

The statistical analysis of former
sea level

Andrew Christopher Parnell

Thesis submitted to the University of Sheffield
for the degree of Doctor of Philosophy

Department of Probability and Statistics

University of Sheffield

November 2005

Acknowledgements

First and foremost, I would like to thank my supervisor Professor Clive Anderson for his expert guidance. His criticism was always kind and fair; his praise always encouraging. I would like to thank my geography supervisor, Professor Ben Horton. He always took an interest in my studies, and kept this thesis practical and relevant. I also owe a debt of gratitude to my archaeological supervisor (soon to be doctor) Alex Bayliss. Her boundless enthusiasm and willingness for me to get my hands dirty during these three years has improved the quality of this tome. I would like to thank Professor Robert van der Noort of Exeter University and Jane Sidell of University College London for their help and guidance in using the archaeological data from the Humber Wetlands Project.

A number of members of staff and students during my three years are also deserving of my gratitude. I would like to thank Dr Caitlin Buck for explaining the finer points of radiocarbon dating to me. I would also like to thank my two office-mates during my time at Sheffield, John Fry and Sammy Rashid, for providing me with excellent conversation. Special mention should go to two PhD students at the University of Durham: Matt Brain and Sarah Woodroffe. These two taught me the basics of coring and the enjoyment of fieldwork. They also provided me with excellent company on two trips to the USA. Finally, I would like to thank all the other PhD students, staff, research assistants and secretaries who all have made my time at Sheffield so enjoyable.

Acknowledgements

First and foremost, I would like to thank my supervisor Professor Clive Anderson for his expert guidance. His criticism was always kind and fair; his praise always encouraging. I would like to thank my geography supervisor, Professor Ben Horton. He always took an interest in my studies, and kept this thesis practical and relevant. I also owe a debt of gratitude to my archaeological supervisor (soon to be doctor) Alex Bayliss. Her boundless enthusiasm and willingness for me to get my hands dirty during these three years has improved the quality of this tome. I would like to thank Professor Robert van der Noort of Exeter University and Jane Sidell of University College London for their help and guidance in using the archaeological data from the Humber Wetlands Project.

A number of members of staff and students during my three years are also deserving of my gratitude. I would like to thank Dr Caitlin Buck for explaining the finer points of radiocarbon dating to me. I would also like to thank my two office-mates during my time at Sheffield, John Fry and Sammy Rashid, for providing me with excellent conversation. Special mention should go to two PhD students at the University of Durham: Matt Brain and Sarah Woodroffe. These two taught me the basics of coring and the enjoyment of fieldwork. They also provided me with excellent company on two trips to the USA. Finally, I would like to thank all the other PhD students, staff, research assistants and secretaries who all have made my time at Sheffield so enjoyable.

Thesis summary

This thesis provides the first template for estimating relative sea level curves and their associated uncertainties. More specifically, the thesis estimates the changing state of sea level in the Humber estuary, UK, over the course of the Holocene. These estimates are obtained through Bayesian methods involving Gaussian processes.

Part of the task involves collating data sources from both archaeologists and geologists which have been collected during frequent study of the region. A portion of the thesis is devoted to studying the nature of the data, and the adjustment of the archaeological information so it can be used in a format suitable for estimating former sea level.

The Gaussian processes are used to model sea-level change via a correlation function which assumes that data points close together in time and space should be at a similar elevation. This assumption is relaxed by incorporating non-stationary correlation functions and aspects of anisotropy. A sequence of models are fitted using Markov chain Monte Carlo. The resultant curves do not pre-suppose a functional form, and give a comprehensive framework for accounting for their uncertainty.

A further complication is introduced as the temporal explanatory variables are stochastic: they arise as radiocarbon dates which require statistical calibration. The resulting posterior date densities are irregular and multi-modal. The spatio-temporal Gaussian process

model takes account of such irregularities via Monte Carlo simulation.

The resultant sea-level curves are scrutinised at a number of locations around the Humber over a selection of time periods. It is hoped that they can provide insight into other areas of sea-level research, and into a broader palaeoclimate framework.

Contents

1 Introduction	11
1.1 Introduction	11
1.2 Why study sea level?	12
1.3 Thesis overview: a statistician's perspective	12
1.4 Thesis overview: a geologist's perspective	14
1.5 Thesis outline	15
1.5.1 Chapter 2	15
1.5.2 Chapter 3	15
1.5.3 Chapter 4	16
1.5.4 Chapter 5	16
1.5.5 Chapter 6	16
1.5.6 Chapter 7	17
2 An introduction to the reconstruction of former sea level (with a focus on the UK and the Humber Estuary)	18
2.1 Introduction	18

2.2	Sea level	19
2.2.1	What is sea level?	19
2.2.2	Eustatic changes	19
2.2.3	Isostatic changes	20
2.2.4	Tectonic change	21
2.2.5	Local scale factors	23
2.3	Reconstructing former sea level	27
2.3.1	Sea-level index points	27
2.3.2	Creating SLIs	27
2.3.3	Indicative meaning	31
2.3.4	Basal index points	33
2.3.5	Intercalated index points	34
2.3.6	Limiting points	34
2.3.7	Tendency	34
2.3.8	Radiocarbon Dating	35
2.3.9	Geophysical modelling	38
2.3.10	Tidal range modelling	41
2.4	Uncertainties in sea-level research	43
2.4.1	Radiocarbon dating and calibration	43
2.4.2	Estimation of index point elevation	45
2.5	Summary	48
3	The Data	49

3.1	Introduction	49
3.2	The Humber estuary	50
3.3	The LOIS database	56
3.4	Archaeological data	58
3.5	The calculation of reference water levels for archaeological data points	60
3.6	Outline of archaeological data	63
3.6.1	Wetland Heritage of the Ancholme and Lower Trent valleys	64
3.6.2	Wetland Heritage of the Vale of York	73
3.6.3	Wetland Heritage of the Hull valley	75
3.6.4	Wetland heritage of the Lincolnshire Marsh	83
3.7	Summary	88
4	Preliminary Data Analysis	90
4.1	Introduction	90
4.2	Age data	91
4.2.1	Radiocarbon dating	91
4.2.2	Period data	98
4.3	Graphical methods	100
4.3.1	Individual index points	100
4.3.2	Plotting data	101
4.4	Cartographic representation	104

4.5	Temporal analysis of sea-level curves	105
4.5.1	Kernel smoothing	106
4.5.2	Endpoint estimation	107
4.5.3	Quantile Regression with check functions	110
4.6	Variograms	112
4.6.1	Temporal variograms	113
4.6.2	Spatial variograms	114
4.7	Tendency data	118
4.8	Summary	119
5	Gaussian processes	123
5.1	Introduction	123
5.2	Gaussian processes for smoothing	124
5.3	Choice of mean	128
5.4	Choice of autocorrelation function	129
5.4.1	Gaussian	130
5.4.2	Powered exponential	130
5.4.3	Matérn	131
5.4.4	Non-stationary correlation	131
5.5	Limiting points and intercalated index points	133
5.6	Gaussian process derivatives	136
5.6.1	Utilising tendency information	139
5.7	Spatial aspects	141

5.8	Predictive distributions	145
5.9	Dating uncertainties	146
5.10	Markov chain Monte Carlo	147
5.10.1	The Metropolis-Hastings Algorithm	148
5.10.2	The Metropolis Algorithm	149
5.10.3	Convergence Diagnostics	150
5.11	Example	152
5.11.1	Introduction	152
5.11.2	Basal index points only	153
5.11.3	Limiting points and intercalated index points	154
5.11.4	Derivative information	157
5.11.5	Non-stationary covariance	160
5.12	Summary	161
6	Modelling RSL	163
6.1	Introduction	163
6.2	Temporal models	164
6.2.1	GeoBUGS	165
6.2.2	Stationary covariance	167
6.2.3	Stationary covariance with tendency information	176
6.2.4	Non-stationary covariance	181
6.2.5	Model comparison	184
6.3	Spatio-temporal models	189

6.3.1	GeoBUGS	189
6.3.2	Spatio-Temporal separable stationary covariance . .	193
6.3.3	Non-stationary spatio-temporal covariance	197
6.3.4	Model comparison	199
6.4	Summary	201
7	Application of results	202
7.1	Introduction	202
7.2	Examining RSL uncertainty	203
7.2.1	Kernel density estimation	203
7.2.2	Functional principal components analysis	206
7.3	Rates of RSL change	212
7.3.1	Differentiated splines	212
7.3.2	Predictive first-derivative distributions	213
7.4	Comparison with geophysical models	215
7.5	Comparison with published estimated RSL rates for the Humber	218
7.6	Compression of intercalated index points	220
7.7	Estimation of isostatic activity in the Humber	226
7.8	The use of archaeological data	228
7.9	Indications of well-studied periods of climate change	231
7.9.1	The 8.2k event	231
7.9.2	The 4k event	234

7.9.3	The little ice age and mediaeval warm period	236
7.10	Summary	238
8	Conclusions and further research	240
8.1	Conclusions	240
8.2	Further research	242
8.2.1	Guidance for archaeologists and geologists	242
8.2.2	Temporal extrapolation and geophysical models	244
8.2.3	Spatial modelling of tendency data	245
8.2.4	Complete spatio-temporal non-stationarity	246
8.2.5	Applying depth and thickness information to intercalated index point shifts	247
A	LOIS Data	249
B	Archaeological Data	255
C	Estimating HAT for archaeological index and limiting points	257
C.1	Estimation of modern HAT and MHWS at non-standard ports	257
C.2	Modelling the ratio of HAT to MHWS	261
C.3	Using the estimated ratio to predict HAT at other locations	262
D	WinBUGS code	265

Chapter 1

Introduction

1.1 Introduction

The aim of this thesis is to develop and apply a technique to derive a sea-level curve with associated error bands. This must accurately represent and take account of all available data. This aim is achieved through the use of non-stationary Gaussian processes. The ideas presented in this thesis are aimed at those with a statistical background. However, much of the data are based upon work in a number of other disciplines, all of which contribute to the measurement of sea level. These ideas are presented in a specific chapter (chapter 2). The data source available for this thesis covers the Humber estuary on the north-east coast of England. This introductory chapter summarises the challenges involved in achieving the work from both a statistical and geological perspective. It ends with a thesis outline for quick reference.

1.2 Why study sea level?

The most obvious effects of change in sea level can be found through sudden rises like that of storm surges or tsunamis, most recently and dramatically shown by the Sumatra earthquake in December 2004 (see, for example Stein and Okal, 2005). Rises in sea level over longer periods have long been linked with increased storm surges (eg Rossiter, 1962). A large proportion of sea-level research, therefore, has been devoted to predicting future sea level.

Pre-historic sea-level change is also studied for a number of reasons. Firstly, it is strongly linked to wider changes in climate so it may assist in the study of the past earth system (Warrick *et al.*, 1993; Woodworth *et al.*, 1992; Mitrovica *et al.*, 2001). Secondly, and consequently, it can assist in the understanding and movement of human and animal life in a region (Fischer, 1995; Bailey, 1978; Bateman *et al.*, 1991). Lastly, it is supposed that past sea level can be measured with reasonable accuracy such that models describing sea-level change can be created.

1.3 Thesis overview: a statistician's perspective

The nature of the problem to be overcome in this thesis is similar to the well-studied topic of regression. There are a set of data points, observed with error, which can be plotted on an (x, y) plane through which a line must be produced. In the x -direction the nature of the error is of non-

standard distributional form but can be easily simulated, allowing the production of all necessary information. On the y -axis, the data can be considered to be Normally distributed with *known* variance. There are three types of y data points: those which the line should pass through (known as *basal index points*), those which the line should almost always pass above (*intercalated index points*) and those which the line should almost always pass below (*limiting points*). A further source of data is available in the form of the *tendency*, z . This is a measure (with uncertainty) of whether the line is rising or falling at particular times and is only given the values 1 and -1 respectively. In chapter 5, this type of data is used by comparing it to derivative information about the fitted line. A final piece of information available is the geographical location at which each (x, y) or (x, z) pair was obtained. Traditionally in sea-level studies this information is not considered; the parameters are assumed to have no spatial variability within arbitrary regional units. The thesis shows how all this information could be used together.

A complication occurs in the form of the set of explanatory variables, x . These are obtained as radiocarbon dates, which are discussed in chapter 4. It is possible to simulate their distribution through the inverse transformation of a Normal distribution via a piece-wise line function with known error bands.

The method proposed in this thesis is that of non-parametric regression based on Gaussian processes. The method has the advantage of taking into account all the available information, though it performs poorly when asked to predict outside the range of the data. The result is given

in the form of a large set of simulated lines: posterior estimates of a random function given the data. This set is analysed using functional data analysis techniques and simple numerical analysis to provide inference for sea-level change and related factors.

1.4 Thesis overview: a geologist's perspective

This thesis outlines a novel way to examine relative sea-level data. Traditionally, index points have been plotted on age-altitude graphs with a view to inferring sea-level change (eg Shennan and Horton, 2002). The uncertainties involved in the data, as well as the different sources of information, make the drawing of a sea-level curve difficult. This thesis presents a statistical method for drawing such curves.

The case study with which the statistical method is examined is that of the Humber estuary. The Humber has been the source of a number of sea-level studies (Gaunt and Tooley, 1974; Long *et al.*, 1998; Shennan *et al.*, 2003; Metcalfe *et al.*, 2000), which provide a well-defined history of sea-level change. It has also been the subject of a number of archaeological surveys (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001). These data have also been used to inform sea-level change, and their usefulness is assessed.

The statistical models used in this thesis integrate all the available

data, weighted according to its information content. The output from the models is used to infer changes, not only in sea level, but also in terms of the rate of change. It is compared with physical models for the region, and is also used to check for compatibility with wider changes in climate.

1.5 Thesis outline

1.5.1 Chapter 2

Chapter 2 contains some geological background to the study of sea level. This chapter contains much of the terminology used throughout the thesis. It first considers the nature of sea level itself, and follows by describing how this has been reconstructed using a wide variety of techniques. This chapter is then linked to the rest of the thesis by discussing the relevant uncertainties in sea-level research.

1.5.2 Chapter 3

Chapter 3 contains a description of the data available for this thesis. Available in two parts, the data arise as collections of studies from two different perspectives; those of archaeologists and geologists. The archaeological data require some pre-treatment to be suitable for combination with the other type. The methods for achieving this are discussed.

1.5.3 Chapter 4

Chapter 4 gives a first look at the nature of the structure underlying the data. The first part is concerned with radiocarbon dating, an aspect previously introduced in chapter 2. The data are then further explored using kernel smoothing methods, together with supplementary techniques for adjusting the intercalated and limiting points. The chapter ends with a discussion on the use of tendency data.

1.5.4 Chapter 5

Chapter 5 introduces the Gaussian process method, initially with a view to smoothing the temporal data. The flexible form allows almost all of the aspects involved in sea-level research to be accounted for and used to improve estimation. The chapter includes an example from which the salient points are highlighted.

1.5.5 Chapter 6

Chapter 6 outlines the models used in this thesis to provide inference on relative sea-level change in the Humber. These models include those with both stationary and non-stationary covariance structures, as well as temporal and spatio-temporal Gaussian processes. The models are compared using the deviance information criterion.

1.5.6 Chapter 7

Chapter 7 examines in more depth the modes of variation for a selection of the models outlined in the previous chapter. This is done using techniques such as kernel density estimation, functional principal components analysis, splines and change-point regression. The chapter ends by considering some questions posed by the sea-level community which may link sea-level change into a more general climate framework.

Chapter 2

An introduction to the reconstruction of former sea level (with a focus on the UK and the Humber Estuary)

2.1 Introduction

This chapter considers the nature of sea-level measurement. The production of data which purport to be estimators of sea level is a complicated process, and contains caveats of which the user of sea-level data should be aware. This chapter initially considers how sea level is measured and the fundamentals of tidal systems, and continues by detailing the methods used to create former sea-level measurements. Following a brief section on some of the physical models which attempt to reconstruct former sea level, it finishes by noting some of the

uncertainties involved in the study.

2.2 Sea level

2.2.1 What is sea level?

Sea level is usually referred to in terms of relative change. Relative sea level (RSL) is made up of a complex mixture of factors. Shennan and Horton (2002) express it formulaically as:

$$\Delta\xi_{\text{rsl}}(\tau, \psi) = \Delta\xi_{\text{eus}}(\tau, \psi) + \Delta\xi_{\text{iso}}(\tau, \psi) + \Delta\xi_{\text{tect}}(\tau, \psi) + \Delta\xi_{\text{local}}(\tau, \psi) \quad (2.1)$$

where τ is time and ψ is location. $\Delta\xi_{\text{rsl}}(\tau, \psi)$ is the change in RSL. $\Delta\xi_{\text{eus}}(\tau, \psi)$ and $\Delta\xi_{\text{iso}}(\tau, \psi)$ are the eustatic and isostatic change respectively. These are discussed below. $\Delta\xi_{\text{tect}}(\tau, \psi)$ is the tectonic change which is considered negligible for the UK, and $\Delta\xi_{\text{local}}(\tau, \psi)$ is local change, made up of that from tidal change and sediment consolidation.

2.2.2 Eustatic changes

Eustatic change is the adjustment of the volume of water in the ocean basins. There are many difference factors and complex inter-workings which cause eustatic change. These include changes in the shape and rotation of the earth (geoidal eustasy, Morner, 1987), changes in the shape of the ocean basins (tectono-eustasy, Devoy, 1987) and density

changes of the water (steric change, Warrick and Oerlemans, 1990). In general, water may be lost to (or obtained from) lakes, the atmosphere, ice, or trapped as groundwater. Estimates of the volumes of these different forms may be found in Pirazzoli (1996).

One of the main causes of eustatic change of particular importance to sea-level studies in the UK over the course of the Holocene is that of glacio-eustatic change caused by melting glaciers. The exact rate of melting is an important consideration in geophysical modelling (Shennan and Horton, 2002; Peltier, 1998a).

2.2.3 Isostatic changes

The measurement of temporal changes in eustatic sea level is hampered by the movement of the landmasses within which it is contained. The change in land concerned with the effects of load on the earth's crust is known as isostasy. In effect, the crusts are floating on the upper mantle of the earth's core. Any movement in the crust will therefore be a function of the viscosity of this liquid, the thickness of the crust (the lithospheric thickness) and the load applied to it (Lambeck, 1993). Problems of this nature were considered by Peltier (1998c) and in all geophysical models (see section 2.3.9). Occurrences such as this are most common during periods of glaciation; this is known as glacio-isostatic change.

Periodic glaciations have affected the shape of the land through glacio-

isostatic change; this is well studied in the UK (Shennan and Horton, 2002). Other forms of isostatic change important to RSL study in the UK are hydro-isostasy and sediment-isostasy . These occur as a result of extra meltwater load in the ocean's basin (Hopley, 1983) or increased sediment loading. Estimated rates of isostatic change in the UK are shown in Figure 2.1. The map shows positive isostatic change centered over Scotland; the location of the body of glacial load. Conversely, southern England (particularly Cornwall) shows negative isostatic change, which indicates a sinking of land level.

In most cases, a separation of isostatic movement and eustatic sea-level change cannot be made. It is for this reason that most work studies the phenomenon of relative sea level (RSL); the relative change of land to sea surface. Furthermore, the tectonic changes, vastly variable across the earth, contribute to the problem of global comparisons of sea-level change (Devoy, 1987). However, in some tectonically stable areas which are not influenced by glacio-isostatic change, such as Barbados, absolute sea level can be measured (Bard *et al.*, 1990). The geophysical models of Peltier (1974) and Lambeck and Chappell (2001) both consider the problem of isostasy in order to form a holistic view of RSL change.

2.2.4 Tectonic change

As mentioned in section 2.2.1, tectonic changes are believed to have had very little effect on RSL in the UK. However, Lambeck (1988) has

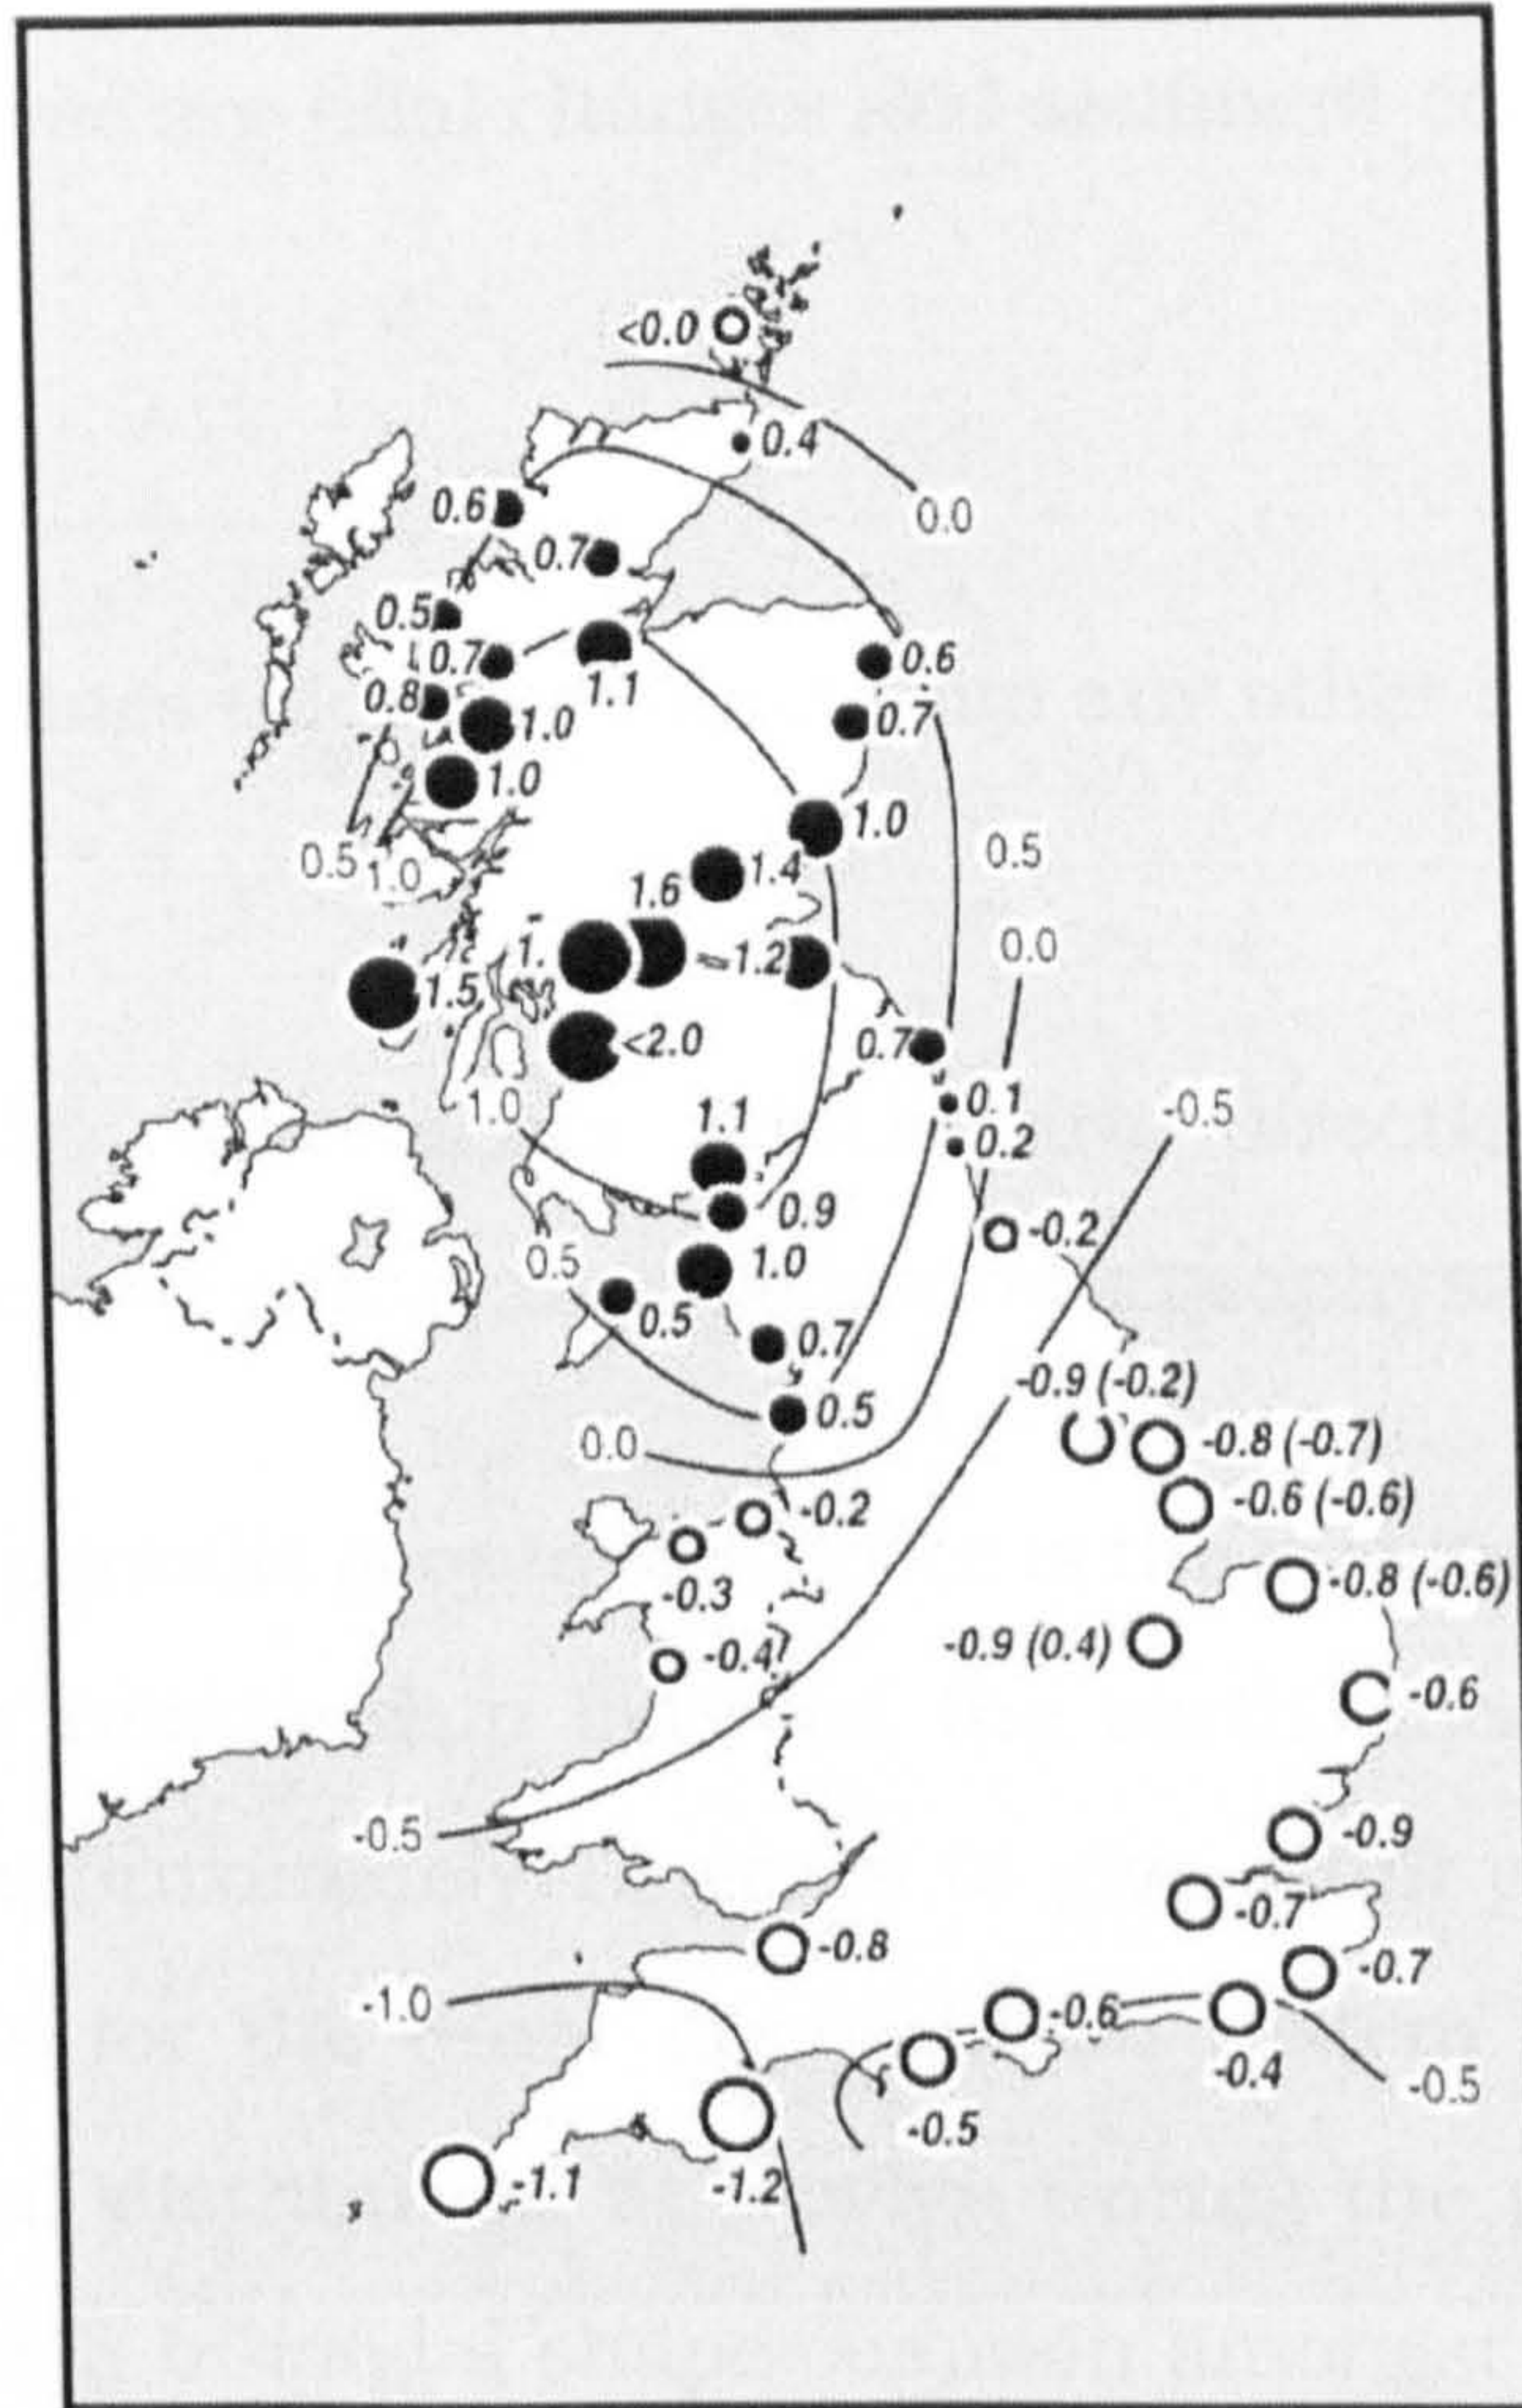


Figure 2.1: Estimates of isostatic change in the UK for the late Holocene in mm/year. From Shennan and Horton (2002).

shown that areas can still be affected where the lithospheric thickness is low.

Figure 2.2 shows an example of a coastal plain where the mean sea level is rising rapidly, and where the highest sea level is the highest. The values of the tide levels are as follows:

Highest (lowest) at Lowest tide (LTL) The highest (lowest) level reached by the tide is the highest (lowest) level reached by the tide.

2.2.5 Local scale factors

The local scale factors were outlined in equation 2.1. The most important of these are tidal changes and sediment compaction.

Tidal changes

Pugh (1987) defines tides as distinct from any other change in sea level to be

“...periodic movements which are directly related in amplitude and phase to some periodic geophysical force.”

The dominant periodic geophysical force is the gravitational field system caused by the relationship between the earth, moon and sun. The period takes approximately 19 years to give a full picture of the tidal system (this is for the lunar gravitational system only). Figure 2.2 shows the tidal distribution at Newlyn during the period 1951-1969. This graph has a bi-modal shape common amongst semi-diurnal tidal systems. Differences occur in places where tides are more simply diurnal or in amphidromic systems where tidal range is zero.

There are two main planetary factors involved in the tidal definitions shown in Figure 2.2. These are spring tides; where the moon and sun are in (or near) alignment, and neap tides; where the opposite is the case. The definitions of the tidal levels are as follows:

Highest (lowest) astronomical tide (HAT) The highest (lowest) level reached under any of the combination of periodical geophysical

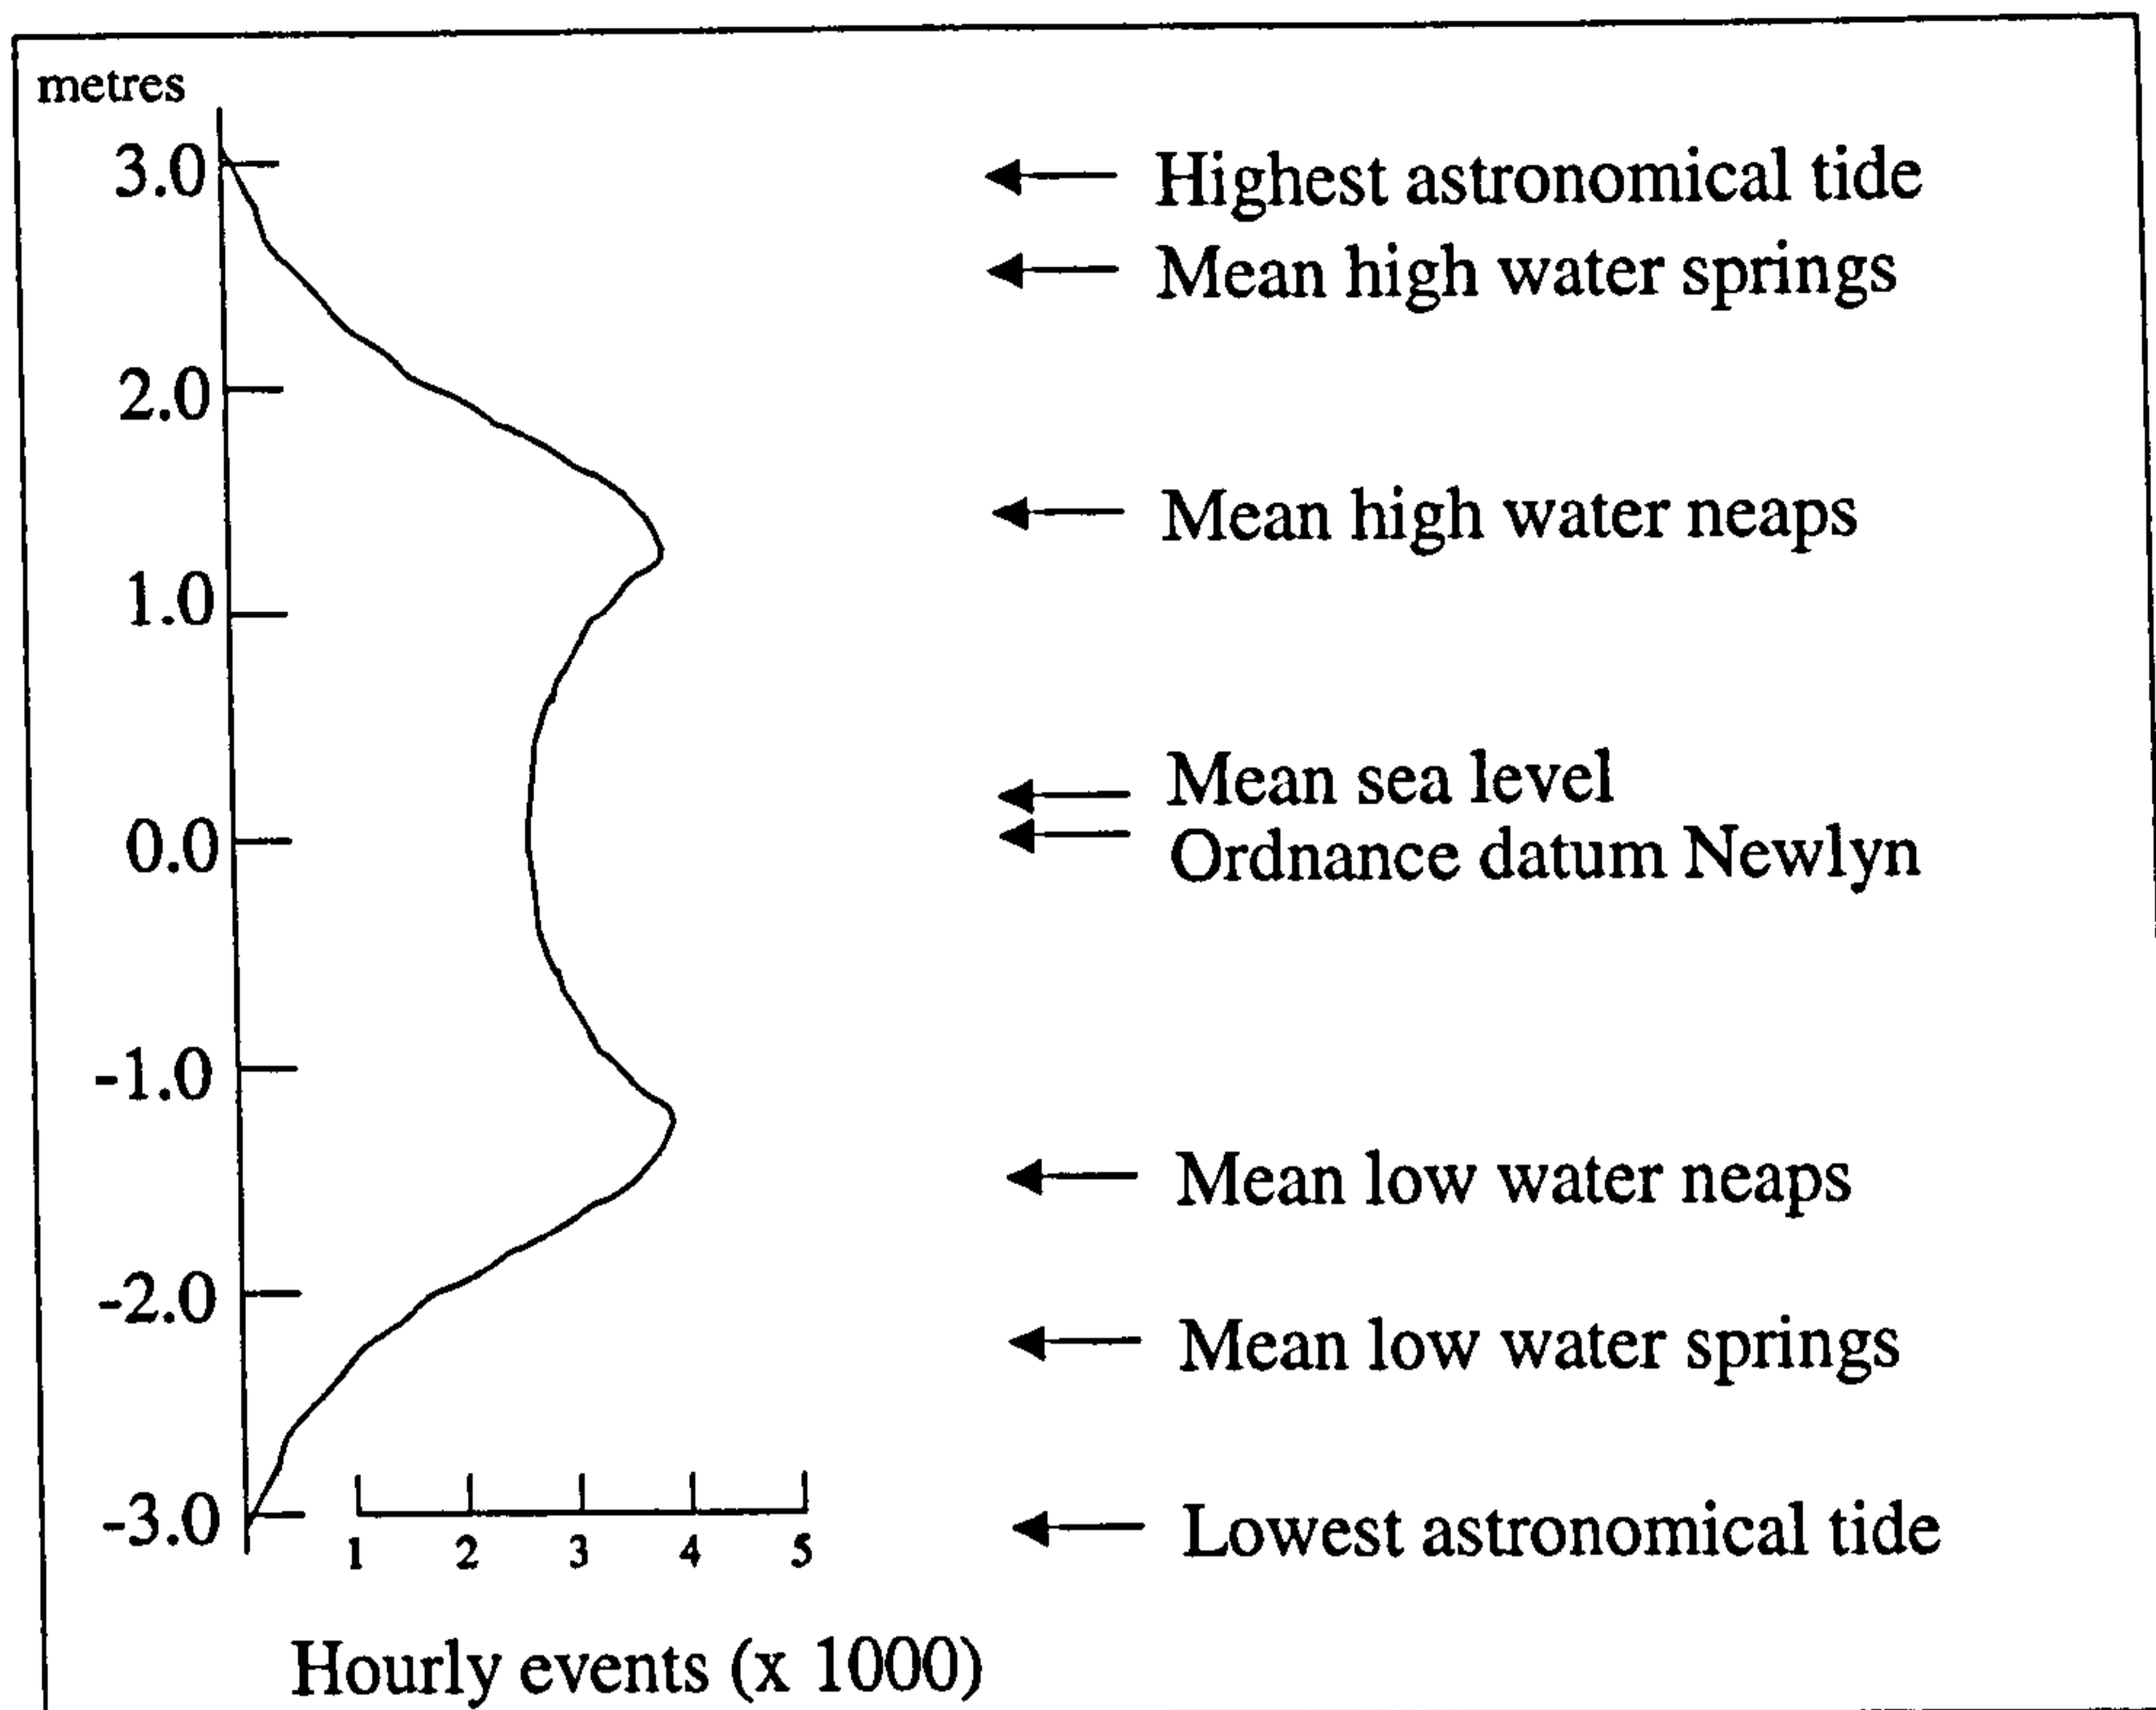


Figure 2.2: Frequency distribution of hourly tidal levels at Newlyn, Cornwall (Pugh, 1987). N.B., not all tidal patterns have two distinct peaks.

forces.

Mean high (low) water springs (MHWS or MLWS) The average high (low) level reached during spring tides.

Mean high (low) water neaps (MHWN or MLWN) The average high (low) level reached during neap tides.

Mean sea level (MSL) The arithmetic mean (preferably computed over the full 19 year cycle).

Mean tide level (MTL) The arithmetic mean of high and low waters. Often used when more detailed data is unavailable, thus is used in many RSL studies. Should only differ from MSL slightly in shallow water systems where tidal distribution is asymmetric.

Ordnance Datum Newlyn (ODN) The benchmark used for the UK as defined in section 2.2.1. Usually given in metres.

It should be noted that, in Figure 2.2, the mark of zero metres O.D. is very close to MSL but they are not identical. The value 0m O.D (Ordnance Datum calculated from Newlyn, Cornwall), which serves as the basis for RSL work in the UK, was obtained from six years of measurement from May 1915 to April 1921. Since relative sea level (see section 2.2.3 for definition) has increased since this time, MSL at Newlyn has increased slightly above this value.

In RSL studies, there are two main factors involving tidal changes. The first is that of the daily change in RSL due to that already mentioned (studied in detail in Pugh, 1987; Brown *et al.*, 1989). The second main factor is that of spatio-temporal tidal range change. This occurs where, for example, the difference in height between HAT and MSL changes in both space and time. This is discussed further in section 2.3.10.

Sediment consolidation

Sediment consolidation or compression (eg Paul and Barras, 1998) is a problem in RSL work and has been under-researched. The problem exists because sedimentary indicators of RSL such as peat can consist of large amounts of water (up to 90%). When pressure is applied to the peat in the form of overlaying sediments this water is removed and the peat contracts. There is also natural pliability in the peat itself from the plant remains from which it is created. The problem can be (partially)

removed by looking only at basal index points (see section 2.3.4).

Occasionally, sea-level data are created from *intercalated* sediments. These are formed within other layers and are far more subject to sediment consolidation. Paul and Barras (1998) determined that compression depends on the thickness of the sequences, the nature of the sediments (ie mass, particle size, etc) and the history of post-depositional loading. They concluded that the nature of the compression is non-linear:

“At shallower depths, the compression is less because the surcharge load is less, whereas at greater depths it is also less because the thickness of underlying compressible sediment is less.”

Thus the compression reaches a maximum level at approximately the centre of the elevation range (depending on the factors outlined above). They estimate this mid-depth compression to be in the order of 5-10% of the layer thickness. Whilst this may appear small considering the maximum peat water composition of 90%, this may impact on RSL studies as the layer thickness is dependent on the size of the tidal range.

2.3 Reconstructing former sea level

2.3.1 Sea-level index points

A sea-level index point (SLI) is a spatio-temporal estimate of the elevation of relative sea level. It has four parts:

- Location
- Age
- Elevation
- Tendency

Each part of an index point should have an associated error (though the location is usually known to such an exact standard that any error term is ignored). The date is usually calculated using radiocarbon techniques and subsequently calibrated, as discussed in section 2.3.8. The elevation is calculated by quantifying the indicative meaning (see below); by finding the reference water level (RWL) with an error given by the indicative range (IR). There are other errors that are also included in the elevation part of an index point. These are outlined in section 2.4.2. The tendency part of an index point is discussed in section 2.3.7.

2.3.2 Creating SLIs

Evidence for change in former sea level can broadly be categorised under the following headings:

- Biological
- Geomorphological
- Sedimentary
- Archaeological

The end result of the analysis of these indicators is the production of an SLI. However, the production of an index point is a complicated process, sometimes requiring detailed stratigraphical and microscopic analysis of sediments (as outlined in van de Plassche, 1986). Usually, many different forms of evidence are amassed to create a SLI.

Biological indicators

The basis of using biological indicators in sea-level research is the assumption of uniformitarianism (Hutton, 1788). For example, a species from many thousands of years ago that is identical to one that only lives within the tidal range today is assumed to have also lived within the tidal range in the past. A full list of possible Biological indicators and their descriptions would require a separate essay. However, a brief description of some of the main indicators is given below:

pollen Used where plant remains are found in the tidal range to discern transitions between fresh and saltmarsh plants (eg Long *et al.*, 1999).

diatoms Unicellular algae used in a similar vein to pollen. Can again be used as sea-level estimators and also for sea-level tendencies (see section 2.3.7). Examples include Shennan (1989); Zong and Horton (1998).

foraminifera Single-celled organisms which can be preserved in sediments. Foraminifera are highly sensitive to changes in temperature and salinity and therefore make excellent indicators of former sea level, eg Horton *et al.* (2000).

Geomorphological indicators

Geomorphological indicators are usually of larger scale than those of biological nature. Some sea-level studies are based solely on geomorphological evidence (eg Firth *et al.*, 1995). They are useful as a broad indication of former sea level and help form a body of evidence in a SLI. Geomorphological signs of former sea level include raised beaches (eg Rees-Jones *et al.*, 2000), deltas, spits, caves (eg Baldini *et al.*, 2002) and coral reefs (eg Henderson, 2005).

Sedimentary indicators

Sedimentary indicators refer to the depositional layers that can be found to offer evidence of change in RSL. The two main forms of sediments are that of clastic, inorganic deposits and biogenic deposits made of plant and animal matter. Sedimentary deposits are used for much of the work on index points (see section 2.3.4 and Shennan 1982) provided they show evidence of being related to former sea level. In particular, it is the transition between clastic and biogenic sediments

that form the basis of many SLIs. A more precise estimate of former RSL can be obtained from looking at the biogenic indicators within the sedimentary layers.

Archaeological indicators

Archaeological remains in the littoral zone or below are in abundance in the UK. Occasionally, the indicative meaning for these artifacts can be quantified. Pirazzoli (1996) categorises these data as follows:

1. Those which must have been located near a shoreline and, according to their use, must have remained above or below the sea.
2. Those belonging to structures partly underwater and depending for their use on tidal fluctuations and marine conditions.

Category 1 includes settlements and other forms of dry-land activity. When there is no evidence of displacement, they can be used as limiting dates for RSL. They are more readily found during the latter Holocene where lithographic or biogenic data is sparse due to human interference and accommodation space.

Category 2 data includes items such as tethered boats or fish traps. Here, the indicative meaning may be quantified, although with varying accuracy. These data types can be found for the Humber in van der Noort and Davies (1993).

Archaeological data can also be useful in estimating former tidal range. For examples, see section 2.3.10.

2.3.3 Indicative meaning

The main crux of the creation of an index point is the study of the indicative meaning. This is described by Shennan (1994) as:

“...the relationship of a sea-level indicator to a tide level, defined in terms of the indicative range and reference water level.”

The study of the indicative meaning is a method for comparing the height at which a sea-level indicator was found in the modern environment with a set of tide levels believed to represent its relationship to former RSL. The result of the application of the indicative meaning is the production of a RWL and an IR. As an example of the use of the indicative meaning for some types of peat, Table 2.1 displays both the IRs and RWLs.

The RWL is given as a function of tide levels. For example, a *Phragmites* peat which lies directly above a saltmarsh deposit is estimated to lie 20cm above a point between highest astronomical tide (HAT) and mean high water spring tide (MHWS). The IR is then an estimate of error around the RWL due to the natural variation in the habitat of the species. It should be noted that the IR varies in this case according to the relationship of the subject to its surrounding material.

Dated material	Indicative range(cm)	Reference water level
<i>Phragmites</i> or monocot peat		
Directly above clastic saltmarsh deposit	±20cm	[(MHWS ± HAT)/2] - 20cm
Directly below clastic saltmarsh deposit	±20cm	MHWS - 20cm
Directly above fen wood deposit	±20cm	MHWS - 10cm
Directly below fen wood deposit	±20cm	[(MHWS + HAT)/2] - 10cm
Directly above and below clastic saltmarsh deposit	±40cm	MHWS
Middle of layer	±70cm	Infer from stratigraphy
Fen wood peat		
Directly above <i>Phragmites</i> or clastic saltmarsh deposit	±20cm	(MHWS + HAT)/2
Directly below <i>Phragmites</i> or clastic saltmarsh deposit	±20cm	MHWS
Basis peat		
Directly below <i>Phragmites</i> or clastic saltmarsh deposit	±20cm	MHWS
Directly below fen wood deposit	±80cm	MTL - MHWS

Table 2.1: Indicative range and reference water-level for commonly dated materials (from Shennan, 1982)

Whilst the indicative meaning forces those studying RSL to think about the relationship between a stratigraphic record and the RSL it purports to show, it is also governed by the laws of uncertainty. It is perfectly feasible for different researchers to arrive at different indicative meanings from the same material (as shown by, for example, Horton *et al.*, 2000). Thus, this derived measure of RSL suffers from its own problem of interpretation.

2.3.4 Basal index points

Most RSL studies, eg Shennan and Horton (2002), refer to there being three types of index point; basal, intercalated and limiting. Basal index points are called so because they are found at the base of that stratigraphic record, therefore denying the possibility of any compression (van de Plassche, 1982, 1986). In this thesis, basal index points are referred to as BIPs. A further distinction concerns that between basal and base-of-basal peats. Base-of-basal refers to when the sediment section used as an SLI lies *directly* above non-compressible material thereby removing even the slightest chance of compression. However, as pointed out by Shennan (1994), basal index points do not always strictly describe RSL. It is plausible for peat to form during a rise in local water-levels, independent of any RSL change.

2.3.5 Intercalated index points

Intercalated index points (IIPs) are found in the middle of layers of compressible sediment, and thus may well be below the true level of RSL for their given age. These are separate from basal index points because, though they can be related to a past tidal level (and therefore have indicative meaning), they do not accurately represent former RSL.

2.3.6 Limiting points

At sites where the stratigraphic records cannot specifically be referred to a past tidal level, limiting points (LPs) can be created. Although not as useful as basal index points are, they provide useful bounds on the RSL curve. In many RSL studies, they are ascribed an indicative meaning with a RWL of around or above MHWS. For examples see Shennan and Horton (2002).

2.3.7 Tendency

The tendency of a SLI describes whether the point records an increase or decrease in water level or salinity (Shennan *et al.*, 1983). The tendency of a SLI may be expressed as either positive (an increase in marine influence) or negative (a decrease in marine influence). Traditionally, they were used to compute sea-level curves, eg Gaunt and Tooley (1974). However, Shennan (1994) cautions on putting too much store by the tendency of a single SLI:

“A tendency of sea-level movement is not a synonym for an altitudinal change of sea-level since incomplete data or the errors associated with the data may preclude such a conclusion.”

It is suggested that a number of similar tendencies (spatio-temporally), known as a *dominant* RSL tendency, are required for any conclusions to be drawn (Allen, 1999).

For the Fenlands, Shennan (1994) devised a nine point code to account for the different strengths of positive and negative tendency. This coding is shown in Table 2.2.

The table shows that not all SLIs generate strong tendencies (this could be considered as the 'error' in the tendency). For code 3, where the base of a basal peat is found, the tendency cannot be calculated as there is no contact with older sediments. This does have the advantage, however, of removing the problem of sediment compaction (see section 2.2.5). For code 8, the possibility of compaction rules out knowledge of the tendency, but does provide chronological information. This is known as a limiting point (see section 2.3.6).

2.3.8 Radiocarbon Dating

A valid SLI must be datable. Whilst there are a wide variety of dating techniques available, the most common used for the dating of

Code	Description
1	Good positive tendency of RSL movement. From a transgressive contact, or a stratigraphic or biostratigraphic change, indicating a positive RSL tendency. There is no indication of erosion.
2	Poor positive tendency of RSL movement. From an apparent transgressive contact, but where either erosion is suspected or possible, or there is no supporting biostratigraphic evidence from the site.
3	Base of basal peat dated.
4	Sample from within the basal peat, but not from the base, or from the transgressive contact (ie not 1, 2 or 3).
5	Good negative tendency of RSL movement. From a regressive contact, or a stratigraphic or biostratigraphic change indicating a negative RSL tendency. There is no indication of a hiatus in sedimentation.
6	Poor negative tendency of RSL movement. From an apparent regressive contact, but where either a hiatus is suspected, or there is no supporting biostratigraphic evidence from the site.
7	Sample from any part of a sequence beyond the limit of marine sedimentation.
8	Sample from an intercalated peat where no single tendency can be defined. It represents a minimum age for a negative tendency followed by a positive tendency.
9	Sample cannot be reliably related to a specific RSL tendency in terms of either age, altitude or indicative meaning.

Table 2.2: Numerical code for RSL tendencies, from Shennan (1994)

Holocene SLI's is that of radiocarbon (first proposed by Libby *et al.*, 1949). This technique has associated uncertainty, but results have been consistently reproduced in numerous trials (Scott, 2003). For more detail about radiocarbon dating than is presented here, as well as other forms of dating, see either Bowman (1990) or Aitken (1994). For more information on how radiocarbon dating can be enhanced with the use of statistical techniques, see Buck *et al.* (1996).

There are three varieties of carbon atom, known as carbon 12, 13 and 14. Of these, only carbon 14 (known as ^{14}C) is radioactive. Its decay takes the form of beta particles, of which half are gone in approximately 5730 years (the half-life). The ^{14}C atoms are formed in the upper atmosphere, where they are mixed with oxygen to form carbon dioxide. This is then absorbed by plants during photosynthesis and subsequently into animals through the food chain. Furthermore, ^{14}C is mixed into the oceans, thus all living plant and animal life is constantly exchanging carbon with the environment, even under the sea. When a plant or animal dies it ceases to exchange carbon with the atmosphere and radioactive decay occurs. The amount of ^{14}C lost by the dead organism is exactly proportional to the radioactive decay it emits.

Extracting the radiocarbon from the sample is a complex chemical process. For sea-level data points used in this thesis, the majority of the samples consist of either peat or ancient wood. One of the main tasks of the laboratory is to remove any contamination in the sample. This is a particular problem with oceanic samples which may have received

carbon of an inappropriate age due to variable mixing rates (section 2.4.1).

For wooden samples, the cellulose is extracted as it is known to be the most reliable (Aitken, 1994). However, the extraction procedure drastically reduces the volume of carbon available to date; thus often large samples are required. A further problem concerns the difference between the date of formation of the carbon and the date of the archaeological event of interest. This is discussed in section 4.2.1.

Peat has the added complication that it can be separated into humins (solid vegetation) and humic or fulvic acids (soluble plant extract), either of which can be dated and either of which can suffer contamination. The humic acids can suffer from washed-through younger carbon whilst the humins may suffer from root intrusion or the inclusion of detrital older material (including coal). Even allowing for these variations, authors such as Shore *et al.* (1995) have discovered remarkably variable dates from different fractions even for stratigraphically related peats.

2.3.9 Geophysical modelling

Geophysical models use the index points for validation. Their idea is that, by taking account of all the physical factors involved in RSL change, any model will be able to accurately predict the RSL to within reasonable (mathematical) proximity of the SLIs. The major advantage

of geophysical models is that isostatic and eustatic change can be estimated independently.


The model proposed by Peltier (1974) is one of the better known geophysical models. A review can be found in Peltier (1998b), but perhaps the most simple non-mathematical explanation is in Peltier (1998a). The model is based on the following set of equations for RSL (though it separates eustatic and isostatic changes), defined here as $S(\delta, \lambda, t)$ where δ and λ are latitude and longitude respectively, and t is time:

$$S(\delta, \lambda, t) = C(\delta, \lambda, t)[G(\delta, \lambda, t) - R(\delta, \lambda, t)] \quad (2.2)$$

$$= C(\delta, \lambda, t) \left[\int_{-\infty}^t dt' \int_{\Omega} d\Omega' L(\delta', \lambda', t') \times \left(\frac{\phi(\gamma, t - t')}{g} - \Gamma(\gamma, t - t') \right) + \frac{\Delta\Phi(t)}{g} \right] \quad (2.3)$$

In equation 2.2, C is the 'ocean' function which is 1 at points where there is ocean and 0 elsewhere. The functions G and R are, respectively, the geoidal perturbation (distance from MSL, the geoid, to the centre of the Earth) and the radial perturbation (the distance from the so-called 'solid surface' of the Earth to the centre) respectively. Figure 2.3 from Mitrovica and Milne (2003) shows the difference between the geoid and MSL, and the resultant ocean function.

When these terms are expanded in equation 2.3, they both depend on the knowledge of $L(\delta', \lambda', t')$. This function describes the load per unit



UNIVERSITY
OF SHEFFIELD
LIBRARY

Figure 2.3: Diagram to show RSL as the difference between the geoid, G , and the solid surface, R . C is the ocean function. The spatial dependence parameters are suppressed. From Mitrovica and Milne (2003).

area, and is given a composite form:

$$L(\delta', \lambda', t') = \rho_I I(\delta', \lambda', t') + \rho_W S(\delta', \lambda', t') \quad (2.4)$$

where I is the ice thickness, ρ_I is the ice density, and similarly for water with ρ_W and S . Thus the RSL term S appears on both sides of the equation. The reader is referred to Peltier (1998c) for more information about the other variables in equation 2.3.

In simple terms, the user inputs detail about the ice loads at the last glacial maximum (LGM), along with lithospheric thickness and estimates of mantle viscosity. This information is then used with the equations outlined above to form a gravitationally consistent history of glacial activity, which in turn reports the RSL history. Many of the updates to the model have come about due to the enhancement

in estimation of the ice loads at the LGM, although more recent enhancements have included the addition of parameters to quantify the changing rotational state of the earth.

The computer modelling required to form such a large and complex model requires detailed integrations at each step. Furthermore, the model is often unable to pick up highly local features that lie apart from the visco-elastic predictions. Such problems occurred in the modelling of glacial forebulges which required adjustment of the model to a more complex viscosity structure (viscosity model 2). Two graphs which show the fit at arbitrary locations in the UK (including the Humber) are to be found in Figure 2.4. It should be noted that the model is only specified in 1000 (sidereal) year slices. The line drawn between the estimates is an interpolation. In the Humber data shown in Figure 2.4, Shennan and Horton (2002) found that correcting the index points for tidal range change brought about a better fit (see section 2.3.10).

2.3.10 Tidal range modelling

Knowledge of the tidal range is of primary importance in the consideration of the indicative meaning. There is strong belief (Austin, 1991; Woodworth *et al.*, 1991; Hinton, 1996; Shennan *et al.*, 2000a) that tidal range has changed drastically over the course of the Holocene. This means that the RWL from the indicative meaning must be corrected both spatially and temporally.




Figure 2.4: A comparison of the geophysical curve of Peltier (dashed line) and RSL index points. Left panel is for the outer Humber, right panel for inner Humber. Both include correction for tidal range change. Red points are limiting dates, green are intercalated and blue are basal. The solid line is the current estimated RSL rate calculated as the gradient between the last good IP to the origin. From Shennan and Horton (2002)

Many models of tidal range change (Gehrels *et al.*, 1995; Hinton, 1996; Austin, 1991) show that tidal range increases with RSL. However, the error in these models is rarely reported, though they are often compared with other models. The main deficiency with these models is lack of understanding of the inputs, particularly that concerning bathymetric change and sediment influx.

Another method for estimating tidal range change is the analysis of the sediment thickness directly from the sedimentary sequence. These can be compared with modern day estimates. Similarly, archaeological information can also be used to predict former tidal range. Sidell (2002) used Roman quays constructed of wood to look at former tidal range in the Thames. Tidal levels were left marked on to the wooden supporting poles. Whilst this does not give detailed measurement, average high/low tide values can be estimated. As with index points, these must be proved not to have suffered any displacement before they can be used as indicators of former tidal range.

2.4 Uncertainties in sea-level research

2.4.1 Radiocarbon dating and calibration

Some of the main uncertainties in radiocarbon dating are listed below. Many of them are taken account of, or reported by, the laboratory.

Isotopic fractionation During photosynthesis plants more readily

take up ^{12}C over ^{14}C (Craig, 1953). This would result in an excess age were it not for correction. The degree of isotopic fractionation varies amongst species.

Mixing rates In the deep ocean, ^{14}C is not replenished as quickly as in the atmosphere and at the earth's surface. This is relevant in terms of RSL modelling as sometimes this ^{14}C depleted water can well up to the surface, making radiocarbon ages on marine samples anomalously old (Stuiver and Braziunas, 1993). Similarly, water containing dissolved limestone can give older ages due to the dissolved geological-age carbon they contain (Fontes, 1992). This can be a problem with sea-level research. Correction factors can be applied for marine samples although the amount of correction varies spatially and temporally (Harkness, 1983).

Variable ^{14}C production The concentration of ^{14}C has not remained constant over time (DeVries and Barendesen, 1954). Throughout the Holocene (and before) levels of ^{14}C have fluctuated due to changes in the earth's magnetic field, or in the cosmic ray background (such as sunspots, eg Damon *et al.*, 1995). Human interaction has also affected ^{14}C levels, through the use of fossil fuels and atomic bombs. These fluctuations can be mapped through the use of a calibration curve, as described in section 4.2.1.

Contamination Any sample measured for radiocarbon activity must contain only material that is relevant to the date required. Contamination with older or younger carbon will lead to an erroneous date. Whilst laboratories will deal with 'the usual'

contaminants during pre-treatment of the sample, this will never correct for a plant that has, for example, fed on older carbon.

Another uncertainty concerning radiocarbon dating is the process of calibration (Bowman, 1990). The radiocarbon clock with which the laboratory dates are reported has not run at a constant pace. Rather it has quickened in places, occasionally stopped altogether, and in many places it even reverses back upon itself. Thus the relationship between a laboratory ^{14}C measurement and the calendar date it purports to show is non-linear. The calibration curve is required to map a radiocarbon date on the radiocarbon timescale on to the calendar timescale (Buck and Blackwell, 2004). Since the method of calibration for radiocarbon dates is a statistical issue, it is dealt with separately in section 4.2.1.

2.4.2 Estimation of index point elevation

Many uncertainties are associated with the nature of data collection in the field. These are included in a final estimate of the elevation error.

The six main components of this error are:

1. The levelling error associated with the sample. This can be highly variable where poor elevation benchmarks exist or are far away from the measurement site.
2. The sample thickness. The thickness of the sample introduces natural variability in the reference water level.
3. The tide level error of the sample. This is assigned in the Admiralty tide tables or must be estimated during tidal range modelling.

4. The altitudinal difference between the two nearest tide gauge stations (unnecessary when tidal range modelling is undertaken).
5. The indicative range. As discussed in section 2.3.4.
6. The compaction error. As discussed in section 2.2.5.

These errors are all considered Normally and independently distributed, and the values they are given are assumed to be estimated standard deviations. If this is the case, then the normal and correct procedure for an estimate of the standard error of the elevation for an SLI is to use the formula:

$$S.e.(h) = \sqrt{\sum_{i=1}^6 e_i^2} \quad (2.5)$$

In Equation 2.5, h is the elevation of the SLI, and e_i is the standard deviation associated with error i outlined in the list above. However, some of the errors have either not been extensively researched or are not calculated in the traditional form of a standard deviation. For example, the IR is usually given as half the difference between the maximum and the minimum where the biogenic indicator was found. This is not a standard deviation, and must be adjusted to correct the total elevation error. However, it is clear that when the original idea of the IR was proposed, in papers such as Shennan (1982), all the errors were given in the form of a 95% confidence interval.

Another problem with this method is that it assumes that covariance is not a factor. It may well be that some of the factors, eg the IR and the sample thickness are related. Finally, the total standard error must

be multiplied by 1.96 if it is to cover the full 95% interval on any age-altitude graph. If the untransformed standard error is simply plotted this will only give a 68% interval, allowing for a high possibility that the true height of the sample lies outside this range. This is true in many sea-level graphs, eg that shown in figure 2.5.

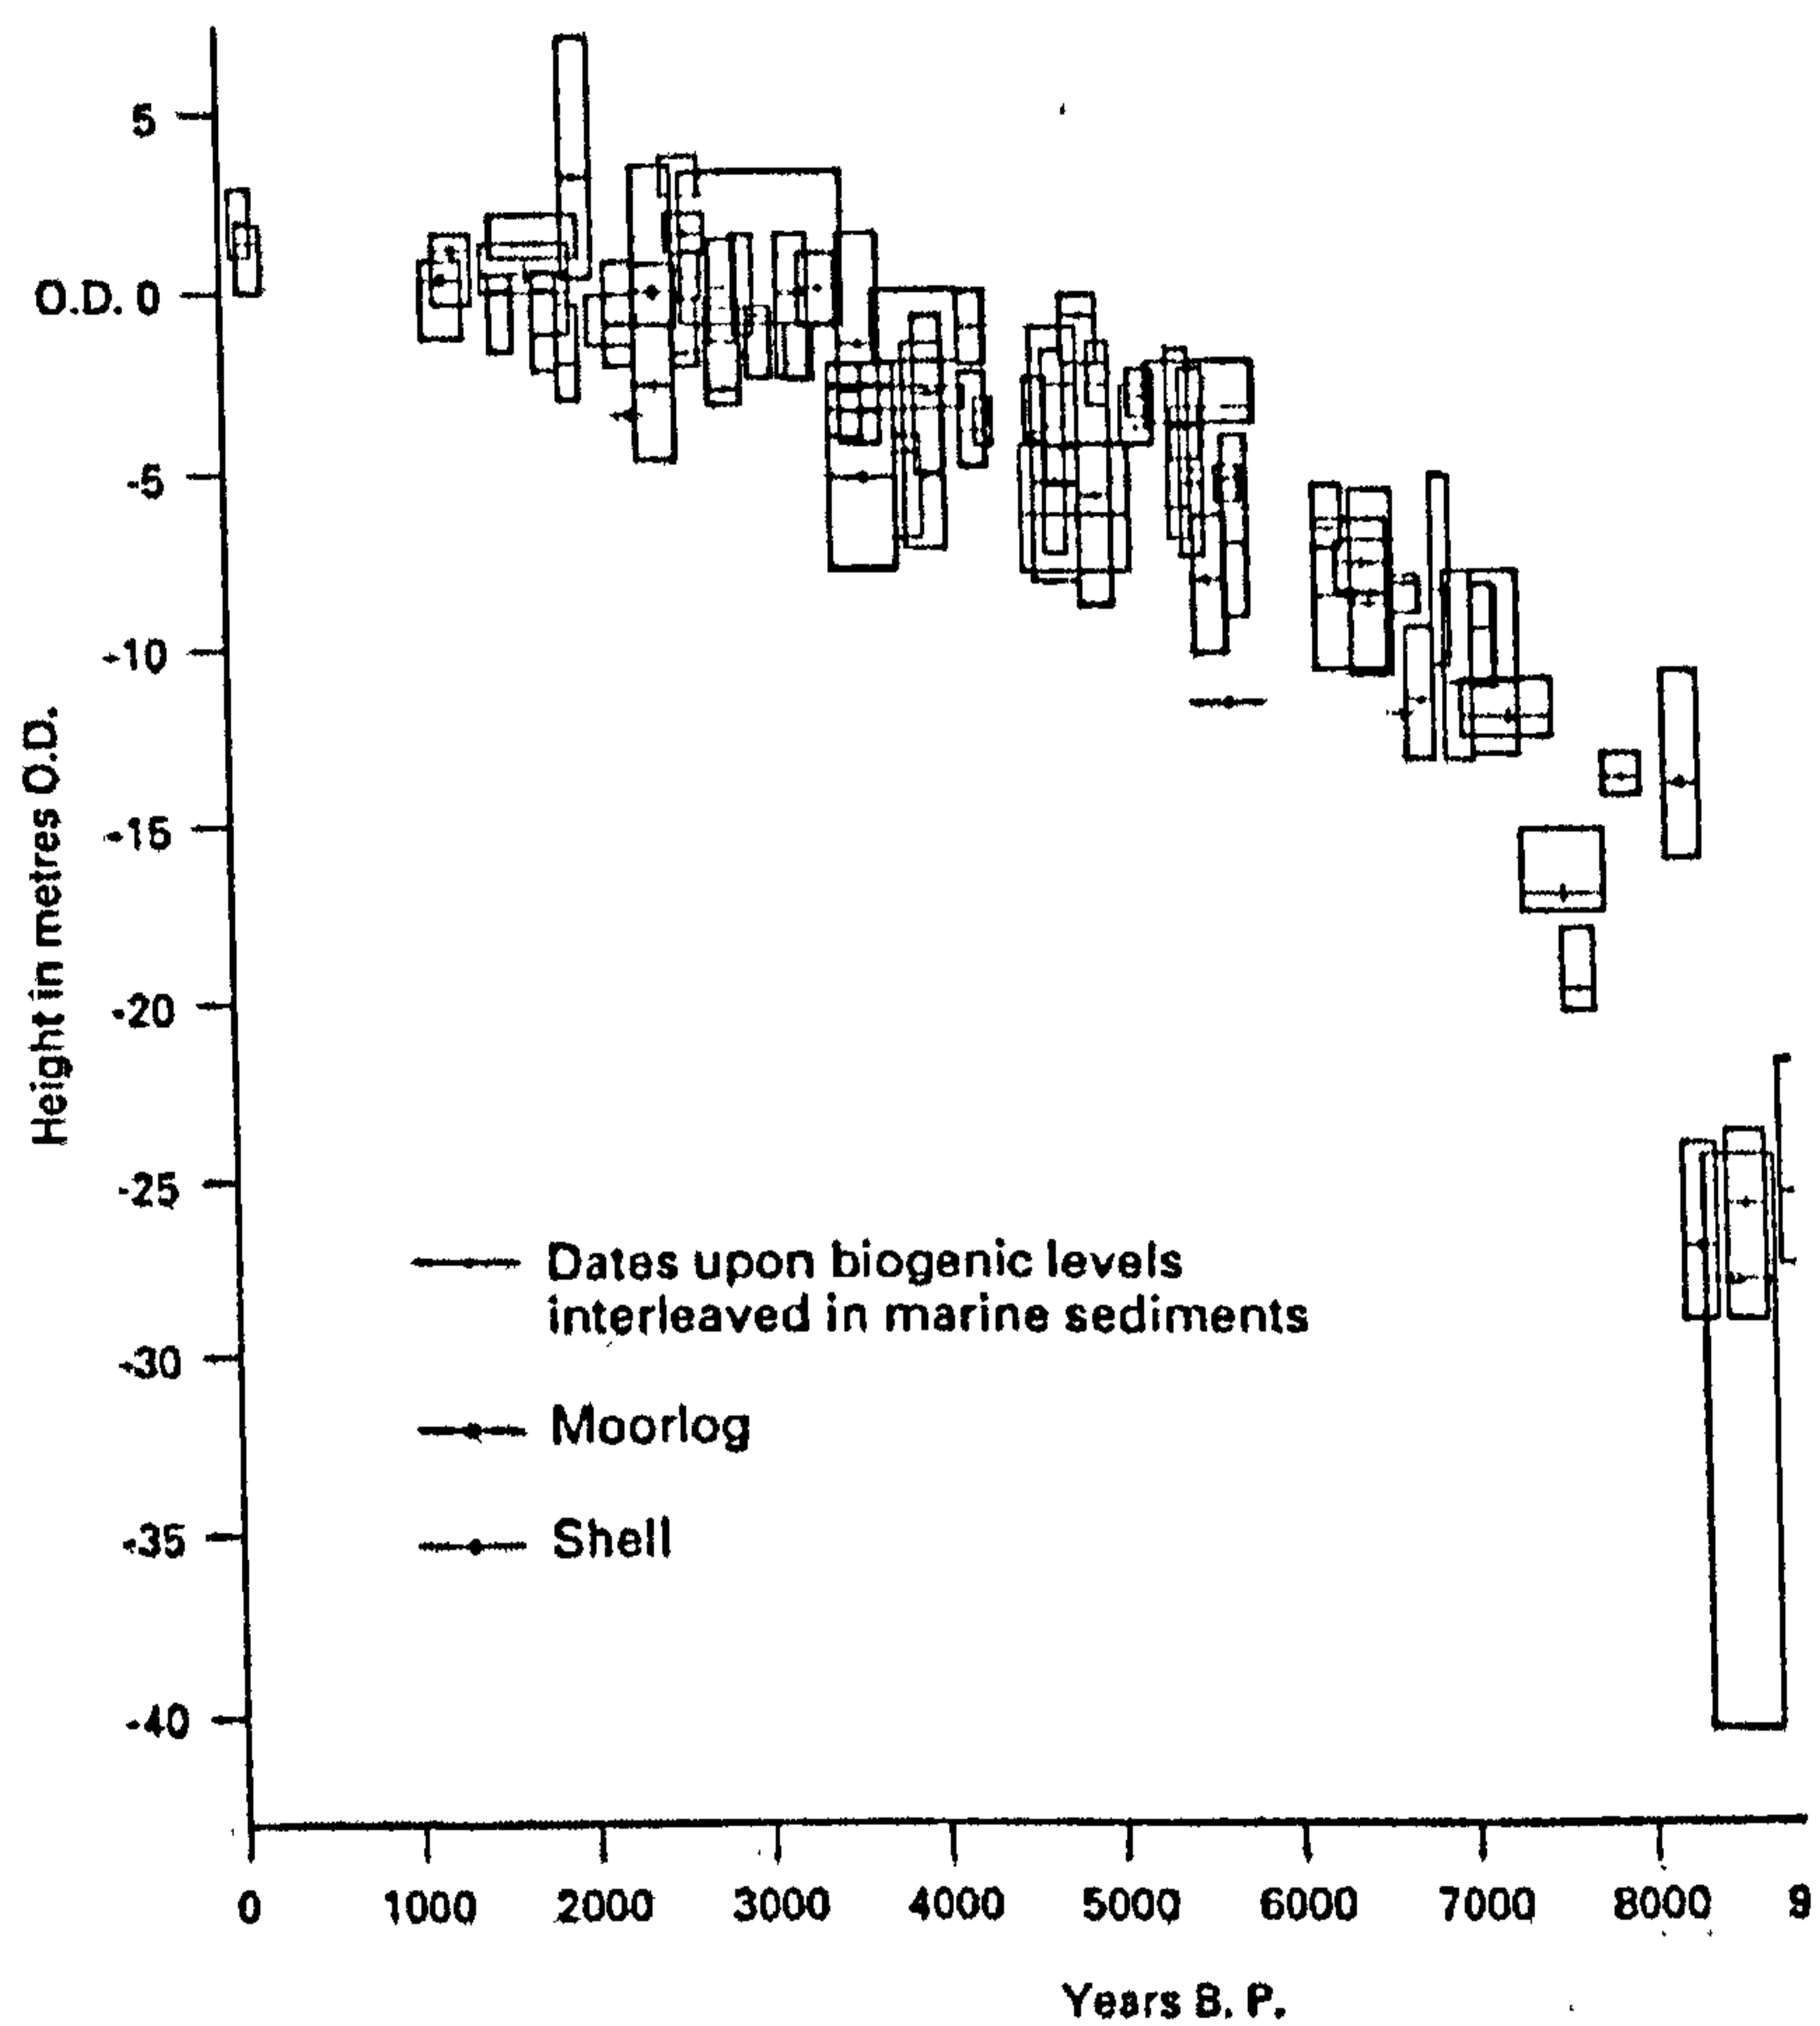


Figure 2.5: Age-altitude plot of SLIs for South East England. The ^{14}C dates are uncalibrated and represented with a 1σ error. The SLIs are represented as boxes to show the range of error. From Devoy (1982).

2.5 Summary

This chapter has dealt with most of the key features involved in reconstructing former RSL. Important concepts introduced include basal index points (BIPs), intercalated index points (IIPs), limiting points (LPs) and the measurement known as the tendency. This chapter has included the examination of the factors that go to make up RSL which are both global and local in their nature. The methods for reconstruction are diverse and complex. They can involve numerous sources for the creation of just one point and rely upon a wealth of different scientific disciplines. However, even though these data points have been created with care, calculated uncertainties still remain.

Chapter 3

The Data

3.1 Introduction

This chapter describes the data set available for this thesis. This is available in two parts, both of which consist of data published over the preceding 50 years. The first part is the Land-Ocean Evolution Perspective Study (LOEPS), contained within the Land-Ocean Interaction Study (LOIS) database (Shennan and Andrews, 2000). This is a list of all the sea-level index points (SLIs) collected for the east coast of England. The other source of data is that collected in the Humber Wetlands Project (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001). This data set is a summary of the archaeological remains in the area. The data requires some complex pre-treatment before it can be used. Indeed, this has become part of the project, see for example Section 2.3.4.

This chapter outlines the contents of both data sources. It begins by studying the nature of the Humber estuary as it came to be. The fields of the LOIS database are then discussed, with notes on the overlap between the Humber Wetlands Project (HWP) and LOIS catalogues. The section on the HWP contains a segment on the tidal distribution of the Humber, for reasons which will be made apparent. This section also contains details of the pre-treatment required to get data points into a suitable format for use by sea level scientists. The chapter concludes with an outline of the data points from the HWP which were deemed worthy of inclusion.

3.2 The Humber estuary

The study area, the Humber estuary, lies on the north east coast of England with the city of Hull on its north bank and Grimsby to the south. A map of the Humber estuary is shown in figure 3.1. The estuary itself takes freshwater from the rivers Trent, Ouse, Derwent, Aire, Don and Hull and discharges them into the North Sea. It is estimated that the Humber estuary has seen a large rise in sea level, approximately 30m (Shennan and Horton, 2002) over the course of the Holocene, mainly due to melt-water from the decline of the last glaciation (Gaunt and Tooley, 1974; Long *et al.*, 1998). This rise is consistent with regions at the edge or beyond the limit of the British ice sheet at the last glacial maximum. The evolution of the Humber estuary over the course of the Holocene is primarily dependent on the retreating glaciers which covered England during the previous ice age.

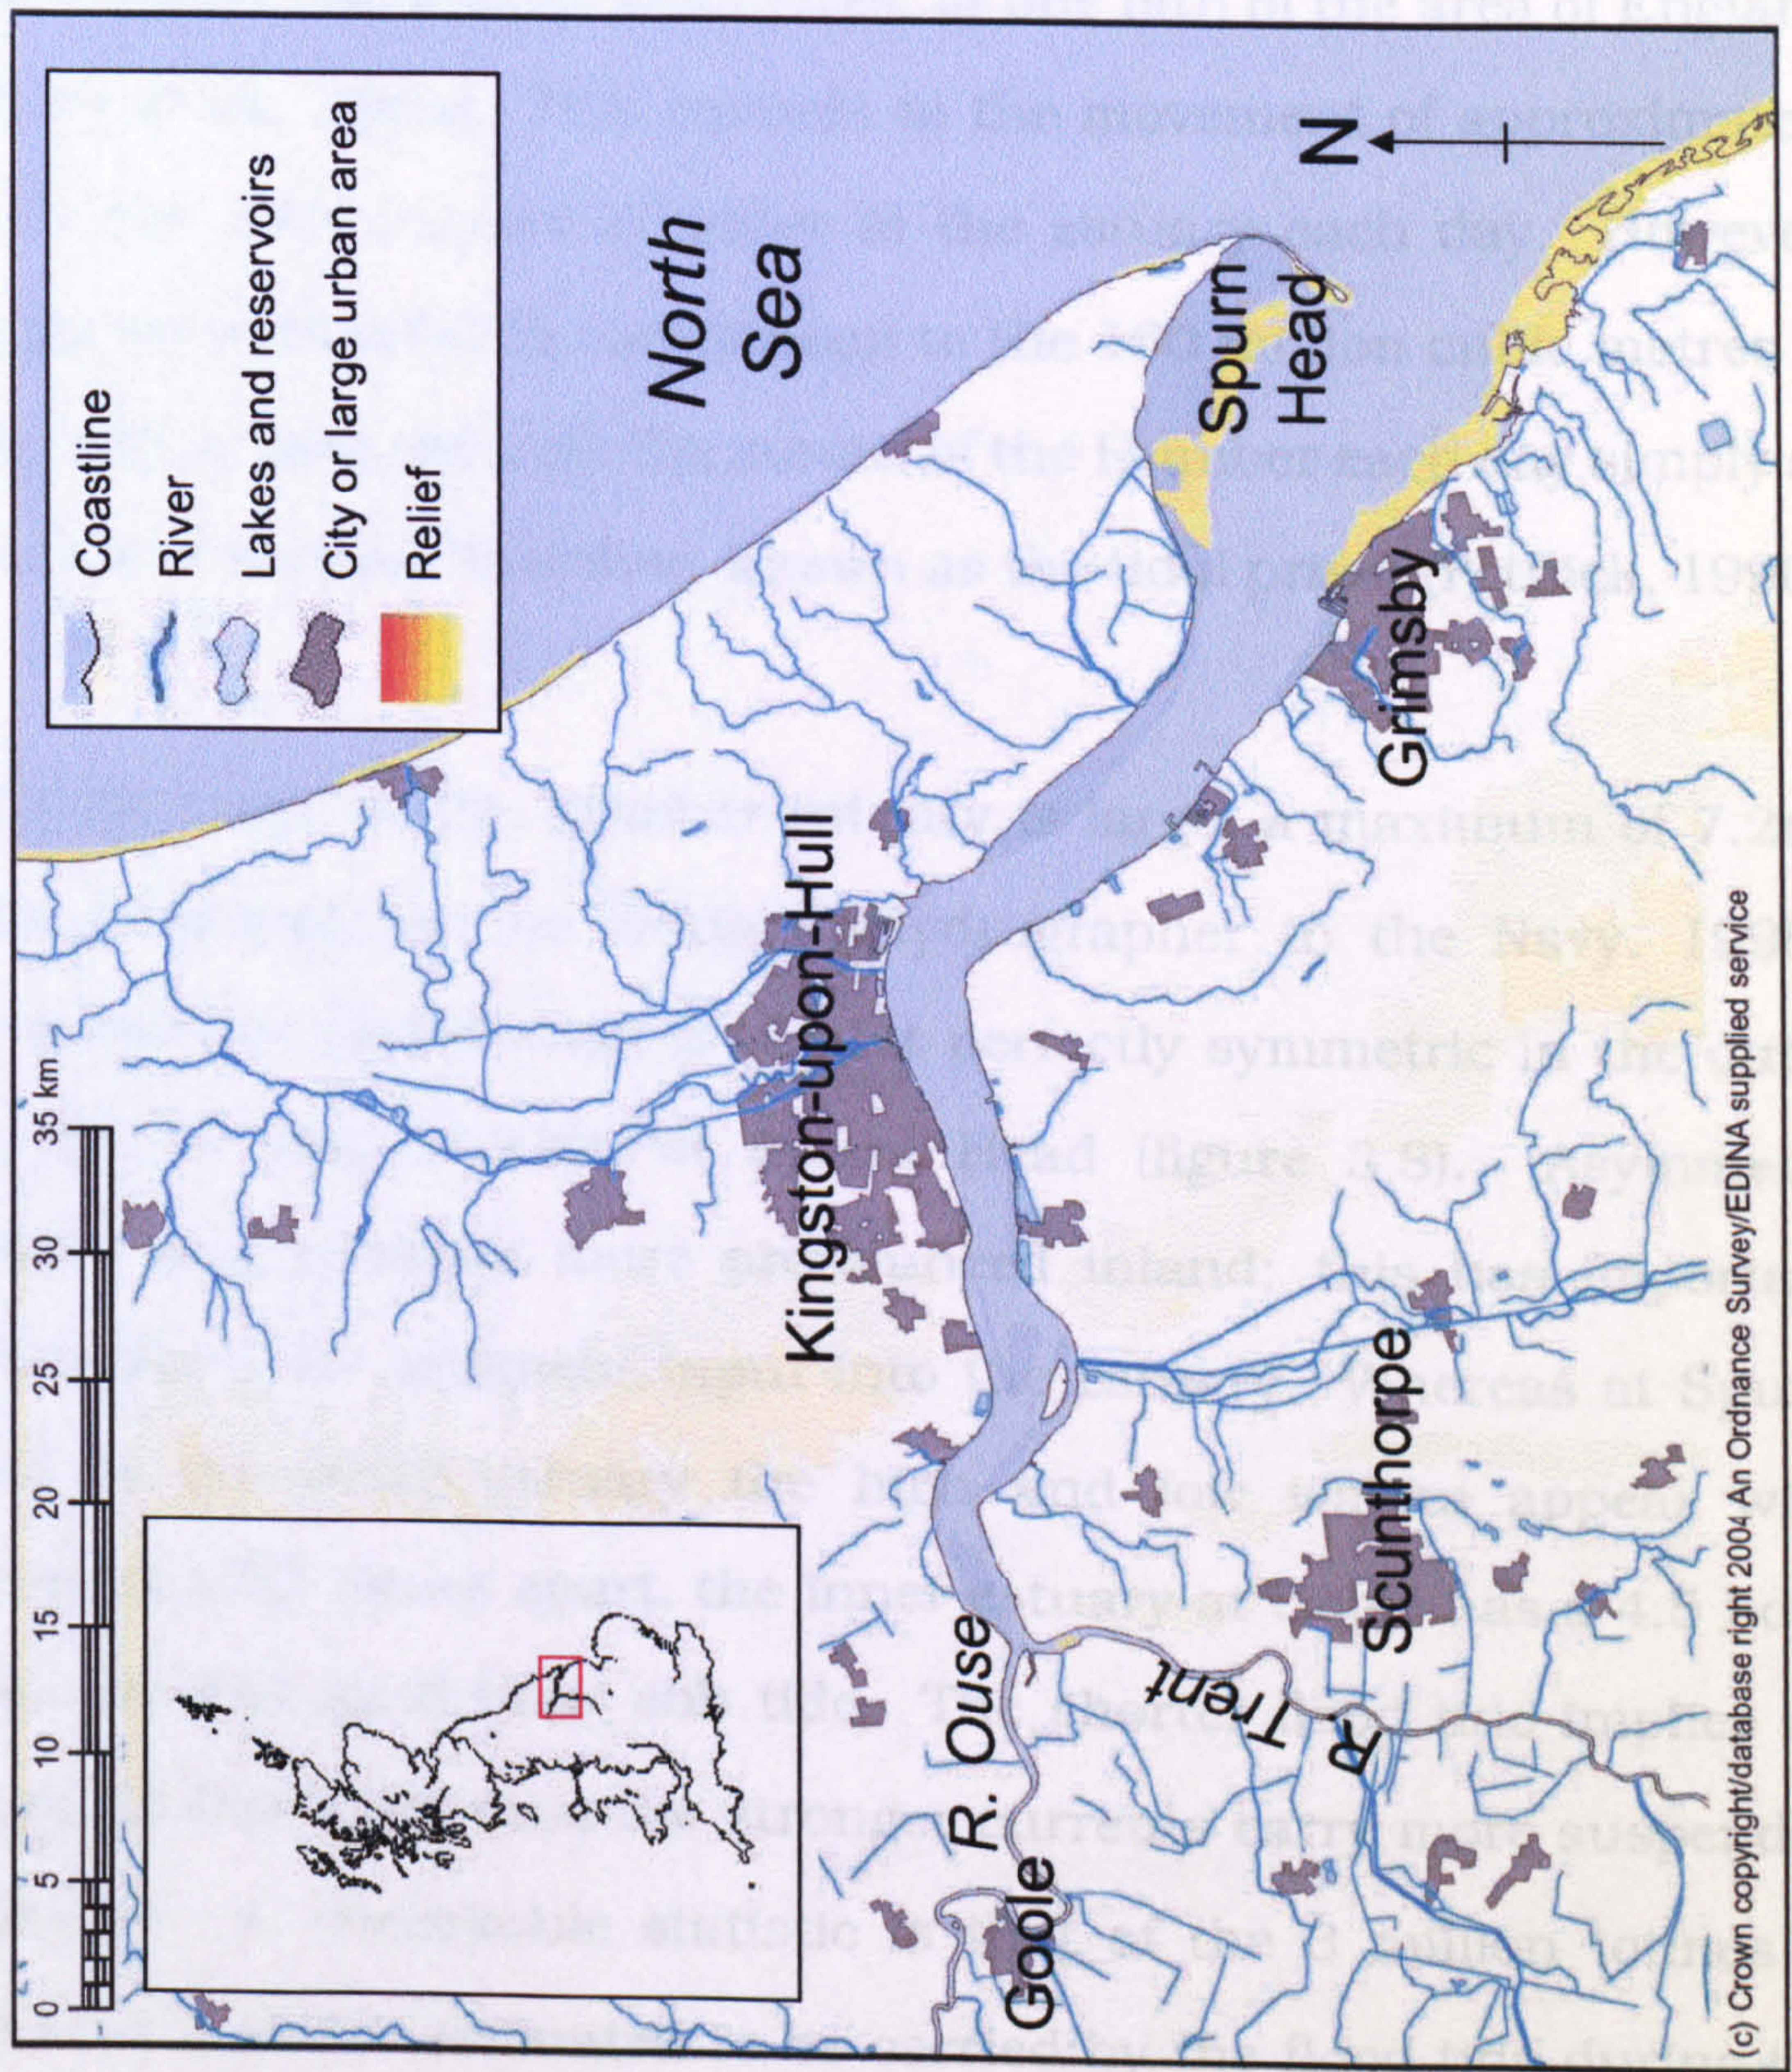


Figure 3.1: The Humber estuary.

