

Stochastic Modelling and Analysis of
Wildfowl *Anatidae*
Monitoring Data from the
Wetland Bird Survey

A THESIS SUBMITTED FOR
THE DEGREE OF DOCTOR OF PHILOSOPHY

Teresa Mary Frost

2010

SCHOOL OF MATHEMATICS, STATISTICS AND ACTUARIAL SCIENCE
UNIVERSITY OF KENT

For Timothy

ABSTRACT

British wetlands are used by many different resident and migratory wildfowl species for some or all of the winter months. The Wetland Bird Survey (WeBS) is a scheme set up with the objectives of assessing the size of waterbird populations, determining trends in numbers and distribution and assessing the importance of individual sites for waterbirds as part of the requirements of international conservation conventions and directives. Over three thousand sites around Britain are included in the survey, and volunteers have been undertaking counts of wildfowl species for over forty years. One of the key features of WeBS is that sites are surveyed at monthly intervals over the winter season.

The analyses for the project used data for twenty-six wildfowl (sub-)species for winter seasons from 1966/67 to 2006/07. The limitations and bias of the sampling methods currently used in the Wetland Bird Survey, are examined. It is shown that post-selection of sites by the proportion of missing values, as is done currently, introduces an additional bias that impacts on reported population trends. A new site selection criteria that minimises additional bias is proposed.

WeBS wildfowl data comprise monthly counts of populations that change over each winter season due to short-term immigration and emigration. Ideas and methods from the field of Functional Data Analysis are used to explore phenological changes (spatiotemporal variation in the seasonal patterns) due to changing species distributions and to select months where the seasonal patterns are most stable.

The Underhill method, of imputing missing values using a site-year-month multiplicative model and the EM algorithm to generate an annual Underhill Index, is reviewed. It is shown that the model is a poor fit to most wildfowl WeBS data sets.

Currently WeBS abundance indices are calculated using the Underhill Index, which treats each species' seasonality as stationary, using the arithmetic mean over months and sites to derive a population index. Using ideas from economics, various alternative indexing approaches to constructing a single yearly index from the monthly counts are compared with simulated examples and WeBS data.

The results have implications beyond the Wetland Bird Survey, to other wildlife monitoring schemes; particularly those that monitor populations which show strong seasonal dynamics.

ACKNOWLEDGEMENTS

Firstly I would like to thank my supervisors, Professor Byron Morgan for his patience, encouragement and expertise and Dr Melanie Kershaw for her inspiration, particularly at the start of the project.

Grateful thanks also to the staff in the WWT Species Conservation department at Slimbridge, particularly Dr Baz Hughes, Richard Hearn, Colette Hall and Dr Mark Trinder and the staff at the WeBS unit at the British Trust for Ornithology; I am especially grateful to Dr Andy Musgrove for providing information and data. I would also like to thank all the hundreds of WeBS volunteers who have collected the data over the past forty years, without whom this research would not have been possible; and to acknowledge the support of the Engineering and Physical Sciences Research Council and the Wildfowl and Wetlands Trust in the form of a CASE studentship.

I am indebted to my colleagues at the Kent and Medway Biological Records Centre, for being so supportive of my research in the latter stages of the project; to Nick Derry for permission to use his beautiful artwork; and to my fellow researchers Rachel, Karen, Dan, Dave, Ash and Eleni for much-needed coffee breaks over the years.

Last, but by no means least, I would not have finished this thesis without the love and support of my husband Timothy, my parents Margaret and Malcolm, my sisters Elaine and Roberta and all my friends and relations who spent so much time encouraging me to succeed. I am eternally grateful to them all.

This work does not necessarily represent the views of the WeBS Partnership.

CONTENTS

Abstract	iii
Acknowledgements	iv
Table of Contents	v
List of Tables	ix
List of Figures	x
List of Programs	xiii
Abbreviations	xiv
Species Names	xv
1. Introduction	1
1.1 Monitoring Overwintering Wildfowl	1
1.1.1 Why count wintering wildfowl?	1
1.1.2 British wintering wildfowl	5
1.2 The Wetland Bird Survey	14
1.2.1 Introduction	14
1.2.2 Methodology of WeBS data collection	15
1.2.3 WeBS data storage and extraction	19
1.3 An Overview of WeBS Data Analysis	25
1.3.1 Introduction	25
1.3.2 Underhill's Indexing Method	31
1.3.3 Criteria for data inclusion in annual Underhill Index	33

1.3.4	Working with WeBS data	35
1.4	Thesis Outline	42
2.	Surveying Wildfowl on a National Scale	45
2.1	Introduction	45
2.2	Detectability and Measurement Error	46
2.2.1	Detectability	46
2.2.2	Optical equipment	47
2.2.3	Measuring counts	49
2.3	WeBS Spatial Coverage	52
2.3.1	Introduction	52
2.3.2	Wetland habitats in Britain	53
2.3.3	Recording effort	58
2.3.4	Case study: WeBS coverage of Kent waterbodies	61
2.4	WeBS Indices Spatial Coverage	66
2.4.1	The completeness criterion	66
2.4.2	Effects of the completeness criterion	70
2.4.3	Top sites criterion	76
2.5	Discussion	85
3.	Seasonality in Wildfowl Populations	89
3.1	Introduction	89
3.2	Intra-winter Wildfowl Movements	95
3.2.1	Introduction	95
3.2.2	Short-term and Weather-related Movements	97
3.2.3	Weather covariates in modelling Wetland Bird Counts	100
3.3	Functional Data Analysis: an Intuitive Way of Looking at Population Dynamics Data	102
3.3.1	Introduction	102
3.3.2	WeBS data as functional data	102
3.3.3	Fitting WeBS population functions	105
3.4	Wildfowl Phenology: Spatiotemporal Variation	112

3.4.1	Phenological change in Shoveler	112
3.5	Identification of Overwintering Wildfowl Populations	118
3.5.1	Introduction	118
3.5.2	Selecting stable months: an FDA approach	121
3.6	Discussion	128
4.	Quantifying Population Change with Incomplete Data	134
4.1	Introduction	134
4.2	Underhill's Site-Year-Month Multiplicative Model	135
4.2.1	Introduction	135
4.2.2	Underhill's Algorithm for fitting the site, year and month factors .	137
4.2.3	Generalised Linear Models	142
4.2.4	Stratified Underhill Model	146
4.3	Assessment of the Underhill S-Y-M model	147
4.3.1	Over-dispersion	147
4.3.2	Assumptions	150
4.3.3	GLM Interactions	155
4.3.4	Consistency intervals	159
4.4	Discussion	168
5.	Indices of population change for seasonal ecological data	172
5.1	Introduction	172
5.2	Aggregation of WeBS counts	174
5.2.1	Introduction	174
5.2.2	Sites	174
5.2.3	Months	176
5.2.4	Phenological implications for reporting trends	180
5.3	Types of index	184
5.3.1	Introduction	184
5.3.2	Non-weighted indices	185
5.3.3	Weighted indices	186
5.3.4	Data-independent weights	188

5.3.5	Data-dependent weights	190
5.3.6	Chain indices	192
5.4	Comparison of Index Performance	194
5.4.1	Summary of example indices	194
5.4.2	Seasonal indices for Shoveler and Shelduck	196
5.5	Discussion	201
6.	Wetland System Simulation	203
6.1	Introduction	203
6.2	Simulation Structure	204
6.2.1	Overview	204
6.2.2	Example simulation	204
6.3	Comparing index estimation in the presence of spatial change, observa- tion error and missing data.	209
6.3.1	Introduction	209
6.3.2	Results	212
6.4	Comparing seasonal population indices using theoretical examples of phenological change	218
6.4.1	Introduction	218
6.4.2	Results	219
6.5	Discussion	223
7.	Conclusion	224
A.	Priority Sites	229
B.	Month Selection	237
	References	265

LIST OF TABLES

1.1	Example of WeBS site sector (count unit) data	21
1.2	Comparison of Ogilvie and Underhill Index methods.	27
1.3	Example of WeBS text file	36
2.1	Completeness of Mallard Counts in Kent	62
2.2	Comparison of imputing with the Underhill model for all sites and sites chosen with the 50% criterion.	75
2.3	Comparison of imputing with the Underhill model for sites chosen with the 50% criterion and the top sites criterion.	84
3.1	Dark-bellied Brent goose counts at The Wash 1997-2000	92
3.2	Months of stable winter populations	127
4.1	Counts imputed by the Underhill Algorithm.	139
4.2	Example Underhill and GLM estimates	144
4.3	Wigeon Underhill and GLM estimates	145
4.4	Dispersion factors	149
4.5	Model Selection	156
4.6	Pintail and Shoveler GLMs with interaction terms	157
4.7	Wigeon GLM with interaction terms	158
4.8	Site consistency interval measures	167
5.1	Shelduck alerts	181
5.2	Summary of prospective indices	194
6.1	Parameters for simple simulation example	207
6.2	Simulation indices	216
B.1	Details of sites used to select months for each species.	238

LIST OF FIGURES

1.1	Number of sites surveyed each year	2
1.2	Mute Swan (<i>Cygnus olor</i>)	6
1.3	Resident species' seasonality plots	8
1.4	Red-breasted Merganser (<i>Mergus serrator</i>)	9
1.5	Mixed Resident and Immigrant species' seasonality plots	10
1.6	Greenland White-fronted Goose (<i>Anser albifrons flavirostris</i>)	11
1.7	Immigrant species' seasonality plots	12
1.8	Number of visits to WeBS counts units by date in 2001/02 and 2002/03.	17
1.9	WeBS core count form	18
1.10	Effect on Underhill Index of incorrect count extraction	22
1.11	Example species account	30
1.12	Iteration time effect of preallocation of MATLAB array	37
2.1	Wintering Waterbirds Indicator: 1975/76-2005/06.	48
2.2	Wigeon (<i>Anas penelope</i>) on ice	50
2.3	Map of WeBS count sector boundaries	52
2.4	Illustrative map showing types of WeBS Sites	53
2.5	Geology map of the British Isles	55
2.6	Map of Great Britain with Mallard sites and elevation.	56
2.7	Maps of Great Britain with Mallard sites and climate.	57
2.8	Map of Great Britain with Mallard sites and urban areas.	59
2.9	WeBS Coverage in 2000/01	60
2.10	Ramsar sites and WeBS counting units in north Kent	64
2.11	Wetlands and WeBS recording effort in Kent	65
2.12	Dark-bellied Brent Goose indices	68

2.13	Mallard indices	69
2.14	Tally of Goldeneye sites by years of non-missing counts	71
2.15	Effect of the 50% completeness criterion on Underhill Indices	73
2.16	Effect of the 50% completeness criterion on Mean Indices	74
2.17	Missing values in Goldeneye sites	77
2.18	Comparison of site criteria for Dark-bellied Brent Goose	80
2.19	Comparison of site criteria for Mute Swan	81
2.20	Comparison of site criteria for Goldeneye	82
2.21	Missing values in Goldeneye sites as added by top sites procedure.	83
3.1	Pintail (<i>Anas Acuta</i>)	96
3.2	Pintail counts at Burry Inlet 2002/03 to 2006/07.	98
3.3	Recent winter counts of Pintail at three sites	99
3.4	Correlation between Bewick swan counts and NAO Index	101
3.5	Example of functional data objects.	103
3.6	B-spline basis functions	107
3.7	Example of m_k , $v(t)$ and $\mathbf{x}(t)$	108
3.8	Effect of changing the smoothing parameter λ	110
3.9	Example of $v'(t)$ and $v''(t)$	111
3.10	Phenological change in Shoveler	113
3.11	FPCA of phenological change in Shoveler	114
3.12	FPCA of phenological change in Shoveler	114
3.13	Shoveler FPCA scores	117
3.14	Kirby method for choosing months	120
3.15	Derivative used to find the stable months.	122
3.16	Bootstrap of derivatives used to find the stable months.	123
3.17	Phase-plane plot of Mallard and European White-fronted Goose	124
3.18	Effect of changing λ on first shoveler principal component	128
3.19	Effect of changing λ on first shoveler principal component score	129
3.20	Effect of changing λ on month choice for Mallard	130
3.21	Effect of changing λ on month choice for European White-fronted Goose	131

4.1	Underhill's Imputing Algorithm	137
4.2	Mallard (<i>Anas platyrhynchos</i>)	141
4.3	Total mallard population on selected sites in December	146
4.4	Month \times year interactions	151
4.5	Month \times site interactions	152
4.6	Fitted values from the Underhill model for Wigeon at the Ribble Estuary .	153
4.7	December Underhill index for Wigeon	153
4.8	Site consistency intervals with 100, 200 & 1000 bootstrap samples	162
4.9	Pintail site consistency intervals	163
4.10	Goldeneye site consistency intervals	164
4.11	Swan and geese site consistency intervals	165
4.12	Duck site consistency intervals	166
4.13	Pintail indices	169
5.1	Presentation of monthly counts	178
5.2	Smooth Shelduck indices	181
5.3	Shelduck month indices	182
5.4	Shoveler month indices	196
5.5	Shelduck indices using varying month aggregation formulae	198
5.6	Shoveler indices using varying month aggregation formulae	199
5.7	Laspeyres indices	200
6.1	Flow diagram of generic wetland system simulation	205
6.2	Flow diagram of timing and location routines	206
6.3	Simple Simulation Example	208
6.4	Spatial change simulations for an increasing population.	210
6.5	Spatial change simulations for a constant population and an invasive pop- ulation.	211
6.6	Observation error.	213
6.7	Precision.	214
6.8	Index results for the invasive population simulation.	215
6.9	Phenological change simulations.	221
6.10	Indices for phenological change simulations.	222

LIST OF PROGRAMS

1.1	MATLAB code to import WeBS data	38
1.2	Vectorising MATLAB code	41
2.1	MATLAB function to extract required sites	67
2.2	Top sites procedure	78
4.1	MATLAB implementation of Underhill's algorithm	140
4.2	R code for fitting generalised linear models	143
4.3	MATLAB code to generate CIs	161
5.1	MATLAB code to calculate indices	195

ABBREVIATIONS AND ACRONYMS

BBS	(UK) Breeding Bird Survey
BoEE	Birds of Estuaries Enquiry
BOU	British Ornithologists' Union
BTO	British Trust for Ornithology
CCW	Countryside Council for Wales
DWS	Dispersed Waterbirds Survey
EN	English Nature (now Natural England)
FDA	Functional Data Analysis
FPCA	Functional Principal Components Analysis
GAM	Generalised Additive Model
GIS	Geographical Information System
GLM	Generalised Linear Model
IOC	International Ornithological Congress
IWC	International Waterfowl Census
IWRB	International Waterfowl and Wetlands Research Bureau (now Wetlands International)
JNCC	Joint Nature Conservation Committee
NE	Natural England
NWC	National Wildfowl Counts (National Waterfowl Counts from 1991)
RSPB	Royal Society for the Protection of Birds
SNH	Scottish Natural Heritage
UK-NEWS	UK Non-Estuarine Coastal Waterbird Survey
WeBS	Wetland Bird Survey
WWT	Wildfowl and Wetlands Trust

SPECIES NAMES

In this thesis, wildfowl species are referred to by their British Ornithologists' Union British English vernacular name, qualified by subspecies or sub-population where this is used in the Wetland Bird Survey. The following table lists the names of species frequently referred to, together with the British Trust for Ornithology's two letter unique identification code, the Scientific name and also the International Ornithological Congress English name, where this differs from British usage.

BTO code	BOU Name	IOC Name	Scientific Name
MS	Mute Swan		<i>Cygnus olor</i>
WS	Whooper Swan		<i>Cygnus cygnus</i>
BS	Bewick's Swan	(Tundra Swan)	<i>Cygnus columbianus bewickii</i>
PG	Pink-footed Goose		<i>Anser brachyrhynchus</i>
EW	European White-fronted Goose	(Greater White-fronted Goose)	<i>Anser albifrons albifrons</i>
NW	Greenland White-fronted Goose		<i>Anser albifrons flavirostris</i>
JI	Icelandic Greylag Goose		<i>Anser anser</i>
JE	Re-established Greylag Goose		<i>Anser anser</i>
CG	Canada Goose	(Greater Canada Goose)	<i>Branta canadensis</i>
YS	Svalbard Barnacle Goose		<i>Branta leucopsis</i>
YE	Naturalised Barnacle Goose		<i>Branta leucopsis</i>
DB	Dark-bellied Brent Goose	(Brant Goose)	<i>Branta bernicla bernicla</i>
SU	Shelduck	(Common Shelduck)	<i>Tadorna tadorna</i>
WN	Wigeon	(Eurasian Wigeon)	<i>Anas penelope</i>
GA	Gadwall		<i>Anas strepera</i>
T	Teal	(Eurasian Teal)	<i>Anas crecca</i>
MA	Mallard		<i>Anas platyrhynchos</i>
PT	Pintail	(Northern Pintail)	<i>Anas acuta</i>
SV	Shoveler	(Northern Shoveler)	<i>Anas clypeata</i>
PO	Pochard	(Common Pochard)	<i>Aythya farina</i>
TU	Tufted Duck		<i>Aythya fuligula</i>
SP	Scaup	(Greater Scaup)	<i>Aythya marila</i>
GN	Goldeneye	(Common Goldeneye)	<i>Bucephala clangula</i>
RM	Red-breasted Merganser		<i>Mergus serrator</i>
GD	Goosander	(Common Merganser)	<i>Mergus merganser</i>
RY	Ruddy Duck		<i>Oxyura jamaicensis</i>

Chapter 1

INTRODUCTION

1.1 Monitoring Overwintering Wildfowl

1.1.1 *Why count wintering wildfowl?*

The British Isles, with a temperate climate due to the Gulf Stream, are home during winter to large numbers of ducks, geese and swans. Some of these birds both breed and winter within the British Isles but many more breed in other regions as far away as Taimyr in Siberia. These waterbirds take advantage of the rarely frozen inland waterbodies and the mild estuaries around the British coast which keep the “larder” open throughout winter. There are some cases where the entire population winters in the British Isles, such as the subspecies Greenland White-fronted Goose *Anser albifrons flavirostris* and the Iceland-Greenland subpopulation of Pink-footed Goose *Anser brachyrhynchus*.

Wildfowl have provided winter food for Europeans for centuries, but in the 20th century it became clear to wildfowler and painter Peter Scott that the populations of some species were declining and needed protection. In 1946 he established the Severn Wildfowl Trust, now the Wildfowl and Wetlands Trust (WWT), to conserve over-wintering wildfowl populations by the creation of nature reserves and refuges where shooting is prohibited, and to create a sound scientific basis to lobby for policies and laws to protect wildfowl from shooting pressures when necessary. Sir Peter Scott, the first person to be knighted for services to conservation, paved the way for today’s conservation industry which employs

thousands of people as reserve wardens, campaigners, educationalists and researchers all seeking to understand the natural world better, to look after it and preserve it for the enjoyment of all.

In response to the same concerns about a possible decline in wildfowl numbers, at about the same time as the Severn Wildfowl Trust was established, a national scheme to count non-breeding wildfowl in Great Britain was created in 1947 (Matless et al., 2005). The purpose of the scheme was to determine trends in numbers, so that the impacts of wildfowling and wetland development on wildfowl numbers could be assessed objectively. Initially organised by the Wildfowl Enquiry committee of the British Section of the International Council for Bird Preservation, the wildfowl counts were administered by the Wildfowl Trust from 1954 onwards and funded by the Nature Conservancy council (Cranswick et al., 1997). Today the scheme has become the Wetland Bird Survey (WeBS) and many other species of waterfowl are counted along with wildfowl; the British Trust for Ornithology (BTO) currently administer the scheme and it is funded by the Joint Nature Conservancy Council (JNCC), WWT, the BTO and the Royal Society for the Protection of Birds (RSPB).

The Wetland Bird Survey is an early example of what is now popularly known as “citizen science”: the use of a network of untrained observers to collect data for scientific




Figure 1.1: Number of sites surveyed each year which recorded Mallard at least once. (Mallard is used as a proxy for site visits, as direct site visit information was unavailable; see Section 1.3.4.) A single site may be split into several sectors, surveyed by different individuals.

research, that both engages the public with the research, and allows more data to be collected than would be possible without the time freely given by the volunteers (Greenwood, 2007). Every year thousands of people across the U.K. take part in the Wetland Bird Survey (Figure 1.1). Current participation is at the level of over 3000 count sectors at about 2000 sites being counted at least once over the winter surveying period (BTO, 200-). Unlike some wildlife monitoring schemes which require special training, high levels of expertise or a government licence, such as small mammal trapping, surveying bat hibernation roosts or bird ringing, any birdwatcher with basic waterfowl identification skills may take part in WeBS.

Continuing monitoring of populations is essential to assess the effectiveness of species conservation and management, to understand the effects of environmental changes such as climate change and pollution and to intimate new management policies. This was recognised on the occasion of WWT's 50th anniversary:

It is in my view wholly appropriate that WWT should now concentrate on the ecological survey and monitoring work on which conservation depends. The long-term data sets, for years dismissed as "mere monitoring", will undoubtedly prove of increasing scientific value in the future, providing unique databases on the long-term history of populations and their response to change, and on which to answer new questions and examine new ideas.

Newton (1997)

The statistical techniques used to analyse basic ecological count or "census" data are often less sophisticated than those used for more experimental ecological data to estimate abundance. This "mere monitoring" data has perhaps in the past been seen as too simple or straightforward to warrant statistical analysis. Many long-term monitoring schemes have evolved over time rather than being designed, in consultation with a statistician, to yield the best results. Nevertheless, they have immense value as a resource, allowing projects such as the UK Phenology network to use data people started collecting just for interest to gain new insight into climate change. Sometimes the true value of long-term

records is not apparent for many years. The Marsham Phenological Record consists of 27 phenological events recorded by successive generations for over 200 years, including Hawthorn leafing and the first swallow of the year. This family tradition has given insights into how species respond to climate over a long period of time (Sparks and Carey, 1995).

Ecological population monitoring schemes encompass a wide variety of surveys with a multitude of aims and methods. One way of differentiating between monitoring programmes which study population changes from year to year is between those where the population is approximately constant for the surveying period and those which have a seasonal component, with the population changing within the year sample. In schemes where within-year changes in the population over the sampling period are considered negligible, often several replicated samples can be obtained to improve parameter estimation. Many of these surveys monitor a resident breeding population, and often will count only the females or the young rather than the whole population.

In schemes such as the Wetland Bird Survey and the U.K. Butterfly Monitoring Scheme, where the seasonal component of the monitoring data is of additional interest to an overall annual population estimate, the extra effort required to obtain a representation of the within-year or within-season change usually precludes true sample replication. However, the spatial and temporal correlations between the samples that have been collected may be turned to advantage in overcoming this limitation.

This thesis will look in detail the Wetland Bird Survey and what models, methods and statistics are appropriate for this type of data. However, many concepts explored are transferable to other ecological monitoring schemes, particularly those with a seasonal component.

1.1.2 *British wintering wildfowl*

In taxonomy, species are classified using a hierarchical structure from order through family, genus, species and sub-species. The wildfowl family, *Anatidae*, comprise three main groups: the swans (genus *Cygnus*, three species with large wintering populations estimated by WeBS data), the geese (*Branta*, three species and *Anser*, three species) and the ducks. Ducks that have sufficient WeBS coverage to construct population indices consist of shelducks (*Tadorna*, one species), dabbling ducks (*Anas*, six species), pochards (*Aythya*, three species), goldeneyes (*Bucephala*, one species), sawbills (*Mergus*, two species) and stiff-tails (*Oxyura*, one species).

Of these twenty-three species which have trends estimated by Wetland Bird Survey data, two, Greylag Goose and Barnacle Goose, have both UK breeding populations and populations from overseas which are treated separately in WeBS. Individuals are assumed to be resident (Naturalised Barnacle Goose and Re-established Greylag Goose) in southern Britain and immigrant (Svalbard Barnacle Goose and Icelandic Greylag Goose) in Scotland. However, due to expansions in range in both the re-established and northwest Scottish populations, it is expected that the distinction in Greylag Goose will become more difficult to maintain in future (Banks et al., 2006). Additionally, White-fronted Goose are split into populations from Greenland and Europe; these sub-species are separable on plumage differences.

The resulting twenty-six wildfowl types are listed for reference on Page xv in a table with their scientific and international English names, and are referred to throughout the text by their British vernacular names. They are grouped below into three broad categories relating to their migratory status. All these species are covered by the Wetland Bird Survey, but the survey's methodology is better suited to those species that spend their lives mostly on water rather than those that spend a lot of time grazing on land away from water. For many of the swan and geese species which have a high fidelity to certain wintering sites, more specialised census surveys are carried out to complement data collected in the Wet-

land Bird Survey. These will, for example, count the birds at the time they are roosting all together on the water at night, rather than the middle of the day. For this reason, most of the methodology examples given throughout the text are of those species for which WeBS is the main source of information on overwintering population trends.

Resident species

The sedentary Mute Swan and Mallard and the (re-)introduced species Greylag Goose, Canada Goose, Barnacle Goose and Ruddy Duck, are species which tend to spend their whole lives in Britain and relatively few individuals, if any, come from populations abroad. Figure 1.3 shows their pattern of occurrence within years as a seasonality plot.

The Mute Swan (illustrated in Figure 1.2) shows a fairly even distribution all year, with a peak in numbers in August, just after the breeding season. This results in an almost circular seasonality plot in Figure 1.3. The three geese species also exhibit fairly even numbers throughout the year, with a dip around March to May when they are nesting. Unlike Mute Swans, they usually nest in wetlands such as grazing marsh, rather than on



Figure 1.2: Mute Swan (*Cygnus olor*). ©Nick Derry, used with kind permission.

the edge of open water, so are less visible to WeBS counters.

Mallard in northwest Europe are largely sedentary or migrate only small distances and so are included here, although some winter immigration to Britain does occur (Wernham et al., 2002). However, although this is a well-known species, there is uncertainty about both the quantity of winter immigrants and birds raised and released for shooting (Kershaw and Frost, 2006). Together with movement from small rivers, ponds and wetlands, (which are not so well represented in WeBS) to larger sites included in the Wetland Bird Survey, this results in apparently greater numbers in the winter months in Figure 1.3.

Similarly, the Ruddy Duck, a non-native species, is widely dispersed on small waterbodies in the breeding season, but congregates on a small number of wintering grounds which are better surveyed by WeBS. This species, like Mallard, thus shows an apparent increase in numbers in winter in the seasonality plot.

Mixed resident and immigrant wintering species

The U.K. over-wintering populations of the sawbills Red-breasted Merganser (Figure 1.4) and Goosander, the goose-like Shelduck, surface-feeding Gadwall and Shoveler and the diving Tufted Duck are from a mixture of British and non-British breeding populations.

Red-breasted Merganser and Goosander are fish-eating gregarious species which move from their breeding grounds into large flocks in winter on the coast and at large waterbodies (Snow and Perrins, 1997). Their movement between different sites and habitats for breeding, moulting and wintering contributes to the unevenness of their seasonality plots. The Shelduck is a coastal species which undertakes major moult migration after breeding (leaving ducklings in the care of a small number of “nannies”); traditionally most Shelduck in northwest Europe have migrated to the German Wadden Sea area for this purpose, but a number moult in the Bridgwater Bay area of the Severn Estuary (Fox and Salmon, 1994). This results in a mini-peak in July in Figure 1.5.

The Tufted Duck, Gadwall and, on a smaller scale, Shoveler, have increasing British

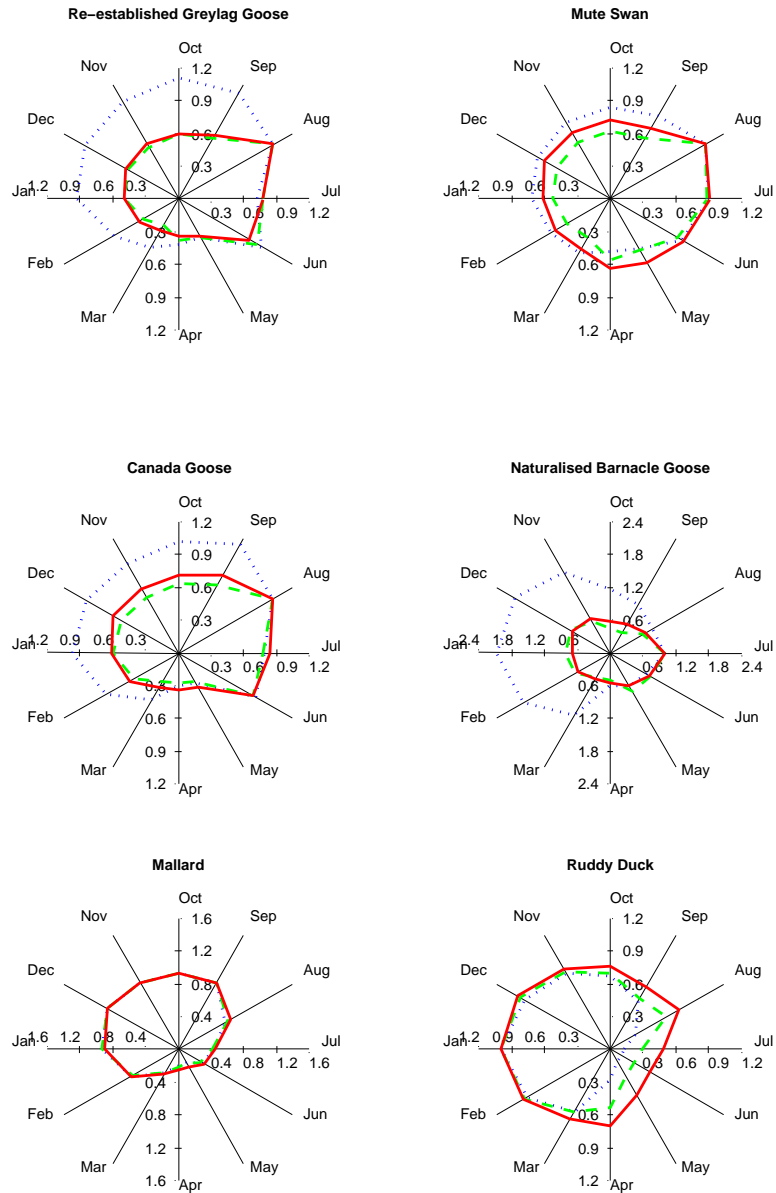


Figure 1.3: Seasonality plots of species whose wintering population originates in the UK breeding population. The solid red line is calculated using the mean of non-missing counts at all sites from 1966-67 to 2004-05 for that month. For the subset of sites selected using 50% criterion, the dashed green line is the mean of the counts and the dotted blue line is the month factor after imputing using the Underhill algorithm (see Section 1.3). In all three cases, the peak month according to the non-imputed counts is used as a base.

breeding populations and the seasonality plots in Figure 1.5 are an average over time. Changes in seasonality will be explored in Chapter 3.

Immigrant wintering species

The summer population is negligible compared with the over-wintering population in the remaining fourteen species, shown in Figure 1.7. In some species, birds occur in spring and autumn as passage migrants en route to other destinations, as well as those that remain all winter.

Bewick and Whooper Swans and Pink-footed, Greenland White-fronted (Figure 1.6) and Icelandic Greylag Geese all arrive in October or November and leave in March or April. Svalbard Barnacle and Dark-bellied Brent Geese stay a little longer, until May, and European White-fronted Geese arrive later, present in large numbers only from December. The ducks all leave in March or April; Pintail and Wigeon arriving back in September, followed by Scaup in October and Goldeneye in November.

Teal and Pochard both breed in the U.K., but the breeding population is dwarfed by wintering birds. In both species the wintering population consists of birds that breed



Figure 1.4: Red-breasted Merganser (*Mergus serrator*). ©Nick Derry, used with kind permission.

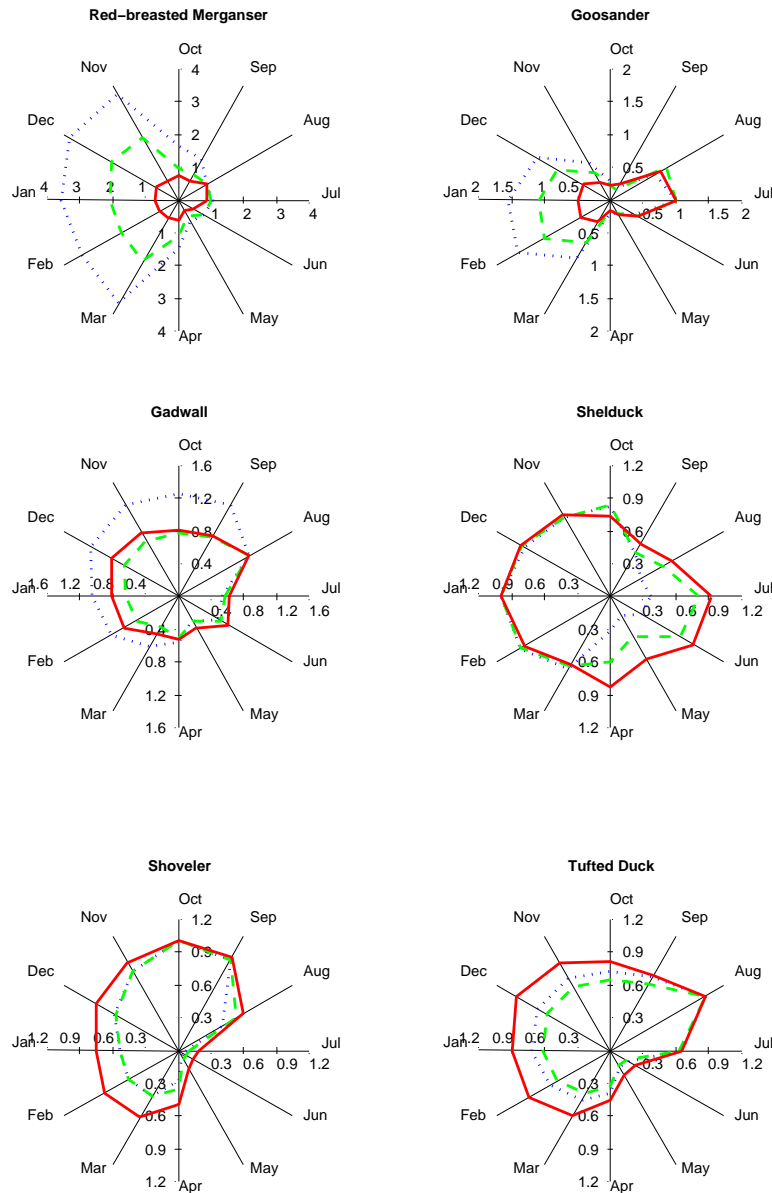


Figure 1.5: Seasonality plots of species whose wintering populations originate both from breeding populations in the UK and overseas. The solid red line is calculated using the mean of non-missing counts at all sites from 1966-67 to 2004-05 for that month. For the subset of sites selected using 50% criterion, the dashed green line is the mean of the counts and the dotted blue line is the month factor after imputing using the Underhill algorithm (see Section 1.3). In all three cases, the peak month according to the non-imputed counts is used as a base.



Figure 1.6: Greenland White-fronted Goose (*Anser albifrons flavirostris*).

©Nick Derry, used with kind permission.

in Denmark, north Poland and Germany, Fenno-Scandia, the Baltic states and Russia (Snow and Perrins, 1997). The proximity of some of the breeding grounds, together with migration after breeding beginning in late June, in the case of Teal, and suspected moult migration to Britain in June and July, in the case of Pochard, result in a less constrained seasonal pattern in Figure 1.7 for these species, with significant numbers present from August to March for Teal and July to March for Pochard (Wernham et al., 2002).

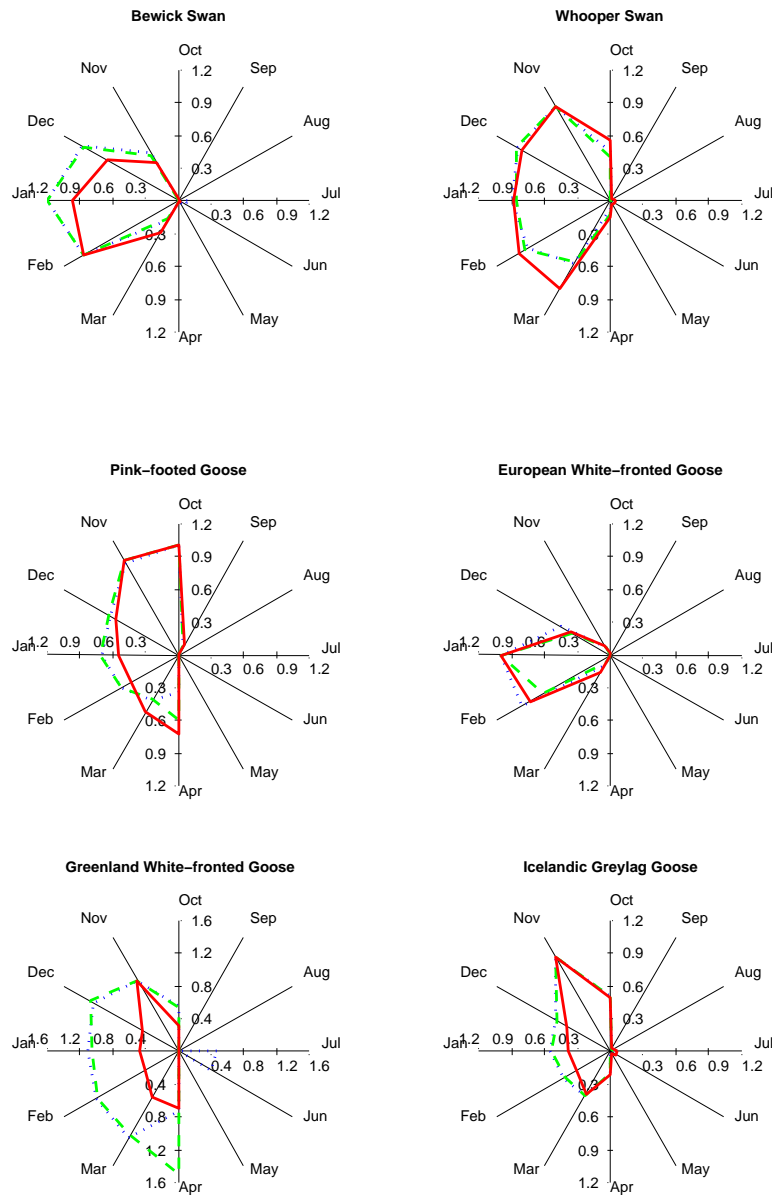


Figure 1.7: Seasonality plots of species where the majority of the overwintering population comes from outside Great Britain. The solid red line is calculated using the mean of non-missing counts at all sites from 1966-67 to 2004-05 for that month. For the subset of sites selected using 50% criterion, the dashed green line is the mean of the counts and the dotted blue line is the month factor after imputing using the Underhill algorithm (see Section 1.3). In all three cases, the peak month according to the non-imputed counts is used as a base. (Figure continues on the next page.)

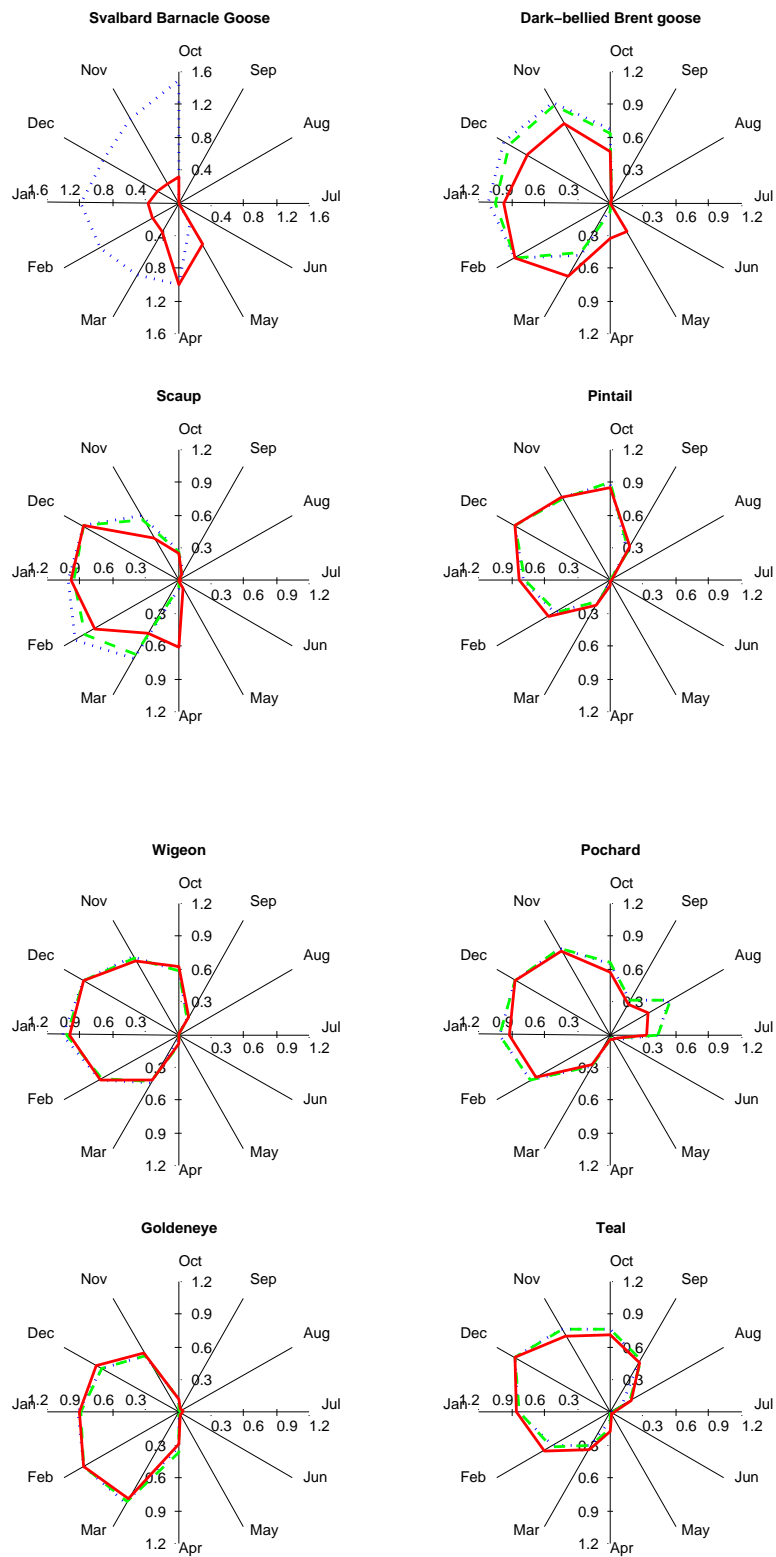


Figure 1.7: Continued. Note that in the case of Svalbard barnacle goose, the 50% criterion is not applied, and the blue dotted line is the month factor after imputing using the Underhill algorithm for all sites.

1.2 The Wetland Bird Survey

1.2.1 Introduction

The Wetland Bird Survey (WeBS) is a major source of information on wetland bird populations in Great Britain and Northern Ireland; an equivalent scheme, the Irish Wetland Birds Survey (I-WeBS) collects information using identical methodology in the Republic of Ireland. Data is collected under the scheme on the two main waterbird groups of wildfowl and waders together with divers, grebes, cormorants, herons, rails and kingfishers. Data recorders also have the option of recording gull and tern numbers. In this thesis we shall be considering the methodology and analysis of WeBS data for wildfowl species only, but most of the conclusions will be valid for the other waterbird taxa covered by the scheme.

WeBS has evolved over time rather than being designed in its final state. It was formed in 1993 by amalgamating two earlier schemes, the National Wildfowl Counts (NWC), designed to cover wildfowl, and the Birds of Estuaries Enquiry (BoEE) which was designed to cover waders and estuarine wildfowl. It is the NWC which has the longer history out of the two original schemes. The first pilot scheme for NWC began in 1947 in response to concerns that some wildfowl populations were decreasing due to hunting pressures. As no quantitative data had been collected in the past, the question of whether hunting had a big impact on the population was unanswerable, and so a monitoring programme was needed. The BoEE followed in 1969 following concern about threats to estuaries from barrage and land reclamation schemes. Wildfowl data from BoEE was routinely added to the NWC database and a joint recording form for the two schemes for coastal sites was used from 1989 and so the formation of WeBS became inevitable (Moser, 1982; Cranswick et al., 1997; BTO, 200-).

When NWC and BoEE were set up, it is doubtful as to whether the pioneers who started thinking about a monitoring scheme considered using the latest statistical sampling techniques. They wanted a quick answer to a simple question, viz. are waterbird populations

stable, increasing or declining? To ornithologists, aware that most waterfowl are concentrated on relatively few key wetland sites, the obvious solution is to count them systematically. In some cases in ecology this would be completely impractical: for example, if you wanted to know how many jellyfish there are in the Atlantic, it would be impossible to count them individually. In the case of British wildfowl populations, however, such a goal, although difficult, is certainly possible in theory. Intuitively, if you could count enough of the individuals at large accessible wetlands, then the populations at smaller, less accessible sites should be negligible.

Since the aim of WeBS is to cover as many species as possible, the survey is not tailored to the ecology and behaviour of one species in particular. Species vary in their tendency to aggregate in large numbers, how widespread they are and their population size and all these factors affect how much the recorded count reflects the actual population. For some species which winter on only a handful of sites, most of the population are present on WeBS sites and so, assuming other factors are constant, the trends recorded on WeBS sites will reflect the wintering population trends. However, for species which are widely dispersed such as mallard, the trends on WeBS sites may or may not reflect trends in the overwintering national population, which is present not only on WeBS sites but also on small ponds, streams and rivers all across the country.

1.2.2 Methodology of WeBS data collection

A Wetland Bird Survey “site” is usually a clear geographical area such as an estuary, complex of gravel pits or lake that the birds are likely to see as one entity and use different parts of on a daily basis. Sites that are too large for one person to cover are divided into count units known as “site sectors”. For large and complex sites a hierarchical structure may be in place where, for example, a sector is part of a large inlet on the south side of the estuary.

WeBS data are collected mostly by volunteer bird-watchers, as well as conservation pro-

professionals such as reserve wardens for some important sites, for example those designated as internationally important for wildfowl. Individuals interested in the scheme contact a local organiser, also a volunteer, whose job it is to allocate local count units, ensuring that important sites in their area are covered first.

Historically, both count sectors and sites were defined by the counters and local organisers themselves. There was no guarantee that when a new counter took over from the previous incumbent, he or she would collect data over exactly the same bounded locality. An early attempt was made to ascertain where exactly counting was taking place, by requesting that all counters make a map of their usual route, so that each count unit was standardised to minimise errors caused by a change in counter and double counting or under counting of birds where two sectors overlapped through an oversight by the local organiser (Bell, 1994a). The information gained from the volunteers formed a wetland register of 1040 sites. This is now being used in conjunction with new research to create a more complete register of WeBS sites within a Geographical Information System (GIS). The provisional WeBS GIS layers can be seen in Figure 2.3 in Chapter 2.

The “look-see” method of counting is used with WeBS data (Bibby et al., 1992). The only special equipment counters need is optical equipment such as binoculars and telescopes and possibly a tally counter to assist in counting large numbers of birds. As technology has advanced, optical equipment has become more effective and more widely used; it is now recommended (although not required) that counters have a telescope with 30x magnification (Gilbert et al., 1998). Section 2.2.3 will address the possibility that apparent increases in some populations over time may be due to counters being able to see and count birds at further distances than previously.

Counters are requested to undertake their count on set “priority count dates”, usually a Sunday in the middle of the month. Ideally the count should take place during the morning, or when tidal conditions are suitable. The priority dates are chosen to maximise the number of estuaries where counts can be undertaken at high tide. Due to differences in the tidal regime around the country, some sites will use a different date so that counts




Figure 1.8: Number of visits to WeBS counts units by date in the submission years 2001/02 and 2002/03. Although each WeBS year for reporting purposes runs from July to June, counters are asked to submit data in April; consequently data for the most recent year in the annual report is incomplete and subject to revision in the following year's report. A major outbreak of foot and mouth disease in February 2001 meant WeBS surveys were suspended and did not return to normal levels until September 2001. From *The Wetland Bird Survey 2001-03* (Cranswick et al., 2005); reproduced with permission.

can be made at high tide. Weather and counter availability will also affect whether the priority count date is adhered to. The local organiser is required to synchronise counts within sites, and preferably also between sites which birds are known to move between frequently. Figure 1.8 shows the distribution of visits by day for two years. Although there is a high degree of synchronisation around the priority dates, indicated by the large peak on the priority dates, a substantial proportion of counts are undertaken slightly early or late.

The counts for each species are recorded on a form together with other information such as observer code, site code, presence of ice and start and finish time of the survey (Figure 1.9). This additional information has not been computerised for historical forms, but is likely to be useful for analysis in the future.

Also recorded is the observer's opinion as to whether all the birds present had been counted successfully ("Count Accuracy OK") or not ("Count Accuracy Low"). If count

The form is divided into several numbered sections:

- 1] COUNTER CODE:** A grid for entering a counter code. Below it, text asks for full name and address if the code has changed.
- 2] COUNT UNIT:** Fields for Central Grid Ref. and Count unit code. A note below asks for site information if part of a larger site.
- 3] DATE:** A grid for entering the date (e.g., 13/10/97). Below it, a section for TIME (Start and Finish) with a 24-hour clock format.
- 4] WATERBIRDS:** A table with columns for bird species and rows for counts. Species include Red-throated Diver (RH), Black-throated Diver (BV), Great Northern Diver (ND), Little Grebe (LG), Great Crested Grebe (GG), Red-necked Grebe (RX), Slavonian Grebe (SZ), and Black-necked Grebe (BN). A 'Please ✓ box if NO waterbirds were present' row is at the top.
- 5] COUNT CONDITIONS:** A section for recording environmental conditions. It includes a grid for % of water area covered by ice and a key for tide states (1 Rising, 2 High, 3 Falling, 4 Low).
- 6] COVERAGE:** A section for recording count accuracy. It includes a grid for 'Count accuracy' (OK / LOW) and checkboxes for 'LOW' reasons: 'I did not cover all of my count area', 'Poor visibility (due to fog, glare, heavy rain etc.)', and 'High levels of disturbance'.
- 7] ADDITIONAL INFORMATION:** A large text area for notes. A sub-section titled 'Additional Counters' asks for names and addresses of other observers.

Figure 1.9: Extract from a WeBS core count form, showing space for counts and additional information.

accuracy is low the observer is asked whether this was due to not covering all the count area, poor visibility or high levels of disturbance. This is used to determine the completeness code for the site as a whole (see Section 1.2.3 below).

Hard copies of the forms are returned by post and then keyed into a database by two professional data inputters, with computer checking identifying discrepancies between the two for correction (Collier et al., 2005). Alternatively, volunteers have been able to submit counts via an online form since 2007 (Musgrove, 2007). Particularly unusual counts (according to a computer algorithm, based on the history of the site) are queried with the original observer.

1.2.3 WeBS data storage and extraction

WeBS data is currently stored in an Oracle database. Data from old forms have also been input to the database back to 1960, but wildfowl counts from 1947 to 1959 have not yet been computerised. When the sector counts have been amalgamated into a site count, associated information is available on the date, the site and the quality of the count.

Date

As shown in Figure 1.8 above, the dates on which the actual surveying takes place varies across the country. When extracting data, the corresponding month and year only is used. The BoEE year was defined to run from July to June, since waders' autumn migration begins in July, and this convention is used in the WeBS database. (For example, in the first line of Table 1.3 on Page 36, "8" in the month column and 2006 in the year column corresponds to surveys undertaken in February 2007.)

Site

The count data for all the count units within a site are summed to give site counts. Due to sites being divided up into different count units over time (as explained above in Section 1.2.2) WeBS data is almost always analysed at the site level to ensure consistency over time. For example, the sector level counts for the Ribble Estuary, site 57901, are shown in Table 1.1; three different divisions of the estuary have been used over time. Some historical counts have only been added to the database at the site level, as is the case here for counts from the 1980s.

When extracting the data, it is important to ensure that all the count sectors associated with a site are included when totalling counts, which is not always straightforward. Indices published in the annual reports *Wildfowl and Wader Counts* for 2001/02-02/03 and 2003/04 for Bewick's Swan and Whooper Swan were erroneous since counts for the

Ouse Washes, Nene Washes and Martin Mere, the most important sites for those species, were incorrectly stored in the database (Banks et al., 2006). Similarly, in the case of the Ribble Estuary, during my research for this thesis I discovered a discrepancy between all wildfowl data extracted from the database in 2004 (counts up to the 2002/03 season) with that extracted in 2006 (2004/05 season) and 2008 (2006/07 season). Further investigation revealed that the counts recorded at the site level had been inadvertently omitted from the summing process, so that counts for the 1975/76 and 1979/80 to 1988/89 seasons were missing completely from the later extractions.

Code	Sector	57901	57401	57402	57403	57405	57407	57451	57452	57453	57454	46430	46431	46432	57455	57456	57457	57458	57459	57460	46433	46434	57901	
		Ribble Estuary	Ribble Estuary - South Outer Marshes	Ribble Estuary - North Outer Marshes	Ribble Estuary - Inner Marshes	Ribble Estuary - Outer	Ribble Estuary - Middle	Banks Marsh East	Hesketh Out-Marsh	Longton Marsh	Freckleton Marsh	Ainsdale Beach	Marshside Beach	Marshside Sand Company	Lytham Beach	St Anne's Beach	Crossens Out-Marsh	Banks Marsh West	Banks Marsh Central	Warton Marsh	Marshside 1	Marshside 2	TOTAL Ribble Estuary	
66-67				150	22																		172	
67-68				140	50																		190	
68-69				850	40																		890	
69-70			750	100	65																		915	
70-71			480	740	140																		1360	
71-72			410	450	95																		955	
72-73			4200	100	190																		4490	
73-74			6000		600																		6600	
74-75			3900																				3900	
75-76		3960																					3960	
76-77			2500	14	1500																		4014	
77-78			1200		350																		1550	
78-79			1218	7	393																		1618	
79-80		118																					118	
80-81		6380																					6380	
81-82		6189																					6189	
82-83		13823																					13823	
83-84		6302																					6302	
84-85		16889																					16889	
85-86		20192																					20192	
86-87		24462																					24462	
87-88		30039																					30039	
88-89		19891																					19891	
89-90					593		40912																41505	
90-91								60	31	70	1						20000	35000	4000		25		59187	
91-92								5000	15670	40	11			15000	3		5500	20000	18600		150	40	80014	
92-93										209	70			12			5000	37000	5000		300	850	48441	
93-94								4000	770	274	21			10000			5700	26000	45000		500	200	92465	
94-95									260	199	19			40000			9100	50000	10000		400	300	110278	
95-96								500	1360	421	111			7500			5000	60000	6500		500	280	1750	83922
96-97									246	151	120			10000			1990	21000	7000		700	2210	43417	
97-98									25	180	200			5000			4620	30000	1100			2000	1675	44800
98-99									130	250	30			15000			30450	40000	10800		45	100	50	96855
99-00									40	18	285			200			3205	9000	11000		10	4710	4730	33198
00-01								1	25	234	25			7	1000		542	28000	10000		3950	9120	52904	
01-02								1000	550	942			26	2200			2515	50000	2110		2288	7030	68661	
02-03								700	2830	550	1180			2600			3843	27000	3650		4385	1000	47738	
03-04								1000	3320	2595	5280		45	7000			1022	18000	7000		5855	7700	58817	
04-05								4000	290	3500	3500		2	300			4920	20000	20000		3470	3250	63232	
05-06								3000		2695	1514			400			2660	45000	11000		1400	4510	7480	79659
06-07									1030	650	900			150				20000	13200		45	8890	12520	57385

Table 1.1: December counts of Wigeon at Ribble Estuary sectors (count units).
Counts have been computerised at a sector level in later years, but 1980s counts have been stored at the site level.

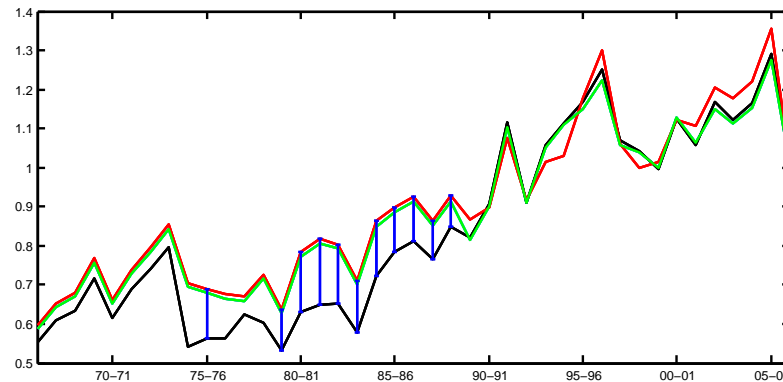


Figure 1.10: Underhill Index for Wigeon (using 50% criterion and September to March counts; see section 1.3.3). The index based on the 2008 extraction excludes the Ribble Estuary as having insufficient complete counts (red; 632 sites). The index based on the 2008 extraction, but with the omitted counts from the Ribble Estuary added from the 2004 extraction, does include the Ribble Estuary (black; 633 sites). Also shown is the index based on the 2008 extraction, including Ribble Estuary counts where available (green; 633 sites). Vertical blue lines indicate years where Ribble counts were incorrectly extracted in 2008.

Figure 1.10 shows the effect that the incorrect extraction in 2008 had on the calculated Underhill Index (see Section 1.3 below). The Ribble Estuary is one of the most important sites in the country for Wigeon, and one undesirable consequence of the incorrect extraction was to exclude the Ribble Estuary counts completely for all years from the Index when basing the index on the winter months September to March, as less than 50% of counts were flagged as complete. If the criterion is overruled, and the data from the 2008 extraction is used including the Ribble Estuary, the index is still higher than it should be in the late 1970s and 1980s. This is caused by the algorithm imputing higher counts at Ribble than were actually present at the time, as Wigeon numbers at the Ribble Estuary have increased dramatically over the forty years of the survey (see Table 1.1 and further discussion in Section 4.3.2).

Each site has a reference code, which contains basic location and habitat information for that site. For example, the site code for Rutland Water is 36156. The last two digits are used to uniquely identify the site. The first two digits indicate the county. This in turn indicates the site is in the region East and Central England. There are ninety-

seven counties allocated to ten regions. The third digit indicates the habitat, in this case reservoir. There are codes for five habitat types: reservoir, estuary, river or marsh, natural lake and coast. In addition, the Ordnance Survey grid reference for the centre of the site is available, as well as provisional GIS polygons of the site boundary.

Quality

As remarked above, counters are asked to mark on their forms whether their count accuracy was “OK” or “low”. If the counter was unable to survey on that occasion then the count is flagged as missing. A low accuracy flag signifies that the counter suspects the count is an underestimate, for example because of fog or stormy conditions or disturbance from shooting or water sports.

In the case of sites consisting of more than one count unit, the following algorithm is then used to determine the quality of coverage code for the site count.

1. Calculate the mean count for each count unit over the complete counts in the preceding four years in the month in question and the two adjacent months (this is to control for seasonal variation and long-term trends).
2. Calculate the contribution of each count unit by dividing the mean count for that unit by the sum of the mean counts for all the count units within the site.
3. For those count units which are flagged as complete on the date in question, calculate the sum of the contributions. If these total 75% or more then the site count is marked as good quality, otherwise it is marked as poor quality.

Indices published in each year’s *Wildfowl and Wader Counts* have used this algorithmic approach to completeness for wildfowl species since the 01-03 report. “Good quality” counts have a quality flag of 1 in Table 1.3 on Page 36, and the count is known as “complete”; “poor quality” counts are flagged as 2 and known as incomplete. Counts may also be missing completely (flagged as 3). As the algorithm is based on a moving window of

counts, the quality flag for a site count may change from year-to-year (Cranswick et al., 2005).

1.3 An Overview of WeBS Data Analysis

1.3.1 Introduction

The National Wildfowl Counts (NWC) started with a handful of sites but as more sites were added more sophisticated data analysis was needed to understand the data. A feature of WeBS data is the number of missing values, which are unlikely to be ignorable as they are not missing completely at random (Schafer, 1997). For example, all the time points before a site was included in WeBS are missing; more sites are counted in January than in other months as this month's data contributes to the International Waterfowl Census; and sometimes external circumstances such as very poor weather or the 2001 Foot and Mouth outbreak in Britain prevent the majority of counts from being carried out. The philosophy behind most of the analyses done on WeBS data is the premise that if there were no missing counts, then the survey would be a complete census of the population of most of the species involved.

Ogilvie Index

Prior to 1992, the chain method developed by Ogilvie (1967) was used to generate a population index for the Wetland Bird Survey. This method was also used at one time for other bird surveys such as the UK Common Bird Census and Constant Effort Sites scheme. Ogilvie's method uses just one count per year; the January WeBS core count was used as January surveys are prioritised as part of the International Waterfowl Census. A missing value for a site means the site is omitted from the index calculation for three years. The method is as follows.

Let,

$$Y_j = \sum_{i=1}^I x_{ij} \delta_{j-1,j}^i \quad (1.1)$$

and

$$Z_j = \sum_{i=1}^I x_{ij} \delta_{j,j+1}^i \quad (1.2)$$

where x_{ij} is the count at site i in year j and we have defined,

$$\delta_{j-1,j}^i = \begin{cases} 0 & \text{if } x_{i(j-1)} \text{ or } x_{ij} \text{ is missing} \\ 1 & \text{otherwise.} \end{cases} \quad (1.3)$$

For year j , the Ogilvie index, O_j , compared to that in a base year, b , is defined as:

$$O_j = \begin{cases} O_{j+1} \frac{Z_j}{Y_{j+1}} & \text{for } 0 \leq j < b \\ 1 & \text{for } j = b \\ O_{j-1} \frac{Y_j}{Z_{j-1}} & \text{for } b < j \leq J. \end{cases} \quad (1.4)$$

Table 1.2 shows the calculation of the Ogilvie Index (O_j) for a small subset of Mute Swan WeBS data.

Note that on a complete data set with no missing data,

$$Y_j = Z_j = \sum_{i=1}^I x_{ij}; \quad (1.5)$$

and consequently,

$$O_j = \frac{Y_j}{Y_{j+1}} \cdot \frac{Y_{j+1}}{Y_{j+2}} \cdot \frac{Y_{j+2}}{Y_{j+3}} \cdots \frac{Y_{b-1}}{Y_b} \quad \text{for } j < b. \quad (1.6)$$

Therefore, on a complete data set, Equation 1.4 reduces to a simple index comparing the current sum over sites with that in the base year:

$$O_j = \frac{\sum_{i=1}^I x_{ij}}{\sum_{i=1}^I x_{ib}}. \quad (1.7)$$

Year j	Site i					Sum		Ogilvie				Underhill	
	1	2	3	4	5	S_j	Y_j	Z_j	ratio	O_j	$U1_j$	$U2_j$	
1998	936	734	432		382	2484	0.81		2484	1.05	0.83	0.81	1.15
1999	976	671	283		439	2369	0.77	2369	2369	0.80	0.79	0.77	0.88
2000	1169	1011	439	341	332	3292	1.07	2951	3292	1.07	0.99	0.96	1.10
2001	807	1110	330	229	597	3073	1.00	3073	2844	0.92	0.92	0.89	0.90
2002	1075	941	555		511	3082	1.00	3082	3082		1.00	1.00	1.00
2003	1368	865	189		672	3094	1.00	3094	3094	1.00	1.00	1.00	1.16
2004	871	801	261	265	414	2612	0.85	2347	2612	0.76	0.76	0.76	0.80
2005	772	883	533	300	467	2955	0.96	2955	2655	1.13	0.86	0.86	0.86
2006	888	1024	422		221	2555	0.83	2555	2555	0.96	0.83	0.83	0.83
2007	920	1164	135	379	205	2803	0.91	2424		0.95	0.79	0.82	0.82

Table 1.2: Comparison of Ogilvie and Underhill Index methods. January Mute Swan counts for the sites ‘Fleet and Wey’(1), ‘Somerset Levels’(2), ‘Ouse Washes’(3), ‘Loch Bee (South Uist)’(4) and ‘Loch of Harray’(5) for 1998 to 2007. Counts at Loch Bee are missing for 1998, 1999, 2002, 2003 and 2006. Counts marked as incomplete are in *italics*. S_j compares the sum of the counts over the sites in year j with the base year 2002, but is affected by the missing counts. Ogilvie’s method for dealing with missing data uses Y_j , Z_j and their ratio to calculate the index O_j (see text). Underhill’s method gives $U2_j$; or $U1_j$ if incomplete counts are assumed complete.

