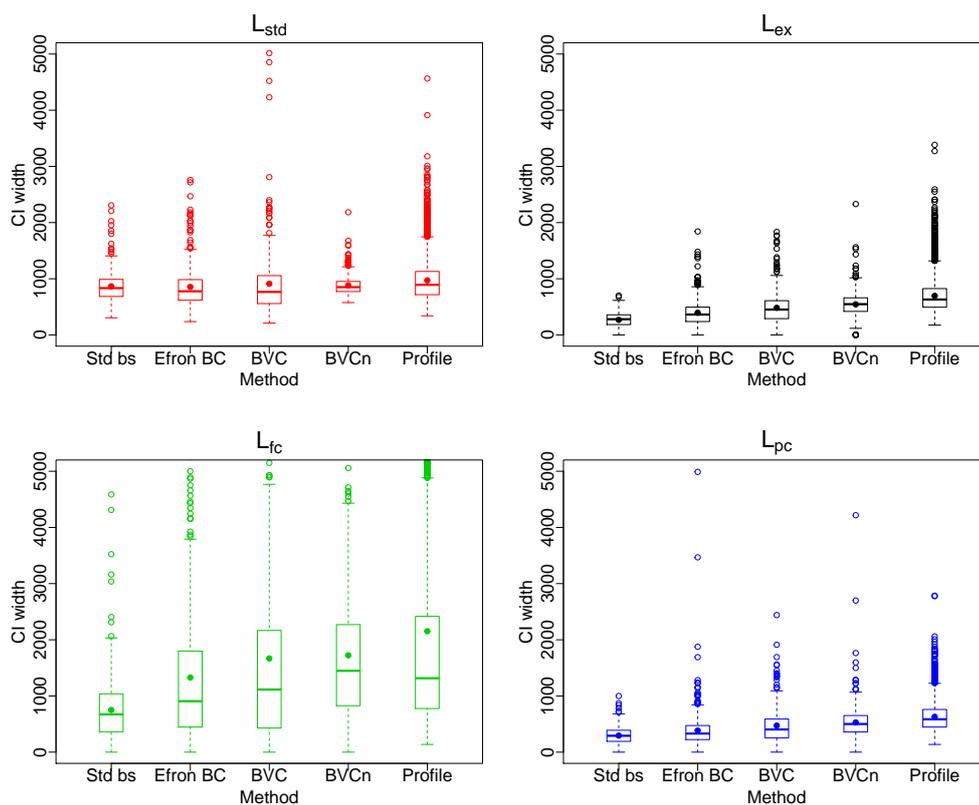


Figure B.2.3: Box plots of confidence interval widths for  $x_{50}$ , the 50-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ . Crosses indicate the mean confidence interval width for each method.

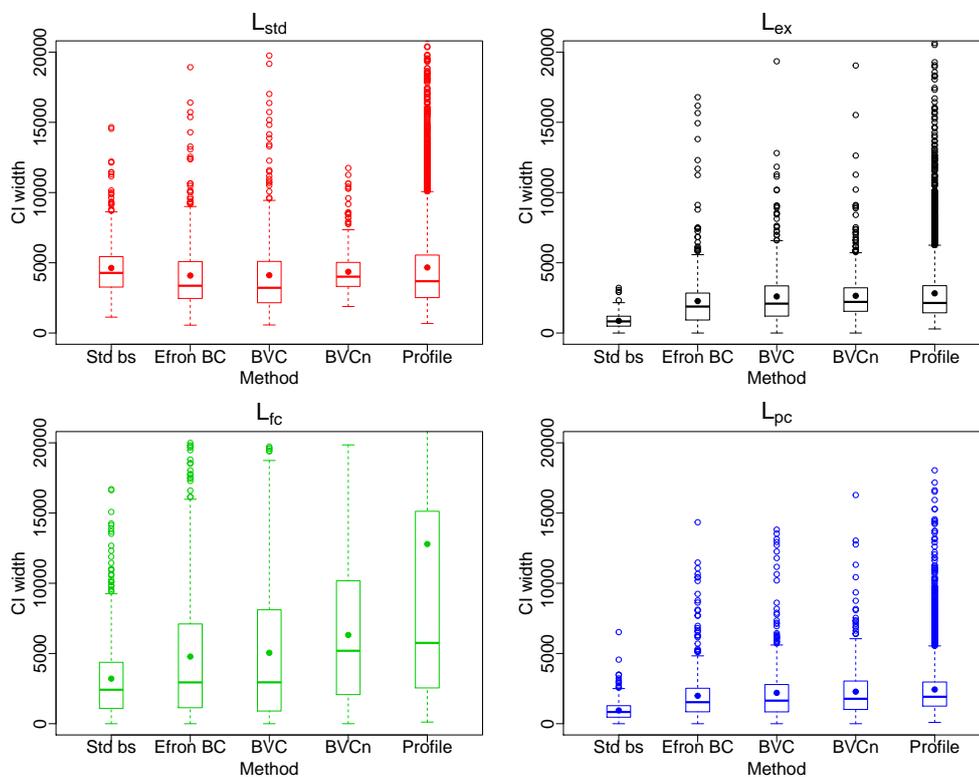


Figure B.2.4: Box plots of confidence interval widths for  $x_{1000}$ , the 1000-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ . Crosses indicate the mean confidence interval width for each method.

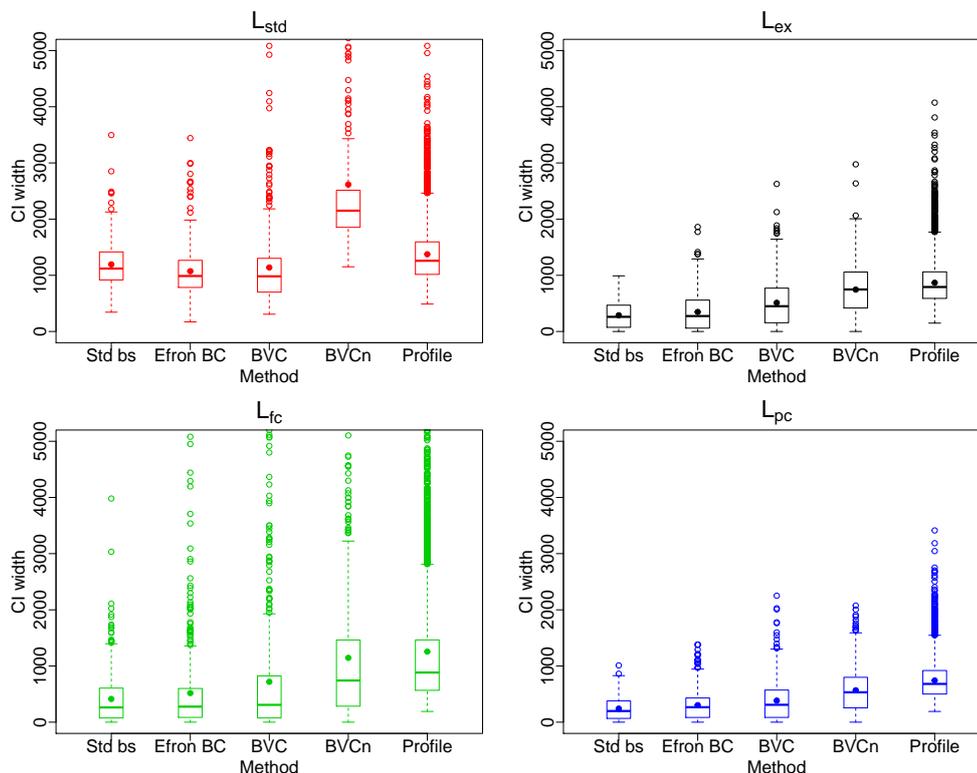


Figure B.2.5: Box plots of confidence interval widths for  $x_{50}$ , the 50-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 500$ . Crosses indicate the mean confidence interval width for each method.

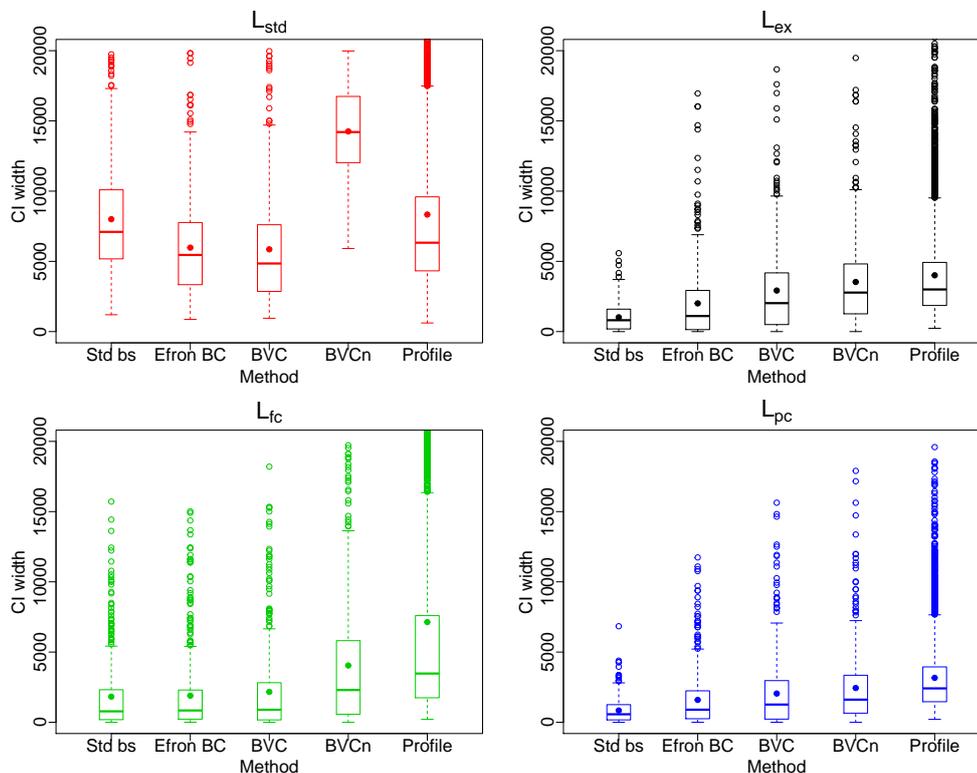


Figure B.2.6: Box plots of confidence interval widths for  $x_{1000}$ , the 1000-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 500$ . Crosses indicate the mean confidence interval width for each method.

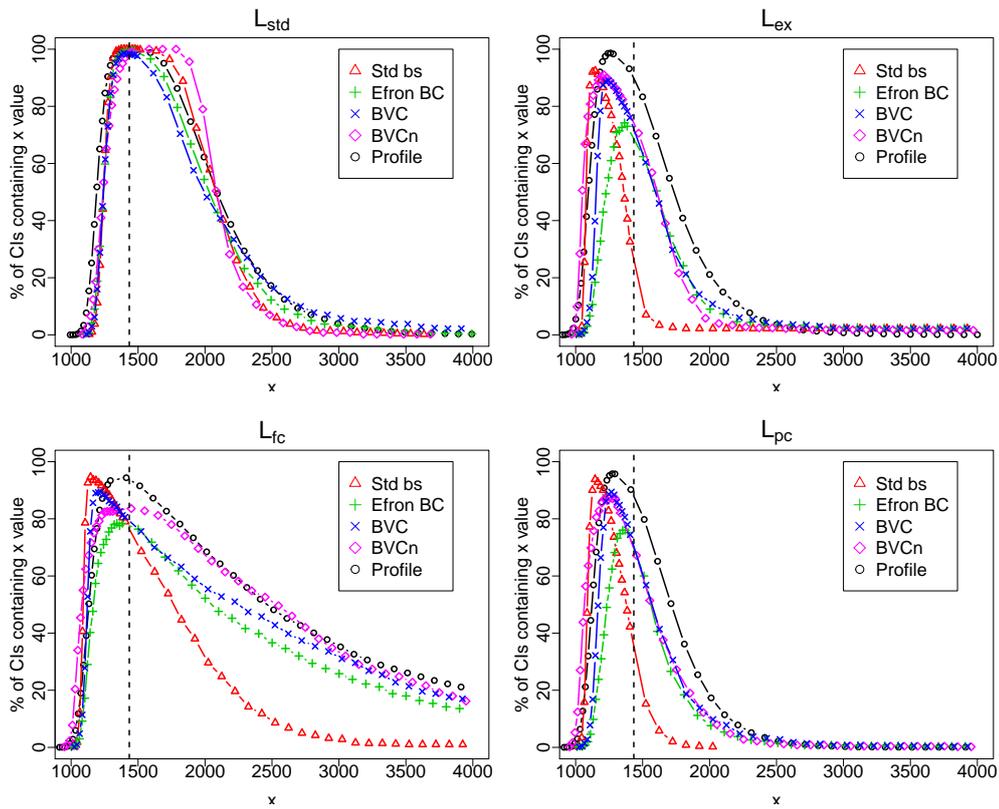


Figure B.2.7: Confidence intervals for  $x_{50}$ , the 50-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ .

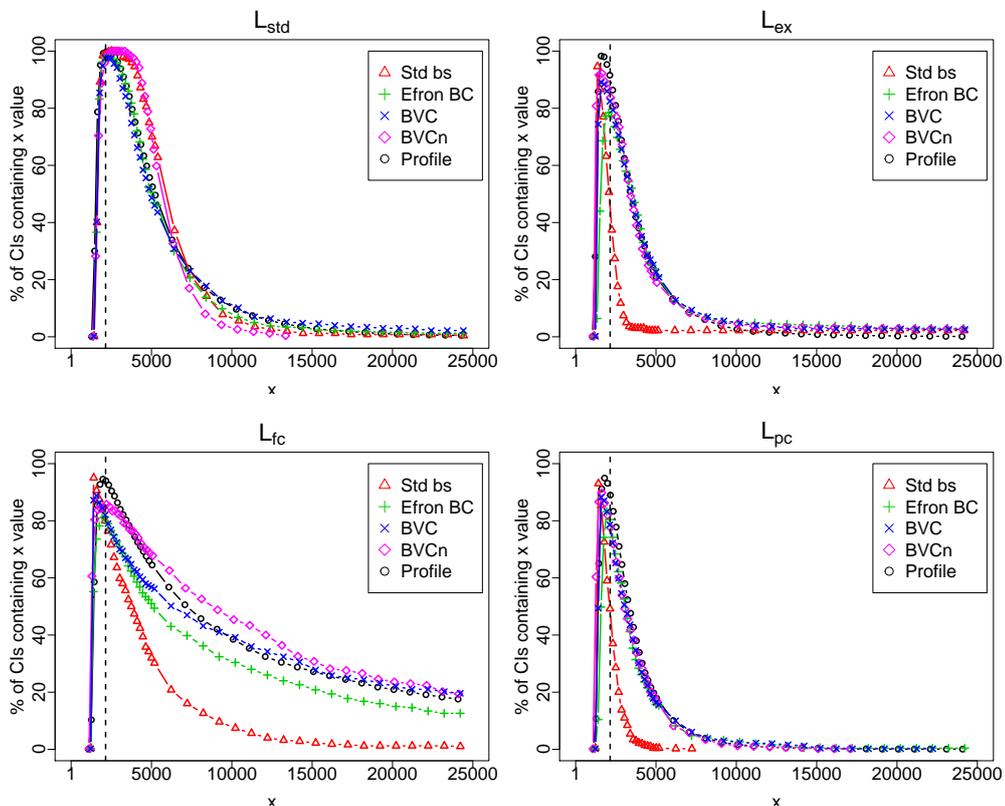


Figure B.2.8: Confidence intervals for  $x_{1000}$ , the 1000-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 50$ .

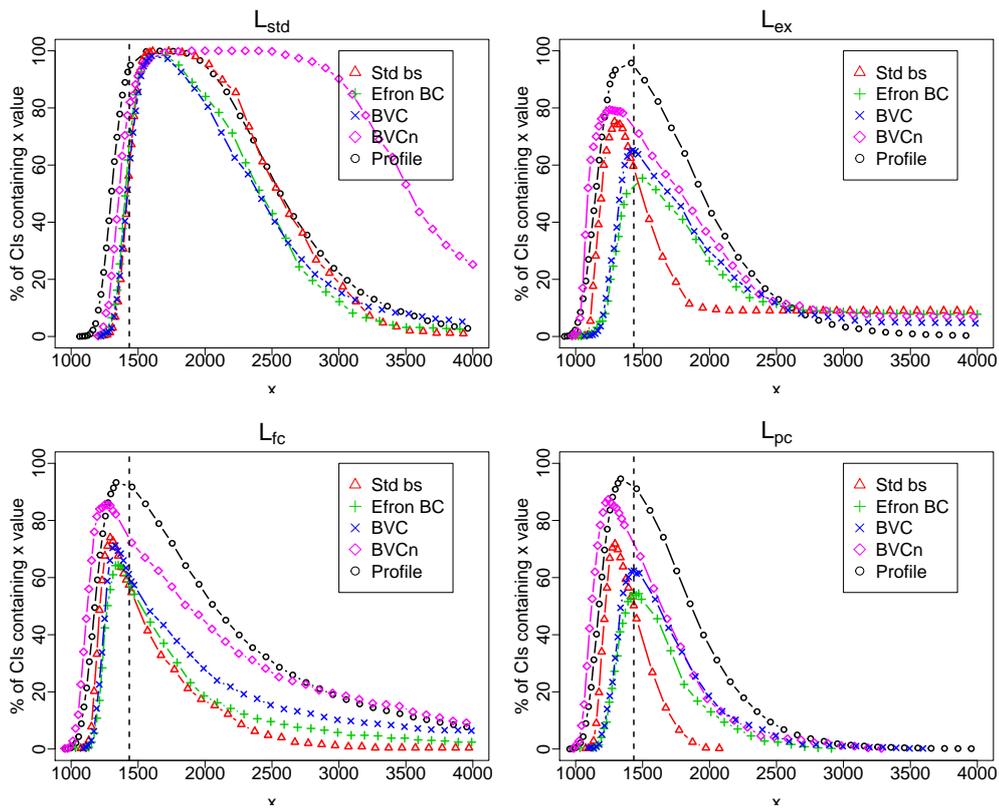


Figure B.2.9: Confidence intervals for  $x_{50}$ , the 50-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 500$ .

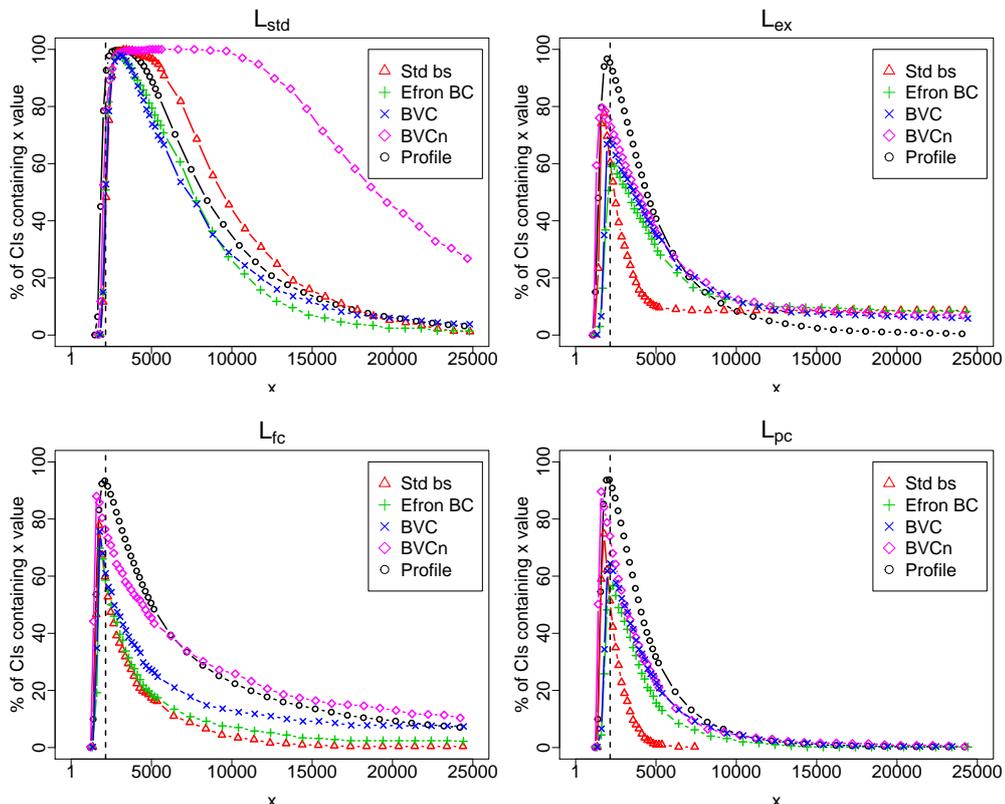


Figure B.2.10: Confidence intervals for  $x_{1000}$ , the 1000-year return level, using the 5 confidence interval methods considered. Primary data sets and bootstrap samples are the same as for Figure 4.3.1 with  $k = 500$ .