One of the most remarkable facts about the Humber estuary is its size. Although the maximum width of 14km and length of 140km show it to be large, its true measure is that of the catchment (shown in figure 3.2), some 24,240 square kilometres, or one fifth of the area of England (Oguchi *et al.*, 2000). This equates to the movement of approximately 13 million cubic metres of water to the estuary each day. However, this figure is dwarfed in comparison to the 160 million cubic metres of water which pass through the mouth of the Humber each day simply as a result of natural tidal flow, known as the tidal prism (Pethick, 1990).

The tidal range in the Humber estuary is large, a maximum of 7.2m, varying throughout the estuary (Hydrographer to the Navy, 1995). The frequency distribution is almost perfectly symmetric in the outer estuary, as can be seen at Spurn Head (figure 3.3). Asymmetry appears and becomes more pronounced inland; this has important ramifications for sediment input into the estuary. Whereas at Spurn Head in the outer estuary the high and low waters appear with regularity 6.25 hours apart, the inner estuary at Goole has a 4.5 hour flood tide and an 8 hour ebb tide. The shorter flood tide implies an increased flow rate, thus the stronger currents carry more suspended sediment. A remarkable statistic is that of the 3 million tonnes of sediment that are estimated to be carried by the flood tide during the winter at any one moment in time (Pethick, 1990). Clearly, such an asymmetry has caused a huge net sediment input into the Humber over the course of the Holocene, causing a narrowing of the channel widths. Wilkinson *et al.* (1973) estimate that the decreasing channel widths also cause the net sediment gain per year to fall. This changing bathymetry

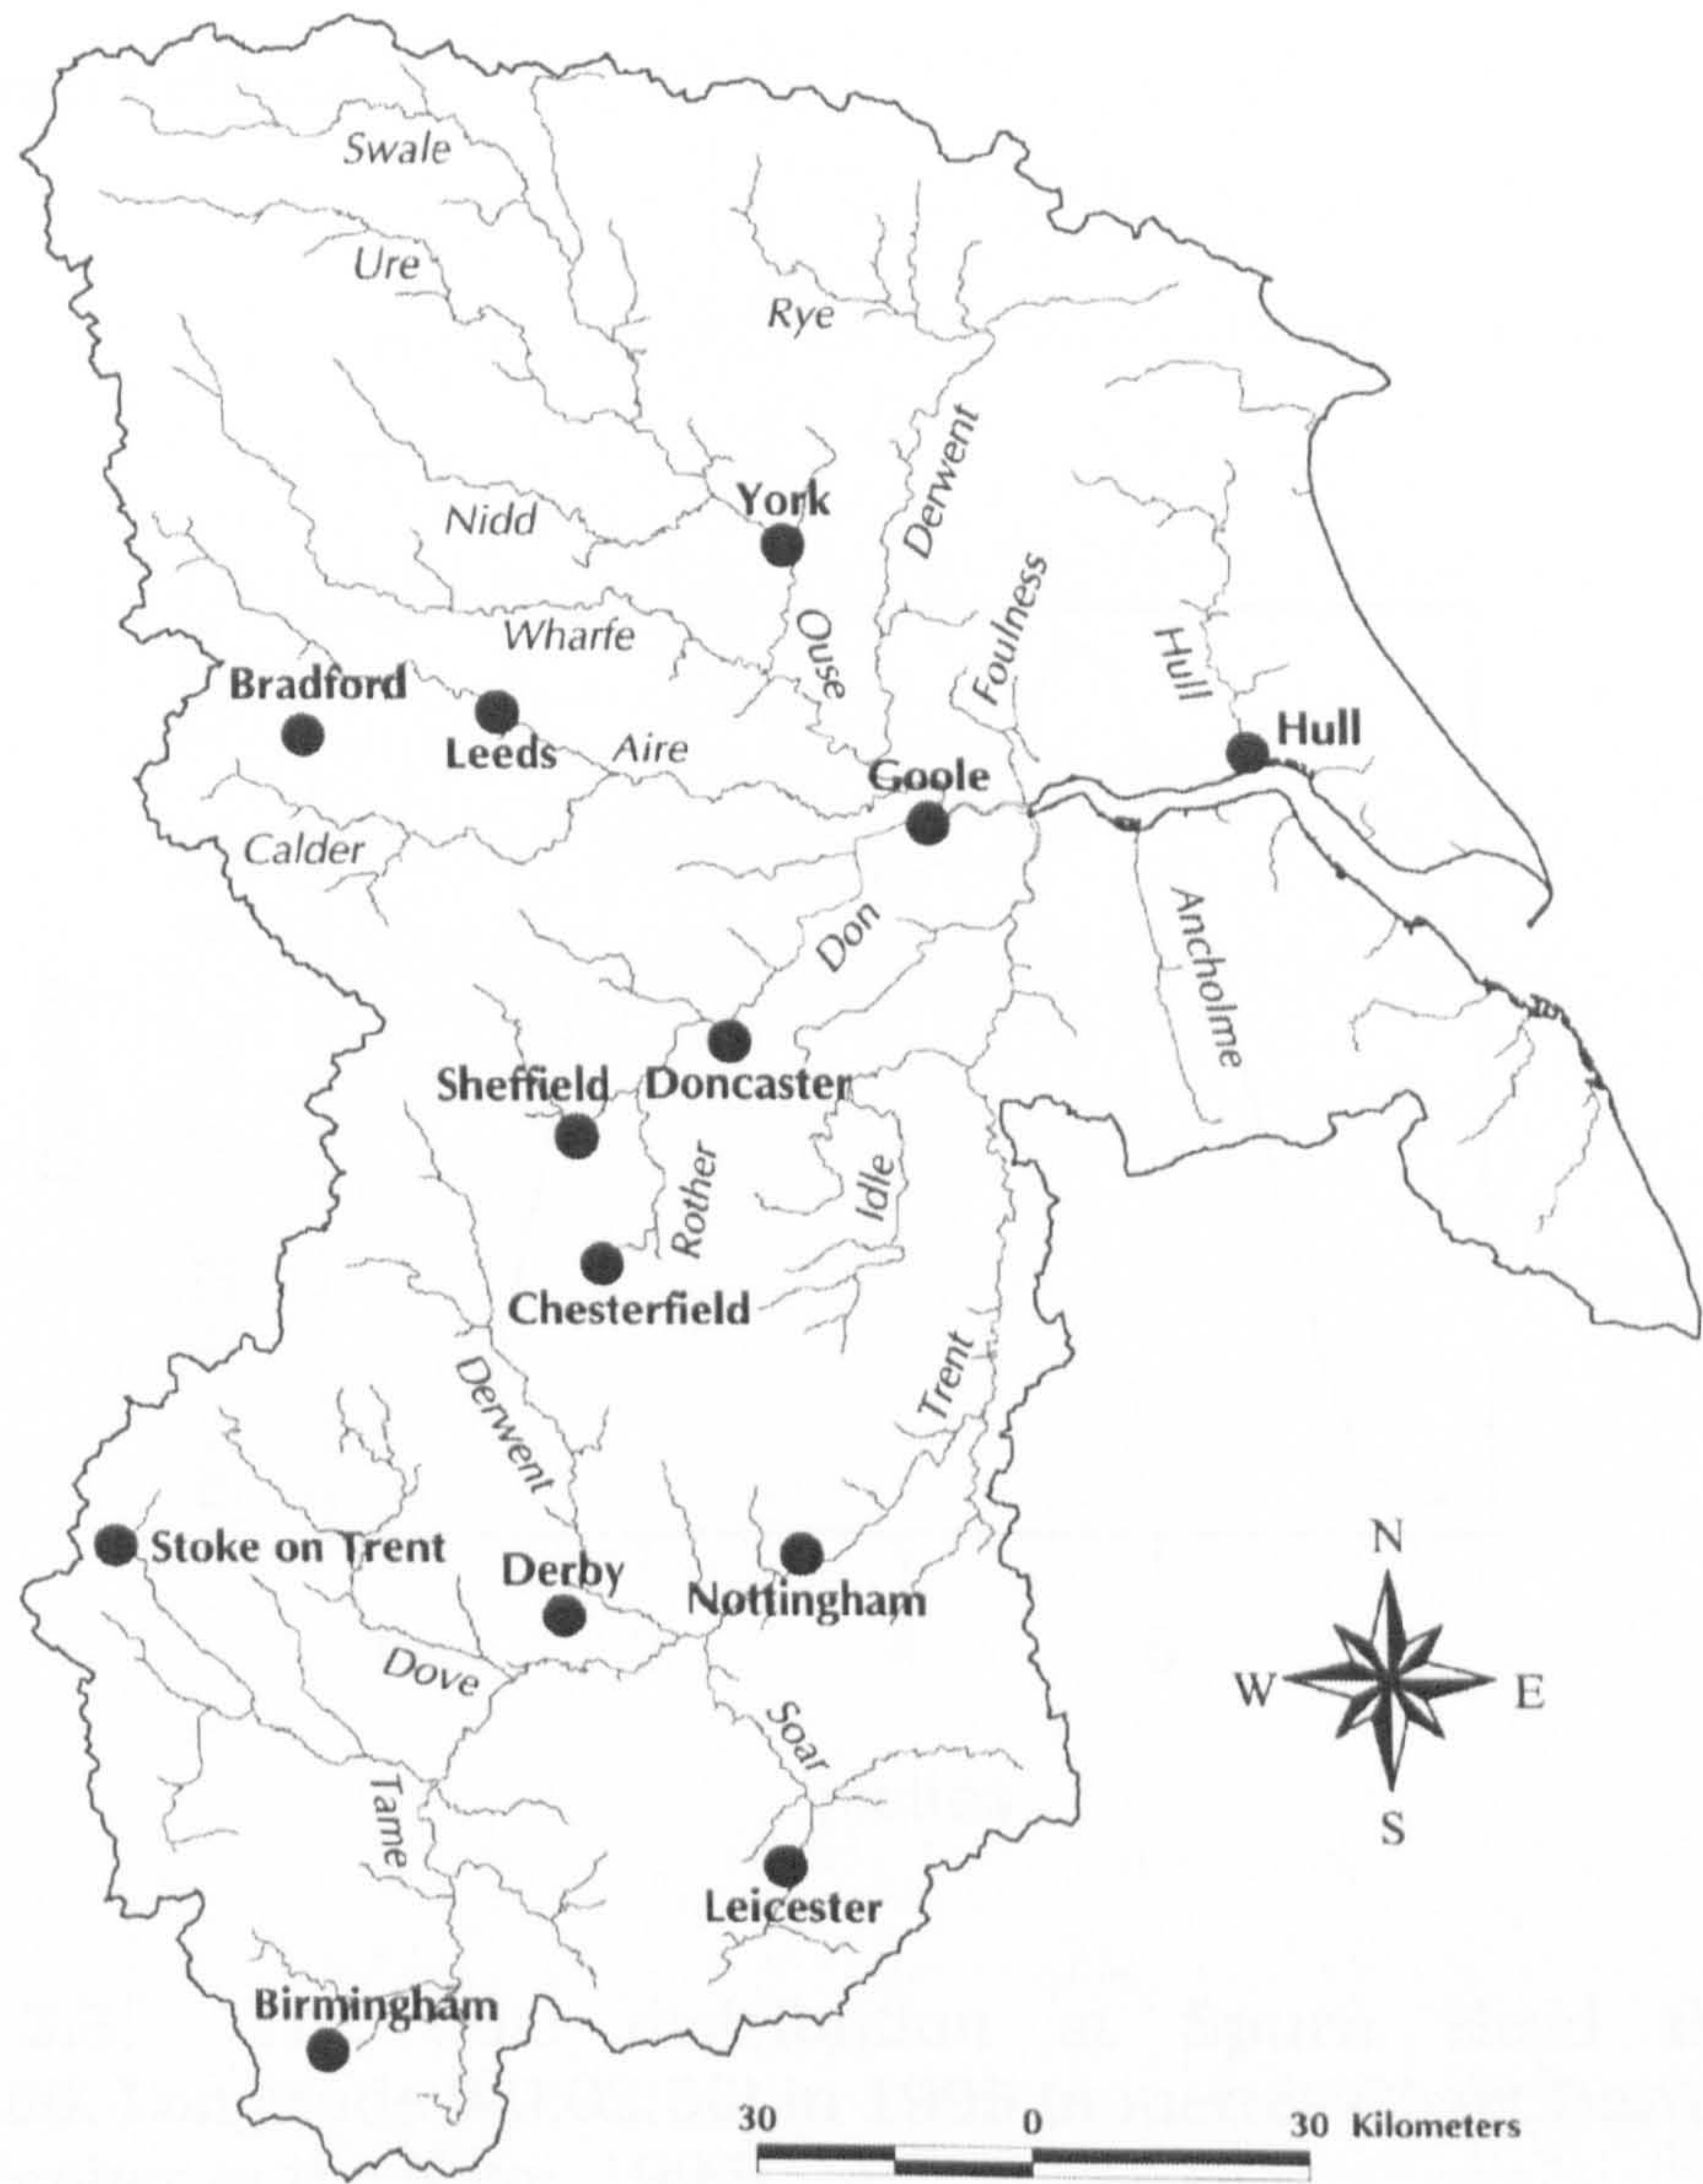


Figure 3.2: The catchment of the Humber estuary (from Oguchi *et al.*, 2000).

over the course of the Holocene would have had a profound effect on the frequency distribution of the tides in the Humber. Primarily, this effect would have shown itself through changing bottom friction and shallow water effects.




Figure 3.3: The tidal distribution at Spurn Head (Latitude N53:33:50, Longitude W0:02:56) in 1995 in metres Chart Datum (from Hydrographer to the Navy, 1995).

The pro-glacial Lake Humber emerged during the Dimlington Stadial (c. 18,000-13,000 BP) and drained around 15,000 BP with the onset of de-glaciation and the infilling of sediment (Gaunt, 1994). It was from this period onwards that the present Humber began to take shape, though there is evidence that it existed in some form during

the Ipswichian period some 120,000 years ago (Pethick, 1990). As the Quaternary period ended and the Holocene began, the picture in the Humber was one of rapid ice-melting and, therefore, sea-level rise (eg Metcalfe *et al.*, 2000). The massive discharge in meltwater would have altered the routes of the Humber tributaries, especially the Ouse and the Trent as their ability to cut through glacial till increased. The sea-level rise, coupled with the tidal prism, would have seen a huge increase in net sediment input to the estuary (Gaunt and Tooley, 1974). On the northern shore of the estuary, the removal of ice was temporarily halted causing the Humber to detour southwards, also leaving behind moraine. This created the marked bend that still appears today (Pethick, 1990).

Evidence of Human settlement in the Humber is well documented in the Humber Wetlands Project (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001) and could possibly have gone back as far as the Mesolithic (eg at Weel, section 3.6.3). Other early finds in the Humber concern timber structures like that at North Ferriby and Brigg, salt production sites like those at Ingoldmells and pottery finds such as those at Marshchapel. Many further sites have been found that contain links to Roman times. The Humber provided a major route to the North East and specifically to York, through Ermine Street. Sections of Roman road have been uncovered and dated precisely, eg South Ferriby. Other sites have given up Roman pottery and coins such as that at Adlingfleet.

More recently the Humber has become a flourishing port and business

centre. The city of Kingston-Upon-Hull now contains almost 250,000 people (ONS, 2002); the town of Grimsby to the south over 150,000 (ONS, 2004). The Humber Bridge was opened in 1981, providing a link between Barton-upon-Humber and south-west Hull. At present, the catchment area of the Humber contains approximately 20% of the population of the UK, around 12 million people (Jarvie *et al.*, 1997).

The future of the Humber estuary will depend very much on the nature of sea-level rise. At present, estimates range from 4.4mm per year to 27.6mm per year, a total rise of between 0.55m to 3.45m by 2100AD (Pethick, 1990). The impact of such a rise on the landscape needs to be balanced with analysis of the quantity of sediment the rise would bring with it. However, there is little doubt that even a modest rise would cause problems for the region, being comprised of such low-lying, flat areas.

3.3 The LOIS database

The LOIS project ran for 7 years and studied the complex workings of air, land and sea systems (Shennan and Andrews, 2000). The LOEPS part of the LOIS study concentrated on the east coast of England. At present the LOIS database contains approximately 2000 radiocarbon dates, the majority of which are sea-level index points. The main fields of the database are:

- Location
- Height of sample and error

- Thickness of sample
- The reference water level and indicative range (see section 2.3.3)
- Type of material dated
- Tidal information (usually modern)
- Dating information
- Reject code (determines the validity of an index point)

To create the data set used in this thesis, only those data points used by Shennan and Horton (2002) were used. They had been calculated following the methods developed during the International Global Correlation Programme (IGCP) projects 61, 200 and 495; outlined in, for example, van de Plassche (1986). The radiocarbon dates were re-calibrated with the latest calibration curve available (Reimer *et al.*, 2004), using the method described in section 4.2.1. The index point elevations had been corrected as outlined in sections 2.3.3 and 2.4. They had also been adjusted using the most recent tidal models run for the Humber, as outlined in Shennan *et al.* (2000a, 2003). The data set had originally been split into two parts, the inner and outer Humber, as there was felt to be a marked difference between the two in terms of tidal distribution. In theory, the tidal modelling should have removed all such differences. Any that remain are the product mainly of isostatic change.

Some data points exist in both the LOEPS and Humber Wetland Project (HWP) databases. The HWP team undertook a number of cores whilst

studying the palaeo-environment of the Humber. A number of these cores contained data which was converted into sea-level index points (Shennan and Horton, 2002). These then found their way into the LOIS database. None of these data points contain any archaeological information.

In total, 102 data points exist for Humber estuary. Of these, 22 are basal index points, 42 are intercalated index points, and 38 are limiting points. Appendix A contains a table of the data.

3.4 Archaeological data

The data created from the Humber Wetlands Project (HWP) is in a similar format. Most of these points are limiting points, but there are some which could be considered index points. Here, the only details given are:

Location Usually a grid square reference.

Age Either as an archaeological period or, when available, an uncalibrated radiocarbon estimate.

Artefact name Occasionally gives details as to the artefact's composition eg a boat, fish trap, etc.

Height Measured to O.D. during the excavation of the site.

These points need to be given the same rigorous treatment as those in the LOIS database. Thus the following need to be collected:

- A reference water level. This can be obtained through expert information from archaeologists. The method for calculating the RWL is given below in section 3.5.
- A tidal elevation estimate for the age/location of the data point in order to calculate the reference water level numerically.
- An estimate of all the appropriate errors involved. This may include levelling errors, tidal estimation errors, etc as discussed in section 2.4.2.

Initially, all 248 sites identified through the HWP were considered for inclusion. A large number (approx. 88%) of these were found via field-walking and were not levelled or dated with any particular accuracy. Their elevations were ascertained from maps but were not found to be of sufficient precision to form a reasonable bound on RSL. Of the remaining 30 points, a further 15 had been dated either through radiocarbon or with sufficient knowledge as to obtain an error range of less than 500 years. However, these data points had not been excavated and so, again, were discarded due to lack of information about their elevations. Finally, 15 data points were deemed to be of sufficient age/elevation quality as to be used as part of the data set. Their names are derived from the locations in which they were found, and from their titles in the HWP set of books (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001). A plot of the locations of the archaeological data points is shown in figure 3.4. A table of the archaeological data points is given in appendix B. The big advantage of the archaeological data is that it corresponds to periods where the geological data is sparse (as

discussed in section 2.3.2).

3.5 The calculation of reference water levels for archaeological data points

A reference water level, as previously defined, is the contemporary tide level represented by a certain stratigraphic sequence. The construction of a RWL is a two-step process. Firstly, the tide level must be obtained in an algebraic form, eg. $(MHWS+HAT)/2$. Secondly, this value must be calculated numerically. Traditionally this has been done by substituting modern tidal data into the formula. However, the tidal distribution of the Humber has changed dramatically, as mentioned in section 2.3.10. For this project, the tidal range model created by Shennan *et al.* (2000a) is used to obtain reference water levels. This model gives estimates of MHWS, MHWN, MSL, MLWS and MLWN in metres O.D. in 1000 year time slices for the period 1000 years before present to 9000 years before present across the Humber estuary. It has a grid resolution of one square per 60th of a degree of longitude and one 90th of a degree latitude. It has an estimated standard error of 22cm.

The archaeological data points used in this thesis are treated as intercalated index points (IIPs). This is because little is known about the stratigraphic record from which they were obtained. The reference water levels for these data points were obtained from expert archaeologists (van der Noort, pers comm.). Unlike index points created through the LOIS project, there is no reason to believe that identical

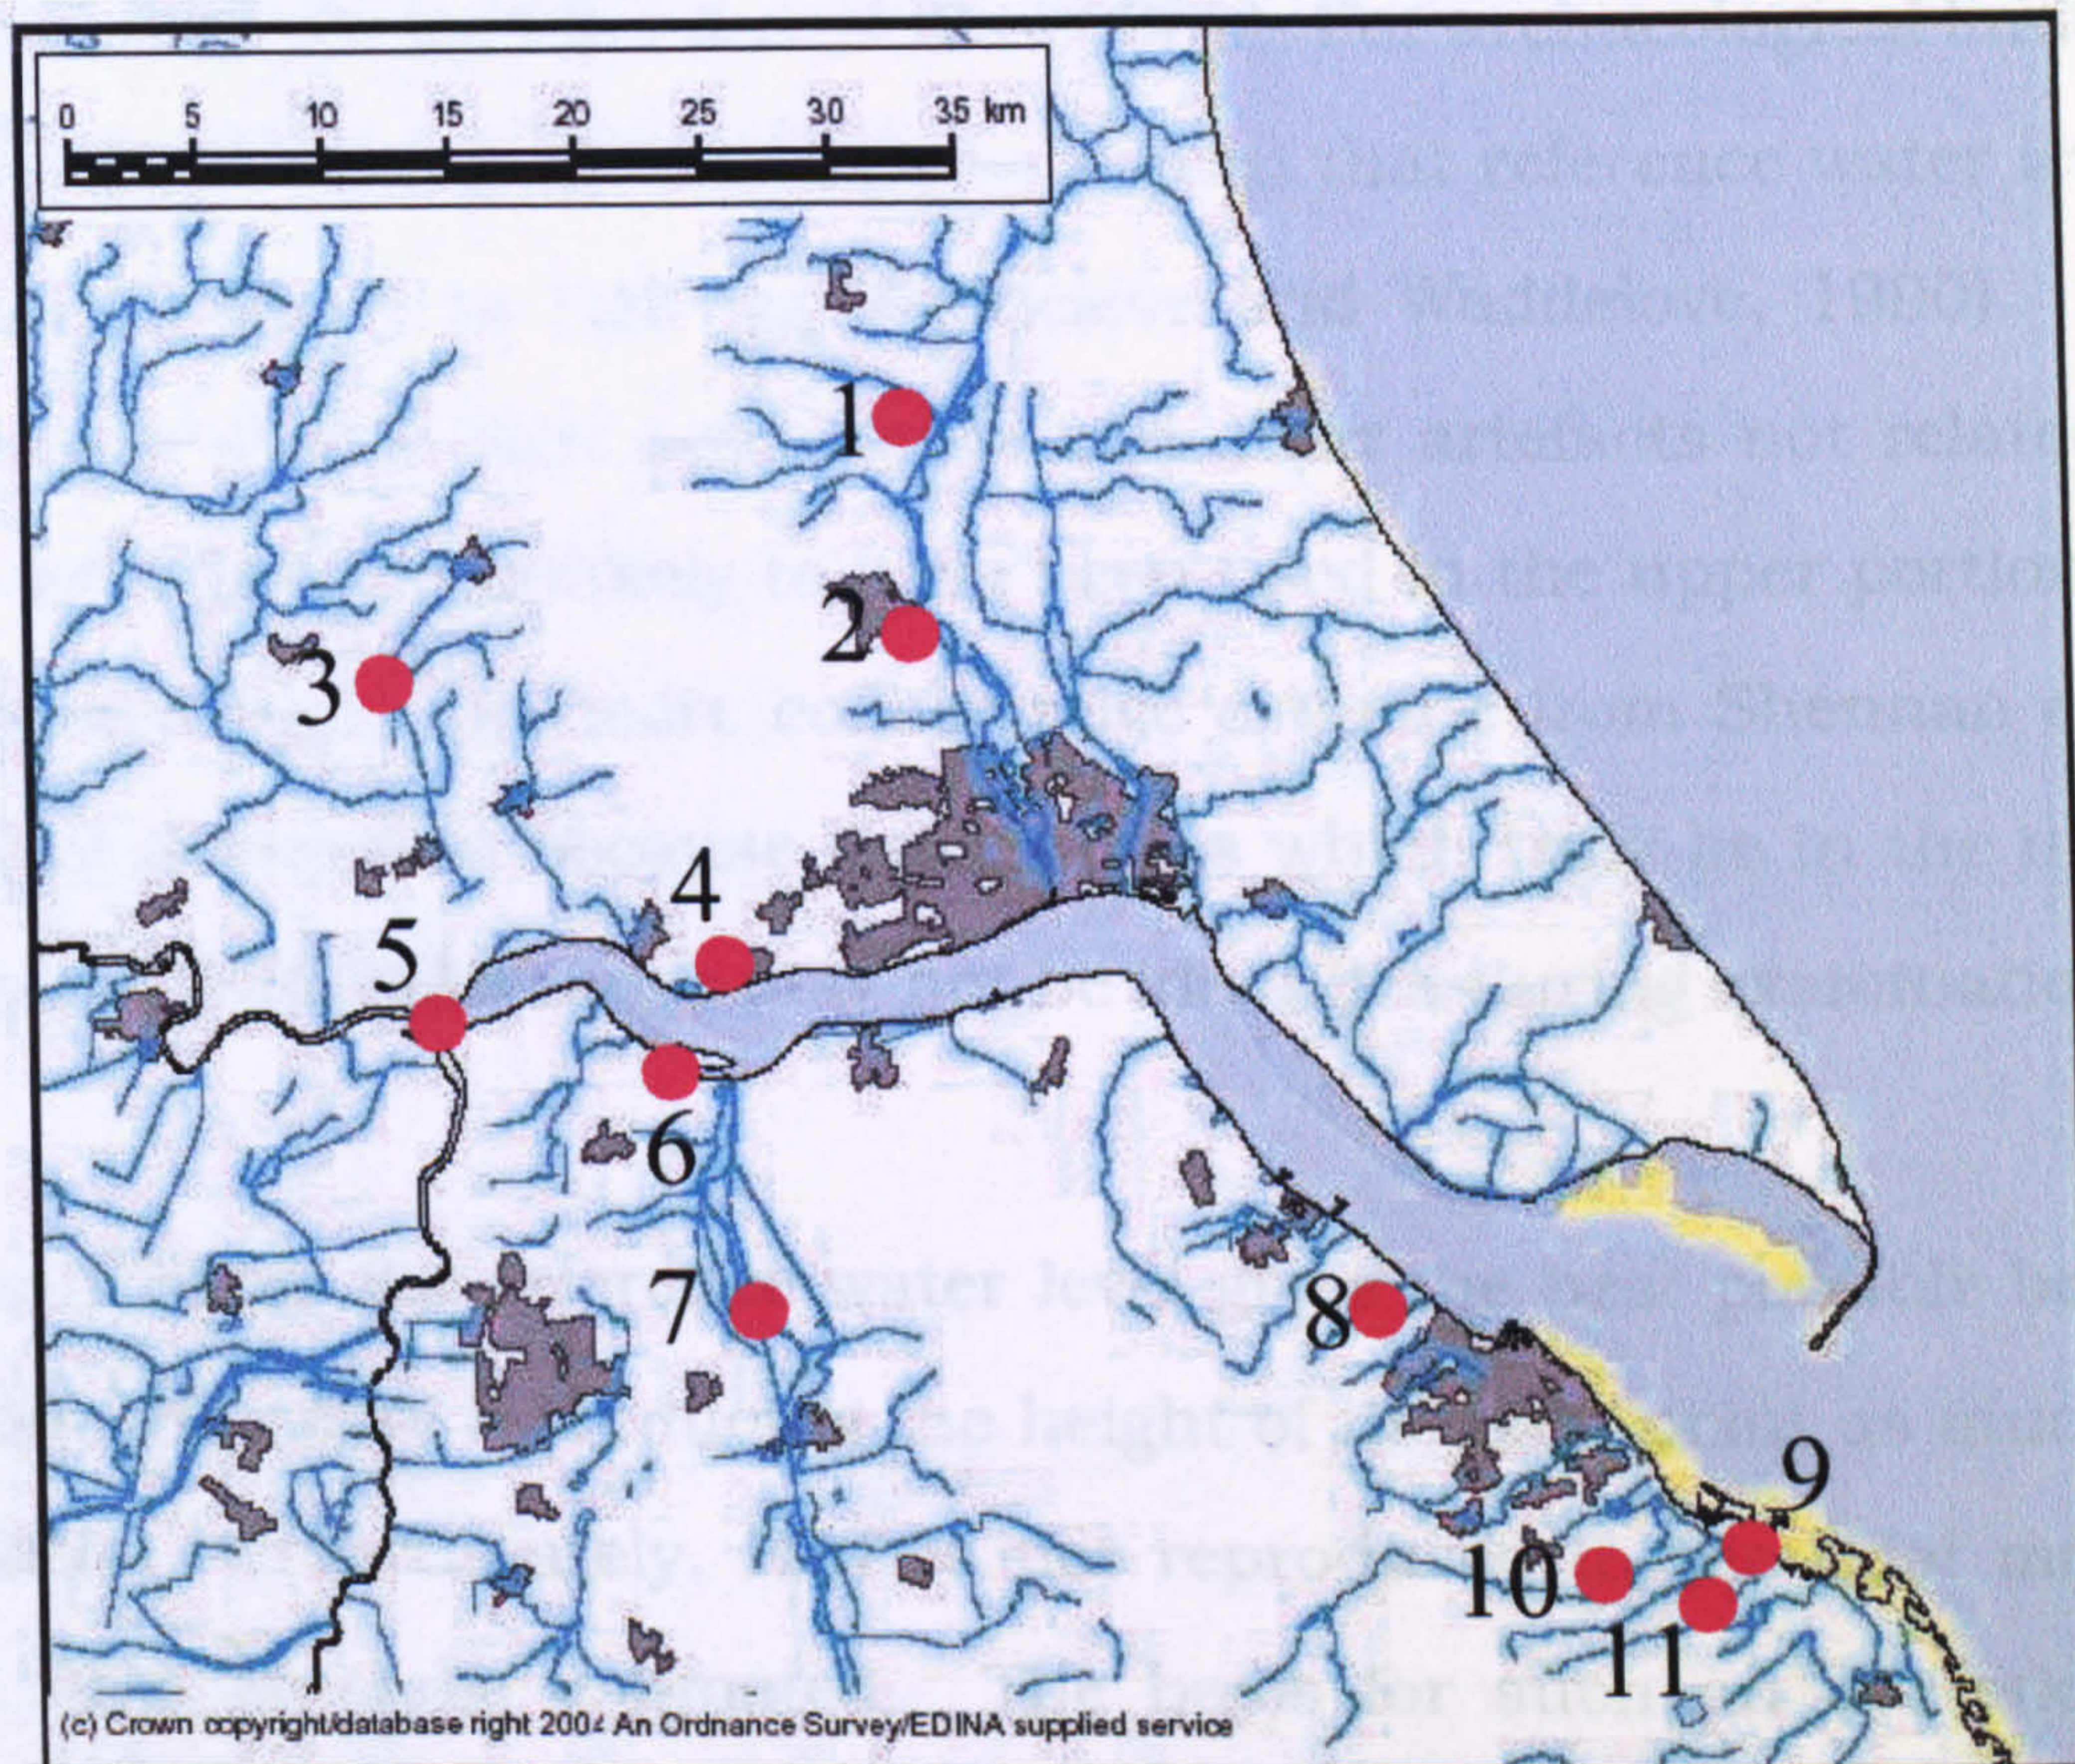


Figure 3.4: Locations of the archaeological data points around the Humber. 1 = Brigham-6, 2 = Weel-2, 3 = Wheldrake-2, 4 = Ferriby-2, 5 = Adlingfleet-2, 6 = South Ferriby-2, 7 = Brigg Raft/Trackway/Boat, 8 = Marsh Chapel-2, 9 = Ingoldmells-12, 10 = Willoughby-13, 11 = Ingoldmells-6.

artefacts will have identical indicative meanings. For this reason, the indicative meaning of each archaeological artefact is discussed individually.

In the LOIS database, limiting points are given a reference water level of $(MHWS+MTL)/2$ (eg Shennan *et al.*, 2000b). For archaeological limiting points however, it could strongly be argued that reference water levels should be closer to HAT (eg Waddelove and Waddelove, 1990). The reason for this is that settlements and other artefacts not related to tidal activities are unlikely to have been used in the upper portions of the tidal range. The more conservative estimate from Shennan *et al.* (2000b) is required because index points which truly lie in the upper portions of the tidal range may not be identified during examination.

Using HAT as the reference water level gives the best possible bound on the RSL curve by reducing the height of the data point as much as possible. Unfortunately, HAT is not reproduced in the tidal models and thus must be estimated. The basis for such an operation is the assumption that the *ratio* of HAT to MHWS is constant over time (though not location). If this assumption is accepted then modern tidal data can be used to obtain the ratio of HAT to MHWS throughout the Humber estuary. Estimates can then be multiplied by the tidal model estimates of MHWS to create an estimate of HAT with an associated error.

The details as to the estimation of HAT for the archaeological locations

are provided in appendix C. A table of the results of this analysis are shown in table 3.1.

Site	Estimated HAT (m)	se(HAT) (m)
Weel-2	2.43	0.41
Marsh Chapel-2	3.16	0.4
Willoughby-13	1.84	0.33
Adlingfleet-2	2.05	0.4
South Ferriby-2	1.75	0.38
Wheldrake-2	2.18	0.42
Kelk-6	3.26	0.44
Brigham-6	3.27	0.33
Ingoldmells-6	2.23	0.38
Ingoldmells-12	2.22	0.38
Ferriby-2	1.76	0.38
Brigg raft	1.74	0.37
Brigg trackway	1.75	0.38
Brigg boat	1.74	0.37

Table 3.1: Table of estimated HAT with predictive standard error for archaeological sites.

3.6 Outline of archaeological data

This section contains a list and discussion of the archaeological data points in turn, taken from the HWP (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001). A table of the data can be found in appendix B

3.6.1 Wetland Heritage of the Ancholme and Lower Trent valleys

Adlingfleet-2

Adlingfleet-2 (figure 3.5) was created after an excavation and geophysical survey in the inner estuary on the west side of the River Trent following on from the discovery of a large scatter of Roman pottery fragments. The archaeologists found a number of Roman-period activities, including the production of salt and pottery. This appears to be part of a much larger settlement which may have lain on the river-side and acted as a port. However, evidence of river transport was not found.

Three trenches were dug to assess the settlement. The lowest of which had a ground surface of 1.68m O.D. Roman pottery, fired clay, charcoal and bone was found approximately 0.8m below this height. Unfortunately, the pottery found was insufficient for dating purposes. In summary, this data point can be used as a limiting point at 0.878m O.D. during the Roman period.

The RWL at this site is 2.05m O.D., thus RSL was given as -1.17m O.D.. It was ascribed a levelling standard error of 2.5cm (corresponding to the standard values given in the LOIS database), a thickness standard error of 10cm (accounting for the uncertainty in the exact depth at which pottery was found) and a tide level standard error of 40.1cm. The total estimated standard error for this data point was 41.4cm.

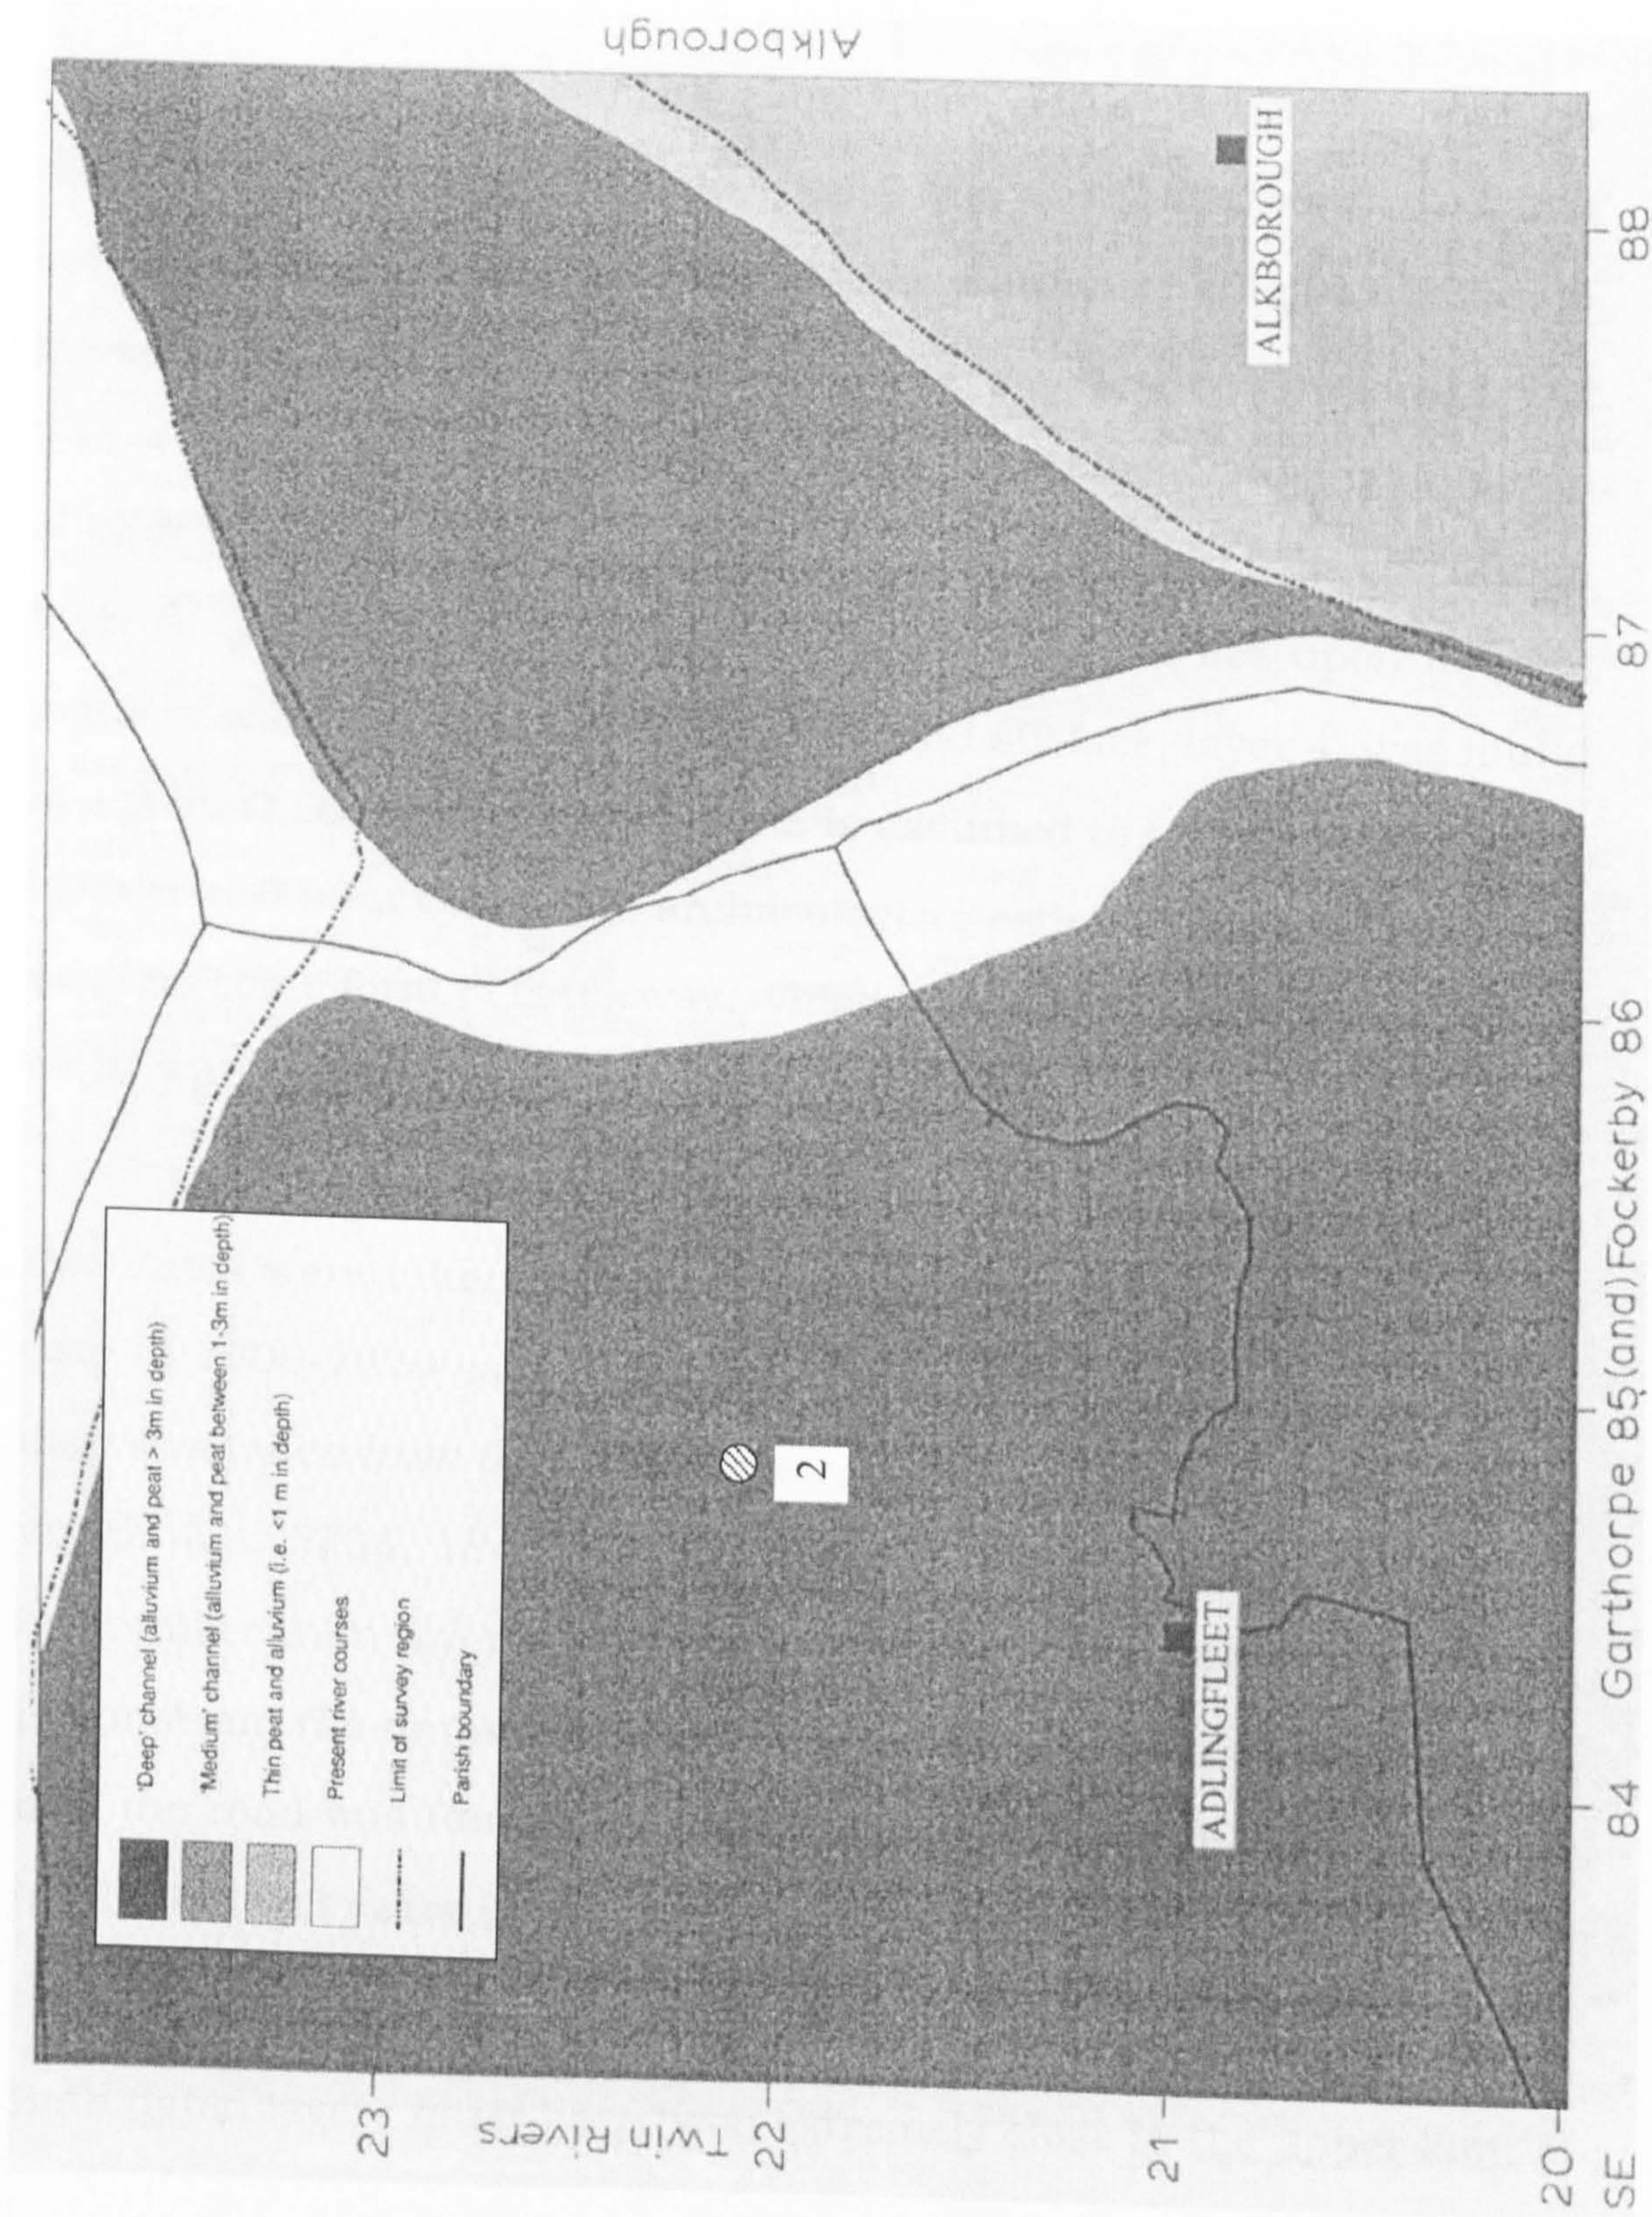


Figure 3.5: Site map of Adlingfleet-2 (hatched circle). From van der Noort and Ellis (1998).

South Ferriby-2

This site, shown in figure 3.6, was the location of a 50 metre length of Roman Road on the southern foreshore of the inner estuary. It is thought to have been built to allow travel across a marshy area. A small part of the road became visible due to rapid erosion and was subject to an excavation in 1996. The excavation revealed a number of layers to the road, some of which could be dated, the uppermost being named layers 4, 5 and 6. Layer 4 is the road surface, being made up of large limestone stones with flint and quartz pebbles. This rests upon layer 5, a thick layer of brushwood, which in turn lies upon layer 6, made of small pebbles and flints. The road surface, layer 4, was found at 1.37m O.D., the height of layer 5 is assumed to be 1.18m O.D., with layer 6 at 0.98m O.D.. The archaeologists estimate that the road was used as some form of causeway, which may have had to be overtopped when water levels rose.

Two dates were taken; one from the brushwood layer 5 which gave a date of 1906-1820 cal yrs BP (GU-5705, 1960 ± 50 yrs BP). A second date was taken from the road surface, layer 4. This gave 1880-1690 cal yrs BP (GU-5704, 1840 ± 60). However, it is known that layer 6 must be younger than layer 5, so the stratigraphic relationship can be used to constrain the dates. Additional precision can be obtained by noting that the road was Roman, so must have been built after the conquest in AD47 (1903 years BP).

Both dated layers must have been extremely close to the upper limit of

the tidal range. The archaeologists noted that the road was overtopped with sediments being deposited for some time. It is unclear whether this was during the use of the road or not. The two data points here are treated as limiting points, though an argument could be made for assigning a reference water level very close to HAT.

Layer 4 (the upper surface) was found at 1.37m O.D. whilst layer 5 was found at 1.18m O.D.. HAT was estimated at this site as 1.75m OD giving final RSL of -0.57m and -0.38m. The associated tidal standard error was estimated at 37.8cm. Levelling standard errors were ascribed at 2.5cm. The road surface layer was found to be 6cm thick whilst the brushwood layer was 60cm thick. These are approximated as 95% intervals for thickness error, thus standard errors of approximately 1.5cm and 15cm were used. The total standard error for the road surface was 38cm; for the brushwood layer, 40.8cm.



Figure 3.6: Site map of South Ferriby-2 (hatched circle). From van der Noort and Ellis (1998).

The Brigg raft, trackway and boat

The three Brigg finds (figure 3.7) are arguably the most important Bronze age discoveries in the Humber. All three were found near the west bank of the River Ancholme between 1884-1887. The raft was said to be 40 feet long and made from a number of wooden strakes, found between 0.8 and 1.03m O.D. It was possibly Bronze Age and made to rise and fall with the tide. An expert archaeologist has given an estimated reference water level and indicative range of within 0.5m of MTL (van der Noort, pers comm.). A number of radiocarbon dates were taken from the raft and were used to provide a detailed chronology of the boat (McGrail, 1981); an estimated date gives 2950-2350 cal yrs BP (a weighted average of Q-1255, 2655±50; Q-1256, 2605±50; Q-1258, 2670±50). The methods used for dating these samples (described in section 4.2.1) allows these three dates to contribute to the same calendar age.

The boat (also known as the dugout canoe) was made from a single piece oak log. It was located approximately 0.9m lower than the raft though the stratigraphic relationship cannot be confirmed due to the location of the deposits. The boat was dated to approximately 3209-2749 cal yrs BP (Q-78, 2784±100). It is also estimated to have lain within 0.5m of MSL. Finally, the trackway was buried 1.2-1.5m O.D. lower than the raft. This was dated to 2899-2299 cal yrs BP (Q-77, 2552±120 yrs BP). The trackway is considered to have lain within 0.2m of HAT.

Taking the objects in turn, the raft was given an elevation of 0.915m

O.D. As already mentioned, the expert archaeologists deemed this object to have lain within 0.5m of MTL with 90% probability. MTL at this location was estimated to be 35cm above MSL, thus the final RSL of the raft is +57cm. As this item was excavated some time ago, the levelling standard error was given as 10cm, larger than many of the other archaeological sea level data points. The thickness standard error here is based upon the range 0.8m to 1.03m over which wooden strakes were found. Again, the range is treated as a 95% interval which provides a standard error of 5.75cm. The tide level standard error for this data point was 22cm (as it came from the supplied tidal estimation model). As this data point is a sea-level index points as opposed to a limiting point, it also has an indicative range based on the 1m range with 90% probability given by the archaeologists. This contributes an extra standard error of 30.4cm. The total standard error for this data point is 39.3cm.

The boat or dugout canoe was given an elevation of 0.015m O.D. This artefact was also given the reference water level of MTL with an indicative range of ± 50 cm with 95% probability. MTL at this location was 0.35m O.D. providing a final elevation of -0.33m. The levelling and thickness standard error were calculated as above to be 10cm each. The tide level standard error was 22cm. The indicative range was again 30.4cm. The total standard error for this data point was 40.1cm.

The Brigg trackway was found at approximately -0.435m O.D. and given a reference water level of HAT with indicative range 20cm at 60% probability. HAT for this location was estimated to be 1.75m. The final

RSL estimate for this data point was -2.18m O.D. As before levelling, thickness and tide level standard errors were calculated as described for the raft. This gave 10cm, 7.5cm and 37.6cm respectively. The indicative range was 23.8cm. The total standard error for this data point was 46.3cm.

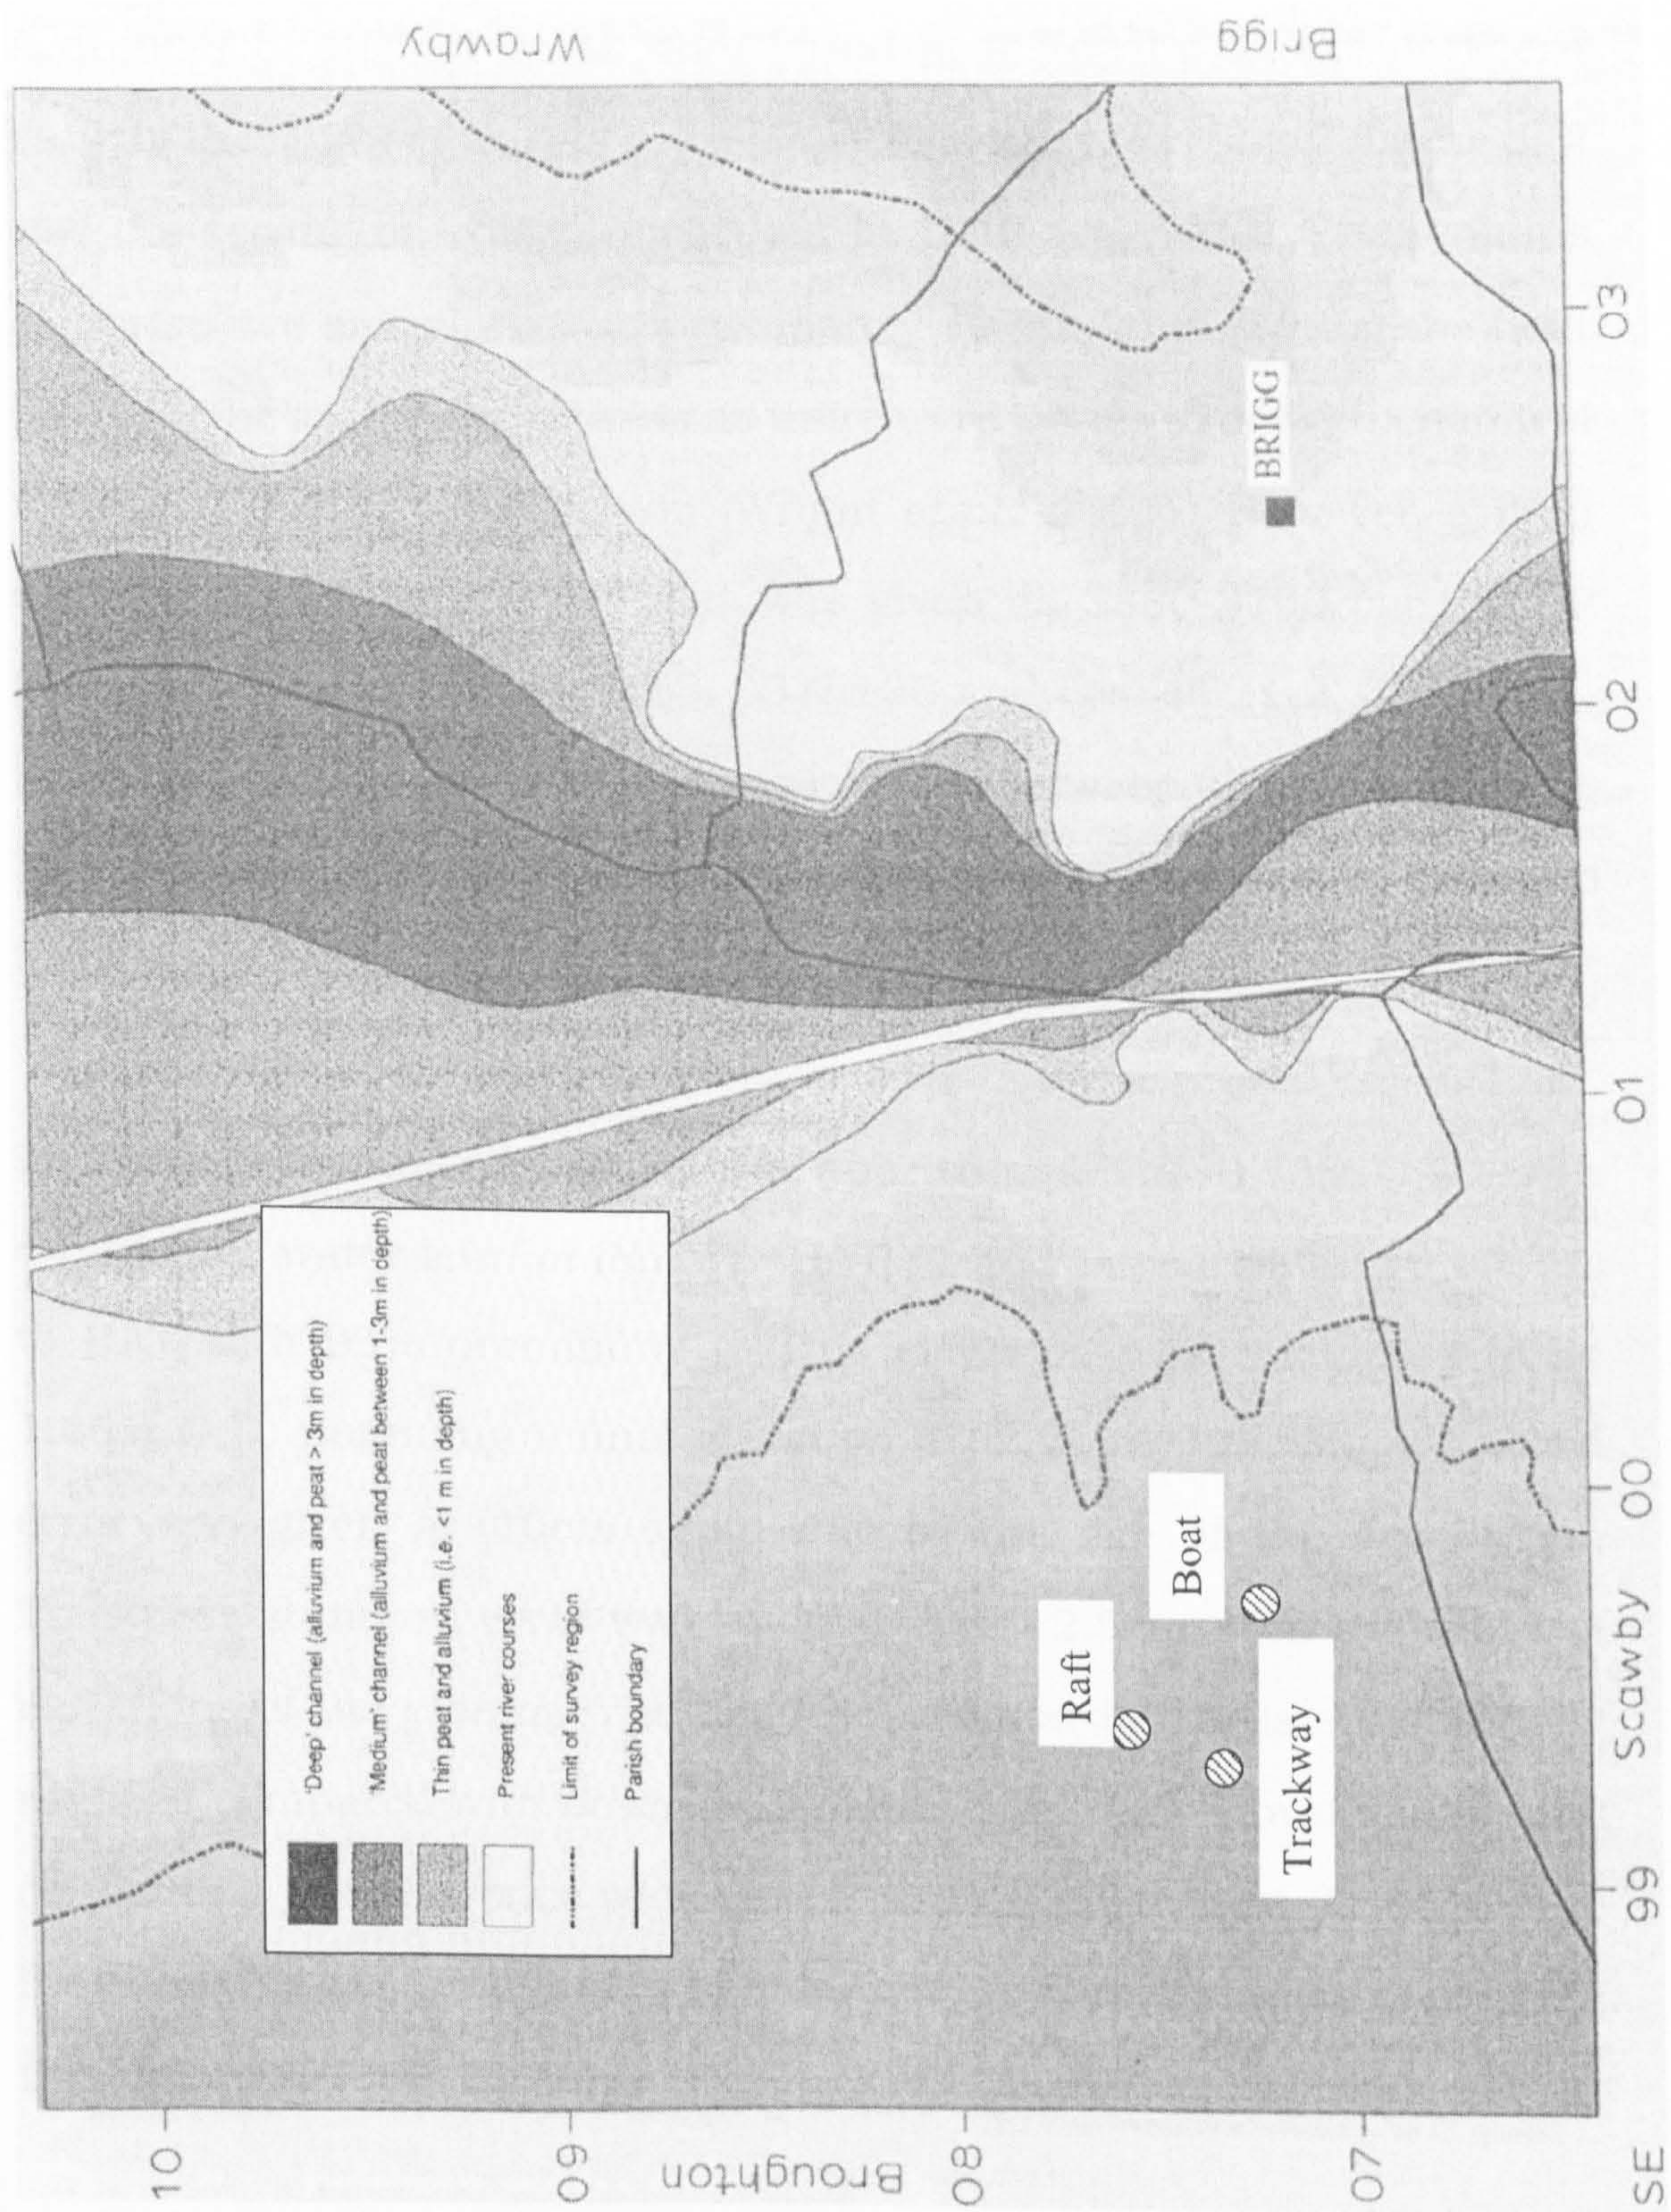


Figure 3.7: Site map of the Brigg boat, raft and trackway (hatched circles). From van der Noort and Ellis (1998).

3.6.2 Wetland Heritage of the Vale of York

Ferriby-2

A number of ancient boats were found near Ferriby on the northern bank of the Humber during the mid-20th century. Ferriby-2 (figure 3.8) was the second of which, unearthed in 1940 and 1946. The remains comprise two oak planks, approximately 11.4m long. Dating the boat proved to be a complex process as there were issues to do with chemical treatment during conservation (Wright *et al.*, 2001). However, a final estimate of the age of the boat was given as 3880-3700 cal yrs BP (a weighted mean of OxA-7458, 3515±30 and OxA-9521-2, 3510±38). Height measurements put the oak planks at between -0.5 and -1m O.D. The archaeologists have ascribed an indicative meaning of between MHWS and HAT.

The initial elevation of the boat was approximated at -0.75m O.D., with a reference water level of (MHWS+HAT)/2 and indicative range MHWS to HAT with 95% probability . This elevation was calculated to be 1.46m O.D. providing a final elevation of -2.21m. Levelling standard error was given as 10cm again due to the age of the excavation. Thickness standard error was calculated at 12.5cm from the range of elevations of the planks. The tide level standard error is substantially reduced compared with other data points due to distributional results concerning the reference water level. This was given at 21.8cm. The indicative range came out to be 15.7cm. The total standard error for this data point was 31.3cm.

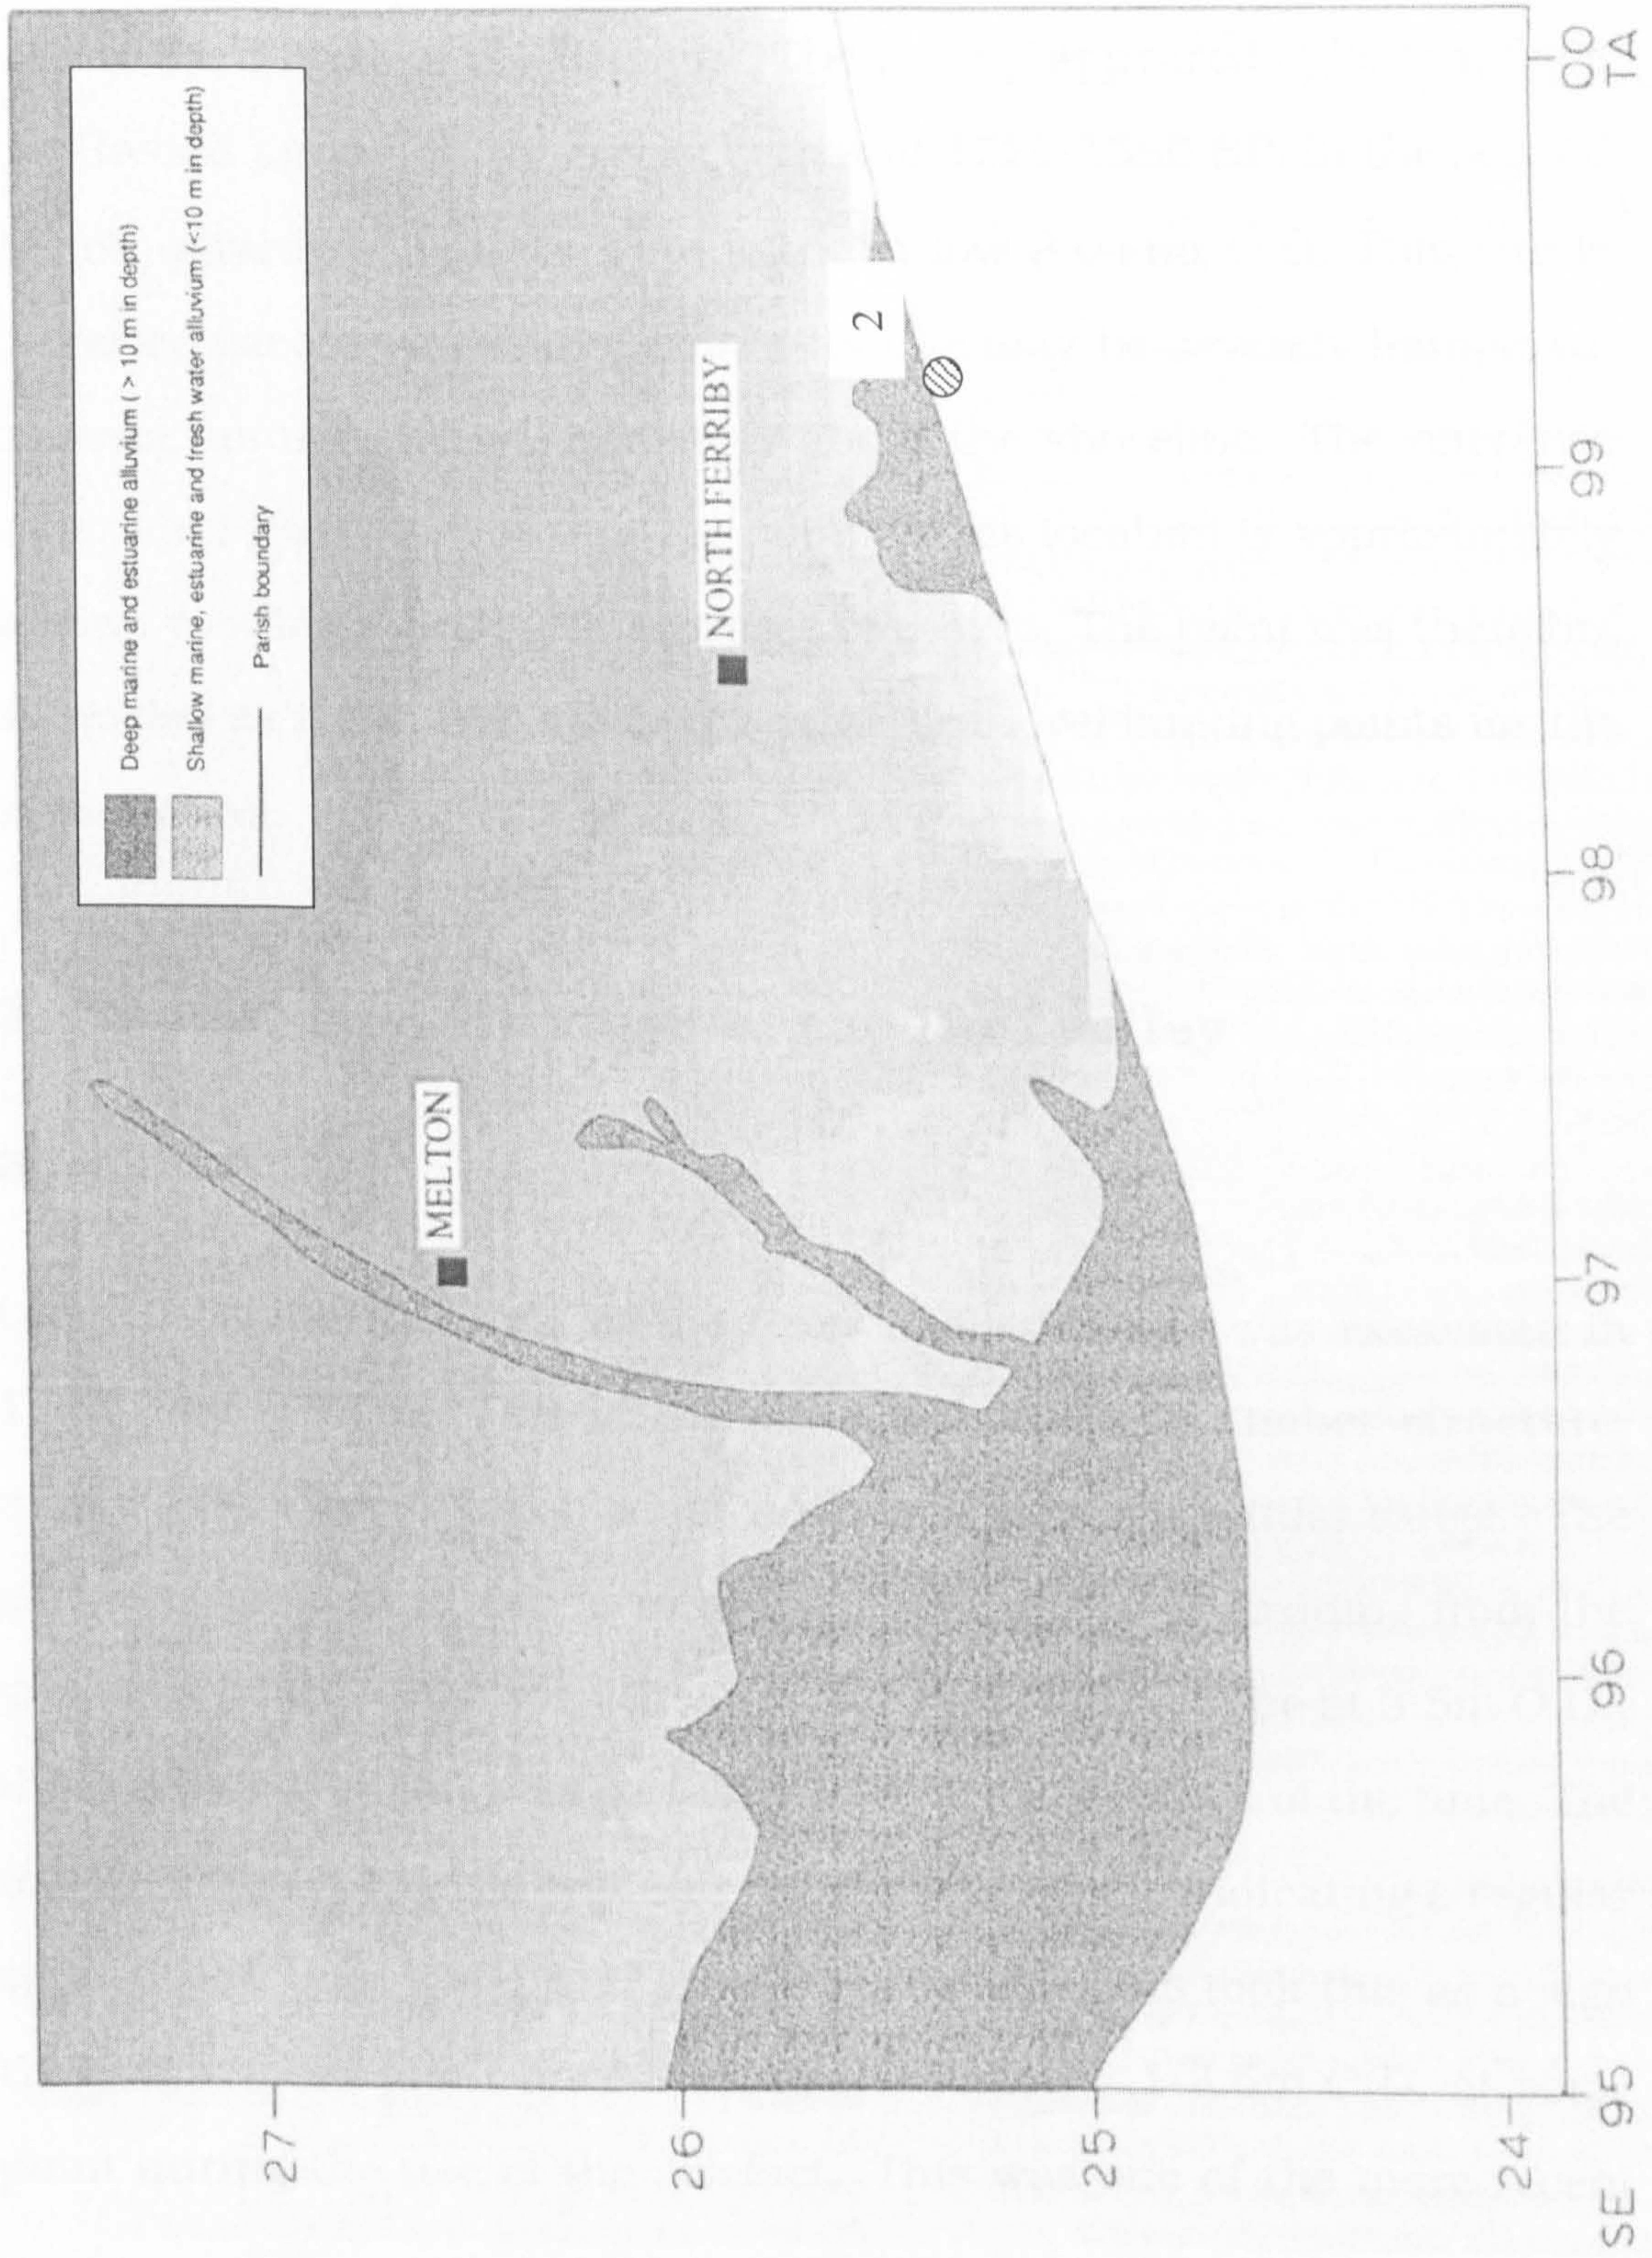


Figure 3.8: Site map of Ferriby-2 (hatched circle). From van der Noort and Ellis (1999).

Wheldrake-2

The site at Wheldrake, to the north-west of the Humber (shown in figure 3.9), was excavated following the discovery of a large pottery scatter next to the course of the Derwent. The pottery appeared to be that from the Roman period, more precisely that of 1720-1580 BP. In the second trench, shards of pottery were found below 9.044m O.D. This site is therefore used as a limiting point. Its use may be severely hampered, however, as the site sits some way above the shoreline. The reference water level (HAT) for the limiting point at this location is approximately 2.18m, yielding a final RSL estimate of 6.87m. The point was therefore discarded as it lies well above the other sea-level limiting points for the same period.