Underhill Index

In the early 1990s, concerns were expressed with some of the characteristics of the Ogilvie method (Ter Braak et al., 1992; Underhill and Prÿs-Jones, 1994). Disadvantages of the Ogilvie method include:

- Data from sites can only be used when the site has been visited in two successive years, so that the method does not use all the available information. The short-term loss of an important site can have a persistent effect on the index number series.
- Index numbers based on paired counts in adjacent years can show spurious trends due to random walk behaviour (Geissler and Noon, 1981).
- Information about the quality (completeness) of counts is not taken into account.
- Only one month of data can be used.

Due to the shortcomings of the chain approach, Underhill (1989) developed an iterative algorithm to impute the missing counts using a multiplicative model of site, year and

month effects. This allowed more than one month's data to be used for each species and also for the quality of the count to be part of the imputing process. Table 1.2 compares the Ogilvie Index (O_j) for a small subset of Mute Swan WeBS data with the Underhill Index ($U2_j$). The Underhill Index ignoring the quality of the counts is also given ($U1_j$). Ogilvie's method is an improvement on the index based on the sum of available counts (S_j). In this particular example O_j and $U1_j$ are very similar, but $U2_j$ is different in years where there are incomplete counts.

As this is the main analysis method currently used with WeBS data an overview is given in Sections 1.3.2 and 1.3.3 below and more detail follows in section 4.2. A recent topic of WeBS methodological research is to use trends in the Underhill population indices to generate "alerts" when there appears to be evidence of a major decline in a species' population. An introduction to the method used is in Section 3.3.2.

Some research has been done on an extension to the Underhill model, splitting sites into different regions or habitat types to enable more realistic estimation of missing counts (Austin et al., 2000; Bell, 1994b) but this type of analysis is not done routinely.

Estimating Population Size

As well as estimating population trends, WeBS data is used in estimating population sizes of waterbird species in conjunction with other survey methods (Kershaw and Cranswick, 2003a,b; Kirby, 1995; Rehfish et al., 2003). A WeBS-based estimate of population size is made using the mean count over five years once missing counts have been imputed using the Underhill method. The mean count is multiplied by a species-specific "extrapolation factor" to generate a population estimate.

The need for absolute population estimates is largely due to their use in designating sites as important on an international basis according to the Ramsar Convention on Wetlands of International Importance by regularly supporting 1% of the individuals of a population of species or subspecies of waterbird (Matthews, 1993). Any site holding more than

20,000 waterbirds is also classified as internationally important. On a national basis, sites holding more than 1% of the estimated population of a species or sub-species of waterbird are classed as nationally important and afforded Special Protection Area (SPA) status under the EC Birds Directive. Presently there are 144 Ramsar sites and 246 SPAs in the UK (Collier et al., 2005).

While estimating wildfowl abundance is very important, the focus of this thesis is on the estimation and quantification of population trends as knowledge of whether populations are stable, increasing or declining is important for their management.

Annual Report

Waterbirds in the UK 2006/07: The Wetland Bird Survey (Austin et al., 2008) is the latest in the series of usually annually published books reporting the latest WeBS results. A synopsis of the survey's history, aims and methods is included, together with species accounts for all species recorded that year, such as the one in Figure 1.11. In addition to a description of the species status, the following statistics are given:

- **Population Index Plot** Time series plots of GB and NI annual indices from 1965-66 to the latest year, accompanied by a smooth trend. This gives an “at a glance” view of the how the species’ population is increasing, decreasing or fluctuating. A smoothed trend aids visual interpretation.
- **Monthly means chart** Bar chart for GB and NI of the last season’s monthly means compared with that calculated over the previous 5 years. This indicates if there has been any unusual phenological changes for the species in that year, such as delayed departure times, which may have affected the population index.
- **Maximum counts** When national totals are calculated no missing values are imputed and incomplete counts are treated as complete counts: these are calculated for all months and presented in a separate table to the species accounts. The maximum over the year is shown in the species account. For some goose species

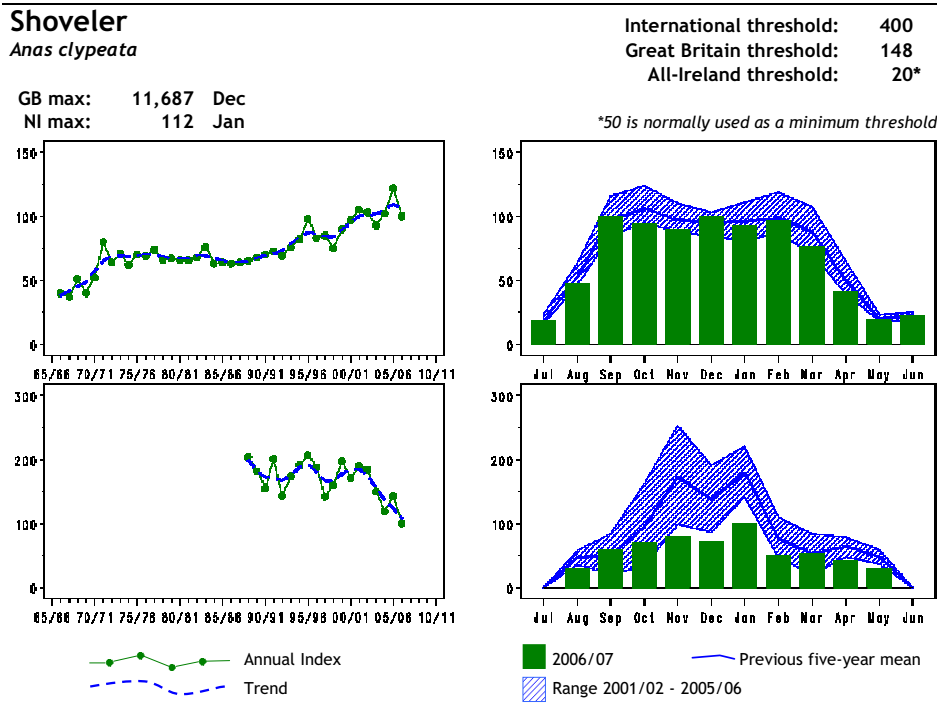


Figure 25.a, Annual indices & trend for Shoveler for GB (above) & NI (below).

Figure 25.b, Monthly indices for Shoveler for GB (above) & NI (below).

The British maximum was around 5% lower than the previous year. This was reflected in the national index, which fell slightly on that of 2005/06, although numbers were similar to the average of the past few years. The monthly indices suggest that average numbers arrived by September, but then were slightly below average for much of the rest of the winter (except in December). The index for Northern Ireland also fell slightly, reaching an all-time low (albeit the numbers of birds involved being far lower than in Britain). The monthly indices for the region suggest that numbers were particularly low between November and February. Numbers at the region's key site, Strangford Lough, were half those of the five-year mean.

Numbers at the Severn Estuary remained high for the second year running resulting in average numbers above those of the

internationally important threshold. In contrast, mean numbers at Dungeness Gravel Pits have fallen below this threshold, although the site remains of national importance for this species. One of the most noticeable counts was at the Ouse Washes where peak numbers were a third those of the five-year mean. This decline is likely to be due to unfavourable water levels at the site, which were particularly suitable for wildfowl during the previous year. Other notably low counts were made at Chew Valley Lake, Abberton Reservoir, the Burry Inlet, Staines Reservoirs and Chichester Gravel Pits. Three sites no longer qualify as supporting nationally important numbers, with the most dramatic decline having been at Wraysbury Gravel Pits; anecdotal evidence suggests that increased disturbance at this site may have been an important factor.

	02/03	03/04	04/05	05/06	06/07	Mon	Mean
Sites of international importance in the UK							
Ouse Washes	1,125	1,104 ¹²	2,725 ¹²	1,548 ¹²	696 ¹²	Nov	1,440
Somerset Levels	(2,190)	784	(902)	845	1,520	Feb	1,335
Rutland Water	504	475	663	680	495	Sep	563
Chew Valley Lake	535	565	395	660	300	Oct	491
Abberton Reservoir	422	488	355	(674)	(152)	Aug	485

Figure 1.11: Example species account from Waterbirds in the UK. Austin et al. (2008); reproduced with permission.

with separate census schemes data from these are used rather than the WeBS core counts.

- **Importance Thresholds** Sites containing numbers greater than the International, Great Britain and All-Ireland thresholds are considered nationally or internationally important for the species and these sites are then listed in the species account. Information from available counts outside of WeBS is also used in assessing site importance.

1.3.2 Underhill's Indexing Method

The Underhill model was developed in 1989 following concerns about the usefulness of Ogilvie's chain index (Underhill, 1989; Underhill and Prÿs-Jones, 1994). It was announced in the 1989/90 Wildfowl and Wader Counts annual report and used to give annual indices in the annual report from 1993 onwards, coinciding with the formation of the Wetland Bird Survey from the BoEE and NWC (Kirby et al., 1990; Waters and Cran-
swick, 1993). The model is used on WeBS data to generate population indices for both wader species (Prÿs-Jones et al., 1994) and wildfowl species (Kirby et al., 1995). The concept behind the index is that for every species a constant unknown proportion, p , of the British overwintering population, B , occurs on the subset of WeBS sites chosen using the site inclusion criterion described in section 1.3.3 below. By expressing the population in year j as a proportion or percentage of that in an arbitrary base year, B_{base} , trends in the population of these sites will reflect trends in the overall British population:

$$\frac{pB_j}{pB_{base}} = \frac{B_j}{B_{base}} = I_j. \quad (1.8)$$

It is not necessary to know what p is, as long as it is constant over the whole time period, to obtain I_j , the population index.

Note that here the "population" B is the mean (or, equivalently, total) population over a

different subset of months and of sites for each species, chosen using the method detailed in Section 1.3.3 below. Chapters 2 and 3 will investigate the most appropriate sets of months and sites to use and Chapter 5 will look at alternative ways of combining monthly counts into a single annual figure.

To estimate the population on the subset of WeBS sites, Underhill developed an algorithm to impute missing values in the data set. A multiplicative model is used where the count of the number of birds at site i in year j and month k is expected to be the product of a site factor, a year factor and a month factor:

$$\mathbb{E}[x_{ijk}] = s_i y_j m_k \quad (1.9)$$

where $i = 1, 2, \dots, I$; $j = 1, 2, \dots, J$ and $k = 1, 2, \dots, K$.

The model is fitted iteratively with a modified EM algorithm that imputes fitted values for missing counts and the larger of the fitted value and the observed value for incomplete counts. Parameters are estimated using Poisson quasi-likelihood estimates; consequently, when the algorithm has converged, the year factors y_j are the index values I_j : that is, the ratio of the sum of counts over all sites and months in year j to that in the base year (usually taken to be the final year). Similarly, the month factors m_k are the ratio of the sum of counts over all sites and years in month k to that in the base month (giving a “month index”). The site factors can be used as a measure of the influence of a site on the resultant year and month indices.

As there are no interaction terms, the model assumes that all sites exhibit the same trend, as do all months, and that seasonality does not vary between sites or over time. Chapter 4 has a more detailed description of the Underhill model and algorithm, and considers whether these assumptions are realistic.

1.3.3 Criteria for data inclusion in annual Underhill Index

50% Completeness Criterion for sites

When the Underhill Index was developed, it was noted that the strong assumptions in the model used for imputing missing values made it undesirable that a high proportion of counts should be missing when producing the population index. Underhill (1989) investigated the performance of the algorithm for missing counts on various subsets of all sites within WeBS. As described in section 1.2.3, each count has a “quality” code associated with it to describe if the count is complete, incomplete or missing. It was recommended that only sites where less than half of the counts are incomplete or missing be used to calculate the index for wader species.

Kirby et al. (1995) researched the effects of different criteria levels for use with wild-fowl species but concluded that there was no reason to use a different criterion to the wader species. They looked at the months September to March over the period 1960/61 to 1991/92. There were 5648 sites and 83.5% of the possible data values were missing. Of the completeness criteria they looked at, the maximum number of sites selected was 657, which were chosen using a criterion of 40% over 32 years and 7 months; this corresponded to only 11.6% of the sites in the database. The 50% criterion over the reduced time period 1966/67 to 1991/92 (26 years) and September to March (7 months) selected 596 sites, 10.6% of the sites in the database. This resulted in a data set in which 23.7% of counts were missing or incomplete.

As the quality code and the number of months used in the index varies according to species, sites may be used to calculate an index for one species and not another. In the production of a population index trend plot, the same set of sites are used. However, the set of sites can vary when recalculating the indices from one year to the next, depending on the precise proportion of missing values over the total data set. The biases introduced by this method of post-selection of sites are investigated in Section 2.4.

Years

Although wildfowl counts have been recorded since the 1940s, wildfowl NWC data from before 1960 has not yet been computerised and so is not included in the indexing process. Kirby et al. (1995) compared the amount of imputation needed for wildfowl data from the period 1960-61 to 1991-92 with the period 1966-67 to 1991-92. The number of sites selected and the proportion of counts imputed were 507 and 25.7% when the first six years were included and 596 and 23.7% when they were not; the average proportion of counts imputed per year for the first six seasons was 51.5% compared with 20.4% for the subsequent 26 seasons. Concluding that the coverage was worse from 1960-61 to 1965-66, they suggested wildfowl indices be based on data from 1966-67 onwards, so that less imputation was necessary.

Months

Annual indices between the 1992/93 WeBS annual report (Waters and Cranswick, 1993) and the 2007/08 WeBS annual report (Musgrove et al., 2007) used the month selections from Prÿs-Jones et al. (1994) for waders and Kirby et al. (1995) for wildfowl. The mid-winter period, defined as December to February, was used for all wader species as being the period when their wintering populations are most stable. For wildfowl, however, the months used for Underhill indexing vary according to species, since seasonality varies from species to species as shown in Figures 1.3, 1.5 and 1.7 in section 1.1.2.

Kirby et al. (1995) chose which months to include for each species by using their seasonality as a guide. For each species, month indices, m_k were produced by the Underhill algorithm on all seven winter months (September to March). The month with the highest index value was used, together with any adjacent months with overlapping consistency intervals, produced by bootstrapping over sites (see section 4.3.4). For example, species whose mean peak occurred in December at all sites would only use December for indexing, whereas in species where the mean peak was in December, January and February in

roughly equal numbers of sites those months would be used for indexing.

The WeBS indexing methodology was changed in 2007. To bring the annual index methodology in line with practice in other uses of WeBS data, WeBS alerts (see Section 3.3.2) and estimates of population size based on five-year means of peak abundance, wader species now use the months November to March inclusive, and wildfowl September to March inclusive (Musgrove et al., 2007). The change for wildfowl was also motivated by concerns over the effect of phenological change on indices. The question of which months to use for an annual index is an interesting one, which will be looked at in detail in Chapters 3 and 5.

Summary of selection of data for indexing

The procedure for data selection from the WeBS database is as follows:

1. Select data from sites in the region(s) required, for example Wales or Great Britain.
2. For the species you are interested in, select data from the appropriate years.
3. Select data from the months appropriate for that species.
4. Select data from sites which have at least half the possible number of complete counts. (The total possible number of counts for each site is the selected number of months multiplied by the selected number of years.)

The data thus selected are used to calculate the Underhill Index. A MATLAB implementation of this procedure is shown in Program 1.1 below.

1.3.4 Working with WeBS data

The British Trust for Ornithology currently maintain the database of Wetland Bird Survey Data. Data was exported for analysis in this project in the form of fixed column width

text files (Table 1.3 shows a few lines of one such file). In the text files there is a line for every count, including ‘real zeros’ (when a site visit was made, but the species was not recorded as present) but not for missing site visits.

59YYY	2006	8	1	45	GN	ENG
59YYY	2006	9	2	57	GN	ENG
59YYY	2006	10	1	28	GN	ENG
59YYY	2006	11	2	0	GN	ENG
59YYY	2006	12	2	0	GN	ENG
60021	1983	4	1	0	GN	WAL
60021	1983	5	1	0	GN	WAL
60021	1983	6	1	0	GN	WAL
60021	1984	4	1	0	GN	WAL
60021	1984	5	1	0	GN	WAL

Table 1.3: Example of part of the 2006 WeBS text file for Goldeneye, ‘gn4wing.2006’. The first column is a unique site identifier that is often numeric but can contain letters. This is followed by the year, month, quality code, count, two letter species code and the region of the UK where the site is located.

When information is needed on all sites, such as in Figure 1.1 on page 2 of the number of sites surveyed each year, and Figure 2.6 on page 56 overlaying site locations with elevation data, Mallard is used an indicator of recording effort. This is because access to the full database of site visits was not possible and so information from the WeBS text files had to be used. Since Mallard is a widespread species it is likely to have occurred at least once on most sites; consequently the Mallard file will contain zeros for other dates when a WeBS count took place but no Mallard were present.

A table of information for each site identifier was available, with the name of the site, the grid reference of the site centroid and, for complex sites, the site identifier for the site above it in the site hierarchy. For example, the site 59YYY in Table 1.3 is Solway Estuary (English counties) at grid reference NY165565 and is part of site 59400, Solway Estuary. Site 60021 is St Pierre Lake at ST515905 and is not part of a larger site. Note that it is important that sites within a complex hierarchy are only used once in an analysis. Provisional geographic information was also available as a GIS polygon layer of sites.

Data manipulation and analysis and the production of figures for this thesis was carried

out in the numerical computing environment and language, MATLAB[®] (The MathWorks, 2008). Program 1.1 shows code to import data from a text file in the format of table 1.3, remove data for sites, years and months that are not of interest, and add extra lines for missing data. To avoid the extraction issue with Ribble Estuary counts (Section 1.2.3) affecting results, where the 2008 extraction data is used in this thesis, the omitted Ribble data for the 1975/76 and 1979/80 to 1988/89 seasons are added from the 2004 extraction data.

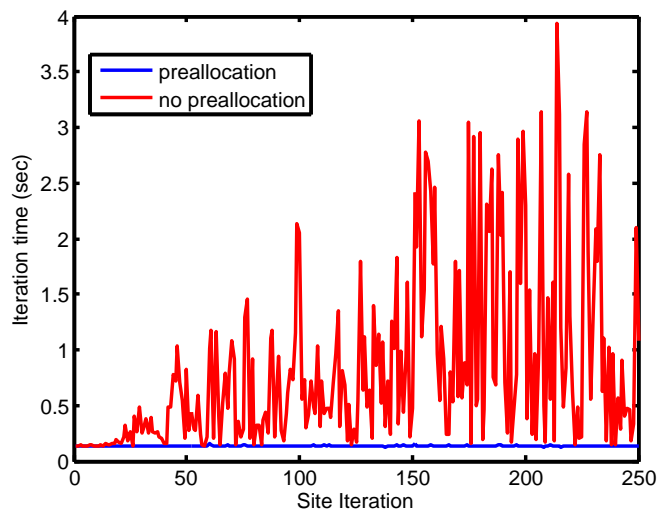


Figure 1.12: Iteration time effect of preallocation of MATLAB array in the code in Program 1.2. Neglecting to preallocate variables causes a time penalty which tends to increase with each iteration.

Program 1.1 MATLAB code to import WeBS data and select the data to be analysed.
(Program continues on the following page.)

```

%Input: User specifies two letter species code. The corresponding textfile
%such as pt4wing.2006, with site number in 1st column, year in 2nd, month
%in 3rd count in 4th and quality in 5th. Parameters setting the months,
%years and regions to be included may be changed in the script.
%
%Filters first on regions, then years, then months. Also filters out sites
%with <50% coverage if specified by user.
%
%Output: "newdata" as a matrix. This expands the original data to give a
%row for each site in each year and each month, with the recorded data
%where it exists and denoting missing values by a -1 and 3 in the count and
%quality columns.
%
%Also creates sitestoinclude, a lookup table of sitecodes and names with
%the row number corresponding to site integer in newdata.
%
%%
%%Set parameters
%%
%%Choose months/years/regions to use:
regtoinclude = {'ENG'; 'SCO'; 'WAL'; 'XXX'};

firstmonth=1; %1-Jul; 2-Aug; 3-Sep; 4-Oct; 5-Nov; 6-Dec;
lastmonth=12; %7-Jan; 8-Feb; 9-Mar; 10-Apr; 11-May; 12-Jun.
startyear = 1966;
endyear = 2006;

nyears=1+endyear-startyear; nmonths=1+lastmonth-firstmonth;

%Asks user for the part of the filename that refers to the species, e.g. tu
%or e_
spcode = input('Please enter the BTO two-letter code for the species: ', 's');
reply = input('If you want to include all sites press A and enter; \n Press C
    for sites with 50% completeness only; \n Press any other key if you want to
    use the variable \n "sitestoinclude" already in the workspace: ', 's');

%%
%Read in data and remove extra data.
location = ['D:\WeBS data\' , spcode, '4wing.2006'];
[site year month quality count species region]=textread(location, '%5s%4n%2n%1n%6
n%2s%3s');
load sitenamesmap sitenamesmap
data = [year month count quality];
clear year month quality count species spcode location
%Select Regions Required:
var1 = ismember(region, regtoinclude);
data = data(var1,:); site = site(var1); region = region(var1);
clear i var1 regtoinclude
%Replace Ribble counts for 1974-1988
[site data region] = getRibblecounts(spcode, site, data, region);
clear spcode

```

Program 1.1 (continued)

```

%Select Years Required:
rows = find(data(:,1)>=startyear & data(:,1)<=endyear);
data = data(rows,:); site = site(rows); region = region(rows);

%Select Months Required:
rows = find(data(:,2)>=firstmonth & data(:,2)<=lastmonth);
data=data(rows,:); site=site(rows); region = region(rows);
%Renumber months so first month = 1:
data(:,2)=data(:,2)-(firstmonth-1);
clear rows firstmonth lastmonth

%%
if exist('sitestoinclude')
    [sitestoinclude site data region] = getsitestoinclude(reply,site,data,region
        ,nyears,nmonths,sitestoinclude);
else
    [sitestoinclude site data region] = getsitestoinclude(reply,site,data,region,
        nyears,nmonths);
end
clear reply
%%
%set up new datamatrix with all sites, years and months ready to fill with data
nsites = length(sitestoinclude)
n = nsites*nyears*nmonths;
newdata=zeros(n,5);
%Generate site, year and month values:
newdata(:,1)= reshape((repmat((1:nsites)',1,nyears*nmonths))',n,1);
newdata(:,2)=repmat(reshape((repmat((startyear:(startyear+nyears-1))',1,nmonths)
    )',nyears*nmonths,1),nsites,1);
newdata(:,3)=(repmat(1:nmonths,1,nsites*nyears))';
%Default count and quality values -1 and 3 respectively:
newdata(:,4)=-1; newdata(:,5)=3;

%Construct unique key for original data and newdata matrices:
[s2,unqsitemax,siten] = unique(site);
clear s2 unqsitemax
key1 = siten*1000000+data(:,1)*100+data(:,2);
key2 = newdata(:,1)*1000000+newdata(:,2)*100+newdata(:,3);
%Check keys are unique:
if length(unique(key1))-length(key1)~=0 || length(unique(key2))-length(key2)~=0
    error('Unique key generation has failed, correct code!')
end

%Fill in newdata for non-missing rows:
rows = ismember(key2,key1);
newdata(rows,4) = data(:,3); newdata(rows,5) = data(:,4);

%Get list of sitecodes and names for included sites:
rows = isKey(sitenamesmap,sitestoinclude);
sitestoinclude(rows,2) = sitenamesmap.values(sitestoinclude(rows));

clear startyear endyear siten nsites nyears nmonths n key1 key2 rows
sitenamesmap

```

The name MATLAB is derived from *matrix laboratory*, and algorithms written in its high-level language work best by vectorising the code, replacing `for` loops wherever possible, since MATLAB is designed for matrix operations. If it is not possible to vectorise code, then it is important to preallocate vectors and arrays as the code will execute slowly if the variable is allowed to grow within the loop. This process is illustrated in the code in Program 1.2; the first, MATLAB-naive, approach takes 207 seconds* on a WeBS data set of 12 months, 41 years and 250 sites. Simply preallocating the matrix in the loop, *newdata*, reduces the time to 33 seconds and removing the loop completely by rewriting the code reduces it to 0.85 seconds, less than 0.5% of the original execution time. The penalty for neglecting to preallocate arrays increases as the the number of iterations increases; the time taken for each site iteration in the loops in Program 1.2 is shown in Figure 1.12, illustrating a generally increasing trend when *newdata* has not been preallocated. Similarly, doubling the number of sites from 250 to 500 takes the loop without preallocation over three and a half times as long, at 746 seconds, but with preallocation the time effect is linear, doubling to 68 seconds. However, vectorising the code results in minimal extra time to execute, with the doubling of the number of sites taking only 10% longer to execute, at 0.93 seconds.

*On a computer with a dual core processor (2.16GHz) and 4 GB DDR2 memory.

Program 1.2 MATLAB code for three methods of transferring information from the original *data* variable to *newdata* in Program 1.1. The first approach uses a loop; the second also uses a loop, but the *newdata* array has been preallocated; and the third vectorises the process and is that used in Program 1.1.

```

%No preallocation or vectorisation:

for i = 1:nsites
    tic
    originalrows = strmatch(sitestoinclude(i),site,'exact');
    for j = 1:length(originalrows)
        currentyr = data(originalrows(j),1);
        currentmth = data(originalrows(j),2);
        newrow = (i-1)*nmonths*nyears + (currentyr-startyear)*nmonths +
            currentmth;
        newdata1(newrow,4) = data(originalrows(j),3);
        newdata1(newrow,5) = data(originalrows(j),4);
    end
    slowtime(i)=toc;
end

%Preallocation
newdata2=zeros(n,5);
medtime = zeros(nsites,1);
for i = 1:nsites
    tic
    originalrows = strmatch(sitestoinclude(i),site,'exact');
    for j = 1:length(originalrows)
        currentyr = data(originalrows(j),1);
        currentmth = data(originalrows(j),2);
        newrow = (i-1)*nmonths*nyears + (currentyr-startyear)*nmonths +
            currentmth;
        newdata2(newrow,4) = data(originalrows(j),3);
        newdata2(newrow,5) = data(originalrows(j),4);
    end
    medtime(i)=toc;
end

%Vectorisation
tic
var = repmat({'N'},length(sitelist),1);
var(ismember(sitelist,sitestoinclude)) = {'Y'};
compinmap = containers.Map(sitelist,var);
rows = strmatch('Y',compinmap.values(site));
data = data(rows,:);
site = site(rows);
[s2,unqsitemax,siten] = unique(site);
newdata3=zeros(n,5);
newdata3(:,1)= reshape((repmat((1:nsites)',1,nyears*nmonths))',n,1);
newdata3(:,2)=repmat(reshape((repmat((startyear:(startyear+nyears-1))',1,nmonths
)')',nyears*nmonths,1),nsites,1);
newdata3(:,3)=(repmat(1:nmonths,1,nsites*nyears))';
key1 = siten*1000000+data(:,1)*100+data(:,2);
key2 = newdata3(:,1)*1000000+newdata3(:,2)*100+newdata3(:,3);
rows = ismember(key2,key1);
newdata3(rows,4) = data(:,3);
newdata3(rows,5) = data(:,4);
fasttime = toc;

```

1.4 Thesis Outline

The Wetland Bird Survey (WeBS) aims to monitor all non-breeding waterbirds in the UK to provide the principal data on which the conservation of their populations and wetland habitats is based. To this end, WeBS has three main objectives:

- to assess the size of non-breeding waterbird populations in the UK;
- to assess trends in their numbers and distribution; and
- to assess the importance of individual sites for waterbirds.

These results also form the basis for informed decision-making by conservation bodies, planners and developers and contribute to the sustainable and wise use and management of wetlands and their dependent waterbirds.

Austin et al. (2008)

This thesis investigates fulfilling the second of these objectives, assessing trends in numbers, in the case of wildfowl. Four components to this problem are addressed:

1. Does current WeBS methodology give an adequate sample of sites for estimating national trends?
2. What within-year seasonality patterns in population abundance are there? Is phenological change occurring? How does this affect assessing trends in non-breeding wildfowl?
3. What about missing data: short runs, caused by recorders not being able to record, and long runs caused by sites entering and leaving the scheme; and how should incomplete counts caused by external factors such as weather or disturbance be dealt with? Should missing data be imputed, and what method should be used?
4. What index aggregation formula is appropriate to combine counts from sites and months?

This chapter began with the motivation for monitoring wintering wildfowl and introduced species variation in seasonal population abundance in Section 1.1. The Wetland Bird Survey as the primary means for collecting information on wintering wetland birds was introduced in Section 1.2.

Section 1.3 gave an overview of the methodology of the Wetland Bird Survey and how the four components to producing an annual index given above are currently handled. In summary:

1. The WeBS sample is constructed by locally informed selection of waterbodies which are known to contain large congregations of waterbirds. For indexing purposes, a second post-survey sample is then made from the available site data, by stipulating that only sites that fulfil the 50% completeness criterion (that at least half of the available counts are of good quality, i.e. neither missing nor incomplete) are included in the index.
2. Seasonality was assessed by Kirby et al. (1995) to obtain stable months, resulting in individual month selections for each wildfowl species. Uncertainty about phenological change motivated a change to use of all data from September to March for all wildfowl species (Musgrove et al., 2007).
3. Missing and incomplete counts are imputed using the Underhill algorithm: a modified EM algorithm for the Underhill model. The Underhill model is a Generalised Linear Model which assumes the count is a product of a month factor, a year factor and a site factor, with no interactions, and can be modelled by a quasi-Poisson distribution. The algorithm is modified to impute incomplete counts (i.e. those flagged as poor quality) only when the modelled count is greater than the observed count.
4. Counts for the months selected by seasonality and sites selected by the 50% completeness criterion, with missing and incomplete counts imputed by the Underhill

algorithm, are summed. Dividing the summed counts in the current year by that in the base year gives the Underhill Index.

The following four chapters consider each of these four factors in producing a useful population index in turn. In Chapter 2 a closer look is taken at the sampling methods used in WeBS. Sources of bias in the survey methods and their effects are discussed. Additionally, the impact of post-selection of sites by the proportion of missing values is investigated, and a new selection procedure proposed.

Chapter 3 uses Functional Data Analysis methods to explore spatiotemporal variation in wildfowl WeBS data. A particular problem considered is whether it is appropriate to use all winter months' data for all species or a subset of months.

Chapter 4 details the Underhill method of imputing missing values using a site-year-month multiplicative model, to give a complete data set from which to calculate a population index. The Underhill method is compared with a Generalised Linear Modelling approach. Results from using the sites and months selected in Chapters 2 and 3 are compared with no restrictions on sites and months by a revised method for estimating site consistency intervals, by refitting the model for bootstrap samples.

Chapter 5 considers indexing approaches for combining site and month counts, inspired by index number theory from economics. Alternative aggregation formulae for monthly counts are presented and compared using WeBS data.

Chapter 6 introduces a WeBS simulation study to further explore some of the key concepts of earlier chapters. Simulated data is used to compare the effect of phenological change on indices with different month aggregation formulae and the impact of missing data on estimating trends is explored.

The thesis concludes with a discussion of the issues raised and presents recommendations for further research into analysis and presentation of annual indices for the Wetland Bird Survey.

Chapter 2

SURVEYING WILDFOWL ON A NATIONAL SCALE

2.1 Introduction

As introduced in Chapter 1, the Wetland Bird Survey (WeBS) is an example of a biological monitoring scheme with a sample that has grown organically over time, rather than being designed using statistical sampling theory. Sites have been added to the scheme over the decades following two implicit criteria. The first is that resources are available to survey the site: there is a volunteer who is willing to commit to visiting the site once a month and who has sufficient waterbird identification skills to be able to count different species. The second is that the site is considered to be worth surveying, on a local level, by the volunteer counter and the local WeBS organiser, as it is known to be used by waterbirds.

In this chapter, we consider what biases may be present in wildfowl population indices due to the current construction of the sample of sites surveyed and the sample used for indexing.

2.2 Detectability and Measurement Error in Wildfowl Counts

2.2.1 *Detectability*

In Chapter 1 the concept of a population index was defined in equation 1.8. The key assumption that makes a population index useful is that a constant proportion of the total population is always recorded, whenever a sample is taken. The proportion, p , of the total British population that is counted in a WeBS sample can be broken down into the proportion of the population that is present on the sites surveyed and the probability that these individuals will be detected. In order for the index to reflect trends in the British population, both of these must remain constant, otherwise the index will be biased (Williams et al., 2002).

Many statistical ecological methods have been developed to overcome non-uniform detection probabilities, and to estimate detection probabilities in order to obtain population estimates, particularly for closed populations. These include manipulating catch effort, mark-recapture, and distance sampling (Borchers et al., 2002).

As well as producing national trends, Wetland Bird Survey data is used to estimate wildfowl national abundance and site abundance, the later being focused on selecting sites for protection under the Ramsar and Special Protection Area designations (Jackson et al., 2004c,b). To obtain national wildfowl estimates Kershaw and Cranswick (2003a) used extrapolation factors, based on the proportion of birds counted intensive regional wetland surveys in the 1990s that were present on WeBS sites (Quinn and Kirby, 1992). However, an estimate of species detection functions was not attempted, beyond noting that some species are more secretive than others. When estimating population sizes, varying species detectability must be taken into account: individuals of secretive species such as Teal are likely to be overlooked when undertaking a WeBS survey, but flocks of mobile species may be counted more than once, particularly if they move between count units (Jackson et al., 2004a).

WeBS national indices are only useful if there is insufficient variation in detection probabilities to affect the index significantly. It is undoubtedly true that detectability varies between WeBS samples. For example, birds often become more secretive when moulting, so detection may vary from month to month. In addition, it is easier to see all the birds present on open sites, with little marginal vegetation, than sites with inlets and areas of open water hidden by reeds. If counts are summed across sites and months (as at present) then so long as these effects are constant from year to year they will not affect the index. However, there are also some factors which seem more likely to confound true year to year changes in abundance: observer skill; optics used; and weather and disturbance affecting counts.

2.2.2 *Optical equipment*

But if, in the old days, the bins were bad, the scopes were even worse. In the early 1970s most people relied solely on their bins, but if they were serious then they had to make a choice between one of only three new telescope models... The Nickel was dreadful, it was like looking down the cardboard tube from a toilet roll. The Hertel, the scope I eventually bought, was much better. It was like looking down the cardboard tube from a kitchen roll... optically the best of the three scope options was the Swift Telemaster. The only disadvantage was it was so short. In those days you didn't have a tripod to stand it on... There was one way to alleviate the problem and that was to obtain an old brass telescope...

Massive improvements in technology during the last twenty years have had a revolutionary effect on birding.

Cocker (2001)

If there has been a trend in detectability over the time period of the survey, this may bias the index for some species. An increase in the index in the first half of the time series, followed by more stable numbers, is noticeable in Shelduck, Shoveler and Tufted Duck, but there are also many species where the effect may be hidden by a real increase or decrease in population (see Figure 2.15 on page 73 for index plots of each species).

Wintering waterbird biodiversity indicator

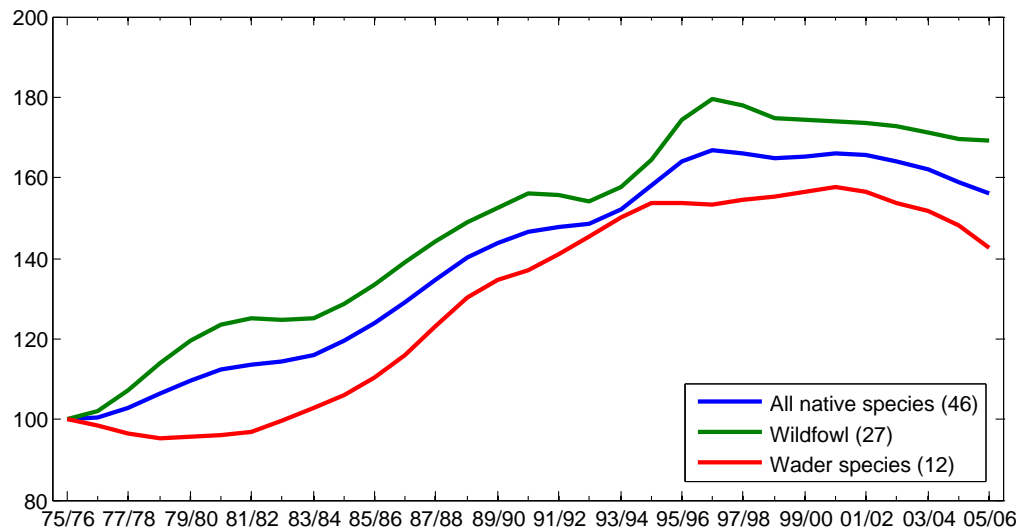


Figure 2.1: Wintering Waterbirds England Biodiversity Indicator: 1975/76-2005/06. Data Source: Defra (2007).

Composite trend indicators such as for birds are seen as a simple way of measuring progress in halting the loss of biodiversity by informing decision makers, with the caveat that it is unclear how representative bird indicators are of other taxa. Wintering bird indicators are particularly useful for showing how changes in climate or habitat impact on numbers and distribution and also for relating to breeding population trends and demographic parameters such as survival (Gregory et al., 2008).

In 2003 the Department for Environment, Food and Rural Affairs published a baseline report for biodiversity indicators, the purpose of which “is to help monitor the implementation of the England Biodiversity Strategy” (Defra, 2003). One of these was a wetland bird indicator, which was later split into two separate indicators, one for breeding wetland birds and one for wintering wetland birds (Defra, 2009). The biodiversity index values up to 2005/06 are available online and are shown in Figure 2.1; the wildfowl indicator is calculated by the geometric mean of individual species’ trends, with the species included varying as research has continued. The overall trend in the indicator has not changed as

more data has been added, however, peaking in the winter of 1996/97 as in Figure 2.1. Defra (2006) hypothesises that the overall rising trend in the wildfowl from the 1970s to the 1990s, “can be attributed to improved protection of wetlands internationally affecting migratory species, enhanced regulation of hunting, and changes in agricultural practices”. Defra (2009) also notes that large increases from some species such as Whooper Swan (+670%) and Gadwall (+568%) masked declines in others such as Mallard (-26%) and Pintail (-24%).

Whilst not dismissing these explanations for the overall pattern in the indicator, it is worth considering other possible factors. Some of the apparent increase from the mid 1970s to the mid 1990s may be the result of a bias, relatively small in individual trends, but more noticeable when brought together in the indicator, relating to the improvement in and greater accessibility of high quality telescopes and binoculars to WeBS observers as described by the quotation from Cocker (2001) above. At least one counter reportedly added as a note to their submitted WeBS count form that the numbers for all species were higher this year, as they had just got their first telescope and could see many more birds than before - and wasn't it wonderful? (WWT Research Staff, 2006). Even those surveyors who always had access to some sort of telescope will have upgraded them over time as technology has improved and become more affordable.

Further research into bird detectability with various ages of optical equipment may produce interesting results. Of all possible violations of the assumptions of non-constancy of detectability, increased probability of detection over time is most likely to confound species trends over time, with a systematic bias in one direction.

2.2.3 *Measuring counts*

Weather and Disturbance

As outlined in Section 1.2.2, acknowledgment is made that conditions such as weather or disturbance can affect detectability of birds, leading to under counting. Robinson



Figure 2.2: Wigeon (*Anas penelope*) landing on ice at Musselburgh, Edinburgh.
©Laurencea (2009)

and Pollitt (2002) found that over 68% of WeBS counters between 1995/96 and 1998/99 reported disturbance during WeBS counts, ranging from dog walkers, to shooters to large aircraft. Infrequently occurring activities are most likely to disturb birds, possibly causing movements and preventing full counting to take place.

The WeBS system of marking counts as “incomplete” when weather or disturbance has led to under counting allows these temporary changes in detectability to be modelled.

Observer

With hundreds of observers taking part in the Wetland Bird Survey each year, there will inevitably be a range of identification and number estimation skills between observers. Link and Sauer (1997) identified two types of observer effects in bird monitoring surveys: among-observer effects (baseline differences in competence), and within-observer effects (changes through time). Within-observer effects can take the form of inexperience when an observer first starts to survey a site and possibly also declines in ability as an observer ages. In the Wetland Bird Survey, observers are encouraged to seek out all the birds, searching areas thoroughly for all flocks. Inevitably observers will only learn over time favoured areas by different species at their site, so may under count to begin with. Incomplete counts are treated as a minimum number of birds present when calculating the Underhill index for WeBS since observers are very aware of the potential for double counting of flocks and studiously avoid it, but may inadvertently miss flocks in

unexpected locations within their count areas (Underhill, 1989).

Within WeBS one observer usually has responsibility for surveying a site for long stretches of time, so it is not as straightforward to attempt to model observer effects on detectability as it is to flag counts as incomplete when birds are hidden by bad weather or disturbed. In addition, the main observer at some count units is sometimes helped by a varying number of assistants, depending on who is available on the day of the count, which can affect the exact area counted as well as the skill of the assistants adding additional uncertainty about detectability. Kery and Schmid (2004) found that species detectability is affected by among-observer effects, but also depended strongly count conditions such as the quantity of birds present and weather conditions. For this reason they recommended that monitoring programmes be designed with estimating detectability in mind. Further research on the impact of detectability on trends and abundance estimates from the Wetland Bird Survey is required.

2.3 WeBS Spatial Coverage: Surveying Wildfowl in Great Britain

2.3.1 Introduction

WeBS is the amalgamation of two surveys which had similar "look-see" methodologies, one for mostly wildfowl species on inland sites (NWC) and one for mostly wader species on coastal sites (BoEE). As wildfowl and wader species occur in all types of wetland habitat the administration of the schemes was brought together in 1993 and the volunteer observers were encouraged to record all waterbirds, including wildfowl, waders, cormorants, kingfishers and gulls, with the last being optional (Cranswick et al., 1997). Few countries have a waterbird monitoring scheme on the scale of WeBS, although many countries count waterbirds once a year in January as part of the International Waterfowl Census (Rose, 1990).



Figure 2.3: Map of WeBS count sector boundaries in Great Britain. Includes all sites that have been part of the WeBS scheme, including some which have been only occasionally sampled as part of special surveys such as UK-NEWS.

Wetlands, wetland birds, and birdwatchers are all unevenly distributed across the British Isles. The WeBS count sector boundaries in Figure 2.3 give an indication of the location of WeBS counts, but only a proportion of these will have been visited in any particular year. In this section, an overview of the geographical, climatic and human reasons for this variation are explored and potential biases discussed.

2.3.2 Wetland habitats in Britain




Figure 2.4: Illustrative map indicating variety in WeBS Sites in Fife, Scotland. WeBS count sector boundaries are indicated in red, and water polygons in blue.

The British Isles are extremely important to waterbirds. This is due to the amount of coastline and temperate climate, giving ideal conditions for many species during the winter. The type and extent of wetland habitat varies considerably around the country, from large estuaries and mudflats to lowland fens, upland boggy areas and inland gravel pits. As mentioned in section 1.2.3, the third digit of WeBS site codes represent a basic habitat classification into: reservoir; estuary; river or marsh; natural lake; and coast. Examples of all these habitats can occur in a small area, as shown in Figure 2.4. One of the basic principles of ecology is that a community of species in one habitat will be quite different to that found in another. Some wildfowl species are generalists and occur in many

types of habitat, whereas others are specialists and are restricted the coast or deep inland waterbodies.

The philosophy of WeBS is to monitor as many species as possible within the same programme of work. If a species occurs mostly on a habitat that isn't well-surveyed within the scheme, or only sporadically, then an index is not published. For example, a WeBS population trend is not published for the Long-tailed Duck (*Clangula hyemalis*), which occurs mostly offshore. Although large numbers are recorded in WeBS surveys (over ten thousand were recorded in February 2007 for example, mostly at one site, the Moray Firth), the influence of weather on the location of the birds and the numbers that are visible from shore means detection probabilities vary wildly, so an annual index of the type used with other WeBS species is meaningless (Austin et al., 2008).

Elphick (2008) remarked that although ecologists seem to have good intuition about the limitations of inferences from ecological counting schemes, "the most serious errors occur where extrapolations had been made beyond the sampled population, such as when the survey was used to estimate population sizes for species that primarily occur in habitats that were not targeted by the design". If indices are produced for unsuitable species there is a danger, from a species conservation standpoint, that this will be misleading and detrimental to the species. If the population at large is declining but the smaller surveyed population is not, then necessary conservation measures for the species may not be taken in time. It is not difficult to imagine a scenario where a species is pushed from its primary habitat, which is not well-surveyed by WeBS, to a secondary habitat which is better surveyed. Take, for example, a species that would normally feed in a dispersed fashion on damp pasture. If large scale changes in farming practice and land management make its preferred habitat unsuitable, some individuals would move to the fringes of large waterbodies, perhaps protected areas such as nature reserves. As there has been an increase on the type of site surveyed by WeBS, counts of that species would increase and so would the WeBS species index. However, in reality, with feeding areas restricted, the national population would actually be decreasing.

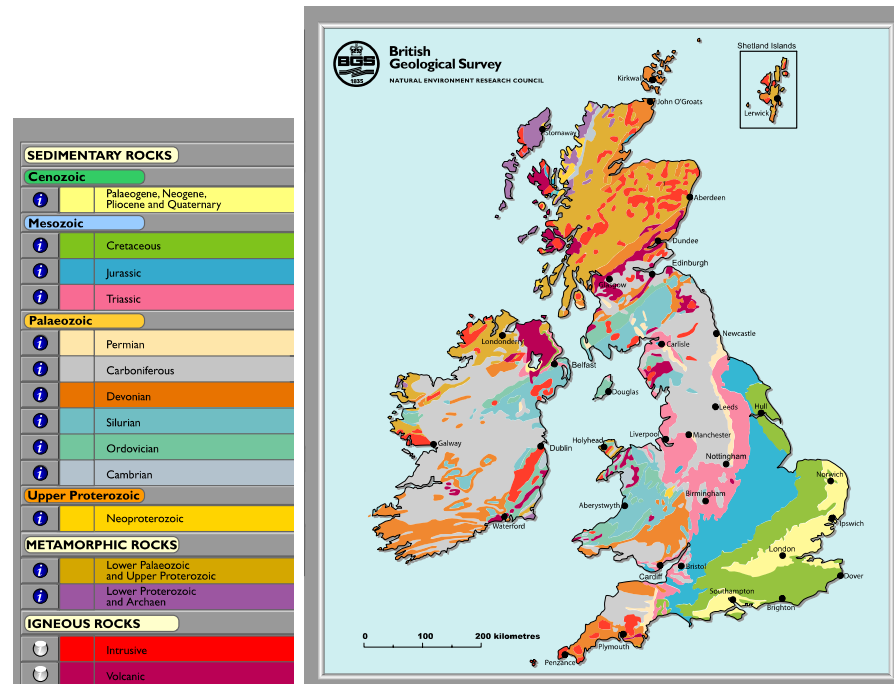


Figure 2.5: Simplified geology map of the British Isles. Reproduced with the permission of the British Geological Survey ©NERC. All rights Reserved.

Habitats are a loose concept, which can be defined by criteria including geology, elevation, prevailing climate, dominant vegetation and human impact. The fundamental geology (Figure 2.5) and topography (Figure 2.6) of the country determines the position of watersheds and the types of wetlands in the area. The prevailing climate (Figure 2.7) also influences wetland type; for example, blanket bogs form in areas of high rainfall and a cool climate, such as in western Scotland.

For comparison of elevation and climate variability with the distribution of WeBS counts in Figures 2.6 and 2.7 and also in Figure 2.8 on page 59, for reasons discussed in Section 1.3.4 site data included in the Mallard file is used as a proxy for locations of WeBS sites.

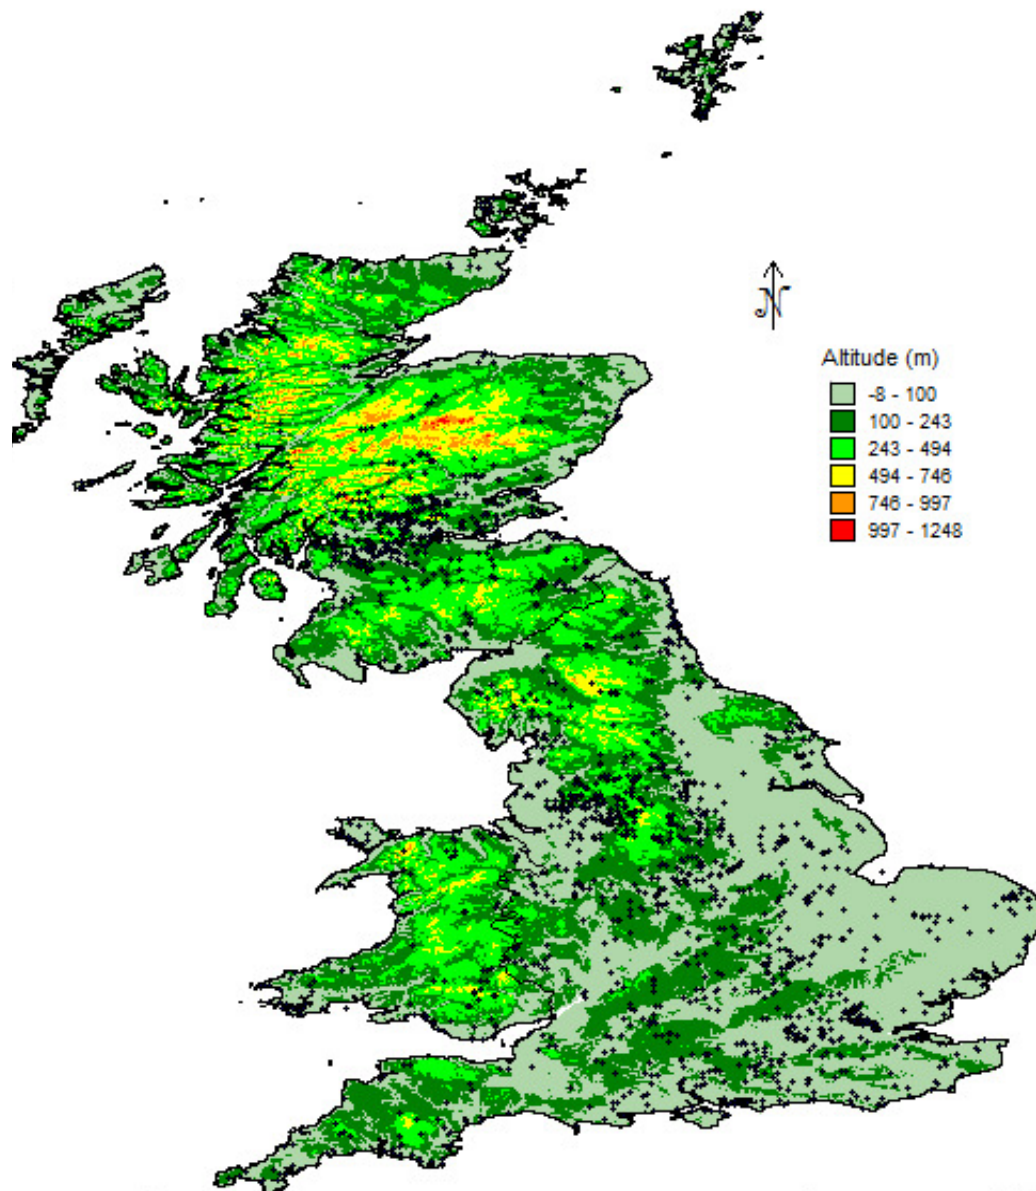


Figure 2.6: Map of Great Britain with sites surveyed between September 2000 and March 2001 overlaid on gridded elevation data. Note map is in geographic coordinates. Produced in and using geographical data from DIVA-GIS (Hijmans et al., 2009).

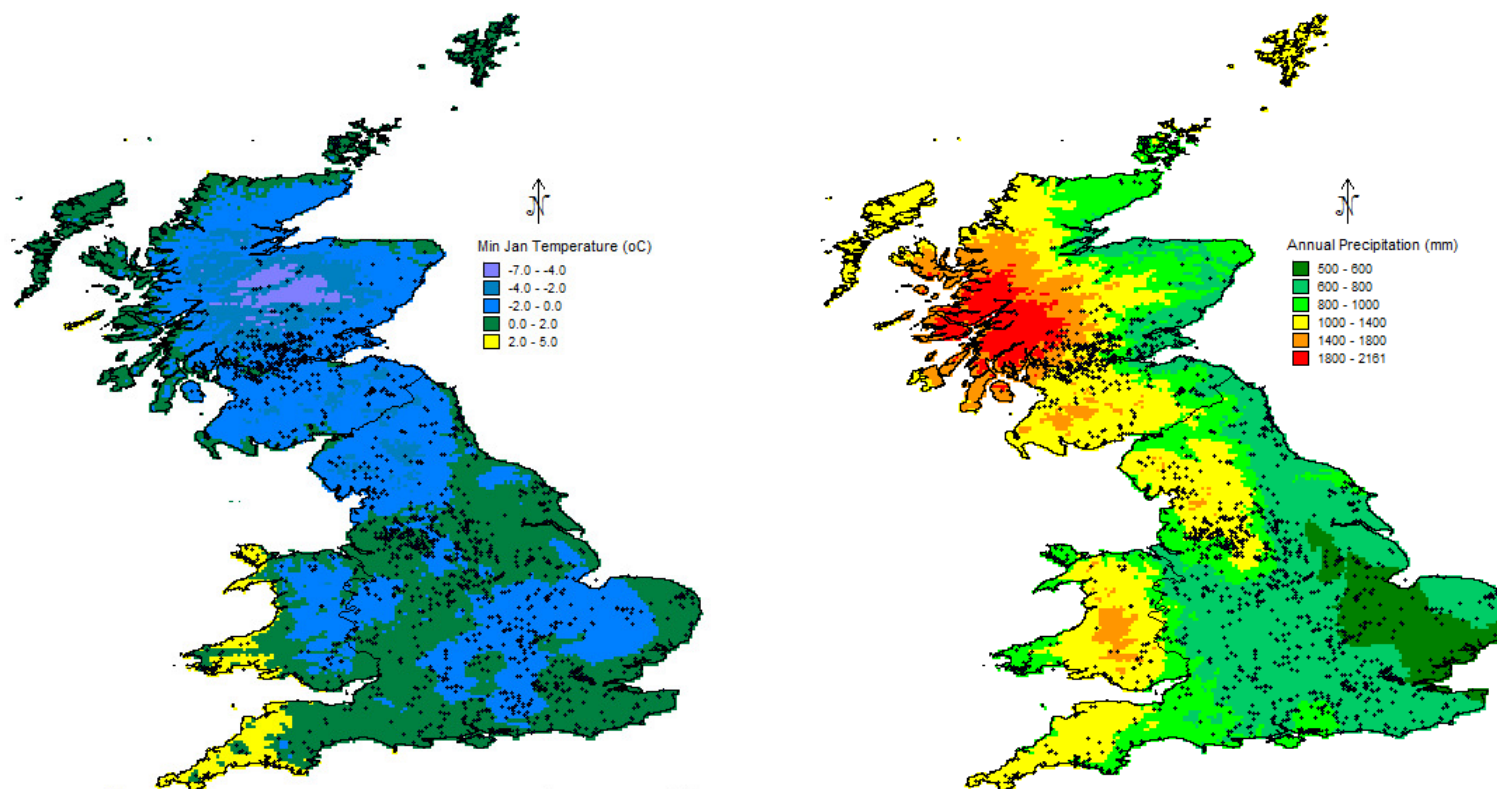


Figure 2.7: Maps of Great Britain with Mallard sites in 2000/01 overlaid on gridded mean minimum January temperature (left) and mean annual precipitation (right). Note map is in geographic coordinates. Produced in DIVA-GIS and using climate data from the WORLDCLIM database (Hijmans et al., 2005, 2009).

2.3.3 *Recording effort*

Both wildfowl and people tend to be attracted to similar winter conditions: preferably mild and sheltered lowland areas rather than more inhospitable upland areas (although this is not true for some species, such as Goldeneye). Nevertheless, some species have requirements closer to humans than others, and it is obviously the case that these will be closely monitored in a scheme such as the WeBS. Pollitt et al. (2003) summarised the coverage of WeBS in the 2000/01 season by representing the density of counts on a map (Figure 2.9). It can be seen in Figures 2.8 and 2.9 that more WeBS counts are conducted near major centres of human population, although many volunteers will travel reasonable distances to nearby waterbodies. It is thought that western Scotland is the most likely place where substantial bird populations are not regularly monitored due to lack of observers (WeBS workshop, 2008).

Humans destroy wetlands (e.g. by diverting rivers underground in towns or draining marshes), change them (e.g. by straightening river courses or creating tidal barrages) and create them (e.g. historical farm ponds, gravel pits and reedbeds from mining, reservoirs, canals and pools aimed at wildlife in nature reserves). Wetland habitats near centres of human population may not therefore be representative of those in more sparsely populated areas; but whether they hold smaller or greater bird populations is not something that will be assessed here.

The U.K Breeding Bird Survey (BBS), organised by the BTO, attempts to maximise volunteer use through a stratified random sampling design. The number of samples, in the form of 1km squares, in the BTO's 83 regions was set according to the potential volunteer resource in that region, and then randomly generated. A minimum level of samples within a region was set, but as some regions have more samples than others the annual counts are weighted by the inverse of the proportion of the area of each region that is surveyed in a given year (BTO, 2008).

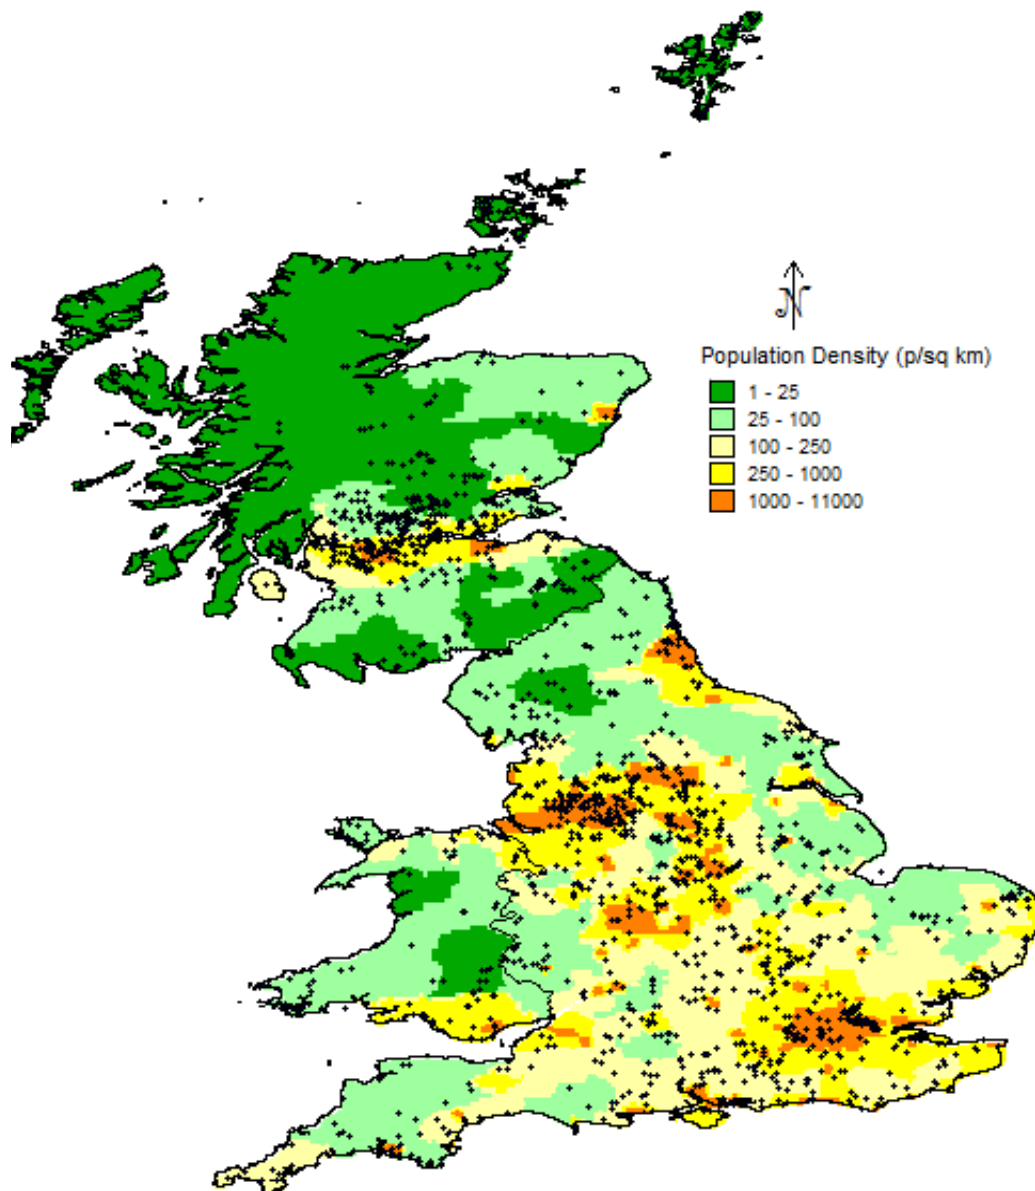


Figure 2.8: Map of Great Britain with sites surveyed between September 2000 and March 2001 overlaid on human population density data. Note map is in geographic coordinates. Produced in and using geographical data from DIVA-GIS (Hijmans et al., 2009).

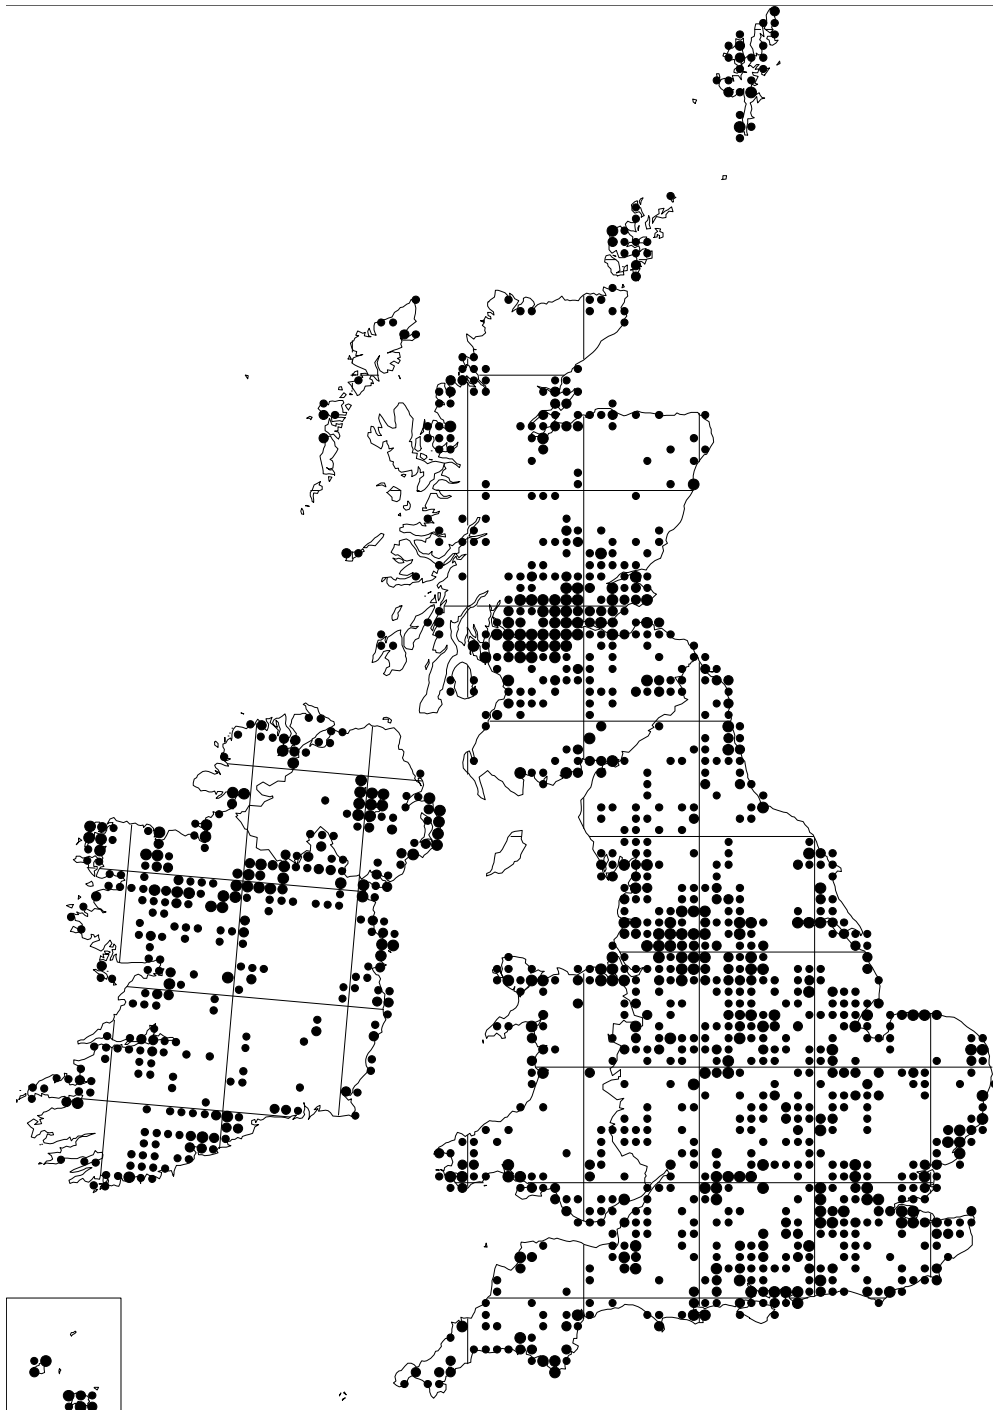


Figure 2.9: Coverage by 10-km grid squares for WeBS Core Counts in the UK, Isle of Man and the Channel Islands and for I-WeBS in the Republic of Ireland in 2000/01. Small dots represent 1-2 count units per 10-km square, medium dots represent 3-4 units and large dots represent five or more units. Pollitt et al. (2003); reproduced with permission.

