

**Seasonal Stability in Time Series of
Zooplankton Abundance**

TRAIANI STARI

DOCTOR OF PHILOSOPHY THESIS

**Department of Mathematics and Statistics
University of Strathclyde
May 2010**

Declaration

‘This thesis is the result of the authors original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.’

‘The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.’

Signed:

Date:

TRAIANI STARI

Acknowledgements

I would like to thank my first supervisor, Professor Eddie McKenzie, for his continuous support and guidance throughout my PhD years. Also, my second supervisor, Dr Steve Hay, who gave me the opportunity to analyse the Stonehaven data and introduced me to the fascinating world of Marine Biology. The very kind and helpful scientists from Marine Scotland, namely Dr. Kathryn Cook, Dr Jens Rasmussen, Dr Eileen Bresnan, Dr Ana Luisa Amorim and many more, that helped me with the interpretation of my results. My family who, though in Greece, made me feel that they were near, and my friends and colleagues who never failed to put a smile on my face in stressing times. Last, but not least, I would like to thank all the people who interacted with me while I was hyperactive due to massive amounts of caffeine.

“ Τα πάντα ρεῖ, μηδέποτε κατά τ’ αυτό μένειν. Ποταμῷ γάρ οὐκ ἔστιν εμβῆναι δῖς τῷ αὐτῷ.”

Ἡράκλειτος, ἀπό Διογένη Λαέρτιο, Ζωές Διακεκριμένων Φιλοσόφων
Ἕλληνας φιλόσοφος (540π.Χ. - 480π.Χ.)

“Nothing endures but change. You cannot step into the same river twice.”

Heraclitus, from Diogenes Laertius, Lives of Eminent Philosophers
Greek philosopher (540 BC - 480 BC)

Contents

Abstract	1
1 Introduction	3
1.1 Climate Change	3
1.2 Phenology	5
1.3 Stonehaven Ecosystem Monitoring	7
2 Seasonal Decomposition	13
2.1 Introduction to Seasonal Decomposition	13
2.2 Trending Seasonality	20
2.3 Brief Description of Seasonal Decomposition Bibliography	23
2.3.1 Weighted Moving Average - WMA	24
2.3.2 Local Regression Smoothing	25
2.3.3 Seasonality Estimation/Extraction	27
2.3.4 Census Bureau methods	27
2.3.5 A Seasonal-Trend Decomposition Based on LOESS - STL	31
2.3.6 Structural Time Series Models	34
2.3.7 TRAMO - SEATS	36
2.4 Discussion	40
3 Testing for Seasonality	41
3.1 Introduction to testing	41
3.1.1 Introduction	41

3.1.2	Preliminary Analyses – Testing for Stable Seasonality	42
3.2	(Seasonal) Unit Root Tests	48
3.2.1	Introduction	48
3.2.2	ARIMA Models	50
3.2.3	Structural Time Series Models	74
3.2.4	Discussion	83
3.2.5	Applications	85
4	Further Tests on Seasonality	90
4.1	Introduction	90
4.2	Resampling Tests for Seasonality	90
4.2.1	Non Seasonal vs Seasonal	91
4.2.2	Stable Seasonal vs Trending Seasonal	93
4.2.3	Non Seasonal vs Trending Seasonal	95
4.2.4	Applications	96
4.2.5	Discussion	97
4.3	Introduction to Generalized Additive Models	97
4.3.1	Testing with GAMs	108
4.3.2	Applications	111
4.3.3	Discussion	121
5	Comparison of Methods	123
5.1	Comparison of Methods	123
5.1.1	Stable Seasonality	126
5.1.2	Systematic Time Shift	129
5.1.3	Random Time Shift	136
5.1.4	Systematic Amplitude Change	138
5.1.5	Random Amplitude Change	140
5.1.6	Systematic Time Shift and Random Amplitude Change	142
5.1.7	Discussion of Results	146
6	Stonehaven Results	149

6.1	The Data	149
6.2	Methodologies	151
6.2.1	Resampling	151
6.2.2	GAMs	152
6.3	Environmental Series Results	153
6.3.1	Temperature	158
6.3.2	Salinity	168
6.3.3	Silicate	180
6.3.4	Phosphate	186
6.3.5	Ammonia	190
6.3.6	Nitrate	201
6.3.7	<i>Chlorophyll a</i>	207
6.4	Phytoplankton Series Results	217
6.4.1	Diatoms	219
6.4.2	Dinoflagellates	225
6.4.3	Total Phytoplankton	229
6.5	Notes on Results	233
7	Stonehaven Zooplankton Results	234
7.1	The zooplankton data	234
7.1.1	Resampling	235
7.1.2	GAMs	236
7.2	Results by species	239
7.2.1	<i>Acartia</i>	239
7.2.2	<i>Centropages</i>	249
7.2.3	<i>Pseudocalanus elongatus</i> copepodite stages 1 to 6	256
7.2.4	<i>Temora longicornis</i>	257
7.2.5	<i>Paracalanus parvus</i> copepodite stages 1 to 6	261
7.2.6	Small Calanoid Copepods copepodite stages 1 to 6	262
7.2.7	<i>Calanus finmarchicus</i>	265
7.2.8	<i>Calanus helgolandicus</i>	274

7.2.9	Juvenile <i>Calanus</i> copepodite stages 1 to 4	290
7.2.10	<i>Oithona</i>	296
7.2.11	<i>Lamellibranchiata</i>	300
7.3	Discussion	303
8	Discussion	305
8.1	Discussion on Results	305
8.2	Areas for Further Work	306
8.2.1	Statistical Interest	306
8.2.2	Biological Interest	309
	Bibliography	310

List of Tables

2.1	Types of Seasonal Patterns	21
3.1	Phillips-Perron Unit Root Test	85
3.2	HEGY Unit Root Test p-values	86
3.3	Canova-Hansen Unit Root Test p-values	86
3.4	Canova-Hansen Unit Root Test Critical Values	87
4.1	Resampling Results	96
4.2	House Data ANOVA	111
4.3	House Data AIC & BIC:	112
4.4	Unemployment Data ANOVA	115
4.5	Unemployment Data AIC & BIC:	116
4.6	Kola Data ANOVA	119
4.7	Kola Data AIC & BIC:	120
5.1	Stable Seasonality - S_{env} Resampling Results	127
5.2	Stable Seasonality - S_{dens} Resampling Results	127
5.3	Stable Seasonality - S_{env} GAMs Results	128
5.4	Stable Seasonality - S_{dens} GAMs Results	128
5.5	Systematic Time Shift by 1/10 of a week - S_{env} Resampling Results . . .	129
5.6	Systematic Time Shift by 1/10 of a week - S_{dens} Resampling Results . .	130
5.7	Systematic Time Shift by 1/40 of a month - S_{env} GAMs Results	130
5.8	Systematic Time Shift by 1/40 of a month - S_{dens} GAMs Results	130
5.9	Systematic Shift 2/10 of a week - S_{env} Resampling Results	131

5.10	Systematic Shift 2/10 of a week - S_{dens} Resampling Results	131
5.11	Systematic Shift 2/40 of a month - S_{env} GAMs Results	132
5.12	Systematic Shift 2/40 of a month - S_{dens} GAMs Results	132
5.13	Systematic Time Shift by 3/10 of a week - S_{env} Resampling Results . . .	133
5.14	Systematic Time Shift by 3/10 of a week - S_{dens} Resampling Results . .	133
5.15	Systematic Time Shift by 3/40 of a month - S_{env} GAMs Results	134
5.16	Systematic Time Shift by 3/40 of a month - S_{dens} GAMs Results	134
5.17	Systematic Time Shift by 4/10 of a week - S_{env} Resampling Results . . .	135
5.18	Systematic Time Shift by 4/10 of a week - S_{dens} Resampling Results . .	135
5.19	Systematic Time Shift by 4/40 of a month - S_{env} GAMs Results	135
5.20	Systematic Time Shift by 4/40 of a month - S_{dens} GAMs Results	136
5.21	Random Time Shift 1/10 of a week - S_{env} Resampling Results	136
5.22	Random Time Shift 1/10 of a week - S_{dens} Resampling Results	137
5.23	Random Time Shift 1/40 of a month - S_{env} GAMs Results	137
5.24	Random Time Shift 1/40 of a month - S_{dens} GAMs Results	137
5.25	Systematic Amplitude Change - S_{env} Resampling Results	138
5.26	Systematic Amplitude Change - S_{dens} Resampling Results	138
5.27	Systematic Amplitude Change - S_{env} GAMs Results	139
5.28	Systematic Amplitude Change - S_{dens} GAMs Results	139
5.29	Random Amplitude Change - S_{env} Resampling Results	140
5.30	Random Amplitude Change - S_{dens} Resampling Results	140
5.31	Random Amplitude Change - S_{env} GAMs Results	141
5.32	Random Amplitude Change - S_{dens} GAMs Results	141
5.33	Systematic Time Shift 1/10 and Random Amplitude Change - S_{env} Re- sampling Results	142
5.34	Systematic Time Shift 1/10 and Random Amplitude Change - S_{dens} Re- sampling Results	143
5.35	Systematic Time Shift 1/40 and Random Amplitude Change - S_{env} GAMs Results	143
5.36	Systematic Time Shift 1/40 and Random Amplitude Change - S_{dens} GAMs Results	143

5.37	Systematic Time Shift 2/10 and Random Amplitude Change - S_{env} Re-sampling Results	144
5.38	Systematic Time Shift 2/10 and Random Amplitude Change - S_{dens} Re-sampling Results	144
5.39	Systematic Time Shift 2/40 and Random Amplitude Change - S_{env} GAMs Results	145
5.40	Systematic Time Shift 2/40 and Random Amplitude Change - S_{dens} GAMs Results	145
6.1	Environmental Series Results	157
6.2	Temperature 1m ANOVA	159
6.3	Temperature 1m AIC & BIC	161
6.4	Temperature 1m ANOVA - Cyclic Component	165
6.5	Temperature 1m AIC & BIC - Cyclic Component	166
6.6	Salinity 1m ANOVA	169
6.7	Salinity 1m AIC & BIC	170
6.8	Salinity 45m ANOVA	174
6.9	Salinity 45m AIC & BIC	175
6.10	Silicate 1m ANOVA	182
6.11	Silicate 1m AIC & BIC	183
6.12	Phosphate 1m ANOVA	186
6.13	Phosphate 1m AIC & BIC	187
6.14	Ammonia 1m ANOVA	190
6.15	Ammonia 1m AIC & BIC	191
6.16	Ammonia 45m ANOVA	195
6.17	Ammonia 45m AIC & BIC	195
6.18	Nitrate 1m ANOVA	201
6.19	Nitrate 1m AIC & BIC	202
6.20	<i>Chlorophyll</i> ANOVA	208
6.21	<i>Chlorophyll</i> AIC & BIC	209
6.22	<i>Chlorophyll</i> AIC & BIC	214

6.23	<i>Chaetoceros</i> ANOVA	220
6.24	<i>Chaetoceros</i> AIC & BIC	221
6.25	Total Diatoms ANOVA	223
6.26	Total Diatoms AIC & BIC	224
6.27	<i>Alexandrium</i> ANOVA	226
6.28	<i>Alexandrium</i> AIC & BIC	227
6.29	Total Dinoflagellates ANOVA	228
6.30	Total Dinoflagellates AIC & BIC	229
6.31	Total Phytoplankton ANOVA	230
6.32	Total Phytoplankton AIC & BIC	231
7.1	Zooplankton Series Results	238
7.2	<i>Acartia clausi</i> C4 ANOVA	240
7.3	<i>Acartia clausi</i> C4 AIC & BIC	241
7.4	<i>Acartia clausi</i> C6m ANOVA	244
7.5	<i>Acartia clausi</i> C6m AIC & BIC	245
7.6	Total <i>Acartia</i> ANOVA	247
7.7	Total <i>Acartia</i> AIC & BIC	248
7.8	<i>C. typicus</i> C1-4 ANOVA	250
7.9	<i>C. typicus</i> C1-4 AIC & BIC	250
7.10	<i>C. hamatus</i> C4 ANOVA	253
7.11	<i>C. hamatus</i> C4 AIC & BIC	254
7.12	<i>T. longicornis</i> C1 ANOVA	257
7.13	<i>Temora</i> C1 AIC & BIC	258
7.14	Small Calanoid Copepods C1-6 ANOVA	262
7.15	Small Calanoid Copepods C1-6 AIC & BIC	263
7.16	<i>C. finmarchicus</i> C5 ANOVA	266
7.17	<i>C. finmarchicus</i> C5 AIC & BIC	270
7.18	<i>C. finmarchicus</i> C6f ANOVA	271
7.19	<i>C. finmarchicus</i> C6f AIC & BIC	272
7.20	<i>C. helgolandicus</i> C5 ANOVA	274

7.21	<i>C. helgolandicus</i> C5 AIC & BIC	276
7.22	<i>C. helgolandicus</i> C6f ANOVA	280
7.23	<i>C. helgolandicus</i> C6f AIC & BIC	281
7.24	<i>C. helgolandicus</i> C6m ANOVA	285
7.25	<i>C. helgolandicus</i> C6m AIC & BIC	286
7.26	Juvenile <i>Calanus</i> C1-4 ANOVA	290
7.27	Juvenile <i>Calanus</i> C1-4 AIC & BIC	291
7.28	<i>Oithona</i> C1-6 ANOVA	297
7.29	<i>Oithona</i> C1-6 AIC & BIC	299
7.30	<i>Lamellibranchiata</i> ANOVA	300
7.31	<i>Lamellibranchiata</i> AIC & BIC	301

List of Figures

1.1	a)Map of Scotland showing Stonehaven. b)The Stonehaven long-term Ecosystem monitoring station. From FRS material.	8
1.2	The circulation patterns of North Sea water. The width of arrows is indicative of volume transport. Red arrows indicate relatively pure Atlantic water. From ICES (2006).	9
1.3	The food web. From FRS material.	11
2.1	The House series.	16
2.2	The House series with the centred moving average estimated trend, the de-trended series with marked Decembers, the stable seasonals and the residuals.	17
2.3	The Unemployment series.	18
2.4	The Unemployment series with the centred moving average estimated trend, the de-trended series, the stable seasonals and the residuals. . . .	18
2.5	The Kola series.	19
2.6	The Kola series with the centred moving average estimated trend, the de-trended series, the stable seasonals and the residuals.	20
2.7	A systematic (increasing) amplitude change, a fluctuation in amplitude, a systematic (earlier) shift in time and a fluctuation in time, using a period of 52 weeks. In the bottom two plots the dotted lines mark the 30th week.	22
3.1	The House, stable seasonal, series estimated p-values.	43
3.2	The Unemployment, trending seasonal, series estimated p-values.	44

3.3	The Kola, stable seasonal, series estimated p-values.	44
3.4	The House estimated p-values Q-Q Plot.	45
3.5	The Unemployment estimated p-values Q-Q Plot.	46
3.6	The Kola estimated p-values Q-Q Plot.	46
3.7	A diagram summarizing the Unit Root Tests that assume an ARIMA structure.	88
3.8	A diagram summarizing the Unit Root Tests that use structural time series models.	89
4.1	The upper left plot shows the House data with the trend from the stable GAM in red. The upper right plot shows the stable seasonals estimated from the stable GAM while in the lower plot the House data are plotted as dots with the fits from the stable (red) and trending (blue) GAMs. . .	112
4.2	The (twelve) trending seasonals from the TS model of the House data with confidence intervals (dotted lines), by row from top left.	113
4.3	Residuals plots from the NS model for the House data. From the upper left corner clockwise: residual QQ plot, residuals plotted against the response, the residuals themselves and the residuals' histogram.	114
4.4	Diagnostic plots for the TS model for the House data.	114
4.5	The diagnostic plots for the SS GAM fitted to the House data.	115
4.6	The Unemployment data with the trend from the SS model in red, the stable seasonals and the data as dots with stable (red) and trending (blue) GAM fits are depicted clockwise from upper left.	116
4.7	The smooth functions of the TS model of the Unemployment data, corresponding to the twelve months.	117
4.8	Diagnostic plots for the TS model of the Unemployment data. Clockwise from top left: a QQ plot of the residuals, the residuals plotted against the response, the residuals plot and histogram.	117
4.9	Diagnostic plots for the TS model with a larger basis dimension, 12, of the Unemployment data. Clockwise from top left: a QQ plot of the residuals, the residuals plotted against the response, the residuals plot and histogram.	118

4.10	The upper left plot depicts the Kola data with the trend from the SS GAM in red. The stable seasonals and the data as dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.	119
4.11	The smooth trending seasonals estimated in the TS model for the Kola data.	120
4.12	Here the residual checking plots for the stable GAM of the Kola data. . .	121
5.1	The seasonal patterns. From the top the pattern corresponding to the environmental data, the mean used for the Negative Binomial distribution and the pattern corresponding to densities, derived as a random Negative Binomial variable with the previous plotted mean.	124
5.2	The computed correlations between 1 : 11 and the 1000 randomised ordered sets of them. The dotted lines mark 0.7 and -0.7	126
6.1	The upper left plot depicts the Temperature at 1m with the trend from the SS GAM as a red line. The stable seasonals and the Temperature 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.	159
6.2	The upper shows the stable seasonals estimated by the SS GAM when using factors while the bottom one when using a cyclic component. . . .	160
6.3	The smooth trending seasonals for weeks 1-36 estimated in the TS model for Temperature 1m.	161
6.4	The smooth functions fitted by the TS GAM, corresponding to the trending seasonals for weeks 37-52 for the Temperature 1m.	162
6.5	The stable seasonal pattern (red) and the trending seasonal pattern (blue) from the TS GAM fitted to the Temperature at 1m series.	162
6.6	The residual checking plots for the SS GAM of the Temperature 1m data. No problems are highlighted.	164
6.7	Residual checking plots for the TS GAM of the of the Temperature 1m data. No violations are noted.	164

6.8	The upper left plot depicts the Temperature at 1m with the trend from the SS GAM as a red line. The stable seasonals and the Temperature 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise. Both GAMs use a cyclic component for stable seasonality.	165
6.9	The residual checking plots for the SS GAM (using a cyclic component) of the Temperature 1m data. No problems are highlighted.	166
6.10	Residual checking plots for the TS GAM (using a cyclic component) of the of the Temperature 1m data. No violations are noted.	167
6.11	The upper left plot depicts Salinity at 1m with the SS model's trend as a red line. The stable seasonals and the Salinity 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.	169
6.12	The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Salinity 1m data.	170
6.13	The smooth trending seasonals estimated in the TS model for the Salinity 1m series, here weeks 37-52.	171
6.14	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Salinity at 1m series.	171
6.15	The residual checking plots for the SS GAM of Salinity 1m show a deviation from normality.	172
6.16	Plots of the standardized deviance residuals from the TS GAM of Salinity 1m show a deviation from normality.	173
6.17	Clockwise from top left the Salinity 45m series with the trend from the SS GAM (red line), the stable seasonals and the Salinity 45m series (dotted) with the stable (red) and trending (blue) fitted GAMs.	175
6.18	The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Salinity 45m data.	176
6.19	The smooth trending seasonals estimated in the TS model for the Salinity 45m, weeks 37-52.	176
6.20	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Salinity at 45m series.	177

6.21	The de-trended fitted values from TS GAM for Salinity 45m. The vertical lines mark weeks 10 and 30. The peak appears to shift earlier in the year.	178
6.22	The residual checking plots for the SS GAM of Salinity 45m show no violations.	178
6.23	The standardized deviance residuals' checking plots from the TS GAM of Salinity 45m show no violations of the assumptions of the model.	179
6.24	The top plot shows the mean value of the Silicate at 1m series for each year against the range of the series for that year. The bottom plot shows the fitted values from the SS GAM plotted against the absolute residuals from that model with a smoother running through them, red line.	181
6.25	Clockwise from top left the Silicate 1m series with the trend from the SS GAM, the stable seasonals and Silicate 1m (dotted) with the stable (red) and trending (blue) fitted GAMs.	182
6.26	The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Silicate 1m data.	183
6.27	The smooth trending seasonals estimated in the TS model for the Silicate 1m, weeks 37-52.	183
6.28	The residual checking plots for the SS GAM of Silicate 1m show no violations.	184
6.29	The standardized deviance residuals' checking plots from the TS GAM of Silicate 1m show no violations of the assumptions of the model.	185
6.30	The Phosphate 1m series is plotted in the top left with the trend from the SS GAM (red line). The stable seasonals and the Phosphate 1m (dotted) with the stable (red) and trending (blue) fitted GAMs follow clockwise.	187
6.31	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Phosphate at 1m series.	188
6.32	The residual checking plots for the SS GAM of Phosphate 1m show no violations.	189
6.33	The standardized deviance residuals' checking plots from the TS GAM of Phosphate 1m show no violations of the assumptions of the model.	189

6.34	Clockwise from top left the Ammonia 1m series with the trend from the SS GAM marked by a red line, the stable seasonals and Ammonia 1m (dotted) with the stable (red) and trending (blue) fitted GAMs.	191
6.35	The smooth fitted functions for weeks 1-36, in the TS model for the Ammonia 1m data.	192
6.36	The smooth trending seasonals estimated in the TS model for the Ammonia 1m, weeks 37-52.	192
6.37	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Ammonia at 1m series.	193
6.38	The residual checking plots for the SS GAM of Ammonia 1m.	194
6.39	The residual checking plots from the TS GAM of Ammonia 1m.	194
6.40	Clockwise from top left the Ammonia 45m series with the trend from the SS GAM, the stable seasonals and Ammonia 45m (dotted) with the stable (red) and trending (blue) fitted GAMs.	196
6.41	The smooth fitted functions for weeks 1-36, in the TS model for the Ammonia 45m data.	197
6.42	The smooth trending seasonals estimated in the TS model for the Ammonia 45m, weeks 37-52.	197
6.43	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Ammonia at 45m series.	198
6.44	The de-trended fitted values from the TS GAM fitted to the Ammonia 45m series. The dotted lines mark weeks 20 and 50.	198
6.45	The residual checking plots for the SS GAM of Ammonia 45m.	199
6.46	The standardized deviance residuals' checking plots from the TS GAM of Ammonia 45m.	200
6.47	Clockwise from top left the Nitrate 1m series with the trend from the SS GAM, the stable seasonals and the Nitrate 1m series marked with dots with the stable (red) and trending (blue) fitted GAMs.	202
6.48	The smooth fitted functions for weeks 1-36, in the TS model for the Nitrate 1m data.	203

6.49	The smooth trending seasonals estimated in the TS model for the Nitrate 1m, weeks 37-52.	203
6.50	The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Nitrate at 1m series.	204
6.51	The de-trended fitted values from TS GAM for Nitrate 1m. The vertical lines mark weeks 20 and 50 in black and red, respectively.	205
6.52	The residual checking plots for the SS GAM of Nitrate 1m.	205
6.53	The standardized deviance residuals' checking plots from the TS GAM of Nitrate 1m.	206
6.54	Clockwise from top left the <i>Chlorophyll</i> series with the trend from the SS GAM, the stable seasonals and <i>Chlorophyll</i> (dotted) with the stable (red) and trending (blue) fitted GAMs.	208
6.55	The smooth fitted functions for weeks 1-36, in the TS model for the <i>Chlorophyll</i> data.	209
6.56	The smooth trending seasonals estimated in the TS model for <i>Chlorophyll</i> , weeks 37-52.	210
6.57	The stable seasonals, red, and the trending seasonals, blue, from the SS and TS GAM, respectively, fitted to the <i>Chlorophyll</i> series.	211
6.58	The de-trended fitted values from TS GAM for <i>Chlorophyll a</i> . Weeks 18 and 42 are marked by vertical lines in black and red, respectively.	211
6.59	The residual checking plots for the SS GAM of <i>Chlorophyll</i> show a deviation from normality at the extremes of the series.	212
6.60	The standardized deviance residuals' checking plots from the TS GAM of <i>Chlorophyll</i> show a deviation from normality at the lower end of the series.	212
6.61	Clockwise from top left the <i>Chlorophyll</i> series with the trend from the SS GAM, the stable seasonals and <i>Chlorophyll</i> (dotted) with the SS (red) and TS (blue) fitted GAMs. The TS GAM does not include the stable cyclic smoother.	213
6.62	The smooth fitted functions for weeks 1-36, in the TS model (without the stable cyclic smoother) for the <i>Chlorophyll</i> data.	214

6.63	The smooth trending seasonals estimated in the TS model (without the stable cyclic smoother) for <i>Chlorophyll</i> , weeks 37-52.	215
6.64	The standardized deviance residuals' checking plots from the TS GAM of <i>Chlorophyll</i> show a deviation from normality at the lower end of the series.	216
6.65	Clockwise from top left <i>Chaetoceros</i> with the trend (red line) from the SS model, the stable seasonals and <i>Chaetoceros</i> (dotted) with the stable (red) and trending (blue) fitted GAMs.	221
6.66	The Logged series of Total Diatoms, after adding 1.	222
6.67	Clockwise from top left the logged Total Diatoms series with the trend (red line) from the SS GAM, the stable seasonals and the series (dotted) with the stable (red) and trending (blue) fitted GAMs.	223
6.68	Diagnostic plots for the Total Diatoms SS model.	224
6.69	Diagnostic check for the Total Diatoms TS model.	225
6.70	The <i>Alexandrium</i> series with the trend (red line) from the SS GAM, the stable seasonals and the series as dots with the stable (red) and trending (blue) fitted GAMs.	226
6.71	The Total Dinoflagellates series with the trend (red line) from the SS GAM, the stable seasonals and the Total Dinoflagellates series marked by dots with the stable (red) and trending (blue) fitted GAMs.	228
6.72	The logged series of Total Phytoplankton.	229
6.73	Clockwise from top left the Total Phytoplankton series with the trend (red line) from the SS GAM, the stable seasonals and the logged series of Total Phytoplankton marked by dots with the stable (red) and trending (blue) fitted GAMs.	230
6.74	Diagnostic check for the Total Phytoplankton SS model.	231
6.75	Diagnostic check for the Total Phytoplankton TS model.	232
7.1	The logged <i>Acartia clausi</i> C4 series with the trend from the SS model, the stable seasonals and the series as dots with the fits from the stable (red) and trending (blue) GAMs. Dotted lines denote the years.	241

7.2	Diagnostic plots for the SS GAM for <i>Acartia clausi</i> 4. The QQ plot shows a slight problem with the Gaussian assumption at the extremes of the series.	242
7.3	Diagnostic plots for the TS GAM for <i>Acartia clausi</i> 4. The QQ plot shows a small problem with the Gaussian assumption at the extremes of the series.	243
7.4	Clockwise from top left corner, the logged <i>Acartia clausi</i> C6m with the trend from the SS model, the stable seasonals and the series as dots with the fitted values from the stable (red) and trending (blue) GAMs.	244
7.5	The smooth fitted functions for weeks 1-36, in the TS model for the <i>Acartia clausi</i> C6m data.	245
7.6	The smooth fitted functions for weeks 37-52, in the TS model for the <i>Acartia clausi</i> C6m data.	246
7.7	The stable seasonals, red, and the trending seasonals, blue, from the SS and TS GAMs, respectively, fitted to the <i>Acartia clausi</i> C6m series. . . .	247
7.8	Total <i>Acartia</i> logged densities with the trend, as a red line, and seasonals from the SS GAM. The lower plot shows the series as dots with the fits from the SS (red) and TS (blue) GAMs.	248
7.9	<i>C. typicus</i> C1-4 with the trend from the SS model, the stable seasonals and the series (marked as dots) with the two fits from the stable (red) and trending (blue) GAMs.	251
7.10	<i>C. typicus</i> C1-4 diagnostic plots for the SS GAM.	252
7.11	<i>C. typicus</i> C1-4 diagnostic plots for the TS GAM.	252
7.12	Clockwise from top left: <i>C. hamatus</i> C4 with the trend from the SS model, the stable seasonals and the series, marked as dots, with the fit from the stable (red) and trending (blue) GAMs.	254
7.13	<i>C. hamatus</i> C4 diagnostic plots for the SS GAM.	255
7.14	<i>C. hamatus</i> C4 diagnostic plots for the TS GAM.	255
7.15	Clockwise from top left: the logged <i>Pseudocalanus elongatus</i> C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs. . .	256

7.16	From top left, clockwise we see the logged series of <i>T. longicornis</i> C1 with the trend from the SS model, the stable seasonals and the logged series marked as dots with the fitted values from the SS (red) and TS (blue) GAMs.	258
7.17	The diagnostic plots for the SS model of <i>T. longicornis</i> C1.	259
7.18	The diagnostic plots for the TS model of <i>T. longicornis</i> C1.	259
7.19	Clockwise from top left: the logged Total <i>T. longicornis</i> with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS(blue) GAMs.	260
7.20	Clockwise from top left: the logged <i>Paracalanus parvus</i> C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS(blue) GAMs.	261
7.21	The logged series of Small Calanoid Copepods C1-6 with the trend from the SS model, the stable seasonals and the logged series plotted as dots with the fitted values from the SS (red) and TS (blue) GAMs.	263
7.22	The diagnostic plots for the SS model of Small Calanoid CopepodsC1-6.	264
7.23	The diagnostic plots for the TS model of Small Calanoid CopepodsC1-6.	264
7.24	The logged series of <i>C. finmarchicus</i> C5 with the trend from the SS model, the stable seasonals and the series as dots with the fit from the SS (red) and TS (blue) GAMs.	267
7.25	The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of <i>C. finmarchicus</i> C5.	268
7.26	The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of <i>C. finmarchicus</i> C5.	268
7.27	The stable (red) and trending (blue) seasonal patterns from the SS and TS GAMs, respectively, for <i>C. finmarchicus</i> C5.	269
7.28	The de-trended fitted values from the TS GAM for <i>C. finmarchicus</i> C5. Weeks 17, 22, 34 and 39 are marked by dotted lines in black, red, green and blue, respectively.	269
7.29	The diagnostic plots for the SS GAM of <i>C. finmarchicus</i> C5.	270
7.30	The diagnostic plots for the TS GAM of <i>C. finmarchicus</i> C5.	271

7.31	The series of <i>C. finmarchicus</i> C6f with the trend from the SS model, the stable seasonals and the series marked by dots with the fit from the SS (red) and TS (blue) GAMs, appear clockwise from top left.	272
7.32	Clockwise from top left: logged <i>C. finmarchicus</i> C6m with the trend from the SS model, the stable seasonals and the series as dots with the fitted values from the SS (red) and TS (blue) GAMs.	273
7.33	Clockwise form top left we see the logged <i>C. helgolandicus</i> C5 series with the trend from the SS model, the stable seasonals and the logged series as dots with the fitted values of the SS (red) and TS (blue) GAMs. . . .	275
7.34	The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of <i>C. helgolandicus</i> C5.	276
7.35	The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of <i>C. helgolandicus</i> C5.	277
7.36	The stable, red, and trending, blue, seasonals estimated by the SS and TS GAMs, respectively, for the <i>C. helgolandicus</i> C5 series.	278
7.37	The de-trended fitted values from the TS model for the <i>C. helgolandicus</i> C5 series. Marked by the black, red, green and blue line are, respectively, weeks 16, 22, 37 and 46.	278
7.38	The diagnostic plots for the SS model of <i>C. helgolandicus</i> C5.	279
7.39	The diagnostic plots for the TS model of <i>C. helgolandicus</i> C5.	279
7.40	The logged <i>C. helgolandicus</i> C6f series with the trend from the SS model, the stable seasonals and the series, this time marked by dots, with the fitted values from the SS (red) and TS (blue) GAMs.	281
7.41	The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of <i>C. helgolandicus</i> C6f.	282
7.42	The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of <i>C. helgolandicus</i> C6f.	282
7.43	The stable (red) and trending (blue) seasonals from the SS and TS models, respectively, for the <i>C. helgolandicus</i> C6f series.	283

7.44	The de-trended fitted values from the TS model for the <i>C. helgolandicus</i> C6f series. Marked by the black and red lines are, respectively, weeks 30 and 46. The pattern shifts earlier in the year.	283
7.45	The diagnostic plots for the SS model of <i>C. helgolandicus</i> C6f.	284
7.46	The diagnostic plots for the TS model of <i>C. helgolandicus</i> C6f.	284
7.47	The logged <i>C. helgolandicus</i> C6m series with the trend from the SS model, the stable seasonals and the series, this time marked by dots, with the fitted values from the SS (red) and TS (blue) GAMs.	286
7.48	The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of <i>C. helgolandicus</i> C6m.	287
7.49	The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of <i>C. helgolandicus</i> C6m.	287
7.50	The stable (red) and trending (blue) seasonal patterns estimated by the SS and TS GAMs, respectively, for the <i>C. helgolandicus</i> C6m series. . . .	288
7.51	The de-trended fitted values from the TS model for the <i>C. helgolandicus</i> C6m series. Marked by the black and the red lines are, respectively, weeks 31 and 41.	288
7.52	The diagnostic plots for the SS model of <i>C. helgolandicus</i> C6m.	289
7.53	The diagnostic plots for the TS model of <i>C. helgolandicus</i> C6m.	289
7.54	Clockwise from top left: the logged Juvenile <i>Calanus</i> C1-4 with the trend from the SS model, the stable seasonals and the logged series plotted as dots with the fitted values of the SS (red) and TS (blue) models.	291
7.55	The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of Juvenile <i>calanus</i> C1-4.	292
7.56	The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of Juvenile <i>calanus</i> C1-4.	292
7.57	The stable (red) and trending (blue) seasonal patterns from the SS and TS models, respectively, for the Juvenile <i>Calanus</i> C1-4.	293
7.58	The de-trended fitted values from the TS model for the Juvenile <i>Calanus</i> C1-4. Vertical lines mark weeks 14, 23, 32 and 41 in colours black, red, green and blue, respectively.	294

7.59	The diagnostic plots for the SS model of Juvenile <i>Calanus</i> C1-4.	294
7.60	The diagnostic plots for the TS model of Juvenile <i>Calanus</i> C1-4.	295
7.61	Clockwise from top left: the logged <i>Oithona</i> C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs.	297
7.62	Diagnostic plots for the stable seasonal GAM for <i>Oithona</i> C1-6.	298
7.63	Diagnostic plots for the trending seasonal GAM for <i>Oithona</i> C1-6.	298
7.64	Clockwise from top left: the logged <i>Lamellibranchiata</i> with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs.	301
7.65	Diagnostic plots for the SS GAM for <i>Lamellibranchiata</i>	302
7.66	Diagnostic plots for the TS GAM for <i>Lamellibranchiata</i>	302

Declaration

‘This thesis is the result of the authors original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.’