3.6.3 Wetland Heritage of the Hull valley

Brigham-6

Located on the east side of the River Hull, this site was excavated in 1982 (figure 3.10). One of the finds was that of a timber structure, most likely a jetty, which must have lain within the tidal range. The jetty was present in the form of multiple timbers protruding from the ground surface. The top of the structure was found to be at 3.5m O.D., though this would almost certainly have been dry most of the time. The top 0.3-0.4m of the timbers showed signs of decay, indicating a regular exposure to both water and air. The archaeologists took this as a sign that high water must have been in the region 3.1-3.5m O.D. at some point during the use of the artefact. This was one of the more recent

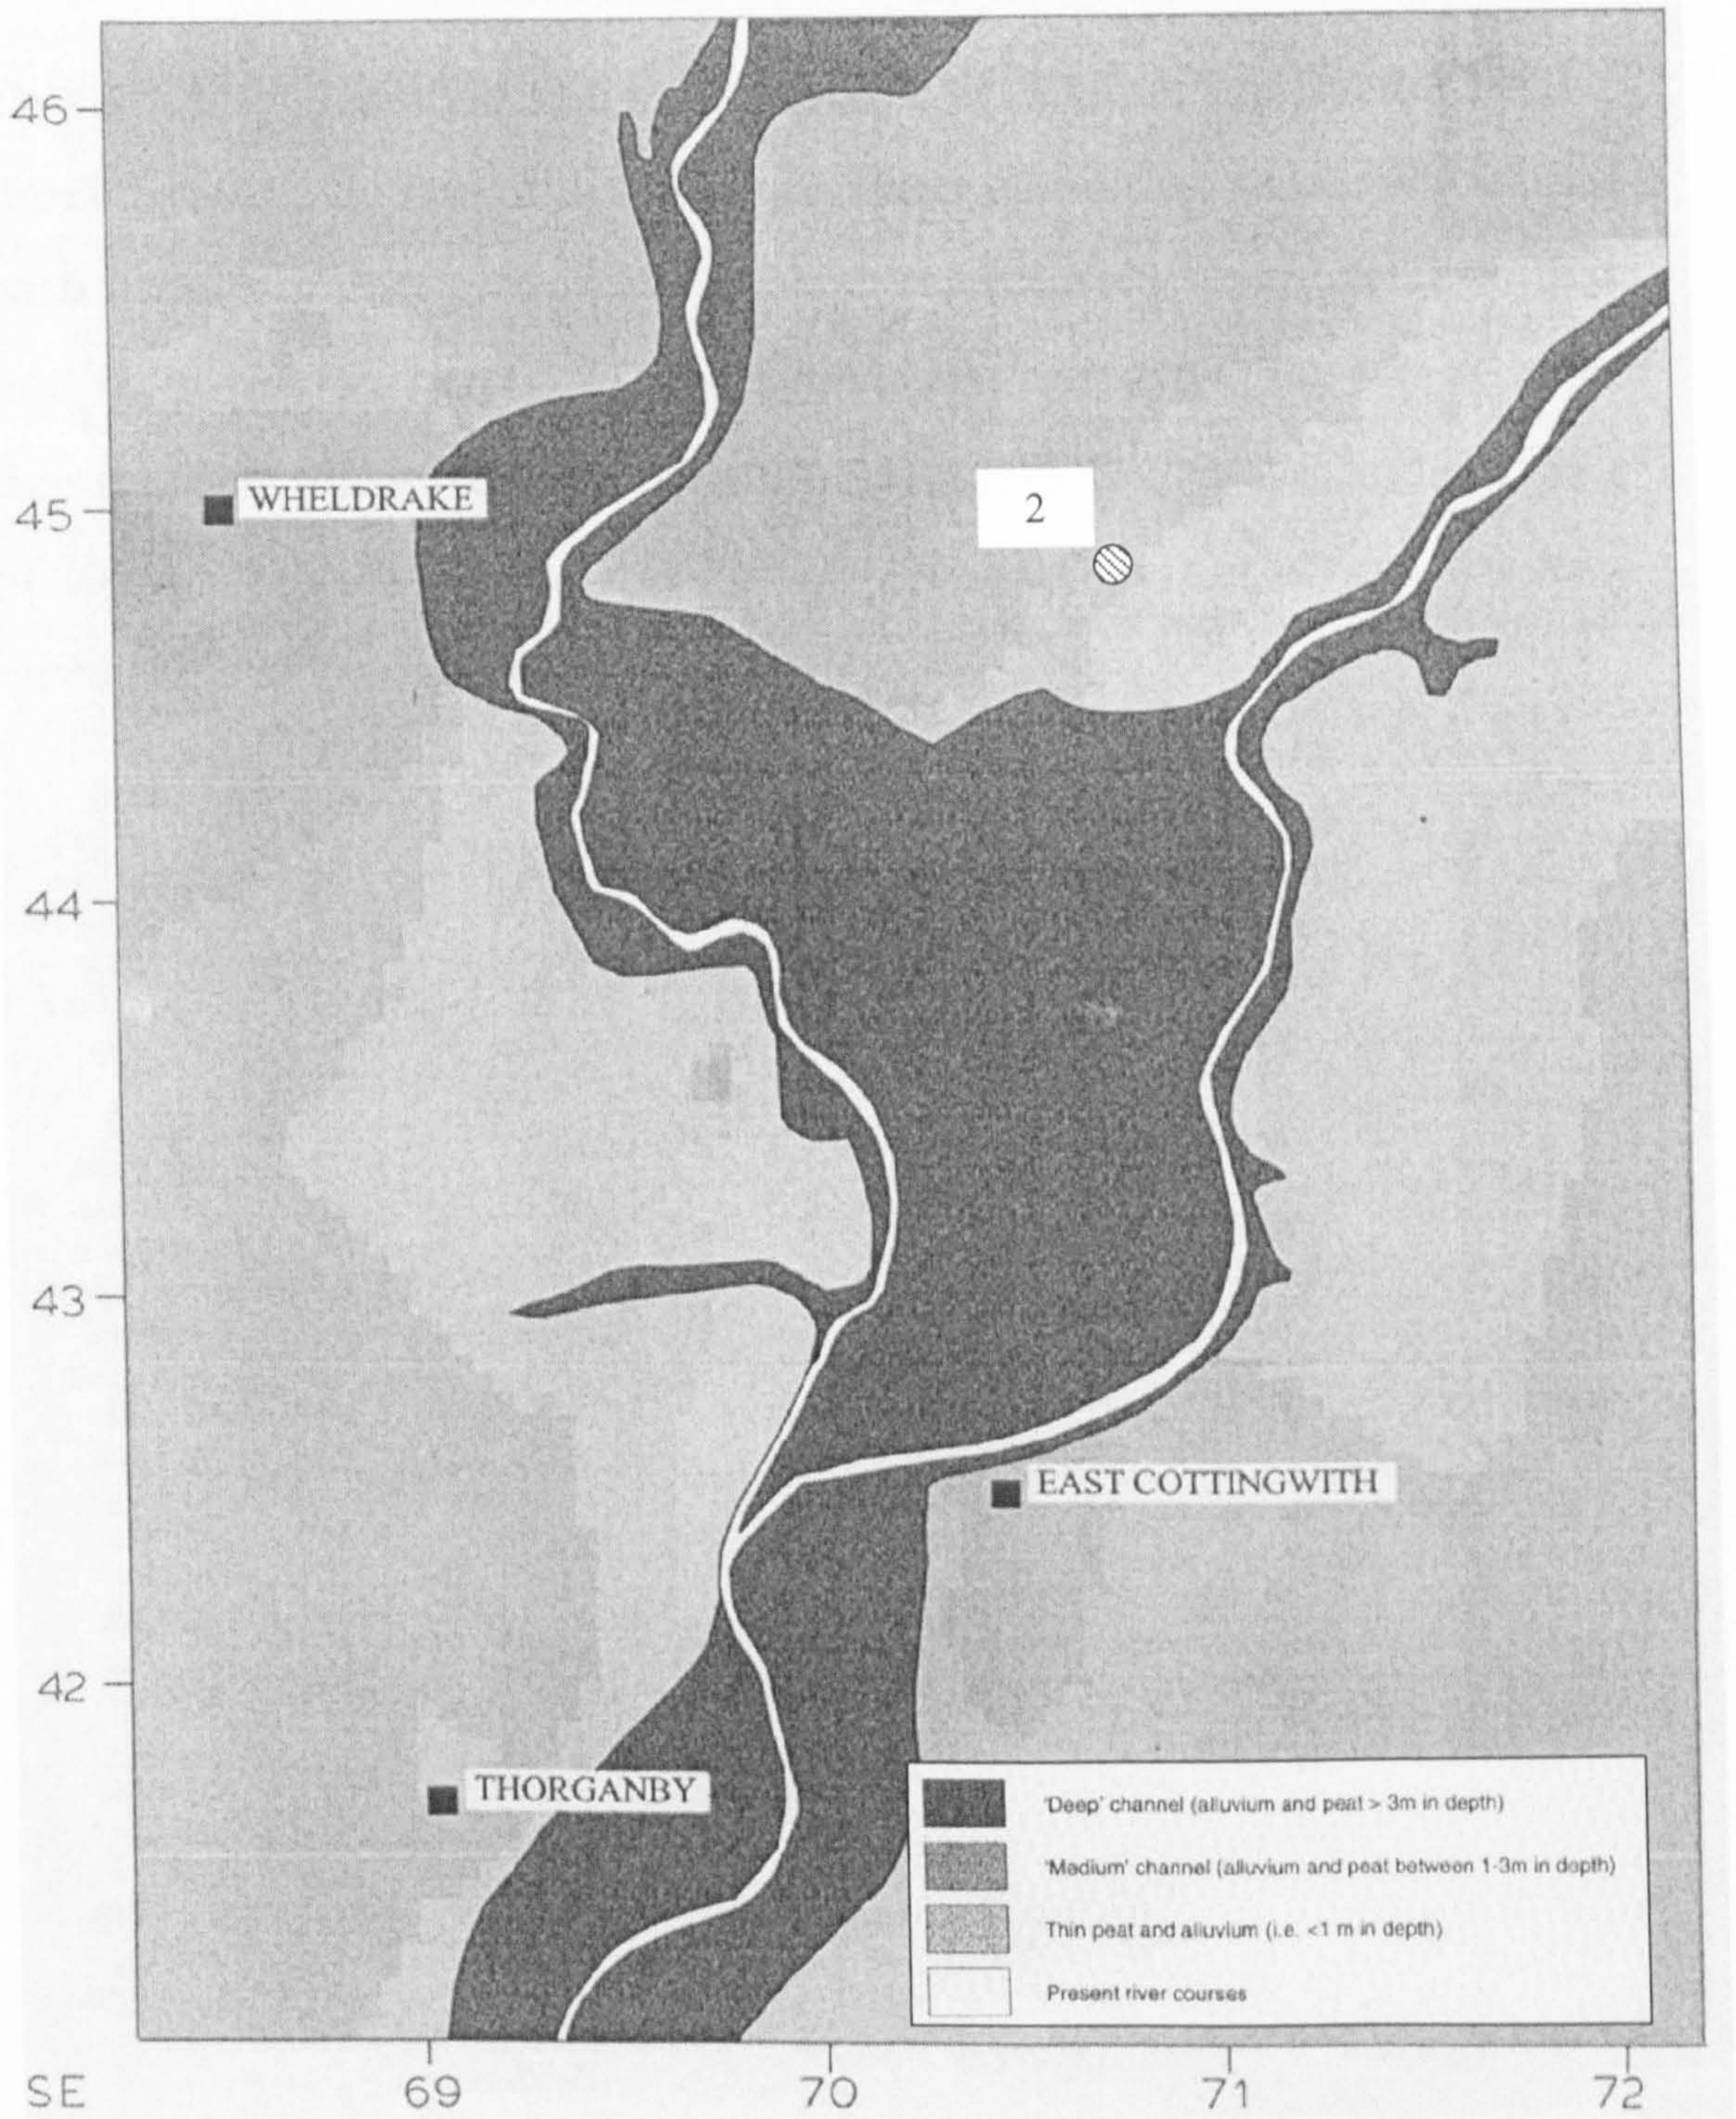


Figure 3.9: Site map of Wheldrake-2 (hatched circle). From van der Noort and Ellis (1999).

artefacts used. It was ascribed a date of 1293-1210 cal yrs BP.

This site was classified as an index point, with a reference water level of $(\text{MHWS}+\text{HAT})/2$. The indicative range was estimated as MHWS to HAT with 95% probability. Levelling standard error here was given as 5cm and thickness s.e. as 10cm. The tidal level standard error was reduced due the distributional form of $(\text{MHWS}+\text{HAT})/2$ to give a standard error of 20cm. The indicative range for the jetty came out at 28cm. The total elevation standard error was 35.9cm.

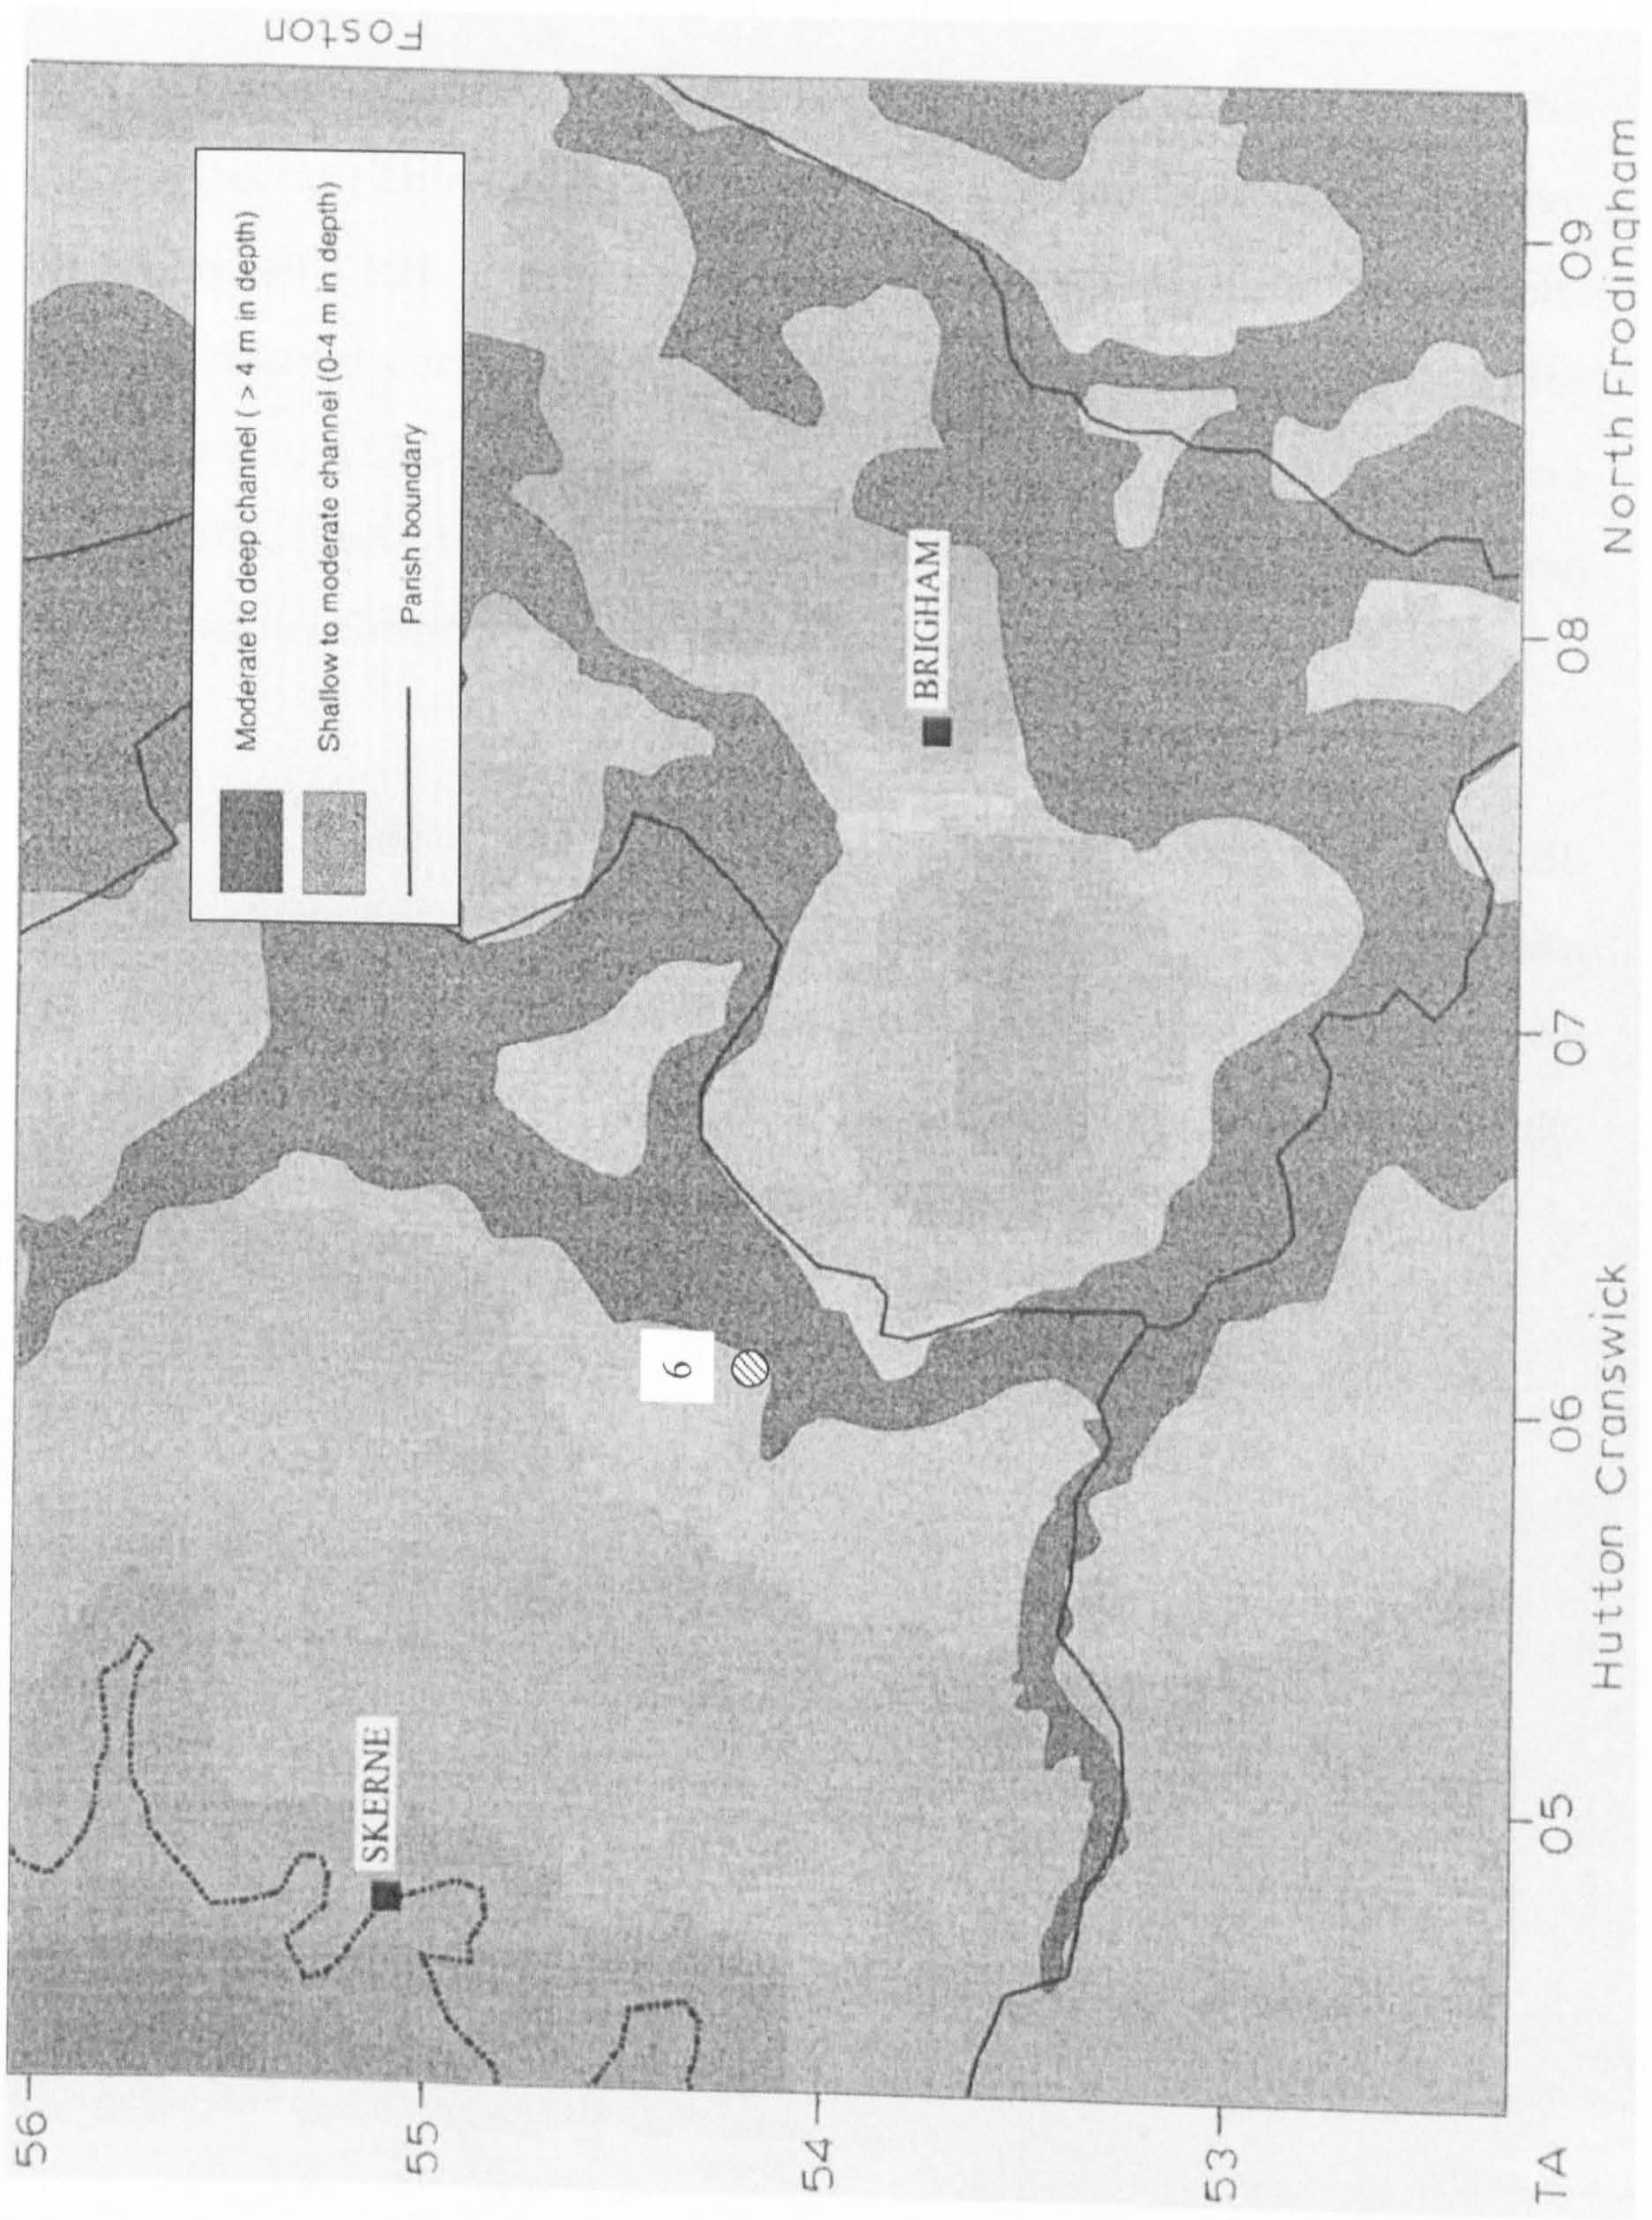


Figure 3.10: Site map of Brigham-6 (hatched circle). From van der Noort and Ellis (2000).

Weel-2

Excavated as part of the Stone Carr excavations in 1999, this site is a late Mesolithic flint production area located to the north of Hull (figure 3.11). A number of flint pieces were found during the excavation, the most important being that of the antlers of a red deer which would have been used as a tool. The site was found on one edge of an excavated pit, the two nearest corners being at 0.754m and 0.802m O.D. The antler was dated to 1880-1520 cal yrs BP (OxA-8954, 3420±45 yrs BP and OxA-9010, 3365±45 yrs BP). Thus this site is classified as a limiting point with an elevation of 0.778m O.D..

HAT at this location was found to be 2.43m O.D., thus the final RSL estimate of the antler was -1.66m. The data points was given elevation standard error 2cm, thickness standard error 1.2cm, and tide level standard error 41.4cm. As a limiting point there is no indicative range. The total standard error the antler was 41.5cm.

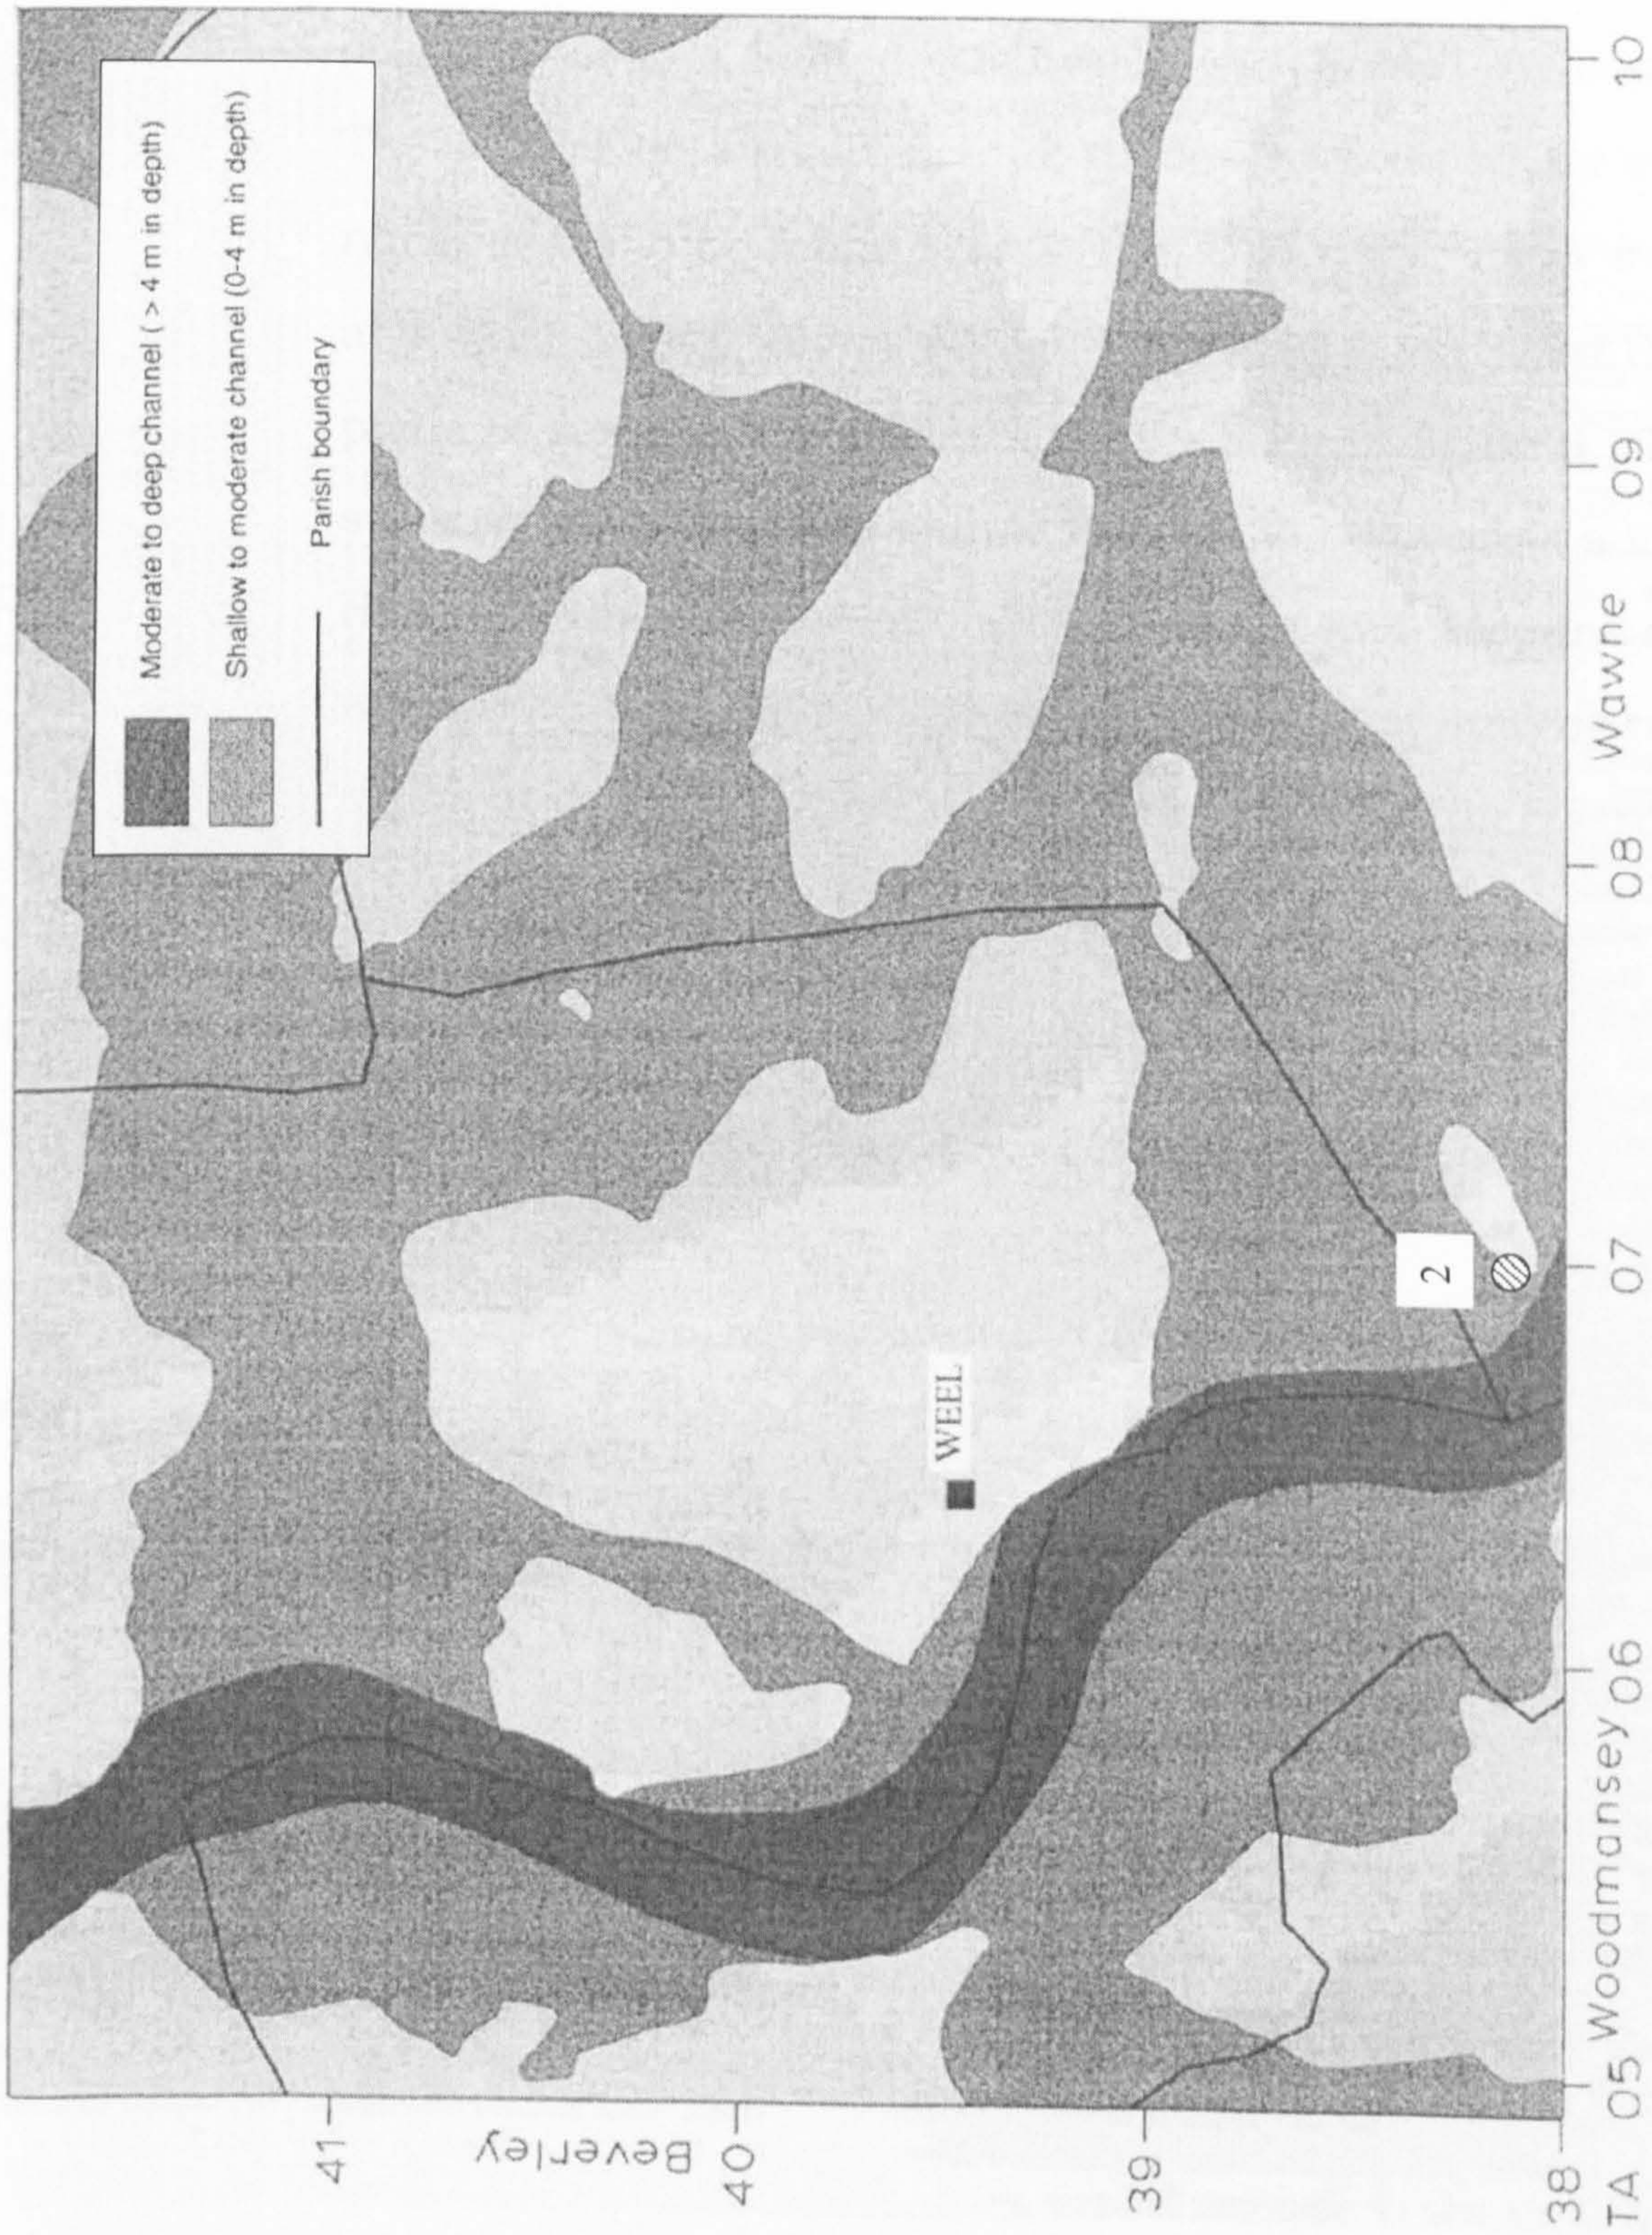


Figure 3.11: Site map of Weel-2 (hatched circle). From van der Noort and Ellis (2000).

Kelk-6

The site at Kelk is a square, late Iron Age settlement located some distance to the north of the Humber (shown in figure 3.12). The site was subject to a limited geophysical survey and excavation in 1999. The site lies on a small mound; the lowest point where archaeological remains are to be found is at 7.15m O.D. This comes from a pit dug at surface height 7.95m, with pottery found 0.8m below the surface. A photo of some of the artefacts can be seen in figure 3.13. No objects suitable for dating were found. This could be used as a limiting point. However, the final elevation, after taking into account the reference water level, was 3.89m. This is well above other limiting points for this region at this date. Therefore the data point was disregarded.

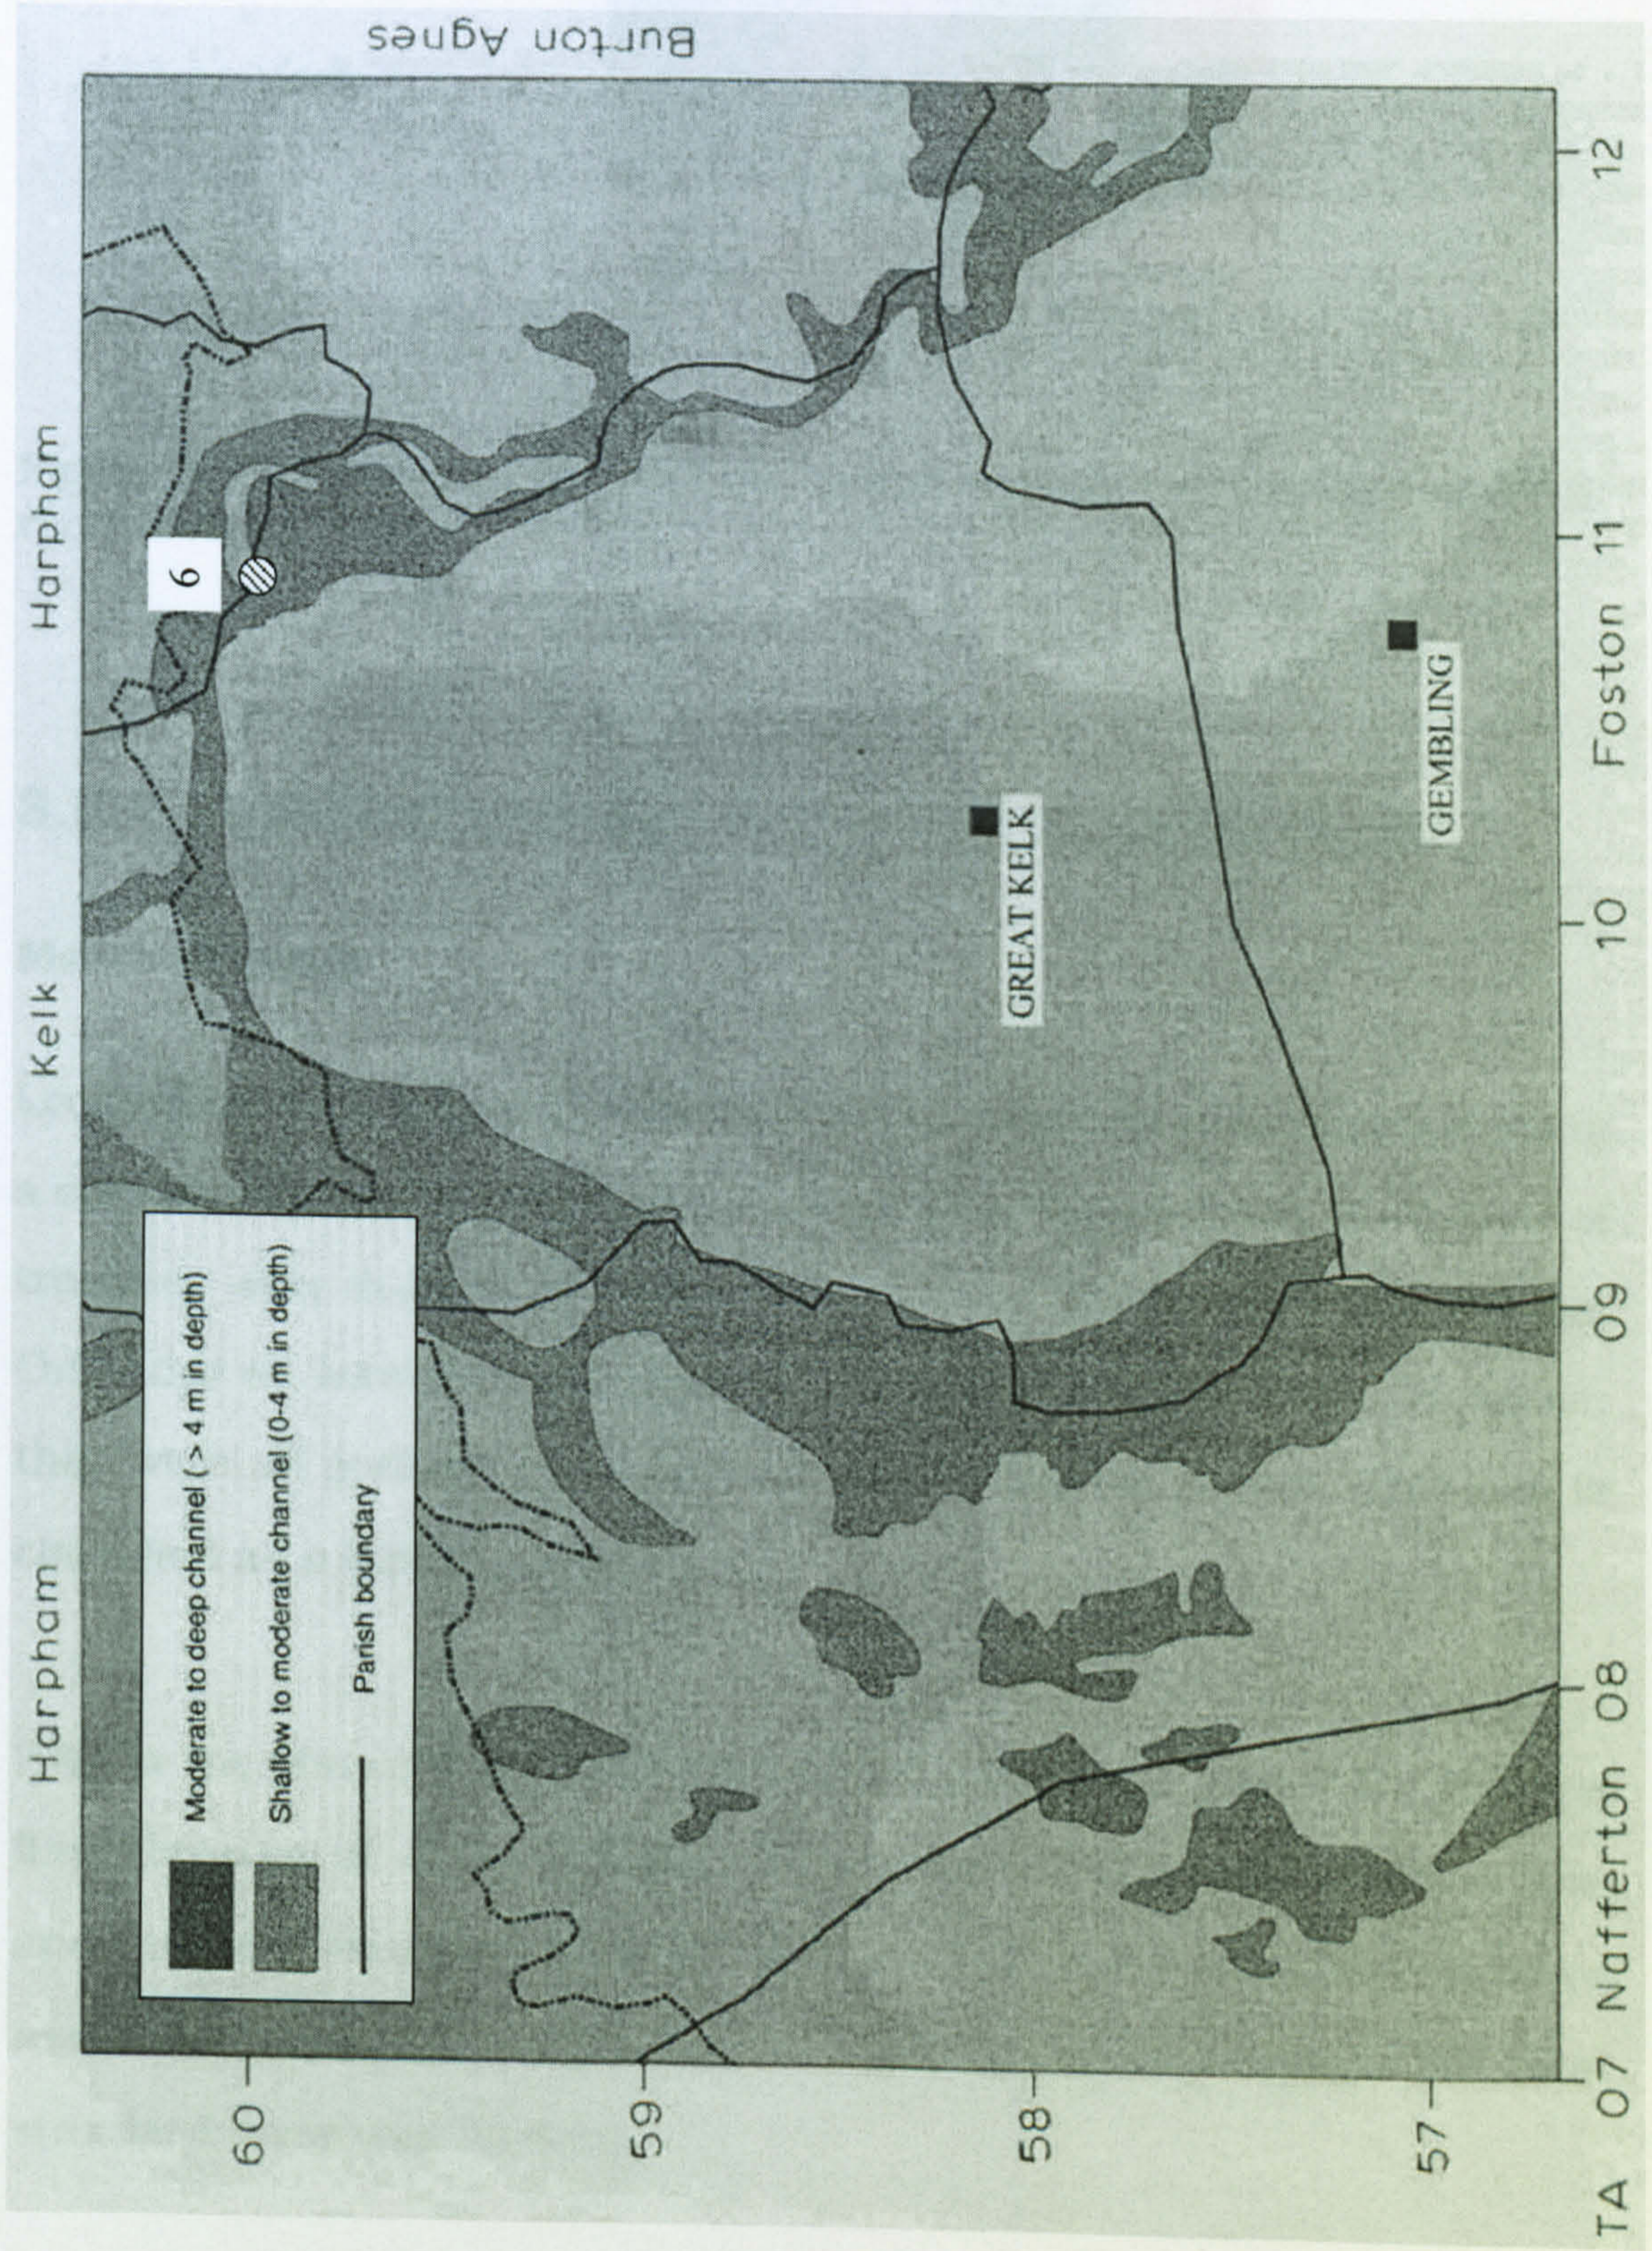


Figure 3.12: Site map of Kelk-6 (hatched circle). From van der Noort and Ellis (2000).



Figure 3.13: A selection of iron age moulds and sprue cups from Kelk-6, recovered in 1999.

3.6.4 Wetland heritage of the Lincolnshire Marsh

Marshchapel-2

Located to the south-east of Grimsby, this site was excavated following a considerable pottery find by the landowner (figure 3.14). A number of trenches were dug, the lowest of which came from a surface of 2.22m O.D., but no information exists as to the depth of the finds. However, they were all considered to be from the Mediaeval period. This site is classified as a limiting point.

HAT at the Marshchapel site was calculated to be 3.16m O.D., giving a final elevation of -0.94m. The levelling standard error was given as 2cm and the tidal standard error as 39.6cm. There was no thickness error associated with this sample as it was levelled at the surface. The total standard error was 39.6cm.

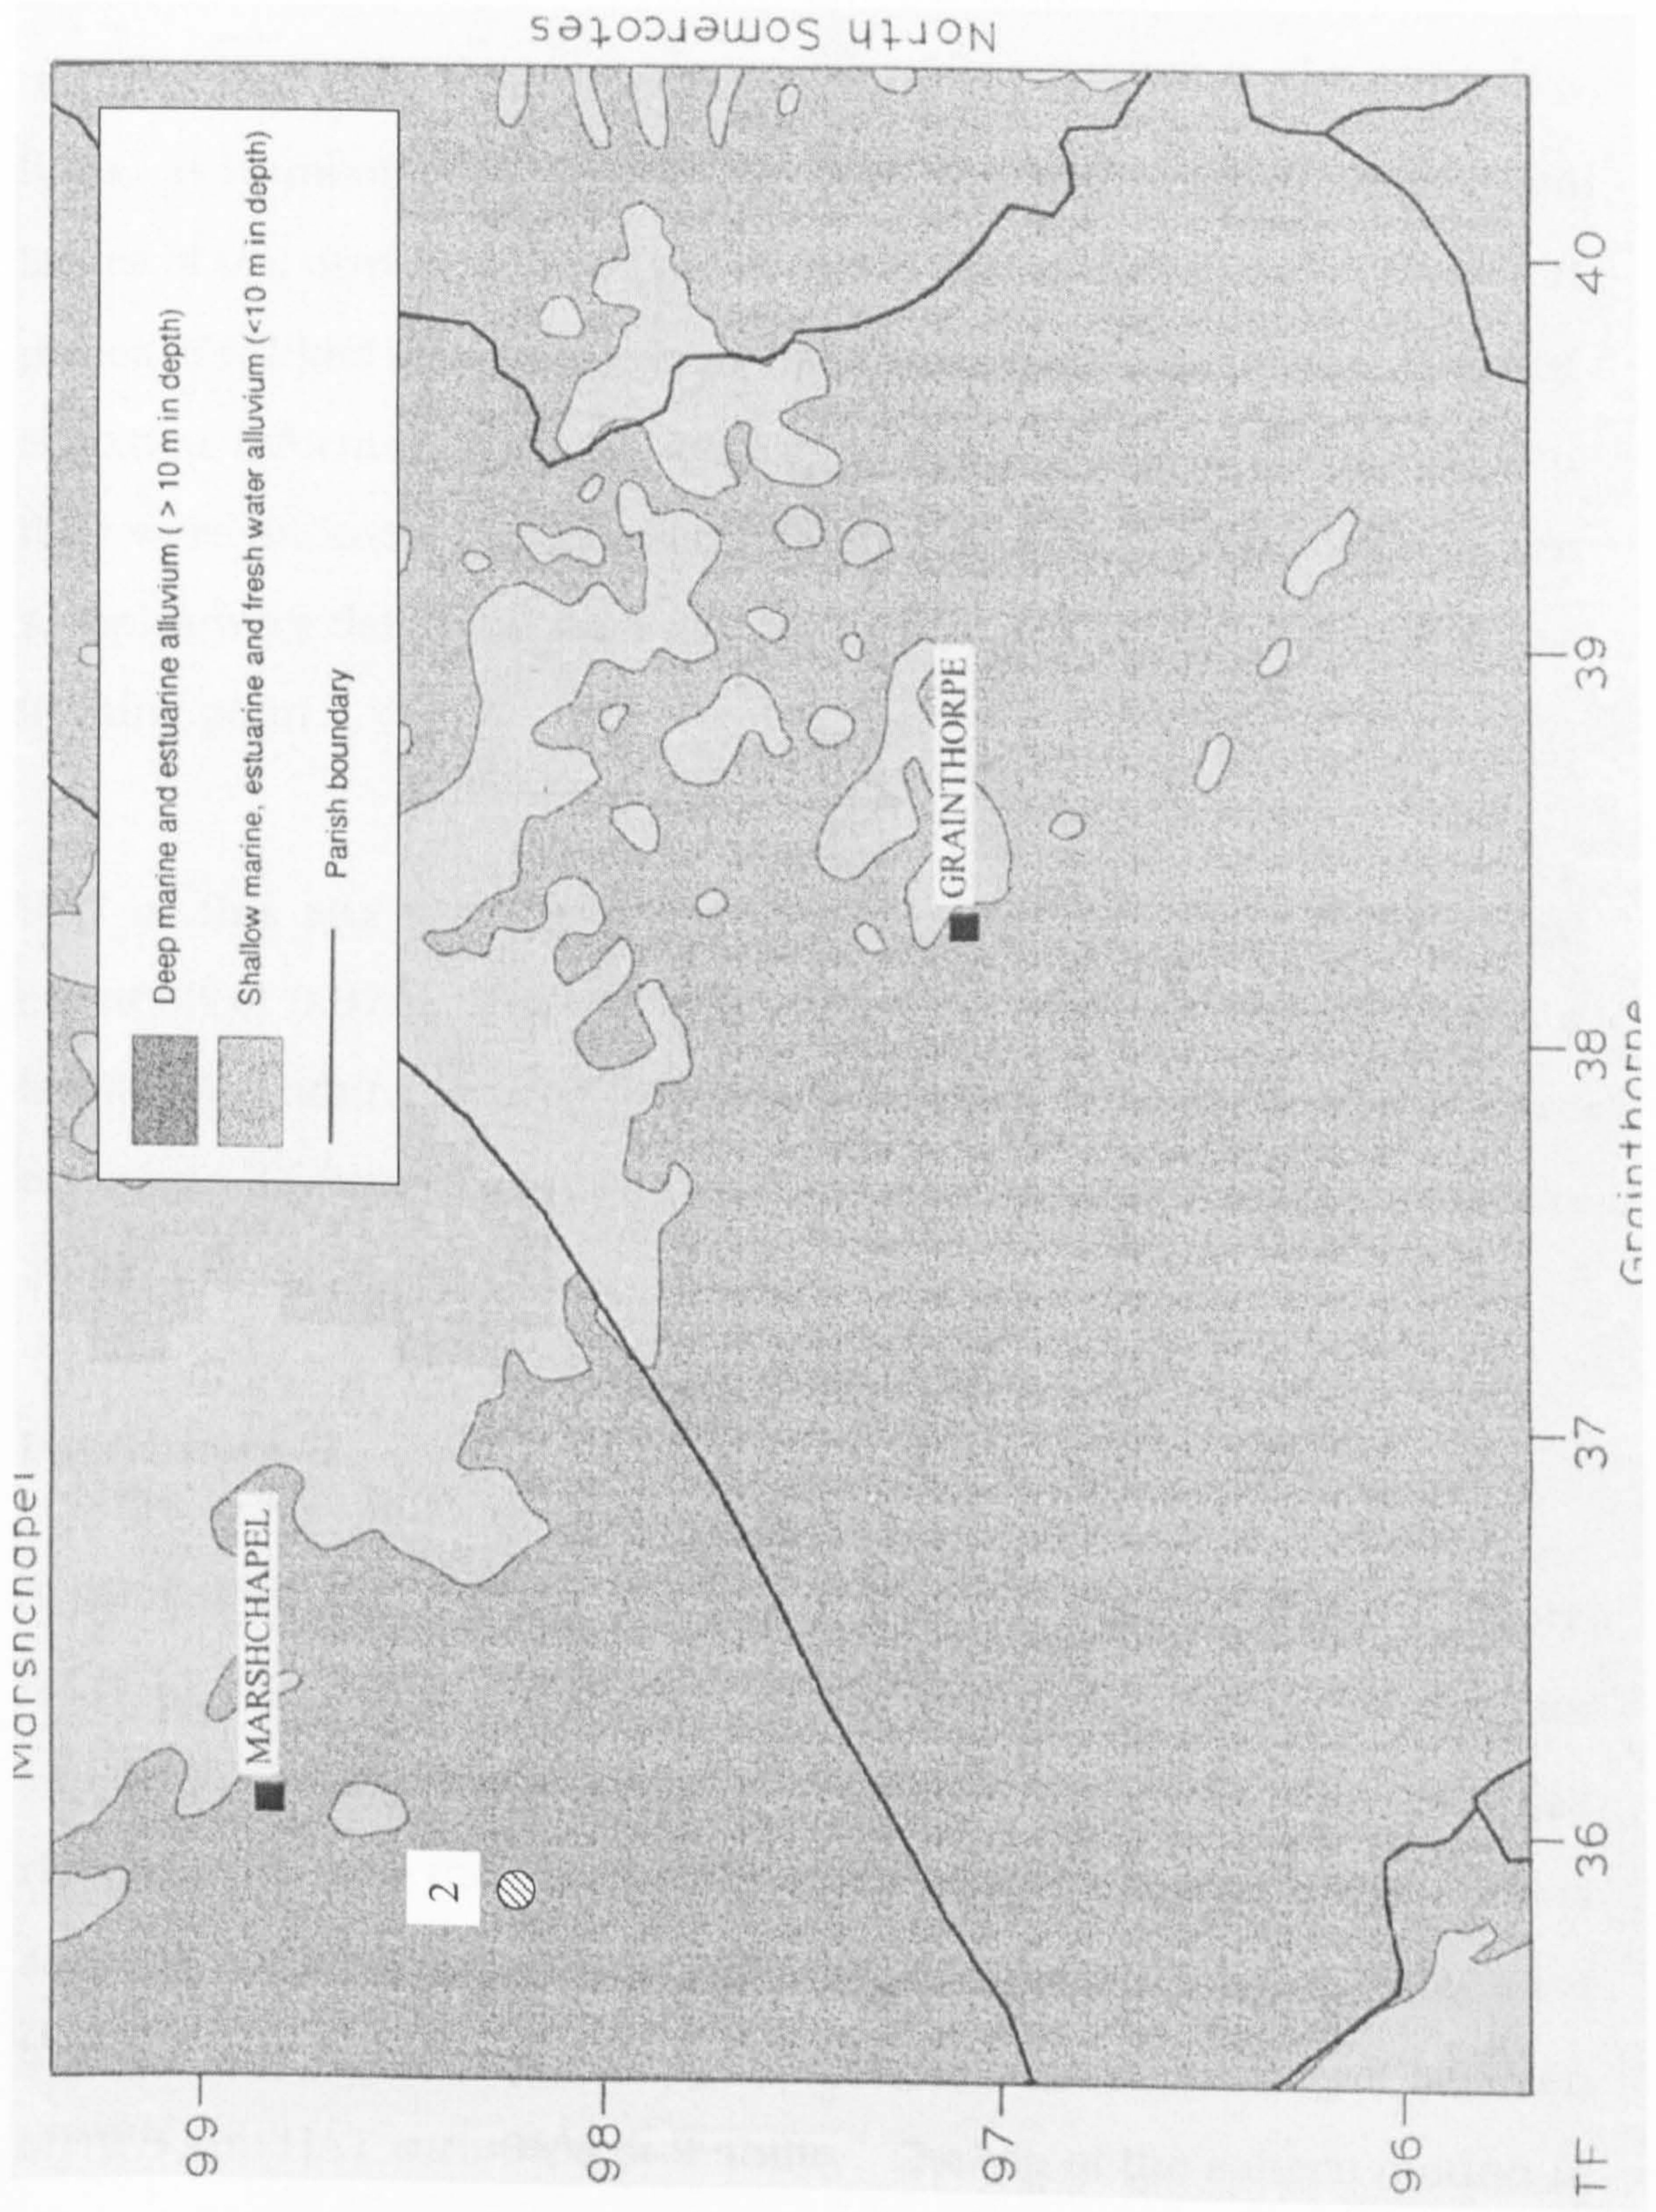


Figure 3.14: Site map of Marshchapel-2 (hatched circle). From van der Noort *et al.* (2001).

Willoughby-13

Willoughby-13 (located to the south-east of the Humber, figure 3.15) is also known as the Butterbump barrow cemetery site and spans a number of archaeological periods due to the variety of archaeological finds. A number of sampling tins were driven through the disturbed layers of soil during excavations in 1999. These turned up a number of pieces of worked wood which may have come from a structure. Detailed elevation information is not available for all the wood fragments but they were all known to have lain below 2.81m O.D. Four of the seven samples were dated, all were middle Bronze Age. This is classified as a limiting point.

HAT at this site was estimated to be 1.84m O.D.. This gave a final elevation of 0.97m. The site was excavated relatively recently, thus a levelling standard error of just 2cm was given. The tide level standard error was 33.3cm. The total standard error for this limiting point was 33.3cm.

Ingoldmells-6

A salt production site, or saltern, of Iron Age/Roman descent (figure 3.17) located to the south-east of the Humber. The archaeologists believe the saltern lay near the shore and was allowed to fill with the rising tide before the water evaporates leaving the salt. This means salterns are ideal as index points and points to a reference water level in the upper tidal range; the archaeologists agreed on a level between MHWS and HAT with 95% probability. The top of the saltern mound at

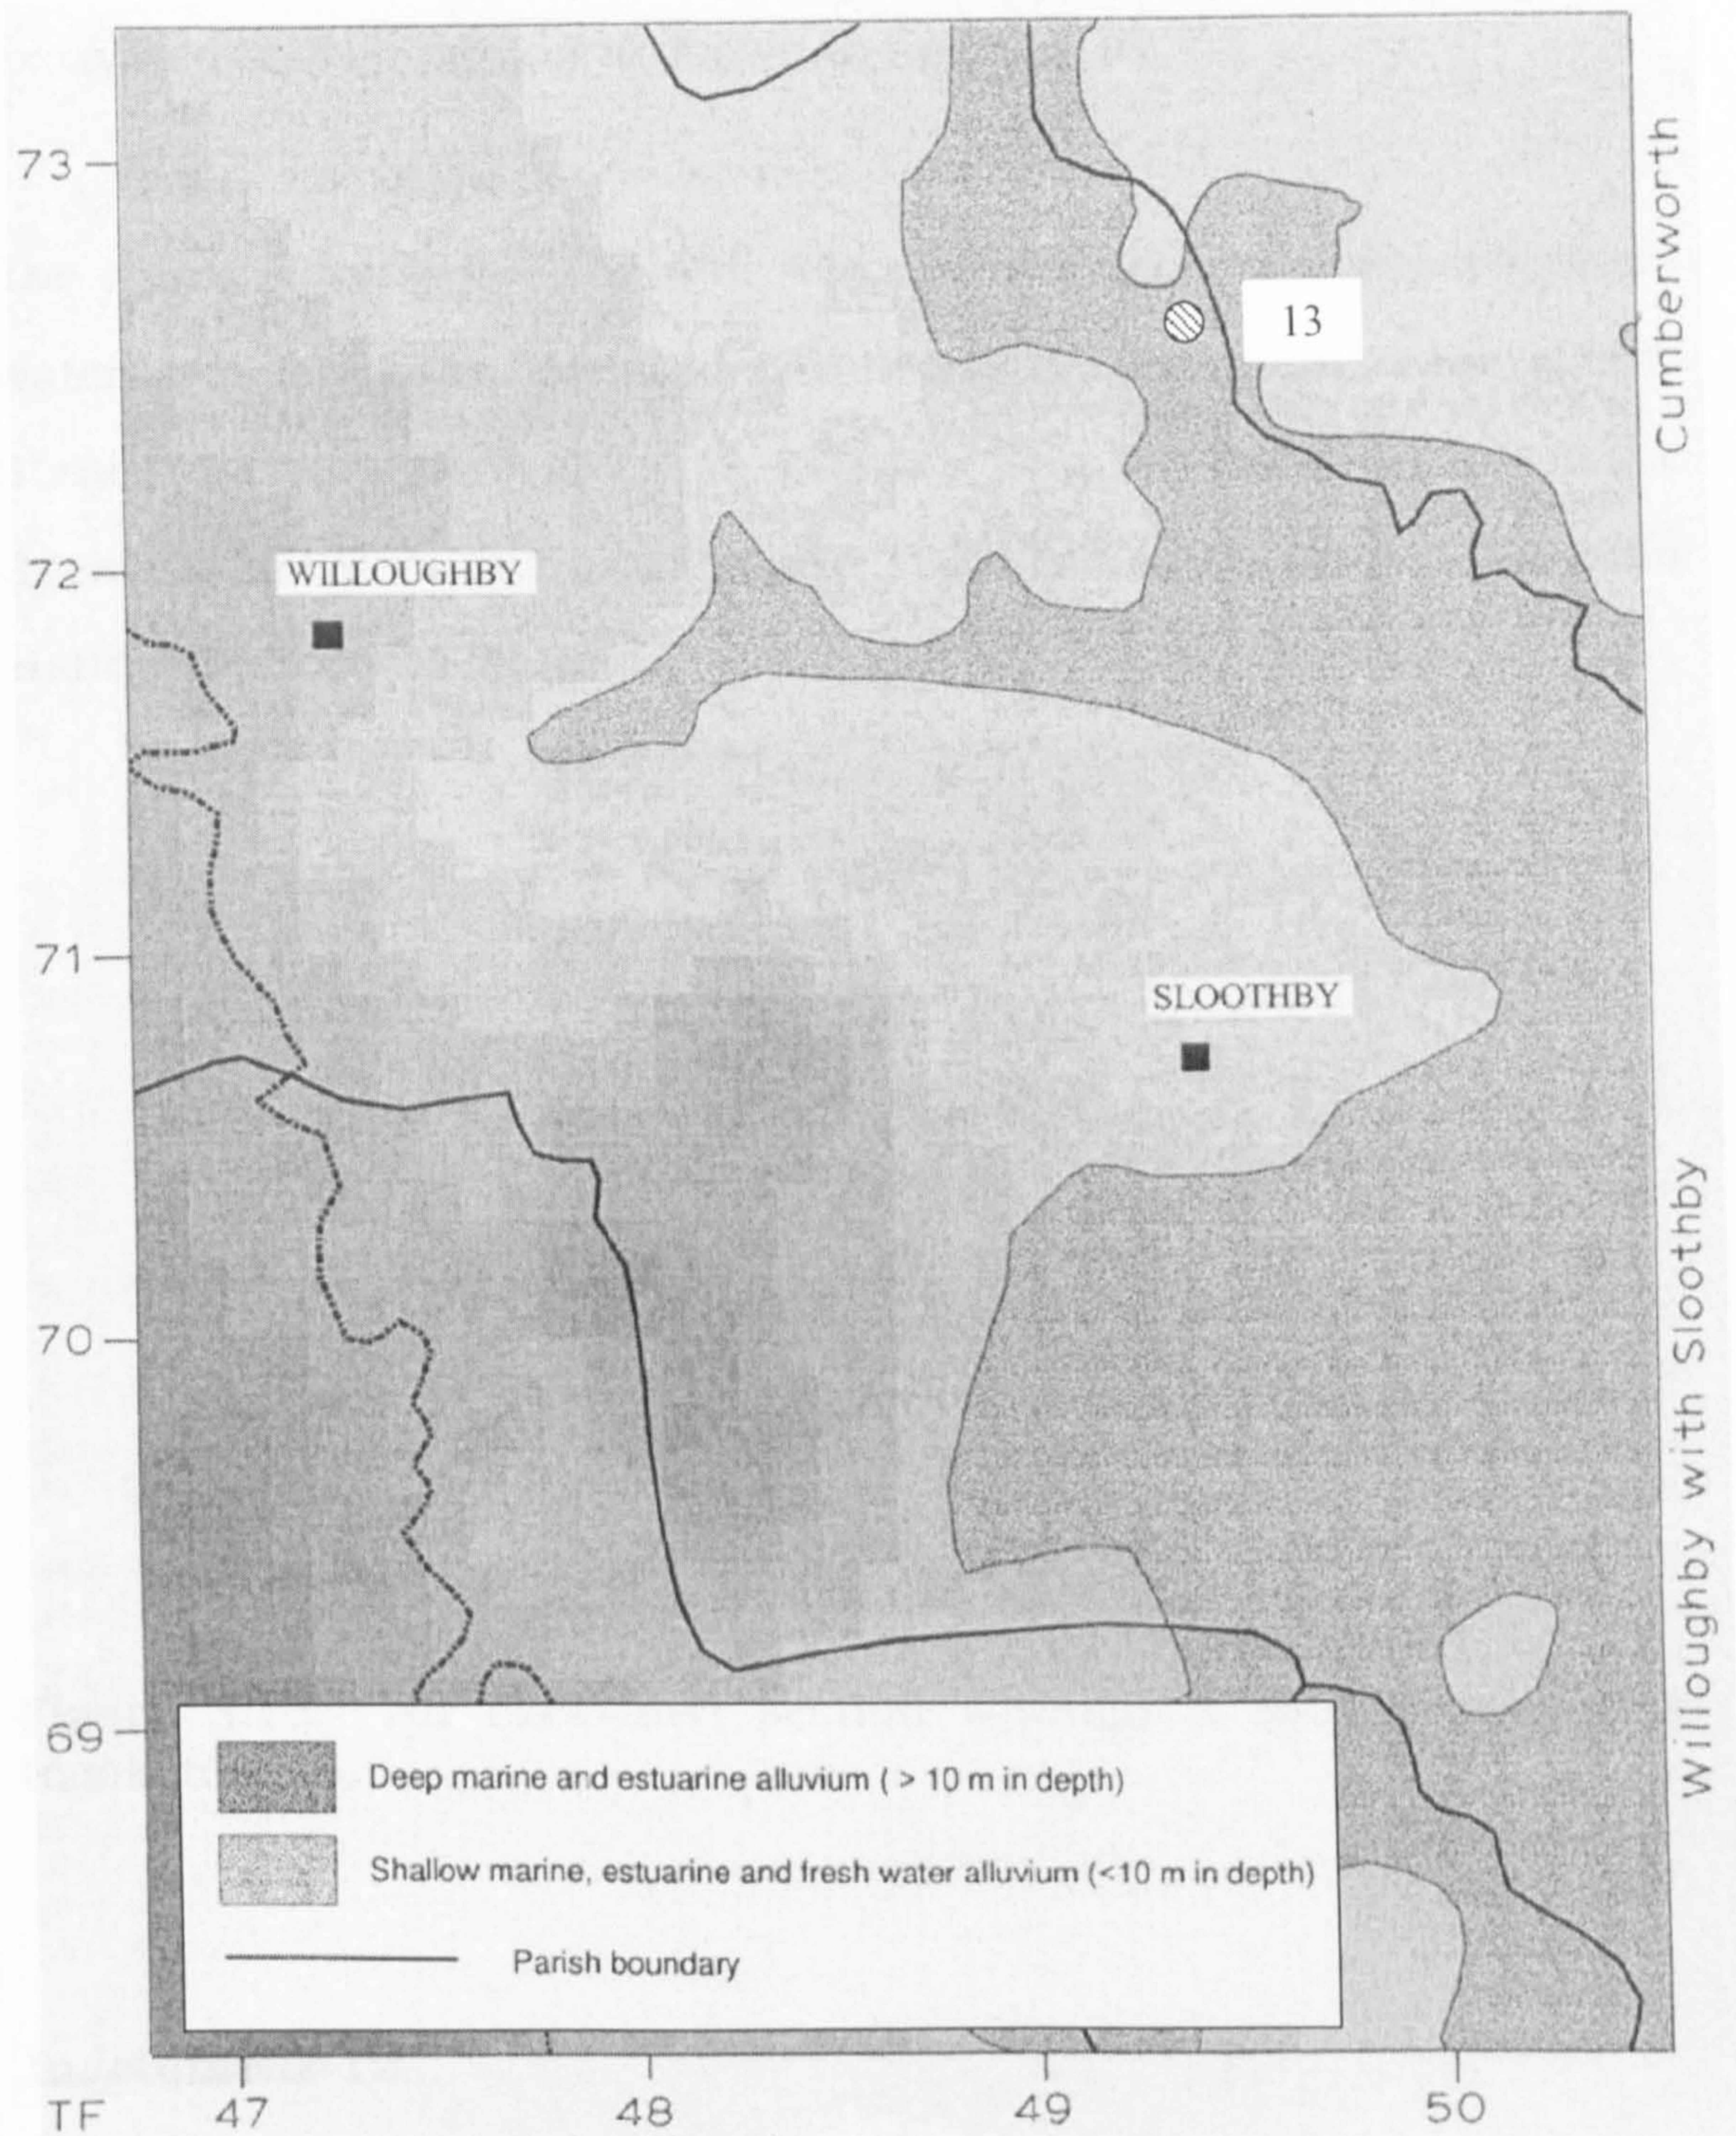


Figure 3.15: Site map of Willoughby-13 (hatched circle). From van der Noort *et al.* (2001).

Ingoldmells-6 was found at 1.18m O.D., the bottom at 0.98m O.D.. No exact dates are available other than period estimates. A picture of an excavated saltern mound is shown in figure 3.16.

The mean elevation of the site was 1.08m O.D., with the reference water level as 2.04m. The final elevation of the data point was -0.96m. Standard errors were given as follows: 2cm for levelling, 2.5cm for thickness, 21.9cm due to tides and 10cm indicative range. The total standard error was 24.1cm.



Figure 3.16: An excavated section through a saltern mound at Ingoldmells-6.

Ingoldmells-12

This saltern is located next to Ingoldmells-6 and is also of Iron Age/Roman descent (figure 3.17). Saltern debris was found between 2.09m O.D. and 1.79m O.D. The same reference water level as

Ingoldmells-6 applies.

A reference water level of 2.04m, identical to that of Ingoldmells-6 was found here, giving a final elevation of -0.25m. Other errors were all identical to Ingoldmells-6 with the exception of thickness standard error, here given as 7.5cm. The final elevation standard error was 25.1cm.

3.7 Summary

This chapter has outlined the two main data sources of this thesis. These comprise the LOIS database and the Humber Wetlands Project. The data from the HWP had not been adjusted for the indicative meaning like that of the LOIS project; this required some basic statistical modelling. The final data set created has a diverse mixture of data points from around the estuary over many time periods. It is hoped that this will provide suitable information for more detailed statistical analysis.

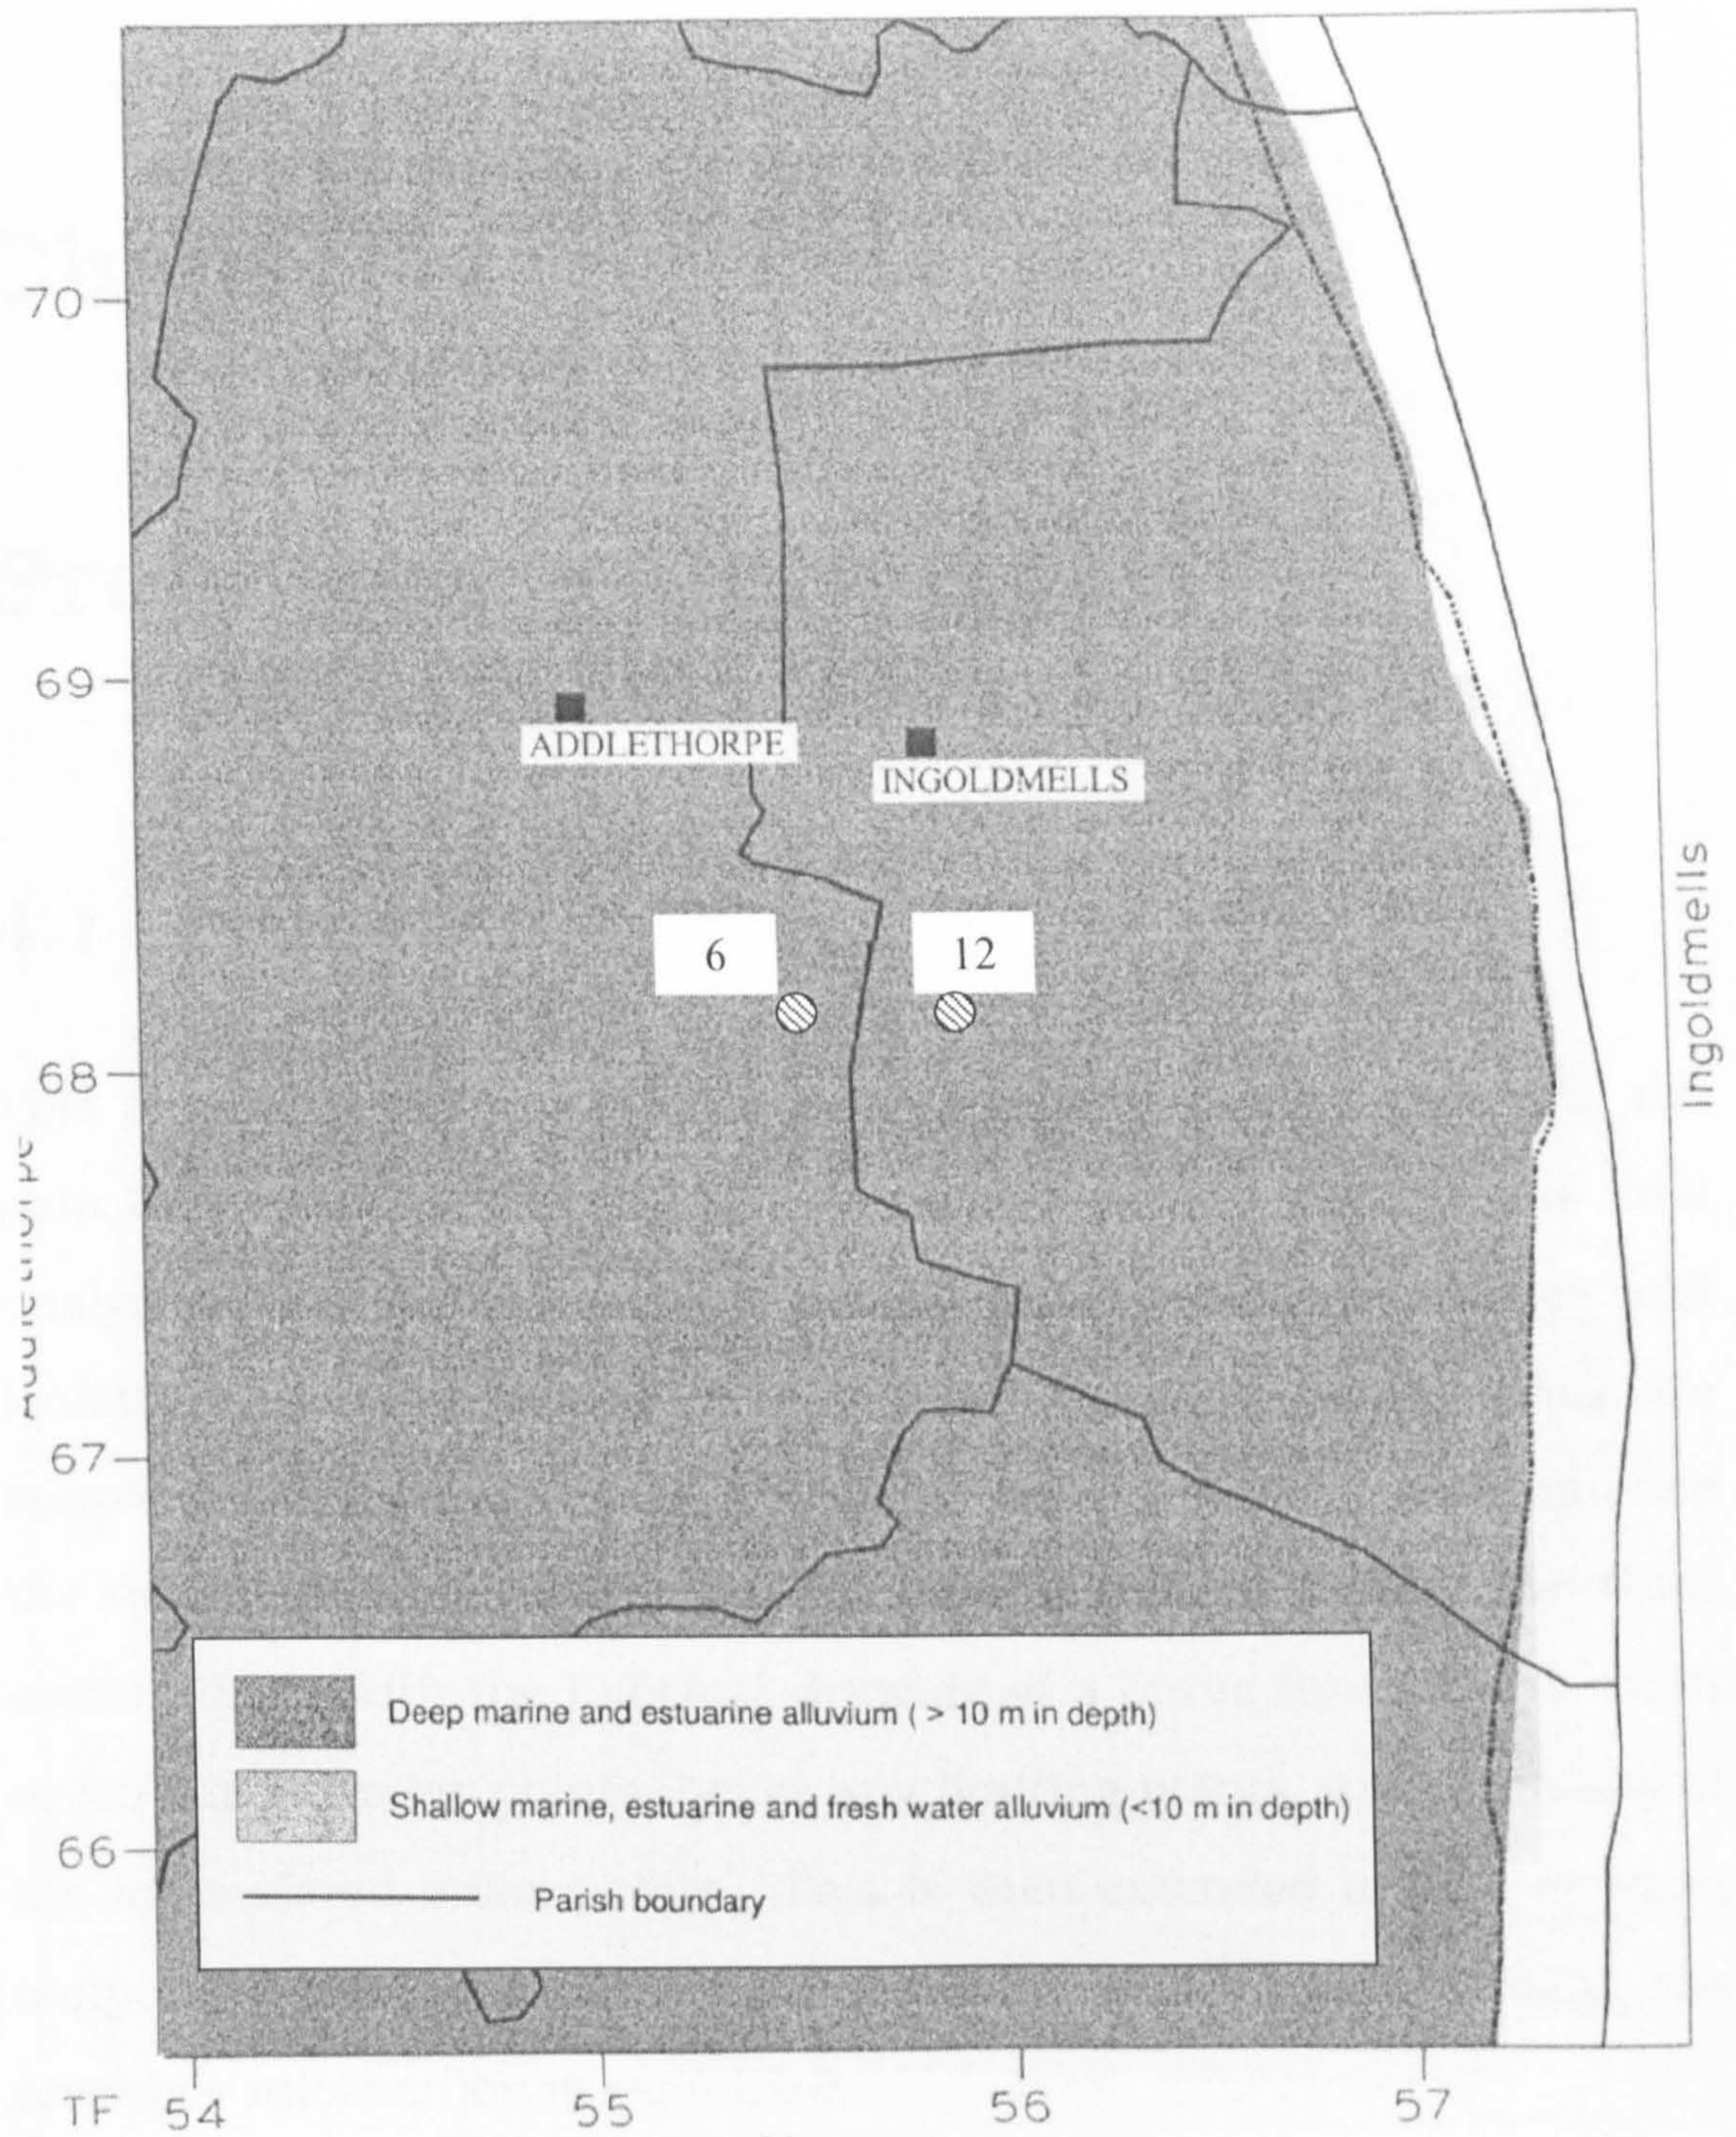


Figure 3.17: Site map of Ingoldmells-6 and 12 (hatched circles). From van der Noort *et al.* (2001).

Chapter 4

Preliminary Data Analysis

4.1 Introduction

This chapter is concerned with gaining a basic understanding of the data by identifying features which may prove important in the final analysis. Initially, this involves calibrating radiocarbon dates and looking at the relevant periods from which the archaeological data are supplied. Subsequently, graphical methods are used to summarise the data with some discussion on cartographic methods. The third section deals with the informal drawing of a curve through the mean or similar for index points, below any limiting points and above any of the intercalated index points. This is then extended to look at extra temporal and spatial correlation structure in the data. Finally, the tendency information is examined.

4.2 Age data

The age data in this project come in two forms:

- Radiocarbon uncalibrated dates,
- Archaeological periods.

The distinctions between the two concern the distribution which must be taken into account by the models involved.

4.2.1 Radiocarbon dating

The need for calibration of radiocarbon dates is outlined in section 2.4.1. To calibrate a radiocarbon date, 3 pieces of information are required:

- A laboratory estimate in radiocarbon years before present (present taken to be the year 1950 AD).
- An estimate of the error around the radiocarbon date, usually expressed as a 1σ standard deviation.
- A curve with error bands comprising the radiocarbon calibration curve.

There currently exists a number of different methods for calibrating radiocarbon dates. A discussion of calibration methods, including a method which uses Gaussian processes, can be found in Aguilar *et al.* (2002). The discrepancy arises because of the way the calibration curve is drawn; also the way in which the error around the curve is accounted

for. The method presented here is based on that proposed by Christen (1994). This method utilises Bayesian analysis.

Suppose that the radiocarbon laboratory return a date of $y \pm \sigma$. The value y requires calibration; it is based on the radiocarbon timescale and corresponds to a calendar date θ . As discussed in section 2.3.8, the laboratory reported value comes from a large sample of count data and it is assumed that the Normal approximation to the Poisson distribution is appropriate. The calibration curve is a function, $r(\cdot)$, which takes the true calendar date as its argument and returns a date on the radiocarbon timescale: it is based upon a set of data points D . The following distributional assumption can be made:

$$y|r(\theta), \sigma \sim N(r(\theta), \sigma^2) \quad (4.1)$$

The task is to find the posterior distribution of θ given the available data. This allows the possibility of setting a prior distribution on theta if such information is available, for example chronological information used in Buck *et al.* (1996); Buck and Millard (2004). To proceed, the calibration curve itself has a distributional assumption:

$$r(\theta) \sim N(\mu(\theta), \sigma^2(\theta)) \quad (4.2)$$

where $\mu(\theta)$ is a piece-wise curve joining known points on the calibration curve, and $\sigma^2(\theta)$ is a quadratic function of the known data points' variances (from D) and their known true dates. It is arrived at through

the following equation:

$$\sigma^2(\theta) = \sigma_i^2 \left(\frac{\theta_{i+1}^{(d)} - \theta}{\theta_{i+1}^{(d)} - \theta_i^{(d)}} \right)^2 + \sigma_{i+1}^2 \left(\frac{\theta - \theta_i^{(d)}}{\theta_{i+1}^{(d)} - \theta_i^{(d)}} \right)^2 + \frac{\lambda^2(\theta - \theta_i^{(d)})(\theta_{i+1}^{(d)} - \theta)}{\theta_{i+1}^{(d)} - \theta_i^{(d)}}, \quad (4.3)$$

where $\theta_i^{(d)} \leq \theta < \theta_{i+1}^{(d)}$, $i = 1, \dots, n$. Here, $\theta_i^{(d)}$ corresponds to the calendar age of the i th sample dated on the calibration curve d , σ_i^2 is the associated error of the uncalibrated version of $\theta_i^{(d)}$, and λ is a value associated with the atmospheric radiocarbon concentration. This is estimated by Christen (1994) to be 19.

The simple result of this analysis is that the posterior for θ can be expressed as:

$$p(\theta|y, \sigma, D) \propto \frac{1}{\sqrt{\sigma^2(\theta) + \sigma^2}} \exp \left(-\frac{1}{2} \frac{(y - \mu(\theta))^2}{\sigma^2(\theta) + \sigma^2} \right) \times p(\theta) \quad (4.4)$$

where $p(\theta)$ is the prior distribution given to the calendar age.

The calibration of radiocarbon determinations requires the use of MCMC. Simulations from the laboratory date distribution are evaluated by the calibration curve many times to produce a suitable set of calendar dates. As an example of the complex distributional shape of posterior calibrated dates, one of the dates from the Humber data set (3170 ± 50) was calibrated individually. The posterior density is shown in figure 4.1.