Section 2.3.2 outlined geographic and climatic factors that are instrumental in determining the extent and type of wetland habitat in different regions of Britain. Any similar weighting of counts in the Wetland Bird Survey to that used in BBS to counter-effect a bias towards areas of high recording effort must reflect the amount and type of wetland habitat within each region or area, rather than the area of land itself. Although challenging, with the continued increase in availability of high quality GIS data such as the Ordnance Survey's digital "Mastermap" of the British Isles and other wetland habitat data such as the Kent Habitat Survey (see Figure 2.11 on page 65), this could be a useful avenue for future research. Digital boundaries for WeBS sites have already begun to be defined (Figure 2.3) and will be further refined. When accurate boundaries and wetland habitat layers are available, it should be possible to undertake more sophisticated sampling and spatial analyses of Wetland Bird Study data.

2.3.4 Case study: WeBS coverage of Kent waterbodies

Kent is a county in southeast England. It has a strong birdwatching tradition, with two coastal bird observatories and several nature reserves. Much of the north Kent coastline is designated under the Ramsar Convention (see Section 1.3.1) as holding internationally important wintering populations of wildfowl; as can be seen in Figure 2.10, there are relatively few WeBS count sectors outside Ramsar protected areas. There are nine wildfowl species which are particularly cited in the Kent Ramsar designations: Dark-bellied Brent Goose; Shelduck; Wigeon; Gadwall; Teal; Pintail; Shoveler; Pochard; and Tufted Duck.

Table 2.1 shows the most surveyed sites in Kent, ordered by a score formed by summing the number of complete counts* and half the number of incomplete counts; thus the Thames Estuary is placed sixth because although it has no counts missing completely, it has a large number of incomplete counts. The top five sites from have "Very Good" recording effort, the next four sites "Good" and the final four sites "Reasonable". Only the sites with "Very Good" recording effort and the two small sites with "Good" recording

*As in Section 2.3.2, Mallard counts are used as a pseudo-representation of all WeBS visits.

effort fulfill the 50% completeness criterion and will be used to create the Mallard index.

Site name	C	I	M	50% C	50% C or I	Sectors	Effort
Stodmarsh NNR and Collards Lagoon	239	12	36	Y	Y	1	VG
Dungeness Gravel Pits	229	23	35	Y	Y	8	VG
Swale Estuary	159	119	9	Y	Y	1	VG
Bewl Water	194	6	87	Y	Y	1	VG
Bough Beech Reservoir	183	28	76	Y	Y	1	VG
Thames Estuary	71	216	0	N	Y	102	G
Medway Estuary	99	155	33	N	Y	42	G
Leybourne and New Hythe Gravel Pits	173	6	108	Y	Y	1	G
Lullingstone Castle Lake	145	6	136	Y	Y	1	G
Pegwell Bay	126	38	123	N	Y	3	R
Thanet Coast	128	16	143	N	Y	8	R
Eastwell Park	117	8	162	N	N	1	R
Chilham and Chartham Gravel Pits	90	36	161	N	N	11	R

Table 2.1: Completeness of Mallard 1966/67 to 2006/07 September to March Counts at important WeBS sites in Kent, England. The number of counts with a quality flag of complete (C); incomplete (I); and missing (M) are shown. Sites that would be included under the 50% completeness criterion are marked **Y** in the 50% C column. Sites that would be included under a 50% non-missing criterion are marked **Y** in the 50% C or I column. The number of sectors associated with the site is also given.

Kent WeBS site outlines are shown in Figure 2.11, coloured according to the recording effort score from Table 2.1. Note that not all the Thames Estuary site is shown, as it extends northwards to Essex and westwards to London. As seen in Figure 2.10, Thames Estuary, Medway Estuary, Pegwell Bay and Thanet Coast are all internationally important wetlands, designated under the Ramsar convention, but they do not have enough complete counts to be used in species indices. Sites which have many sectors (count units), such as the Thames and Medway Estuaries, tend to have more incomplete counts. This is likely to be related to the challenge of finding enough volunteer observers to survey the site at the same time to avoid duplicate counts from birds moving within the site. Loosening the criterion for site inclusion to be 50% of complete or incomplete counts would allow four of Kent's internationally important wetlands to also be included in species indices, including those species which are particularly mentioned in the citation for designating the site.

There are three extensive areas of coastal marshes in Kent. The north Kent marshes are well surveyed within the the Thames, Swale and Medway Ramsar sites discussed above. The marsh that once separated the "Isle of Thanet", in the northeast of the county, and

Romney Marsh, created by the silting up of coastal lagoons behind the shingle spit of Dungeness point to the south, have both been drained by Man in historical times. This has left the network of ditches visible as dense blue areas in Figure 2.11. Ditches like these, together with the riverine habitat all over Kent (except the chalk downland areas) are not well represented in the WeBS scheme.

Relative to many parts of the U.K., Kent is well populated and does not have many wetlands, so should have excellent coverage. However, there are still areas of water, particularly rivers, that are not surveyed at all. What is particularly striking is that despite reasonable recording effort, only seven sites have enough good quality counts for September to March to be included in the September to March Mallard index according to the 50% completeness criterion (Table 2.1). Furthermore, three internationally designated Ramsar wetlands are excluded, despite having many years of data. In the following section, the 50% completeness criterion is assessed and an alternative criterion proposed.

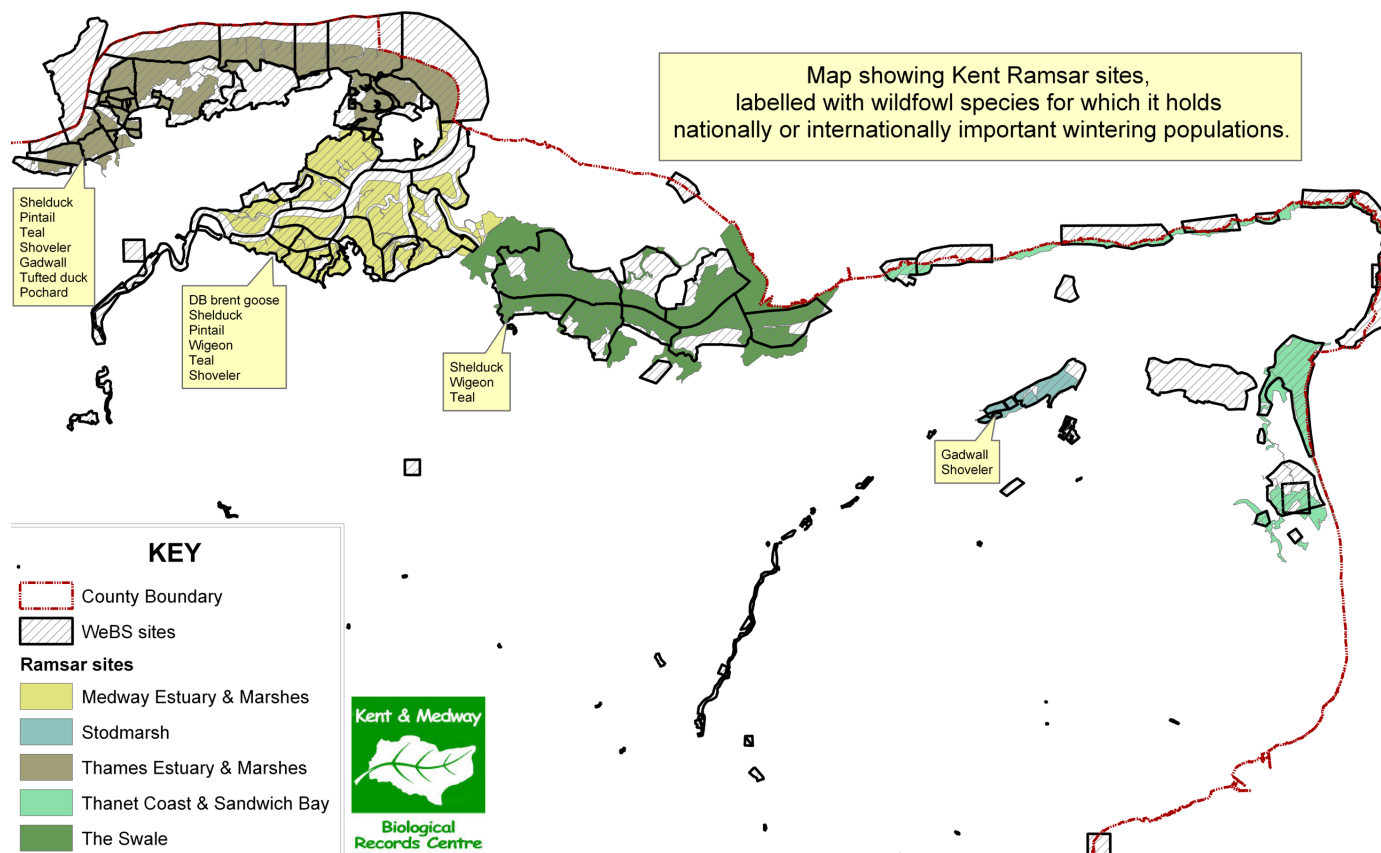


Figure 2.10: Ramsar site boundaries and WeBS counting units (site sectors outlined in black) in north Kent. *Acknowledgements:* Natural England and Kent and Medway Biological Records Centre for Ramsar boundaries; BTO for WeBS boundaries.

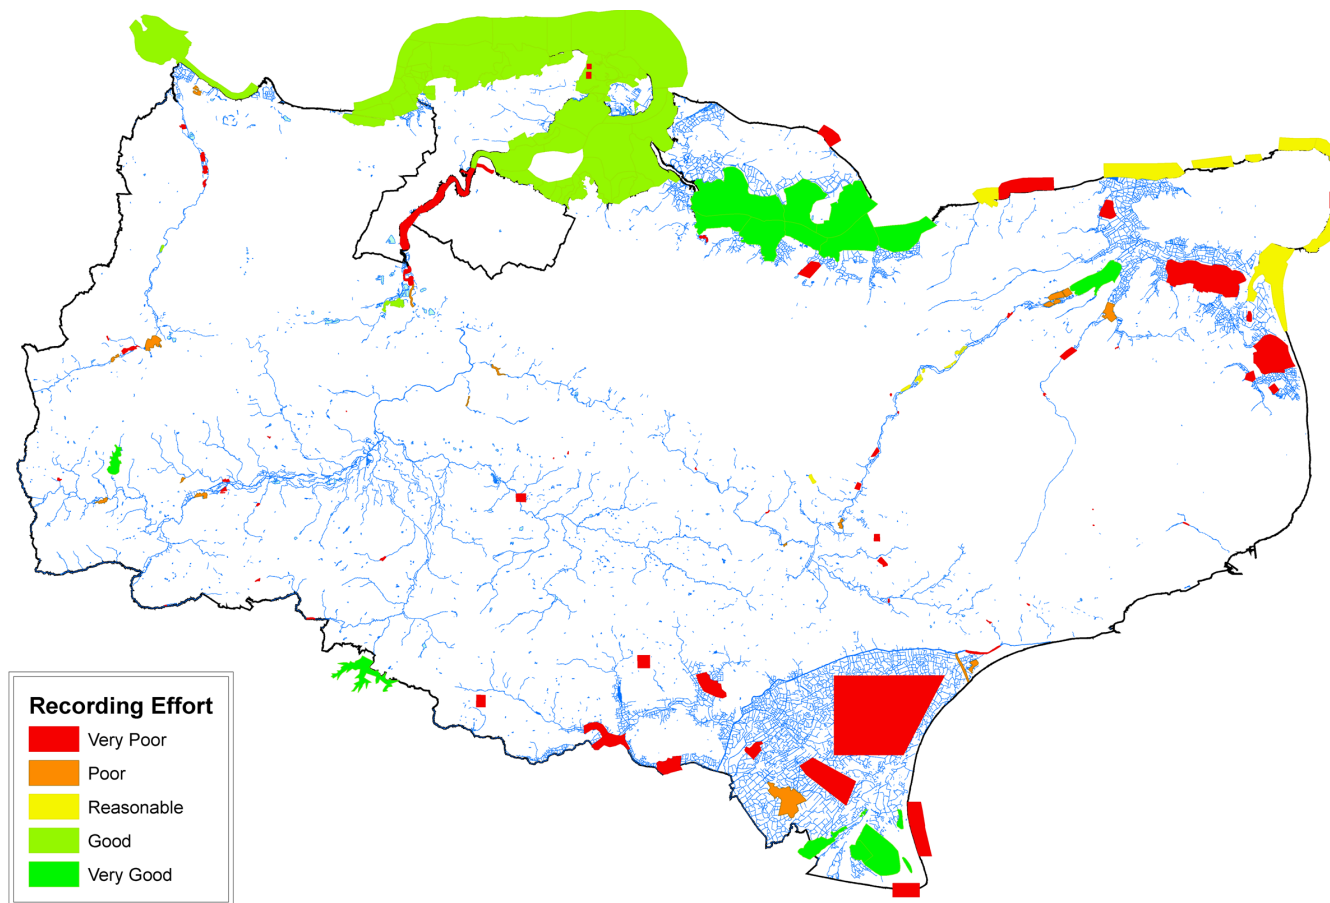


Figure 2.11: Wetlands and WeBS recording effort in Kent. Recording effort was measured by a score calculated by summing the number of complete counts plus half the number of incomplete counts (see Table 2.1). *Acknowledgements:* Kent and Medway Biological Records Centre for water data from the Kent Habitat Survey 2003; BTO for WeBS boundaries.

2.4 WeBS Indices Spatial Coverage: the 50% Completeness Criterion for Including Sites

2.4.1 *The completeness criterion*

As outlined in the first chapter, a key concept currently used when analysing WeBS data is that of count **completeness**, expressed as a quality code associated with the count. The quality of a site count is defined as “good” (complete) or “poor” (incomplete) depending on the observer flag; or, for more complex sites, the proportion of count units within the site which the observers have flagged as good quality, according to the algorithm in Section 1.2.3. A count can also be missing altogether.

As described in Sections 1.3.2 and 1.3.3, the completeness flags are used twice to produce a standard Underhill Index. Firstly, the 50% completeness criterion is applied to the data set to determine which sites should be included in the index. The criterion stipulates that only sites which have 50% of counts marked as complete (i.e. neither missing nor incomplete) should be used; Program 2.1 is a MATLAB function to find the sites that satisfy the criterion. Secondly, the flags are used in the Underhill algorithm, which imputes missing counts and incomplete counts that are less than expected by the Underhill model. In this section, the first of these, using the completeness counts to determine which sites are included in a national annual index, is considered.

Wildfowl species are unevenly distributed across Great Britain, responding to the geographical and climatic factors discussed in Section 2.3.2 according to their differing ecological requirements. Some species are widespread, appearing at many sites, whereas others have a geographically restricted distribution. Similarly, some species occur in large flocks and rarely in small groups, whereas others frequently occur in small numbers (although may still congregate at particularly attractive sites).

The Dark-bellied Brent Goose is a clumped species, occurring in large numbers on some sites but on many sites it is not present at all. Its distribution is also restricted, occurring

Program 2.1 MATLAB function to extract required sites: all those available, those that fulfill the 50% completeness criterion, or a list of site codes supplied by the user. Used in Program 1.1, Preparedata.

```

function [sitestoinclude site data region] = getsitestoinclude(reply,site,data,
    region,nyears,nmonths,sitestoinclude)
    %The last input variable is only needed when giving a list of sites to
    %include. Assumes data is in site order.

if strcmp('C',reply)
    %User requires 50% criterion to be applied.
    sitelist = unique(site);
    %Tally the number of complete counts for each site.
    %Default to zero, as some sites have no complete counts:
    sitetally = zeros(length(sitelist),1);
    %Find number of complete counts (quality=1) for each site:
    sitel = site(data(:,4)==1);
    [sitelist1,uniqsitemax] = unique(sitel);
    [sitelist1,uniqsitemin] = unique(sitel,'first');
    comptally = uniqsitemax-uniqsitemin+1;
    sitelistmap = containers.Map(sitelist1,num2cell(comptally));
    Var1= isKey(sitelistmap,sitelist);
    %N.B., when iskey is false, we know the site has no complete counts.
    sitetally(Var1) = cell2mat(sitelistmap.values(sitelist(Var1)));

    %Select sites which have at least 50% complete counts:
    sitestoinclude = sitelist(sitetally>=floor(0.5*nyears*nmonths));

    %Select data where the site fulfils the 50% complete requirement:
    var = repmat({'N'},length(sitelist),1);
    var(ismember(sitelist,sitestoinclude)) = {'Y'};
    compincmap = containers.Map(sitelist,var);
    rows = strmatch('Y',compincmap.values(site));
    data = data(rows,:);    site = site(rows);    region = region(rows);

elseif strcmp('A',reply)
    %User does not require 50% criterion to be applied.
    sitestoinclude = unique(site);

else
    %User requires data from a particular set of sites.
    if exist('sitestoinclude', 'var') == 0
        error('You didn't answer A or C correctly when asked about completeness
            and sitestoinclude doesn't exist - Aborting.')
    end
    sitelist = unique(site);
    %Revise sitestoinclude to those sites in species file
    %(for cases where sitestoinclude in the workspace includes site
    %codes not in the species file).
    %Also sort to be the same order as sites in species file, so that
    %correct site names are found later.
    [tf loc]=ismember(sitestoinclude,sitelist);
    [s1 loc2]=sort(loc);
    zar=sitestoinclude(loc2);
    zar2=tf(loc2);
    sitestoinclude=zar(zar2);

    %Select data where the site is in the user's list:
    var = repmat({'N'},length(sitelist),1);
    var(ismember(sitelist,sitestoinclude)) = {'Y'};
    compincmap = containers.Map(sitelist,var);
    rows = strmatch('Y',compincmap.values(site));
    data = data(rows,:);    site = site(rows);    region = region(rows);
end

```

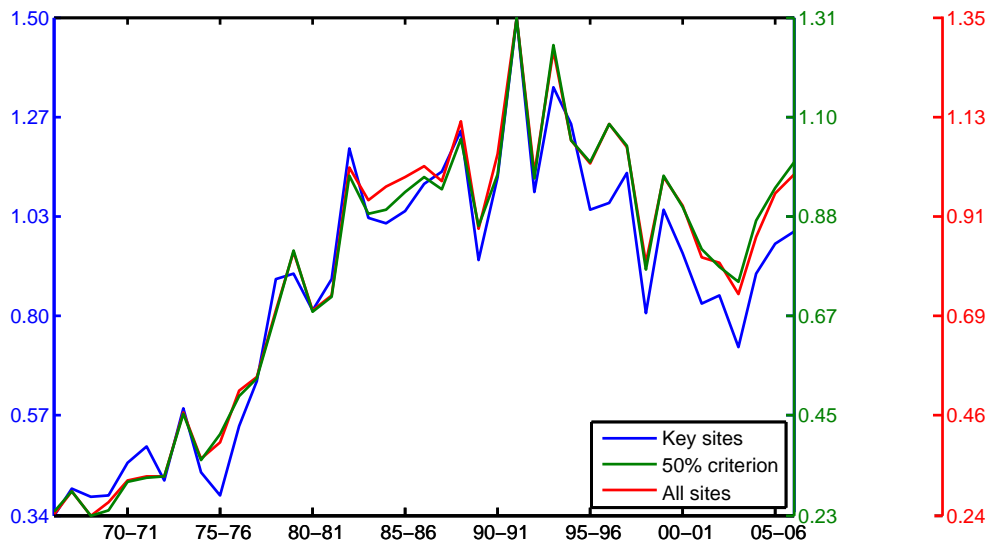


Figure 2.12: Underhill indices based on all Dark-bellied Brent Goose sites (red; 338 sites), those that fulfill the 50% criterion (green; 115 sites) and the six most influential sites for the species (blue), for the months December to February.

mostly in the southeast of Britain, from the Humber estuary to the Severn estuary (Collier et al., 2005). In species with this kind of clumped distribution the Underhill index is strongly influenced by the handful of sites which hold large numbers of the species; it can be seen in Figure 2.12 that the index looks much the same when based on only the six most influential Dark-bellied Brent Geese sites as when using over 300 sites.

In contrast, the Mute Swan population is more thinly spread across Great Britain. In widespread species, basing the index on a small number of sites only does not so closely mirror the trend using more sites, as seen in Figure 2.13. The influence of a site on the Underhill index is measured by the site factors in the Underhill algorithm are a measure of the influence of a site; the top six sites for Dark-bellied Brent Goose account for 62% of the all-site index, whereas the top six Mute Swan sites have a 10% influence on the final index. (The indices are shown on different scales to better display shape similarities.)

The Wetland Bird Survey is ambitiously aimed at monitoring species which are widespread and geographically restricted, and which occur in small groups and vast flocks. Recent research into surveying in WeBS has focussed on estimating waterbird population sizes, particularly for widespread species. The UK Non-Estuarine Coastal Waterbird

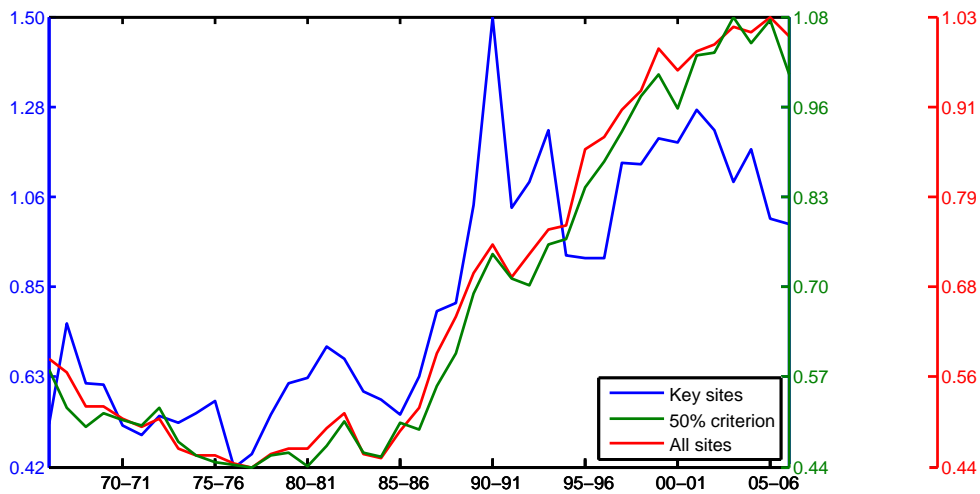


Figure 2.13: Underhill indices based on all Mute Swan sites (red; 3994 sites), those that fulfill the 50% criterion (green; 689 sites) and the six most influential sites for the species (blue), for the months September to March.

Survey (UK-NEWS) aims to supplement counts for coastal waterbird species (mostly waders). Similarly, the Dispersed Waterbird Survey (DWS) has been designed to supplement WeBS data with a stratified random sampling based approach for estimating the numbers of Mute Swan, Greylag Goose, Canada Goose, Teal, Mallard and Tufted Duck. However, national population estimates using the Dispersed Waterbird Survey for Pochard and Goldeneye were lower than expected. Pochard tend to congregate at large waterbodies which were not as likely to be captured in the DWS as in the Wetland Bird Survey core counts. A large proportion of Goldeneye occur in upland areas in winter, which were excluded from DWS analyses (Kershaw, 1997; Armitage et al., 2001; Jackson et al., 2006).

The costs of carrying out DWS and UK-NEWS surveys, which require professional counters to supplement volunteers to attain full coverage of randomly chosen sampling areas, mean that these surveys will be periodic at best, perhaps taking place every six to ten years. This type of intensive periodic survey will therefore not replace the annual WeBS core counts on which indices will continue to be based, but may provide greater clarity on whether annual indices based on WeBS sites accurately reflect national trends for

dispersed species (WeBS workshop, 2008).

The problem of selecting the sample of sites on which to base the national indices for species of varying aggregations and distributions is considered in the rest of this chapter. An appraisal of the 50% criterion is given, followed by an alternative method which avoids some of the biases inherent in the 50% method.

2.4.2 Effects of the completeness criterion on spatial coverage of WeBS indices

In Section 2.3.3 it was seen that although there are many WeBS sites across the country, the count sectors which are regularly surveyed tend to be close to centres of human population; the 50% completeness criterion may, therefore, introduce additional spatial bias. Analysis of sites included by the 50% criterion in Kent showed that sites which hold a significant proportion of the population, including those holding internationally important numbers, may not be included, even when some data is available (Section 2.3.4). These two major issues with using the 50% criterion are joined by a third, which is that any new site entering WeBS today will not be used for indexing until at least as many years have past since the start of the scheme (i.e. currently over forty), even assuming no future missing or incomplete counts.

One solution is to remove any post-survey sampling completely, modelling all missing counts within the data set. However, using the Underhill algorithm, this can have the effect of exaggerating trends at the major sites: smaller sites have very little impact on the overall index, as seen in Figures 2.12 and 2.13. Figure 2.15 illustrates the difference in the annual Underhill index plots for all species and Figure 2.16 illustrates the difference in the index plots for all species using the mean of available counts (good or poor quality), without using the Underhill imputing model. For most species, as with Mute Swan and Dark-bellied Brent Goose, there is very little difference between the index with the 50% criterion applied or not.

Since it apparently has little effect, why remove any sites from the index at all? Figure

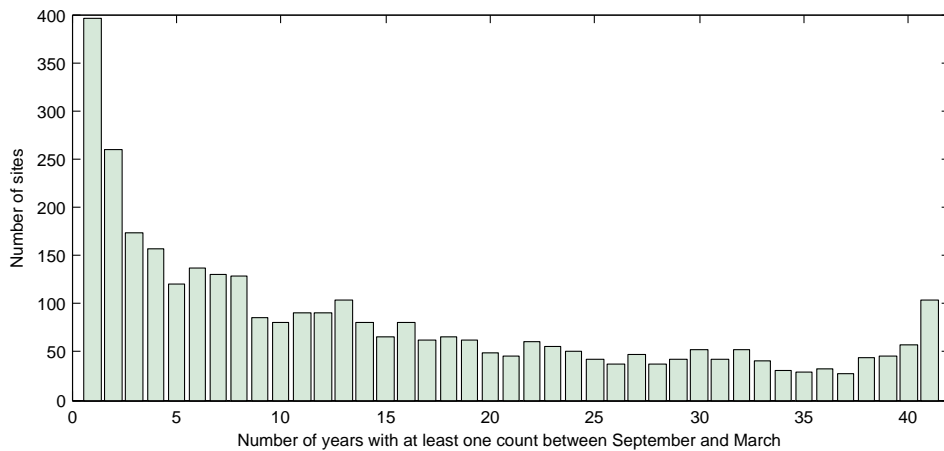


Figure 2.14: Tally of Goldeneye sites by years of non-missing counts.

2.14 tallies the number of sites which have complete or incomplete counts in at least one of the September to March surveys over the 41 year period. Almost 400 sites (11% of the total) have counts for only a single year, 47% have less than ten years and 72% of sites have less than 22 years with any counts recorded. The 50% criterion stipulates that over half the counts (i.e. seven counts a year in this case) must be complete; 83% of sites do not meet this criterion. Basing an index on 4000 sites rather than 600 may at first seem more “statistically” valid, but that assumes the extra sites are adding more information. If that is not the case, then the apparent sample size is misleading, and may encourage greater confidence in the representative of the index for national populations than should be the case.

Table 2.2 compares some statistics on using the 50% criterion with using data from all available sites for the twenty-six wildfowl types of interest. As expected, there are significantly fewer missing and incomplete counts when the 50% criterion is applied: an average of 30.6% rather than 73.7%. The number of iterations it takes the Underhill algorithm to converge is much less when the criterion is applied; the convergence criterion of 10^{-5} hadn't been reached after 500 iterations for any species when all sites are used. Convergence is slow for all species; for example using all 3393 sites for Goldeneye the convergence criterion is reached after 2706 iterations and for Bewick's Swan 647 sites it

is reached after 13443 iterations.

Kirby et al. (1995) recommended the use of the 50% completeness criterion for wildfowl data from the NWC based on work by Underhill and Prÿs-Jones (1994). However, the 50% criterion was originally developed for estuaries in the BoEE scheme, which resulted in 112 of 219 sites being retained, which held between 75% and 90% of the population present on the 219 sites. As counts in the NWC scheme took place on a wider variety of wetlands which had been surveyed for a more widely varying time period, for NWC counts, this corresponded to 596 sites (about 10% of the sites in the database at the time). The overall proportion of values imputed was then 23.7%; as already noted from Table 2.2 the proportion is higher now at around 30% and 772 sites used for Mallard. However, an analysis was not presented of the proportion of wildfowl population on the sites chosen with the 50% criterion.

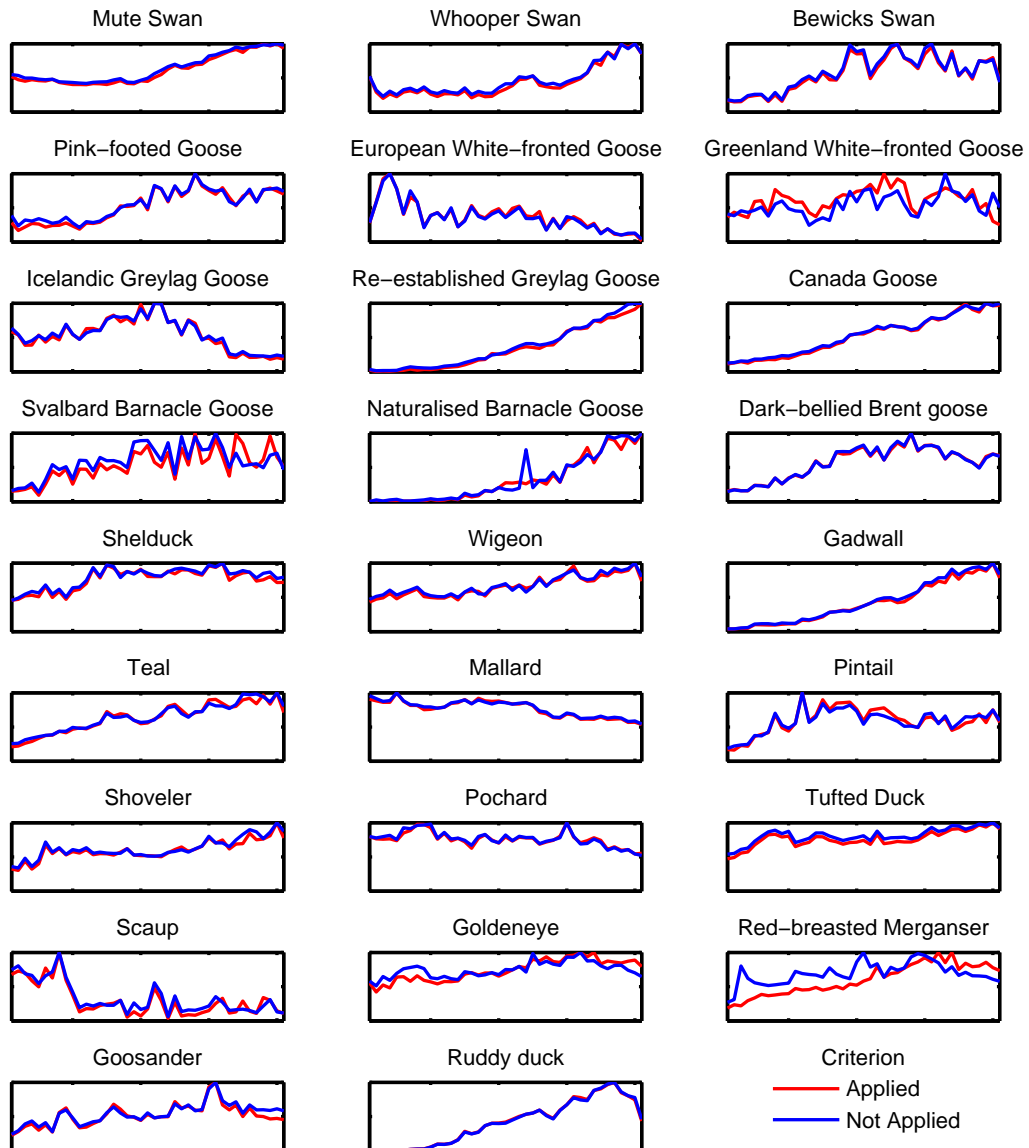


Figure 2.15: Annual Underhill Index plots for the 26 wildfowl types, produced using monthly counts from September to March and sites selected with and without the 50% completeness criterion. Each plot is scaled so that the peak is set equal to 1, with y-axis limits of 0 and 1 and x-axis limits of 1966/67 and 2006/07.

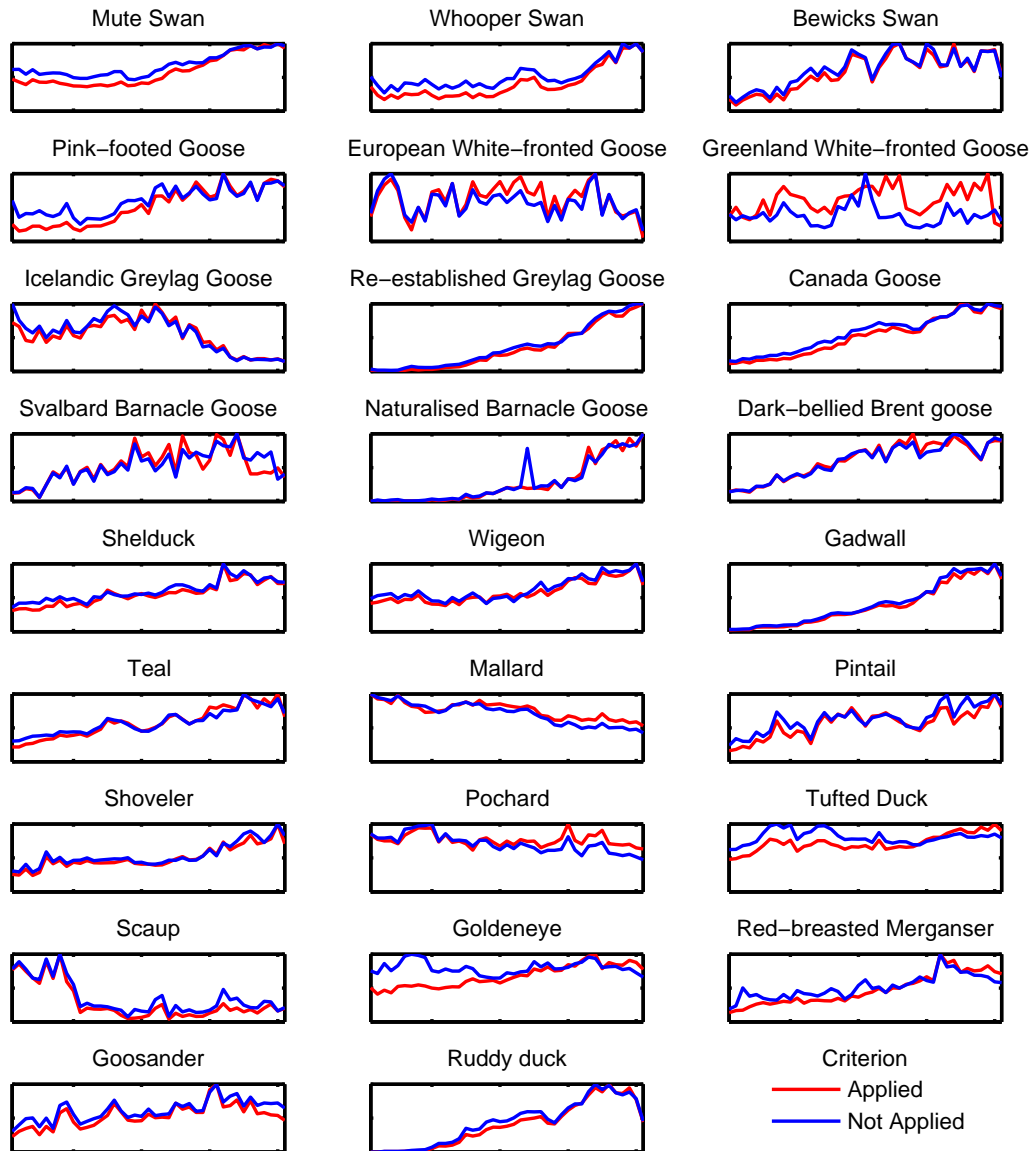


Figure 2.16: Annual Mean Index plots for the 26 wildfowl types, produced using monthly counts from September to March and sites selected with and without the 50% completeness criterion. Each plot is scaled so that the peak is set equal to 1, with y-axis limits of 0 and 1 and x-axis limits of 1966/67 and 2006/07.

	50% Criterion Sites					All Sites				
	#sites	%mod	%miss	#iter	convergence	#sites	%mod	%miss	#iter	convergence
Mute Swan	689	29.0	(27.1)	19	9×10^{-06}	3994	75.9	(74.6)	500	2×10^{-04}
Whooper Swan	283	31.8	(29.5)	24	9×10^{-06}	1716	76.1	(74.8)	500	2×10^{-04}
Bewicks Swan	139	31.9	(29.0)	20	8×10^{-06}	647	68.9	(66.9)	500	7×10^{-04}
Pink-footed Goose	206	31.1	(27.8)	26	9×10^{-06}	1220	75.1	(73.7)	500	4×10^{-03}
European White-fronted Goose	94	35.3	(31.9)	61	1×10^{-05}	510	71.5	(69.7)	500	2×10^{-04}
Greenland White-fronted Goose	44	35.1	(33.5)	30	9×10^{-06}	319	78.8	(77.3)	500	6×10^{-04}
Icelandic Greylag Goose	173	31.0	(28.8)	24	8×10^{-06}	958	76.5	(75.4)	500	5×10^{-03}
Re-established Greylag Goose	323	29.4	(27.1)	29	8×10^{-06}	1426	70.6	(68.7)	500	6×10^{-04}
Canada Goose	587	28.6	(26.7)	23	1×10^{-05}	2899	73.1	(71.8)	500	5×10^{-04}
Svalbard Barnacle Goose	45	32.5	(29.0)	82	1×10^{-05}	208	72.7	(71.1)	500	8×10^{-03}
Naturalised Barnacle Goose	189	30.7	(27.4)	22	9×10^{-06}	835	69.6	(67.7)	500	9×10^{-05}
Dark-bellied Brent Goose	81	27.1	(21.2)	16	9×10^{-06}	362	68.1	(64.4)	500	9×10^{-05}
Shelduck	313	30.0	(27.2)	22	1×10^{-05}	1823	75.5	(73.6)	500	5×10^{-05}
Wigeon	633	29.0	(26.9)	23	7×10^{-06}	3384	75.0	(73.7)	500	3×10^{-04}
Gadwall	419	28.3	(26.2)	24	8×10^{-06}	1824	70.1	(68.3)	500	2×10^{-04}
Teal	678	29.7	(27.6)	18	9×10^{-06}	4002	76.6	(75.4)	500	7×10^{-04}
Mallard	772	29.4	(27.4)	24	1×10^{-05}	5778	80.4	(79.3)	500	1×10^{-03}
Pintail	328	30.2	(27.0)	26	6×10^{-06}	1473	70.6	(68.8)	500	9×10^{-05}
Shoveler	482	28.9	(26.7)	15	5×10^{-06}	2118	71.2	(69.6)	500	3×10^{-05}
Pochard	676	29.7	(27.9)	19	7×10^{-06}	3326	73.4	(72.3)	500	1×10^{-04}
Tufted Duck	717	29.4	(27.6)	20	7×10^{-06}	4118	75.7	(74.7)	500	7×10^{-05}
Scaup	206	32.0	(28.6)	54	8×10^{-06}	1063	72.5	(70.9)	500	5×10^{-05}
Goldeneye	581	30.4	(28.2)	16	8×10^{-06}	3393	76.2	(75.0)	500	9×10^{-05}
Red-breasted Merganser	190	32.8	(29.6)	17	6×10^{-06}	1495	79.6	(77.6)	500	7×10^{-05}
Goosander	450	32.4	(30.6)	16	9×10^{-06}	2539	74.5	(73.3)	500	6×10^{-05}
Ruddy Duck	277	29.8	(28.0)	21	7×10^{-06}	1074	68.5	(67.0)	500	3×10^{-05}

Table 2.2: Comparison of imputing with the Underhill model for all sites and sites chosen with the 50% criterion. The number of sites, the percentage of counts imputed by the Underhill model (including some incomplete counts) and the percentage of missing counts for September to March, 1966/67 to 2006/07 are listed. Also given is the number of iterations for the Underhill algorithm to converge and the last value of the algorithm's convergence measure (see Section 4.2.2). The maximum number of iterations was set at 500 and the convergence criterion at 10^{-5} . Cases where the convergence measure at the 500th iteration was greater than 10^{-4} are in bold.

2.4.3 *Top sites criterion*

The rationale behind limiting the number of sites for which data is imputed is sound. To address the three issues with the 50% completeness criterion in the previous chapter, an alternative criterion should:

- be less restrictive of the proportion of missing values from any one site, to make it possible for new sites entering the WeBS scheme to be included in the index in a reasonable amount of time;
- prioritise the inclusion of sites which are important for the species in question;
- minimise extra spatial bias to centres of human population to that already present in the data set as a whole.

Ranking observations is a useful approach for many environmental problems; for example, ranked set sampling can be used where the cost of measurement exceeds that of collecting a sample (Patil, 2002). The amount of information a site adds to the index depends on the proportion of the population the site holds and the amount of missing data at the site. A site which has many missing counts and only ever contains small numbers of the species for which it is wished to construct an index is less useful than a site which is important for the species and has few missing and incomplete counts. The site factors from the Underhill model, obtained by running the Underhill algorithm on all the potential sites for each species, are a measure of the influence of a site on the year factors (Underhill index) and month factors. These can be used to rank sites.

A new “top sites” criterion for site inclusion is proposed that attempts to balance the influence and amount of missing data of the sites to be included.

- **To be included, a site must have at least one count in any month for a run of ten years or more.** This excludes sites for which there is little information on which to assess the significance of the site (i.e. the site factor), but is not so stringent as to preclude flexibility in sites entering and leaving the scheme.

- **The set of sites chosen contain the maximum proportion of the population on WeBS sites, subject to the other criteria.** Sites are ranked according to their Underhill site factors, calculated on all the remaining potential sites for that species.
- **The number of sites included shall be the higher of (i) 5% of the total number of potential sites and (ii) the number of sites which correspond to a data set with no more than 25% of missing values.** The overall proportion of missing values should be no higher than the original missing rate in Kirby et al. (1995), i.e. 25%, unless the species is particularly widespread, in which case a higher missing rate is permissible.
- **When calculating the criterion, missing values are used rather than missing and incomplete values.** Counts at large sites such as estuaries are often marked as incomplete due to not all parts of the estuary being surveyed; however there is still valid information on which to assess the significance of the site.




Figure 2.17: Proportion of missing values in set of Goldeneye data as sites are added in site factor order in the top sites procedure.

Program 2.2 is a MATLAB implementation of this criterion. Figure 2.17 shows how the missing rate changes as the data set grows site by site. The site factor is modified so that any sites which do not have at least ten years with at least one monthly count are set to

Program 2.2 MATLAB function to choose which sites to use for indexing according to the top sites criterion.

```

function [newsitestoinclude, effective_missing, influence, n_toomanymiss]...
    = topsites(data, sitestoinclude, sitefactor, missingrate, use5rule)
%Returns sites that fulfil the criterion, the overall missing rate of those
%sites, what proportion of the original site influence in the new set of sites
%and the number of sites excluded because there were not ten years of data.
%Find sites with greatest amount of birds present
%but limit amount of missing data.
%Defaults to a missing rate of 25% and use 5% rule equal to 1 (TRUE).
if nargin < 4, missingrate=0.25; end, if nargin < 5, use5rule=1; end
%%
nsites = length(sitestoinclude); nyears=1+max(data(:,2))-min(data(:,2));
nmonths=max(data(:,3));          x = zeros(nsites,4);
miss=data(data(:,5)==3,:);
miss2=miss(:,1:2);               miss2(:,2)=miss(:,2)-min(data(:,2))+1;

%Tally number of years with at least one count in for each site:
g=nmonths-accumarray([miss2(:,1) miss2(:,2)],miss(:,1),[],@length);
g(find(g))=1;    x(1:max(miss(:,1)),1)=sum(g,2);
%Tally number of missing values for each site:
x(1:max(miss(:,1)),4)=accumarray(miss(:,1),miss(:,1),[],@length);

%Sites must have data for at least one month in each of ten years:
sitefactor2=sitefactor;
IX = x(:,1)<10;          sitefactor2(IX) = 0;
n_toomanymiss = [sum(IX) length(IX)-sum(IX)];
[B IX] = sort(sitefactor2,'descend');
criterion = (1:nsites)*nyears*nmonths*missingrate;
y = cumsum(x(IX,4));
%
nsitestouse=find(criterion-y > 0,1,'last');
%%
%Want to have at least 5% of the original number of sites:
if isempty(nsitestouse)
    if use5rule==0, disp('Sorry having to use 5% rule!'), end
    newsitestoinclude = sitestoinclude(IX(1:floor(0.05*nsites)));
    effective_missing = y(floor(0.05*nsites))...
        /(criterion(floor(0.05*nsites))/missingrate);
    warning('Not enough sites were included at missing rate "%s", \n using 0
        .05 criterion with effective missing rate "%s".',...
        missingrate, effective_missing)
    influence = sum(B(1:floor(0.05*nsites)))/sum(sitefactor);
elseif nsitestouse<floor(0.05*nsites) && use5rule==1
    newsitestoinclude = sitestoinclude(IX(1:floor(0.05*nsites)));
    effective_missing = y(floor(0.05*nsites))...
        /(criterion(floor(0.05*nsites))/missingrate);
    warning('Not enough sites were included at missing rate "%s", \n using 0
        .05 criterion with effective missing rate "%s".',...
        missingrate, effective_missing)
    influence = sum(B(1:floor(0.05*nsites)))/sum(sitefactor);
else
    newsitestoinclude = sitestoinclude(IX(1:nsitestouse));
    effective_missing=y(nsitestouse)/(criterion(nsitestouse)/missingrate);
    influence = sum(B(1:nsitestouse))/sum(sitefactor);
end
end

```

have a site factor of zero (sites 1763 to 3393); the site with the largest site factor is added first and so on. In this case a missing rate of 25% corresponds to 166 sites, and 5% of 3393 (rounded down) is 169, so 169 sites will be used.

Figures 2.18, 2.19 and 2.20 demonstrate how the new criterion addresses the issues with the 50% criterion. By prioritising important sites for the species, the index is very close to that produced by using all the sites; it is nearly identical in the case of Dark-bellied Brent Goose, a clumped species (Figure 2.18). The index from the site inclusion criterion is also representative for widespread species such as Mute Swan (Figure 2.19). Goldeneye (Figure 2.20) demonstrates the improvement in avoiding spatial bias over the 50% criterion, with the spatial mean of the 50% site selection being further south than it should be for this northern species, as sites in southern Britain tend to have better coverage than those in north Scotland. It can be seen that there has been an upwards trend over time at southern sites in this species whereas numbers at northern sites have fluctuated without showing an such an evident trend. The southern bias results in the overall index on sites chosen using the 50% criterion evidencing an upwards trend which is not as pronounced when using the proposed site inclusion criterion.

There are more missing counts in the autumn and spring than in mid-winter; coverage is highest in January when the International Waterfowl Census takes place. Basing the index on an appropriate set of months increases the set of sites chosen by this criterion. For example, Figure 2.21 compares the missing rate as sites are added according to the **top sites** site inclusion algorithm. In both cases there are not sufficient sites included by the 25% criterion, so 5% of the total available sites are used. If counts are used between September and March this corresponds to 166 sites with a missing rate of 29.4%, whereas if counts from November to March are used then the top sites criterion selects 169 sites and the missing rate is reduced to 26.3%. To avoid unnecessary imputation whilst keeping a representative data set, in the next chapter the selection of appropriate months for each species is explored.

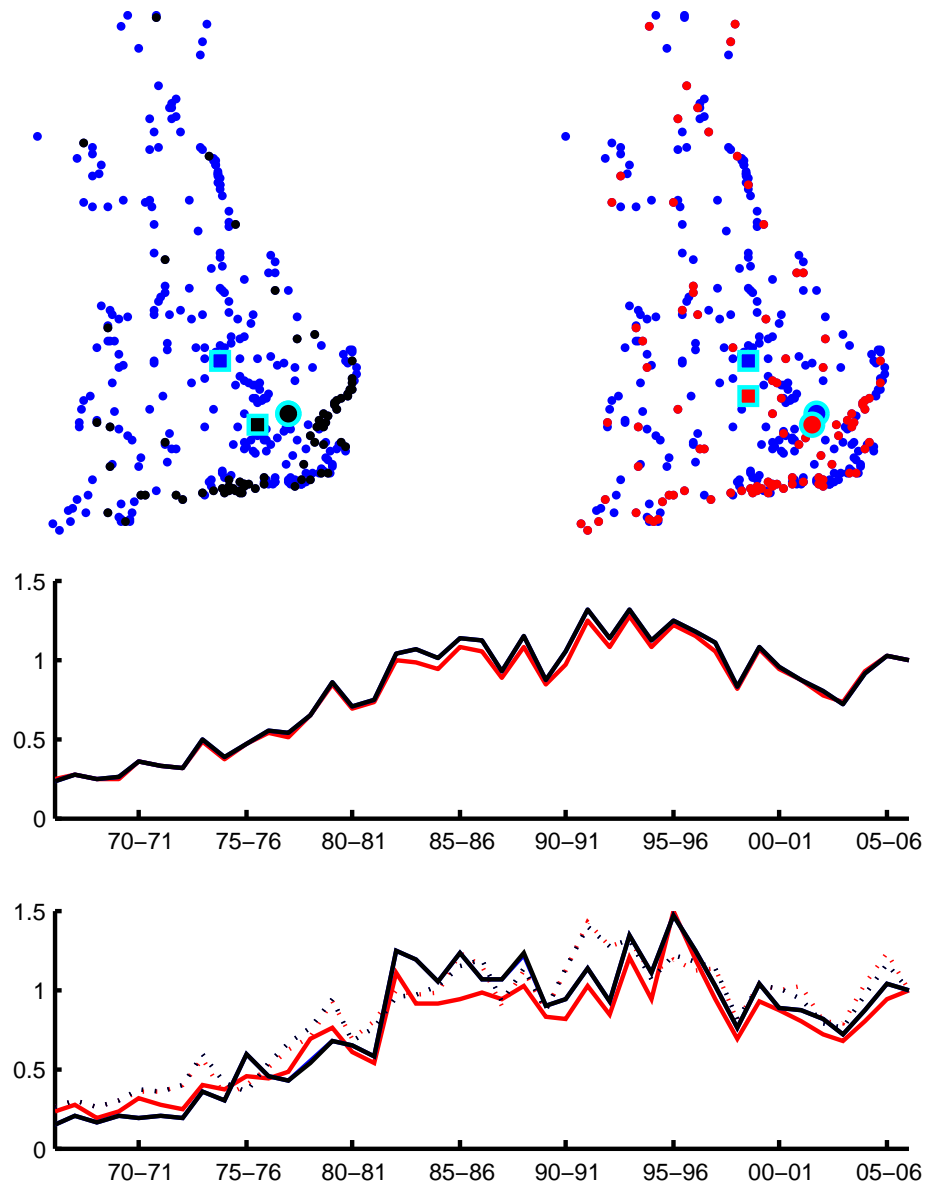


Figure 2.18: Comparison of site criteria for Dark-bellied Brent Goose counts for November to March. Top: Map of all sites (blue) and sites selected by the proposed site inclusion criterion (left; black) and the 50% completeness criterion (right; red). The spatial mean of the site locations (squares) and site locations weighted by the site factors (large circles) for each set of sites are also shown. Middle: Underhill index using all sites (blue), proposed site inclusion criterion (black) and 50% completeness criterion (red). Bottom: As middle, but indices produced for sites north of the weighted mean of all sites (solid) and south (dotted).

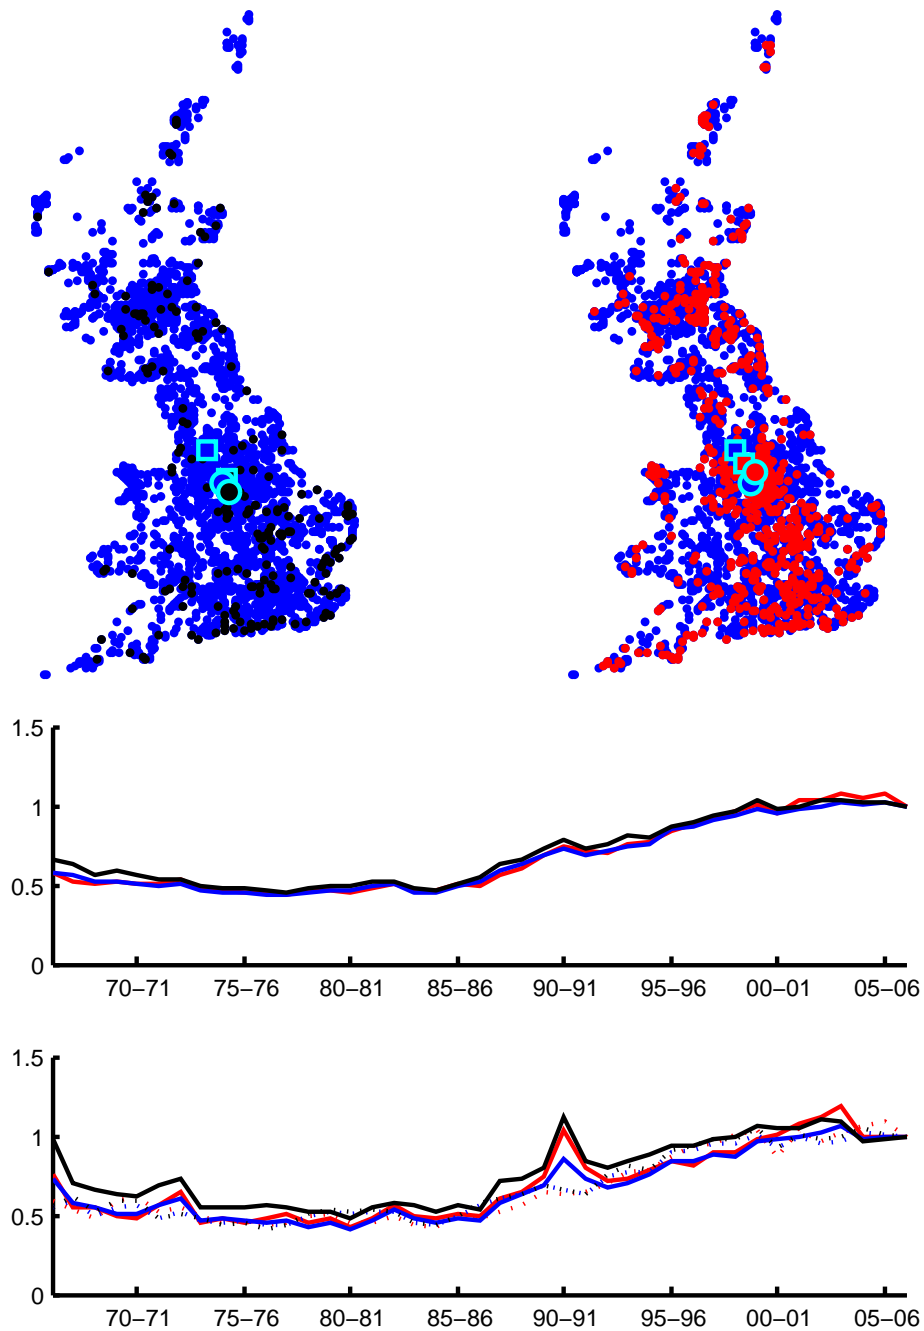


Figure 2.19: Comparison of site criteria for Mute Swan counts for September to March. Top: Map of all sites (blue) and sites selected by the proposed site inclusion criterion (left; black) and the 50% completeness criterion (right; red). The spatial mean of the site locations (squares) and site locations weighted by the site factors (large circles) for each set of sites are also shown. Middle: Underhill index using all sites (blue), proposed site inclusion criterion (black) and 50% completeness criterion (red). Bottom: As middle, but indices produced for sites north of the weighted mean of all sites (solid) and south (dotted).

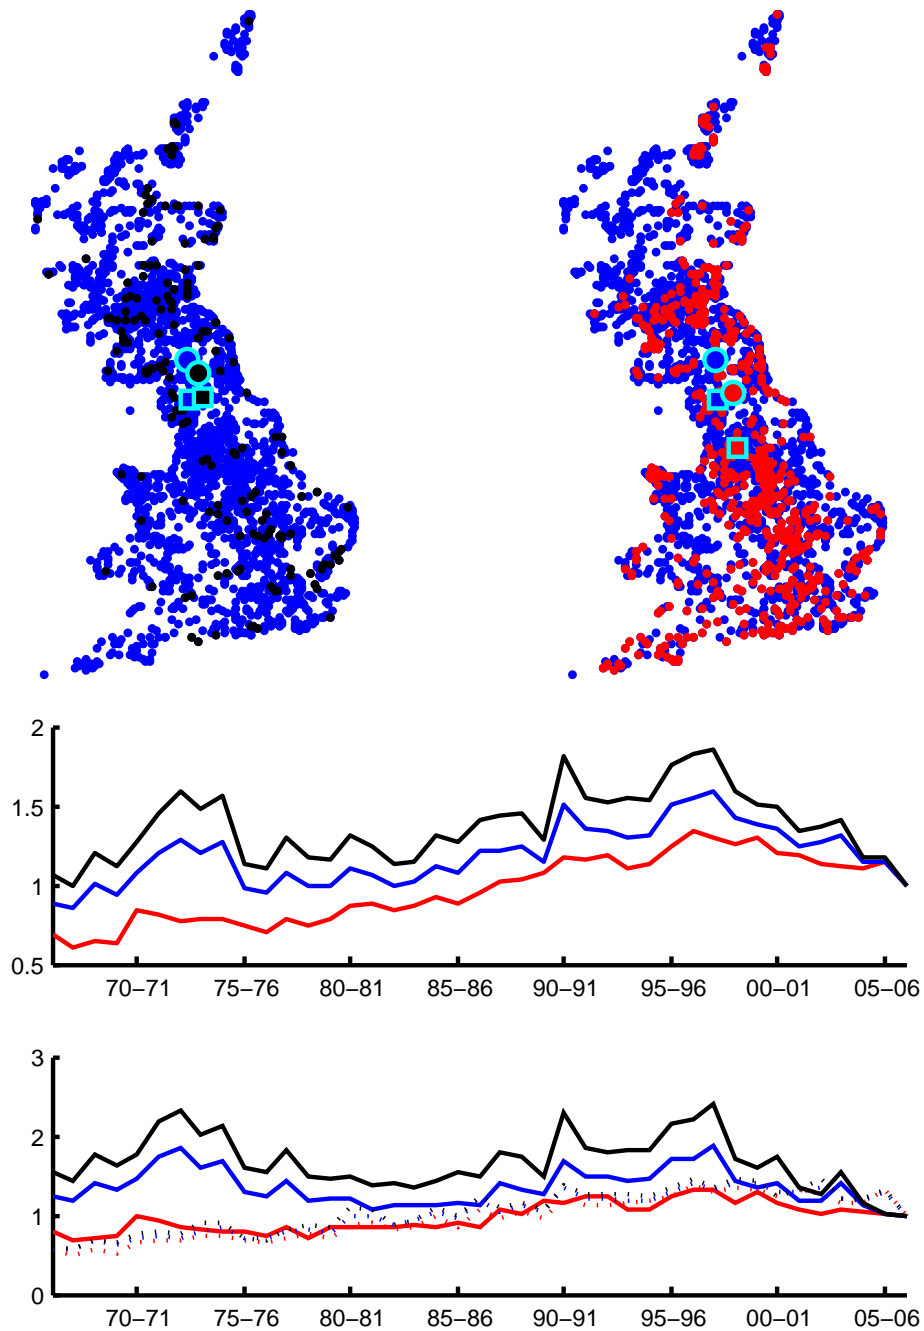


Figure 2.20: Comparison of site criteria for Goldeneye counts for November to March. Top: Map of all sites (blue) and sites selected by the proposed site inclusion criterion (left; black) and the 50% completeness criterion (right; red). The spatial mean of the site locations (squares) and site locations weighted by the site factors (large circles) for each set of sites are also shown. Middle: Underhill index using all sites (blue), proposed site inclusion criterion (black) and 50% completeness criterion (red). Bottom: As middle, but indices produced for sites north of the weighted mean of all sites (solid) and south (dotted).

The months chosen in Section 3.5.2 are used in Table 2.3 for comparison of the top sites criterion with the 50% completeness criterion. The Underhill algorithm takes longer to converge when the top sites criterion is used than the 50% criterion, because there can be individual sites with many missing values using the top sites criterion.

The ranking system used in the top sites criterion can also be used to highlight sites that are particularly influential in determining the twenty-six species indices. Usually influential sites will be well known waterbird “hotspots”; Appendix A contains a league table of sites using this method. Allocating resources (volunteer or professional counters) to the sites higher up the league table will limit the impact of missing and incomplete values on annual WeBS indices.




Figure 2.21: Proportion of missing values in set of Goldeneye data as first 300 sites are added in site factor order in the top sites procedure when using counts September to March (blue) and November to March (red).