‘The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.’

Signed:

Date:

TRAIANI STARI

Acknowledgements

I would like to thank my first supervisor, Professor Eddie McKenzie, for his continuous support and guidance throughout my PhD years. Also, my second supervisor, Dr Steve Hay, who gave me the opportunity to analyse the Stonehaven data and introduced me to the fascinating world of Marine Biology. The very kind and helpful scientists from Marine Scotland, namely Dr. Kathryn Cook, Dr Jens Rasmussen, Dr Eileen Bresnan, Dr Ana Luisa Amorim and many more, that helped me with the interpretation of my results. My family who, though in Greece, made me feel that they were near, and my friends and colleagues who never failed to put a smile on my face in stressing times. Last, but not least, I would like to thank all the people who interacted with me while I was hyperactive due to massive amounts of caffeine.

“ Τα πάντα ρεῖ, μηδέποτε κατά τ’ αυτό μένειν. Ποταμῷ γάρ οὐκ ἔστιν εμβῆναι δῖς τῷ αὐτῷ.”

Ἡράκλειτος, ἀπό Διογένη Λαέρτιο, Ζωές Διακεκριμένων Φιλοσόφων
Ἑλληνας φιλόσοφος (540π.Χ. - 480π.Χ.)

“Nothing endures but change. You cannot step into the same river twice.”

Heraclitus, from Diogenes Laertius, Lives of Eminent Philosophers
Greek philosopher (540 BC - 480 BC)

Abstract

We investigate ways to test for various aspects of seasonality in marine time series. In particular, we analyze series from the Stonehaven Ecosystem long-term monitoring station, run by Marine Scotland, Aberdeen. Our main focus is in identifying any systematic changes in seasonality as these may be attributed to climate change. We explore existing methodologies and introduce simple and robust Resampling tests. Generalized Additive Models (GAMs) are used to seasonally decompose each series and test the significance of individual components.

We study the seasonal patterns of some environmental series but also phytoplankton and zooplankton species series from Stonehaven. Climate change will have a direct effect on phytoplankton communities as they are influenced by changes in the currents, the water temperature and salinity, as well as the availability of nutrients. The zooplankton species feed on the phytoplankton biomass and are themselves prey to larger fish and other top predators. Evaluating the seasonal patterns of phytoplankton and zooplankton populations is in itself important because plankton species are at the heart of the marine food web.

Of the environmental series from Stonehaven which we examined, Salinity, Ammonia and Nitrate have patterns that are systematically shifting earlier in the year. *Chlorophyll a*, which is an indicator of the phytoplankton biomass, has a pattern whose amplitude is gradually decreasing while it is also shifting slightly earlier in the year. Amongst the zooplankton series analyzed, *Acartia clausi* copepodite stage

6 males (C6m), *Calanus finmarchicus* C5 and *Calanus helgolandicus* C5 have patterns systematically shifting earlier in time. *Calanus helgolandicus* C6f and C6m and Juvenile *Calanus* copepods C1 - 4 have seasonal patterns whose amplitude is systematically increasing as well as shifting earlier in the year. These zooplankton species are important because of their role as prey in the life cycle of the most commercially important fish. Our results indicate the nature and magnitude of the effect that climate change has on marine life and provide a basis for further analyses.

Chapter 1

Introduction

1.1 Climate Change

In the summer of 2003 a heat wave was responsible for many deaths in Europe. Anderson & Bausch (2006), in a briefing note on climate change and natural disasters to the European Parliament, report that at least 22,146 people lost their lives due to the unprecedented heat wave. In the same note they mention a number of other natural disasters that have struck Europe and the rest of the world and are linked to climate change.

Climate change is generally regarded as the single most serious problem facing the world today. Climate is a generic name for the average weather conditions experienced over a long period of time. The climate is always changing as a result of various natural causes but it is believed that during the last century human factors have been responsible for other, greater than natural, climate change. Over the last hundred years, the temperature of the Earth has risen by 0.74°C , while 0.4°C of this warming has occurred since the 1970s (Department for Environment, Food & Rural Affairs, DEFRA, 2008).

In the UK alone, many changes over recent years have been observed. According to the Inter-Agency Committee on Marine Science and Technology (IACMST) some notable examples include:

- The annual mean Central England Temperature has increased by 0.5°C during the last century.
- The mean of annual mean temperatures over 30 years in Northern Ireland and Scotland increased by 0.3°C from 1873-1902 to 1961-1990.
- There is a tendency towards wetter winters in north-east England and drier summers in south-east England.
- Mean annual sea surface temperature of the UK coastline has increased by 0.5°C .
- The mean sea level around the UK coast has increased by about 1 mm per year during the last century.

These are all considered manifestations of climate change by IACMST (2005). The phrase 'climate change' is used to describe the man-attributed, rather than natural change in the climatological conditions.

A major contributor towards climate change is the usage of fossil fuels in human societies. Fossil fuels, such as coal and oil, are used in homes, factories, trains, planes and cars. These emit carbon dioxide which is considered a harmful 'greenhouse' gas. Greenhouse gases are the gases that when released into the atmosphere surrounding the Earth, keep it warm. Due to the large scale deforestation since the Industrial Revolution there are not as many trees to absorb the increasing quantities of this extra carbon dioxide. Additionally, other types of greenhouse gases are released into the atmosphere and thus, the mixture of gases in earth's atmosphere is changing. This greenhouse effect is very important when we talk about climate change as these greenhouse gases prohibit the Earth from cooling. It is these extra greenhouse gases which humans have released that are thought to pose the greatest threat to

the environment and the current climatological conditions, according to the UK Phenology Network (2009). Nature's Calendar, the name of the web-site of the UK's Phenology group, has a good discussion and material on climate change.

Eutrophication, defined as 'an increase in the rate of supply of organic matter to an ecosystem' by Nixon (1995), is considered a global concern, too, and is attributed to man-made causes. Eutrophication is partly a consequence of increased system productivity due to nutrient enrichment, say Carpenter *et al.* (1998) and triggers significant changes in coastal ecosystems. The presence of eutrophication has many undesirable outcomes such as toxic algal blooms and low dissolved-oxygen levels (hypoxia) with disastrous effects on the ecosystem (National Research Council, 2000).

Nonetheless, it is not only the phenomenon of climate change itself that is worrying but also the ways these changes in climate influence the course of life on earth. Plant life is affected greatly by the existing climatological conditions and thus, a change in them would result in a change in the plant life (Sherry *et al.*, 2007). Animal and human life are in turn affected directly and indirectly (through plants) by climate change. Therefore, climate change is not to be taken lightly. In modern society, scientists and politicians are concerned with finding ways to tackle this. A first step in battling climate change is assessing the current state. A lot of research has been done to estimate the extent of change by examining indicators of seasonal changes, as they occur in nature, (IACMST, 2005; Heath *et al.*, 1999; Menzel *et al.*, 2005). In science, an event is considered seasonal when it is periodic, i.e. it repeats itself after a set period of time. Thus, by examining the timing and succession of seasons in nature one can monitor changes in the climate.

1.2 Phenology

This interest in the succession of seasons is not new. In Japan and China the time of blossoming of some trees (e.g. cherry trees) is associated with ancient festivals and thus, some of these dates can be traced back to the eighth century. The study of

the times of recurring natural phenomena, especially in relation to climate, is called phenology. Examples include recording when the first cuckoo was heard each year or when the blackthorn blossom was first spotted. Robert Marsham was Britain's first phenologist and started recording his 'Indications of Spring' as early as 1736. The first individual records found are from a Weather Diary from Egioke (near Redditch) from March 1703 (UK Phenology Network, 2009).

These individual records can then be compared with other ones allowing us to compare, for example, the rates of change in the succession of seasons. The UK Phenology Network (UKPN) is run by the Woodland Trust and the Centre for Ecology and Hydrology. It has over 25,000 volunteers who send in observations from their local areas. From their web-site, Nature's Calendar, the UKPN urges observers to record a variety of events by the time of year. These include, for example, the sighting of specific, migrating, bird species or the blossoming of specific plants.

In addition, many scientific monitoring projects / stations have been running since the last century. For example, the Sir Alister Hardy Foundation for Ocean Science (SAHFOS) runs the Continuous Plankton Recorder (CPR) Survey, a marine monitoring programme that has been collecting data from the North Atlantic and the North Sea on the ecology and biogeography of plankton since 1931. Another example is Plymouth's Marine Laboratory which is involved in various monitoring projects including the CPR Survey and long-term and decadal time series monitoring stations in the English Channel (Plymouth Marine Sciences Partnership, 2009).

As climate change has gained much publicity in recent years so has phenology (Cook *et al.*, 2005; Bell, 2009). Bell (2009) talks about the history of phenology and presents findings from her analysis of the data on flowering at the Kew Gardens. She notes shifts, earlier in the year, in the flowering date of many species and draws attention to the consequences of climate change.

An interesting neologism in phenology is the phrase 'season creep' used to refer to observed changes in the timing of seasons. A gradual shift of a season to earlier or later in the year is called season creep. It is especially used to describe the movement

of spring to earlier in the year due to climate change, (McFedries, 2006). We will refer to season creep later when we try to detect the effects of climate change in marine data sets.

1.3 Stonehaven Ecosystem Monitoring

Marine life is affected by climate change and it is the possibility of detecting and identifying season creep in marine life that is the main focus of this thesis. We are investigating ways to evaluate seasonal stability in data from the Stonehaven long-term Ecosystem monitoring station, located in the eastern North Sea ($56^{\circ}57.8'N$ $02^{\circ}06.2'W$), in a water depth of around 50m, see Figure 1.1. This monitoring station is run by the Fisheries Research Services (FRS) in Aberdeen¹.

The Stonehaven data consist of weekly samples (weather permitting) since January 1997. The database includes measurements of environmental variables (such as temperature and salinity) and nutrients (e.g. nitrate) but also counts of many phytoplankton and zooplankton species.

Analyzing the data, we examine the nature of seasonal patterns focusing mainly on zooplankton time series. This relates to other work at various monitoring sites around Scotland, the UK and the North Atlantic generally (Beaugrand & Reid, 2003; Heath *et al.*, 1999). In Figure 1.2, we see the ocean currents in the North Sea. Stonehaven is in the northwestern North Sea and is influenced by the inflow of mixed Scottish coastal and oceanic water that flows south along the Scottish east coast.

Furthermore, studying the seasonal patterns of phytoplankton and zooplankton populations is important because plankton species are at the heart of the marine food web, see Figure 1.3. Climate change influences the course and strength of currents, the water temperature and salinity and the thermal stratification of the sea, which

¹On April 1 2009, Fisheries Research Services (FRS) was merged with the Scottish Fisheries Protection Agency (SFPA) and the Scottish Government Marine Directorate to form Marine Scotland - Scotlands new marine management organisation.

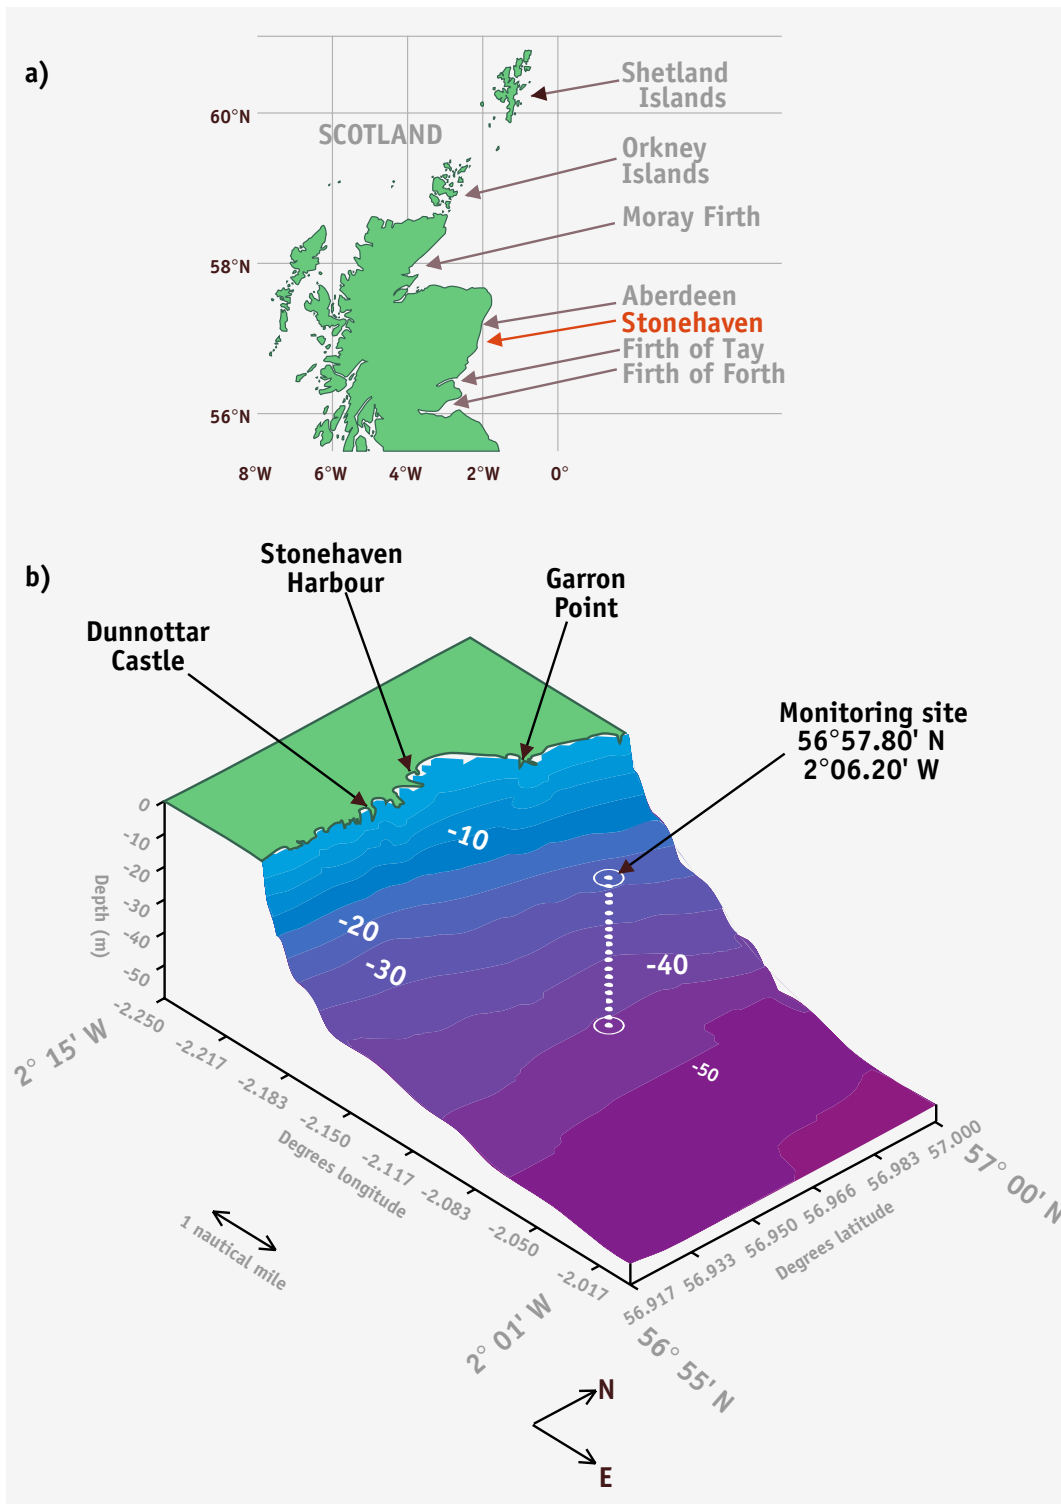


Figure 1.1: a)Map of Scotland showing Stonehaven. b)The Stonehaven long-term Ecosystem monitoring station. From FRS material.




Figure 1.2: The circulation patterns of North Sea water. The width of arrows is indicative of volume transport. Red arrows indicate relatively pure Atlantic water. From ICES (2006).

all affect plankton communities and chemical nutrient supply. These changed conditions influence the physiology and population dynamics of the phytoplankton species that absorbs the energy from the sun and chemical nutrients to provide the biomass that feeds the zooplankton species. Larger fish such as cod and haddock feed on zooplankton, and thus a change in the seasonality of phytoplankton can alter the timing and balance of productivity in the sea and affect commercially important fish and other top predators such as seabirds and cetaceans.

We use methods based on Classical Seasonal Decomposition (CSD) to test for stability in the seasonal pattern of the series. A seasonal pattern can exhibit changes in the size of the seasonals (amplitude change) or in the timing of occurrence (phase change). Furthermore, these changes can happen in a systematic or a random way. A random change (fluctuation) in the seasonal cycle can often be attributed to environmental factors (e.g. an unusually cold period) while systematic changes are mostly related to climatological conditions and their effects on the environment (e.g. days become gradually colder). Once the instability of the seasonal pattern is verified we explore the nature of it. We are interested in systematic changes in either the size or the timing of the seasonal pattern. Thus, we are actually looking for evidence that climate change influences marine life.

There exists a vast bibliography of studies on seasonal stability, many of which come from econometrics such as Said & Dickey (1984); Hylleberg *et al.* (1990). These studies have produced a number of highly sophisticated, complex and elaborate techniques e.g. (seasonal) unit root tests, that aim to address issues arising in the context of seasonal (in)stability in economic time series. Nonetheless, these methods are based on theoretical hypotheses and should be applied with great care. Furthermore, they were created for use with monthly or quarterly time series, and thus cannot be applied to higher frequency (weekly) biological data without problems. We are interested in testing the stability of a pattern and do not necessarily require the detailed output of some of these unit root tests. Furthermore, the Stonehaven data are not equi-spaced in time, with many missing values and duplicate measurements, and include erroneous values due to equipment malfunctions. Therefore, it is important to

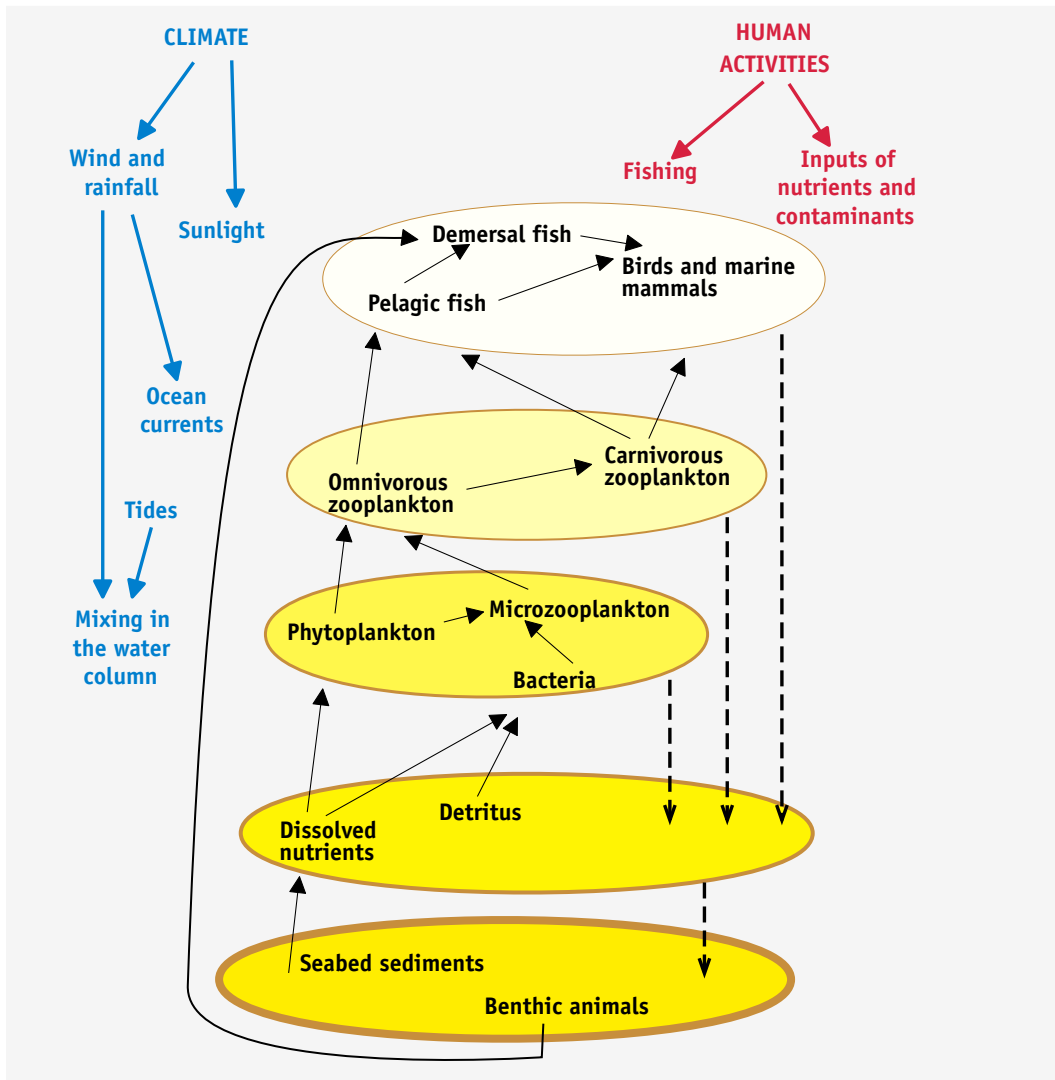


Figure 1.3: The food web. From FRS material.

find straightforward, robust and simple procedures for testing scientific hypotheses about changes in the seasonal cycle that can be used to analyse the Stonehaven data set but also can easily be applied by marine biologists in the future. We consider two main approaches: Resampling and Generalized Additive Modelling.

In the following chapters we firstly talk about CSD techniques and use three time series as illustrative examples. Then we review some testing methodologies for the different types of seasonality and present some new simple approaches. A sensitivity analysis of the newly presented methods is conducted. Finally the results from our analysis of the Stonehaven data sets are discussed along with possibilities for further work.

Chapter 2

Seasonal Decomposition

2.1 Introduction to Seasonal Decomposition

A time series is defined as a sequence of observations occurring in time. Typically, in time series, adjacent observations are dependent and time series analysis consists of methodologies to analyse this dependence. A time series is said to be seasonal with period s when similarities in the series occur every s basic time intervals (Box & Jenkins, 1976).

One of the oldest approaches to the analysis of a time series, X_t with period s , is to consider the series as the union of three components; a long term, relatively slow, movement of the data referred to as trend; a periodic component of period s , corresponding to seasonality; and the remainder or error. Hence, the series can be decomposed into those three components, the trend component, T_t , the seasonal component, S_t and the remainder, or the error term, R_t . Then the individual components can be further examined to allow more detailed interpretation of the data or even projected to the future to create forecasts.

$$X_t = f(S_t, T_t, R_t) \tag{2.1}$$

The trend in a series, which can be thought of as a measure of its local level, exhibits slow movement and can be linear or non linear. A linear trend can be easily projected in the future to give an insight as to what the behaviour of the series will probably be a few time steps ahead.

The seasonal component is the periodic movement of the series, and thus can be viewed in terms of amplitude (size) and phase (timing) similar to a sinusoid. By examining the time of occurrence and the nature of the seasonal pattern one can get an appreciation of the different effects influencing the behaviour of the series. The seasonal component in a series can behave in a number of different ways as explained in the next section. It is of great importance in biological applications as animal life cycles are often strongly seasonal and so it may yield insight into the underlying mechanisms that drive the process examined.

Furthermore, both the trend and the seasonal components can be modelled as either deterministic or stochastic. The above methodology is known as seasonal decomposition. The decomposition of a time series can take a variety of forms but the two most common and useful ones are the i) additive, with the underlying assumption that adding these components together will produce the original series, thus:

$$X_t = T_t + S_t + R_t \quad (2.2)$$

or ii) multiplicative, in which the original series is the product of the components, thus:

$$X_t = T_t * S_t * R_t. \quad (2.3)$$

The selection of the model depends on the nature of the time series. One easy approach is to fit a trend to the series and look at the derived discrepancies between the fitted trend and the seasonal troughs and peaks in the series. If the size of the discrepancies appears to be stable over changes in trend then the additive decomposition model may be applicable, while if they are increasing or decreasing with the trend, the multiplicative model is preferred. In the latter case, however, instead

of employing the multiplicative model one may choose to take the logarithm of the analyzed series and use the additive decomposition model:

$$\log(X_t) = \log(T_t * S_t * R_t) \quad (2.4)$$

$$\log(X_t) = \log(T_t) + \log(S_t) + \log(R_t). \quad (2.5)$$

A vast variety of techniques used to decompose series exists. The most common method used for trend extraction is the moving average (MA). The MA method is based on the reasoning that neighbouring observations are likely to be close in value (Makridakis *et al.*, 1998). The value of the moving average of order q , MA(q), at time t is the average of the observation at time t and the $m = \frac{q-1}{2}$ points on either side of it:

$$T_t = \frac{1}{q} \sum_{j=-m}^{j=m} X_{t+j}, \quad (2.6)$$

where X_{t+j} is the observation at time $t + j$, i.e. j observations after the one we are estimating the trend for. In the above definitions q is odd and there are no MA(q) values for the first or last $(q - 1)$ time points. The higher the order of the MA, the more terms and information from the data are lost. There are different ways to overcome the problem of missing values introduced by such a MA. A simple one is to use a lower order MA for the end points ('end filters') hence, reducing the number of missing values.

Choosing the order of a MA equal to the suspected periodicity, s , eliminates the seasonality as it is creating averages with length equal to the periodicity, negating thus, the seasonal variation in the series. Therefore, a MA(s) may be used to capture the slow long term movement in the data. If the periodicity is an even number then a centred MA($2 \times s$) is employed to ensure that the estimated values correspond to meaningful observation times. A centred MA($2 \times s$) is simply a MA(s) followed by a MA(2).

Once a trend has been estimated it is removed from the time series to produce a de-trended series that is then examined for seasonal effects. A stable seasonal cycle is estimated by averaging for each period across all years. These are called seasonal factors. Commonly the seasonal factors are standardized by subtracting from each of them their overall mean, thus ensuring that they sum to zero. The term ‘seasonals’ will always be used to refer to standardized estimates of stable seasonal factors. However, it is not always clear that the seasonality present in the series is stable over time. This is illustrated using three example time series below.

1. The monthly U.S. housing starts of privately owned single-family structures from January 1965 to January 1976, is referred to as the House series (Abraham & Ledolter, 1983). The House series is plotted in Figure 2.1.




Figure 2.1: The House series.

Using a centred moving average for trend extraction and averaging for the seasonal component, the House series is decomposed, figure 2.2.

The de-trended series in figure 2.2 indicates a stable seasonal pattern. However, the December values in the de-trended series (marked with red circles) tend to fall over time questioning the stability of the seasonals. A change in the December size of the seasonals could be connected to unusually heavy winters for some years or a change




Figure 2.2: The House series with the centred moving average estimated trend, the de-trended series with marked Decembers, the stable seasonals and the residuals.

in the behaviour of the market during the holidays period.

2. The monthly unemployment rate in the UK from January 1986 to December 1996, referred to as the Unemployment series (National Statistics Online, 1997). This series is shown in Figure 2.3.

In Figure 2.4 the de-trended series of the Unemployment data shows two peaks that appear to change in relative size indicating that the nature of seasonality may not be stable. Also there is still some trend in the de-trended series so a further trend extraction should be considered.

3. Finally a monthly temperature time series from the Kola peninsula from January 1967 to January 2001 is also examined, the Kola series (Marshall, pers. comm., 2004). This series, shown in Figure 2.5, is a part of the observed temperatures in the Kola section that have been monitored for more than 100 years (Ottersen *et al.*, 2005). The Russian State Hydro-meteorological University (RSHU) and the Polar Research Institute of Marine Fisheries and Oceanography (PINRO) are involved in the collection process and in collaboration with the Institute of Marine Research in Norway are analyzing the data (Bochkov, 2005; Averkiev *et al.*, 2005). The Kola




Figure 2.3: The Unemployment series.

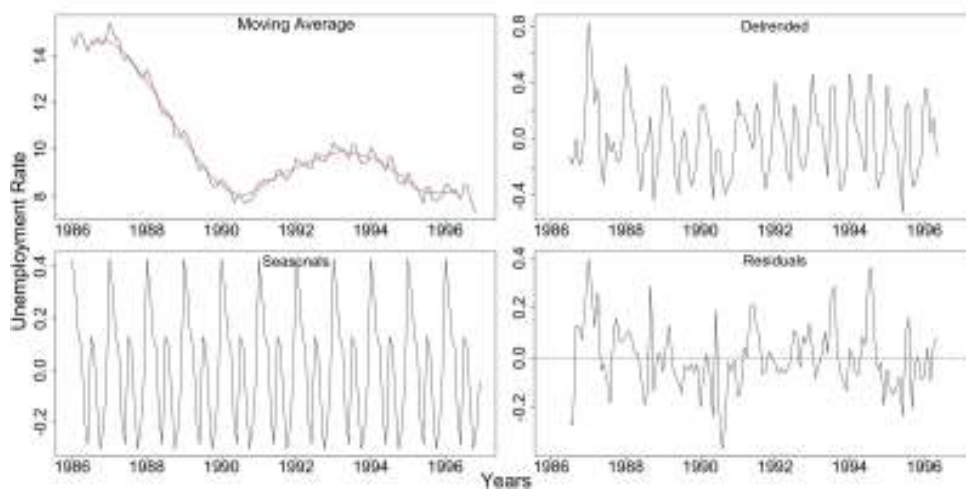


Figure 2.4: The Unemployment series with the centred moving average estimated trend, the de-trended series, the stable seasonals and the residuals.

section temperatures are considered to be representative of fluctuations in the climate in the Barents Sea, (Ottersen *et al.*, 2005). Yndestad *et al.* (2008) examine relationships between the Kola section and the Faroe-Shetland Channel time series. For us, the Kola series is of particular interest as part of our analysis involves looking at temperature series from Stonehaven.