Calendar age distributions (when compared to radiocarbon age distributions) are generally notable for their irregular shape and multi-

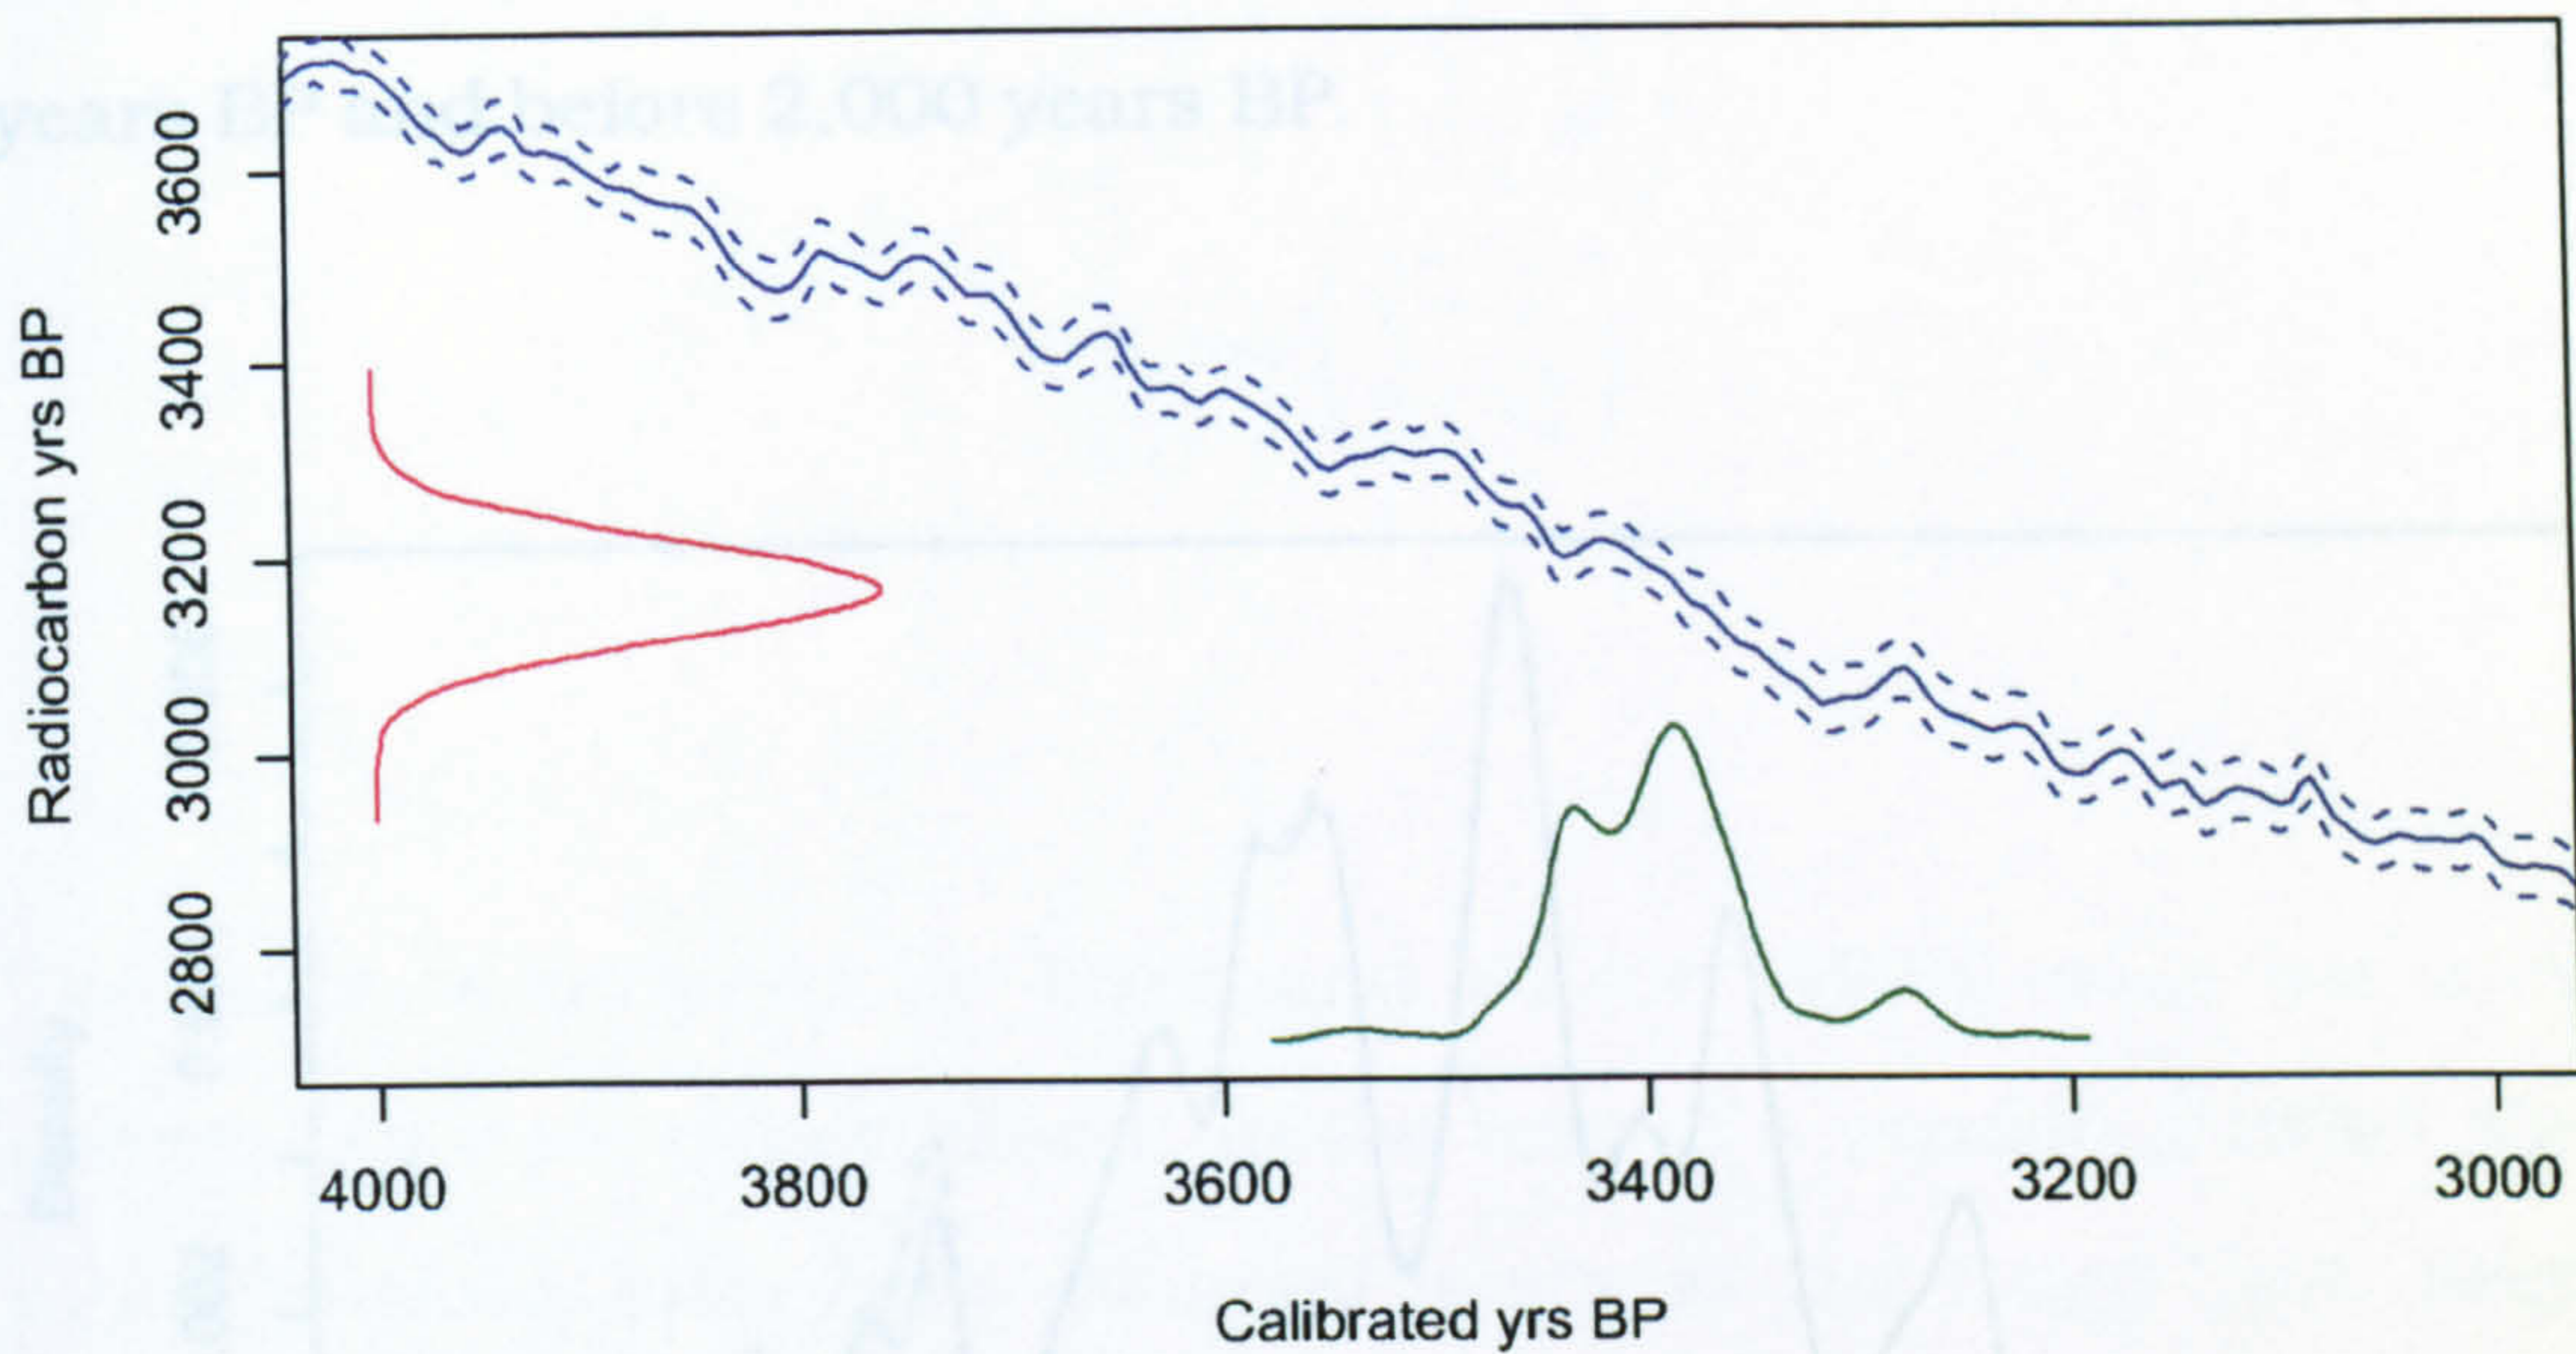


Figure 4.1: The process of calibration - the calibration curve is shown in blue with dotted error lines, the uncalibrated date in red and the calibrated date in dark green.

modality. It is probably for this reason that many RSL curves which feature radiocarbon dates ignore the uncertainty in the dates. Even when dates with errors are produced, they are often reported as continuous ranges rather than highest posterior density regions which identify the probability ranges more clearly.

The Humber LOIS data set contains over 100 single radiocarbon dates. Additionally, there are 4 single archaeological radiocarbon dates in the Humber. These were all calibrated as described above. As an illustration, the densities have been summed in figure 4.2. This is designed to give an indication of the expected size of uncertainty if the radiocarbon dates were to be used as covariates. As can be seen from

the plot, the majority of dates lie between 3,000 and 7,000 years BP. There is comparatively little information on dates older than 10,000 years BP and before 2,000 years BP.

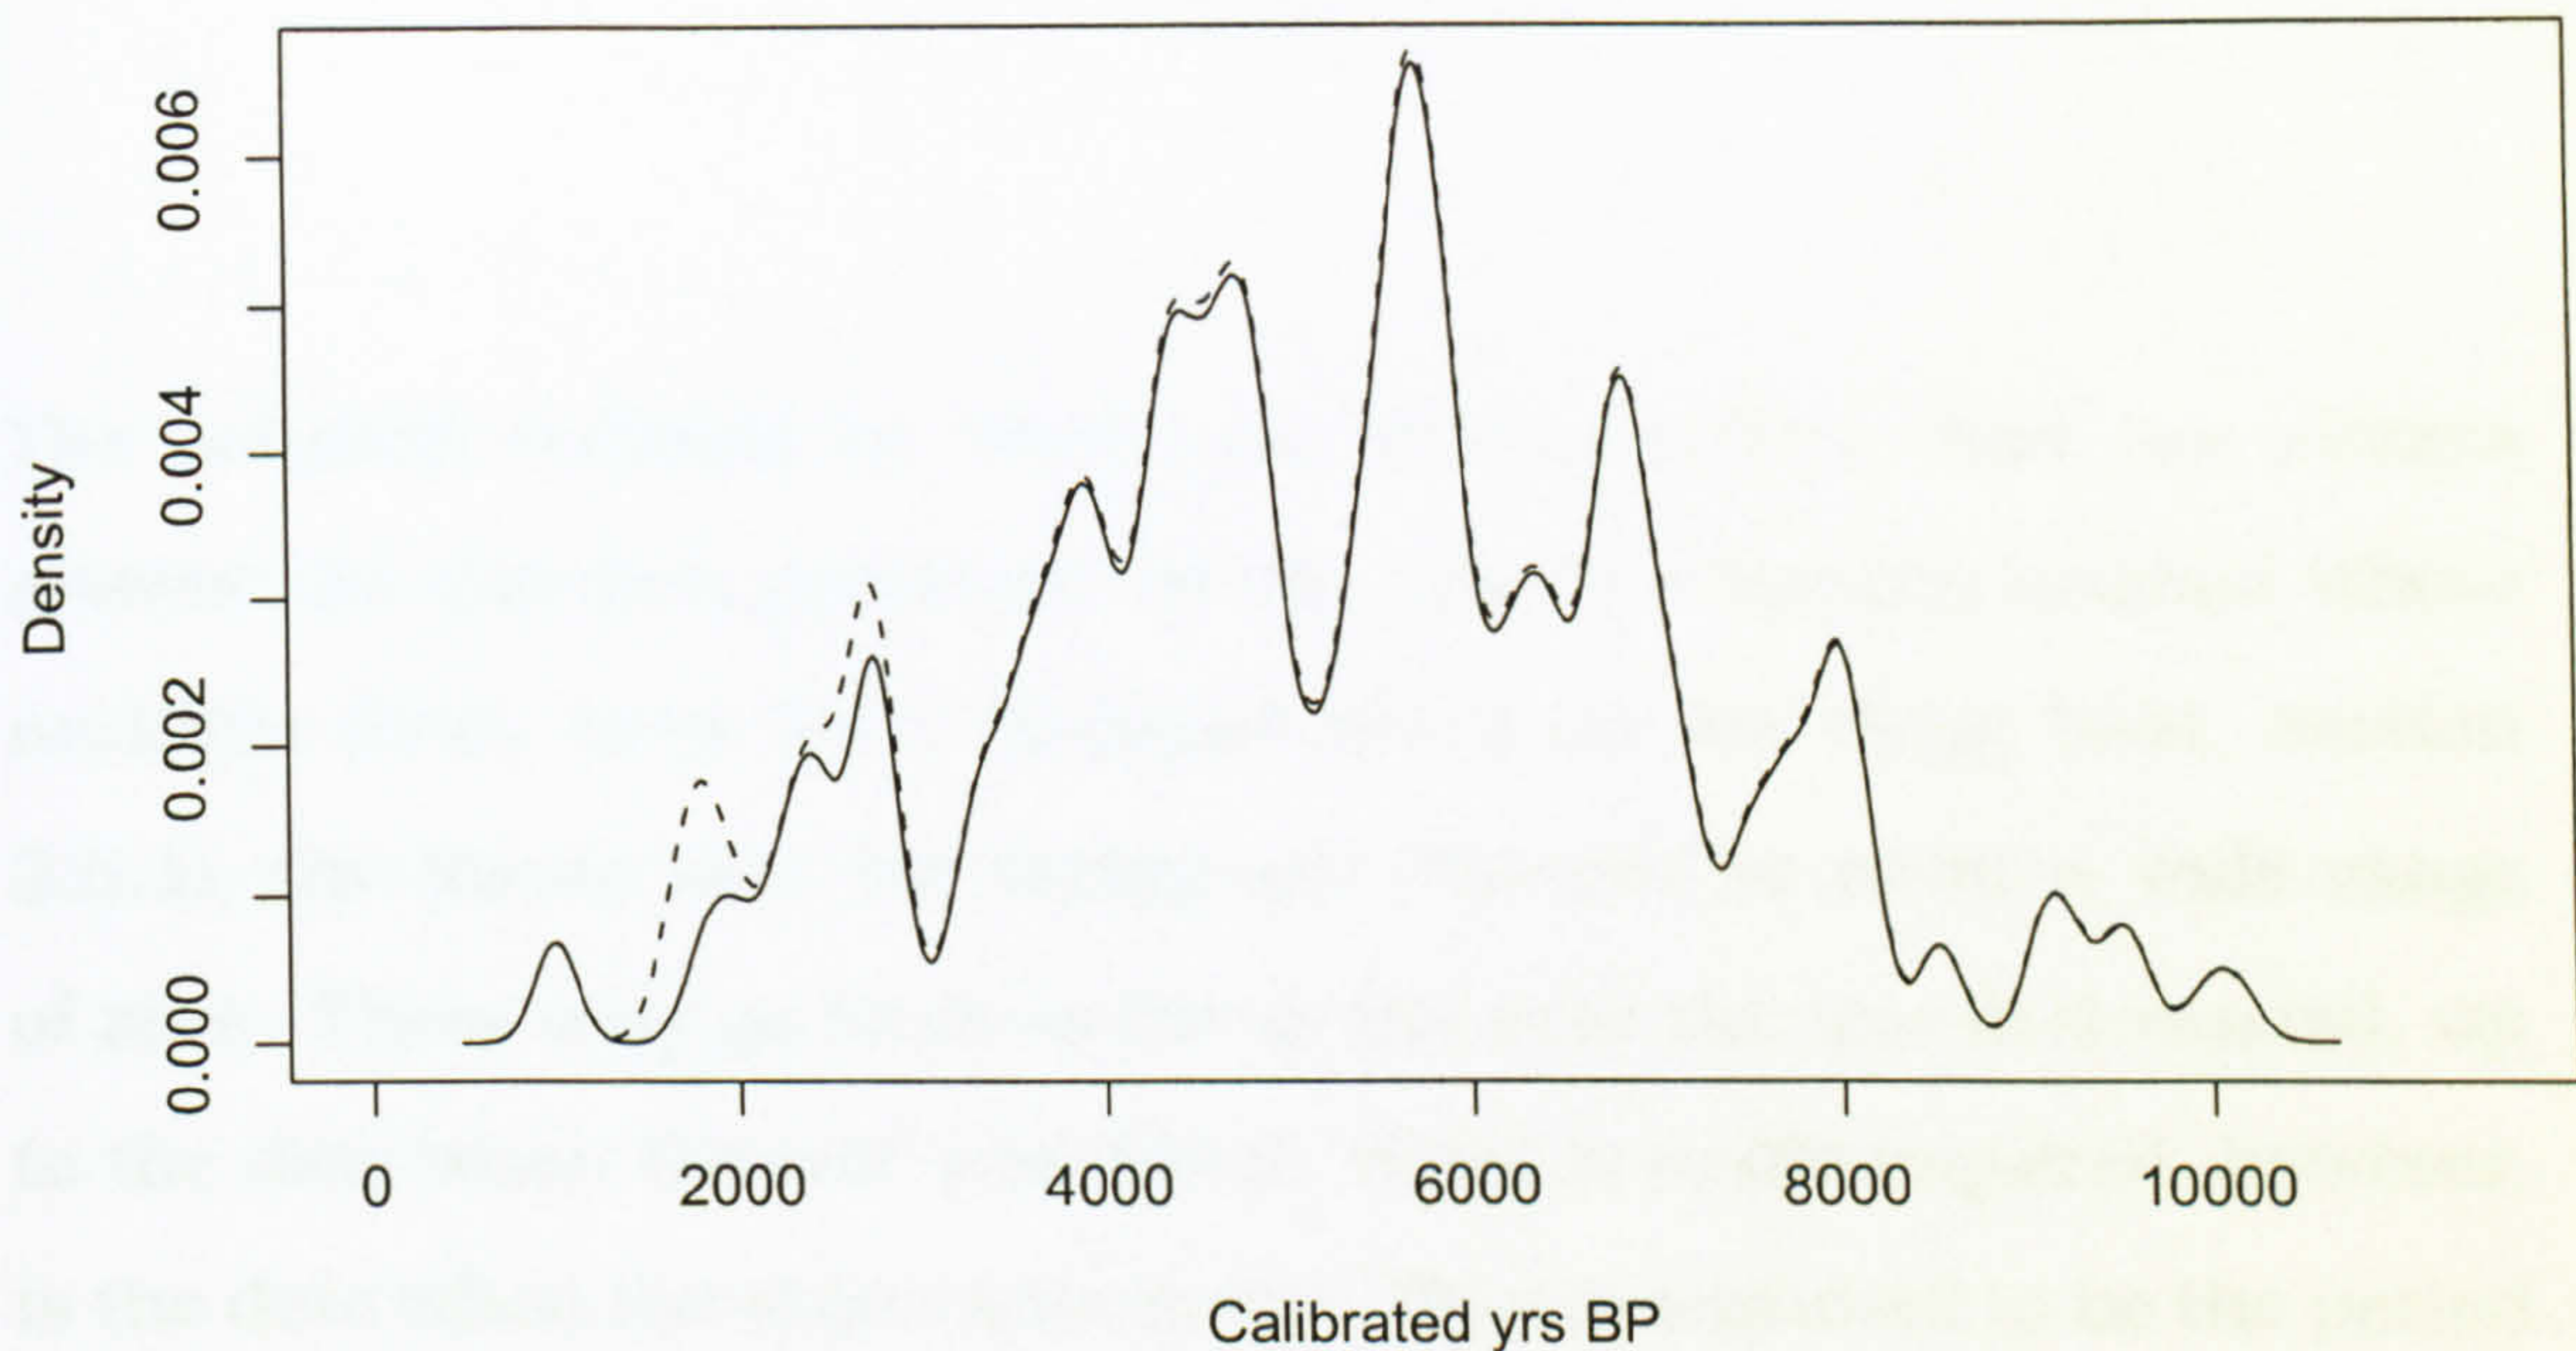


Figure 4.2: A plot of overlaid calibrated dates for the Humber. The addition of the archaeological dates for the Humber is shown with the dotted line.

Multiple radiocarbon ages

Often, multiple radiocarbon ages are obtained for the same artifact. This is certainly the case with the archaeological data outlined in section 3.4. The method to construct a posterior distribution in this case has traditionally followed the method of Ward and Wilson (1978) which is based on weighted means. Given a set of radiocarbon determinations, they assume a global distribution for the set of

radiocarbon dates and calibrate them as a pooled mean and variance:

$$x_p = \left(\sum_{i=1}^n \frac{x_i}{\sigma_i^2} \right) / \left(\sum_{i=1}^n \frac{1}{\sigma_i^2} \right), \quad (4.5)$$

$$V(x_p) = \left(\sum_{i=1}^n \frac{1}{\sigma_i^2} \right)^{-1}. \quad (4.6)$$

The solution outlined by Ward and Wilson (1978) does not always answer the question proposed. In the case of a wooden artefact where multiple dates have been obtained (such as the Brigg boat, section 3.6.1), the pieces sent for dating are *expected* to cover a wide range of ages. These may go back as far as the year the tree first existed, up to the date when the tree was felled. What is really required, however, is the date when the object was in use. This is assumed to be the period almost immediately after the tree was felled.

Such a problem is dealt with by Buck *et al.* (1996). As before, a set of radiocarbon determinations are given in the form $x_1 \pm \sigma_1, x_2 \pm \sigma_2, \dots, x_n \pm \sigma_n$. The calendar date of each sample is denoted by $\theta_i, i = 1, \dots, n$. A prior distribution is given to the set of calendar dates corresponding to the birth and death of the wood. If there is no reason to believe that dates are more likely to appear in one part of this range than another then a Uniform prior distribution is appropriate:

$$\theta_i | \alpha_\theta, \beta_\theta \sim U(\alpha_\theta, \beta_\theta) \quad (4.7)$$

β_θ now corresponds to the date of interest; when the wood was felled

and used in the artefact. Prior distributions can also be specified for the hyper-parameters:

$$\begin{aligned}
 p(\alpha_\theta, \beta_\theta) &\propto I(\alpha_\theta, \beta_\theta) \\
 &= \begin{cases} 1 & \text{if } \alpha_\theta > \beta_\theta \\ 0 & \text{otherwise} \end{cases} \quad (4.8)
 \end{aligned}$$

and also for θ itself:

$$I_{(\alpha_\theta, \beta_\theta)}(\theta_i) = \begin{cases} 1 & \text{if } \alpha_\theta \geq \theta_i \geq \beta_\theta \\ 0 & \text{otherwise} \end{cases} \quad (4.9)$$

Thus a posterior for the set of parameters can be created given the dates and errors:

$$\begin{aligned}
 p(\theta, \alpha_\theta, \beta_\theta | \mathbf{x}, \sigma^2) &\propto I(\alpha_\theta, \beta_\theta)(\alpha_\theta - \beta_\theta)^{-1} \\
 &\times \prod_{i=1}^n \exp \left\{ -\frac{(x_i - r(\theta_i))^2}{2(\sigma^2(\theta) + \sigma_i^2)} I_{(\alpha_\theta, \beta_\theta)}(\theta_i) \right\} \quad (4.10)
 \end{aligned}$$

The estimated posterior distribution of the date of the artefact is now that of β_θ . It is noted in Buck *et al.* (1996) that these date distributions are more likely to be unimodal than traditional calibrated radiocarbon dates, but unfortunately tend to have longer tails.

In the Humber Wetlands Project data sets, there are four such sets of multiple radiocarbon dates. The one which most closely mirrors that of the situation above is that of Willoughby-13. A histogram of the dates for Willoughby-13 is shown in figure 4.3. Here, 4 radiocarbon dates are supplied; 3220 ± 59 , 3210 ± 50 , 3180 ± 50 and 3010 ± 80 , providing


information about calendar dates θ_1 , θ_2 , θ_3 and θ_4 . Further information is also available as the site is known to date between the late Mesolithic to the Roman periods. This provides further bounds on the parameter of interest β_θ . The plot shows that the site was estimated to have been used at some point between 1,600 and 3,450 cal years BP. Where archaeological radiocarbon determinations are present (whether multiple or single), it is the parameter β_θ which is carried forward as the date estimate for that particular location and elevation.

4.2.2 Period data

Many of the data points created from the Humber Wetlands Project have not been radiocarbon dated. However, some of these can be attributed to various time periods through the expert knowledge of the archaeologists. The periods for which archaeological data are available are shown in table 4.1. In the statistical models that follow, the dates are taken to follow a uniform distribution over the period given. It is clear that for very old dates, the variability is large.

Name	Dates	Length
Mesolithic	10000 - 4000 BC	6000
Neolithic	4000 - 2500 BC	1500
Bronze age	2500 - 600 BC	1900
Iron age	600 BC - 43 AD	644
Roman	43 - 410 AD	367
Early middle ages	410 - 1066 AD	656
Mediaeval	1066 - 1540 AD	474
Post-mediaeval	1540 AD - present	465

Table 4.1: Table of archaeological periods



4.3.1 Individual index points

Figure 4.3: Posterior densities for the worked wood at Willoughby (section 3.6.4). The first four panels are the radiocarbon dates, the bottom two for the bounding parameters α_θ and β_θ .

There are seven such dates from the Outer Humber; Brigham-6, Marsh Chapel-2, Ingoldmells-6, Ingoldmells-12, Adlingfleet-2, Wheldrake-2 and Kelk-6. Their contribution to the overall date uncertainty is shown in figure 4.4.




Figure 4.4: Overlaid dates for the Humber. The red line shows the effect of the inclusion of the archaeological period dates.

4.3 Graphical methods

4.3.1 Individual index points

The error in the data points in both age and elevation is an important factor to be considered. In many RSL studies (Shennan and Horton, 2002; Shennan *et al.*, 2000a) whilst the error is recognised, it is not used in any meaningful way in the production of RSL curves. In order to correct this, it is important to gain an insight as to the distributional shape of an individual index point. A plot of the joint density (in elevation and age) for an individual index point with radiocarbon date 3170 ± 50 and elevation -2.91m (standard error 0.011m) is shown in figure 4.5. The age uncertainty is calculated by calibrating a

radiocarbon date, whilst elevation uncertainty is assumed Normally distributed as discussed in section 2.4.2. In this thesis, the elevation uncertainty and the age uncertainty are both assumed known. It can be seen from figure 4.5 that the joint distribution can be multi-modal and highly irregular, as a direct consequence of the radiocarbon calibration curve.

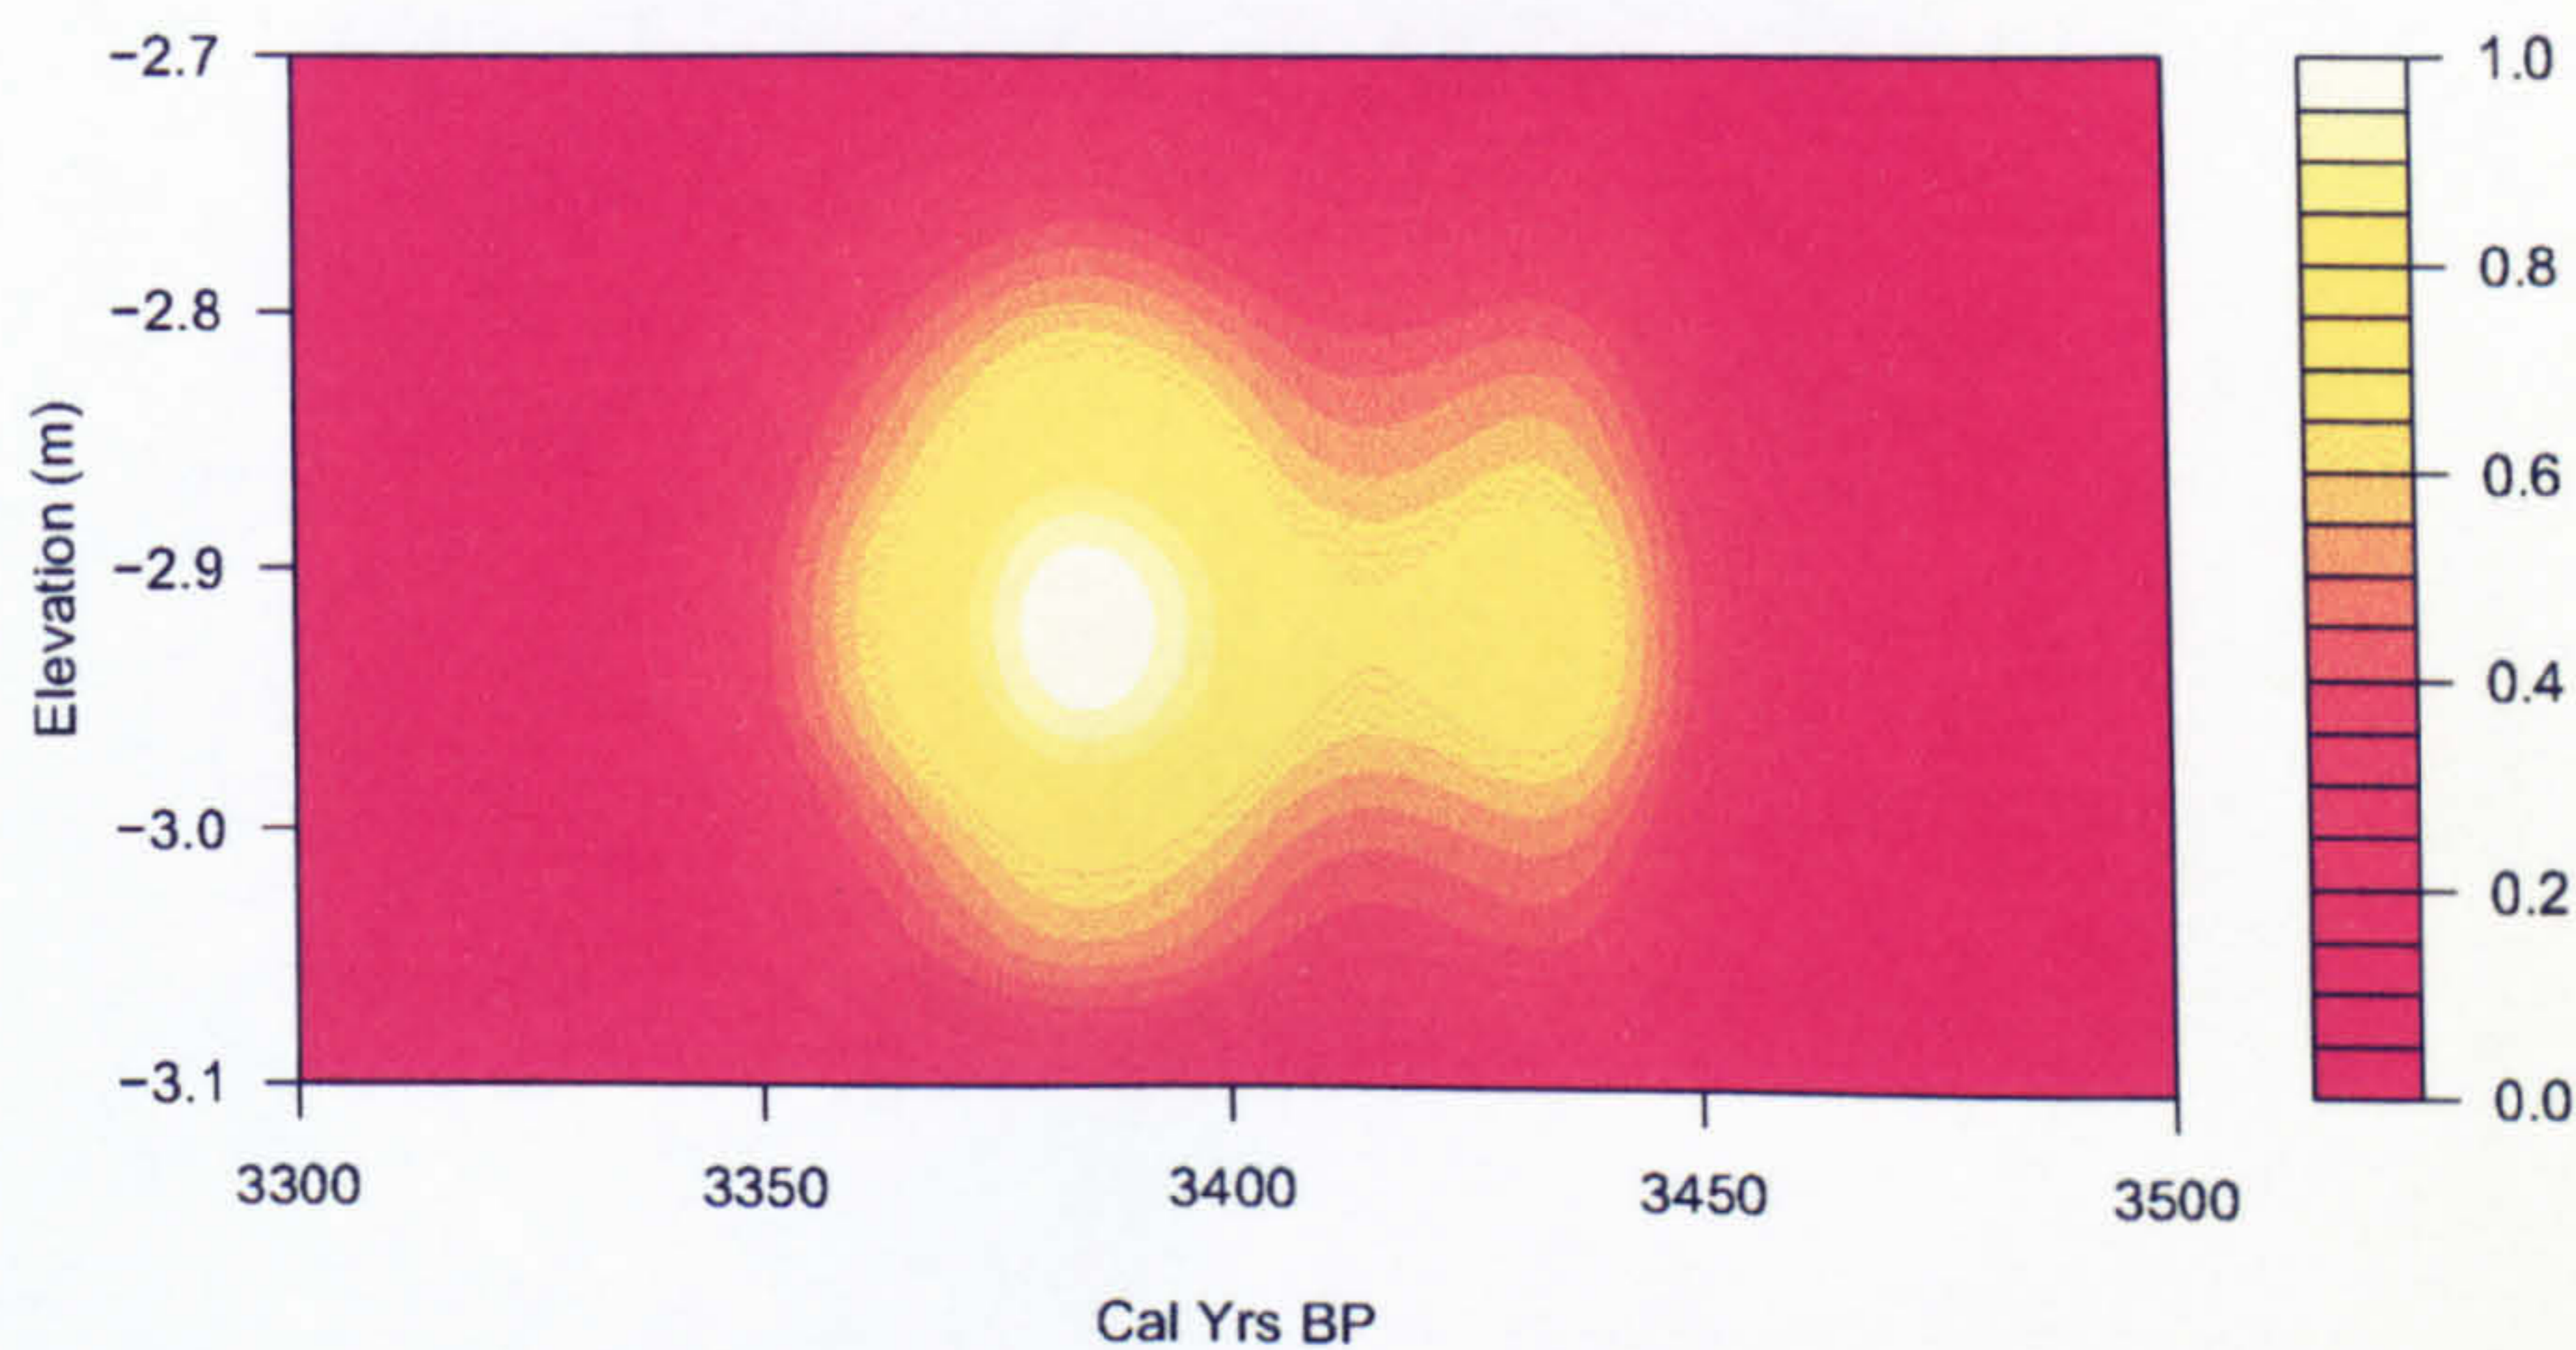


Figure 4.5: Joint distribution of elevation and age for a single index point.

4.3.2 Plotting data

The traditional method of plotting RSL data has been on an age-altitude graph. These graphs ignore the location facet of the data. The error in both age and altitude is usually presented with either a "+" symbol (eg Shennan and Horton, 2002) or an unfilled box (eg Devoy, 1982), as shown in figure 2.5 for the South East of England. In either case, the bi-directional error is represented by the vertical or horizontal length of the box or plus symbol. The box method is preferred in this thesis

as it more adequately displays the nature of the error; in particular the plausibility of the true height/age being located towards the edge of a box which is not adequately displayed when using a plus symbol. A superior method would be to use ellipses to show a 95% interval. This is not used here due to the problems in computing such ellipses. An age-altitude plot of the Humber data using the box method is shown in figure 4.6.




Figure 4.6: An age-altitude plot of the Humber basal index points (BIP), intercalated index points (IIP), limiting points (LP), older limiting points (LP2) and similar archaeological data with 95% errors presented with error boxes.

4.4 Cartographic representation

An important aspect of the data is that of the location of the data points. There are a number of ways to present this type of data graphically. In this thesis, many of the locations are given in Latitude/Longitude pairs. However, these are merely angles relating to the equator and the Greenwich meridian. The distance between degrees latitude and longitude is dependent on the location of the data points. To obtain real differences in suitable units a projection is required. For more details see, for example, Banerjee *et al.* (2004).

This section considers three different methods for measuring distances between data points. They are considered in the context of two data points which are the most distant. For clarity their decimal locations in degrees longitude have been truncated. Point 1 lies at latitude 53.2° , longitude 0.3° ; point 2 at latitude 53.6° , longitude -0.3° . The three methods considered are (in order of complexity):

1. Using the latitude and longitude with Euclidean distance: assuming one degree of latitude represents 111.23km and one degree of longitude represents 71.70km. These are approximate results for latitude 50° .
2. Using standard Ordnance Survey (OS) 6 digit grid references with Euclidean distance.
3. Using the Mercator projection and calculating the arc distance between the two data points.

Method 1 gives an estimated distance of 103.28km. Method 2

requires transformation into OS coordinates. Now, point 1 becomes (553729,369460) and point 2 (466274,411957). This is now given in x - y pairs where an increment in x or y of 1 corresponds to a single metre. The Euclidean distance is now 97.23km. Finally, the arc separation is calculated from:

$$D = R \arccos[\cos(\theta_1) \cos(\theta_2) \cos(\lambda_1 - \lambda_2) + \sin(\theta_1) \sin(\theta_2)], \quad (4.11)$$

where R is the radius of the Earth in km (taken here to be 6371), θ_i is the latitude of point i in radians and λ_i is the longitude of point i in radians. This gives a distance of 96.98km.

Method 1 differs most from methods 2 and 3. This is largely due to the changing distance between degrees at different latitudes and longitudes. Methods 2 and 3, however, produce highly similar results, with method 2 requiring much less complexity. For this reason, method 2 was used for calculating all distances between sites.

4.5 Temporal analysis of sea-level curves

Statistical approaches to drawing RSL curves have tended to either ignore the location aspect or to divide it into supposedly homogeneous subsets. For this reason, much of the focus has been on temporal sea-level change and often linear regression. This has the advantage of being relatively simple and quick to apply but ignores many of the salient features of the RSL curve and throws away much of the

importantly collected data. In particular, the use of regression models forces a globally smooth structure on the RSL curve which may be unrealistic. Similar ideas have problems with varying uncertainty and most attempts have no method of dealing with uncertainty in either date or age. Finally, intercalated index points and limiting points are ignored provided they do not interfere with the curve.

This section aims to extend these ideas under the assumption that the data are believed to be from a homogeneous set of locations (eg Metcalfe *et al.*, 2000). Two methods for dealing with smoothness are discussed: the degree of smoothness in RSL curves is important to the geologists and archaeologists (See debate in Fairbridge, 1961; Shepard, 1963). These methods are coupled with techniques for dealing with IIPs and LPs

4.5.1 Kernel smoothing

Kernel smoothing (eg Hastie and Tibshirani, 1995) is a technique for summarising a response in order to reduce the variability in the response measurement. In a kernel smoother, the response variable is averaged locally with weights prescribed by a kernel function. The value of the smoothed curve at any date, x , is given by:

$$S(x) = \frac{\sum_{i=1}^n d\left(\frac{x-x_i}{h}\right) y_i}{\sum_{i=1}^n d\left(\frac{x-x_i}{h}\right)} \quad (4.12)$$

where $d(z)$ is any decreasing function in $|z|$; it is often a Gaussian density function. The parameter h is known as the bandwidth. This

technique is also known as locally-constant kernel-weighted regression as when $h \rightarrow \infty$ the value of $S(x) \rightarrow \bar{y}$. The disadvantages of locally-constant regression are biased endpoint behaviour and the under- or over-estimate of stationary points.

A natural extension to the locally-constant kernel smoother is that of the locally-linear kernel smoother. The linear kernel smoother can partially overcome some of the locally-constant problems especially those involving endpoint behavior (Yu and Jones, 1998). Such a model can be fitted to produce a smoothed estimate:

$$S(x) = \arg \min_x \sum_{i=1}^n (Y_i - a - b(X_i - x))^2 d\left(\frac{x - X_i}{h}\right) \quad (4.13)$$

This minimisation will lead to an estimate of the mean line function. The models are highly dependent on the bandwidth. Results presented in this section are obtained using optimal bandwidth selectors. These are discussed further in section 7.2.1.

4.5.2 Endpoint estimation

One starting method for utilising the IIPs and LPs is to transform them such that they can be treated as basal index points (BIPs). It is possible then to use a simple curve fitting approach, such as that described above. The nature of endpoint estimates have been discussed in papers such as Robson and Whitlock (1964); Hall (1982); Cooke (1979). The approach used here is that of Robson and Whitlock (1964) whereby the upper endpoint of a distribution of sample size n can be estimated from

the simple function:

$$T = X_n + (X_n - X_{n-1}), \quad (4.14)$$

where X_k is the k -th largest order statistic of the sample. For the LPs, Equation 4.14 is adjusted to estimate the lowest point in the sample. This is calculated for a number of points along the horizontal axis. The transformed values are then used with the BIPs and a kernel regression in date is performed. It should be noted that there is a difference between estimating an end-point and that of a basal index point. It is therefore acknowledged that this approach is contentious.

The major disadvantage of this approach is the slope in which the neighbourhood of LPs are transformed. In cases where there is a gradient present, the function will tend only to use the values at one end of the neighbourhood as these will be the extreme order statistics. To counteract this problem, a kernel smoother is initially used on the LPs and IIPs. The residuals from these models can then be used as the order statistics and the centroid point can be adjusted down locally by the transformed residual value by using Equation 4.14. The opposite equivalent method is used on intercalated index points.

The local linear kernel-smoother with endpoint estimation was applied to the Humber data set. The elevation and age errors were ignored, and the median of the radiocarbon calibrated date used. The optimal bandwidth estimate for the Humber data set is 0.69. A plot of the data with the bandwidth estimate is shown in figure 4.7.

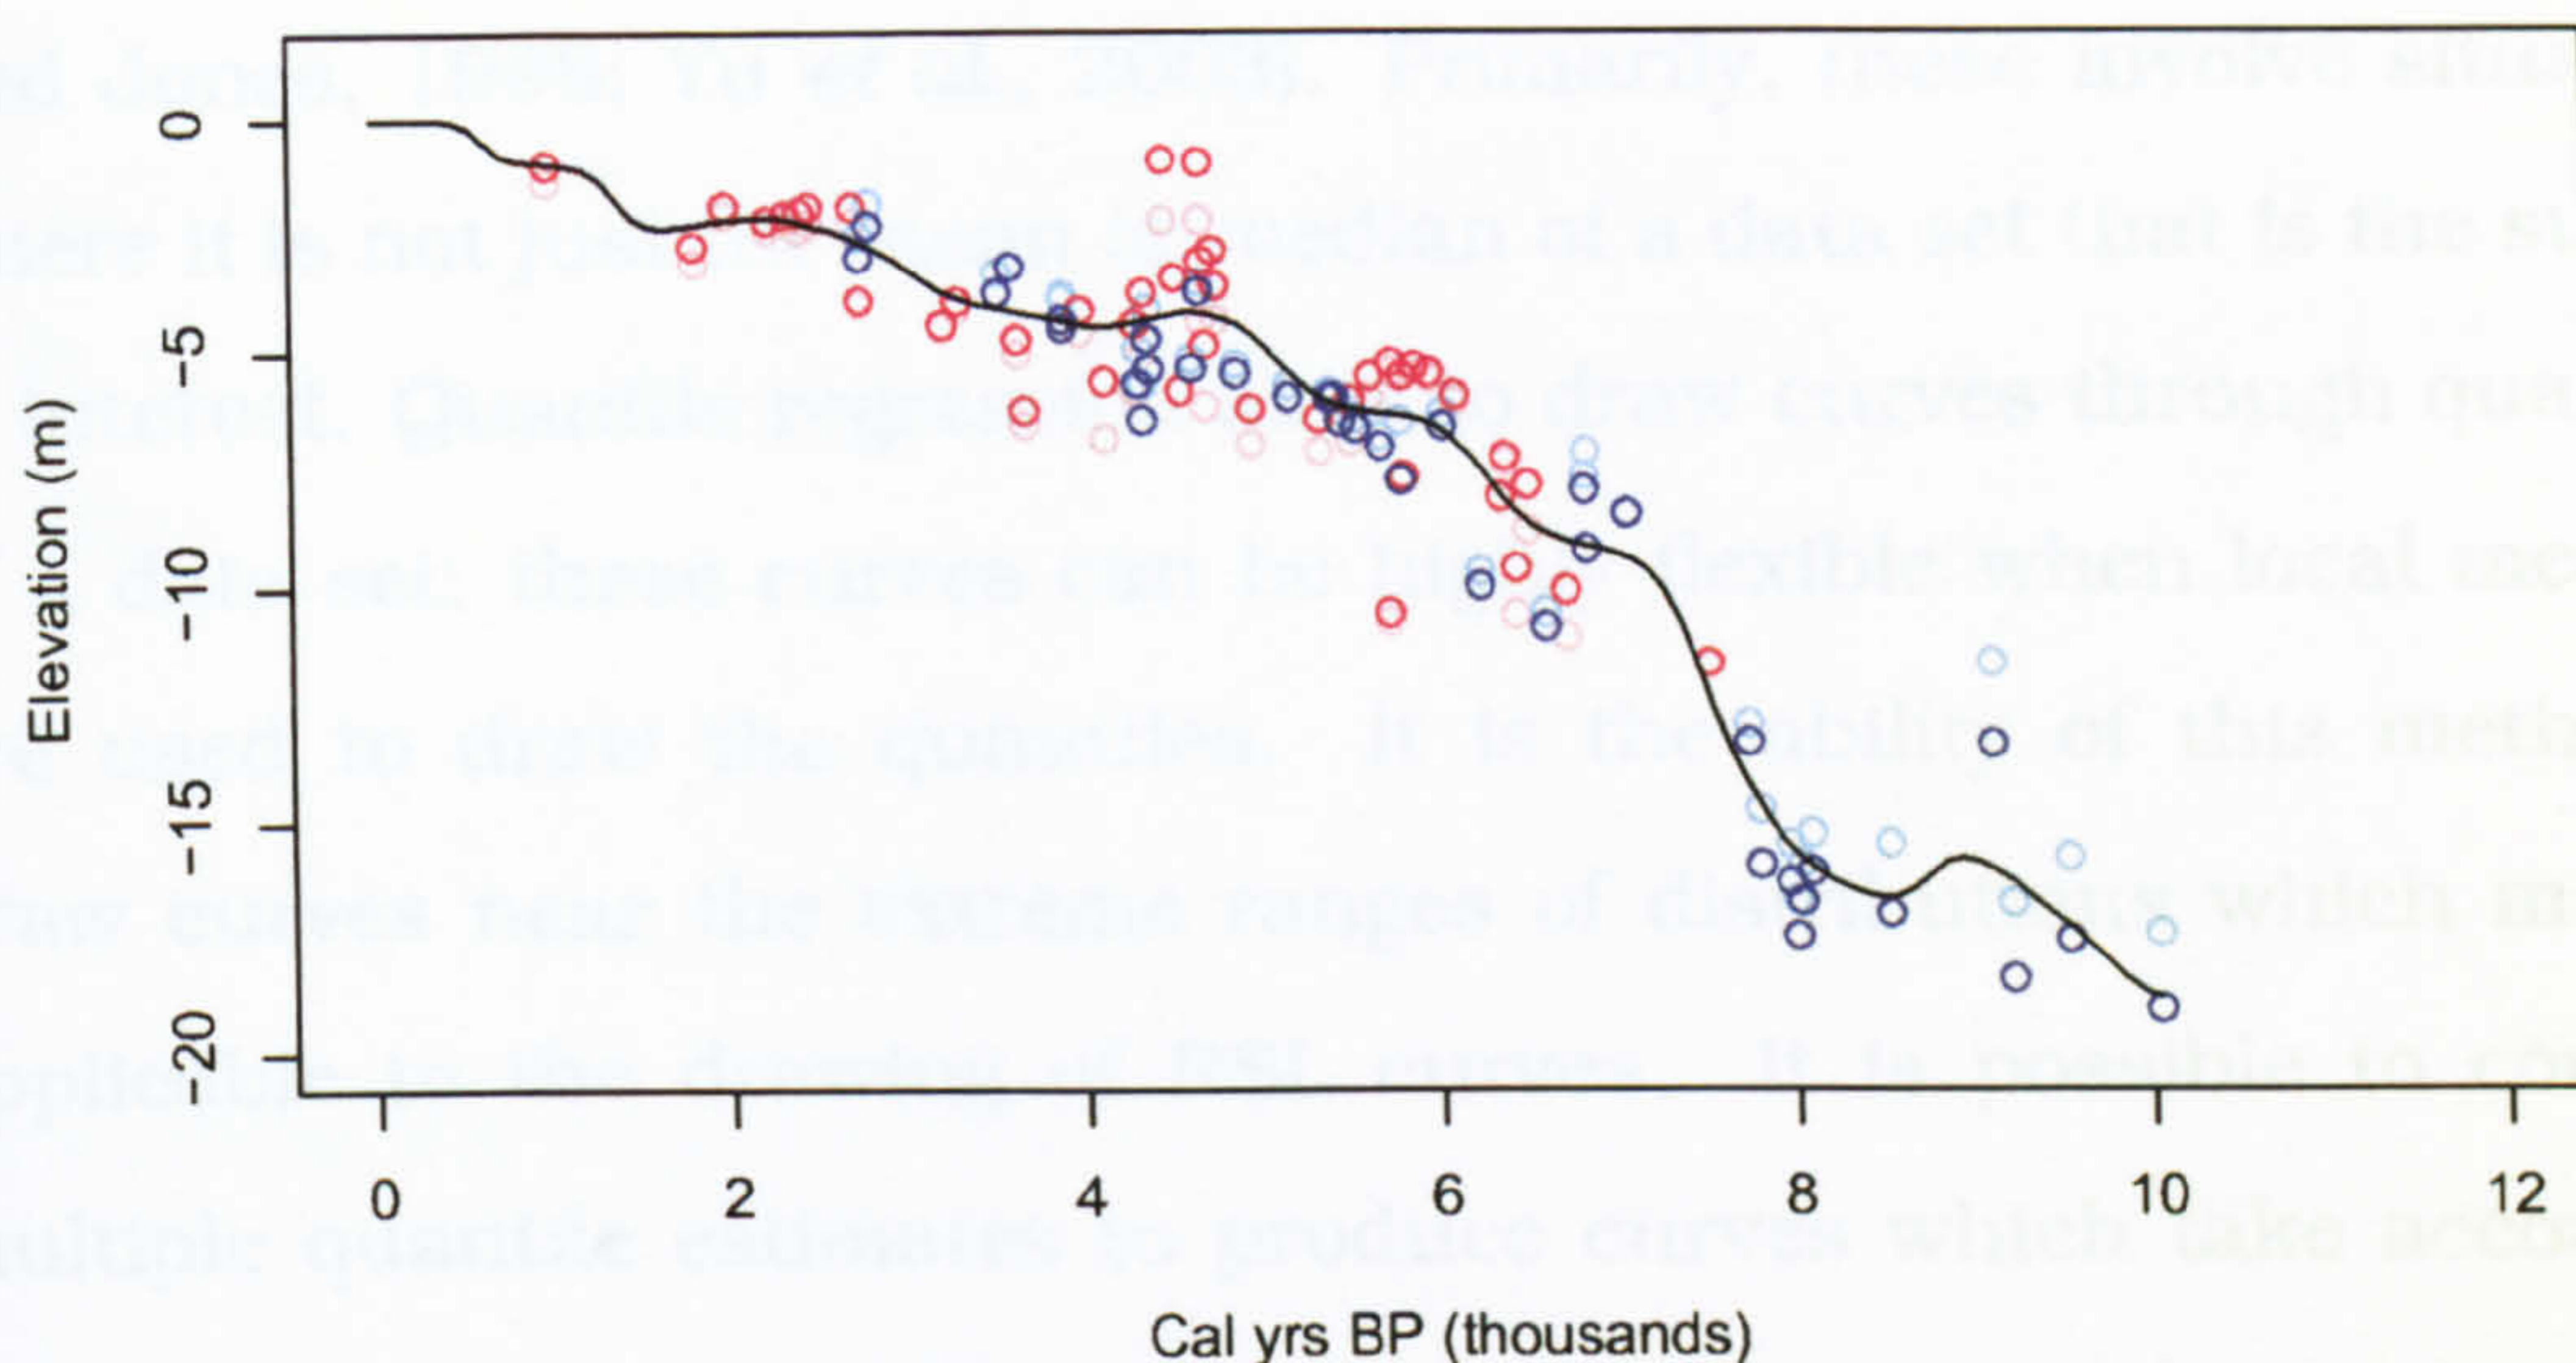


Figure 4.7: An estimated line for MSL via the local linear kernel smoothing/end-point estimation method with estimated optimal bandwidths. Basal index points are in black, adjusted limiting points in blue, adjusted intercalated index points in red. The original data points are shown in lighter colours.

It is clear from figure 4.7 that the set of limiting points aged 8000 years or older suggest a strong rise in estimated mean RSL. This rise continues to oscillate with decreasing amplitude up until around 0.5k years BP. It appears that the adjustment process has worked reasonably well as there is an apparently even spread in the data for the period where there are most data points. However, the older data points may require some more detailed treatment as geophysical models in the Humber show a rapid rise during this period (eg Shennan and Horton, 2002), not a levelling off.

4.5.3 Quantile Regression with check functions

Quantile regression is used in a wide number of applications (see Yu and Jones, 1998; Yu *et al.*, 2003). Primarily, these involve situations where it is not just the mean or median of a data set that is the subject of interest. Quantile regression aims to draw curves through quantiles of a data set; these curves can be highly flexible when local methods are used to draw the quantiles. It is the ability of this method to draw curves near the extreme ranges of distributions which make it applicable to the drawing of RSL curves. It is possible to combine multiple quantile estimates to produce curves which take account of different quantiles for different data sets simultaneously.

Locally-linear quantile regression can be considered an extension to that outlined in equation 4.13. Now, a quantile line function can be produced by minimising:

$$\sum_{i=1}^n \rho(Y_i - a - b(X_i - x)) d\left(\frac{x - X_i}{h}\right) \quad (4.15)$$

where ρ is known as a check function. To produce a median estimate the desired check function is $\rho(z) = |z|$. This idea can be extended further by making ρ a function of p , a fixed probability in the range $(0, 1)$. A check function of the form:

$$\rho_p(z) = \frac{|z| + (2p - 1)z}{2}. \quad (4.16)$$

will produce estimates that are quantiles, in relation to p , of the locally-linear estimated distribution at a point x . Using this approach (from Yu

and Jones, 1998), the height of the curve at all points is now given by \hat{a} - the equivalent of finding $\mathbb{E}[\rho_p(Y - a)|X = x]$.

To incorporate BIPs, LPs and IIPs, different values of p are required. This requires a mixture check function penalty for the different data points:

$$S(x) = \arg \min_x \sum_{j=1}^m \sum_{i=1}^n \rho_{p_j}(Y_i - a - b(X_i - x))d\left(\frac{x - X_i}{h}\right) \quad (4.17)$$

for m different groups of points.

Using the Humber data set, the values of the percentile, p , given in Equation 4.16 were changed for the different types of points. The value 0.5 (the median) was used for the basal index points, with 0.05 for LPs and 0.95 for IIPs. The graph in Figure 4.8 was produced using this method.

Figure 4.8 again shows a rapid sea-level rise for the period up to 7000 years BP when compared with figure 4.7 (produced using endpoint estimates with kernel regression). Again, a levelling off after this date is apparent. The differences in the two graphs appear to lie in the older limiting points and the use of bandwidth. The quantile produced graph shows much less respect to the extreme dates; there is little sign of a levelling off around 8000 years BP, and none of the oscillations.

A disadvantage of the quantile model is the arbitrary nature of the

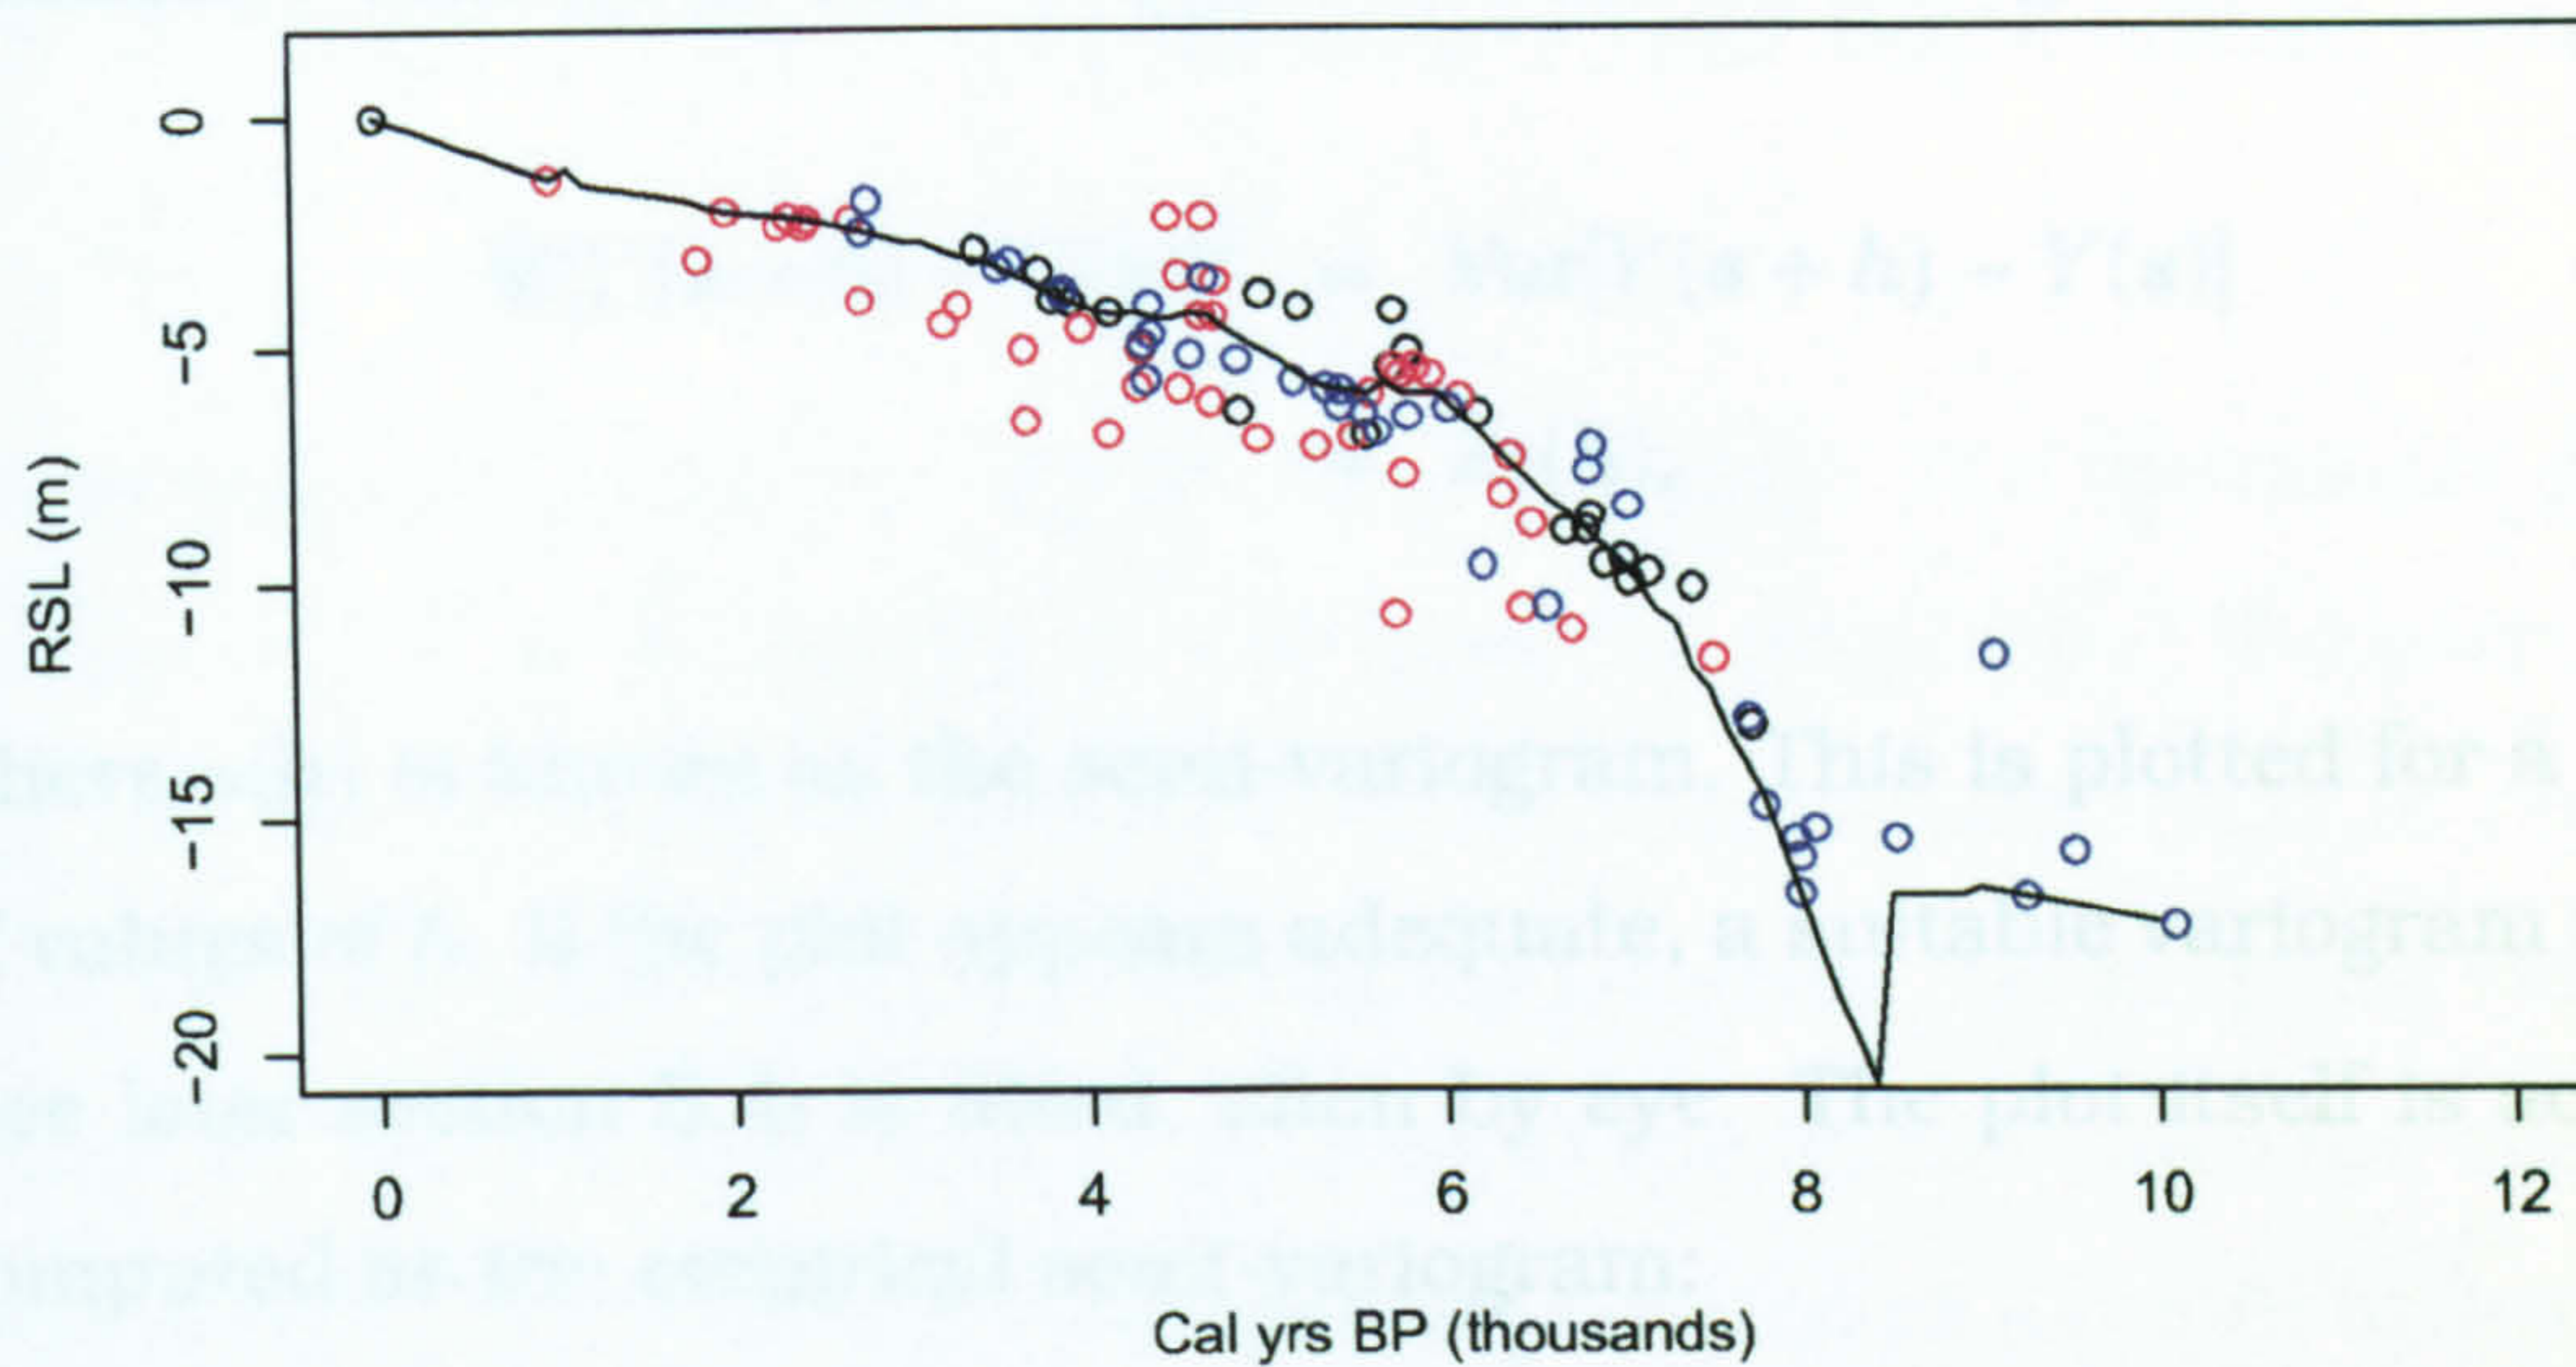


Figure 4.8: Graph of index points (black), intercalated index points (red) and limiting points (blue). Lines show local the quantile regression fit for percentiles on intercalated and limiting points of 5% and 95%.

values of p_j and the bandwidth as before. Yu and Moyeed (2001) adapt the model for Bayesian use. In a Bayesian framework, it would then be possible to place priors on the p_j which may produce more realistic line estimates. However, this method does not point forward to ways of including location, radiocarbon error and tendency measurements. Therefore the method was not taken forward.

4.6 Variograms

Variograms are traditionally used to fit correlation functions to stationary data sets. The aim of this section is to find mean trends which allow a good fit to a suitable correlation function. The estimated autocorrelation can then be used to improve inference on the desired

set of parameters. Formally, a process $Y(s)$ is defined as intrinsically stationary with $\mathbb{E}[Y(s+h) - Y(s)] = 0$, if the following holds:

$$\begin{aligned}\mathbb{E}[Y(s+h) - Y(s)]^2 &= \text{Var}[Y(s+h) - Y(s)] \\ &= 2\gamma(h),\end{aligned}\tag{4.18}$$

where $\gamma(h)$ is known as the semi-variogram. This is plotted for a range of values of h . If the plot appears adequate, a suitable variogram model (see later section 5.4) is fitted, often by eye. The plot itself is actually computed as the empirical semi-variogram:

$$\hat{\gamma}(t) = \frac{1}{2N(t)} \sum_{s_i, s_j \in N(t)} [Y(s_i) - Y(s_j)]^2,\tag{4.19}$$

where $N(t)$ is the set of pairs where $\|s_i - s_j\| = t$. In practice the observations are often 'binned' so that $\|s_i - s_j\| \approx t$.

4.6.1 Temporal variograms

The radiocarbon dates used to form the data set cause problems for temporal variogram estimation. Unlike many usual cases in spatio-temporal modelling, the time variable here is random and has a known distribution from which it can be sampled. In this section, the medians of the calendar dates were used in conjunction with a number of different means with which to identify overall temporal autocorrelation. Four methods were tried:

1. Subtracting a constant mean.
2. Subtracting a linear trend.

3. Subtracting a quadratic trend.

4. Subtracting a quantile smoothing fit like that outlined in section 4.5.3.

These were computed for the Humber Estuary. The results are shown in figure 4.9. It is possible to see from the graph that the correlation appears to be most well-defined when a quadratic trend is removed from the data. When a quantile trend is removed, no temporal variation remains. The residual sum of squares of the fitted values are shown in table 4.6.1. The values in this table are created as a result of fitting the Matérn correlation function (see section 5.4.3). The values show that the quantile fits best though the quadratic and linear are not too far behind.

Trend subtracted	Sum of Squares
Constant	82.16
Linear	2.74
Quadratic	0.79
Quantile	0.44

Table 4.2: Error sum of squares values for the Matérn variogram with various temporal trends removed.

4.6.2 Spatial variograms

Variograms were also used to assess any spatial dependence in the data set. As above, temporal trends were subtracted and the semi-variance examined in order to assess correlation at various lags. Possible temporal structures include:

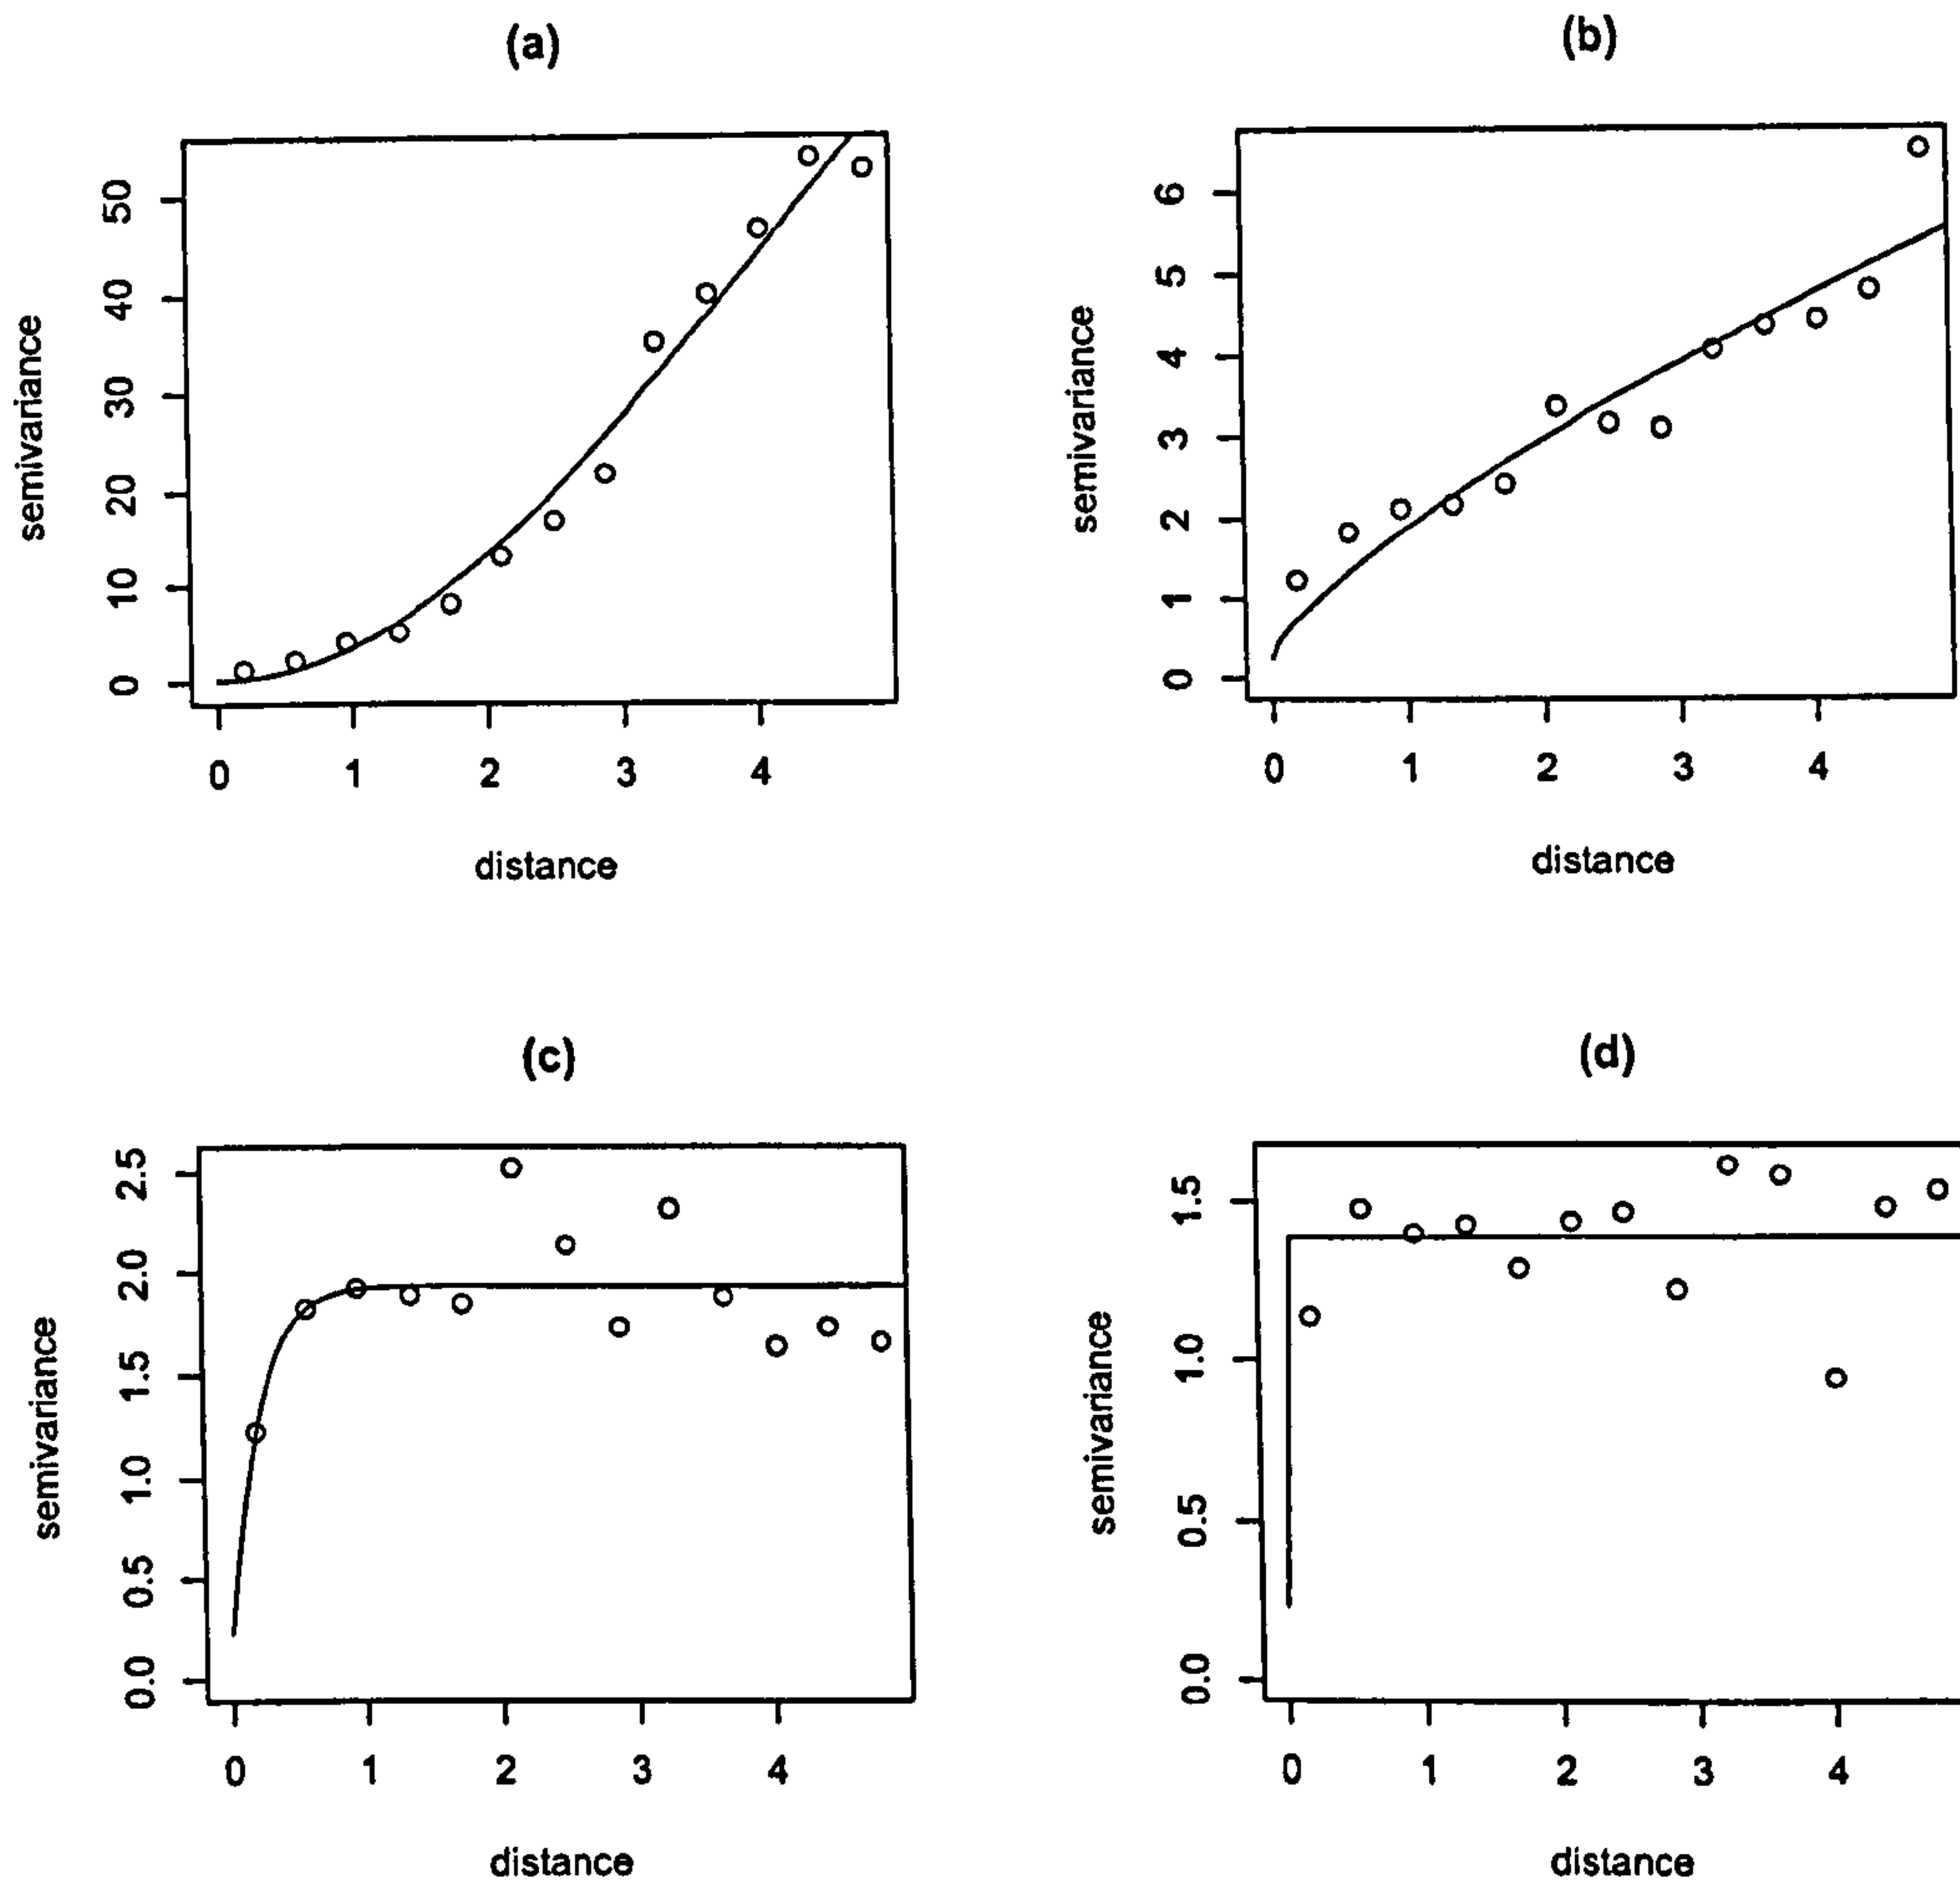


Figure 4.9: Temporal variograms for the Humber. Graph (a) has just the mean subtracted, (b) a linear trend, (c) a quadratic trend and (d) a multi-quantile local-linear fit. Distance is measured in thousands of years BP. A Matérn correlation function (outlined in section 5.4.3) was fitted to each variogram and is shown with the solid line.

1. A mean subtracted (independent of time).
2. A linear trend subtracted.
3. A quadratic trend subtracted.
4. A multi-quantile smooth subtracted.
5. A quantile smooth subtracted with a E-W trend.
6. A quantile smooth subtracted with E-W, N-S and an interaction term.
7. A quantile smooth subtracted with an extra variable indicating distance up-Estuary (where the start of the Estuary is deemed to be at Spurn Head).

The resulting variograms are shown in figure 4.10. The graph shows that including an extra location factor in the mean removes almost all spatial variation. It appears this is an effect that can be accounted for by using either an east-west trend or distance up-estuary as an explanatory variable. It also appears there may also be a slight north-south effect (with possible interaction) as this variogram also fits the data well.

A table of the fitted values is given in table 4.6.2. The table shows that the variogram fit for panel (a) in figure 4.10 is poor. Those involving a spatial aspect appear to fit slightly better than those involving a quadratic or linear trend.