	50% Criterion Sites					Top Sites				
	#sites	%mod	%miss	#iter	convergence	#sites	%mod	%miss	#iter	convergence
Mute Swan	404	36.6	(34.9)	31	8×10^{-06}	194	55.4	(44.4)	151	1×10^{-05}
Whooper Swan	273	32.0	(29.9)	19	6×10^{-06}	72	50.3	(42.3)	190	1×10^{-05}
Bewicks Swan	247	25.8	(23.1)	27	8×10^{-06}	57	37.5	(24.4)	500	7×10^{-05}
Pink-footed Goose	146	35.4	(32.5)	32	8×10^{-06}	56	54.5	(44.9)	202	1×10^{-05}
European White-fronted Goose	135	30.2	(26.7)	101	1×10^{-05}	23	39.9	(23.7)	144	1×10^{-05}
Greenland White-fronted Goose	46	34.1	(32.8)	29	9×10^{-06}	14	67.3	(52.3)	191	1×10^{-05}
Icelandic Greylag Goose	172	30.8	(28.4)	23	8×10^{-06}	44	51.3	(42.7)	237	1×10^{-05}
Re-established Greylag Goose	137	40.5	(39.2)	32	7×10^{-06}	58	60.8	(52.7)	272	1×10^{-05}
Canada Goose	253	40.0	(38.2)	37	8×10^{-06}	118	57.3	(48.8)	164	1×10^{-05}
Svalbard Barnacle Goose	42	37.0	(34.0)	117	9×10^{-06}	10	53.7	(41.6)	184	1×10^{-05}
Naturalised Barnacle Goose	88	39.0	(36.1)	24	1×10^{-05}	37	51.2	(44.8)	223	1×10^{-05}
Dark-bellied Brent Goose	107	27.3	(22.9)	19	7×10^{-06}	76	38.6	(24.9)	73	9×10^{-06}
Shelduck	332	29.2	(26.2)	24	6×10^{-06}	116	39.6	(24.8)	52	1×10^{-05}
Wigeon	697	28.5	(26.3)	23	7×10^{-06}	155	39.5	(25.0)	500	4×10^{-05}
Gadwall	254	37.1	(35.3)	29	9×10^{-06}	87	53.3	(42.1)	71	9×10^{-06}
Teal	660	30.4	(28.3)	21	8×10^{-06}	163	41.8	(28.6)	192	1×10^{-05}
Mallard	305	39.9	(38.3)	57	1×10^{-05}	229	58.2	(49.4)	500	5×10^{-04}
Pintail	335	29.5	(26.7)	27	9×10^{-06}	103	38.2	(24.8)	201	1×10^{-05}
Shoveler	280	36.8	(34.6)	36	8×10^{-06}	102	48.6	(37.4)	106	1×10^{-05}
Pochard	693	29.0	(26.9)	20	9×10^{-06}	134	35.1	(26.4)	229	1×10^{-05}
Tufted	420	37.3	(35.5)	48	9×10^{-06}	200	51.7	(46.2)	171	1×10^{-05}
Scaup	247	30.8	(27.9)	45	9×10^{-06}	43	42.5	(28.5)	112	1×10^{-05}
Goldeneye	693	29.3	(27.1)	17	7×10^{-06}	142	37.9	(26.6)	128	9×10^{-06}
Red-breasted Merganser	119	35.8	(32.2)	22	9×10^{-06}	67	53.5	(31.8)	165	1×10^{-05}
Goosander	487	30.0	(28.3)	18	8×10^{-06}	105	47.1	(39.0)	174	1×10^{-05}
Ruddy Duck	216	36.5	(35.1)	20	7×10^{-06}	50	43.7	(38.4)	106	1×10^{-05}

Table 2.3: Comparison of imputing with the Underhill model for sites chosen with the 50% criterion and the top sites criterion. The number of sites, the percentage of counts imputed by the Underhill model (including some incomplete counts) and the percentage of missing counts for the months chosen in Section 3.5.2, 1966/67 to 2006/07 are listed. Also given is the number of iterations for the Underhill algorithm to converge and the last value of the algorithm's convergence measure (see Section 4.2.2). The maximum number of iterations was set at 500 and the convergence criterion at 10^{-5} . Cases where the convergence had not been achieved by the 500th iteration are in bold.

2.5 Discussion

This chapter has considered the practicalities of surveying wildfowl across a country on a regular basis and the consequences for constructing national population indices.

Section 2.2 discussed sampling considerations that could impact on population trends. Although WeBS counts are often classed as bird census data, in reality the number of birds recorded is only an estimate of the number of birds present. The proportion of birds detected on a WeBS count depends on many factors, including the weather conditions, the amount of human disturbance and the number and skill of the observers. These factors may change systematically over time; for example, as a WeBS observer becomes more experienced and his or her knowledge of their site increases, the probability of finding flocks of birds may increase.

Estimating animal detection probabilities is an important part of estimating biological populations; although WeBS data is used to estimate population size at the national and site level, no formal assessment of individual species detection functions for WeBS has been undertaken. For population indices derived from WeBS data, if the proportion of birds detected by the WeBS methodology remains constant over time, then no knowledge of the underlying detection function is needed; however, if the probability of observing birds changes over time then the population index will be affected. Possible mechanisms for systematic change in detection over time include short-term effects (e.g. changes in behaviour due to moulting) and medium-term effects (e.g. habitat changes at a site such as increased marginal vegetation). However, it is long-term effects on a large spatial scale that would impact most on annual population indices. A possible mechanism for this is improvements over the past forty years in the capabilities and affordability of telescopes used for surveying wildfowl data. A comparison of wildfowl counts carried out at the same time with a range of ages of equipment would be a useful exercise to ascertain if this is likely to affect conclusions on long-term wildfowl population trends.

Some British ecological monitoring schemes, such as the British Breeding Bird Survey

(BBS), are moving towards a random sampling approach so that trends from the survey can be taken to represent the national situation. Section 2.3 discussed some of the complicating factors of designing a sampling scheme for surveying overwintering wetland birds. WeBS sites (particularly those which are regularly surveyed) tend to be large waterbodies which hold significant wildfowl and/or wader populations, close to urban areas. They are not chosen by a random sampling scheme, but rather by local knowledge and surveyor availability. A concern often expressed about the locations of WeBS sites is that they are concentrated close to human centres of population; whilst this is undoubtedly the case, if wildfowl are also concentrated in these areas then the bias introduced into population trends will be small. Wetlands are unevenly distributed across Britain and vary in type due to local geography. Wetland birds are also unevenly distributed in response to the type and amount of wetland habitat available and other factors such as climate and the historical range of a particular species. Any future sampling design for WeBS should take into account the amount of different wetland habitats available in an area, rather than being based on a grid-based system such as selecting 1km squares in the BBS. In recent years there has been a great increase in the availability and accuracy of geographic information; this should allow for a sophisticated sampling design using water GIS layers in the future.

Many of the sites in the WeBS database have only been surveyed a few times and are of little use for estimating long-term trends. Section 2.4 showed that the current method of site inclusion in indices using the 50% completeness criterion can omit important sites from the index even when there is significant data available for the site and can further emphasise a spatial bias towards the south-east of Britain which gives potentially misleading trends for species such as Goldeneye.

The more species a survey attempts to cover, the less the survey can be targeted at the ecology of each species. The Wetland Bird Survey covers not only wildfowl, but other taxa including herons, rails, cormorants and waders; even within the ducks, geese and swans the habitat requirements of species vary greatly. Perhaps of most importance for

consideration when designing a sampling scheme is the balance between surveying a representative sample of waterbodies used by widespread generalists and targeting the areas where those species who tend to winter at a small number of sites are to be found.

The optimum sampling strategy for producing population trends is not necessarily the same as that for producing national population estimates. To estimate abundance, the species' utilisation of different habitats must be considered, as well as ensuring the sample is spatially representative. However, if the population proportion on each habitat can be assumed to be unchanging (equivalently, population trends on different types of habitat can be assumed identical) then a representative sample of each habitat type is not required for producing population trends. This is a strong assumption and may not be ecologically valid; for example, if small-scale habitats such as ditches and farm ponds are being lost or degraded, populations may decrease or be forced to relocate to protected nature reserves. However, if trends for widespread species on those waterbodies where the more specialised species occur do reflect those on other wetland habitats, then surveying these sites gives the best return on survey effort. Another reason for concentrating effort on waterbodies which contain large numbers of individuals and species is that it encourages retention of volunteer observers.

Section 2.4.3 proposes that a new "top sites" criterion is used for deciding which sites to base population indices on, and that consequently survey effort should be directed to those waterbodies which either contain a large proportion of the population of one species, or significant populations of several species. A suggested prioritisation of survey effort by site is given in Appendix B. Switching to the use of the top sites criterion would result in trends only being representative of populations on wetlands with high concentrations of waterbirds. However, because of the construction of the Underhill index this is already true and is simply being made more explicit (this is why an index produced using the small number of sites chosen using the top sites criterion is very similar to that using all available sites).

The policy of what species trends should be prioritised in WeBS must be discussed and

decided at a strategic level. Is it more important for WeBS trends to be accurate for relatively common and widespread species such as Mallard and Mute Swan, or more localised species such as Dark-bellied Brent Goose and Pintail? If the latter than it is recommended that the top sites criterion is used for site inclusion in indices and that observer effort is directed at sites with large influence (as listed in Appendix B). If the former, then a formal GIS-based sampling scheme, together with the use of weighted habitat indices (see Chapter 5), should be a priority for future WeBS research.

This chapter has considered the question of which sites should be included when generating a population index. The next chapter considers a similar question for a wildfowl population index: which months should be included?

Chapter 3

SEASONALITY IN WILDFOWL POPULATIONS

3.1 Introduction

The notion of seasonality, or periodic fluctuations in data, has long been used to properly account for and model many processes. The most obvious field where seasonal fluctuations are of interest is meteorology; indeed, the United Kingdom's national weather service, the Met Office, now uses a mixture of statistical and physical computer models to create forecasts for entire seasons, such as the probability that coming winter will be warmer or wetter than usual (Met Office, 200-). The planet's seasons impact on animal, human and plant behaviour, resulting in a need for seasonal models in such diverse applications as population dynamics, retail sales, hospital admissions and carbon dioxide levels in the atmosphere. As a discipline, statistics has had to develop appropriate methods of dealing with the myriad consequences of Earth's axial tilt; with seasonal considerations being given particular emphasis in times series analysis and functional data analysis.

Population Dynamics

There are many areas in ecology where seasonality is of interest. Life histories vary considerably from species to species: some exhibiting migration patterns or breeding strategies completely defined by the timing of the seasons, whereas in other species, including *Homo sapiens*, seasonal effects are more subtle, although still present (Foster

and Roenneberg, 2008). As a result, the population* of a species changes throughout the year.

Fluctuations in populations over time, caused by birth and death processes and immigration and emigration, are modelled mathematically in the field of population dynamics. Incorporating stochasticity into mathematical models is a topic of current research. However, collaboration between the empiricists and the theoreticians - mathematicians, statisticians and biologists - is something that is in relative infancy (Ives, 2000). Interesting research is beginning to emerge: studies in the laboratory on flour beetles (*Tribolium*) have successfully shown in real population data chaotic and non-chaotic behaviour as predicted by mathematical models using modern statistical methods (Costantino et al., 1997; Perry et al., 1997).

Coulson et al. (2008) consider that within statistical ecology itself, traditionally there have been two separate approaches to collecting and analysing data related to population dynamics: modelling demographic processes such as productivity, senescence and survival from data collected by mark-recapture-recovery or other methods; and phenomenological time series models, where the population at time $t + 1$ is related to that at t by some function, perhaps involving covariates such as weather or carrying capacity. They were able to bridge the two strands of data analysis in work on Soay sheep population on a Scottish island, where an annual census of the islands population takes place in August each year. Their time series model is of the form

$$\begin{aligned} Population_{t+1} &= Population_t - deaths_t + recruits_t \\ &= survivors_t + recruits_t \end{aligned}$$

The number of survivors and recruits can be modelled with age-structured individual and environmental covariates (complete life histories of individuals are available for a sub-population), linking the demographic and time series approaches.

*Note that the term *population* is used here in the sense of a discrete or self-contained sub-population of a species.

The collection and analysis of Wetland Bird Survey data falls into the second tradition, recording month to month and year to year changes in population, but not collecting demographic data. Waterbird counts are an obvious choice for combined analyses of count data with demographic data since both are available. Besbeas et al. (2002, 2003) improved estimates of productivity in Lapwing, *Vanellus vanellus*, and the Grey Heron, *Ardea cinerea*, by integrating mark-recapture-recovery and “bird census data”, this being an index of the total population produced by the Common Bird Census scheme in the case of the Lapwing and an estimate of the total number of breeding pairs in the case of the Heron. The statistical analyses needed to produce the index from the lapwing census data are treated as a separate problem, so Besbeas and Freeman (2006) developed a method of using the demographic and survey data simultaneously, by noting that the year effects are correlated as they are a consequence of abundance, fecundity and survival in previous years.

Research is continuing into similar analyses, in part using WeBS data for Teal (Besbeas, 2007). However, it is important to note that in previous applications the census data comprises one value per year per site, whereas with WeBS data there are several monthly values; and unlike with studies of breeding populations and colonial nesters such as herons, month to month and year to year abundances are not necessarily correlated, since WeBS is surveying open populations in many cases. Weather has a large influence on the proportions of European wintering populations of wildfowl which are in the United Kingdom, apart from resident species such as Mute Swan and those geese and swan species where an entire sub-population winters together at traditional areas, such as Bewick’s Swan.

Collecting dynamic population data

When ecologists are attempting to estimate or monitor populations in the field, generally the aim is to estimate population parameters at a single point in time: a snapshot of the abundance of a species or “census” value. As discussed in Section 2.2, much re-

Year	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
1997/98	5	8	36	4262	17595	23797	16871	14225	11580	9962	10140	7
1998/99	5	3	3	13446	17736	11668	16630	12060	8949	8723	7552	16
1999/00	4	3	26	3802	28811	25002	17849	16321	14033	8721	8190	32

Table 3.1: Recorded counts of Dark-bellied Brent Goose at the Wash, a large estuary in eastern England; a well watched site, with counts available outside the core winter months. Note the abrupt changes in population as birds arrive at the site, with only 3 birds present in September 1998 but over 13,000 the following month.

search in the phenomenological branch of statistical ecology has consequently focused on overcoming errors in counts due to biases in detection probabilities. Often confidence around estimates is improved by replicating counts: concentrating resources in a short time frame to achieve good point estimates.

Wildfowl counts collected monthly as part of the Wetland Bird Survey are relatively unusual, as a sequence of point estimates of populations which are *expected* to be constantly fluctuating due to short-term immigration and emigration. Furthermore, these fluctuations are not necessarily relatively small, as would be the case with changes due to births and deaths: it is not unusual to go from having ten individuals to ten thousand individuals at a single site from one monthly count to the next (see Table 3.1 for an example).

There is only pseudo-replication of counts, in the form of a series of monthly values; but since we expect numbers to fluctuate, these are not real replications at all. Unfortunately, due to the lack of replicated counts at each site, these structural seasonal movements can be confounded by even shorter-term movements. For example, birds moving off-site to separate feeding grounds such as surrounding fields can make estimating numbers problematic, particularly when there is uncertainty as to whether the movement has been to a neighbouring WeBS site. Birds' feeding patterns may also change throughout the winter; for example, Light-bellied Brent Geese (*Branta bernicla hrota*) at a site in Northern Ireland followed a sequential pattern of habitat use, feeding on mudflats in autumn and early winter before moving to predominantly saltmarsh and farmland in late winter and early spring (Tinkler et al., 2009).

We have already noted that the most direct impact of seasons is on the weather. The

weather in turn affects survival and productivity of species: wildfowl for example will have lower survival in cold, frosty weather and lower productivity in droughts in summer, as they find it more difficult to access food from the water as well as being more vulnerable to predators. Low or high counts in a particular year could be merely the result of abnormal weather, but if climate change makes the “abnormal” increasingly “normal”, then long term abundance and distribution of species may be affected. The boundary between weather and climate is blurred: one hot summer does not indicate anthropogenic climate change, but a long run of hot summers may constitute evidence for it. Correspondingly, the spectrum of investigating climatic effects in statistical ecology ranges from incorporating weather into survival models to studying changes in seasonal animal behaviour in the field of phenology (Catchpole et al., 1999; Sparks and Carey, 1995).

Phenology

Phenology is a branch of ecology wholly concerned with the impact of seasons on the timing of natural events, investigating changes in an organism’s seasonality over time. Phenological processes in nature have gained more public interest in recent times and also have risen in scientific prominence, as both signals of climate change and of its potential consequences. There has been concern among conservationists of the negative impacts of phenological change on species already under pressure from mankind; for example Pearce-Higgins et al. (2005) showed changes in the timing of Golden Plover (*Pluvialis apricaria*) egg-laying dates and the emergence of their prey, putting the two out of step with each other, could result in poorer breeding success. Similarly, Both et al. (2006) found that declines in Dutch Pied Flycatcher *Ficedula hypoleuca* populations were much greater in areas where breeding was mistimed, not coinciding with the time of maximum food abundance.

Given the wide range of natural events for which there is evidence of climate change related trends in phenology, it would be surprising if there was no changes in the monthly phenology of overwintering birds over time. Maclean et al. (2008), using mid January

counts of waders from the Wetland Bird Survey and similar surveys from other parts of northwest Europe, established that population change are correlated with changes in temperature and that the weighted geographical centroid of seven species had shifted by up to 115km in response to climate change.

Chapter Overview

In this chapter, our starting point is the concept that we have a species population that is fluctuating in time and space as a consequence of regular seasonal drivers. The population in question may be on a site, regional or national scale. In Section 3.3 the relatively recently developed philosophy of Functional Data Analysis (FDA) is introduced. This allows us to expand on the initial idea that we are sampling a population function of time. We also see that a recently published analysis of WeBS data, to produce population alerts when there are stark population changes, can be considered within a Functional Data Analysis framework.

We then take a fresh approach to analysing wildfowl counts from the Wetland Bird Survey using some FDA techniques. In Section 3.4, we explore phenological change in Pintail and Shoveler. We then return to the problem of generating a representative annual index in Section 3.5, where we attempt to identify when the population fluctuations are small enough on a national scale to consider the population in some sense “stable”, so that we can treat the monthly counts as replicates if needed for a particular statistical method or analysis.

3.2 Intra-winter Wildfowl Movements

3.2.1 Introduction

In Section 2.2, an overview of observational error in detecting wildfowl present at sites was given. In this section, some of the challenges of surveying a dynamic population are considered, using a case study to illustrate why estimating annual indices is difficult, even if there is no measurement error.

In Section 1.1.2, the wildfowl species covered by this thesis were broadly categorised into **resident species**, **mixed resident and immigrant wintering species** and **immigrant wintering species** according to their migratory status. However, this considerably simplifies the different scales of migration that species undergo. For many wildfowl species, the birds present in the winter months in Great Britain will be a mixture of resident birds, passage migrants and wintering birds. Even within the same species, some individuals that have bred in Britain will stay at the same site for most or all of the winter, whereas some will move to other sites within the country or go abroad. Some sites will have passage migrants that are heading for other sites in Britain whereas others will only be present in the U.K. for a short while, before moving to other parts of Europe or Africa. Birds that come from the north to winter in north-west Europe sometimes spend some of their time in the Netherlands before moving west to England, and may move between sites and regions within Great Britain. To add to an already confusing picture, many wildfowl species spend some time feeding away from waterbodies, for example on grazing marsh, and in wet periods may move to flooded fields which will not be counted by WeBS.

To illustrate and explore some of these issues, and the impact they can have on the monthly counts and the final annual index, the Pintail (*Anas Acuta*) is used as an example. Pintail is a highly clumped species in northwest Europe, with most birds being aggregated in a small number of estuaries. The top sites algorithm (see Section 2.4.3) selects 114 sites from the available 1473 sites for the months September to March, which

have a combined site factor influence of 88%. An intensive survey of wetlands in north-west England found an increase of only 0.1% in the number of Pintail counted compared to just the WeBS sites and it is likely that WeBS coverage of the species is almost total (Quinn and Kirby, 1992; Cranswick et al., 1997).

Although only a small proportion of the world population of Pintail winter in Britain, about half of the Pintail in north-west Europe in midwinter do so. Recorded population declines on both European breeding and wintering grounds, possibly due to loss of wetland habitat and hunting pressure, have designated European Pintail as having 'unfavourable conservation status' and so, as with many wildfowl species, conservationists are keen to monitor the status of the British overwintering population.

Pintail was classified as a immigrant wintering species in Section 1.1.2 . A small number of Pintail breed in Britain, with an estimated breeding population of fifty pairs, which are thought to remain in the British Isles over winter. Most of the British overwintering population comes from further east, from Scandinavia, the Baltic States and Russia. Within-winter recoveries of birds ringed in Britain and Ireland in autumn indicate that there is some passage of birds that travel onwards to the Mediterranean and North Africa; it is unclear whether these individuals return in the Spring by the same route. It is thought the entire Icelandic population, at about five hundred pairs, overwinter in the



Figure 3.1: Male Pintail (*Anas Acuta*) at Martin Mere, Lancashire. ©Gidzy (2009)

British Isles; two Icelandic-ringed Pintail have been recovered in Britain in October, but it is unknown when they arrive in general. In addition there are sometimes cold weather movements from Germany, Denmark and the Netherlands to Britain during the winter (Wernham et al., 2002; Snow and Perrins, 1997).

3.2.2 *Short-term and Weather-related Movements*

Despite being well-represented on WeBS sites, Pintail is a highly mobile species during winter. Short-term changes in Pintail distribution can disguise seasonal patterns. The duration of these movements include:

- **Hours** Pintail at some sites feed on nearby farmland and so the timing of the count during the day will affect the proportion of birds recorded: at the Ouse washes site, two feeding visits to farmland were recorded, one from dawn to just after sunrise and the other from mid-afternoon to dusk (Kershaw, 1998). An apparent change in the Pintail population of a site may be due to the survey taking place at a different time of day.
- **Days** Frequent counts at some WeBS sites have shown that local changes in distribution due to Pintail mobility causes numbers at a site to fluctuate tremendously over a period of days (Kershaw, 1998). As well as moving between nearby sites within the scheme, Pintail will use habitats temporarily available due to flooding (Brown and Smart, 2002). Thus a drop in recorded birds in one monthly count may be related to wet weather, rather than movement away from the U.K., since temporarily flooded sites are not surveyed by WeBS.
- **Weeks** Numbers of Pintail present will also be affected by severe winter weather, with movements of birds wintering in the Netherlands to Britain (Snow and Perrins, 1997). Weather movements may thus obscure the underlying spatial pattern of Pintail wintering in North-west Europe, due to short-term variations in the proportion recorded in the monthly counts in each region.

Monthly counting under the WeBS scheme cannot capture this information, with all of these movements having the potential to greatly affect the recorded count for any one month. Aside from general feeding forays, it is clear that the weather has a large influence on where Pintail are to be found. Cold weather on the continent may result in higher numbers on U.K. WeBS sites, but if the freezing weather occurs in Britain as well it may also mean fewer Pintail on non-coastal sites, as these will be more affected by ice and snow. Wet weather may also result in fewer Pintail at WeBS sites if there is more food to be found on flooded farmland; but this will depend on the extent and severity of the wet weather.

At an individual site level, there is considerable variation between years in the pattern of counts (for example, see Figure 3.2). Some of these apparent changes may be due to the site's population having temporarily moved for feeding or due to weather but others are probably indicative of the flexibility of Pintail overwintering and migration habits, which mean sites are used at different times and different extents from one year to another, both within Britain and in the rest of the Western Palearctic.




Figure 3.2: Pintail counts at Burry Inlet 2002/03 to 2006/07.

Movement between neighbouring sites is known to exist from frequent counts (Kershaw, 1998). The county of Kent has over a hundred sites that have been surveyed as part of the Wetland Bird Survey, but only twenty of these have ever recorded Pintail and three sites have significantly more records than the others, these being the Thames, Swale and Medway estuaries. As can be seen from the map in Figure 2.10 on page 64, these sites are all adjacent to each other and are complex sites, made up of multiple count areas. Although the infrequency of WeBS counts does not lend itself to tracking movements between sites, it is reasonable to assume that birds move frequently between these three areas; looking at the breakdown of Pintail between the three sites in Figure 3.3 there is a suggestion that there are occasions, such as mid-winter 2005, where the overall numbers of Pintail in the general area is relatively unchanged, but the proportions on the Swale and Medway sites fluctuate.

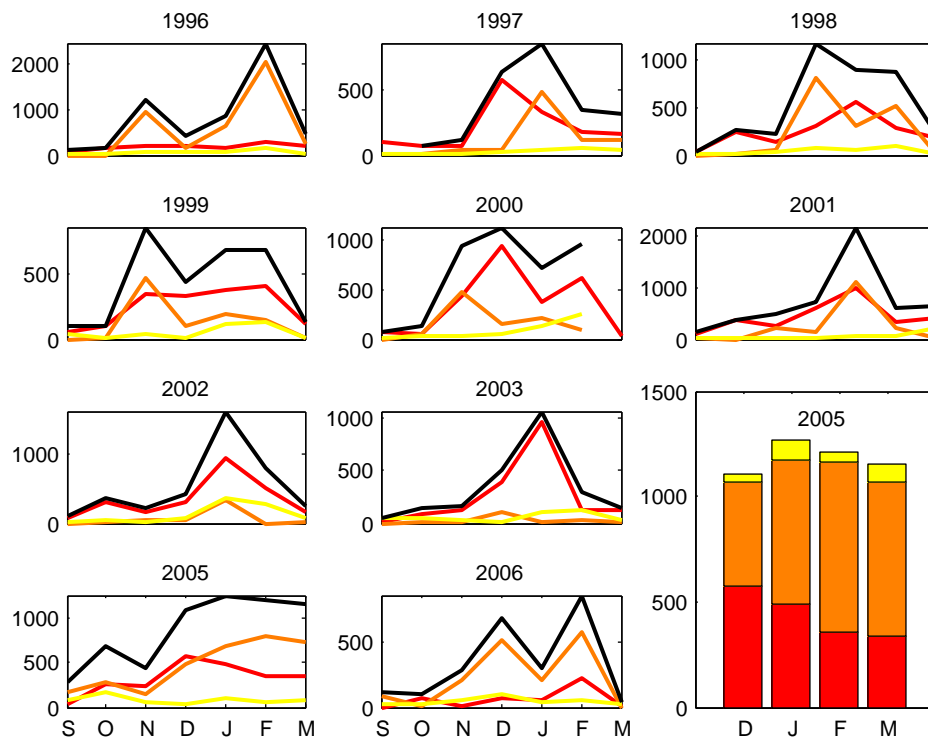


Figure 3.3: Recent winter counts of Pintail for the Thames (yellow), Medway (orange) and Swale (red) estuaries. The total count for the three estuaries is indicated in black. Counts for December to March in 2005 are also shown as a stacked bar chart.

Species vary in extent of their short-term and weather-related movements depending on their ecology; but most wildfowl species exhibit some degree of winter mobility in the same way as Pintail.

3.2.3 *Weather covariates in modelling Wetland Bird Counts*

There is considerable scope for further analysis of the impact of weather on WeBS counts, and possible incorporation of weather covariates in future modelling of counts. Appropriate weather covariates will depend on what is known of the individual species' biology. The winter North Atlantic Oscillation Index (NAO) is useful as an indicator of the general climate; it is calculated as the mean of the the pressure difference between Gibraltar and Iceland for December to March inclusive (Osborn et al., 1999). The winter NAO values are published online for 1994/95 onwards and can be calculated using downloaded monthly NAO data for 1966/67 to 1993/94 (Osborn, 2009; Salmon, 2004). The NAO index can be used to explain variation in many areas of ecology, including bird migration phenology (Rainio et al., 2006).

The winter NAO index has potential for modelling abundance in WeBS counts for species whose populations in Britain or in parts of Britain are strongly dependent on large scale weather patterns. For example, a regular wintering population of Bewick's Swans at Slimbridge in Gloucestershire, England (part of the Severn Estuary site) was established in 1955/56 and its growth was thought to be related to the cold winters of 1961/62 and 1962/63 bringing the birds further west (Evans, 1979). The correlation between the winter weather, as summarised by the winter NAO index, on the numbers of Bewick's swans at Slimbridge can be seen in Figure 3.4. Peak counts tend to be lower in years with a high NAO index, which usually indicates wet, mild conditions in north-west Europe, and higher in years with a low NAO index, corresponding to colder, drier conditions which tend to force the birds westwards to the Severn Estuary.

As well as the NAO Index, when modelling counts at an individual site level, data from




Figure 3.4: Correlation between the winter North Atlantic Oscillation weather index and peak count of Bewick Swan at the Severn Estuary (English counties). Pearson correlation $r = -0.38$ ($p = 0.015$).

nearby Met Office weather stations is available for many sites, which can return weather conditions of the day of the count itself (Met Office, 2006). Analyses using local data such as this could give greater insight to the effect of immediate weather conditions to recorded numbers and the subsequent effects on the national index.

3.3 Functional Data Analysis: an Intuitive Way of Looking at Population Dynamics Data

3.3.1 Introduction

Functional Data Analysis (FDA) is a collection of statistical techniques suitable for use on functional data. Functional data are typically of the form of a sample of n observations of an underlying continuous function and occur in a wide variety of applications, from analysing growth rates in children to weather data to analysing handwriting. The fundamental philosophy behind FDA, introduced in the books by Ramsay and Silverman (2002, 2005), is to consider the data functions as single entities rather than as a series of related observations; many of the techniques use ideas from time series analysis, multivariate analysis and longitudinal analysis, all set in a functional framework.

3.3.2 WeBS data as functional data

Why FDA for WeBS data?

Functional data often come to us as a single long time series spanning many days, months, years or other time units. The variation in data such as these is usually multilevel in nature. There is usually a clear annual, diurnal or other cycle over the basic time unit called the *season* of the data, combined with longer-term trends that span many time units. Moreover, the seasonal cycle may also show some evolution over the time spanned by the series.

Ramsay and Silverman (2002)

In Chapter 2 it was seen that WeBS observations have been made at many different sites around Great Britain at monthly intervals for over forty years. Often, wildlife counts are treated as discrete entities; indeed the results of bird counts are sometimes referred to as

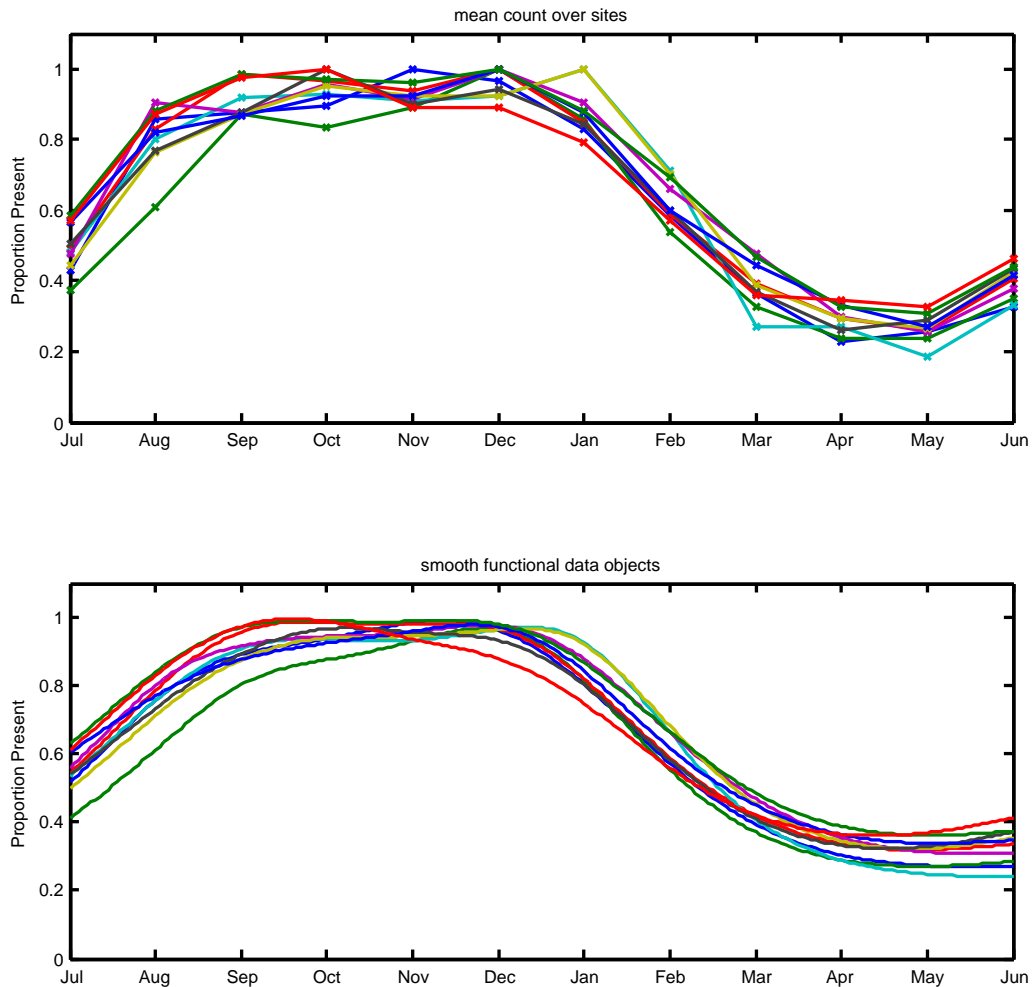


Figure 3.5: Example of functional data objects. Top: Mean monthly counts (ignoring missing data) of Mallard data over 349 sites (chosen by the algorithm in Section 2.4.3) for 1997/98 to 2006/07. Each year is a different colour. The data points are standardised by dividing by the maximum of the monthly means for that year. Bottom: Functional data curves of the data using the `smooth_pos` routine (see Section 3.3.3).

bird “census” data, implying that it is possible to confidently estimate the entire population (Bibby et al., 1992). In the introduction to this chapter, however, it was established that the wildfowl populations studied are expected to be constantly changing as birds move in and out of the country. Monthly counts are not independent estimates of an annual population, but correlated estimates of a population that is varying constantly over time. WeBS data are thus neatly encompassed in the concept of time series functional data introduced in the quotation from Ramsay and Silverman (2002) at the beginning of

this section, and, like some other examples of functional data, although there are not independent replications, there is repetition of information that can be exploited (Ramsay and Silverman, 2005).

For WeBS wildfowl counts, the underlying theoretical population function is the total population of a species present in Great Britain at a particular instant in time. In practice, this cannot be directly measured; so instead the total population on WeBS sites at a particular instant in time is estimated by summing site counts (for example, see Figure 3.5). Although it is tempting to consider the population at individual sites to themselves be a population function, there is unlikely to be much to be gained from treating the site counts as functional rather than multivariate data. Due to most species' mobility in winter, numbers can change abruptly from one time point to the next (for example, see Section 3.2 and Figure 3.2). National populations will vary in a more continuous fashion from one day to the next, with the population changing gradually as migrants cross the sea; the monthly counts are a snapshot of this process. However, there are undoubtedly individual species and sites where a functional approach would be appropriate at a site-level; such as Bewick's Swans at Slimbridge in Gloucestershire, where numbers build up steadily and there is little turnover in individuals over the winter.

WeBS Alerts: FDA in current WeBS analysis

One purpose of smoothing functional data that exhibits variation on multilevel timescales is to remove variation that is too short-range to be relevant to the current enquiry. This is conceptually different to smoothing to reduce the effects of measurement error in data: the short-term fluctuations may be real, but a distraction to a longer-term pattern. For example, wildfowl abundance naturally varies from year-to-year, particularly between mild and cold winters. Smoothing away these short-term fluctuations removes variation that is not giving useful information when trying to ascertain structural changes in abundance over time; with the amount of smoothing needed depending on whether one is interested in short-term, medium-term or long-term trends. To this end, a system for generating

conservation “alerts” to highlight population declines of at least 50% over various time scales has been developed (Austin et al., 2004; Atkinson et al., 2006).

Atkinson et al. (2006) introduces analysis of WeBS data using a Generalised Additive Models (GAM) to obtain a smooth trend in annual WeBS indices. The technique does not explicitly mention Functional Data Analysis, but generates curves using a GAM equivalent to the Underhill model, with the unconstrained year factors replaced by a smooth year curve. The methodology is a two-stage process: firstly the standard WeBS Underhill Generalised Linear Model (GLM) is applied, imputing missing values by the Underhill modified EM algorithm (see Section 1.3.2); then the GAM is applied to the data after imputing.

Since Atkinson et al. (2006) are applying the Underhill model before the GAM, the spirit of the “WeBS Alerts” method is in some ways closer to that of FDA than the more formal methodology of GAMs. By fitting the Underhill model first to impute missing and incomplete counts, the year factors from the GLM will determine the resulting smooth year fit in the GAM, and the month and site factors fitted by the GAM will also be affected by those fitted by the GLM. The GAM routine used fits the smoothed index curve non-parametrically, using smoothing splines, in a very similar way to that done in this chapter to monthly counts in individual years.

It is interesting to observe that data from the Wetland Bird Survey is already being considered in a functional way. In the remainder of this chapter a functional data approach is used to better understand seasonality in WeBS data.

3.3.3 *Fitting WeBS population functions*

The first step in a functional data analysis is to convert observed values, m_k , into a function, $v(t)$, that is computable for any desired argument, t . If the discrete values m_k are assumed to be exact, then it is appropriate to interpolate between the data points. For example, in the top plot in Figure 3.5 on page 103, showing the mean count over 349

sites for Mallard in each of ten years, the straight lines between the month values are a simple method of interpolating between data points. However, each of the month means are subject to errors arising from observation error at each site (as discussed in Section 2.2) as well as uncertainty from missing and incomplete information* and birds moving on, off and between sites (as seen in Section 3.2). It is, therefore, appropriate to perform a certain amount of smoothing to the data, as has been done in the bottom plot in Figure 3.5.

There are many methods for fitting a smooth curve v , such as fitting a Fourier series to the data or by polynomial smoothing. Functions may be represented by basis functions, ϕ_n , which can approximate another function, \mathbf{x} , by a linear combination of some number N of the basis functions:

$$\mathbf{x}(t) = \sum_{n=1}^N c_n \phi_n(t). \quad (3.1)$$

The choice of basis, such as Fourier series or spline functions, depends on the application, with a Fourier basis often being used for periodic data and a B-spline basis for non-periodic data. A spline function is constructed by splitting the interval over which v is to be approximated into subintervals; the separation points are called “breakpoints”. A polynomial of order κ (corresponding to a degree of $\kappa - 1$) is used for each subinterval, and adjacent polynomials are constrained to join smoothly at the breakpoints, as are their derivatives up to order $\kappa - 2$ (Ramsay and Silverman, 2002).

Functional Data Analysis functions for R, MATLAB and S-PLUS are freely available, and the MATLAB functions were used to analyse WeBS wildfowl data (Ramsay, 2006; Ramsay et al., 2009). Although twelve monthly values are shown in Figure 3.5 for the resident Mallard, observations are only available for September to March at many sites, and most species are only present in large numbers in the winter. For a twelve month period a Fourier basis may be used, but when fitting a curve to seven monthly values (or nine, as used in Section 3.5) a simple smoothing spline is appropriate. To ensure that the

*Note that the mean of available counts is used in this chapter, to avoid introducing strong assumptions about missing data, such as those in the Underhill model discussed later in Section 4.3.2.




Figure 3.6: The nine B-spline basis functions of order four defined by the breakpoints shown as red dotted lines. As WeBS counts take place close to the middle of the month, 1 corresponds to mid-September and 7 to mid-March.

function is positive the curve, $v(t)$, based on the month values, m_k , was defined as the exponential of a function $\mathbf{x}(t)$:

$$v(t) = e^{\mathbf{x}(t)}. \quad (3.2)$$

$\mathbf{x}(t)$ can be expanded in terms of basis functions $\phi(t)$ as in Equation 3.1 above. The MATLAB function used was `smooth_pos.m` (Ramsay, 2006) which uses a roughness penalty approach (Green and Silverman, 1994) to fit the data. Piecewise cubic splines ($\kappa = 4$) were used as a basis, with breakpoints at each month as illustrated in Figure 3.6. The large number of breakpoints (equal to the number of data points) allows more flexibility in the spline, needed because of the curvature required for abrupt changes in the population from month to month. Since here $\kappa = 4$, the function and its first and second derivatives are required to be continuous at the break points. The roughness penalty prevents rapid fluctuation in $\mathbf{x}(t)$ by constraining the integrated squared second derivative, a

global measure of roughness. The fitting criterion is the penalised mean square error:

$$PENSSE(\lambda) = \sum \mathbf{x}_k [m_k - v(t_k)]^2 + \lambda \cdot \int [D^2 \mathbf{x}]^2. \quad (3.3)$$

Figure 3.7 shows an example of the observed month values m_k , the fitted function $\mathbf{x}(t)$ obtained by setting $\lambda = 0.01$, stored as a linear combination of the basis functions in Figure 3.6, and the resultant fitted curve $v(t)$ which is always positive.




Figure 3.7: Bottom: Observed mean of available monthly counts of Gadwall data from September 1971 to March 1972 over 91 sites, chosen by the algorithm in Section 2.4.3. The data points, m_k , are standardised by dividing by the maximum of the monthly means to give values between 0 and 1 (blue crosses). The fitted smoothed functional data curve $v(t) = e^{\mathbf{x}(t)}$ of the observed values is shown as a black line. Top: $\mathbf{x}(t)$ fitted by the `smooth_pos` routine with $\lambda = 0.01$.

Smoothing Parameter

The parameter λ determines the smoothness of the fit, as can be seen in Figure 3.8 comparing three values of λ . Small values of λ result in a function $v(t)$ that closely follows the observed values m_k . Increasing λ gives a smoother curve that has a reduced tendency to follow short-range fluctuations. A value of $\lambda = 0.01$ balances retaining enough of the seasonal change to allow analysis of phenological change with reducing the influence of outliers, which may be the result of a brief spell of inclement weather, or missing data from key sites. $\lambda = 0.01$ is used for the analyses that follow in Sections 3.4 and 3.5 and the robustness of the following results to the choice of λ is discussed in Section 3.6.

Derivatives

One of the main characteristics of functional data analysis is the possibility of using rates of change from smooth functional curves to illuminate problems and gain extra information. The curve \mathbf{x} and hence v is twice differentiable to give v' and v'' . To ensure that the second derivative, v'' , is smooth (i.e. has no rapid fluctuation), basis functions with an order $\kappa = 6$ can be used. However, as can be seen in Figure 3.9, because in this case we have a large number of break points, using $\kappa = 4$ gives sufficiently smooth curves v' and v'' to be able to plot them and see the WeBS function from a different perspective by seeing how the derivatives change over time. By plotting them against each other in a *phase-plane plot* an alternative view of the seasonal dynamics is obtained. The first and second derivatives of the WeBS functions are used in Section 3.5 to explore the notion of finding months where the population is in some sense stable.

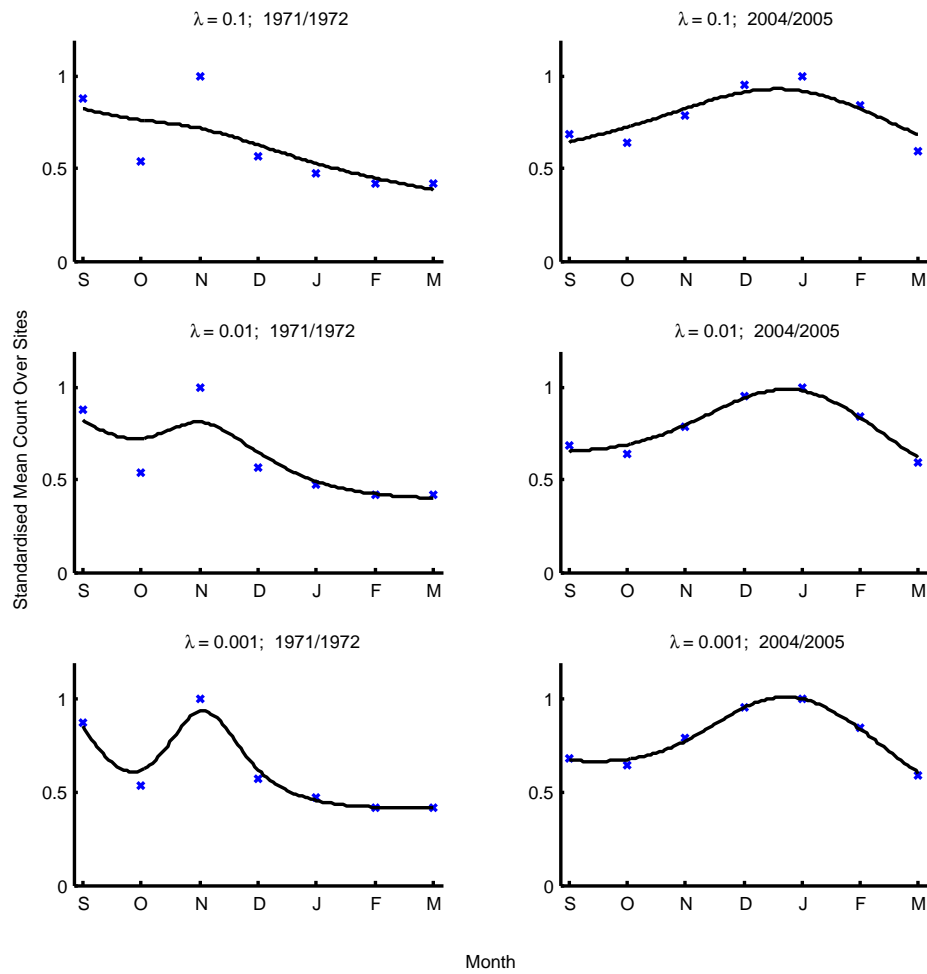


Figure 3.8: Effect of changing the smoothing parameter λ . Mean of available monthly counts of Gadwall data over 91 sites (chosen by the algorithm in Section 2.4.3) for 1971/72 and 2004/05 (blue crosses). The data points are standardised by dividing by the maximum of the monthly means for that year. Also shown are smoothed functional data curves of the data using the `smooth_pos` routine (black lines) with $\lambda = 0.1$ (top), $\lambda = 0.01$ (middle) and $\lambda = 0.001$ (bottom).

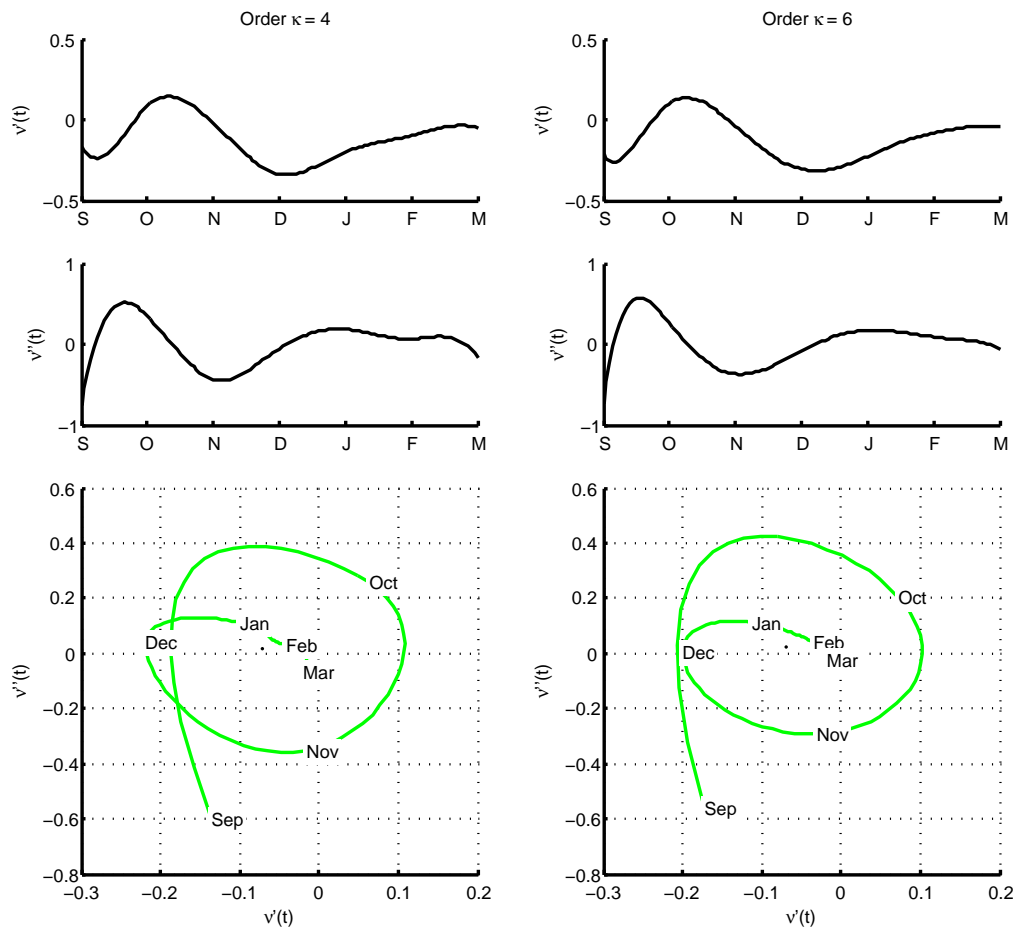


Figure 3.9: First ($v'(t)$) and second ($v''(t)$) derivatives of the fitted smoothed functional data curve $v(t)$ from Figure 3.7 against time (top and middle) and each other as a phase plane plot with time indicated by month labels (bottom). The basis functions used to fit $x(t)$ have an order of $\kappa = 4$ in the left-hand plots and $\kappa = 6$ in the right-hand plots.

3.4 Wildfowl Phenology: Spatiotemporal Variation

3.4.1 *Phenological change in Shoveler*

As introduced earlier in this chapter, phenology is the study of changes in the seasonality of a species over time. Smoothing monthly counts for each year to create functional data curves using the method described above in Section 3.3.3, helps highlight changes over time in species such as Shoveler. Inspection of the fitted Shoveler functions in Figure 3.10 suggests that in more recent years the Shoveler population is steady from September to March, contrasting with the pattern in earlier years of larger numbers in autumn than winter and spring.

A principal components analysis (PCA) of the monthly mean counts can be used to reduce the dimension from seven monthly counts to one or two combinations of the monthly counts that contain most of the variability. An equivalent technique can be used for functional data (Ramsay and Silverman, 2005). Functional principal components analysis finds curves (“harmonics”) that summarise the variation in the data. Routines for functional PCA are available for MATLAB (Ramsay, 2006); the `pca.fd` procedure was used on the 41 fitted functional data curves shown in Figure 3.10, subtracting the mean from each observation, to summarise the variability in a smaller number of curves. The first seven principal components contain 99.99% of the variation and the first four components 98.89%; these four can be seen in Figures 3.11 and 3.12.

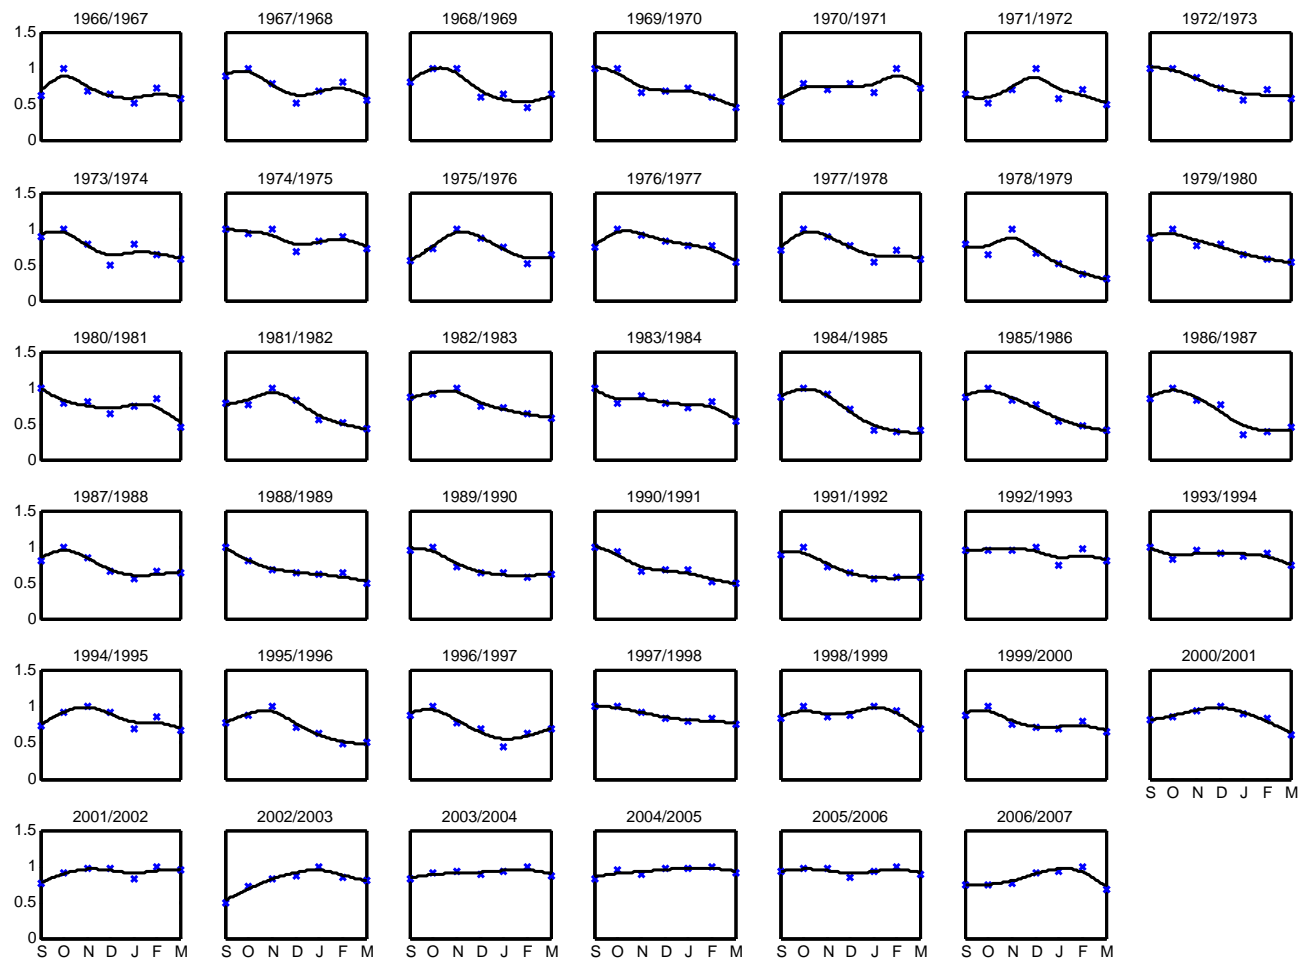


Figure 3.10: Phenological change in Shoveler. Mean monthly counts (ignoring missing data) of Shoveler data over 118 sites (chosen by the algorithm in Section 2.4.3) for 1966/67 to 2006/07 (blue crosses). The data points are standardised by dividing by the maximum of the monthly means for that year. Also shown are Functional data curves of the data using the `smooth_pos` routine (black lines).




Figure 3.11: Functional Principal Component Analysis of phenological change in Shoveler. Harmonics 1 (blue, 70.7% of the variation), 2 (green, 14.8%), 3 (red, 9.9%) and 4 (turquoise, 3.4%).

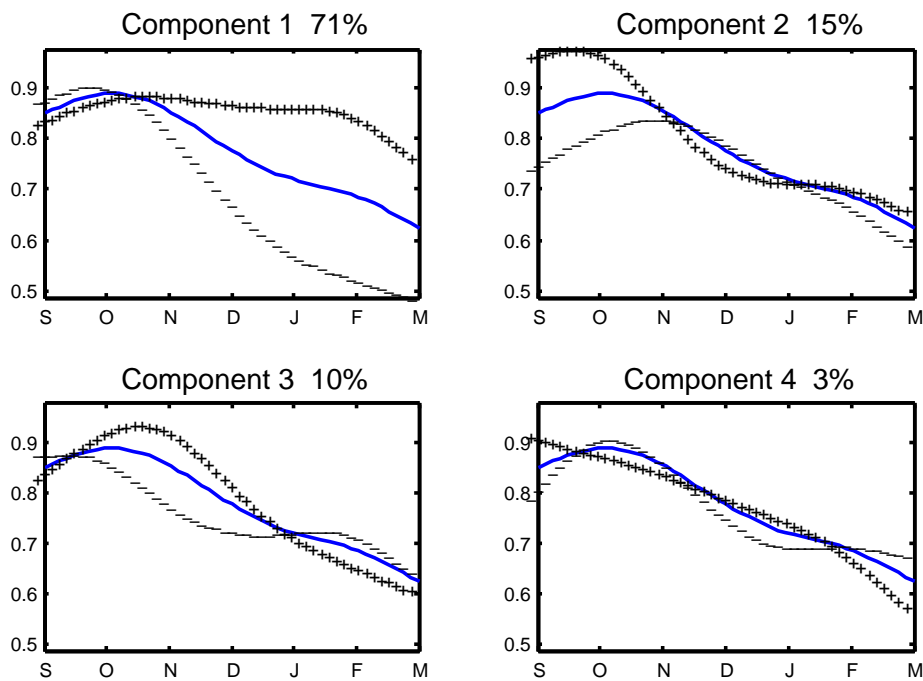


Figure 3.12: Functional Principal Components Analysis of phenological change in Shoveler. A multiple of each of the four harmonics from Figure 3.11 is added (++) and subtracted (--) from the mean function (blue).

From Figure 3.11 it can be seen that a year which has a lower than average proportion of birds present in September/October and a much higher proportion than average in birds present in November to March (particularly February) will have a high score on the first harmonic. Figure 3.12 illustrates the shapes where the annual curves from Figure 3.10 will have high and low scores on the first four principal components. A year where the number of birds remains at a similar level throughout the winter to the population present in September/October will have a strongly positive score on the first component, such as in 2004/05. A year where there are many more birds present in the autumn compared with the rest of the winter will have a strongly negative score, such as 1984/85. The first principal component thus measures the “flatness” of the within year population functions and accounts for 71% of the variability in the curves. A year has a high score on this component when the population stays level from September to March and a low score when the population declines over the winter from an October peak.

The second principal component accounts for 15% of the variability and years have a high score when the standardised population is particularly high in September and October. 1971/72 has a very low score on this component as there were relatively few birds present in the autumn compared to mid-winter. The third component has a high score when peak numbers of Shoveler are present in October and November with the population declining steadily to a minimum in February and March and a low score when peak numbers present in September decline to December and remain steady for two months before declining further. This accounts for a further 10% of the variability in the population functions; together the three components contain 96% of the variability between the population curves.

The scores of the first four principal components are plotted against time in Figure 3.13, providing an immediate summary of how Shoveler seasonality has changed over the past forty years. The first two principal components are also plotted against each other and the years grouped into two time periods 1966/67 to 1991/92 and 1992/93 to 2006/07. The first principal component shows a trend over time towards the population remaining

steady from September to March rather than declining from an autumn peak to a spring low, whereas there is no apparent trend in the second, third or fourth components which capture additional year to year variability. The first principal component is thus a convenient way of summarising the pattern which could be seen by inspecting each yearly curve individually in Figure 3.10.

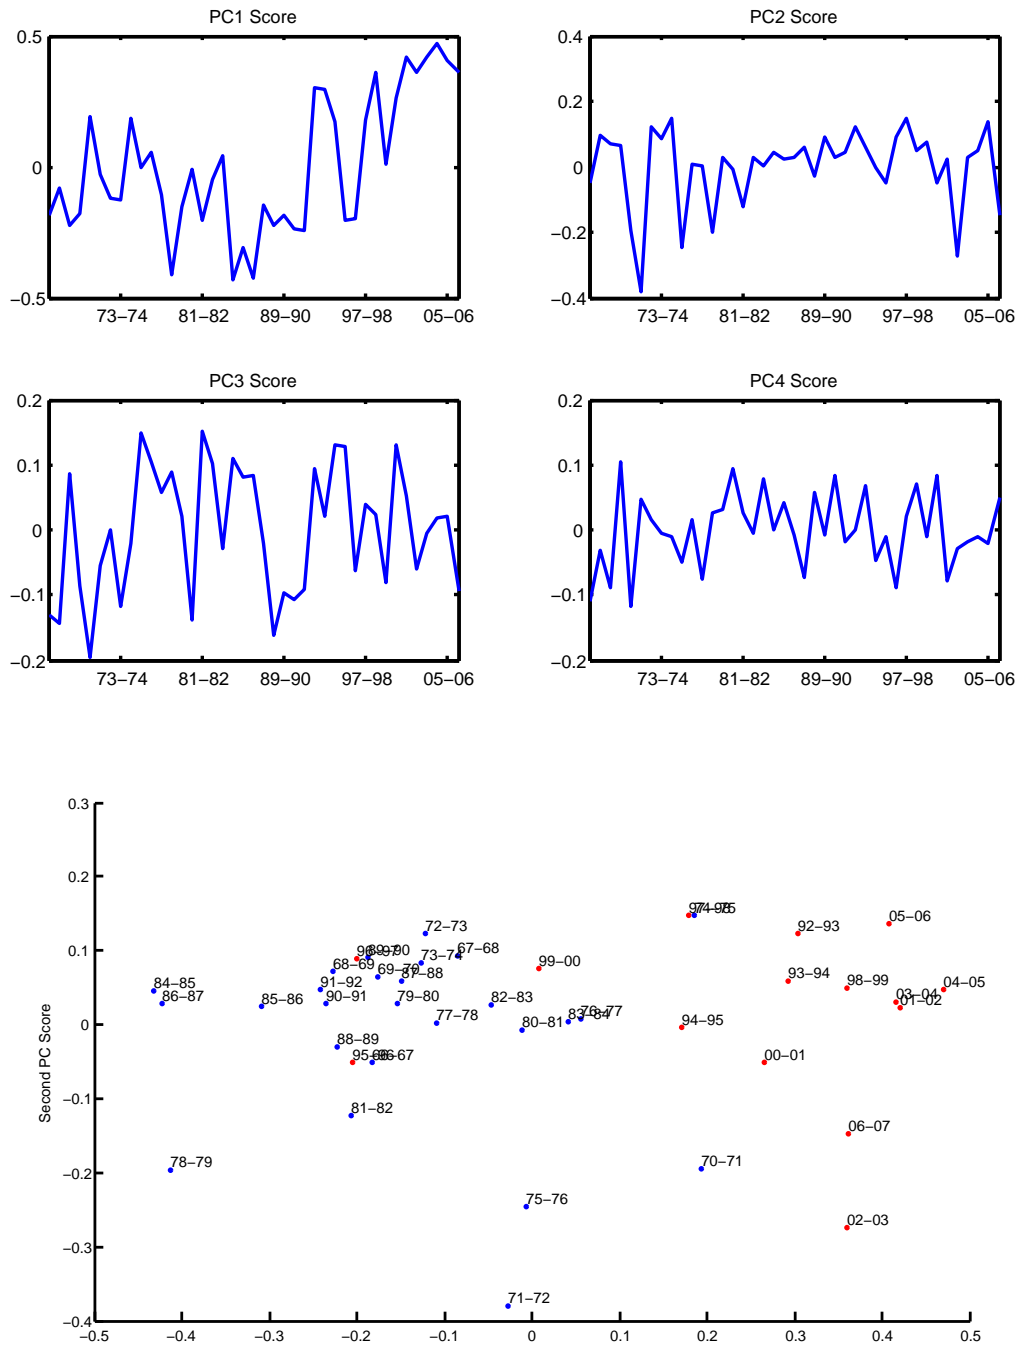


Figure 3.13: Shoveler principal component scores. First four principal components against years (top) and PC1 against PC2, with red markers for scores for years 1966/67 to 1991/92 and blue markers for scores for years 1992/93 to 2006/07 (bottom).

3.5 Identification of Overwintering Wildfowl Populations

3.5.1 *Introduction*

It is generally considered desirable that the WeBS annual index numbers can be considered representative of overall numbers of wintering individuals (e.g. Atkinson et al., 2006). Since multiple counts are available, indices based on several counts from each site during each winter are less likely to be influenced by unusual counts caused by short-term movements. However, a blanket approach of using all available months may also be inadvisable. More counts are available for mid-winter months than in autumn and spring (and summer), particularly historical counts, so by restricting the months to those important for the species in question unnecessary imputing of missing counts may be avoided. Using mid-winter counts only will also avoid annual indices being unduly influenced by influxes of passage migrants, or wintering birds arriving earlier or leaving later than usual which would increase the mean count over the months, resulting in an inflated index number. However, this may be a desirable quality in an index; the question of how to combine monthly counts into a single index number is returned to in Chapter 5.

The Underhill model for imputing missing counts assumes that the expected count is the product of a site factor, a month factor and a year factor for each species (Section 1.3.2). The month factors can be thought of as the mean count across all sites and years in the given month, after imputing missing values, divided by the mean count in the base month. The month factors provide a measure of the long-term pattern of arrival and departure of wintering waterbirds. Underhill (1989) suggested that “one wishes to base a winter index on the group of months in which the population may be considered stable” and that exploratory analysis of the month factors might be used to do this.