Figure 2.5: The Kola series.

Figure 2.6 shows the Kola series decomposed by the aforementioned methodology. We notice that although the original series in Figure 2.5 displays frequent and substantial changes, the de-trended series in Figure 2.6 indicates that these are probably changes in trend as there are no signs of any changes in seasonality.

A strength or weakness of the available decomposition methods is their ability or inability to cope with missing values in the series and this is a problem with our data from Stonehaven, so it must be considered. A time series can have missing values at the beginning or at the end, or even internally. There are two direct solutions to this problem. One is to use techniques that allow for the existence of missing values and do not get affected by them, while the second is to replace the missing values in some way and then use “standard” procedures. Thus, one could perform a seasonal decomposition that allows for missing values and then extrapolate or interpolate the estimated components accordingly to substitute the




Figure 2.6: The Kola series with the centred moving average estimated trend, the de-trended series, the stable seasonals and the residuals.

missing entries, or use extrapolation (or interpolation) on the time series before conducting the decomposition.

The last section of this chapter is a brief review of the existing literature on seasonal decomposition. Further discussion on seasonality follows as it is important to explain the terminology used for this central component of seasonal decomposition.

2.2 Trending Seasonality

The idea of a trending seasonal pattern may seem peculiar at first but it appears to be so for many time series. For example, in Nature even though our intuition may be that seasonality should be stable, e.g. colder winters and warmer summers, for temperature data, the phenomenon of global warming with all its consequences raises the question of whether that seasonal pattern stays stable or whether there is a shift/change of the seasonality as time passes, i.e. coldest or warmest temperature occurs earlier or later in the year. This is just one example of how a trending seasonal pattern may arise in an environmental context, but there are many more occurrences

in other areas. In this section we discuss briefly what kinds of trend are possible in an existing seasonal pattern and explain the reasons for considering one specific type of change more important in the context of this thesis.

We may define a trending seasonal pattern as one which exhibits change from year to year. This change can be in the size of the pattern (amplitude change), in the time of the occurrence of the pattern (phase change) or a combination of the above. Furthermore, regardless of the nature of the change (amplitude or phase) the change can take place in a systematic way - meaning an average drift in one direction, either increase or decrease - or as random fluctuations. Thus, a systematic shift in time (later or earlier each year) is a systematic change of phase while a systematic increase or decrease in the size of the seasonals is a systematic amplitude change. A random movement in the timing of peaks and troughs is referred to as a fluctuation in time and a random change of the size of the seasonals as fluctuating amplitude change. Considering the combinations of the possible nature and types of change, there are nine different possibilities for a seasonal pattern that are demonstrated in Table 2.1.

Table 2.1: Types of Seasonal Patterns

		Phase		
		No Change	Systematic Change	Fluctuation
Amplitude	No Change	AN & PN	AN & PS	AN & PF
	Systematic Change	AS & PN	AS & PS	AS & PF
	Fluctuation	AF & PN	AF & PS	AF & PF

The different types of change that a seasonal pattern may exhibit are shown in Figure 2.7. Clockwise from top left we see a systematic increase and a fluctuation in amplitude and a systematic earlier shift and a fluctuation in phase.

In Nature and in particular, in marine biology, commonly series exhibit seasonals changing in one way or another. However, even though all types of changing seasonality may be scientifically of interest, we can often differentiate between the causes




Figure 2.7: A systematic (increasing) amplitude change, a fluctuation in amplitude, a systematic (earlier) shift in time and a fluctuation in time, using a period of 52 weeks. In the bottom two plots the dotted lines mark the 30th week.

of these types of change. While fluctuations in the size or timing of a seasonal pattern are associated with climatological factors that can themselves fluctuate rapidly, a systematic change is mostly associated with large scale environmental factors, i.e. the underlying long term change in the environment, and thus are realised as systematic -usually small- changes over time. In our analysis of the data from Stonehaven that follows, we are particularly interested in those slow, long term dynamics, and thus our focus is mainly to find ways to identify and estimate systematic change in either size or time. Clearly, a systematic time shift can be related to climate change and global warming. Initially, however, all types of seasonality will be considered for the series analysed.

Not all time series have a seasonal pattern and amongst those that exhibit seasonal behaviour, some seasonal patterns stay stable while others change over time. A seasonal pattern will be referred to as a *stable seasonal pattern* or *stable seasonality* if it remains stable over time, as a *pseudo-trending seasonal pattern* or *pseudo-trending seasonality* if it exhibits any kind of fluctuations, and as a *trending seasonal pattern* or *trending seasonality* if it is changing smoothly with time, i.e. systematic shift in time

and systematic amplitude change. Thus, when identifying the kind of seasonality in a series only systematic changes will be considered as a trending seasonality while attempts to identify random changes are noted, even though considered ‘pseudo-trending’ as mentioned above.

2.3 Brief Description of Seasonal Decomposition Bibliography

In this section a brief review of existing methodologies for (seasonal) decomposition of a time series is presented. As noted, these methodologies try to extract two separate components, a trend-cycle and seasonality, from a time series and thus distinguish them from the randomness. A quick look at the history of decomposition follows based on Makridakis *et al.* (1998).

The first attempts to decompose a series originated around the start of the 20th century. It was believed that to study the serial correlation within or between variables one had to remove any spurious correlation caused by trend. Poynting (1884) and Hooker (1901) attempted to eliminate the trend and seasonal fluctuation by averaging over several years. Later Spencer (1904) and Anderson and Nochmals (1914) generalized the procedure of trend extraction to include polynomials of higher order.

At the same time economists were working towards the same goal, trying to predict economic depressions. Separating the elements of economic activity would isolate the changes in the business cycle from seasonal and other effects. In 1911 a committee was appointed in France to analyse the causes of the economic crisis in 1907. The group attempted to separate the trend from the cycle of the series. In the United States, Copeland (1915) tried to extract the seasonal fluctuation from the other components of the series.

Time series decomposition as known today was first employed by F. R. Macauley of

the National Bureau of Economic Research in the 1920s. He introduced the ratio-to-moving average approach to time series decomposition that is the basis of Census II. Shiskin (1957) created Census II, a computer programme for decomposition of time series which is now the most widely used of the decomposition packages.

A brief discussion of the most commonly used decomposition ideas and techniques follows. There is a variety of different ways to estimate or extract the trend and seasonal components and many possible combinations. We start with some simple ways to estimate these components individually and go on to look at more elaborate algorithms.

2.3.1 Weighted Moving Average - WMA

Averaging neighbouring values can be used to estimate the trend in a series. An extension on the MA(q) is to attribute (unequal) weights to the observations according to their distance from the centre of the MA. Meaning that the closer the observation is to the time point of estimation the more influence it will have on that estimate. This method is known as a weighted moving average (WMA). The size of the weights can be decided based on the nature of the series. Procedures that use weights that decrease exponentially are called exponential smoothing procedures (Makridakis *et al.*, 1998). A general form of a WMA at time t is:

$$T_t = \sum_{j=-m}^{j=m} \alpha_j X_{t+j}, \quad (2.7)$$

where α_j are the weights and $m = \frac{q-1}{2}$. When all α_j are set equal to $1/q$ we have the simple MA(q), (2.6). Two conditions commonly apply to the weights: they should

sum to one and they should be symmetric:

$$\sum_{j=-m}^m \alpha_j = 1$$

$$\alpha_j = \alpha_{-j}$$

There are, however, situations that asymmetric weights are more appropriate. When the weights do not sum to one (2.7) is divided by their sum.

The benefit from using a WMA is that since the observations are slowly down-weighted and do not leave the average abruptly the result is much smoother than a simple MA. Nevertheless, there are two main disadvantages. At the end points of the series, relative to the order used, missing values occur as they cannot be estimated by the WMA. Trying to overcome this issue, one could employ smaller order WMA at the ends, known as ‘end filters’. This, however, further results in flat estimates near the end of the series even when a pronounced trend is obvious (Makridakis *et al.*, 1998).

Furthermore, using WMA leaves the trend sensitive to outliers and extreme values. A modified smoothing technique that behaves better when outliers are present is to use running (moving) medians. Again an order needs to be chosen that corresponds to the number of data points included but now the median and not the mean of these points is used. This procedure is robust to outliers but produces ‘rough’ fits. One could use a smoother on the trend estimated by running medians to produce a smooth trend (Faraway, 2006). Similarly robust to outliers, running medians are defined, where medians are the means once the highest and lowest observation are removed (Makridakis *et al.*, 1998).

2.3.2 Local Regression Smoothing

Using local regression smoothing one can avoid the above issues created by using a (W)MA. Instead of averaging a window of the series one now fits a straight line

through them,

$$T_t = a_t + b_t t. \quad (2.8)$$

The two parameters a_t and b_t represent the intercept and slope of the line at time t , respectively. These can be estimated by Ordinary Least Squares (OLS), minimizing the sum of squared errors. Hence, fitting a series of straight lines to the data we get estimates for each data point. Using local regression to estimate the trend allows us to estimate a local trend instead of a global linear trend (Chatfield, 2000). Similar to WMA we can attribute weights, α_j to the observations and thus fit by minimizing the weighted sum of squares. This method improves the bias observed when using (W)MA at the end of the series (Makridakis *et al.*, 1998).

In order to use local regression smoothing one has first to decide the value of the smoothing parameter, q , which in analogy to the order of a MA, represents the number of data points considered in each fit. A large value of q would produce a very smooth picture but could lose some of the existing pattern while a small value could incorporate noise in our estimates.

Cleveland's Locally Weighted Scatterplot Smoother (LOWESS or LOESS) is an implementation of weighted local regression smoothing with protection against outliers that can influence the fitted lines or curves. It fits a locally-fitted polynomial of degree d to the data with common values of d equal to 1, locally-linear, or 2, locally-quadratic (Cleveland *et al.*, 1990). LOWESS starts with an initial weighted local regression. The residuals from this regression are computed. Then the weights are adjusted so that the observations that resulted in large residuals have smaller contribution to the regression. The local regression is then repeated with the new weights. The new residuals are calculated and the weights adjusted again, accordingly. The procedure repeats down-weighting the points that give large residual values, as suspected outliers. When the estimated trend converges the procedure stops (Cleveland, 1981).

Constantly down-weighting data points with large produced residuals makes LOWESS

more robust than local weighted regression. In fact, if there are no outliers, the resulting trends from the two procedures would be almost the same while when there are extreme values LOWESS produces better results (Makridakis *et al.*, 1998).

2.3.3 Seasonality Estimation/Extraction

If a stable seasonal pattern is assumed, one can average each season's observations to produce seasonal factors. This, nonetheless, can easily be affected by extreme values. Taking seasonal medians or medians can help deal with outliers. Nevertheless, the above methods assume a seasonal pattern that does not change in time.

MA can be used to extract the seasonal pattern of a series that is believed to change in time. Seasonal averaging to estimate seasonal factors would not capture a changing pattern. Using a (W)MA for each season would, however, allow for changes in time. Alternatively, one could use local regression to estimate a seasonal pattern that changes in time. For each season's points local regression can be employed to provide smooth estimates of this periodic component (Makridakis *et al.*, 1998).

To extract and not estimate the seasonal pattern one could also use a linear filter. Though, many possibilities exist, seasonal differencing is a simple commonly used way. The seasonal difference is written as $1 - B^s$ where s is the periodicity of the series and B is the back-shift operator, $Bx_t = x_{t-1}$ so that $(1 - B^s)x_t = (x_t - x_{t-s})$ (Chatfield, 2000).

2.3.4 Census Bureau methods

The U.S. Bureau of Census has developed, over the years, a number of intricate methods for seasonal adjustment of, mainly economic, time series. The most recent variant is X-12-ARIMA (Findley *et al.*, 1998). This is an extension of the X-11-ARIMA, Census II method, in an effort to suitably seasonally adjust more types of

economic data (Makridakis *et al.*, 1998). The following discussion on X-12-ARIMA is based on Findley *et al.* (1998); Makridakis *et al.* (1998).

The Census II decomposition uses a number of different WMAs. As mentioned before, this leads to a loss of data points at the beginning and end of the series. The X-12-ARIMA uses end-filters to avoid losing many data points at the end points of the series. Additionally, it gives the user the ability to create forecasts and backcasts based on an ARIMA model and thus, extend the series to accommodate for the WMAs.

The X-12-ARIMA method has twelve steps and extends its predecessor, X-11, to include alternative seasonal, trade-day and holiday effect adjustments with a variety of seasonal and trend filter options. Furthermore, the diagnostics for assessing the quality of the seasonal adjustments are more methodical and it allows extensive modelling for linear regression models with ARIMA errors (regARIMA). Finally, it uses an interface which can facilitate batch processing of large numbers of series.

This latest Census II method allows for additive (A) and multiplicative (M) decomposition and also a pseudo-additive (PA) one. The last one is used to seasonally adjust series with small (even zero) values in the same season every year. A PA decomposition has the form:

$$X_t = T_t * (S_t + R_t - 1) \quad (2.9)$$

where, X_t is the examined time series, S_t the seasonals, T_t the trend and R_t the remainder component.

Assuming a seasonal time series X_t with periodicity s , the algorithm begins with a centred MA($2 \times s$) to estimate the trend, T_t^1 . For monthly data a MA(2×12) is used:

$$T_t^1 = \frac{1}{24}X_{t-6} + \frac{1}{12}X_{t-5} + \dots + \frac{1}{12}X_t + \dots + \frac{1}{12}X_{t+5} + \frac{1}{24}X_{t+6} \quad (2.10)$$

The estimated trend is then removed from the series - forecasts and backcasts are used at the end points- to create the de-trended series, $Detr_t^1$. The de-trended series is formed by subtraction in the additive case or division in the multiplicative one, of the trend from the original series:

$$A : Detr_t^1 = X_t - T_t^1 \quad (2.11)$$

$$M, PA : Detr_t = X_t/T_t^1 \quad (2.12)$$

From the de-trended series the initial estimate of the seasonal component, \hat{S}_t^1 , is computed using a $MA(3 \times 3)^1$ seasonal MA for monthly periodicity:

$$\hat{S}_t^1 = \frac{1}{9}Detr_{t-24}^1 + \frac{2}{9}Detr_{t-12}^1 + \frac{3}{9}Detr_t^1 + \frac{2}{9}Detr_{t+12}^1 + \frac{1}{9}Detr_{t+24}^1 \quad (2.13)$$

To remove any noise included in this initial estimate of the seasonals, \hat{S}_t^1 , a centred $MA(2 \times s)$ (2×12 for monthly) is applied to it and it is then removed from the initial estimate to produce seasonal factors, S_t^1 :

$$A : S_t^1 = \hat{S}_t^1 - \left(\frac{1}{24}\hat{S}_{t-6}^1 + \frac{1}{12}\hat{S}_{t-5}^1 + \dots + \frac{1}{12}\hat{S}_t^1 + \dots + \frac{1}{12}\hat{S}_{t+5}^1 + \frac{1}{24}\hat{S}_{t+6}^1 \right),$$

or

$$M, PA : S_t^1 = \frac{\hat{S}_t^1}{\frac{1}{24}\hat{S}_{t-6}^1 + \frac{1}{12}\hat{S}_{t-5}^1 + \dots + \frac{1}{12}\hat{S}_t^1 + \dots + \frac{1}{12}\hat{S}_{t+5}^1 + \frac{1}{24}\hat{S}_{t+6}^1}$$

Thus, the initial remainder, R_t^1 can now be computed, removing the estimated seasonal factors from the de-trended.

$$A : R_t^1 = Detr_t^1 - S_t^1 \quad (2.14)$$

$$M : R_t^1 = Detr_t^1/S_t^1 \quad (2.15)$$

$$PA : R_t^1 = Detr_t^1 - S_t^1 + 1 \quad (2.16)$$

¹double MA(3)

The points at which this remainder takes large values are identified as extreme observations and the seasonal component is re-estimated without them. Hence, any points that fail to follow the pattern of the rest of the data are replaced by estimates. These seasonals are removed from the original series to give a de-seasonalized series, Des_t^1 :

$$A : Des_t^1 = X_t - S_t^1 \quad (2.17)$$

$$M : Des_t^1 = X_t/S_t^1 \quad (2.18)$$

$$PA : Des_t^1 = X_t - T_t^1 * (S_t^1 - 1) \quad (2.19)$$

The trend is now estimated from the de-seasonalized series using a Henderson's WMA that removes the noise from the trend. The Henderson's filters have a choice of lengths of 9, 13 or 23 with the possibility of it being automatically chosen, see Findley *et al.* (1998). A second estimate of the seasonal factors is computed similarly to before but with a MA(3×5) used for monthly series. Then a second de-seasonalized series is computed, as in (2.17), (2.18) and (2.19), and a final trend is extracted from it using Henderson's filters. These steps are iterated to ensure a more successful estimation of the components by replacing extreme values and eliminating any noise that enters the estimation.

Systematic diagnostic checks are used to evaluate the decomposition. Spectrum estimates can be examined to determine whether seasonal effects are present in the residual series. A detectable residual seasonal pattern can be attributed to inappropriate seasonal adjustments of the model or to a trend in the seasonals. X-12-ARIMA includes two types of stability diagnostics; 1) sliding spans which uses overlapping sub-spans to analyze the difference between the largest and smallest adjustments made to a point and 2) revision histories which considers the difference between the earliest adjustment of a point and the latest one based on all future available points.

Furthermore, X-12-ARIMA allows for a variety of regressors to be included in the

regARIMA models. These can be used to model sudden changes in the level or the seasonality of the series. Additionally, seasonal effects such as leap years, holidays, trading-days and others can be modelled this way. There are automatic model identification and selection procedures employed in the X-12-ARIMA which help utilize the regARIMA models. Constructing regARIMA models with specific terms the user can test for the existence of effects such as changes in the seasonal pattern. Nevertheless, X-12-ARIMA allows only for quarterly or monthly seasonal periods.

2.3.5 A Seasonal-Trend Decomposition Based on LOESS - STL

An alternative to the aforementioned Census II methods is a Seasonal-Trend Decomposition Based on LOESS (STL) (Cleveland *et al.*, 1990). It consists of a sequence of LOESS applications and thus is a resistant to extreme values decomposition technique. It can be applied to time series with missing values and contrary to the X-12-ARIMA the STL allows any seasonal period greater than 1. It cannot accommodate, however, the many seasonal effects (such as trading-day etc.) that X-12-ARIMA can and it has no built-in diagnostics for the seasonal adjustment. There is only an additive version of STL but the logarithm of the data can be used for the multiplicative case (Makridakis *et al.*, 1998). The following description is based on Cleveland *et al.* (1990).

This method consists of two recursive procedures, one nested within the other. During each iteration of the inner loop the trend and seasonal components are estimated. One iteration of the outer loop consists of multiple runs of the inner and then computation of robustness weights to be used in the next run of the inner loop, to reduce the contribution of the extreme values to the overall estimation procedure. For the initial run of the outer loop these robustness weights are all set to 1.

The inner loop starts by de-trending the series, (2.11), and thus, an initial estimate of the trend, T_t^1 is needed. This is usually set to be zero. Wherever there are missing

values in the original series there will be ones also in the de-trended one as STL does not substitute them by estimates.

The de-trended values for each season are collected to form seasonal ‘sub-series’. For example, for monthly data twelve seasonal sub-series are created. Each sub-series is then smoothed by LOESS with parameter q_s and $d = 1$. The smooth is estimated at all internal points of the series, including any missing values, but it is also extrapolated one point before and one ahead of the end points of the sub-series, i.e. one year backcast and forecast of the seasonals. The smoothed sub-series are put together to create the seasonal component, \hat{S}_t^1 .

A low-pass filter, $MA(3 \times s \times s)$, is used on this seasonal component, consisting of three MAs followed by a LOESS with parameter q_l and $d = 1$ to create L_t^1 . Data points are lost at both ends, s at each end, from the usage of MAs but as the sub-series were extended proactively earlier this does not cost any of the original time points. The low frequency component L_t^1 is subtracted from the seasonals \hat{S}_t^1 (i.e. de-trending of the seasonal) to ensure that no low frequency movement enters the seasonal estimation:

$$S_t^1 = \hat{S}_t^1 - L_t^1 \quad (2.20)$$

where S_t^1 is the final seasonal component of this iteration. A de-seasonalized series is created as in (2.17) and LOESS is used on it to extract the trend, T_t^1 . This LOESS has parameter q_t and $d = 1$. The trend is estimated at all points of the original series, including any missing values. This trend will be used as the initial estimate for the next iteration of the inner loop.

The outer loop begins after one or two iterations of the inner one. The remainder component is computed by subtracting the estimated trend and seasonals from the original series. A large value in the remainder is considered evidence of an extreme value in the original series. Once these are identified, robustness weights are computed giving small or zero weights to them. Let $h = 6 \times \text{median}(|R_t|)$, the robustness

weight at time t is:

$$w_t = B(|R_t|/h), \quad \text{where} \quad (2.21)$$

$$B(t) = \begin{cases} (1 - t^2)^2, & 0 \leq t \leq 1 \\ 0, & t \geq 1 \end{cases} \quad (2.22)$$

Therefore, when the inner loop is repeated after one outer loop iteration the new robustness weights are used for two LOESS applications, for the trend and seasonal estimations. The two LOESS used to estimate the seasonal component and the trend component are used on the appropriate series after it has been multiplied by the newly calculated robustness weights.

STL allows us to choose a number of parameters according to the data examined. Firstly, the number of inner and outer loop iterations has to be decided. Commonly one pass through the outer loop consists of one or two iterations of the inner loop while a total of 10 or 20 outer loop iterations take place (Makridakis *et al.*, 1998). If we use one inner loop iteration per outer loop iteration then a total of 5 passes through the outer are sufficient while 10 provide a near certainty of convergence².

Furthermore, the q_s , q_l and q_t parameters for the LOESS used have to be chosen. Cleveland *et al.* (1990) suggest the value of the least odd integer greater than or equal to s for the q_l parameter and an odd value of at least 7 for the q_s parameter, based on an eigenvalue analysis of the inner loop filters. Parameter q_l controls the level of smoothness of the trend that entered the seasonal estimation and is then removed from it. The value of the q_s parameter determines the amount of change allowed in the estimated seasonal component. A small value of q_s allows substantial changes from year to year while a large ensures that the change is slow.

Finally, the q_t parameter sets the amount of variation from the data to be included

²criterion used for convergence by Cleveland *et al.* (1990) is $\max_t |U_t^k - U_t^{k+1}| / (\max_t U_t^k - \min_t U_t^k) < 0.01$ where U^k is the trend or seasonal component estimated in iteration k .

in the trend estimation. A high value includes less variation, smoother estimate of trend, while a small one more variation, wigglier estimate. Thus, a trade-off is created as we want a high value for q_t so that only the slow, long-term variation is included in the estimation but we also want a low enough value so that no low-frequency effects remain in the remainder to be treated as extreme values. An odd value is proposed for this parameter, too by Cleveland *et al.* (1990) and they also suggest that the following inequality should hold:

$$q_t \geq \frac{1.5s}{1 - 1.5q_s^{-1}}. \quad (2.23)$$

If the series examined has two or more seasonal components, STL can be used successively to estimate firstly the seasonal component with the shortest period and then, after that is subtracted from the data, the other ones till the longest-period component is estimated.

Therefore, STL offers choices in specifying the amount of variation in the trend and seasonal components, making it very flexible. Nevertheless, the estimated trend is unstable at the end of the series since LOESS is used (Makridakis *et al.*, 1998).

2.3.6 Structural Time Series Models

Structural time series models (STMs) are linear Gaussian state-space models for time series based on a decomposition of the series into various components (Harvey, 1990). These components are referred to as ‘unobserved’ components because they are not observed directly but are assumed to have ARMA representations (Harvey, 1993). STMs commonly include three components; a trend, a seasonal and a remainder (or error) component. The following description is based on Harvey (1990, 1993).

Let $x_t, (t = 1, 2, \dots, T)$, be the observed time series. A simple STM that includes a

local trend (local level model) is:

$$x_t = m_t + v_t, \quad v_t \sim N(0, \sigma_v^2) \quad (2.24)$$

$$m_{t+1} = m_t + u_t, \quad u_t \sim N(0, \sigma_u^2) \quad (2.25)$$

where m_t is the level/trend at time point t . Though the trend can take many forms, in this case it is modelled as a random walk by (2.25). When $\sigma_u^2 = 0$, x_t has a constant level while when $\sigma_v^2 = 0$ x_t is a random walk. A signal-noise ratio ($q = \sigma_u^2/\sigma_v^2$) can be used to describe this model.

A basic structural model that includes a trend and a seasonal component is:

$$x_t = m_t + s_t + v_t, \quad v_t \sim N(0, \sigma_v^2) \quad (2.26)$$

$$m_{t+1} = m_t + b_t + u_t, \quad u_t \sim N(0, \sigma_u^2) \quad (2.27)$$

$$b_{t+1} = b_t + z_t, \quad z_t \sim N(0, \sigma_z^2) \quad (2.28)$$

$$s_{t+1} + s_t + \cdots + s_{t-s+2} = w_t, \quad w_t \sim N(0, \sigma_w^2) \quad (2.29)$$

where m_t is the trend and s_t is the seasonal dummy at time t when x_t has period s . When $\sigma_w^2 = 0$, x_t has a stable seasonal pattern.

From the above models we can see that STMs can be regarded as regression models in which the explanatory variables are functions of time and the parameters are time varying.

Changes in the seasonal pattern can be incorporated in these models with a stochastic trigonometric seasonal component:

$$y_t = \mu_t + \gamma_t + \epsilon_t \quad t = 1, \dots, T \quad (2.30)$$

where $\gamma_t = \sum_{j=1}^{\lfloor s/2 \rfloor} \gamma_{j,t}$ and each $\gamma_{j,t}$ is generated by

$$\begin{bmatrix} \gamma_{j,t} \\ \gamma_{j,t}^* \end{bmatrix} = \begin{bmatrix} \cos\lambda_j & \sin\lambda_j \\ -\sin\lambda_j & \cos\lambda_j \end{bmatrix} \begin{bmatrix} \gamma_{j,t-1} \\ \gamma_{j,t-1}^* \end{bmatrix} + \begin{bmatrix} \omega_{j,t} \\ \omega_{j,t}^* \end{bmatrix} \quad \begin{array}{l} j=1, \dots, \lfloor s/2 \rfloor \\ t=1, \dots, T. \end{array} \quad (2.31)$$

In the above equation $\lambda_j = 2\pi j/s$ is frequency, in radians, and $\omega_{j,t}$ and $\omega_{j,t}^*$ are two mutually uncorrelated white noise disturbances with zero means and common variance which is the same for all j .

Once a model has been put in state space form, the Kalman filter may be applied for prediction and smoothing. In addition, for a Gaussian model the Kalman filter can be used to construct the likelihood function by a prediction error decomposition.

2.3.7 TRAMO - SEATS

Contrary to the U.S. Bureau of Census, in Europe, Eurostat uses Time series Regression with ARIMA noise, Missing observations and Outliers (TRAMO) and Signal Extraction in ARIMA Time Series (SEATS) to seasonally adjust time series with monthly or lower frequency. A short description of these methods follows and is based on Maravall (2002).

TRAMO can be used for estimation, forecasting and interpolation of regARIMA models. Additionally, it can be run in an entirely automatic manner. It performs a regARIMA in the presence of missing values and outliers. Let x_t be our time series with n observations. TRAMO fits the regression:

$$x_t = \mathbf{y}_t \boldsymbol{\beta}_t + z_t \quad (2.32)$$

where $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_m)^T$ is a vector of regression coefficients, $\mathbf{y}_t = (y_{1t}, y_{2t}, \dots, y_{mt})^T$

is a vector of m regression variables and z_t follows an ARIMA process:

$$\phi(B)\varphi(B)z_t = \theta(B)a_t, \quad (2.33)$$

where B is the backshift operator (i.e. $Bx_t = X_{t-1}$), $\phi(B)$, $\varphi(B)$, $\theta(B)$ are finite polynomials in B and a_t is white noise. All the unit roots from differencing are contained in $\varphi(B)$ so that $\phi(B)$ is the stationary autoregressive polynomial and $\theta(B)$ is the moving average polynomial.

We can specify which regression variables to be included in the model or let the program generate them. These regression variables include trading day effects, holiday effects and intervention variables as dummy variables or sequences of ones and zeros.

The parameters in (2.32) and (2.33) are initially estimated by Maximum Likelihood (ML) or Ordinary Least Squares (OLS). Then the Kalman filter (Harvey, 1990) and the QR algorithm (Francis, 1961, 1962) are used to get new parameter estimates. Constructed t-tests for outliers are computed for each datum to check for the existence of outliers.

The detected outliers are not included in the estimation. Similarly to stepwise regression, multiple regressions are used to detect outliers and select the “best” regression equation. Three types of outliers are detected; additive outlier, level shift and transitory change. Nevertheless, the user can also specify outliers. Furthermore, seasonal effects, such as trading day effects, are considered and tested for, when applicable.

Missing observations are dealt in two ways; interpolation or assigning an arbitrary value and an additive outlier. In the latter case the interpolator is computed as the difference between the assigned value and the estimated regression parameter. Mean Squared Errors (MSE) are reported for all interpolators.

Finally, TRAMO can automatically identify and fit an ARIMA model to the series after the outliers, missing values and trading effects have been accommodated properly. Firstly the non-stationary polynomial, $\varphi(B)$, of (2.33) is identified by iterating a sequence of AR and ARMA(1,1) models³. Then an ARMA model is chosen for the stationary series, within the range of ARMA models with $0 \leq (p, q) \leq 3$, $0 \leq (P, Q) \leq 2$.

Hence, TRAMO adjusts a time series, estimating the deterministic effects of outliers, trading day or holiday effects and the ARIMA model that was identified. This information is then passed to SEATS for further analysis. SEATS decomposes (additive or multiplicative) a series that follows this ARIMA into four components (trend, seasonal, transitory and remainder) using signal extraction techniques. The trend component has a spectral peak at 0 frequency, the seasonal has spectral peaks at seasonal frequencies while the remainder (irregular) component captures white noise behaviour and thus has a flat spectrum. Contrary to the remainder, the zero-mean stationary transitory component picks up fluctuations that are not white noise but do not belong in the trend or the seasonal components. These components depend on the structure of the identified ARIMA and if no outliers and no trading day, holiday etc. effects are present in a series, then SEATS can be used directly to identify and decompose the ARIMA model that fits the series.

The ARIMA models for the components are established by partitioning the spectrum into additive spectra, one for each component. Under the assumption that only the remainder component includes noise ('canonical' property) a unique decomposition of the spectrum is identified. Minimizing the MSE, estimators of the components are computed filtering the series after extending it first with forecasts and backcasts. Revisions are used, as in X-12-ARIMA, till convergence is achieved. Along with these estimated components, SEATS provides their forecasts for several years with standard errors for all. A comparison of the differences between the theoretical and the empirical moments for these estimators can be used for diagnostic checks while spectral diagnostics are also available.

³Differencing and seasonal differencing are considered up to order $(1 - B)^2(1 - B^s)$.

SEATS, thus, decomposes the series, as modelled by TRAMO, into four components. If the ARIMA model identified and estimated by TRAMO cannot be decomposed appropriately then SEATS identifies a different model to work with. By default additive outliers and transitory changes are included in the remainder component by SEATS while level shifts to the trend. All trading day, holiday and generally all calendar effects are added to the seasonal component. The regression variables in the regARIMA can be added to any desired component but by default are treated as a separate one.