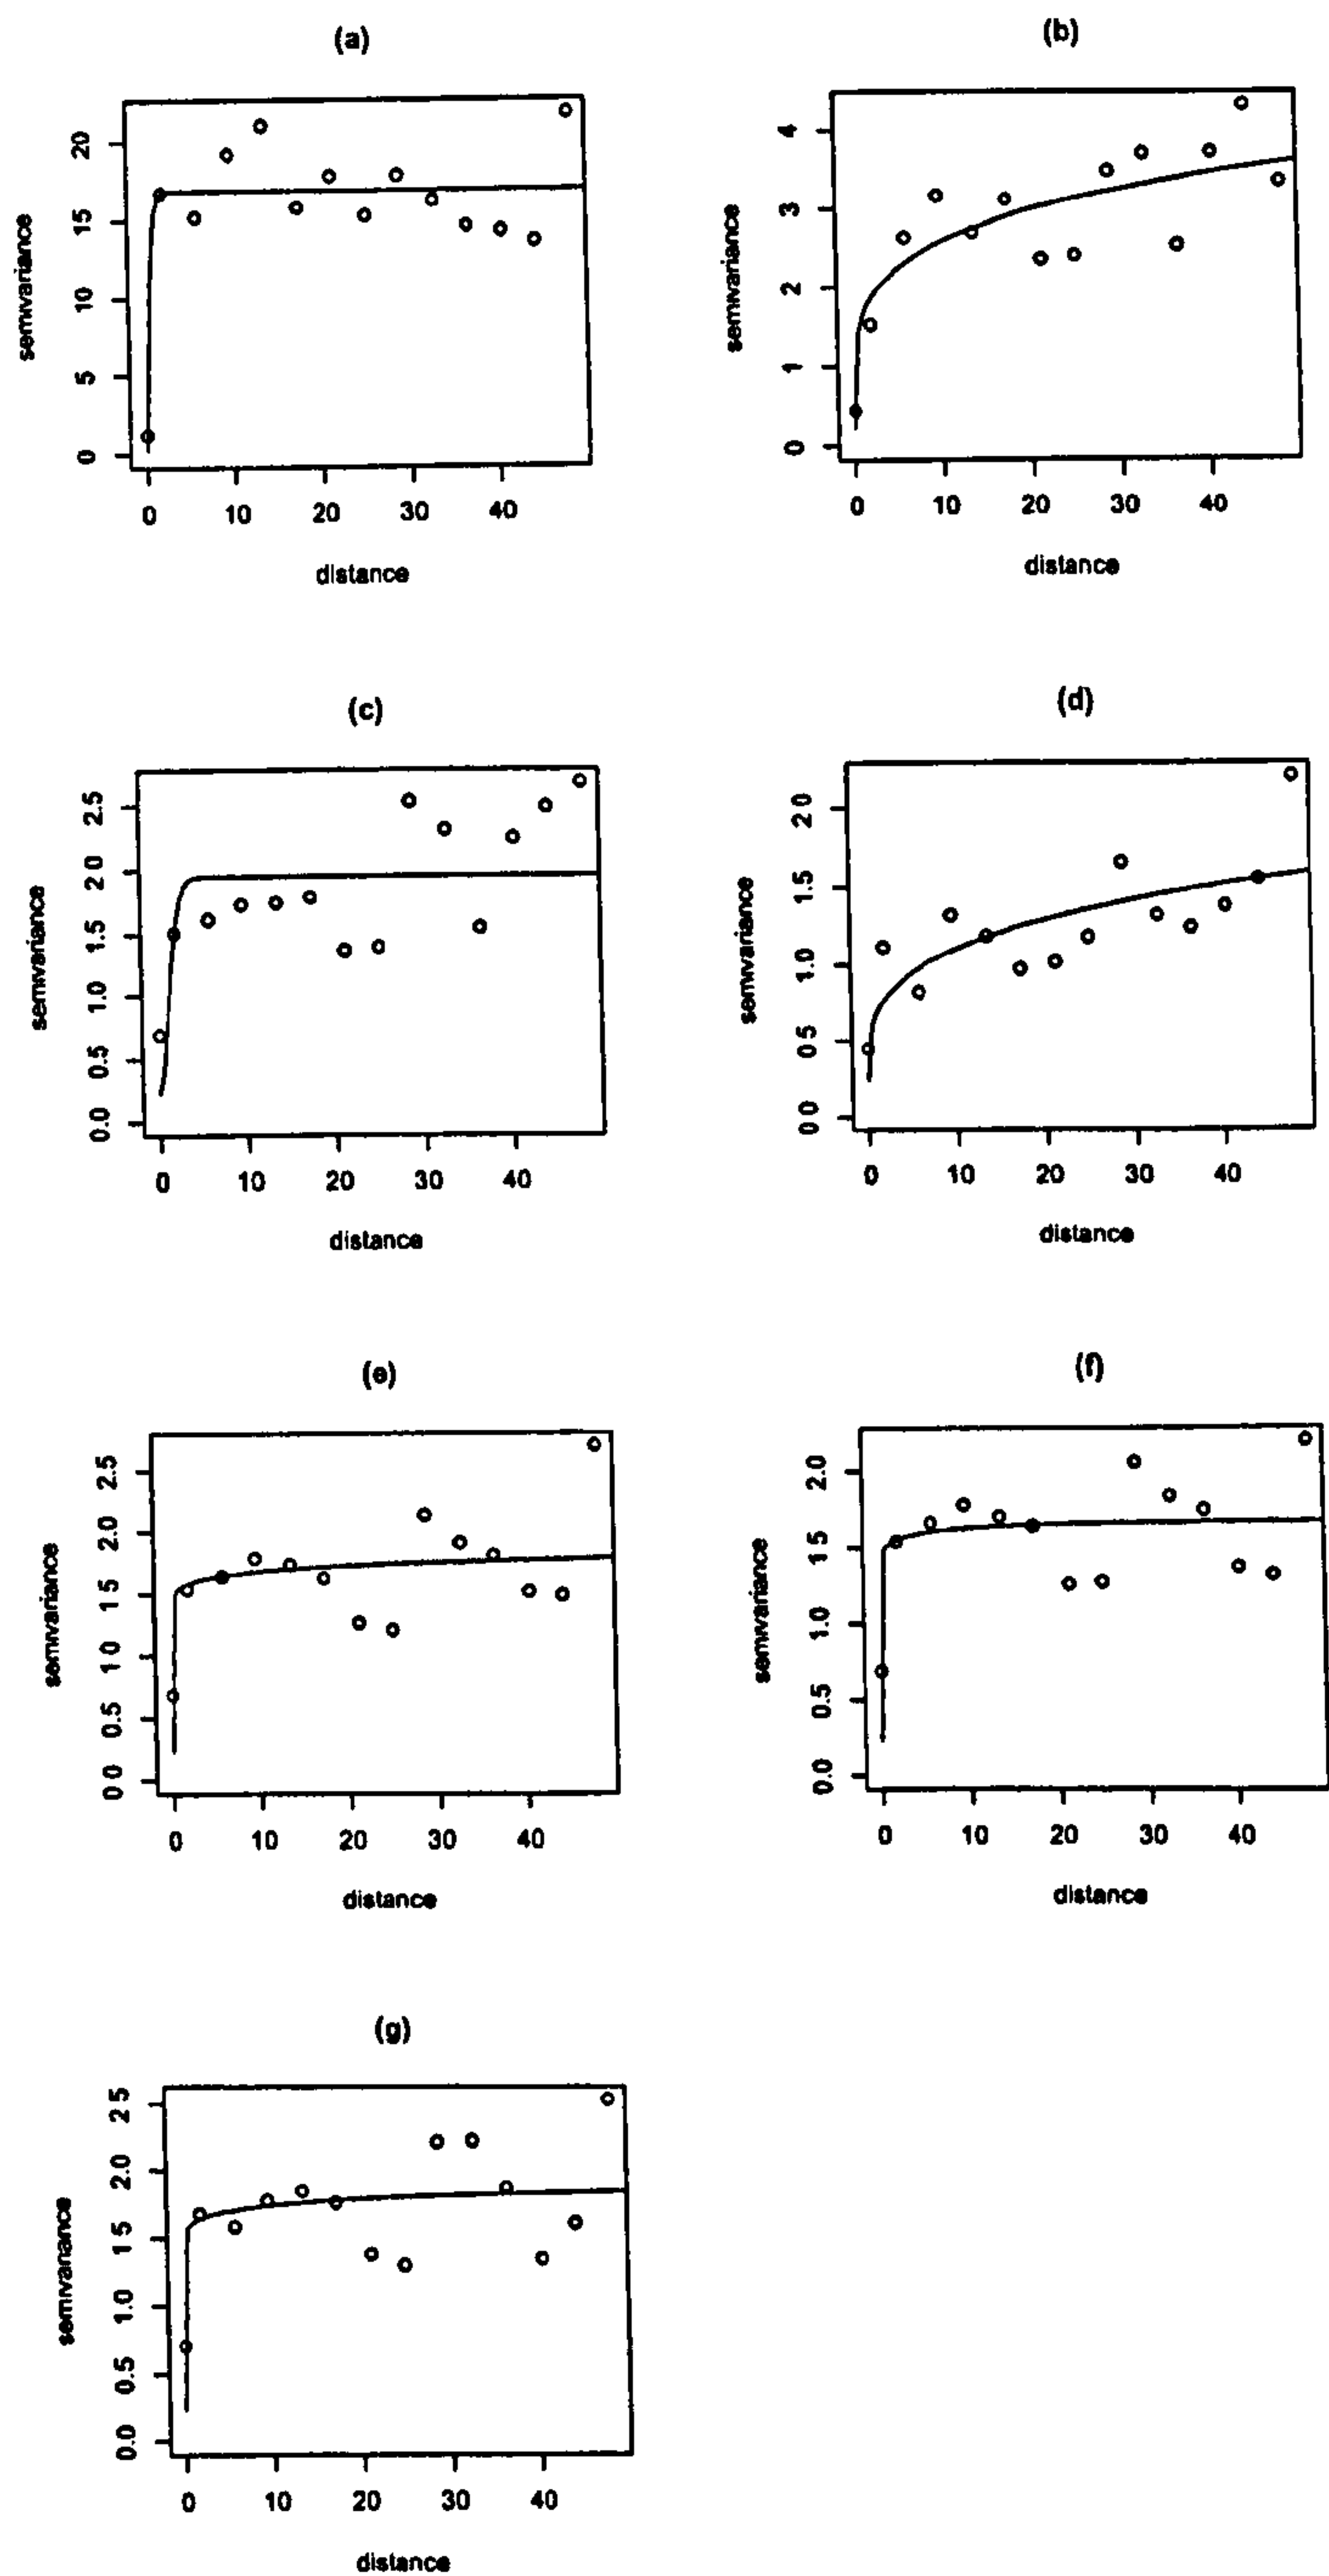


Figure 4.10: Spatial variograms with temporal trends subtracted. (a) is with a mean only removed, (b) is with a linear trend removed, (c) a quadratic trend, (d) a multi-quantile trend, (e) a multi-quantile trend with an E-W trend, (f) a multi-quantile trend with an E-W and N-S trend and (g) a multi-quantile trend with distance up-channel. A fitted Matérn correlation function is shown in the solid line. Distance is in kilometres.

Trend removed	Sum of Squares
Constant	80.37
Linear	3.23
Quadratic	2.58
Quantile	0.98
Quadratic + E/W	1.66
Quadratic + E/W and N/S	1.00
Quadratic + Dist up channel	1.56

Table 4.3: Error sums of squares for the Matérn fitted spatial variogram with various temporal trends removed.

4.7 Tendency data

The tendency part of the sea-level data set is perhaps the most awkward of the four pieces of information that are part of an individual index point. For details on the meaning of the tendency of an index point, see section 2.3.7 and in particular table 2.2. It should be noted that a positive or negative tendency does not mean that RSL was definitely rising or falling respectively, but it may mean that there is more of a chance that it may do so. An initial plot of the tendencies for the Humber are shown in figure 4.11.

It is clear from figure 4.11 that if the tendency information was taken at face value, the sea-level curve would not be very smooth! Indeed, such sea-level curves produced by relying on tendency information show many fluctuations (Gaunt and Tooley, 1974). Modern methods for looking at tendencies have focussed on its spatial aspect (Horton *et al.*, 2004). It is thought that if similar tendencies are grouped together in space-time they may provide evidence of RSL change. A graph of the location of the tendencies is shown in figure 4.12. This graph on its

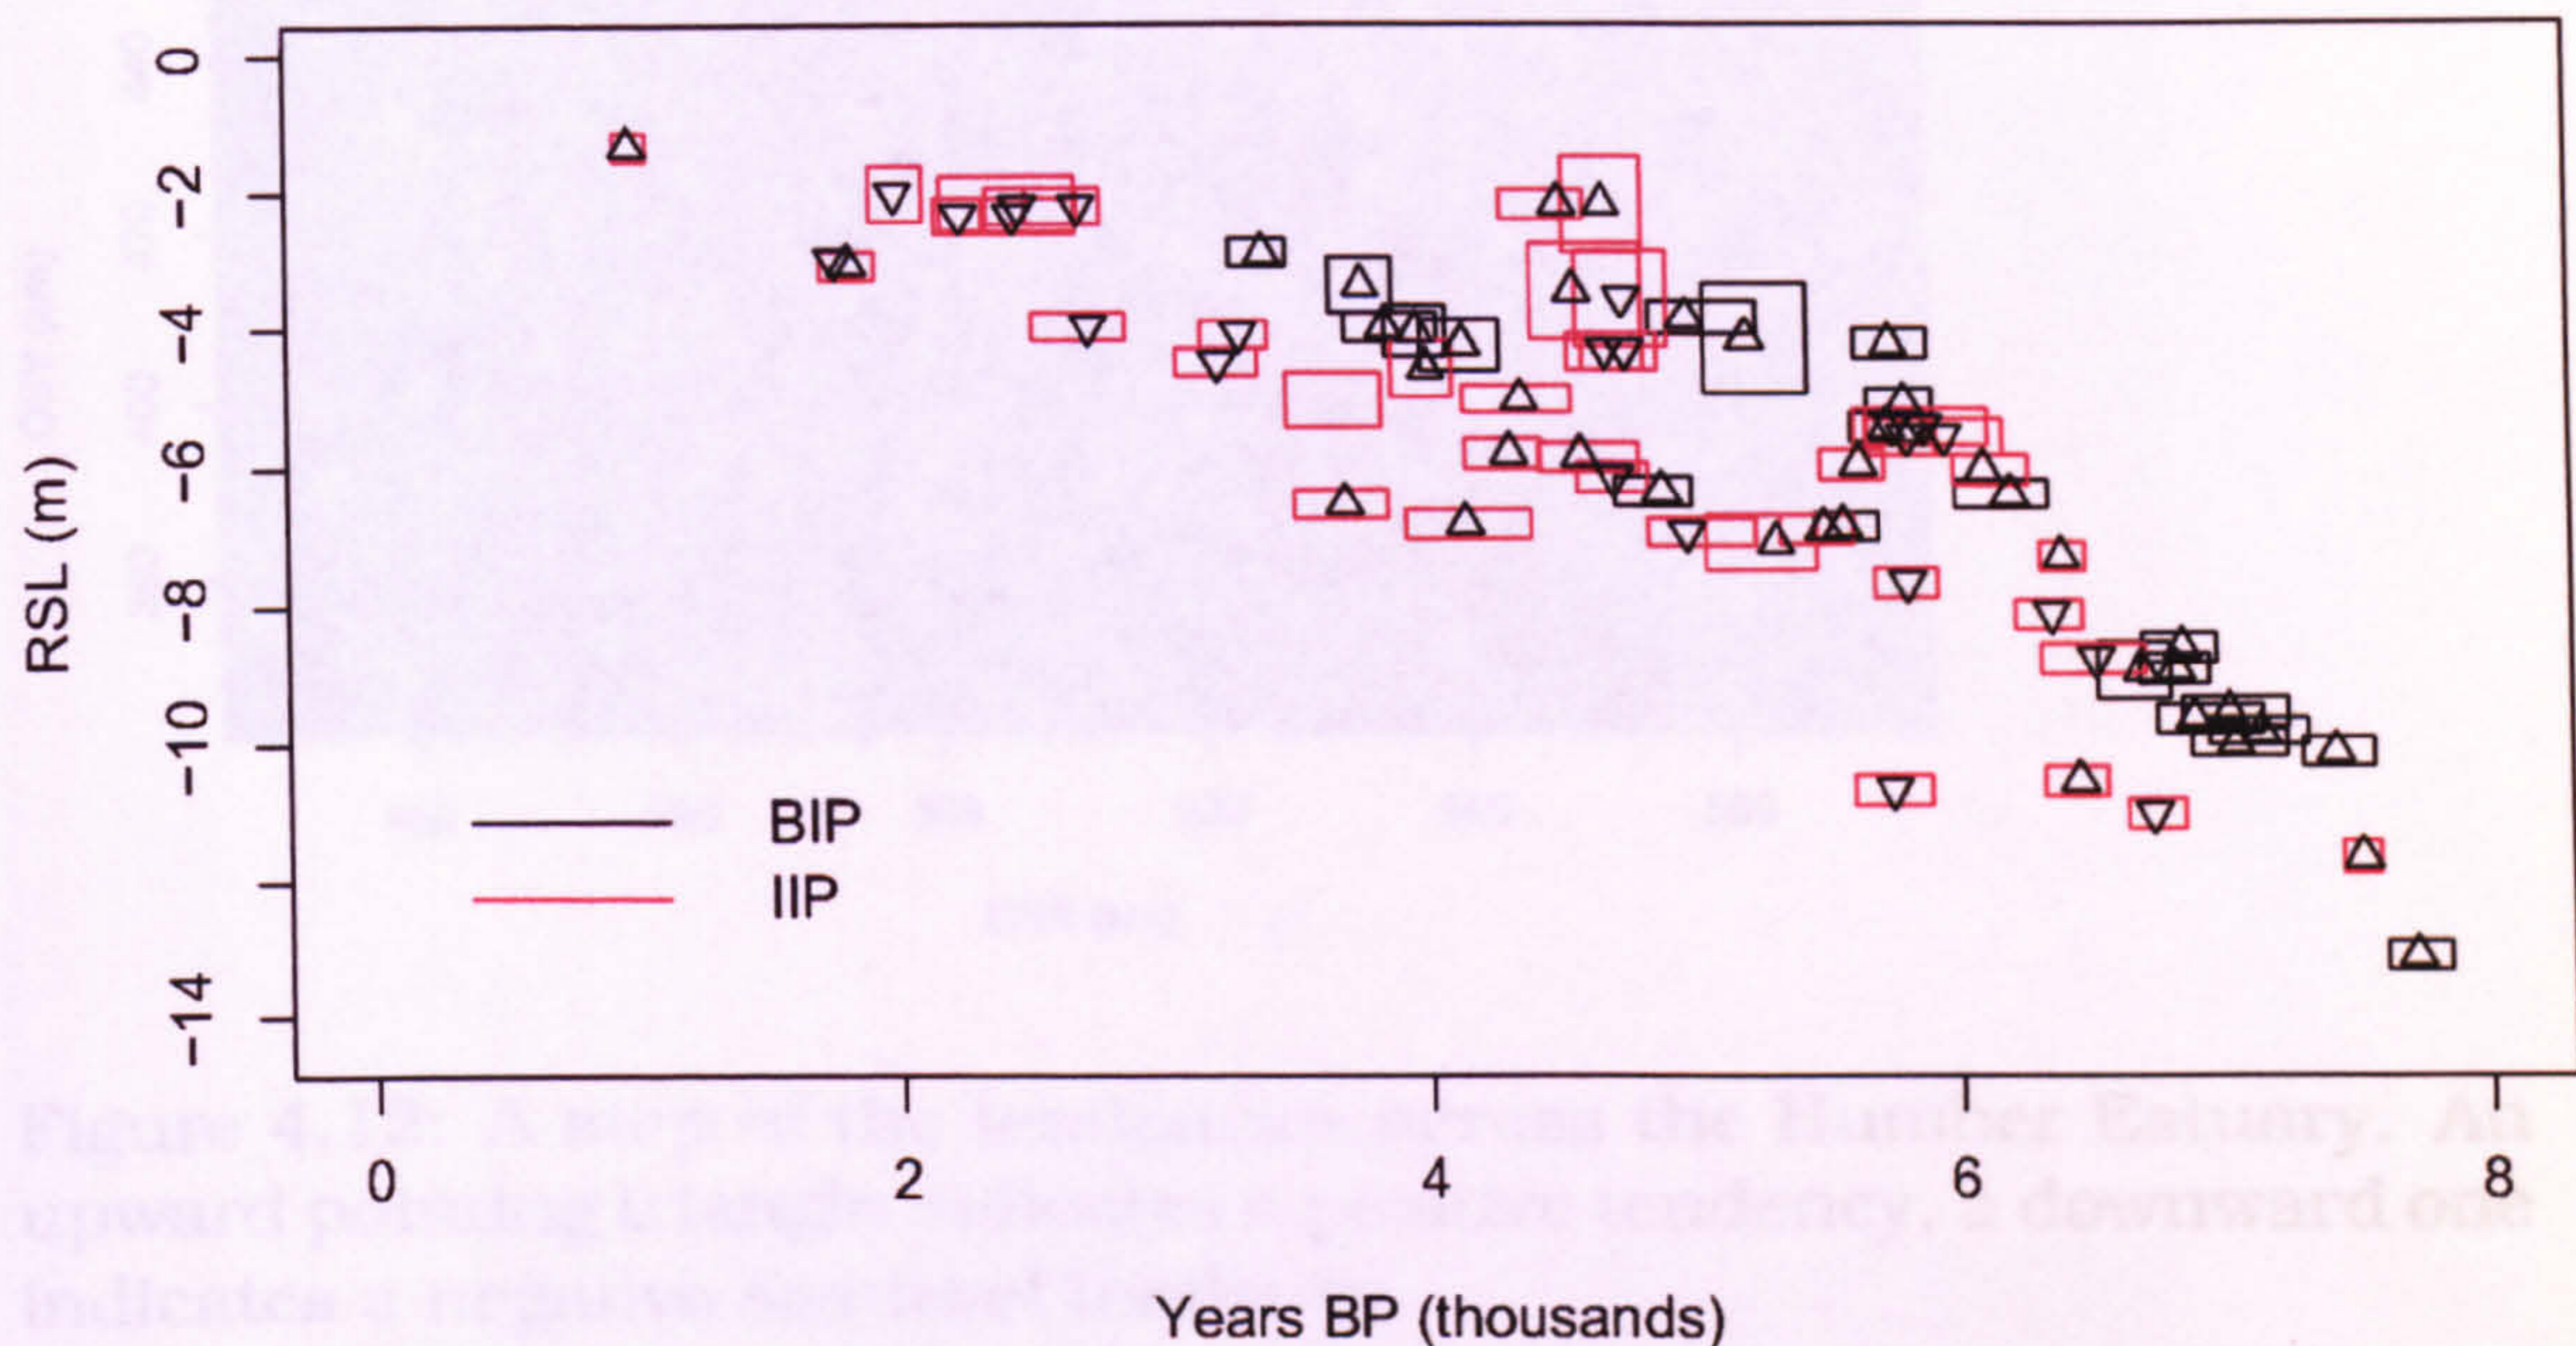


Figure 4.11: Age-altitude plot showing tendencies for the Humber. An upward pointing triangle indicates a positive tendency, a downward one indicates a negative sea-level tendency.

own is uninformative; no clear pattern emerges. There appears to be many places where there are both negative and positive tendencies. It is more informative to plot the tendencies over time. This is shown in figure 4.13. It may be possible then to discern that, for the early Outer Humber at least (8-6ka yrs BP), there is a pattern of positive sea-level tendencies.

4.8 Summary

This chapter has outlined some of the ways in which index point data can be analysed. Specifically, this analysis has involved the technique of radiocarbon calibration which demonstrates how raw radiocarbon

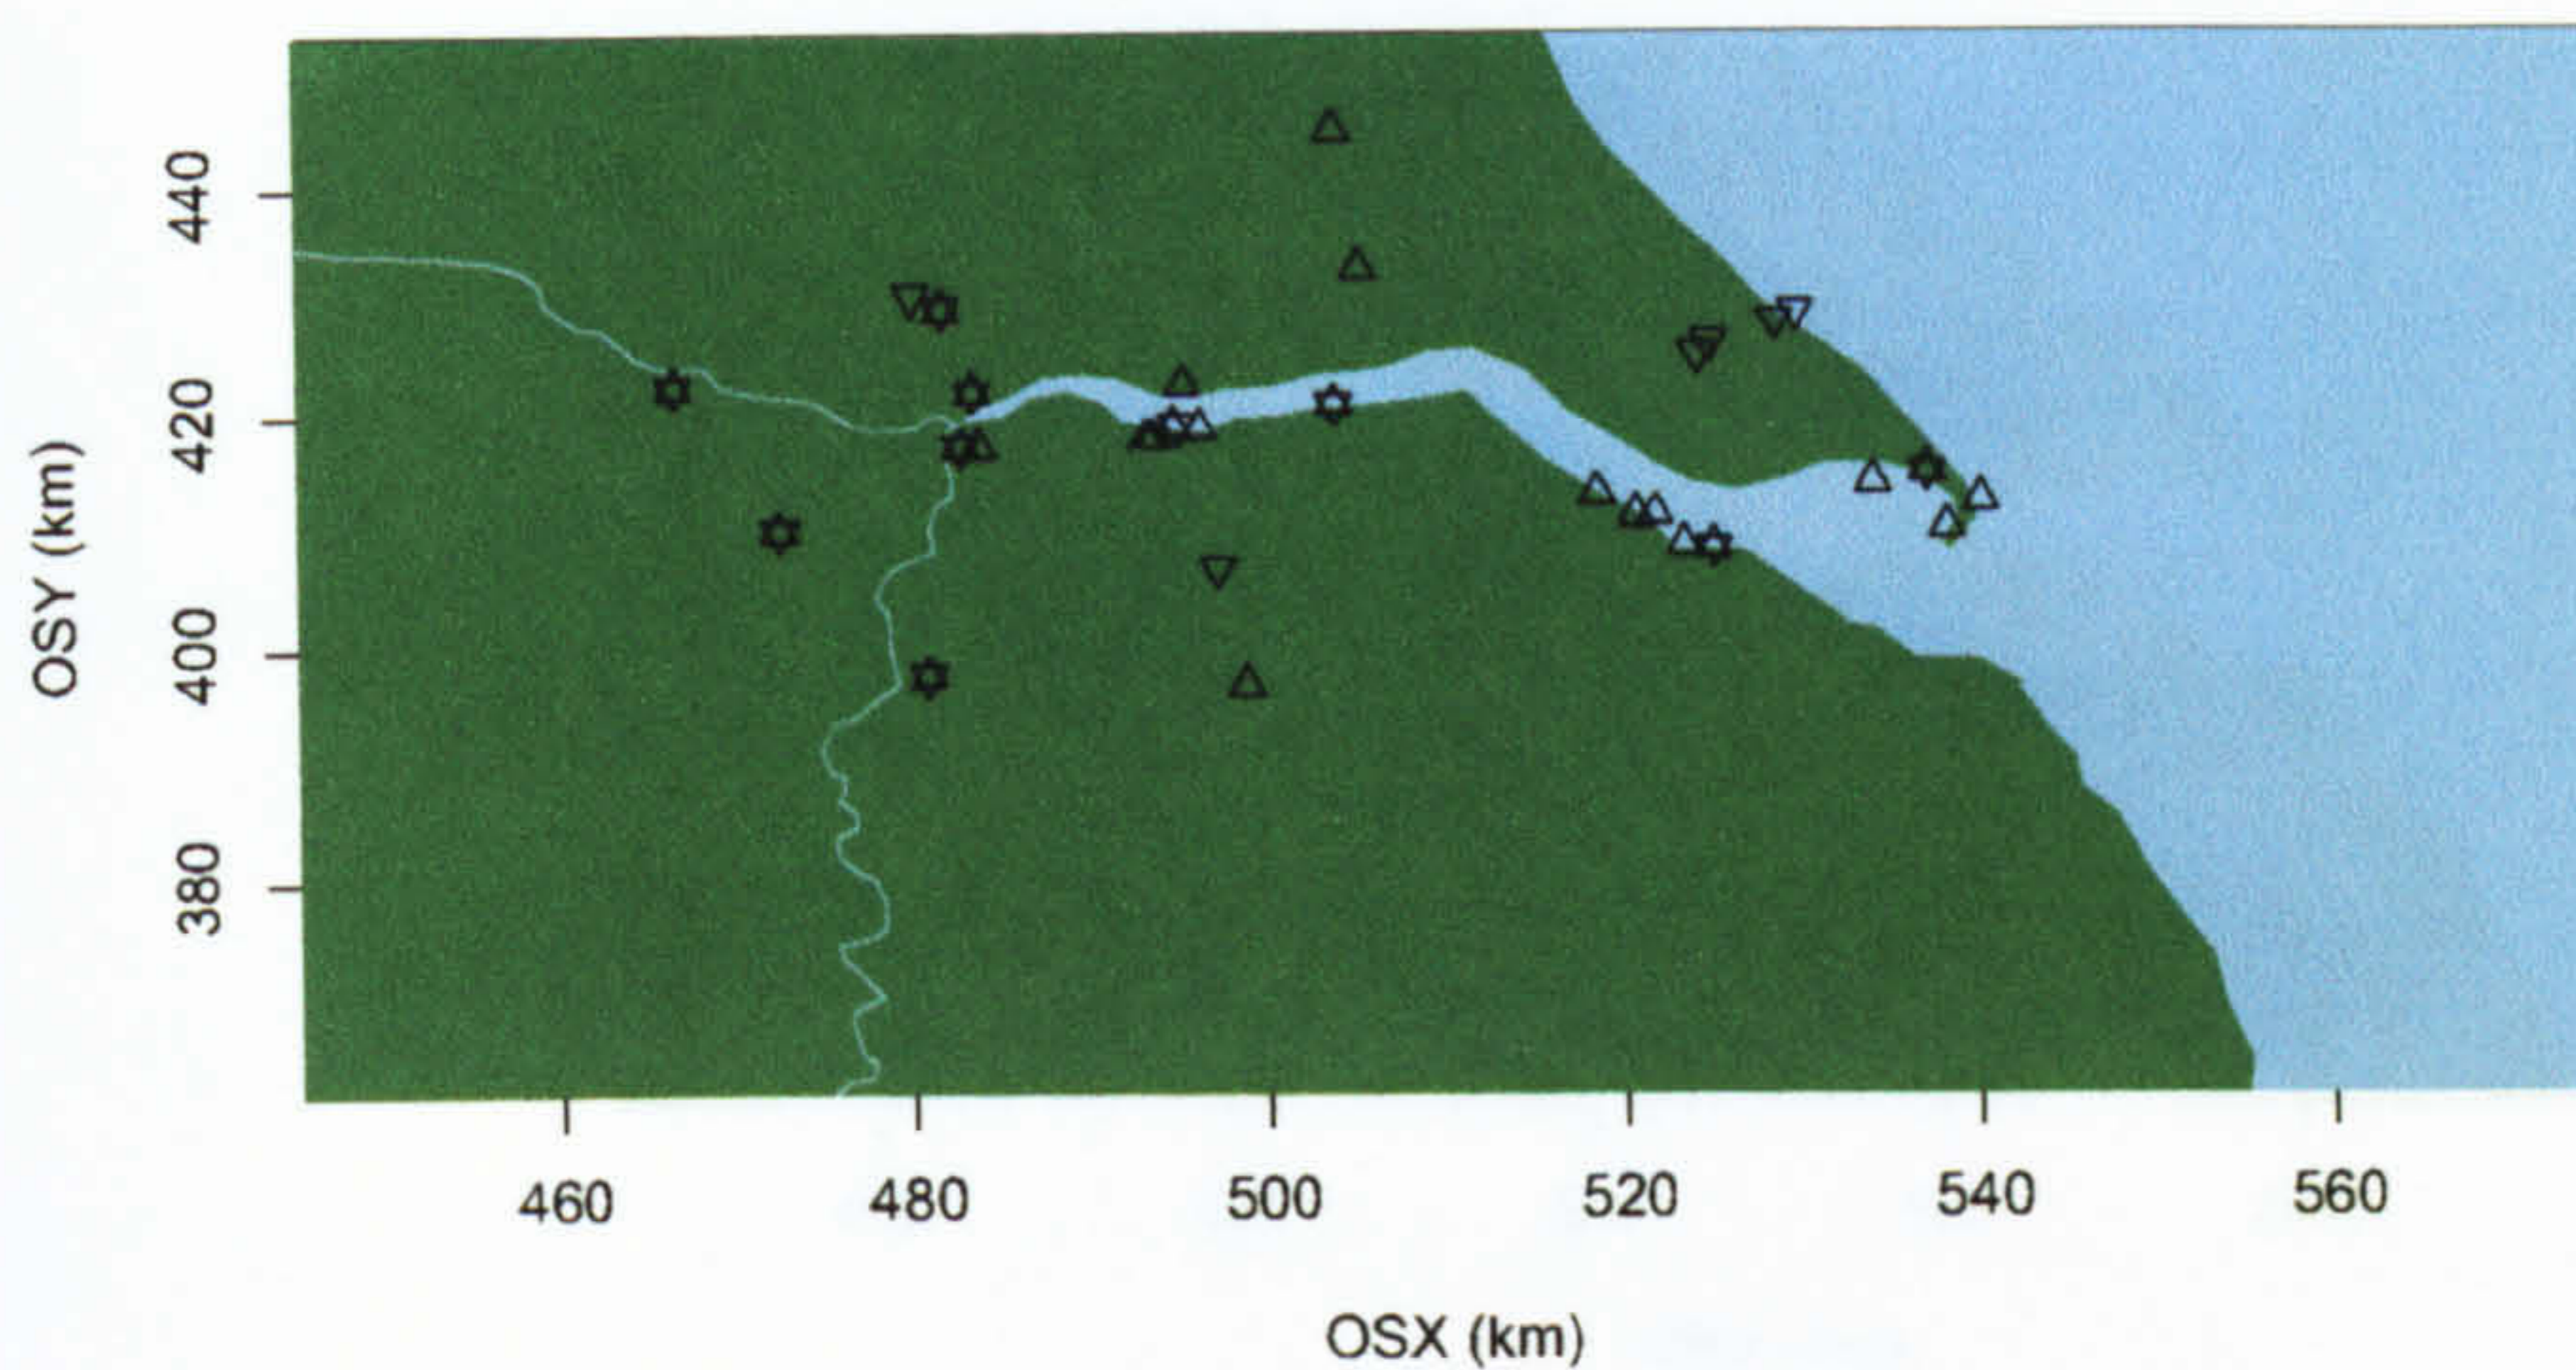


Figure 4.12: A map of the tendencies across the Humber Estuary. An upward pointing triangle indicates a positive tendency, a downward one indicates a negative sea-level tendency.

estimates can be mapped on to the calendar timescale using Bayesian techniques. This produces highly irregular posterior density estimates for the dates involved, which is why they are often ignored in many RSL studies. This chapter has also dealt with multiple radiocarbon determinations and the dating techniques preferred for archaeological dates.

To adjust limiting points and intercalated index points, initially smoothing techniques were employed alongside endpoint estimates. This idea was extended by using local quantile regression to force any RSL smoothed estimate to lie on one side of the local distribution of LPs or IIPs. The results showed a broadly rising sea-level curve with minor fluctuations.

Temporal variograms showed that correlation structure was best

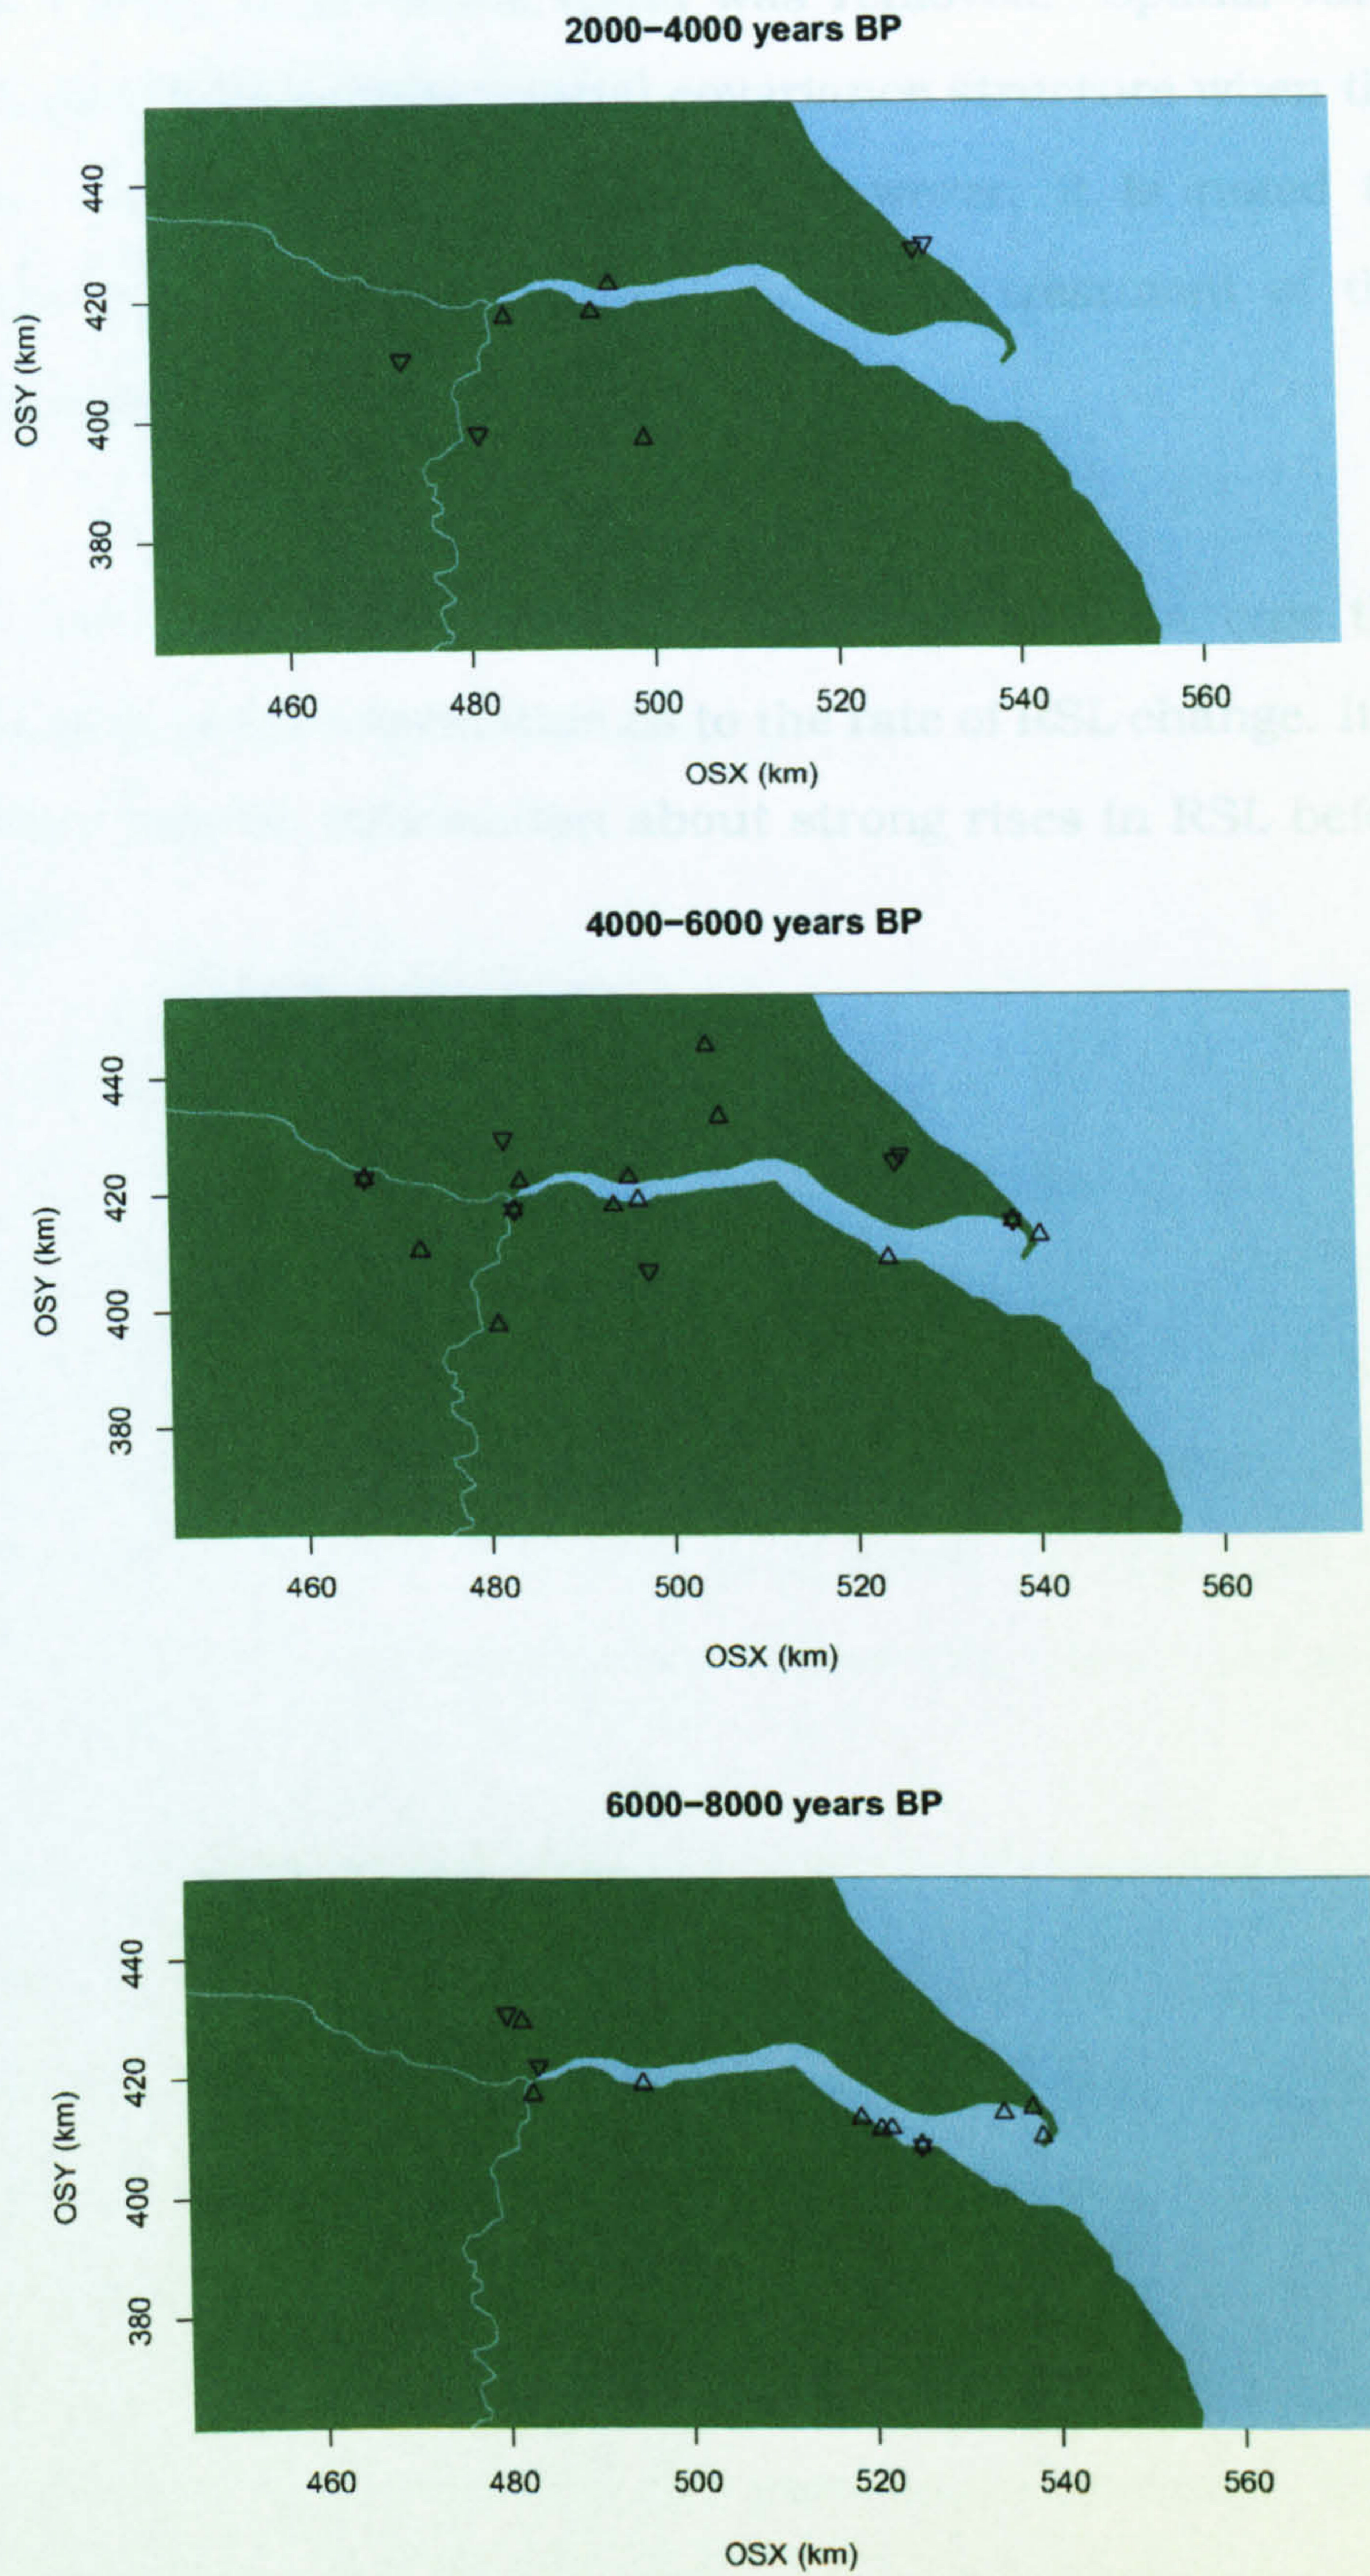


Figure 4.13: Maps of the tendency over different time periods.
121

identified when a quadratic trend was removed. Spatial variograms showed that there is little spatial covariance structure when the trend includes an East-West component. However, it is noted that the variogram model is highly unrealistic in its treatment of the data; ignoring errors in date and elevation.

Lastly, the tendency was noted as an unusual data type that may provide some extra information as to the rate of RSL change. It appears that there may be information about strong rises in RSL before 6000 years BP.

Chapter 5

Gaussian processes

5.1 Introduction

In this chapter the Gaussian process (GP) method for non-parametric regression is reviewed in the context of relative sea-level reconstruction. This theory has been dealt with in various fields, for example O'Hagan (1978); Cressie (1991). The idea is to estimate a function, $\eta(\theta)$, at a given date θ . This random function represents relative sea level (RSL).

The first section contains an introduction to the GP model. In subsequent sections, the model is extended to deal with limiting points and intercalated index points, as well as tendency information and location, and by considering more flexible covariance structures including aspects of non-stationarity. Due to the complexity of the model, a simulation approach is required to gain knowledge of the posterior distributions of the parameters involved. The chapter therefore contains a brief section on aspects of Markov chain Monte

Carlo. The chapter ends with an example showing how the technique can be used to practical effect.

5.2 Gaussian processes for smoothing

Suppose there is interest in a random function $\eta(\theta)$. A data set is observed, y , in the form:

$$y_i | \eta(\theta_i), \sigma_i \sim N(\eta(\theta_i), \sigma_i^2) \text{ for } i = 1, \dots, n. \quad (5.1)$$

The y_i s are conditionally independent with $\eta(\cdot)$ unknown. There is usually strong information about the σ_i^2 s; these may or may not be distinct. In the RSL literature, the set of y_i s are the basal index points, whilst $\eta(\theta)$ is an estimate of RSL at a date θ after taking into account the indicative meaning (see section 2.3.3). The values of σ_i^2 are assumed known as they are well-studied (eg Shennan, 1994).

A Gaussian process approach would be to assume some prior knowledge about the form of $\eta(\cdot)$. Firstly:

$$\mathbb{E}[\eta(\theta) | \beta] = \mu(\theta), \quad (5.2)$$

The random function $\eta(\cdot)$ is then given the conditional distribution:

$$\eta(\theta) | \mu(\theta), \sigma^2 \sim N(\mu(\theta), \sigma^2). \quad (5.3)$$

Furthermore, $\eta(\cdot)$ is prescribed a form whereby $\eta(\theta)$ and $\eta(\theta')$ are similar if θ and θ' are close. This is achieved by applying a covariance structure

to $\eta(\theta)$ and $\eta(\theta')$:

$$\text{Cov}(\eta(\theta), \eta(\theta') | \sigma^2, \phi) = \sigma^2 c_\phi(\theta, \theta'), \quad (5.4)$$

where $c_\phi(\theta, \theta')$ is a function in $|\theta - \theta'|$ which depends upon a set of covariance parameters ϕ . There is more detail about the nature of $c_\phi(\cdot, \cdot)$ and ϕ itself in section 5.4. Together, the mean and variance information assumed about $\eta(\theta)$ are known as the Gaussian process prior.

Overall, the distributional form for y (assuming conditional independence between errors and a Gaussian process prior for $\eta(\cdot)$) can be written as:

$$y | \mu(\theta), \Sigma, \phi, \sigma \sim N(\mu(\theta), \Sigma + \sigma^2 A), \quad (5.5)$$

where

$$\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_n^2),$$

$$A = \begin{bmatrix} 1 & c_\phi(\theta_1, \theta_2) & \dots & c_\phi(\theta_1, \theta_n) \\ c_\phi(\theta_2, \theta_1) & 1 & & \vdots \\ \vdots & & \ddots & \\ c_\phi(\theta_n, \theta_1) & \dots & & 1 \end{bmatrix}.$$

This is known as the marginal model. In the spatial statistics literature (eg Cressie, 1991) σ^2 is known as the 'sill', whilst the individual diagonal elements of Σ are known as 'nuggets'.

Alternatively, the model can be set out in a hierarchical fashion, known

as the conditional model:

$$\mathbf{y}|\eta(\boldsymbol{\theta}), \Sigma \sim N(\eta(\boldsymbol{\theta}), \Sigma) \quad (5.6)$$

$$\eta(\boldsymbol{\theta}) \sim N(\boldsymbol{\mu}(\boldsymbol{\theta}), \sigma^2 \mathbf{A}). \quad (5.7)$$

This format allows greater flexibility as the entire generalised linear modelling scheme allows transformation of the original data vector \mathbf{y} to other members of the exponential family. It does, however, have the disadvantage that each element of the unknown mean $\eta(\boldsymbol{\theta})$ requires estimation (it is now a latent Gaussian process) and so dramatically increases the number of parameters. For this reason, and because the data are treated as already Normally distributed, this style is not carried forward. More information about the marginal and conditional formats can be found in Banerjee *et al.* (2004).

Using the marginal model, the posterior estimate for the parameters can be obtained from the Bayes' equation:

$$p(\boldsymbol{\mu}, \phi, \sigma^2 | \mathbf{y}) \propto p(\boldsymbol{\mu}, \phi, \sigma^2) \times p(\mathbf{y} | \boldsymbol{\mu}, \phi, \sigma^2). \quad (5.8)$$

In the case of prediction for a value θ_0 , the task is to estimate the value of the function $\eta(\theta_0)$ given the data \mathbf{y} . This is done by initially evaluating the joint distribution of \mathbf{y} with $\eta(\theta_0)$:

$$\begin{pmatrix} \eta(\theta_0) \\ \mathbf{y} \end{pmatrix} \sim N \left(\begin{pmatrix} \boldsymbol{\mu}(\theta_0) \\ \boldsymbol{\mu}(\boldsymbol{\theta}) \end{pmatrix}, \begin{pmatrix} \sigma^2 & \mathbf{T}(\theta_0)^T \\ \mathbf{T}(\theta_0) & \sigma^2 \mathbf{A} + \boldsymbol{\Sigma} \end{pmatrix} \right), \quad (5.9)$$

where $\mathbf{T}(\boldsymbol{\theta})^T = \sigma^2(c_{\phi}(\theta_0, \theta_1), \dots, c_{\phi}(\theta_0, \theta_n))$.

To obtain the distribution of $\eta(\theta_0)|y$, a standard result is used. Suppose x , a p -dimensional random variable, has a multivariate normal distribution and can be partitioned into two parts so that $x = (x_1 \ x_2)^T$. Suppose further that the mean and variance matrix of x can also be partitioned and written as $\mu = (\mu_1 \ \mu_2)^T$ and $\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$. The conditional distribution of x_1 given that $x_2 = a$ now has a multivariate Normal distribution with mean:

$$\mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(a - \mu_2) \quad (5.10)$$

and variance:

$$\Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}. \quad (5.11)$$

A proof of this result is given in Krzanowski (1988).

Applying this result to the joint distribution in 5.9:

$$\eta(\theta_0)|y, \mu, \Sigma, \phi, \sigma^2 \sim N(m^*(\theta_0), c^*(\theta_0, \theta_0)), \quad (5.12)$$

where

$$m^*(\theta_0) = \mu(\theta_0) + T(\theta_0)^T C^{-1}(y - \mu(\theta)), \quad (5.13)$$

$$c^*(\theta_0, \theta_0) = \sigma^2 - T(\theta_0)^T C^{-1}T(\theta_0), \quad (5.14)$$

with

$$C = \Sigma + \sigma^2 A. \quad (5.15)$$

5.3 Choice of mean

In this thesis, the prior mean of the Gaussian process (following Oakley, 2002) is given the form :

$$\mathbb{E}[\eta(\theta)|\beta] = f(\theta)^T \beta, \quad (5.16)$$

where $f(\theta)$ is a vector of length q of known functions of an independent variable, θ , and β is a vector of the same length of unknown parameters. The random function $\eta(\cdot)$ now has the conditional distribution:

$$\eta(\theta)|\beta, \sigma^2 \sim N(f(\theta)^T \beta, \sigma^2). \quad (5.17)$$

The choice of mean indicates prior structure about $\eta(\cdot)$. The function f may be constant, linear, quadratic, or more complicated if necessary. Clearly, increasing q (and therefore the dimension of f) introduces extra parameters in the vector β . Other approaches, such as that by Host (1999), consider more complicated mean structures.

It has been shown by Oakley (2002) that using the version of the prior mean in equation 5.16 and a simple form for the covariance leads to explicit forms for the posterior distributions of the parameters β and σ^2 under conjugacy. Indeed, some of these methods are

used where appropriate (chapter 6). For many of the parameters, however, only non-conjugate priors are appropriate. In these cases simulation techniques such as MCMC are required to evaluate posterior distributions.

5.4 Choice of autocorrelation function

There is a wide range of choices for the autocorrelation function $c_\phi(t, t')$ and its parameters as described in equation 5.4. Some of the choices are discussed here. The parameters ϕ , which govern autocorrelation, can be categorised into three types:

- Scale parameters which attempt to standardise the distance between t and t' .
- Range parameters which set the distance at which correlation decreases to zero. For correlation functions with infinite range, an effective range is given for which the correlation decreases beyond an arbitrary value (in this section given as c).
- Smoothness parameters which determine the speed of correlation decay.

Also of concern is the degree to which the correlation function can be differentiated. As will be shown in section 5.6, the Gaussian process estimate $\eta(\cdot)$ can only be differentiated to the lowest order of which its mean and covariance can be differentiated. This section assumes an

infinitely differentiable mean, so that only the covariance determines the differentiability of $\eta(\cdot)$.

5.4.1 Gaussian

The Gaussian correlation model is so named because of its form:

$$c_\phi(t, t') = \exp(-\phi(t - t')^2), \quad \phi > 0. \quad (5.18)$$

Here, ϕ is a scale parameter; the correlation has effective range $\sqrt{-\phi \log c}$. Raising the distance measure $|t - t'|$ to the power of 2 enables the function to be infinitely differentiable which can be useful (see section 5.6). However, it also forces resulting estimates to be very smooth which is often unrealistic. This criticism and others are outlined in Stein (1999).

5.4.2 Powered exponential

This is a generalisation of the Gaussian model where the distance measure is raised to the power ϕ_2 rather than 2:

$$c_\phi(t, t') = \exp(-\phi_1 |t - t'|^{\phi_2}), \quad \phi_1 > 0, \quad \phi_2 \in (0, 2). \quad (5.19)$$

Here, ϕ_1 is a scale parameter and ϕ_2 a smoothness parameter; the effective range is $(-\phi \log c)^{\frac{1}{\phi_2}}$. Such models are used by the Bayesian inference package WinBUGS and briefly by Sacks *et al.* (1989). Clearly, when ϕ_2 is less than 2 this is non-differentiable.

5.4.3 Matérn

This is a class of autocorrelation function more flexible than either the powered exponential or the Gaussian. It was originally suggested by Matérn (1960). The correlation function is given by:

$$c_{\phi}(t, t') = \frac{1}{2^{\phi_2-1}\Gamma(\phi_2)} \left(\frac{2\phi_2^{1/2}|t-t'|}{\phi_1} \right)^{\phi_2} \mathcal{K}_{\phi_2} \left(\frac{2\phi_2^{1/2}|t-t'|}{\phi_1} \right), \quad (5.20)$$

where \mathcal{K} is a modified Bessel function of the third type. The Matérn covariance function again has scale parameter ϕ_1 and smoothness parameter ϕ_2 . The range cannot be computed in closed form but, for $\phi_2 = 1.5$, it can be approximated as $4.744/\phi_1$. The advantage of the Matérn form is flexibility; the Gaussian covariance is recovered if $\phi_2 \rightarrow \infty$ and the function is m times differentiable if $\phi_2 > m$. More information can be found in Stein (1999). A plot of the Matérn correlation function with different values of ϕ_1 and ϕ_2 is given in figure 5.1.

5.4.4 Non-stationary correlation

By forcing a global set of covariance parameters on a model, the analyst is saying that the correlation between $\eta(\theta_i)$ and $\eta(\theta_j)$ is independent of θ . If this is not the case, or if it is not known, it is wise to allow any covariance parameters to vary as they move around θ -space.

The non-stationary covariance function proposed by Paciorek and

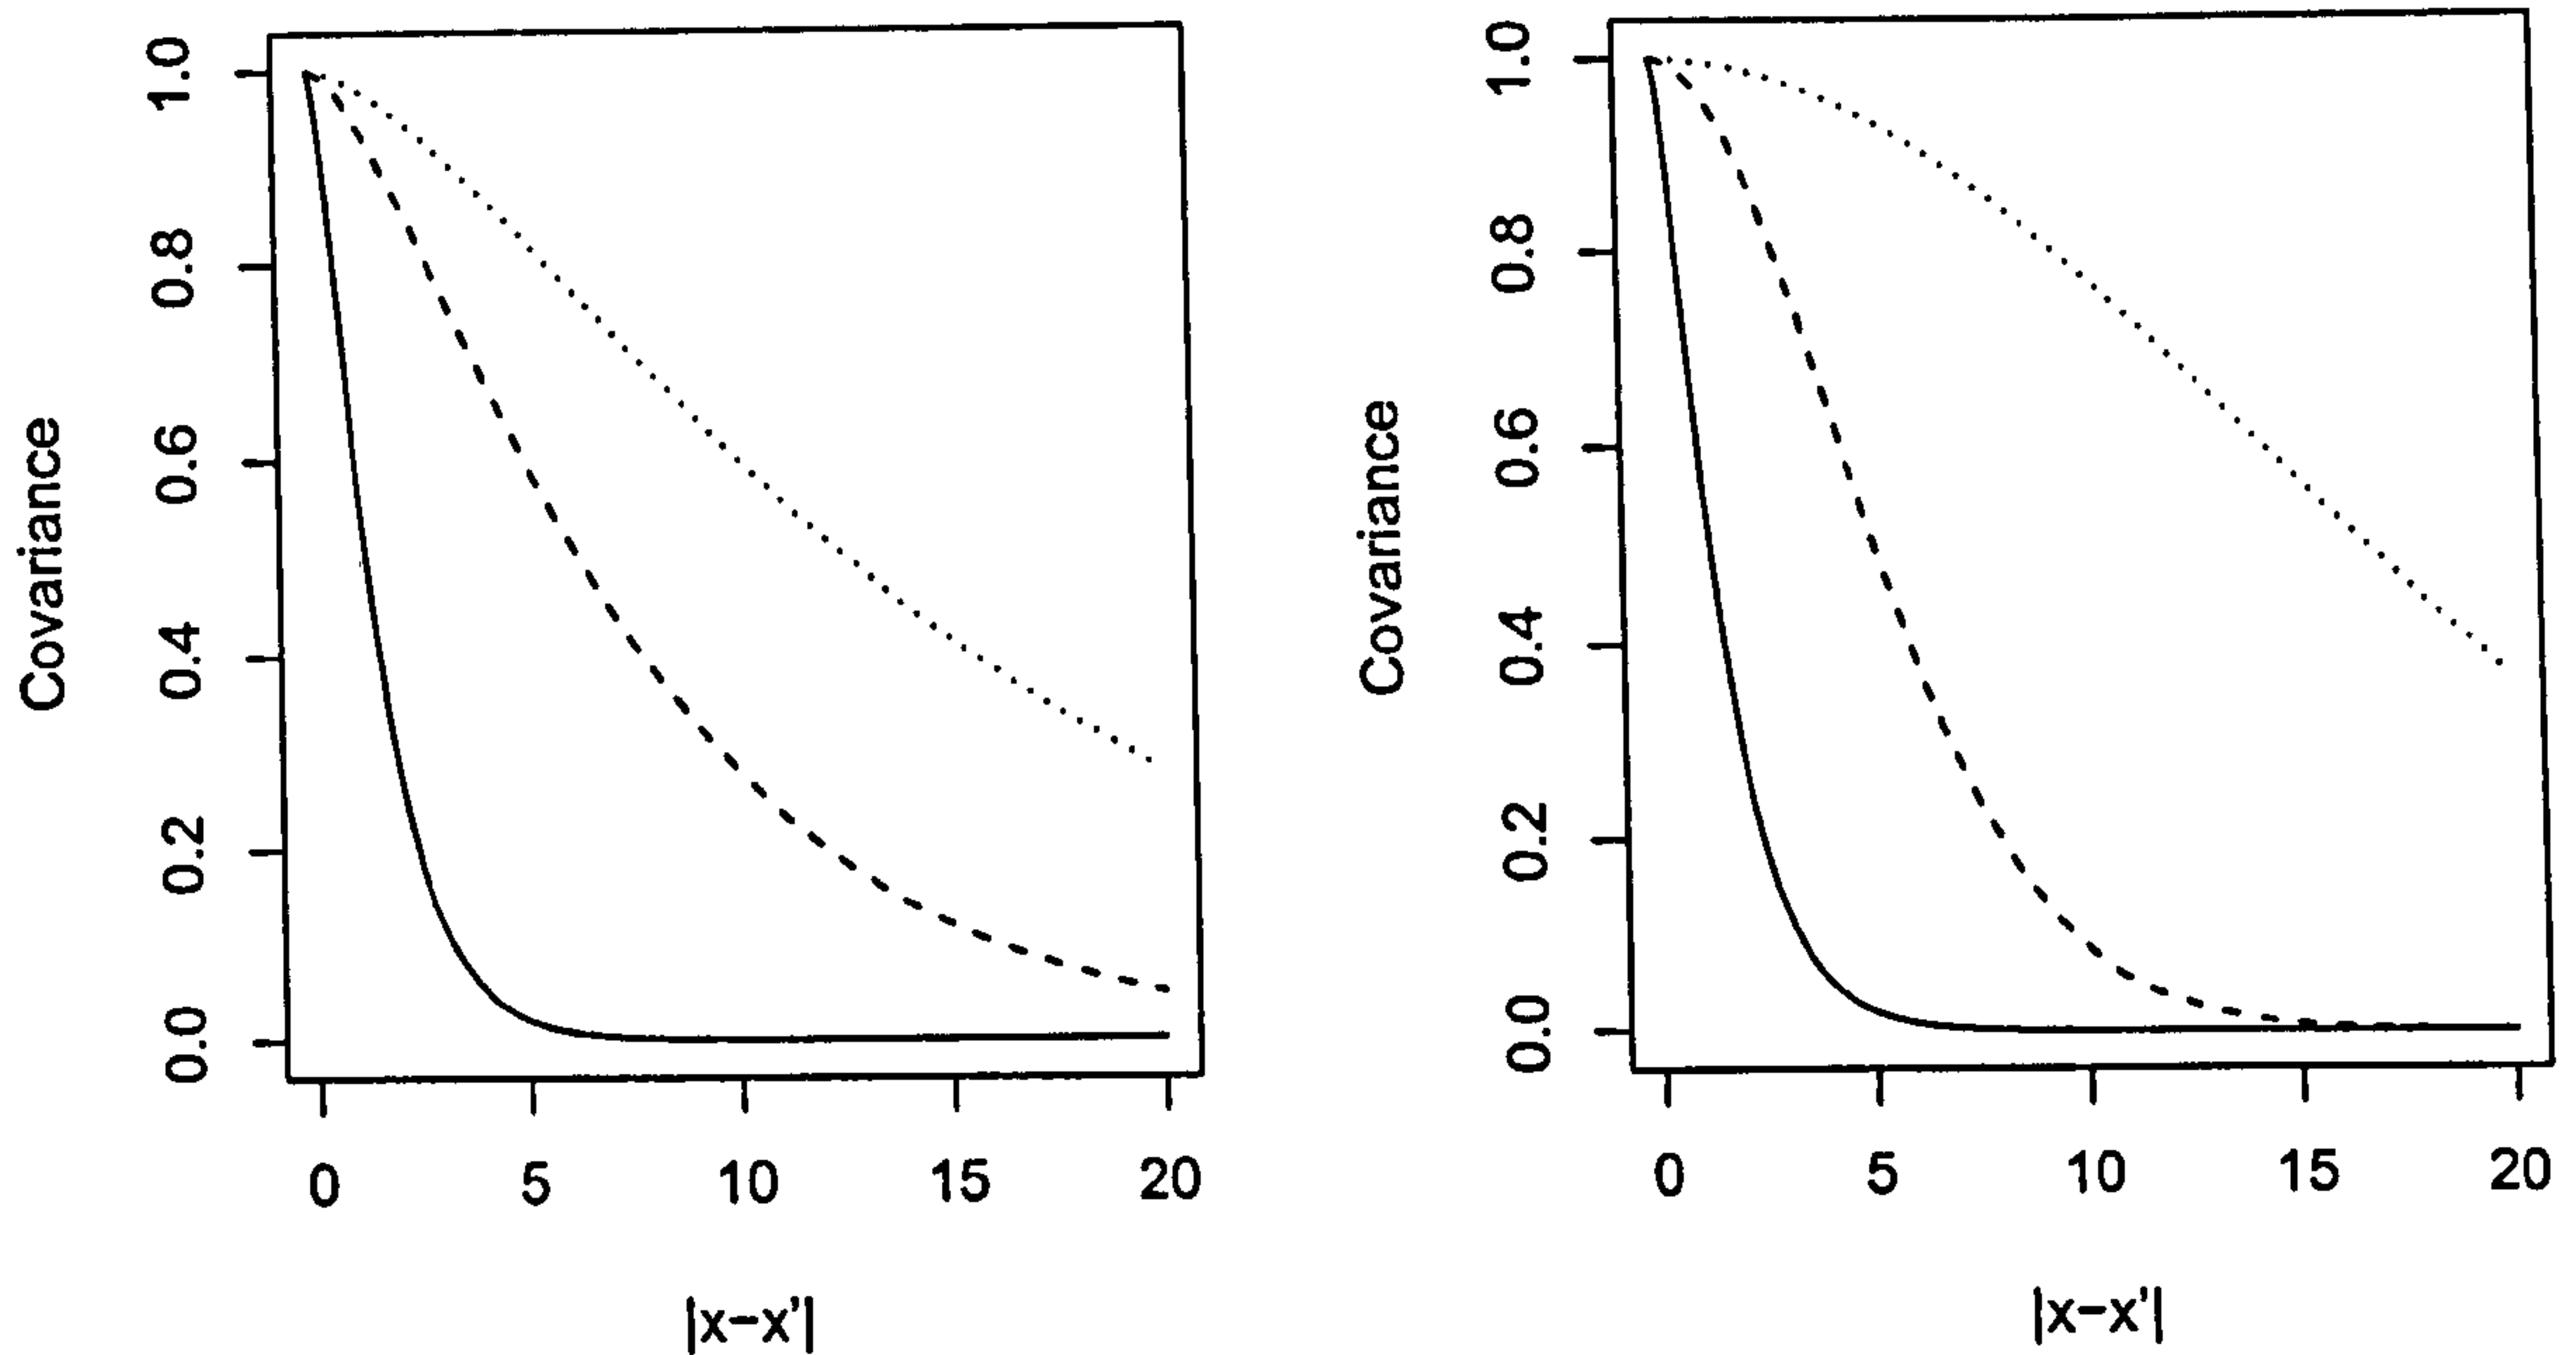


Figure 5.1: A plot of the Matérn correlation function. In the first panel ϕ_1 is set to the values 1 (solid line), 5 (dashed) and 10 (dotted) with ϕ_2 fixed at 1. In the second panel ϕ_2 is given the values 1 (solid), 10 (dashed) and 100 (dotted) with ϕ_1 fixed at 1.

Schervish (2004) has the following form:

$$C^{NS}(t_i, t_j) = \sigma^2 |\Omega_i|^{\frac{1}{4}} |\Omega_j|^{\frac{1}{4}} |\Omega_i + \Omega_j|^{-\frac{1}{2}} \rho_\phi(t_i - t_j), \quad (5.21)$$

where $\Omega_i = \Omega(t_i)$ is a kernel function centered at t_i . In the one dimensional case, each Ω_i is just a scalar and is allowed to vary by giving $\log(\Omega)$ a Gaussian process prior with a stationary covariance function. For the Matérn covariance function, Paciorek and Schervish (2004) propose:

$$C^{NS}(t_i, t_j) = |\Omega_i|^{\frac{1}{4}} |\Omega_j|^{\frac{1}{4}} |\Omega_i + \Omega_j|^{-\frac{1}{2}} \frac{1}{\Gamma(\phi) 2^{\phi-1}} \left(\sqrt{2\phi Q_{ij}} \right)^\phi \mathcal{K}_\phi \left(\sqrt{2\phi Q_{ij}} \right), \quad (5.22)$$

where

$$Q_{ij} = (t_i - t_j)^T ((\Omega_i + \Omega_j)/2)^{-1} (t_i - t_j). \quad (5.23)$$

This approach allows for more flexibility in the smoothness and variability of $\eta(x)$. The resulting Bayes' equations are now adjusted to take account of the extra parameters. If the sets of parameters are defined with $\tau^T = (\beta, \sigma^2)$ and $\tau_\Omega^T = (\beta_\Omega, \sigma_\Omega^2, \phi_\Omega)$ then the posterior distribution of the parameters can be written as:

$$p(\tau, \tau_\Omega, \Omega | y) \propto p(\tau)p(\Omega | \tau_\Omega)p(\tau_\Omega)p(y | \Omega, \tau_y). \quad (5.24)$$

The first three terms are the prior distribution; the last term is the likelihood.

5.5 Limiting points and intercalated index points

For points that lie either above or below the estimate of RSL, $\eta(\theta)$, an extra term is introduced. For more information on limiting points and intercalated index points, see section 2.3.6. To extend the model, these data points are now given the following form:

$$y_i = \eta(\theta_i) + \gamma_i + e_i \text{ for } i = 1, \dots, m$$

where each of these terms are random variables with distributions:

$$\eta(\theta_i) \sim N(\mathbf{f}(\theta_i)^T \boldsymbol{\beta}, \sigma^2) \quad (5.25)$$

$$\text{Cov}(\eta(\theta_i), \eta(\theta_j)) = \sigma^2 \rho(|\theta_i - \theta_j|) \quad (5.26)$$

$$e_i \sim N(0, \sigma_i^2) \text{ i.i.d.} \quad (5.27)$$

$$\gamma_i \sim N(h, \kappa^2) \text{ i.i.d.} \quad (5.28)$$

The parameters γ_i act as 'shifts' for each y_i to take account of the fact that limiting points, for example, may lie away from the sea-level curve.

Ignoring intercalated index points for the moment, let the combined data points \mathbf{y} be partitioned into two parts $\begin{pmatrix} \mathbf{y}_{BIP} \\ \mathbf{y}_{LP} \end{pmatrix}$ where \mathbf{y}_{BIP} are the set of basal index points and \mathbf{y}_{LP} are the set of limiting points. The joint distribution is:

$$\begin{pmatrix} \mathbf{y}_{BIP} \\ \mathbf{y}_{LP} \end{pmatrix} \sim N \left(\begin{pmatrix} \mathbf{F}^T \boldsymbol{\beta} \\ \mathbf{F}^T \boldsymbol{\beta} + \mathbf{h} \end{pmatrix}, \begin{pmatrix} \sigma^2 \mathbf{A}_{BIP} + \boldsymbol{\Sigma}_{BIP} & \sigma^2 \mathbf{A}_{BIP,LP} \\ \sigma^2 \mathbf{A}_{BIP,LP} & \sigma^2 \mathbf{A}_{LP} + \boldsymbol{\Sigma}_{LP} + \boldsymbol{\Psi} \end{pmatrix} \right) \quad (5.29)$$

where all parameters are defined as before and $\boldsymbol{\Psi} = \kappa^2 \mathbf{I}_m$ where \mathbf{I}_m is an $m \times m$ identity matrix. These have a joint normal distribution with $\eta(\theta)$:

$$\begin{pmatrix} \eta(\theta) \\ \mathbf{y}_{BIP} \\ \mathbf{y}_{LP} \end{pmatrix} \sim N \left(\begin{pmatrix} \mathbf{f}(\theta)^T \boldsymbol{\beta} \\ \mathbf{F}^T \boldsymbol{\beta} \\ \mathbf{F}^T \boldsymbol{\beta} + \mathbf{h} \end{pmatrix}, \begin{pmatrix} \sigma^2 & \mathbf{T}_{BIP}(\theta)^T & \mathbf{T}_{LP}(\theta)^T \\ \mathbf{T}_{BIP}(\theta) & \sigma^2 \mathbf{A}_{BIP} + \boldsymbol{\Sigma}_{BIP} & \sigma^2 \mathbf{A}_{BIP,LP} \\ \mathbf{T}_{LP}(\theta) & \sigma^2 \mathbf{A}_{BIP,LP} & \sigma^2 \mathbf{A}_{LP} + \boldsymbol{\Sigma}_{LP} + \boldsymbol{\Psi} \end{pmatrix} \right), \quad (5.30)$$

where the subscript on the matrices \mathbf{T} , \mathbf{A} and $\boldsymbol{\Sigma}$ determine whether the set of index or limiting points are used.

The posterior for $\eta(\theta)$ can be obtained as before:

$$\eta(\theta)|\mathbf{y}, \beta, \Sigma, \phi, \Psi, \sigma^2, \kappa^2, h \sim N(m^{**}(\theta), c^{**}(\theta, \theta)), \quad (5.31)$$

with

$$m^{**}(\theta) = f(\theta)^T \beta + \begin{pmatrix} T_{BIP}(\theta)^T & T_{LP}(\theta)^T \end{pmatrix} \begin{pmatrix} \sigma^2 A_{BIP} + \Sigma_{BIP} & \sigma^2 A_{BIP,LP} \\ \sigma^2 A_{BIP,LP} & \sigma^2 A_{LP} + \Sigma_{LP} + \Psi \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{y}_{BIP} - F^T \beta \\ \mathbf{y}_{LP} - F^T \beta - h \end{pmatrix}, \quad (5.32)$$

$$c^{**}(\theta, \theta) = \sigma^2 - \begin{pmatrix} T_{BIP}(\theta)^T & T_{LP}(\theta)^T \end{pmatrix} \begin{pmatrix} \sigma^2 A_{BIP} + \Sigma_{IP} & \sigma^2 A_{BIP,LP} \\ \sigma^2 A_{BIP,LP} & \sigma^2 A_{LP} + \Sigma_{LP} + \Psi \end{pmatrix}^{-1} \begin{pmatrix} T_{BIP}(\theta) \\ T_{LP}(\theta) \end{pmatrix}. \quad (5.33)$$

For intercalated index points, an extra term λ_i is introduced which is also Gaussian, with parameters h_2 and κ_2^2 . The difference between the λ and γ terms is that the posterior expectation of h_1 (formerly h) in the limiting points would be expected to be negative, whilst for the intercalated index points the posterior expectation of h_2 should be positive. Obviously, prior information could force this to be the case. In subsequent sections, the subscripts of 1 and 2 will be dropped. Now, \mathbf{y} contains the set of basal, limiting and intercalated index points. The vector \mathbf{x} indicates the dates of the basal, intercalated and limiting points. The mean of \mathbf{y} will be referred to as $F^T \beta + I(\mathbf{x})h$ where I is an indicator function that is 0 for basal index points and 1 for intercalated

and limiting points. h is the vector:

$$h = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ h_1 \\ \vdots \\ h_1 \\ h_2 \\ \vdots \\ h_2 \end{pmatrix}. \quad (5.34)$$

The variance of y will be referred to as $\sigma^2 A + \Sigma + \Psi$ where:

$$\Psi = \begin{pmatrix} 0 & 0 \\ 0^T & \kappa^2 I_m \end{pmatrix}. \quad (5.35)$$

I_m is defined as before.

5.6 Gaussian process derivatives

One advantage of Gaussian processes is the behaviour of derivatives with respect to the input vector θ . This has been discussed, for example, in Cramér and Leadbetter (1967), and was introduced in the context of function estimation by O'Hagan (1992). The result arises

from first principles:

$$\frac{d}{dx}\eta(\theta) = \lim_{h \rightarrow 0} \left(\frac{\eta(\theta + h) - \eta(\theta)}{h} \right). \quad (5.36)$$

Since $\eta(\theta)$ is Normally distributed and the derivative is simply a linear function of $\eta(\theta)$, the derivative is also Normal with mean and variance given by:

$$\text{Cov} \left(\eta(\theta_i), \frac{\partial}{\partial \theta_j} \eta(\theta_j) | \sigma^2, \phi \right) = \frac{\partial}{\partial \theta_j} \sigma^2 c(\theta_i, \theta_j) = \sigma^2 c_2(\theta_i, \theta_j), \quad (5.37)$$

$$\text{Cov} \left(\frac{\partial}{\partial \theta_i} \eta(\theta_i), \frac{\partial}{\partial \theta_j} \eta(\theta_j) | \sigma^2, \phi \right) = \frac{\partial}{\partial \theta_i} \frac{\partial}{\partial \theta_j} \sigma^2 c(\theta_i, \theta_j) = \sigma^2 c_{12}(\theta_i, \theta_j) \quad (5.38)$$

where l in $c_l(\theta_i, \theta_j)$ denotes which element(s) of the covariance function to differentiate with respect to. These derivatives can be calculated explicitly for the Gaussian covariance function and those similar to it. However, for the Matérn covariance function it is simpler when they are numerically evaluated. Nonetheless, having observed $\frac{\partial}{\partial \theta_i} \eta(\theta_i) = \eta'(\theta_i)$ the posterior information about $\eta(\theta)$ can be formed by considering the joint distribution:

$$\begin{pmatrix} \eta(\theta) \\ \mathbf{y} \\ \eta'(\boldsymbol{\theta}_T) \end{pmatrix} \sim N \left(\begin{pmatrix} \mathbf{f}(\theta)^T \boldsymbol{\beta} \\ \mathbf{F}^T \boldsymbol{\beta} + \mathbf{I}(\theta) \mathbf{h} \\ \frac{\partial}{\partial \boldsymbol{\theta}_T} \mathbf{F}^T \boldsymbol{\beta} \end{pmatrix}, \begin{pmatrix} \sigma^2 & \mathbf{T}(\theta)^T & \mathbf{T}_2(\theta)^T \\ \mathbf{T}(\theta) & \sigma^2 \mathbf{A} + \boldsymbol{\Sigma} + \boldsymbol{\Psi} & \mathbf{C}_2(\boldsymbol{\theta}_T, \theta)^T \\ \mathbf{T}_2(\theta) & \mathbf{C}_2(\boldsymbol{\theta}_T, \theta) & \mathbf{C}_{12}(\boldsymbol{\theta}_T) \end{pmatrix} \right) \quad (5.39)$$

where $\mathbf{T}_2(\mathbf{x})^T = \sigma^2 (\partial c(\theta, \theta_1) / \partial \theta_1, \dots, \partial c(\theta, \theta_n) / \partial \theta_n)$, $\boldsymbol{\theta}_T$ is the row subset of

θ which have derivatives, and

$$C_l(\theta_T, \theta) = \sigma^2 \begin{bmatrix} c_l(\theta_1, \theta_1) & c_l(\theta_1, \theta_2) & \dots & c_l(\theta_1, \theta_m) \\ c_l(\theta_2, \theta_1) & c_l(\theta_2, \theta_2) & & \vdots \\ \vdots & & \ddots & \\ c_l(\theta_n, \theta_1) & \dots & & c_l(\theta_n, \theta_m) \end{bmatrix}. \quad (5.40)$$

Again, the posterior for $\eta(\theta)$ can be obtained as:

$$\eta(\theta) | y, \eta'(\theta), \Sigma, \phi, \beta, \Psi, \sigma^2, \kappa^2, h \sim N(m^{**}(\theta), c^{**}(\theta, \theta)), \quad (5.41)$$

with

$$m^{**}(\theta) = f(\theta)^T \beta + \begin{pmatrix} T(\theta)^T & T_2(\theta)^T \end{pmatrix} \begin{pmatrix} \sigma^2 A + \Sigma + \Psi & C_2(\theta)^T \\ C_2(\theta) & C_{12}(\theta) \end{pmatrix}^{-1} \\ \times \begin{pmatrix} y - F^T \beta - I(\theta_{LP})h \\ \eta'(\theta) - \frac{\partial}{\partial \theta} F^T \beta \end{pmatrix}, \quad (5.42)$$

$$c^{**}(\theta, \theta) = \sigma^2 - \begin{pmatrix} T(\theta)^T & T_2(\theta)^T \end{pmatrix} \begin{pmatrix} \sigma^2 A + \Sigma + \Psi & C_2(\theta)^T \\ C_2(\theta) & C_{12}(\theta) \end{pmatrix}^{-1} \\ \times \begin{pmatrix} T(\theta) \\ T_2(\theta) \end{pmatrix}. \quad (5.43)$$

The addition of derivative information changes the inference about the set of parameters, τ , because now the values of η' are jointly distributed with the data y conditional on the same parameters. The form is identical to that of equation 5.8 but with y replaced by the set y, η' . However, this new setup assumes that the values of η' are observed directly, which is not the case in the sea-level setting.

5.6.1 Utilising tendency information

In practice, the set of derivatives η' are observed with error, with only two possible responses; positive and negative. This data type is known as the tendency. The challenge is then to form a posterior distribution for the set of parameters given both the original data and the tendency. It should be noted that a positive tendency indicates that the gradient, $\eta'(\theta)$, should be negative as time is always plotted in reverse order. The data are observed as:

$$z_i = I(\eta'(\theta_i) + \zeta_i) = \begin{cases} 1 & \text{if } \eta'(\theta_i) + \zeta_i < 0 \\ -1 & \text{if } \eta'(\theta_i) + \zeta_i > 0 \end{cases} \quad (5.44)$$

$$\zeta \sim N(0, \delta^2). \quad (5.45)$$

Here, ζ and δ control the amount of uncertainty in the tendency information. A large value of δ would lead to a higher probability of z_i being positive when the true gradient was positive.

A form for the distribution $\tau, \eta'(\theta_T), \delta | y, z$, the posterior estimates of the parameters given the data, is required. It can be re-written in Bayes' form as:

$$p(\tau, \eta'(\theta_T), \delta | y, z) \propto p(\tau, \eta'(\theta_T), \delta) \times p(y, z | \tau, \eta'(\theta_T), \delta). \quad (5.46)$$

Recall that $\eta(\theta)$ has already a prior form; $N(f(\theta)^T \beta, \sigma^2 c_\phi(\theta, \theta))$, so that a prior for $\eta'(\theta)$ is already specified; $N(\frac{\partial}{\partial \theta} f(\theta)^T \beta, \sigma^2 \frac{\partial}{\partial \theta} c_\phi(\theta, \theta))$.

As y and z are not independent, a form for $p(y, z | \tau, \eta'(\theta_T), \delta)$ is required.

However, this can be obtained by noting that y and z are *conditionally independent* given $\tau, \eta'(\theta_T)$ and δ (written $y \perp\!\!\!\perp z | \eta'(\theta), \tau, \delta$). Formally, it is required to show that:

$$p(y, z | \eta', \tau, \delta) = p(y | \eta', \tau, \delta) \times p(z | \eta', \tau, \delta) \quad (5.47)$$

Taking the left hand side first:

$$\begin{aligned} p(y, z | \eta', \tau, \delta) &= p(y = a, z = b | \eta' = c, \tau = d, \delta = e) \\ &= p(y = a, I(\eta' + \zeta) = b | \eta' = c, \tau = d, \delta = e) \\ &= p(y = a, I(c + \zeta) = b | \eta' = c, \tau = d, \delta = e) \\ &= p(y = a | \eta' = c, \tau = d, \delta = e) \\ &\quad \times p(I(c + \zeta) = b | \eta' = c, \tau = d, \delta = e) \\ &= p(y | \eta', \tau, \delta) \times p(z | \eta', \tau, \delta) \end{aligned} \quad (5.48)$$

since y and ζ are independent, and conditional independence is a consequence of standard independence. As already shown in equation 5.41, once η' has been observed (and is known to be Normally distributed), the distribution of $y | \eta' \dots$ is also Normally distributed.

A form for the second term on the right hand side of equation 5.48 can

be written out, for example, as:

$$\begin{aligned}
 p(\mathbf{z}|\boldsymbol{\eta}'(\boldsymbol{\theta}_T), \boldsymbol{\tau}, \delta) &= p \left(\begin{array}{c} \eta'(\theta_{T,1}) + \zeta_1 < 0 \\ \eta'(\theta_{T,2}) + \zeta_2 > 0 \\ \eta'(\theta_{T,3}) + \zeta_3 > 0 \\ \eta'(\theta_{T,4}) + \zeta_4 < 0 \\ \vdots \end{array} \middle| \boldsymbol{\tau}, \delta \right) \\
 &= \prod_{i=1}^k p(\zeta > -\eta'(\theta_i)) \prod_{j=k+1}^n p(\zeta < -\eta'(\theta_j)) \\
 &= \prod_{i=1}^k \Phi_{(0,\delta)}(\eta'(\theta_i)) \prod_{j=k+1}^n \Phi_{(0,\delta)}(-\eta'(\theta_j)) \quad (5.49)
 \end{aligned}$$

where k is the number of positive tendencies from a set of size n . Thus, using this framework, the tendency information can be utilised as part of the Gaussian processes framework.

5.7 Spatial aspects

Besides the data already mentioned, Ordnance Survey co-ordinates for each data point are also available. It appears from plots like figure 4.6 that RSL change across the Humber is relatively small compared with that over time. However, the spatial information can be added into the mean and covariance of the Gaussian process.

The technical change is with respect to the RSL height y which before depended only on the date θ . The elevation vector y now becomes a function of coordinates, λ and ω , which together will be denoted by Λ .

The changes can be seen as:

$$\mathbb{E}[y_i|x_i, \Lambda_i] = f(\theta_i, \Lambda_i)^T \beta,$$

and

$$A = \begin{bmatrix} 1 & c(\theta_1, \theta_2) \times c(\Lambda_1, \Lambda_2) & \dots & c(\theta_1, \theta_n) \times c(\Lambda_1, \Lambda_n) \\ c(\theta_2, \theta_1) \times c(\Lambda_2, \Lambda_1) & 1 & & \vdots \\ \vdots & & \ddots & \\ c(\theta_n, \theta_1) \times c(\Lambda_n, \Lambda_1) & \dots & & 1 \end{bmatrix},$$

where $c(.,.)$ is defined as before. Here, A has been given a separable covariance structure. This allows the covariance matrix to stay positive definite whilst incorporating elements in both time and space independently. The conditional results for $\eta(\theta, \Lambda)$ (formerly $\eta(\theta)$) are also the same.

It should be noted that the covariance in the model corresponding to location is now calculated as a function of distance between a two-dimensional location rather than just a one-dimensional date. A standard metric for calculating the distance between points Λ_1 and Λ_2 is given as:

$$\|\Lambda_1 - \Lambda_2\| = (\Lambda_1 - \Lambda_2)B(\Lambda_1 - \Lambda_2)^T, \quad (5.50)$$

where B is a positive definite matrix with, in this case, dimension 2. When $B = cI$, the covariance structure is known as isotropic, and the separation vector h in the covariance function $c(h)$ is calculated

based purely on the scaled Euclidean distance between sites. More complicated structure for B allows the covariance function to become anisotropic; the covariance function depends upon both distance and direction. A graph showing the effect of different B is shown in figure 5.2, with $B = I$ in the left panel and $B = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}$ in the right panel.




Figure 5.2: Isotropic correlation function (left panel). Anisotropic correlation (right panel).

Bayesian methods for computing the posterior distribution of B have been discussed in Ecker and Gelfand (1999). They impose a Wishart prior on B so that:

$$p(\mathbf{B}) = |\mathbf{B}|^{(p-n-1)/2} \exp\left(-\frac{1}{2}\text{trace}(p\mathbf{B}\mathbf{R}^{-1})\right), \quad (5.51)$$

with \mathbf{R} a scale matrix and $p \geq n$ a precision parameter. The posterior elements of B are relatively uninformative individually, so Ecker and Gelfand (1999) propose a selection of statistics that capture the information in B . These include:

- The range r_ι in the direction ι , given by

$$r_\iota = \frac{c}{\left(\tilde{\mathbf{h}}B\tilde{\mathbf{h}}^T\right)^{1/2}}.$$

Here, $\tilde{\mathbf{h}} = (\cos \iota, \sin \iota)$ is a unit vector in direction ι , and c is a constant determined by the point at which covariance reaches a defined limit, eg 0.05.

- The orientation of the ellipse, χ , of the anisotropy matrix is given by

$$\chi = \arg \max_{\iota} r_\iota.$$

- The ratio of anisotropy (or affinity), λ , given by the ratio of the major axis of the ellipse to the minor ellipse axis:

$$\lambda = \frac{r_\chi}{r_{\pi-\chi}}.$$

Estimates of these parameters provide information on the anisotropy present in the data. For example, for the right panel of figure 5.50, where $B = \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix}$, χ is given as 135° , r_χ is 1.47 in this direction (assuming a Matérn correlation with unit scale parameter and smoothness $\nu = 3/2$) whilst the ratio of anisotropy is 1.73. A graphical method for summarising the data is given in the rose diagram of figure 5.3. The rose diagram gives an indication of both the direction of anisotropy and its relative strength.