Kirby consistency interval method for choosing months

As outlined in Section 1.3.3, Kirby et al. (1995) developed a method for choosing which of the winter months to use when calculating the yearly index value for wildfowl species. Site consistency intervals around the month factors from the Underhill algorithm were used to choose which months out of September to March to include for each species. “Consistency intervals” are analogous to confidence intervals and calculated using a bootstrap method (further details follow in Section 4.3.4). Those months whose consistency intervals overlapped with the peak month were chosen. The Kirby consistency interval analysis in Kirby et al. (1995) was not repeated and the chosen months were used for published indices in *Wildfowl and Wader Counts* from 1993 until 2007, when concerns over phenological change and for increased simplicity prompted a move to using all months from September to March for all wildfowl species (Waters and Cranswick, 1993; Musgrove et al., 2007).

Consistency intervals were calculated around the month factors by selecting subsets of sites on which to calculate the month factors and using bootstrapping over sites to obtain approximate 90% confidence intervals. For each species a subset of months were selected comprising the month with the largest month factor plus any other months where the 90% consistency intervals for the month factor overlapped with those of the peak month. A technical difficulty with the method as detailed in Kirby et al. (1995) is that computing limitations at the time meant bootstrapping was carried out after the Underhill model has been fitted, which resulted in misleading consistency intervals (see Section 4.3.4).

Figure 3.14 compares the analysis in Kirby et al. (1995) which resulted in the months October to January being used for Pintail, with a similar analysis including more recent data, where only the months of December and January would have been chosen. The calculated month factors are very similar for both, illustrating the arbitrary aspect of the consistency interval method.

For some species only one month was chosen by Kirby et al. (1995) on which to base

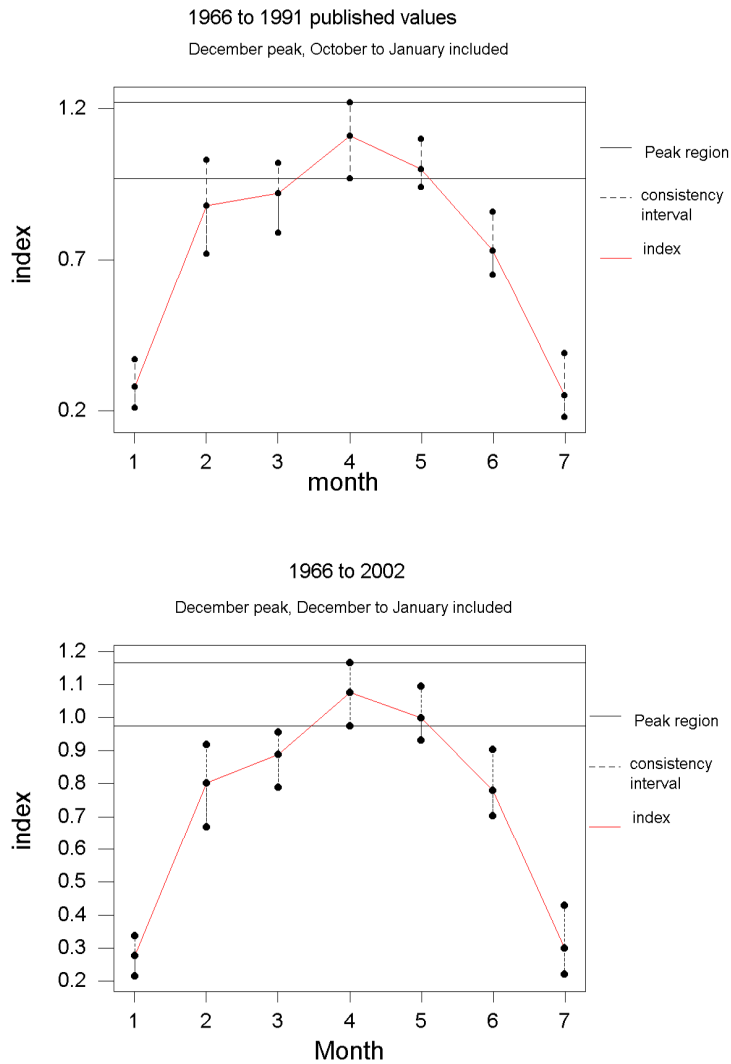


Figure 3.14: Kirby et al. (1995) method for choosing months for Pintail using overlapping site consistency intervals of the month indices (factors). The month index is set so month 5 (January) is equal to 1. Analysis as done in Kirby et al. (1995), using data from 1966/67 to 1991/92 (top) and repeated with data updated to 2002/03 (bottom).

the index. For example, in the case of Mallard, the largest month factor occurred in December. As none of the consistency intervals in the other months overlapped with the December 90% consistency intervals, only December was selected as the indexing month. Where only one month of data is used in the generation of indices, a low index value can occur if the peak count occurs in one of the other months, when arguably the index should reflect the broader picture of how “good” the winter was for that species.

For example, Mallard peak counts have occasionally been in January or November rather than December.

Given that we are interested in a national index, it is of little importance if there are regional differences in seasonality differs across sites and regions in Britain as the index is designed to reflect the national population. Of more concern is temporal variation such as phenological change; it would be more intuitive to use consistency intervals calculated by bootstrapping over years rather than sites. Bootstrapping over years is used in the ‘BTS’ method in Section 3.5.2 below.

In selecting which months to use in the index a balance must be struck between only using months where the population is stable and having enough months to deal with phenological changes. In Section 3.5.2 an exploratory functional data analysis approach is used to ascertain when the population function can be considered approximately stable for each species.

3.5.2 Selecting stable months: an FDA approach

An alternative method to that used by Kirby (1995) is a more exploratory approach of the seasonal dynamics of each species. The positive smoothing method described in Section 3.3.3 uses a roughness penalty approach which ensures that both the first and second derivatives of the population functions are smooth. These can then be used for selecting the months where the population is most stable for each species.

One method is to use months where the first derivative is arbitrarily close to zero (Frost, 2006; Kershaw and Frost, 2006). The mean recorded count was calculated for each month between September and March in each year across the set of sites chosen by the 50% completeness criterion (see Section 2.4.1) for the years 1966/67 to 2002/03. The mean counts over sites for each year were then standardised by dividing by the largest mean in that year to give monthly proportions; this ensures the phenology pattern for each year is equally weighted and not affected by population change from year to year. For each year




Figure 3.15: Values of the derivatives of the smoothed mean curves at the point mid-way between month pairs for each year from 1966/67 to 2002/03 (black crosses) and the mean over years (blue curve) for Mallard. The mean is between -0.1 and 0.1 midway between September and October through to December and January.

the month proportions were smoothed using a positive smoothing algorithm (see Section 3.3.3). The smooth curves were then differentiated to give population change curves and the overall mean of the population change curves is found.

The population is considered “stable” between a pair of months when the overall mean curve at the point halfway between the month pair is greater than -0.1 and less than 0.1; that is, close to zero (see Figure 3.15). If all the stable months are adjacent, they should all be used for indexing. If no month pairs are found to be stable then only the peak month, that with the highest overall mean, is used. If stable pairs are not adjacent, then the individual year curves should be inspected to decide what would be most appropriate: if the species’ phenology has changed over time then all the stable months should be included, as should all the intervening months; whereas if the species’ phenology has not changed over time, so that the break in stable months is due to regularly occurring stable periods interspersed with periods where the population changes substantially, then the stable months that include the peak month should be used.




Figure 3.16: Values of the derivatives of the smoothed mean curves at the point mid-way between month pairs for each year from 1966/67 to 2006/07 (green crosses) and the 95% CI of the mean over years (black curves) for Mallard. The 95% CI of the mean midway between the month pairs overlaps the -0.1 and 0.1 region between September and October through to December and January.

The results of the analysis from Frost (2006) are in Table 3.2, labelled 'M50'. An analysis was also carried out using a similar methodology to that detailed above, but instead of the sites used being selected by the 50% criterion, the top sites criterion from Section 2.4.3 was used; and instead of using the sites where the mean derivative curve was between -0.1 and 0.1, a 95% confidence interval of the mean derivative was used (bootstrapping over years). Data from 2002/03 to 2006/07 was also included and data from August and April was used so that the derivatives for September and March could be better estimated (Figure 3.16). The results of this analysis are in Table 3.2, labelled 'BTS', and plots of the 95% CIs for each species, can be found in Appendix B.

As there are so many wildfowl species of interest, with varying ecologies, blindly applying the 'M50' or 'BTS' method is inadvisable; it is better to look at the curves for each

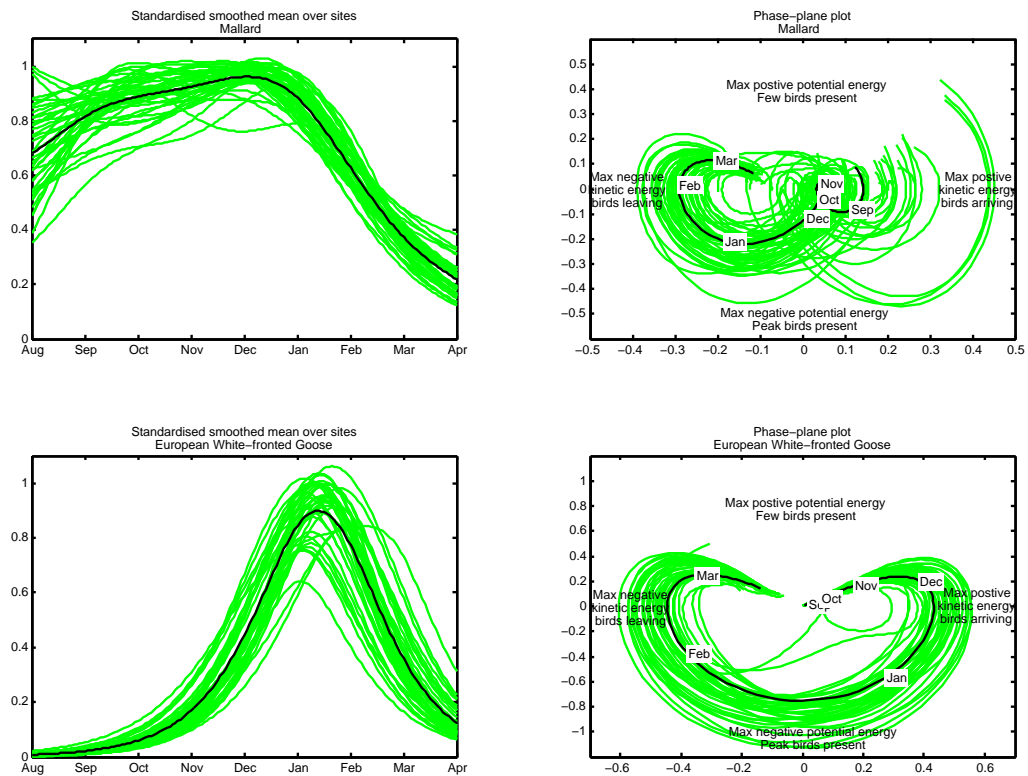


Figure 3.17: Smoothed seasonality plots (left) and phase-plane plots (right) for Mallard and European White-fronted Goose. A phase-plane plot is a plot of acceleration (y -axis; the second derivative of the smooth curves in the seasonality plot) against velocity (x -axis; the first derivative). Each individual year is shown as a green curve and the overall mean as a black curve.

species individually for any unique features. As well as inspecting the smooth seasonality curves v themselves (Figure 3.17; left-hand plots), the first derivative v' and second derivative v'' are also useful. A peak in v has $v' \approx 0$ and $v'' < 0$, whereas a minimum has $v'' > 0$. When the population is changing most rapidly, $|v'|$ is large and $v'' \approx 0$. By plotting v'' against v' in a phase-plane plot and labelling time on the curve, as in the right-hand plots in Figure 3.17, different aspects of the dynamics are represented in different parts of the graph. Periods where birds are arriving in Britain are to the east of the graph (centre right); when they are leaving are to the west (centre left); if there is a definite minimum in the population over the time period it is to the north (top centre) and a population peak is to the south (bottom centre). Flat, stable parts of the population curve have both $v' \approx 0$ and $v'' \approx 0$ so are close to the origin in the phase-plane plot (but may correspond

to periods of high or low abundance). When choosing “stable” months for indexing, here we mean that there is a population present (if the population is staying steady at close to 0 it is stable but not of great interest) and the numbers aren’t changing much, so we are interested in months where v' is close to zero and $v'' \leq 0$. Furthermore, we shall insist that the population peak month(s) are always included. When inspecting the phase-plane plot this means the two months where v' changes from positive to negative whilst v'' is negative are always chosen.

Although we could look directly at the left-hand plots in Figure 3.17 and see the months where these conditions are true (September to January for Mallard and January to February for European White-fronted Goose) the added convenience of the phase-plane plots is that it is not necessary to read the months off the axis, since the months are labelled on the plot itself. In the case of Mallard, a species with a mixture of resident and immigrant birds, the change in v' from positive to negative happens between December and January. December is part of a cluster of months in the stable area close to the origin, where the population dynamics are not changing much, prompting us to use the set of months September to January for this species. The seasonal dynamics of European White-fronted Goose are quite different, reflecting their ecology as an immigrant species. There is a clear pattern in the phase-plot, with birds arriving in December-January and leaving in February-March. The plot is less “messy” than that of the Mallard, indicating less year to year variation and thus less of a need for conservative choice of months, since there is more certainty of when the peak in the population will occur. The point at which $v' = 0$ is between January and February and so these two months are used, with no adjacent stable months. There is not much change in the dynamics in September-November but as there are few birds present these are not used. Phase plane plots of each wildfowl type can be found in Appendix B.

The positive smoothing method used in this analysis tends to push the smooth function away from zero, so caution needs to be applied to derivatives in regions where there are actually very few birds present. In addition, when subjectively selecting months the

variation across years should be taken into account so that if a species is undergoing phenological change, the months used incorporate all months where the peak may occur in different years. For example, Kirby et al. (1995) chose the months September to October for Shoveler, based on analysis of data from 1966/67 to 1991/92; however it was shown in Section 3.4 that there has been a change in seasonal patterns in Shoveler in the years 1992/93 to 2006/07 to that for 1966/67 to 1991/92 and hence it is now recommended that all months from September to March are used for this species.

Recommended months to use for general analysis for each species are given in Table 3.2 (labelled 'Rec.')

based on the results of the bootstrapping and mean derivative methods (BTS and M50) and visual inspection of the plots in Appendix B. The original month selections for species where the analysis was published in Kirby et al. (1995) are also given in Table 3.2 for comparison. However, it is also recommended that months selection for each species is repeated regularly, so that any phenological changes can be examined and if necessary the months used can be changed accordingly.

	First month				Last month			
	Rec.	BTS	M50	Kirby	Rec.	BTS	M50	Kirby
Mute Swan	Sep	Sep	Sep	Sep	Mar	Mar	Mar	Mar
Whooper Swan	Nov	Nov	Nov	Nov	Feb	Feb	Mar	Dec
Bewicks Swan	Jan	Jan	Jan	Jan	Feb	Feb	Feb	Feb
Pink-footed Goose	Oct	Jan	-	-	Mar	Mar	-	-
European White-fronted Goose	Jan	Jan	-	-	Feb	Feb	-	-
Greenland White-fronted Goose	Nov	Nov	-	-	Mar	Mar	-	-
Icelandic Greylag Goose	Nov	Nov	-	-	Mar	Mar	-	-
Re-established Greylag Goose	Sep	Sep	-	-	Jan	Jan	-	-
Canada Goose	Sep	Sep	Oct	Sep	Jan	Jan	Jan	Sep
Svalbard Barnacle Goose	Oct	Oct	-	-	Mar	Mar	-	-
Naturalised Barnacle Goose	Sep	Oct	-	-	Feb	Feb	-	-
Dark-bellied Brent Goose	Dec	Dec	Dec	Dec	Feb	Feb	Feb	Feb
Shelduck	Dec	Dec	Jan	Jan	Feb	Feb	Feb	Feb
Wigeon	Dec	Dec	Dec	Jan	Jan	Jan	Jan	Jan
Gadwall	Sep	Sep	Sep	Sep	Mar	Mar	Jan	Mar
Teal	Nov	Nov	Dec	Dec	Jan	Jan	Jan	Dec
Mallard	Sep	Sep	Sep	Dec	Jan	Jan	Jan	Dec
Pintail	Nov	Nov	Dec	Oct	Jan	Jan	Jan	Jan
Shoveler	Sep	Sep	Sep	Sep	Mar	Mar	Mar	Oct
Pochard	Dec	Dec	Dec	Nov	Jan	Jan	Jan	Jan
Tufted Duck	Sep	Sep	Sep	Nov	Mar	Feb	Feb	Feb
Scaup	Dec	Dec	-	-	Feb	Feb	-	-
Goldeneye	Jan	Jan	Jan	Feb	Feb	Feb	Feb	Feb
Red-breasted Merganser	Oct	Oct	Sep	Oct	Mar	Feb	Mar	Mar
Goosander	Dec	Dec	Jan	Dec	Feb	Feb	Feb	Feb
Ruddy Duck	Oct	Nov	-	-	Feb	Feb	-	-

Table 3.2: Months of stable winter populations, chosen using the Kirby et al. (1995) month consistency method (**Kirby**); the mean derivative method with sites that fulfill the **50%** criterion (**M50**); the bootstrapping derivative method with sites that fulfill the **top sites** criterion (**BTS**); and **recommended** months from a subjective judgement of the BTS and other information shown in Appendix B (**Rec.**). Further details of each method may be found in the text.

3.6 Discussion

Robustness of results to choice of λ




Figure 3.18: Effect of changing λ in a Functional Principal Components Analysis of phenological change in Shoveler. A multiple of each of the first Principal Component harmonics is added (+++) and subtracted (---) from the mean function (blue), where the smoothing parameter λ used to fit the year functions is set to 0.1 (top); 0.01 (middle); and 0.001 (bottom).

As discussed in Section 3.3.3, the choice of smoothing parameter λ controls the trade-off between the smoothness of the fitted functions and their adherence to the observed monthly standardised counts. The results of the analyses in Sections 3.4 and 3.5 consequently depend in part on what value of λ has been chosen.

In the principal components analysis of shoveler data in Section 3.4, the first principal

component captured 71% of the variation in the curves and the scores of each individual year for this component exhibited a trend over time, reflecting phenological change in the species, when a value of $\lambda = 0.01$ was used. Figure 3.18 shows how the mean curve and first principal component changes if the magnitude of λ is increased or decreased by a factor of ten. Increasing the smoothing parameter results in smoother curves whereas decreasing it allows more roughness in the curves. Increasing λ thus decreases the potential amount of variability between curves, allowing the first principal component to more easily capture a greater proportion of the variability; when λ is increased to 0.1 it captures 79% of the variation, whereas decreasing it to 0.001 results in the variation captured dropping to 64%. The scores of the individual years on the first principal component remain very similar whichever of the three values of λ are used (Figure 3.19). The conclusion that shoveler phenology has changed over time is thus not affected by which of these three values of λ is used.

Table 3.2 in Section 3.5.2 recommends subsets of months to use for different species, based on their seasonal dynamics by inspecting seasonality curves and their phase-plane




Figure 3.19: First principal component score for each year in Shoveler functional principal components analysis with $\lambda = 0.1$ (green); 0.01 (blue); and 0.001 (red).

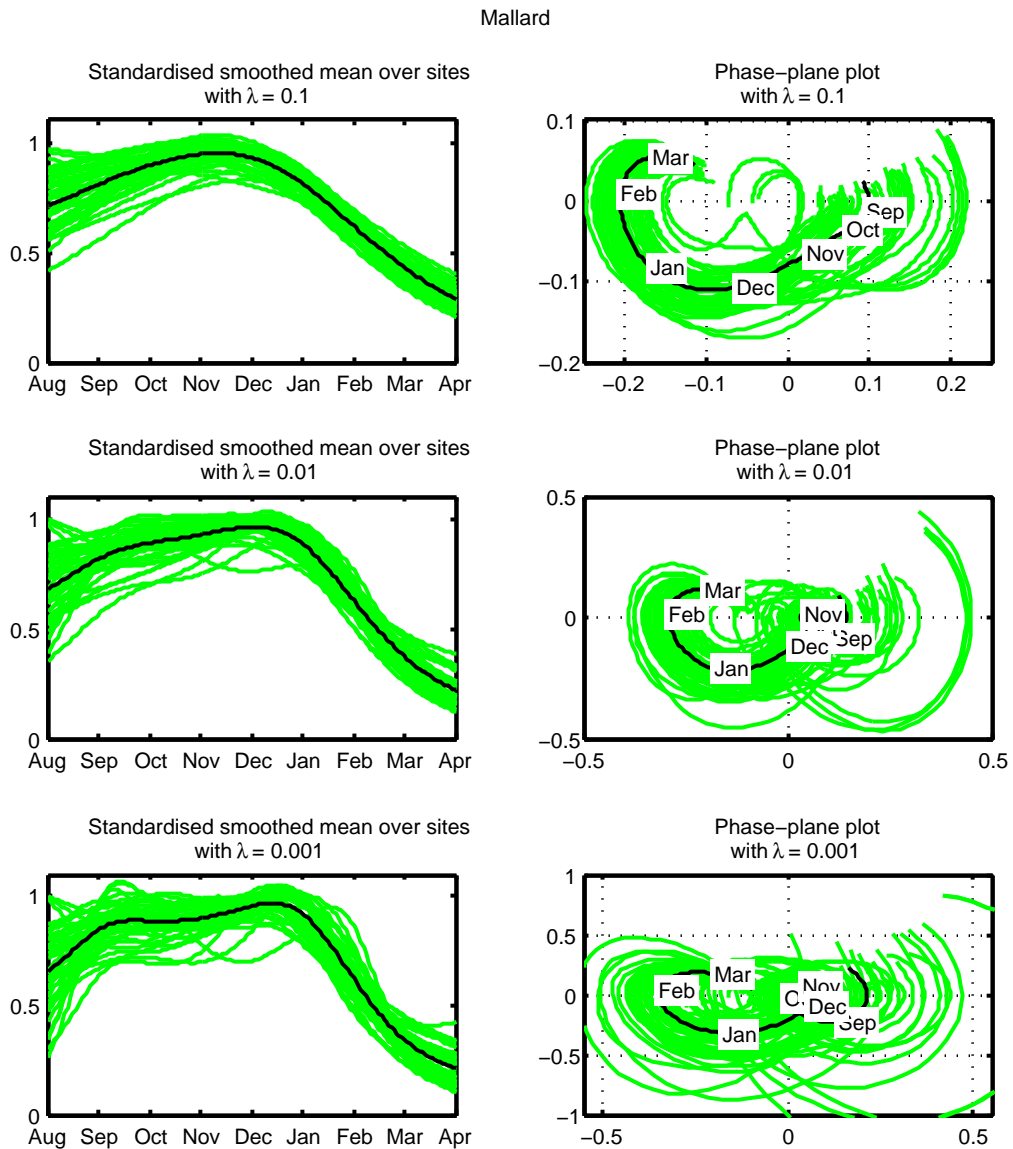


Figure 3.20: Smoothed seasonality plots (left) and phase-plane plots (right) for Mallard. Each individual year is shown as a green curve and the overall mean as a black curve, where the smoothing parameter λ used to fit the year functions is set to 0.1 (top); 0.01 (middle); and 0.001 (bottom).

plots such as those in Figure 3.17 for Mallard and European White-fronted Goose. The value of $\lambda = 0.01$ that was used in Figure 3.17 is compared with $\lambda = 0.1$ and 0.001 in Figures 3.20 and 3.21. The difference in smoothness does not much affect the apparent dynamics of either species, and the same months would be chosen (September to January for Mallard; January to February for European White-fronted Goose). However, there

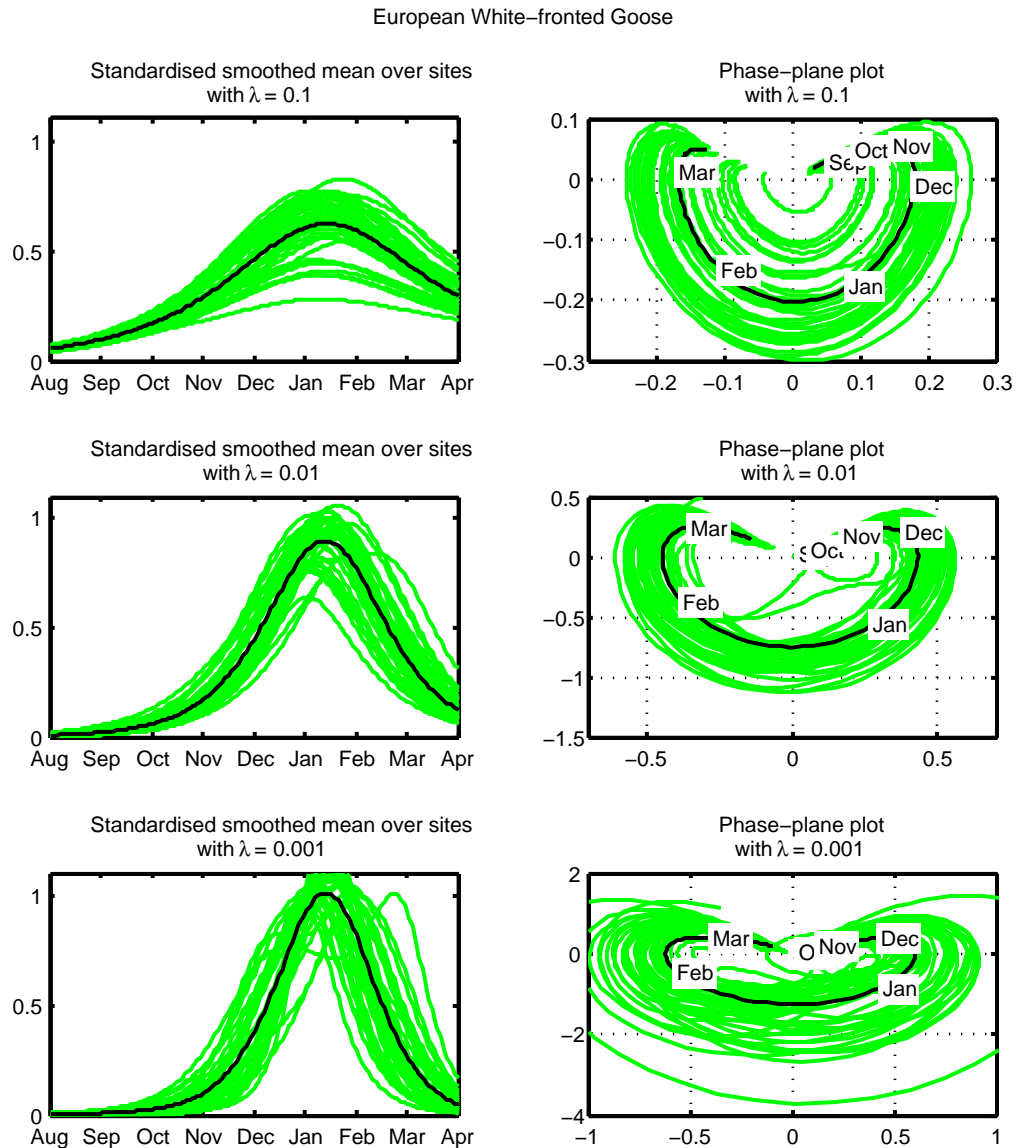


Figure 3.21: Smoothed seasonality plots (left) and phase-plane plots (right) for European White-fronted Goose. Each individual year is shown as a green curve and the overall mean as a black curve, where the smoothing parameter λ used to fit the year functions is set to 0.1 (top); 0.01 (middle); and 0.001 (bottom).

are changes in detail, particularly when $\lambda = 0.1$ is used. In Figure 3.20 the peak in the mean function for Mallard is between December and January for $\lambda = 0.01$ and 0.001, whereas it is between November and December for $\lambda = 0.1$. It looks like the underlying data may have been smoothed too much with $\lambda = 0.1$, with the fitted functions not giving a representative picture of the true complexity of the seasonal dynamics. A similar issue

can be seen in Figure 3.21 with the peak in the smoothed mean functions of European White-fronted Goose all around 0.5 rather than 1 (recall that the underlying data was standardised so the peak monthly count in each year was set to 1). Smoothing with $\lambda = 0.1$ has not allowed the sharp peak in the numbers of European White-Fronted Goose to be accurately represented.

Seasonality in wildfowl populations

This chapter has focused on one of the special characteristics of Wetland Bird Survey data: sequences of counts which exhibit seasonal variation as well as trends over a period of years. Although there are many ecological monitoring schemes, only a few attempt to collect within-year population dynamics data as well as information on population change from year to year in the way WeBS does.

In Section 3.2 the confounding nature of intra-winter bird movements were considered. As there are no replicates of counts for any particular monthly site count and most wildfowl are unconstrained to territories in winter, moving from site to site and country to county for feeding and roosting or in response to weather such as prolonged periods of wet weather or a long cold winter, it is difficult to disentangle true movements with the sources of measurement error discussed in Section 2.2. Consequently, imputing missing values is a challenging proposition.

A major advantage of collecting data over the whole winter period is that it allows analysis of changes in seasonality over time. Section 3.3 introduced the set of statistical techniques known as functional data analysis (FDA), and showed that each monthly WeBS count could be thought of as a snapshot of a long population time series on which FDA techniques can be used. In Section 3.4 year functions were fitted by smoothing the monthly counts of Shoveler. Using functional Principal Components Analysis enabled us to show a trend over time from a pattern of high Shoveler counts in the autumn decreasing over the winter, to a seasonal curve where the population remained at similar levels

from September through to March.

For some further analyses, it is useful to be able to think of the monthly counts as replicates of an estimate of the year population level. Section 3.5 used the derivatives made available by taking a functional data approach to choose a subset of “stable” months. If an analysis requires imputing of missing counts, it is beneficial to work with a subset of months, particularly if these are concentrated in the mid-winter as there are more missing counts at the season ends due to a decrease in observer effort. The previous analysis used to choose the subset of months used for calculating annual indices for each species used variability between sites to choose the months. The functional data approach in this section by contrast chooses the months by the form of the fitted population functions themselves, whilst considering variability between years to minimise the effects of phenological change.

We will draw on the greater understanding of the seasonal effects in wildfowl monitoring data we gain in this chapter in subsequent chapters. In Chapter 4 the stable months from this chapter will be used for generating annual indices. One of the limitations of the Underhill index discussed in the following chapter is that it does not allow for phenological changes in wildfowl populations. In Chapter 5 alternative methods of compiling a population index to reflect the seasonal effects discussed in this chapter will be explored and what we have learned about the form of wildfowl population functions will allow us to simulate WeBS data.

Chapter 4

QUANTIFYING POPULATION CHANGE WITH INCOMPLETE DATA

4.1 Introduction

Assessment of trends in the number of non-breeding waterbirds in the UK is one of the main objectives of the Wetland Bird Survey. As outlined in section 1.3, since the mid-1990s changes in population size have been tracked by the Underhill index. The Underhill indexing method can be broken down into two concepts:

1. Impute missing values using an EM-type algorithm based on site, month and year information to obtain a complete data set.
2. Obtain a single yearly index number from the complete data set by finding the ratio of the total number of birds counted over all months and all sites in each year compared to a base year.

In this chapter the underlying Underhill model for imputing missing data to obtain a complete data set is assessed. Alternative methods for formulating an index number from a complete data set are investigated in Chapter 5.

4.2 Underhill's Site-Year-Month Multiplicative Model

4.2.1 Introduction

For a species data set of I sites over K months and J years, the Underhill multiplicative model was developed, as introduced in Section 1.3.2. The model assumes that the number of birds at a site at a particular time depends on the site itself, the population that year and the time of year. It also assumes that the site and month effects are constant over time and that the year and month effects are the same for all sites. The expected value of the count of the number of birds at site i in year j and month k can then be expressed as the product of a site factor, a year factor and month factor:

$$\mathbb{E}[x_{ijk}] = s_i y_j m_k \quad (4.1)$$

where, $i = 1, 2, \dots, I$; $j = 1, 2, \dots, J$; and $k = 1, 2, \dots, K$.

Underhill and Prŷs-Jones (1994) assume that since the x_{ijk} s are counts they are likely to have a Poisson-like distribution. The probability mass function for the Poisson distribution is,

$$\mathbb{P}(X_{ijk} = x_{ijk}) = \frac{e^{-s_i y_j m_k} (s_i y_j m_k)^{x_{ijk}}}{x_{ijk}!}; \quad (4.2)$$

and the Poisson log-likelihood function is:

$$l = \sum_{ijk} [x_{ijk} \ln(s_i y_j m_k) - s_i y_j m_k - \ln(x_{ijk}!)]. \quad (4.3)$$

However, because of the clumped spatial distribution of waterbirds, it is likely that the counts will be over-dispersed (the variance of the counts will be much larger than the mean). The dispersion factor is assumed to be a constant multiple of the mean, $D > 1$, so that,

$$\text{Var}(x_{ijk}) = D \mathbb{E}(x_{ijk}) = D s_i y_j m_k. \quad (4.4)$$

Quasi-likelihood methods are used when a relationship between the mean and variance, such as that in Equation 4.4, is assumed, but no assumption is made about the distribution of the data. Suppose that there are no missing or incomplete counts. The estimates of s_i , y_j and m_k are those that minimise the deviance function,

$$A = \sum_{ijk} \frac{(x_{ijk} - s_i y_j m_k)^2}{D s_i y_j m_k}. \quad (4.5)$$

Underhill and Prŷs-Jones (1994) note the close analogy between Equation 4.5 and the Pearson chi-squared statistic. The estimates of s_i , y_j and m_k found by minimising the equation 4.5 are the same as the maximum likelihood estimates from the Poisson log-likelihood in equation 4.3 (McCullagh and Nelder, 1989). For this reason the counts are said to have a quasi-Poisson distribution.

Two constraints on the parameters s_i , y_j and m_k are needed to determine the estimates uniquely. A base year and a base month are chosen so that $y_b = 1$ and $m_b = 1$. The estimates are then,

$$\begin{aligned} \hat{s}_i &= \frac{\sum_{j,k} x_{ijk} \cdot \sum_{i,k} x_{ibk} \cdot \sum_{i,j} x_{ijb}}{\sum_{i,j,k} x_{ijk}^2} \\ \hat{y}_j &= \frac{\sum_{i,k} x_{ijk}}{\sum_{i,k} x_{ibk}} \\ \hat{m}_k &= \frac{\sum_{i,j} x_{ijk}}{\sum_{i,j} x_{ijb}}. \end{aligned} \quad (4.6)$$

Underhill defines the population index to be the total population over sites and months relative to that in the base year,

$$I_j = \frac{\sum_{i,k} x_{ijk}}{\sum_{i,k} x_{ibk}}. \quad (4.7)$$

The year factors y_j are thus the required index values I_j .

4.2.2 Underhill's Algorithm for fitting the site, year and month factors

The Underhill algorithm (Figure 4.1) which is used to fit the model in equation 4.1 is similar to the well-known expectation-maximisation (EM) algorithm (Dempster et al., 1977). An additional step in the algorithm is the treatment of poor quality, incomplete counts, where the greater of the expected count and actual incomplete count are used.

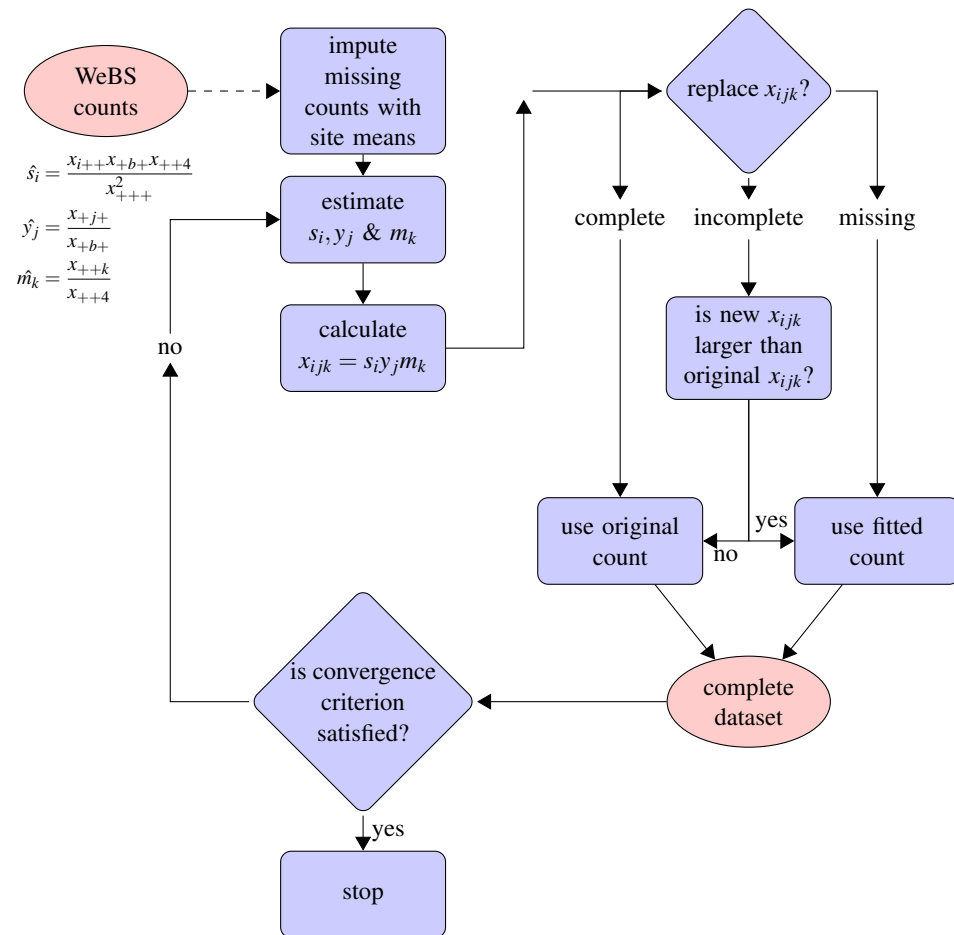


Figure 4.1: Flow diagram of Underhill's Imputing Algorithm for WeBS Counts. "+" denotes summation over that subscript

The algorithm is as follows:

1. Find the mean of the good quality counts for each site. These are used as starting values to allow the algorithm to quickly converge.
2. Estimate the site, year and month parameters by equation 4.6.
3. Calculate the fitted counts by equation 4.1.
4. Replace each count flagged as missing with the appropriate fitted count.
5. Replace the counts flagged as poor quality with the larger of the actual count and the fitted count.
6. If the convergence criterion has not been reached, return to step 2.

Programs that fit the model using this algorithm include the UINDEX computer program for Windows (Bell, 1995) and Wing, a quicker program written in Fortran. The algorithm can easily be programmed in other languages: currently a SAS implementation is used when generating index plots for annual reports. In this project, a MATLAB implementation of Underhill's algorithm was used; the main code can be seen in Program 4.1 (Section 1.3.4 introduced MATLAB).

Table 4.1 shows the values imputed by the Underhill algorithm for the small example set of Mute Swan counts from Section 1.3.1. If no distinction is made between the quality of counts, so that incomplete counts are treated as complete, only the five missing counts at Loch Bee are imputed (Table 4.1(B)). The Underhill algorithm compares the fitted count and the observed count for each of the eight incomplete counts; in this example the fitted count is the greater for all eight so the observed counts are replaced.

(A) Original Counts					
	1	2	3	4	5
1998	936	<i>734</i>	432		382
1999	976	<i>671</i>	283		439
2000	1169	<i>1011</i>	439	341	332
2001	807	1110	330	229	597
2002	1075	941	555		511
2003	1368	<i>865</i>	189		672
2004	871	<i>801</i>	261	265	414
2005	772	883	533	300	467
2006	888	1024	422		221
2007	920	1164	135	379	205

(B) Underhill Modelled Counts ignoring incompleteness					
	1	2	3	4	5
1998	936	734	432	284.5	382
1999	976	671	283	271.3	439
2000	1169	1011	439	341	332
2001	807	1110	330	229	597
2002	1075	941	555	352.9	511
2003	1368	865	189	354.3	672
2004	871	801	261	265	414
2005	772	883	533	300	467
2006	888	1024	422	292.6	221
2007	920	1164	135	379	205

(C) Underhill Modelled Counts					
	1	2	3	4	5
1998	1214.3	1332.0	432	387.4	557.3
1999	976	1026.3	283	298.5	439
2000	1169	1276.5	439	341	534.1
2001	807	1110	330	229	597
2002	1075	941	555	337.7	511
2003	1368	1347.3	189	391.8	672
2004	871	931.0	261	265	414
2005	772	883	533	300	467
2006	888	1024	422	279.9	221
2007	920	1164	135	379	205

Table 4.1: Counts imputed by the Underhill Algorithm. January Mute Swan counts for the sites 'Fleet and Wey'(1), 'Somerset Levels'(2), 'Ouse Washes'(3), 'Loch Bee (South Uist) '(4) and 'Loch of Harray'(5) for 1998 to 2007. In the original data in Table A, counts at Loch Bee are missing for 1998, 1999, 2002, 2003 and 2006 and counts marked as incomplete (poor quality) are in *italics*. If the WeBS counts are all treated as 'complete' then the Underhill algorithm imputes missing data only, resulting in Table B. Underhill's method of using the larger of the fitted and original count for incomplete counts, the standard procedure for WeBS data, gives Table C. (See Table 4.2 on page 144 for Underhill Index results for this example.)

Program 4.1 MATLAB implementation of Underhill's algorithm.
 (Program continues on the following page.)

```

function [yearfactor, monthfactor, sitefactor,...
  modelledcounts, fits, delta] = Underhillfn(data,beta)

%Underhill algorithm function. Input: matrix generated by preparedata and a
%structure of parameters. Returns year, month and site factors and, if
%required, the counts with missing/incomplete counts imputed in
%modelledcounts, and the fitted values. This function is usually called
%from Underhill (when bootstrap CIs are not required) or UnderhillCI.

site=data(:,1); year=data(:,2); month=data(:,3); count=data(:,4);
year1 = year-min(year)+1; baseyear1 = beta.baseyear-min(year)+1;
missingrows = data(:,5)==3;
incompleterows = data(:,5)==2;
completerows = data(:,5)==1;

%Initialise some variables.
criterion = zeros(beta.niter,1);
yearfactor = zeros(beta.nyears,1); monthfactor = zeros(beta.nmonths,1);
sitefactor = zeros(beta.nsites,1); fits = zeros(beta.n,1);

%Underhill algorithm.

%STEP 1 Set modelledcounts to be counts/site means.
%This should reduce the time it takes for the algorithm to converge.
sitemean=accumarray(data(~missingrows,1),data(~missingrows,4), [], @mean);
modelledcounts=count;
modelledcounts(missingrows)=sitemean(site(missingrows));

%Set up a test for convergence. The criterion vector holds the convergence
%criterion values for each iteration - this is the average absolute
%difference in fitted values between iterations. These values are held in
%the two vectors convergence1 for the previous iteration and convergence2
%for the current iteration.

convergence2 = modelledcounts - count;
criterion(1) = mean(abs(convergence2));
%Iterate STEPS 2 - 6
for iterstep = 1:beta.niter
  convergence1 = convergence2;
  %STEP 2 Estimate the values of the site, year and month factors.
  %(accumarray(month,modelledcounts,[],@sum) etc. is too slow)
  sortedy=modelledcounts(beta.sortyindex);
  for i = 1:beta.nyears
    yearfactor(i)=sum(sortedy(beta.nmonths*beta.nsites*(i-1)...
      +1:beta.nsites*beta.nmonths*i));
  end
  yearbasesum = yearfactor(baseyear1);
  yearfactor = yearfactor./yearbasesum;

  sortedm = modelledcounts(beta.sortmindex);
  for i = 1:beta.nmonths
    monthfactor(i) =sum(sortedm(beta.nsites*beta.nyears*(i-1)...
      +1:beta.nsites*beta.nyears*i));
  end
  monthbasesum=monthfactor(beta.basemonth);
  monthfactor = monthfactor./monthbasesum;

  for i = 1:beta.nsites
    sitefactor(i)=sum(modelledcounts(beta.nmonths*beta.nyears*(i-1)...
      +1:beta.nmonths*beta.nyears*i));
  end
  sitefactor = (sitefactor.*monthbasesum.*yearbasesum./...
    ((sum(modelledcounts)*sum(modelledcounts))));

```

Program 4.1 (continued)

```
%STEP 3 Calculate the fitted counts
fits = sitefactor(site).*yearfactor(year1).*monthfactor(month);
%STEP 4 Impute fit for each missing count
modelledcounts(missingrows)=fits(missingrows);
%STEP 5 Impute the larger of the fit & observed for incomplete counts
modelledcounts(incompleterows)=...
    max(fits(incompleterows),count(incompleterows));
%STEP 6 Has the convergence criterion been satisfied?
convergence2 = modelledcounts - count;
criterion(iterstep+1) = mean(abs(convergence2 - convergence1));
if criterion(iterstep+1) <= beta.convcrit
    break
end
end
if nargout==6, getdelta, end %Calculate model statistics
```



Figure 4.2: Mallard (*Anas platyrhynchos*). ©Nick Derry, used with kind permission.

4.2.3 Generalised Linear Models

The multiplicative model in equation 4.1 is equivalent to a Generalised Linear Model with a quasi-Poisson error distribution and log link function, since the model may be written as:

$$\mathbb{E}[\log(x_{ijk})] = \log s_i + \log y_j + \log m_k = S_i + Y_j + M_k. \quad (4.8)$$

Underhill and Prŷs-Jones (1994) expressed concern that using the model formulation in equation 4.8 with estimates for the year factors of the type,

$$Y_j = \frac{1}{IK} \sum_{i=1}^I \sum_{k=1}^K \log(x_{ijk}) \quad (4.9)$$

resulted in indices based on the geometric mean over sites and months rather than the arithmetic mean (or equivalently, total) since

$$\frac{e^{Y_j}}{e^{\bar{Y}_b}} = \left(\prod_{i=1}^I \prod_{k=1}^K \log(x_{ijk}) \right)^{1/IK} \bigg/ \left(\prod_{i=1}^I \prod_{k=1}^K \log(x_{ibk}) \right)^{1/IK}. \quad (4.10)$$

Geometric means over sites do not give satisfactory index numbers for national populations of wildfowl, since the resultant index numbers do not track annual changes in the total population, which is distributed unevenly over sites. The implications of geometric means as an alternative way of combining monthly counts is returned to in Section 5.3.2.

However, fitting the GLM in equation 4.8 by maximum likelihood does produce estimates for S_i , Y_j and M_k which can be used to obtain s_i , y_j and m_k by taking the exponential. The statistical computing package R contains the package 'glm' to fit generalised linear models which can fit both the Poisson model and quasi-Poisson model (R Development Core Team, 2009). Code for fitting GLMs in R to WeBS data is given in Program 4.2. The model can also be extended to have site and year interactions (individual site trends),

$$\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + S_i \times Y_j, \quad (4.11)$$

Program 4.2 R code for fitting generalised linear models to the WeBS data matrix “new-data” created by Program 1.1 in MATLAB , where missing counts have been recoded from -1 in the count column to ‘NaN’ and exported as a text file.

```
x<-read.table(file="Path/Name.dat",sep=",",na.strings="NaN")
colnames(x) <- c("st", "yr", "mth", "ct", "qty")
basey=max(x$yr); basem=max(x$mth);
#Tell R that the site, year and month columns are category vectors
x$st<-as.factor(x$st); x$yr<-as.factor(x$yr); x$mth<-as.factor(x$mth)
#Set the base year and month
x$yr<-relevel(x$yr, as.character(basey))
x$mth<-relevel(x$mth, as.character(basem))
#Fit GLMs and display results
xqGLM<-glm(x$ct~x$st+x$yr+x$mth-1, family=quasipoisson(link=log))
xGLM<-glm(x$ct~x$st+x$yr+x$mth-1, family=poisson(link=log))
phenGLM<-glm(x$ct~x$st*x$yr*x$mth-1, family=quasipoisson(link=log))
sitetrendGLM<-glm(x$ct~x$st*x$yr+x$mth-1, family=quasipoisson(link=log))
summary(xqGLM); summary(xGLM); summary(phenGLM); summary(sitetrendGLM)
```

site and month interactions (site seasonality),

$$\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + S_i \times M_k, \quad (4.12)$$

and month and year interactions (phenology effects)

$$\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + Y_j \times M_k. \quad (4.13)$$

Note, however, that adding interaction terms involving sites increases the number of parameters to be estimated by a large amount. For example, the number of parameters in the model without interaction terms (equation 4.8) has 239 parameters for the Wigeon data in Table 4.3. Adding phenology effects (equation 4.13) increases the parameter count to 279; site seasonality (equation 4.12) has 436 parameters; and individual site trends (equation 4.11) increases the number of parameters to 8119. With 4034 counts out of 16236 missing there will be sites where there is not any data on which to estimate some of the site \times year interaction terms. The lack of interaction terms in the Underhill model will be considered further in Sections 4.3.2 and 4.3.3.

Tables 4.2 and 4.3 compare the results from using the R glm package to the Underhill method. If the quality flags for the counts are ignored, the point estimates for the year

Year	Underhill		xqGLM		xGLM	
	$U1y_j$	$U2y_j$	Y_j	y_j	Y_j	y_j
1997	0.99	1.40	-0.012 (0.178)	0.99	-0.012 (0.028)	0.99
1998	0.94	1.08	-0.06 (0.18)	0.94	-0.06 (0.028)	0.94
1999	1.17	1.34	0.161 (0.165)	1.17	0.161 (0.026)	1.17
2000	1.10	1.10	0.092 (0.168)	1.10	0.092 (0.026)	1.10
2001	1.23	1.22	0.203 (0.169)	1.23	0.203 (0.026)	1.23
2002	1.23	1.42	0.207 (0.168)	1.23	0.207 (0.026)	1.23
2003	0.93	0.98	-0.071 (0.175)	0.93	-0.071 (0.027)	0.93
2004	1.05	1.05	0.053 (0.169)	1.05	0.053 (0.026)	1.05
2005	1.02	1.01	0.016 (0.177)	1.02	0.016 (0.027)	1.02
2006	1.00	1.00	0.00	1.00	0.00	1.00

Table 4.2: Underhill and GLM estimates. Year factor estimates for January Mute Swan counts for five sites for 1998 to 2007. Eight of the original observations were flagged as incomplete and five were missing. The Underhill algorithm results in the year factors $U2y_j$; or $U1y_j$ if completeness is ignored (i.e. incomplete counts are assumed complete). The glm routine in R assuming a Poisson distribution (xGLM) and quasi-Poisson distribution (xqGLM) give year factor estimates Y_j (standard error in parentheses) which are transformed to give $y_j = e^{Y_j}$. (See Table 4.1 for the observed counts.)

factors are the same whether the Underhill algorithm is used or the glm routine in R with a Poisson or quasi-Poisson distribution assumed for the counts. The site and month factor estimates are also unchanged and thus so are the predicted values for the missing counts. It is the treatment of incomplete counts which is the important feature of the Underhill algorithm: 15% of the Wigeon counts in Table 4.3 are incomplete, resulting in year factors which are particularly higher in the 1960s than if completeness is not taken into account.

Year	Underhill		xqGLM		xGLM	
	$U1y_j$	$U2y_j$	Y_j	y_j	Y_j	y_j
1966-67	0.44	0.62	-0.83 (0.083)	0.44	-0.83 (0.003)	0.44
1967-68	0.49	0.77	-0.71 (0.082)	0.49	-0.71 (0.003)	0.49
1968-69	0.50	0.71	-0.702 (0.079)	0.50	-0.702 (0.002)	0.50
1969-70	0.53	0.79	-0.641 (0.077)	0.53	-0.641 (0.002)	0.53
1970-71	0.52	0.80	-0.653 (0.077)	0.52	-0.653 (0.002)	0.52
1971-72	0.50	0.70	-0.687 (0.078)	0.50	-0.687 (0.002)	0.50
1972-73	0.54	0.81	-0.621 (0.077)	0.54	-0.621 (0.002)	0.54
1973-74	0.60	0.84	-0.512 (0.074)	0.60	-0.512 (0.002)	0.60
1974-75	0.46	0.51	-0.775 (0.081)	0.46	-0.775 (0.003)	0.46
1975-76	0.44	0.47	-0.82 (0.082)	0.44	-0.82 (0.003)	0.44
1976-77	0.55	0.58	-0.599 (0.075)	0.55	-0.599 (0.002)	0.55
1977-78	0.44	0.66	-0.818 (0.081)	0.44	-0.818 (0.003)	0.44
1978-79	0.57	0.63	-0.554 (0.073)	0.57	-0.554 (0.002)	0.57
1979-80	0.52	0.56	-0.655 (0.075)	0.52	-0.655 (0.002)	0.52
1980-81	0.60	0.62	-0.507 (0.072)	0.60	-0.507 (0.002)	0.60
1981-82	0.72	0.71	-0.323 (0.068)	0.72	-0.323 (0.002)	0.72
1982-83	0.58	0.58	-0.553 (0.072)	0.58	-0.553 (0.002)	0.58
1983-84	0.43	0.48	-0.835 (0.079)	0.43	-0.835 (0.002)	0.43
1984-85	0.80	0.85	-0.227 (0.066)	0.80	-0.227 (0.002)	0.80
1985-86	0.82	0.83	-0.204 (0.065)	0.82	-0.204 (0.002)	0.82
1986-87	0.78	0.79	-0.254 (0.066)	0.78	-0.254 (0.002)	0.78
1987-88	0.78	0.77	-0.248 (0.066)	0.78	-0.248 (0.002)	0.78
1988-89	0.76	0.77	-0.271 (0.067)	0.76	-0.271 (0.002)	0.76
1989-90	0.81	0.76	-0.217 (0.065)	0.81	-0.217 (0.002)	0.81
1990-91	0.82	0.78	-0.196 (0.064)	0.82	-0.196 (0.002)	0.82
1991-92	1.02	0.95	0.024 (0.061)	1.02	0.024 (0.002)	1.02
1992-93	0.94	0.88	-0.065 (0.062)	0.94	-0.065 (0.002)	0.94
1993-94	1.08	1.03	0.08 (0.06)	1.08	0.08 (0.002)	1.08
1994-95	1.17	1.13	0.157 (0.059)	1.17	0.157 (0.002)	1.17
1995-96	1.15	1.13	0.142 (0.059)	1.15	0.142 (0.002)	1.15
1996-97	1.22	1.17	0.198 (0.059)	1.22	0.198 (0.002)	1.22
1997-98	1.04	0.97	0.043 (0.06)	1.04	0.043 (0.002)	1.04
1998-99	1.21	1.13	0.193 (0.059)	1.21	0.193 (0.002)	1.21
1999-00	1.03	1.00	0.032 (0.061)	1.03	0.032 (0.002)	1.03
2000-01	1.20	1.09	0.18 (0.059)	1.20	0.18 (0.002)	1.20
2001-02	1.16	1.05	0.147 (0.059)	1.16	0.147 (0.002)	1.16
2002-03	1.22	1.14	0.196 (0.059)	1.22	0.196 (0.002)	1.22
2003-04	1.36	1.25	0.304 (0.057)	1.36	0.304 (0.002)	1.36
2004-05	1.19	1.08	0.171 (0.059)	1.19	0.171 (0.002)	1.19
2005-06	1.36	1.23	0.31 (0.057)	1.36	0.31 (0.002)	1.36
2006-07	1.00	1.00	0	1.00	0	1.00

Table 4.3: Underhill and GLM estimates. Year factor estimates for December and January Wigeon counts for 198 sites chosen by the top sites criterion. 2361 of the original observations were flagged as incomplete and 4034 were missing; the remaining 9841 counts were flagged as complete. The Underhill algorithm results in the year factors $U2y_j$; or $U1y_j$ if completeness is ignored (i.e. incomplete counts are assumed complete). The glm routine in R assuming a Poisson distribution (xGLM) and quasi-Poisson distribution (xqGLM) give year factor estimates Y_j (standard error in parentheses) which are transformed to give $y_j = e^{Y_j}$.

4.2.4 Stratified Underhill Model




Figure 4.3: Total mallard population on selected sites in December after imputing missing values. The graph compares the population on large inland waterbodies (blue), small inland waterbodies (green) and coastal sites (red). For ease of interpretation, although data from September to January was used to generate missing values, the total number of mallard in December only is shown. Due to the nature of the Underhill imputing algorithm, this is a rough guide to the population size and not an accurate population estimate. Originally in Kershaw and Frost (2006).

The Underhill model is usually fitted to all sites together. However, due to variation in population and seasonality within species on different habitats and in different regions of the U.K., a stratified approach can be used, fitting the model to sites within each habitat or region stratum separately (Bell, 199-). In Kershaw and Frost (2006) the method was used to model Mallard (Figure 4.2) data to establish that the Mallard population on large inland waterbodies has approximately halved over the last 35 years, whereas the population on small inland WeBS sites has remained relatively stable (Figure 4.3). The coastal mallard population appears to have almost doubled from 1966 to 1986, but to have then returned to 1966 levels by 2002. The increase at coastal sites in the early 1980s disguised a continual decline on large inland waterbodies throughout the whole recording period in the non-stratified Underhill index.

4.3 Assessment of the Underhill S-Y-M model

4.3.1 Over-dispersion

The dispersion factor in equations 4.4 and 4.5 is not needed to estimate the site, year and month effects but, following McCullagh and Nelder (1989) it may be estimated by:

$$\hat{D} = \sum_{ijk} \frac{(y_{ijk} - s_i y_j m_k)^2}{s_i y_j m_k} \cdot \frac{1}{N - p}. \quad (4.14)$$

A large D indicates that over-dispersion is present and that the quasi-Poisson model is preferable to a Poisson model. When no data is missing, the degrees of freedom

$$N - p = IJK - (I + J + K), \quad (4.15)$$

where I, J , and K are the number of sites, years and months. For an unbalanced GLM, where missing data are excluded, then the sum in equation 4.14 is over the available values and,

$$N - p = IJK - n_M - (I + J + K), \quad (4.16)$$

where n_M is the number of missing values.

In the Underhill algorithm, incomplete values are treated as missing or non-missing on each iteration depending on their value compared to the fitted value, with the larger of the modelled count and observed count used. Bell (1995) estimated the dispersion factor using all complete and incomplete observations in the sum in equation 4.14 and the degrees of freedom with n_M in equation 4.16 as the number of missing values in the original data. This estimate is denoted as D_1 below. However, since some of the incomplete values are treated as missing in the final iteration, an alternative estimate for the dispersion factor, D_2 , treats the incomplete counts as missing or non-missing according to their status in the final iteration. Only those incomplete observations that are not treated as missing are used, together with the the complete observations, in the sum in equation 4.14. The

degrees of freedom are amended accordingly, with n_M in equation 4.16 being the number of missing observations plus those incomplete observations treated as missing in the final iteration of the Underhill algorithm.

In the small example data set of Mute Swan counts at five sites, ten years and one month in Table 4.1 on page 139, all the eight incomplete counts are replaced by the fitted counts by the Underhill algorithm. Five counts are missing. The estimates for the dispersion factor are,

$$D_1 = \frac{1841.98}{50 - 5 - 16} = 63.52$$

and

$$D_2 = \frac{1009.05}{50 - 5 - 8 - 16} = 48.05.$$

Table 4.4 compares the dispersion factors using these two methods, with D_1 treating all incomplete counts as complete and D_2 treating incomplete counts as complete or missing depending on how they are used in the final iteration of the Underhill algorithm. The dispersion factor for each species is given for all sites and winter months and the sites and months chosen for each species using the top sites criterion from Section 2.4.3 and the stable months assessment in Section 3.5.2. The two methods produce similar estimates, D_1 and D_2 , of the dispersion factor. The relative dispersion factors between species is similar for all sites and months and the subset of sites and months. Species which can occur in large congregations but also occur in small groups at other sites and times of year, such as Pink-footed Goose and Teal have large dispersion factors whereas species with smaller variability in numbers which tend to always occur as individuals or small flocks, such as Mute Swan, have smaller dispersion factors. The estimated dispersion factors are much greater for the top sites and stable months data; this is likely to be because the sites and months were chosen to be those that tend to have larger numbers of birds present.

	Peak months and top sites				September to March and all sites			
	<i>I</i>	<i>K</i>	<i>D</i> ₁	<i>D</i> ₂	<i>I</i>	<i>K</i>	<i>D</i> ₁	<i>D</i> ₂
Mute Swan	199	7	35	35	3994	7	9.9	9.3
Whooper Swan	83	4	100	110	1716	7	25	25
Bewick's Swan	44	2	85	96	647	7	40	42
Pink-footed Goose	60	6	3900	4400	1220	7	590	590
European White-fronted Goose	23	2	310	250	510	7	63	63
Greenland White-fronted Goose	15	5	88	85	319	7	45	45
Icelandic Greylag Goose	47	5	2900	3200	958	7	850	870
Re-established Greylag Goose	65	5	120	130	1426	7	38	39
Canada Goose	134	5	160	170	2899	7	59	60
Svalbard Barnacle Goose	10	6	2200	2000	208	7	460	440
Naturalised Barnacle Goose	41	6	28	33	835	7	8.8	9.3
Dark-bellied Brent Goose	75	3	320	270	362	7	200	180
Shelduck	133	3	220	160	1823	7	61	46
Wigeon	198	2	940	870	3384	7	200	180
Gadwall	91	7	64	69	1824	7	19	19
Teal	188	3	380	370	4002	7	84	78
Mallard	278	5	230	200	5778	7	69	64
Pintail	133	3	230	230	1473	7	54	51
Shoveler	118	7	76	79	2118	7	21	21
Pochard	180	2	220	230	3326	7	53	53
Tufted Duck	205	7	140	140	4118	7	38	37
Scaup	52	3	250	240	1063	7	32	29
Goldeneye	176	2	41	35	3393	7	12	11
Red-breasted Merganser	85	6	49	47	1495	7	14	12
Goosander	122	3	36	37	2539	7	9.7	9.9
Ruddy Duck	67	5	56	59	1074	7	12	13

Table 4.4: Dispersion factors to 2 significant figures with the number of sites, *I*, and months, *K*, using top sites method from Section 2.4.3 and stable months from Section 3.5.2 (left) and all available sites for September to March (right). The number of years, *J* is 41 for both methods. Two estimates of the dispersion factor are given, *D*₁ includes all incomplete counts and *D*₂ includes only incomplete counts which are treated as non-missing (see text).

4.3.2 Assumptions

Three major assumptions of the Underhill model are:

Annual trends are identical for all months The model assumes there are no month \times year interactions i.e. there is no phenological change.

Seasonal patterns are the same at all sites The model assumes there are no month \times site interactions i.e. birds arrive and leave all areas of the U.K. at around the same time of year.

Annual trends are identical for all sites The model assumes there are no site \times year interactions i.e. population trends are the same across all regions, habitats and individual sites.

Phenological change

Section 3.4 explored spatiotemporal variation in WeBS data, showing that there can be marked changes in seasonality over time in species such as Shoveler. A consequence of no month \times year interactions is that the mean of the model fits for each month will be the same, as can be seen in the bottom left plot of Figure 4.4. Notice in the middle right plot that after imputing using the Underhill algorithm, each month index is forced much closer to the overall index.