Overall, TRAMO-SEATS provides the user with an automated tool for seasonal adjustment, however, there are a few points that merit some consideration. SEATS prefers ARIMA models where the total AR order, including differencing, is equal to the total MA order (balanced models) and if that is not the case for the chosen one by TRAMO, SEATS may use many revisions. Thus, SEATS is more suited for balanced models. Furthermore, SEATS should only be used to seasonally adjust time series with a seasonal pattern as it can induce seasonality into the adjustment of non seasonal series (Hood, 2002). Spectral diagnostics can be used to check for seasonality in the original series as well as in the remainder component. Findley *et al.* (2005) propose modifications of SEATS diagnostics used to detect over- or underestimation of the components. The diagnostics of SEATS are computed using variance estimates under the assumption of an applied infinitely long filter which results in bias towards underestimation. Findley *et al.* (2005) advocate that usage of time-varying variances associated with the actual finite length filters that are employed leads to unbiased versions of the current diagnostics.

We should note that a new version of the X-12-ARIMA is now based on TRAMO while another is developed that includes SEATS. TRAMO-SEATS and X-12-ARIMA have been implemented in a single interface called ‘DEMETERA’⁴.

⁴This is freely available at <http://circa.europa.eu/irc/dsis/eurosam/info/data/demetra.htm>

2.4 Discussion

In this chapter we have talked about the different types of seasonality that may be present in a series. Furthermore, we have presented a number of commonly used techniques to decompose a series into a trend and seasonals. Since, however, there are many types of seasonality we should identify the correct type before we attempt to extract or estimate it. Additionally, for many of these methods to work properly the existence of a seasonal pattern has firstly to be verified.

The discussed methodologies are complex and require a well behaved quarterly or monthly times series. As discussed before, the Stonehaven data are higher frequency data and include missing values and duplicate measurements. Therefore, we would have to use monthly aggregates, risking that small changes in the pattern could go unnoticed. Additionally, in the context of this thesis we are mainly interested in tests between the different types (including absence) of seasonality. Hence, the detailed output of the above procedures is not required. These tests should be simple, robust and fairly automated so that they can be implemented easily on the Stonehaven data and help us draw accurate conclusions about the seasonality in them but also in order to be of use to marine biologists in the future. The next chapter is a brief introduction to testing techniques for seasonality in a series.

Chapter 3

Testing for Seasonality

3.1 Introduction to testing

3.1.1 Introduction

As mentioned before, in classical seasonal decomposition of time series, the series is decomposed into three basic components, the trend, the seasonal and the residuals. In common practice the seasonals are assumed to be stable over time, i.e. that the shape, size and timing of the seasonal cycle does not change with time. The previous chapter presented the possible types of seasonality in a time series and discussed techniques available to model these components in a seasonal decomposition. Nevertheless, one needs to decide the nature of the components to be included in the seasonal decomposition model.

In addition, in many sciences the purpose of a decomposition is to develop an understanding of the underlying dynamics that drive the process observed and therefore, an assumption of stability is beyond scope. For example in biological science the study of biological series is not intending to register the emerging patterns and simply extrapolate them in to the future but rather investigate what drives the series;

how and why the patterns are derived. In that context one may wish to conduct tests to check the nature of the seasonals and then estimate them accordingly. Estimating a stable seasonal component when in fact there is a trend in its shape, size or timing will distort the picture one has of the data and thus, it may lead one to draw wrong conclusions or simply miss the existence of some features.

Testing the existence of a seasonal pattern and furthermore trying to evaluate the kind of seasonality is of great interest to theoretical and applied scientists. In theory testing for existence and type of seasonals is used to verify theories and expand knowledge of phenomena. In practice it is of utmost importance to test whether what a scientist thinks he is observing corresponds to reality.

This chapter discusses different ways to test for seasonality. There are tests for the existence of seasonality and tests for the type of seasonality, once the existence has been verified. Firstly, old preliminary analyses for a stable seasonal pattern are presented and then formal testing frameworks for the detection of seasonality, from econometrics, are reviewed. The House, Unemployment and Kola time series are used as examples.

3.1.2 Preliminary Analyses – Testing for Stable Seasonality

Suppose that during the analysis of a discrete time series $\{X_t\}$ there are strong indications of an existing stable seasonal pattern and the objective is to assess whether this is true. A number of ways to explore and verify the existence of stable seasonality in a time series exist.

A popular first approach is to fit a straight line to each seasonal factor of the detrended series across the years. The trend is removed by a moving average of order equal to the periodicity of the series. The p-values are then weighed against a pre-defined significance level to evaluate the fit, thus testing the null hypothesis that

the seasonality is stable, the regression coefficient is equal to zero, ($H_0 : b = 0$) against the alternative of non stable seasonality, the regression coefficient is not zero, ($H_A : b \neq (0)$). Therefore, a p-value smaller than the predefined significance level rejects the null hypothesis in favor of the alternative while a greater accepts it.




Figure 3.1: The House, stable seasonal, series estimated p-values.

However, in the above case we are conducting multiple testing and thus a few issues arise. Using a 5% significance level implies that there is a 5% chance that the p-value will in fact be insignificant when it appears significant. Considering that when dealing with monthly data we are conducting twelve such tests while for weekly, as the Stonehaven data, fifty two, it becomes evident that the interpretation of the outcome of these tests is at best ambiguous if not misleading and a simpler way of testing for stability is needed.

From figure 3.1 we can see that two of twelve p-values for the House data are close to 0.05 but none of them are significant, while figure 3.2 shows that four p-values are significant for the Unemployment data. The Kola p-values are all insignificant, figure 3.3. Hence, the p-values coincide with our previous observation of the behaviour of the series. However, inference is not always straightforward.




Figure 3.2: The Unemployment, trending seasonal, series estimated p-values.




Figure 3.3: The Kola, stable seasonal, series estimated p-values.




Figure 3.4: The House estimated p-values Q-Q Plot.

In addition, one could check whether the collected p-values follow a Uniform distribution by plotting a Quantile-Quantile (Q-Q) plot of them.

The Q-Q plots in figures 3.4, 3.5 and 3.6 all indicate a non stable seasonal pattern. This can be due to a violation of the assumption of independence.

Another approach is to sum the double negative logarithm of the p-values and check if it follows a Chi-squared distribution. If the p-values are insignificant then the sum of their the double negative logarithm follows a Chi-squared distribution with degrees of freedom equal to twice the number of observations in each cycle of the seasonal component, i.e. twenty four degrees of freedom for monthly data. Assuming it does, we conclude that the seasonality is stable.




Figure 3.5: The Unemployment estimated p-values Q-Q Plot.




Figure 3.6: The Kola estimated p-values Q-Q Plot.

	$X = \sum_{i=1}^{12} (-2\ln(p_i))$	$\text{prob}(X \sim \mathcal{X}^2(24))$
House Data Set	39.103	2.7×10^{-2}
Unemployment Data Set	73.288	6.871×10^{-07}
Kola Data Set	9.675	9.958×10^{-1}

From the table above, only the Kola data set is deemed stable. The Unemployment data set is identified as trending while the House series is also deemed to have a trending seasonal pattern.

Additionally, there are a number of assumptions for the above tests to show accurate results. The most commonly violated one being the assumption that the p-values are independent to each other, while assumptions of homoscedasticity and normality of the logarithm of the p-values are also made. Often time series data is made from monthly or even weekly observations and therefore there is high concern that the p-values may in fact be correlated. More often than not, one will notice that there is high autocorrelation between some values and that could be responsible for the failure of the tests to verify the suspicion of a stable seasonal pattern.

Looking at p values when fitting harmonics or straight lines can be helpful. However, the problem of multiplicity arises from the sequential tests, since for monthly data one would need to look at twelve of these p values while for weekly fifty two! Q-Q plots and Fisher's test can be applied to deal with the multiplicity problem but then the assumption of independence is forced. An assumption that is commonly violated. The accuracy of these tests depends on the p values being independent but in time series they are highly dependent.

In the following section formal testing methodologies are presented.

3.2 (Seasonal) Unit Root Tests

3.2.1 Introduction

The early analysis of time series consisted of a (seasonal) decomposition of a series into trend and seasonal components. These two estimated components contain important information and reward careful examination and interpretation.

In many sciences the information contained in the aforementioned components is considered of essence in order to understand the behavior and dynamics of the series but particularly in economics a number of methods have been devised to estimate long-run economic relationships and therefore the majority of articles cited in this chapter will be of economic areas.

Many different issues might occur when one performs a regression with different orders of integrated variables. To avoid this problem, one has to be very careful during the identification stage of the model to identify the correct order of integration. In the context of ARIMA models, identifying the order of integration is equivalent to determining the parameter d in the ARIMA (p, d, q) model. The Box–Jenkins approach (Box & Jenkins, 1976) involved the use of graphs of the autocorrelation function for determining the parameter d . The recent developments of unit root tests is simply the use of formal statistical tests in place of the visual inspection of the graphs of the autocorrelation function (Maddala & Kim, 1998).

Phillips & Perron (1988) state that:

“One major field of application where the hypothesis of a unit root has important implications is economics. This is because a unit root is often a theoretical implication of models which postulate the rational use of information that is available to economic agents. Formal statistics tests of the unit root hypothesis are of additional interest to economists because they can help to evaluate the nature of the nonstationarity that most macroeconomic data exhibit. In particular, they help in determining whether the

the trend is stochastic, through the presence of a unit root, or deterministic, through the presence of a polynomial time trend.”

There is no question that unit root tests can provide useful results to contribute to the investigation and evaluation of the trend and seasonal components; however, they do come with limitations as a number of different assumptions need to be made before one can use them effectively. In all likelihood, the most important step before conducting a unit root test is the correct specification of the deterministic trend in a model. As mentioned before, most time series models include both a deterministic and a stochastic trend and the miss-specification of the deterministic trend will affect the identification of an underlying stochastic one by shadowing or overemphasizing it. If a time series has a unit root then this is the indication of an existing stochastic trend; thus, it is of great importance that the deterministic trend, one incorporates in the model to be tested, is in fact appropriate. The critical values given for the different unit root tests also differ according to the various forms of deterministic trends in the tested models. Due to this fact, the unit root testing procedures are not as straightforward as conventional testing procedures (Maddala & Kim, 1998).

Over recent years there has been a continuous development of new unit root testing frameworks in an effort to overcome drawbacks and limitations of the earlier ones, and different paths have been created. Early on, the tests were developed to distinguish between a stationary and a nonstationary process; testing for a zero frequency unit root but were then extended to analyze seasonal patterns; testing for unit roots at all or some seasonal frequencies. In this section we will try to give a broad overview of the best known and most commonly used tests along with some of their development over the years; the purpose being to use unit root tests to evaluate trend and seasonal components in biological series and thus expand their application in an area other than economics. An effort was made to sort the tests by first arranging them in a chronological order - not very strictly so as to separate between unit root tests and seasonal unit root tests - and also classify them in two broad categories according to the approach they use towards time series analysis; whether the analysis uses (seasonal) ARIMA models or the structural time series models. Since a full review

of all existing unit root testing techniques is beyond the scope of this chapter, the reader is referred to Maddala & Kim (1998) for a more detailed discussion. At the end of this chapter two diagrams are summarizing the discussed tests, Figure 3.7 the ones using ARIMA models and Figure 3.8 the ones that use structural time series models.

3.2.2 ARIMA Models

Dickey - Fuller Test

The first unit root testing framework was introduced by Dickey & Fuller (1979). The realization that the number of existing unit roots is equal to the parameter d in the ARIMA (p, d, q) model, motivated Dickey and Fuller to replace the somewhat intuitive visual inspection of the sample autocorrelation function with a formal testing procedure for unit roots. The test, known as the standard DF test, is of the null hypothesis that a first-order autoregressive, AR(1), process contains a unit root against the alternative that it is stationary and is based on the assumption of independently and identically distributed (iid) errors.

As a basis for the test the following autoregressive model is used:

$$X_t = \rho X_{t-1} + \epsilon_t, \quad t = 1, 2, \dots, \quad (3.1)$$

where $X_0 = 0$, ρ is a real number and $\{\epsilon_t\}$ is a sequence of independent normal random variables with mean zero and variance σ^2 [i.e., $\epsilon_t \sim \text{NID}(0, \sigma^2)$].

For the time series X_t to converge (as $t \rightarrow \infty$) to a stationary time series $|\rho| < 1$ is needed. Given n observations, the Maximum Likelihood Estimator (MLE) of ρ is the OLS estimator

$$\hat{\rho} = \left(\sum_{t=1}^n X_{t-1}^2 \right)^{-1} \sum_{t=1}^n X_t X_{t-1}. \quad (3.2)$$

Then a t test statistic is formed:

$$\hat{t}_\rho = \frac{\hat{\rho} - 1}{SE(\hat{\rho})} \quad (3.3)$$

where $SE(\hat{\rho})$ is the standard error of the OLS estimator. The asymptotic distribution of the statistic is a functional of the Wiener process,

$$\hat{t}_\rho \rightarrow \frac{W(1)^2 - 1}{2 \left\{ \int_0^1 W(r)^2 dr \right\}^{\frac{1}{2}}} \quad (3.4)$$

where W is a standard Wiener process.

Wiener Process or Brownian Motion:

A stochastic process $\{X(t), t \geq 0\}$ is said to be Brownian Motion process if:

1. $X(0) = 0$;
2. $\{X(t), t \geq 0\}$ has stationary and independent increments;
3. for every $t > 0$, $X(t) \sim N(0, \sigma^2 t)$.

In the above model no drift or linear trend was included and the derived asymptotic distribution, for the null hypothesis, is under the assumption that the data generating process (DGP) had no drift and no linear trend in it. The DF test can, however, be used when a drift and/or linear trend are included in the model and depending on whether the DGP has a drift and/or linear trend in it the critical values change. The test statistic stays the same for all cases but for the case of the linear trend the asymptotic distribution of the statistic, under the hypothesis of a unit root, is a different functional of the Wiener process. Fuller in 1976 provided tables with

the various critical values needed for the above mentioned test statistics. The null hypothesis of a unit root is rejected when the test statistic is smaller than the given critical value (Dickey & Fuller, 1979).

This is a parametric test for a zero frequency unit root but can also be used to test the validity of a given co-integrating relationship (Harvey, 2001). An issue with this test is that the normal test significance level is affected when the error terms ϵ_t in (3.1) are autocorrelated, as is often the case with time series data. In addition, as Dickey and Fuller mention : *“There is evidence that these tests are biased tests, accepting the null hypothesis more than 95 percent of the time for ρ close to, but less than, one.”*

In a recent paper (da Silva Lopes, 2004) investigates the effects of deterministic seasonality in DF tests and concludes that: *“...the common perception that deterministic seasonality has nothing to do with testing the long-run properties of the data is incorrect”* and suggests ways to account for it.

Augmented Dickey - Fuller Test

For a wide class of errors which allows some heterogeneity and serial correlations in errors, a different regression which contains lagged differences was suggested:

$$\Delta X_t = \alpha + \rho X_{t-1} + \sum_{i=1}^p \beta_i \Delta X_{t-i} + \epsilon_t, \quad (3.5)$$

where the first differences series $\{\Delta X_t\}$ has a stationary $AR(p)$ representation with a known p . This is the augmented Dickey-Fuller test (ADF) (Maddala & Kim, 1998). Augmenting the regression with lagged differences intends to remove the serial correlation from the disturbance (Chambers & McGarry, 2002).

The nominal significance level of 5% under the null hypothesis ($\rho = 1$) is held better, if p is larger (Maddala & Kim, 1998). However, it has been noted that then the power

of the test decreases, i.e. the tests can no longer distinguish between a process with unit roots and a stationary process with ρ values close to, but less than, one- much like the DF test . Thus, there is a trade-off between the validity and the power of the test in choosing p . Some authors suggest using information based rules such as the Akaike Information Criterion (AIC) or the Schwarz Bayesian Information Criterion (BIC) for the selection of p , while others argue using sequential rules (see Maddala & Kim (1998)).

Furthermore, the ADF test only works well in the presence of seasonal unit roots in the DGP if the test regression is *sufficiently* augmented with lags; otherwise, over-rejections of the null will occur (da Silva Lopes, 2004). In the same article Lopes also mentions that the ADF test does not work very well with small samples.

Finally (Busetti & Taylor, 2003) discuss the presence of unattended structural breaks in the sample; *“a process which is stochastically stationary about a deterministic component subject to structural breaks can display properties very similar to a unit root process. . . .the conventional ADF tests cannot reject the unit root null hypothesis, even asymptotically, where a broken trend exists.”*

Said - Dickey (SD) Test

The ADF test was extended to the more general case where the series of first difference in (3.5) are of the general $ARMA(p, q)$ form with the parameters p and q unknown and serially correlated errors. Said & Dickey (1984) showed that it is possible to approximate an $ARIMA(p, d, q)$ model by an autoregression whose order is a function of the sample size n . By the use of OLS the coefficients in this autoregressive approximation are estimated and then used as statistics whose limit distributions coincide with the ones tabulated by Dickey and listed by Fuller (1976).

The procedure will be given, here, for $p = d = 1$ but in the same paper it is extended to the general p, q case with iid errors. The model considered is as in (3.1) where now:

$$\epsilon_t = \alpha\epsilon_{t-1} + e_t + \beta e_{t-1} \quad (t = \dots, -2, -1, 0, 1, 2, \dots), \quad (3.6)$$

where it is assumed that $|\alpha| < 1$, $|\beta| < 1$, $X_0 = 0$ and $\{e_t\}$ is a sequence of iid random variables. The null hypothesis to be tested is $H_0 : \rho = 1$. Then,

$$e_t = \sum_{j=0}^{\infty} (-\beta)^j (\epsilon_{t-j} - \alpha\epsilon_{t-j-1}) \quad (3.7)$$

and it follows that

$$X_t - X_{t-1} = (\rho - 1)(X_{t-1}) + (\alpha + \beta)(\epsilon_{t-1} - \beta\epsilon_{t-2} + \beta^2\epsilon_{t-3} - \dots) + e_t. \quad (3.8)$$

Under the null hypothesis of a unit root, $\epsilon_t = X_t - X_{t-1}$. This motivates us to estimate the coefficients in (3.8) by regressing the first difference, ΔX_t , on X_{t-1} , ΔX_{t-1} , ΔX_{t-2} , \dots , ΔX_{t-k} where k is a suitable integer. For the estimators of the coefficients in (3.8) to be consistent it is essential to let k be a function of the number of observations, n . In addition, it is assumed that $n^{-1/3}k \rightarrow 0$ and that there exist $c > 0$, $r > 0$ such that $ck > n^{1/r}$, Said and Dickey (1984).

When the stated conditions are met the limiting distribution of the t-statistics of the coefficient on the lagged dependent variable X_{t-1} has the same Dickey & Fuller (1979) distribution as when the errors are iid.

Since, this test is an extension of the ADF test it suffers from the same drawbacks. In addition (Maddala & Kim, 1998) highlight that the actual size of the test may deviate greatly from the nominal size if the order or the autoregressive correction is not increased as the sample size increases; in order to accommodate the additional effect of the correlation structure of the residuals. Thus, if there are important moving average components in the structure of the series $\{X_t\}$ a large number of

nuisance parameters may be needed in the estimation. Additionally, one needs to remember that an effective observation is lost for each extra lag of ΔX_t introduced, leading one to deduce that this approach may in fact have significantly lower power when moving average terms are more important than if the errors were iid.

Phillips Test

All the testing frameworks mentioned above are confined to the case where the sequence of innovations driving the model are independent with a common variance. Furthermore, it is frequently assumed that these innovations ϵ_t are iid $(0, \sigma^2)$ or, even further, that they are iid and follow the normal distribution. However, independence and homoscedasticity are strong assumptions that are often violated when working with time series data (Phillips, 1987). In addition, from economic theory, these assumptions are false in the context of aggregate time series that maybe characterized as a random walk. Thus (Phillips, 1987) stated that : *“For both empirical and theoretical considerations, therefore, it is important to develop tests for unit roots that do not depend on these conditions.”*

Consequently (Phillips, 1987) provided asymptotic theory for the least squares regression estimator and the associated regression t statistic which allows for quite general weakly dependent and heterogeneously distributed innovations; performing non-parametric adjustments to the existing test statistics.

Let $\{X_t\}_{t=1}^{\infty}$ be a stochastic process generated in discrete time according to (3.1) and

$$\rho = 1 \tag{3.9}$$

Under (3.9) the representation $X_t = S_t + X_0$ is formed in terms of the partial sum $S_t = \sum_1^t \epsilon_j$ of the innovation sequence $\{\epsilon_j\}$ in (3.1) and the initial condition X_0 . Phillips defines $S_0 = 0$ and assumes that X_0 has a certain specific distribution.

To be precise it is assumed that $\{\epsilon_t\}_t^\infty$ is a sequence of random variables that satisfy the following conditions:

1.

$$E(\epsilon_t) = 0, \text{ all } t;$$

2.

$$\sup_t E|\epsilon_t|^\beta < \infty \text{ for some } \beta > 2;$$

3.

$$\sigma^2 = \lim_{T \rightarrow \infty} E(T^{-1}S_T^2) \text{ exists and } \sigma^2 > 0;$$

4.

$\{\epsilon_t\}_t^\infty$ is strong mixing with mixing coefficients α_m that satisfy:

$$\sum_t^\infty \alpha_m^{1-2/\beta} < \infty.$$

These conditions allow for both temporal dependence and heteroscedasticity in the process $\{\epsilon_t\}_t^\infty$; while condition (4) controls the extent of temporal dependence so that, although there may be substantial dependence amongst recent events, events which are separated by long intervals of time are almost independent (Phillips, 1987). For the definition of strong mixing and mixing coefficients α_m the reader is referred to White (1984).

Then the new test statistics are:

$$Z_\rho = T(\hat{\rho} - 1) - (1/2)(s_{Tl}^2 - s_\epsilon^2) / \left(T^{-2} \sum_1^T X_{t-1}^2 \right) \quad (3.10)$$

and

$$Z_t = (\hat{\rho} - 1) \left(\sum_1^T X_{t-1}^2 \right)^{t-2} / s_{Tl} - (1/2)(s_{Tl}^2 - s_\epsilon^2) \left[s_{Tl} \left(T^{-2} \sum_1^T X_{t-1}^2 \right)^{1/2} \right]^{-1} \quad (3.11)$$

where T is the sample size, l is the lag truncation number and $\hat{\rho} = \sum_1^T X_t X_{t-1} / \sum_1^T X_{t-1}^2$ is the OLS estimator of ρ . Furthermore, s_ϵ^2 and s_{Tl}^2 , under the null hypothesis (3.9), are the consistent estimators:

$$s_\epsilon^2 = T^{-1} \sum_1^T (X_t - X_{t-1})^2 = T^{-1} \sum_1^T \epsilon_t^2 \quad (3.12)$$

and

$$s_{Tl}^2 = T^{-1} \sum_1^T \epsilon_t^2 + 2T^{-1} \sum_{\tau=1}^l \sum_{t=\tau+1}^T \epsilon_t \epsilon_{t-\tau} \quad (3.13)$$

of σ_ϵ^2 , the variance of the innovation process $\{\epsilon_t\}$, and $\sigma_T^2 = \text{var}(T^{-1/2}S_T)$, respectively.

An interesting feature of the new test statistics is that their limiting distributions are identical to those found in earlier work under the assumption of iid errors. In particular, as T tends to infinity,

$$Z_\rho \implies \frac{(W(1)^2 - 1)/2}{\int_0^1 W(t)^2 dt} \quad (3.14)$$

and

$$Z_t \implies \frac{(W(1)^2 - 1)/2}{\{\int_0^1 W(t)^2 dt\}^{1/2}} \quad (3.15)$$

where once again critical values are given by Fuller (1976).

The above test can be used in a broad category of models. Phillips (1987) says:

“It applies, for example, to virtually any ARMA model with a unit root and even ARMAX systems with a unit root and with stable exogenous processes that admit a Wold decomposition.” One of the benefits of the Phillips test is that it is not necessary to identify the model in order to apply the test and consistently estimate the existence of a unit root in the series. Furthermore, it can easily be performed on models with a drift and a time trend (Phillips, 1987).

However (Chambers & McGarry, 2002) highlight that the t-ratio of this test suffers nuisance parameter dependencies that are a reflection of the dynamics in the disturbance making the inference of the outcome difficult.

Phillips - Perron (PP) test

Phillips & Perron (1988) extended the above study of Phillips to the cases where (a) a drift, and (b) a drift and a linear trend are included in the specification of the model. These extensions are important for practical applications, where the presence of a nonzero drift is very common. Moreover, in many cases and, particularly with economic time series, the main competing alternative to the presence of a unit root is deterministic linear time trend. It is therefore important that regression tests for unit roots allow for this possibility.

The models considered are driven by a sequence of innovations denoted by $\{\epsilon_t\}$ which satisfies the four conditions as in Phillips test. Let $\{X_t\}$ be a time series generated by (3.1) and (3.9). Initial conditions are set at $t = 0$ and y_0 may be any random variable, including a constant, whose distribution is fixed and independent of the sample size T .

The two least squares regression equations are considered

$$X_t = \hat{\mu} + \hat{\rho}X_{t-1} + \hat{\epsilon}_t \quad (3.16)$$

$$X_t = \tilde{\mu} + \tilde{\beta}(t - \frac{1}{2}T) + \tilde{\rho}X_{t-1} + \tilde{\epsilon}_t, \quad (3.17)$$

where $(\hat{\mu}, \hat{\rho})$ and $(\tilde{\mu}, \tilde{\beta}, \tilde{\rho})$ are the conventional least squares regression coefficients. The regression t statistics are formed and their limiting distributions derived; all depending upon the nuisance parameters σ^2 and σ_ϵ^2 , which is an obstacle to inference in the case where $\sigma^2 \neq \sigma_\epsilon^2$. Using the consistent estimators from Phillips (1987) - replacing the $\{\epsilon_t\}$ for the residuals from (3.16) and (3.17)- transformations of the test statistics are used; which eliminate the nuisance parameters asymptotically.

Furthermore, since the estimator in (3.13) is not constrained to be nonnegative, as it was defined can be negative when there are large negative sample covariances, simple modifications to it are suggested. For example the weighted variance estimators:

$$\hat{\sigma}_{Tl}^2 = T^{-1} \sum_{t=1}^T \hat{\epsilon}_t^2 + 2T^{-1} \sum_{s=1}^l w_{sl} \sum_{t=s+1}^T \hat{\epsilon}_t \hat{\epsilon}_{t-s}, \quad (3.18)$$

$$\tilde{\sigma}_{Tl}^2 = T^{-1} \sum_{t=1}^T \tilde{\epsilon}_t^2 + 2T^{-1} \sum_{s=1}^l w_{sl} \sum_{t=s+1}^T \tilde{\epsilon}_t \tilde{\epsilon}_{t-s}, \quad (3.19)$$

where $w_{sl} = 1 - s/(l + 1)$, is the triangular window, while other windows could be used.

The limit distributions of the new test statistics developed here are expressed as functionals of standard Brownian motion and are the same as those tabulated by Fuller (1976). This means that their tests may be used with existing tabulations even though they allow for much more general time series specifications. In the same paper they also study the asymptotic local power properties of their tests using the theory of near-integrated processes and provide some simulation evidence on the finite sample performance of the new tests Phillips & Perron (1988).

Variance ratio tests

Cochrane (1988) proposed a variance ratio (VR) statistic to evaluate the presence of a unit root by measuring the degree of persistence in a time series. The variance ratio is the variance of the k th lag difference of the series divided by k -times the

variance of its first-difference:

$$VR_k = \frac{V_k}{V_1} \quad (3.20)$$

where $V_k = \text{var}(X_t - X_{t-k})/k$. The VR_k can be expressed as the weighted sum of the correlations r_j between ΔX_t and ΔX_{t-j} :

$$VR_k = 1 + 2 \sum_{j=1}^k \left[1 - \frac{j}{k+1}\right] r_j \quad (3.21)$$

As an estimator for the variance ratio Cochrane used:

$$\hat{V}R_k = \frac{k^{-1} \text{var}(X_t - X_{t-k})}{\text{var}(X_t - X_{t-1})} \left(\frac{T}{T - k + 1} \right) \quad (3.22)$$

which is an unbiased estimator and as $T \rightarrow \infty$, $k \rightarrow \infty$, and $k/T \rightarrow 0$, has a limiting normal distribution with mean VR_k and variance $4kVR_k^2/3T$.

In practice one considers VR_k for different values of k and rejects models when at least some of the VR statistics provide evidence against it. However, further papers from different authors, (see Maddala & Kim (1998)) suggest that this penalizes the null hypothesis since the tests are correlated and thus argue that there are substantial size distortions with the use of the asymptotic approximations and furthermore with the use of sequential testing procedure with different values of k . Instead one could use a joint test that takes under consideration the different VR_k statistics and by using Monte Carlo methods derives the critical values for them.

Dickey, Hasza and Fuller (DHF) test

Dickey *et al.* (1984) researched regression estimators of coefficients in seasonal autoregressive models. They expanded the DF test to look for unit roots in all the seasonal and zero frequencies.

Let the time series X_t satisfy

$$X_t = \rho_d X_{t-d} + \epsilon_t, \quad t = 1, 2, \dots, \quad (3.23)$$

where $X_{-d+1}, X_{-d+2}, \dots, X_0$ are initial conditions and the ϵ_t are iid $(0, \sigma^2)$ random variables. The above model is a simple seasonal time series model in which d is equal to the period of the present seasonality.

In their paper they consider several regression-type estimators of ρ_d and compute percentiles of their distributions under the hypothesis that $\rho_d = 1$. The null hypothesis for this test is that there are unit roots at the 0 and all seasonal frequencies with the alternative of no unit roots (Maddala & Kim, 1998).

The first estimator of ρ_d is the OLS estimator defined as

$$\hat{\rho}_d = \left(\sum_{t=1}^n X_{t-d}^2 \right)^{-1} \sum_{t=1}^n X_{t-d} X_t. \quad (3.24)$$

If the initial conditions are fixed and ϵ_t are normal, $\hat{\rho}_d$ is the MLE. The Studentized regression statistic for testing the hypothesis $H_0 : \rho_d = 1$ is

$$\hat{\tau}_d = (\hat{\rho}_d - 1) \left[\left(\sum_{t=1}^n X_{t-d}^2 \right)^{-1} S^2 \right]^{-1/2}, \quad (3.25)$$

where

$$S^2 = (n-1)^{-1} \sum_{t=1}^n (X_t - \hat{\rho}_d X_{t-d})^2; \quad (3.26)$$

are included in the standard output from a computer regression of $X_t^* = X_t - X_{t-d}$ on X_{t-d} .

An alternative model for seasonal data is the stationary model in which the observations satisfy

$$X_t = \rho_d X_{t-d} + \epsilon_t, \quad |\rho_d| < 1. \quad (3.27)$$

Where, for normal stationary X_t satisfying (3.27),

$$X_t = \rho_d X_{t+d} + \epsilon_t, \quad \epsilon_t \sim NID(0, \sigma^2). \quad (3.28)$$

An alternative estimator of ρ_d , which they call the symmetric estimator is

$$\tilde{\rho}_d = \left(2 \sum_{t=1}^n (X_t X_{t-d}) \right) / \left(\sum_{t=1}^n (X_t^2 + X_{t-d}^2) \right), \quad (3.29)$$

and define the associated Studentized statistic as

$$\tilde{\tau}_d = 2^{\frac{1}{2}} (\tilde{\rho}_d - 1) \left[\left\{ \sum_{t=1}^n (X_t^2 + X_{t-d}^2) \right\}^{-1} S^2 \right]^{-1/2}, \quad (3.30)$$

where

$$S^2 = (2n - 1)^{-1} \sum_{t=1}^n [(X_t - \tilde{\rho}_d X_{t-d})^2 + (X_{t-d} - \tilde{\rho}_d X_t)^2]. \quad (3.31)$$

Models (3.23) and (3.27) both imply that the time series X_t has a zero mean, something that is rarely encountered in practice. Therefore, an alternative regression model is considered

$$X_t = \sum_{i=1}^d \theta_i \delta_{it} + \rho_d X_{t-d} + \epsilon_t, \quad t = 1, 2, \dots, \quad (3.32)$$

where

$$\delta_{it} = \begin{cases} 1, & \text{if } t = i(\text{mod } d) \\ 0, & \text{otherwise} \end{cases} \quad (3.33)$$

and $\{\epsilon_t\}$ is a sequence of iid $(0, \sigma^2)$ random variables. The regression of X_t on $\delta_{1t}, \delta_{2t}, \dots, \delta_{dt}, X_{t-d}$ for $t=1, 2, \dots, n$, produces coefficients $\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_d, \hat{\rho}_{\mu d}$. The Studentized regression statistic associated with $\hat{\rho}_{\mu d} - 1$ is denoted by $\hat{\tau}_{\mu d}$.