## B.3 Optimisation issues near the boundary

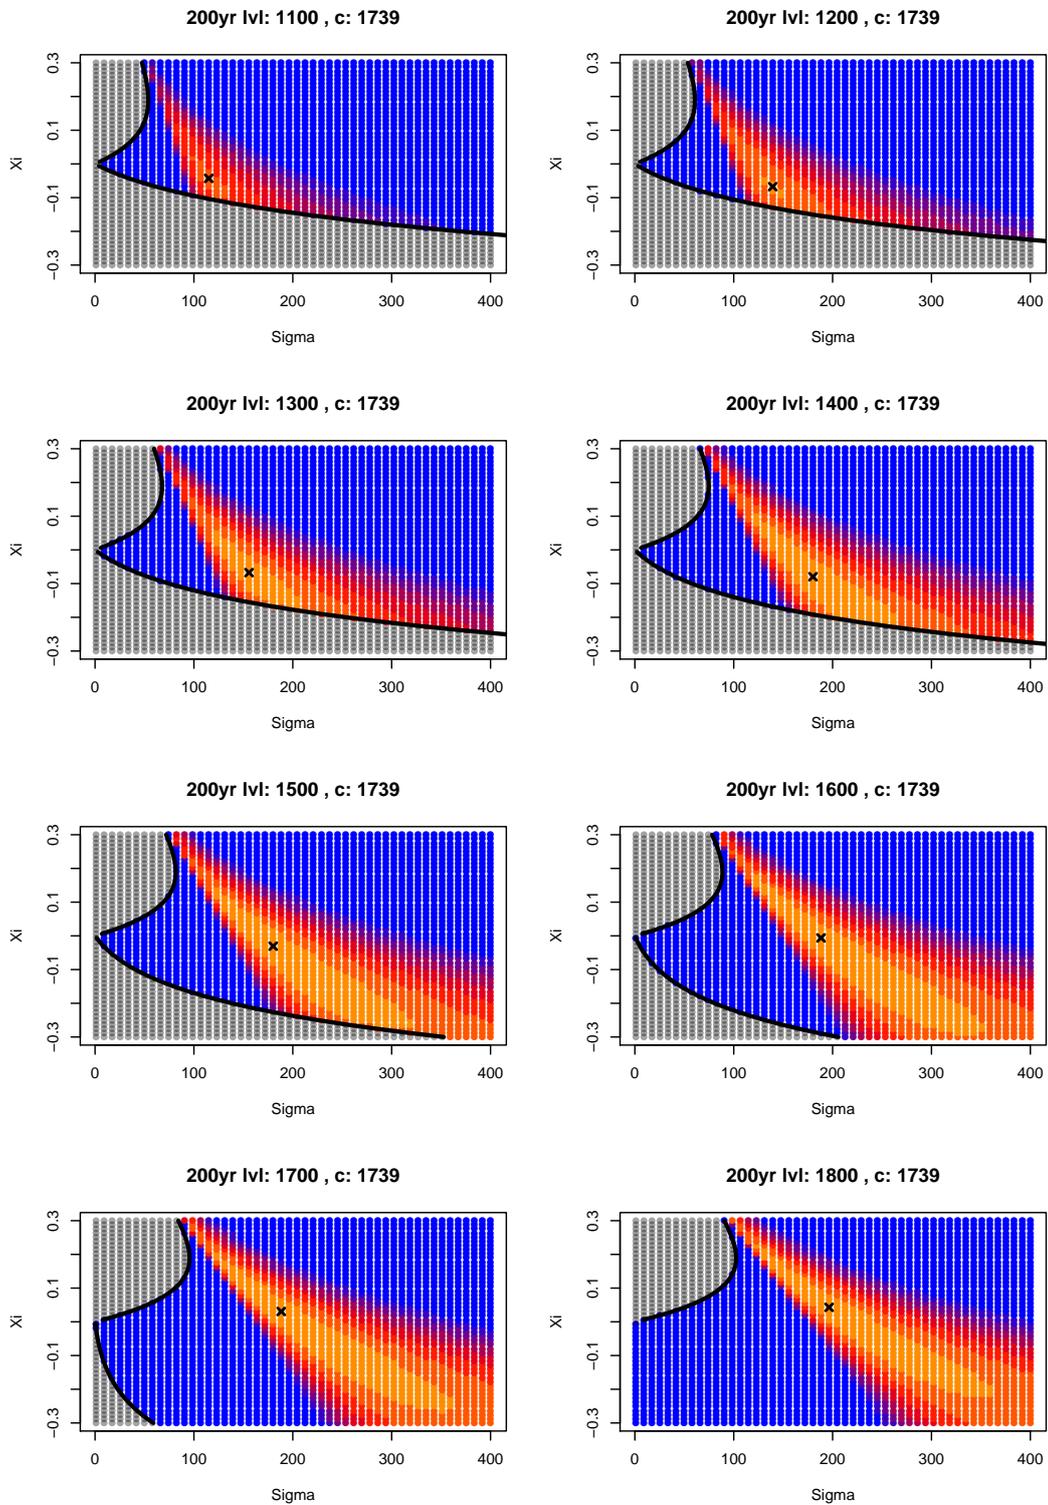


Figure B.3.1: Standard likelihood for each combination of  $\sigma$  and  $\xi$  given fixed  $x_{200}$  and  $c = 1739$  over a range of  $x_{200}$  values. Low/high likelihood regions are coloured in blue/yellow with grey being outside the parameter space. The black/red crosses are the MLEs using the optimisation/grid method respectively.

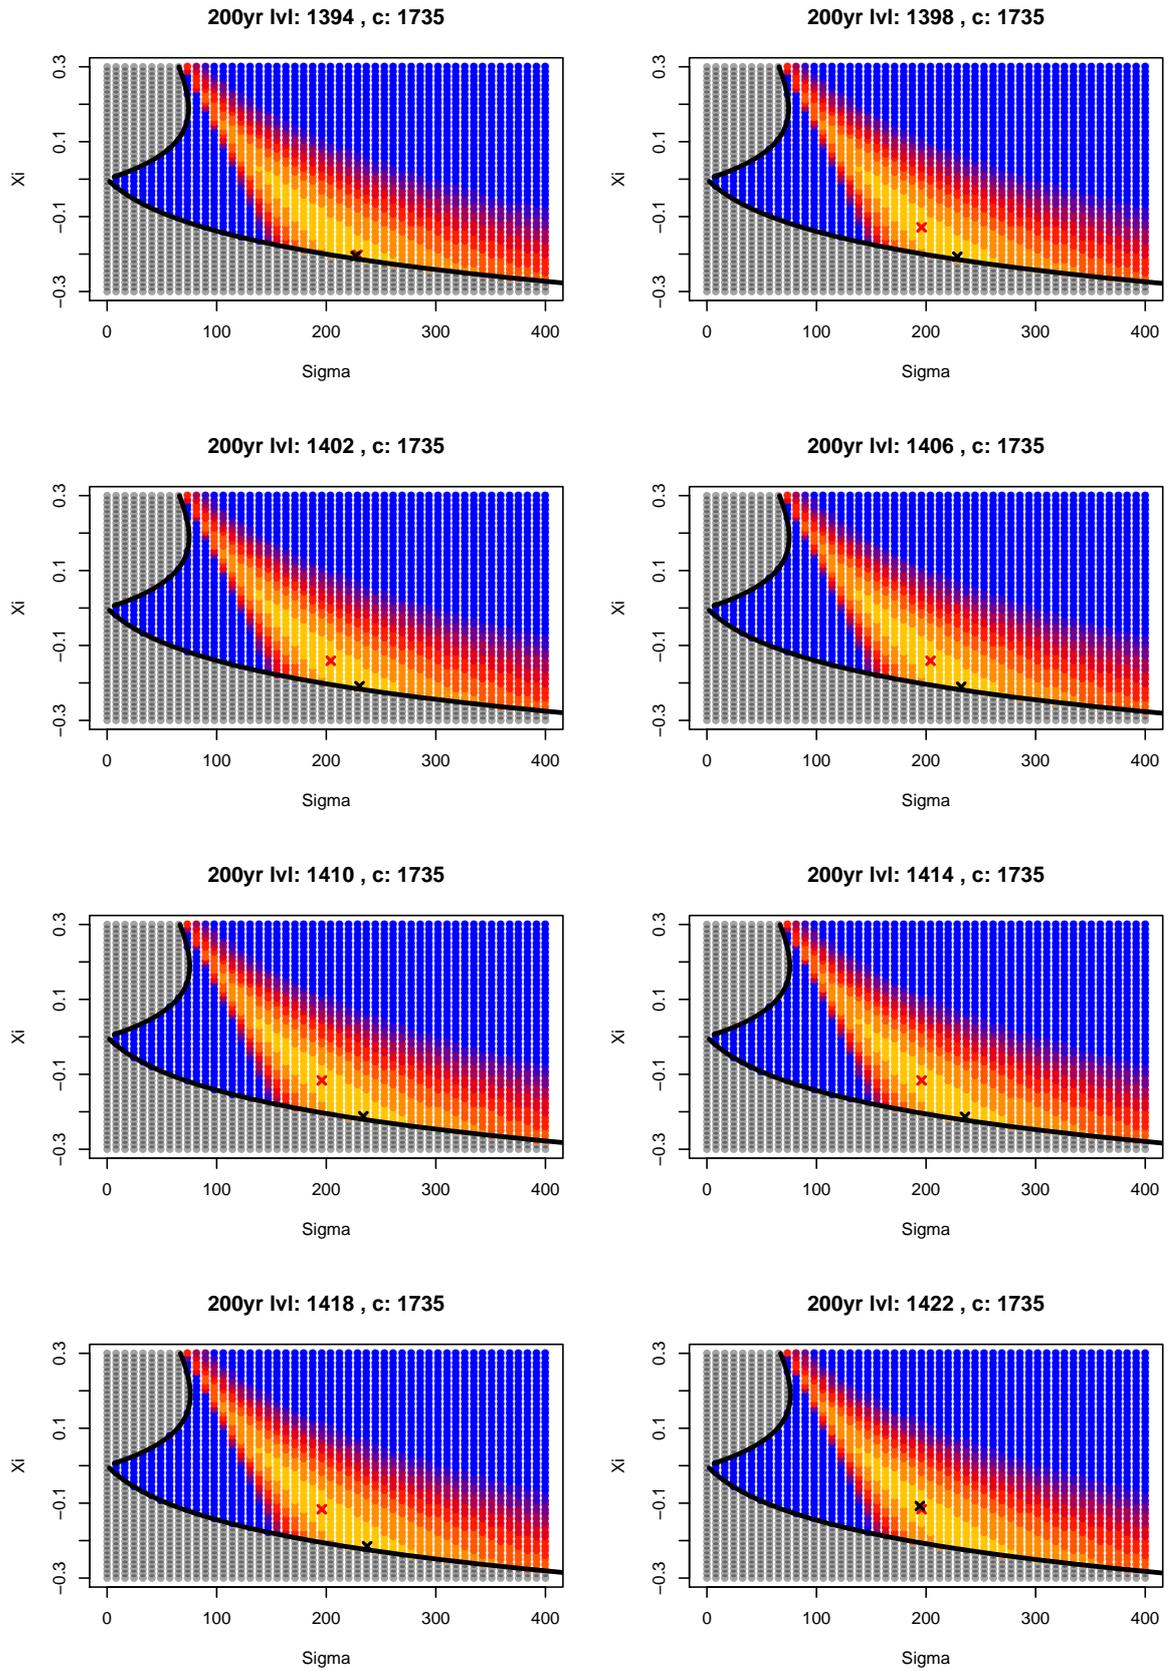


Figure B.3.2: As Figure 4.5.1 but for  $x_{200}$  values around the discontinuity.

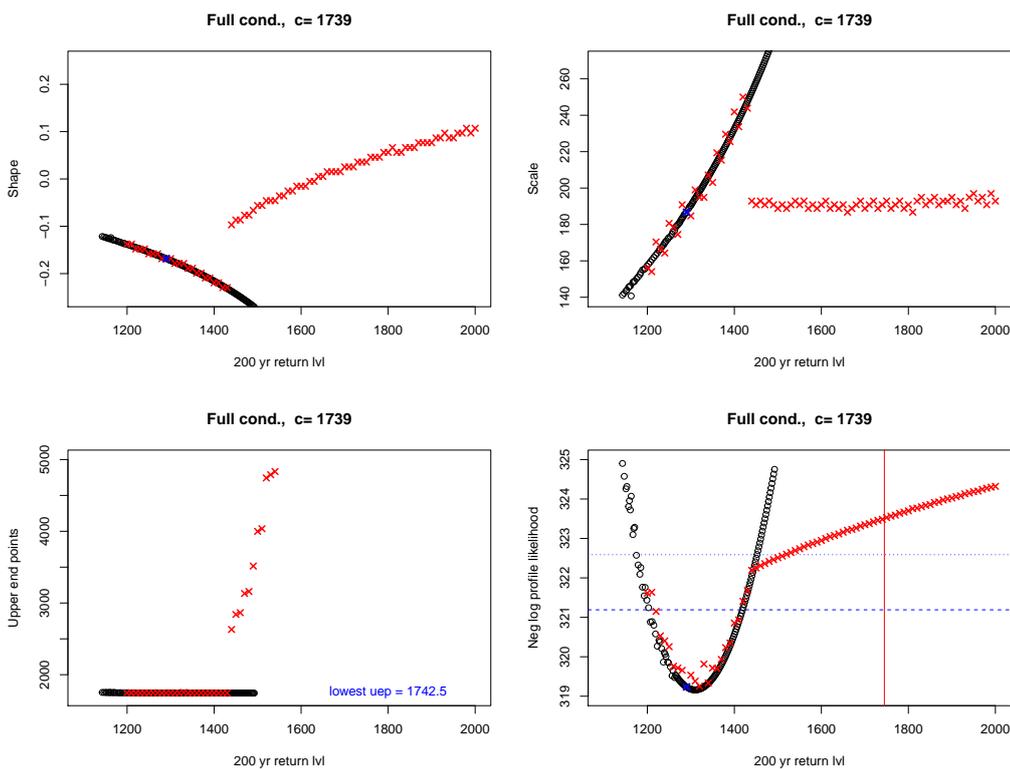


Figure B.3.3: Top: MLE of scale(left) and shape(right) parameters given the  $x_{200}$  value on the  $x$ -axis. Bottom: Estimated upper end point (left) and profile likelihood (right) over different  $x_{200}$  values. The black/red points refer to the optimisation/grid method respectively with the full conditional likelihood with  $c = 1739$ . The blue crosses are at the MLE using the grid method and the red vertical line on some of the profile likelihood plots is the estimated upper end point at the MLE. The blue dotted and dashed lines are the thresholds to obtain the 99% and 95% confidence intervals respectively.

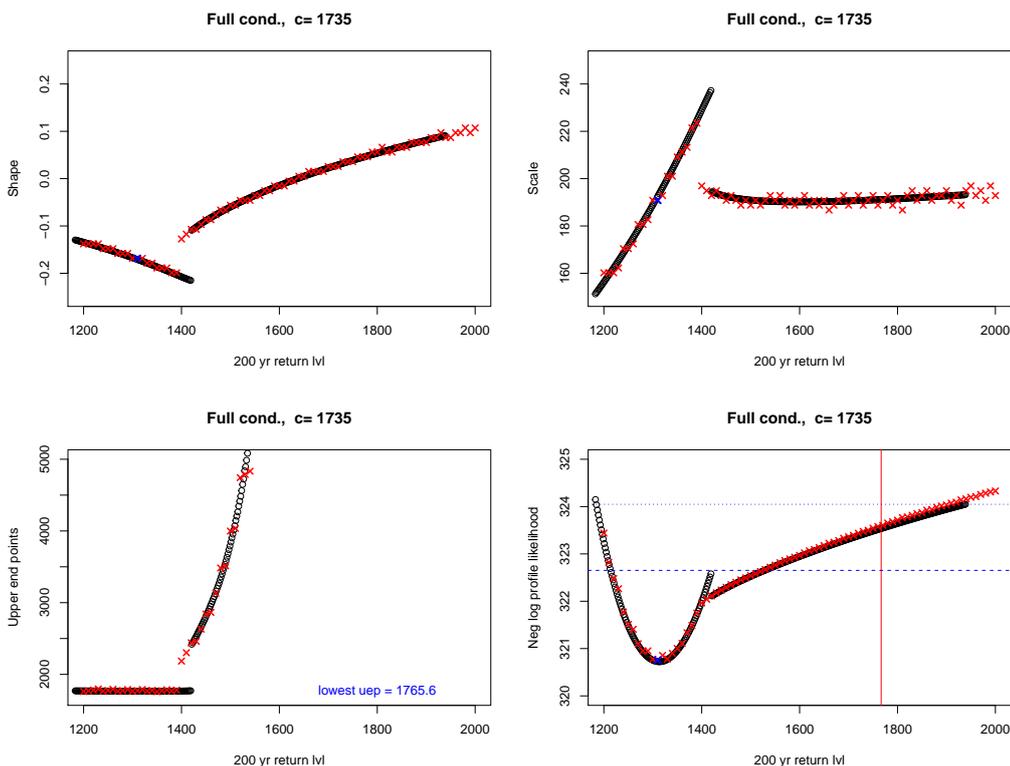


Figure B.3.4: As Figure B.3.3 but with  $c = 1735$ .

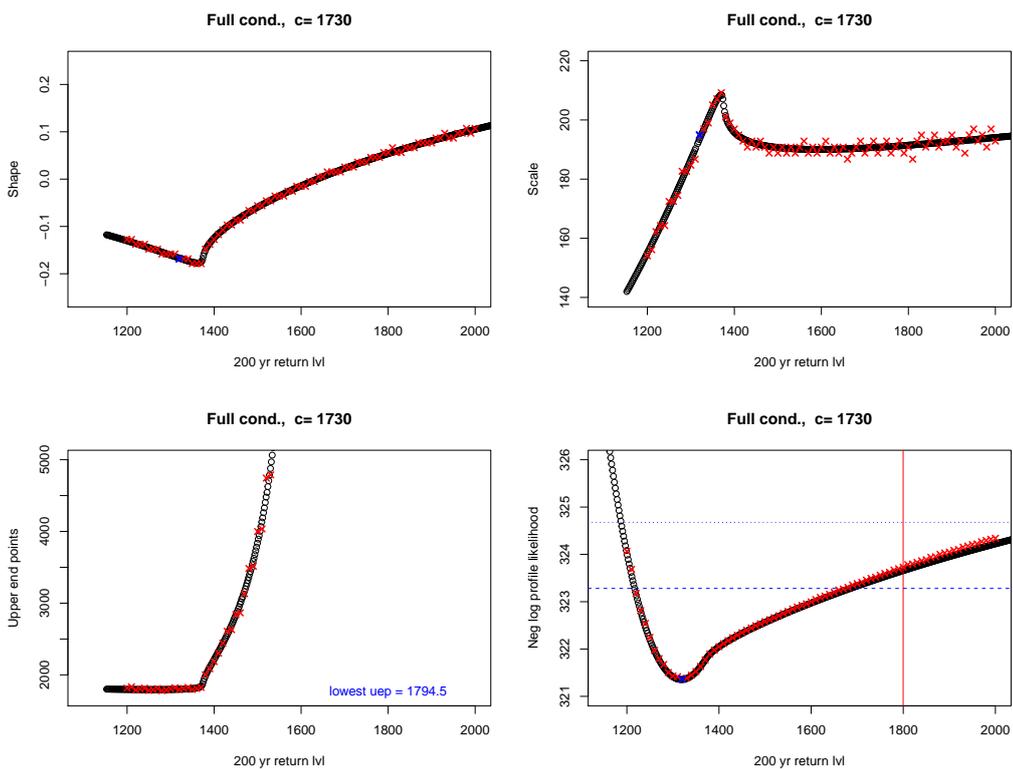


Figure B.3.5: As Figure B.3.3 but with  $c = 1730$ .

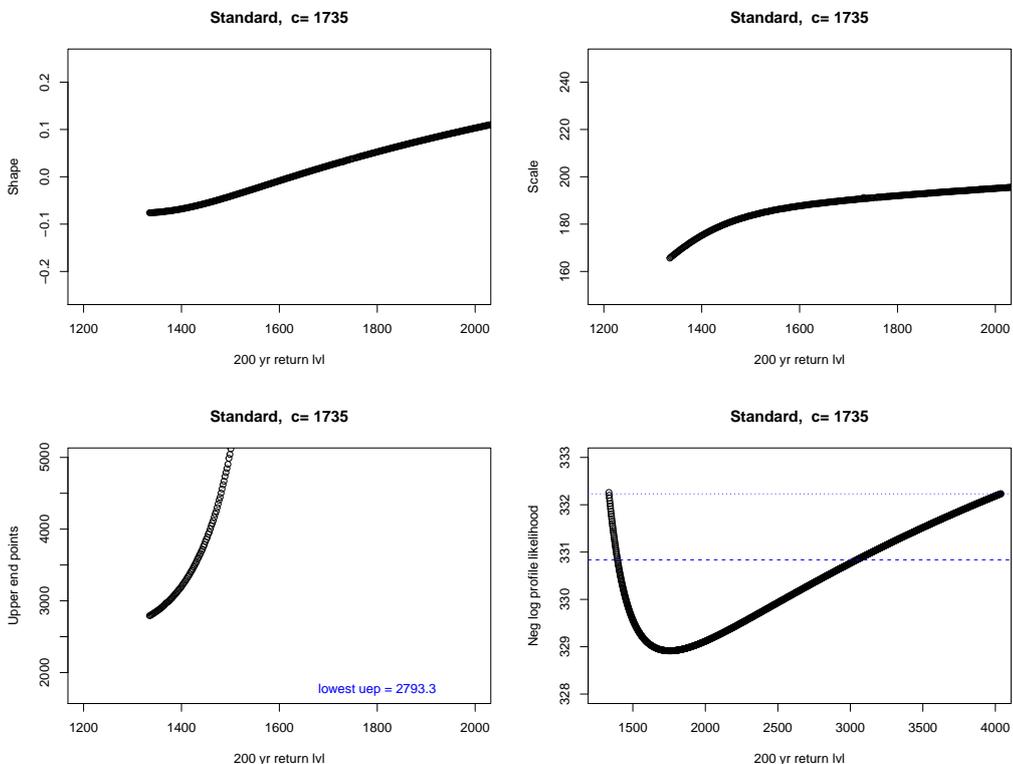


Figure B.3.6: Top: MLE of scale(left) and shape(right) parameters given the  $x_{200}$  value on the  $x$ -axis. Bottom: Estimated upper end point (left) and profile likelihood (right) over different  $x_{200}$  values.