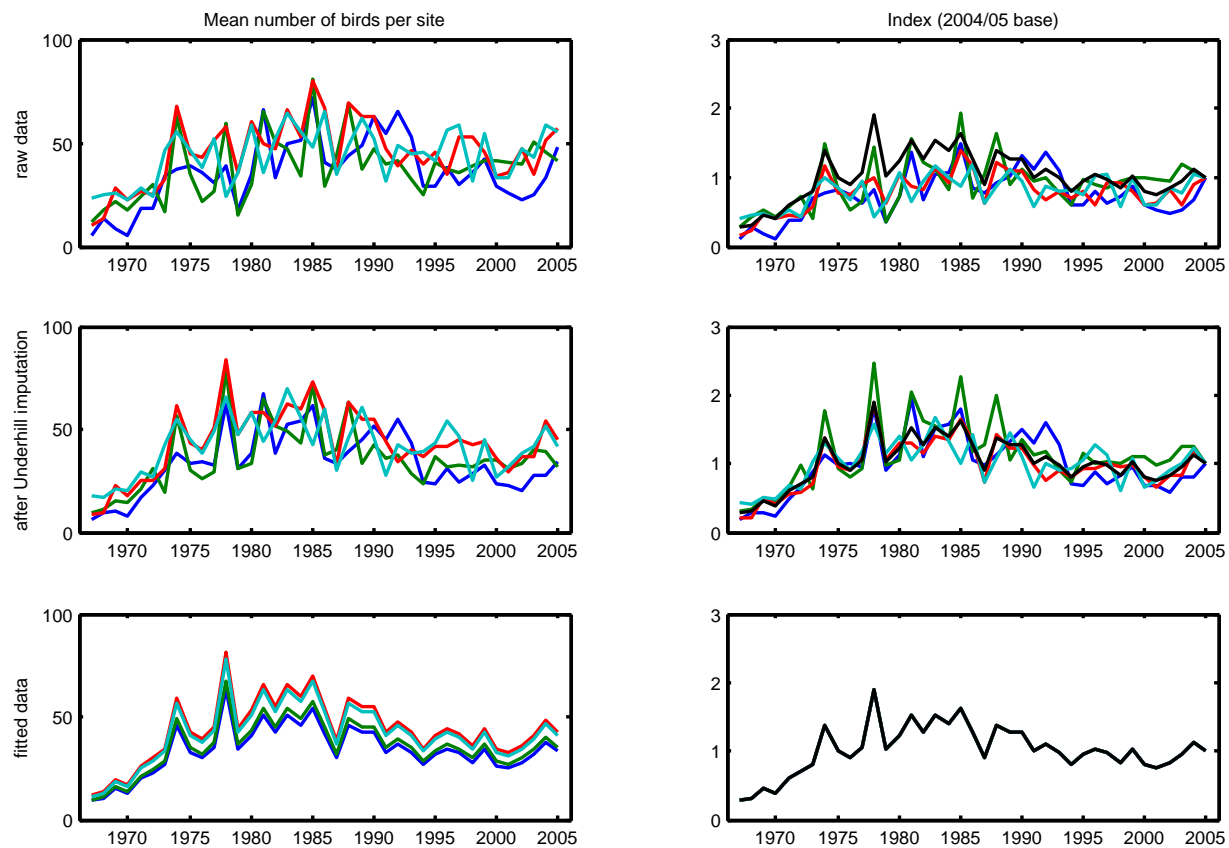


Figure 4.4: Month \times year interactions. The left hand plots show the mean count of Pintail over sites selected by the 50% criterion for October (blue) November (Green) December (Red) and January (turquoise) for 1966-67 to 2004-05. The raw data omitting missing counts is in the top plot, the data with missing counts imputed in the middle and the model fits at the bottom. Dividing each mean count by the value in the the final (2004/05) to obtain an index gives the results on the right-hand plots.

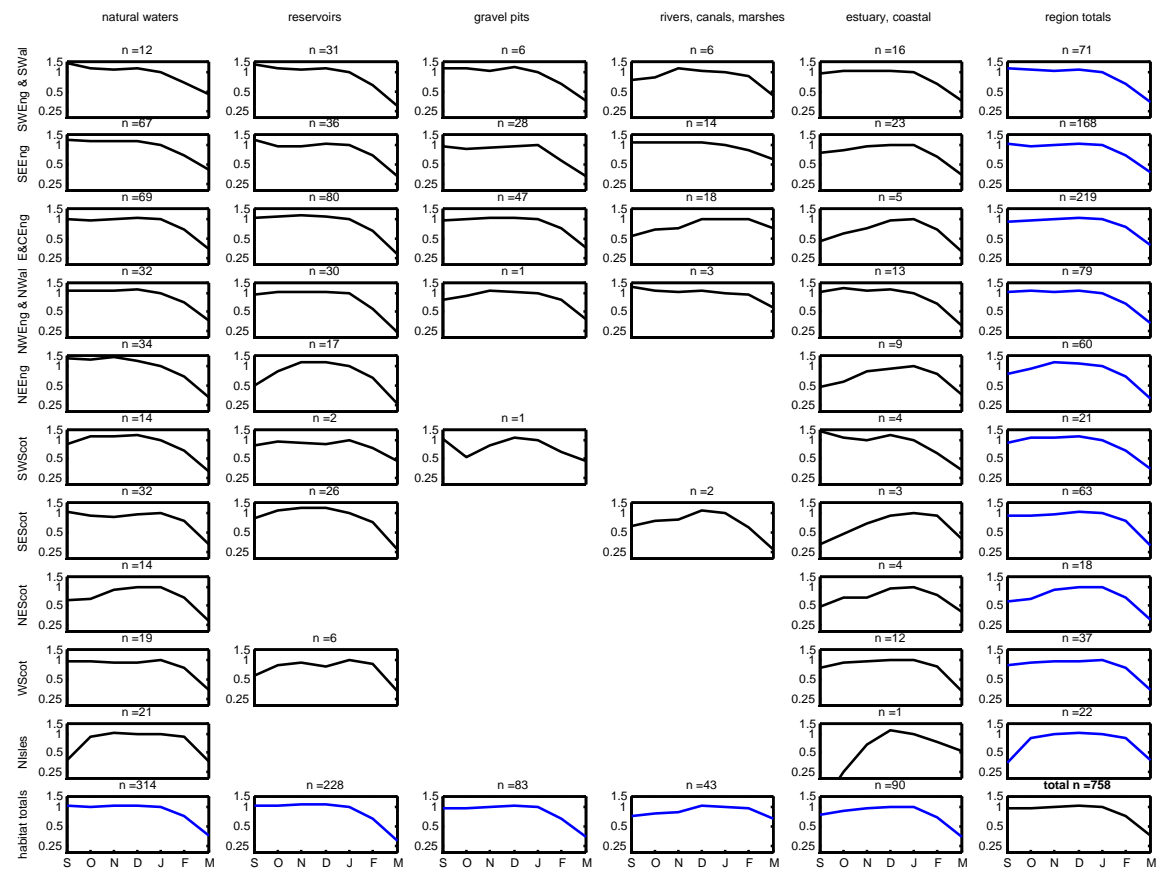


Figure 4.5: Month index plots for mallard counts for region and habitat strata. The index for December is set to be 1 and a log scale is used. The number of sites used to produce each plot is shown (based on the 50% completeness criterion and data from 1966/67 to 2002/03). Blank plots indicate there were no sites in a particular region-habitat combination. The overall month index is also shown.




Figure 4.6: Counts (blue) and Fitted counts (black) from the Underhill model based on December counts only, for Wigeon at the Ribble Estuary. Incomplete counts in the years 1966/67 to 1973/74; 1977/78; 1994/95; 1999/2000; 2000/01; 2003/04 and 2006/07 are indicated by stars. The Underhill algorithm uses the larger of the incomplete count and the fitted count for the index. The values used to calculate the index are indicated by red circles.




Figure 4.7: December Underhill index for Wigeon. Based on all 2693 sites.

Site seasonality

This assumption is more likely to be violated in some species than others; species which migrate from abroad may be expected to arrive at sites nearest their country of origin before moving on to other regions of the country. Sites may be grouped into strata based on their location (region) and type (habitat). Even resident species such as Mallard show some variation in their seasonal abundances on some site strata, for example the numbers on rivers, canals and marshes do not decrease proportionally in spring as they do on other habitat types (Figure 4.5).

Site trends

The Mallard stratified Underhill example in Figure 4.3 on page 146 showed that different types of site can have important differences in their population trends. This is also true on an individual site basis. The Underhill model assumes that the count can be decomposed into a site factor, a year factor and a month factor. Influential sites (i.e. those with a relatively large site factor) will greatly affect the values the algorithm converges to for the year and month factors. For clumped species Section 2.4.1 showed that a few sites can almost completely determine the shape of the annual index plot and missing values in influential sites have a particularly pronounced effect.

As may be expected, when there is not enough information to guide the algorithm, the uncertainty can give some unexpected results. For example, the Ribble estuary is the most influential site out of 2693 sites for December Wigeon counts, with a site influence of over 11%. However, there seems to be a strong individual site trend at the Ribble Estuary (Figure 4.6). There have been significant changes to the site over the time period; since it was purchased by the British government in 1979 and put into management as a National Nature Reserve there has been a decrease in wildfowling and disturbance and an increase in the area of salt marsh.

Assuming there has been a real tenfold increase in Wigeon at the Ribble Estuary from

under 10000 to around 100000 in the 1990s, there are two consequences. One is that in early years where there are incomplete counts too many birds have probably been imputed, with up to 30000 “imaginary” birds added (Figure 4.6). The second, related, problem, is that in years where imputing does **not** take place at Ribble in the period 1966/67 to 1989/90, i.e. before the Ribble population rose dramatically, the annual index will be more depressed than it should be. This is due to the index as a whole depending highly on Ribble, and can be seen in Figure 4.7 where the large fluctuations in the first half of the time series are closely related to whether data from Ribble was marked incomplete and therefore imputed (compare with Figure 4.6).

4.3.3 *GLM Interactions*

One way of dealing with violations in the three assumptions discussed in Section 4.3.2 is to add corresponding interaction terms giving extended models, as in equations 4.11, 4.12 and 4.13 in Section 4.2.3. Tables 4.6 and 4.7 give annual index estimates for Pintail, Shoveler and Wigeon produced from the predicted values from GLMs with interaction terms included, assuming a quasi-Poisson distribution for the counts.

Over-dispersion can be real or apparent, where apparent over-dispersion can be due to missing covariates or interactions (Zuur et al., 2009). Tables 4.6 and 4.7 show that adding interaction terms decreases the estimates of the over-dispersion parameter D , reducing some of the apparent over-dispersion. However, the flocking behaviour of wildfowl results in a large amount of over-dispersion remaining.

In practice, it is not possible to fit site \times year interactions in R unless the number of sites is greatly reduced. For example, by only modelling ten Wigeon sites it is possible to fit a model contain site interactions with both month and year factors (Table 4.7). This example shows that adding individual site trends delivers much greater reduction in the deviance and over-dispersion than interactions for site seasonality or phenological change. An element of caution must be taken in interpreting the results of the this model

Term omitted	df	Deviance	F statistic	$\mathbb{P}(F)$
Pintail $\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + S_i \times M_k + Y_j \times M_k$				
None		2301675		
$Y_j \times M_k$	80	2361929	3.861	$< 2 \times 10^{-16}$
$S_i \times M_k$	263	2419764	2.302	$< 2 \times 10^{-16}$
Shoveler $\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + S_i \times M_k + Y_j \times M_k$				
None		1230783		
$Y_j \times M_k$	240	1263859	2.726	$< 2 \times 10^{-16}$
$S_i \times M_k$	702	1588138	10.07	$< 2 \times 10^{-16}$
Wigeon $\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k + S_i \times Y_j + S_i \times M_k + Y_j \times M_k$				
None		316117		
$Y_j \times M_k$	40	414947	2.704	$< 2 \times 10^{-16}$
$S_i \times M_k$	9	510500	23.64	$< 2 \times 10^{-16}$
$S_i \times Y_j$	359	3716313	10.37	$< 2 \times 10^{-16}$

Table 4.5: Model Selection. Results of Analysis of Deviance for omitting one interaction term using the drop1 routine in R for the most complex models for Pintail and Shoveler (see Table 4.6) and ten Wigeon sites (see Table 4.7).

and those with interactions in Table 4.6 as some of the interaction terms are not estimable as there is no non-missing data available for the combination of factors. A decision on whether to drop extra interactions from the complex models can make use of the ‘drop1’ routine in R for the fitted quasi-Poisson GLM, which uses an analysis of deviance to compare nested models using an F statistic, dropping one term at a time. For each of the three example species data sets, all the interaction terms were found to be significant (Table 4.5).

	Pintail				Shoveler			
	S $+Y+M$	$+Y \times M$	$+S \times M$	$+Y \times M$ $+S \times M$	S $+Y+M$	$+Y \times M$	$+S \times M$	$+Y \times M$ $+S \times M$
1966-67	0.266	0.261	0.264	0.260	0.320	0.323	0.308	0.312
1967-68	0.341	0.335	0.339	0.334	0.266	0.267	0.258	0.258
1968-69	0.378	0.378	0.377	0.376	0.405	0.407	0.392	0.395
1969-70	0.387	0.386	0.386	0.386	0.318	0.321	0.309	0.313
1970-71	0.507	0.506	0.503	0.501	0.425	0.424	0.414	0.414
1971-72	0.495	0.492	0.494	0.491	0.728	0.729	0.710	0.713
1972-73	0.597	0.593	0.595	0.592	0.537	0.543	0.517	0.524
1973-74	1.149	1.153	1.148	1.150	0.607	0.612	0.596	0.601
1974-75	0.959	0.960	0.960	0.963	0.519	0.523	0.512	0.515
1975-76	0.761	0.755	0.761	0.755	0.547	0.547	0.539	0.541
1976-77	0.846	0.846	0.845	0.846	0.557	0.559	0.548	0.550
1977-78	0.896	0.948	0.905	0.967	0.583	0.587	0.575	0.579
1978-79	0.544	0.529	0.547	0.532	0.526	0.535	0.516	0.523
1979-80	0.865	0.859	0.859	0.854	0.545	0.550	0.538	0.541
1980-81	0.983	0.983	0.980	0.979	0.561	0.563	0.557	0.559
1981-82	0.916	0.913	0.917	0.914	0.584	0.587	0.585	0.590
1982-83	1.063	1.063	1.065	1.065	0.613	0.616	0.613	0.616
1983-84	0.925	0.924	0.925	0.924	0.667	0.672	0.661	0.664
1984-85	1.267	1.264	1.275	1.270	0.567	0.570	0.568	0.568
1985-86	1.087	1.067	1.094	1.073	0.605	0.608	0.605	0.608
1986-87	0.755	0.755	0.757	0.756	0.585	0.587	0.588	0.587
1987-88	1.231	1.231	1.235	1.233	0.628	0.632	0.630	0.632
1988-89	1.065	1.065	1.067	1.067	0.628	0.632	0.630	0.632
1989-90	1.087	1.087	1.088	1.087	0.687	0.691	0.684	0.685
1990-91	0.863	0.864	0.865	0.864	0.696	0.700	0.695	0.697
1991-92	1.056	1.056	1.056	1.055	0.740	0.744	0.740	0.742
1992-93	0.842	0.841	0.842	0.841	0.678	0.681	0.681	0.684
1993-94	0.846	0.846	0.847	0.846	0.742	0.745	0.745	0.748
1994-95	0.899	0.899	0.900	0.899	0.838	0.841	0.842	0.846
1995-96	1.002	1.000	1.002	1.000	0.941	0.943	0.944	0.946
1996-97	1.121	1.121	1.120	1.118	0.790	0.795	0.785	0.791
1997-98	0.867	0.867	0.866	0.866	0.884	0.889	0.875	0.879
1998-99	1.098	1.099	1.097	1.097	0.769	0.772	0.765	0.769
1999-00	0.796	0.797	0.796	0.796	0.898	0.903	0.898	0.903
2000-01	0.867	0.868	0.866	0.866	1.061	1.055	1.050	1.040
2001-02	1.053	1.051	1.053	1.051	1.128	1.126	1.120	1.119
2002-03	1.058	1.061	1.058	1.058	1.073	1.073	1.062	1.063
2003-04	1.083	1.082	1.081	1.081	0.920	0.923	0.907	0.910
2004-05	1.047	1.047	1.047	1.047	1.060	1.061	1.056	1.057
2005-06	1.071	1.072	1.070	1.071	1.176	1.183	1.172	1.179
2006-07	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
D	256	249	238	230	84	81	63	61
Deviance $\times 10^5$	24.79	24.20	23.62	23.02	16.24	15.88	12.64	12.31

Table 4.6: Pintail and Shoveler GLMs with interaction terms. Annual Index estimates for Pintail for November to January for 133 sites and Shoveler counts for 118 sites for September to March. Sites and months chosen using the top sites method from Section 2.4.3 and stable months from Section 3.5.2. Indices were calculated from the predicted values from the glm routine in R, assuming a quasi-Poisson distribution using the model without interaction terms ($\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k$) and then with additional interaction terms as indicated. Also given is the over-dispersion estimate, \hat{D} , and the deviance for each model.

	S				+Y × M
	+Y	+Y × M	+S × M	+S × Y	+S × M
	+M				+S × Y
1966-67	0.273	0.272	0.270	0.250	0.249
1967-68	0.323	0.315	0.316	0.334	0.300
1968-69	0.284	0.283	0.282	0.271	0.270
1969-70	0.416	0.414	0.413	0.403	0.398
1970-71	0.423	0.419	0.415	0.398	0.389
1971-72	0.490	0.491	0.486	0.468	0.466
1972-73	0.461	0.459	0.455	0.510	0.464
1973-74	0.554	0.555	0.553	0.546	0.544
1974-75	0.444	0.445	0.443	0.404	0.403
1975-76	0.344	0.346	0.342	0.328	0.327
1976-77	0.449	0.450	0.448	0.443	0.441
1977-78	0.382	0.382	0.381	0.377	0.375
1978-79	0.404	0.404	0.403	0.398	0.397
1979-80	0.451	0.452	0.450	0.445	0.444
1980-81	0.582	0.583	0.581	0.574	0.572
1981-82	0.623	0.624	0.622	0.614	0.612
1982-83	0.647	0.648	0.645	0.638	0.636
1983-84	0.420	0.421	0.419	0.414	0.413
1984-85	0.545	0.545	0.543	0.537	0.535
1985-86	0.678	0.679	0.676	0.668	0.666
1986-87	0.603	0.604	0.605	0.577	0.575
1987-88	0.838	0.839	0.836	0.826	0.823
1988-89	0.667	0.668	0.665	0.657	0.655
1989-90	0.934	0.935	0.932	0.921	0.917
1990-91	0.879	0.880	0.876	0.866	0.863
1991-92	1.174	1.175	1.171	1.158	1.153
1992-93	1.048	1.049	1.046	1.034	1.030
1993-94	1.221	1.222	1.218	1.204	1.199
1994-95	1.416	1.418	1.413	1.396	1.391
1995-96	1.344	1.345	1.340	1.325	1.320
1996-97	1.255	1.256	1.252	1.237	1.233
1997-98	1.050	1.051	1.047	1.035	1.031
1998-99	1.407	1.409	1.404	1.388	1.383
1999-00	1.017	1.018	1.014	1.003	0.999
2000-01	1.263	1.265	1.260	1.246	1.241
2001-02	1.230	1.231	1.227	1.213	1.208
2002-03	1.233	1.234	1.230	1.215	1.211
2003-04	1.516	1.517	1.512	1.494	1.489
2004-05	1.173	1.174	1.170	1.157	1.152
2005-06	1.433	1.435	1.430	1.413	1.408
2006-07	1.000	1.000	1.000	1.000	1.000
<i>D</i>	5697	5814	5347	1421	867
Deviance × 10 ⁴	399.15	390.34	380.63	59.73	31.61

Table 4.7: Wigeon GLM with interaction terms. Annual Index estimates for Wigeon counts for the ten most influential sites for December and January. Sites chosen using a similar methodology to the top sites criterion. Indices were calculated from the predicted values from the glm routine in R, assuming a quasi-Poisson distribution using the model without interaction terms ($\mathbb{E}[\log(x_{ijk})] = S_i + Y_j + M_k$) and then with additional interaction terms as indicated. Also given is the over-dispersion estimate, \hat{D} , and the deviance for each model.

4.3.4 Consistency intervals

Underhill and Prŷs-Jones (1994) developed a concept of ‘consistency intervals’, analogous to confidence intervals, using a bootstrapping procedure. The bootstrap samples are samples of sites with replacement and measure the consistency of change in population across sites. Therefore, they provide a partial assessment of the assumption that there are no site \times year interactions. They are termed consistency intervals rather than confidence intervals because they do not reflect all uncertainty about the index, such as that due to measurement error or unusual monthly counts.

Due to computing limitations at the time, the method used to generate the intervals did not at first refit the Underhill model for each sample of sites; instead the counts were imputed and bootstrap samples were then taken. Computing power has considerably increased in the intervening period; Program 4.3 is a MATLAB routine to take bootstrap samples and refit the Underhill model for the bootstrap sample each time. The number of bootstrap samples used for this thesis was 200, as this is sufficient to estimate the intervals and took a more reasonable amount of computer time than 1000 samples (see Figure 4.8). The total time taken for 200 bootstrap samples for all twenty-six species was just under 30 hours using all sites and months for all species, and less than 1 hour using the site and month selections from Sections 2.4.3 and 3.5.2 respectively*.

It can be seen in Figures 4.9 (page 163) and 4.10 (page 164) and Table 4.8 (page 167) that the consistency intervals calculated by refitting the the model are significantly different from those using the original method, where missing counts were not imputed. It is at first surprising to see that the 95% site consistency intervals do not always contain the point index estimates, for example the Goldeneye index in 1971/72, as shown in Figure 4.10. However this does not necessarily mean the index is wrong, rather it indicates that the index is particularly sensitive to a small number of sites at that particular point, in this case the most influential site for Goldeneye, the Forth estuary.

*On a computer with a dual core processor (2.66GHz) and 2 GB DDR2 memory.

Figures 4.11 and 4.12 (pages 165 and 166) compare the bootstrap 95% consistency intervals of the indices using all sites and months and the **stable month** and **top site** selections from Sections 2.4.3 and 3.5.2. The results are very similar for most species; the most striking difference is for Icelandic Greylag Goose. The proportion of observed birds on the sites and months selected by the top sites criterion for this species dropped from between 45% and 77% in the 1966-67 to 1995-96 period to 19% and 35% in the 1996-97 to 2006-7 period. This is because a new site (WWC07 Orkney), which has lately been the most important site for this species, only has data from 1998 onwards and was thus omitted by the top sites criterion for not having ten years of data available. In addition the complication of dividing Icelandic and resident Greylag goose in Scotland means caution is urged in interpreting any results for Icelandic Greylag Goose in this thesis; see Austin et al. (2008) for recent trends. In general however, it is recommended that the site and months selections from Sections 2.4.3 and 3.5.2 are used in future due to the benefits of faster computing time and added clarity as to which sites and months the index is tracking.

Program 4.3 MATLAB code to generate bootstrap Confidence Intervals over sites.

```

%UnderhillCI
%This script calculates Bootstrap CIs that accurately reflect the
%uncertainty by refitting the Underhill index for the bootstrap sample of
%sites.

%Default to 100 bootstrap samples and standard beta parameters.
if ~exist('nboot','var'), nboot = 100; end, iota.nboot=nboot; clear nboot
setdefaultbeta
%Fit Underhill model.
[yearfactor, monthfactor, sitefactor, fits] = Underhillfn(newdata,beta);

%Obtain nboot bootstrap samples of sites.
[maxs, iota.bootsamples] = bootstrp(iota.nboot,'max',(1:beta.nsites));

%Reshape data to a 3D array - site*row*(year month count quality)
site3darray = zeros(beta.nyears*beta.nmonths,4,beta.nsites);
for i = 1:beta.nsites
    site3darray(:, :, i) = newdata(((beta.nyears*beta.nmonths)...
        *(i-1)+1):((beta.nyears*beta.nmonths)*i),2:5);
end
%Fit Underhill model for each bootstrap sample. In sampledata the original
%unique site blocks in newdata are replaced by the data from the sites in
%the bootstrap sample. The site numbers as left as "1 to beta.nsites"
%rather than replaced by the original site number as there can't be
%duplicate numbers for Underhillfn.
iota.timer = zeros(iota.nboot,1);
yearbootsum = zeros(beta.nyears,1); monthbootsum = zeros(beta.nmonths,1);
iota.yearsample = zeros(iota.nboot,beta.nyears);
iota.monthsample = zeros(iota.nboot,beta.nmonths);

h=waitbar(0,'Bootstrapping Progress...');
for i = 1:iota.nboot
    waitbar(i/iota.nboot)
    sampledata = newdata;
    for j = 1:beta.nsites
        sampledata(((beta.nyears*beta.nmonths)*(j-1)+1)...
            :((beta.nyears*beta.nmonths)*j),2:5) = ...
            site3darray(:, :, iota.bootsamples(j,i));
    end
    tic, [s, y, m, data] = Underhillfn(sampledata,beta);
    iota.timer(i) = toc;
    %Compute the proportions of birds in the bootstrap sample for each year
    %and month.
    sortedy=data(beta.sortyindex);    sortedm=data(beta.sortmindex);
    for j = 1:beta.nyears
        yearbootsum(j) = sum(sortedy(beta.nsites*beta.nmonths*(j-1)...
            +1:beta.nsites*beta.nmonths*j));
    end
    for j = 1:beta.nmonths
        monthbootsum(j) = sum(sortedm(beta.nsites*beta.nyears*(j-1)...
            +1:beta.nsites*beta.nyears*j));
    end
    iota.yearsample(i, :) = yearbootsum/sum(yearbootsum,1);
    iota.monthsample(i, :) = monthbootsum/sum(monthbootsum,1);
end, close(h)
A = sort(iota.yearsample);
B=repmat(A(end,:),41,1);    A=A./B;
yearlower = C(:,0.05*200); yearupper = C(:,0.95*200);

```





Figure 4.8: Comparison of site consistency intervals for Dark-bellied Brent Goose (top) and Red-breasted Merganser (bottom) with 100 (black), 200 (blue) and 1000 (red) bootstrap samples. Indices based on data for all sites from September to March.




Figure 4.9: Site consistency intervals (CIs) for Pintail, using data from November to January for 133 sites chosen using the top sites method from Section 2.4.3 and stable months from Section 3.5.2 (bottom) and all available sites for September to March (top). CIs were calculated by bootstrap sampling of sites and then fitting the Underhill model to impute missing counts (left) and by fitting the model and imputing missing counts and then undertaking bootstrap sampling of sites (right).




Figure 4.10: Site consistency intervals (CIs) for Goldeneye, using data from January to February for 176 sites chosen using the top sites method from Section 2.4.3 and stable months from Section 3.5.2 (bottom) and all available sites for September to March (top). CIs were calculated by bootstrap sampling of sites and then fitting the Underhill model to impute missing counts (left) and by fitting the model and imputing missing counts and then undertaking bootstrap sampling of sites (right).

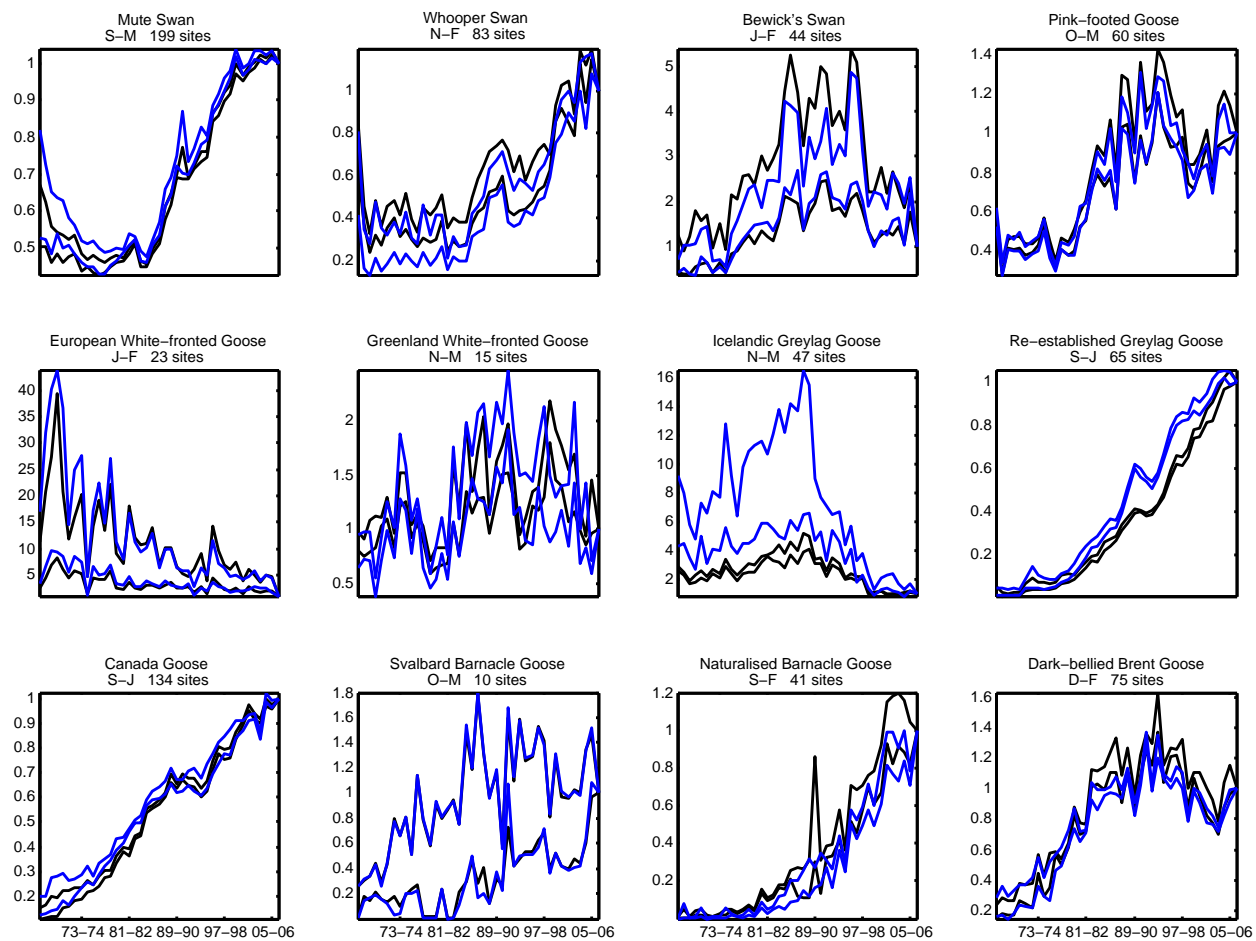


Figure 4.11: Site consistency intervals (CIs) for the Underhill index for twelve swan and geese species, using top sites method from Section 2.4.3 and stable months from Section 3.5.2 (blue) and all available sites for September to March (black).

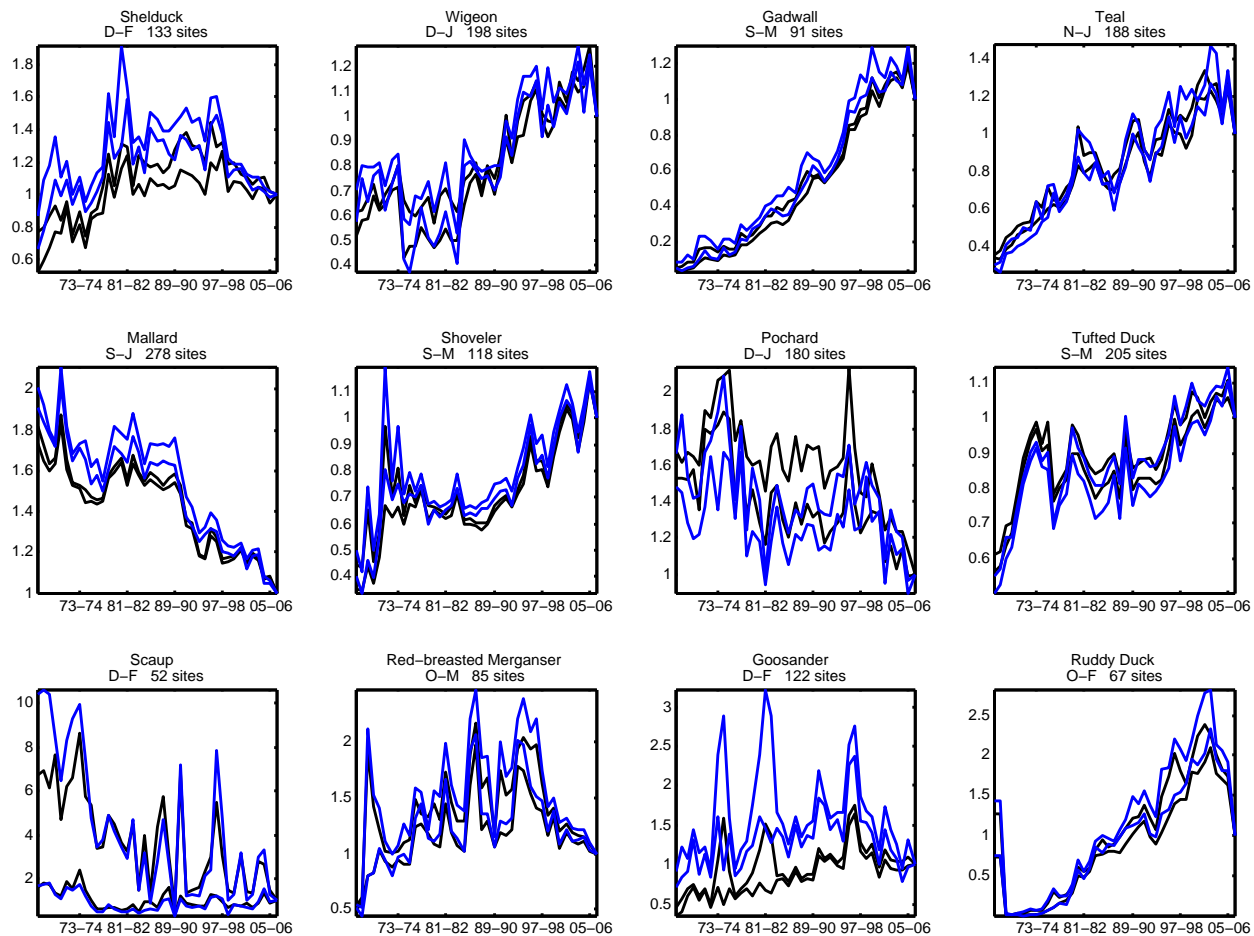


Figure 4.12: Site consistency intervals (CIs) for the Underhill index for twelve ducks species, using top sites method from Section 2.4.3 and stable months from Section 3.5.2 (blue) and all available sites for September to March (black). The plots for Pintail and Goldeneye may be found in Figures 4.9 and 4.10 respectively.

	Peak months and top sites				September to March and all sites			
	CI refit		CI imputed		CI refit		CI imputed	
	Lower	Upper	Lower	Upper	Lower	Upper	Lower	Upper
Mute Swan	1.9	1.8	3.0	3.5	1.1	1.3	1.4	1.5
Whooper Swan	3.7	14.4	9.8	18.4	1.6	11.2	5.3	6.3
Bewick's Swan	3.1	23.6	10.6	23.4	2.4	49.1	6.3	17.7
Pink-footed Goose	2.0	4.2	8.2	11.4	1.4	6.0	5.8	7.1
European White-fronted Goose	19.0	23.0	23.7	59.6	18.9	28.7	17.4	32.7
Greenland White-fronted Goose	9.1	4.6	6.8	7.2	4.1	5.2	2.5	2.6
Icelandic Greylag Goose	7.8	20.1	9.8	13.5	2.5	3.1	3.5	4.2
Re-established Greylag Goose	6.5	7.1	6.1	6.9	6.2	4.7	3.3	3.4
Canada Goose	3.6	2.6	5.0	5.8	2.5	1.9	2.0	2.1
Svalbard Barnacle Goose	26.1	5.1	6.2	5.8	25.1	5.2	5.0	4.9
Naturalised Barnacle Goose	15.2	11.7	17.2	23.3	7.4	11.3	5.1	7.1
Dark-bellied Brent Goose	4.6	3.0	5.8	6.6	5.1	2.2	6.3	7.3
Shelduck	2.9	1.4	4.5	4.8	2.4	2.5	4.5	5.1
Wigeon	2.7	2.5	5.5	6.7	2.6	1.8	3.7	3.9
Gadwall	5.6	4.0	6.1	7.1	4.2	3.8	3.7	3.8
Teal	2.0	2.9	5.1	6.0	1.2	2.6	2.7	3.0
Mallard	1.0	0.7	2.4	2.5	0.3	0.4	0.9	0.9
Pintail	7.9	2.7	12.9	18.1	6.0	2.9	9.7	13.0
Shoveler	1.5	2.3	4.5	5.4	1.6	1.4	2.8	3.2
Pochard	2.1	3.6	7.1	9.4	2.4	4.2	4.5	4.9
Tufted Duck	1.4	1.5	4.0	4.3	0.8	1.2	2.0	2.0
Scaup	24.2	16.5	22.9	51.4	22.2	17.6	16.0	27.4
Goldeneye	6.7	6.6	8.9	11.5	4.0	4.4	3.5	4.1
Red-breasted Merganser	5.0	3.7	7.0	7.2	5.7	1.8	3.5	3.8
Goosander	5.2	3.2	7.1	8.2	4.0	3.1	2.2	2.6
Ruddy Duck	10.8	6.3	15.6	20.5	9.1	5.5	11.6	14.4

Table 4.8: Bootstrap site consistency interval (CI) measures, using top sites method from Section 2.4.3 and stable months from Section 3.5.2 (left) and all available sites for September to March (right). The measure indicates the relative width of the upper and lower 95% CI bounds to the index, calculated by the distance from bound to index, as a proportion of the index, summed over all 41 years from 1966/67 to 2006/07.

4.4 Discussion

This chapter has taken a detailed look at the model currently used to impute missing counts in Wetland Bird Survey data to produce a complete data set which is then used to generate national annual population indices.

Section 4.2 detailed the site-year-month multiplicative model and the algorithm Underhill (1989) developed for fitting it. Underhill's method was compared to the equivalent Generalised Linear Model and differences in the annual indices were due to the Underhill algorithm's treatment of counts marked as poor quality (incomplete), where the larger of the fitted and observed count is used.

A closer look at the model and its limitations was taken in Section 4.3. Table 4.4 showed that over-dispersion is present in all species. Some of the over-dispersion will be 'real' (the variance in the number of birds present is greater than the mean number of birds present). However, some apparent over-dispersion may be due to limitations in the model such as a lack of interaction effects and covariates such as weather (see Section 3.2.3).

Following on from Chapters 2 and 3 we see that the assumptions implicit in the Underhill model, that there are no differences in trends between sites, no phenological change and no spatial effects on seasonality patterns, are biologically suspect. Ignoring incomplete counts and using GLMs with interaction terms showed that the related interaction parameters were significant. The Underhill algorithm's treatment of incomplete counts has a greater impact on the resultant index values than adding interaction terms involving months in a GLM framework (Table 4.6 and Figure 4.13).

It was not possible to fit a model with site \times year interactions for a full data set. Results from using a small number of sites for Wigeon indicated that the site \times year interactions were more significant than the site \times month or year \times month interactions (Tables 4.5 and 4.7). For clumped wildfowl species, it is the assumption that year effects are the same for all sites which seems most biologically implausible, and the assumption that will impact annual indices most severely. Sometimes species which favour traditional




Figure 4.13: Pintail indices, using data from November to January for 133 sites chosen using the top sites method from Section 2.4.3 and stable months from Section 3.5.2. Indices were produced after fitting counts with the Underhill incompleteness method (black); a GLM with no interactions (blue) and a GLM with site \times month and year \times month interactions (red). The latter two indices are very similar (see Table 4.6).

wintering grounds do start favouring a new area, and decline dramatically at the old area. If observers then stop visiting the site as there are no longer any birds to count, the ensuing missing data will be imputed with all the birds that have moved to the new site, possibly effectively being counted twice. It could be argued then that it would be better to have a model that assumed low numbers of birds at one site is likely to be a consequence of high numbers at a nearby site. Future research could concentrate on fitting individual site trends for the most influential sites for a species. An advantage of the top sites criterion for selecting sites to index on (see Section 2.4.3) is that by using a smaller subset of the available sites in the database, exploration of individual site effects becomes more attainable.

It is unclear that the Underhill method of taking the larger of the fitted and observed count for incomplete counts produces better index estimates, particularly due to the impact of individual site trends. The reasoning behind using the larger of the fitted and

original count for poor quality observations is that a count is marked as incomplete when a partial count is made of a site, implying a minimum of the number of birds present. When presenting the case for using this method, using a set of 112 estuaries for Grey Plover *Pluvialis squatarola* populations, Underhill and Prÿs-Jones (1994) argued that the introduced bias would be small, since there are relatively few incomplete counts and that, assuming the fitted values to be unbiased, the error in replacing a correct small observed count with an incorrect larger one would on average occur half the time. However, in Chapter 2 we saw that on large sites made up of several count units it is often the case that the entire site is not surveyed, leading to proportionally more counts being marked as incomplete than on smaller sites. Since these are the sites that tend to contain larger populations and consequently have large site factors, this means the Underhill treatment of incomplete counts tends to amplify high index numbers (e.g. Figure 4.13). This has implications for wetland bird conservation, since the upward bias may obscure the extent of a population decline.

Obtaining meaningful measures of confidence around the annual indices is not straightforward. Underhill and Prÿs-Jones (1994) used bootstrapping to obtain approximate confidence intervals (CIs). One of the assumptions of bootstrapping is that the observations are independent and identically distributed; Underhill and Prÿs-Jones (1994) note that effectively when bootstrapping over sites it is assumed that the sites are a random sample of an infinite set of wetlands, which is not the case with WeBS sites (see Chapter 2) and hence call the resulting CIs site Consistency Intervals, since they measure the consistency of the change in populations across sites. Previously site consistency intervals for Underhill indices have been produced by fitting the Underhill model and then using bootstrap samples of sites. Here for the first time site consistency intervals are generated by refitting the Underhill model for each sample of sites. This can result in intervals which do not include the index estimate when the index depends very strongly on one or two sites (e.g. Figure 4.10).

This chapter has used modelling techniques to understand the data better and to impute

missing counts to obtain a complete data set. The Underhill index definition in equation 4.7 is not the only way of producing an index for the resultant complete data set; the following chapter looks at alternative index formulae.

Chapter 5

INDICES OF POPULATION CHANGE FOR SEASONAL ECOLOGICAL DATA

5.1 Introduction

Population indices are widely used in ecological site-based surveys to indicate trends in the population size without estimating the population size itself, with the assumption that trends on the surveyed sites will mirror those of the entire population from which the sample was taken. Plotting the indices over time facilitates year-to-year comparisons of population size and evaluation of evidence of increasing or decreasing trends.

In statistical applications the term *index* or *index number* refers to the results of a formula that is calibrated to produce a number that is relatively straight-forward for a non-specialist to understand and extract key information from. Index numbers are used in a wide variety of applications; for example, the strength of the sun's ultraviolet radiation is given in weather forecasts in the form of the World Health Organisation's Solar UV index scale, and the Gunning Fog Index of a document indicates how many years of education someone needs to understand the text. Often time series of index numbers are produced, where a statistic at the current time point is calibrated by comparing it to the statistic at a chosen base time point; population indices in biology and share indices in economics are examples of time series indices.

The development of index number theory has been particularly focused on developing measures of general price levels and consequently monetary inflation in an economy. For

comparable commodities, such as pints of milk, the price can be directly measured and the price index of milk may be simply a ratio of the current price to its price at a chosen reference time. However, if we want to know the general price level, it is not possible to measure it directly, as it is made up of a large number of different commodities, measured in different units. Economists are interested in how the inflation in each commodity can be aggregated to indicate the general rate of inflation.

In this chapter, the potential for transferring ideas from economic index number theory to ecological population indices is explored. The chapter begins with an overview of indexing schemes and some key ideas from the biological and economics literature on indexing. Different seasonal index formulae are proposed and compared with some example WeBS data.

5.2 Aggregation of WeBS counts

5.2.1 Introduction

A primary objective of the Wetland Bird Survey is to assess trends in the numbers and distribution of non-breeding waterbird populations in the UK. Annual population indices are used because:

1. Since a national census recording the entire UK population of any wildfowl species is not possible in practice, especially for widespread species which occur on small waterbodies, a sample of sites is used.
2. The “population” varies over the recording season due to death and migration processes.

WeBS counts are recorded for multiple sites and months, so in designing an index formula for WeBS data consideration must be given to the most appropriate way to combine the information. The aim of this chapter is to consider how best to aggregate a complete Wetland Bird Survey data set to give a representative yearly index number.

5.2.2 Sites

In ecological applications it has generally been the case that a simple arithmetic mean non-weighted index of sites has been used, as is currently the case with Wetland Bird Survey data. The most intuitive formula for combining site counts for the whole population is to sum over sites (or equivalently take the arithmetic mean over sites). This is particularly appealing when the population is mobile between sites, as is the case with overwintering wildfowl. It is not the only possible aggregation formula, however: in many economic applications more complicated indices are used such as weighted arithmetic means or geometric means.

The Wetland Bird Survey attempts to collect as much monitoring information as possible on multiple species and relies on unpaid volunteers for data collection. As discussed in Chapter 2, neither waterbodies nor bird populations are randomly distributed throughout Great Britain. It is important to use available data as effectively as possible and other applications indicate that indexing methods may improve the way that population trends are calculated so as better to reflect national population trends.

Parallels with house price indices

Like wildfowl counts, house prices vary spatially, so that property prices in central Scotland may be behaving very differently from prices in southern England, in the same way as wildfowl populations may be doing better in some parts of the country than others. Wetlands vary in habitat classification and in size, so that a site classified as an inland waterbody may be a small fishing pool or a large lake; in ecological terms, the physical size of the wetland may be indicative of its carrying capacity. Similarly, houses are heterogeneous, varying in size and type. Equivalent data that may be collectable for housing is the number of bedrooms as an indication of size and classification into flat or terraced, semi-detached or detached house.

None of the commonly used British house price indices accounts for the size of properties, but some “mix-adjust” types of housing. If no mix-adjustment takes place, a decline in the index may be caused by fewer large detached properties being sold and an increase in the quantity of small flats being sold rather than a general decrease in prices of all property types. Mix-adjustment attempts to correct for low transaction volumes in particular house types by weighting the category averages by the proportion of the total housing stock in that category.

An inventory-adjusted index could be constructed using WeBS data by weighting wetland habitat types by the amount of that habitat available in the U.K. using a national wetland inventory in a similar manner to mix-adjusted house price indices. Work currently being

carried out using GIS and Ordnance Survey habitat data may make this feasible shortly. A clear advantage of an inventory-adjusted index over changing the sampling scheme is that it can be used with the forty years of data already collected.

The government executive agency Land Registry publishes a house price index report monthly that also gives house price indicators stratified by property type and county. As most people are interested in their local property market this is useful information. However, the more specific the category, the fewer the transactions to base the index on and so the data is often noisier. In Figure 4.5 on Page 152, a similar approach was taken with WeBS data, allowing the user to see how a species is faring in different areas and on different habitats.

It is likely that with further advances in the wetland register being constructed by the BTO it would be appropriate to do further research in this area, so that in the medium-term WeBS indices could be constructed to be less-biased and more representative of national population trends on all UK wetland sites.

5.2.3 *Months*

As discussed in Section 3.1, it is often the case in mathematical and statistical ecology that a year is the natural time period for studying a population: perhaps the variable of interest is the number of females at the beginning of the breeding season, for example. However, there are many applications in ecology where the seasonality of the population is important and non-ignorable. The within-season dynamics of animal populations is an important area of research in mathematical and statistical ecology. Animal populations have a seasonal component due to birth and death processes and inward and outward migration. Modelling seasonal dynamics may be especially important with migratory species, as conditions in their wintering grounds can affect their productivity in the subsequent breeding season as well as survival (Norris and Marra, 2007).

An ecological population index typically consists of the ratio of the sample count at the

current time point to the sample count at a base time point. It is assumed that the sample count is at a particular instant in time. In many ecological surveys, the sample counts may be taken over a time interval but the period of data capture is such that the effects of birth, death and migration are negligible. In contrast, WeBS counts are taken over a period of months where, for most species, these effects are not negligible. We have explored the seasonality of over-wintering wildfowl populations using functional data analysis in Chapter 3 and shown the phenology to be of intrinsic interest to ecologists and that an assumption of a stationary over-wintering population is flawed for many species.

In the previous section an analogy was drawn between house prices and WeBS data, with both needing consideration of differing regional and type specific trends. Another problem common to both applications is seasonality. In property, there is usually more trading activity in the spring and early summer, that tends to result in greater house price inflation at those times of year. Many of the house price indices are “seasonally adjusted”, which attempts to smooth out this seasonal effect. Seasonal adjustment is done by dividing each value by an appropriate seasonal index. A potential disadvantage of seasonal adjustment is when the effect is more pronounced in some years than others, making seasonally-adjusted house price indices sometimes difficult to interpret.

House price indices are generally reported on a monthly basis, whereas WeBS indices are reported yearly. Figure 5.1 illustrates some possible approaches to presenting WeBS monthly count data. The species considered here, Dark-bellied Brent Goose, is migratory and does not breed in Britain. Although only a small amount of counts are available for the summer months, they are enough to indicate that the species is not present. A plot showing all the counts (Figure 5.1 (c)) is confusing to the eye as the seasonal fluctuations in population disguise the underlying trend. A basic seasonal adjustment of the figures is of little assistance as if one or two birds are recorded in one of the summer months the seasonally adjusted figure is corresponding very high, again painting a confused picture (figure 5.1 (e) and (f)). Only displaying winter records as in figure 5.1 (b) and (d) results in a less cluttered diagram. The seasonally adjusted plot (d) makes it easy to pick out

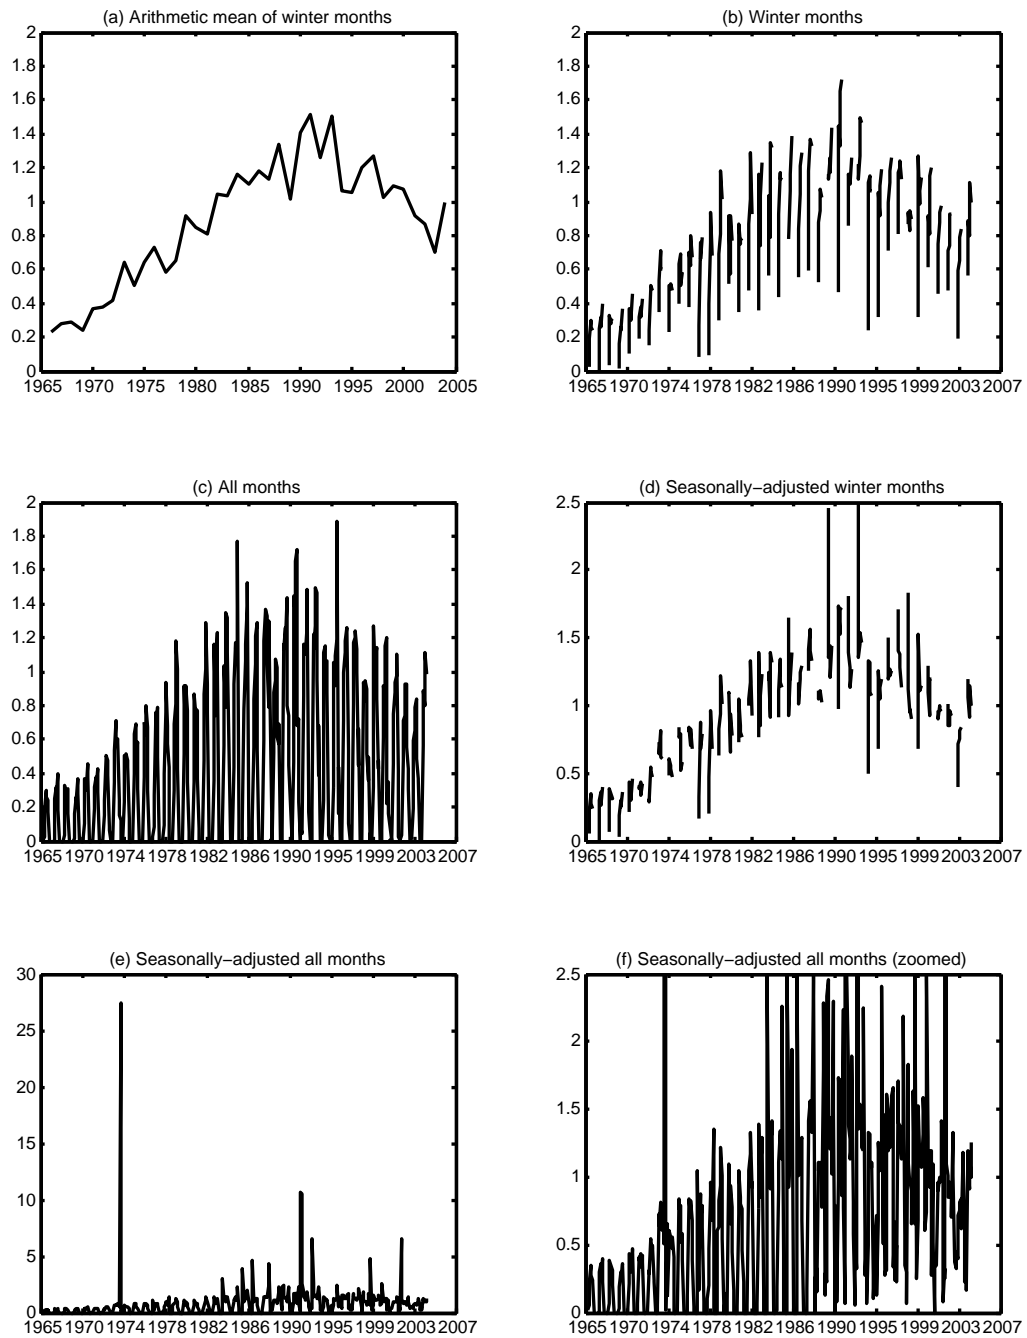


Figure 5.1: Comparison of methods for presenting monthly WeBS counts. The mean recorded counts of dark-bellied Brent goose over all sites in the database are used. (a) A single index value for each year, calculated by the arithmetic mean of the counts from September to March with the final year as the base year. (b) An index value for each month from September to March, with the February count of the final year as the base month. (c) As (b) with all 12 monthly values shown. (d) As (b) but with seasonal adjustment by the monthly means across the entire time series. (e) As (d) with all 12 monthly values shown. (f) As (e) but showing a restricted portion of the y-axis.

unusual monthly counts and also to see the overall trend in the population. Although a plot such as figure 5.1 (d) may be of great use to scientific researchers using the data, it is likely that communication of the results to a wider public will continue to rely on simple annual index plots such as figure 5.1 (a).

Seasonality is also an issue when compiling general price indices, particularly when the seasonal variation is associated with variation in the properties of the product. For example, economists have questioned whether a strawberry consumed in June, grown outside locally, should be treated as the same product as a strawberry grown abroad in a hothouse in the middle of winter. The United Nations guidelines on price and quantity statistics suggest that in general, it is better to treat them as different products (United Nations, 1977).

The strawberry example can inform the problem of obtaining a single annual index for Wetland Bird Survey data when a time series plot such as that in figure 5.1 (a) is required. By considering the monthly counts as different goods, the problem becomes the estimation of the underlying population change from several, sometimes conflicting, data points in a similar way to estimating underlying general inflation from the price movements of individual products. Choosing which monthly counts to include in the index, such as mid-winter counts, September to March counts or all twelve monthly counts of the year is a similar problem to choosing a basket of goods. In Section 5.3 we use ideas from constructing indices to measure general price inflation with seasonal population data.

5.2.4 Phenological implications for reporting trends

One of the current conservation uses for WeBS population indices is to generate alerts when populations are declining over the short, medium and long term. As introduced in Section 3.3.2, WeBS alerts are designed to flag up possible population declines. Full details of the methodology may be found in Atkinson et al. (2006); but the key parts of the method are:

1. Use the Underhill algorithm to impute missing and incomplete counts.
2. Fit a GAM version of equation 4.8 assuming the counts have a Poisson distribution,

$$\mathbb{E}[\log(x_{ijk})] = f(k) + S_i + M_k, \quad (5.1)$$

where f is a smooth function over years which takes the value $f(k)$ in year k and $f(b) = 1$ in the most recent year (the base year).

3. Compare the value of $f(k)$ with that in the base year to assess the percentage decline in the smooth counts. ‘Medium Alerts’ are triggered by a decline of between 25% and 50% ($1.25 \leq f(k) < 1.5$) and ‘High Alerts’ are triggered by a decline of greater than 50% ($f(k) \geq 1.5$) over the short- (5-year), medium- (10-year) and long- (25-year) term.
4. Apply a ‘biological filter’ so that alerts are not issued for species where the decline is within the normal fluctuation in the index for that species.

In the most recent WeBS Alerts report using data from 1966-67 to 2003-04, the alert results for Shelduck were based on data from January and February and the ‘biological filter’ analysis recommended that short-term alerts should not be issued for this species. The Shelduck smoothed index had declined 11% over the short-term, 17% over the medium-term and 20% over the long-term and hence did not flag up any alerts (Maclean and Austin, 2008).

	Sep-Mar	Jan-Feb
Long-term	22.0%	58.3%
Medium-term	25.2%	53.1%
Short-term	7.3%	15.4%

Table 5.1: Shelduck alerts from the smooth trends in Figure 5.2 based on September to March counts and January to February counts. Green indicates no alert, orange a medium alert and red a high alert.

Figure 5.2 shows the smoothed indices from Shelduck counts from September to March, compared to January to February, imputing missing data with the Underhill algorithm and then fitting a GAM similar to that described in Atkinson et al. (2006) by using the ‘gam’ routine from the ‘mgcv’ package in R (Wood, 2006). Table 5.1 gives the percentage change in the smooth Shoveler indices from Figure 5.2 over the three alert time periods. As noted in Section 1.1.2, traditionally most Shelduck in northwest Europe have migrated to the German Wadden Sea area in autumn to moult, but there has been a trend towards the use of the Bridgwater Bay area of the Severn Estuary by some individuals (Fox and Salmon, 1994). As a result, numbers present in the autumn are increasing, whereas the




Figure 5.2: Shelduck Underhill indices based on September to March and January and February with smoothed trends produced by a GAM (see text). Data for 111 sites chosen by the top sites criterion for September to March.

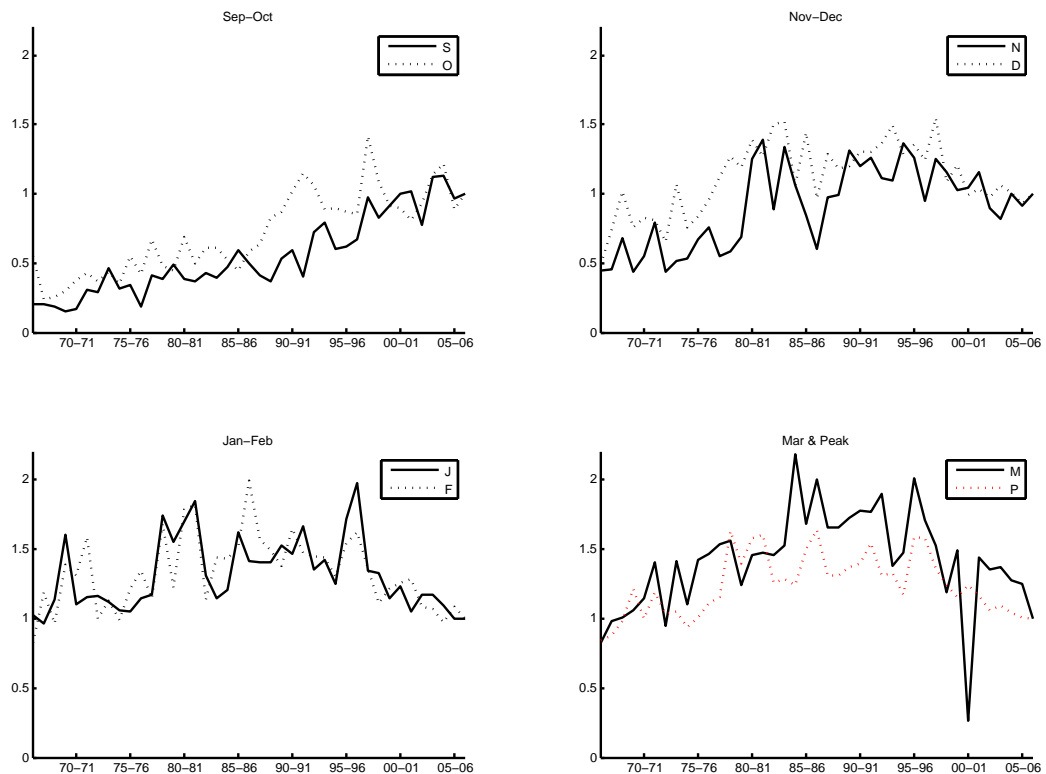


Figure 5.3: Underhill indices for Shelduck based on one month only as indicated. Peak based on imputing using the Underhill index for September to March and then using data from the month with the maximum number of birds present each year. Data for 111 sites chosen by the top sites criterion for September to March.

peak number of birds present is decreasing (Figure 5.3). Consequently, a WeBS alerts-style analysis using data from September to March will result in a medium alert for the medium-term, whereas if data from January and February only is used high alerts will be issued for both the medium- and long-term (Table 5.1).

Chapter 3 considered the problem of selecting subsets of months on which to construct an index and the example of Shelduck shows that deciding to limit the months can have consequences on trend reporting, potentially affecting the conservation policy for a species. As discussed in Chapter 3, some WeBS analyses use all the months from September to March whereas others use “stable” winter months that are designed to exclude passage periods, as little is known about site turnover of birds in these periods. In effect, the winter months (e.g. specified by one of the methods in Table 3.2 on page 127) are

given a weight of 1, and the other months a weight of 0: birds present in mid-winter are worth something in the index, whereas passage birds are not. In this chapter, alternative weighting schemes to this binary system are considered, so for example a bird present in mid-winter is still worth more than a bird in a passage period, but migrating birds do still contribute to the index to a lesser degree.

5.3 Types of index

5.3.1 Introduction

There are at least eight house price indices that are commonly reported in the British media (Briscoe, 2004). Some are based on all sales, some on repeat sales, and some on mortgages from a particular lender (which are often spatially biased to one part of the country); others are based on asking prices or surveyors' opinions on price movements. Seasonality has resulted in a proliferation of comparisons, with (annualised) monthly, (annualised) quarterly and yearly inflation all being reported, usually depending on the argument the writer wishes to make in the report. The differences in their methodologies mean that they never agree exactly and this in turn can lead the general public to be dismissive of them: informed that house prices increased by 0.7% in September 2007 in a newspaper report on one index and then one week later in the same newspaper that they decreased by 0.6% in September according to another index, the reader is understandably confused (O'Grady, 2007a,b).

The precise methodology used also has a large impact on general price indices. For example, in the U.K. two inflation indices are often referred to: the consumer price index, CPI, and the retail price index, RPI. Price changes in a basket of goods representing "average" household expenditure are tracked. Goods can exhibit price changes in differing directions and magnitude: for example from September 2006 to September 2007 the cost of renting a house increased by 3% and the cost of dairy products by 10%, whereas the price of photographic, cinematographic and optical equipment decreased by 28%. For this period the CPI was 1.8% and the RPI was 3.9% (National Statistics, 2007a).

The difference between the indices is due to both the make-up of the basket of goods and the index formulae used. The RPI includes some housing costs and taxes which are not used to compile the CPI, but omits some educational and financial product costs which the CPI includes. Both CPI and RPI are Laspeyres-type indices (see Section 5.3.3 below) but individual sample prices from different outlets are combined into elementary

aggregates by a geometric mean when calculating the CPI and an arithmetic mean for the RPI (National Statistics, 2007b).

The heterogeneity of house prices and goods and services means that it is not the case that one index is correct and the others not, but rather indices may be legitimately defined differently. The concept of how much more expensive properties are this year compared to last year, or how many more birds there are this year than last year seems simple but the impact of seasonal data from multiple sites makes the answer complicated. Section 5.2.4 gave a WeBS example where the choice of index leads to different conclusions.

For simplicity, the index formulae defined in this section for combining WeBS monthly counts assume that multiple site counts are summed to give a national total for each time point, $x_{jk} = \sum_i x_{ijk}$, where i is the site, j is the year and k is the month. Other methods of aggregating the site counts may be useful, as discussed in Section 5.2.2, and the index formulae discussed in this section may also be used in future to combine site information.

We consider a count matrix, X , of size $J \times K$ where J is the number of years and K is the number of months in the data set. For each year, j , we have a vector of monthly counts x_j of length K . We investigate how to obtain a single index value, $I_j = \frac{f(x_j)}{f(x_b)}$ for year j from the K monthly counts with reference to a base year, b , using some formula f .

5.3.2 Non-weighted indices

As introduced in Section 1.3, the Ogilvie Index method was used to estimate population trends from Wetland Bird Survey data until 1992, from when the Underhill Index method has been used. The Underhill and Ogilvie methods use the same approach to combining counts from multiple sites. Although site counts are combined using a chain index approach in the Ogilvie index when there are missing data, on a complete data set the Ogilvie index reduces to the equivalent of a sum over sites.