Figure 5.3: Rose diagram to illustrate anisotropy. The red line is χ , the direction of anisotropy.

5.8 Predictive distributions

The problem remains to estimate the distribution of $\eta(\cdot)$ conditional only on the observed data \mathbf{y} and the derivative information \mathbf{z} . Let $\boldsymbol{\tau}$ be the set of parameters $(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \phi, \sigma^2)$. The predictive distribution of $\eta(\cdot)$ given the data can be obtained from:

$$p(\eta(\cdot)|\mathbf{y}, \mathbf{z}) = \int p(\eta(\cdot), \boldsymbol{\tau}|\mathbf{y}, \mathbf{z}) d\boldsymbol{\tau}. \quad (5.52)$$

The integrand can be re-arranged as:

$$p(\eta(\cdot), \boldsymbol{\tau}|\mathbf{y}, \mathbf{z}) = p(\eta(\cdot)|\boldsymbol{\tau}, \mathbf{y}, \mathbf{z}) p(\boldsymbol{\tau}|\mathbf{y}, \mathbf{z}), \quad (5.53)$$

which is the posterior density of the parameters given the data multiplied by the likelihood of $\eta(\cdot)$. This is analogous to finding a posterior predictive density for a new observation given the data.

To simulate values from the posterior predictive density it is possible to sample from $\eta(\cdot)|y, z, \tau$ (from, for example, equation 5.41) using suitably chosen values from $\tau|y, z$ as obtained from equation 5.8. Thus using a simulation approach it is now possible to simulate from the desired distribution. Summary statistics can be calculated using normal Monte Carlo estimates.

5.9 Dating uncertainties

As described in section 2.3.8, the procedure for dealing with the radiocarbon age is stochastic. Whilst much of the preceding notation has suppressed explicit recognition of conditionality on the dates θ , the uncertainty in them needs to be taken into account for the overall distribution of $\eta(\cdot)$. The technique proposed here is to use simulations from the posterior density of the dates as non-stochastic inputs into the models outlined above. The Monte Carlo Markov chain would then be run for each of the simulations from the posterior date distribution. This is, unfortunately, computationally expensive.

The theoretical justification of this approach is obtained through simple conditional probability. Section 5.8 shows how a distribution for $p(\eta(\cdot)|y, z)$ can be obtained. With the conditionality upon the dates θ , this is $p(\eta(\cdot)|y, z, \theta)$. Expanding the notation slightly, let the set of *uncalibrated* radiocarbon dates be given as x , with associated errors ϵ . Now let the set of calibrated calendar dates be $\theta|x, \epsilon$. As in section 5.8,

the conditionality on θ needs to be removed. Note that:

$$p(\eta()|y, z, x, \epsilon) = \int p(\eta(), \theta|y, z, x, \epsilon)d\theta, \quad (5.54)$$

where the integrand can be re-arranged as:

$$\begin{aligned} p(\eta(), \theta|y, z, x, \epsilon) &= p(\eta()|\theta, y, z) \times p(\theta|y, z, x, \epsilon) \\ &= p(\eta()|\theta, y, z) \times p(\theta|x, \epsilon), \end{aligned} \quad (5.55)$$

since θ , and (y, z) are independent. Thus the posterior calibrated date distribution can easily be taken into account by simulating from it whilst examining $\eta(\cdot)$.

5.10 Markov chain Monte Carlo

Markov chain Monte Carlo (MCMC) is a technique for sampling from posterior distributions. It is especially useful when the distribution is only known up to a constant of proportionality. This is when the denominator of the following equation is hard to calculate:

$$P(\Delta|D) = \frac{P(\Delta)P(D|\Delta)}{\int P(\Delta)P(D|\Delta)d\Delta}, \quad (5.56)$$

where Δ is the set of parameters and D is the data.

Suppose the requirement is to simulate values from a distribution $\pi(\cdot)$. The method works by setting up a Markov Chain moving from state X_0 to X_n after n iterations based on the probabilities $P(X_{t+1}|X_t)$. After

a sufficient number of iterations with the correct setup the transition probabilities should be almost independent of the starting values (ie $P(X_n|X_0) \approx 0$), and the values of X_{t+1}, X_{t+2}, \dots should be samples from $\pi()$. The first t iterations here are known as the burn-in period. Occasionally, the state vectors X_i exhibit strong autocorrelation. The autocorrelation is removed by thinning the iterations by a factor k ; the sample used as representing $\pi()$ is that of $X_k, X_{2k}, X_{3k}, \dots$

The results of the simulations, X_i can then be summarised using Monte Carlo integration, for example:

$$\mathbb{E}[f(X)] \approx \frac{1}{n} \sum_{t=1}^n f(X_t). \quad (5.57)$$

Constructing Markov Chains which converge to the required distribution, $\pi()$, is surprisingly easy. Two such methods are discussed below.

5.10.1 The Metropolis-Hastings Algorithm

The Metropolis-Hastings (MH) algorithm comes initially from the work of Metropolis *et al.* (1953), and subsequently Hastings (1970). The technique requires the setting up of a proposal distribution, $q()$, from which it is easy to sample. The value sampled from this proposal distribution may be accepted or rejected by a rule concerning the probability of this value being from the required stationary distribution. For example, suppose the most recently accepted value sampled from the proposal distribution was X_t . The next value sampled is Y_{t+1} . If this

value is accepted then $X_{t+1} = Y_{t+1}$, and X_{t+1} is the latest sampled value. If it is rejected then X_{t+1} is set to X_t and the process is started again. The acceptance rate is governed by the ratio:

$$\alpha(X, Y) = \frac{\pi(Y)q(X|Y)}{\pi(X)q(Y|X)}. \quad (5.58)$$

Here, the candidate point from the proposal distribution is Y and the most recently accepted value is X . The value $\pi(Y)$ is the value of the likelihood multiplied by the prior such as that in equation 5.8. This does not need to be a probability distribution, since the ratio of $\pi(Y)/\pi(X)$ in $\alpha(X, Y)$ allows normalising factors to cancel out. In practice, the user would sample a uniform random variable on the range $(0, 1)$ and compare it with the value of $\alpha(X, Y)$. If the numerator considers that the value of Y is preferable over X then Y will be accepted.

The proposal distribution is often quoted in the form $q(.|.)$ and is evaluated with the conditional value as its mean. This does not need to be the case; the proposal distribution can be independent of everything apart from its argument. In this case the proposal distribution is known as an independence sampler.

5.10.2 The Metropolis Algorithm

This is a subset of the MH algorithm for symmetric proposal distributions, ie $q(Y|X) = q(X|Y)$. This occurs, for example, when $q(Y|X)$ is normal density with mean X . Now the acceptance ratio

reduces to:

$$\alpha(X, Y) = \frac{\pi(Y)}{\pi(X)} \quad (5.59)$$

If improper flat priors are used for the chosen variable, the acceptance probability will always accept values of higher likelihood, and will accept values with lower likelihood with a probability depending on the simulated uniform random variable it is compared to. The choice of proposal distribution will have a strong effect on the speed at which the MCMC converges.

5.10.3 Convergence Diagnostics

Convergence diagnostics are required when using Markov Chain Monte Carlo to assess whether the proposed chain has reached stationarity. As a by-product of this they can also estimate the length of the burn-in and the amount of thinning required, if at all. Many of the tests are simply informal methods and justification for what can be done quite informally by simply running multiple chains and looking at plots of parameter values versus iterations. All of the tests outlined below can be performed using the R package `boa` (Smith, 2005).

Brooks, Gelman and Rubin

The Brooks, Gelman and Rubin diagnostic, suggested by Gelman and Rubin. (1992) and refined by Brooks and Gelman (1998), is used when multiple chains are run with differing starting values. The method

begins by comparing the within- and between-chains variance for each of the parameters. This is adjusted to yield a Scale Reduction Factor (SRF) based on the ratio of an unbiased estimate of the true variance with the estimated mean of within-chain variances. The SRF is then corrected to obtain the Corrected Scale Reduction Factor (CSRf) by taking account of the sampling variability in the chains. If the CSRf is large (> 1) it suggests that the estimate of variance can be decreased by more simulations. If the CSRf is close to one, it is expected that the distribution has attained convergence. The Brooks, Gelman and Rubin diagnostic is used as standard in the package WinBUGS.

Geweke

The convergence diagnostic devised by Geweke (1991), requires only a single chain for its calculation. The method proceeds by creating two 'windows'; one at the start of the chain, another at the end. A Z-statistic is produced comparing the mean of the two samples, based on the asymptotic standard error of the difference. As the number of iterations increases, the value of the Z-statistic approaches the standard normal. A large value of Z may be considered to be evidence against convergence.

Heidelberger and Welch

Heidelberger and Welch (1983) proposed a two-stage test for single chains. The first test uses the Cramer-von-Mises test statistic (a test for stationarity) to compare a section (usually 10%) of the iterations to a hypothetical stationary distribution. The test continues to discard the

same proportion until the test is passed (in which case the discarded iterations are the burn-in period) or it has less than 50% of iterations left (in which case the test fails). If the parameter has passed, the remaining iterations are subjected to a further test known as the half-width test. This test compares the ratio of the mean and the half-width of the confidence interval for the chain with a pre-specified value (based upon the power of the test). If the ratio exceeds this value the test records a failure. If either test is failed then there is evidence against convergence.

Raftery and Lewis

The Raftery and Lewis (1992) method concentrates on the convergence of specified quantiles rather than means or variances as in the other tests. It will also estimate the length of the burn-in period and whether any further thinning is required. The user specifies a desired quantile and accuracy. The test then converts these values into a binary Markov chain and produces a test statistic known as the dependence factor. If this value is large (> 5) then the chain is thought not to have converged.

5.11 Example

5.11.1 Introduction

To illustrate some of the theory set out in this chapter, an example will now be considered. 20 points were simulated from the data set y where:

$$y_i = 5 - x_i - \cos(x_i) + e_i, \quad (5.60)$$

with $e_i \sim N(0, 0.2^2)$ and $x \in (0, 4)$. A plot of the 20 data points with the true line is shown in figure 5.4.




Figure 5.4: A plot of the 20 example data points simulated from $5 - x - \cos(x)$ (solid line).

5.11.2 Basal index points only

Initially, the method described in section 5.2 is used to formulate predictions for $\eta(x)$. As is the case in the full RSL model, the values of σ_i^2 , $i = 1, \dots, 20$ are assumed known to be 0.2^2 .

This setup requires a number of choices. Firstly, the dimension of f in equation 5.2 is required. Secondly priors must be specified for the various parameters. For the majority, vague normal distributions will suffice with no prior information available. For the variance parameter σ^2 the inverse gamma prior is used as it can be reasonably vague yet

remains proper. Also, as is noted by Stein (1999, pg. 224), improper priors on this parameter will lead to improper posteriors. In the first instance, the powered exponential correlation function was used. It should be noted that in this case the Gaussian covariance function would most likely provide equally good results (with 1 fewer parameter) as the function to be estimated is known to be very smooth.

The MCMC was run with both a constant and a linear prior on the data y using the WinBUGS package. More details about how WinBUGS can be used with this type of data are given in section 6.2.1. Convergence was attained very quickly; a plot of the densities for the parameters σ and ϕ is given in figure 5.5. It is clear that, in this case, adding a linear trend reduces the underlying sill variance, σ^2 . However, in the linear case this reduction is coupled with less certainty about ϕ_1 and less smoothness. Using the posterior modes for the entire set of parameter values, it is possible to plot the mean values of $\eta(x)$ for a variety of x with 95% intervals, and is shown in figure 5.6. The graph was created from the technique outlined in section 5.8 which summarises the multiple lines.

5.11.3 Limiting points and intercalated index points

The next step is to introduce points that are above and below the line. Three limiting points are created in the form:

$$y_i = 5 - x - \cos(x_i) + e_i + \gamma_i, \quad (5.61)$$

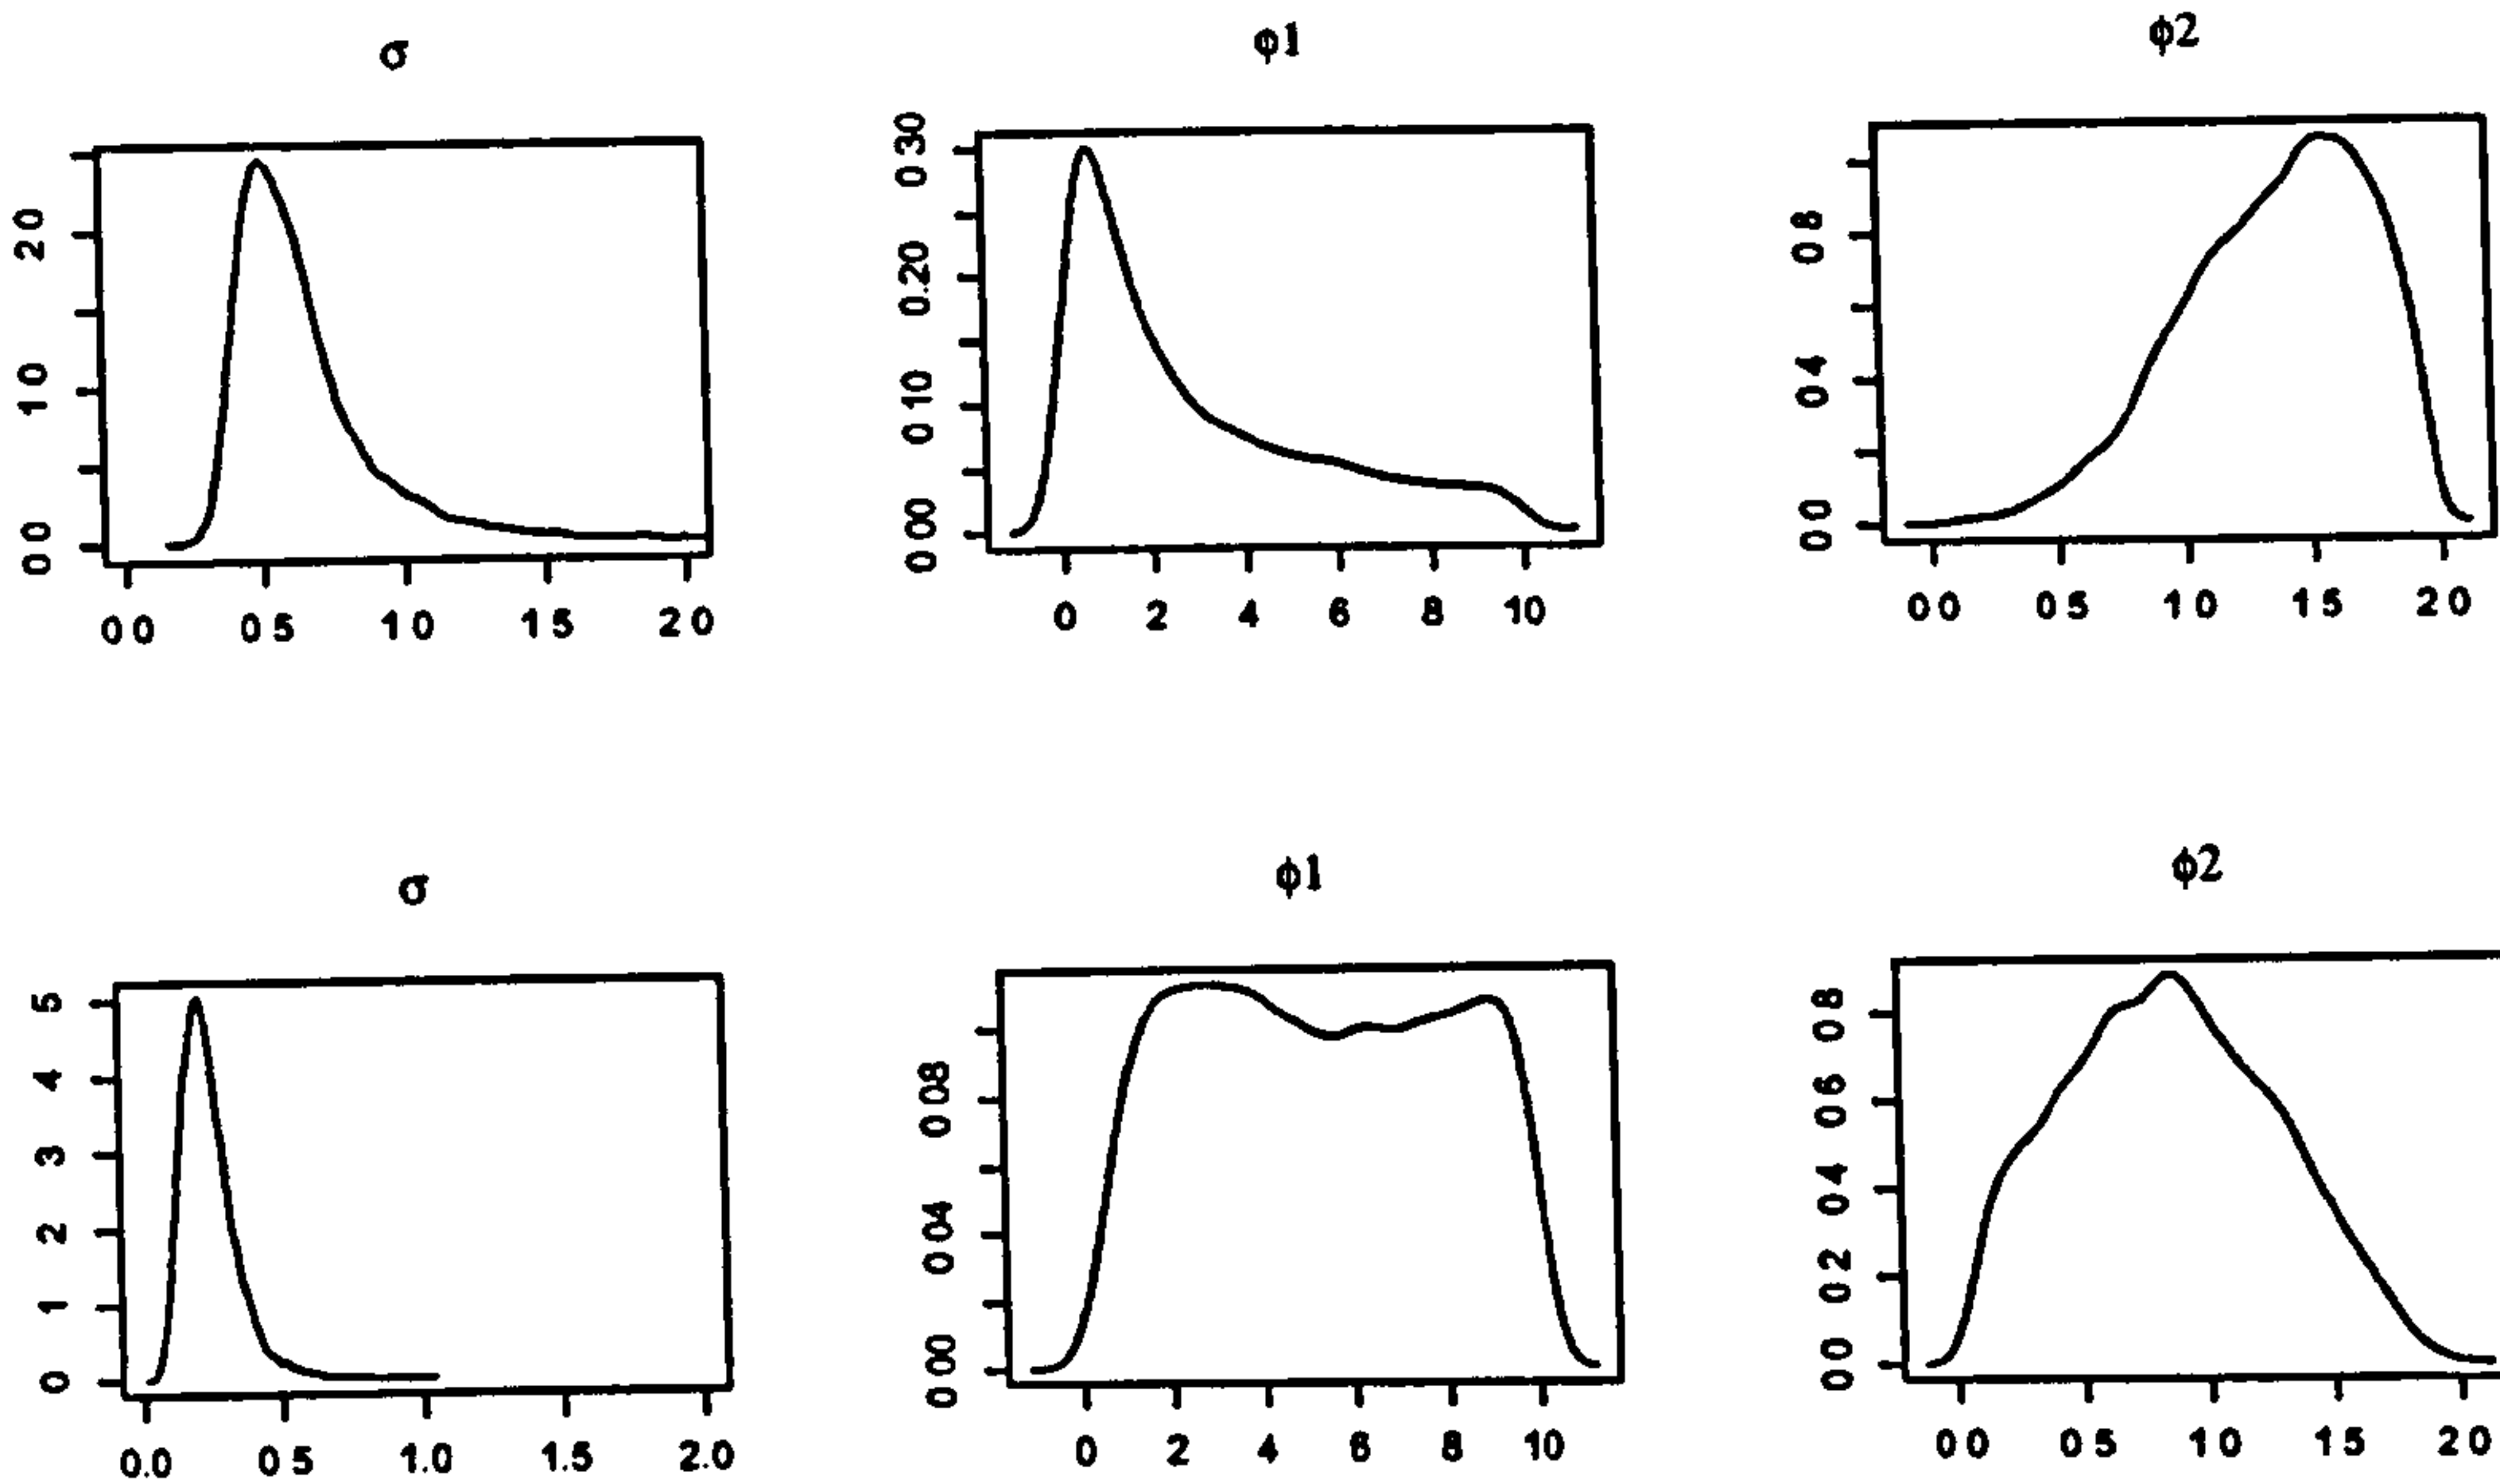


Figure 5.5: Posterior density plots of the parameters σ , ϕ_1 and ϕ_2 . The upper three graphs are for the constant prior, the lower three for the linear prior.




Figure 5.6: A plot of the Gaussian process estimates for the linear prior. The true line is shown in blue. The GP estimate is in black with dotted lines for 95% intervals.

where e_i is as before and $\gamma_i \sim N(0.5, 0.5^2)$. Three intercalated index points are also created with $\gamma_i \sim N(-0.5, 0.5^2)$. A graph of the new data

points is shown in figure 5.7.




Figure 5.7: A plot of the 6 new points added in to the model. The limiting points are in blue, the intercalated index points are in red.

The MCMC was again run in WinBUGS, this time using the model outlined in section 5.5. Again, the model converged in a short time, approximately 500 iterations. A graph of the standard deviations along the range of x is shown in figure 5.8. It is clear that the standard deviations are, as expected, reduced slightly with the incorporation of the limiting points.

points is shown in figure 5.7.




Figure 5.8: A comparison of the standard error of $\eta(x)$. The blue and red (dotted) lines show the standard error before and after inclusion of the limiting points and intercalated index points respectively.

Figure 5.7: A plot of the 6 new points added in to the model. The limiting points are in blue, the intercalated index points are in red.

The MCMC was again run in WinBUGS, this time using the model outlined in section 5.5. Again, the model converged in a short time, approximately 500 iterations. A graph of the standard deviations along the range of x is shown in figure 5.8. It is clear that the standard deviations are, as expected, reduced slightly with the incorporation of the limiting points.

The results are shown in figure 5.9. The graph clearly deviates from the path when the derivative information is included in a desirable fashion.

Derivatives may also be used when they are known only with error




Figure 5.8: A comparison of the standard error of $\eta(x)$. The blue and red (dotted) lines show the standard error before and after inclusion of the limiting points and intercalated index points respectively.

5.11.4 Derivative information

Information about the derivative of the curve is now included in the model, as described in section 5.6. In the first instance, derivatives are given as known at the points $x = 1, x = 2.5$ and $x = 3.5$. The Matérn correlation function is now used for its desirable derivative properties. The parameter ϕ_2 is forced to be greater than 1 to enable the function to be at least once differentiable. The results are shown in figure 5.9. The graph clearly deviates from the path when the derivative information is included in a desirable fashion.

Derivatives can also be used when they are known only with error




Figure 5.9: A comparison of $\eta(x)$ with and without derivative information. The red line is the true line. The line without derivative information is shown in blue, the line with derivative information is shown in green. The points where derivative information is given are marked with an X.

variance δ^2 (outlined in equation 5.46). As would be expected, the larger the value of δ^2 , the less influence a data point has over the shape of $\eta(x)$. In the extreme case that $\delta \rightarrow 0$, the true derivative information is obtained and a graph similar to that in figure 5.9 is produced.

A problem remains in the joint posterior distribution of $(\eta'(x_i), \delta^2 | z_i)$. When improper priors are used this posterior is also improper. Thus simulated values of δ^2 tend to become very large, so values of $\eta'(x)$ are

ignored. The problem is similar, yet slightly less severe, when priors on these parameters are proper but vague. The parameters only start to inform when reasonably tight priors are given. An example is shown in figure 5.10. It would appear that tendency information, which is similar to derivative information with large error, has little influence on line estimates.




Figure 5.10: A plot of the example data set with different strengths of prior derivative information. The true line is shown in red. The line with vague prior information is in blue. The line with strong prior information is shown in black. Data points where derivative information is given are shown in green.

5.11.5 Non-stationary covariance

The non-stationary covariance function proposed by Paciorek and Schervish (2004) is inappropriate for this example as the true curve displays constant smoothness by design. To display the usefulness of this technique the example data was changed slightly to include variable smoothness:

$$y_i = \begin{cases} e_i & \text{if } x_i < 0 \\ \sin(x_i) + e_i & \text{if } x_i \geq 0 \end{cases} \quad (5.62)$$

with $e_i \sim N(0, 0.1^2)$. The Gaussian process prior on $\log(\Omega)$ was given a constant mean and a stationary Matérn covariance structure. It is clear from the structure of the data (and figure 5.11) that, were a horizontal line fitted to the data, no autocorrelation function would be required for $x < 0$. However, for $x > 0$ a reasonable level of autocorrelation would be expected. Thus any parameters representing stationarity would be required to change around $x = 0$.

The model was run in R for a large number of iterations to give estimates of non-stationarity. A plot of the line fits with 95% intervals is shown in figure 5.12. It is clear that the non-stationarity GP parameters show a marked change in level; the resulting line estimate fits the data well.




Figure 5.11: A plot of the data used to illustrate non-stationarity. The function from which the data points were sampled from is shown in blue.

5.12 Summary

This chapter has shown how Gaussian processes can overcome the problems presented by sea-level data. These problems include; the drawing of accurate sea-level curves, the adjustment of limiting points and intercalated index points, the use of spatial information, the incorporation of point-wise errors, the use of calibrated radiocarbon dates and the use of the tendency measure through derivative information.

Understanding and describing the results produced by the GP technique remain the main statistical hurdle to be overcome. The output shown in the figures of this chapter have focussed on point-wise




Figure 5.12: A plot of point-wise means with 95% error bands from the non-stationary GP estimate. The log GP for the non-stationarity parameters is shown in the top left-hand corner. The true function is shown in blue, the data points in red.

A selection of models are examined each with increasing posterior means and variances. However, the MCMC output produces complete distributions which should be utilised to fully describe the behaviour of the sea-level curves. As is shown in chapter 7, simply looking at point-wise means and variances is insufficient; more complex ways are examined.

Chapter 6

Modelling RSL

6.1 Introduction

This chapter applies the theory outlined in chapter 5 to the data sets provided. A selection of models are examined each with increasing complexity. They are of two kinds; firstly those containing the temporal information only (for comparison with the traditional method for producing RSL curves assuming no spatial component), secondly utilising the spatial information to examine whether any spatio-temporal variation exists within the Humber. The models are compared (within the two kinds) using the deviance information criterion (DIC).

The output of the models fitted in this chapter is expected to produce estimates with uncertainty concerning the RSL over the course of the Holocene within the Humber Estuary. The functional output, $\eta(x)$, is only analysed summarily in this chapter. This analysis includes:

- An overview of the model fit and covariance parameters.

- An outline of the MCMC convergence with estimated burn-in times and details of thinning.
- A general description of the manner of relative sea-level change over the last 10,000 years.
- An assessment of the importance of the location in the estimation of sea level.
- An estimate of the DIC of the particular model.

More detailed analysis which may be of more use to those interested in geological aspects of RSL are given in chapter 7.

6.2 Temporal models

Initially, the Humber data was modelled using the Gaussian process technique outlined in chapter 5. The full set of index points, limiting points, intercalated index points and all the archaeological data were used. Different types of model were considered:

- A GeoBUGS model utilising only the modal calibrated dates and a powered exponential covariance structure.
- A stationary Matérn GP model.
- A stationary Matérn GP model incorporating the tendency information.
- A non-stationary Matérn GP model.

6.2.1 GeoBUGS

A GeoBUGS Gaussian process model was run on the Humber data.

This first model contains the data in three parts:

- The elevations in metres, y ,
- The median calibrated dates, x ,
- A diagonal matrix of elevation variances, Σ .

The Gaussian process model can be set out in terms of a hierarchical multivariate Normal distribution:

$$y \sim N(\eta(x) + h, \Sigma + \Psi) \quad (6.1)$$

$$\eta(x) \sim GP(F\beta, \sigma^2 A) \quad (6.2)$$

where

$$F = \begin{pmatrix} 1 & x_1 & x_1^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & x_n^2 \end{pmatrix}, \quad \beta = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \end{pmatrix}, \quad h = Kh^*, \quad h^* = \begin{pmatrix} 0 \\ h_{\text{IIP}} \\ h_{\text{LP}} \end{pmatrix}, \quad K = \begin{pmatrix} 0 & 0 & 0 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ \vdots & \vdots & \vdots \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 1 \end{pmatrix}$$

$$A = [a_{ij}] = \exp(-\phi_1(x_i - x_j)^{\phi_2}), \Psi = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \kappa_{IIP}^2 I_{nIIP} & 0 \\ 0 & 0 & \kappa_{LP}^2 I_{nLP} \end{pmatrix}$$

A quadratic prior through the origin is specified for the mean of the Gaussian process, corresponding to prior knowledge that the rising rate of sea-level change has not remained constant. A quadratic regression through the origin is usually the analysis tool used for many sea-level data sets (eg Metcalfe *et al.*, 2000) so is used here as prior structure. The intercalated index points and limiting points (including the archaeological data points) are ‘shifted’ via a Normal distribution with parameters h_{IIP} , κ_{IIP} and h_{LP} , κ_{LP} respectively. This was discussed in section 5.5. The matrix A is governed by a powered exponential covariance function with parameters ϕ_1 and ϕ_2 : the superior Matérn covariance function is not available in GeoBUGS. The model yields a set of 10 parameters, $(\beta, h, \kappa, \phi, \sigma^2)$, for which inference is required. Vague priors were given to all with the exception of h , for which h_{IIP} was restricted to be negative, h_{LP} was restricted to be positive.

The model was implemented using the `spatial.exp` function. This is designed to work for two dimensional spatial data but can be ‘fooled’ by using $(x, 0)$ as the (x, y) -coordinates. The code for the models is available in appendix D. Due to the limitations of the GeoBUGS software, the date errors could not be taken into account. This was also the case with the tendency data. Convergence was achieved after approximately 3000 iterations with no thinning requirement. This was checked by using the Brooks-Gelman diagnostic on multiple MCMC

runs with differing starting values. The posterior densities for the covariance parameters σ, ϕ_1 and ϕ_2 are in figure 6.1. It is clear these parameters are well-defined by the data.

As discussed in section 5.8, predictive distributions can be obtained for the underlying function of interest $\eta(t)$ from the posterior parameters estimates for any time t . A graph of the mean with 95% intervals for a set of t -points is shown in figure 6.2. It is clear that the curve employs many deviations from the quadratic prior.

The GeoBUGS model was not taken further due to current restrictions with the software. The more flexible Matérn covariance function could not be used, similarly the radiocarbon date uncertainties. Further extensions to the model were dealt with in R.

6.2.2 Stationary covariance

The GeoBUGS model was then extended to include a stationary Matérn autocovariance function and to take account of the error in the radiocarbon dates, now known as θ . For the archaeological data, the parameter β_θ is used to denote the age distribution of the use of the artifacts (outlined in section 4.2.1). The ages in radiocarbon years are now x with associated $1-\sigma$ errors, ϵ .

As before, the data are shifted according to their status. In this model, the archaeological IIPs and LPs were given their own Normally






Figure 6.1: Posterior densities for σ, ϕ_1 and ϕ_2 from the GeoBUGS model.




Figure 6.2: Point-wise mean curves with 95% intervals for the Humber using the GeoBUGS model.

distributed shifts. An extra group was also used from the older limiting points collected during the LOIS project. These limiting points are especially old, and often represent the lowest possible limit of coring. This makes them poor indicators of former sea level compared to the more recent limiting points, which provide a stronger bound for the RSL curve. Here, the older limiting points are included with a prior bound on the quantile at which they pass through, to ensure that they almost certainly lie above the RSL curve.

In summary, six groups of index points arise:

- Basal index points from the LOIS database (requiring no offset)

- Intercalated index points from the LOIS database (offset distribution $N(h_{IIP}, \kappa_{IIP}^2)$)
- Limiting points from the LOIS database (offset distribution $N(h_{LP}, \kappa_{LP}^2)$)
- Older limiting points (offset distribution $N(h_{LP2}, \kappa_{LP2}^2)$)
- Intercalated index points from the Humber Wetlands Project (offset distribution $N(h_{IIPa}, \kappa_{IIPa}^2)$)
- Limiting points from the Humber Wetlands Project $N(h_{LPa}, \kappa_{LPa}^2)$.

Mathematically, the model is exactly as that set out in equation 6.1. However, some of the matrices change slightly with the introduction of the radiocarbon dates and the extra archaeological shifts:

$$F = \begin{pmatrix} 1 & \theta_1 & \theta_1^2 \\ \vdots & \vdots & \vdots \\ 1 & \theta_m & \theta_m^2 \\ 1 & \beta_{\theta, m+1} & \beta_{\theta, m+1}^2 \\ \vdots & \vdots & \vdots \\ 1 & \beta_{\theta, n} & \beta_{\theta, n}^2 \end{pmatrix}, \quad h^* = \begin{pmatrix} 0 \\ h_{IIP} \\ h_{LP} \\ h_{LP2} \\ h_{IIPa} \\ h_{LPa} \end{pmatrix}$$

$$A = [a_{ij}] = \frac{1}{2^{\phi_1 - 1} \Gamma(\phi_1)} \left(\frac{2\phi_1^{1/2} |\theta_i - \theta_j|}{\phi_2} \right)^{\phi_1} \mathcal{K}_{\phi_1} \left(\frac{2\phi_1^{1/2} |\theta_i - \theta_j|}{\phi_2} \right),$$

$$\Psi = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \kappa_{IIP}^2 I_{nIIP} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \kappa_{LP}^2 I_{nLP} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \kappa_{LP2}^2 I_{nLP2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \kappa_{IIPa}^2 I_{nIIPa} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \kappa_{LPa}^2 I_{nLPa} \end{pmatrix}$$

The parameter set, $(\theta, \beta_\theta, \beta, h, \kappa, \phi, \sigma^2)$ now contains $16+n$ parameters where n is the total number of dates. The structure of the model can best be seen by looking at a Directed Acyclic Graph (DAG). The DAG for the model above is given in figure 6.3. Here, for convenience, define $\mu = F\beta + h$ and $C = \sigma^2 A + \Sigma + \Psi$.

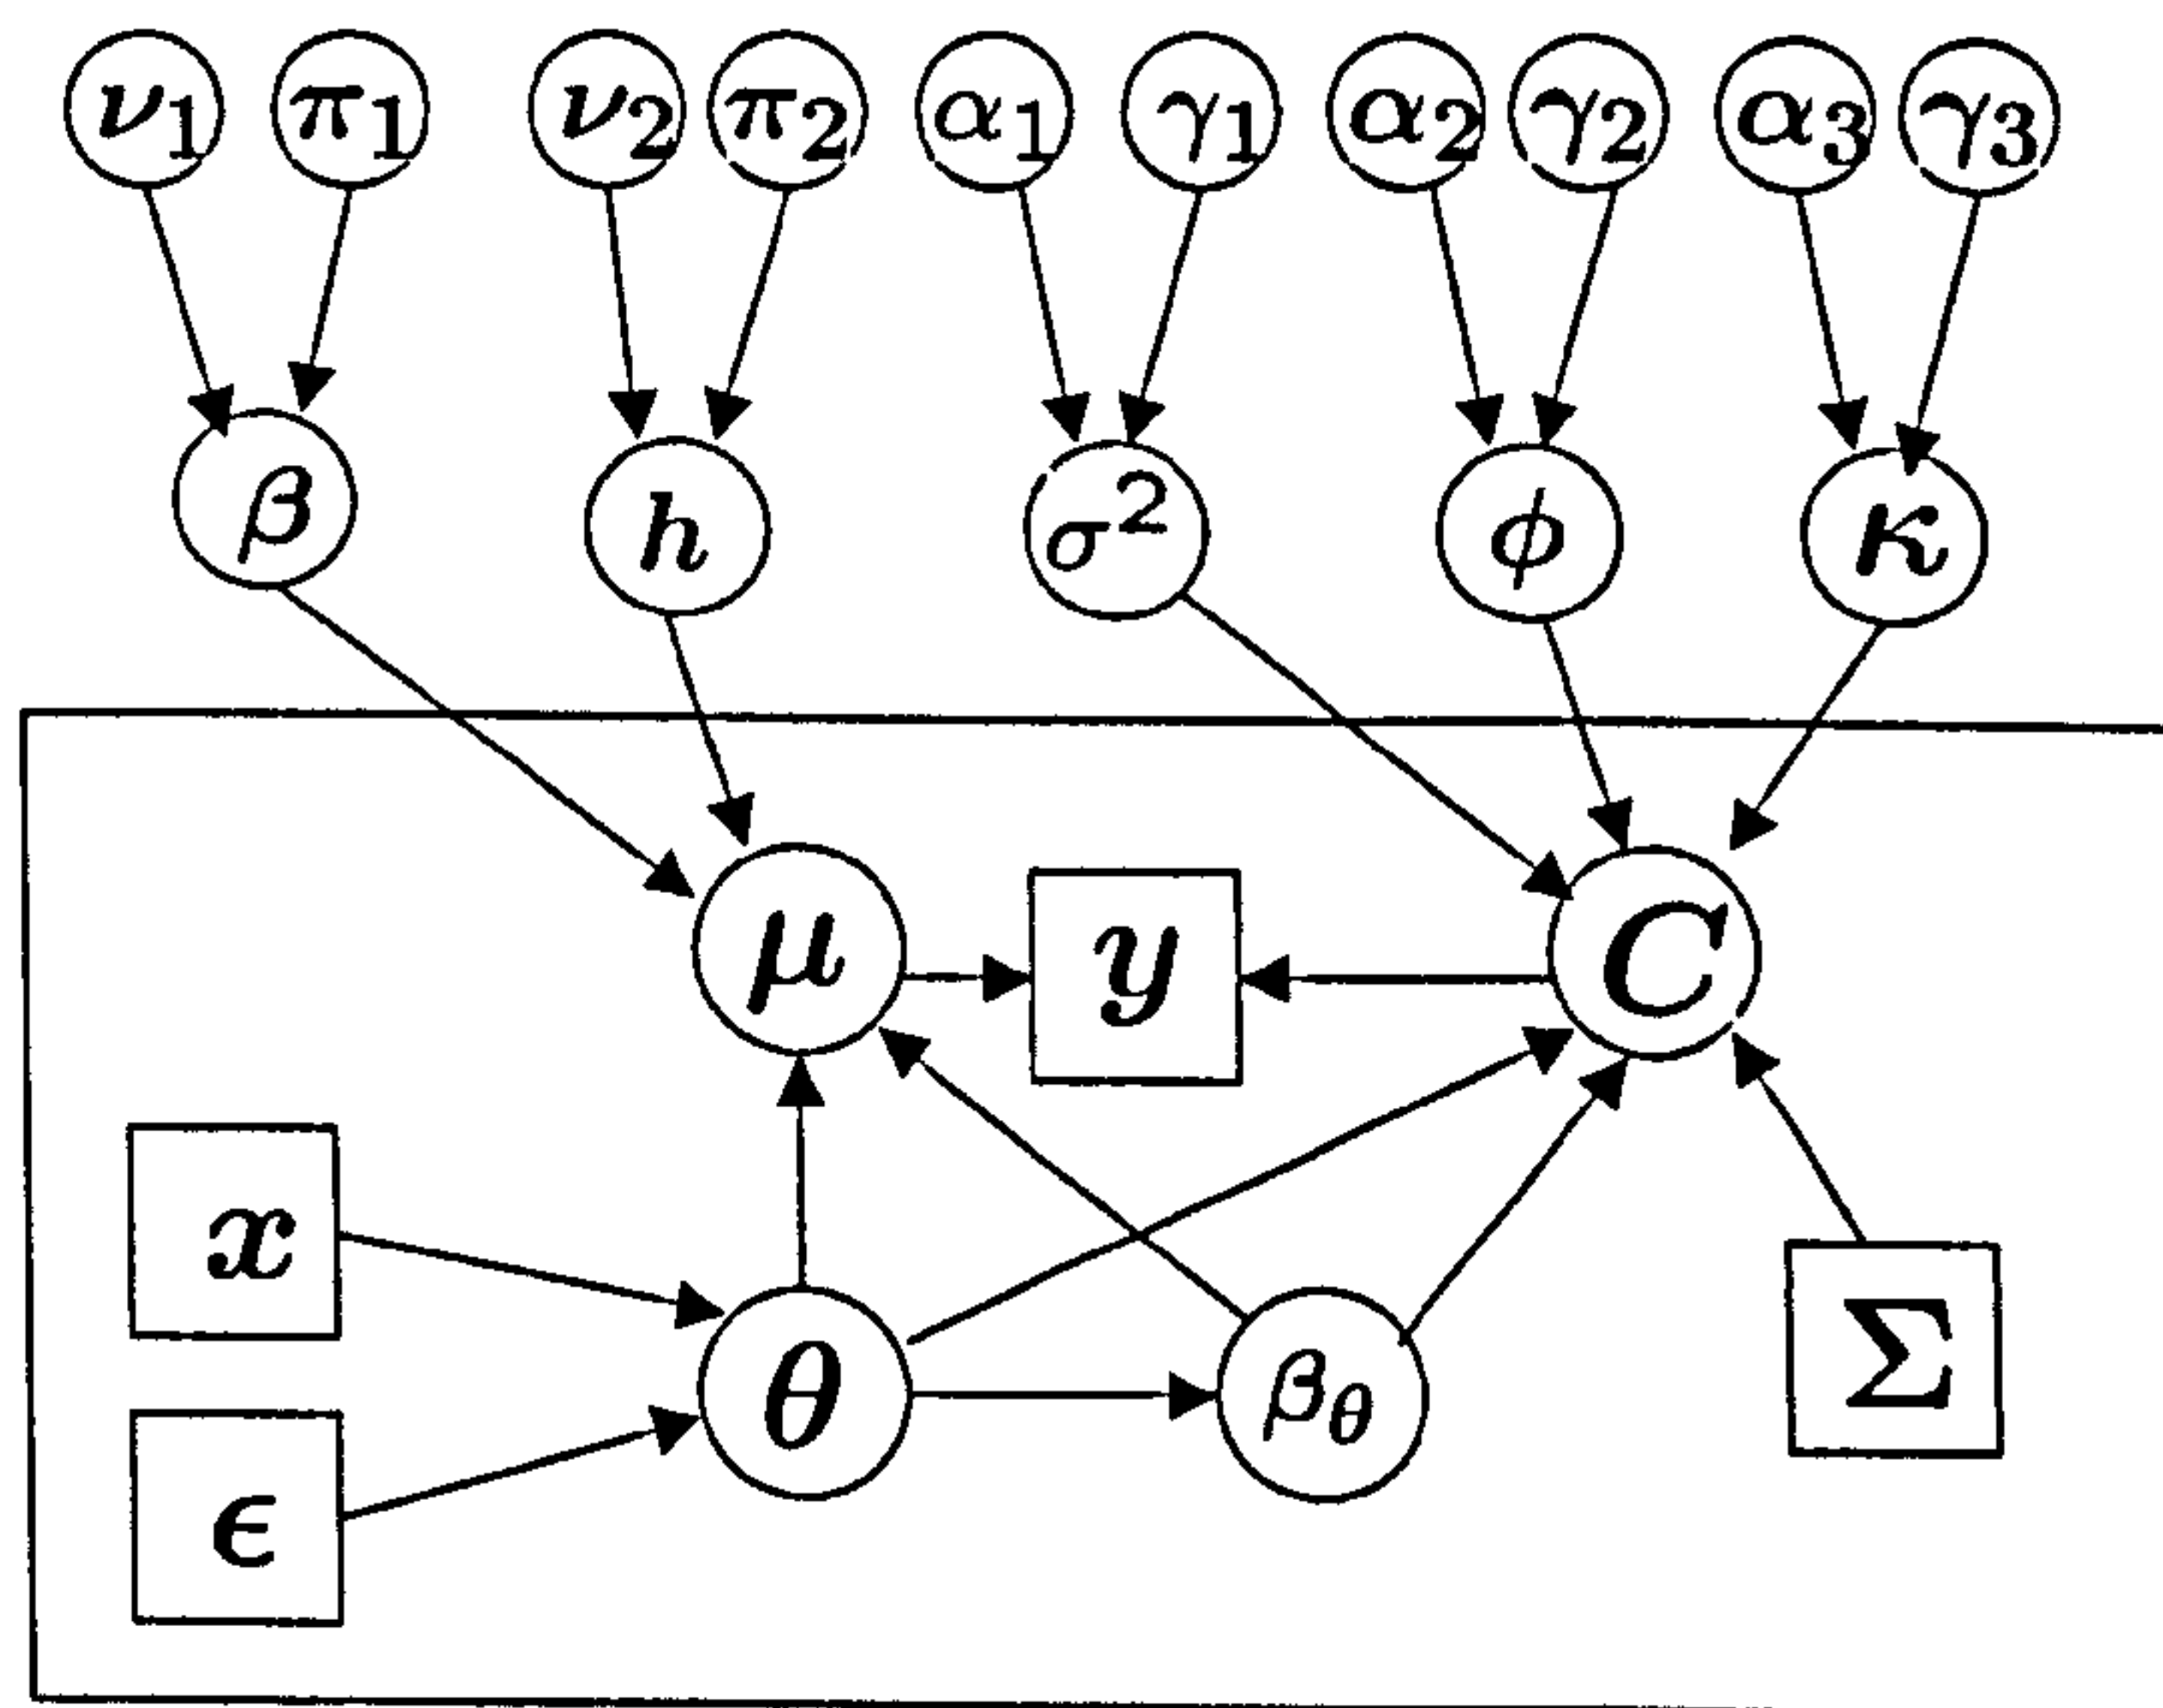


Figure 6.3: A Directed Acyclic Graph for the stationary Matérn model.

Since the distribution of y is considered known, the full or complete conditional distributions of the parameters given the data can be calculated. This is required for the MCMC. The set of parameters are

grouped together as $\tau = (\beta, \sigma^2, \phi, \kappa^2, h, \kappa^2)$. The joint distribution can now be read off the DAG in figure 6.3:

$$\begin{aligned}
p(\tau, \theta, \beta_\theta, y, x, \epsilon, \Sigma) &= p(y|\mu, C) \times p(\mu|\beta_0, \beta_1, h, \theta, \beta_\theta) \\
&\times p(C|\sigma^2, \phi, \kappa, \Sigma, \theta, \beta_\theta) \times p(\theta|x, \epsilon) \times p(\beta_\theta|\theta) \\
&\times p(\beta_0|\nu_1, \pi_1) \times p(\beta_1|\nu_2, \pi_2) \times p(h_{IIP}|\nu_3, \pi_3) \times p(h_{LP}|\nu_4, \pi_4) \\
&\times p(h_{LP2}|\nu_5, \pi_5) \times p(h_{IIPa}|\nu_6, \pi_6) \times p(h_{LPa}|\nu_7, \pi_7) \times p(\sigma^2|\alpha_1, \gamma_1) \\
&\times p(\phi|\alpha_2, \gamma_2) \times p(\kappa_{IIP}|\alpha_3, \gamma_3) \times p(\kappa_{LP}|\alpha_4, \gamma_4) \times p(\kappa_{LP2}|\alpha_5, \gamma_5) \\
&\times p(\kappa_{IIPa}|\alpha_6, \gamma_6) \times p(\kappa_{LPa}|\alpha_7, \gamma_7),
\end{aligned}$$

with $p(\mu|\beta_0, \beta_1, h, \theta, \beta_\theta) = 1$ and $p(C|\sigma^2, \phi, \kappa, \Sigma, \theta, \beta_\theta) = 1$.

For many of the parameters, it is impossible to obtain posterior distributions in closed form. It is worth noting, however, that posteriors for β and h can be expressed in simple Normal distributions conditional on the data (denoted as D here) and the other parameters. These parameters can be updated using Gibbs' steps for computational simplicity.

The parameter β is given the multivariate prior:

$$\beta \sim N(\mu_\beta, \Sigma_\beta). \quad (6.3)$$

This leads to the multivariate posterior:

$$\beta|D \sim N(M_\beta, V_\beta) \quad (6.4)$$

with

$$V_{\beta} = (\Sigma_{\beta}^{-1} + F^T C^{-1} F)^{-1}, \quad (6.5)$$

$$M_{\beta} = V_{\beta}(\Sigma_{\beta}^{-1} \mu_{\beta} + F^T C^{-1} (y - h)) \quad (6.6)$$

The vector h can be re-written as Kh^* where h^* and K are as defined previously. A joint prior is specified for h^* :

$$h^* \sim N(\mu_h, \Sigma_h). \quad (6.7)$$

This leads to the multivariate posterior:

$$h^* | D \sim N(M_h, V_h) \quad (6.8)$$

with

$$V_h = (\Sigma_h^{-1} + K^T C^{-1} K)^{-1}, \quad (6.9)$$

$$M_h = V_h(\Sigma_h^{-1} \mu_h + K^T C^{-1} (y - F\beta)) \quad (6.10)$$

The rest of the parameters are updated using single component Metropolis-Hastings. It would be possible to obtain an inverse gamma distribution for σ^2 if the model were set out in the conditional form as in equation 5.7. However, as mentioned previously, this format would dramatically increase the number of parameters. As with most Bayesian models, it is helpful to marginalise wherever possible. The posterior is calculated for σ^2 (and similarly for the other parameters), in

the form:

$$p(\sigma^2 | \mathbf{y}, \mathbf{x}, \Sigma, \epsilon, \mu, \mathbf{C}) \propto p(\sigma^2) \times p(\mathbf{y} | \mu, \mathbf{C}). \quad (6.11)$$

The prior for σ^2 is inverse gamma.

This model was implemented in the package R (R Development Core Team, 2005) with vague priors specified for the parameters. The Matérn correlation parameters were restricted so as not to obtain computationally singular covariance matrices. The MCMC chain for the Humber passed both Geweke and Raftery-Lewis convergence diagnostics after discarding a burn-in period of approximately 4000 iterations and thinning by 5. A plot of the densities of ϕ_1 , ϕ_2 and σ are shown in figure 6.4. It is clear that the scale parameter ϕ_1 is less well-defined than in the GeoBUGS model. σ also appears to be larger.

As before, simulations of $\eta(t)$ are created to produce an envelope of lines. Means and 95% intervals for the Humber are shown in figure 6.5. It is clear that there is most uncertainty in the period before 8k years cal BP and after around 2k years BP. The mean also predicts a number of fluctuations.






Figure 6.4: Posterior densities for σ, ϕ_1 and ϕ_2 for the stationary Matérn model.




Figure 6.5: Point-wise means and 95% intervals for the Humber using a stationary Matérn correlation function.

6.2.3 Stationary covariance with tendency information

The above model is now extended to include tendency information, but retains a stationary Matérn autocovariance function. A description of the origin of tendency information can be found in section 2.3.7, whilst its statistical use can be found in section 5.6.1. The t extra data points have dates θ_T and tendency values z_T . Recall that these tendencies can only take the value 1 or -1, with each z_i being defined as $I(\eta'(\theta_i) + \zeta_i)$ with $\zeta \sim N(0, \delta^2)$. As outlined in section 5.6.1, the conditional independence between the elevations y and the tendencies z_T changes the MCMC

process slightly. It can be written out as:

$$p(\tau, \eta'(\theta_T), \theta, \beta_\theta, \delta | y, z_T, x, \epsilon, \Sigma) \propto p(y, z_T, x, \epsilon, \Sigma | \tau, \eta'(\theta_T), \theta, \beta_\theta, \delta) \times p(\tau, \eta'(\theta_T), \theta, \beta_\theta, \delta) \quad (6.12)$$

$$= p(y, x, \epsilon, \Sigma | \tau, \eta'(\theta_T), \theta, \beta_\theta, \delta) \times p(z_T, x_T, \epsilon_T | \tau, \eta'(\theta_T), \delta) \times p(\tau, \eta'(\theta_T), \theta, \beta_\theta, \delta) \quad (6.13)$$

The first term in equation 6.13 is a conditional Normal distribution identical in form to that outlined in equations 5.10 and 5.11. The joint distribution of y and $\eta'()$ can be written out as:

$$\begin{pmatrix} y \\ \eta'(\theta_T) + \epsilon \end{pmatrix} \sim N \left[\begin{pmatrix} F\beta + h \\ F_T\beta \end{pmatrix}, \begin{pmatrix} \sigma^2 A + \Sigma + \Psi & \sigma^2 A_T^T \\ \sigma^2 A_T & \sigma^2 A_{TT} + \text{diag}(\delta^2) \end{pmatrix} \right], \quad (6.14)$$

From this a conditional normal distribution can be formed for $y | \eta'()$ with mean and variance:

$$\mu^* = F\beta + h + \sigma^2 A_T^T (\sigma^2 A_{TT} + \text{diag}(\delta^2))^{-1} (\eta'(\theta_T) - F_T\beta) \quad (6.15)$$

$$C^* = \sigma^2 A + \Sigma + \Psi - \sigma^2 A_T^T (\sigma^2 A_{TT} + \text{diag}(\delta^2))^{-1} \sigma^2 A_T \quad (6.16)$$

with:

$$\begin{aligned}
 F_T &= \begin{pmatrix} \frac{\partial}{\partial \theta_T} 1 \\ \frac{\partial}{\partial \theta_T} \theta_T \end{pmatrix}, \\
 A_T &= \begin{pmatrix} \frac{\partial}{\partial \theta_{n+1}} c_\phi(\theta_{n+1}, \theta_1) & \cdots & \frac{\partial}{\partial \theta_{n+1}} c_\phi(\theta_{n+1}, \beta_{\theta, n}) \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial \theta_{n+t}} c_\phi(\theta_{n+t}, \theta_1) & \cdots & \frac{\partial}{\partial \theta_{n+t}} c_\phi(\theta_{n+t}, \beta_{\theta, n}) \end{pmatrix}, \\
 A_{TT} &= \begin{pmatrix} \frac{\partial^2}{\partial \theta_{n+1} \partial \theta_{n+1}} c_\phi(\theta_{n+1}, \theta_{n+1}) & \cdots & \frac{\partial^2}{\partial \theta_{n+1} \partial \theta_{n+t}} c_\phi(\theta_{n+1}, \theta_{n+t}) \\ \vdots & \ddots & \vdots \\ \frac{\partial^2}{\partial \theta_{n+t} \partial \theta_{n+1}} c_\phi(\theta_{n+t}, \theta_{n+1}) & \cdots & \frac{\partial^2}{\partial \theta_{n+t} \partial \theta_{n+t}} c_\phi(\theta_{n+t}, \theta_{n+t}) \end{pmatrix}.
 \end{aligned}$$

Other variables are left unchanged. The second term in equation 6.13 can be calculated as outlined in equation 5.49:

$$p(\mathbf{z}_T, \mathbf{x}_T, \epsilon_T | \tau, \eta'(\theta_T), \delta) = \prod_{i=1}^k \Phi_{(0, \delta)}(-\eta'(\theta_i)) \prod_{j=k+1}^t \Phi_{(0, \delta)}(\eta'(\theta_j)) \quad (6.17)$$

for k positive tendencies with tendency variance δ^2 . A new DAG for this model is shown in figure 6.6.

Again, complete conditional distributions are simply obtained from equation 6.13, including the extra date parameters θ_T and the tendency variance parameter δ . Closed distributional form can be obtained for the mean parameters β and h . The parameter β is again given the multivariate prior:

$$\beta \sim N(\mu_\beta, \Sigma_\beta). \quad (6.18)$$

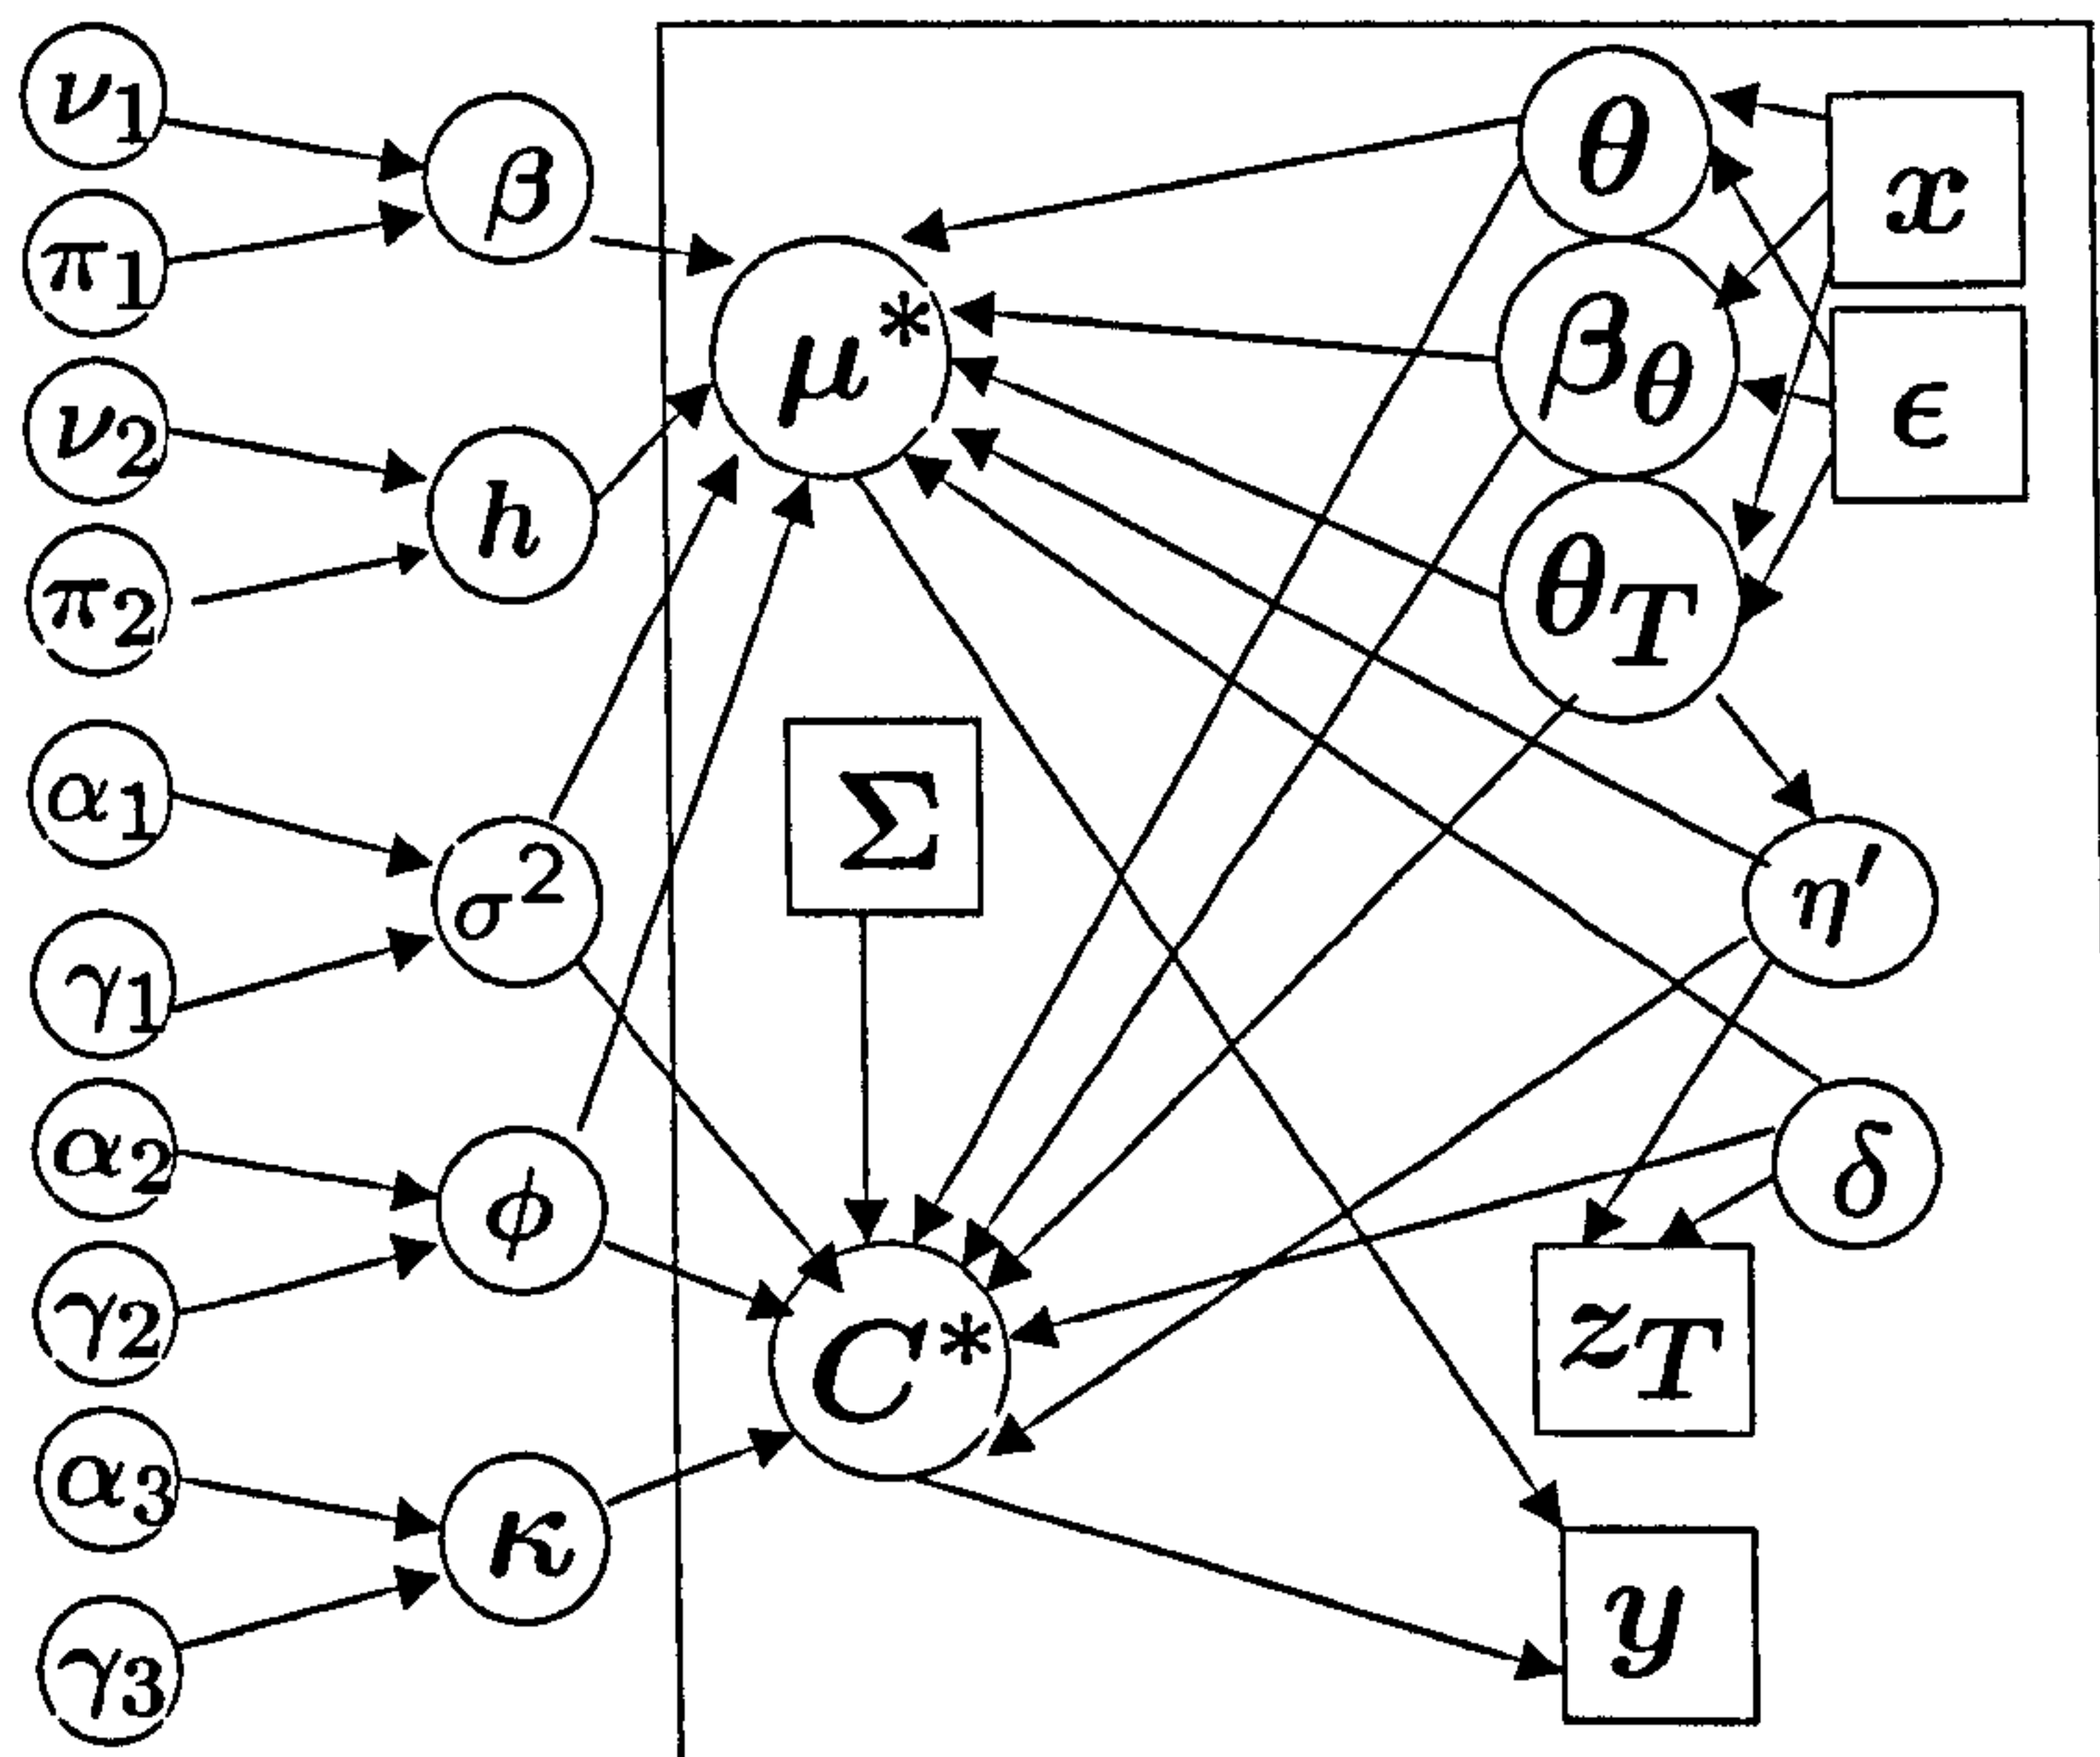


Figure 6.6: A Directed Acyclic Graph (DAG) for the model including tendency information with a stationary Matérn correlation matrix.

This leads to the multivariate posterior:

$$\beta|D \sim N(M_\beta, V_\beta) \quad (6.19)$$

with

$$V_\beta = (F^T C^{*-1} [F - J F_T] + F_T^T J^T C^{*-1} [J F_T - F] + V^{-1})^{-1}, \quad (6.20)$$

$$M_\beta = V_\beta (F^T C^{*-1} [y - h - J \eta'] + F_T^T J^T C^{*-1} [h - y] - V^{-1} \mu_\beta) \quad (6.21)$$

and

$$J = \sigma^2 A_T^T (\sigma^2 A_{TT} + \delta^2 I)^{-1}. \quad (6.22)$$

As before, the vector h can be re-written as KH . A joint prior is specified for H :

$$H \sim N(\mu_H, \Sigma_H). \quad (6.23)$$

This leads to the multivariate posterior:

$$H|D \sim N(M_H, V_H) \quad (6.24)$$

with

$$V_H = (\Sigma_H^{-1} + K^T C^{*-1} K)^{-1}, \quad (6.25)$$

$$M_H = V_H [\Sigma_H^{-1} \mu_H + K^T C^{*-1} \{y - F\beta - J(\eta' - F_T\beta)\}] \quad (6.26)$$

The remaining parameters are updated using single-component Metropolis-Hastings.

The model was required a burn-in of around 3000 iterations and was thinned by 6. Whilst the derivative information appears to inform the parameters well, the GP estimates are practically unchanged from the previous model. Figure 6.7 shows the GP estimate with 95% intervals. This model was not taken forward due to concerns that individual tendencies should not be used to inform temporal changes in sea level (see section 2.3.7). More information on how tendency information might be taken forward is given in section 8.2.3.




Figure 6.7: Point-wise means and 95% intervals for the Humber using a stationary Matérn correlation function with the tendency information.

6.2.4 Non-stationary covariance

The next model ignores the tendency information but is now extended to incorporate the non-stationarity elements of the Matérn autocorrelation function following the model of Paciorek and Schervish (2004) as outlined in section 5.4.4. The logarithm of the non-stationarity parameters Ω are given a Gaussian process prior with parameters $\beta_\Omega, \sigma_\Omega, \phi_\Omega$, and a stationary Matérn autocorrelation function:

$$\log \Omega | \theta \sim N(F\beta_\Omega, \sigma_\Omega^2 A_\Omega) \quad (6.27)$$

with

$$A_{\Omega} = \begin{pmatrix} c_{\phi_{\Omega}}(\Omega_1, \Omega_1) & \dots & c_{\phi_{\Omega}}(\Omega_1, \Omega_n) \\ \vdots & \ddots & \vdots \\ c_{\phi_{\Omega}}(\Omega_n, \Omega_1) & \dots & c_{\phi_{\Omega}}(\Omega_n, \Omega_n) \end{pmatrix} \quad (6.28)$$

The new parameters contribute additional layers of complexity into the model. A DAG for this model is shown in figure 6.8.

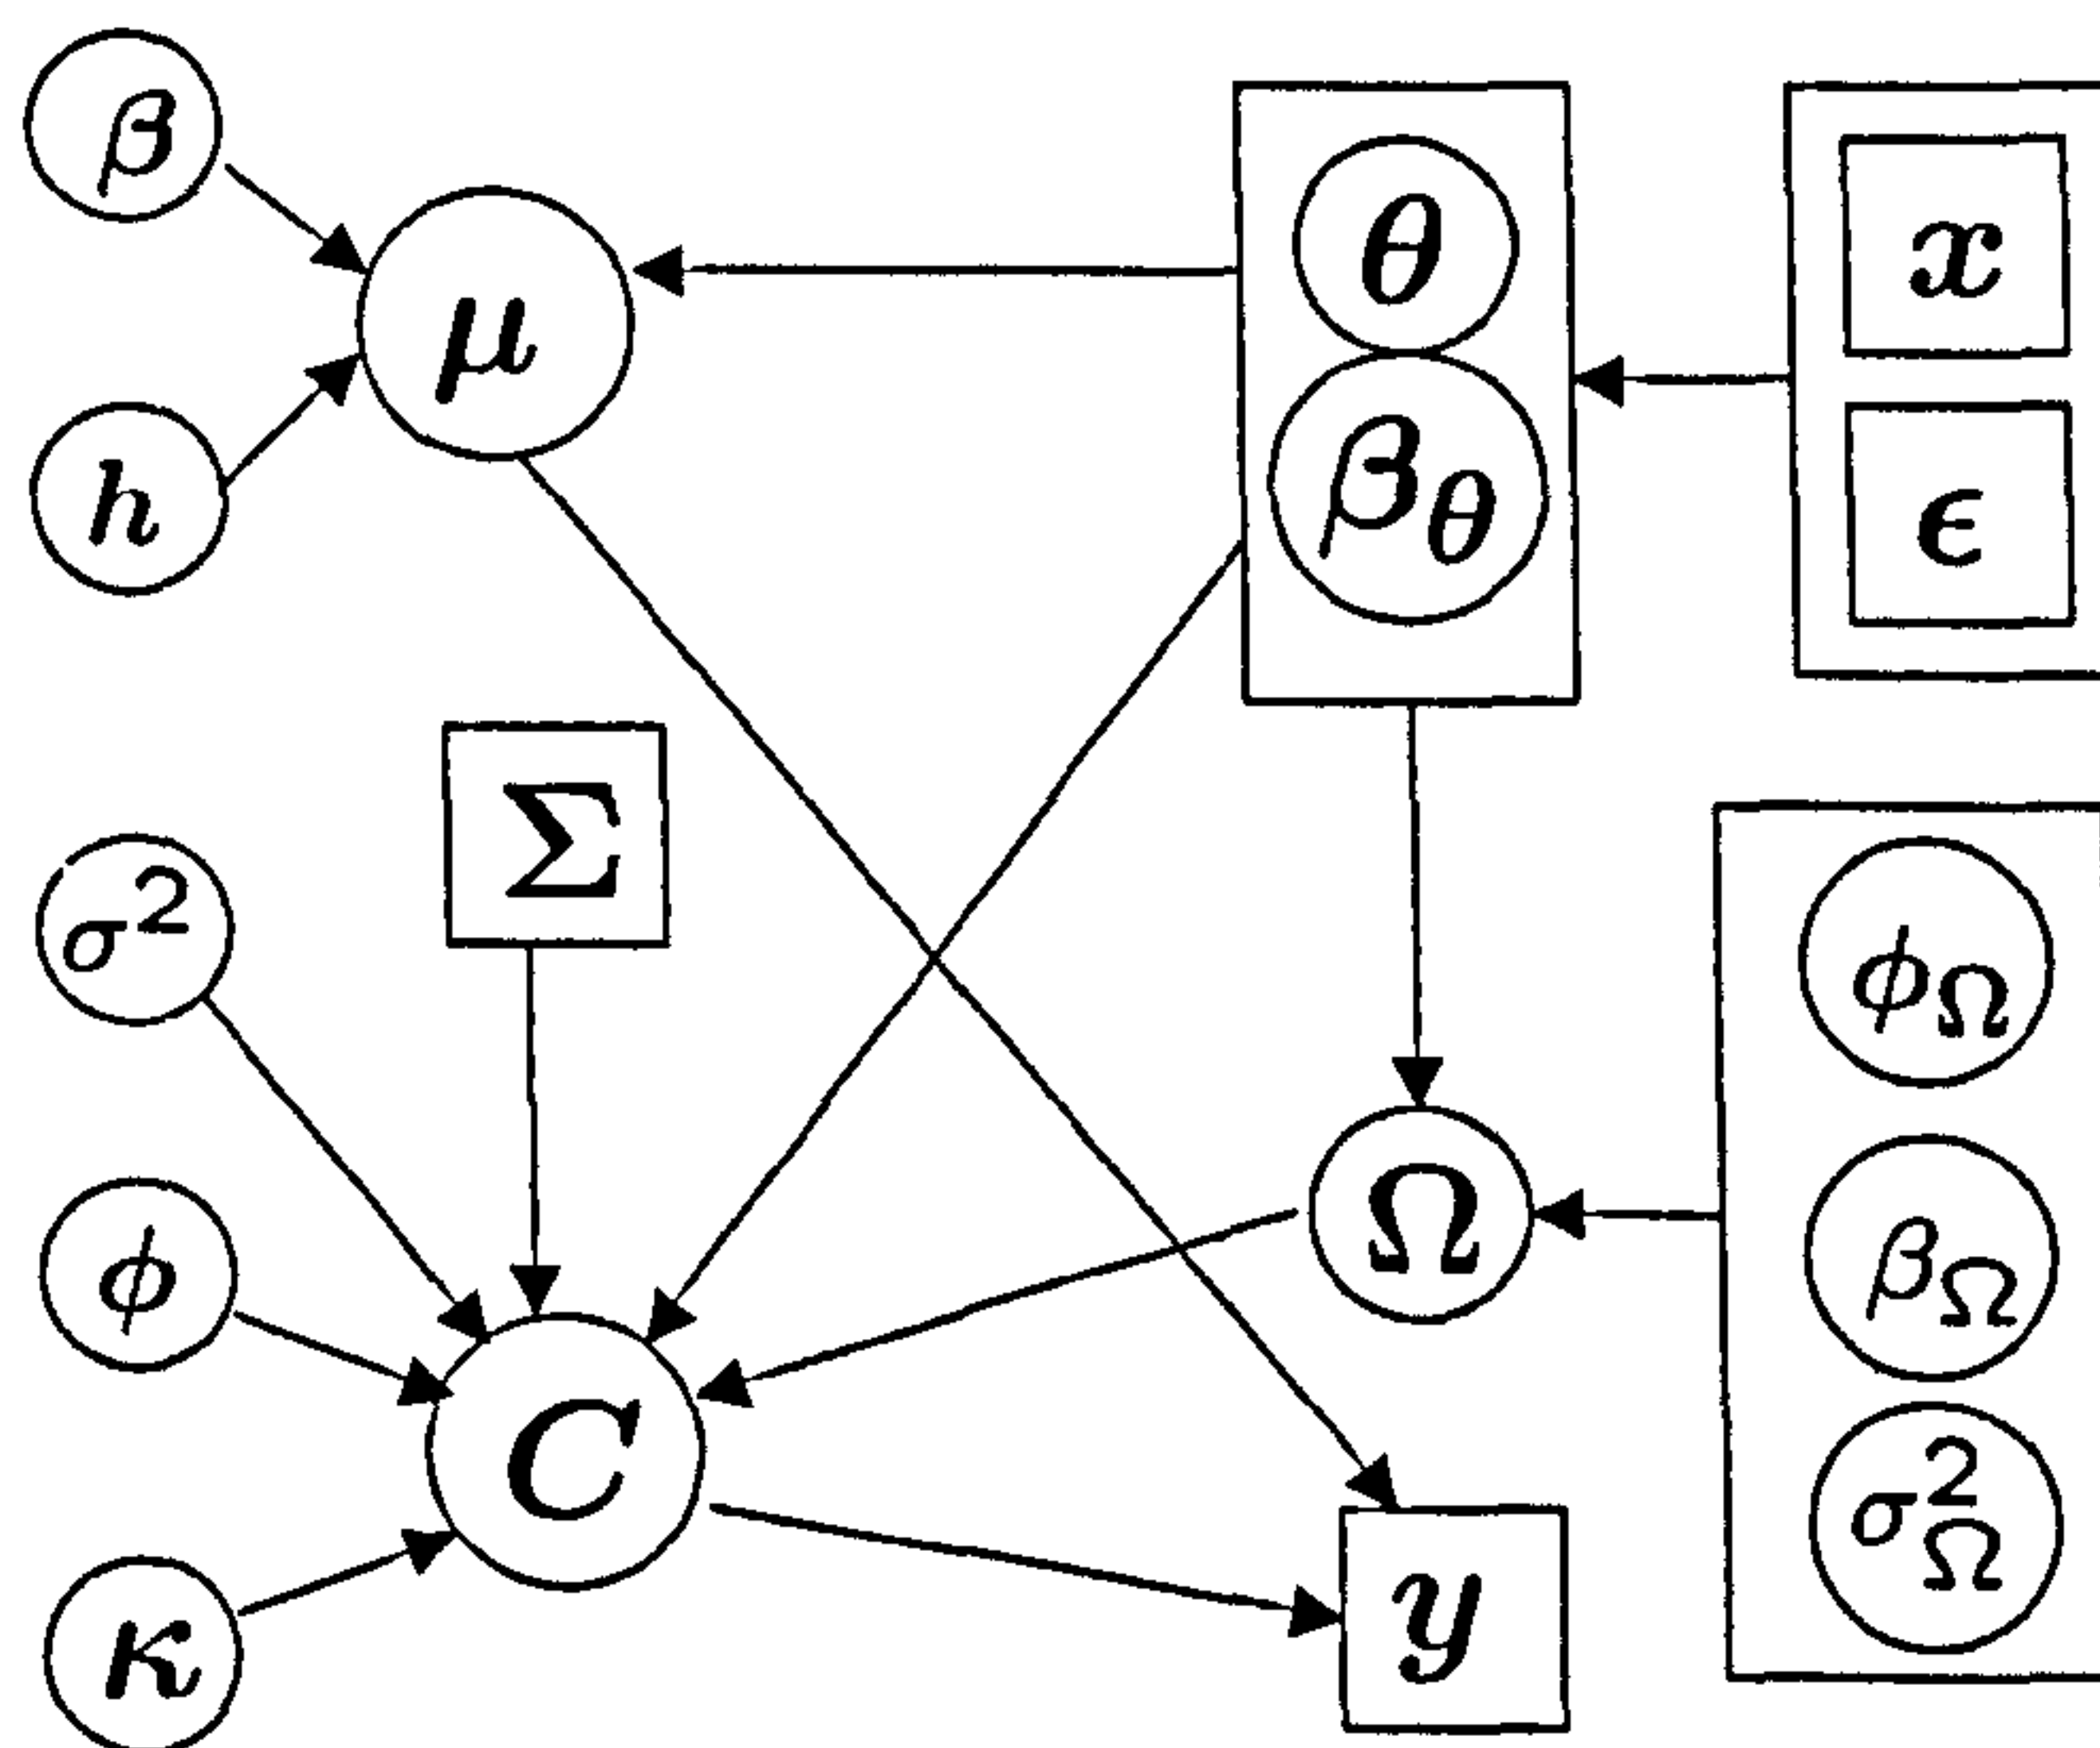


Figure 6.8: A Directed Acyclic Graph (DAG) for the non-stationary Matérn model. The hyper-parameters are ignored here for the purpose of clarity.

Superficially, the complete conditional distributions for the set of parameters in this model remain largely the same as that of stationary

model. The difference lies in the autocorrelation function where now:

$$c_{\phi, \Omega}^*(\theta_i, \theta_j) = |\Omega_i|^{\frac{1}{4}} |\Omega_j|^{\frac{1}{4}} |\Omega_i + \Omega_j|^{-\frac{1}{2}} \frac{1}{\Gamma(\phi) 2^{\phi-1}} \left(\sqrt{2\phi Q_{ij}} \right)^{\phi} \mathcal{K}_{\phi} \left(\sqrt{2\phi Q_{ij}} \right), \quad (6.29)$$

with

$$Q_{ij} = (\theta_i - \theta_j)^T ((\Omega_i + \Omega_j)/2)^{-1}. \quad (6.30)$$

Let $D = (y, z, x, \epsilon, \Sigma)$ and $\tau_{\Omega} = (\beta_{\Omega}, \sigma_{\Omega}^2, \phi_{\Omega})$. The extra parameters $\Omega, \beta_{\Omega}, \sigma_{\Omega}^2, \phi$ now have conditional distributions:

$$p(\Omega|D) \propto p(\Omega|\tau_{\Omega}) \times p(\tau_{\Omega}) \times p(y|\mu^*, C^*) \quad (6.31)$$

$$p(\beta_{\Omega}|D) \propto p(\beta_{\Omega}) \times p(\Omega|\tau_{\Omega}) \quad (6.32)$$

$$p(\sigma_{\Omega}^2|D) \propto p(\sigma_{\Omega}^2) \times p(\Omega|\tau_{\Omega}) \quad (6.33)$$

$$p(\phi_{\Omega}|D) \propto p(\phi_{\Omega}) \times p(\Omega|\tau_{\Omega}) \quad (6.34)$$

Note that the distribution $\Omega|\tau_{\Omega}$ is log-normal. As with the top-level Gaussian process, the posterior of the parameter β_{Ω} can be updated using Gibbs sampling.

This model was run for a larger number of iterations than those previously. A total of 5000 iterations were discarded as burn-in and the iterations were thinned by a factor of 8. A plot of the posterior Gaussian process for $\log(\Omega)$, representing the non-stationary covariance structure, is shown in figure 6.9. A plot of the resulting RSL estimates are shown in figure 6.10. It appears that the non-stationarity parameters represented by $\log(\Omega)$ indicate that some non-stationarity exists. Including the non-stationarity aspect appears to decrease the

uncertainty in RSL when compared with the stationary model for the majority of the time period involved. The mean also appears slightly smoother. This is shown in figure 6.11.




Figure 6.9: Plot of the log non-stationarity Gaussian process with 95% error bands.

6.2.5 Model comparison

The comparison of models, as with all hypothesis testing in a Bayesian setting, is yet to be made an exact science. This is largely due to two factors. Firstly, how to judge the importance of inferential differences; secondly, how to define the number of parameters in cases where models are hierarchically structured. In this section, two model comparison methods are outlined; Bayes factors and the deviance information criterion (DIC). Both methods have disadvantages; they are




Figure 6.10: A plot of the mean RSL curve with 95% point-wise intervals for the non-stationary Matérn model.

used here only as indicators of relative model fit. For more information on model comparison see, for example, chapter 9 of Gilks *et al.* (1996).

Bayes factors

A Bayes factor (BF) represents the relative weight in favour of one model over another. It is created from the equation:

$$\text{BF} = \frac{f(D|M_1)}{f(D|M_2)}. \quad (6.35)$$

Here, D denotes all the available data, whilst M_i represents the model under consideration. The density $f()$ is the marginal density of the observed data, sometimes known as the prior predictive density. It also




Figure 6.11: Comparison of the stationary and non-stationary temporal models. The solid lines show the mean estimates corresponding to the left axis, the dotted lines are the standard errors as shown on the right axis.

appears as the normalising constant in the standard Bayes equation and can be calculated from:

$$f(\mathbf{D}) = \int f(\mathbf{D}|\Theta) f(\Theta) d\Theta, \quad (6.36)$$

with Θ the set of parameters with prior density $f(\Theta)$.