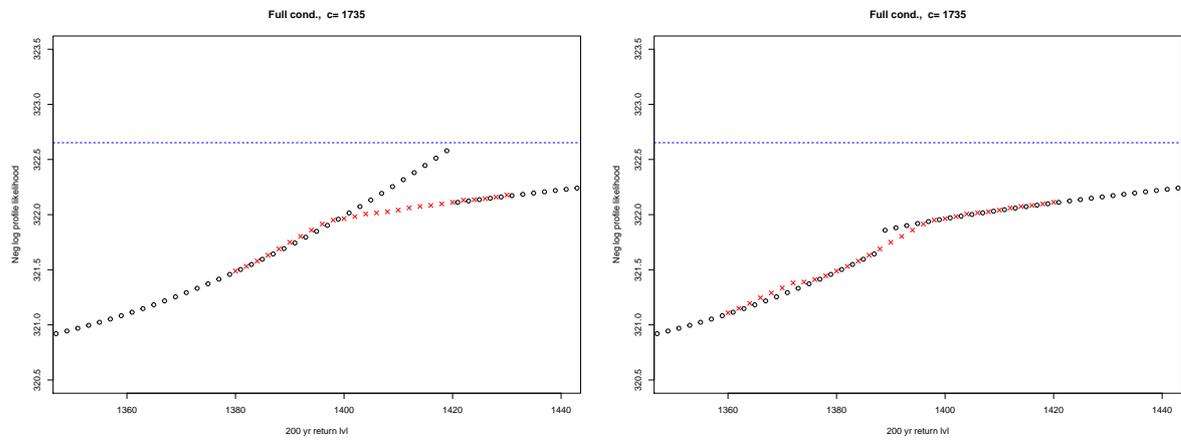


Figure B.3.7: Full conditioning profile likelihood,  $Pl(x_{200})$ , for values of  $x_{200}$  around the discontinuity without/with a shift away from boundary in the initial parameters (left/right). The black/red points refer to the optimisation/grid method respectively.

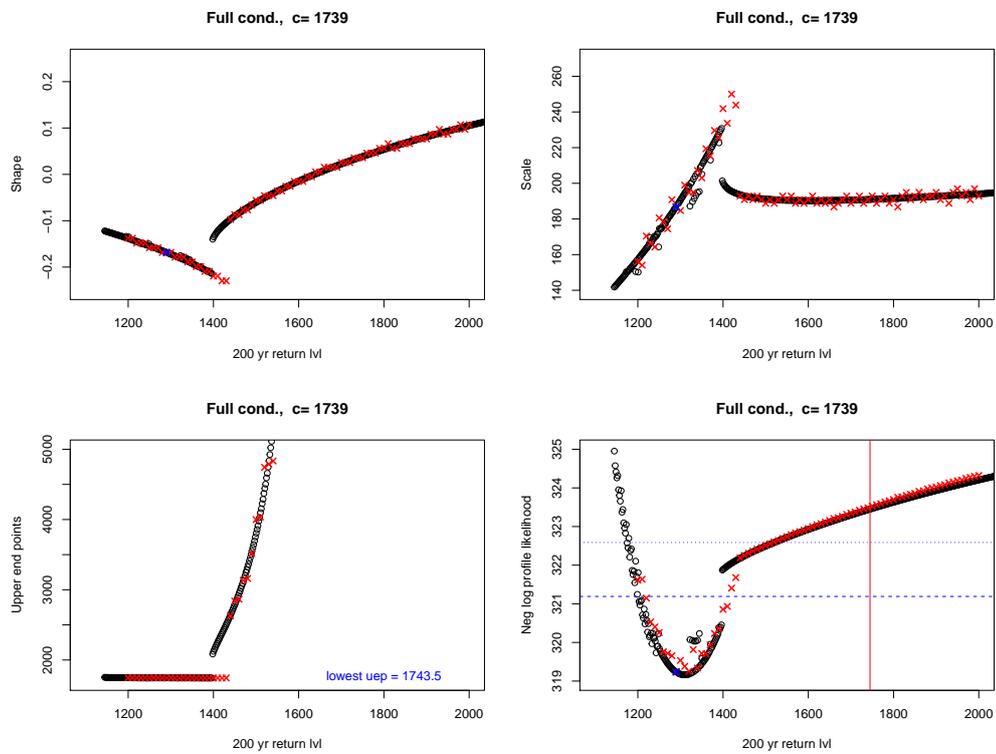


Figure B.3.8: As Figure B.3.3 but with initial parameters shifted away from the boundary in the optimisation method.

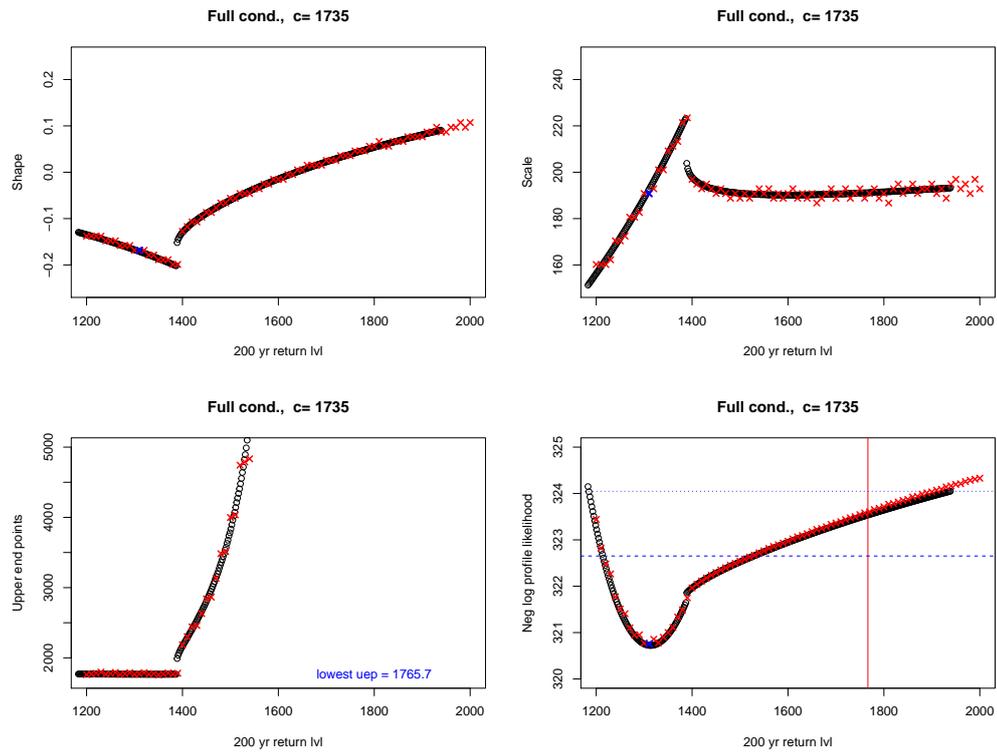


Figure B.3.9: As Figure B.3.8 but with  $c = 1735$ .

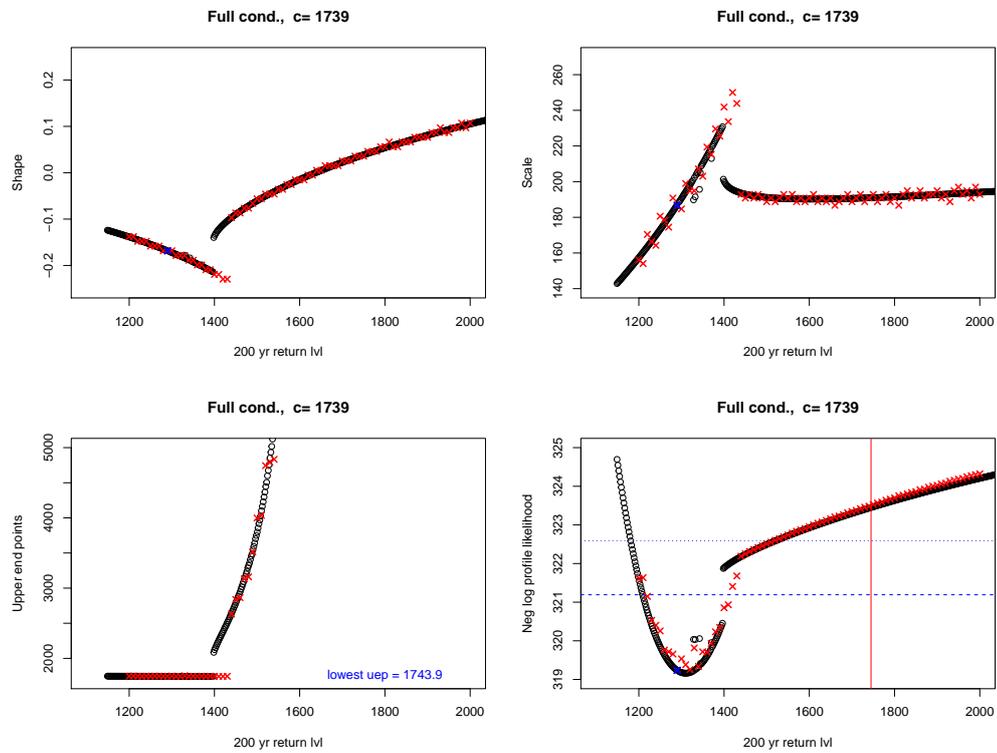


Figure B.3.10: As Figure B.3.8 but also using the GP approximation when  $G(c)$  close to 1.

# Appendix C

## Appendix to Chapter 5

### C.1 Proof of Hoeffding's Inequality (Theorem 5.3.4.1)

**Proof** For some fixed  $\lambda > 0$  (chosen later) we have by Chernoff's inequality (5.3.7):

$$\begin{aligned}\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \exp(-\lambda nt) \prod_{i=1}^n \mathbb{E}[\exp(\lambda(X_i - \mathbb{E}[X_i]))] \\ &= \exp(-\lambda nt) \prod_{i=1}^n \mathbb{E}[\exp(\lambda Y_i)].\end{aligned}\tag{C.1.1}$$

Now we prove an upper bound on  $\mathbb{E}[e^{\lambda Y}]$  known as Hoeffding's Lemma (as in proof of Theorem 2 in Hoeffding (1963)). For any real valued random variable,  $Y = X - \mathbb{E}[X]$ , with  $b \leq X \leq c$ :

$$\begin{aligned}\mathbb{E}[e^{\lambda Y}] &\leq \frac{c - \mathbb{E}[X]}{c - b} e^{\lambda(b - \mathbb{E}[X])} + \frac{\mathbb{E}[X] - b}{c - b} e^{\lambda(c - \mathbb{E}[X])} \quad ((5.3.2) \text{ applied to } Y) \\ &= \exp\left(\log\left[e^{\lambda(b - \mathbb{E}[X])} (1 - \omega) + \omega e^{\lambda(c - b)}\right]\right) \quad (\text{where } \omega = (\mathbb{E}[X] - b)/(c - b)) \\ &= \exp(g(z)),\end{aligned}$$

where  $g(z) = -\omega z + \log(1 - \omega + \omega e^z)$  and  $z = \lambda(c - b)$ . The Taylor expansion of  $g(z)$  around 0 is  $g(z) = g(0) + zg'(0) + \frac{z^2}{2}g''(\eta)$  for some  $\eta$  between 0 and  $z$ . Now  $g(0)$  and  $g'(0)$  are 0 and  $g''(\eta) \leq \frac{1}{4}$  so  $g(z) \leq \frac{z^2}{8}$ . Thus,

$$\mathbb{E}[e^{\lambda Y}] \leq \exp\left(\frac{\lambda^2(c - b)^2}{8}\right).\tag{C.1.2}$$

So, following on from (C.1.1):

$$\begin{aligned}\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) &\leq \exp(-\lambda nt) \prod_{i=1}^n \left(1 + \frac{\lambda^2(c_i - b_i)^2}{8}\right) \quad (\text{Hoeffding's Lemma (C.1.2)}) \\ &\leq \exp(-\lambda nt) \exp\left(\frac{\lambda^2 \sum_{i=1}^n (c_i - b_i)^2}{8}\right),\end{aligned}$$

using the inequality  $1 + x \leq e^x$ . The right hand side is minimised at  $\lambda = 4nt / \sum_{i=1}^n (c_i - b_i)^2$  and thus we arrive at (5.3.9).  $\blacksquare$

## C.2 Proof and exploration of Proposition 5.3.9.1

Define  $h(x) = (1 + x) \log(1 + x) - x$ ,

$$p_B(t) := \exp\left(-\frac{\text{Var}(S_n)}{a_{\max}^2} h\left(\frac{na_{\max}t}{\text{Var}(S_n)}\right)\right) \quad \text{and} \quad p_{CH}(t) := \exp\left(-nD_{KL}\left(\frac{\bar{c}\bar{p} + t}{c_{\max}} \middle| \middle| \bar{p}\right)\right).$$

Then we have the following two inequalities:

### Bennett

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq p_B(t) \quad 0 < t < \bar{a} = \bar{c} - \bar{c}\bar{p}. \quad (\text{C.2.1})$$

### Chernoff-Hoeffding

$$\mathbb{P}(S_n \geq \mathbb{E}[S_n] + nt) \leq p_{CH}(t) \quad \bar{p}c_{\max} - \bar{c}\bar{p} < t < c_{\max} - \bar{c}\bar{p}. \quad (\text{C.2.2})$$

For small  $p$  and when all  $c_i = c$  the Bennett bound is almost equal to *but slightly larger* than the Chernoff-Hoeffding bound.

**Proposition** (Recalling Proposition 5.3.9.1)

Let  $p_{CH}(t)$  and  $p_B(t)$  as defined in equations (C.2.2) and (C.2.1). Define  $\rho := (1 - p_{\min}) / (1 - \bar{p}^2 / \bar{p})$ . If  $c_i = c \forall i$  then

$$-\frac{1}{n} \log p_{CH}(t) = \bar{p}h\left(\frac{t}{\bar{c}\bar{p}}\right) + \frac{(t/c)^2}{2(1 - \bar{p})} + \mathcal{O}\left(\left[\frac{t}{c}\right]^3\right).$$

for  $t \leq \epsilon \bar{a}$  for some  $0 < \epsilon < 1$ . Furthermore, if  $t / [\bar{c}\bar{p}] > b$  for any fixed  $b > 1$  then

$$-\frac{1}{n} \log p_B(t) = \bar{p}h\left(\frac{t}{\bar{c}\bar{p}}\right) \times \frac{1}{1 - p_{\min}} \rho^{\{t/[\bar{c}\bar{p}] - \log(1+t/[\bar{c}\bar{p}])\}/h(t/[\bar{c}\bar{p}])} \times \{1 + \mathcal{O}([\log \rho]^2)\}.$$

Both bounds are  $\bar{p}h(t/[\bar{c}\bar{p}])$  modulo a small correction. For the correction term in  $p_B$ ,  $p_{\min}\bar{p} \leq \bar{p}^2 \leq p_{\max}\bar{p}$ , so  $1 \leq \rho \leq (1 - p_{\min}) / (1 - p_{\max})$  and  $\rho = 1$  when  $p_i = p \forall i$ . In the latter case  $-\frac{1}{n} \log p_B = \bar{p}h(t/(\bar{c}\bar{p})) / (1 - p)$ . Further,  $[x - \log(1 + x)] / h(x)$  is positive and decreasing from 1 (when  $x = 0$ ) to 0 as  $x \rightarrow \infty$ , so the correction term in  $-\frac{1}{n} \log p_B(t)$  becomes irrelevant

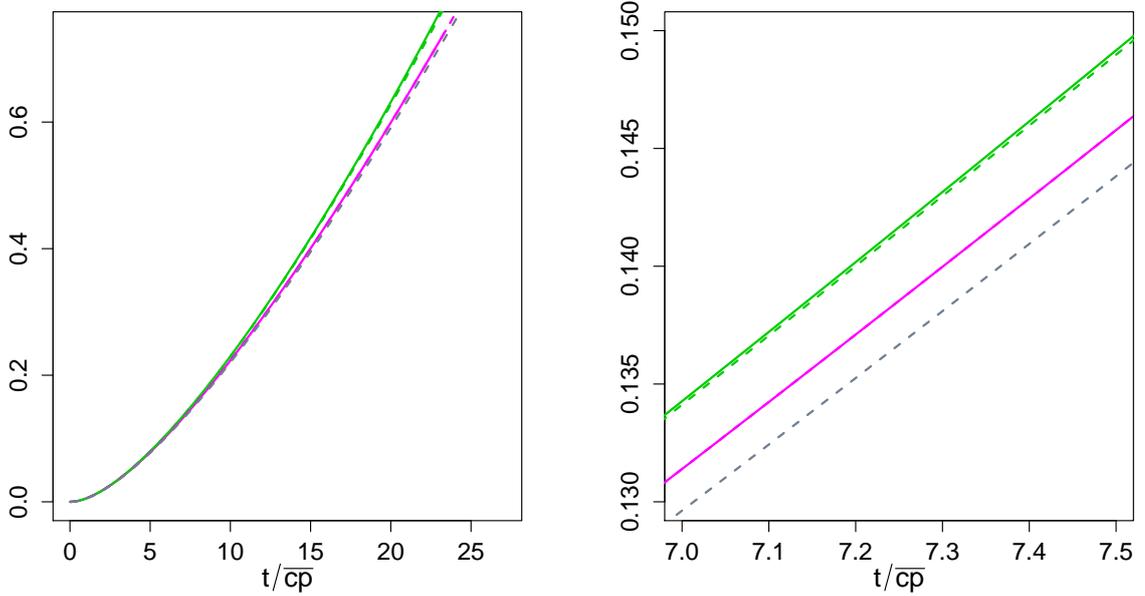


Figure C.2.1: *Bounds and approximations for our example with fixed  $p = 0.0135$ . The dark grey dashed line is  $\bar{p}h(t/[c\bar{p}])$ , the solid lines are the actual CH (green) and Bennett (magenta) bounds. The dashed lines are as in the proposition but without the  $\mathcal{O}$  term with colour corresponding to the relevant concentration inequality.*

as  $t/[c\bar{p}]$  becomes large. Thus, when  $t$  and  $c$  are fixed and  $p \rightarrow 0$ , the Bennett correction becomes negligible since  $\rho \rightarrow 1$  and the exponent of  $\rho \rightarrow 0$ . In contrast the positive correction term for  $-(1/n) \log p_{CH}$  only becomes negligible as  $t/c$  becomes small and is larger than that for  $p_B$  when  $p_i = p \forall i$  or  $t$  is large enough.

Figures C.2.1 and C.2.2 show the bounds and approximations based on the proposition (including  $\bar{p}h(t/[c\bar{p}])$  in dark grey) and Figure C.2.3 shows the difference:  $-\frac{1}{n} \log(p_{CH}) - (-\frac{1}{n} \log(p_B))$ . The solid and dashed green lines in Figures C.2.1 and C.2.2 are  $-\frac{1}{n} \log p_{CH}(t)$  and  $-\frac{1}{n} \log p_{CH}(t)$  without the  $\mathcal{O}$  term:  $ph(t/(c\bar{p})) + \frac{1}{2}(t/c)^2/(1-\bar{p})$ , respectively. The solid and dashed magenta lines are  $-\frac{1}{n} \log p_B(t)$  and  $-\frac{1}{n} \log p_B(t)$  without the  $\mathcal{O}$  term:  $\bar{p}h(t/(c\bar{p})) \times \rho^{\{t/[c\bar{p}] - \log(1+t/[c\bar{p}])\}h(t/[c\bar{p}])} / (1 - p_{\min})$ , respectively.