Both the Underhill and Ogilvie indices are non-weighted indices. The indices differ on a complete set of data solely in how monthly counts are utilised: the Ogilvie index uses

a single month of data, whereas the Underhill Index uses data from a species-specific selection of months. The months used for the Underhill index are chosen so that the population may be treated as approximately stable for those months (see Section 3.5.2).

Following the notation above, the Underhill index has the formula $f(x_j) = \sum_k x_{jk}$, or equivalently, $f(x_j) = \frac{\sum_k x_{jk}}{K}$. The Ogilvie index has the formula $f(x_j) = x_{ja}$ where a is the monthly count used (usually January). The indices can be interpreted as the ratio of the total population on the sites of interest in the current year compared to the base year, where the population estimate is based on the arithmetic mean over the particular months used for that species for the Underhill index and a single month for the Ogilvie Index.

A similar index to the Underhill approach of the arithmetic mean, is to use the geometric mean of the monthly counts, so that $f(x_j) = (\prod_k x_{jk})^{1/K}$. Another possible index is a peak index, where instead of using the count from the same single month each year, as in the Ogilvie index, the month used can be varied each year such that $f(x_j) = x_{ja}$ where $x_{ja} = \max_k(x_{jk})$.

5.3.3 Weighted indices

The most active application research area in index number theory is in cost of living and inflation indices. There is a long history of weighting a selection of commodities to obtain an inflation index, stretching back to work by William Fleetwood in 1707 and developed by Joseph Lowe in 1823 (Diewert, 1993).

A weighted index for the current year, j , with respect to the base year, b , is defined as $I_j = \frac{\sum_k x_{jk} w_k}{\sum_k x_{bk} w_k}$ for some set of weights w_k . The value of goods purchased by a household can be calculated by multiplying together the price and quantity of the goods purchased. Each price index has a corresponding quantity index; the quantities are used as weights for the price index and the prices are used as weights for the quantity index. Although wildfowl counts cannot be split into “price” and “quantity”, weights may be defined in some other way.

Since consumers change their spending patterns over time, for example buying greater amounts of a commodity as a result of price reduction, when defining a weighted price index economists must decide from which time period to choose quantities to use as weights. When comparing prices in year j with those in year b , if quantity information from the base year is used to weight the prices then the index is known as a Laspeyres-type index. Alternatively, quantity information from the current year are used as weights in a Paasche-type index. If we denote the weights for a particular year y as w_{iy} then the Laspeyres Index, L_j is

$$L_j = \frac{\sum_k x_{jk} w_{bk}}{\sum_k x_{bk} w_{bk}}$$

and the Paasche Index P_j is

$$P_j = \frac{\sum_k x_{jk} w_{jk}}{\sum_k x_{bk} w_{jk}}$$

A third commonly used index is the Fisher index, F_j , which is the geometric mean of the Laspeyres and Paasche indices,

$$F_j = (L_j P_j)^{1/2} = \left[\frac{\sum_k x_{jk} w_{bk}}{\sum_k x_{bk} w_{bk}} \frac{\sum_k x_{jk} w_{jk}}{\sum_k x_{bk} w_{jk}} \right]^{1/2}.$$

Diewert (1993) summarises various tests that have been suggested for evaluating the effectiveness of index numbers. The Fisher index is known as *Fisher's ideal index* as it satisfies three tests that Fisher considered important (namely, commodity reversal, factor reversal and time reversal). However, the test approach to index number research has been controversial, as it has been shown that no index number formula can satisfy all the proposed tests and there is no general agreement as to which tests to use. Many of these tests have importance specific to economic price indices: for example, the factor reversal test that the Fisher index satisfies is that the price index multiplied by the corresponding quantity index should equal the total change in value from the base period. In a weighted biological population index the product of the weight and the count is not a straightforward concept such as "value" and there is no "quantity" index so this test is not relevant.

The Fisher index is an example of a superlative price index as it uses weights from both the current and base period. There are other superlative index formulae including the Walsh index and the Törnqvist index but these three formulae give similar numerical results and so only the Fisher index will be considered here (Hill, 2006b). In practice in economics, up-to-date data on quantity is not usually as readily available as data on price. Laspeyres indices are often used as it is only necessary to have quantity information for the base year. However, with modern technology it is becoming feasible to collect information on quantities and thus move to a Paasche or Fisher type index (Hill, 2006a).

In the next section we shall consider choosing weights independently of the count data and then in Section 5.3.5 the count data are used to determine weights for each year from which we can construct Laspeyres-type, Paasche-type and Fisher-type indices.

5.3.4 *Data-independent weights*

Section 5.2.4 introduced the idea of a WeBS population index where the contribution of counts in individual months are not all equal. By specifying a constant weight vector chosen independently of the data, more importance can be attached to counts in some months than others. For the Wetland Bird Survey for example, it may be appropriate to construct an index which weights mid-winter counts more than the spring and autumn migration months, as the over-wintering population are the individuals we are most interested in. A set of counts that do not follow a prescribed seasonal pattern are penalised so that if the total number of birds recorded is the same in two years, then the year which more closely matches the distribution of weights will have a higher index than the other.

It is worth noting the effect that the choice of base year has on the index values under this scenario. In the simple non-weighted case, the population in the base year determines the magnitude of the index in the other years. When constant weights are used the magnitude of the index values for all years depends on the seasonal pattern as well as the population size in the base year, as shown by the following example.

Example

Consider three years of seven WeBS monthly counts, where identical counts are recorded in each year, but the month in which they are recorded changes. The observed count matrix is,

$$X = \begin{pmatrix} 5 & 7 & 5 \\ 10 & 10 & 10 \\ 150 & 10 & 10 \\ 10 & 10 & 150 \\ 10 & 150 & 10 \\ 7 & 10 & 10 \\ 10 & 5 & 7 \end{pmatrix}$$

and the weight vector is arbitrarily defined as $w = [0.075 \ 0.1 \ 0.2 \ 0.25 \ 0.2 \ 0.1 \ 0.075]$.

Since only the order of the recorded counts changes, the non-weighted arithmetic mean index is equal to 1 for all three years, regardless of the base year chosen. We obtain the following indices, when the given year is used as the base year:

Base Year	Index Value		
	Year 1	Year 2	Year 3
1	1.000	1.002	1.190
2	0.998	1.000	1.187
3	0.841	0.842	1.000

The order of the indices is determined by how closely the seasonal distribution matches the weight vector: the third year always has the highest index because the peak count of 150 occurs in the fourth month, which is the month with the highest weighting.

5.3.5 Data-dependent weights

In Chapter 3, species' seasonality over time was used to limit the months used for indexing to the stable mid-winter months. In a similar way, the seasonal patterns can be used to give less weight to passage periods, when fewer birds are present, than mid-winter when the population is high. If the seasonal pattern is assumed to be stationary then a constant set of weights can be defined; alternatively if phenological change is known to be present then seasonal weights can be allowed to change over time.

Constant weights

Rather than specifying a weight vector independently, the count data set can be used to define a constant weight vector. For example, all years could be used to specify a set of weights, $w_k = \frac{\sum_j x_{jk}}{\sum_{j,k} x_{jk}}$. The index using these weights would be greater in years when the yearly seasonal pattern closely matched the long-term average.

Annual weights

We can assign weights to each individual year rather than keeping them constant over time, using information about the phenology within the data, so that for each year j we define a weight vector,

$$w_{jk} = \frac{x_{jk}}{\sum_k x_{jk}}.$$

Recall from Section 5.3.3 that the Laspeyres index, L_j uses weights from the base year to compare the base year and the current year, the Paasche index, P_j uses the weights from the current year and the Fisher index, F_j is the geometric mean of the Laspeyres and Paasche indices. Substituting the weight equation defined by the seasonal pattern in the current year gives:

$$L_j = \frac{\sum_k x_{jk} w_{bk}}{\sum_k x_{bk} w_{bk}} = \frac{\sum_k x_{bk} x_{jk}}{\sum_k x_{bk}^2},$$

$$P_j = \frac{\sum_k x_{jk} w_{jk}}{\sum_k x_{bk} w_{jk}} = \frac{\sum_k x_{jk}^2}{\sum_k x_{bk} x_{jk}}, \text{ and}$$

$$F_j = \left(\frac{\sum_k x_{bk} x_{jk}}{\sum_k x_{bk}^2} \frac{\sum_k x_{jk}^2}{\sum_k x_{bk} x_{jk}} \right)^{1/2} = \left(\frac{\sum_k x_{jk}^2}{\sum_k x_{bk}^2} \right)^{1/2}.$$

No separate source of information about seasonality is available so it is necessary to use the data to obtain annual weights; this is apparent in the Fisher index which simplifies to the square root of the ratio of the sum of the squares of the counts in the current year compared to the equivalent in the base year.

Moving average weights

The movements of wildfowl populations are highly weather-dependent and can exhibit large swings from year to year due to late or early arrival at wintering grounds or temporary migration to escape freezing weather which inhibits feeding. However, over a number of years a seasonal pattern emerges for each species. A moving average approach to specifying annual weights will have a smoothing effect on this type of weather-influenced noise, whilst still allowing for long-term phenological changes.

As indices are published yearly it is more practical to use a moving average of previous years rather than one that uses data both forwards and backwards in time. For example, weights based on a five year moving average could be defined,

$$w_{yk} = \begin{cases} \frac{\sum_{j=y-4}^y x_{jk}}{\sum_{j=y-4}^y \sum_k x_{jk}} & \text{if } y \geq 5 \\ w_{5k} & \text{if } y < 5. \end{cases}$$

5.3.6 Chain indices

When a time series of more than two index numbers is required Forsyth and Fowler (1981) and Silver (1984) argue that the chain approach to indexing is more appropriate than binary comparisons in the case of economic price indices. The premise of a chain index is that the index of a price change from the base year to the current year should be path-dependent: it should depend on price movements in the intervening years.

The chain index I_c comparing the current year, c , with the base year b is found by multiplying the indices for each intervening pair of years,

$$I_c = \begin{cases} \prod_{j=b}^{c-1} \frac{f(x_{j+1})}{f(x_j)} & \text{if } c > b \\ 1 & \text{if } c = b \\ \left(\prod_{j=c}^{b-1} \frac{f(x_{j+1})}{f(x_j)} \right)^{-1} & \text{if } c < b, \end{cases} \quad (5.2)$$

where f denotes the index aggregation formula used.

A chain index may be appropriate for wildfowl population indices, particularly when phenological changes are taking place. For example, consider a species whose peak population in Britain was gradually getting later from mid-October to mid-November due to climate change. Suppose that exactly the same counts for each month were recorded thirty years apart and that the observed peak was in October. This was normal for the species in the earlier case, but thirty years later the peak is normally in November. With a direct binary index then the index numbers for both years would be identical (assuming that moving-average weights were not being used). With a path-dependent chain index, however, the index numbers for the two years would be different because of phenological changes in the intervening period.

In some cases the chain index is identical to the non-chain index. For example, if a

constant set of weights is used then chaining has no effect:

$$I_c = \frac{\sum_k w_k x_{ck}}{\sum_k w_k x_{(c-1)k}} \frac{\sum_k w_k x_{(c-1)k}}{\sum_k w_k x_{(c-2)k}} \cdots \frac{\sum_k w_k x_{(b-1)k}}{\sum_k w_k x_{bk}}.$$

Similarly, the Fisher index in Section 5.3.5 based on weights derived from the current year also simplifies to the binary index:

$$\begin{aligned} F_c &= \left(\frac{\sum_k x_{ck}^2}{\sum_k x_{(c-1)k}^2} \right)^{1/2} \left(\frac{\sum_k x_{(c-1)k}^2}{\sum_k x_{(c-2)k}^2} \right)^{1/2} \cdots \left(\frac{\sum_k x_{(b-1)k}^2}{\sum_k x_{bk}^2} \right)^{1/2} \\ &= \left(\frac{\sum_k x_{ck}^2}{\sum_k x_{(c-1)k}^2} \frac{\sum_k x_{(c-1)k}^2}{\sum_k x_{(c-2)k}^2} \cdots \frac{\sum_k x_{(b-1)k}^2}{\sum_k x_{bk}^2} \right)^{1/2}. \end{aligned}$$

A much discussed aspect of chain indices is their liability to drift. Forsyth and Fowler (1981) show that the price of not exhibiting drift in the binary indices is loss of representativity and that choosing the aggregation formula appropriately will minimise drift. They warn that where price oscillation is present particular care must be taken in choosing the aggregation formula. The difference or spread between the Paasche and Laspeyres indices is a measure of the sensitivity to the choice of index formula (Hill, 2001). Hill (2006a) showed that chaining does not always reduce the Paasche-Laspeyres spread, even when price oscillation is not present.

5.4 Comparison of Index Performance

5.4.1 Summary of example indices

In the following section the different types of index defined in Section 5.3 will be compared for some example wildfowl observed data from the Wetland Bird Survey. The index formulae we will be considering can be categorised as binary (non-weighted) indices, weighted indices and chain indices. The proposed indices are summarised in Table 5.2 and calculated using the code in Program 5.1.

Description	Weight	Formula	Type
Arithmetic mean ¹	1	$\sum_k x_{jk}$	binary
Geometric mean ¹	1	$(\prod_k x_{jk})^{1/K}$	binary
Weighted mean ²	$\frac{\sum_j x_{jk}}{\sum_{j,k} x_{jk}}$	$\sum_k w_k x_{jk}$	constant
Peak ³	1	x_{ja}	binary
One month ⁴	1	x_{ja}	binary
Annual seasonal weights ⁵	$\frac{x_{jk}}{\sum_k x_{jk}}$	$\sum_k w_{bk} x_{jk}$	Laspeyres
	$\frac{x_{jk}}{\sum_k x_{jk}}$	$\sum_k w_{jk} x_{jk}$	Paasche
	$\frac{x_{jk}}{\sum_k x_{jk}}$	$\left(\sum_k w_{bk} x_{jk} \sum_k w_{jk} x_{jk} \right)^{1/2}$	Fisher
5 year seasonal weights ⁵	$\frac{\sum_{j=y-4}^y x_{jk}}{\sum_{j=y-4}^y \sum_k x_{jk}}$	$\sum_k w_{bk} x_{jk}$	Laspeyres
	$\frac{\sum_{j=y-4}^y x_{jk}}{\sum_{j=y-4}^y \sum_k x_{jk}}$	$\sum_k w_{jk} x_{jk}$	Paasche
	$\frac{\sum_{j=y-4}^y x_{jk}}{\sum_{j=y-4}^y \sum_k x_{jk}}$	$\left(\sum_k w_{bk} x_{jk} \sum_k w_{jk} x_{jk} \right)^{1/2}$	Fisher

¹ Possibly $k \in \Phi$ where Φ is some subset of months.

² Alternatively the weights may be specified in advance.

³ Where $x_{ja} = \max_k(x_{jk})$.

⁴ a is specified in advance (Ogilvie index) or where $\sum_j x_{ja} = \max_k(\sum_j x_{jk})$.

⁵ Also chain version, see equation 5.2.

Table 5.2: Summary table of prospective indices for seasonal wildfowl population counts. Binary type indices give equal weight to the months used but exclude other months. Constant type indices use the same weights for all years. Laspeyres type indices use the weight from the base year and Paasche type indices the weight from year j ; the geometric mean of these gives Fisher-type indices.

Program 5.1 MATLAB code to calculate indices.

```

%Indices are calculated for the nmonths*nyears array "counts". This
%should be in the workspace, and will usually consist of a set of WeBS
%counts, summed over all sites. All indices use a base year, either set in
%the workspace as basey or set to the final year.
[nmonths,nyears] = size(counts);
if exist('basey') == 0, basey = nyears; end

arithmean = mean(counts); arithmean=arithmean/arithmean(end);
geomeanind = geomean(counts); geomeanind=geomeanind/geomeanind(end);
wt = mean(counts')/max(mean(counts'));
wtarithmean = wt*counts; wtarithmean = wtarithmean/wtarithmean(basey);
peak = max(counts); peak = peak/peak(end);
[o1 o2]=find(sum(counts')==max(sum(counts')));
onemonth = counts(o2,:)/counts(o2,end);

%ANNUAL WEIGHTS
%Calculate quantities needed to compute indices:
qcur = counts./repmat(sum(counts),nmonths,1); pqmat = counts'*qcur;
pkk1 = zeros(nyears-1,1); %P_{k,k+1} for k = 1,2,...
for i = 1:(nyears-1), pkk1(i)=pqmat((i+1),i)/pqmat(i,i); end
pkk2 = zeros(nyears-1,1); %P_{k,k+1} for k = 1,2,... with Paasche def.
for i = 1:(nyears-1), pkk2(i)=pqmat((i+1),(i+1))/pqmat(i,(i+1)); end
Laspeyres = (pqmat(:,basey)/pqmat(basey,basey))';
for i=1:nyears
    Paasche(i) = ((counts(:,i))*qcur(:,i))/((counts(:,basey))*qcur(:,i));
end
Fisher = geomean([Laspeyres;Paasche]);
Laschain = zeros(nyears,1); Laschain(basey)=1;
for i = (basey+1):nyears Laschain(i)=prod(pkk1(basey:(i-1))); end
for i=1:(basey-1), Laschain(i)=1/prod(pkk1(i:basey-1)); end
Pachain = zeros(nyears,1); Pachain(basey)=1;
for i = (basey+1):nyears, Pachain(i)=prod(pkk2(basey:(i-1))); end
for i=1:(basey-1), Pachain(i)=1/prod(pkk2(i:basey-1)); end
Fishchain = geomean([Pachain; Laschain]);

%5 YEAR WEIGHTS
malength = 5; movingweight = zeros(nmonths,nyears);
for i=1:nmonths, for j = malength:nyears
    n1 = sum(counts(i,j-(malength-1):j));
    d1 = sum(sum(counts(:,j-(malength-1):j)));
    movingweight(i,j) = n1/d1; end
for j = 1:malength-1, movingweight(i,j)=movingweight(i,malength); end end
pqmat = counts'*movingweight;
pkk1 = zeros(nyears-1,1);
for i = 1:(nyears-1),pkk1(i)=pqmat((i+1),i)/pqmat(i,i); end
pkk2 = zeros(nyears-1,1);
for i = 1:(nyears-1),pkk2(i)=pqmat((i+1),(i+1))/pqmat(i,(i+1)); end
winLaspeyres = (pqmat(:,basey)/pqmat(basey,basey))';
for i=1:nyears, winPaasche(i) = ((counts(:,i))*movingweight(:,i))/((counts(:,
    basey))*movingweight(:,i)); end
winFisher = geomean([winLaspeyres;winPaasche]);
winLaschain = zeros(nyears,1); winLaschain(basey)=1;
for i = (basey+1):nyears, winLaschain(i)=prod(pkk1(basey:(i-1)));end
for i=1:(basey-1), winLaschain(i)=1/prod(pkk1(i:basey-1)); end
winPachain = zeros(nyears,1); winPachain(basey)=1;
for i = (basey+1):nyears, winPachain(i)=prod(pkk2(basey:(i-1))); end
for i=1:(basey-1), winPachain(i)=1/prod(pkk2(i:basey-1)); end
winFishchain = geomean([winPachain; winLaschain]);

```

5.4.2 Seasonal indices for Shoveler and Shelduck

Section 5.2.4 showed that restricting the months used in the calculation of Shelduck indices to the mid-winter months resulted in significantly different conclusions to when all the months were used due to an increase in the number of birds present in September, October and November. Figure 5.5 compares indices using the month aggregation formulae from Table 5.2. The contrast between the increasing trend in the autumn months with the other months that could be seen in Figure 5.3 on page 182 can be seen in the difference between the arithmetic mean over all the months and the indices based on the peak each year and the overall peak month (in this case January), particularly in the period up to the mid-90s.

In Section 3.4 it was shown that Shoveler phenology has also changed over the WeBS survey period. Figure 5.4 shows that the biggest different in trends in individual months in Shoveler was in the 1980s, where the indices based on December, January, February and March were lower than those based on September, October and November.




Figure 5.4: Underhill indices for Shoveler based on one month only as indicated. Data for 118 sites chosen by the top sites criterion for September to March.

5.6 shows that the gap between the indices based on averaging over all the months and the peak and the overall peak month (in this case October) indices is largest in this period.

In both Figures 5.5 and 5.6 a tendency to drift can be seen in the Laspeyres and Paasche chain indices when using annual seasonal weights; the further from the base year the more the drifting effect, so that the indices for the late 1960s are dissimilar to both the average-based indices and the peak-based indices. However, using moving average seasonal weights based on data from five years removes this undesirable effect.

Due to the construction of the Fisher indices with seasonal weights derived from the data, the Fisher indices are very similar to the arithmetic mean index (see Section 5.3.5). In fact it can be seen that most of the Fisher-, Laspeyres- and Paasche-type indices give similar results to the arithmetic mean. The index which is most promising as a compromise between the arithmetic mean of all the months and the peak index is the 5-year Laspeyres chain index. Recall that this index weights each count using a weight calculated using five-year moving average for the base year (so when the base year is the final year in the index, the weights are the average seasonal pattern over the most recent five years). Each index is calculated by multiplying the index for the intervening years between it and the base year.

Figure 5.7 compares the arithmetic mean index, the peak index and the 5-year Laspeyres chain index for the Shelduck and Shoveler, and two other species which do not necessarily show phenological change, Mute Swan and Ruddy Duck. The index behaves well in balancing the pattern in the overall mean with the pattern in the peak number of birds present.

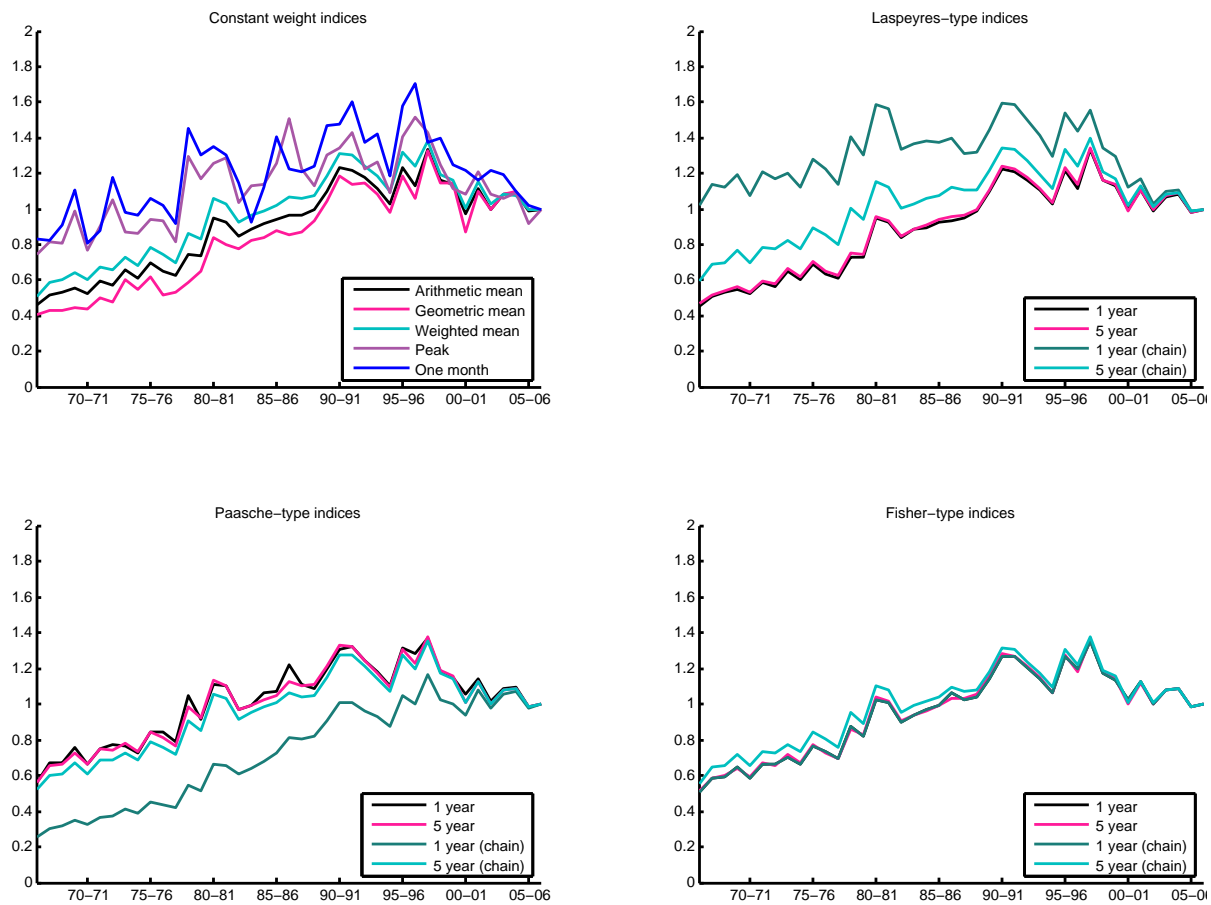


Figure 5.5: Shelduck indices using varying month aggregation formulae (see Table 5.2).



Figure 5.6: Shoveler indices using varying month aggregation formulae (see Table 5.2).

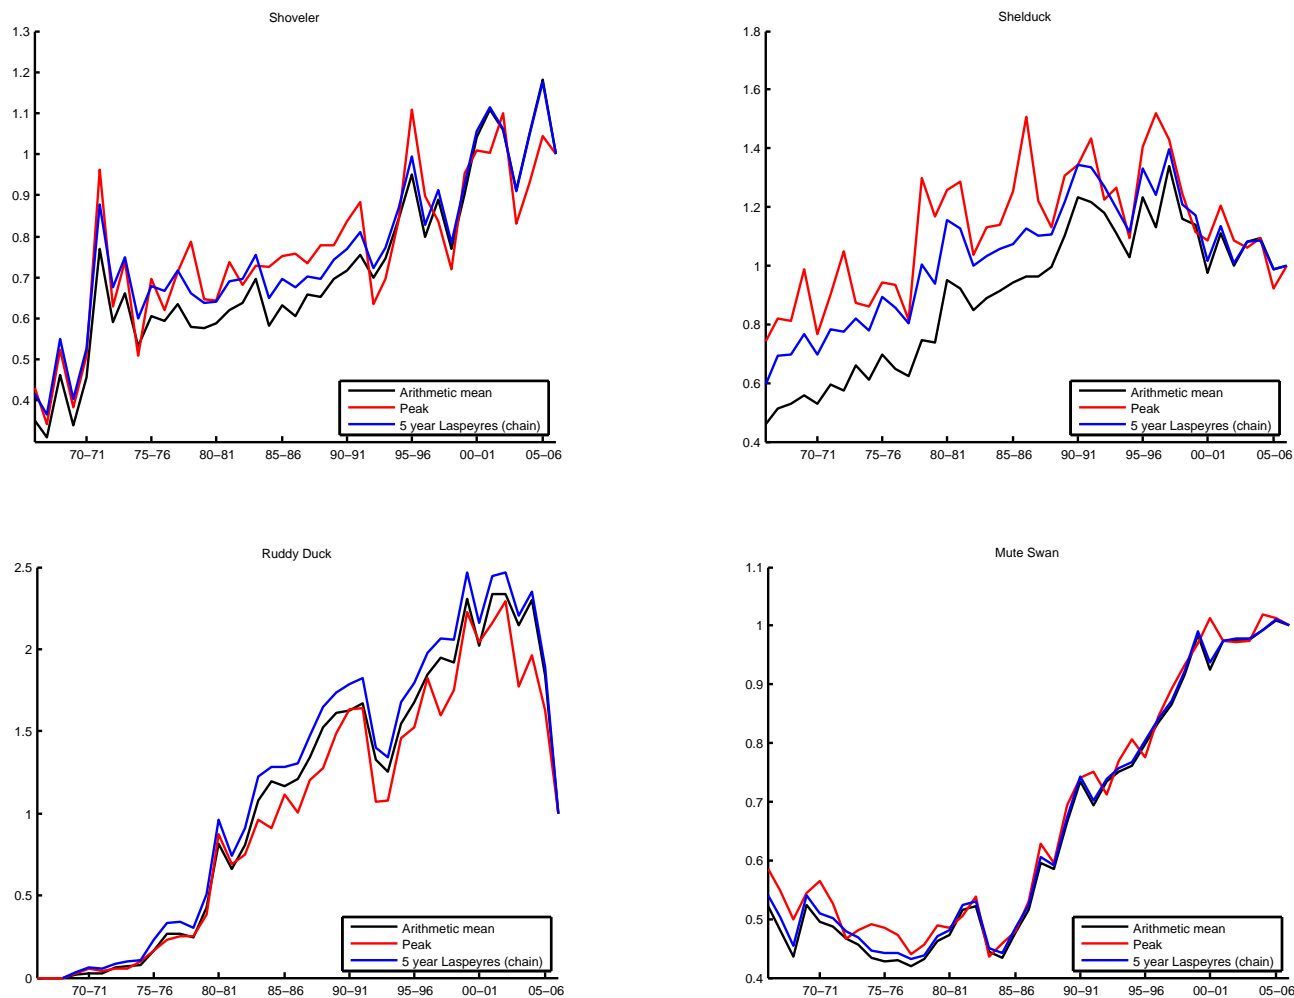


Figure 5.7: Laspeyres indices. Comparison of Laspeyres indices with the arithmetic mean for counts from September to March for Shoveler, Shelduck, Ruddy Duck and Mute Swan.

5.5 Discussion

This chapter began with a discussion of what an index number is. An index based on WeBS data must aggregate data from both sites and months. Section 5.2.2 followed on from Chapter 2 in discussing how a site aggregation formula could weight sites of different habitats or in different locations so that the index could reflect the country as a whole. There are interesting parallels with the problem of devising house price indices, as houses vary in price by location, type. House prices and wildfowl numbers also vary seasonally and Section 5.2.3 introduced combining monthly counts to obtain a single index number. In Section 5.2.4 we saw that phenological change can result in different conclusions about the extent of a decline in a species, and hence the amount of concern about the decline.

In Chapter 3 the approach to seasonality in index numbers was to restrict the months used to those where the population is not changing very much and omitting the other months from the index completely. Section 5.3 suggested aggregation formulae for monthly data that used data from all the months, but emphasised those months where the population is larger through weighting schemes. Ideas from economic index theory were used to generate possible ways of using seasonal information to weight the counts. The proposed aggregation formulae included the arithmetic mean of subsets of months, which is the approach in the Underhill index, using the monthly maximum for each year to give a peak index, and an index where counts were weighted by the average seasonality across all years. Weights based on the seasonality in the base year (Laspeyres-type indices), the year being indexed (Paasche-type indices) and the geometric mean of these (Fisher-type indices) were suggested. It was proposed to use weights based on the seasonal pattern within the data using one year only or an average over the previous five years. Chain indices were also proposed, where the index depends not only on the counts in the current year and the base year, but also on the indices for the year in between.

Section 5.4.1 summarised the proposed formulae in Table 5.2 and the resultant indices

for some example species using WeBS data were presented in Section 5.4.2. If the desire is a compromise between the maximum number of birds present and the average number of birds present then the 5-year Laspeyres chain index gives the best results. This uses the most recent five years of data to obtain seasonal weights, but then uses chaining so that phenological change in the intervening years is taken into account. In the following chapter, we will return to comparing the results of different month aggregation formulae for indices by using simulated data, allowing us to compare how different index formulae may be more appropriate under different biological assumptions.

Chapter 6

WETLAND SYSTEM SIMULATION

6.1 Introduction

Simulation is a valuable tool for comparing the efficacy, advantages and disadvantages of different statistical models. Computer simulations can be used to construct biologically plausible wetland systems, where birds enter into the system at a particular site and then move between sites throughout the season before exiting the system (through death or outward migration). At any time the locations of each individual are known and observer counts can be simulated. In this chapter, stochastic wetland system simulations are constructed and used to compare methods from previous chapters to illustrate index estimation for Wetland Bird Survey data.

6.2 Simulation Structure

6.2.1 Overview

The wetland system simulation distributes a population of N birds among M sites, with birds arriving into the system, moving around the sites and then departing the system according to random variables generated from specified probability distributions. Figure 6.1 illustrates the basic structure of the simulation model, where movement events and their times are generated and then measurement error is introduced to simulate WeBS counts. The basic procedures for generating the times at which a bird considers moving and the sites that the bird moves to at those times are given in Figure 6.2.

Many wildfowl species exhibit some form of flocking behaviour and often travel in groups. In some species, for example Bewick's Swan, family groups remain together throughout the overwintering season whereas in other species it is likely that although an individual is always in a flock, the members of that flock changes over the winter period. The simulation model can include flocking behaviour, by using the flock as the base unit rather than the individual, with the size of the flock being an additional random variable.

6.2.2 Example simulation

To illustrate the simulation model, let us consider a simple example, for N birds in a system comprising M sites. In this example the birds behave independently, that is they do not exhibit flocking behaviour.

The arrival times for the individuals into the system are assumed to be independent exponential variables, so that bird i arrives at time t_{i1} , where

$$t_{i1} \sim \text{expo}(\alpha), i = 1, \dots, N.$$

The sites at which the birds arrive are also independently distributed, according to a

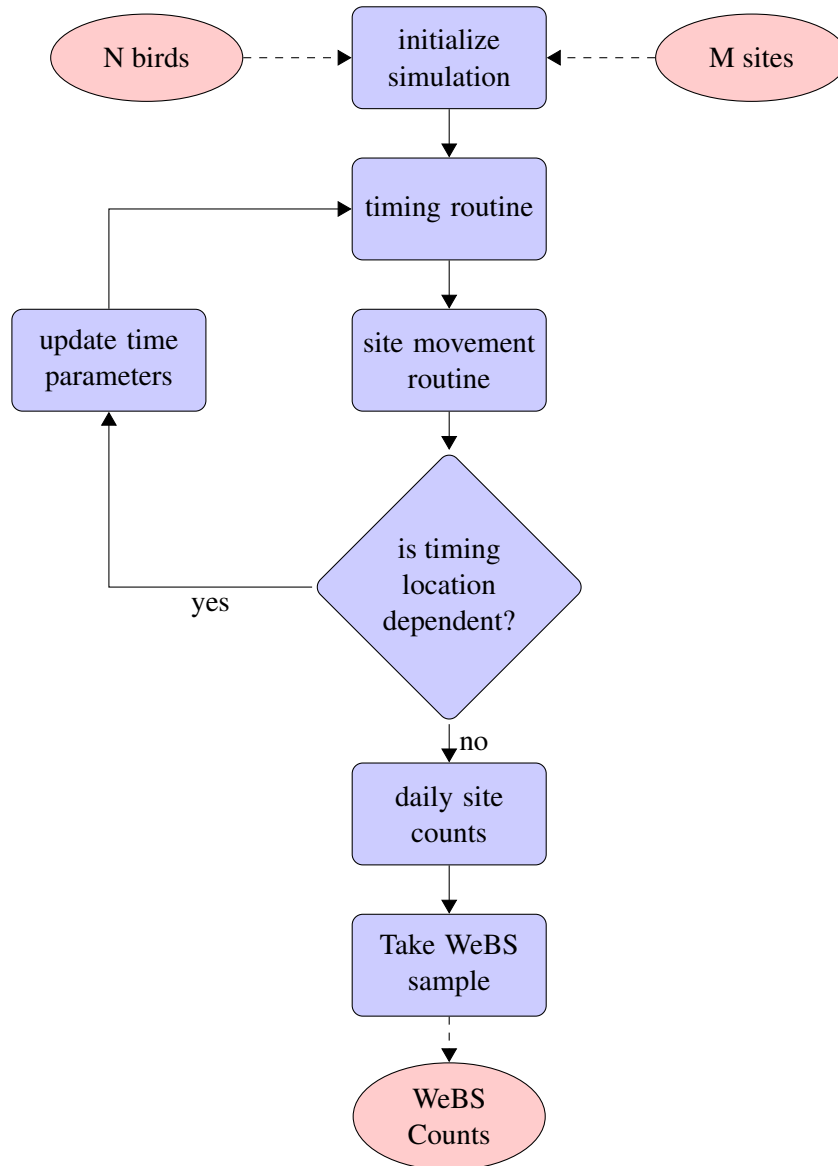


Figure 6.1: Flow diagram of generic wetland system simulation

multinomial distribution, so that bird i arrives at site s_{i1} , where

$$s_{i1} \sim \text{mult}(1, \mathbf{p}_{i1})$$

and the vector of probabilities \mathbf{p}_{i1} are defined by a set of site weights,

$$p_{i1k} = w_k, k = 1, \dots, M.$$

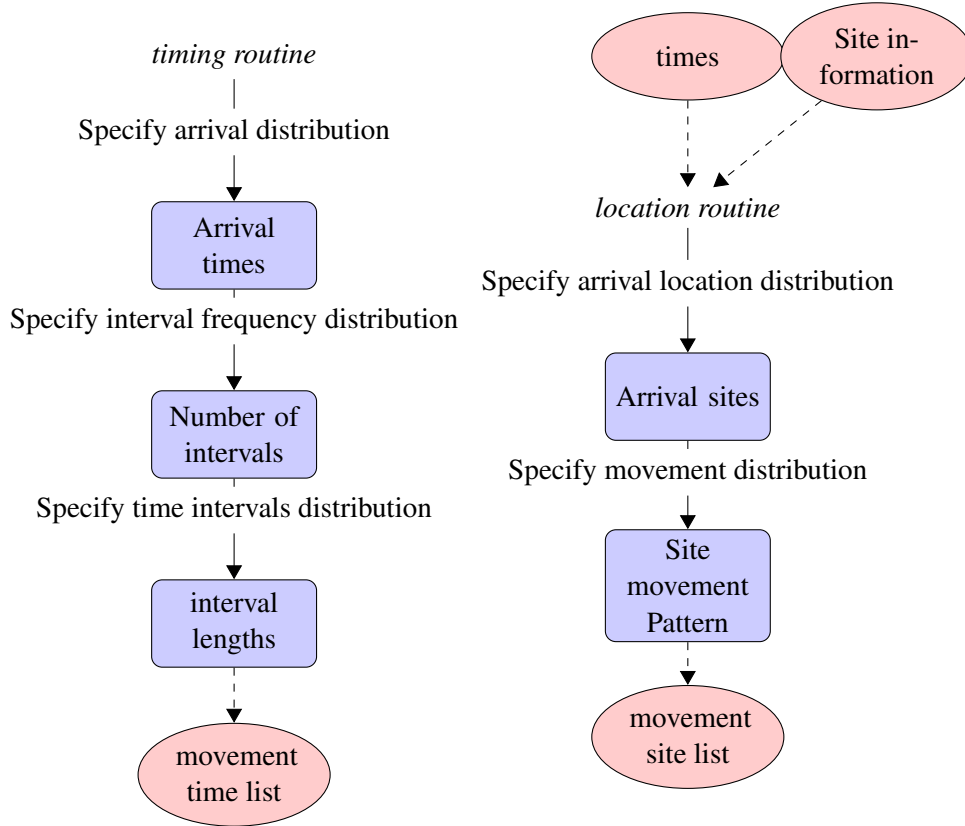


Figure 6.2: Flow diagram of timing and location routines

The number of occasions T_i when bird i becomes restless and considers moving to another site is assumed to follow a discrete uniform distribution with a specified minimum β_1 and maximum β_2 , so that

$$T_i \sim \text{UniD}(\beta_1, \beta_2).$$

The times $t_{i,j+1}$ at which bird i considers moving are calculated using a gamma distribution to simulate the lengths of the time intervals τ_{ij} . That is,

$$t_{i,j+1} = t_{ij} + \tau_{ij},$$

where the lengths of the time intervals between events are given by

$$\tau_{ij} \sim \text{gamma}(\gamma_1, \gamma_2), j = 1, \dots, T_i.$$

At time t_{ij+1} bird i chooses location s_{ij+1} according to a multinomial distribution where the probabilities have been modified by an inertia factor, ξ , which means that a bird would rather stay where it is than go to the bother of moving sites, so that

$$s_{ij+1} \sim \text{mult}(1, \mathbf{p}_{ij+1}),$$

where the vector of probabilities \mathbf{p}_{ij+1} are defined by a set of site weights,

$$p_{ij+1k} = \begin{cases} \xi w_k & \text{if } k = s_{ij} \\ (1 - \xi)w_k & \text{otherwise.} \end{cases}$$

On the final occasion when the bird becomes restless, T_i , the individual departs the system.

Figure 6.3 shows the results of this simulation using the parameters specified in Table 6.1. Note that as no measurement error has been included in this simulation, the WeBS counts are identical to the actual number of birds present.

N	1000
M	4
α	20
γ_1	3.5
γ_2	2
w	[0.15 0.2 0.25 0.4]
β_1	15
β_2	20
ξ	0.5

Table 6.1: Parameters for simple simulation example




Figure 6.3: Example results from the simple wetland system simulation in Table 6.1. The coloured lines are the numbers of birds present at each of the four sites in the simulation on each day. The solid black line indicates the total number of birds present in the system. WeBS counts are indicated by circles connected by dotted lines.

6.3 Comparing index estimation in the presence of spatial change, observation error and missing data.

6.3.1 Introduction

Data simulated as described in Section 6.2 allows experimentation with the possible effects of missing values, regional trends (perhaps caused by climate change) and precision and observation error.

The simulation movement parameters ($\alpha, \gamma_1, \gamma_2, \beta_1, \beta_2$ and ξ) for all the simulations were set the same as those in the example simulation, Table 6.1 on page 207. The movement of flocks of birds rather than individuals was simulated and the number of individuals at a site at a particular time was simulated as a negative binomial random variable with parameters corresponding to a mean of $20 \times$ the number of flocks and a variance of $40 \times$ the number of flocks.

Three example population trends were simulated: a constant population, $N = 10000$, where N is the number of flocks; an increasing population based on the WeBS index for Shoveler; and a population trend that increases from a low level, grows steadily, then abruptly starts falling, based on the index for Ruddy Duck (where the drop was caused by the commencement of a culling programme). For both these simulations $N = 10000 \times I_j$ where I_j is the population index for Shoveler and Ruddy Duck respectively.

Data was simulated for the three population trends for a set of sites, $M = 100$, assuming that all sites were equally attractive to birds and there was no change in site status over time, so the site attractiveness weight for each site was 0.01. Data was also simulated for the three population trends assuming that some sites held more flocks than others, and that this changed over time. An arbitrary set of site factors, f^* , were defined using the first 25 Shoveler site factors. Four regions each containing 25 sites with these site factors were modelled for the fifty years of the simulation. It was assumed that over time birds

* $f = [19; 23; 24; 24; 27; 30; 30; 32; 34; 39; 42; 44; 48; 49; 51; 54; 72; 77; 81; 84; 86; 91; 128; 150; 432]$.

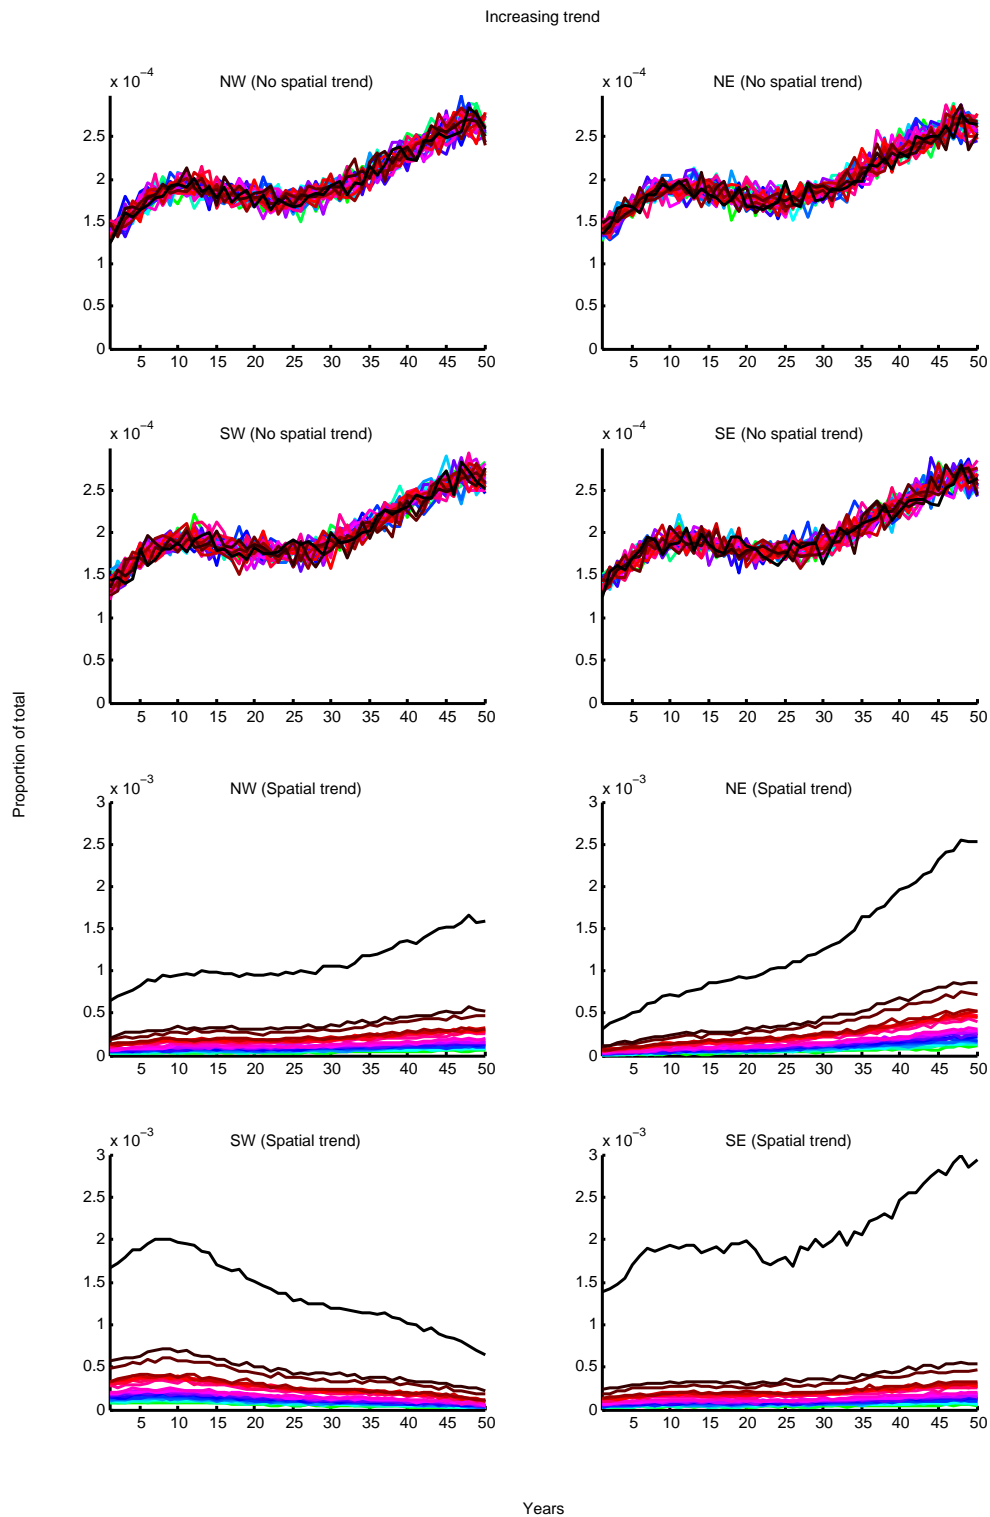


Figure 6.4: Spatial change simulations. Constant spatial distribution simulations (top four plots) and spatial change simulations (bottom four plots) over a fifty year period for an increasing population over four spatial regions. Each coloured line represents one site. For simulation details see text.

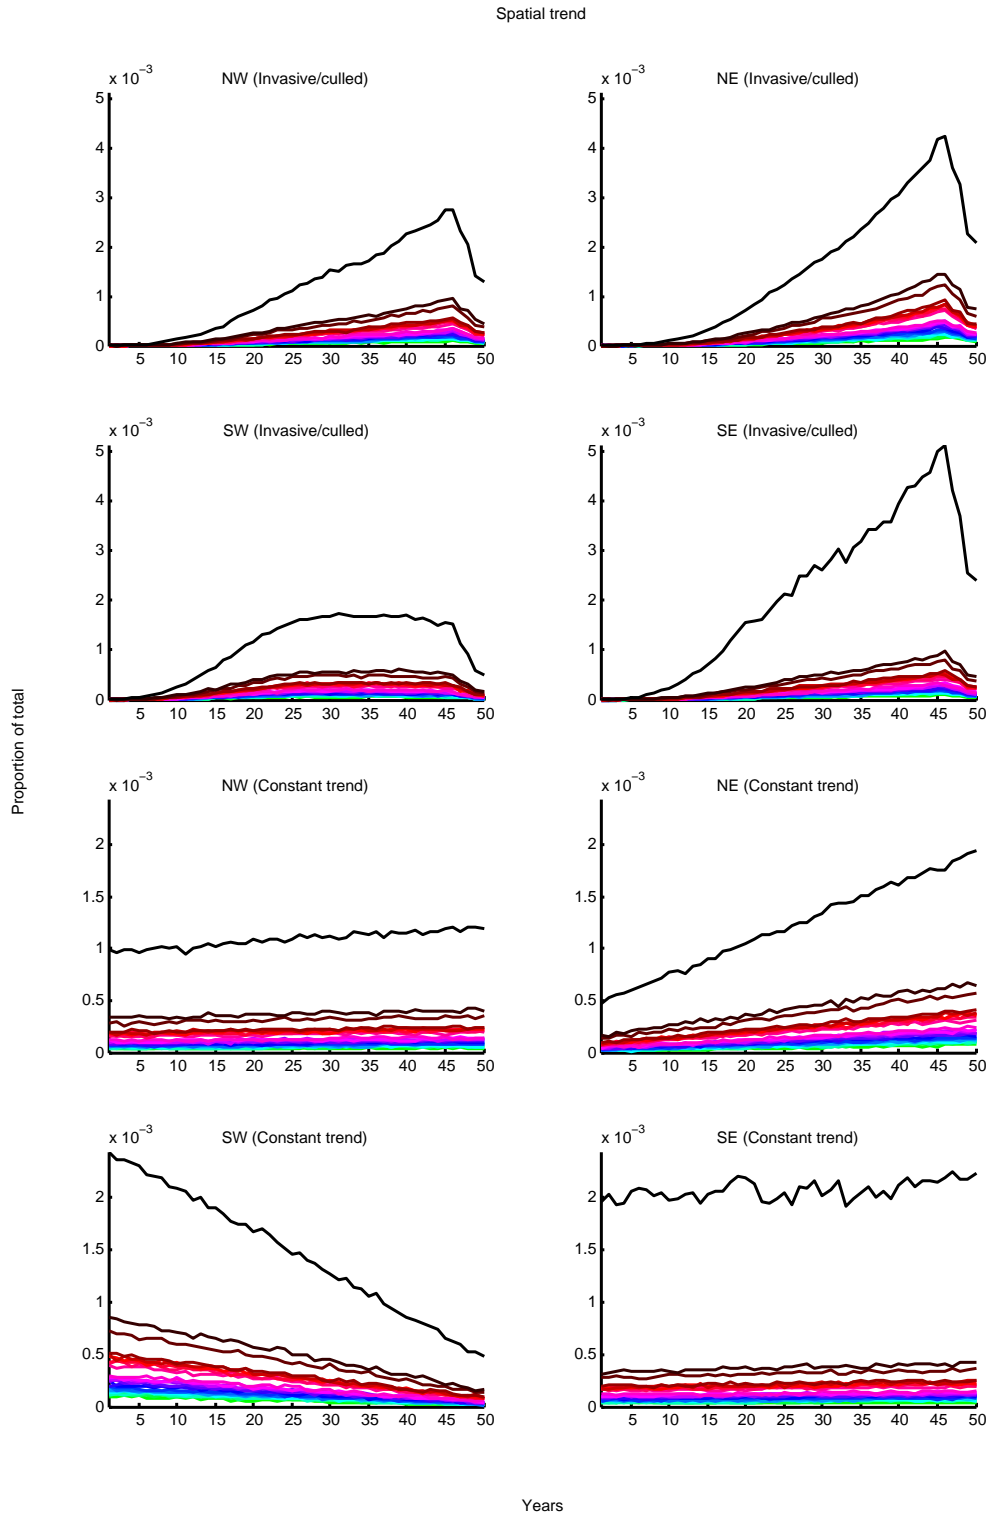


Figure 6.5: Spatial change simulations. Spatial change scenarios over a fifty year period for an increasing, then decreasing, invasive population (top four plots) and a constant population (bottom four plots) over four spatial regions. Each coloured line represents one site. For simulation details see text.

were tending to move away from the south-west (SW) region and towards the south-east (SE) and north-west (NW) regions and most particularly the north-east (NE) region, so that the SW region went from holding 50% of the population in year 1 to 10% in year 50, the SE and NW regions changed from 20% to 25% each and the population in the NE increased from 10% to 40% of the total.

The first four plots in Figure 6.4 show the site proportions of the total number of birds (calculated over each day of the simulation) plotted against time for each of the 100 sites, with 25 sites in each region when there is no spatial trend (each site is equally attractive as the others in all years). The bottom four plots in Figure 6.4 and the plots in Figure 6.5 show the site proportions where there is uneven distribution over sites and a spatial trend away from the SW region. Note that there is one site weighted much more highly than the others in each region. As the simulation is stochastic and includes site inertia, the most highly weighted site in the SE in the invasive simulation happens to contain a higher population at the peak in year 45 than the NE (Figure 6.5).

6.3.2 Results

Observation error was added to the simulated data during each run using symmetrical and unsymmetrical error (Figure 6.6). Beta errors were defined on an interval with a lower bound of 55% of the site population and an upper bound of 120% of the count and normal errors so that 95% of the values from the distribution of a population fall within 87.5% of the population. Values were rounded to the nearest integer and it was assumed that zero counts were always correct (i.e. if an individual of the species was present, at least one individual would be observed). The distribution of observed values for a site population of 100 are shown in Figure 6.6.

Since for large congregations of birds observers will usually record a rounded estimate, numbers were rounded to according to the number of birds observed (with observation error). Where this was less than 50 no rounding was applied, between 51 and 300 the

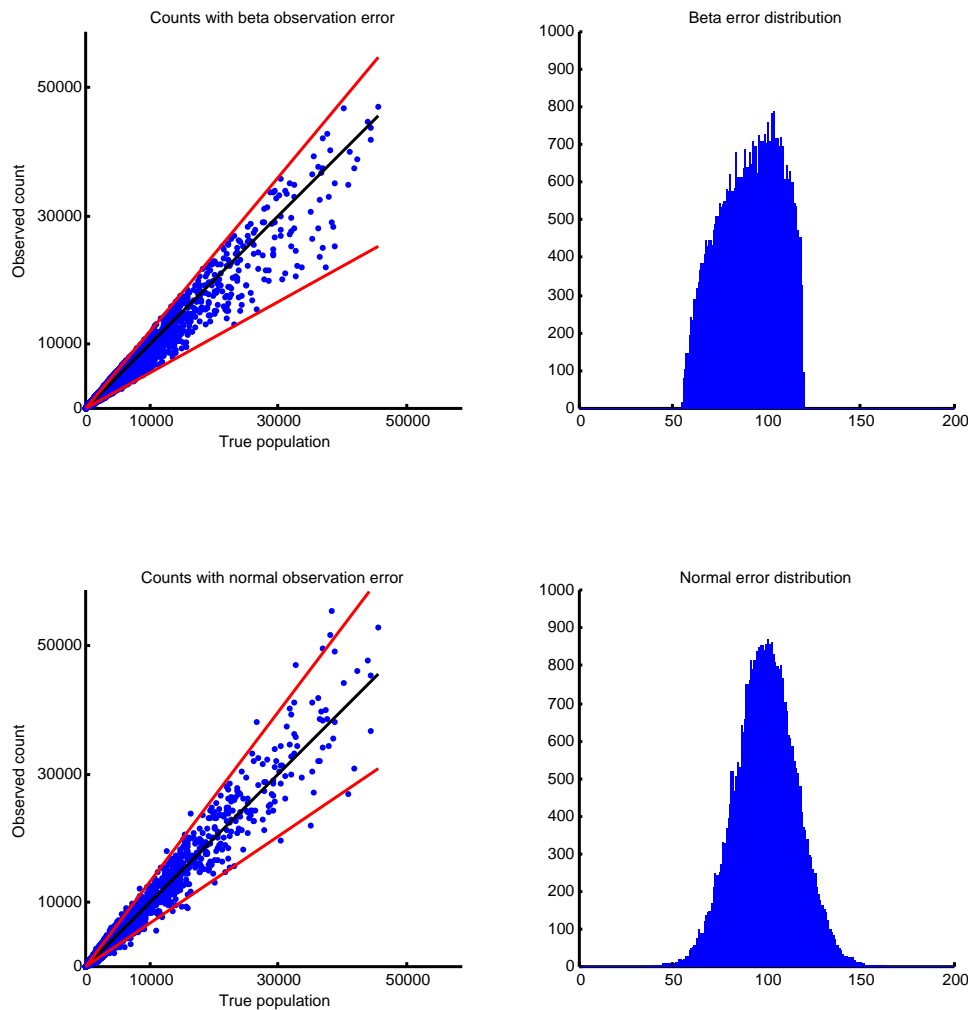


Figure 6.6: Simulating observation error for the invasive population simulation (see Figure 6.5). Left-hand plots show observations with beta or normal measurement error added. An underestimate of the true population is more likely with beta errors (top) and overestimation and underestimation are equally likely with normal errors (bottom). The black line shows where the observation and true population are equal, and the red lines the lower and upper bounds of the error distributions (for details see text). The error distributions for a count of 100 based on 35,000 samples are shown in the right-hand plots.

observation was rounded to the nearest multiple of 5, from 301 to 1000 to the nearest 50 from 1001 to 10000 the nearest 500 and from 10000 onwards to the nearest 2000 (Figure 6.7). The counts were then marked as incomplete or missing. All counts were equally likely to be missing but counts from sites with higher site factors were more likely to be flagged as incomplete.




Figure 6.7: Simulating precision effects for the invasive population simulation (see Figure 6.5) with beta errors added (see Figure 6.6).

Table 6.2 summarises the results in the indices produced on the simulated data sets with various combinations of observation error, rounding and missing flags added. An example index plot of one of the simulation runs is in Figure 6.8. Neither adding observation error, rounding flock sizes nor marking up to 45% of the data as missing and 21% as incomplete made a significant impact on the estimated indices. Simulating spatial trends also did not affect the estimated index to a great degree.




Figure 6.8: Index results for the invasive population simulation with spatial trends (see Figure 6.5). The index before error is added is in red ('Real'). Observation error where observations are more likely to under count than over count was added (see Figure 6.6) and the observations rounded (see Figure 6.7) and the index recalculated ('No missing'; blue). 45% of the observations were set to be missing and 21% incomplete. The Underhill index (black) was then calculated together with the equivalent index which ignores incompleteness ('GLM'; green).

Trend	Spatial	Error	Precision	Missing %	No missing data		Generalised Linear Model		Underhill Model	
					X	Sign	X	Sign	X	Sign
RY	CL	Normal	Precise	Low	0.612 (0.000)	0.620 (0.000)	0.936 (0.369)	0.631 (0.217)	1.231 (0.529)	0.878 (0.114)
RY	CL	Normal	Rounded	Low	0.580 (0.000)	0.580 (0.000)	0.939 (0.357)	0.605 (0.213)	1.177 (0.460)	0.863 (0.136)
RY	CL	Beta	Precise	Low	0.645 (0.000)	0.820 (0.000)	0.940 (0.388)	0.663 (0.204)	0.917 (0.224)	0.662 (0.172)
RY	CL	Beta	Rounded	Low	0.563 (0.000)	0.780 (0.000)	0.971 (0.421)	0.648 (0.212)	0.942 (0.261)	0.630 (0.164)
RY	CL	Beta	Rounded	High	0.563 (0.000)	0.780 (0.000)	1.219 (0.568)	0.611 (0.232)	1.567 (0.417)	0.737 (0.156)
RY	CO	Normal	Precise	Low	0.360 (0.000)	0.220 (0.000)	0.411 (0.108)	0.331 (0.150)	0.449 (0.114)	0.497 (0.154)
RY	CO	Normal	Rounded	Low	0.342 (0.000)	0.440 (0.000)	0.434 (0.073)	0.411 (0.151)	0.470 (0.093)	0.515 (0.148)
RY	CO	Beta	Precise	Low	0.345 (0.000)	0.380 (0.000)	0.438 (0.100)	0.381 (0.166)	0.497 (0.125)	0.526 (0.162)
RY	CO	Beta	Rounded	Low	0.332 (0.000)	0.380 (0.000)	0.442 (0.104)	0.379 (0.158)	0.500 (0.141)	0.509 (0.155)
RY	CO	Beta	Rounded	High	0.332 (0.000)	0.380 (0.000)	0.557 (0.156)	0.442 (0.208)	0.785 (0.246)	0.528 (0.192)

Table 6.2: Results for the simulations. The first column gives the trend type from Figures 6.4 and 6.5. ‘RY’ denotes the invasive population trend, ‘SV’ an increasing population and ‘CO’ a constant population. ‘CO’ in the second column, denotes an evenly distributed, unchanging spatial population and ‘CL’ an unevenly distributed, spatially changing population. Observation error where observations are more likely to under count than over count was added are denoted by ‘Beta’ and when both are equally likely denoted by ‘Normal’ in the third column (see Figure 6.6). Column four gives observations as being ‘Rounded’ or ‘Precise’ (see Figure 6.7). Column five the level of missing data applied: a missing rate of 26% and an incomplete rate of 16% is denoted as ‘Low’ and a missing rate of 45% and an incomplete rate of 21% is denoted as ‘High’. The sum of the difference over the 50 years between the index calculated before error is added and three indices was calculated for 100 runs of counts flagged as missing at the appropriate rate. The mean and standard error for each is shown in the columns headed ‘X’. The mean and standard error for proportion of time each index was greater than the error-free index is given in the columns headed ‘Sign’. The three indices are the Underhill index, where incomplete counts are imputed; the equivalent GLM index which ignores incompleteness; and the index with all counts flagged as complete (‘No missing data’). Note that the latter is the same for each run of simulating missing data. *Continues on next page.*

Table 6.2 continued...

Trend	Spatial	Error	Precision	Missing %	No missing data		Generalised Linear Model		Underhill Model	
					X	Sign	X	Sign	X	Sign
SV	CL	Normal	Precise	Low	0.647 (0.000)	0.820 (0.000)	1.017 (0.450)	0.679 (0.284)	0.953 (0.313)	0.361 (0.240)
SV	CL	Normal	Rounded	Low	0.621 (0.000)	0.800 (0.000)	0.964 (0.438)	0.673 (0.247)	0.959 (0.326)	0.363 (0.237)
SV	CL	Beta	Precise	Low	0.449 (0.000)	0.800 (0.000)	0.944 (0.452)	0.613 (0.308)	0.980 (0.286)	0.289 (0.221)
SV	CL	Beta	Rounded	Low	0.467 (0.000)	0.780 (0.000)	0.847 (0.350)	0.618 (0.272)	0.981 (0.297)	0.298 (0.220)
SV	CL	Beta	Rounded	High	0.467 (0.000)	0.780 (0.000)	1.217 (0.522)	0.573 (0.312)	1.591 (0.591)	0.208 (0.151)
SV	CO	Normal	Precise	Low	0.280 (0.000)	0.200 (0.000)	0.359 (0.112)	0.303 (0.182)	0.394 (0.133)	0.322 (0.215)
SV	CO	Normal	Rounded	Low	0.495 (0.000)	0.060 (0.000)	0.523 (0.162)	0.134 (0.110)	0.501 (0.168)	0.198 (0.142)
SV	CO	Beta	Precise	Low	0.288 (0.000)	0.800 (0.000)	0.371 (0.093)	0.710 (0.184)	0.411 (0.122)	0.702 (0.193)
SV	CO	Beta	Rounded	Low	0.307 (0.000)	0.700 (0.000)	0.391 (0.101)	0.651 (0.192)	0.472 (0.138)	0.651 (0.219)
SV	CO	Beta	Rounded	High	0.307 (0.000)	0.700 (0.000)	0.513 (0.182)	0.627 (0.233)	0.674 (0.253)	0.623 (0.236)
CO	CO	Normal	Precise	Low	0.715 (0.000)	0.960 (0.000)	0.754 (0.218)	0.928 (0.073)	0.814 (0.271)	0.914 (0.104)
CO	CO	Normal	Rounded	Low	0.797 (0.000)	0.960 (0.000)	0.827 (0.283)	0.925 (0.083)	0.852 (0.331)	0.880 (0.152)
CO	CO	Beta	Precise	Low	0.311 (0.000)	0.460 (0.000)	0.451 (0.101)	0.529 (0.216)	0.497 (0.127)	0.516 (0.221)
CO	CO	Beta	Rounded	Low	0.383 (0.000)	0.320 (0.000)	0.514 (0.113)	0.358 (0.201)	0.564 (0.135)	0.372 (0.226)
CO	CO	Beta	Rounded	High	0.383 (0.000)	0.320 (0.000)	0.608 (0.145)	0.435 (0.232)	0.746 (0.194)	0.476 (0.252)
CO	CL	Normal	Precise	Low	0.985 (0.000)	0.100 (0.000)	1.464 (0.710)	0.239 (0.261)	2.126 (0.750)	0.094 (0.100)
CO	CL	Normal	Rounded	Low	1.044 (0.000)	0.080 (0.000)	1.534 (0.728)	0.204 (0.240)	2.100 (0.790)	0.099 (0.114)
CO	CL	Beta	Precise	Low	0.766 (0.000)	0.280 (0.000)	1.222 (0.459)	0.278 (0.229)	1.910 (0.749)	0.129 (0.152)
CO	CL	Beta	Rounded	Low	0.904 (0.000)	0.200 (0.000)	1.368 (0.525)	0.268 (0.267)	2.144 (0.847)	0.104 (0.153)
CO	CL	Beta	Rounded	High	0.904 (0.000)	0.200 (0.000)	1.665 (0.763)	0.346 (0.282)	3.789 (1.180)	0.073 (0.119)

6.4 Comparing seasonal population indices using theoretical examples of phenological change

6.4.1 Introduction

Simulated examples can be used to investigate how the index formulae summarised in Table 5.2 on page 194 perform under different possible scenarios. As discussed in Section 5.4.2, the 5-year Laspeyres chain index appears to offer a useful compromise index between indices based on the average number of birds present in different month and those related to the maximum number of birds present.