Assuming that $|\rho_d| < 1$, a reparameterized version of model (3.32) is

$$X_t - \sum_{i=1}^d \delta_{it} \mu_i = \rho_d \left(X_{t-d} - \sum_{i=1}^d \delta_{it} \mu_i \right) + \epsilon_t, \quad (3.34)$$

where

$$\theta_i = (1 - \rho_d) \mu_i, \quad i = 1, 2, \dots, d. \quad (3.35)$$

Under the model (3.34) the hypothesis $\rho_d = 1$ implies that $\theta_i = 0$ regardless of the value of μ_i . Thus, specifying $\rho_d = 1$ in the model (3.34) allows μ_i to assume any value. Under the alternative of $|\rho_d| < 1$, however, μ_i is an identified parameter and should be estimated (Dickey *et al.*, 1984).

Two estimators for μ_i are then considered for the stationary model. The first is that defined by the regression estimators for (3.32), and the second is the seasonal mean $\tilde{\mu}_i$ defined by

$$\tilde{\mu}_i = (n_i + 1)^{-1} \sum_{j=0}^{n_i} X_{-d+i+dj}, \quad i = 1, 2, \dots, d, \quad (3.36)$$

where n_i is the greatest integer not exceeding $(n + d - i)/d$. The estimator $\tilde{\mu}_i$ can be used to define a symmetric estimator of ρ_d , analogous to (3.29), as follows

$$\tilde{\rho}_{\mu d} = \left\{ \sum_{t=1}^n (x_t^2 + x_{t-d}^2) \right\}^{-1} 2 \sum_{t=1}^n x_t x_{t-d}, \quad (3.37)$$

where

$$x_t = X_t - \sum_{i=1}^d \tilde{\mu}_i \delta_{it}. \quad (3.38)$$

However, if the initial conditions are fixed, the θ_i (or μ_i) should be estimated using the regression model.

The percentiles of the distributions for time series that have unit roots at the seasonal lag are computed in the same paper (Dickey *et al.*, 1984), in which the above method was introduced, by Monte Carlo integration for finite samples and by analytic techniques and Monte Carlo integration for the limit case. The tabled distributions may be used to test the hypothesis that a time series has a seasonal unit root.

A major drawback of this test is that it does not allow for unit roots at some but not all of the seasonal frequencies and that the alternative has a very particular form, namely that all the roots have the same modulus (Hylleberg *et al.*, 1990).

Hylleberg, Engle, Granger and Yoo (HEGY) Test

Hylleberg *et al.* (1990) developed a test that follows the Dickey-Fuller framework and looks at unit roots at all the seasonal frequencies as well as the zero frequency. In the same paper they provide tables of the critical values for the limiting distributions of the developed test statistics.

To test the hypothesis that the roots of $\varphi(B)$ lie on the unit circle against the alternative that they lie outside the unit circle, it is convenient to rewrite the autoregressive polynomial in the form:

$$\varphi(B) = \sum_{k=1}^p \lambda_k \Delta(B)(1 - \delta_k(B))/\delta_k(B) + \Delta(B)\varphi^*(B), \quad (3.39)$$

where the λ_k are a set of constants, $\varphi^*(B) = \varphi^{**}(B) + \sum \lambda_k$ with $\varphi^{**}(B)$ a (possibly infinite or rational) polynomial and

$$\delta_k(B) = 1 - \frac{1}{\theta_k} B \quad , \quad \Delta(B) = \prod_{k=1}^p \delta_k(B). \quad (3.40)$$

In this representation $\varphi(0) = \varphi^*(0)$ which is normalized to unity.

It is clear that the polynomial $\varphi(B)$ will have a root at θ_k if and only if $\lambda_k = 0$. Thus testing for unit roots can be carried out equivalently by testing for parameters $\lambda = 0$ in an appropriate expansion.

For example, to test for seasonal unit roots in quarterly data one would expand a polynomial $\varphi(B)$ about the roots $+1, -1, +i$, and $-i$ as $\theta_k, k = 1, \dots, 4$. Then, from (3.39),

$$\begin{aligned} \varphi(B) = & -\pi_1 B(1 + B + B^2 + B^3) - \pi_2(-B)(1 - B + B^2 - B^3) \\ & - (\pi_4 + \pi_3 B)(-B)(1 - B^2) + \varphi^*(B)(1 - B^4), \end{aligned} \quad (3.41)$$

where $\pi_1 = -\lambda_1$, $\pi_2 = -\lambda_2$, $2\lambda_3 = -\pi_3 + i\pi_4$, and $2\lambda_4 = -\pi_3 - i\pi_4$. Since $\varphi(B)$ is real, λ_3 and λ_4 had to be complex conjugates.

The data are assumed to be generated by a general autoregression of the form

$$\varphi(B)x_t = \epsilon_t, \quad (3.42)$$

and (3.41) is used to replace $\varphi(B)$, giving

$$\varphi^*(B)y_{4t} = \pi_1 y_{1t-1} + \pi_2 y_{2t-1} + \pi_3 y_{3t-2} + \pi_4 y_{3t-1} + \epsilon_t, \quad (3.43)$$

where

$$\begin{aligned} y_{1t} &= (1 + B + B^2 + B^3)x_t = S(B)x_t, \\ y_{2t} &= -(1 - B + B^2 - B^3)x_t, \\ y_{3t} &= -(1 - B^2)x_t, \\ y_{4t} &= (1 - B^4)x_t = \Delta_4 x_t. \end{aligned}$$

Equation (3.43) can be estimated by OLS, possibly with additional lags of y_4 to whiten the errors. To test the hypothesis that $\varphi(\theta_k) = 0$, where θ_k is either $+1, -1, +i$ or $-i$, one needs simply to test that λ_k is zero. For the root 1 this is simply a test

for $\pi_1 = 0$, and for -1 it is $\pi_2 = 0$. For the complex roots $+i$ and $-i$, λ_3 will have absolute value of zero only if both π_3 and π_4 equal zero which suggests a joint test. There will be no seasonal roots if π_2 and either π_3 or π_4 are different from zero, which therefore requires the rejection of both a test for π_2 and a joint test for π_3 and π_4 . To conclude that a series has no unit roots at all and is therefore stationary, one would have to establish that each of the π 's is different from zero (save possibly either π_3 or π_4).

The hypothesis tests can be amended to include a constant, seasonal dummies and a time trend. Then (3.43) becomes

$$\varphi^*(B)y_{4t} = \mu + \beta t + \sum_{k=2}^4 S_{kt} + \pi_1 y_{1t-1} + \pi_2 y_{2t-1} + \pi_3 y_{3t-2} + \pi_4 y_{3t-1} + \epsilon_t, \quad (3.44)$$

which can be estimated by OLS and statistics on the π 's can be used for inference. The asymptotic and finite sample distributions change (Beaulieu & Miron, 1993).

F-type statistics for the joint null hypotheses are computed, F_{1234} , F_{234} , F_{34} , against the alternative that they are not all equal to zero. Furthermore, t-statistics are computed for the one-sided 't' tests on π_1 , π_2 and π_3 and the two-sided 't' test on $\pi_4 = 0$. Critical values for the above tests are given in Hylleberg *et al.* (1990), where they present the aforementioned testing framework.

Nonetheless, the lag augmentation that is employed is only a partial solution to the problem caused by serial correlation between the innovations driving the process (Burrige & Taylor, 2001a,b). In addition, the HEGY test is sensitive to periodic heteroscedasticity (PH) amongst the innovations and may have less power for the zero frequency unit root than the DF test (da Silva Lopes, 2004).

Chambers McGarry

As demonstrated by Phillips (1987), testing for a unit root at the zero frequency can be carried out using a straightforward regression of a variable on its lag, with any neglected (stationary) dynamics captured by the disturbance process. However, the limiting distribution of the resulting t-ratio suffers from nuisance parameter dependencies that are a reflection of the dynamics in the disturbance, thus making inference difficult. A method of eradicating nuisance parameters from the limiting distributions is to perform the regression in the frequency domain and by doing so it is not necessary to be concerned with selecting the appropriate number of lagged variables to include in the regression, which always takes the same form in the frequency domain. Another advantage is that the setup allows for unit root testing in autoregressive moving average (ARMA) models as well as purely autoregressive models (Chambers & McGarry, 2002).

Consider a univariate process X_t having the autoregressive representation of (3.42), where $\epsilon_t \sim IID(0, \sigma_\epsilon^2)$, B denotes the lag operator, T denotes sample size, and $\varphi(z) = 1 - \sum_{j=1}^p \varphi_j z^j$ is a polynomial of order p , where $p \geq s$, the number of seasons. The polynomial $\varphi(z)$ can be expressed as the product of two lower-order polynomials in the form $\varphi(z) = a_s(z)b(z)$. In this representation $a_s(z)$ is a polynomial of order s having $0 \leq s_1 \leq s$ roots on the unit circle and $s - s_1$ roots outside the unit circle, while $b(z)$ is a polynomial of order $p - s$ that has all its roots outside the unit circle. It is, therefore, the polynomial $a_s(z)$ that captures possible seasonal integration in the process X_t . Attention will be focused on testing the number of roots of $a_s(z)$ that lie on the unit circle when $s = 4$, which is appropriate for testing for seasonal unit roots with quarterly data.

With the factorisation of $\varphi(z)$ described above and $s = 4$ it is possible to write (3.42) as

$$a_4(B)X_t = u_t \tag{3.45}$$

where $u_t = \phi(B)\epsilon_t$ is a stationary random disturbance in which $\phi(z) = 1 + \sum_{j=1}^{\infty} \phi_j z^j = b(z)^{-1}$ that satisfies $\sum_{j=1}^{\infty} j^2 \phi_j^2 < \infty$ and $\epsilon_t \sim \text{IID}(0, \sigma_\epsilon^2)$ with $\sigma_\epsilon^2 < \infty$.

Following HEGY, the polynomial $a_4(z) = 1 + \sum_{j=1}^4 a_j z^j$ can be decomposed as

$$a_4(z) = -\pi_1 z \alpha_1(z) - \pi_2 z \alpha_2(z) - \pi_3 z \alpha_3(z) - \pi_4 z \alpha_4(z) + (1 - z^4), \quad (3.46)$$

where $\alpha_1(z) = 1 + z + z^2 + z^3$, $\alpha_2(z) = -(1 - z + z^2 - z^3)$, $\alpha_3(z) = -z(1 - z^2)$ and $\alpha_4(z) = -(1 - z^2)$. The coefficients π_1, π_2, π_3 and π_4 correspond to the roots $1, -1, i$ and $-i$ respectively in the sense that if $\pi_j = 0$ then $a_4(z)$ possesses the corresponding root. Matching the coefficients in (3.46) with those of $a_4(z)$ yields the relationships $a_1 = -\pi_1 + \pi_2 + \pi_4$, $a_2 = -\pi_1 - \pi_2 + \pi_3$, $a_3 = -\pi_1 + \pi_2 - \pi_4$, and $a_4 = -\pi_1 - \pi_2 - \pi_3 - 1$. Alternatively, solving these expressions for the π_j in terms of the a_j yields $\pi_1 = -a_4(1)/4$, $\pi_2 = -(1 - a_1 + a_2 - a_3 + a_4)/4$, $\pi_3 = -(1 - a_2 + a_4)/2$ and $\pi_4 = (a_1 - a_3)/2$. Note that the term $(1 - z^4)$ is not multiplied by another polynomial, as in HEGY, in which a representation of $\varphi(z)$, rather than just $a_4(z)$, is sought and which results in a term of the form $a^*(z)(1 - z^4)$, where $a^*(z)$ is a finite polynomial. This term is required in HEGY to account for the additional terms in the polynomial involving $b(z)$, which in this approach has already been incorporated within the disturbance term u_t .

Let $y_t = (1 - B^4)X_t$ and define the variables

$$y_{jt} = \alpha_j(B)X_t \quad (j = 1, \dots, 4). \quad (3.47)$$

Then, using the representation (3.46) for $a_4(z)$, (3.45) may be written as the regression model

$$y_t = \sum_{j=1}^4 \pi_j y_{j,t-1} + u_t, \quad t = 1, 2, \dots, T, \quad (3.48)$$

which is in the form of the HEGY regression¹ but with one notable exception their representation of $\varphi(z)$ yields a variable on the left hand side which may be denoted $a^*(B)y_t$ and which incorporates the dynamics that are effectively associated with $b(z)$, in order to produce a white noise disturbance. In the spectral regression approach adopted here such dynamics are assigned to the disturbance term u_t and are treated non-parametrically via the use of appropriate spectral density estimates.

Defining the vectors $x_t = [y_{1,t-1}, y_{2,t-1}, y_{3,t-1}, y_{4,t-1}]'$ and $\pi = [\pi_1, \pi_2, \pi_3, \pi_4]'$ (3.48) may be written as the regression model

$$y_t = x_t' \pi + u_t, \quad t = 1, 2, \dots, T. \quad (3.49)$$

The frequency domain tests of the restrictions $\pi_j = 0$ ($j = 1, \dots, 4$) considered here are based on the frequency domain regression estimator of π defined by

$$\hat{\pi} = \left[\frac{1}{2M} \sum_{j=-M+1}^M \hat{f}_{xx}(\omega_j) \hat{f}_{\hat{u}\hat{u}}(\omega_j)^{-1} \right]^{-1} \left[\frac{1}{2M} \sum_{j=-M+1}^M \hat{f}_{xy}(\omega_j) \hat{f}_{\hat{u}\hat{u}}(\omega_j)^{-1} \right]. \quad (3.50)$$

In the above definition of $\hat{\pi}$, $\hat{f}_{ab}(\omega)$ denotes a nonparametric estimate of the spectral density function of two (possibly vector) random processes a_t and b_t , given by

$$\hat{f}_{ab}(\omega) = \frac{1}{2\pi} \sum_{n=-M}^M k\left(\frac{n}{M}\right) C_{ab}(n) e^{-in\omega}, \quad (3.51)$$

where $k(\cdot)$ is a bounded, even, kernel function satisfying $k(0) = 1$ and $k(x) = 0$ for $x \notin [-1, 1]$, M is a bandwidth parameter, $\omega_j = \pi j/M$, and

¹The variable y_{4t} is denoted y_{3t} by HEGY and y_{3t} is HEGY's y_{3t-1} .

$$C_{ab}(n) = \begin{cases} T^{-1} \sum_{t=1}^{T-n} a_t b'_{t+n}, & n \geq 0 \\ T^{-1} \sum_{t=|n|+1}^T a_t b'_{t-|n|}, & n < 0. \end{cases} \quad (3.52)$$

In addition, when $T \rightarrow \infty$, the bandwidth parameter $M \rightarrow \infty$ but in a way such that $M/T^{1/2} \rightarrow 0$.

The variable $\hat{u}_t = y_t - x_t' \hat{\pi}_{OLS}$ denotes the residual from a time domain regression of y_t on x_t , where $\hat{\pi}_{OLS}$ is the OLS estimator. The estimated asymptotic covariance matrix of $\hat{\pi}$ is

$$V_T = \frac{1}{T} \left[\frac{1}{2M} \sum_{j=-M+1}^M \hat{f}_{xx}(\omega_j) \hat{f}_{\hat{u}\hat{u}}(\omega_j)^{-1} \right]^{-1}. \quad (3.53)$$

Two types of test statistic are considered for testing the individual null hypothesis that $\pi_j = 0$ ($j = 1, \dots, 4$). The first type of statistic is simply $T\hat{\pi}_j$ ($j = 1, \dots, 4$) while the second type is the t-ratio defined by

$$t_j = \frac{\hat{\pi}_j}{V_{T,jj}^{1/2}}, \quad j = 1, \dots, 4, \quad (3.54)$$

where $V_{T,jj}$ denotes the j 'th diagonal element of V_T . The hypotheses that $\pi_3 = \pi_4 = 0$ and that $\pi_1 = \pi_2 = \pi_3 = \pi_4 = 0$ can be tested using Wald statistics constructed from the unrestricted estimator $\hat{\pi}$. They will be denoted J_{34} and J_{1234} respectively, and are defined by

$$J_i = \hat{\pi}' R_i' [R_i V_T R_i']^{-1} R_i \hat{\pi}, \quad i = 1234, 34. \quad (3.55)$$

where $R_{1234} = I_4$, $R_{34} = [0_2, I_2]$, 0_n denotes an $n \times n$ matrix of zeros, I_n denotes an $n \times n$ identity matrix and V_T is the covariance matrix defined in (3.53). The

limiting distributions of these statistics, under the stated assumptions, are given by Chambers & McGarry (2002) when they propose their method.

Bootstrapping the HEGY tests

Burridge & Taylor (2003) argue that the effects on the sampling distributions of the HEGY statistics induced by shocks which are serially correlated, periodically heteroscedastic, and possibly asymmetric, can be successfully accommodated by the use of a bootstrap. In addition, the bootstrap corrects the adverse effects of data-dependent lag selection seen in the conventional augmented HEGY tests and stops inflation of test significance levels above their nominal levels. Finally, the bootstrap delivers estimated tail probabilities which are the quantities required for inference and so there is no issue of the unreliability of tabulated critical values.

The model considered is a quarterly time series which can be written as

$$a(B)X_{4t+s} = u_{4t+s} + \mu_s + \beta_s(4t + s), \quad (3.56)$$

$$\phi(B)u_{4t+s} = \epsilon_{4t+s}, \quad (3.57)$$

where a is a fourth order polynomial in the usual lag operator, B . This allows for periodic intercepts and time trends through μ_s and β_s respectively. The shocks $\{u_{4t+s}\}$ are an $AR(m)$ process, in which the m roots of $\phi(z) = 0$ all lie outside the unit circle. In addition, they are allowed to be periodically heteroscedastic, and to have an asymmetric distribution, that is, we define the annualised vector innovation process $\epsilon_t = (\epsilon_{4t-3}, \epsilon_{4t-2}, \epsilon_{4t-1}, \epsilon_{4t})'$ and assume that $\epsilon_t \sim IID(0, -)$ with $- = \text{diag}(\sigma_{-3}^2, \dots, \sigma_0^2)$, and with finite fourth moments. The innovations are otherwise unrestricted. The shocks, $\{\epsilon_{4t+s}\}$, and hence $\{X_{4t+s}\}$, display PH unless $\sigma_s = \sigma$, for all s , Burridge and Taylor (2003).

In order to derive the HEGY tests, the polynomial $a(B)$ is factorised at the seasonal frequencies $\omega_k \equiv 2k\pi/4$, $k=0, 1, 2$, and expanded around the seasonal unit roots $\exp(\pm 2k\pi i/4)$, $k = 0, 1, 2$, to obtain the auxiliary regression equation

$$\Delta_4 X_{4t+s} = \mu_s^* + \beta_s^*(4t + s) + \sum_{j=1}^4 \pi_j y_{j,4t+s-1} + \sum_{j=1}^m \phi_j \Delta_4 X_{4t+s-j} + \epsilon_{4t+s}, \quad (3.58)$$

which may be estimated along $4t + s = m + 5, \dots, 4T$. The inclusion of seasonal level and trend dummies in (3.58), whose parameters μ_s^* and β_s^* , respectively, are linear mappings of μ_s and β_s of (3.56), $s = -3, \dots, 0$, ensures that the sampling distributions of the estimated coefficients on the transformed level variables, $y_{j,4t+s}$, $j = 1, \dots, 4$, and their associated t- and F-statistics are unaffected by the $\boldsymbol{\mu} \equiv (\mu_{-3}, \dots, \mu_0)$ and $\boldsymbol{\beta} \equiv (\beta_{-3}, \dots, \beta_0)$ parameters. The transformed level variables which correspond to the seasonal frequencies $\omega_k = 2k\pi/4$, are given by (3.47).

The tests of interest are the regression t-statistics, t_1, t_2, t_3 (one-sided) and t_4 (two-sided), together with the F-statistics, F_{34}, F_{234} and F_{1234} .

The bootstrap algorithm begins with estimation of (3.58), with the lag length, m , and the intervening lags to be retained, selected using the sequential elimination procedure advocated by Beaulieu & Miron (1993, pp.318-319). That is, a maximum lag, m_{max} , and the deterministics are specified, and the test equation estimated. Thereafter, if any lagged fourth differences have t-statistics smaller than 1.65 in absolute value (i.e. insignificant at an approximate level of 10%), the least significant lag is removed and the equation re-estimated. This continues until all the included lags are significant, at which point their estimated coefficients and the seven unit root test statistics are recorded, and their residual vector is stored.

The residuals for each quarter are stored separately, and a sample from each of their empirical distributions is drawn. These four independently drawn samples are merged, preserving the seasonal ordering, into the vector ϵ^* , which is then used

to construct a bootstrapped observed sequence under the null hypothesis via the equation,

$$\hat{\phi}(B)\Delta_4 X_{4t+s}^* = \epsilon_{4t+s}^* \quad (3.59)$$

initialized at $\mathbf{0}$.

The full fitting algorithm, using the same maximum lag, selection method, and the deterministic is then applied to the X_{4t+s}^* series and the resulting test statistics compared with the originals. This procedure is repeated for a large number of bootstrap samples and record how many of the bootstrapped statistics are more extreme (in the relevant tail/s) than the original, thus locating the latter in the bootstrap null cumulative density function (cdf).

An important feature of the fitted equation is that the test statistics and the residual vector are unaffected by the deterministic parameters under both the null and the alternatives. Thus, in simulating the null distribution by the bootstrap sampling from the seasonal residual empirical cdfs we need not incorporate the fitted deterministic parameters in the bootstrap samples, provided they are included in the test equation fit to those samples (Burrige & Taylor, 2003).

Pons

Pons (2004) asserts that temporal aggregation has important implications for seasonal time-series analysis, since this data transformation confuses non-observable seasonal cycles with observable ones. This sampling effect known as aliasing implies that when a particular quarterly seasonal unit root is detected, it is not possible to state whether the unit root is present at a monthly frequency with the same period, or at another monthly frequency with a period not observable at the quarterly interval.

In his paper (Pons, 2004) proposes to control effects from temporal aggregation by combining monthly and quarterly seasonal unit root tests and applies it to the HEGY test. It is argued, first, theoretically and then demonstrated with a Monte

Carlo experiment that it is possible to improve substantially the performance of the seasonal unit root test that only uses monthly information. Furthermore, the proposed methodology can be extended to all seasonal unit root tests that test for integration or no integration at the separate frequencies.

3.2.3 Structural Time Series Models

Unit root testing is normally carried out within a framework of autoregressive models. Auto-regressions are popular, especially in economics, because they are easy to fit. An alternative approach is to use structural time series models (STMs). These models are formulated in terms of unobserved components such as trends, seasonals, and cycles. The most interesting testing issues in structural time series models concern testing the null hypothesis that a particular component is deterministic against the alternative that it is stochastic and non-stationary. The non-stationarity in question appears under the alternative rather than under the null hypothesis. In the unit root tests, such as augmented Dickey-Fuller, the situation is reversed, but for unobserved components models this is not the natural way to proceed (Harvey, 2001).

Kwiatkowski, Phillips, Schmidt and Shin (KPSS) Test

Kwiatkowski *et al.* (1992) propose a test of the null hypothesis that an observable series is stationary around a deterministic trend. The series is expressed as the sum of deterministic trend, random walk and stationary error, and the test is the Lagrange Multiplier (LM) test of the hypothesis that the random walk has zero variance.

Under the additional assumptions that the random walk is normal and that the stationary error is normal white noise, the one-sided LM statistic for the trend stationarity hypothesis is the same as the Locally Best Invariant (LBI) test statistic. However, the assumption that the error is white noise is not credible in many applications, since it implies that under the null hypothesis the variable should have iid

deviations from the trend. Therefore, they derive the asymptotic distribution of the statistics under general conditions on the stationary error, and propose a modified version of the LM statistic that is valid asymptotically under these general conditions. The asymptotic distribution is nonstandard, involving higher-order Brownian bridges.

Let $X_t, (t = 1, 2, \dots, T)$, be the observed series to be tested for stationarity. It is assumed that it can be decomposed into the sum of a deterministic trend, a random walk, and a stationary error:

$$X_t = \beta t + r_t + \epsilon_t. \quad (3.60)$$

Here r_t is a random walk:

$$r_t = r_{t-1} + u_t, \quad (3.61)$$

where the u_t are iid $(0, \sigma_u^2)$. The initial value r_0 is treated as fixed and serves the role of an intercept. The stationarity hypothesis is simply that $\sigma_u^2 = 0$. Since ϵ_t is assumed to be stationary, under the null hypothesis X_t is trend-stationary. Kwiatkowski, Phillips, Schmidt and Shin also consider the special case of the model (3.60) in which case under the null hypothesis X_t is stationary around a level (r_0) rather than around a trend.

Let $e_t, (t = 1, 2, \dots, T)$, be the residuals from the regression of X on an intercept and time trend. Let $\hat{\sigma}_\epsilon^2$ be the estimate of the error variance from this regression (the sum of squared residuals, divided by T). Define the partial sum process of the residuals:

$$S_t = \sum_{i=1}^t e_i, \quad t = 1, 2, \dots, T. \quad (3.62)$$

Then the LM -and the LBI- statistic is

$$LM = \sum_{t=1}^T S_t^2 / \hat{\sigma}_\epsilon^2. \quad (3.63)$$

The 'long-run variance' of the model is defined as

$$\sigma^2 = \lim_{T \rightarrow \infty} T^{-1} E(S_T^2), \quad (3.64)$$

which will enter into the asymptotic distribution of the test statistic. A consistent estimator of σ^2 , say $s(l)^2$, can be constructed from the residuals e_t . Specifically, KPSS use an estimator of the form

$$s^2(l) = T^{-1} \sum_{t=1}^T e_t^2 + 2T^{-1} \sum_{s=1}^l w(s, l) \sum_{t=s+1}^T e_t e_{t-s} \quad (3.65)$$

as in Phillips (1987) or Phillips & Perron (1988).

Here $w(s, l)$ is an optional weighting function that corresponds to the choice of a spectral window. In the construction of the test, KPSS use the Bartlett² window $w(s, l) = 1 - s/(l + 1)$ which guarantees the non-negativity of $s^2(l)$. For consistency of $s^2(l)$, it is necessary that the lag truncation parameter $l \rightarrow \infty$ as $T \rightarrow \infty$. The rate $l = o(T^{1/2})$ will usually be satisfactory under both the null and the alternative (Kwiatkowski *et al.*, 1992).

For the tests of both the level-stationary and trend-stationary hypotheses, the denominator of the LM statistic in (3.63) is $\hat{\sigma}_\epsilon^2$, which converges in probability to σ_ϵ^2 . However, when the errors are not iid, the appropriate denominator of the test statistic is an estimate of σ^2 instead of σ_ϵ^2 .

First the level-stationary case is considered. The model is as in (3.60) with β set to zero, so that the residuals e_t are from a regression of y on intercept only; that is, $e_t = X_t - \bar{X} = X_t - \mu$. S_t is then the partial sum process of the residuals e_t as in (3.62). So the test statistic is

²In Phillips & Perron (1988) is mentioned as the triangular window.

$$\hat{\eta}_\mu = T^{-2} \sum_{t=1}^T S_t^2 / s^2(l) \quad (3.66)$$

with

$$\hat{\eta}_\mu \rightarrow \int_0^1 W_1(r)^2 dr, \quad (3.67)$$

under the assumptions of Phillips and Perron (1988). Here $W_1(r)$ is a standard Brownian bridge:

$$W_1(r) = W(r) - rW(1),$$

where $W(r)$ is a Wiener process (Brownian motion).

The analysis of the trend-stationary case is very similar to that of the level-stationary case. The model is now exactly as in (3.60). Let e_t be the residuals from a regression of X_t on intercept and trend, and let S_t be the partial sum process of the e_t as in (3.62). Furthermore, let $\hat{\eta}_\tau$ be the new test statistic, where the subscript τ indicates that we have extracted a mean and a trend from X , and serves to distinguish the trend-stationary case from the level-stationary case. Then

$$\hat{\eta}_\tau = T^{-2} \sum_{t=1}^T S_t^2 / s^2(l) \quad (3.68)$$

and its asymptotic distribution is

$$\hat{\eta}_\tau \rightarrow \int_0^1 W_2(r)^2 dr \quad (3.69)$$

where the second-level Brownian bridge $W_2(r)$ is given by

$$W_2(r) = W(r) + (2r - 3r^2)W(1) + (-6r + 6r^2) \int_0^1 W(s) ds. \quad (3.70)$$

The upper tail critical values for both derived limiting distribution are calculated via

a direct simulation and tabulated by the authors of these tests in the paper where the tests were introduced (Kwiatkowski *et al.*, 1992). In addition, KPSS say that the above tests are intended to complement unit root tests, such as the Dickey-Fuller tests and can be extended to allow for nonlinear trends.

Canova and Hansen (CH) Test

In the same sense that HEGY generalized the DF framework from the zero frequency to the seasonal frequencies (Canova & Hansen, 1995) generalize the KPSS framework from the zero frequency to the seasonal frequencies. They describe a set of tests to examine the structural stability of seasonal patterns over time. The tests are built on the null hypothesis of unchanged seasonality and can be tailored to test for unit roots at seasonal frequencies or for time variation in seasonal dummy variables. Thus, they propose a test for whether the seasonal pattern changes sufficiently over time to warrant a seasonal unit root, or whether a stable seasonal pattern is more appropriate. A discussion of the way they propose to test for non-constant seasonal patterns follows.

They start from a linear time series model with stationary seasonality, a seasonal dummy model:

$$X_i = y_i' \beta + d_i' \alpha + \epsilon_i, \quad i = 1, 2, \dots, T. \quad (3.71)$$

In (3.71), X_i is real valued, y_i is a $k \times 1$ vector of explanatory variables, d_i is an $s \times 1$ vector of seasonal dummy indicators and α is an $s \times 1$ parameter vector, where s is the period of the seasonal component. In addition, $\epsilon_i \sim (0, \sigma^2)$ is an error uncorrelated with y_i and d_i . Notice that there is no intercept μ included in the above model to achieve identification. The advantage of this formulation is that the coefficients α represent seasonal effects. To study whether the seasonal intercepts, α , have changed over time, Canova and Hansen modify the above conventional seasonal dummy model by entering α_i in place of α .

There are many forms of potential nonstationarity for α_i that could be considered; in their paper they consider stochastic variation of a martingale form:

$$A' \alpha_i = A' \alpha_{i-1} + u_i, \quad (3.72)$$

where α_0 is fixed and $\{u_i, F_i\}$ is a martingale difference sequence with covariance matrix $E(u_i u_i') = \tau^2 G$, $G = (A' \hat{\Omega} A)^{-1}$. The $s \times a$ matrix A selects the elements of α that we allow to stochastically vary under the alternative hypothesis. Note that when $\tau = 0$ the coefficient vector is fixed at α_0 for the entire sample.

The LM test for $H_0 : \tau = 0$ against $H_1 : \tau \neq 0$ is given by the statistic:

$$L = \frac{1}{T^2} \sum_{t=1}^T \hat{D}_t' A (A' \hat{\Omega} A)^{-1} A' \hat{D}_t = \frac{1}{T^2} \text{tr} \left[(A' \hat{\Omega} A)^{-1} A' \sum_{t=1}^T \hat{D}_t \hat{D}_t' A \right], \quad (3.73)$$

where $\hat{D}_t = \sum_{i=1}^t d_i \hat{\epsilon}_i$, $\text{tr}(Q)$ is the trace of Q and $\hat{\Omega}$, the consistent kernel estimate of the long-run covariance matrix of $d_i \epsilon_i$, Ω , is defined as:

$$\hat{\Omega} = \sum_{k=-m}^m w\left(\frac{k}{m}\right) \frac{1}{n} \sum_i d_{i+k} \hat{\epsilon}_{i+k} d_i' \hat{\epsilon}_i. \quad (3.74)$$

In the above equation m is the bandwidth and $w(\cdot)$ is any kernel function that produces positive semidefinite covariance matrix estimates, such as the Bartlett, Parzen, or the quadratic spectral.