The problems of calculating this density, and the subsequent Bayes factor, are well studied. They include:

1. The use of improper priors for $f(\Theta)$ causes problems in estimation of $f(\mathbf{D})$ which must necessarily be improper too.

2. The integral may be of very high dimension causing problems of computation.
3. If priors are sufficiently vague and models are nested, the BF may exhibit the 'Lindley paradox'. The density for the larger model will be necessarily small forcing the BF to reject the larger model in spite of what the data may point towards. In this case the BF is said to be overly parsimonious.

It is for these reasons that the Bayes factor is not used here. Many of the priors on parameters are vague, and those used for the dates are improper. The integral is also of very high dimension requiring as many steps as that of the MCMC itself.

DIC

The deviance information criterion (DIC), proposed by Spiegelhalter *et al.* (2002), aims to produce an easily obtainable measure of model fit whilst retaining reasonable parsimony in a Bayesian setting. Specifically, the DIC estimates an *effective* number of parameters, p_D , for a hierarchical model. This statistic is of use because similar measures (such as AIC, BIC, Akaike, 1973; Schwarz, 1978) require the number of parameters to be known precisely, which can be a matter for debate in Bayesian hierarchical models. The DIC is calculated as:

$$DIC = \overline{Dev(\Theta)} + p_D, \quad (6.37)$$

where $Dev(\Theta)$ is the deviance for parameter set Θ and $\overline{Dev(\Theta)}$ is a posterior average. In cases of non-symmetric parameters, this may be a median or mode rather than a mean. p_D is estimated from:

$$p_D = \overline{Dev(\Theta)} - Dev(\bar{\Theta}), \quad (6.38)$$

the posterior average of the deviance minus the deviance of the posterior average of the parameters. From this, DIC is created as a measure of fit penalised by the complexity of the model.

The advantage of this form is that it is easily calculated from the parameter output of a model. However, many disadvantages still exist. There is still no idea as to the size of differences in DIC which may indicate 'better' models. Another disadvantage is that of the focus of the model. This is defined as the lowest level of parameter used to calculate the deviance. In this thesis, as in WinBUGS, this constitutes the hyper-parameters at their lowest level. Varying the focus will alter the DIC. Numerous other concerns remain; see the discussion in Spiegelhalter *et al.* (2002).

The DIC was computed for the two models run on the same data: stationary covariance and non-stationary covariance. The GeoBUGS model did not include stochastic dates whilst the model with tendency included was dropped for reasons mentioned previously. The details of the DIC for the other two models are outlined in table 6.1. The table indicates that the non-stationary covariance models perform better, even after taking into account the extra parameters.

Model type	$\overline{Dev(\Theta)}$	$Dev(\overline{\Theta})$	p_D	DIC
Stationary covariance	376.89	326.23	50.66	427.55
Non-stationary covariance	296.27	214.45	81.82	378.10

Table 6.1: Table of DIC values for the temporal models.

6.3 Spatio-temporal models

As mentioned previously, there is a spatial aspect to the data. This section attempts to explore the option of additional spatial variation in RSL across the Humber. Now, location in terms of Ordnance Survey (x, y) coordinates in kilometres are used as extra information about the overall Humber data set.

As with the temporal models, prediction of RSL is required. This involves the input of suitable time-points over which to estimate. With spatio-temporal models, it also requires locations at which RSL curves are desired. This section therefore contains predictions at three arbitrary locations in the Humber. They correspond to the modern tidal stations at Spurn Head, Kingston-upon-Hull and Goole.

6.3.1 GeoBUGS

In section 4.6.2, it was suggested that incorporating spatial terms into the mean trend reduced spatial and temporal autocorrelation. This

section extends that idea further. Here, x again refers to the median calibrated dates. Using GeoBUGS, a first model was fitted with:

$$y \sim N(\mu + h, \sigma^2 A + \Sigma + \Psi), \quad (6.39)$$

where

$$A = [a_{ij}] = \exp(-\phi_1(x_i - x_j)^{\phi_2}), \mu_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 \lambda_i + \beta_4 \omega_i,$$

where λ is the OS x -coordinate location, and ω is the y -coordinate. Again, the parameters ϕ_1 and ϕ_2 control the nature of decay of temporal autocorrelation. Thus this model examines whether any spatial variation can be uncovered in its entirety by using simple linear functions of the x - and y -coordinates.

Vague priors were given to all the parameters (with the same exceptions for h_{IIP} and h_{LP} as in section 6.2.1). The model was run for a total of 5000 iterations (removing 2000 for a burning-in period) using two different sets of starting values. The Brooks-Gelman statistic was used to check for convergence (covered in section 5.10.3). Posterior densities for β_3 and β_4 are shown in figure 6.12. β_3 is clearly positive, indicating a general slight increase in RSL towards the outer estuary. β_4 is clearly negative, indicating a lower RSL towards the north of the Humber. However, though β_4 appears to be bigger than β_3 , it must be remembered that the vast majority of the variability in the data lies in the East-West direction when compared to the North-South direction (the variance is almost three times as big). Thus it may be that the β_3 parameter has more of an effect on RSL.




Figure 6.12: Densities for β_3 and β_4 in the spatial GeoBUGS model.

It is also interesting to note the effect the introduction of these two parameters has on the covariance parameters, σ , ϕ_1 and ϕ_2 . A plot of the densities of these is shown in figure 6.13 for comparison with figure 6.1. The introduction of the spatial component appears to have increased the overall variation, σ , whilst the scale parameter ϕ_1 and smoothness parameter ϕ_2 are now given more precisely.

Finally, estimates of the RSL curve with error are given in figure 6.14. It is clear that the curves are generally smooth with only minor fluctuations easily accounted for by uncertainty. As expected, this uncertainty increases for the period of 8k years BP and before.

As with the temporal-only models, the models can only be extended by dropping the use of GeoBUGS because of its limitations. In this case, an additional disadvantage is that GeoBUGS cannot handle models where






Figure 6.13: Posterior densities for σ, ϕ_1 and ϕ_2 for the GeoBUGS model with spatial mean.




Figure 6.14: Point-wise means for the Humber from the GeoBUGS model with spatial mean. The dotted black line shows the standard error as given by the right axis.

both mean and covariance are spatio-temporal. This is considered in further sections using R.

6.3.2 Spatio-Temporal separable stationary covariance

The model presented in this section is a small extension over the stationary Matérn model presented in section 6.2.2. Again, the mean is allowed to vary with the Ordnance Survey x - and y -coordinates but is now extended to allow spatio-temporal covariance too. It can be set out

as:

$$y \sim N(\eta(x, \Lambda) + h, \Sigma + \Psi) \quad (6.40)$$

$$\eta(x, \Lambda) \sim GP(F\beta, \sigma^2 A) \quad (6.41)$$

with

$$F = \begin{pmatrix} 1 & \theta_1 & \theta_1^2 & \lambda_1 & \omega_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \theta_m & \theta_m^2 & \lambda_m & \omega_m \\ 1 & \beta_{\theta, m+1} & \beta_{\theta, m+1}^2 & \lambda_{m+1} & \omega_{m+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \beta_{\theta, n} & \beta_{\theta, n}^2 & \lambda_{m+n} & \omega_{m+n} \end{pmatrix}.$$

The covariance structure used here is that of a separable Matérn form where now:

$$A = [a_{ij}] = \rho_{B, \phi_{2L}}(\Lambda_i - \Lambda_j) \rho_{\phi_{1t}, \phi_{2t}}(\theta_i - \theta_j). \quad (6.42)$$

The topic of anisotropy was discussed in section 5.7. Here, this is tested for the Humber data. This involves two possible forms for B from equation 5.50; $B = \begin{pmatrix} B_{11} & 0 \\ 0 & B_{11} \end{pmatrix}$ and $B = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix}$, corresponding to isotropy and anisotropy respectively. For the isotropic model, an inverse gamma prior was given to B_{11} , whilst for the anisotropic model the matrix B was given a Wishart prior, as discussed in section 5.7.

The isotropic model was run for a total of 10000 iterations. Convergence was obtained after discarding 2000 as a burn-in and thinning by 7. A plot of the location covariance parameters, B_{11} and

ϕ_{2L} , is shown in figure 6.15. Note that the location scale parameter B_{11} is well-defined, and that there is a strong degree of smoothness, as ϕ_{2L} is large. This is as expected as there is unlikely to be large-scale changes over the Humber.




Figure 6.15: Densities for B_{11} (left panel) and ϕ_{2L} (right panel) in the isotropic model.

A plot of the mean lines from the Gaussian process output across the Humber is shown in figure 6.16. It is interesting to note that RSL in the inner Humber appears to have been lower up until very recently. However, the uncertainty is also large over the last 2k years BP. Before 8k years BP, the uncertainty is so great as to make almost any inference impossible.

The anisotropic model required a burn-in of 3000 iterations and thinning by 6. The nature of the anisotropy can be examined through the use of the summary statistics outlined in section 5.7; these are the orientation of the ellipse, χ , and the ratio of anisotropy, λ . The




Figure 6.16: Point-wise means for the Humber from the isotropic spatio-temporal model. The dotted lines shows the standard error as given by the right axis.

ratio of anisotropy has 95% HPD interval of (1.00,8.35); not a strong relationship. The orientation of the covariance ellipse is shown for the median estimate of B in the rose diagram of figure 6.17, together with a density plot. The orientation seems to peak at around 85° though there is some uncertainty around this figure. This indicates there may be an increased range of covariance along this orientation. It is interesting to note that the direction is approximately east-west across the Humber. This is discussed further in section 7.7.

The RSL estimates for the anisotropic model are shown in figure 6.18. The lines follow much the same pattern as that of the isotropic version,




Figure 6.17: Density plot for χ (left panel). Rose diagram for χ taken from median estimate for B (right panel).

though now the uncertainties for the different locations appear to vary more.

6.3.3 Non-stationary spatio-temporal covariance

Finally, the model is further extended to incorporate features of temporal non-stationarity as discussed in section 6.2.4. Now:

$$A = [a_{ij}] = \rho_{\phi_1}^{(1)} (\Lambda_i - \Lambda_j) \rho_{\Omega, \phi_2}^{(2)} (\theta_i - \theta_j), \quad (6.43)$$

with $\rho^{(2)}$ as outlined in equation 6.29.

The model required a burn-in period of around 5,000 iterations. Multiple chains were run to check for convergence. As with previous models, the non-stationarity aspect of the model can be examined; a




Figure 6.18: RSL for the Humber from the separable spatio-temporal GP model with anisotropic stationary Matérn covariance. The solid lines follow the left axis and give the mean estimate of RSL, the dotted lines follow the right axis and give the standard errors.

graph of the Gaussian process for $\log(\Omega)$ is shown in figure 6.19. This is more well-defined than that of the temporal-only model. It is clear that the smoothness is most clearly identified during the period 3k to 6k years BP. A plot of the direction of anisotropy, χ , is given in figure 6.20, together with a rose diagram for the median estimate of B . χ is clearly bi-modal, with peaks at around 80° and 150° . With such a distribution, the rose diagram is clearly less useful, but shows again how the main direction of anisotropy lies in the east-west direction. The north-south variation of the anisotropic stationary model is still shown but to a lesser extent. The ratio of anisotropy had a 95% HPD interval of (1.00, 19.05). Finally, a plot of the mean lines with standard

errors for Spurn Head, Hull and Goole are shown in figure 6.21. The uncertainties remain reasonably similar across the estuary, but the mean lines again show some discrepancy.




Figure 6.19: Plot of the mean of the Gaussian process for $\log(\Omega)$ with 95% intervals (dotted).

6.3.4 Model comparison

As before, the DIC is used to compare models. In this case, the models involving spatio-temporal correlation are scrutinised. The models compared here are the model with separable spatio-temporal stationary covariance (isotropic and anisotropic) and the model with non-stationary separable spatio-temporal covariance. A table of the DIC values is shown in table 6.2. It is clear that there is very little to choose




Figure 6.20: Density plot for χ (left panel). Rose diagram for χ taken from median estimate for B (right panel).




Figure 6.21: RSL for the Humber from the separable spatio-temporal GP model with anisotropic non-stationary Matérn covariance. The solid lines follow the left axis and gives the mean estimate of RSL, the dotted lines follow the right axis and give the standard errors.

between the isotropic and anisotropic models, even after taking into account the extra complexity. However, the non-stationary anisotropic model exhibits the lowest DIC and is thus carried forwards as the finally chosen model.

Model	$\overline{Dev(\Theta)}$	$Dev(\bar{\Theta})$	p_D	DIC
Isotropic stationary covariance	349.63	330.84	18.79	368.42
Anisotropic stationary covariance	347.99	327.53	20.46	368.45
Anisotropic non-stationary covariance	326.57	309.12	17.45	344.02

Table 6.2: Table of DIC values for the spatio-temporal models.

6.4 Summary

This chapter has examined a number of models which attempt to explain the shape and variation in RSL data. The models have fallen into two broad categories; those that model the Humber as a homogeneous unit and those that use the location as an extra explanatory variable in the model. Of the temporal-only models, the non-stationary Matérn model outperformed the rest in term of the DIC. However, the spatio-temporal models are favoured as they use the data more naturally. In particular the non-stationary spatio-temporal Gaussian process model gave the lowest DIC score for this group; this model is preferred over all others and is carried over to the next chapter for further analysis.

Chapter 7

Application of results

7.1 Introduction

This chapter examines in detail the nature of the lines produced by the models proposed in chapter 6. Previously, the output from the Gaussian process models was analysed by simply looking at point-wise means and variances. Here, the treatment is extended to deal with some of the more subtle features that may not be immediately apparent. Initially, this is done through kernel density estimation and functional principal components analysis (FPCA). The chapter then considers rates of RSL change; two methods are presented for its calculation. The rates are also compared to published estimates for the Humber. The chapter contains brief digressions on the use of the archaeological data, the compression of intercalated index points and estimation of isostatic activity in the Humber. Finally, the RSL model is used to look for evidence of wider changes in climate.

The model used in this chapter is that of the non-stationary spatio-temporal Gaussian process with a Matérn autocorrelation function. Unless otherwise specified, this model has been used to produce all the data in this chapter's graphs and tables.

7.2 Examining RSL uncertainty

7.2.1 Kernel density estimation

Kernel methods have already been discussed in the context of smoothing in section 4.5.1. Here, they are used for density estimation in two ways. Firstly in a point-wise context, with a desire to look at the variation in elevation across a set of chosen dates. Secondly, the density is examined by treating the discretised Gaussian process output as an independent set of point-referenced data in age-elevation space.

The standard kernel density estimator (Silverman, 1996) at point t , denoted $\hat{f}(t)$, is given by:

$$\hat{f}(t) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{t - x_i}{h}\right), \quad (7.1)$$

where K is known as the kernel function defining weights usually centred on t , with observed data x_1, \dots, x_n . h is known as the bandwidth and controls the amount of smoothing. Results presented in this section (and section 4.5.1) are obtained using the Gaussian kernel with

optimal bandwidth selectors. The approach used here for choosing h is from Silverman (1996):

$$h_{OPT} = 0.9An^{-1/5} \quad (7.2)$$

where A is the minimum of the standard deviation and an adjusted interquartile range. The idea behind this estimate is to minimise the mean integrated squared error of the density estimate. This includes roughness penalties (of the form $\int f''(x)^2 dx$) for the design density. Unfortunately, the calculation of the roughness penalty for the design density requires knowledge of the nature of the density in the first place. The value settled on by Silverman (1996) is a conservative estimate designed to work for a wide range of densities.

Initially, point-wise kernel estimates were calculated for the RSL model. The density on the elevation axis was calculated at dates of 0, 2k, 4k, 6k, 8k, and 10k years BP. These are shown in figure 7.1. There appears to be little evidence of multi-modality. As expected, uncertainty increases with age. The least uncertainty is attained at around 4k years BP. This time-point corresponds to the period where there are the most basal index points. Goole and Spurn Head appear to differ most in the period 6k to 0k years BP, where the uncertainty at Spurn Head, by contrast, is reasonably constant.

The kernel method can be extended to work in multiple dimensions.

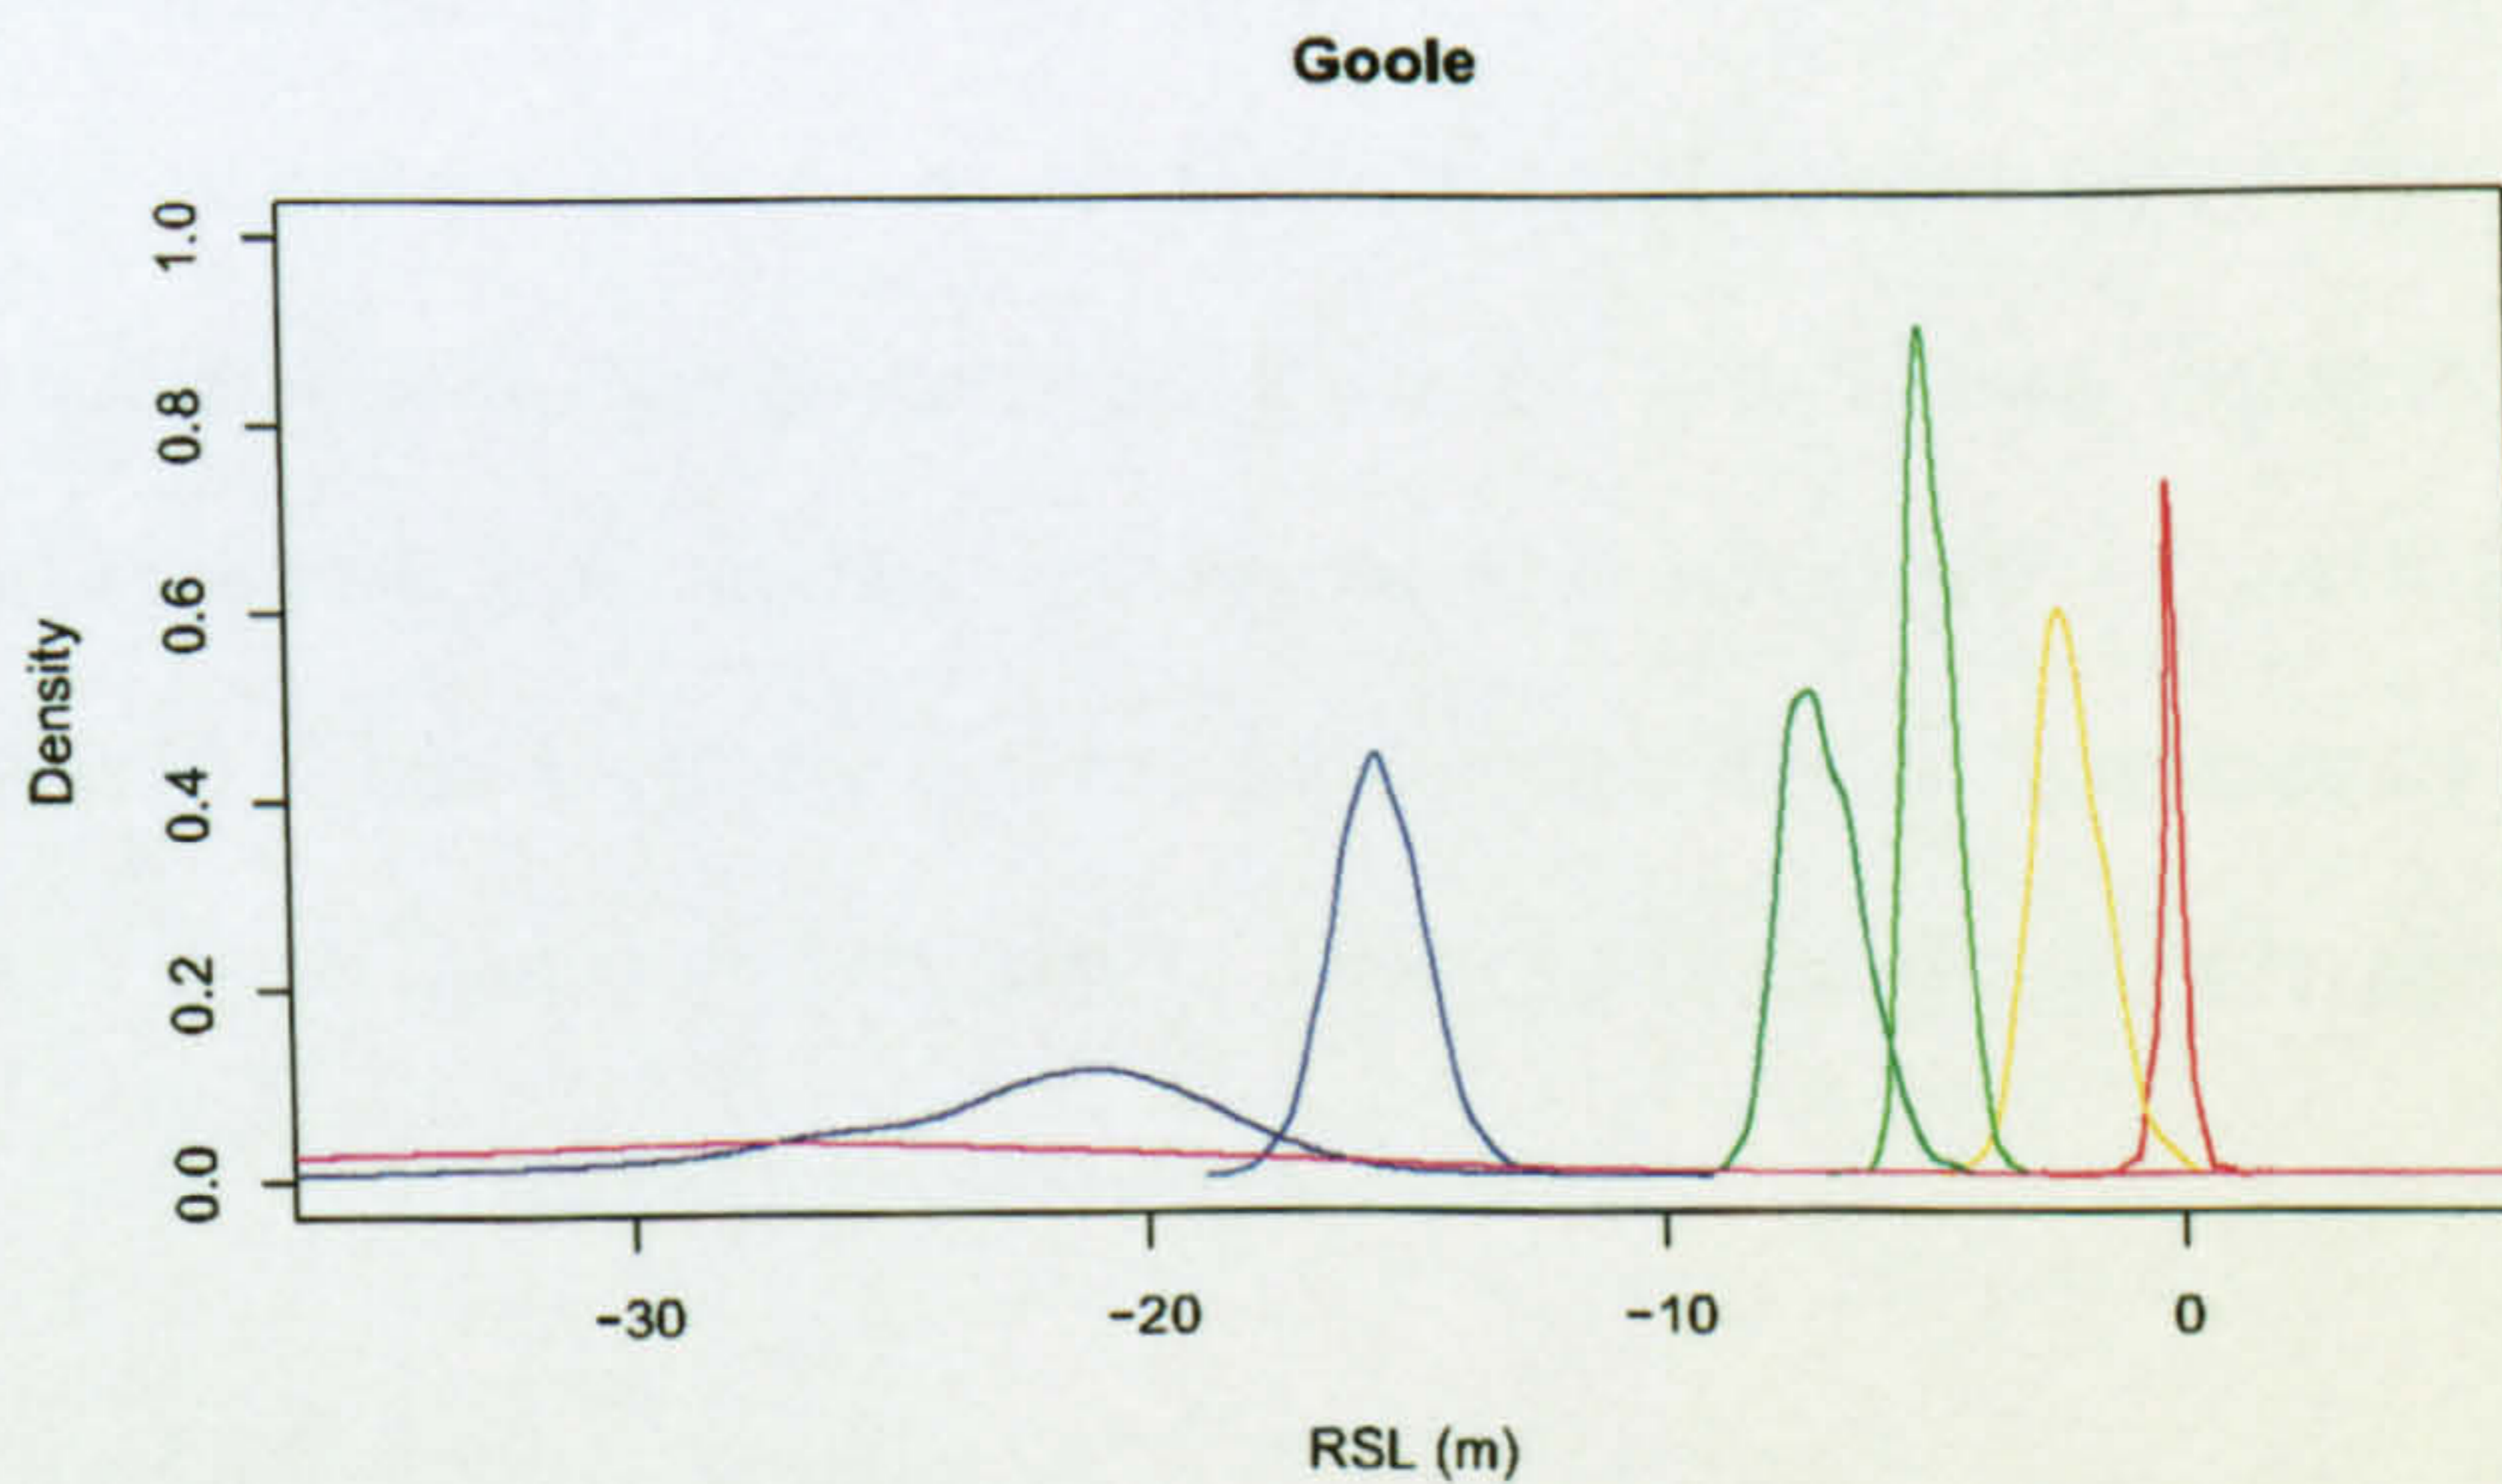




Figure 7.1: Point-wise elevation densities for locations around the Humber for 2k year steps for the non-stationary spatio-temporal model. Furthest right density is 0 years cal BP, furthest left is 12k years cal BP.

7.2.2 Functional principal components analysis

Functional multivariate analysis is a useful qualitative tool for analysing functional data. Functional principal components analysis (FPCA) is analogous to standard PCA but used on functional data. Here, the functional data are represented by successive lines (estimates of $\eta(t)$) created in the previous chapter through the Gaussian process framework. These are 'discretised' as they are only evaluated for a discrete set of inputs. In this case, the inputs are pre-chosen dates in thousands of years cal BP. The methods presented here follow the work of Ramsay and Silverman (1997).

Firstly, consider the original aims of PCA; find a linear combination of variables and data, such that variation in the linear combination is maximised in a number of orthogonal directions. The discretised functions can be written using inner product notation as:

$$f_i = \langle \xi, x_i \rangle, \quad i = 1, \dots, N \quad (7.3)$$

with ξ a p -vector of unknown weighting coefficients applied to the observed data x_i . Here, p is the number of dimensions and N the number of observations. The method starts by finding the weight vector for which $f_{i1} = \langle \xi_1, x_i \rangle$ has the largest possible mean square subject to the constraint $\|\xi_1\| = 1$. This continues for subsequent ξ_j , $j > 1$, with the additional constraints for the m th step:

$$\|\xi_k, \xi_m\| = 0, \quad \text{for } k < m. \quad (7.4)$$

The weights ξ are now a set of principal components (PCs) which display the modes of variation in f . The sets of f_i , $i = 1, \dots, N$ are known as the principal components scores; they display how each case tallies on each of the principal components. Finally, the mean square for each principal component is scaled to show the amount of variation each has 'explained'. This usually becomes small after a number of principal components have been calculated. Principal components are usually discarded once the mean square is small.

Consider the problem now in the functional domain. The data, previously x_{ij} with p dimensions and N data points now become N

functions evaluated over the continuous measure t : $x_i(t)$. The linear combination required now is:

$$f_i = \int_S \xi(s)x_i(s)ds. \quad (7.5)$$

The now continuous weights $\xi(s)$ and $x_i(s)$ are discretised as they are only observed at suitable values $t \in S$. Thus the values t become the dimensions used in standard PCA. All the same restrictions on the ξ still apply.

The problem remains that the discretised values of ξ require estimation. This is done through eigenanalysis of the sample covariance matrix $(N - 1)^{-1}X^T X$ where X is now the functional data vector with zero mean. The justification for using eigenanalysis for computing principal components is given in Mardia *et al.* (1979). Other PCA techniques such as rotation of components to enable better interpretation are also possible. In particular, the Varimax rotation (Kaiser, 1958) finds an orthogonal transformation which causes the principal components to have a small number of large values and a large number of small values.

As an example of the functional data set, five simulated lines representing RSL from Spurn Head are plotted in figure 7.2. The first three varimax rotated PCs were extracted from a larger set of lines (over 1000). These are shown in figure 7.3 for three locations in the Humber. Functional data for the period 8k-12k years BP were discarded as they dominated the principal components. For Spurn Head and Hull, the first principal component evidently represents the uncertainty present

in the period between 2k and 0k years BP. The second represents that of between 6k and 8k years BP. The third looks at the uncertainty between 2k and 6k years BP. The order of the three components is reversed for the inner-most RSL curve at Goole. These three principal components account for 60.8%, 60.4% and 71.4% of the variation at Spurn Head, Hull and Goole respectively.

One possible interpretation of the functional principal components is that they represent, in order of decreasing importance, where RSL uncertainty is greatest. The first principal component would therefore suggest that, for future data collection purposes, it would be wise to target more recent finds at the outer Humber, whilst the best method for reducing uncertainty in the inner Humber is to collect more data from the period 2k to 6k years BP.




Figure 7.2: Five sample lines for Spurn Head. Produced using the non-stationary spatio-temporal Matérn model outlined in section 6.3.3.

7.3 Rates of RSL change

7.3.1 Differentiated splines

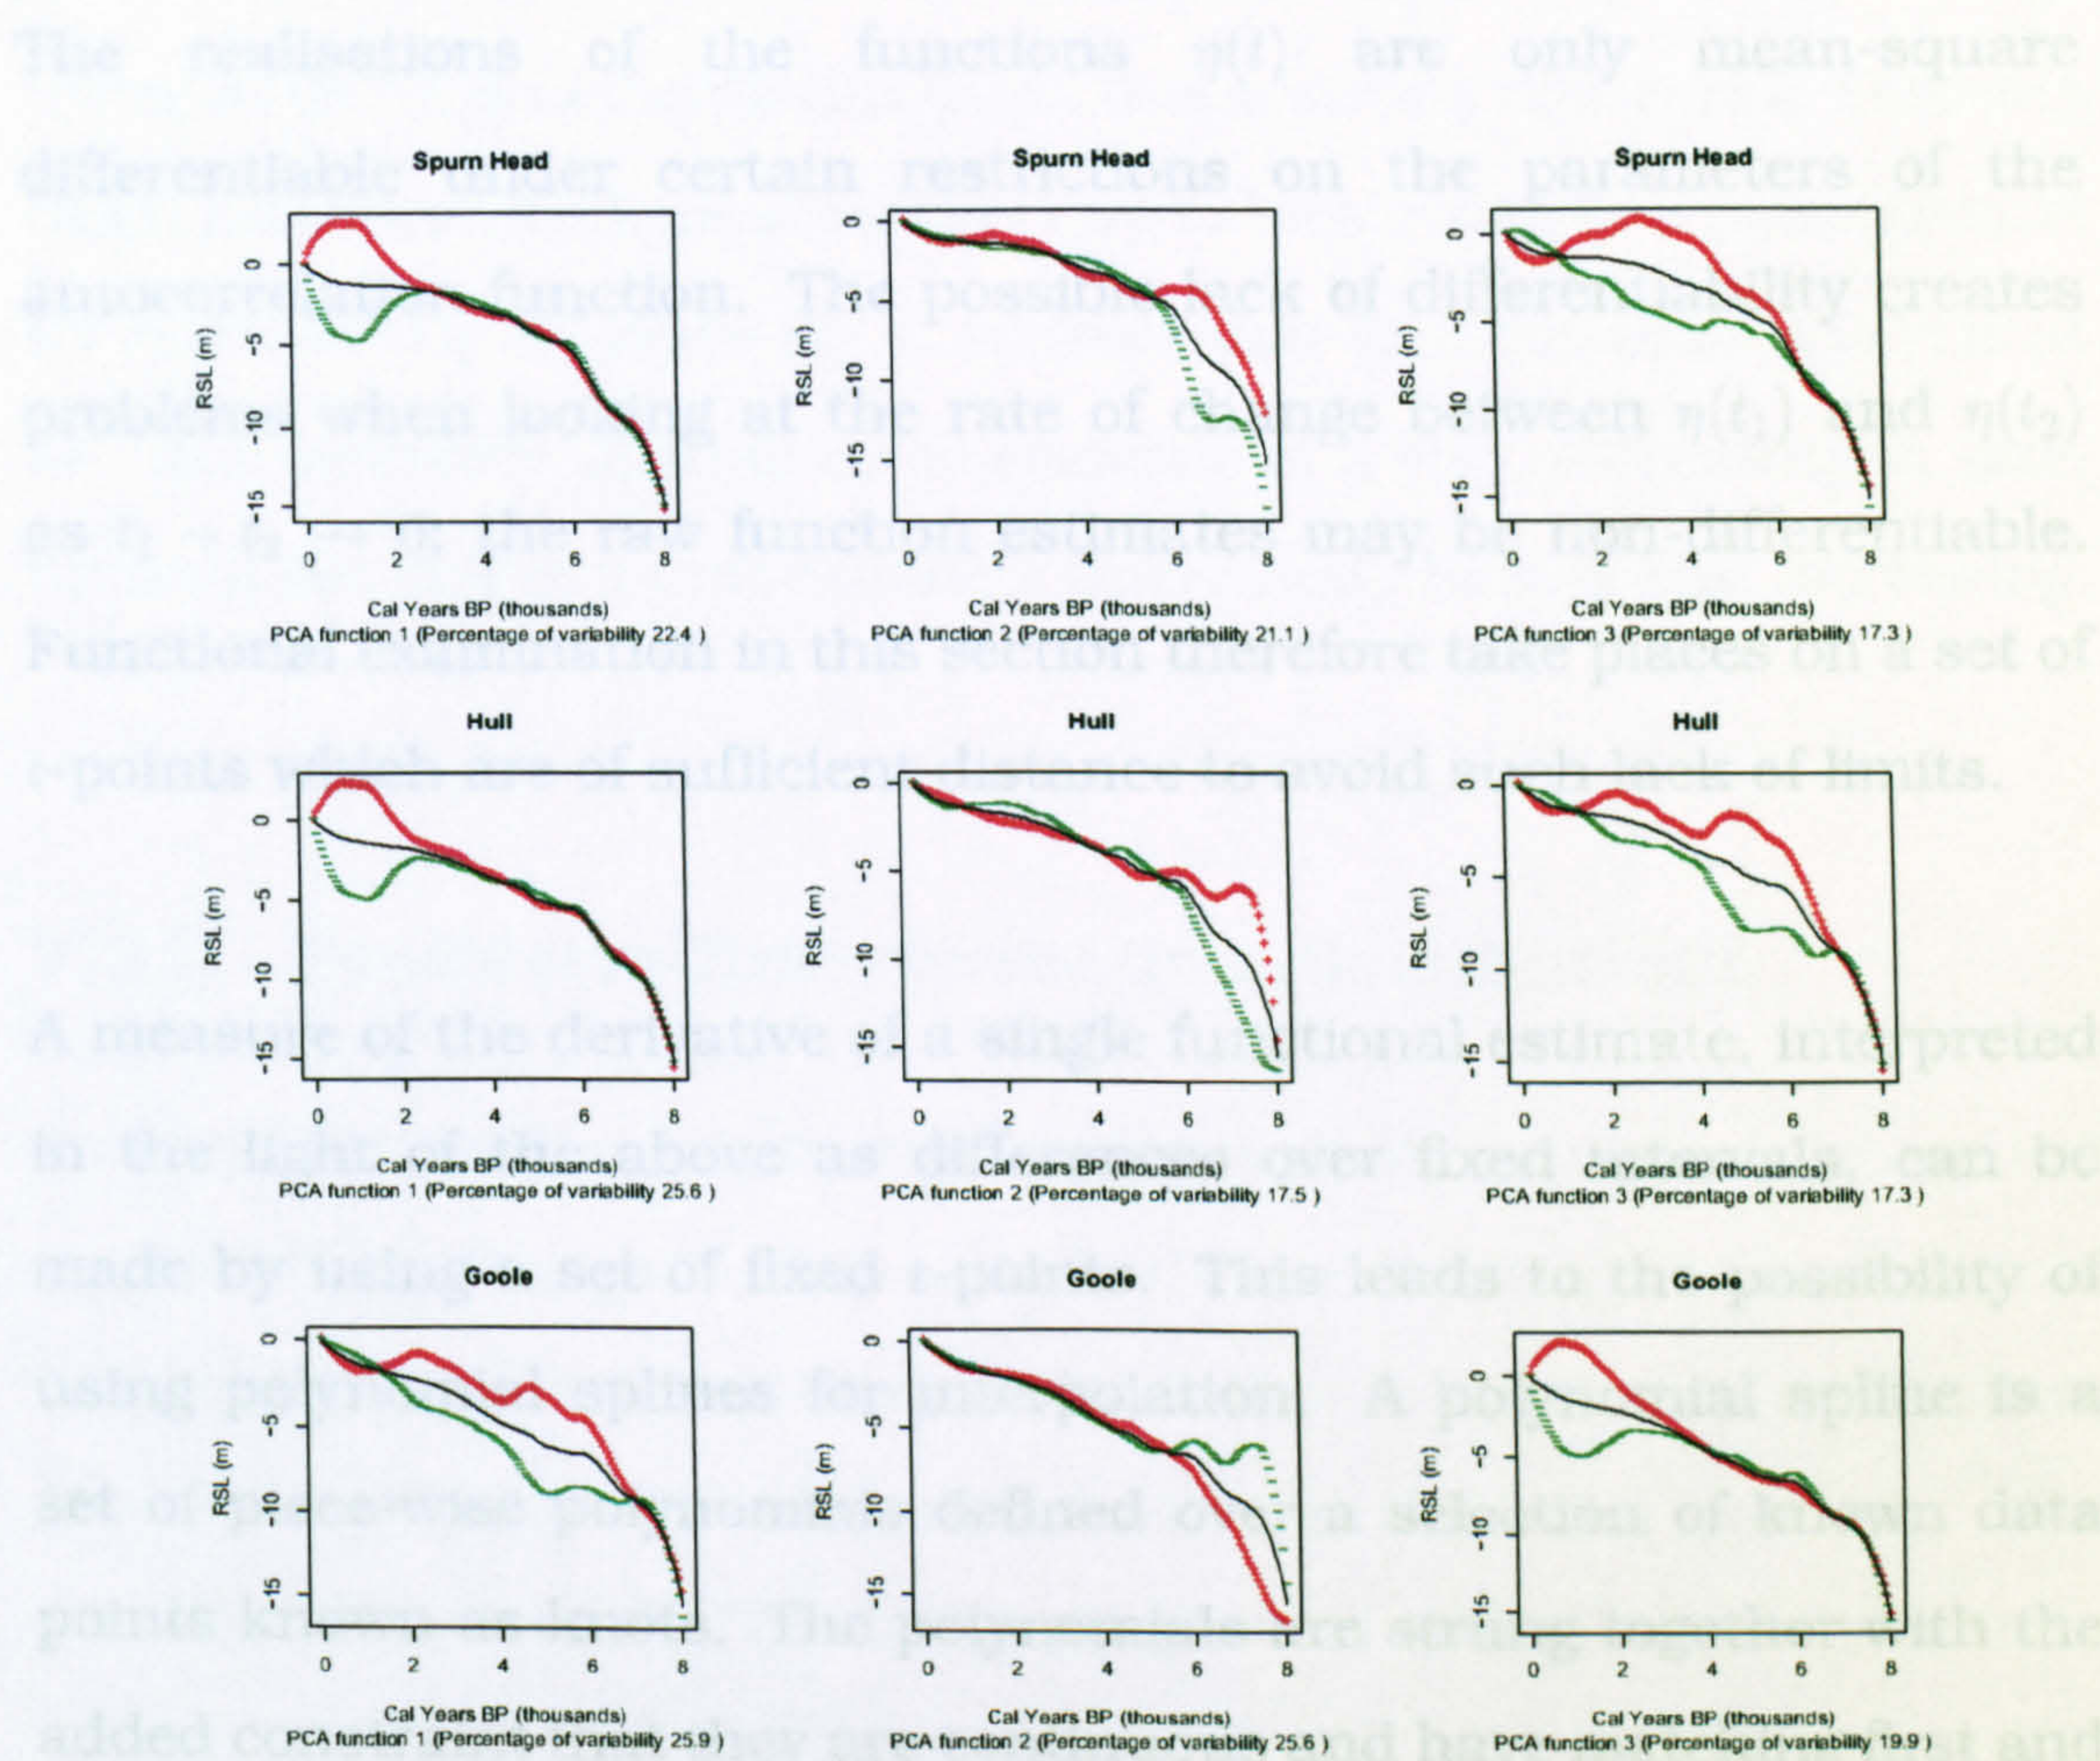


Figure 7.4: First three functional principal components from the non-stationary spatio-temporal Matérn model for Spurn Head (top row), Hull (second row) and Goole (third row). Thick line is the mean, the "+" line and "-" line are the mean plus and minus a multiple of that principal component respectively.

7.3 Rates of RSL change

7.3.1 Differentiated splines

The realisations of the functions $\eta(t)$ are only mean-square differentiable under certain restrictions on the parameters of the autocorrelation function. The possible lack of differentiability creates problems when looking at the rate of change between $\eta(t_1)$ and $\eta(t_2)$ as $t_1 - t_2 \rightarrow 0$; the raw function estimates may be non-differentiable. Functional examination in this section therefore take places on a set of t -points which are of sufficient distance to avoid such lack of limits.

A measure of the derivative of a single functional estimate, interpreted in the light of the above as differences over fixed intervals, can be made by using a set of fixed t -points. This leads to the possibility of using polynomial splines for interpolation. A polynomial spline is a set of piece-wise polynomials defined over a selection of known data points known as knots. The polynomials are strung together with the added constraint that they are continuous and have matching first and second derivatives at the knot values. It is therefore usual to make the polynomials of order 3. The smooth spline curve is calculated as a regression on an appropriate set of basis functions. The basis function used here is a B-spline basis; a recursive function built up from step functions over the range of the knots. The details can be found in de Boor (1978).

The spline interpolants have well-defined first derivatives. A plot of

the first derivative of RSL (RSL speed) is shown in figure 7.5 with 95% pointwise HPD intervals for Spurn Head, Hull and Goole. These plots show interesting features of the data. Though there appears to be evidence of a strong rise before 8k years BP, this period has large uncertainty. It is not until this point that a strong rise in RSL becomes clear. All three locations show this rise to be slowing to a point around 6k years BP. By 2k years BP, it appears that RSL has ceased to rise at all. Subsequently there is some evidence of a rise, increasing in rate, up until the present; unfortunately uncertainty here is also increasing.

7.3.2 Predictive first-derivative distributions

As shown in section 5.6, derivative observations can inform about the nature of RSL change. In this section, the conditioning of equation 5.39 is reversed so that the posterior distribution of the set of parameters generate estimates of the derivative of the Gaussian process. The posterior distribution of the Gaussian process derivative can be written out as:

$$\eta'(t)|\mathbf{y} \sim N(m^{**}(t), c^{**}(t, t)), \quad (7.7)$$

$$m^{**}(t) = \frac{\partial}{\partial t} f(t)^T + C_2(\theta, t)(\sigma^2 A + \Sigma + \Psi)^{-1}(\mathbf{y} - F\beta - \mathbf{h}), \quad (7.8)$$

$$c^{**}(t, t) = C_{12}(t) - C_2(\theta, t)(\sigma^2 A + \Sigma + \Psi)^{-1}C_2(\theta, t)^T. \quad (7.9)$$

All elements are as defined in section 5.6.

The predictive density of $\eta'()$ was obtained with the constraint that only




Figure 7.5: Plots of RSL speed for various locations around the Humber with 95% error bands using the spline method.

MCMC iterations with $\phi_{2t} \geq 1$ are used to draw samples of $\eta'()$. The constraint has the effect that only differentiable covariance matrices are used. Plots of $\eta'()$ for the usual three locations around the Humber are shown in figure 7.6. It is clear that the uncertainties produced via this method are larger than those of the spline method. The rise at around 8k years BP is more pronounced, but otherwise the main features of the rate of change are preserved.

7.4 Comparison with geophysical models

Geophysical models were introduced in section 2.3.9. The task of this section is to discern whether the geophysical model of Peltier *et al.* (2002) differs from the Gaussian process output. The geophysical model estimates RSL at chosen locations at 1k year intervals. A version of this model has been applied to the Humber at two different locations corresponding to the inner and outer Estuary. A plot of the two resultant RSL curves is shown in figure 7.7. The geophysical model produces no uncertainty estimate.

The inner and outer Humber are defined as locations east and west of around 0.4° longitude, corresponding approximately to the location of a buried chalk cliff line (Gaunt and Tooley, 1974) and also the location of the Humber bridge. The geophysical model is therefore tested against two locations corresponding to the average between the Humber bridge and Goole, and the average between the Humber bridge and Spurn Head. These two locations are used to represent the inner and outer

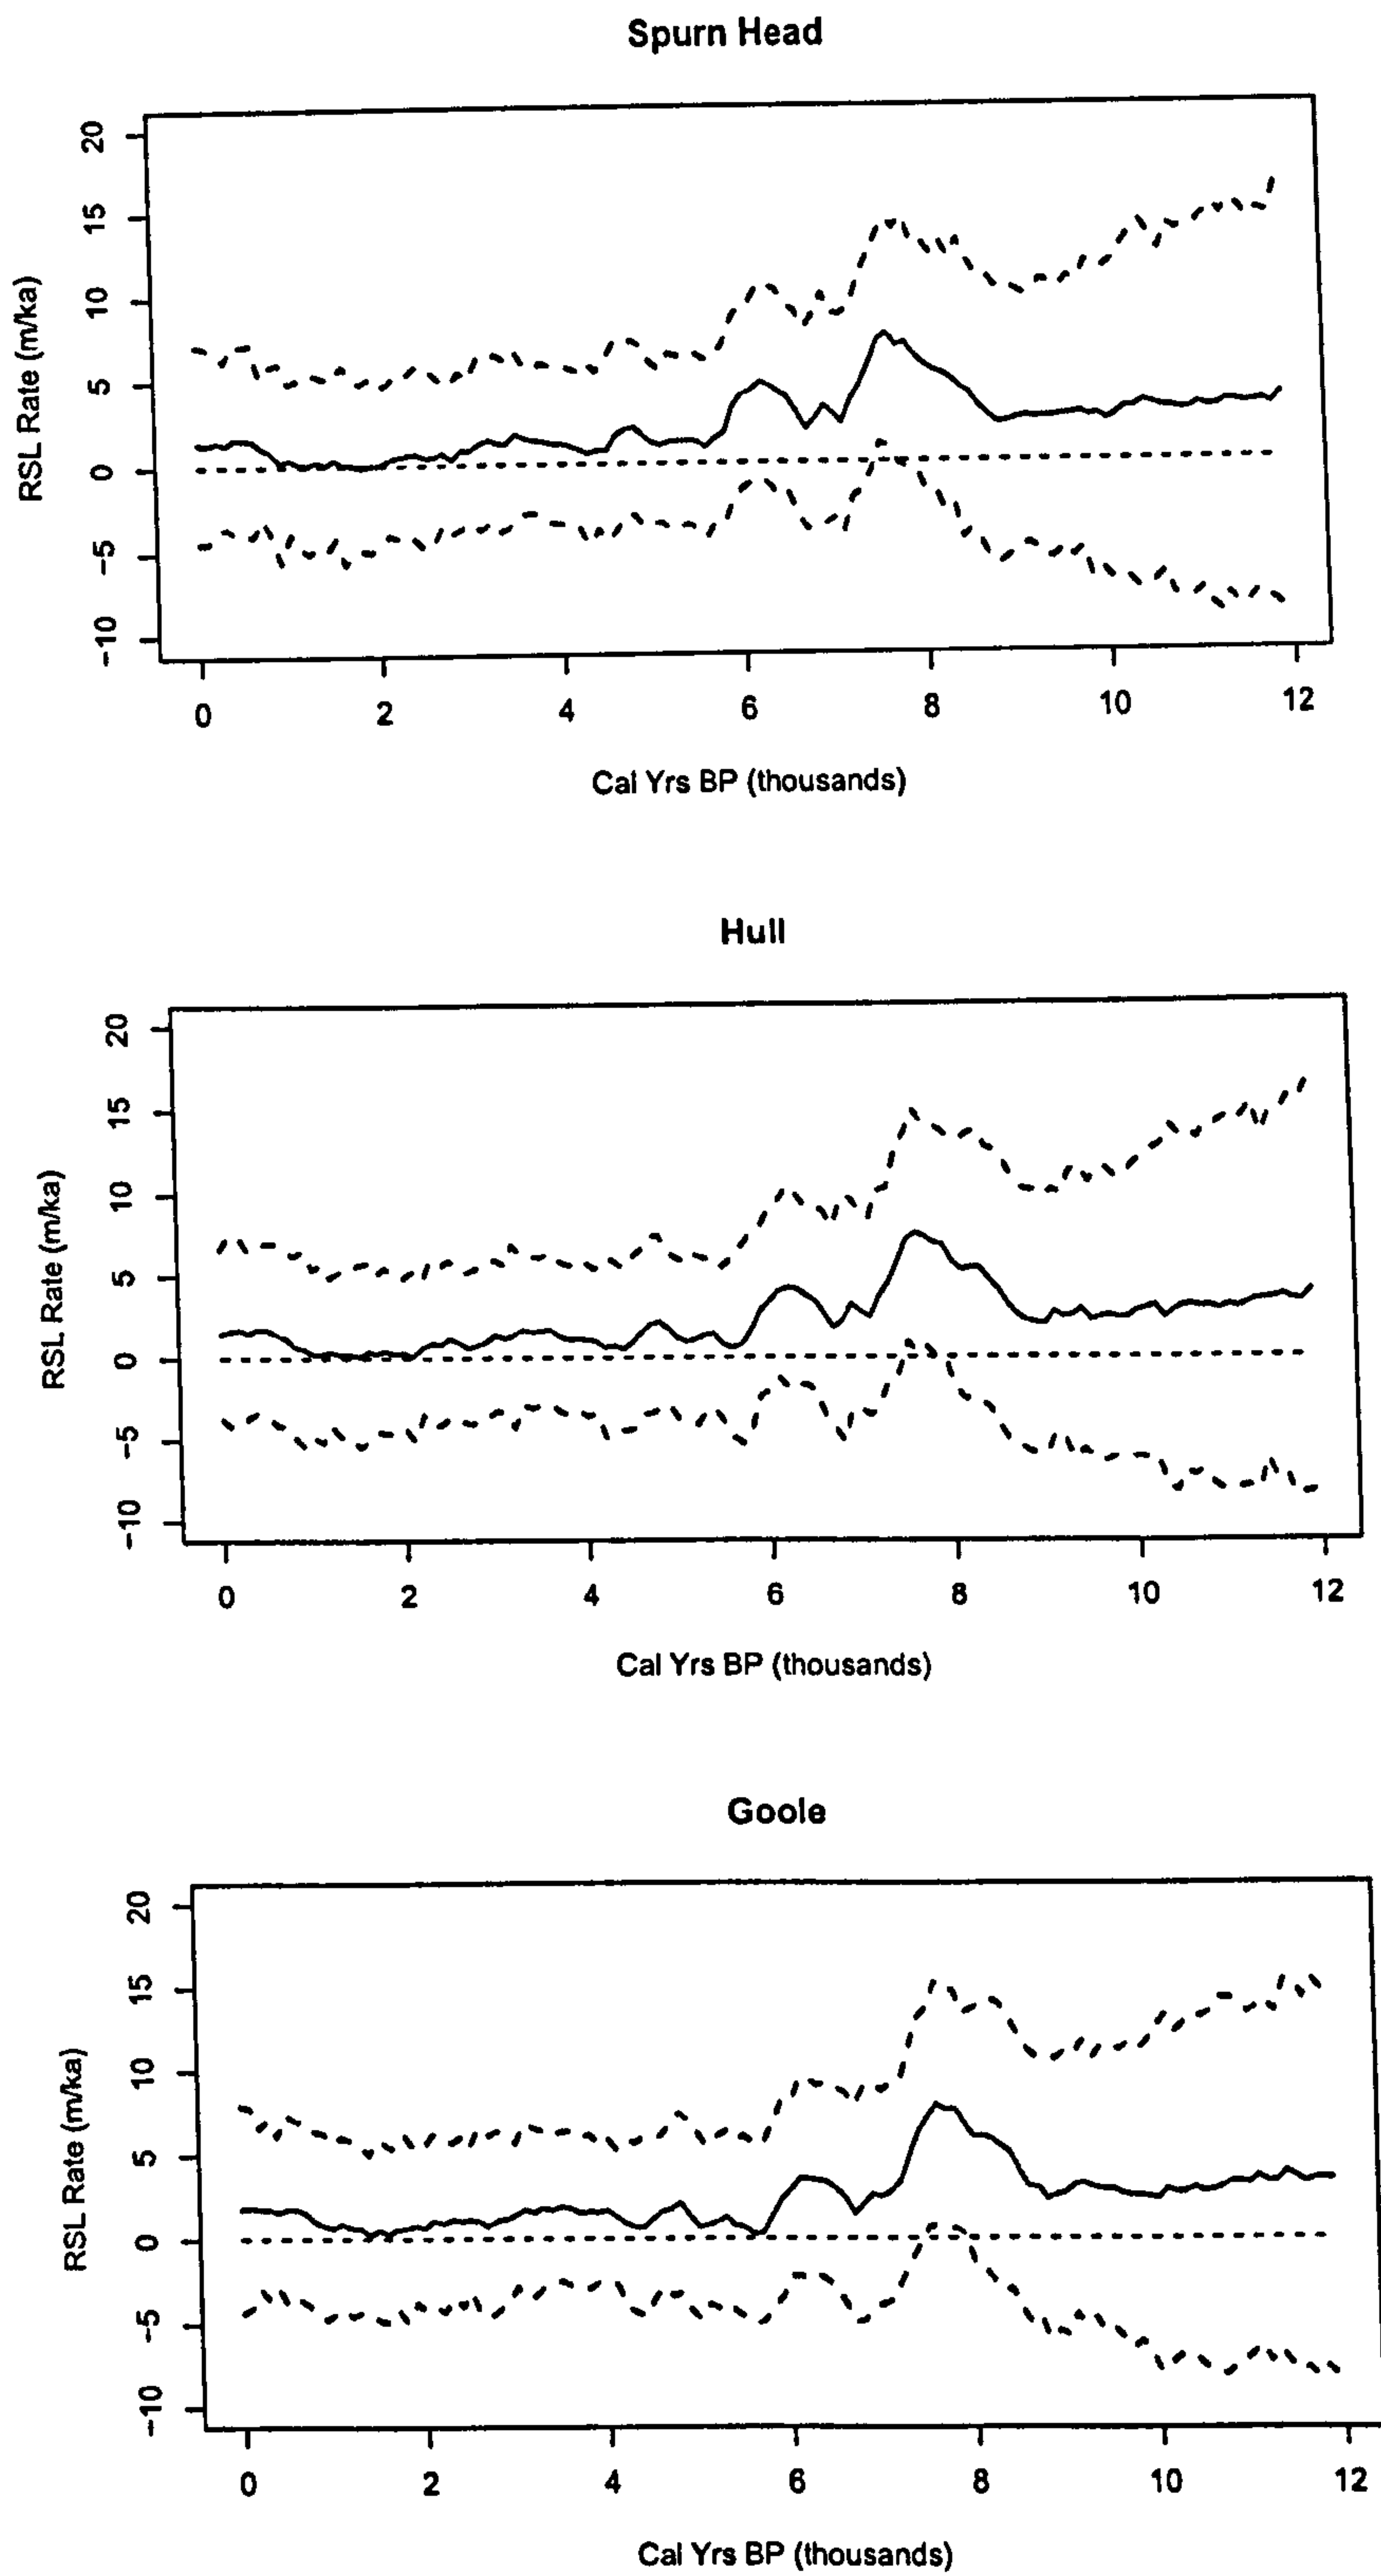


Figure 7.6: Plots of RSL speed for various locations around the Humber with 95% error bands using the GP method.




Figure 7.7: Plots of the ICE-4G model at two locations; the inner and outer Humber.

estuary respectively. The results are shown in figure 7.8. It is clear that, for both the inner and outer Humber, discrepancies occur with the geophysical model. For the inner Humber, the biggest discrepancy occurs around 5k - 3k years BP, where RSL is over-estimated by the geophysical model compared with the GP. For the outer Humber, the biggest discrepancy occurs before 6k years BP where, again, the rate is over-estimated by the geophysical curve in comparison with that estimated from the index points. Plausible reasons for the discrepancy include the lack of local factors used by the geophysical model, a lack of temporal resolution of the geophysical model, or some bias in the index point data such as poor tidal range correction.




Figure 7.8: Comparison between the spatio-temporal GP model and the geophysical model of Peltier *et al.* (2002).

7.5 Comparison with published estimated RSL rates for the Humber

This section concerns itself with the average rate at which RSL has risen over longer periods of time. Gaunt and Tooley (1974) estimate a rise of circa 20m from around 9k years BP (to the present), and approximately 15m from around 8k years BP. Long *et al.* (1998) do not draw a sea-level curve, though they estimate MHWS to have risen 9m between 7.5k and 4k years BP, an average rise of 2.57m/kyr. However, they do not have any basal index points before 6k years BP. Shennan *et al.* (2003) are mainly concerned with the adjustment of index points for tidal range change. However, once this is taken into account, they suggest a rise of around 13m since 8k years BP. Metcalfe *et al.* (2000) do not confine themselves to estimates, but fit a quartic regression to the data. The details of the fit are not made clear but the line suggests a rise of around 10m between 8k and 6k BP, followed by a rise of around 7.5m between 6k and 0k BP. Finally, Shennan and Horton (2002) give an estimate of

1.21m/kyr for the inner Estuary and 1.03m/kyr for the outer Estuary for the period 0k to 4k years BP. They give standard errors of 0.09m and 0.08m respectively (inner and outer Estuary having been defined above). None of the other papers mentioned above include an estimate of uncertainty around the rate.

Estimates of uncertainty around rates of RSL rise can be easily obtained using the Gaussian process models. Results for the GP model and comparisons with the aforementioned papers are given in table 7.1. It is clear that the GP model agrees with most of the published data when uncertainty is taken into account. However, there are some discrepancies, even with the most recent estimates. In particular, the rise estimated by Metcalfe *et al.* (2000) from 6k-0k years BP appears to be an over-estimate, whereas that estimated by Long *et al.* (1998) may be an under-estimate. Unfortunately, the lack of uncertainty estimates on the published data creates difficulty in assessing the agreement with the model. The paper Shennan and Horton (2002), which does give estimates of uncertainty around the rates, also shows some disagreement with the GP model. The model suggests the rise over the period 0k-4k years BP is lower than that suggested by Shennan and Horton (2002).

Source	Period	Estimated rate (m/kyr)	95% Interval
Gaunt and Tooley (1974)	9k-0k	2.22	-
ST GP (at Spurn Head)		2.04	(1.37,3.25)
ST GP (at Hull)		2.13	(1.55,3.44)
ST GP (at Goole)		2.18	(1.61,3.41)
Long <i>et al.</i> (1998)	7.5k-4k	2.57	-
ST GP (at Spurn Head)		2.40	(1.84,2.94)
ST GP (at Hull)		2.15	(1.64,2.66)
ST GP (at Goole)		1.86	(1.11,2.61)
Shennan <i>et al.</i> (2003)	8k-0k	1.63	-
ST GP (at Spurn Head)		1.89	(1.63,2.20)
ST GP (at Hull)		1.93	(1.69,2.18)
ST GP (at Goole)		1.94	(1.74,2.19)
Metcalf <i>et al.</i> (2000)	8k-6k	5.00	-
ST GP (at Spurn Head)		4.67	(3.48,5.89)
ST GP (at Hull)		4.64	(3.52,5.93)
ST GP (at Goole)		4.28	(3.23,5.32)
Metcalf <i>et al.</i> (2000)	6k-0k	1.25	-
ST GP (at Spurn Head)		0.95	(0.75,1.14)
ST GP (at Hull)		1.00	(0.82,1.21)
ST GP (at Goole)		1.16	(0.89,1.38)
Shennan and Horton (2002)	4k-0k	1.21 (inner)	(1.03, 1.39)
ST GP (inner)		0.95	(0.76,1.11)
Shennan and Horton (2002)	4k-0k	1.03 (outer)	(0.87, 1.19)
ST GP (outer)		0.84	(0.66,1.03)

Table 7.1: Comparison of published results on RSL rise in the Humber and those of the spatio-temporal Gaussian process model.

7.6 Compression of intercalated index points

The shift applied to the intercalated index points was given a Normal distribution with mean h and variance κ^2 . In terms of predicting RSL, h and κ were essentially nuisance parameters as they contained no information as to how RSL changed over the course of the Holocene. However, the parameters can give some information as to the levels of compaction (section 2.2.5) which the data points have undergone. It should be noted that, because compaction was not the main subject of the study, the Normal distribution was used as a convenience

to enable the use of a marginal Gaussian process. A more refined analysis of the distribution of compacted sediments could be made by using alternative distributions together with the conditional form of the Gaussian process.

A study of this sort was made by Shennan *et al.* (2000b). They tried to identify local processes by comparing a quartic regression model (based upon the Humber data) with elevations of intercalated index points. Their estimates did not incorporate the uncertainty associated with the index point, or that of the line used to fit the data. It was also undertaken before the detailed tidal range adjustment modelling of Shennan *et al.* (2003). However, the residuals from their analysis were then compared with:

- Age - older data points which were deeper than expected would lead to the conclusion that compression was a function of age.
- Distance up estuary - points further up or down estuary deeper than expected would imply differential sediment sequences.
- Depth of overburden - heavy loads on top of data points may point towards increased compression.
- Depth to base of Holocene sequence - increased sediment below the index point may lead to increased compression
- Total thickness - long sequences of sediment may allow more compression to occur.

The goal of this section is to determine which, if any, of the above factors influence the distance of an intercalated index point from the RSL curve. An improvement on the method of Shennan *et al.* (2000b) for deriving the distance of an index points to the RSL curve can be made by following these steps for each intercalated index point (IIP) in turn:

1. Simulate an age/elevation from the joint distribution of the IIP.
2. Calculate a single RSL estimate for this location/age.
3. Calculate the elevation distance from the RSL estimate to the simulated elevation.
4. Repeat.

Whilst computationally expensive, the model will produce an estimate of the distance between the IIP and the RSL curve. These can then be plotted against the relevant factors to assist in determining their importance.

A plot of the calculated difference against the possible explanatory variables is shown in figure 7.9. Shennan *et al.* (2000b) found that distance up estuary, depth to base and total thickness were relevant by calculating linear correlations on the mean values. The same analysis on this data would conclude that depth to base, total thickness and depth of overburden are of importance, and would traditionally be called significant. It is interesting to note that distance up estuary is no longer considered important, most likely because of the spatial nature

of the RSL model.

An improved method for looking at the variation in the differences can be found via the Gaussian process method. Individual GPs were fitted to each of the data sets with a linear prior and powered exponential covariance. Such models can be implemented in GeoBUGS. The resultant mean estimates of the GP with 95% limits are shown in figure 7.10. The GPs show no clear trend for either age or distance up estuary (though distance up estuary may require a more even selection of data points). However, depth of overburden, depth to base and total thickness all show increasing trends. Interestingly for depth of overburden, the trend only starts to appear above around 4m, and apparently tails off when the IIP is below 8m of sediment. Depth to base also shows a trailing off of the trend after around 8-9m. This is in agreement with work by Paul and Barras (1998).

The extra knowledge gained from this section could be used to build more detailed RSL models, where the shift proposed for IIPs is adjusted for the depth and thickness information. This is discussed further in section 8.2.5.

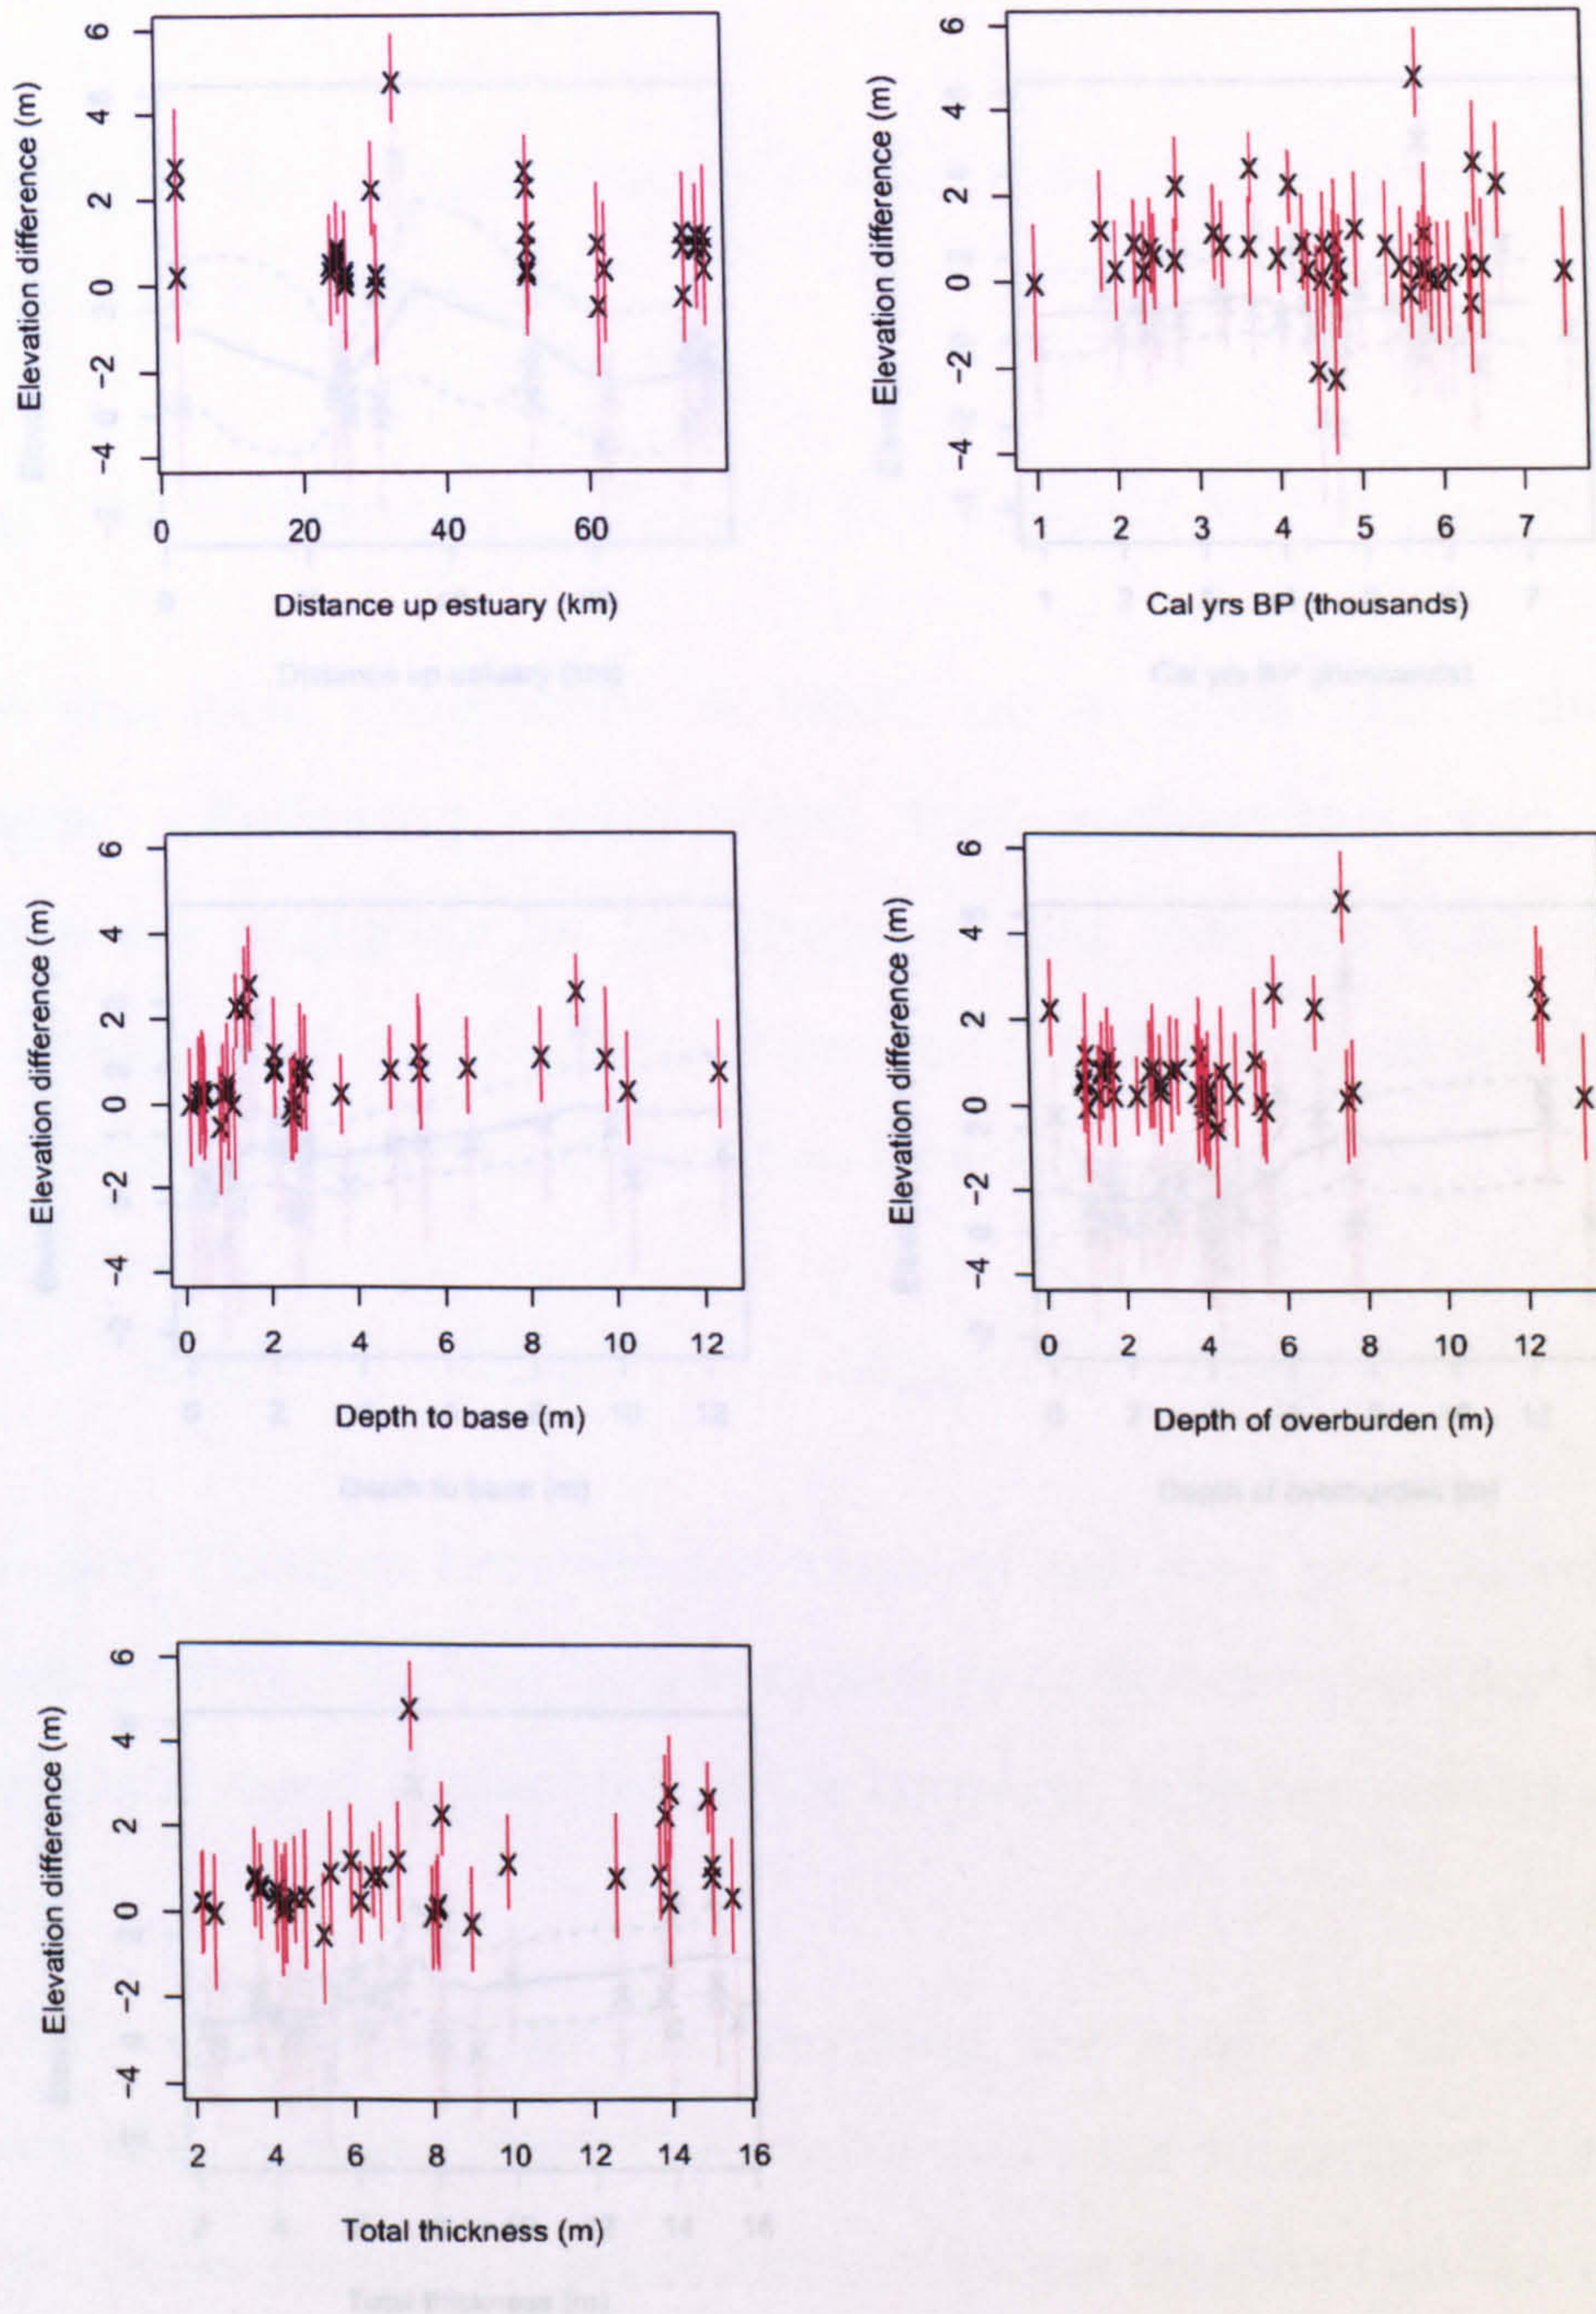


Figure 7.9: Plots of the differences in elevation between the spatio-temporal RSL curve and the intercalated index points after taking into account age and elevation uncertainty. The lines are the 95% HPD regions, whilst the X indicates the median difference.

7.7 Estimation of isostatic activity in the Humber

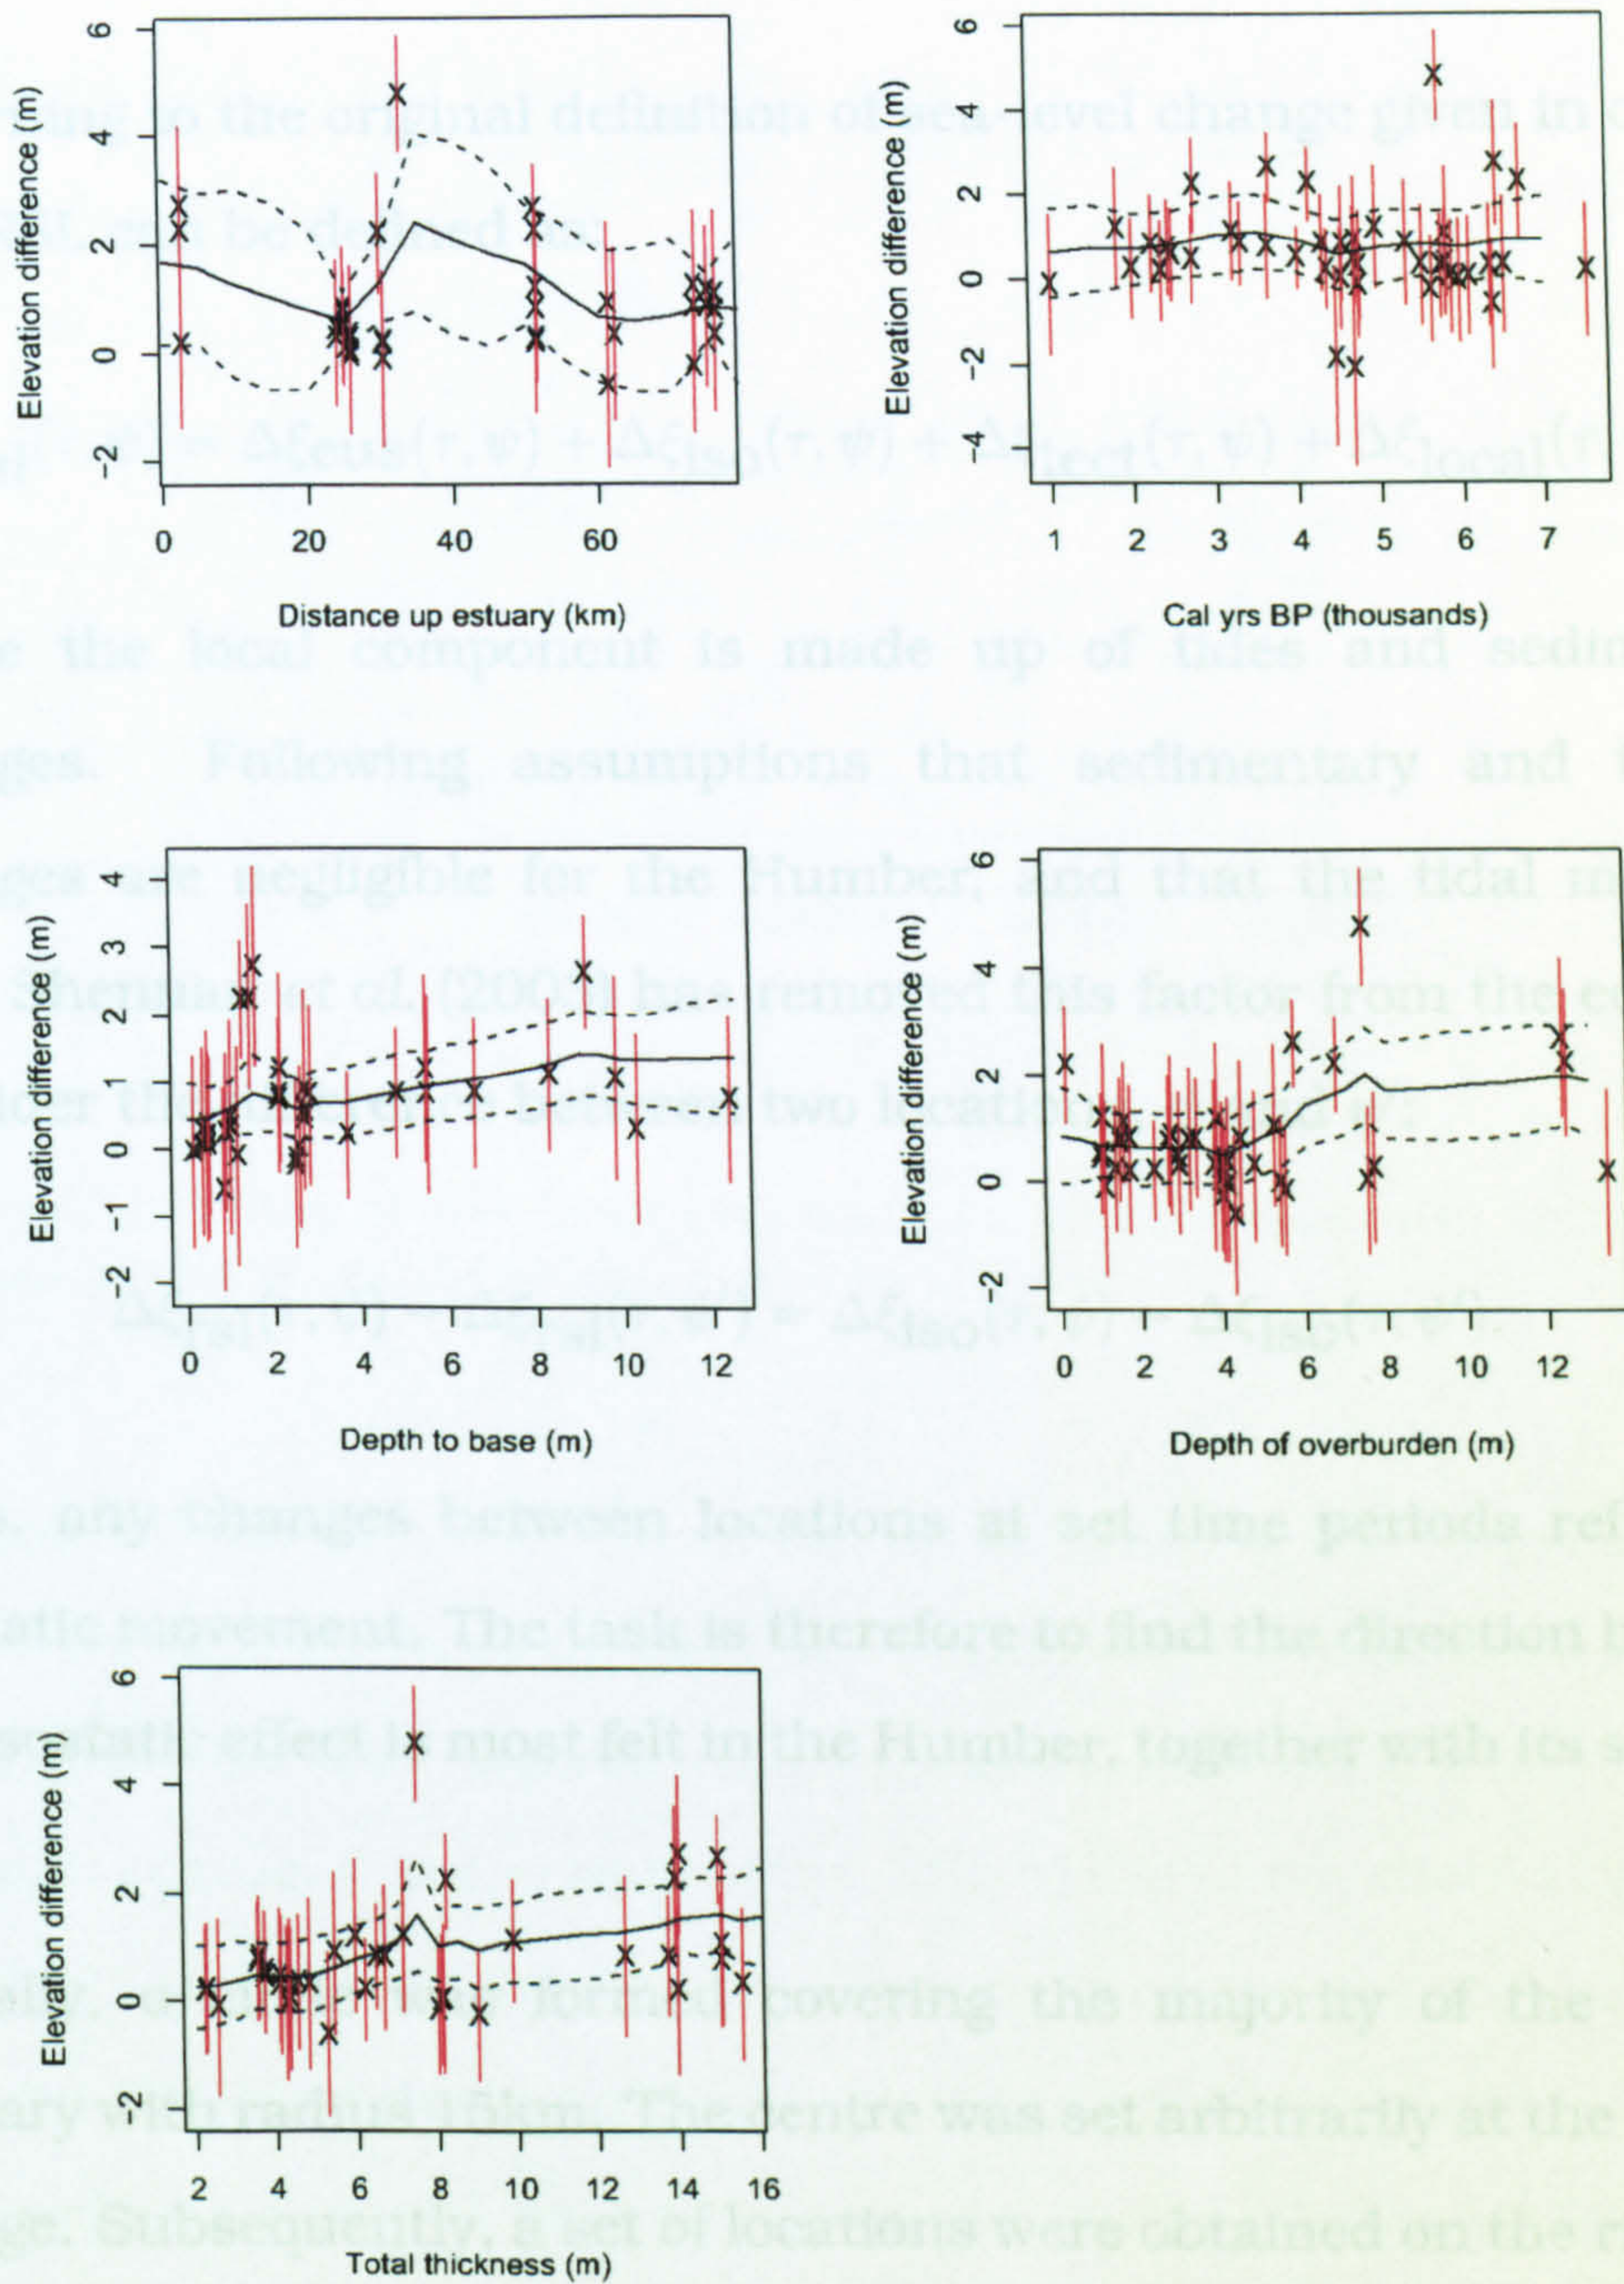


Figure 7.10: Plots of the Gaussian process with 95% error limits overlaid on figure 7.9.