**Example (Figure C.2.1):**  $c = 1$ ,  $p_i = 0.0135 \forall i$ .

*In this case  $p_{\min} = \bar{p} = p$  so  $\rho = 1$  and the Bennett bound is  $\frac{p}{1-p}h\left(\frac{t}{p}\right)$  (magenta lines in Figure C.2.1 are identical). The correction term  $\frac{(t/c)^2}{2(1-\bar{p})}$  for the Chernhoff-Hoeffding bound is much larger than the correction of the Bennett bound; this difference increases with increasing  $t/p$  (Figure C.2.3 left panel).*

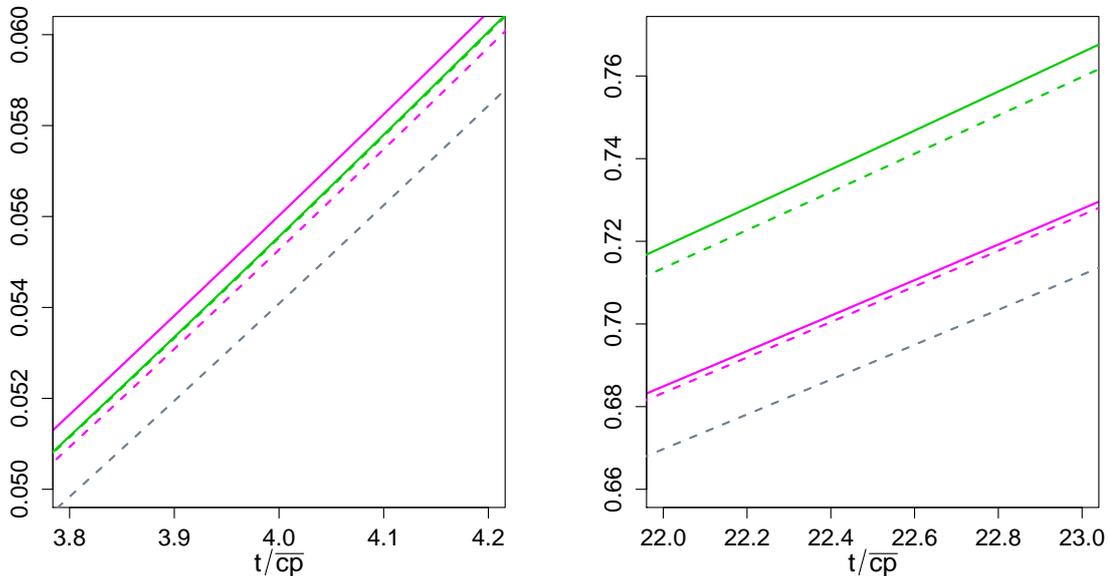


Figure C.2.2: *Bounds and approximations for our example with  $P \sim \text{Beta}(0.3, 22)$ . The dark grey dashed line is  $\bar{p}h(t/\bar{c}p)$ , the solid lines are the actual CH (green) and Bennett (magenta) bounds. The dashed lines are as in the proposition but without the  $\mathcal{O}$  term with colour corresponding to the relevant concentration inequality.*

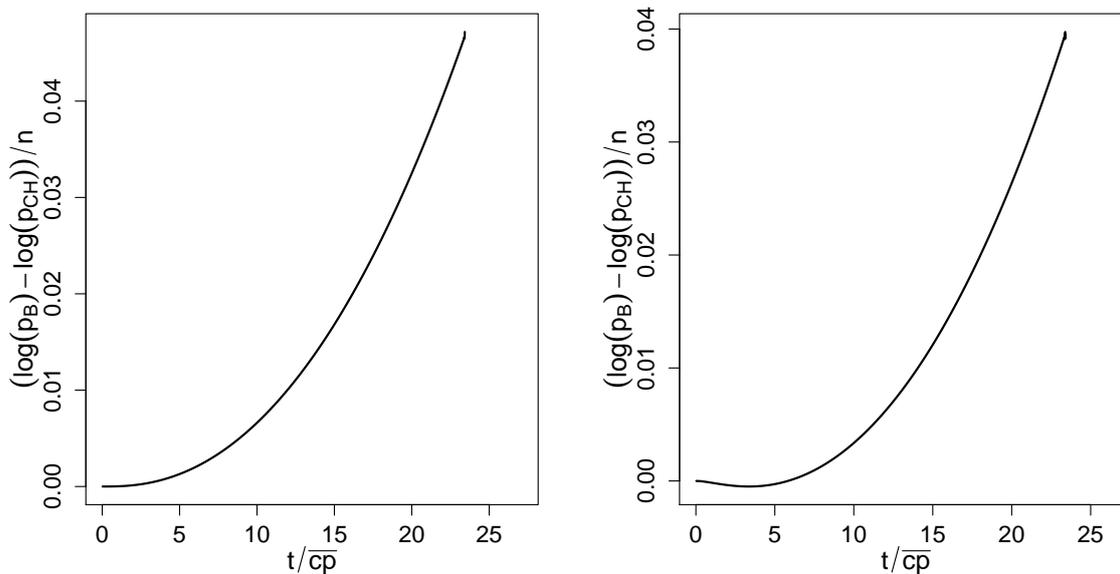


Figure C.2.3:  *$(-\log(p_{CH}) - (-\log(p_B)))/n$  against  $t/\bar{c}p$  with  $c = 1$ . Left: Fixed  $p = 0.0135$ , Right:  $P \sim \text{Beta}(0.3, 22)$ . Positive values indicate that the CH bound is smaller than the Bennett bound.*

**Example (Figure C.2.2):**  $c = 1, P \sim \text{Beta}(0.3, 22)$

We have  $\bar{p} \approx 0.0135, \bar{p}^2 \approx 0.00075, p_{\min} = 5 \times 10^{-12}$  so  $\rho \approx 1.06$  and  $0 < t < 25c\bar{p} = \bar{a}$ . For large  $t$  we have the same behaviour as for fixed  $p$  (Figure C.2.2 right) but for  $t < 6$  the correction for CH is less than than for Bennett making the former bound larger for small  $t$  (Figure C.2.3 right and Figure C.2.2 left).

**Proof** [Proposition 5.3.9.1] Let  $\ell_B := -\frac{1}{n} \log p_B(t)$  and  $\ell_{CH} := -\frac{1}{n} \log p_{CH}(t)$ , where we suppress the dependency on  $t$  for simplicity of presentation. Firstly, since  $c_i = c$  for  $i = 1, \dots, n$ ,  $a_{\max} = c(1 - p_{\min})$ , and

$$\frac{\text{Var}(S_n)}{na_{\max}} = \frac{1}{nc(1 - p_{\min})} c^2 \sum_{i=1}^n p_i(1 - p_i) = \frac{c(\bar{p} - \bar{p}^2)}{1 - p_{\min}} = \frac{c\bar{p}}{\rho}.$$

Thus,

$$\ell_B = \frac{1}{\rho(1 - p_{\min})} \bar{p} h\left(\rho \frac{t}{c\bar{p}}\right).$$

Set  $g(y) = \log h(\exp(y)) = \log[\{1 + \exp(y)\} \log\{1 + \exp(y)\} - \exp(y)]$ . Then, after some algebra,

$$g'(y) = \exp(y) \frac{\log\{1 + \exp(y)\}}{h(\exp(y))} = \frac{x \log(1 + x)}{h(x)},$$

$$g''(y) = \frac{\exp(y) [\log\{1 + \exp(y)\}]^2 - \exp(3y)/\{1 + \exp(y)\}}{[h(\exp(y))]^2} = \frac{x [\log(1 + x)]^2 - x^3/(1 + x)}{h(x)^2},$$

where  $x = \exp(y)$ . Now  $|g''(y)|$  is bounded for  $y \geq 0$ . Define  $x_0 = t/(c\bar{p})$  and  $y_0 = \log(x_0)$ , then

$$\log h(\rho x_0) = g(y_0 + \log \rho) = g(y_0) + \log \rho \times g'(y_0) + r,$$

where the remainder term,  $r = \mathcal{O}((\log \rho)^2)$  provided  $x_0 \geq b > 1$ . Thus,

$$h(\rho x_0) = h(x_0) \times \rho^{g'(y_0)} \times \exp(r) = h(x_0) \times \rho^{x_0 \log(1+x_0)/h(x_0)} \exp(r).$$

Bringing this all together we obtain:

$$\begin{aligned} \ell_B &= \frac{1}{1 - p_{\min}} \bar{p} h\left(\frac{t}{c\bar{p}}\right) \rho^{x_0 \log(1+x_0)/h(x_0) - 1} \exp(r) \\ &= \frac{1}{1 - p_{\min}} \bar{p} h\left(\frac{t}{c\bar{p}}\right) \rho^{[x_0 - \log(1+x_0)]/h(x_0)} \exp(r), \end{aligned}$$

as required.

Since  $c_i = c$  for  $i = 1, \dots, n$ ,

$$\begin{aligned} \ell_{CH} &= D_{KL} \left( \frac{c\bar{p} + t}{c_{\max}} \middle| \middle| \bar{p} \right) = D_{KL} \left( \bar{p} + \frac{t}{c} \middle| \middle| \bar{p} \right) \\ &= \left( \bar{p} + \frac{t}{c} \right) \log \left( 1 + \frac{t}{c\bar{p}} \right) + \left( 1 - \bar{p} - \frac{t}{c} \right) \log \left( \frac{1 - \bar{p} - t/c}{1 - \bar{p}} \right). \end{aligned}$$

Now  $h'(x) = \log(1+x)$  and  $h''(x) = 1/(1+x)$  so  $h'(0) = 0$  and  $h''(0) = 1$ . For  $x \geq -\varepsilon$  where  $\varepsilon < 1$ ,  $|h'''| \leq 1/(1-\varepsilon)^2$ ; hence, the Taylor expansion gives  $h(x) = x^2/2 + \mathcal{O}(x^3)$ . Setting  $x = -(t/c)/(1-\bar{p})$  we have that for  $x \geq -\varepsilon$ , *i.e.*,  $t \leq \varepsilon\bar{a}$

$$\begin{aligned} \left( 1 - \bar{p} - \frac{t}{c} \right) \log \left( \frac{1 - \bar{p} - t/c}{1 - \bar{p}} \right) &= (1 - \bar{p})(1+x) \log(1+x) = (1 - \bar{p})[h(x) + x] \\ &= (1 - \bar{p}) \left[ x + x^2/2 + \mathcal{O}(x^3) \right] \\ &= (1 - \bar{p}) \left[ -\frac{(t/c)}{2(1 - \bar{p})} + \frac{(t/c)^2}{(1 - \bar{p})^2} + \mathcal{O}((t/c)^3) \right] \\ &= -\frac{t}{c} + \frac{(t/c)^2}{2(1 - \bar{p})} + \mathcal{O} \left( \frac{t^3}{c^3} \right). \end{aligned}$$

So

$$\ell_{CH} = \bar{p}h \left( \frac{t}{c\bar{p}} \right) + \frac{(t/c)^2}{2(1 - \bar{p})} + \mathcal{O} \left( \frac{t^3}{c^3} \right).$$

■

### C.3 Issues with Jebara's Bennett refinement

The inequality of Jebara (2018) appears to be strictly tighter than Bennett's for fixed  $c$  and  $p$  as claimed, however, as we saw in §5.3.8, when  $c$  and/or  $p$  are not fixed with some  $p$  being extremely small there is a considerable number of cases for which the bound does not perform as well as Bennett and may even do much worse. If  $\lambda_Q^*$  (5.3.25) were the minimiser of (5.3.24), the bound would indeed be strictly tighter, however, in some settings this is not the case. This is because the true minimum of (5.3.24), which we will denote  $\lambda^*$ , is not the same as the minimum of the sum of the quadratic bounds,  $\lambda_Q^*$ .

Looking more in depth at the performance of *Jebara's* bound for different  $n$ ,  $p$ ,  $c$  and  $t$  values we can construct many more examples where *Jebara* is not performing optimally, particularly when  $p$  is small. For simplicity we consider cases where  $p$  is fixed and let  $t = \bar{a}/k$

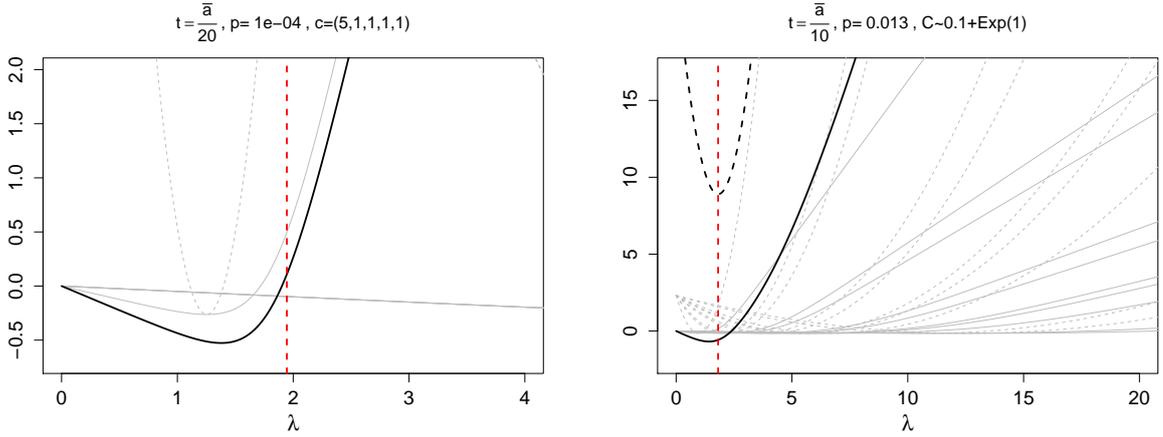


Figure C.3.1: *Examples of Jebara's bound being sub-optimal. The curve  $\sum_{i=1}^n b_i(\lambda)$  (solid black) and curves  $b_i(\lambda), i = 1, \dots, n$  (solid grey) and their quadratic bounds (dashed grey) against  $\lambda$ . The sum of the quadratic bounds are in dashed black and the dashed red line indicates  $\lambda_Q^*$ . Left (Example 1):  $n = 5, p = 10^{-4}, c_1 = 5, c_i = 1, i = 2, \dots, 5, t = \bar{a}/20$ . Right (Example 2):  $n = 10, p = 0.013, t = \bar{a}/10$ .*

for some  $k > 1$  so Jebara's bound is

$$\mathbb{P}(S_n > \mathbb{E}[S_n] + nt) \leq \exp\left(\sum_{i=1}^n b_i(\lambda_Q^*)\right) \quad (\text{C.3.1})$$

where

$$b_i(\lambda) = \log\left(\frac{p}{1-p} \left[e^{\lambda c_i(1-p)} - 1 - \lambda c_i(1-p)\right] + 1\right) - \frac{\lambda c_i(1-p)}{k},$$

$$\lambda_Q^* = \frac{\sum_{i=1}^n c_i^2 \lambda_i}{\sum_{i=1}^n c_i^2} = c_1 \lambda_1 \frac{\sum_{i=1}^n c_i}{\sum_{i=1}^n c_i^2}$$

and

$$\lambda_i = \frac{1}{c_i} \left\{ \frac{k-1}{1-p} + \frac{1}{p_i} - \frac{1}{1-p} W\left(\exp\left[k - 2 + \frac{1}{p} + \log(k-1)\right]\right) \right\}.$$

The quadratic bounds (5.3.26) are, for each  $i$ :

$$\frac{c_i^2(1-p)^2(\lambda - \lambda_i)^2}{1 - e^{1-1/p}} + b_i(\lambda_i).$$

We can also write  $nt$  as some multiple,  $n_{sd}$ , of the total standard deviation,  $\sigma = \sqrt{p(1-p) \sum c_i^2}$ :

$$n_{sd} = \frac{\bar{c}}{k} \sqrt{\frac{n(1-p)}{c^2 p}}.$$

**Example 1:**

First, we consider a simple if extreme example, with  $n = 5$ , for which Jebara's refinement results in a probability bound larger than 1. Consider 5 binary variables,  $X_i, i = 1, \dots, 5$  with  $0 \leq X_1 \leq c_1 = 5$ ,  $0 \leq X_i \leq c_i = 1$  for  $i = 2, 3, 4, 5$  and  $p = 10^{-4}$ . Suppose we want to find an upper bound on the probability:  $\mathbb{P}(S_5 - \mathbb{E}[S_5] \geq 5t)$ , where  $S_5 = \sum_{i=1}^5 X_i$  and  $t = \bar{a}/20$  (so  $5t \approx 8.4\sigma$ ). Figure C.3.1 (right panel) shows the individual  $b_i(\lambda)$  against  $\lambda$  (solid grey), the quadratic bound on  $b_1(\lambda)$  (dashed grey), the sum,  $\sum_{i=1}^5 b_i(\lambda)$ , to be minimised is in black and the minimiser of the sum of the quadratic bounds,  $\lambda_Q^*$  (red dashed). The sum of the quadratic bounds is not shown on the plot since it is 42 at its minimum,  $\lambda_Q^*$ . The curve,  $\sum_{i=1}^5 b_i(\lambda)$ , is minimised at  $\sim 1.38$  leading to a bound of  $\sim 0.591$ . However  $\lambda_Q^* \approx 1.95$  and inserting this into (C.3.1) leads to a value greater than 1.

**Example 2:**

Our next example is of a more realistic setting similar to that of the simulations in §5.3.9 Figure 5.3.1. We consider 10 binary random variables  $X_i \sim c_i \text{Bern}(p)$  with  $p = 0.013$  and  $C_i \sim \text{Exp}(1)$  and let  $t = \bar{a}/10$  (which results in  $n_{sd} \approx 2.1$ ). Figure C.3.1 (left panel) shows the individual  $b_i(\lambda)$  against  $\lambda$  (solid grey), their quadratic bounds (dashed grey), the sum,  $\sum_{i=1}^{10} b_i(\lambda)$ , to be minimised is in black and the red dashed line indicates  $\lambda_Q^*$ , the minimiser of the sum of the quadratic bounds (black dashed). It is clear that  $\sum_{i=1}^{10} b_i(\lambda_Q^*) > \sum_{i=1}^{10} b_i(\lambda^*)$ . This results in an upper bound on  $\mathbb{P}(S_{10} \geq \mathbb{E}[S_{10}] + 10t)$  of approximately 0.561 using  $\lambda_Q^*$  compared to approximately 0.505 at the optimal  $\lambda^*$ . Bennett's bound in this setting is 0.553 so Jebara's Bennett refinement is slightly looser than Bennett.

In the loss estimation setting we will have a range of  $c$  values and many small  $p$  values and so it is likely that Jebara's Bennett refinement will not provide tight bounds reliably. Also, in all cases considered, when the refinement is tighter, it is only a slight improvement on the Bennett refinement we introduced in §5.4.2. For these reasons we do not consider Jebara further.

# Appendix D

## Appendix to Chapter 6

### D.1 Investigating the effect of event set size

In §6.1 we briefly discussed the effect of the number of years,  $n_y$ , in the event set on the discard sets calculated with our exclude method. In particular, the percentage of discarded years generally reduces in size when the number of years in the event set is increased, especially for low return periods. We now investigate this observation by looking in detail at the behaviour of the two terms,  $b_1(u; t)$  (6.2.3) and  $b_2(u; t)$  (6.2.4), of the non-exceedance probability for a set of  $n_y$  years,  $\mathcal{Y}$ , and subset of  $\tilde{n}_y$  years,  $\mathcal{Y}_{\tilde{n}_y} \subset \mathcal{Y}$ . We illustrate with an example of the estimation of the 5-year return level using the exclude method with Bennett's inequality for  $\tilde{n}_y = 500$  and  $n_y = 1000$ . We choose our subset,  $\mathcal{Y}_{500}$ , of years from the 1000 year subset by ordering the years by expected loss and choosing every other year in this sequence to be in the subset  $\mathcal{Y}_{500}$ . In this way we ensure that our subset of years has similar characteristics to the full set of years, in particular we should have a similar percentage of years which are prone to high losses etc.