Testing the stability of the a th seasonal intercept, (where $1 \leq a \leq s$), can be achieved by choosing A to be the unit vector with a 1 in the a th element and zeros elsewhere. This produces the test statistic:

$$L_a = \frac{1}{\hat{\Omega}_{aa} T^2} \sum_{t=1}^T \hat{D}_{at}^2, \quad (3.75)$$

where \hat{D}_{at} is the a th element of \hat{D}_t , and $\hat{\Omega}_{aa}$ is the a th diagonal element of $\hat{\Omega}$. Under H_0 the limiting distribution of this test statistic is given by:

$$L_a \rightarrow_d CvM(1) \quad \text{for each } a = 1, \dots, s, \quad (3.76)$$

where $CvM(p) = \int_0^1 W_p(r)' W_p(r) dr$, ' \rightarrow_d ' denotes convergence in distribution and W_p denotes a vector standard Brownian bridge of dimension p . When $p = 1$, the distribution of $CvM(p)$ simplifies to that known as the Cramér-von Mises goodness-of-fit distribution widely used in the statistical literature, so we will refer to $CvM(p)$ as the generalized Cramér-von Mises distribution with p degrees of freedom.

Here, the statistics L_a are essentially the KPSS statistic applied to the seasonal subseries (only the observations from the a th seasonal are used). Thus, the KPSS test is for instability in the average level of the series, but the L_a tests are for instability in the seasonal subseries.

One straightforward test statistic for testing instability in all the seasonal intercepts can be obtained by taking $A = I$, yielding:

$$L_J = \frac{1}{T^2} \sum_{t=1}^T \hat{D}_t' \hat{\Omega}^{-1} \hat{D}_t. \quad (3.77)$$

Standard analysis shows that under H_0 , $L_J \rightarrow_d VM(s)$. Note that L_J is a test for instability in any of the seasonal intercepts, so that it will have power against zero-frequency movements in X_t . In other words, L_J is a joint test for instability at the zero frequency as well as at the seasonal frequencies. To cope with this problem one could test for variation in the joint seasonal intercept process that keeps the overall mean constant. Specifically, decompose the seasonal intercepts a into an overall mean and deviations from the mean. Then test the joint stability of those deviations from the mean.

Critical values for the pre-described test statistics, the generalized Von Mises distribution, are given in the same paper where they propose their tests Canova & Hansen (1995).

One shortcoming of the above framework is that the test for a unit root at a particular frequency is seriously affected by the existence of other “unattended” unit roots which may exist at the same or other frequencies. In particular, the stationarity tests distort below nominal size under the null and display an associated loss of power under the alternative as argued by Hylleberg (1995). Furthermore, even though the above test takes care of both stationary autocorrelation and heteroscedasticity, is not well suited to handle non-stationary dynamics in the residuals. Finally, it can be argued that, in the same way CH advise against second differencing, because it can absorb at least one of the annual unit roots, it is possible to lose semi-annual unit roots by the first differencing they are proposing (Hylleberg, 1995).

Taylor

As an answer to the problem raised by the unattended unit roots by Hylleberg (1995), Taylor (2003) suggests that before testing for a unit root at a particular frequency one should first transform the data by applying a differencing filter (prefilter), that reduces the order of integration at each one of the remaining (unattended) frequencies by one.

Taylor proposes the following prefilters:

$$F_0 = (1 - B^d)/(1 - B) \text{ for testing at the zero frequency,} \quad (3.78)$$

$$F_\pi = (1 - B^d)/(1 + B) \text{ for testing at the Nyquist, } \pi, \text{ frequency,} \quad (3.79)$$

where d is even, and

$$F_k(B) = (1 - B^d)/(1 - 2\cos(2k\pi/d)B + B^2) \text{ for testing at the harmonic frequencies} \quad (3.80)$$

where $k = 1, \dots, d^*$ and $d^* \equiv (d/2) - 1$, if d is even or $[d/2]$, if d is odd.

An advantage of the above technique is that it does not alter the limiting distribution theory for the tests vis-à-vis the case in which there are no unattended unit roots (Taylor, 2003).

(Un)Attended Structural Breaks - Busetti and Harvey, Busetti and Taylor

Busetti & Harvey (2003) argue that the KPSS and CH stationarity tests are likely to be oversized if there are structural breaks in the seasonal pattern and propose modifications to the test statistics to overcome that effect. However, these modifications change the limiting null distribution of the statistics from $CvM(s_1)$ to $CvM(2s_1)$, where s_1 is the number of seasonal intercepts included in the test. In contrast to the seasonal breaks, breaks in the trend leave the asymptotic distribution of the statistics unaffected as long as they are correctly modelled.

In the context of a regression of a time-series variable on a set of zero and seasonal frequency spectral indicator variables the null hypothesis, of interest, is that of fixed parameters against the alternative that (at least one of) the parameters on a given subset, say J_1 , of the spectral indicators evolve as random walks, revealing thus that the process has unit root(s) at (at least one of) the spectral frequencies included in J_1 . If some or all of the parameters on the spectral frequency regressors included in J_1 display a structural break these breaks are termed *attended* while otherwise *unattended* (Busetti & Taylor, 2003).

Busetti and Taylor, referring to the KPSS and CH tests, state: *‘From a practical perspective, the impact of unattended breaks and unattended unit roots on the stationarity tests of Section 2 are just as important as those arising from attended breaks.’*

Busetti & Taylor (2003) put forward two ways of dealing with processes that are stochastically stationary about a deterministic component subject to structural breaks

and thus, can display properties very similar to a unit root process. The first, for structural breaks is bias correcting the original stability tests while the second, effective against both structural breaks and unattended unit roots is achieved by running stability tests on pre-filtered data. Their work extends the suggestions of Busetti & Harvey (2003) on dealing with structural breaks.

Suppose that there is a structural break at an unknown point in the sample, $a_0 \in (0, 1)$. They propose to replace in the KPSS and CH test statistics the biased estimator -in the presence unattended breaks- of the error variance by an asymptotically unbiased one, which is obtained by minimizing the sum of squared residuals over all the possible break dates, Busetti and Harvey (2003), Busetti and Taylor (2003). Specifically, under the null, $a^* = \arg \inf_a \hat{\sigma}^2(a)$ is a superconsistent³ estimator of a_0 and thus, $\hat{\sigma}^2(a) = T^{-1} \sum_{t=1}^T \epsilon_t(a)^2$ is the estimator they are proposing for the denominator of the statistics; where $\epsilon_t(a)$ denotes the OLS residuals from the fitted regression.

To deal with both structural breaks and unattended unit roots they use the differencing filter $\nabla_2 = (1 + B)(1 + B^2)$ to reduce by one the order of integration at all of the spectral frequencies where a structural breaks is suspected.

Both suggested methodologies recover (under mixing conditions, see Busetti & Taylor (2003)) the usual limiting null distribution of the Cramér-von Mises family, appropriate to the case where there are no breaks or unattended unit roots.

3.2.4 Discussion

It is evident that the question of which testing procedure one should use is not readily answered because the reason that one decides to test for unit roots in the first place is crucial. However, the DF, ADF and PP tests should be avoided since they lack power against meaningful alternatives. In addition, as for the variance ratio test, in small samples, using the asymptotic distributions results in substantial

³It converges to the true value at rate T .

size distortions and it is advised to obtain small sample critical values using Monte Carlo methods.

Generally, one of the issues arising when conducting unit root tests is the specification of the deterministic trend in a model, as mentioned at the beginning of this chapter, as it is closely related to the power and size of the tests. It is important to include as many deterministic regressors as there are deterministic components in the trend function of the DGP. Otherwise the test will at best lose finite sample power or at worst have power that goes to zero as the sample size increases. Nevertheless, one does not want to include more deterministic regressors than necessary as that will decrease the power of the test, Maddala & Kim (1998). Some authors propose sequential testing to overcome the above problem, but that then leads to new issues as the asymptotic distributions for general trends have yet to be derived. Finally for the majority of the tests mentioned above a trade off arises from the fact that if the number of observations, n is too small the tests are biased when there is autocorrelation while for a large value of n the tests lose power.

A large number of tests have been devised as modifications to the older ones, trying to deal with some of the arising issues. Chambers & McGarry (2002), state that there are different ways to deal with nuisance parameter dependencies. Firstly, one can augment the regression with lagged differences in an attempt to remove the serial correlation from the disturbance, i.e. ADF test, or, extended to seasonal unit root tests, i.e. HEGY, include a sufficient number of lagged (seasonal) differences as additional regressors in order to whiten the residuals. However, the outcome of such tests will be affected by the number of lagged differences employed.

Secondly, one can make non-parametric adjustments to the test statistics, thus achieving the desired limiting distributions. Lastly, by carrying out the regression in the frequency domain instead of the time domain, one can use non-parametric spectral density estimators to deal with the heteroscedasticity attributed to serial correlation. Furthermore, by employing the bootstrap method or Monte Carlo one can be sure of the limiting distributions not being affected by violated assumptions.

These four ways are the main ways by which the mentioned tests were modified to accommodate more general conditions. For more information about unit root and seasonal unit root tests the reader is referred to Maddala and Kim (1998).

3.2.5 Applications

Three of the aforementioned tests are applied to the House, Unemployment and Kola data sets. We start by applying the Phillips-Perron test which tests the null hypothesis of a unit root at 0 frequency against the alternative hypothesis of stationarity.

Table 3.1: Phillips-Perron Unit Root Test

Series	p-value
House Data	0.0140
Unemployment Data	0.826
Kola Data	0.01

The results are presented in Table 3.1. The test concludes that only Unemployment has a stochastic trend; unit root at 0 frequency.

The HEGY test tests the null hypothesis of a unit root at all frequencies against the alternative of stationarity. No deterministic components were included in the models. In Table 3.2 the results from this test are displayed. All series are found to have a seasonal unit root and no root at the long-run (0) frequency.

The Canova-Hansen test tests for instability in the seasonal pattern. The null hypothesis is of a stable seasonal pattern against the alternative of a unit root. The following tests are for stationary cycles at all seasonal frequencies without including a first order lag and a linear trend. The Canova-Hansen test accepts the null hypothesis of a stable seasonal pattern, i.e. no unit root, see Tables 3.3 and 3.4.

Table 3.2: HEGY Unit Root Test p-values

Freq.	House	Unemployment	Kola
0	0.100	0.100	0.100
$\pi/6$	0.100	0.010	0.100
$\pi/3$	0.100	0.100	0.010
$\pi/2$	0.010	0.100	0.010
$2\pi/3$	0.048	0.100	0.010
$5\pi/6$	0.010	0.010	0.010
π	0.010	0.100	0.010

Table 3.3: Canova-Hansen Unit Root Test p-values

Series	L-statistic
House Data	1.117
Unemployment Data	1.088
Kola Data	0.656

It is seen that the interpretation of the test results is difficult as a number of decisions have to be made before the tests are performed. The power and the size of these tests is closely related to the model chosen to represent the data. One has to decide on the appropriate order of integration in the model to be tested. The order of integration is important as it can affect the outcome of the tests directly, through the creation of the test statistic, and indirectly, through the appointed critical values. Furthermore, the deterministic parts of the model have to be chosen correctly.

The specification of the deterministic trend in the model, as mentioned before, is crucial. One should include as many deterministic regressors as there are deterministic components in the trend function of the DGP. Otherwise the power of the test may converge to zero as the sample size increases. On the other hand, including more deterministic regressors than necessary can still have a detrimental effect to

Table 3.4: Canova-Hansen Unit Root Test Critical Values

L-statistic	p-value
2.49	0.10
2.75	0.05
2.99	0.025
3.27	0.01

the power of the test. The critical values used for inference differ according to the various forms of the specified deterministic trend. Through the above three examples we see that it is hard to decide on the above features, especially when the purpose of conducting these tests is to explore the nature of the series examined as is in the context of this thesis.

In addition, these tests come with limitations as a number of different assumptions need to be made before one can use them effectively. Assumptions of independence and Normality are often required while heteroscedasticity and autocorrelation can also hinder the correct interpretation of the test results.

Due to this fact, the unit root testing procedures are not as straightforward as conventional testing procedures (Maddala & Kim, 1998). For the purpose of our analyses we require methods that do not impose many assumptions on the series and whose interpretation is fairly straightforward. Furthermore, as discussed in the previous chapter we are concerned with the different types of seasonality focusing mainly on systematic changes in either phase or amplitude. Thus, the above sophisticated testing techniques emphasize on many aspects that are not necessary in our studies while they fail to answer our main concerns regarding the type of seasonality. Thus, there is a need for a simple and robust way to test for the various types seasonality.

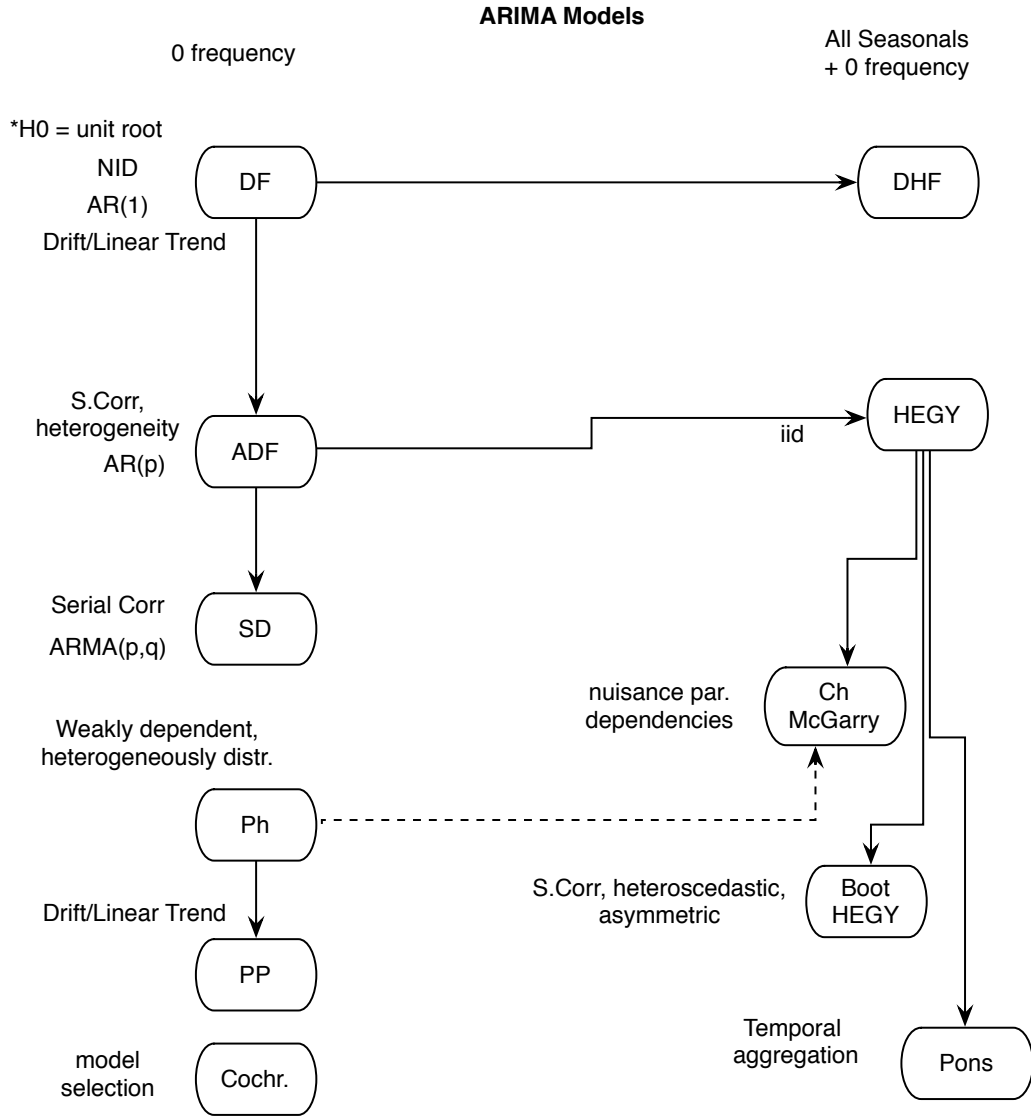


Figure 3.7: A diagram summarizing the Unit Root Tests that assume an ARIMA structure.

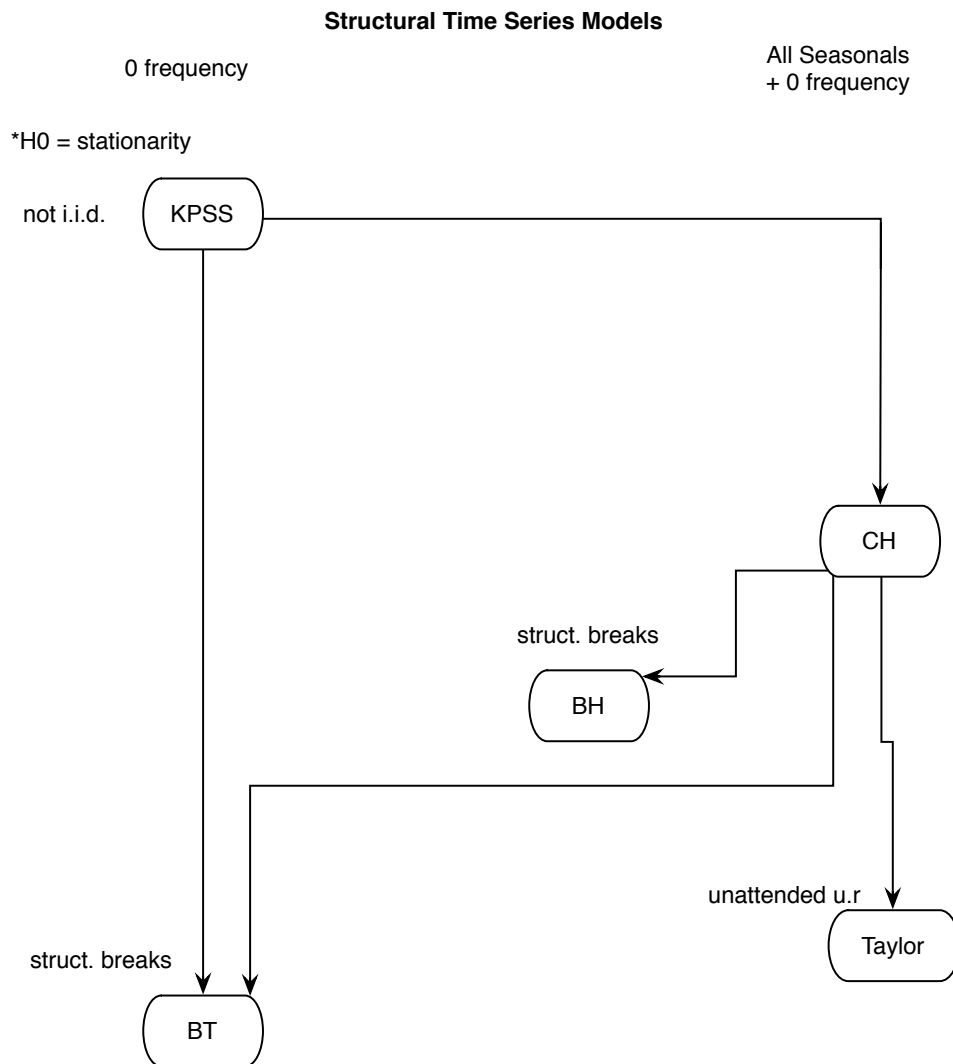


Figure 3.8: A diagram summarizing the Unit Root Tests that use structural time series models.

Chapter 4

Further Tests on Seasonality

4.1 Introduction

As we have seen from the last chapter conducting seasonality tests is not always straightforward and never without difficult inference. Furthermore, these elaborate and sophisticated testing frameworks focus on aspects that, even though interesting, may fall outside the main scope of our study. As an alternative to the above mentioned techniques simple, robust and fairly automated procedures are presented in this chapter. The first section introduces Resampling Tests based on Classical Seasonal Decomposition (CSD) while the second looks at Generalized Additive Models (GAMs) and how they can be utilized to test for types of seasonality.

4.2 Resampling Tests for Seasonality

The first procedure presented is used to test the absence of a seasonal structure against the alternative of a stable seasonal cycle, the second the presence of a stable seasonal cycle against the alternative of a trending one and the third the absence of

a seasonal structure against the presence of a trending one.

4.2.1 Non Seasonal vs Seasonal

This testing procedure is beneficial when one needs to test the existence of a stable seasonal pattern in the examined time series. The starting hypothesis is that of no seasonal pattern ($H_0 : \textit{Non Seasonal}$) with the alternative hypothesis of a stable one ($H_A : \textit{Stable Seasonal}$). Thus, one should be cautious when applying this method to a series that is suspected of having a trending seasonal pattern as this test may fail to indicate it. A series with a highly trending seasonal cycle that is buried in noise may in fact be deemed “non seasonal” using this testing algorithm. For those series, a different test, one between no seasonal pattern and a trending one, might be more appropriate and follows.

The procedure starts by employing an additive approach, assumes no seasonal pattern exists, and thus, a trend component, T_t , is extracted from the data, $\{X_t\}$, using a (centred, if necessary) moving average of order s , $\text{MA}(s)$, where s is the hypothetical periodicity:

$$X_t = T_t + R_t \quad (4.1)$$

Stable seasonal indices are evaluated by averaging R_t across the years for each season and then standardizing them. Thus:

$$R_t = S_t + R_t^*$$

The sum of the squared residuals, R_t^* is saved, $\text{SSR} = \sum_t R_t^{*2}$.

The residuals, $\{R_t\}$, are randomized to create the series $\{R_{t1}\}$. Then seasonal indices, S_t , are evaluated from R_t . The sum of the squared residuals, (SSR_1) is then noted and the procedure of randomizing the errors from the first decomposition, $\{R_t\}$, is repeated nine hundred and ninety eight times, (998), to generate the sums of squared

residuals, SSR_1, \dots, SSR_{999} . Thus, for $k=1, \dots, 999$ we have:

$$R_{tk} = S_{tk} + R_{tk}^* \quad (4.2)$$

and $SSR_k = \sum_t R_{tk}^{*2}$.

Inference is based on the fact that if the series really has a stable seasonal pattern, then SSR will be significantly smaller than the sum of squared residuals evaluated from the randomized residuals, R_{tk} , i.e. SSR_k . Hence, the lower tail of the distribution is used and a p-value is computed as follows¹:

$$\text{p-value} = \frac{\sum_{i=1}^{999} I(SSR > SSR_i)}{1000}, \quad (4.3)$$

where

$$I(\text{TRUE}) = 1 \quad \text{and} \quad I(\text{FALSE}) = 0 \quad (4.4)$$

For a given significance level, too small a p-value indicates a stable seasonal pattern while, an acceptable p-value is evidence of non-seasonality. This algorithm offers a simple test but without the need for any assumptions about distribution or the serial (in)dependence of the data making it robust and useful in practical time series analysis.

One should note, however, that using MAs means that if we test a non seasonal series we may find it seasonal as the MAs can induce seasonality to a non seasonal series during trend extraction. It may, thus, be better to use a different smoothing technique to estimate the trend. In a later chapter we will modify this algorithm to use Friedman's super-smoother which is described later.

¹A further test on whether the seasonality is stable or not, once its existence has been verified, will be conducted.

4.2.2 Stable Seasonal vs Trending Seasonal

Using the previous algorithm one can verify the existence of seasonality, but only in the form of a stable seasonal pattern. When seasonality is clearly present one wants to test whether that pattern is truly stable and not trending. In this case, the starting hypothesis is that of a stable seasonal pattern ($H_0 : \textit{Stable Seasonality}$) with the alternative hypothesis of a trending one ($H_A : \textit{Trending Seasonality}$).

It becomes crucial to distinguish between the trend in the series and any trend in the seasonal pattern, as over-smoothing the series will underestimate the existing trend, while if one allows the trend to pick up too much of the variation in the series it could actually extract some of the seasonality in it, and even pick up noise. The latter may not seem important if the purpose of our analysis is to forecast, since the variation may still be modelled correctly (though not by the seasonal component) and thus, the forecasts could be successful. However, since the purpose of this study is to understand and describe the behaviour of the series and the underlying dynamics in it, it is very important to correctly recognize and distinguish between the long term trend and any trend in the seasonals.

Using a stable CSD procedure, a stable seasonal pattern and a trend component are extracted from the data. This stable CSD is performed using a MA(s) where s is the periodicity of the series, X_t , to extract a first trend, T_t^1 , and produce the first de-trended series, $Detr_t^1$:

$$Detr_t^1 = X_t - T_t^1. \quad (4.5)$$

Then the first seasonals, S_k^1 , $k = 1, \dots, s$, are estimated in the usual way. These are then repeated to create the series of first seasonals, S_t^1 , of equal length to the starting time series, X_t . This in turn is subtracted from the data to produce a deseasonalized series, Des_t^1 :

$$Des_t^1 = X_t - S_t^1. \quad (4.6)$$

A MA(s) is again employed to smooth the de-seasonalized series, Des_t^1 , and yield a

second trend, T_t^2 , and thus, a second de-trended series, $Detr_t^2$:

$$Detr_t^2 = X_t - T_t^2. \quad (4.7)$$

We now generate a second set of seasonals, S_k^2 , $k = 1, \dots, s$ which form a second/final time series of seasonals, S_t^2 . The overall mean of the second seasonals is added to the second trend to form the final trend, T_t^3 :

$$X_t = T_t^3 + S_t^2 + R_t. \quad (4.8)$$

The values in each season of the residuals, $\{R_t\}$, from this decomposition are separately smoothed across the years to evaluate trending seasonals, S_t^* . Then we remove these trending seasonals from the series to generate residuals, R_t^* ,

$$R_t = S_t^* + R_t^*. \quad (4.9)$$

The smoother in this case is Friedman's super-smoother, a variable span scatter-plot smoother. Friedman's super-smoother is a running lines smoother that chooses between three different spans for the lines. The running lines smoothers are symmetric, using k data points for each predicted value, $k/2$ data points on each side, with values of k as $0.5 * n$, $0.2 * n$ and $0.05 * n$, where n is the number of data points in the time series. Then, by cross-validation for each prediction, the best of the three fitted smoothers is chosen. Finally, the best spans are smoothed by a running lines smoother and the final prediction is chosen by linear interpolation (Friedman, 1984).

The above smoother was selected as the most appropriate from others available in literature because it uses cross-validation to choose from different spans, making it automatic, more robust and more reliable to extract the present trending seasonals whilst not over-smoothing the data.

The sum of the squared residuals is saved, SSR. The residuals, R_t , are now randomized to create the series $\{R_{t1}\}$. Then the super-smoother is again used on the $\{R_{t1}\}$ series for each season across the years to extract a trending seasonal pattern, S_{t1}^* , and generate residuals, R_{t1}^* :

$$R_{t1} = S_{t1}^* + R_{t1}^*.$$

The sum of the squared residuals (SSR_1) from this decomposition is then noted and the procedure of randomizing the $\{R_t\}$ series, from which the trending seasonals are extracted, is repeated nine hundred and ninety eight times to generate the sums of squared residuals, SSR_1, \dots, SSR_{999} , which are saved, as before.

If the series does not truly have trending seasonality then applying this procedure to the randomized data should be essentially the same as applying it to the original data, whereas if the series does exhibit trending seasonality then the application to the original data will yield a significantly better fit as measured by the sum of squared residuals.

Thus, the lower tail of the distribution is used for inference and the p-value is as in (4.3). Too small a p-value is taken to indicate trending seasonality at the appropriate level, whereas an acceptable p-value indicates stable seasonality.

As mentioned in the previous chapter, there may be different types of changing seasonals. A sensitivity analysis of the above procedure will follow later. Finally we recall that the definition of a trending seasonal pattern is that the seasonals are systematically shifted in time, while a seasonal pattern that is stable or fluctuating in time is considered a stable seasonal pattern, i.e. regardless of other changes.

4.2.3 Non Seasonal vs Trending Seasonal

In this section we discuss a test between the starting hypothesis of no seasonal pattern ($H_0 : \textit{Non Seasonal}$) and the alternative hypothesis of a trending seasonal pattern ($H_A : \textit{Trending Seasonality}$). This testing procedure is similar to the two already

mentioned with the obvious difference that the non seasonal decomposition, (only a trend is extracted) employed in the first test is used on the data to generate the residuals that will now be used in exactly the same way as in the second test.

Under the null hypothesis of no seasonality present, a trend, T_t , is extracted, using a MA(s), from the data, $\{X_t\}$, where s is again the putative periodicity, as in (4.1). The residuals, R_t , from this non seasonal decomposition are then smoothed, each season separately, as in (4.9). Repeated randomisation of the residuals, as before, allows us to compute appropriate residual sums of squares and perform our test.

When the series does not really have trending seasonality applying this method to the randomized data is the same as applying it to the original data. If the series, however, does exhibit trends in the seasonals then the fit to the original data will be significantly better than the randomized ones. Therefore, inference uses the lower tail of the distribution with p-value from equation (4.3).

This testing procedure should be valuable if we wish to examine the behaviour of a time series buried in noise, i.e. the signal of the trending seasonal pattern is too small or too trending to be detected as a stable one using the first testing procedure.

4.2.4 Applications

The above Resampling tests are used to draw conclusions on the seasonality of the three example time series. The p-values from the tests for the House, Unemployment and Kola time series are presented in Table 4.1.

Table 4.1: Resampling Results

	NS vs SS	SS vs TS	(NS vs TS)
House Data Set	0	0.317	0.134
Unemployment Data Set	0	0	0.027
Kola Data Set	0	0.994	0.738

The House time series clearly has stable seasonality with a zero p-value for the

Non Seasonal vs. Seasonal test and a non significant p-value of 0.317 for the Stable Seasonal vs. Trending Seasonal test. Similarly, the Unemployment time series clearly exhibits trending seasonality with both p-values equal to zero for the same tests. Finally, the results for the Kola data indicate a stable seasonal pattern with zero and 0.994 p-values. Note that the latter p-value is close to 1 as the evaluated stable seasonal model gives a very good fit, meaning that the SSR of the first CSD has a low value so that the subsequent decompositions do not generate smaller SSRs. In practice, since they all tested seasonal, the Non Seasonal vs Trending Seasonal test would not be applied; nevertheless, the Unemployment series is identified as having a trending seasonal even against the alternative of no seasonality.

4.2.5 Discussion

The above tests do not require any assumptions of a specific distribution or of serial (in)dependence of the time series examined, making them very robust. Consequently they are easy to implement and have a more reliable interpretation than other more elaborate techniques that focus on many different aspects of seasonality. The main purpose of this thesis is to analyze and examine biological time series that are serially correlated and exhibit different types of seasonality that need to be identified and explored, so the above tests provide a good basis. Along with the GAM testing procedures that follow, these are the main tools used in the data analysis of these biological time series in the following chapters.

4.3 Introduction to Generalized Additive Models

Generalized Additive Models (GAMs) are often used in literature to build models of biological data (Venables & Ripley, 2002) (Guisan *et al.*, 2002) (Barry & Welsh, 2002). In this thesis we use GAMs to test for different types of seasonality. We start

with a brief explanation of GAM theory and then we construct three different types of model, under assumptions about seasonality. Examples are used to illustrate the framework described. The description that follows is based on Wood (2006).

We start with a linear model:

$$X_i = a + by_i + \epsilon_i, \quad i = 1, 2, \dots, n \quad (4.10)$$

where X_1, \dots, X_n are independent r.v.s following the Normal distribution with mean $\mu_i = a + by_i$ ($i = 1, 2, \dots, n$) and variance σ^2 . The y_1, \dots, y_n are the explanatory variables of the model and $\epsilon_1, \dots, \epsilon_n$ are the errors. The errors are assumed to be iid with $\epsilon_i \sim N(0, \sigma^2)$. Parameters a and b are unknown and need to be estimated.