Figure 6.9 shows four theoretical biological scenarios of birds arriving and departing a wetland system. Seven counts are taken, in the middle of each month, corresponding to Wetland Bird Survey practice. Figure 6.10 compares the simulated mean and maximum number of birds present with the indices based on the mean of the seven monthly counts and the maximum of the seven monthly counts. Also shown is the the 5-year Laspeyres chain index of the seven monthly counts, and the total number of birds simulated.

In the first, “Phenological shift” scenario it is assumed that a population of 2000 individuals does not change over 50 years, but the period that the birds spend within the U.K. is changing over time. The birds arrive from abroad (i.e. outside the simulated system) on average one day later each year, but also leave on average one day later so that the average length of time spent within the system does not change. The second example, “Autumn phenological change”, assumes a situation of smooth phenological change where birds in the simulation are arriving on average one day later a year, but are not changing their departure date. This could represent a situation where the species is staying longer in other countries due to a later onset in hard winter conditions before being forced to migrate to the U.K., but where the trigger for leaving for breeding areas in spring is increasing daylight hours, which is not affected by climate. In both these examples the number of individuals stays constant over the fifty year time period.

In the second two examples, the population changes over time but the individual's seasonality remains the same. The simulation "No phenological trend" is of a simple population that increases steadily from 1000 to 2000 individuals over time. "Increasing spring population" is an example of a mixed population, perhaps from two different breeding grounds, which both winter in the U.K. The first sub-population peaks in autumn and is not changing over time, remaining steady at 1000 individuals. The second sub-population peaks in spring and increases from 500 individuals to 1000 individuals over the fifty year time period. The time the two sub-populations spend in the system overlap in mid-winter.

6.4.2 Results

All the indices give similar results for the "Phenological shift" and "No phenological trend" as in these examples the underlying mean, maximum and number of individuals present all exhibit the same trend. When there is real or apparent phenological change, as with the "Autumn phenological change" and "Increasing spring population" examples, the mean, maximum and number of individuals present depend on the underlying behaviour of the individuals and hence the indices based on the mean count and maximum count show different trends. The 5-year Laspeyres chain index is a compromise between these two, but this does not mean it necessarily closely match the number of individuals present.

Figure 6.10 highlights the necessity of being explicit about what the index is intended to represent: the maximum number of birds present in the U.K. at any one time; the total number of individuals that spend any proportion of the winter in the U.K.; or the mean number of individuals present over the desired period (e.g. September to March). The choice of which month aggregation formula to use depends on which of these is considered desirable. The choice for a particular species may require ornithological field research; if nothing was known about the behaviour of the sub-populations, it would not be possible to distinguish the monthly counts from the "Increasing spring population" scenario with a scenario of phenological change where individuals were staying later in

the spring over time.




Figure 6.9: Phenological change simulations. Simulated phenological change scenarios over 50 years (for simulation details see text).

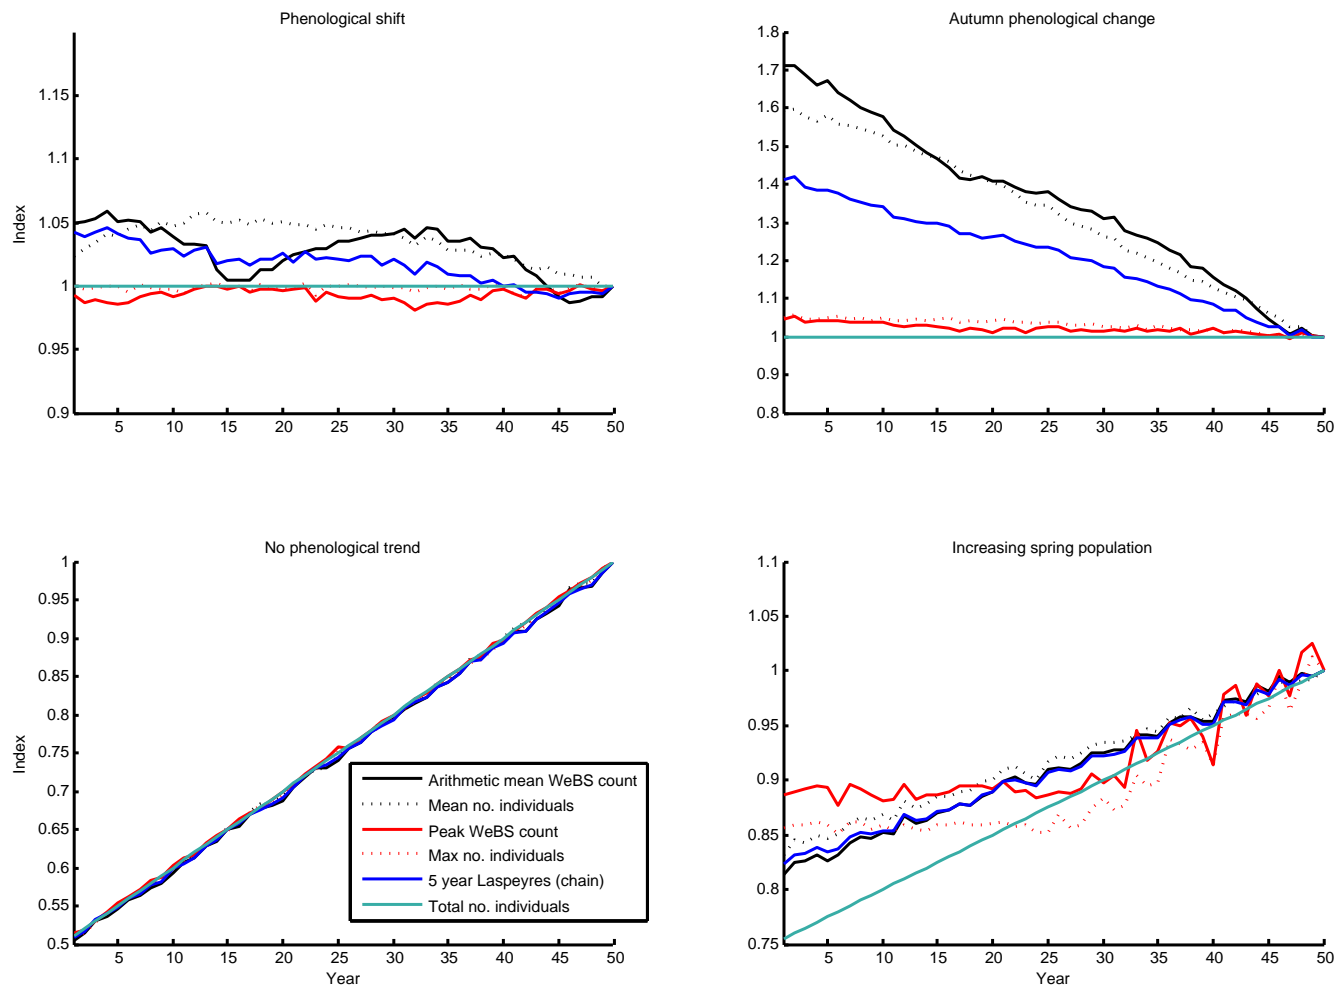


Figure 6.10: Indices for the phenological change simulations in Figure 6.9.

6.5 Discussion

This chapter began by building a simulation model for wintering wildfowl movements, which simulated the location of each individual bird or flock for each day over a winter season at a system of sites. The population was assumed to be wholly migratory, arriving in autumn and departing again in spring.

Section 6.3 added observation error, rounding effects, uneven spatial distribution and missing data to simulations of a constant population trend, an increasing population and an invasive population which began to be culled. Indices based on the mean of available counts over all sites which are currently used in WeBS were compared and found to be robust to large amounts of error and missing data at estimating the underlying population trend.

Different trends in the intra-year seasonal pattern, kept constant in the previous section, were simulated in Section 6.4. The simulations highlighted the difficulty of deciding on an appropriate index when it is not known if seasonal changes are the result of individual birds increasing or decreasing their arrival or departure dates or true population increases or decreases in sub-populations of wintering birds.

Chapter 7

CONCLUSION

The Wetland Bird Survey (WeBS), introduced in Chapter 1, is an ecological monitoring data set on a scale which presents unique challenges. This thesis looked in turn at four aspects to the WeBS objective of assessing trends in wildfowl populations: what sites to use in the index (Chapter 2); what months to use in the index (Chapter 3); how to model missing counts (Chapter 4); and how to aggregate site and month counts into a national annual index (Chapter 5). Some of the methods were then compared using simulated wildfowl count data (Chapter 6).

Sites

Although it is desirable that annual wildfowl indices reflect changes in population at all waterbodies, Chapter 2 discussed some of the difficulties in designing a survey approach to take account of variations in waterbody and waterfowl distributions across the country. Section 5.2.2 discussed the possibility of a weighted index that takes into account the distribution of wetland habitat types to produce a representative index. This approach would require accurate geospatial data for both WeBS sites and all wetlands, together with ecological assessment of wildfowl distributions: how closely can wildfowl numbers at a wetland be modelled by the underlying waterbody, specifically its location and wetland habitat type?

In the absence of this detailed assessment of wildfowl distributions, it is important to emphasise that indices based on WeBS data only represent trends on WeBS sites. The

sites used are not selected by a random sampling scheme, but rather on the availability of observers and prior knowledge of the number of waterfowl found there. Confidence intervals produced by methods such as the bootstrap sample of sites approach (Section 4.3.4) are thus not an estimate of confidence in the trend in the national population of a species.

Due to the uneven distribution of waterfowl over sites, any index based on the arithmetic mean over sites, such as the Underhill index, is influenced mainly by the sites with the largest populations (for example, see Figure 2.12 on page 68). Section 2.4 proposed that a ‘top sites’ criterion replace the ‘50% completeness’ criterion currently used to determine site inclusion in national indices. The proposed method prioritises sites which hold the most individuals for each species and hence results in indices which look similar to indices using every site in the WeBS database. Advantages of restricting the number of sites to those in the ‘top sites’ criterion include faster convergence of the Underhill algorithm and making feasible computer intensive procedures, for example fitting more complicated models such as those in Section 4.3.3, and the bootstrapping methods in Section 4.3.4. In addition, there is greater clarity in which sites are determining the trend than nominally including hundreds of additional sites which do not affect the resultant trend, and surveying effort can then be directed at these sites to minimise missing data. Unlike the current 50% completeness criterion, the top sites criterion for site selection ensures that data from internationally important wetlands is not ignored when constructing national trends (Section 2.3).

Simulations appear to show that indices based on mean counts over sites such as the Underhill index are robust to different site trends (Section 6.3). However, Chapter 4 showed that the assumption that has the greatest effect on annual indices is that all sites have an identical trend. A consequence of this assumption is that missing data at one or two key sites for very clumped species such as Wigeon can have a dramatic effect on the annual index. It is recommended that future research focuses on developing methods for modelling or imputing counts non-parametrically for relatively few influential sites

rather than attempting to model missing values at all WeBS sites.

Months

Chapter 3 considered the seasonal aspects of WeBS data. As well as marked inter-species variation in the seasonal patterns of abundance, there is considerable spatiotemporal intra-species variation. We saw that functional data analysis is particularly useful for analysing dynamic wildlife counts. The potential of smoothing for research into population change using large scale abundance monitoring programs is starting to be revealed (e.g. Moussus et al., 2009) but it is thought by the author that the functional data approach presented here is completely novel. There is scope for more research in how to use ideas from functional data analysis in ecological modelling data. For example, phase plane plots could be used to investigate phenological change in other ecological monitoring data such as weekly butterfly counts, as they help reveal the intra-year population dynamics.

Following on from the investigation into phenological change in Chapter 3, it was shown in Chapter 4 that the key assumptions in the currently used Underhill model for imputing missing counts, namely that annual trends are identical for all months and that seasonal patterns are the same at all sites, are not realistic. To minimise the effect of seasonality on the index, mid-winter months were chosen in Section 3.5 to be those where the population was at its highest and not changing greatly. One of the advantages of restricting the number of months used is to avoid imputing missing data in the spring and autumn, when there tends to be more missing and incomplete counts. An alternative approach is to improve the modelling of counts in spring and autumn by including phenological change and regional seasonality and not exclude any months.

Previously, most analysis with the objective of determining trends has focused on dealing with missing data and combining information from different sites to create a yearly wintering population index. Little attention has been paid to the challenges of combining

monthly data to obtain a single yearly index figure, as was explored in Chapter 5. Indices based on the mean over the whole winter period were contrasted with those based on the winter population peak; a compromise between these two can be achieved using a weighted index, such as the Laspeyres-type chain index where counts are weighted using the average seasonality in the base year and the four years preceding it. As it is a chain index, each index number depends on the intervening index values as well as the current and base year counts.

Only by being clear about what the index is supposed to be tracking, can a decision be made about which index is best. Measuring trends in wintering wildfowl is important, because they are used in determining conservation policy. For example, national population trends are needed to put local trends on nationally and internationally designated sites, such as Sites of Special Scientific Interest and Special Protection Areas, into context, as this information is used to monitor the site condition. Trends in non-native species such as the Ruddy Duck can be easier to determine using winter counts than breeding surveys, since the Ruddy Duck congregate on large waterbodies in the winter which are easier to survey; large increases can highlight the need to take measures to control the population size, as it did for this species. Decreases in national trends in quarry species may lead to protection under the EC Directive on wild birds, restricting hunting outside the breeding season. Declines in wintering numbers using trends from the Wetland Bird Survey are one of the criteria for listing birds as 'Red' or 'Amber' in 'Birds of Conservation Concern', although it is acknowledged that declines may be due to a movement away from the UK as opposed to decline in the international population (Eaton et al., 2009).

This is the heart of the matter: does it make sense to produce national trends in open populations, such as wintering wetland birds, in the way that population indices are produced for closed breeding populations? If so, what should the trend be based on: the peak number of birds present, or the average number of birds present over the whole season (so that phenological change can affect the increase or decrease in the index.) It may be considered desirable for the index to indicate the UK's annual success as a winter "holi-

day destination”, tracking the total number of individual birds that visit these shores, but that requires detailed knowledge about bird movements which is not currently collected.

Conclusion

The oft-repeated rider in Wetland Bird Survey results that WeBS indices are most representative for clumped species, whose populations are concentrated on large waterbodies that are well surveyed by WeBS, is at variance with some of the methods currently used for analysing the data. The philosophy behind choosing which sites and months to include in an index, the method used to impute or model data and the aggregation formulae used to construct an annual index, should all reflect the underlying clumped distribution of wildfowl populations.

The method for producing indices by Underhill (1989) was a large step forward for the Wetland Bird Survey, but only small changes to analyses have been made in the past twenty years, despite the availability of many new statistical methods and advances in technology opening up many new and exciting possibilities for better using the data collected so diligently by volunteers over the past sixty years. I hope that this thesis has given a glimpse of the possibilities of this long term monitoring data set, and how it can be used more effectively to answer ornithologists' questions.

Appendix A

PRIORITY SITES

This league table of WeBS sites is intended as a guide to how significant sites are to constructing indices for the twenty-six wildfowl types. The top sites procedure from Section 2.4.3 using the stable months for each species from Section 3.5.2 was used to select sites for each species; 739 unique sites appeared amongst the data sets. The sum of the influence of each site over the twenty-six species (multiplied by 100) was used to rank the sites. The number of species for which each site was chosen by the top sites routine is also given in the following table. The table lists the most significant 338 sites, which have an influence measure of at least one. Incomplete and missing counts should be avoided at high ranking sites if possible in the future.

Rank	WeBS Code	Site name	Influence	no. species
1	32355	Ouse Washes	96.19	14
2	83420	Forth Estuary	90.88	16
3	15YYY	Severn Estuary (English counties)	81.44	16
4	70YYY	Solway Estuary (Scottish counties)	73.83	14
5	35901	The Wash	50.09	15
6	34905	North Norfolk Coast	48.84	17
7	25901	Thames Estuary	44.94	17
8	57910	Morecambe Bay	44.27	19
9	36156	Rutland Water	42.15	17
10	90920	Inner Moray and Inverness Firth	40.91	15
11	57901	Ribble Estuary	37.91	10
12	45YYY	Dee Estuary (English counties)	35.02	9
13	57011	WWT Martin Mere	33.88	13
14	22450	Swale Estuary	31.95	16
15	45421	Mersey Estuary	30.45	7
16	38950	Humber Estuary	30.44	20
17	70080	WWT Caerlaverock (Inland)	30.09	7
18	88091	Loch of Strathbeg	28.56	13
19	86001	Loch Leven	28.13	16
20	25121	Abberton Reservoir	27.3	16
21	38051	Hornsea Mere	27.01	13
22	59YYY	Solway Estuary (English counties)	26.87	13
23	13300	Somerset Levels	24.67	11
24	76311	Rhunahaorine	22.62	1
25	14102	Chew Valley Lake	21.39	14
26	77938	Loch Indaal (Bruichladdich to Laggan)	20.69	5
27	43156	Blithfield Reservoir	20.16	14
28	90950	Cromarty Firth	19.43	12
29	76305	Machrhanish (Geese)	18.76	1
30	70380	Mersehead RSPB Reserve	18.58	2
31	90970	Dornoch Firth Consolidated	17.44	13
32	12901	Fleet and Wey	17.42	15
33	WWC06		16.6	2
34	25948	Blackwater Estuary	16.45	13
35	55481	Lindisfarne	16.39	13
36	41751	Middle Tame Valley Gravel Pits	16.19	11
37	71121	Loch Ken	15.87	9
38	WWC05		15.79	2
39	26701	Lee Valley Gravel Pits	15.2	11
40	22460	Medway Estuary	14.79	12
41	WWC11		14.37	1
42	WWC08		13.93	2
43	58XXX	Duddon Estuary	13.89	9
44	17743	Eversley Cross and Yateley Gravel Pits	13.69	6
45	17304	River Avon: Fordingbridge to Ringwood	13.59	11
46	86013	Dupplin Lochs	12.73	2
47	960H1	Loch of Harray	12.68	8
48	34301	Middle Yare Marshes	12.56	8
49	75040	Loch Lomond	12.44	10
50	15220	Cotswold Water Park (West)	12.29	12
51	20401	Chichester Harbour	12.24	9
52	22291	Dungeness Gravel Pits	11.85	17
53	30141	Pitsford Reservoir	11.39	14
54	17431	Langstone Harbour	10.88	8
55	12421	Poole Harbour	10.6	15

Rank	WeBS Code	Site name	Influence	no. species
56	17073	Stratfield Saye	10.52	4
57	88021	Loch of Skene	10.3	9
58	33480	Alde Complex	10.22	12
59	63400	Burry Inlet	10.03	10
60	51003	Fairburn Ings	10	12
61	63499	Dyfi Estuary	9.76	9
62	28700	Wraysbury Gravel Pits	9.64	8
63	25481	Stour Estuary	9.59	10
64	90096	Loch Eye	9.53	8
65	34901	Breydon Water and Berney Marshes	8.95	12
66	38301	Lower Derwent Ings	8.81	11
67	77438	Loch Indaal: Blackrock to Bowmore	8.66	5
68	35700	Baston and Langtoft Gravel Pits	8.64	9
69	87461	Montrose Basin	8.56	9
70	81318	Lower Teviot Valley	8.5	6
71	81840	River Tweed: Kelso to Coldstream	8.43	5
72	17805	River Avon: Ringwood to Christchurch	8.4	9
73	25101	Hanningfield Reservoir	8.29	12
74	43111	Belvide Reservoir	8.06	9
75	WWC01		7.96	1
76	20800	Arun Valley	7.88	12
77	88501	Lochs Davan and Kinord	7.79	6
78	52901	Tees Estuary	7.36	14
79	17022	Baffins Pond	7.25	1
80	32201	Little Paxton Gravel Pits	7.06	8
81	71592	Stranraer Lochs	7	2
82	23103	Staines Reservoirs	6.97	5
83	38151	Tophill Low Reservoirs	6.95	10
84	14101	Blagdon Lake	6.91	10
85	30251	Thrapston Gravel Pits	6.9	10
86	25953	Colne Estuary	6.8	9
87	17912	Southampton Water	6.73	11
88	69420	Welsh Dee	6.68	7
89	32803	Nene Washes	6.63	9
90	41071	Packington Park	6.53	6
91	58031	Windermere	6.45	8
92	25971	Hamford Water and Naze Combined	6.45	9
93	22381	Walland Marsh	6.43	6
94	25931	Crouch-Roach Estuary	6.4	10
95	11450	Exe Estuary	6.29	9
96	930F1	Loch Bee (South Uist)	6.15	2
97	75999	Inner Firth of Clyde	6.01	9
98	63342	River Tywi: Dryslwyn	6	4
99	26702	Colne Valley Gravel Pits	5.95	7
100	12001	Abbotsbury	5.9	11
101	33071	Minsmere Levels	5.75	10
102	71491	Loch Ryan	5.69	6
103	85971	Tay Estuary	5.6	8
104	33082	Benacre Broad	5.5	3
105	33902	Orwell Estuary	5.5	11
106	45381	Woolston Eyes	5.46	7
107	51239	Nosterfield Gravel Pits	5.44	4
108	34005	Hickling Broad	5.42	9
109	82101	West Water Reservoir	5.4	2
110	60YYY	Severn Estuary (Welsh counties)	5.38	11

Rank	WeBS Code	Site name	Influence	no. species
111	15741	Cotswold Water Park (East)	5.34	4
112	950K3	Loch Heilen	5.14	3
113	29290	Lower Windrush Valley Gravel Pits	5.1	7
114	44076	Ellesmere Lakes	5.09	6
115	22211	Sevenoaks Exp Wildfowl Reserve	5.03	2
116	88561	Slains Lochs (Meikle and Sand and Cotehill)	5.02	6
117	33352	North Warren and Thorpness Mere	5	9
118	75552	Lochs Quien Fad and Dhu (Bute)	4.95	9
119	36151	Eyebrook Reservoir	4.84	10
120	87101	Loch of Lintrathen	4.82	7
121	960H0	Loch of Stenness	4.81	6
122	45805	Dee Flood Meadows	4.78	6
123	WWC07		4.62	1
124	67412	Lavan Sands	4.6	9
125	32207	Fen Drayton Gravel Pits	4.59	9
126	51251	Bolton-on-Swale Gravel Pits	4.59	6
127	37262	Sutton and Lound Gravel Pits	4.58	6
128	86041	Carsebreck and Rhynd Lochs	4.5	7
129	36101	Stanford Reservoir	4.47	7
130	86031	Drummond Pond	4.4	4
131	63943	Cleddau Estuary	4.37	7
132	26121	Tring Reservoirs	4.31	9
133	33101	Alton Water	4.21	5
134	51031	Castle Howard Lake	4.17	5
135	55491	Tweed Estuary	4.17	3
136	32101	Grafham Water	4.09	7
137	34025	Gunton Park	4.07	4
138	24142	Walthamstow Reservoirs	4.05	5
139	30700	Stanwick Gravel Pits Consolidated	3.96	8
140	88303	River Dee: Kincardine O'Neil (Geese)	3.89	1
141	80001	Hirsel Lake	3.89	4
142	17901	North West Solent	3.73	9
143	50189	Eccup Reservoir	3.71	4
144	960J2	Loch of Boardhouse	3.68	3
145	30242	Earls Barton Gravel Pits	3.68	7
146	33911	Deben Estuary	3.66	9
147	27261	Emberton Gravel Pits	3.59	2
148	45271	Farmwood Pool	3.56	2
149	87880	Lour Area (Geese)	3.54	1
150	89018	Loch Spynie	3.54	4
151	26101	Hilfield Park Reservoir	3.47	3
152	22072	Stodmarsh NNR and Collards Lagoon	3.46	6
153	953M1	Broubster Leans	3.43	2
154	21111	Barcombe Mills Reservoir	3.4	1
155	32232	Ferry Meadows (Nene Park)	3.34	3
156	57081	Leighton Moss	3.21	4
157	48019	Kedleston Park Lake	3.2	1
158	20412	Pagham Harbour	3.15	6
159	44071	Shavington Park	3.15	5
160	21151	Bewl Water	3.15	5
161	WWC04		3.14	1
162	72415	Girvan to Turnberry	3.12	4
163	15201	Frampton Pools	3.09	3
164	37202	Holme Pierrepont Gravel Pits	3.07	5
165	960V2	Milldam and Balfour Mains Pools	3.05	6

Rank	WeBS Code	Site name	Influence	no. species
166	88071	Haddo House Lakes	3.04	2
167	950K2	Loch Watten	3.04	6
168	17421	Portsmouth Harbour	3.01	5
169	25441	Dengie Flats	2.94	6
170	83023	Duddingston Loch	2.94	2
171	27253	Linford Gravel Pits	2.91	6
172	WWC02		2.9	2
173	73316	River Clyde: Carstairs Junction Bridge...	2.88	2
174	85061	Kilconquhar Loch	2.88	5
175	68121	Llyn Alaw	2.85	5
176	68070	Llynau Y Fali	2.84	4
177	33201	Lackford GPs	2.83	5
178	27132	Willen Lake	2.83	6
179	20201	Chichester Gravel Pits	2.82	5
180	40101	Bittell Reservoirs	2.76	2
181	85962	St Andrews Bay	2.75	2
182	55081	Holburn Moss	2.74	4
183	74021	Lochwinnoch (Aird and Barr and Castle Semple)	2.73	5
184	71421	Wigtown Bay	2.7	6
185	23102	King George VI Reservoir	2.67	5
186	21420	Rye Harbour and Pett Level	2.67	10
187	48190	Carsington Water	2.65	6
188	37201	Attenborough Gravel Pits	2.62	5
189	37077	Clumber Park Lake	2.61	3
190	68024	Llyn Traffwll	2.6	2
191	86337	River Tay: Haughs of Kercock	2.57	2
192	50011	Swillington Ings	2.56	5
193	59031	Derwent Water	2.55	5
194	WWC10		2.5	1
195	24101	Hampton and Kempton Reservoirs	2.5	4
196	35266	Kirkby-on-Bain Gravel Pits	2.49	5
197	59026	Ullswater	2.43	2
198	18402	Newtown Estuary	2.43	6
199	11XX1	Taw-Torridge Estuary	2.42	8
200	77022	Loch Skerrols	2.41	3
201	13121	Cheddar Reservoir	2.4	5
202	40210	Bodenham Gravel Pit	2.32	1
203	32602	Buckden and Stirtloe Pits	2.28	4
204	50061	Harewood Lake	2.28	3
205	50220	Winterset Country Park Lake	2.28	5
206	30727	Clifford Hill Gravel Pits Consolidated	2.27	5
207	85161	Cameron Reservoir	2.25	3
208	10460	Tamar Complex	2.24	8
209	87034	Kinnordy Loch	2.24	3
210	87351	Padanaram	2.23	1
211	34038	Wroxham Broad	2.22	1
212	73042	Hamilton Low Parks and Strathclyde Park	2.22	4
213	73052	Hogganfield Loch	2.21	5
214	45057	Rostherne Mere	2.2	7
215	36141	Swithland Reservoir	2.19	4
216	69075	Hanmer Mere	2.16	3
217	90994	Loch Fleet Complex	2.15	5
218	83139	Gladhouse Reservoir	2.15	4
219	49046	Wath Main Ings	2.13	6
220	41112	Draycote Water	2.13	5

Rank	WeBS Code	Site name	Influence	no. species
221	28228	Dinton Pastures	2.08	4
222	85461	Eden Estuary	2.07	6
223	34309	St Benet's Levels	2.06	3
224	23211	Thorpe Water Park	2.04	4
225	970B1	Loch of Spiggie	2.03	2
226	80011	Hule Moss	2.01	2
227	57021	Marton Mere (Blackpool)	1.98	5
228	50054	Pugneys Country Park Lakes	1.97	6
229	70341	Waterside Mains Keir near Thornhill	1.97	1
230	28245	Burghfield Gravel Pits	1.96	4
231	24121	Brent Reservoir	1.95	4
232	70171	Glenkiln Reservoir	1.95	1
233	24115	Barn Elms Reservoirs	1.92	4
234	28242	Theale Gravel Pits	1.89	4
235	43225	Croxall Pits	1.88	5
236	30052	Blatherwyke Lake	1.87	3
237	47364	River Irwell	1.87	4
238	35268	Tattershall Pits	1.86	3
239	43061	Aqualate Mere	1.84	3
240	23213	Shepperton Gravel Pits	1.84	4
241	970U1	Easterloch Uyeasound (Unst)	1.81	2
242	72425	Ayr to North Troon	1.81	5
243	68901	Inland Sea and Beddmanarch Bay and Alaw Estuary	1.8	4
244	29221	Dorchester Gravel Pits	1.78	4
245	30246	Ditchford Gravel Pits	1.77	4
246	72470	Arran	1.76	1
247	35001	Grimsthorpe Lake	1.75	2
248	49007	Rother Valley Country Park	1.74	4
249	74341	Black Cart Water (Gryfe-White Cart)	1.73	2
250	81062	Hoselaw Loch	1.71	2
251	55201	Caistron Quarry	1.71	5
252	78082	Loch A'phuill (Tiree)	1.69	4
253	83001	Linlithgow Loch	1.68	5
254	22272	Seaton Gravel Pits and River	1.67	3
255	73309	River Clyde: Lamington	1.67	4
256	75055	Loch Ascog (Bute)	1.67	3
257	43262	Barton Pits	1.65	5
258	34006	Horse Mere	1.63	4
259	34030	Hardley Flood	1.63	4
260	71093	White Loch (Lochinch)	1.63	1
261	47112	Audenshaw Reservoirs	1.62	2
262	75054	Dhu Loch (Bute)	1.62	1
263	47503	Wigan Flashes (Bryn Marsh Scotsmans...)	1.6	4
264	70023	Castle Loch (Lochmaben)	1.6	3
265	16800	Avon Valley: Salisbury to Fordingbridge	1.56	3
266	12431	Christchurch Harbour	1.56	4
267	45001	Combermere	1.54	1
268	17405	Beaulieu Estuary	1.53	5
269	72427	Irvine and Garnock Estuary	1.53	6
270	43221	King's Bromley Gravel Pits	1.53	2
271	42111	Chasewater	1.51	3
272	31281	Blunham Gravel Pit	1.51	3
273	24023	St James's Park	1.5	2
274	88917	Don Mouth to Ythan Mouth	1.5	3
275	25301	Colchester Hythe	1.49	1

Rank	WeBS Code	Site name	Influence	no. species
276	90041	Loch Insh and Spey Marshes	1.49	3
277	10413	Fal Complex	1.47	4
278	950M1	Loch Scarmclate	1.46	3
279	88321	River Don: Kemnay (Fetternear) to Inverurie	1.46	2
280	85012	Loch Ore	1.46	3
281	24022	Serpentine and Long Water	1.45	3
282	32221	Meadow Lane Gravel Pits St Ives	1.44	3
283	45251	Fiddlers Ferry Power Station Lagoons	1.44	5
284	37212	Hoveringham and Bleasby Gravel Pits	1.42	3
285	22101	Bough Beech Reservoir	1.41	3
286	45058	Tatton Mere	1.4	3
287	24151	William Girling Reservoir	1.4	4
288	32003	Guppy's Pond	1.39	3
289	37242	Besthorpe and Girton Gravel Pits and Fleet	1.37	4
290	34012	Ludham How Hill	1.34	4
291	950K1	Loch of Wester	1.34	2
292	51223	Hay-a-Park Gravel Pits	1.32	2
293	WWE09		1.31	1
294	32243	Bainton Pits	1.3	3
295	45075	Farmwood Pool	1.3	1
296	28226	Twyford Gravel Pits	1.29	2
297	41051	Coombe Country Park	1.28	3
298	31245	Radwell Gravel Pits	1.28	2
299	36202	Watermead Gravel Pits	1.26	2
300	71009	Auchenreoch Loch	1.24	2
301	51131	Leighton and Roundhill Reservoirs	1.24	2
302	42515	Sandwell Valley Country Park	1.23	2
303	23101	Wraysbury Reservoir	1.23	2
304	53400	Durham Coast	1.22	2
305	34075	Stanford Training Area	1.21	2
306	21320	Pevensey Levels	1.21	5
307	88418	Ythan Estuary	1.21	8
308	30122	Hollowell Reservoir	1.21	5
309	70409	Islesteps	1.21	1
310	35823	River Welland: Spalding to Borough Fen	1.2	2
311	960H4	Loch of Skail	1.19	3
312	48201	Drakelow Gravel Pit	1.19	3
313	23121	Walton Reservoirs	1.18	2
314	66441	Traeth Bach	1.16	6
315	29370		1.16	3
316	23111	Queen Mary Reservoir	1.15	2
317	59270		1.15	3
318	71903	Orchardton and Auchencairn Bays	1.15	4
319	67421	Foryd Bay	1.14	5
320	43304	Doxey Marshes SSSI	1.14	2
321	72421	Doon Estuary	1.13	3
322	29234		1.13	2
323	23118	Knight and Bessborough Reservoirs	1.12	2
324	68303	Malltraeth RSPB	1.12	2
325	88910	Dee Estuary (Scotland)	1.12	2
326	59370	Rivers Eamont and Eden: Honeypot to Edenhall	1.11	1
327	11103	Venford Reservoir	1.1	1
328	17072	Elvetham	1.09	1
329	83058	Fala Flow (Confidential)	1.09	1
330	86331	River Tay: Almond Mouth to Perth Bridge	1.09	2

Rank	WeBS Code	Site name	Influence	no. species
331	25131	Ardleigh Reservoir	1.07	3
332	63423	Carmarthen Bay	1.06	6
333	34056	Bayfield Park	1.06	1
334	84126	Gartmorn Dam	1.05	3
335	950K4	St Johns Loch	1.03	2
336	88131	Fedderate Reservoir	1.03	1
337	24371	Beddington Sewage Farm	1.01	2
338	59351	Confluence of Rivers Eden and Eamont	1.01	1





Appendix B

MONTH SELECTION





For each of the twenty six wildfowl types listed on page xv, a functional data analysis was undertaken to assess which months may be considered stable. For each species, the top sites routine (see Section 2.4.3) was used for counts for August to April inclusive; although there are more missing counts in August and April, being outside the main September to March counting period, these were used to better assess the population function behaviour in September and March. The number of sites used, their “influence” (the sum of the site factors of the sites used, divided by the total site factor) and the proportion of missing values are given in the Table B.1 below. In each figure that follows, the mean of available counts for each of the years from 1966/67 to 2006/07 are in the top left plot; the seasonality curves for each year and the mean over years in the top right plot; bootstrapped 95% confidence intervals of the derivative of the mean seasonality curve is shown in the bottom left plot; and the phase plane plot of the first derivative against the second derivative is shown in the bottom right plot. For further details see Section 3.5.

	#sites	%miss.	%infl.
Mute Swan	200	41	42
Whooper Swan	86	53	41
Bewicks Swan	32	43	74
Pink-footed Goose	61	48	68
European White-fronted Goose	25	50	82
Greenland White-fronted Goose	16	60	32
Icelandic Greylag Goose	48	54	36
Re-established Greylag Goose	72	45	47
Canada Goose	146	43	35
Svalbard Barnacle Goose	10	48	81
Naturalised Barnacle Goose	42	45	14
Dark-bellied Brent Goose	35	25	99
Shelduck	93	30	89
Wigeon	169	36	74
Gadwall	91	38	52
Teal	200	42	57
Mallard	290	43	35
Pintail	73	32	86
Shoveler	106	35	56
Pochard	166	39	55
Tufted Duck	207	42	47
Scaup	53	40	66
Goldeneye	170	41	49
Red-breasted Merganser	75	35	52
Goosander	127	50	33
Ruddy Duck	55	37	70





Table B.1: Details of sites used to select months for each species. Number of sites chosen by top sites criterion over the months August to April; the percentage missing values in the data for the sites chosen; and the total influence of the sites chosen measured by the sum of site factors as a percentage of the total available sites.







Mute Swan







Whooper Swan







Bewick's Swan




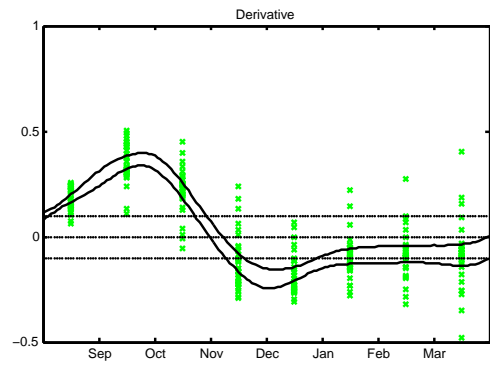


Pink-footed Goose







European White-fronted Goose







Greenland White-fronted Goose







Icelandic Greylag Goose







Re-established Greylag Goose







Canada Goose







Svalbard Barnacle Goose



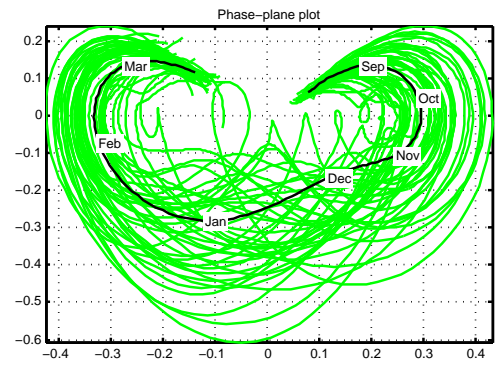



Naturalised Barnacle Goose







Dark-bellied Brent Goose







Shelduck







Wigeon






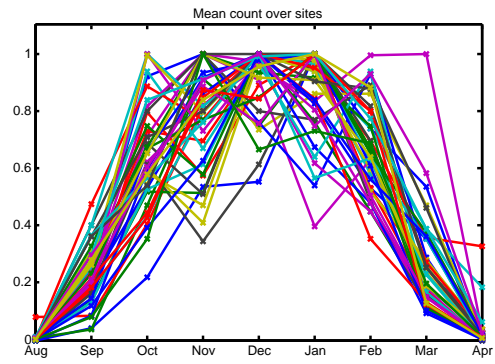
Gadwall







Teal



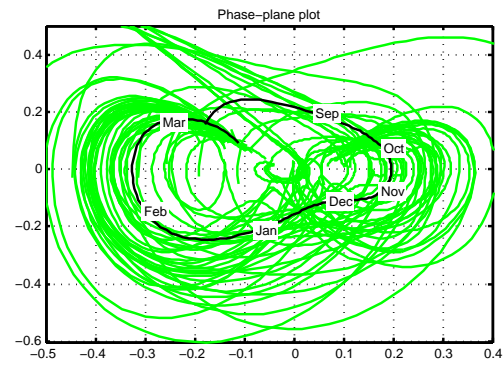



Mallard







Pintail







Shoveler







Pochard







Tufted







Scaup







Goldeneye



Red-breasted Merganser



Goosander



Ruddy Duck

REFERENCES

- M. J. S. Armitage, G. E. Austin, S. J. Holloway, and M. M. Rehfisch. The Pilot Dispersed Waterbirds Survey. Technical Report 271, BTO, 2001.
- P. W. Atkinson, G. E. Austin, M. M. Rehfisch, H. Baker, P. A. Cranswick, M. Kershaw, J. A. Robinson, R. H. W. Langston, R. H. W. Stroud, C. Van Turnhout, and I. M. D. Maclean. Identifying declines in waterbirds: The effects of missing data, population variability and count period on the interpretation of long-term survey data. *Biological Conservation*, 130(4):549–559, 2006.
- G. E. Austin, I. Peachel, and M. M. Rehfisch. Regional trends in coastal wintering waders in Britain. *Bird Study*, 47:352–371, 2000.
- G. E. Austin, S. S. F. Jackson, and H. J. Mellan. WeBS Alerts 2000/2001: Changes in numbers of wintering waterbirds in the United Kingdom, its Constituent Countries, Special Protection Areas (SPAs) and Sites of Special Scientific Interest (SSSIs). Technical report, British Trust for Ornithology, 2004.
- G. E. Austin, M. P. Collier, N. A. Calbrade, C. Hall, and A. J. Musgrove. *Waterbirds in the UK 2006/07: The Wetland Bird Survey*. BTO/WWT/RSPB/JNCC, Thetford, 2008.
- A. N. Banks, M. P. Collier, G. E. Austin, R. D. Hearn, and A. J. Musgrove. *Waterbirds in the UK 2004/05: The Wetland Bird Survey*. BTO/WWT/RSPB/JNCC, Thetford, 2006.
- M. C. Bell. Index numbers for waterbird populations. 4. Methods for the analysis of regional trends. Wwt internal report, The Wildfowl and Wetlands Trust, Slimbridge., 199-.
- M. C. Bell. Modelling winter waterfowl numbers in relation to wetland characteristics: lessons from the old wetland register, 1994a.
- M. C. Bell. The analysis of regional trends in British waterfowl populations: extensions to the Underhill method. Internal report, The Wildfowl and Wetlands Trust, 1994b.
- M. C. Bell. UINDEX4: A computer programme for estimating population index numbers by the Underhill method. User instructions. Manual, The Wildfowl and Wetlands Trust, 1995.
- P. Besbeas. Discussion regarding use of WeBS data for integretate population modelling. Personal Communication [Conversation]. University of Kent, 2007.
- P. Besbeas and S. N. Freeman. Methods for joint inference from panel survey and demographic data. *Ecology*, 87(5):1138–1145, 2006.

- P. Besbeas, S. N. Freeman, B. J. T. Morgan, and E. A. Catchpole. Integrating mark-recapture-recovery and census data to estimate animal abundance and demographic parameters. *Biometrics*, 58(3):540–547, 2002.
- P. Besbeas, J. D. Lebreton, and B. J. T. Morgan. The efficient integration of abundance and demographic data. *Journal Of The Royal Statistical Society Series C-Applied Statistics*, 52:95–102, 2003.
- C. J. Bibby, N. D. Burgess, and D. A. Hill. *Bird Census Techniques*. Academic Press, London, 1992.
- D. L. Borchers, W. Zucchini, and S. T. Buckland. *Estimating Animal Abundance: Closed Populations*. Springer-Verlag, London, 2002.
- C. Both, S. Bouwhuis, C. M. Lessells, and M. E. Visser. Climate change and population declines in a long-distance migratory bird. *Nature*, 441(7089):81–83, 2006.
- S. Briscoe. House price indices. *Significance*, 1(1):18–20, 2004.
- L. Brown and M. Smart. Pintail in the Severn Vale. *Worcestershire Record*, 12, April 2002 2002.
- BTO. *The History and Future of the Wetland Bird Survey (WeBS)*. British Trust for Ornithology. URL: <http://www.bto.org/webs/about/history.htm>. [Date Accessed: October 2009], 200-.
- BTO. *BBS Research: Methodology and survey design*. British Trust for Ornithology. URL: <http://www.bto.org/bbs/research/methodology.htm>. [Date Accessed: October 2009], 2008.
- E. A. Catchpole, B. J. T. Morgan, S. N. Freeman, and W. J. Peach. Modelling the survival of British Lapwings *Vanellus vanellus* using ring-recovery data and weather covariates. *Bird Study*, 46:5–13, 1999.
- M. Cocker. *Birders: Tales of a Tribe*. Jonathan Cape, London, 2001.
- M. P. Collier, A. N. Banks, G. E. Austin, T. Girling, R. D. Hearn, and A. J. Musgrove. *The Wetland Bird Survey 2003/04: Wildfowl and Wader Counts*. BTO/WWT/RSPB/JNCC, Thetford, 2005.
- R. F. Costantino, R. A. Desharnais, J. M. Cushing, and B. Dennis. Chaotic dynamics in an insect population. *Science*, 275(5298):389–391, 1997.
- T. Coulson, T. H. G. Ezard, F. Pelletier, G. Tavecchia, N. C. Stenseth, D. Z. Childs, J. G. Pilkington, J. M. Pemberton, L. E. B. Kruuk, T. H. Clutton-Brock, and M. J. Crawley. Estimating the functional form for the density dependence from life history data. *Ecology*, 89(6):1661–1674, 2008.
- P. A. Cranswick, J. S. Kirby, D. G. Salmon, G. L. Atkinson-Willes, M. S. Pollitt, and M. Owen. A history of wildfowl counts by the Wildfowl and Wetlands Trust. *Wildfowl*, 47:216–229, 1997.

- P. A. Cranswick, J. Worden, R. M. Ward, H. E. Rowell, C. Hall, A. J. Musgrove, R. D. Hearn, S. J. Holloway, A. N. Banks, G. E. Austin, L. R. Griffin, B. Hughes, M. Kershaw, M. J. O'Connell, M. S. Pollitt, E. C. Rees, and L. E. Smith. *The Wetland Bird Survey 2001-03: Wildfowl & Wader Counts*. BTO/WWT/RSPB/JNCC, Slimbridge, 2005.
- Defra. A biodiversity strategy for England. Measuring progress: baseline assessment. Technical report, Department for Environment, Food and Rural Affairs, 2003.
- Defra. Working with the grain of nature taking it forward: Volume II. Measuring progress on the England Biodiversity Strategy: 2006 assessment. Technical report, Department for Environment, Food and Rural Affairs, 2006.
- Defra. *Key Facts about: Wildlife. Population of wintering waterbirds: 1975/1976 - 2005/2006*. Department for Environment, Food and Rural Affairs. URL: <http://www.defra.gov.uk/environment/statistics/wildlife/kf/wdkf03.htm> [Date Accessed: 6th September 2009], 2007.
- Defra. *England Biodiversity Strategy Indicators: W1(b) Populations of wintering water and wetland birds in England*. Department for Environment, Food and Rural Affairs. URL: <http://www.defra.gov.uk/wildlife-countryside/biodiversity/indicator.htm> [Date Accessed: 6th September 2009], 2009.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood From Incomplete Data Via the EM Algorithm. *Journal Of The Royal Statistical Society Series B-Methodological*, 39(1):1–38, 1977.
- W. E. Diewert. The Early History of Price Index Research. In W. E. Diewert and A. O. Nakamura, editors, *Essays in Index Number Theory, Volume 1*. North Holland, Amsterdam, 1993.
- M. A. Eaton, A. F. Brown, D. G. Noble, A. J. Musgrove, R. D. Hearn, N. J. Aebischer, D. W. Gibbons, A. Evans, and R. D. Gregory. Birds of Conservation Concern 3: the population status of birds in the United Kingdom, Channel Islands and the Isle of Man. *British Birds*, 102:296341, 2009.
- C. S. Elphick. How you count counts: the importance of methods research in applied ecology. *Journal Of Applied Ecology*, 45(5):1313–1320, 2008.
- M. E. Evans. Effects Of Weather On The Wintering Of Bewicks Swans *Cygnus-Columbianus-Bewickii* At Slimbridge, England. *Ornis Scandinavica*, 10(1):124–132, 1979.
- F. G. Forsyth and R. F. Fowler. The theory and practice of chain price-index numbers. *Journal of the Royal Statistical Society Series A - Statistics in Society*, 144(2):224–246, 1981.
- R. G. Foster and T. Roenneberg. Human responses to the geophysical daily, annual and lunar cycles. *Current Biology*, 18(17):R784–R794, 2008.

- A. D. Fox and D. G. Salmon. Breeding And Moulting Shelduck (*Tadorna tadorna*) of the Severn Estuary. *Biological Journal of the Linnean Society*, 51(1-2):237–245, 1994.
- T. M. Frost. WeBS Steering Group Update: Choosing which months to use in calculating the yearly Underhill index for wildfowl species. Unpublished report, University of Kent, 2006.
- P. H. Geissler and B. R. Noon. Estimates of avian population trends from the North American breeding bird survey. *Studies in Avian Biology*, 6:42–51, 1981.
- Gidzy. *Pintail Skimming the water, Martin Mere, March 2009*. flickr®; Creative Commons Licence: Attribution 2.0 Generic. URL: <http://www.flickr.com/photos/gidzy/3320392172/> [Date Accessed: 11th October 2009], 2009.
- G. Gilbert, D. W. Gibbons, and J. Evans. *Bird monitoring methods: a manual of techniques for key UK species*. RSPB, in association with British Trust for Ornithology, The Wildfowl & Wetlands Trust, Joint Nature Conservation Committee, Institute of Terrestrial Ecology and The Seabird Group, Sandy, Bedfordshire, 1998.
- P. J. Green and B. W. Silverman. *Nonparametric Regression and Generalized Linear Models: A roughness penalty approach*. Monographs on Statistics and Applied Probability 58. Chapman and Hall, London, 1994.
- J. J. D. Greenwood. Citizens, science and bird conservation. *Journal Of Ornithology*, 148:S77–S124, 2007.
- R. D. Gregory, P. Voříšek, D. G. Noble, A. Van Strien, A. Klvaňová, M. Eaton, A. W. Gmelig Meyling, A. Joys, R. P. B. Foppen, and I. J. Burfield. The generation and use of bird population indicators in Europe. *Bird Conservation International*, 18 (SupplementS1):S223–S244, 2008.
- R. J. Hijmans, S. E. Cameron, J. L. Parra, P. G. Jones, and A. Jarvis. Very high resolution interpolated climate surfaces for global land areas. *International Journal Of Climatology*, 25(15):1965–1978, 2005.
- R. J. Hijmans, L. Guarino, A. Jarvis, R. O'Brien, C. Bussink, P. Mathur, M. Cruz, I. Barrantes, and E. Rojas. *DIVA-GIS Version 7.1.6. A geographic information system for the analysis of species distribution data*. Manual for version 5.2 (2005) available at <http://www.diva-gis.org>, 2009.
- R. J. Hill. Measuring inflation and growth using spanning trees. *International Economic Review*, 42(1):167–185, 2001.
- R. J. Hill. When does chaining reduce the Paasche-Laspeyres spread? An application to scanner data. *Review of Income and Wealth*, 52(2):309–325, 2006a.
- R. J. Hill. Superlative index numbers: not all of them are super. *Journal of Econometrics*, 130(1):25–43, 2006b.
- A. R. Ives. Stochasticity and statisticians in environmental biology. *Trends In Ecology & Evolution*, 15(12):485–486, 2000.

- S. E. Jackson, M. Kershaw, and K. J. Gaston. Size matters: the value of small populations for wintering waterbirds. *Animal Conservation*, 7:229–239, 2004a.
- S. F. Jackson, M. Kershaw, and K. J. Gaston. The buffer effect and the selection of protected areas for waterbirds in Britain. *Biological Conservation*, 120(1):137–143, 2004b.
- S. F. Jackson, M. Kershaw, and K. J. Gaston. The performance of procedures for selecting conservation areas: waterbirds in the UK. *Biological Conservation*, 118(2):261–270, 2004c.
- S. F. Jackson, G. E. Austin, and M. J. S. Armitage. Surveying waterbirds away from major waterbodies: implications for waterbird population estimates in Great Britain. *Bird Study*, 53:105–111, 2006.
- M. Kershaw. Validation of WeBS methodology: planning study for stratified sampling programme. Technical report, WWT, 1997.
- M. Kershaw. Long-term population trends in wintering pintail (*Anas acuta*) in Great Britain 1966-95. Technical report, WWT, 1998.
- M. Kershaw and P. A. Cranswick. Numbers of wintering waterbirds in Great Britain, 1994/1995- 1998/1999: I. Wildfowl and selected waterbirds. *Biological Conservation*, 111(1):91–104, 2003a.
- M. Kershaw and P. A. Cranswick. Deriving population estimates for wintering wildfowl in Great Britain. *Ornis Hungarica*, 12-13:75–87, 2003b.
- M. Kershaw and T. M. Frost. Investigating the decline in Mallard *Anas platyrhynchos* numbers in Great Britain: an integrated approach. Unpublished report, WWT, 2006.
- M. Kery and H. Schmid. Monitoring programs need to take into account imperfect species detectability. *Basic And Applied Ecology*, 5(1):65–73, 2004.
- J. S. Kirby. Winter Population Estimates For Selected Waterfowl Species In Britain. *Biological Conservation*, 73(3):189–198, 1995.
- J. S. Kirby, R. J. Waters, and R. P. Prÿs-Jones. *Wildfowl and Wader Counts 1989-90*. WWT, Slimbridge, 1990.
- J. S. Kirby, D. G. Salmon, G. L. Atkinson-Willes, and P. A. Cranswick. Index Numbers For Waterbird Populations.3. Long-Term Trends In The Abundance Of Wintering Wildfowl In Great-Britain, 1966/67- 1991/92. *Journal Of Applied Ecology*, 32(3): 536–551, 1995.
- Laurencea. *IMG9499: not a soft landing*. (Tags: *edinburgh; musselburgh*.) flickr®; Creative Commons Licence: Attribution-No Derivative Works 2.0 Generic. URL: <http://www.flickr.com/photos/laurencea/3163736793/> [Date Accessed: 11th October 2009], 2009.

- W. A. Link and J. R. Sauer. Estimation of population trajectories from count data. *Biometrics*, 53(2):488–497, 1997.
- I. M. D. Maclean and G. E. Austin. Wetland Bird Survey Alerts 2004/2005 (Release 2). Technical report, BTO Research Report No. 492, 2008.
- I. M. D. Maclean, G. E. Austin, M. M. Rehfisch, J. Blew, O. Crowe, S. Delany, K. Devos, B. Deceuninck, K. Gunther, K. Laursen, M. Van Roomen, and J. Wahl. Climate change causes rapid changes in the distribution and site abundance of birds in winter. *Global Change Biology*, 14(11):2489–2500, 2008.
- D. Matless, P. Merchant, and C. Watkins. Animal landscapes: otters and wildfowl in England 1945-1970. *Transactions of the Institute of British Geographers*, 30(2):191–205, 2005.
- G. V. T. Matthews. *The Ramsar Convention on Wetlands: its History and Development*. Ramsar Convention Bureau, Gland, Switzerland, 1993.
- P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman and Hall, London, 2nd ed. edition, 1989.
- Met Office. *Seasonal forecasting*. URL: <http://www.metoffice.gov.uk/science/creating/monthsahead/seasonal>. [Date Accessed: 18 June 2009], 200-.
- Met Office. *MIDAS Land Surface Stations data (1853-current) [Internet]*. British Atmospheric Data Centre. URL: <http://badc.nerc.ac.uk/data/ukmo-midas>. [Date Accessed: October 2009], 2006.
- M. E. Moser. The Birds of Estuaries Enquiry. *Wader Study Group Bulletin*, 36:7–8, 1982.
- J.-P. Moussus, F. Jiguet, J. Clavel, and R. Julliard. A method to estimate phenological variation using data from large-scale abundance monitoring programmes. *Bird Study*, 56:198–212, 2009.
- A. J. Musgrove. WeBS Online is Here! *WeBS News: Newsletter of the Wetland Bird Survey*, 23, 2007.
- A. J. Musgrove, M. P. Collier, A. N. Banks, N. A. Calbrade, R. D. Hearn, and G. E. Austin. *Waterbirds in the UK 2005/06: The Wetland Bird Survey*. BTO/WWT/RSP-B/JNCC, Thetford, 2007.
- National Statistics. Consumer Price Indices First Release - September 2007. Technical report, National Statistics, 16 October 2007 2007a.
- National Statistics. Consumer Price Indices Technical Manual - 2007 Edition. Technical report, National Statistics, 2007b.
- I. Newton. Fifty years of scientific research by The Wildfowl & Wetlands Trust. *Wildfowl*, 47:1–8, 1997.

- D. R. Norris and P. P. Marra. Seasonal interactions, habitat quality, and population dynamics in migratory birds. *Condor*, 109(3):535–547, 2007.
- M. A. Ogilvie. Population changes and mortality of the mute swan in Britain. *Wildfowl*, 18:64–73, 1967.
- S. O’Grady. Double whammy for home owners; House prices slump by 0.6% - But Bank won’t cut loan rates. *Daily Express*, 5th October 2007, 2007a.
- S. O’Grady. House prices beat the gloom. *Daily Express*, 28th September 2007, 2007b.
- T. J. Osborn. *Tim Osborn: North Atlantic Oscillation index data*. URL: http://www.cru.uea.ac.uk/~timo/projpages/nao_update.htm [Date Accessed: 20/08/09], 2009.
- T. J. Osborn, K. R. Briffa, S. F. B. Tett, P. D. Jones, and R. M. Trigo. Evaluation of the North Atlantic Oscillation as simulated by a coupled climate model. *Climate Dynamics*, 15(9):685–702, 1999.
- G. P. Patil. Ranked set sampling. In A.H. El-Shaarawi and W.W. Piergorsch, editors, *Encyclopedia of Environmetrics*, volume 3, pages 1684–1690. Wiley, Chichester, 2002.
- J. W. Pearce-Higgins, D. W. Yalden, and M. J. Whittingham. Warmer springs advance the breeding phenology of golden plovers *Pluvialis apricaria* and their prey (Tipulidae). *Oecologia*, 143(3):470–476, 2005.
- J. N. Perry, I. P. Woiod, R. H. Smith, Morse, David, R. A. Desharnais, R. F. Costantino, J. M. Cushing, Dennis, and Brian. Estimating Chaos in an Insect Population. *Science*, 276(5320):1881–1882, 1997.
- M. S. Pollitt, C. Hall, S. J. Holloway, R. D. Hearn, P. Marshall, A. J. Musgrove, J. A. Robinson, and P. A. Cranswick. *The Wetland Bird Survey 2000-01: Wildfowl and Wader Counts*. BTO/WWT/RSPB/JNCC, Slimbridge, 2003.
- R. P. Prýs-Jones, L. G. Underhill, and R. J. Waters. Index Numbers For Waterbird Populations.2. Coastal Wintering Waders In The United-Kingdom, 1970/71-1990/91. *Journal Of Applied Ecology*, 31(3):481–492, 1994.
- J. L. Quinn and J. S. Kirby. An ornithological evaluation of North West Water’s reservoirs and implications for their future management. Volume 1: Main analyses. Unpublished report, WWT, 1992.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2009. URL <http://www.R-project.org>. ISBN 3-900051-07-0.
- K. Rainio, T. Laaksonen, M. Ahola, A. V. Vahatalo, and E. Lehikoinen. Climatic responses in spring migration of boreal and arctic birds in relation to wintering area and taxonomy. *Journal Of Avian Biology*, 37(5):507–515, 2006.
- J. O. Ramsay. *Matlab functional data analysis functions [Internet]*. URL: <ftp://ego.psych.mcgill.ca/pub/ramsay/FDAfuns/Matlab/> [Date Accessed: 19/08/09], 2006.

- J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer series in statistics. Springer-Verlag, New York, 2nd ed. edition, 2005.
- J. O. Ramsay, G. Hooker, and S. Graves. *Functional Data Analysis with R and MATLAB*. Springer series in statistics. Springer-Verlag, New York, 2009.
- John. O. Ramsay and B. W. Silverman. *Applied Functional Data Analysis: Methods and Case Studies*. Springer series in Statistics. Springer-Verlag, New York, 2002.
- M. M. Rehfisch, G. E. Austin, M. J. S. Armitage, P. W. Atkinson, S. J. Holloway, A. J. Musgrove, and M. S. Pollitt. Numbers of wintering waterbirds in Great Britain and the Isle of Man (1994/1995-1998/1999): II. Coastal waders (Charadrii). *Biological Conservation*, 112(3):329–341, 2003.
- J. A. Robinson and M. S. Pollitt. Sources and extent of human disturbance to waterbirds in the UK: an analysis of Wetland Bird Survey data, 1995/96 to 1998/99. *Bird Study*, 49:205–211, 2002.
- P. Rose. Manual for international waterfowl census coordinators. Technical report, IWRB, 1990.
- M. Salmon. *Climatic Research Unit: Data. North Atlantic Oscillation (NAO)*. URL: <http://www.cru.uea.ac.uk/cru/data/nao.htm> [Date Accessed: 20/08/09], 2004.
- J. L. Schafer. *Analysis of Incomplete Multivariate Data*. Monographs on statistics and applied probability. Chapman & Hall, London, 1997.
- M. S. Silver. Criteria for choosing between alternative consumer price-index number formulas with special reference to chained indexes. *The Statistician*, 33(2):229–237, 1984.
- D. Snow and C. Perrins, editors. *The Birds of the Western Palearctic*. Oxford University Press, Oxford, consise edition, 1997.
- T. H. Sparks and P. D. Carey. The Responses Of Species To Climate Over 2 Centuries - An Analysis Of The Marsham Phenological Record, 1736-1947. *Journal of Ecology*, 83(2):321–329, 1995.
- C. J. F. Ter Braak, A. J. Van Strien, R. Meijer, and T. J. Verstrael. Analysis of monitoring data with many missing values: which method? In E. J. M. Hagemeyer and T. J. Verstrael, editors, *Bird numbers 1992: Distribution, monitoring and ecological aspects. Proceedings of the 12th International Conference of IBCC and EOAC*, Noordwijkerhout, The Netherlands, 1992.
- The MathWorks. Matlab (R2008b). Computer Program, 2008.
- E. Tinkler, W. I. Montgomery, and R. W. Elwood. Foraging ecology, fluctuating food availability and energetics of wintering brent geese. *Journal of Zoology*, 278(4):313–323, 2009.

- L. G. Underhill. Indices for Waterbird Populations. Technical report, British Trust for Ornithology, 1989.
- L. G. Underhill and R. P. Prÿs-Jones. Index Numbers For Waterbird Populations.1. Review And Methodology. *Journal of Applied Ecology*, 31(3):463–480, 1994.
- United Nations. Guidelines on Principles of a System of Price and Quantity Statistics. Technical Report 59, United Nations Department of Economic and Social Affairs, 1977.
- R. J. Waters and P. A. Cranswick. *The Wetland Bird Survey 1992-93: Wildfowl and Wader Counts*. BTO/WWT/RSPB/JNCC, Slimbridge, 1993.
- WeBS workshop. Stratification and Possible Development of the Scheme. Personal Communication [Meeting]. Thetford, 25th April 2008 2008.
- C. V. Wernham, M. P. Toms, J. H. Marchant, J. A. Clark, G. M. Siriwardena, and S. R. Baillie, editors. *The Migration Atlas: movements of the birds of Britain and Ireland*. T. & A.D. Poyser, London, 2002.
- B. K. Williams, J. D. Nichols, and M. Conroy. *Analysis and Management of animal populations*. Academic Press, London, 2002.
- S.N. Wood. *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2006.
- WWT Research Staff. Discussion on impact of greater availability of optics to WeBS counters. Personal Communication [Conversation]. Slimbridge, 2006.
- A. F. Zuur, E. N. Ieno, N. J. Walker, A. A. Saveliev, and G. M. Smith. *Mixed Effects Models and Extensions in Ecology with R*. Springer, New York, 2009.