The discard set is only present in term  $b_2(u; t)$  of the non-equivalence probability bound; for a fixed threshold,  $u$ , this term increases as the discard set size increases whereas for a fixed discard set,  $\mathcal{D}$ , it is monotonically decreasing in  $u$ . Thus the largest discard set possible is restricted by how large the threshold,  $u$ , can be. Figure D.1.1 shows  $b_2(u; t)$ , with  $\varepsilon_{+,y}^u$  determined by Bennett's inequality, plotted against the threshold for a range of discard set sizes (indicated by different coloured curves). The dashed line indicates the restriction on  $b_2(u; t)$  such that  $b_1(u; t) + b_2(u; t) < \varepsilon_\circ$  with  $\varepsilon_\circ = 0.001$ . The restriction on  $u$  comes from the term  $b_1(u; t)$  in the non-equivalence probability bound; the term  $b_1(u; t)$  will only be smaller than  $\varepsilon_\circ$  as long as the threshold  $u$  is small enough such that all  $\varepsilon_{-, (k)}^u$  in the product are less than  $\varepsilon_\circ$ . We see that a slightly larger percentage of years (approximately 27.5% compared with 23.5%) can be discarded from the subset  $\mathcal{Y}_{500}$  than from the full 1000-year set  $\mathcal{Y}$ . This is

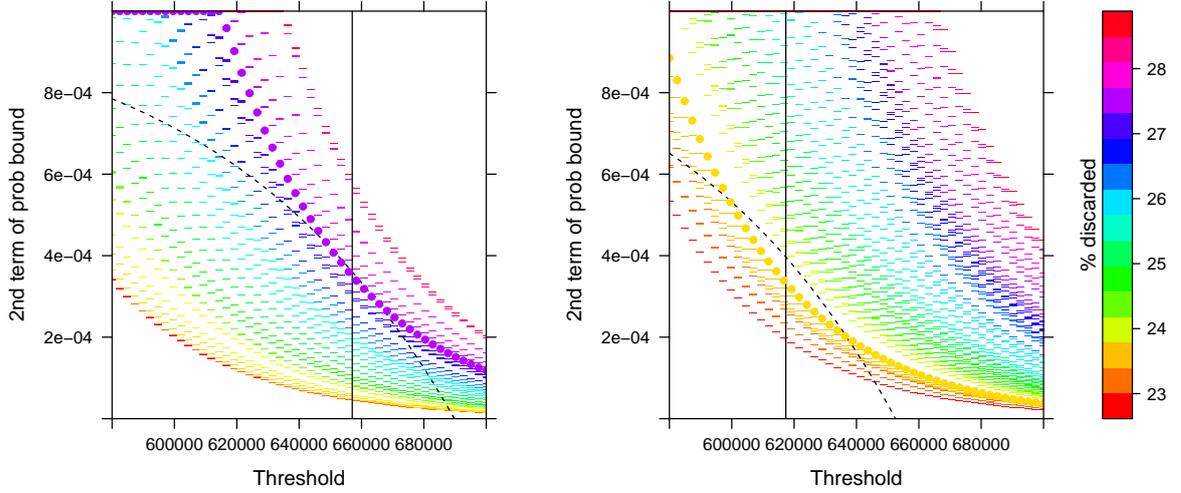


Figure D.1.1: The second term of the non-equivalence probability bound,  $b_2(u; t)$ , for thresholds in the neighbourhood of the ‘optimal’ threshold. Each unique colour corresponds to a discard set size. Left:  $n_y = 500$  with 28 discard set sizes from 115 (red) to 143 (pink), right:  $n_y = 1000$  with 55 discard set sizes from 230 (red) to 285 (pink). The black dashed line indicates the bound  $\varepsilon_0 - b_1(u; t)$  as in Figure D.1.2 bottom row. The coloured curve which represents the maximum discard set size is emphasised with filled circles.

due to the restriction imposed by the dashed line and that for a given percentage of discards the curve is higher for the full 1000-year set (in fact we will see that in our example  $b_2(u; t)$  for  $\mathcal{Y}$  is approximately double  $b_2(u; t)$  for  $\mathcal{Y}_{500}$ ). For insight we now discuss in detail how  $b_1(u; t)$  and  $b_2(u; t)$  change with differing event set size.

### Bound $b_1(u; t)$

As in §6.1 we sort the yearly losses of the years in  $\mathcal{Y}$  into order of descending expectation and denote the ordered set of random variables by  $\{\mathbb{S}\}_{k=1, \dots, n_y}$ . Let  $T_{\tilde{n}_y}^t$  denote the indices of  $\mathbb{S}_k$  which correspond to the top  $\lceil \frac{\tilde{n}_y+1}{t} \rceil$  years, in terms of expectation, of the subset  $\mathcal{Y}_{\tilde{n}_y}$  and similarly denote the indices of the top  $\lceil \frac{n_y+1}{t} \rceil$  years, in terms of expectation, of  $\mathcal{Y}$  by  $T_{n_y}^t$  (note that  $T_{n_y}^t = \{1, \dots, \lceil \frac{n_y+1}{t} \rceil\}$  by definition). For our construction of the subset  $\mathcal{Y}_{500}$  we have  $T_{500}^t = \{1, 3, \dots, 2\lceil \frac{501}{t} \rceil - 1\}$ . First we consider the case where the ‘top’ years of the smaller event set form a subset of the ‘top’ years of the large event set, *i.e.*,  $T_{\tilde{n}_y}^t \subset T_{n_y}^t$ . This is true for our example due to the way we have chosen the years in the subset  $\mathcal{Y}_{500}$ . For the

large event set term is:

$$\begin{aligned}
 b_1^{n_y}(u; t) &= 1 - \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} (1 - \varepsilon_{-, (k)}^u) \\
 &= 1 - \prod_{k \in T_{n_y}^t \setminus T_{\tilde{n}_y}^t} (1 - \varepsilon_{-, (k)}^u) \prod_{k \in T_{\tilde{n}_y}^t} (1 - \varepsilon_{-, (k)}^u) \\
 &> 1 - \prod_{k \in T_{\tilde{n}_y}^t} (1 - \varepsilon_{-, (k)}^u) = b_1^{\tilde{n}_y}(u; t).
 \end{aligned}$$

That is, for a fixed threshold  $u$ ,  $b_1(u; t)$  is larger for the event set of  $n_y$  years than for the smaller set of years,  $\mathcal{Y}_{\tilde{n}_y}$ . Consequently, since  $b_1(u; t)$  is monotonically increasing in  $u$ , the maximum threshold for which  $b_1(u; t) < \varepsilon_o$  is smaller for larger  $n_y$ . The top row of Figure D.1.2 shows  $b_1(u; 5)$  against threshold for finding the 5-year return-level estimate (*i.e.*,  $t = 5$ ) when  $\tilde{n}_y = 500$  (left) and  $n_y = 1000$  (right). The dashed horizontal line is  $\varepsilon_o = 0.001$  and the solid horizontal line is the non-equivalence probability bound at the optimal threshold (vertical line) and discard set combination. It can be seen that the the first term of the non-equivalence probability bound reaches  $\varepsilon_o$  at a lower threshold for  $n_y = 1000$  ( $\sim 652000$ ) than for  $\tilde{n}_y = 500$  ( $\sim 670000$ ).

For the subset,  $\mathcal{Y}_{500}$ , used here we have  $T_{500}^t \subset T_{1000}^t$ , however, for another subset of years this may not be true. If  $T_{\tilde{n}_y}^t \not\subset T_{n_y}^t$  we cannot say whether  $b_1^{n_y}(u; t) > b_1^{\tilde{n}_y}(u; t)$  or vice versa, however, if we assume that the subset  $\mathcal{Y}_{\tilde{n}_y}$  is representative of the full set  $\mathcal{Y}$  we can deduce relations between  $b_1^{n_y}(u; t)$  and  $b_1^{\tilde{n}_y}(u; t)$ . In particular we assume that the average  $\varepsilon_{-, (k)}^u$  over  $k \in T_{\tilde{n}_y}^t$  and over  $k \in T_{n_y}^t$  are approximately equal, *i.e.*,  $\frac{1}{|T_{\tilde{n}_y}^t|} \sum_{k \in T_{\tilde{n}_y}^t} \varepsilon_{-, (k)}^u \approx \frac{1}{|T_{n_y}^t|} \sum_{k \in T_{n_y}^t} \varepsilon_{-, (k)}^u$ . Now  $\varepsilon_{-, (k)}^u \leq 1$  for all  $u > 0$  and  $k \in \{1, \dots, n_y\}$  and, moreover, for all acceptable thresholds and  $k \in T_{n_y} \cup T_{\tilde{n}_y}$  we have  $\varepsilon_{-, (k)}^u \leq \varepsilon_o$ , so we can approximate  $1 - \prod_{k \in T_{n_y}} (1 - \varepsilon_{-, (k)}^u) \approx \sum_{k \in T_{n_y}} \varepsilon_{-, (k)}^u$ . This approximation is more accurate the smaller the threshold  $u$  is. Then under the assumption on the means of  $\varepsilon_{-, (k)}^u$  we have:

$$b_1^{n_y}(u; t) \approx \sum_{k \in T_{n_y}} \varepsilon_{-, (k)}^u \approx \frac{|T_{n_y}^t|}{|T_{\tilde{n}_y}^t|} \sum_{k \in T_{\tilde{n}_y}^t} \varepsilon_{-, (k)}^u \approx \frac{\lceil \frac{n_y+1}{t} \rceil}{\lceil \frac{\tilde{n}_y+1}{t} \rceil} b_1^{\tilde{n}_y}(u; t). \quad (\text{D.1.1})$$

In our example the mean assumption is valid due to the way the years were selected so  $b_1^{1000}(u; 5) \approx \frac{\lceil \frac{1001}{5} \rceil}{\lceil \frac{501}{5} \rceil} b_1^{\tilde{n}_y}(u; t) \approx 2b_1^{500}(u; 5)$ . In the  $\tilde{n}_y = 500$  (top left) plot of Figure D.1.2 we see given the threshold the points are almost double the corresponding points in the

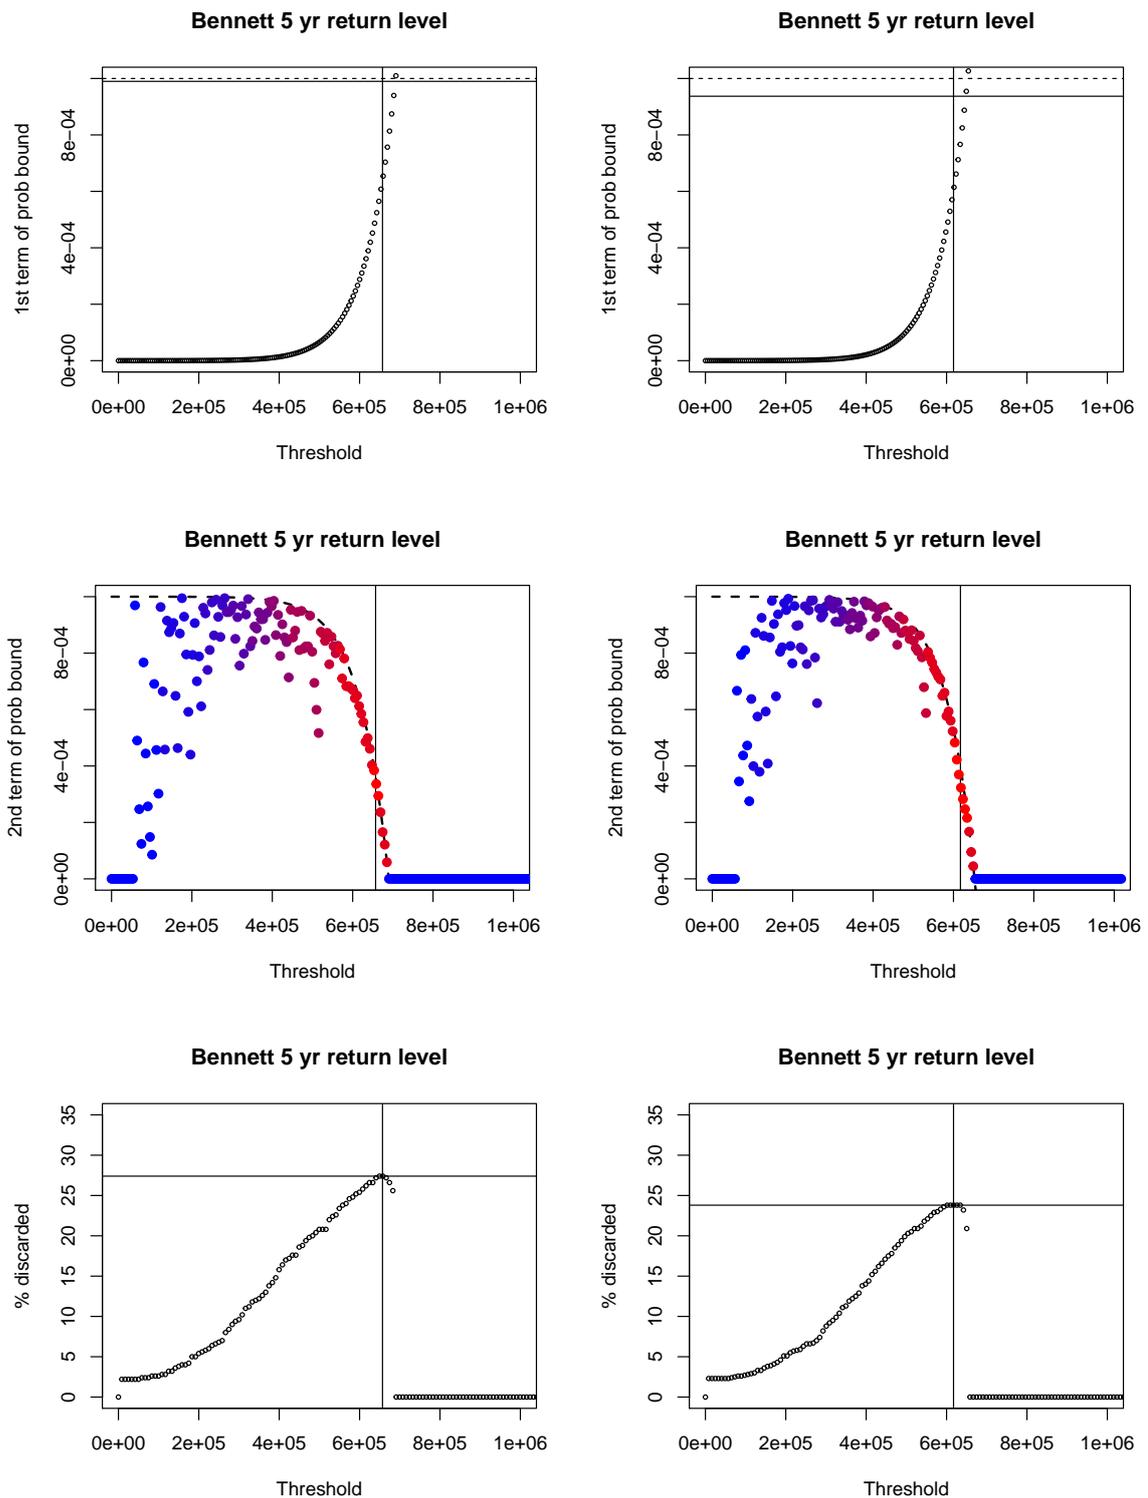


Figure D.1.2: Investigating the bound  $b_1(u; 5) + b_2(u; 5)$  (6.2.5) for the exclude method with Bennett's inequality and  $\epsilon_0 = 0.001$ . Left: 500-year event set, right: 1000-year event set. The vertical lines indicate the 'optimal' threshold. **Top:**  $b_1(u; t)$  from (6.2.3) against threshold,  $u$ . The horizontal lines indicate  $\epsilon_0$  (dashed) and  $b_1(u^*; 5) + b_2(u^*; 5)$  at the optimal combination of threshold,  $u^*$ , and discard set (solid). **Middle:**  $b_2(u; t)$  from (6.2.4) at the maximum number of discards possible given  $u$  against the threshold,  $u$ . The dashed line indicates the bound  $\epsilon_0 - b_1(u; t)$ . The points are coloured according to the percentage of discarded years, with blue indicating low numbers and red indicating high numbers. **Bottom:** Maximum number of discards such that the upper bound on non-equivalence probability is less than  $\epsilon_0$  against threshold. The horizontal line indicates the maximum percentage of discarded years.

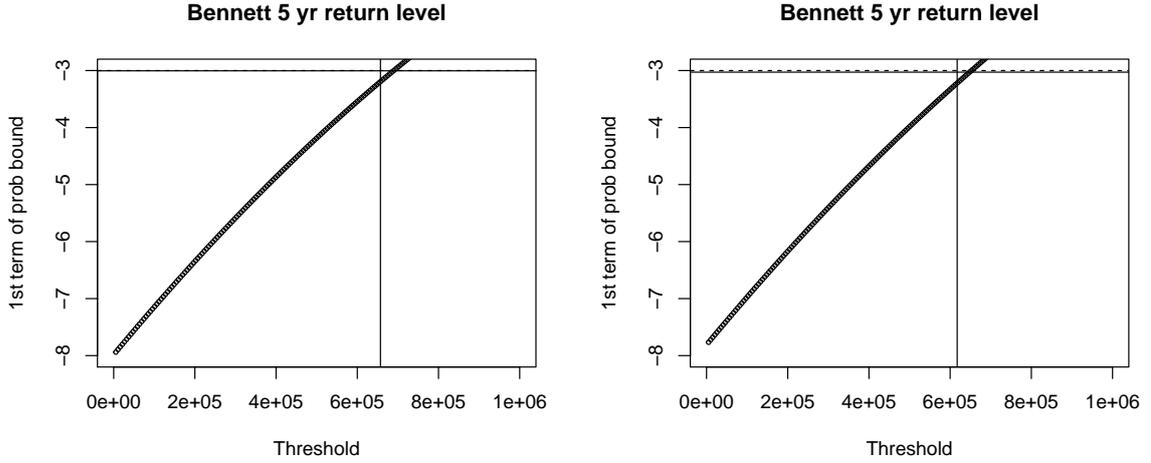


Figure D.1.3:  $\log_{10}(b_1(u; 5))$   $V$   $u$  when estimating the 5-year return level using the exclude method with Bennett's inequality and  $\varepsilon_o = 0.001$ . Left: Simulating 500 years, right: 1000 years. The vertical lines indicate the 'optimal' threshold. The horizontal lines indicate  $\varepsilon_o$  (dashed) and the non-equivalence probability bound, (6.2.5), at the optimal combination of threshold and discard set (solid).

$n_y = 1000$  (top right) plot. In Figure D.1.3  $\log_{10}(b_1(u; 5))$  is shown against threshold, the two curves have almost the same gradient but are shifted by  $\sim 0.25$  which is slightly less than the  $\log_{10}(2) \approx 0.3$  shift under approximation (D.1.1). For larger  $t$   $b_1^{n_y}(u; t)$  is still larger than  $b_1^{\tilde{n}_y}(u; t)$  but by a smaller factor.

So if  $T_{\tilde{n}_y}^t \subset T_{n_y}^t$  and/or  $\mathcal{Y}_{\tilde{n}_y}$  is representative of  $\mathcal{Y}$  in terms of lower tail probabilities then  $b_1^{n_y}(u; t) > b_1^{\tilde{n}_y}(u; t)$  given  $u$  and the maximum possible threshold will be larger for the subset  $\mathcal{Y}_{\tilde{n}_y}$ .

### Bound $b_2(u; t)$

We now discuss how the discard set size is affected by  $b_1(u; t)$  and the number of years in the event set. The middle row of Figure D.1.2 shows the second term of the non-equivalence probability bound,  $b_2(u; t)$  (6.2.4), at the maximum number of discards possible given the threshold indicated on the x-axis. The dashed line indicates the bound due to  $b_1(u; t)$ :

$$\sum_{y \in \mathcal{D}} \varepsilon_{+,y}^u < \varepsilon_o - b_1(u; t) = \varepsilon_o - 1 + \prod_{k=1}^{\lceil \frac{n_y+1}{t} \rceil} (1 - \varepsilon_{-, (k)}^u),$$

which is smaller the larger the threshold is. The points are coloured from blue through to red by increasing % of discarded years. The straight, almost vertical, line patterns seen in

some places (*e.g.*, around  $5 \cdot 10^5$  on the left plot) is due to a range of thresholds resulting in the same optimal discard set; within this range  $b_2(u; t)$  decreases as the threshold increases and the discard set is fixed. All the points on the plot correspond to an acceptable threshold and discard set (such that the bound (6.2.5) on the non-equivalence probability is less than  $\varepsilon_\circ$ ) but the largest discard set is found at the threshold indicated by the vertical line.

With a similar argument as used for the bound  $b_1(u; t)$  we find that  $b_2^{n_y}(u; t) \approx \frac{|\mathcal{D}|}{|\tilde{\mathcal{D}}|} b_2^{\tilde{n}_y}(u; t)$ . The assumption made here is that the mean of the upper tail probabilities,  $\varepsilon_{+,y}^u$ , of the low loss years (those which are discarded) are similar for  $y \in \mathcal{D}$  and  $y \in \tilde{\mathcal{D}}$ . If one discards the same percentage of years from both  $\mathcal{Y}$  and  $\mathcal{Y}_{\tilde{n}_y}$  then the discard sets are related as  $|\mathcal{D}| = \frac{n_y}{\tilde{n}_y} |\tilde{\mathcal{D}}|$  thus the bound  $b_2^{n_y}(u; t)$  will be  $\frac{n_y}{\tilde{n}_y}$  times larger than  $b_2^{\tilde{n}_y}(u; t)$  at a given threshold  $u$  (as seen in Figure D.1.1).

### Maximum percentage of years discarded

The percentage of discarded years increases as threshold increases while  $\varepsilon_\circ - b_1(u; t)$  is close to  $\varepsilon_\circ$  but begins to decrease when  $u$  is so large that  $\varepsilon_\circ - b_1(u; t)$  is close to 0. This behaviour can be seen in the bottom row of Figure D.1.2 which shows the maximum percentage of years discarded against threshold; the percentage of discarded years steadily increases as threshold increases while the term  $b_1(u; t) \ll \varepsilon_\circ$  (Figure D.1.2 top row) and, as  $b_1(u; t)$  approaches  $\varepsilon_\circ$ , the percentage reaches a peak followed by a steep drop to 0.