7.7 Estimation of isostatic activity in the Humber

Returning to the original definition of sea-level change given in chapter 2.2, RSL can be defined as:

$$\Delta\xi_{\text{rsl}}(\tau, \psi) = \Delta\xi_{\text{eust}}(\tau, \psi) + \Delta\xi_{\text{iso}}(\tau, \psi) + \Delta\xi_{\text{tect}}(\tau, \psi) + \Delta\xi_{\text{local}}(\tau, \psi),$$

where the local component is made up of tides and sedimentary changes. Following assumptions that sedimentary and tectonic changes are negligible for the Humber, and that the tidal modelling from Shennan *et al.* (2003) has removed this factor from the equation, consider the difference between two locations, ψ and ψ' :

$$\Delta\xi_{\text{rsl}}(\tau, \psi) - \Delta\xi_{\text{rsl}}(\tau, \psi') = \Delta\xi_{\text{iso}}(\tau, \psi) - \Delta\xi_{\text{iso}}(\tau, \psi').$$

Thus, any changes between locations at set time periods reflect any isostatic movement. The task is therefore to find the direction by which the isostatic effect is most felt in the Humber, together with its strength.

Initially, a circle was formed covering the majority of the Humber estuary with radius 15km. The centre was set arbitrarily at the Humber Bridge. Subsequently, a set of locations were obtained on the rim of the circle, to be used in the Gaussian process model, at $\pi/8^c$ intervals. GP estimates of the RSL at these sites were obtained, and subtracted from those diametrically opposite. It should be expected that those with the greatest difference would inform about the nature of isostatic change.

A plot of the circle used is shown in figure 7.11.




Figure 7.11: Plot of the locations used to test for isostatic activity. The outline of the Humber estuary is also shown.

For each of the points on the circle, an estimate of RSL was calculated at 2k, 4k, 6k, and 8k years BP. The distribution for each point was then compared with that diametrically opposite. In theory, if any isostatic change is present at any time point, RSL to the north should be slightly higher than that to the south. This is because the ice-load which covered Scotland during the last glacial maximum should be causing the northern shore of the estuary to rise slightly faster (or fall slightly slower) than the southern estuary.


Plots of the difference in each direction are given in figure 7.12. It is clear there is little north-south effect, apart from a small north-

north-east versus south-south-west at 6k years BP. The larger effects all appear to be in the region of 6k years BP, and involve an east-west effect (west RSL being lower than east); most likely that of tidal range change or even the effect of the Fennoscandian ice sheet (as discussed by Peltier, 1998b). It may be that isostatic activity cannot be estimated for the Humber due to the uncertainties involved in RSL measurement. Furthermore, it appears that the assumption of tidal variation being removed from the Humber is questionable for the period around 6k years BP.


7.8 The use of archaeological data

This section tests whether the archaeological data has had a positive effect on the production of RSL curves. A positive effect is defined as having reduced the uncertainty in the RSL estimates. To this end, the temporal stationary Matérn model was run without the extra archaeological data. Plots of the differences between the upper and lower 95% intervals and the mean in the resulting RSL estimate are shown in figure 7.13. It is clear that the addition of the archaeological data has some effect of reducing the uncertainty. The size of this effect appears to be up to around 50cm in places. Peaks still occur at around 0.4k years BP and 1.4k years BP, possibly indicating that the inclusion of archaeological data points from the early middle ages up until the early post-mediaeval would improve RSL estimates even further.


North-South 2k-8k yrs BP




North north east-South south west 2k-8k yrs




North east-South west 2k-8k yrs BP




East north east-West south west 2k-8k yrs BP




East-West 2k-8k yrs BP




East south east-West north west 2k-8k yrs BP



South east-North west 2k-8k yrs BP



South south east-North north west 2k-8k yrs



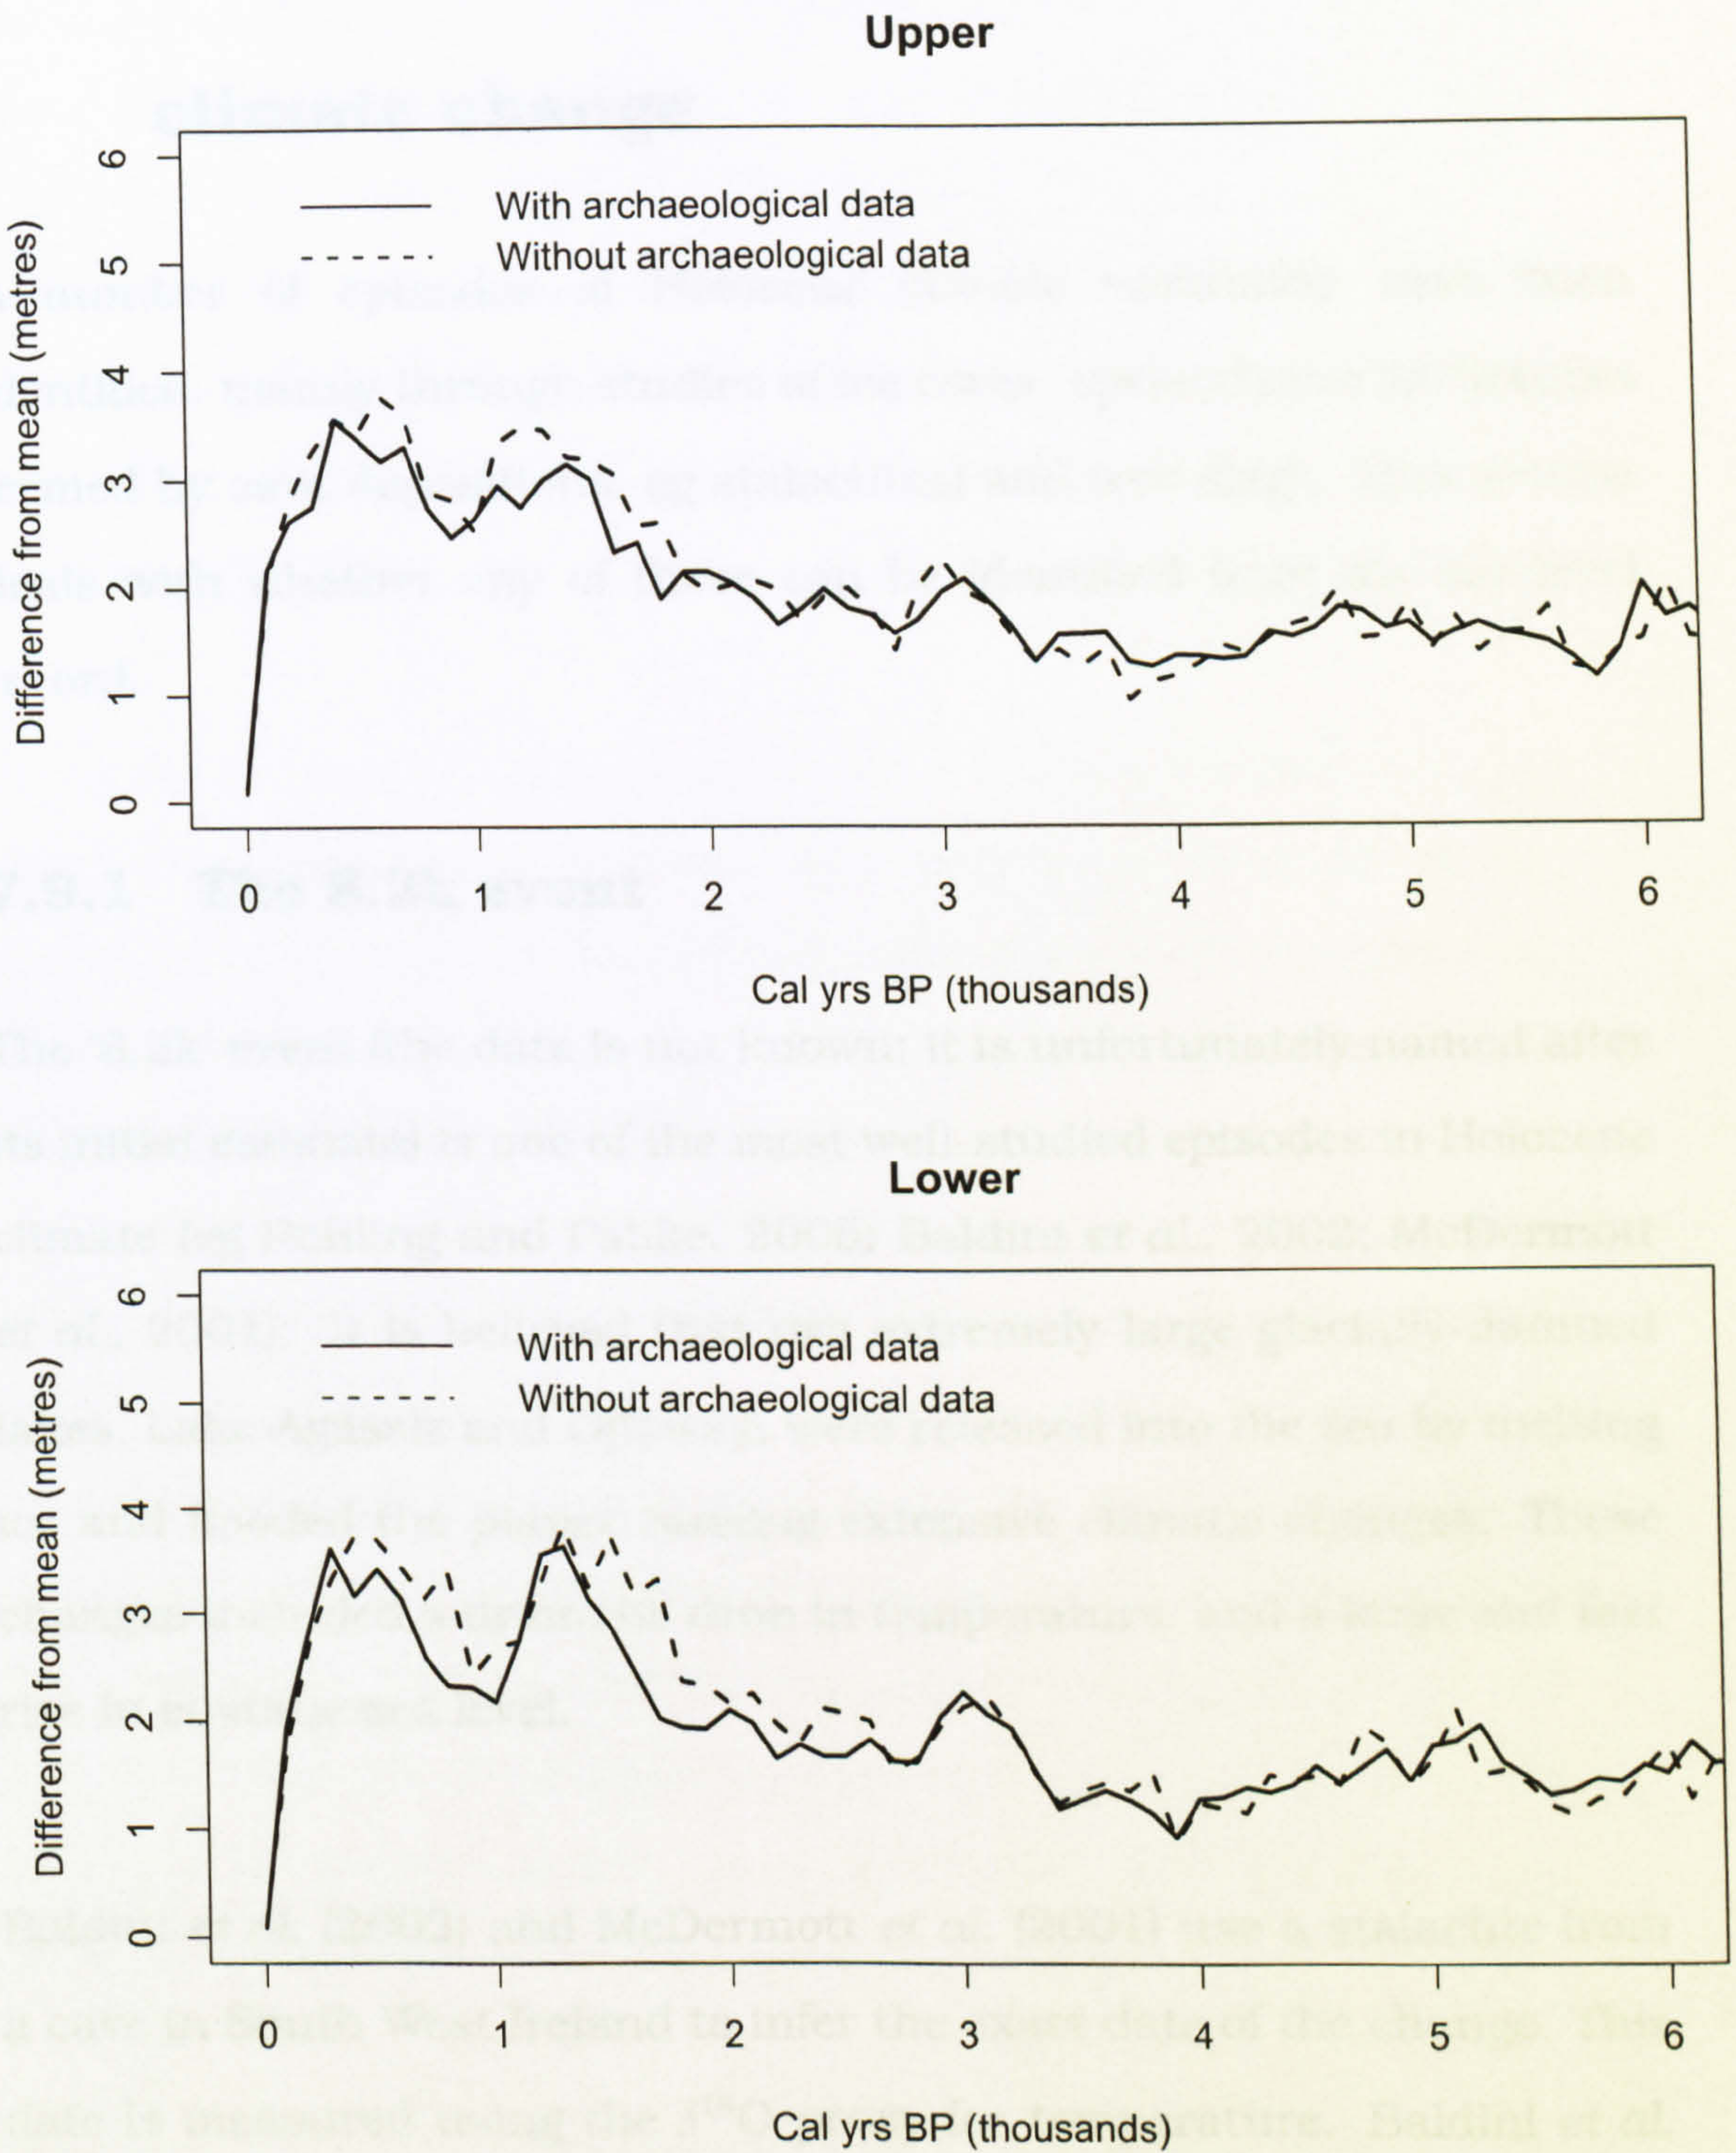


Figure 7.13: Plot of the difference between the mean of the GP estimate with (solid line) and without (broken line) the archaeological data. The top panel shows the difference between the mean and the upper 95% limit, the bottom panel the difference between the mean and the lower 95% limit.

7.9 Indications of well-studied periods of climate change

A number of episodes of Holocene climate variability have been identified, mainly through studies of ice cores, speleothems (structures formed by cave depositions, eg stalactites) and tree-rings. This section deals with whether any of these can be identified from the sea-level record.

7.9.1 The 8.2k event

The '8.2k' event (the date is not known; it is unfortunately named after its initial estimate) is one of the most well-studied episodes in Holocene climate (eg Rohling and Pälike, 2005; Baldini *et al.*, 2002; McDermott *et al.*, 2001). It is believed that two extremely large glacially-dammed lakes, Lake Agassiz and Ojibway, were released into the sea by melting ice and flooded the planet causing extensive climatic changes. These changes included a dramatic drop in temperature, and a large and fast rise in eustatic sea level.

Baldini *et al.* (2002) and McDermott *et al.* (2001) use a stalactite from a cave in South West Ireland to infer the exact date of the change. This date is measured using the $\delta^{18}\text{O}$ proxy for temperature. Baldini *et al.* (2002) noticed an abrupt change centred on 8330 ± 80 years cal BP; McDermott *et al.* (2001) put it at 8320 ± 120 years cal BP. They estimate the event lasted around 37.5 years.

Rohling and Pälike (2005) take a more conservative approach. They argue that the information from the Greenland ice cores show more variability than a single jump at 8.2k years BP. They estimate that any sudden change in Holocene climate could have occurred as early as 8.6k years BP, whilst there is evidence that these changes were still in effect at 8k years BP. They conclude that the extent of these events are yet to be determined.

In summary, there are two questions to be answered when looking at this event. Firstly, it must be seen as to whether there is any change around the region of 8.2k years BP; secondly, if any such change is to be found, the region over which its effect is felt must be identified. There are various tools with which to achieve these aims; the plots of RSL over the Humber, the plots of RSL speed, and those of the functional principal components. First consider the problem of how to identify an event such as the 8.2k event. If such an event occurred, a large sea-level rise must have taken place in the Humber at some point. A plot of the RSL for the Humber over the period 7k to 9k years BP is shown in figure 7.13. It appears that the uncertainty associated with RSL before 8k years BP precludes the possibility of drawing conclusions about abrupt changes in RSL during this period. However, there does appear to be a levelling off at around 7.5k years BP.

The second part of the investigation into the 8.2k event concerns its chronology; its start, duration, and end. To this end, change-point linear regression was used to establish the times when the event might




Figure 7.14: Plots of mean RSL for Spurn Head with 95% HPD intervals for the period 7k - 9.2k years BP.

have taken place. The model was formulated as:

$$y_i | \mu_i, \sigma \sim N(\mu_i, \sigma) \quad (7.10)$$

$$\mu_i = \alpha_1 + \beta_1(x_i - x_{c1}) \text{ if } x_i \leq x_{c1} \quad (7.11)$$

$$\mu_i = \alpha_1 + \beta_2(x_i - x_{c1}) \text{ if } x_{c1} < x_i < x_{c2} \quad (7.12)$$

$$\mu_i = \alpha_1 + \beta_2(x_{c2} - x_{c1}) + \beta_3(x_i - x_{c2}) \text{ if } x_i \geq x_{c2}. \quad (7.13)$$

Here, x is the set of design dates chosen at which to output the Gaussian process, y is their association elevation. The model is forced to have two change-points, x_{c1} and x_{c2} , corresponding to the start and end of the event respectively. The model is based on a set of parameters $(\alpha_1, \beta_1, \beta_2, \beta_3, x_{c1}, x_{c2})$. Prior information is used on the β s

so that $\beta_1, \beta_3 \geq \beta_2$, and on the change-points so that $x_{c1} > x_{c2}$. The prior on the β s causes RSL to rise faster whilst the event is occurring. Finally, the duration of the event can be calculated from $x_{c1} - x_{c2}$.

The change-point model was initially run for the Spurn Head Humber data of the non-stationary spatio-temporal Matérn model using MCMC. The data given to the model only involved that between 7k and 9k years BP so as to not to pick up other possible events occurring in the Holocene. Mixing was slow given the small set of parameters. However, convergence was achieved after around 3000 iterations and involved a small degree of thinning. Posterior distributions for x_{c1} and x_{c2} are shown in figure 7.15. It may be possible to conclude that there is evidence of an event, possibly related to global flooding, which affected the Humber, starting at around 8.5k years and ending at around 7.4k years cal BP. A 95% HPD interval for the duration of the event gives 0.99k to 1.37k year. The start date requires some degree of skepticism as there is large uncertainty and no basal index points before 8k years BP.

7.9.2 The 4k event

The ICE-5G model of Peltier *et al.* (2002) argues that all glacial melting had stopped by 4000 calendar years BP. It could therefore be argued that thereafter the eustatic component of sea-level change should have remained constant. In terms of the effect on relative sea-level change in the Humber, the point at which melting must have stopped will cause




Figure 7.15: Plots of the density for the start and end of an event at Spurn Head in the outer Humber. The solid line is the mean estimate of RSL, with 95% intervals (dotted). The densities at the bottom show the estimated change-point densities.

a slow down in sea-level rise. However, negative isostatic change would have continued throughout this period, thus still showing RSL rise. The task for this section is to note if there is a marked drop-off in RSL change around 4k years BP.

A plot of the relevant period of RSL for Spurn Head are shown in figure 7.16. It appears that there is a small amount of levelling off. The exact date of this is unclear; there is certainly a slowing down of RSL between 6k and 2k years BP. As before, a change-point linear regression was used on the Gaussian process output to test where in the range the change-points might lie. Only one change-point was implemented, which meant the model could be run in WinBUGS. The results are

shown in figure 7.16. It is clear that the one change point is very clearly defined at around 5.5k years BP, well before the estimated 4k date.




Figure 7.16: Mean lines with 95% intervals for Spurn Head. The density of the change-point is also shown.

7.9.3 The little ice age and mediaeval warm period

A number of other smaller and more recent climate events have been noticed by climatologists and are discussed in, for example, Maasch *et al.* (2005); Matthews and Briffa (2005); Esper *et al.* (2002). This section considers only two of the most well-studied; the little ice age and the mediaeval warm period.

Matthews and Briffa (2005) look at tree ring densities over the period of the second millennium and deduce there to have been temperature anomalies between 1570-1900 AD (380-50 BP), with glacial anomalies

between 1300-1950 AD (650-0 BP), which they call the 'Little Ice Age' (LIA). They believe these changes were centered over the Alps but may have been felt more widely. Their estimates for temperature change across the northern hemisphere suggest little change in temperature around the Humber, though the glacial effect may be more apparent in studies of RSL.

The mediaeval warm period (MWP) was studied by Esper *et al.* (2002). Using long-term tree ring records they deduce a period, replicated over 14 sites in the northern hemisphere, where temperature has increased. They estimate this period to have occurred between 1300 and 1000 years BP.

Maasch *et al.* (2005) consider both the MWP and LIA. Using comparisons between atmospheric ^{14}C , they estimate that these events were felt on a global scale. They consider the MWP period to occur between 800 and 1200 AD (1150-750 BP), and to have been represented by temperatures slightly higher than those of the present day. They estimate the LIA to have overtaken the MWP by around 1400AD \pm 40 (550BP \pm 50). They further estimate a widespread cooling of around 0.5-1°C.

The effect that any of these climate fluctuations would have on RSL is arguable. The LIA may have seen a small drop in RSL, whilst the MWP may have seen a small rise. Thus any change-point model should pick up a fall then rise approximately over the period of 2k to 0k

years BP. A model with 2 change-points was run for the appropriate period. However, the two change-points tended to converge. Figure 7.17 shows a model with one change-point. It is poorly-defined, a 95% HPD interval gives 0.64k to 1.72k BP. This may provide weak evidence for a mediaeval warm period.




Figure 7.17: Mean lines with 95% intervals for Spurn Head. The density of the change-point is also shown.

7.10 Summary

This chapter has examined the output of the RSL Gaussian process method with a view to making inference about the nature of RSL change in the Humber. The kernel density estimates gave an idea as to the variable uncertainty in the RSL estimates. Functional principal components were used to assess the variability in this uncertainty

across location. Rates of RSL change were calculated using splines, and the original GP method. The geophysical models were found to disagree with the GP model for certain time periods. Intercalated index points were shown to be linked most strongly with information regarding their depth. The archaeological data was found to be important in reducing the uncertainty in the RSL estimates. Finally, the RSL model showed up a few of the wider changes in climate outlined during the Holocene, though these were not necessarily in accord with previous temporal estimates.

Chapter 8

Conclusions and further research

This chapter briefly reviews the findings of the thesis. It also considers some statistical and non-statistical aspects which would lead to improved sea-level research. Topics discussed include: temporal extrapolation and geophysical models, spatial modelling of the tendency data, spatio-temporal non-stationarity and improved modelling of intercalated sediments.

8.1 Conclusions

This thesis has provided the first attempt at drawing RSL curves which take account of all the available information; most importantly the uncertainties in the data points. The method could be used as an extension to any RSL paper to produce appropriate RSL curves. The

results presented in this thesis clearly show promise in estimating both the smoothness and variability in the Humber RSL data.

In summary, the other conclusions can be laid out as follows:

- The archaeological data used in this thesis has reduced uncertainty in the more recent RSL record for the Humber. However, with more data points (particularly from the post-mediaeval period) this uncertainty could be even further reduced. The Humber data set of SLIs still lacks data from the early middle ages to the post-mediaeval period. Another problem with the archaeological record is that of the quality of the data points. The relative value of basal index points over intercalated or limiting points is clear. If sea-level-associated archaeological artefacts could be found above non-compressible material their value would be much greater.
- The uncertainty in the RSL record before 8k years BP is very large. This is because there are only limiting points for this period. Any conclusions drawn about RSL change before 8k years BP must be regarded with strong skepticism.
- Clear spatial effects are present in the data. These tend to follow an east-west trend, and point toward the conclusion that the tidal modelling has not completely removed a tidal range trend in the RSL record. There appears to be no observable isostatic effect.
- The geophysical models show some discrepancies with the Gaussian process RSL model. The source of this discrepancy may

lie with the poor resolution of the geophysical models, with their lack of ability to pick up highly local features, or with incomplete tidal range correction in the GP model.

- The tendency data when used temporally does not change the estimated RSL record for the Humber. The extra information is largely ignored by the GP.
- The intercalated index point residuals show a positive relationship with depth of overburden, depth to base and total thickness of sediment. These are as expected.
- The RSL estimate for the Humber shows some evidence of an increased rise between 8.5k and 7.4k years BP. This rise may be related to a wider change in climate during this period such as the 8.2k event. However, it appears to last for a longer period than those outlined in relevant publications, and the start of the event lies in a period with high uncertainty.
- There is also some evidence of an increase in rate of RSL change at around the time expected of the Mediaeval warm period. The exact date of this change is subject to large uncertainty.

8.2 Further research

8.2.1 Guidance for archaeologists and geologists

Future data may still be collected during field trips to the Humber estuary to find either archaeological sites of interest or further index

points. If these data are to be used as part of an RSL study there are several factors of which it would be useful to note.

For the geologists looking for future index points and looking to core in the Humber, it is worth noting that the uncertainty produced by the GP model between 8k and 2k years BP is small compared with elsewhere. The outside of this range is therefore worth concentrating on, particularly before 8k years BP, although this may involve a substantive coring element. It would be especially useful if basal or intercalated index points could be found before this date as they would provide an excellent bound on RSL at a period of great scientific interest. Sea-level research would also benefit with a clearer understanding of the different methods of measuring uncertainty. It is imperative that errors, particularly in elevation, are calculated as standard errors or similar. The confusion as to what the uncertainty represents is a very real barrier to reducing RSL uncertainty.

Many more archaeological studies could have been incorporated in this thesis had the data been more clearly presented in a fashion suitable to create index points and limiting points. Unfortunately, this is not the main task of the archaeologist collecting the data. Improvements could be made by more thorough levelling of the base of the site, and noting more prominently the sources or error involved in this calculation. The archaeological data has an inherent advantage in that it usually dates to the more recent past and would be therefore be more effective in reducing RSL uncertainty. A selection of well-levelled, well-dated sites on the Humber foreshore would most definitely improve the RSL model

in the latter Humber.

8.2.2 Temporal extrapolation and geophysical models

The prediction of *future* sea-level change is of much interest both to those researching sea level and those studying the wider field of climate change. This thesis has concentrated on interpolating between modern and prehistoric sea-level indicators in order to provide an accurate estimate of relative sea level with associated uncertainty. Estimating future sea level corresponds to extrapolating beyond the modern set of data points. One facet of the Gaussian process method is that any extrapolated prediction will be strongly determined by the mean of the Gaussian process. This is a direct consequence of the autocovariance function, $\gamma(t)$, decaying to zero as $t \rightarrow \infty$. In Bayesian terms, the absence of data forces the model to consider only the prior specification.

A possibility for better extrapolation lies with geophysical modelling (section 2.3.9). As seen in chapter 6, parameters in the mean of the Gaussian process have estimable posterior densities. In this thesis these parameters have little or no physical meaning. If, however, the mean of the Gaussian process was given a form similar to that of one of the current set of geophysical models, extrapolation may be possible (eg Kennedy and O'Hagan, 2001).. Furthermore, posterior estimates of the parameters in the mean function may now have physical meaning.

Traditionally, methods for fitting complex geophysical models have

not used the data from SLIs. However, there is no reason why the estimation of the physical parameters could not be accomplished in a likelihood setting.

Define the relative sea-level elevation to be:

$$RSL(X, t) = f(X, t, \theta) \quad (8.1)$$

where X is a location, t a time, and θ to be a set of unknown physical parameters contained within the function $f(\cdot)$. This comprises the geophysical model. By giving this function a Gaussian process prior it becomes possible to both fit and estimate parameter uncertainty:

$$Y \sim GP(f(\cdot) + h, V + \Sigma + \Psi) \quad (8.2)$$

As before, V is an appropriate autocovariance function, Σ is the known elevation variances of the index points, h and Ψ deal with shifts for intercalated and limiting points, and Y is the set of elevations.

8.2.3 Spatial modelling of tendency data

As mentioned in section 6.2.3, the temporal use of tendency data is often downplayed due to uncertainties when comparing sea level movement and coastal change (Shennan, 1994). However, it is believed possible that the occurrence of similar tendencies found at similar times and similar locations may add weight to arguments for relative sea-level change. The task is then to define, firstly, how close each tendency must be in order to inform RSL change and, secondly, to

determine the size of the effect on RSL given an indication of 'true' tendency at a particular location and time.

A first step can be considered by separating positive and negative tendencies (an assumption of independence here would be a matter for debate), and describing each as a spatio-temporal point process (Cressie, 1991; Schabenberger and Gotway, 2005). The tendency data does not neatly fit in to the types mentioned in the literature:

- A space-time shock point process, where observations occur instantly over both space and time.
- A space-time survival point process, where events are 'born' at a random location and time, and then 'live' for a random length of time.
- A space-time explosion point process, where the point process is two-stage: firstly temporally and then spatially.

The tendency data corresponds most to the survival process. However, instead of being born and then living for a set period of time, tendency data points could be considered marked by the temporal density. This would create an arbitrarily large data set. Evidence of clustering would then inform RSL change.

8.2.4 Complete spatio-temporal non-stationarity

The non-stationary Matérn autocorrelation function of Paciorek and Schervish (2004) can be extended to work in multiple dimensions. With

such an extension it would then be possible to model the spatial aspect of the data (which is in two dimensions) as a non-stationary Gaussian Process. Each element, Ω_i , in equation 5.22 is now a square matrix of dimension two. In the one-dimensional case, these were given a log-Gaussian process with stationary Matérn covariance. Paciorek and Schervish (2004) use the spectral decomposition of Ω :

$$\Omega = \Gamma D \Gamma^T. \quad (8.3)$$

Here, Γ is an eigenvector matrix and D is a diagonal matrix of eigenvalues. To ensure smoothness in the spatial variability of the covariance, these elements are forced to vary smoothly themselves. They give the set of eigenvectors and the logarithm of each eigenvalue a GP with a stationary anisotropic Matérn covariance.

8.2.5 Applying depth and thickness information to intercalated index point shifts

The intercalated index points were shown to have links with their depth in the stratigraphic record. The extra data could be used to better constrain the RSL record through the Gaussian process. Following the notation of section 5.5, an adjustment to γ_i , where i is in the set of IIPs would be required:

$$\gamma_i \sim N(d_{IIP} \beta_{IIP}, \kappa^2) \quad (8.4)$$

where d_{IIP} represents the depth information and β_{IIP} is a set of unknown parameters. Alternatively, if such a relationship were deemed

too simple, it would be possible to assign a separate GP to the IIP shifts based on the depth data.

Appendix A

LOIS Data

This appendix contains all the basal index points, intercalated index points and limiting points used as the data set for this thesis, taken from the LOIS database. The terms reference water level (RWL) and indicative range (IR) are as defined in section 2.3.3. The radiocarbon ages are calibrated as outlined in section 4.2.1. They are summarised in the table using 95% HPD intervals which assume uni-modality.

Site	Lab code	Type	Location <i>x</i> <i>y</i>	Alt (m.O.D.)	RWL	IR	¹⁴ C	Age Error	Cal 95% range min max	RSL	Error
Paradise Farm, PA94-12	CAM41317	Basal	50100 39862	0.49	3.6	0.2	3170	50	3520 3260	-2.91	0.21
South Ferriby, HMB3	AA25577	Basal	49685 42112	-4.66	4.2	0.2	6000	55	6990 6680	-8.66	0.21
Winterton Ings, WT195-6	CAM41321	Basal	49580 42003	-2.99	3.6	0.2	4300	50	5030 4680	-6.39	0.21
Winterton Ings, WT195-3	CAM41319	Basal	49541 42002	-0.6	3.6	0.2	3560	60	3990 3700	-4	0.21
Garthorpe Farm, GF96-1B	AA24138	Basal	48492 41925	-2.91	4.2	0.2	4845	55	5700 5370	-6.91	0.21
East Clough-1, Melton	SRR4748	Basal	49740 42475	-0.35	3.9	0.2	3640	45	4090 3830	-4.05	0.37
Paradise Farm, PA94-12	CAM41320	Basal	50100 39862	0.22	3.6	0.4	3490	50	3890 3630	-3.38	0.4
South Ferriby, HMB3	AA25578	Basal	49685 42112	-4.76	4.2	0.2	5985	55	6970 6680	-8.96	0.21
Garthorpe Farm, GF96-1B	AA27584	Basal	48492 41925	-4.38	4.2	0.4	5870	55	6830 6510	-8.98	0.4
East Clough-2, Melton	SRR4749	Basal	49740 42475	-0.56	3.9	0.2	3770	45	4300 3990	-4.26	0.37
South Marsh, HMB5	AA23431	Basal	52289 41339	-6.48	3.4	0.2	6050	70	7150 6720	-9.68	0.21
South Marsh, HMB5	AA23430	Basal	52289 41339	-6.37	3.4	0.2	6135	75	7230 6790	-9.57	0.21
Immingham, HMB10	AA23432	Basal	52076 41517	-6.66	3.4	0.2	6245	80	7340 6910	-9.86	0.21
Immingham, HMB10	AA23433	Basal	52076 41517	-6.77	3.4	0.2	6520	75	7570 7270	-10.17	0.21
Pyewipe Pumping Station	AA27586	Basal	52570 41080	-2.38	3.1	0.2	4985	55	5890 5610	-5.48	0.21
Lockham, HMB12	AA23434	Basal	53932 41720	-3.45	3	0.2	5425	70	6380 6000	-6.45	0.21
Pyewipe Pumping Station	AA27585	Basal	52570 41080	-2.22	3.1	0.2	5065	60	5920 5660	-5.12	0.21
Kilnsea Warren	UB3900	Basal	54230 41470	-1.29	2.88	0.8	4562	59	5450 5000	-4.17	0.81
Kilnsea Warren	UB3901	Basal	54230 41470	-1.19	2.88	0.2	4384	54	5260 4850	-3.87	0.25
Dunswell, HMB2A	AA25592	Basal	50746 43503	-0.65	3.6	0.2	5000	55	5900 5610	-4.25	0.21
Old Den, HI3	AA27583	Basal	54043 41231	-10.31	3	0.2	6875	60	7840 7590	-13.11	0.21
Humber Bight, HI1	AA26378	Basal	53629 41623	-7.23	3	0.2	6140	65	7220 6810	-10.03	0.21

Table A.1: Table of LOIS data. -99 denotes missing data. Reference water level includes tidal correction.

Site	Lab code	Type	Location		Alt (m.O.D.)	RWL	IR	¹⁴ C	Age	Cal 95% range		RSL	Error
			x	y						min	max		
Hasholme-430	HAR7007	Intercalated	48222	43261	-4.41	4.2	0.2	5710	100	6730	6310	-8.81	0.21
Dirtness Levels, DL961B	AA26376	Intercalated	47465	41215	-2.94	4.2	0.4	4610	60	5560	5060	-7.14	0.4
Newlands, HMB19	AA23439	Intercalated	46868	42464	-2.84	4.3	0.2	4800	55	5640	5330	-6.94	0.21
East Ferry, EF97-1	AA27615	Intercalated	48294	39945	-2.01	4.2	0.2	4890	55	5750	5490	-6.01	0.21
Dirtness Levels, DL961	AA26377	Intercalated	47465	41215	-1.84	4.2	0.2	3880	55	4440	4120	-5.84	0.21
Newlands, HMB19	AA23440	Intercalated	46868	42464	-3.25	4.3	0.2	5075	55	5930	5670	-7.73	0.21
Garthorpe, HMB16	AA23437	Intercalated	48594	41923	2.56	4.2	0.2	3425	65	3850	3490	-6.56	0.21
Newlands, HMB19	AA27141	Intercalated	46867	42463	-0.76	4.3	0.4	3395	75	3840	3470	-5.06	0.4
Dirtness Levels, DL961	AA26375	Intercalated	47465	41215	-0.12	4.2	0.2	3020	60	3360	3000	-4.5	0.21
Garthorpe Farm, GF96-1B	AA23887	Intercalated	48492	41925	-1.89	4.2	0.2	4060	60	4810	4410	-5.89	0.21
Garthorpe Farm, GF96-1B	AA23886	Intercalated	48492	41925	-1.05	4.2	0.2	3920	75	4560	4120	-5.05	0.21
East Ferry, EF95-2	AA23438	Intercalated	48294	39945	0.29	4.2	0.2	3090	50	3430	3140	-4.09	0.21
East Ferry, EF95-2	AA22678	Intercalated	48294	39945	0.93	4.2	0.2	1880	45	1920	1700	-3.07	0.21
Garthorpe Farm, GF96-1B	AA23888	Intercalated	48492	41925	-2.6	4.2	0.2	4365	65	5260	4650	-6.98	0.21
Sandholme Lodge	SRR4894	Intercalated	48390	43150	-3.35	4.2	0.2	5615	45	6490	6300	-7.35	0.22
Sandholme Lodge	SRR4743	Intercalated	48390	43150	-1.78	4.2	0.2	4170	45	4830	4540	-6.18	0.22
Thornton Lands, TL96-1	AA26379	Intercalated	48564	42413	-2.86	4.2	0.2	3760	80	4390	3890	-6.86	0.21
Thornton Lands, TL96-1	AA26380	Intercalated	48564	42413	-3.8	4.2	0.2	5560	65	6480	6210	-8.18	0.21
South Ferryby	OXA7052	Intercalated	49840	42090	-0.4	4.2	0.4	3670	40	4140	3870	-4.6	0.41
Brigg	OXA7137	Intercalated	49940	40880	-6.34	4.2	0.2	4990	75	5910	5600	-10.72	0.21
Wilfholme Landing, WL971	AA34284	Intercalated	50617	44700	-99	-99	-99	3995	50	4640	4300	-2.2	0.21
Wilfholme Landing, WL971	AA34285	Intercalated	50617	44700	-99	-99	-99	4145	55	4830	4480	-2.2	0.7
Barrowhaven	SRR4897	Intercalated	50602	42307	1.77	3.6	0.4	2040	40	2120	1890	-2.03	0.41
Barrowhaven	SRR1374	Intercalated	50602	42307	1.43	3.6	0.4	2325	60	2700	2150	-2.17	0.41
Lockham, HMB12	AA22672	Intercalated	53932	41720	-3.3	3	0.2	5325	50	6270	5950	-6.1	0.21

Table A.2: Table of LOIS data (continued). -99 denotes missing data. Reference water level includes tidal correction.

Site	Lab code	Type	Location		Alt (m.O.D.)	RWL	IR	¹⁴ C	Age	Cal 95% range	RSL	Error
			x	y				Error	min	max		
Sand-le-Mere, SM95-5	AA23821	Intercalated	53210	43111	-0.81	3	0.2	2630	60	2860	-3.97	0.21
Lockham, HMB12	AA23891	Intercalated	53932	41720	-1.27	3	0.7	4235	60	4940	-3.57	0.7
Lockham, HMB12	AA23890	Intercalated	53932	41720	-1.16	3	0.7	4040	65	4800	-3.46	0.7
Barrowhaven	SRR1373	Intercalated	50602	42307	2.07	3.6	0.2	1080	40	1080	-1.33	0.21
Union Dock-3, Grimsby	SRR4746	Intercalated	52730	41060	-8.78	3.1	0.2	6645	45	7610	-11.68	0.21
Union Dock-2, Grimsby	SRR4745	Intercalated	52730	41060	-7.78	3.1	0.2	5900	45	6850	-11.05	0.21
Union Dock-1, Grimsby	SRR4744	Intercalated	52730	41060	-7.68	3.1	0.2	5665	45	6600	-10.58	0.21
Thirtle Bridge	GU5489	Intercalated	53090	43040	0.79	2.88	0.2	2420	50	2710	-2.22	0.27
Roos Drain West	GU5483	Intercalated	52720	42840	-2.46	2.88	0.2	5010	70	5910	-5.47	0.27
Thirtle Bridge	GU5480	Intercalated	53090	43040	0.66	2.88	0.2	2370	50	2710	-2.35	0.27
Thirtle Bridge	GU5490	Intercalated	53090	43040	0.79	2.88	0.2	2590	50	2790	-2.22	0.27
Roos Drain West	GU5478	Intercalated	52720	42840	-2.61	2.88	0.2	5200	60	6170	-5.62	0.27
Roos Drain West	GU5477	Intercalated	52720	42840	-2.61	2.88	0.2	5080	60	5940	-5.62	0.27
Halsham Carrs	GU5476	Intercalated	52650	42740	-1.34	2.88	0.2	4150	50	4830	-4.35	0.27
Halsham Carrs	GU5475	Intercalated	52650	42740	-1.34	2.88	0.2	4240	50	4890	-4.35	0.27
Thirtle Bridge	GU5479	Intercalated	53090	43040	0.66	2.88	0.2	2250	50	2350	-2.35	0.27
Roos Drain West	GU5484	Intercalated	52720	42840	-2.46	2.88	0.2	5140	90	6170	-5.47	0.27

Table A.3: Table of LOIS data (continued). -99 denotes missing data. Reference water level includes tidal correction.

Site	Lab code	Type	Location		Alt (m.O.D.)	RWL	IR	¹⁴ C	Age		Cal 95% range		RSL	Error
			x	y					Error	min	max			
South Ferriby	GU5706	Limiting	49840	42090	-1.12	4.2	1.58	3890	60	4500	4120	-4.94	1.58	
South Ferriby	OXA7054	Limiting	49820	42090	-2.99	4.2	1.58	4960	40	5860	5600	-6.81	1.58	
South Ferriby	OXA7053	Limiting	49840	42090	0.53	4.2	1.58	3310	40	3640	3440	-3.29	1.58	
Whitton Ness, HMB13	AA22674	Limiting	49232	42426	-1.88	4.2	1.58	3910	45	4490	4170	-5.7	1.58	
Whitton Ness, HMB13	AA23436	Limiting	49232	42426	-2.6	4.2	1.58	4830	70	5710	5340	-6.42	1.58	
South Ferriby	OXA7056	Limiting	49780	42070	-5.87	4.2	1.58	5440	45	6310	6040	-9.69	1.58	
South Ferriby	OXA7055	Limiting	49830	42090	-2.07	4.2	1.58	4650	50	5570	5150	-5.89	1.58	
Winterton Ings, WTI95-6	CAM41318	Limiting	49580	42003	-3.84	4.2	1.58	5990	60	6990	6670	-7.66	1.58	
Winterton Ings, WTI95-3	CAM41322	Limiting	49541	42002	-1.36	4.2	1.58	4100	50	4810	4450	-5.18	1.58	
South Ferriby	OXA7057	Limiting	49800	42080	-3.32	4.2	1.58	6000	50	6980	6690	-7.14	1.58	
Brigg	GU5699	Limiting	49940	40880	-2.67	4.2	1.58	5040	170	6200	5350	-6.49	1.58	
South Ferriby	OXA7067	Limiting	49840	42090	2.01	4.2	1.58	2690	70	2950	2590	-1.81	1.58	
Redbourne Hayes	GU5703	Limiting	50080	40010	-0.05	4.2	1.58	3590	50	4070	3720	-3.87	1.58	
Redbourne Hayes	GU5702	Limiting	50080	40010	0.67	4.2	1.58	3370	60	3810	3470	-3.15	1.58	
Redbourne Hayes	GU5701	Limiting	50080	40010	-0.26	4.2	1.58	3940	70	4580	4160	-4.08	1.58	
Winterton Carrs, HMB7	AA34280	Limiting	49644	41892	-99	-99	-99	5800	55	6730	6450	-10.53	2.11	
Garthorpe, HMB16	AA25585	Limiting	48594	41923	-11.43	4.2	1.58	7265	60	8180	7950	-15.26	1.58	
Ousefleet, HMB20	AA25589	Limiting	48225	42315	-12.08	4.2	1.58	7205	85	8180	7480	-15.9	1.58	
Ousefleet, HMB20	AA25590	Limiting	48225	42315	-12.85	4.2	1.58	7190	60	8160	7870	-16.67	1.58	
Ousefleet, HMB20	AA25591	Limiting	48225	42315	-12.92	4.2	1.58	8240	110	9500	8850	-16.74	1.58	
Ousefleet, HMB20	AA34283	Limiting	48225	42315	-99	-99	-99	8930	70	10210	9760	-17.38	1.58	
South Ferriby	GU5707	Limiting	49830	42090	-2.12	4.2	1.58	4700	80	5590	5130	-5.94	1.58	
Redbourne Hayes	OXA7065	Limiting	50080	40010	0.06	4.2	1.58	3595	50	4070	3730	-3.76	1.58	
Brigg	GU5700	Limiting	49940	40880	-2.46	4.2	1.58	4730	100	5650	5100	-6.28	1.58	
Redbourne Hayes	OXA7064	Limiting	50080	40010	-1.92	4.2	1.58	4494	31	5290	4990	-5.74	1.58	
Brigg	OXA7136	Limiting	49940	40880	-4.61	4.2	1.58	6170	90	7260	6800	-8.43	1.58	
Brigg	OXA7091	Limiting	49940	40880	-0.92	4.2	1.58	3940	45	4520	4240	-4.74	1.58	
Brigg	OXA7090	Limiting	49940	40880	-1.47	4.2	1.58	4300	50	5030	4680	-5.29	1.58	

Table A.4: Table of LOIS data (continued). -99 denotes missing data. Reference water level includes tidal correction.

Site	Lab code	Type	Location		Alt (m.O.D.)	RWL	IR	¹⁴ C	Age		Cal 95% range		RSL	Error
			x	y					Error	min	max			
South Ferriby	OXA7066	Limiting	49840	42090	1.33	4.2	1.58	2635	45	2850	2560	-2.49	1.58	
Garthorpe, HMB16	AA25586	Limiting	48594	41923	-11.65	4.2	1.58	7745	60	8640	8410	-15.48	1.58	
Hull	IGS99	Limiting	51003	42851	-11.55	3.6	1.7	6970	100	7970	7610	-14.77	1.7	
Stoneferry, Yorks	IGS97	Limiting	51036	43214	-4	3.6	0.7	5240	100	6300	5750	-6.31	2.11	
South Farm, HMB8	AA25582	Limiting	52572	41753	-11.95	4.2	1.58	8555	65	9710	9430	-15.77	1.58	
South Farm, HMB8	AA25581	Limiting	52572	41753	-11.67	4.2	1.58	7145	60	8110	7800	-15.49	1.58	
Kilnsea, HMB1	AA25575	Limiting	54212	41456	-0.87	3	1.43	4130	50	4820	4460	-3.5	1.43	
Union Dock-4, Grimsby	SRR4747	Limiting	52730	41060	-8.91	3.1	1.48	8170	45	9300	9020	-11.63	1.48	
Hull	IGS100	Limiting	51003	42851	-9.73	3.6	1.7	6890	100	7930	7560	-12.95	1.7	

Table A.5: Table of LOIS data (continued). -99 denotes missing data. Reference water level includes tidal correction.

Appendix B

Archaeological Data

This appendix contains the archaeological data created especially for use in this thesis, taken from the Humber Wetlands Project set of books (van der Noort and Davies, 1993; van der Noort and Ellis, 1995, 1997, 1998, 1999, 2000; van der Noort *et al.*, 2001). Lab codes are given for data points with radiocarbon dates. For the dates of the periods concerned, see table 4.1.

Site	Lab code	Location		Alt (m.O.D.)	RWL	Period	¹⁴ C Age		Cal range		RSL	Error
		x	y				Error	min	max			
Brigham-6		506050	454350	3.3	2.725	657-740AD			0.575	0.359		
Weel-2	OxA-8954 OxA-9010	507000	438100	0.778	2.434	Bronze Age	3420 3365	45 45	3820 3720	3510 3470	-1.656	0.415
Marsh Chapel-2		535800	398300	2.22	3.165	AD1000-1200			-0.945	0.396		
Willoughby-13	GU-5813 GU-5814 GU-5815 GU-5816	549300	372400	2.81	1.841	Late Mesolithic to Roman	3220 3210 3180 3010	50 50 50 80	3570 3560 3540 3380	3350 3330 3270 2950	0.969	0.333
Ingoldmells-6		555500	368200	1.08	2.039	600BC-400AD			-0.959	0.241		
Ingoldmells-12		555900	368200	1.79	2.037	600BC-400AD			-0.247	0.251		
Ferriby-2	Q-715 Q-3044 Q-3123	499100	425200	-0.75	1.456	Bronze Age	3120 3095 3120	105 40 45	3600 3400 3450	3000 3180 3210	-2.206	0.313
Adlingfleet-2		484800	422100	0.878	2.052	Roman			-1.174	0.414		
South Ferriby-2	GU-5705	497100	421200	1.18	1.753	Roman	1960	50	2040	1770	-0.573	0.379
South Ferriby-2	GU-5704	497100	421200	1.37	1.753	Roman	1840	60	1910	1610	-0.383	0.408
Brigg raft	Q-1255 Q-1256 Q-1257 Q-1258 Q-1263	499300	407620	0.915	0.35	Bronze Age	2655 2605 2595 2670 2570	50 50 50 75 50	2870 2840 2810 2950 2770	2620 2490 2490 2510 2430	0.565	0.393
Brigg trackway	Q-77	499230	407490	-0.435	1.747	Bronze Age	2552	120	2900	2350	-2.182	0.463
Brigg boat	Q-78	499750	407370	0.015	0.348	Bronze Age	2784	100	3200	2740	-0.333	0.401
Wheldrake-2		470800	444800	9.044	2.178	Roman - AD230-370			6.866	0.421		
Kelk-6		510800	460100	7.15	3.258	Iron Age			3.892	0.442		

Table B.1: Table of archaeological data. Reference water level includes tidal correction.

Appendix C

Estimating HAT for archaeological index and limiting points

C.1 Estimation of modern HAT and MHWS at non-standard ports

The proposed method is to estimate the ratio of HAT to MHWS as a function of latitude and longitude. This first requires estimates of the ratio and, hence, modern estimates of HAT and MHWS. Ideally these would be obtained from the full 18.6 year metonic cycle.

The Admiralty Tide Tables used in this section (Hydrographer to the Navy, 1982, 1988, 1993, 1995) give two possible methods for tidal prediction. The first is an interpolation approach which is too simple

to be used here because it does not allow the possibility that there may be tides higher than MHWS. The second approach is known as the 'simplified method of tidal prediction' or NP159. For this project, data were obtained for the 2 major and 10 minor ports around the Humber; Spurn Head, Immingham, Skegness, Inner Dowsing Light Tower, Bull Sand Fort, Hull (King George Dock), Hull (Albert Dock), Humber Bridge, Burton Stather, Blacktoft and Goole.

The full details of how the NP159 method are not included here; the reader is referred to Pugh (1987). However, a brief description of its working are included to illuminate the findings. The sea-level height (over short periods) at time t can be represented as a time series:

$$X(t) = Z_0(t) + T(t) + S(t), \quad (\text{C.1})$$

where Z_0 is mean sea level, T is the tidal variation and S is any surge component. This can further be extended using ideas from Fourier analysis:

$$\begin{aligned} X(t) &= Z_0 + \sum_{m=1}^{M/2} A_m \cos m\sigma t + \sum_{m=1}^{M/2} B_m \sin m\sigma t \\ &= Z_0 + \sum_{m=1}^{M/2} H_m \cos(m\sigma t - g_m), \end{aligned} \quad (\text{C.2})$$

with

$$H_m = (A_m^2 + B_m^2)^{\frac{1}{2}}$$

and

$$g_m = \arctan(B_m/A_m).$$

The coefficients A_m and B_m can be evaluated from sampled data over the M chosen harmonics. H_m and g_m are known as the amplitude and phase lag of the m th harmonic. The NP159 method uses four such tidal harmonics, known as M_2 , S_2 , K_1 and O_1 which represent the lunar semi-diurnal tide, the solar semi-diurnal tide, and two representing the lunar declination effect respectively. If required, two further harmonics are used to evaluate any bottom-friction effect. Each of these harmonics has an angle and speed, depending on their location, and the date and time required. The NP159 method lists all of these parts allowing the user to simply sum together the appropriate harmonic constituents.

As mentioned previously, the ideal situation would be to run the NP159 method for 18.6 years for each of the locations. However, this presented a number of problems:

- 19 years of tidal angles and factors would require a very long period of data input.
- The most recent tide tables available were from 1995, thus 1977 would be required to complete the cycle. However, the data are presented differently before 1982 rendering the NP159 method unusable.
- Periods of harmonic data for ports were missing from the tide tables, as was some of the individual port data.

To circumvent these problems, the modelling was undertaken within each individual year for whichever ports were available for four years; 1982, 1988, 1993 and 1995. This was not an ideal solution, but it

was felt that the time periods were of reasonable spread to provide satisfactory solutions.

As an example of the NP159 method, a sample of 300 hours of predictions during a spring-neap cycle, created using this method for Goole in 1995 is shown in figure C.1. From these data sets, levels of HAT and MHWS were derived. For HAT, this is simply the maximum over the entire set of calculated values. For MHWS, this can be calculated as the sum of the M2 and S2 tidal harmonics (Pugh, 1987).




Figure C.1: 300 hours of tidal elevations for Goole (Latitude N53:42:13, Longitude W0:52:50) in 1995. Note that the heights are given in metres Chart Datum.

Output from the NP159 program is shown in table C.1. Ratios are calculated directly from this data.

C.2 Modelling the ratio of HAT to MHWS

Once the ratio has been obtained, an interpolation model is required to estimate the ratio for a given location. The data given in table C.1 provides 33 data points for such a purpose. A number of different models were tested:

1. A standard linear regression in latitude and longitude,
2. A linear regression model with errors proportional to latitude,
3. A Gaussian Random Field (GRF) non-parametric regression.

The first model gave excellent results in terms of adjusted R^2 value (0.9903) and small standard error (0.1414), when only latitude and longitude were included in the model (with no constant part). However, when the distribution of residuals was examined, some small anomalies were uncovered. There was some evidence of heteroscedasticity in latitude plots and possible sequential curvature in longitude plots. When plotted spatially, there was a slight indication of possible spatial dependence.

The second model included a heteroscedastic variance term as a function of latitude. This introduction of this term had negligible effect. The third model gave no clear indication of spatial correlation once latitude and longitude were included in a linear fashion. In view of the very small size residuals and in spite of the evidence of structure in them, the initial model was used as adequate for the current set of

data. The model can simply be written as:

$$\frac{\text{HAT}}{\text{MHWS}} = 0.025 \times \text{Lat} - 0.385 \times \text{Long} + e_i.$$

$$e_i \sim N(0, 0.141^2)$$

C.3 Using the estimated ratio to predict HAT at other locations

Predictive values of the ratio can be found using the predictive distribution:

$$Y^* \sim N(X^*\beta, \hat{\sigma}^2(I + X^*(X^T X)^{-1} X^{*T}))$$

where X is the design matrix containing the column of latitudes and longitudes supplied through the tidal modelling, Y^* is the vector of unknown desired ratios, X^* is the design matrix of desired latitudes and longitudes and $\hat{\sigma}^2$ is the estimated standard error.

The predictive distribution of the estimated ratio can be applied to the MHWS level at any point from any date in the tidal predictions to product an estimated value for HAT at that location. The distribution of this estimate is a product Normal distribution (as both the predictive distribution of the ratio and the distribution of the tidal estimate of MHWS are Normal). Thus the standard error of the estimated HAT was calculated using Monte Carlo estimation for simplicity.

One complication remains: the tidal predictions from Shennan *et al.* (2000a) do not give tidal values of MHWS at locations beyond the modern or prehistoric shoreline. However, many of the archaeological data points are located some distance away from this point. The problem was overcome by using a 5% trimmed mean on the closest (in Euclidean distance) set of tidal heights given. A trimmed mean was used as some data cells contain extreme results. Once the trimmed mean had been calculated, the value was linearly interpolated between the two nearest thousand year data sheets. This was done to match the mean date of the archaeological point of interest.

The result of these calculations, the tidal standard error associated with each of the new archaeological sea-level data points is given in table 3.1.

Location	HAT (metres O.D.)				MHWS (metres O.D.)			
	1982	1988	1993	1995	1982	1988	1993	1995
Skegness	0	3.86	3.89	3.58	0	2.88	2.88	2.88
Inner Dowsing Light Tower	3.13	3.43	3.45	3.16	2.65	2.65	2.65	2.65
Spurn Head	0	0	3.73	3.46	0	0	2.85	2.85
Bull Sand Fort	3.44	3.62	3.64	3.37	2.85	2.79	2.79	2.83
Immingham	3.62	3.86	4.01	3.7	3.01	3	3.03	3.05
Hull (King George Dock)	3.98	4.3	4.34	3.98	3.2	3.2	3.2	3.18
Hull (Albert Dock)	0	0	0	4.1	0	0	0	3.2
Humber Bridge	0	0	0	4.34	0	0	0	3.24
Burton Stather	0	4.4	4.45	4.04	0	2.46	2.46	2.46
Blacktoft	0	4.47	4.46	4.19	0	2.69	2.69	2.65
Goole	3.96	4.26	4.45	4.2	2.29	2.29	2.29	2.42

Table C.1: Table of HAT and MHWS for tidal stations around the Humber. A zero denotes missing data.

Appendix D

WinBUGS code

The WinBUGS code used in sections 6.2.1 and 6.3.1 is included here.

The temporal only model:

```
model {
  for (i in 1:N) {
    mu[i] <- beta1*dates[i]+beta2*dates[i]*dates[i]
    Y[i] ~ dnorm(mean[i],Tau[i])
    Tau[i] <- 1/Sigma[i]
    mean[i] <- mu[i]+w[i]+gamma[i]
    zeros[i] <- 0
  }
  for(i in 1:NBIP) {
    gamma[i] ~ dnorm(0,10000)
  }
  for(i in 1:NIIP) {
    gamma[NBIP+i] ~ dnorm(BhIIP[i],PsiIIP[i])
  }
}
```

```

    BhIIP[i] <- -hIIP
    PsiIIP[i] <- Kappa1inv
}
for(i in 1:NLP) {
    gamma[NBIP+NIIP+i] ~ dnorm(BhLP[i],PsiLP[i])
    BhLP[i] <- hLP
    PsiLP[i] <- Kappa2inv
}

w[1:N] ~ spatial.exp(zeros[],dates[],zeros[],tausig,phi1,phi2

hIIP ~ dunif(0,20)
hLP ~ dunif(0,20)

Kappa1inv ~ dgamma(0.01,0.01)
Kappa2inv ~ dgamma(0.01,0.01)

Kappa1 <- 1/Kappa1inv
Kappa2 <- 1/Kappa2inv

beta1 ~ dnorm(0,0.01)
beta2 ~ dnorm(0,0.01)

phi1 ~ dunif(0,10)
phi2 ~ dunif(0.05,1.95)

```

```

    tausig ~ dgamma(0.01,0.01)
    sigmasq <- 1/tausig
}

```

The spatial model:

```

model {
  for (i in 1:N) {
    mu[i] <- b1*dates[i]+b2*dates[i]*dates[i]+b3*OSX[i]+b4*OSY[i]
    Y[i] ~ dnorm(mean[i],Tau[i])
    Tau[i] <- 1/Sigma[i]
    mean[i] <- mu[i]+w[i]+gamma[i]
    zeros[i] <- 0
  }
  for(i in 1:NBIP) {
    gamma[i] ~ dnorm(0,10000)
  }
  for(i in 1:NIIP) {
    gamma[NBIP+i] ~ dnorm(BhIIP[i],PsiIIP[i])
    BhIIP[i] <- -hIIP
    PsiIIP[i] <- Kappalinv
  }
  for(i in 1:NLP) {
    gamma[NBIP+NIIP+i] ~ dnorm(BhLP[i],PsiLP[i])
    BhLP[i] <- hLP
    PsiLP[i] <- Kappa2inv
  }
}

```

```
w[1:N] ~ spatial.exp(zeros[], dates[], zeros[], tausig, phi1, phi2
```

```
hIIP ~ dunif(0,20)
```

```
hLP ~ dunif(0,20)
```

```
Kappa1inv ~ dgamma(0.01,0.01)
```

```
Kappa2inv ~ dgamma(0.01,0.01)
```

```
Kappa1 <- 1/Kappa1inv
```

```
Kappa2 <- 1/Kappa2inv
```

```
b1 ~ dnorm(0,0.01)
```

```
b2 ~ dnorm(0,0.01)
```

```
b3 ~ dnorm(0,0.01)
```

```
b4 ~ dnorm(0,0.01)
```

```
phi1 ~ dunif(0,10)
```

```
phi2 ~ dunif(0.05,1.95)
```

```
tausig ~ dgamma(0.01,0.01)
```

```
sigmasq <- 1/tausig
```

```
}
```

Bibliography

- Aguilar, G. P., Litton, C. D. and O'Hagan, A. (2002). A new piecewise linear radiocarbon calibration curve with more realistic variance. *Radiocarbon*, **44**, 195–212.
- Aitken, M. (1994). *Science-based Dating in Archaeology*. Longman, London.
- Akaike, H. (1973). Information theory and an extensions of the maximum likelihood principle. In B. Petrov and F. Csáki (eds.), *Proceedings of the 2nd international symposium on information theory*, Akadémiai Kiadó, 267–281.
- Allen, J. (1999). Geological impacts on coastal wetland landscapes: some general effects of sediment autocompaction in the Holocene of northwest Europe. *The Holocene*, **9**, 1–12.
- Austin, R. (1991). Modelling Holocene tides on the NW European continental shelf. *Terra Nova*, **3**, 276–288.
- Bailey, G. (1978). Shell middens as indicators of postglacial economies: a territorial perspective. In P. Mellars (ed.), *The Early Postglacial Settlement of Northern Europe*, Duckworth, London.
- Baldini, J., McDermott, F. and Fairchild, I. (2002). Structure of the

- 8200-year cold event revealed by a speleothem trace element record. *Science*, **296**, 2203–2206.
- Banerjee, S., Carlin, B. and Gelfand, A. (2004). *Hierarchical modelling and analysis for spatial data*. Chapman and Hall.
- Bard, E., Hamelin, B. and Fairbanks, R. (1990). U-Th ages obtained by mass spectrometry in corals from Barbados: sea level during the past 130,000 years. *Nature*, **346**, 456–458.
- Bateman, I., Bateman, S., Brown, D., Doktor, P., Karas, J. H. W., Maher, A. and Turner, R. (1991). *Economic Appraisal of the Consequences of Climate-Induced Sea-level Rise: A Case Study of East Anglia*. University of East Anglia.
- Bowman, S. (1990). *Radiocarbon Dating*. British Museum Publications Ltd, London.
- Brooks, S. and Gelman, A. (1998). General Methods for Monitoring Convergence of Iterative Simulations. *Journal of computational and graphical statistics*, **7**, 434–455.
- Brown, J., Colling, A., Park, D., Phillips, J., Rothery, D. and Wright, J. (1989). *Waves, tides and shallow water processes*. Pergamon.
- Buck, C. and Blackwell, P. (2004). Formal statistical models for estimating radiocarbon calibration curves. *Radiocarbon*, **46**, 1093–1102.
- Buck, C. and Millard, A. (eds.) (2004). *Tools for Constructing Chronologies: crossing disciplinary boundaries*. Springer-Verlag.

- Buck, C. E., Cavanagh, W. G. and Litton, C. D. (1996). *Bayesian Approach to Interpreting Archaeological Data*. John Wiley and Sons, Chichester.
- Christen, J. A. (1994). *Bayesian interpretation of ^{14}C results*. Ph.D. thesis, University of Nottingham, Nottingham, UK.
- Cooke, P. (1979). Statistical inference for bounds of random variables. *Biometrika*, **66**, 367–374.
- Craig, H. (1953). The geochemistry of the stable carbon isotopes. *Geochimica et Cosmochimica Acta*, **3**, 53–92.
- Cramér, H. and Leadbetter, M. (1967). *Stationary and Related Stochastic Processes*. Wiley.
- Cressie, N. (1991). *Statistics for Spatial Data*. John Wiley and Sons.
- Damon, P., Kaimei, D., Kocharov, G., Mikheeva, I. and Peristykh, A. (1995). Radiocarbon Production by the Gamma-Ray Component of Supernova Explosions. *Radiocarbon*, **37**, 599–604.
- de Boor, C. (1978). *A practical guide to splines*. Springer-Verlag.
- Devoy, R. (1982). Analysis of the geological evidence for Holocene sea-level movements in South East England. *Proceedings of the Geologists Association*, **93**, 65–90.
- Devoy, R. (1987). Hydrocarbon exploration and biostratigraphy: the application of sea-level studies. In R. Devoy (ed.), *Sea Surface Studies*, Croom Helm.
- DeVries, H. and Barendesen, G. (1954). Measurements of age by the carbon-14 technique. *Nature*, **174**, 1138–1141.

- Ecker, M. and Gelfand, A. (1999). Bayesian modelling and inference for geometrically anisotropic spatial data. *Mathematical Geology*, **31**, 67–83.
- Esper, J., Cook, E. and Schweingruber, F. (2002). Low-frequency signals in long tree-ring chronologies for reconstructing past temperature variability. *Science*, **295**, 2250–2253.
- Fairbridge, R. (1961). Eustatic changes in sea level. *Physics and chemistry of the earth*, **5**, 99–185.
- Firth, C., Smith, D., Hansom, J. and Pearson, S. (1995). Holocene spit development on a regressive shoreline, Dornoch Firth, Scotland. *Marine Geology*, **124**, 203–214.
- Fischer, A. (1995). *Man and Sea in the Mesolithic: coastal settlement above and below present sea level*. Oxbow, London.
- Fontes, J. (1992). Chemical and isotopic constraints on ^{14}C dating of groundwater. In R. Taylor, A. Long and R. Kra (eds.), *Radiocarbon After Four Decades*, Springer-Verlag, 242–261.
- Gaunt, G. (1994). *Geology of the country around Goole, Doncaster and the Isle of Axholme. Memoir for one-inch sheets 79 and 88 (England and Wales)*. HMSO.
- Gaunt, G. and Tooley, M. (1974). Evidence for Flandrian sea-level changes in the Humber estuary and adjacent areas. *Bulletin of the Institute of Geological Sciences*, **48**, 25–41.
- Gehrels, W., Belknap, D., Pearce, B. and Gong, B. (1995). Modeling the contribution of M2 tidal amplification to the Holocene rise of mean

- Ecker, M. and Gelfand, A. (1999). Bayesian modelling and inference for geometrically anisotropic spatial data. *Mathematical Geology*, **31**, 67–83.
- Esper, J., Cook, E. and Schweingruber, F. (2002). Low-frequency signals in long tree-ring chronologies for reconstructing past temperature variability. *Science*, **295**, 2250–2253.
- Fairbridge, R. (1961). Eustatic changes in sea level. *Physics and chemistry of the earth*, **5**, 99–185.
- Firth, C., Smith, D., Hansom, J. and Pearson, S. (1995). Holocene spit development on a regressive shoreline, Dornoch Firth, Scotland. *Marine Geology*, **124**, 203–214.
- Fischer, A. (1995). *Man and Sea in the Mesolithic: coastal settlement above and below present sea level*. Oxbow, London.
- Fontes, J. (1992). Chemical and isotopic constraints on ^{14}C dating of groundwater. In R. Taylor, A. Long and R. Kra (eds.), *Radiocarbon After Four Decades*, Springer-Verlag, 242–261.
- Gaunt, G. (1994). *Geology of the country around Goole, Doncaster and the Isle of Axholme. Memoir for one-inch sheets 79 and 88 (England and Wales)*. HMSO.
- Gaunt, G. and Tooley, M. (1974). Evidence for Flandrian sea-level changes in the Humber estuary and adjacent areas. *Bulletin of the Institute of Geological Sciences*, **48**, 25–41.
- Gehrels, W., Belknap, D., Pearce, B. and Gong, B. (1995). Modeling the contribution of M2 tidal amplification to the Holocene rise of mean

- high water in the Gulf of Maine and the Bay of Fundy. *Marine Geology*, **124**, 71–85.
- Gelman, A. and Rubin, D. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, **2**, 457–472.
- Geweke, J. (1991). Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In A. D. J.M. Bernardo, J.O. Berger and A. Smith (eds.), *Bayesian Statistics 4*, Oxford University Press, 169–194.
- Gilks, W., Richardson, S. and Spiegelhalter, D. (1996). *Markov chain Monte Carlo in practice*. Chapman and Hall.
- Hall, P. (1982). On estimating the endpoint of a distribution. *Annals of Statistics*, **2**, 556–568.
- Harkness, D. (1983). The extent of the natural ^{14}C deficiency in the coastal environment of the United Kingdom. *Journal of the European Study Group on Physical, Chemical and Mathematical Techniques Applied to Archaeology*, **PACT8 (IV.9)**, 351–364.
- Hastie, T. and Tibshirani, R. (1995). *Generalized Additive Models*. Chapman and Hall.
- Hastings, W. (1970). Monte Carlo sampling using Markov Chains and their applications. *Biometrika*, **57**, 97–109.
- Heidelberger, P. and Welch, P. (1983). Simulation run length control in the presence of an initial transient. *Operations Research*, **31**, 1109–1144.

- Henderson, G. (2005). Coral clues to rapid sea-level change. *Science*, **308**, 361–362.
- Hinton, A. (1996). Tides in the northeast atlantic: considerations for modelling water depth changes. *Quaternary Science Reviews*, **15**, 873–894.
- Hopley, D. (1983). Deformation of the North Queensland continental shelf in the Late Quaternary. In D. Smith and A. Dawson (eds.), *Shorelines and Isostasy*, Academic Press, 347–368.
- Horton, B., Edwards, R. and Lloyd, J. (2000). Implications of a microfossil-based transfer function in Holocene sea-level studies. In I. Shennan and J. Andrews (eds.), *Holocene Land-Ocean Interaction and Environment Change around the North Sea*, The Geological Society Special Publication, 166.
- Horton, B., Innes, J., Shennan, I., Lloyd, J. and McArthur, J. (2004). Holocene coastal change in East Norfolk, UK: palaeoenvironmental data from Somerton and Winterton Holmes, near Horsey. *Proceedings of the Geologists' Association*, **115**, 209–220.
- Host, G. (1999). Kriging by local polynomials. *Computational statistics and data analysis*, **29**, 295–312.
- Hutton, J. (1788). An Investigation of the Laws Observable in the Composition, Dissolution and Restoration of the Globe. *Transactions of the Royal Society of Edinburgh*, **1**, 209–304.
- Hydrographer to the Navy (1982). *Admiralty Tide Tables: European Waters including the Mediterranean Sea*.

- Hydrographer to the Navy (1988). *Admiralty Tide Tables: European Waters including the Mediterranean Sea.*
- Hydrographer to the Navy (1993). *Admiralty Tide Tables: European Waters including the Mediterranean Sea.*
- Hydrographer to the Navy (1995). *Admiralty Tide Tables: European Waters including the Mediterranean Sea.*
- Jarvie, H., Neal, C. and Robson, A. (1997). The Geography of the Humber catchment. *The science of the total environment*, **194/195**, 87-89.
- Kaiser, H. (1958). The varimax criterion for analytic rotation in factor analysis. *Psychometrika*, **23**, 187-200.
- Kennedy, M. and O'Hagan, A. (2001). Bayesian calibration of computer models (with discussion). *Journal of the Royal Statistical Society, Series B*, **63**, 425-464.
- Krzanowski, W. (1988). *Principles of multivariate analysis.* Oxford University Press.
- Lambeck, K. (1988). *Geophysical Geodesy.* Clarendon Press.
- Lambeck, K. (1993). Glacial Rebound of the British Isles - I. Preliminary model results. *Geophysical Journal International*, **115**, 941-959.
- Lambeck, K. and Chappell, J. (2001). Sea-level change through through the last glacial cycle. *Nature*, **292**, 679-686.
- Libby, W., Anderson, E. and Arnold, J. (1949). Age determinations by radiocarbon content: Worldwide assay of natural radiocarbons. *Science*, **109**, 227-228.

- Long, A., Innes, J., Kirby, J., Lloyd, J., Rutherford, M., Shennan, I. and Tooley, M. J. (1998). Holocene sea-level change and coastal evolution in the Humber estuary, eastern England: an assessment of rapid coastal change. *The Holocene*, **8**, 229–247.
- Long, A., Scaife, R. and Edwards, R. (1999). Pine pollen in intertidal sediments from Poole Harbour, U.K.; implications for late-Holocene sediment accretion rates and sea-level rise. *Quaternary International*, **55**, 3–16.
- Maasch, K., Mayewski, P., Rohling, E., Stager, J., Karlen, W., Meeker, L. and Meyerson, E. (2005). A 2000-year context for modern climate change. *Geografiska Annaler*, **87A**, 7–15.
- Mardia, K., Kent, J. and Bibby, J. (1979). *Multivariate analysis*. Academic Press.
- Matérn, B. (1960). Spatial Variation. *Meddelanden från Statens Skogsforskningsinstitut*, **49**, 100–144.
- Matthews, J. and Briffa, K. (2005). The 'little ice age': re-evaluation of an evolving concept. *Geografiska Annaler*, **87A**, 17–36.
- McDermott, F., Matthey, D. and Hawkesworth, C. (2001). Centennial-scale Holocene climate variability revealed by a high-resolution speleothem $\delta^{18}\text{O}$ record from SW Ireland. *Science*, **294**, 1328–1331.
- McGrail, S. (1981). The Brigg 'Raft' and her prehistoric environment. *British Archaeological Report, International Series*, **89**.
- Metcalfe, S., Ellis, S., Horton, B., Innes, J., MacArthur, J., Mitlehner, A., Parkes, A., Pethick, J., Rees, J., Ridgway, J., Rutherford, M.,

- Shennan, I. and Tooley, M. (2000). The Holocene evolution of the Humber Estuary: reconstructing change in a dynamic environment. In I. Shennan and J. Andrews (eds.), *Holocene Land-Ocean Interaction and Environmental change around the North Sea*, Geological Society of London, 97–118.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. and Teller, E. (1953). Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, **21**, 1087–1092.
- Mitrovica, J. and Milne, G. (2003). On post-glacial sea level: I. General theory. *Geophysical Journal International*, **154**, 253–267.
- Mitrovica, J., Tamisiea, M., Davis, J. and Milne, G. (2001). Recent mass balance of polar ice sheets inferred from patterns of global sea-level change. *Nature*, **409**, 1026–1029.
- Morner, N. (1987). Pre-Quaternary long-term changes in sea level. In R. Devoy (ed.), *Sea Surface Studies*, Croom Helm, 242–263.
- Oakley, J. (2002). Eliciting Gaussian process priors for complex computer codes. *The Statistician*, **51**, 81–97.
- Oguchi, T., Jarvie, H. and Neal, C. (2000). River water quality in the Humber catchment: an introduction using GIS-based mapping and analysis. *The science of the total environment*, **251**, 9–26.
- O'Hagan, A. (1978). Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society (Series B)*, **40**, 1–40.
- O'Hagan, A. (1992). Some Bayesian numerical analysis. In J. M.

- Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds.), *Bayesian Statistics 4*, Oxford University Press, 345–363.
- ONS, N. E. L. U. (2004). *Resident Population Estimates Mid-2003*.
URL <http://www.nelincs.gov.uk/council/populationstatistics/mid2003>
- ONS, P. E. U. (2002). *Registrar General's Mid-Year Estimate of Population Kingston upon Hull*.
URL http://www.hullcc.gov.uk/statistics/download/population_estimates
- Paciorek, C. and Schervish, M. (2004). Nonstationary covariance functions for Gaussian process regression. *Advances in Neural Information Processing Systems*, **16**, 273–280.
- Paul, M. and Barras, B. (1998). A geotechnical correction for post-depositional sediment compression: examples from the Forth valley, Scotland. *Journal of Quaternary Science*, **13**, 171–176.
- Peltier, W. (1974). The impulse response of a Maxwell earth. *Reviews of Geophysics and Space Science*, **12**, 649–669.
- Peltier, W. (1998a). Mechanisms of relative sea-level change and the geophysical responses to ice-water loading. In R. Devoy (ed.), *Sea surface studies*, Croom Helm, 165–197.
- Peltier, W. (1998b). Postglacial variations in the level of the sea: implications for climate dynamics and solid-earth geophysics. *Reviews of Geophysics*, **36**, 603–689.
- Peltier, W. (1998c). The inverse problem for mantle viscosity. *Inverse Problems*, **14**, 441–478.

- Peltier, W., Shennan, I., Drummond, R., and Horton, B. (2002). On the postglacial isostatic adjustment of the British Isles and the shallow viscoelastic structure of the Earth. *Geophysical Journal International*, **148**, 443–475.
- Pethick, J. (1990). The Humber Estuary. In *Humber Perspectives: a region through the ages*, Hull University Press, 54–70.
- Pirazzoli, P. (1996). *Sea-level changes: the last 20,000 years*. Wiley and Sons.
- Pugh, D. (1987). *Tides, surges and mean sea-level*. Wiley and Sons.
- R Development Core Team (2005). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
URL <http://www.R-project.org>
- Raftery, A. and Lewis, S. (1992). How many iterations in the Gibbs sampler? In J. Bernardo, J. Berger, A. Dawid and A. Smith (eds.), *Bayesian Statistics 4*, Oxford University Press.
- Ramsay, J. and Silverman, B. (1997). *Functional Data Analysis*. Springer-Verlag.
- Rees-Jones, J., Rink, W., Norris, R. and Litchfield, N. (2000). Optical luminescence dating of uplifted marine terraces along the Akatore Fault near Dunedin, South Island, New Zealand. *New Zealand Journal of Geology and Geophysics*, **43**, 419–424.
- Reimer, P. J., Baillie, M. G., Bard, E., Bayliss, A., Beck, W. W., Bertrand, C. J., Blackwell, P. G., Buck, C. E., Burr, G. S., Cutler,

- K. B., Damon, P. E., Edwards, L. L., Fairbanks, R. G., Friedrich, M., Guilderson, T. P., Hogg, A. G., Hughen, K. A., Kromer, B., McCormac, G., Manning, S., Ramsey, C. B., Reimer, R. W., Remmele, S., Southon, J. R., Stuiver, M., Talamo, S., Taylor, F., van der Plicht, J. v. d. and Weyhenmeyer, C. E. (2004). Intcal04 terrestrial radiocarbon age calibration, 0-26 cal bp. *Radiocarbon*, **46**, 1029–1058.
- Robson, D. and Whitlock, J. (1964). Estimation of a Truncation Point. *Biometrika*, **51**, 33–39.
- Rohling, E. and Pälike (2005). Centennial-scale climate cooling with a sudden cold event around 8,200 years ago. *Nature*, **434**, 975–979.
- Rossiter, J. (1962). Long-term variations in sea level. In N. Hill (ed.), *The Sea 1*, Interscience Publishers, 590–610.
- Sacks, J., Welch, W., Mitchell, T. and Wynn, H. (1989). Design and analysis of computer experiments. *Statistical Science*, **4**, 409–423.
- Schabenberger, O. and Gotway, C. (2005). *Statistical methods for spatial data analysis*. Chapman and Hall.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, **6**, 461–464.
- Scott, E. (2003). The fourth international radiocarbon inter-comparison. *Radiocarbon*, **45**, 135–408.
- Shennan, I. (1982). Interpretation of Flandrian sea-level data from the Fenland, England. *Proceedings of the Geologists' Association*, **93**, 53–63.

- Shennan, I. (1989). Holocene crustal movements and sea-level changes in Great Britain. *Journal of Quaternary Science*, **4**, 77–89.
- Shennan, I. (1994). Coastal Evolution. In M. Waller (ed.), *The Fenland Projects, Number 9: Flandrian Environmental Change in the Fenland*, East Anglian Archaeology, 47–84.
- Shennan, I. and Andrews, J. (eds.) (2000). *Holocene Land-Ocean Interaction and Environmental Change around the North Sea*. Geological Society of London.
- Shennan, I., Coulthard, T., Flather, R., Horton, B., Macklin, M., Rees, and Wright, M. (2003). Integration of shelf evolution and river basin models to simulate Holocene sediment dynamics of the Humber Estuary during periods of sea-level change and variations in the catchment sediment supply. *The Science of the total environment*, **314-316**, 737–754.
- Shennan, I. and Horton, B. (2002). Holocene land and sea level changes in Great Britain. *Journal of Quaternary Science*, **17**, 511–526.
- Shennan, I., Lambeck, K., Flather, R., Horton, B., MacArthur, J., Innes, J., Lloyd, J., Rutherford, M. and Wingfield, R. (2000a). Modelling western North Sea palaeogeographies and tidal changes during the Holocene. In I. Shennan and J. Andrews (eds.), *Holocene Land-Ocean Interaction and Environmental change around the North Sea*, Geological Society of London, 299–319.
- Shennan, I., Lambeck, K., Horton, B., Innes, J., Lloyd, J., McArthur, J. and Rutherford, M. (2000b). Holocene isostasy and relative sea-level changes on the east coast of England. In I. Shennan and J. Andrews

- (eds.), *Holocene Land-Ocean Interaction and Environmental change around the North Sea*, Geological Society of London, 275–298.
- Shennan, I., Tooley, M., Davis, M. and Haggart, B. (1983). Analysis and interpretation of Holocene sea-level data. *Nature*, **302**, 404–406.
- Shepard, F. (1963). Thirty-five thousand years of sea level. In T. Clements (ed.), *Essays in honour of K.O. Emery*, University of Southern California Press, 1–10.
- Shore, J., Bartley, D. and Harkness, D. (1995). Problems encountered with the ^{14}C dating of peat. *Quaternary Science Reviews*, **14**, 373–383.
- Sidell, J. (2002). *The fluvial regime of the river Thames*. Ph.D. thesis, University of Durham.
- Silverman, B. (1996). *Density Estimation*. Chapman and Hall.
- Smith, B. J. (2005). *boa: Bayesian Output Analysis Program (BOA) for MCMC*. R package version 1.1.4-3.
URL <http://www.public-health.uiowa.edu/boa>
- Spiegelhalter, D., Best, N., Carlin, B. and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society (Series B)*, **64**, 583–639.
- Stein, M. (1999). *Interpolation of Spatial Data*. Springer-Verlag.
- Stein, S. and Okal, E. (2005). Speed and size of the Sumatra earthquake. *Science*, **434**, 581–582.

- Stuiver, M. and Braziunas, T. (1993). Modeling atmospheric ^{14}C influences and ^{14}C ages of marine samples to 10,000 BC. *Radiocarbon*, **35**, 137–189.
- van de Plassche, O. (1982). Sea-level and water-level movements in the Netherlands during the Holocene. *Geological Survey of the Netherlands*, **36**, 1–93.
- van de Plassche, O. (1986). *Sea level research: A manual for the collection and evaluation of data*. Geo Books.
- van der Noort, R. and Davies, P. (1993). *Wetland heritage. An archaeological assessment of the Humber wetlands*. Hull University Press.
- van der Noort, R. and Ellis, S. (1995). *Wetland Heritage of Holderness*. Hull University Press.
- van der Noort, R. and Ellis, S. (1997). *Wetland Heritage of the Humberhead levels*. Hull University Press.
- van der Noort, R. and Ellis, S. (1998). *Wetland heritage of the Ancholme and Lower Trent valleys*. Hull University Press.
- van der Noort, R. and Ellis, S. (1999). *Wetland Heritage of the vale of York*. Hull University Press.
- van der Noort, R. and Ellis, S. (2000). *Wetland Heritage of the Hull valley*. Hull University Press.
- van der Noort, R., Ellis, S., Fenwick, H. and Lillie, M. (2001). *Wetland Heritage of the Lincolnshire Marsh*. Hull University Press.

- Waddelove, A. and Waddelove, E. (1990). *Archaeology and Research into Sea-level during the Roman Era: Towards a methodology Based on Highest Astronomical Tide*, vol. 21. Britannia.
- Ward, G. and Wilson, S. (1978). Procedures for comparing and combining radiocarbon age determinations: a critique. *Archaeometry*, **20**, 19-31.
- Warrick, R., Barrow, E. and Wigley, T. (eds.) (1993). *Climate and Sea Level Change: Observations, Projections and Implications*. Cambridge University Press.
- Warrick, R. and Oerlemans, J. (1990). Sea Level Rise. In G. J. J.T. Houghton and J. Ephraums (eds.), *Climate Change. The IPCC Scientific Assessment*, Cambridge University Press, 257-281.
- Wilkinson, H., de Boer, G. and Thunder, A. (1973). *A cartographic analysis of the changing bed of the Humber*. Tech. rep., University of Hull.
- Woodworth, P., Pugh, D., De Ronde, J., Warrick, R. and Hannah, J. (eds.) (1992). *Sea Level Changes: Determination and Effects*. American Geophysical Union.
- Woodworth, P., Shaw, S. and Blackman, D. (1991). Secular trends in mean tidal range around the British Isles and along the adjacent European coastline. *Geophysical Journal International*, **104**, 593-609.
- Wright, E., Hedges, R., Bayliss, A. and Van de Noort, R. (2001). New AMS radiocarbon dates for the North Ferriby boats - a contribution to

dating prehistoric seafaring in northwestern Europe. *Antiquity*, **75**, 726–734.

Yu, K. and Jones, M. (1998). Local linear quantile regression. *Journal of the American Statistical Association*, **93**, 228–237.

Yu, K., Lu, Z. and Stander, J. (2003). Quantile Regression: applications and current research areas. *The Statistician*, **52**, 331–350.

Yu, K. and Moyeed, R. (2001). Bayesian quantile regression. *Statistics and Probability Letters*, **54**, 437–447.

Zong, Y. and Horton, B. (1998). A Diatom-based transfer function as an aid in reconstructing the Quaternary history of sea-level movements in the UK. *Journal of Quaternary Science*, **14**, 153–167.