The degrees of freedom (d.f.) for the above linear model are $(n - 2)$ since it involves two parameters, a and b . In general, if a model with p parameters is being fitted to n data points then its residual d.f. are $(n - p)$.

Such a more general model is more easily written in matrix form, using that $X_i \sim N(\mu_i, \sigma^2)$, as:

$$E(\mathbf{X}) = \boldsymbol{\mu} = \mathbf{y}\boldsymbol{\beta} \quad (4.11)$$

where $E(\cdot)$ denotes expectation, $\mathbf{X} = (X_1, \dots, X_n)^T$, is the response variable, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)^T$ is the mean vector, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$, a vector of unknown parameters and

$$\mathbf{y} = \begin{pmatrix} 1 & y_{11} & y_{12} & \cdots & y_{1p-1} \\ 1 & y_{21} & y_{22} & \cdots & y_{2p-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & y_{n1} & y_{n2} & \cdots & y_{np-1} \end{pmatrix}$$

is a matrix of the explanatory variables and is known as the model matrix. The linear component of the model is given by $\mathbf{y}\boldsymbol{\beta}$.

Two approaches are commonly used for estimation; OLS and ML estimation. In

OLS we minimise, with respect to (w.r.t.) $\boldsymbol{\beta}$, the mean squared error (MSE):

$$\frac{1}{n}(\mathbf{x} - \mathbf{y}\boldsymbol{\beta})(\mathbf{x} - \mathbf{y}\boldsymbol{\beta})^T \quad (4.12)$$

where $\mathbf{x} = (x_1, \dots, x_n)^T$ is a realisation of X . Hence, the estimated parameter values minimise the errors in our model in this sense.

The ML estimation is based on maximising the likelihood function of \mathbf{X} w.r.t. $\boldsymbol{\beta}$. Thus, we find the parameter estimates that maximise the likelihood of the data given the chosen model. Let X_i have a probability (density or mass) function $f(x_i)$ then the Likelihood (L) is defined as:

$$L(\boldsymbol{\beta}) = \prod_{i=1}^n f(x_i). \quad (4.13)$$

where $L(\boldsymbol{\beta})$ implies that for given values of x_i the likelihood function can be seen as a function of the unknown parameters.

Maximisation of the log-likelihood, denoted by l , yields the same results and is often computationally more efficient,

$$l(\boldsymbol{\beta}) = \log \left(\prod_{i=1}^n f(x_i) \right) = \sum_{i=1}^n \log(f(x_i)) = \sum_{i=1}^n l_i \quad (4.14)$$

The MLE and the OLS estimator of $\boldsymbol{\beta}$ are the same for the linear model when the errors are normal and it is given by:

$$\hat{\boldsymbol{\beta}} = (\mathbf{y}^T \mathbf{y})^{-1} \mathbf{y}^T \mathbf{x} \quad (4.15)$$

provided that $(\mathbf{y}^T \mathbf{y})$ is invertible. The matrix $\mathbf{H} = \mathbf{y}(\mathbf{y}^T \mathbf{y})^{-1} \mathbf{y}^T$ is called the ‘hat matrix’ and the elements of its main diagonal, H_{ii} will appear again later.

Generalized Linear Models (GLMs) allow for a response distribution other than Normal and for a degree of non-linearity in the model structure. The independent r.v.s X_i now follow an exponential family² distribution with p.d.f.:

$$f(x) = \exp \left\{ \frac{x\theta - b(\theta)}{\alpha(\tau)} + c(x, \tau) \right\} \quad (4.16)$$

where $\alpha(\cdot)$, $b(\cdot)$ and $c(\cdot)$ are arbitrary functions, θ is the canonical parameter and τ is the scale (dispersion) parameter. For more information on exponential family properties see Dobson (2002).

The general structure of a GLM is:

$$g(\boldsymbol{\mu}) = \mathbf{y}\boldsymbol{\beta} \quad (4.17)$$

where, since X_i follows an exponential family distribution, $E(\mathbf{X}) = \boldsymbol{\mu} = b'(\theta)$ and

$$\text{Var}(\mathbf{X}) = b''(\theta)\alpha(\tau). \quad (4.18)$$

Assuming $\alpha(\tau)$ can be written as $\alpha(\tau) = \tau/\omega$, where ω is a known constant, we can write the variance of \mathbf{X} in relation to $\boldsymbol{\mu}$ as:

$$\text{Var}(\mathbf{X}) = V(\boldsymbol{\mu})\tau \quad (4.19)$$

where $V(\cdot)$ a function such that $V(\boldsymbol{\mu}) = b''(\theta)/\omega$. The relationship between the linear predictor (4.17) and the mean $\boldsymbol{\mu}$ is given by g , a monotonic, differentiable ‘link function’, while $\boldsymbol{\beta}$ is a vector of unknown parameters (Dobson, 2002). Due to this generalization the model fitting has to be iterative and large sample limiting results are needed for inference. The estimation and inference are based on ML estimation but now this requires an iterative LS approach.

²Examples include the Poisson, the Binomial, the Normal and the Gamma distributions.

The log-likelihood function for each of the observed x_i 's is:

$$l_i = \{x_i\theta - b(\theta)\} / \alpha(\tau) + c(x_i, \tau) \quad (4.20)$$

and for all the observed x_i 's :

$$l = \sum_{i=1}^n l_i = \sum_{i=1}^n x_i\theta / \alpha(\tau_i) - \sum_{i=1}^n b(\theta_i) / \alpha(\tau_i) + \sum_{i=1}^n c(x_i, \tau_i) \quad (4.21)$$

The above model's terms are parametric and thus many plausible non-parametric (n.p.) functions are excluded from the model fit. Also we have to decide the appropriate form of the model while it would be easier, yet computationally more costly, to let the modelling approach choose the appropriate form of the model. GAMs provide a framework for n.p. components in GLMs.

“A generalized additive model is a generalized linear model with a linear predictor involving a sum of smooth functions of covariates.” state Hastie & Tibshirani (1990). So a general structure is:

$$g(\boldsymbol{\mu}) = \mathbf{y}\boldsymbol{\beta} + f_1(y_1) + f_2(y_2) + f_3(y_3) + \cdots + f_m(y_m) \quad (4.22)$$

where the f_j are smooth, twice differentiable n.p. functions of the covariates, y_k , the other parts are as defined earlier. Thus, the relationship between the mean of the response variable and the covariates is specified in terms of smooth functions and not only parametric terms. The range of potential fits to the data is much larger than the parametric approach (Faraway, 2006).

INFERENCE

To estimate f_1, f_2, \dots, f_m one needs to represent them in such a way that (4.22) becomes a linear model. This is achieved by choosing a basis for each one of them, defining the space of functions of which each of the f_1, f_2, \dots, f_m is an element. The basis selection amounts to choosing some basis functions that are treated as

completely known: if $b_{ji}(y)$ is a set of such basis functions for $f_j(y)$, $i = 1, 2, \dots, q_j$, then $f_j(y)$ can be represented as:

$$f_j(y) = \sum_{i=1}^{q_j} \beta_{ji} b_{ji}(y), \quad (4.23)$$

for some values of the unknown parameters, β_{ji} . The f_j are smooth n.p. functions of the covariates y_j .

So, for example, one could use a cubic spline basis. A cubic spline is a curve that consists of sections of cubic polynomials joined together at points, known as knots, in such a way that they are continuous in value, as well as the first and second derivatives. Generally, a spline is a polynomial in each interval between two knots and a piecewise polynomial overall. A disadvantage of a single polynomial would be that it does not perform well over the whole domain but only locally. Spline bases are not faced with the same problem, but need more than n parameters to be estimated for a fit to n data points. This computational cost can be addressed using regression splines. For regression splines the smoothing problem is solved for a small subset of the data and the implied basis is then used for model construction of the whole data set. For the above bases one has to choose the knot locations. The knots are often evenly spaced through the range of the data but specific knots can be chosen (Wood, 2003).

Choosing basis functions can restrict the model as other ones might be better suited. Using thin plate regression splines (t.p.r.s.) one does not have to choose basis functions and knot locations as these emerge from the minimization of the constructed objective function, see Wood (2003). A t.p.r.s. basis has the benefits of thin-plate splines but requires a smaller number of parameters to be estimated. Instead of choosing number of knots and knot locations in t.p.r.s. one has only to choose the basis dimension, later when we refer to the number of knots for a t.p.r.s. we mean the value of the basis dimension. Another benefit from using t.p.r.s. is that contrary to the above bases, t.p.r.s. can successfully represent smooths of more than one

variable.

Using OLS or ML to select f_j would result, almost certainly, in a very wiggly line that passes through most points of the data. Therefore, we need a way to control the smoothness of these functions (splines). One way would be to repeatedly change the basis dimension by altering the number of knots³ but that could lead to various issues with uneven knot spacing and/or dependence on the location of the knots. Furthermore, if a large number of basis functions is used to approximate the real underlying function, the model will probably overfit while if the basis dimension is kept small it will be too restricted to approximate the truth (Wood & Augustin, 2002). An alternative is to keep the dimension of the basis fixed at a sufficiently large value and add a ‘wiggleness’ penalty to the log-likelihood to be maximized or to the objective function of the OLS.

A common way to measure the wiggleness of a function is to use its second derivative. This gives the rate of change of the slope, known as curvature, of that function. For example, Silverman & Green (1993) state that for a given set of increasing points $\{x_i, y_i : i = 1, \dots, n\}$ from all the continuous functions on $[y_1, y_n]$, with continuous first derivatives, the smoothest function can be defined as the one that minimizes:

$$\int_{y_1}^{y_n} f''(y)^2 dy. \quad (4.24)$$

Assuming that f has a basis expansion $f(y) = \sum_i^q \beta_i b_i(y)$ we can write

$$\int_{y_1}^{y_n} f''(y)^2 dy = \boldsymbol{\beta}^T \mathbf{S} \boldsymbol{\beta} \quad (4.25)$$

³i.e. with backward selection we would start with many knots and then drop knots sequentially.

since $f''(y) = \boldsymbol{\beta}^T \mathbf{z}(y)$, where $z_i(y) = b_i''(y)$ and thus,

$$\int_{y_1}^{y_n} f''(y)^2 dy = \int_{y_1}^{y_n} \boldsymbol{\beta}^T \mathbf{z}(y) \mathbf{z}^T(y) \boldsymbol{\beta} dy = \boldsymbol{\beta}^T \mathbf{S} \boldsymbol{\beta}, \quad \mathbf{S} = \int_{y_1}^{y_n} \mathbf{z}(y) \mathbf{z}^T(y) dy \quad (4.26)$$

GAMs maximize a penalized log-likelihood which can be used as a simple generalisation of the penalized OLS:

$$(\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{x} - \boldsymbol{\mu}) + \sum_j \lambda_j J_j \quad (4.27)$$

where J_j is a function measuring the wiggleness of f_j and λ_j^4 is the smoothing parameter relating to f_j . Smoothing parameters are used to balance the minimization of the badness of the fit, measured by the first part, and the model wiggleness, as measured by the second part (Wood, 2003).

Thus, instead of maximizing the log-likelihood, $l(\boldsymbol{\beta})$ a penalized log-likelihood, $l_p(\boldsymbol{\beta})$ is maximized, using Penalized Iterative Re-weighted Least Squares (P-IRLS) :

$$l_p(\boldsymbol{\beta}) = l(\boldsymbol{\beta}) - \frac{1}{2} \sum_j \lambda_j \boldsymbol{\beta}^T \mathbf{S}_j \boldsymbol{\beta} \quad (4.28)$$

where $\boldsymbol{\beta}$ is a vector of the unknown parameters of the corresponding penalized GLM once the bases are chosen and the term $\boldsymbol{\beta}^T \mathbf{S}_j \boldsymbol{\beta}$ is used to measure the wiggleness of the j^{th} smooth function with \mathbf{S}_j a matrix of known coefficients⁵. The λ_j are smoothing parameters, controlling the trade-off between goodness of fit of the model and model smoothness.

For known λ_j the l_p can be maximized to find $\hat{\boldsymbol{\beta}}$ but since the λ_j need themselves to be estimated iterative steps for convergence are employed, see Wood (2006). The automatic methods of choosing smoothing parameters are explained later in this

⁴A spline whose smoothness is controlled by smoothing parameters is called a ‘smoothing spline’.

⁵The majority of wiggleness measures, e.g. $\int f_j''(x)^2 dx$, can be written in the form $\boldsymbol{\beta}^T \mathbf{S}_j \boldsymbol{\beta}$.

section.

In GAMs the d.f. cannot be defined as in the linear model case. To measure the flexibility of the model Wood defines effective d.f. as the trace⁶ of the hat matrix of the GAM, $\text{tr}(\mathbf{H})$. Thus, the residual d.f. of the GAM in (4.22) can be written as $(n - \text{tr}(\mathbf{H}))$, where the hat matrix, \mathbf{H} , is:

$$\mathbf{H} = \mathbf{y}(\mathbf{y}^T \mathbf{y} + \boldsymbol{\lambda} \mathbf{S})^{-1} \mathbf{y}^T, \quad (4.29)$$

where \mathbf{S} and $\boldsymbol{\lambda}$ as above.

When the link function is the identity function and the errors are normal, then analogous to the linear parametric model case the residual variance, σ^2 , can be estimated by the residual sum of squares divided by the residual d.f.:

$$\hat{\sigma}^2 = \left(\sum_{i=1}^n (x_i - \hat{\mu}_i)^2 \right) / (n - \text{tr}(\mathbf{H})) \quad (4.30)$$

where \mathbf{H} is the influence (hat) matrix such that $g(\hat{\boldsymbol{\mu}}) = \mathbf{H}\mathbf{x}$. The above estimator for the additive model is unbiased (see Wood (2006)). From equations (4.19) and (4.30), the scale parameter for the GAM can be estimated by the Pearson-type scale estimator,

$$\hat{\tau} = \left(\sum_{i=1}^n V(\hat{\mu}_i)^{-1} (x_i - \hat{\mu}_i)^2 \right) / (n - \text{tr}(\mathbf{H})). \quad (4.31)$$

In the package “mgcv”, created by Simon Wood, available in R, the modeller can choose the number of knots to be used for basis construction, but this is only setting an upper bound on the flexibility of a term. The smoothing parameters, λ_j , control the effective d.f. and this is part of model specification. Therefore, the actual fit retains some insensitivity to the chosen basis dimension as long as it is not set restrictively low. The only difference between two sufficiently large basis dimensions

⁶Trace is defined as the sum of the elements on the main diagonal of a square matrix.

is that a function space with the larger one will contain more functions with the estimated effective d.f. than the smaller one.

Un-Biased Risk Estimator

The smoothing parameters $\boldsymbol{\lambda}$, that control the level of smoothing, are needed for the estimation of the model coefficients and can be estimated in two ways. When the scale parameter is known, then attempting to minimize the expected MSE leads to minimizing an estimate of it given by the Un-Biased Risk Estimator (UBRE) which is also closely related to Mallows's C_p ⁷ :

$$V_u(\boldsymbol{\lambda}) = (\mathbf{x} - \mathbf{H}\mathbf{x})^T (\mathbf{x} - \mathbf{H}\mathbf{x}) / n - \sigma^2 + 2tr(\mathbf{H})\sigma^2/n \quad (4.32)$$

which depends on the smoothing parameters through \mathbf{H} . Nevertheless this is inappropriate when the scale parameter has to be estimated.

Cross Validation

When the scale parameter is unknown, minimizing the mean square prediction error (P) is preferred. P is the average squared error in predicting a new observation x using the fitted model and can be written as:

$$P = \sigma^2 + \text{MSE}, \text{ where } \text{MSE} = \frac{tr(\mathbf{H})}{n} \sigma^2 \quad (4.33)$$

by substituting (4.30) into (4.32). Since P is directly dependent on σ^2 , any criteria based on P will resist over-smoothing much more than criteria based on MSE alone, as the σ^2 estimate would then be inflated.

To estimate P , cross validation is used. Excluding one datum, x_i , from the model fitting process implies that the model fitted to the remaining data is independent of the omitted x_i . Thus, the squared error in predicting x_i can be estimated, and, by omitting all data in turn we get the ordinary cross validation (OCV) score estimate

⁷ $C_p = (\mathbf{x} - \mathbf{H}\mathbf{x})^T (\mathbf{x} - \mathbf{H}\mathbf{x}) / \sigma^2 - n + 2tr(\mathbf{H})$.

of P :

$$V_o = \frac{1}{n} \sum_{i=1}^n \left(x_i - \hat{\mu}_i^{|-i|} \right)^2 \quad (4.34)$$

where $\hat{\mu}_i^{|-i|}$ denotes the prediction of $E(x_i)$ obtained from the model fitted without observation x_i . While the right hand side (r.h.s.) of (4.34) requires n model fits, it can be shown, see Wood (2006), that V_o can be computed from one fit of the original model using the following formula:

$$V_o = \frac{1}{n} \sum_{i=1}^n \frac{(x_i - \hat{\mu}_i)^2}{(1 - H_{ii})^2} \quad (4.35)$$

Nevertheless, OCV is hard to compute for a GAM where there are several smoothing parameters. In practice, the weights $1 - H_{ii}$, are commonly replaced by the mean weight, $\text{tr}(\mathbf{I} - \mathbf{H})/n$. Then the corresponding OCV score of (4.34) can be re-written as:

$$V_g = \frac{n}{[n - \text{tr}(\mathbf{H})]^2} \sum_{i=1}^n (x_i - \hat{\mu}_i)^2, \quad (4.36)$$

which is known as the Generalized Cross Validation (GCV) score. The above equation can be modified to include a parameter γ , as in (4.37), which can be used to inflate the model d.f. in the GCV score and result in smoother fits (Kim & Gu, 2004). For the traditional GCV score, as in equation(4.36), this parameter should be set to 1 while values greater than 1 favour smoother fits.

$$V_g^* = \frac{n \sum_{i=1}^n (x_i - \hat{\mu}_i)^2}{[n - \gamma \text{tr}(\mathbf{H})]^2} \quad (4.37)$$

Simulation results by Kim & Gu (2004) suggest a value of 1.4 for γ . This value gives smooth fits with good performance.

The estimation of the smoothing parameters, $\boldsymbol{\lambda}$, is done using iterative minimization

of the GCV score or UBRE score, see Wood (2006). Once a basis is chosen for a GAM this is transformed in a linear model or a GLM with one or more associated penalties. Hence the GAM is fitted, once a basis and wiggleness penalty are chosen, by penalized likelihood maximization which is achieved using Penalized Iterative Re-weighted Least Squares (P-IRLS).

In the “mgcv” package, estimation of a GAM consists of two iterative loops. The outer loop optimizes a smoothness selection score minimising it, w.r.t. the parameters, using Newton’s method or other non-linear optimisation algorithm. Thus, the smoothing parameters are estimated. Once the outer loop converges, the inner loop uses P-IRLS to estimate the model given the smoothing parameters. This package allows us to choose the maximum number of iterations per loop and the optimisation algorithm adopted in the outer loop. We are using Newton’s method in several dimensions, backed up by steepest descent to iteratively adjust the smoothing parameters for each penalty.

4.3.1 Testing with GAMs

The function ‘gam’ from the “mgcv” package in R is used to build GAMs. We develop three different GAMs to allow us to test directly for our different forms of seasonality. The first model is non seasonal and includes only a smooth function of time to incorporate the long term trend in the series. The second one models stable seasonality using factors, and a trend using a smooth function over time, as before. Finally, our third model includes, in addition, trending seasonality, in the form of smooth functions of time per season used as a factor. Hence, the models are nested making formal comparison easier. Considering the nature and form of the time series examined an appropriate exponential family distribution is chosen.

Constructing hypotheses models:

Non Seasonal Model, Model-NS contains only a smooth function of time.

$$g(\mu_t) = f_1(y_{1t}) = s(t), \quad t : \text{time} \quad (4.38)$$

Stable Seasonal Model, Model-SS contains a smooth function of time and incorporates factors for the seasons.

$$g(\mu_t) = f_1(y_{1t}) + fac = s(t) + fac, \quad t : \text{time}, fac : 1, 2, \dots, s \quad (4.39)$$

Trending Seasonal Model, Model-TS contains a smooth function of time, factors for seasons and smooth functions of time per season as factor.

$$g(\mu_t) = f_1(y_{1t}) + fac + f_2(y_{2t}) + \dots + f_s(y_{st}) = s(t) + fac + s(t, by = \mathbf{ff}),$$

$$t : \text{time}, fac : 1, 2, \dots, s, \mathbf{ff}_{ij} = \begin{cases} 1, & j = i - [i/j] \\ 0, & \text{otherwise} \end{cases}, i = 1, \dots, n \text{ and } j = 1, \dots, s. \quad (4.40)$$

In the above equation \mathbf{ff} is an $n \times s$ matrix, where n is the length of the data and s the suspected periodicity. The (i, j) element in \mathbf{ff} is equal to 1 when i is the j th season of the current year.

Thin plate regression splines (TPRS) are used for the trend and the trending seasonals while seasonal factors are used for the stable seasonals.

We consider two ways to test between models. When the scale parameter, σ^2 , is unknown we use F-ratio tests to check whether the increase in the residual sum of squares from dropping the extra term(s) is significant. When the scale parameter is known we use the scaled deviances of the two models to construct a log-likelihood ratio statistic. The scaled deviance, D^* of a model is the deviance, D divided by the scale parameter, τ . In the following examples the scale parameter is unknown and thus, ANOVA is used between non-seasonal versus seasonal and stable seasonal

versus trending seasonal models.

We use the following criteria to evaluate and compare the model fits:

1) Akaike's Information Criterion (AIC):

$$-2l + 2p \tag{4.41}$$

2) (Schwartz) Bayesian Information Criterion

$$-2l + p \ln(n) \tag{4.42}$$

where l the log-likelihood, p the number of parameters and n the length of the data. Furthermore, the GCV score can be used for comparison between nested models with the lowest score preferable. Since we are comparing nested models we are using the GCV score as another criterion.

To avoid over-fitting and to favour smoother fits we reduce the model d.f. thus inflating the GCV score using $\gamma = 1.4$ instead of the default value 1 as a constant multiplier in the construction of the model in equation(4.37) (Kim & Gu, 2004). This is done because the nature of the models we are examining calls for smooth fitted terms for all GAMs fitted. The default size of basis dimension, $k = 10$, is used for all model terms.

Finally, one has to check the validity of the GAMs in question. We look at the residuals of these models creating diagnostic plots. For this purpose standardised deviance residuals are created, \hat{d}_{std} ⁸. To check the normality assumption we examine a QQ (quantile-quantile) plot of these residuals and their histogram. To check the constant variance assumption of the response variable we plot the residuals against the response while a plot of the residuals themselves would reveal any remaining

⁸ $\hat{d}_{std,i} = (\hat{d}_i - \bar{\hat{d}}) / (\hat{\sigma} \sqrt{(1 - H_{ii})})$, where H_{ii} are the elements on the leading diagonal of the hat matrix, H , and $\bar{\hat{d}}$ the average of the deviance residuals \hat{d}_i .

pattern that was not incorporated in the model.

4.3.2 Applications

We demonstrate the above described procedure using the House, Unemployment and Kola data sets from previous chapters. For these examples the Gaussian family with the identity link function are chosen. Criteria are reported along with plots to support the choice of seasonality for each series. For the first series diagnostic plots for all fitted models are discussed while later only the ones from the chosen model are presented.

House Data

The GCV score can be used for comparison between nested models. The smallest score, in this case the SS model, is preferred, which corresponds to the results of our Resampling test, Table 4.1.

GCV Score: Non-Seas: 39378.33, **Stable:** 5129.542 & Trending: 45951.786

Table 4.2: House Data ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	124.047	4358058				
SS	111.295	412375	12.752	3945682	83.508	< 2.2e-16
TS	97.065	311305.536	14.23	101069.75	2.215	0.012

The results from the ANOVA between the models are displayed in Table 4.2. The first row of the table refers to the NS model while the second one gives the results of the ANOVA between the NS and the SS model and the third one the results between the SS and TS GAMs. These indicate that the series is seasonal with a trending seasonal pattern, conflicting with GCV and Resampling.

Figure 4.1 shows the two most likely fits for the House data; the SS model fit is depicted in red while the TS in blue. The trend and the stable seasonals from the SS model are plotted. The stable seasonal pattern is very smooth, considering that




Figure 4.1: The upper left plot shows the House data with the trend from the stable GAM in red. The upper right plot shows the stable seasonals estimated from the stable GAM while in the lower plot the House data are plotted as dots with the fits from the stable (red) and trending (blue) GAMs.

twelve factors were estimated.

The smooth functions fitted by the TS model are plotted in Figure (4.2). Their confidence intervals are given by the dotted lines. For the first eleven months the confidence intervals (C.I) of the smooths include the zero value for all the x-axis. December, however, shows a significant deviation from zero, which agrees with our previous thoughts.

Table 4.3: House Data AIC & BIC:

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	8.953	21.705	35.935
AIC	1765.93	1480.198	1471.545
BIC	1791.739	1542.768	1575.138

The AIC criterion, Table 4.3 identifies, like the ANOVA and possibly for the same reasons, the TS as the best model. BIC is more conservative and points to a stable pattern.




Figure 4.2: The (twelve) trending seasonals from the TS model of the House data with confidence intervals (dotted lines), by row from top left.

It seems unlikely that the House data have a systematically trending seasonal pattern, since only the Decembers display apparent change in the sense of fluctuations in size. This change could be due to unusually cold Decembers or to more people going away on holidays. In other words, it may simply be a fluctuation rather than systematic change in December values. The GAMs are clearly sensitive to any kind of change in the seasonals. For example, the change in the size of the seasonals in the December months in the series may be the reason for this result. Since we are only interested in systematic shifts and not random changes the House series is found to have a stable seasonal pattern.

Figures: (4.3) and (4.4) show the diagnostic plots from the NS and TS models. The NS QQ plot and histogram show a violation in the Gaussian assumption while there is some remaining pattern in the residuals. Even though this is far from satisfactory it is to be expected since this series is suspected and tested positive for presence of seasonality. The TS model's diagnostic plots are better. There is a slight deviation from normality shown in the QQ plot and in the histogram of the residuals. The other plots, however, show no problems so the fit was found acceptable.

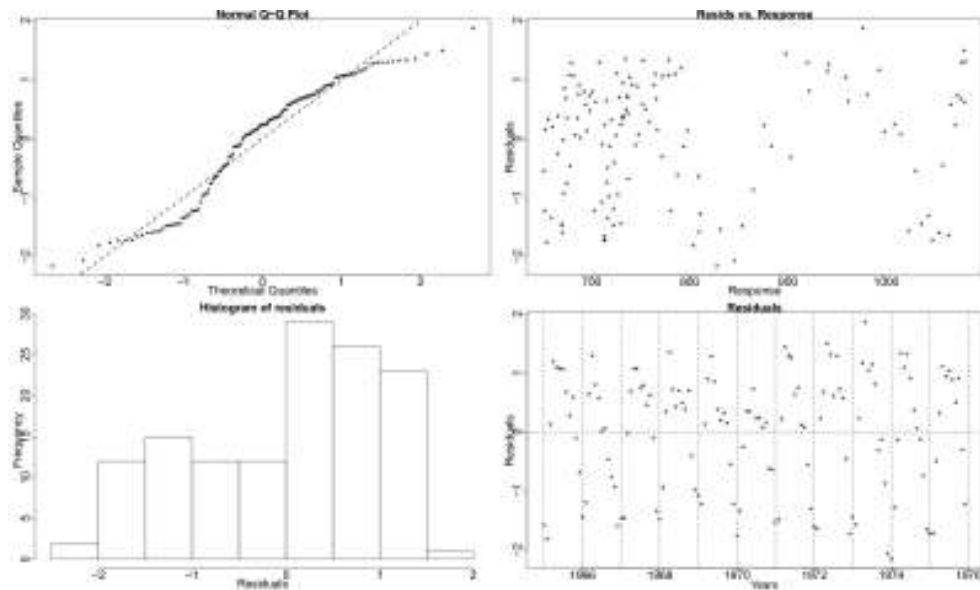


Figure 4.3: Residuals plots from the NS model for the House data. From the upper left corner clockwise: residual QQ plot, residuals plotted against the response, the residuals themselves and the residuals' histogram.

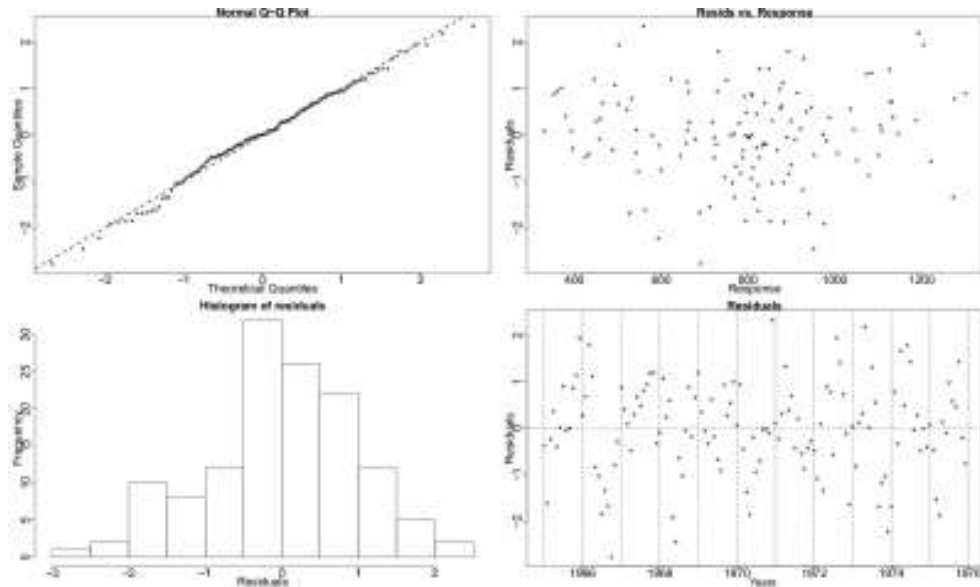


Figure 4.4: Diagnostic plots for the TS model for the House data.

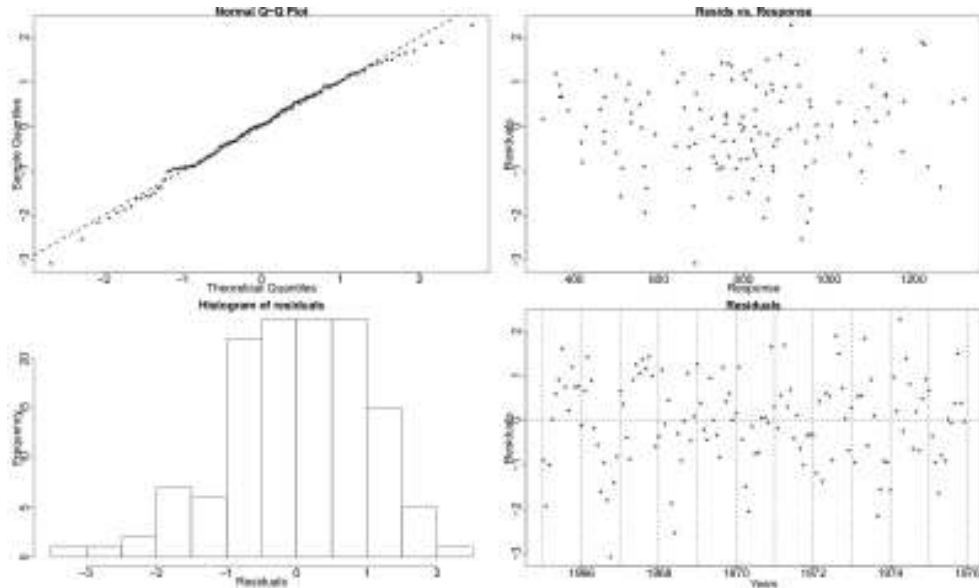


Figure 4.5: The diagnostic plots for the SS GAM fitted to the House data.