We have seen that there are two main effects which can lead to a higher percentage of discarded years for the smaller subset of years: the bound  $b_1(u; t)$  restricting the possible thresholds and the bound  $b_2(u; t)$  being larger for the larger set of years. In our example with  $t = 5$  the bound  $b_1(u; t)$  has a substantial effect on the possible discard set. For  $n_y = 1000$  the threshold cannot exceed  $6.5 \times 10^5$  resulting in at most 23.8% of years discarded whereas for  $\tilde{n}_y = 500$  we can discard a higher percentage of years (27%) since a much higher threshold ( $6.9 \times 10^5$ ) is possible. For larger return periods the bound  $\varepsilon_\circ - b_1(u; t)$  is more similar for both sets of years (so on the plots the dashed black lines would be around the same place) so then the difference in maximum percentage of discards is due to  $b_2^{n_y}(u; t) > b_2^{\tilde{n}_y}(u; t)$ .

For some more insight, we revisit Figure D.1.1. It is clear from the curves that  $b_2(u; t)$  is monotonically decreasing in  $u$  when  $\mathcal{D}$  is fixed and that, given  $u$ ,  $b_2(u; t)$  is larger when  $|\mathcal{D}|$  is larger, hence the coloured curves have no overlap. We can see how the bound  $\varepsilon_\circ - b_1(u; t)$  restricts the number of discarded years; the curves for larger discard sets do not cross the

bound and when the threshold is too large (*e.g.*, greater than  $\sim 637000$  on the right plot) the curve corresponding to the maximum discard set is larger than the bound. Note that for  $n_y = 1000$  any threshold in the range  $600000 - 637551$  with  $|\mathcal{D}| = 238$  would be acceptable, the threshold  $617346$  is chosen to obtain the smallest possible  $b_1(u; 5) + b_2(u; 5)$  bound with this discard set  $\mathcal{D}$ . We also see on these plots that  $b_2^{1000}(u; 5) \approx 2b_2^{500}(u; 5)$  (for a particular curve on the left plot the curve with the same colour on the right plot is doubly large in the  $y$ -direction).

In conclusion, if  $T_{\tilde{n}_y}^t \subset T_{n_y}^t$  or the subset  $\mathcal{Y}_{\tilde{n}_y} \subset \mathcal{Y}$  is a ‘good representation’ of the years in  $\mathcal{Y}$  then the maximum possible threshold is larger for the subset  $\mathcal{Y}_{\tilde{n}_y}$  and the maximum percentage of discarded years is larger for the smaller set of years. This behaviour is stronger the smaller the return period  $t$  is. On the other hand if  $T_{\tilde{n}_y}^t \not\subset T_{n_y}^t$  and the subset is ‘skewed’ towards years with higher or lower losses then it is not clear whether the percentage of years discarded will be larger for the subset or the full set. More investigation would be needed to evaluate the impact of an increased number of years in the event set on the performance of the exclude method.

## D.2 Exclude procedure results

Return period	No. of kept years								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	1000	1000	1000	1000	1000	1000	1000	1000	703
5	1000	1000	977	977	874	941	1000	1000	310
10	1000	1000	929	762	486	474	922	1000	163
20	1000	1000	778	340	296	280	730	1000	89
50	1000	1000	451	110	98	97	370	907	26
75	1000	1000	367	66	62	57	267	535	20
100	1000	1000	327	44	39	38	226	469	15
150	1000	1000	195	24	24	24	122	293	9
200	1000	1000	175	24	24	24	116	224	9
250	1000	1000	168	21	20	20	100	208	7
500	1000	1000	168	20	19	19	92	170	7

Table D.2.1: Number of years kept out of 1000 years when using the exclude method for each concentration inequality with  $\varepsilon_\circ = 0.0001$ .

Return period	No. of kept years								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	500	500	500	500	500	500	500	500	355
5	500	500	487	489	363	357	488	500	148
10	500	500	415	239	201	192	403	500	73
20	500	500	329	124	110	107	294	500	45
50	500	500	186	40	38	37	140	267	14
75	500	500	133	20	20	20	93	151	10
100	500	500	130	17	16	15	82	151	10
150	500	500	72	10	10	10	39	82	4
200	500	500	63	8	7	7	33	52	4
250	500	500	63	8	7	7	33	52	4

Table D.2.2: Number of years kept out of 500 years when using the exclude method for each concentration inequality with  $\varepsilon_o = 0.001$ .

Return period	No. of kept years								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	500	500	500	500	500	500	500	500	344
5	500	500	489	489	421	414	489	500	156
10	500	500	456	347	238	234	454	500	82
20	500	500	380	164	142	128	353	500	46
50	500	500	219	52	49	47	178	393	14
75	500	500	166	29	25	24	121	224	12
100	500	500	159	20	20	20	109	224	8
150	500	500	98	12	12	12	55	140	4
200	500	500	81	10	10	10	48	84	4
250	500	500	81	10	10	10	48	84	4

Table D.2.3: Number of years kept out of 500 years when using the exclude method for each concentration inequality with  $\varepsilon_o = 0.0001$ .

Return period	Percentage of simulations performed								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	95.6
5	100.0	100.0	100.0	100.0	99.2	99.9	100.0	100.0	70.2
10	100.0	100.0	99.8	97.2	85.0	84.2	99.8	100.0	50.6
20	100.0	100.0	97.6	73.3	68.7	67.2	96.4	100.0	35.1
50	100.0	100.0	82.5	40.2	37.2	37.0	76.0	99.6	15.6
75	100.0	100.0	75.8	28.6	27.7	26.0	65.5	87.8	13.0
100	100.0	100.0	72.0	21.8	20.3	19.9	60.3	83.8	10.5
150	100.0	100.0	55.9	14.9	14.9	14.9	42.9	68.5	7.3
200	100.0	100.0	52.7	14.9	14.9	14.9	41.4	60.0	7.3
250	100.0	100.0	51.6	13.4	13.0	13.0	37.8	57.7	6.1
500	100.0	100.0	51.6	13.0	12.4	12.4	35.6	52.1	6.1

Table D.2.4: Simulated subrisk and event combinations for the exclude procedure as a percentage of the standard procedure time with 1000 years for each concentration inequality with  $\varepsilon_o = 0.0001$ .

Return period	Percentage of simulations performed								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	93.2
5	100.0	100.0	99.8	99.8	98.0	99.3	100.0	100.0	66.3
10	100.0	100.0	99.1	95.5	81.3	80.5	99.0	100.0	46.8
20	100.0	100.0	96.0	69.5	64.8	63.1	94.3	100.0	32.1
50	100.0	100.0	78.8	37.2	34.4	34.2	72.2	98.6	14.9
75	100.0	100.0	71.9	26.3	25.5	23.9	61.5	84.4	12.3
100	100.0	100.0	68.2	20.6	19.4	19.1	56.3	80.2	10.3
150	100.0	100.0	51.9	14.2	14.2	14.2	39.6	64.5	7.6
200	100.0	100.0	48.8	14.2	14.2	14.2	38.3	56.0	7.6
250	100.0	100.0	47.6	12.8	12.3	12.3	34.9	53.8	6.4
500	100.0	100.0	47.6	12.3	11.9	11.9	32.8	48.0	6.4

Table D.2.5: Simulated subrisk and event combinations for the exclude procedure as a percentage of the standard procedure time with 500 years for each concentration inequality with  $\varepsilon_o = 0.0001$ .

Return period	Percentage of simulations performed								
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++	BC-MC
2	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	95.7
5	100.0	100.0	99.9	99.9	96.1	95.8	99.9	100.0	71.4
10	100.0	100.0	98.2	85.4	80.4	79.0	97.8	100.0	50.1
20	100.0	100.0	93.7	65.2	61.9	61.0	90.8	100.0	38.1
50	100.0	100.0	78.1	35.4	34.3	33.9	69.7	88.5	18.4
75	100.0	100.0	67.6	22.4	22.4	22.4	57.4	72.0	5.1
100	100.0	100.0	67.0	20.6	20.1	19.2	53.5	72.0	15.1
150	100.0	100.0	49.9	15.1	15.1	15.1	34.7	53.5	7.8
200	100.0	100.0	46.2	12.9	11.5	11.5	31.5	41.6	7.8
250	100.0	100.0	46.2	12.9	11.5	11.5	31.5	41.6	7.8

Table D.2.6: Simulated subrisk and event combinations for the exclude procedure as a percentage of the standard procedure time with 500 years for each concentration inequality with  $\varepsilon_o = 0.001$ .

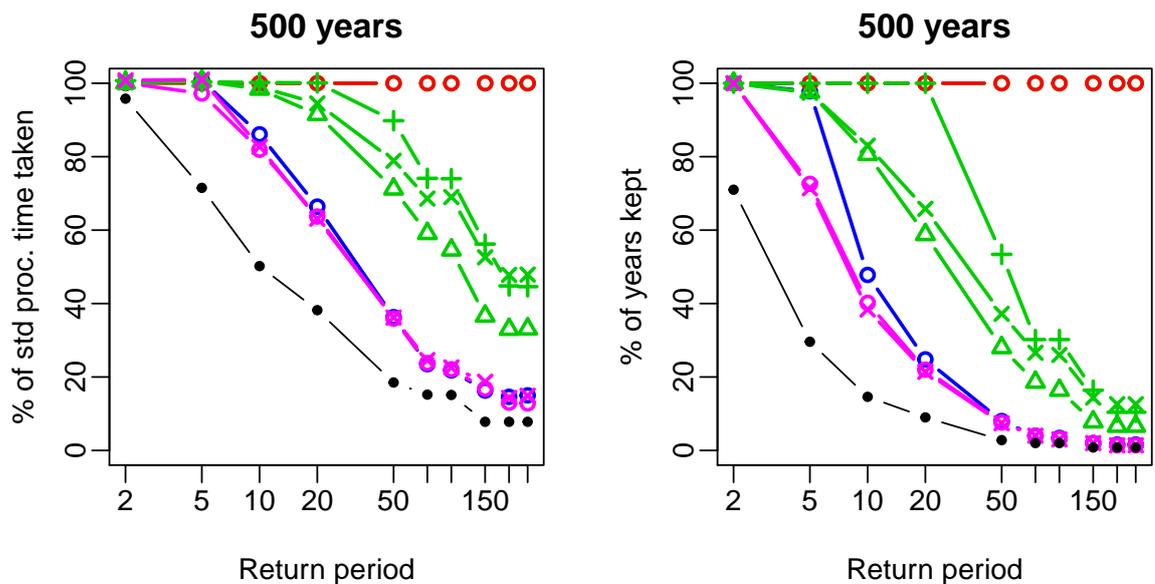


Figure D.2.1: Comparing performance of our procedure with various concentration inequalities against the standard procedure with a 500 year event set and  $\varepsilon_o = 0.001$ . Left:  $T$ , all as a percentage of  $T^{std}$  and right: % of years kept. Usual color scheme for concentration inequalities (Red: Cantelli, Green + : C-H, Green  $\Delta$  : C-H+, Green  $\times$  : C-H++, Blue: Bernstein, Pink o : Bennett, Pink  $\times$  : Bennett+). The black filled circles is the lower bound BC-MC.

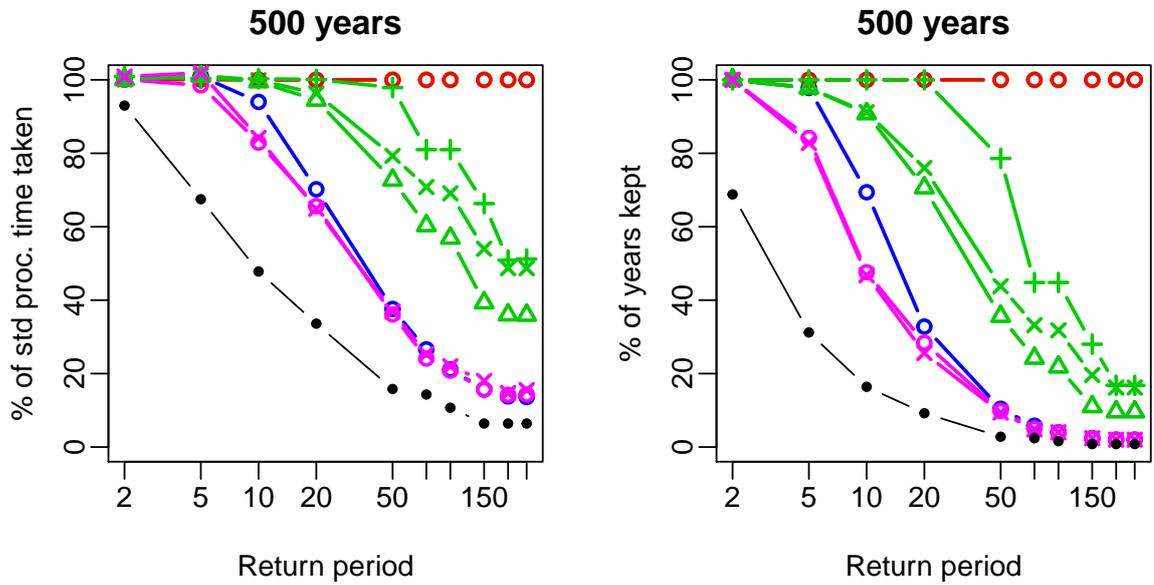


Figure D.2.2: Comparing performance of our procedure with various concentration inequalities against the standard procedure with a 500 year event set and  $\varepsilon_o = 0.0001$ . Left:  $T$ , all as a percentage of  $T^{std}$  and right: % of years kept. Usual color scheme for concentration inequalities (Red: Cantelli, Green + : C-H, Green  $\Delta$  : C-H+, Green  $\times$  : C-H++, Blue: Bernstein, Pink  $\circ$  : Bennett, Pink  $\times$  : Bennett+). The black filled circles is the lower bound BC-MC.

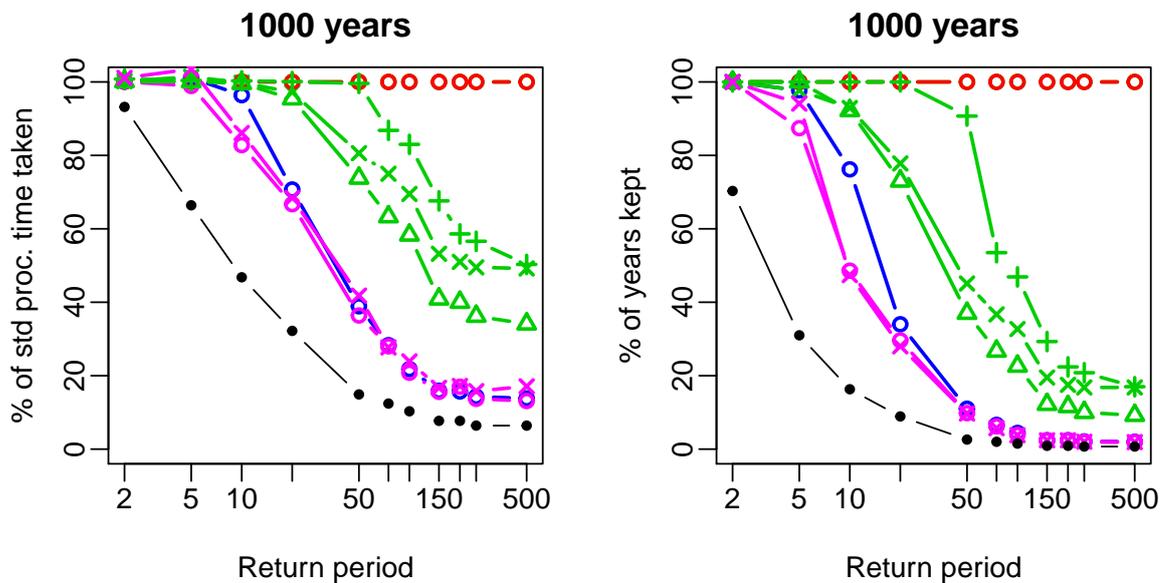


Figure D.2.3: Comparing performance of our procedure with various concentration inequalities against the standard procedure with a 1000 year event set and  $\varepsilon_o = 0.0001$ . Left:  $T$ , all as a percentage of  $T^{std}$  and right: % of years kept. Usual color scheme for concentration inequalities (Red: Cantelli, Green + : C-H, Green  $\Delta$  : C-H+, Green  $\times$  : C-H++, Blue: Bernstein, Pink  $\circ$  : Bennett, Pink  $\times$  : Bennett+). The black filled circles is the lower bound BC-MC.

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	0
5	0	0	0.03	0.05	0.2	0.2	0.04	0
10	0	0	$9 \cdot 10^{-5}$	$1 \cdot 10^{-4}$	$9 \cdot 10^{-5}$	$8 \cdot 10^{-5}$	$9 \cdot 10^{-5}$	0
20	0	0	$4 \cdot 10^{-7}$	$2 \cdot 10^{-7}$	$2 \cdot 10^{-8}$	$9 \cdot 10^{-9}$	$3 \cdot 10^{-7}$	0
50	0	0	$2 \cdot 10^{-13}$	$1 \cdot 10^{-13}$	$1 \cdot 10^{-19}$	$4 \cdot 10^{-21}$	$3 \cdot 10^{-13}$	0
75	0	0	$2 \cdot 10^{-15}$	$2 \cdot 10^{-16}$	$1 \cdot 10^{-29}$	$2 \cdot 10^{-31}$	$4 \cdot 10^{-16}$	0
100	0	0	$7 \cdot 10^{-21}$	$5 \cdot 10^{-19}$	$2 \cdot 10^{-35}$	$2 \cdot 10^{-37}$	$8 \cdot 10^{-23}$	0
150	0	0	$2 \cdot 10^{-42}$	$5 \cdot 10^{-27}$	$2 \cdot 10^{-56}$	$9 \cdot 10^{-59}$	$1 \cdot 10^{-42}$	0
200	0	0	$1 \cdot 10^{-15}$	$3 \cdot 10^{-29}$	$2 \cdot 10^{-62}$	$9 \cdot 10^{-65}$	$2 \cdot 10^{-50}$	0
250	0	0	$2 \cdot 10^{-48}$	$3 \cdot 10^{-29}$	$2 \cdot 10^{-62}$	$9 \cdot 10^{-65}$	$2 \cdot 10^{-50}$	0
500	0	0	$2 \cdot 10^{-125}$	$2 \cdot 10^{-47}$	$5 \cdot 10^{-116}$	$3 \cdot 10^{-119}$	$7 \cdot 10^{-114}$	0
No. yrs kept	1000	1000	929	762	486	474	922	1000
% simulations	100.0	100.0	99.8	97.2	85.0	84.2	99.8	100.0

Table D.2.7: Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 10-year return level non-equivalence probability to be less than  $\varepsilon = 0.0001$  (3rd row of Table D.2.1).

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	0
5	0	0	0.6	1	1	1	0.7	0
10	0	0	0.007	0.03	0.06	0.07	0.008	0
20	0	0	$1 \cdot 10^{-4}$	$1 \cdot 10^{-4}$	$1 \cdot 10^{-4}$	$9 \cdot 10^{-5}$	$9 \cdot 10^{-5}$	0
50	0	0	$6 \cdot 10^{-10}$	$9 \cdot 10^{-11}$	$4 \cdot 10^{-12}$	$1 \cdot 10^{-12}$	$1 \cdot 10^{-9}$	0
75	0	0	$4 \cdot 10^{-12}$	$2 \cdot 10^{-13}$	$1 \cdot 10^{-15}$	$1 \cdot 10^{-16}$	$8 \cdot 10^{-12}$	0
100	0	0	$2 \cdot 10^{-13}$	$6 \cdot 10^{-15}$	$1 \cdot 10^{-18}$	$4 \cdot 10^{-21}$	$2 \cdot 10^{-13}$	0
150	0	0	$2 \cdot 10^{-20}$	$9 \cdot 10^{-22}$	$6 \cdot 10^{-33}$	$3 \cdot 10^{-37}$	$1 \cdot 10^{-17}$	0
200	0	0	$1 \cdot 10^{-15}$	$6 \cdot 10^{-24}$	$3 \cdot 10^{-37}$	$6 \cdot 10^{-42}$	$2 \cdot 10^{-21}$	0
250	0	0	$1 \cdot 10^{-23}$	$6 \cdot 10^{-24}$	$3 \cdot 10^{-37}$	$6 \cdot 10^{-42}$	$2 \cdot 10^{-21}$	0
500	0	0	$6 \cdot 10^{-68}$	$4 \cdot 10^{-42}$	$2 \cdot 10^{-76}$	$3 \cdot 10^{-85}$	$4 \cdot 10^{-54}$	0
No. yrs kept	1000	1000	778	340	296	280	730	1000
% simulations	100.0	100.0	97.6	73.3	68.7	67.2	96.4	100.0

Table D.2.8: Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 20-year return level non-equivalence probability to be less than  $\varepsilon = 0.0001$  (4th row of Table D.2.1).

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	1
5	0	0	1	1	1	1	1	1
10	0	0	1	1	1	1	1	0.3
20	0	0	0.2	1	1	1	0.1	0.02
50	0	0	$6 \cdot 10^{-5}$	$3 \cdot 10^{-5}$	$3 \cdot 10^{-5}$	$2 \cdot 10^{-5}$	$6 \cdot 10^{-5}$	$1 \cdot 10^{-4}$
75	0	0	$2 \cdot 10^{-6}$	$2 \cdot 10^{-7}$	$1 \cdot 10^{-7}$	$5 \cdot 10^{-8}$	$2 \cdot 10^{-6}$	$9 \cdot 10^{-7}$
100	0	0	$2 \cdot 10^{-7}$	$5 \cdot 10^{-9}$	$2 \cdot 10^{-9}$	$6 \cdot 10^{-10}$	$1 \cdot 10^{-7}$	$3 \cdot 10^{-7}$
150	0	0	$1 \cdot 10^{-10}$	$5 \cdot 10^{-13}$	$2 \cdot 10^{-14}$	$3 \cdot 10^{-15}$	$3 \cdot 10^{-10}$	$4 \cdot 10^{-9}$
200	0	0	$4 \cdot 10^{-11}$	$4 \cdot 10^{-14}$	$1 \cdot 10^{-15}$	$1 \cdot 10^{-16}$	$6 \cdot 10^{-11}$	$3 \cdot 10^{-9}$
250	0	0	$4 \cdot 10^{-11}$	$3 \cdot 10^{-14}$	$8 \cdot 10^{-16}$	$1 \cdot 10^{-16}$	$5 \cdot 10^{-11}$	$3 \cdot 10^{-9}$
500	0	0	$7 \cdot 10^{-25}$	$8 \cdot 10^{-30}$	$3 \cdot 10^{-43}$	$5 \cdot 10^{-47}$	$1 \cdot 10^{-18}$	$2 \cdot 10^{-13}$
No. yrs kept	1000	1000	451	110	98	97	370	907
% simulations	100.0	100.0	82.5	40.2	37.2	37.0	76.0	99.6

Table D.2.9: Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 50-year return level non-equivalence probability to be less than  $\varepsilon = 0.0001$  (5th row of Table D.2.1).

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	0
5	0	0	0.2	0.4	1	1	0.2	0
10	0	0	0.0009	0.0009	0.0008	0.0008	0.0009	0
20	0	0	$8 \cdot 10^{-6}$	$2 \cdot 10^{-6}$	$3 \cdot 10^{-7}$	$2 \cdot 10^{-7}$	$5 \cdot 10^{-6}$	0
50	0	0	$1 \cdot 10^{-11}$	$1 \cdot 10^{-12}$	$5 \cdot 10^{-16}$	$6 \cdot 10^{-17}$	$2 \cdot 10^{-11}$	0
75	0	0	$6 \cdot 10^{-14}$	$2 \cdot 10^{-15}$	$1 \cdot 10^{-24}$	$4 \cdot 10^{-26}$	$8 \cdot 10^{-14}$	0
100	0	0	$2 \cdot 10^{-15}$	$4 \cdot 10^{-17}$	$6 \cdot 10^{-30}$	$1 \cdot 10^{-31}$	$2 \cdot 10^{-15}$	0
150	0	0	$6 \cdot 10^{-29}$	$4 \cdot 10^{-25}$	$8 \cdot 10^{-49}$	$3 \cdot 10^{-51}$	$9 \cdot 10^{-27}$	0
200	0	0	$1 \cdot 10^{-15}$	$2 \cdot 10^{-27}$	$2 \cdot 10^{-54}$	$7 \cdot 10^{-57}$	$7 \cdot 10^{-32}$	0
250	0	0	$3 \cdot 10^{-33}$	$2 \cdot 10^{-27}$	$2 \cdot 10^{-54}$	$7 \cdot 10^{-57}$	$7 \cdot 10^{-32}$	0
500	0	0	$1 \cdot 10^{-90}$	$2 \cdot 10^{-45}$	$1 \cdot 10^{-103}$	$1 \cdot 10^{-107}$	$3 \cdot 10^{-74}$	0
No. yrs kept	1000	1000	862	517	415	398	835	1000
% simulations	100.0	100.0	99.0	86.9	79.8	78.5	98.7	100.0

Table D.2.10: *Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 10-year return level non-equivalence probability to be less than  $\varepsilon_0 = 0.001$  (3rd row of Table 6.1.1).*

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	0
5	0	0	1	1	1	1	1	0
10	0	0	0.04	0.2	0.3	0.4	0.05	0
20	0	0	0.001	0.0009	0.0009	0.0009	0.0009	0
50	0	0	$2 \cdot 10^{-8}$	$1 \cdot 10^{-9}$	$8 \cdot 10^{-11}$	$2 \cdot 10^{-11}$	$4 \cdot 10^{-8}$	0
75	0	0	$2 \cdot 10^{-10}$	$4 \cdot 10^{-12}$	$3 \cdot 10^{-14}$	$5 \cdot 10^{-15}$	$4 \cdot 10^{-10}$	0
100	0	0	$1 \cdot 10^{-11}$	$1 \cdot 10^{-13}$	$4 \cdot 10^{-16}$	$7 \cdot 10^{-18}$	$1 \cdot 10^{-11}$	0
150	0	0	$1 \cdot 10^{-15}$	$1 \cdot 10^{-19}$	$1 \cdot 10^{-29}$	$3 \cdot 10^{-33}$	$8 \cdot 10^{-15}$	0
200	0	0	$1 \cdot 10^{-15}$	$9 \cdot 10^{-22}$	$1 \cdot 10^{-33}$	$1 \cdot 10^{-37}$	$9 \cdot 10^{-16}$	0
250	0	0	$1 \cdot 10^{-16}$	$9 \cdot 10^{-22}$	$1 \cdot 10^{-33}$	$1 \cdot 10^{-37}$	$8 \cdot 10^{-16}$	0
500	0	0	$1 \cdot 10^{-51}$	$9 \cdot 10^{-40}$	$1 \cdot 10^{-71}$	$7 \cdot 10^{-79}$	$4 \cdot 10^{-40}$	0
No. yrs kept	1000	1000	690	274	242	228	618	1000
% simulations	100.0	100.0	95.1	66.4	62.4	60.5	91.9	100.0

Table D.2.11: *Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 20-year return level non-equivalence probability to be less than  $\varepsilon = 0.001$  (4th row of Table 6.1.1).*

Return period	Upper bound on non-equivalence probability							
	Cantelli	Hoeffding	C-H	Bernstein	Bennett	Benn+	C-H+	C-H++
2	0	0	1	1	1	1	1	1
5	0	0	1	1	1	1	1	1
10	0	0	1	1	1	1	1	0.9
20	0	0	0.6	1	1	1	0.5	0.1
50	0	0	0.0007	0.0004	0.0004	0.0004	0.0007	0.001
75	0	0	$3 \cdot 10^{-5}$	$3 \cdot 10^{-6}$	$2 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	$3 \cdot 10^{-5}$	$2 \cdot 10^{-5}$
100	0	0	$3 \cdot 10^{-6}$	$8 \cdot 10^{-8}$	$3 \cdot 10^{-8}$	$2 \cdot 10^{-8}$	$2 \cdot 10^{-6}$	$5 \cdot 10^{-6}$
150	0	0	$4 \cdot 10^{-9}$	$9 \cdot 10^{-12}$	$7 \cdot 10^{-13}$	$2 \cdot 10^{-13}$	$8 \cdot 10^{-9}$	$1 \cdot 10^{-7}$
200	0	0	$2 \cdot 10^{-9}$	$8 \cdot 10^{-13}$	$3 \cdot 10^{-14}$	$1 \cdot 10^{-14}$	$2 \cdot 10^{-9}$	$8 \cdot 10^{-8}$
250	0	0	$1 \cdot 10^{-9}$	$6 \cdot 10^{-13}$	$3 \cdot 10^{-14}$	$8 \cdot 10^{-15}$	$2 \cdot 10^{-9}$	$8 \cdot 10^{-8}$
500	0	0	$3 \cdot 10^{-19}$	$2 \cdot 10^{-27}$	$7 \cdot 10^{-39}$	$6 \cdot 10^{-41}$	$3 \cdot 10^{-15}$	$2 \cdot 10^{-11}$
No. yrs kept	1000	1000	376	81	78	75	299	621
% simulations	100.0	100.0	76.5	32.9	31.9	31.1	69.2	92.1

Table D.2.12: *Non-equivalence probability (rounded to 1 s.f.) when simulating the number of years necessary for the 50-year return level non-equivalence probability to be less than  $\varepsilon = 0.001$  (5th row of Table 6.1.1).*

# Appendix E

## Appendix to Chapter 7

### E.1 Heffernan and Tawn's conditional multivariate extremes model

The Heffernan and Tawn model is based upon an assumption on the asymptotic joint distribution conditioned on an extreme component. Let  $\mathbf{X} = (X_1, \dots, X_d)$  be a random vector and consider a set,  $A$ , where at least on one component of  $\mathbf{X}$  is extreme. This set can be partitioned as

$$A = \bigcup_{j=1}^d A_j,$$

where  $A_j$  is the subset of  $A$  in which  $X_j$  is the largest component of  $\mathbf{X}$  in terms of the quantiles of its marginal distribution. Then, using this decomposition the probability of lying in this extreme set is

$$\mathbb{P}(\mathbf{X} \in A) = \sum_{j=1}^d \mathbb{P}(\mathbf{X} \in A_j, X_j > \nu_j) = \sum_{j=1}^d \mathbb{P}(\mathbf{X} \in A_j | X_j > \nu_j) \mathbb{P}(X_j > \nu_j),$$

where  $\nu_j$  is the smallest  $x_j$  in the set  $A_j$  and so  $\nu_j$  is large. The last term of the sum,  $\mathbb{P}(X_j > \nu_j)$ , can be easily estimated using the generalised Pareto distribution above some high threshold.

The original theory of Heffernan and Tawn (2004) assumes  $\mathbf{X}$  has standard Gumbel margins (which can be achieved using the probability integral transform) so each  $X_j$  has an exponential upper and lower tail. Here we formulate the theory in Laplace margins, as described in Keef et al. (2012), since this makes the modelling of negative dependence cases more parsimonious. Consider the limiting behaviour of  $\mathbb{P}(\mathbf{X}_{\setminus j} \leq \mathbf{x}_{\setminus j} | X_j > \nu_j)$  where  $\mathbf{X}_{\setminus j}$  is the random vector  $\mathbf{X}$  excluding the  $j$ th component. In a similar manner to that described in §2.1 for univariate extremes,  $\mathbf{X}_{\setminus j}$  is normalised to avoid degeneracy in the limit as  $\nu_j \rightarrow \infty$ .

So, we consider the asymptotic distribution of

$$\mathbf{Z}_{|j} = \frac{\mathbf{X}_{|j} - \mathbf{a}_{|j}(X_j)}{\mathbf{b}_{|j}(X_j)},$$

where  $\mathbf{a}_{|j}(X_j)$  and  $\mathbf{b}_{|j}(X_j)$  are chosen vector functions so that  $\lim_{\nu_j \rightarrow \infty} \mathbb{P}(\mathbf{Z}_{|j} \leq \mathbf{z}_{|j} | X_j > \nu_j) = G_{|j}(\mathbf{z}_{|j})$  where  $G_{|j}(\mathbf{z}_{|j})$  are non-degenerate in all margins.

This assumption leads to the following result for all fixed  $y > 0$ :

$$\lim_{\nu_j \rightarrow \infty} \mathbb{P}(\mathbf{Z}_{|j} \leq \mathbf{z}_{|j}, X_j - \nu_j > y | X_j > \nu_j) = G_{|j}(\mathbf{z}_{|j}) \exp(-y). \quad (\text{E.1.1})$$

Therefore,  $X_j - \nu_j$  and the residuals,  $\mathbf{Z}_{|j}$ , are independent conditional on  $X_j > \nu_j$  as  $\nu_j \rightarrow \infty$ . Furthermore, the limit distribution of the residuals is  $G_{|j}(\mathbf{z}_{|j})$  and the limit distribution of  $X_j - \nu_j$  is exponential. This asymptotic independence is an important aspect of the method for inference of the Heffernan and Tawn model.

In Laplace margins  $\mathbf{a}_{|j}(x)$  and  $\mathbf{b}_{|j}(x)$  can be simplified to  $\boldsymbol{\alpha}x$  and  $x^\beta$  respectively with  $\boldsymbol{\alpha} = (\alpha_{i|j}, i = 1, \dots, d, i \neq j)$  and  $\beta = (\beta_{i|j}, i = 1, \dots, d, i \neq j)$  where  $-1 \leq \alpha_{i|j} \leq 1$  and  $\beta_{i|j} \leq 1$  for  $i \neq j$ . For example, for the multivariate extreme value distribution with the logistic model (7.3.5) but here with dependence parameter  $\gamma$ , for all  $0 \leq \gamma \leq 1$ ,  $\boldsymbol{\alpha} = \mathbf{1}$ ,  $\beta = 0$  and

$$G_{|j}(\mathbf{z}_{|j}) = \left[ 1 + \sum_{i \neq j} \exp\left(-\frac{z_{i|j}}{\gamma}\right) \right]^{\gamma-1}.$$

The dependence model Heffernan and Tawn 2004 derive from (E.1.1) is a semi parametric model:

$$\mathbf{X}_{|j} = \mathbf{a}_{|j}(x_j) + \mathbf{b}_{|j}(x_j)\mathbf{Z}_{|j} \quad X_j > \nu_j,$$

where  $\nu_j$  is some high threshold,  $\mathbf{a}_{|j}(x)$  and  $\mathbf{b}_{|j}(x)$  follow a parametric model and the residuals  $\mathbf{Z}_{|j}$  are modelled non-parametrically. The residuals are modelled non-parametrically since we have no theory to specific the margins or the  $d - 1$ -dimensional dependence structure. All four classes of dependence as described in §7.3.4 can be attained with this model. For example,  $(X_i, X_j)$  are asymptotically dependent if they tend to infinity at the same rate - this occurs when  $\alpha_{i|j} = 1$  and  $\beta_{i|j} = 0$  since then  $X_i = X_j + Z_{i|j}$  for  $X_j > \nu_j$ .

### E.1.1 Inference in the bivariate setting

There are two steps to inference: estimating the normalisation functions parametrically and non-parametrically modelling the distribution of  $Z$ . Here we describe the method for the bivariate setting,  $\mathbf{X} = (X, Y)$ , with the simplified normalisation functions,  $a(X) = \alpha X$  and  $b(X) = X^\beta$ .

Due to the independence in (E.1.1),  $X$  and  $Z$  can be simulated independently. First samples for  $X$ , which we denote  $x_{sim}$ , are obtained by adding  $\nu$  to simulations from the standard exponential distribution. Second, if  $\alpha$  and  $\beta$  are known,  $Z$  can be simulated from the empirical distribution of  $Z_{obs} = (Y_{obs} - \alpha X_{obs}) / X_{obs}^\beta$  where  $X_{obs}$  and  $Y_{obs}$  are the observed  $(X, Y)$  pairs with  $X_{obs} > \nu$ . This method leads to simulated  $X_{sim}$  and  $Y_{sim} = \alpha X_{sim} + X_{sim}^\beta Z_{sim}$  values occurring along loci of decreasing density. Since  $Z_{sim}$  are from the empirical distribution, the simulated  $Y_{sim}$  values are simply a translation of an observed point, along a loci determined by  $\alpha$  and  $\beta$ , to the simulated  $X$  value. In the simplest case when  $\beta = 0$  the rays are parallel to the line  $Y = \alpha X$ , whereas when  $\beta > 0$  the rays fan outwards with increasing  $X$ . One way to avoid the unrealistic simulation along rays is to smooth the empirical distribution so samples of  $Z$  will not just correspond to observed points (Towe et al., 2016). This smoothing can be done using a kernel density estimate. Nevertheless this empirical modelling of the residuals suffers from the curse of dimensionality; Towe et al. (2019) propose a model-based copula to replace the non-parametric empirical estimate to reduce this issue in the model.

Following Heffernan and Tawn (2004) inference for the parameters,  $(\alpha, \beta)$ , is based on the working assumption<sup>1</sup> that  $Z$  is a normal random variable with mean  $\mu$  and variance  $\sigma^2$ . Under this assumption we have the following model:

$$Y|X > \nu \sim N(\alpha X + X^\beta \mu, X^{2\beta} \sigma^2),$$

from which we can estimate the parameters,  $\mu$ ,  $\sigma$ ,  $\alpha$  and  $\beta$  using likelihood inference with the observed  $(X, Y)$  values for  $X > \nu$ . There are also constraints on  $\alpha$  and  $\beta$  that must be taken into account, in particular  $-1 \leq \alpha \leq 1$  and  $0 \leq \beta < 1$ . A value of  $\beta$  less than 0 would lead to strange behaviour as the simulated rays would then converge to the line  $Y = \alpha X$ .

---

<sup>1</sup>Note that this assumption is only used for the estimation of  $(\alpha, \beta)$  and does not affect the rest of the modelling procedure

Keef et al. (2012) propose additional joint constraints, restricting the combinations of  $\alpha$  and  $\beta$  possible.

### E.2 ARMAX simulation

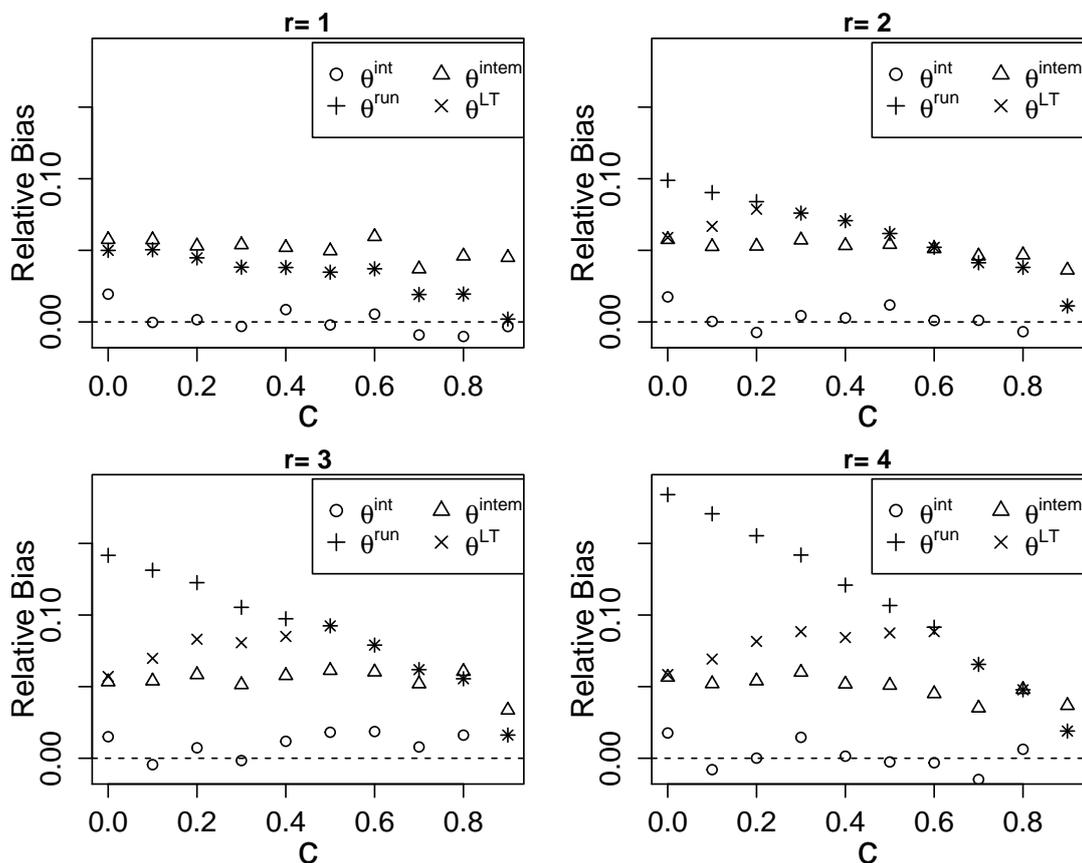


Figure E.2.1: Relative bias of the extremal index estimators: the intervals estimator,  $\hat{\theta}^{int}$ , and the empirical estimators,  $\hat{\theta}^{intem}$ ,  $\hat{\theta}^{run}$  and  $\hat{\theta}^{LT}$ , when applying to simulated data the intervals method of Ferro and Segers (2003), the runs method with  $r = 1, 2, 3, 4$  and Laurini and Tawn's method with threshold equal to the 80% quantile of the data respectively. Data consists of (for each  $r$  and  $c$ ) 100 simulated ARMAX sequences of length 10000 with unit Fréchet common distribution and  $c$  as indicated on the x-axis with threshold equal to the 95% quantile of the data.

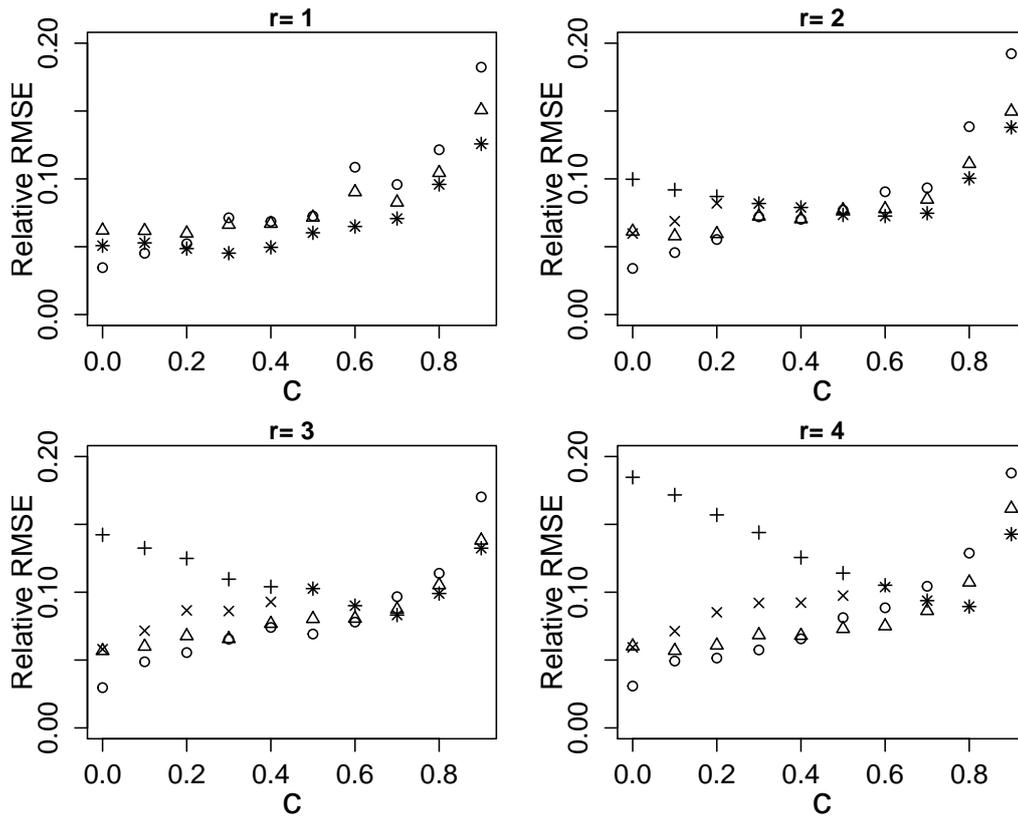


Figure E.2.2: RMSE of the extremal index estimators: the intervals estimator,  $\hat{\theta}^{int}$ ; and the empirical estimators,  $\hat{\theta}^{intem}$ ,  $\hat{\theta}^{run}$  and  $\hat{\theta}^{LT}$ , when applying to simulated data the intervals method of Ferro and Segers (2003), the runs method with  $r = 1, 2, 3, 4$  and Laurini and Tawn's method with threshold equal to the 80% quantile of the data respectively. Data consists of (for each  $r$  and  $c$ ) 100 simulated ARMAX sequences of length 10000 with unit Fréchet common distribution and  $c$  as indicated on the x-axis with threshold equal to the 95% quantile of the data.

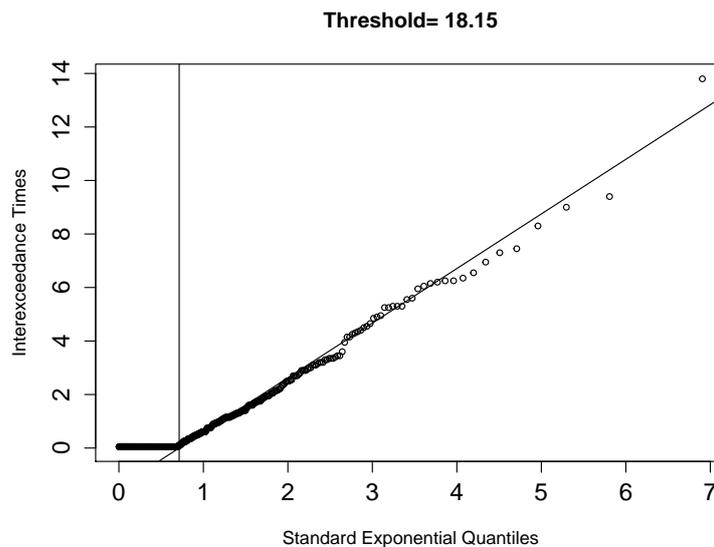


Figure E.2.3: Quantile-quantile plot of interexceedance times against standard exponential quantiles when applying the intervals method of Ferro and Segers (2003) to a simulated ARMAX sequence with unit Fréchet common distribution and  $c = 0.5$  with threshold equal to the 95% quantile of the data. The vertical line corresponds to the  $(1 - \hat{\theta})$  quantile, where  $\hat{\theta}$  is the intervals estimator, and the diagonal line has gradient  $1/\hat{\theta}$ .

# Bibliography

- Alpuim, M. T. (1989). An Extremal Markovian Sequence. *Journal of Applied Probability*, 26(2):219–2.
- Andrews, D. W. . K. . (1999). Estimation When a Parameter is on a Boundary. *Econometrica*, 67(6):1341–1383.
- Asadi, P., Davison, A. C., and Engelke, S. (2015). Extremes on river networks. *Annals of Applied Statistics*, 9(4):2023–2050.
- Asmussen, S. (1987). *Applied Probability and Queues*. Wiley, Chichester.
- Asmussen, S. (2018). Conditional Monte Carlo for sums, with applications to insurance and finance. *Annals of Actuarial Science*, 12(2):269–295.
- Asmussen, S., Jensen, J. L., and Rojas-Nandayapa, L. (2016). Exponential Family Techniques for the Lognormal Left Tail. *Scandinavian Journal of Statistics*, 43(3):774–787.
- Avramidis, A. N. and Wilson, J. R. (1998). Correlation-induction techniques for estimating quantiles in simulation experiments. *Operations research*, 46(4):574–591.
- Barlow, A. M., Sherlock, C., and Tawn, J. (2020). Inference for extreme values under threshold-based stopping rules. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 69(4):765–789.
- Barndorff-Nielsen, O. E. and Cox, D. R. (1984). The effect of sampling rules on likelihood statistics. *International Statistical Review*, 52(3):309–326.
- Barry, J. and Coggan, R. (2010). The visual fast count method: Critical examination and development for underwater video sampling. *Aquatic Biology*, 11(2):101–112.
- Bauer, P., Koenig, F., Brannath, W., and Posch, M. (2010). Selection and bias—two hostile brothers. *Statistics in Medicine*, 29(1):1–13.
- Beck, J. L. and Zuev, K. M. (2005). Rare Event Simulation. *Probability in the Engineering and Informational Sciences*, 20(01):45–66.
- Ben Rached, N., Haji-Ali, A.-L., Rubino, G., and Tempone, R. (2021). Efficient Importance Sampling for Large Sums of Independent and Identically Distributed Random Variables. *arXiv e-prints*, page arXiv:2101.09514.
- Bennett, G. (1962). Probability inequalities for the sum of independent random variables.

- Journal of the American Statistical Association*, 57(33-45).
- Bernstein, S. N. (1946). *The Theory of Probabilities*. Gostehizdat Publishing House, Moscow.
- Berry, A. C. (1941). The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136.
- Billingsley, P. (1995). *Probability and Measure*. Wiley Series in Probability and Statistics. Wiley.
- Boucheron, S. (2013). *Concentration inequalities : a nonasymptotic theory of independence*. Oxford University Press, Oxford, England.
- Boucheron, S., Lugosi, G., and Bousquet, O. (2004). Concentration inequalities. In Bousquet, O., von Luxburg, U., and Rätsch, G., editors, *Advanced Lectures on Machine Learning: ML 2003. Lecture Notes in Computer Science*, volume 3176, pages 208–240. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Bowden, J. and Glimm, E. (2008). Unbiased estimation of selected treatment means in two-stage trials. *Biometrical Journal*, 50(4):515–527.
- Carreras, M. and Brannath, W. (2013). Shrinkage estimation in two-stage adaptive designs with midtrial treatment selection. *Statistics in Medicine*, 32(10):1677–1690.
- UK Centre for Ecology and Hydrology (1999). *Flood Estimation Handbook*, volume 3. Wallingford HydroSolutions.
- Cérou, F. and Guyader, A. (2007). Adaptive Multilevel Splitting for Rare Event Analysis. *Stochastic Analysis and Applications*, 25(2):417–443.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer-Verlag London.
- Coles, S., Heffernan, J., and Tawn, J. (1999). Dependence Measures for Extreme Value Analyses. *Extremes*, 2(4):339–365.
- Coles, S. G. and Tawn, J. a. (1991). Modelling Extreme Multivariate Events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 53(2):377–392.
- Cox, D. R. (1952). A note on the sequential estimation of means. *Mathematical Proceedings of the Cambridge Philosophical Society*, 48(03):447.
- Crossley, A., Lamb, R., and Waller, S. (2010). Fast solution of the Shallow Water Equations using GPU technology. <http://www.jflow.co.uk/>. JBA Consulting, South Barn, Broughton Hall, Skipton, North Yorkshire, BD23 3AE.
- Davison, A. and Smith, R. (1990). Models for exceedances over high thresholds (with discus-

- sion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 52:393–442.
- Dieker, A. and Mandjes, M. (2005). On Asymptotically Efficient Simulation of Large Deviation Probabilities. *Advances in Applied Probability*, 37(2):539–552.
- Dong, H. and Nakayama, M. K. (2017). Quantile estimation using conditional Monte Carlo and Latin hypercube sampling. In *2017 Winter Simulation Conference (WSC)*, pages 1986–1997.
- Eastoe, E. F. (2019). Nonstationarity in peaks-over-threshold river flows: A regional random effects model. *Environmetrics*, 30(5):1–18.
- Eastoe, E. F. and Tawn, J. A. (2009). Modelling Non-Stationary Extremes with Application to Surface Level Ozone. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(1):25–45.
- Eastoe, E. F. and Tawn, J. A. (2012). Modelling the distribution of the cluster maxima of exceedances of subasymptotic thresholds. *Biometrika*, 99(1):43–55.
- Efron, B. (1981). Nonparametric Standard Errors and Confidence Intervals. *Canadian Journal of Statistics*, 9:139–172.
- Efron, B. (1987). Better bootstrap confidence intervals. *Journal of the American Statistical Association*, 82(397, Theory and Methods):171–185.
- Efron, B. (1990). More efficient bootstrap computations. *Journal of the American Statistical Association*, 85:79–89.
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Number 57 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, Boca Raton, Florida, USA.
- Environmental Agency (2020). National flood and coastal erosion risk management strategy for england: executive summary. <https://www.gov.uk/government/publications/national-flood-and-coastal-erosion-risk-management-strategy-for-england--2/national-flood-and-coastal-erosion-risk-management-strategy-for-england-executive-summary>.
- Esseen, C.-G. (1942). On the Liapunoff limit of error in the theory of probability. *Arkiv for matematik, astronomi och fysik.*, 28A(2):1–19.
- Esseen, C. G. (1956). A moment inequality with an application to the central limit theorem. *Scandinavian Actuarial Journal*, 1956(2):160–170.

- Ferreira, H. (1994). Multivariate extreme values in T-periodic random sequences under mild oscillation restrictions. *Stochastic Processes and their Applications*, 49(1):111–125.
- Ferreira, M. and Ferreira, H. (2013). Extremes of multivariate ARMAX processes. *Test*, 22(4):606–627.
- Ferro, C. A. T. and Segers, J. (2003). Inference for clusters of extreme values. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):545–556.
- From, S. G. and Swift, A. W. (2013). A refinement of Hoeffding’s inequality. *Journal of Statistical Computation and Simulation*, 83:977–983.
- Glasserman, P., Heidelberger, P., Shahabuddin, P., and Zajic, T. (1999). Multilevel Splitting for Estimating Rare Event Probabilities. *Oper. Res.*, 47(4):585–600.
- Gollini, I. and Rougier, J. (2015). Rapidly bounding the exceedance probabilities of high aggregate losses. *arXiv e-prints*, page arXiv:1507.01853.
- Gumbel, E. J. (1960). Distributions de valeurs extrêmes en plusieurs dimensions. *Publ. Inst. Statist. Paris*, 9:171–173.
- Hall, P. and Tajvidi, N. (2000). Distribution and dependence-function estimation for bivariate extreme-value distributions. *Bernoulli*, 6(5):835–844.
- Hammersley, J. M. (1956). Conditional Monte Carlo. *Journal of the ACM (JACM)*, 3(2):73–76.
- Hannaford, J. and Marsh, T. J. (2008). High-flow and flood trends in a network of undisturbed catchments in the uk. *International Journal of Climatology*, 28(10):1325–1338.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):497–546.
- Hertz, D. (2020). Improved Hoeffding’s Lemma and Hoeffding’s Tail Bounds. *arXiv e-prints*, page arXiv:2012.03535.
- Hill, B. M. . (1975). A Simple General Approach to Inference About the Tail of a Distribution. *The Annals of Statistics*, 3(5):1163–1174.
- Hoeffding, W. (1963). Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Hoorfar, A. and Hassani, M. (2008). Inequalities on the Lambert W Function and Hyperpower Function. *JIPAM*, 9.
- Hsing, T. (1989). Extreme Value Theory for Multivariate Stationary Sequences. *Journal of*

- Multivariate Analysis*, 29(2):274–291.
- Hsing, T., Hüsler, J., and Leadbetter, M. R. (1988). On the exceedance point process for a stationary sequence. *Probability Theory and Related Fields*, 78(1):97–112.
- Hyndman, R. J. and Fan, Y. (1996). Sample Quantiles in Statistical Packages. *American Statistician*, 50(4):361–365.
- Impagliazzo, R. and Kabanets, V. (2010). Constructive proofs of concentration bounds. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 6302 of *Lecture Notes in Computer Science*, pages 617–631. Springer Berlin Heidelberg.
- Jebara, T. (2018). A refinement of Bennett’s inequality with applications to portfolio optimization. *IDEAS Working Paper Series from RePEc*, 1:1–24.
- Joe, H., Smith, R. L., and Weissman, I. (1992). Bivariate Threshold Methods for Extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 54(1):171–183.
- Keef, C., Papastathopoulos, I., and Tawn, J. A. (2012). Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the heffernan and tawn model. *Journal of Multivariate Analysis*, 115:396 – 404.
- Keef, C., Svensson, C., and Tawn, J. A. (2009). Spatial dependence in extreme river flows and precipitation for Great Britain. *Journal of Hydrology*, 378(3-4):240–252.
- Keef, C., Tawn, J. A., and Lamb, R. (2013). Estimating the probability of widespread flood events. *Environmetrics*, 24(1):13–21.
- Kenward, M. G. and Molenberghs, G. (1998). Likelihood based frequentist inference when data are missing at random. *Statistical Science*, 13(3):236–247.
- Kotz, S. and Nadarajah, S. (2000). *Extreme value distributions : theory and applications*. Imperial College Press, River Edge, NJ.
- Kutin, S. (2002). Extensions to McDiarmid’s inequality when differences are bounded with high probability. Technical report, Dept. Comput. Sci., Univ. Chicago, Chicago, IL, Tech., Chicago.
- Laurini, F. and Tawn, J. A. (2003). New estimators for the extremal index and other cluster characteristics. *Extremes*, 6(3):189–211.
- Leadbetter, M. (1983a). *Extremes and related properties of random sequences and processes*. Springer series in statistics. Springer-Verlag, New York.

- Leadbetter, M. R. (1983b). Extremes and local dependence in stationary sequences. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(2):291–306.
- Leadbetter, M. R. (1991). On a basis for 'Peaks over Threshold' modeling. *Statistics and Probability Letters*, 12(4):357–362.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for Near Independence in Multivariate Extreme Values. *Biometrika*, 83(1):169–187.
- Ledford, A. W. and Tawn, J. A. (1997). Modelling dependence within joint tail regions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):475–499.
- Makkonen, L. and Pajari, M. (2014). Defining sample quantiles by the true rank probability. *Journal of Probability and Statistics*, 2014.
- Martins, A. P. and Ferreira, H. (2005). Measuring the extremal dependence. *Statistics and Probability Letters*, 73(2):99–103.
- McLeish, D. L. and Men, Z. (2015). Extreme value importance sampling for rare event risk measurement. In Glau, K., Scherer, M., and Zagst, R., editors, *Innovations in Quantitative Risk Management*, pages 317–335, Cham. Springer International Publishing.
- Molenberghs, G., Kenward, M. G., Aerts, M., Verbeke, G., Tsiatis, A. A., Davidian, M., and Rizopoulos, D. (2014). On random sample size, ignorability, ancillarity, completeness, separability, and degeneracy: Sequential trials, random sample sizes, and missing data. *Statistical Methods in Medical Research*, 23(1):11–41.
- Mulzer, W. (2018). Five proofs of chernoff's bound with applications. *Bulletin of the EATCS (BEATCS)*, Issue 124, February 2018.
- Nadarajah, S. (2001). Multivariate declustering techniques. *Environmetrics*, 12(4):357–365.
- Nakayama, M. K. (2007). Quantile Estimation When Applying Conditional Monte Carlo. In *2014 4th International Conference On Simulation And Modeling Methodologies, Technologies And Applications (SIMULTECH)*, pages 280–285. SCITEPRESS.
- Nandagopalan, S. (1994). On the multivariate extremal index. *Journal of Research of the National Institute of Standards and Technology*, 99(4):543.
- O'Brien, G. L. . (1987). Extreme Values for Stationary and Markov Sequences. *The Annals of Probability*, 15(1):281–291.
- Pawitan, Y. (2013). *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. OUP Oxford.
- Pickands, J. (1981). Multivariate extreme value distributions. *Proc. of the 43rd session of*

- the International Statistical Institute*, 49:859–878.
- Ridder, A. and Rubinstein, R. (2007). Minimum cross-entropy methods for rare-event simulation. *Simulation (San Diego, Calif.)*, 83(11):769–784.
- Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer-Verlag New York, 2nd edition.
- Robert, C. Y. (2008). Estimating the multivariate extremal index function. *Bernoulli*, 14(4):1027–1064.
- Robinson, M. E. and Tawn, J. A. (2000). Extremal Analysis of Processes Sampled at Different Frequencies. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1):117–135.
- Ross, S. M. (1996). *Stochastic Processes*. John Wiley and Sons, 2nd edition.
- Sadowsky, J. (1993). On the optimality and stability of exponential twisting in monte carlo estimation. *IEEE transactions on information theory*, 39(1):119–128.
- Scarrott, C. and Macdonald, A. (2012). A Review of Extreme Value Threshold Estimation and Uncertainty Quantification. *REVSTAT – Statistical Journal*, 10(1):33–60.
- Schlather, M. and Tawn, J. A. (2003). A dependence measure for multivariate and spatial extreme values: Properties and inference. *Biometrika*, 90(1):139–156.
- Shevtsova, I. G. (2010). An improvement of convergence rate estimates in the lyapunov theorem. *Doklady. Mathematics*, 82(3):862–864.
- Siegmund, D. (1976). Importance sampling in the monte carlo study of sequential tests. *The Annals of statistics*, 4(4):673–684.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marge. *Publ. Inst. Statist. Univ. Paris (in French)*, 8:229–231.
- Smith, R. L. and Weissman, I. (1994). Estimating the extremal index. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 56(3):515–528.
- Smith, R. L. and Weissman, I. (1996). Characterization and estimation of the multivariate extremal index. Technical report, Univ. North Carolina.
- Stallard, N. and Todd, S. (2005). Point estimates and confidence regions for sequential trials involving selection. *Journal of Statistical Planning and Inference*, 135(2):402 – 419.
- Swain, R. and Swallow, D. (2015). The prudential regulation of insurers under solvency ii. *European Economics: Macroeconomics and Monetary Economics eJournal*.
- Todd, S., Whitehead, J., and Facey, K. (1996). Point and interval estimation following a

- sequential test. *Biometrika*, 83(2):453–461.
- Towe, R., Tawn, J. A., Lamb, R., Sherlock, C., and Liu, Y. (2016). Improving statistical models for flood risk assessment. *FLOODRisk conference*.
- Towe, R. P., Tawn, J. A., Lamb, R., and Sherlock, C. G. (2019). Model-based inference of conditional extreme value distributions with hydrological applications. *Environmetrics*, 30(8).
- Turkman, K. F., Amaral Turkman, M. A., and Pereira, J. M. (2010). Asymptotic models and inference for extremes of spatio-temporal data. *Extremes*, 13(4):375–397.
- UK Centre for Ecology and Hydrology (2018). National River Flow Archive (NRFA). Data retrieved from the National River Flow Archive, <https://nrfa.ceh.ac.uk/data/station/info/72004>.
- Wadsworth, J. L. (2016). Exploiting structure of maximum likelihood estimators for extreme value threshold selection. *Technometrics*, 58(1):116–126.
- Wald, A. (2004). *Sequential Analysis (reprint)*. Dover Phoenix Editions. Dover Publications.
- Whitehead, J. (1986). On the bias of maximum likelihood estimation following a sequential test. *Biometrika*, 73(3):573–581.
- Zhang, Z. (2002). *Multivariate Extremes, Max-Stable Process Estimation and Dynamic Financial Modeling*. PhD thesis, University of North Carolina.
- Zhang, Z. and Smith, R. L. (2004). The behavior of multivariate maxima of moving maxima processes. *Journal of Applied Probability*, 41(4):1113–1123.
- Zheng, S. (2017). A refined Hoeffding’s upper tail probability bound for sum of independent random variables. *Statistics & Probability Letters*, 131(C):87–92.