In Figure 4.5 we see some diagnostic plots for the chosen GAM with the stable seasonal pattern. The QQ plot shows that there is no serious violation of the Gaussian assumption which is further verified by the histogram, which is not ideal but satisfactory. The plot of the residuals against the response does not show a violation of the constant variance assumption. Finally, there are no patterns apparent in the plot of the residuals.

Unemployment Data

The GCV score suggests the TS model as most suitable for the Unemployment data, as with Resampling, Table 4.1.

GCV Score: Non-Seas: 0.104, Stable: 0.049 & **Trending:** 0.046

Table 4.4: Unemployment Data ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	121.857	11.102				
SS	110.306	3.877	11.551	7.225	17.796	< 2.2e-16
TS	96.287	2.397	14.019	1.48	4.24	< 2.2e-16

Table 4.5: Unemployment Data AIC & BIC:

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	10.143	21.694	35.713
AIC	68.725	-45.992	-80.941
BIC	97.965	16.548	22.013

The ANOVA, Table 4.4, between the models indicate that the series is seasonal with a trending seasonal pattern.

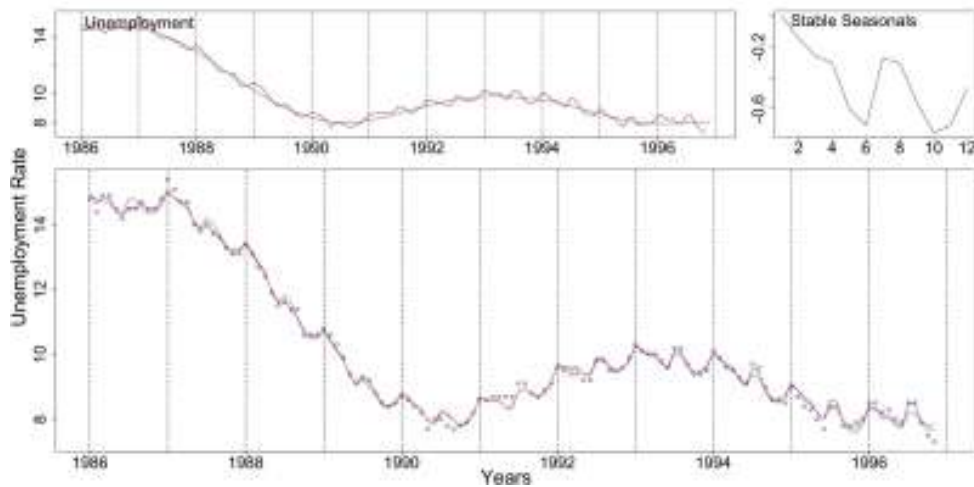


Figure 4.6: The Unemployment data with the trend from the SS model in red, the stable seasonals and the data as dots with stable (red) and trending (blue) GAM fits are depicted clockwise from upper left.

The SS and TS fits in Figure (4.6) indicate that the TS model is better. Also, in the same Figure, the terms of the SS model are plotted.

The twelve trending seasonals for the Unemployment data are shown in Figure (4.7). February, July and August seem to deviate from a straight line with their C.I. excluding zero for some points.

The BIC, Table 4.5, conservatively points to a stable seasonal pattern while the AIC recognizes a trending one. The BIC seems to favour smoother models while the AIC more closely fitted ones.




Figure 4.7: The smooth functions of the TS model of the Unemployment data, corresponding to the twelve months.

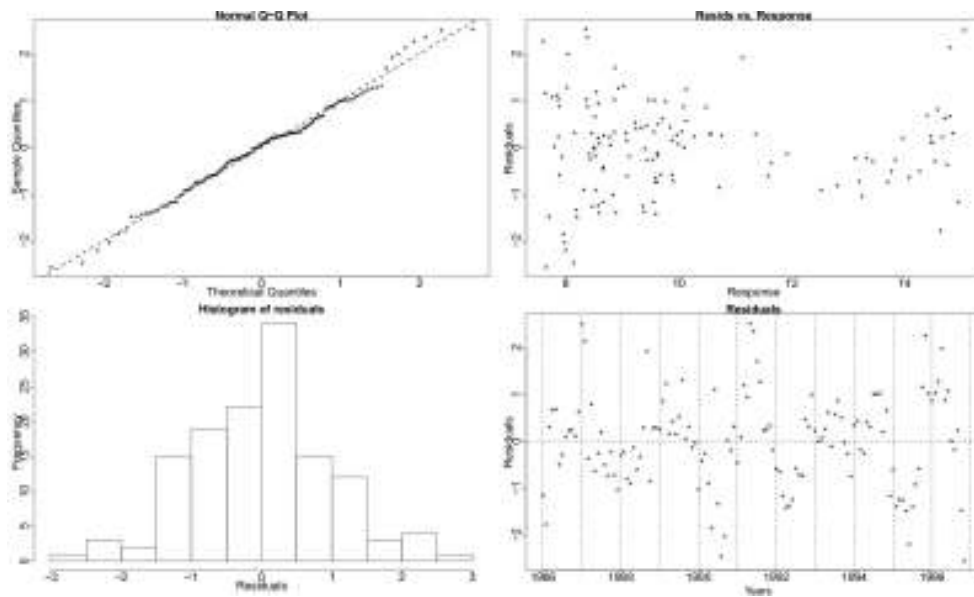


Figure 4.8: Diagnostic plots for the TS model of the Unemployment data. Clockwise from top left: a QQ plot of the residuals, the residuals plotted against the response, the residuals plot and histogram.

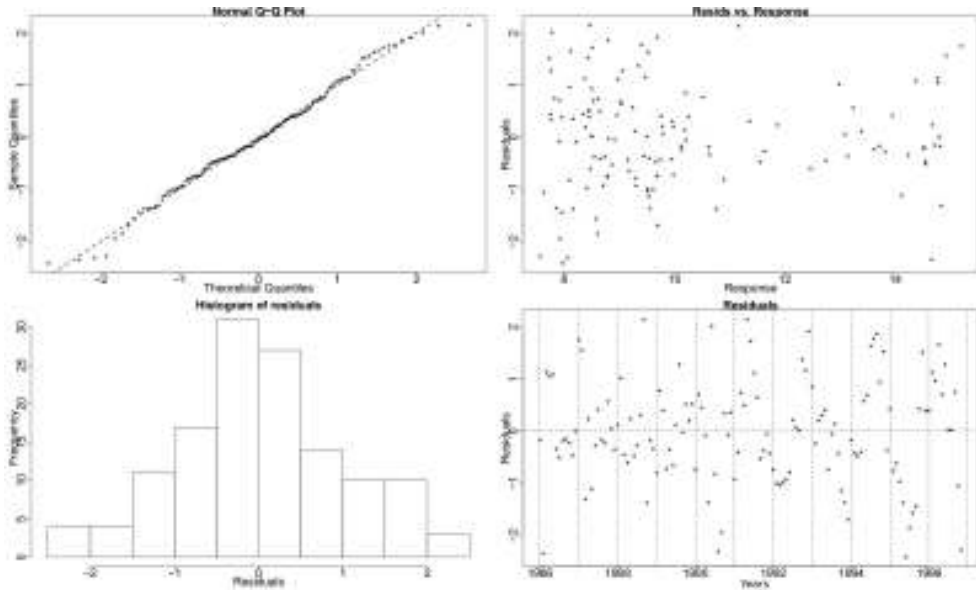


Figure 4.9: Diagnostic plots for the TS model with a larger basis dimension, 12, of the Unemployment data. Clockwise from top left: a QQ plot of the residuals, the residuals plotted against the response, the residuals plot and histogram.

The residual plots for the chosen TS model in Figure (4.8) show a slight deviation from the normality assumption in the QQ plot and histogram of the residuals. This is mainly evident in the tails of the distribution. In addition, there is some remaining pattern in the residuals as the model does not go high enough in the peaks and low enough in the troughs.

Trying to accommodate the remaining pattern we fitted the TS model with basis dimension equal to 12. The checking plots in Figure (4.9), however, are not satisfactory. The QQ plot shows a deviation from normality which is further verified by the histogram which shows that the distribution of the residuals is left skewed. The skewness is caused by the fact that the new model is going high in to the peaks of the series but not low enough in to the troughs thus creating more negative residuals than positive ones. The pattern in the residuals now is not so profound at the beginning but continues the same in the middle and end of the series. Thus, since none of our results were affected by this change, we decided that the first TS model, in

Figure (4.8), is more satisfactory.

Kola Data

The Kola data are deemed to have a stable pattern according to the GCV score, which agrees with the Resampling result in Table 4.1.

GCV Score: Non-Seas: 0.691, **Stable:** 0.155 & Trending:0.168

Table 4.6: Kola Data ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	121.857	11.102				
SS	110.306	3.877	11.551	7.225	17.796	< 2.2e-16
TS	376.094	54.382	10.998	0.322	0.203	0.997

The ANOVA, Table 4.6, between the models further reinforces the opinion that the series is seasonal with a stable seasonal pattern.




Figure 4.10: The upper left plot depicts the Kola data with the trend from the SS GAM in red. The stable seasonals and the data as dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.

In Figure (4.10) we see the fits of the SS and TS models and the terms of the former. The stable seasonals look a lot like a sinewave which is expected since as a

temperature series it should follow a smooth pattern of change. The trend reflects the changes in the level of the series.




Figure 4.11: The smooth trending seasonals estimated in the TS model for the Kola data.

The trending seasonals can further be examined by looking at Figure (4.11). No month has a slope that is significantly different from zero.

Table 4.7: Kola Data AIC & BIC:

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	10.553	21.908	32.906
AIC	1001.01	381.864	401.449
BIC	1043.34	469.743	533.442

BIC and AIC, Table 4.7, recognize a stable seasonal pattern. The Kola data are temperature data so they exhibit a very well behaved stable pattern that does not change in time in any way and thus, even the GAMs, which are sensitive to any kind of change, do not pick up a trend in the seasonals.

Since the SS model is chosen for this series, its residual checking plots are shown in Figure (4.12). There seem to be no violations of the assumptions but the residuals still show a pattern in them. This is because the SS GAM, as well as the TS, fails

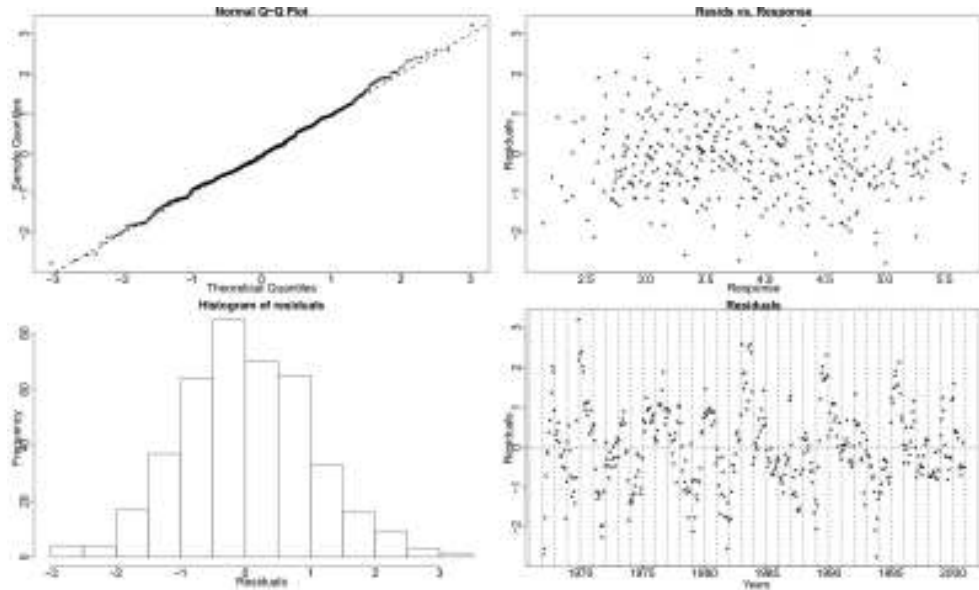


Figure 4.12: Here the residual checking plots for the stable GAM of the Kola data.

to capture all the variability in the series. Thus, they do not go as high in the peaks of the data or as low in the troughs, hence leaving some seasonality not included in the model.

4.3.3 Discussion

The GAMs offer a simple and straightforward way to construct hypothesis about the nature of the seasonal components in the series. The results from tests between the constructed models are easily interpreted. However, as seen through the examples, the GAMs are very sensitive and are quite close to performing twelve tests for unit roots, when dealing with monthly series as the above examples, or fifty-two when dealing with weekly series as the data sets from Stonehaven. Nevertheless, the only assumption made in the construction of these models is a distributional one for the response variable and thus are more robust when testing between different kinds of seasonality. The following chapter presents a sensitivity analysis and a comparison of the two testing methodologies discussed in this chapter. Furthermore, ways to

control various kinds of change in the seasonals are investigated as sometimes the presence of one change in the seasonals can interfere with the identification of a systematic shift in time.

Chapter 5

Comparison of Methods

5.1 Comparison of Methods

In order to compare the methodologies presented in the previous chapter and to evaluate their shortcomings and advantages a brief sensitivity analysis is conducted and the summary of the results is given in this chapter. We create two different seasonal patterns to correspond to the environmental series and to the densities of the zooplankton and phytoplankton species. We then simulate different types of changes and try to measure the ability of the Resampling and the GAMs to distinguish between a stable and a trending seasonal pattern.

The first pattern we will examine is the seasonal pattern of the Temperature 45m series from Stonehaven and is denoted by S_{env} . This is estimated using Friedman's super-smoother on the de-trended series and is a sinewave with amplitude of 6.7. The second set of seasonals, S_{dens} , corresponds to the densities of the zooplankton and phytoplankton species. In order to construct density data with zero values we generate Negative Binomial random variables with means equal to the seasonal pattern of the *C. finmarchicus* C5 series from Stonehaven as this is estimated by CSD methods using Friedman's super-smoother. These random variables are divided by a

set volume of 200ml to generate densities. The levels of noise in the generated series are controlled by the θ parameter of the Negative Binomial distribution. Figure 5.1 shows the created patterns.




Figure 5.1: The seasonal patterns. From the top the pattern corresponding to the environmental data, the mean used for the Negative Binomial distribution and the pattern corresponding to densities, derived as a random Negative Binomial variable with the previous plotted mean.

We manipulate these seasonal patterns accordingly to construct different types of seasonality as presented in section 2.2 and use the Resampling and GAMs to test the patterns. Once the seasonals from the S_{env} pattern have the appropriate behaviour we add to them white noise of different variances. The TS GAM fitted to the Temperature 45m series had residuals with s.d. = 0.46 giving a signal to noise ratio of $3.35/0.46 = 7.283$. Thus, we create noise with s.d. equal to 1, 0.5 and 0.25 and then depending on how well or badly the method performs we also use noise with s.d. equal to 2 or 0.1, respectively.

The *C. finmarchicus* C5 series has a mean $\hat{\mu} = 1.75$ and a variance of $\hat{\sigma}^2 = 33.07^1$. Assuming the Negative Binomial distribution, the θ parameter can be calculated as

$$^1 \hat{\sigma}_{res}^2 = 17.14$$

$\hat{\theta} = \hat{\mu}^2 / (\hat{\sigma}^2 - \hat{\mu})$ since $\hat{\sigma}^2 = \hat{\mu} + \hat{\mu}^2 / \hat{\theta}$. This gives a value of 0.098² for θ . Therefore, for the S_{dens} pattern we use values of 0.5 and 0.03 for the θ parameter. We note that the variance (i.e. noise) for a Negative Binomial variable decreases as the value of the θ parameter increases. When the method performs badly we also use a value of 1 for the parameter while if the method performs well we also use $\theta = 0.01$. For the rest of this chapter, we denote by S_{env} and S_{dens} the patterns after the appropriate manipulations according to the type of change examined are performed.

We should note that when creating random changes, systematic ones may be created by chance. Since we generate 11 years of observations for the simulations, we used a sequence of values, 1 : 11, to estimate the random occurrence of systematic changes. Thus, we created 1000 random ordered samples without replacement from the set of 1 : 11. Then the correlations between the randomly created sets and the original ordered sequence were computed. From those, 29 correlations were found greater in absolute value from 0.7. A histogram of them is shown in Figure 5.2. In 1000 randomly created sets 29 were found highly correlated. This in the context of this chapter means that when we are constructing random changes there is roughly a chance of 2.9% that a systematic one will occur.

We avoid running weekly GAMs because they are computationally intensive and instead create monthly series for GAMs while we use weekly for the Resampling. Both weekly and monthly series run for eleven years and the monthly ones are simply monthly aggregates of the weekly series, assigning four weeks to each month with 13 months within a year.

For each type of seasonality simulated we create 100 random instances that are examined by the two methods. For random changes in the S_{env} pattern we create 100 random changes and add to them randomly generated noise while for systematic ones, where only one pattern can be created, random instances are created by adding 100 different random noise sequences. Random instances for the S_{dens} are easily created as we generate 100 different sets of observations from a Negative Binomial

²0.199 with $\hat{\sigma}_{res}^2$




Figure 5.2: The computed correlations between 1 : 11 and the 1000 randomised ordered sets of them. The dotted lines mark 0.7 and -0.7 .

distribution. Furthermore, while for systematic changes the change in each year is relative to the previous one, e.g. when shifted earlier in time this is relative to the previous year, for random changes the change is always relative to the starting year, e.g. shifting randomly in time by an equal amount each year having as starting point the time in the first year.

For model choice for the GAMs we note the GCV, AIC and BIC scores and the ANOVA between the SS and TS models. All Resampling tests below are between SS and TS. For the S_{dens} created series we use the logarithm of the series plus 1 for the Resampling and a log-Normal distribution for the GAMs while the original series and a Normal distribution are used for the S_{env} pattern.

5.1.1 Stable Seasonality

We start by examining the ability of the Resampling and the GAMs to identify a stable seasonal pattern that exhibits no changes. Thus, we keep the patterns

unchanged and start by testing with Resampling between SS and TS. Tables are produced to summarise the percentage of correct results for each series. A correct result in this case is stable seasonality.

In Table 5.1 we see that the Resampling identifies the stable seasonality every time when the added noise has $\text{s.d.} \leq 1$ and more than 90% of the time with $\text{s.d.} = 2$ for the S_{env} pattern.

Table 5.1: Stable Seasonality - S_{env} Resampling Results

Seasonals	Noise s.d.			
	2	1	0.5	0.25
S_{env}	96%	100%	100%	100%

Table 5.2 shows the Resampling results for the S_{dens} pattern. Similar to the results for the S_{env} one, the Resampling performed on the S_{dens} pattern recognises stable seasonality all the time, even when $\theta = 0.01$.

Table 5.2: Stable Seasonality - S_{dens} Resampling Results

Seasonals	Theta		
	0.5	0.03	0.01
S_{dens}	100%	100%	100%

The GAMs do well in identifying a SS pattern in both occasions, Tables 5.3 and 5.4. Therefore, we also use noise with s.d. equal to 2 for the S_{env} pattern and $\theta = 0.01$ for the S_{dens} . The GAM results for the S_{env} pattern improve as the size of the noise is reduced and overall are satisfactory.

For the S_{dens} pattern the GAM results are also very good, see Table 5.4. As the value of the θ parameter increases the percentage of correct results increases, too.

Table 5.3: Stable Seasonality - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	77%	76%	100%
	1	100%	85%	82%	100%
	0.5	100%	93%	90%	100%
	0.25	100%	100%	100%	100%

Table 5.4: Stable Seasonality - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	100%	78%	81%	100%
	0.5	97%	75%	79%	99%
	0.03	93%	71%	77%	96%
	0.01	90%	67%	73%	95%

5.1.2 Systematic Time Shift

Now we examine the ability of the resampling and the GAMs to identify a systematic shift in time as trending seasonality. To simulate small systematic changes in time we fit a spline to the seasonals and then predict 520 points within a year. Thus, we can now move the patterns by tenths of a week per year. We examine the situations in which the pattern moves by one, two, three or four tenths of a week per year. For our analysis with GAMs we use monthly data with the equivalent monthly shifts, e.g. 1/40 of a month per year. Then we evaluate the ability of the two methods to identify the trending patterns.

Systematic Time Shift by 1/10 of a week per year

We simulate a systematic shift in time for the patterns, shifting them by one tenth of a week earlier every year. In table 5.5 we see the result from the Resampling for the S_{env} created series. The pattern is correctly identified as trending only about a quarter of the time when the noise has a s.d. equal to 0.5. When the noise is reduced further (s.d. = 0.25 and 0.1) the Resampling identifies the time shift 80% and 100% of the time, respectively. This shift is very small and therefore we did not expect it to be easily detectable after we add noise to the pattern. Nonetheless, there is an improvement as the size of the noise is reduced and when the noise is small the Resampling test always identifies the pattern as trending.

Table 5.5: Systematic Time Shift by 1/10 of a week - S_{env} Resampling Results

Seasonals	Noise s.d.			
	1	0.5	0.25	0.1
S_{env}	10%	24%	80%	100%

Table 5.6 summarises the results for the S_{dens} created series. The Resampling does not perform very well for the S_{dens} , with correct results less than 55% of the time for all attempted θ values.

Table 5.6: Systematic Time Shift by 1/10 of a week - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	52%	17%	3%

The GAM results for the S_{env} are presented in Table 5.7. It is interesting that for the S_{env} pattern the ANOVA and the AIC, when the size of the noise is equal to 1, almost always detect a TS pattern while the GCV score and the BIC detect a stable pattern. As the s.d. of the noise is reduced the results improve and when the s.d. of the noise is 0.25 all criteria always point to a trending pattern.

Table 5.7: Systematic Time Shift by 1/40 of a month - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	1	25%	92%	94%	4%
	0.5	100%	100%	100%	91%
	0.25	100%	100%	100%	100%

In Table 5.8, the criteria for the GAMs mostly fail to identify the change for the S_{dens} for all θ parameter values examined. The ANOVA and AIC point to TS more often than not when $\theta = 1$.

Table 5.8: Systematic Time Shift by 1/40 of a month - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	7%	84%	72%	1%
	0.5	6%	47%	36%	1%
	0.03	6%	23%	14%	1%

Systematic Time Shift by 2/10 of a week per year

Here we simulate again a systematic shift in time for the patterns, but shifting them by two tenths of a week earlier every year. This is twice the shift attempted earlier.

The Resampling results for the S_{env} , Table 5.9 have improved from the previous 1/10 of a week shift results. We also note the great improvement with the reduction of the size of the noise. The shift is now always identified when the s.d. is equal to 0.25 instead of 0.1 that was previously. Nonetheless, the shift remains mostly unidentified when a noise with s.d. greater than 0.5 is used.

Table 5.9: Systematic Shift 2/10 of a week - S_{env} Resampling Results

Seasonals	Noise s.d.		
	1	0.5	0.25
S_{env}	18%	77%	100%

In Table 5.10 we see that the results for the S_{dens} pattern have not improved greatly.

Table 5.10: Systematic Shift 2/10 of a week - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	60%	29%	16%

Looking at the results from the GAMs, Tables 5.11 and 5.12, we notice that they have improved overall for the S_{env} series and now the systematic time shift is always identified by the GCV score, the ANOVA and the AIC when the noise has s.d. equal to 1. The BIC identifies TS 92% of the time with s.d.= 1 and always for smaller noise. Hence, we attempted adding noise with s.d.= 2. The ANOVA and AIC still identify a trending pattern more than 90% of the time.

The results for the S_{dens} have also improved. For all attempted values of θ the

Table 5.11: Systematic Shift 2/40 of a month - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	30%	94%	95%	0%
	1	100%	100%	100%	92%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

trending seasonality remains undetected by the GCV score and the BIC most of the time. When $\theta = 1$, though, the ANOVA and AIC point to a TS.

Table 5.12: Systematic Shift 2/40 of a month - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	16%	88%	75%	8%
	0.5	7%	60%	38%	4%
	0.03	5%	34%	21%	3%

Systematic Time Shift by 3/10 of a week per year

Now we create series based on the two patterns after we shift them by 3/10 of a week per year.

We notice that the percentage of correct Resampling results for the S_{env} pattern, Table 5.13 has doubled from the previous section for noise with s.d. greater than 0.5.

Table 5.13: Systematic Time Shift by 3/10 of a week - S_{env} Resampling Results

Seasonals	Noise s.d.		
	1	0.5	0.25
S_{env}	44%	100%	100%

The Resampling results for the S_{dens} series have also improved, Table 5.14. The trending pattern is identified by the Resampling 82% of the time for $\theta = 1$. For $\theta < 1$, however, the pattern is mistaken for stable.

Table 5.14: Systematic Time Shift by 3/10 of a week - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	82%	51%	26%

The percentages of correct results from the criteria used on GAM selection are presented in Tables 5.15 and 5.16. For the S_{env} series there is a significant improvement with all criteria identifying the shift 100% of the time when the signal to noise ratio is $3.35/1 = 3.35$. For added noise with s.d.= 2 the GCV score, the ANOVA and the AIC point to a TS pattern more than 90% of the time while the conservative BIC only 40%.

There is a slight improvement for the S_{dens} series but still the shift is mostly undetected when $\theta < 1$. The results improve as the level of noise is decreased.

Table 5.15: Systematic Time Shift by 3/40 of a month - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	91%	99%	99%	40%
	1	100%	100%	100%	100%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

Table 5.16: Systematic Time Shift by 3/40 of a month - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	25%	90%	82%	4%
	0.5	13%	71%	42%	3%
	0.03	5%	40%	29%	1%

Systematic Time Shift by 4/10 of a week per year

Finally, we create a shift of a 4/10 of a week per year, equivalent to 4/40 of a month per year. The Resampling results for the S_{env} , Table 5.17, are further improved and now identify trending seasonality more than 50% when used on a series with added noise with s.d.= 1. The results for the S_{dens} have also improved and now is identified as trending seasonality 93% of the time with $\theta = 1$.

Table 5.17: Systematic Time Shift by 4/10 of a week - S_{env} Resampling Results

Seasonals	Noise s.d.		
	1	0.5	0.25
S_{env}	52%	100%	100%

Table 5.18: Systematic Time Shift by 4/10 of a week - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	93%	62%	32%

The criteria used for the GAMs created for the S_{env} series almost always identify the trend in the series, Table 5.19. The percentage of correct results has improved and now the GCV score also points to a TS pattern always when the noise has s.d. = 2. In Table 5.20 we see that the results for the S_{dens} have improved, as well, and the TS is identified more than half the time by the ANOVA and the AIC with $\theta = 1$.

Table 5.19: Systematic Time Shift by 4/40 of a month - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	100%	100%	89%
	1	100%	100%	100%	100%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

Table 5.20: Systematic Time Shift by 4/40 of a month - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	52%	97%	91%	19%
	0.5	21%	82%	48%	7%
	0.03	9%	58%	34%	5%

5.1.3 Random Time Shift

We simulate a random change in time, shifting the patterns by one tenth of a week. The starting point of the shift is always the timing of the peak in the first year. The direction is chosen randomly, allowing also the pattern to not move, i.e. stay at the starting point. A pattern that exhibits random changes in time is considered a stable pattern, as it is pseudo-trending.

The percentage of correct Resampling test results for trending seasonal pattern in the S_{env} series are given in Table 5.21. The random change in time is never confused for a systematic change by the Resampling algorithm.

Table 5.21: Random Time Shift 1/10 of a week - S_{env} Resampling Results

Seasonals	Noise s.d.			
	2	1	0.5	0.25
S_{env}	100%	100%	100%	100%

In Table 5.22 we see that the Resampling results for the S_{dens} also do not confuse the random shift in time with a systematic one. For all attempted values of θ for the generated Negative Binomial variables more than 90% of them are found stable.

In Tables 5.23 and 5.24 we see the percentage of correctly identified SS by the different criteria used on the constructed GAMs. For the S_{env} series the GAMs correctly identify a stable pattern most of the time. As the size of the noise is decreased the

Table 5.22: Random Time Shift 1/10 of a week - S_{dens} Resampling Results

Seasonals	Theta		
	0.5	0.03	0.01
S_{dens}	96%	96%	93%

ANOVA and AIC results improve while the GCV and BIC identify it as stable even when adding a noise with s.d. equal to 2.

Table 5.23: Random Time Shift 1/40 of a month - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	65%	63%	100%
	1	100%	79%	77%	100%
	0.5	100%	81%	80%	100%
	0.25	100%	80%	78%	100%

For the S_{dens} series the GAMs also, more often than not, say that it is stable. The above imply that the random shift in time is not treated by the GAMs as a trend in the seasonals, which would not necessarily be wrong but according to our definition of trending seasonality it is undesirable.

Table 5.24: Random Time Shift 1/40 of a month - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	87%	59%	79%	100%
	0.5	96%	49%	65%	99%
	0.03	93%	65%	73%	98%
	0.01	93%	54%	67%	100%

5.1.4 Systematic Amplitude Change

We simulate a systematic amplitude change (increase) by a 1/10 of the amplitude size per year. We consider this trending seasonality as the change is of a systematic nature.

The Resampling test results for the S_{env} pattern are presented in Table 5.25. The resampling always detects the trending pattern.

Table 5.25: Systematic Amplitude Change - S_{env} Resampling Results

Seasonals	Noise s.d.			
	2	1	0.5	0.25
S_{env}	98%	100%	100%	100%

The Resampling results for the S_{dens} series are presented in Table 5.26. For $\theta = 1$ the Resampling always identifies a TS pattern.

Table 5.26: Systematic Amplitude Change - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	100%	60%	25%

In Table 5.27 the results from the GAMs for the S_{env} series are summarised. The results look good and the size of the noise does not appear to influence the ability of the GAMs to identify trending seasonality.

The results for the S_{dens} are at about the same level as the ones for the systematic time shift by 4/40 of a month. The ANOVA and the AIC are the two criteria that more often recognise the trend in the seasonals.

Table 5.27: Systematic Amplitude Change - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	100%	100%	100%
	1	100%	100%	100%	100%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

Table 5.28: Systematic Amplitude Change - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	48%	92%	87%	13%
	0.5	13%	67%	56%	2%
	0.03	5%	42%	26%	0%

5.1.5 Random Amplitude Change

We now create random amplitude changes in the two seasonal patterns. The change is by 1/10 of the size of the pattern and the direction (increase or decrease) is randomly chosen. This type of pattern is considered pseudo-trending and thus stable.

The Resampling results for the S_{env} series are presented in Table 5.29. Even when the noise is big the pattern is always identified as stable while when the noise is significantly reduced it is identified as stable only 97% of the time. This could be attributed to the fact that since the direction of the changed is randomly chosen systematic changes (i.e. in one direction) may in fact be created. We, nonetheless, believe that the Resampling does not recognise a random amplitude change as TS.

Table 5.29: Random Amplitude Change - S_{env} Resampling Results

Seasonals	Noise s.d.			
	2	1	0.5	0.25
S_{env}	100%	100%	100%	97%

The results from the Resampling on the S_{dens} series are summarised in Table 5.30. The S_{dens} series is more than 90% of the time recognised as SS when in the presence of noise, even when using $\theta = 0.01$. We thus, believe that the fluctuation in the amplitude is identified as stable by the Resampling.

Table 5.30: Random Amplitude Change - S_{dens} Resampling Results

Seasonals	Theta		
	0.5	0.03	0.01
S_{dens}	96%	94%	93%

The GAM results for the S_{env} are presented in Table 5.31. The ANOVA and AIC results deteriorate when the noise is reduced. This can be translated to mean that these two criteria identify random amplitude change as TS but in the presence of a big noise it remains undetected. Therefore, we also use noise with s.d.= 2 and s.d.=0.01

to have a clearer picture of the behaviour of the results. We believe that the ANOVA and AIC identify a random amplitude change as trending seasonality.

Table 5.31: Random Amplitude Change - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	65%	57%	100%
	1	100%	27%	16%	100%
	0.5	99%	1%	1%	100%
	0.25	100%	0%	0%	100%
	0.01	100%	0%	0%	100%

The results for the S_{dens} are similar to the ones for the S_{env} pattern. As the noise decreases the results from the GCV score and the BIC improve while the AIC and ANOVA detect more frequently a TS pattern. Again, this verifies that the ANOVA and AIC treat a random amplitude change as trending seasonality.

Table 5.32: Random Amplitude Change - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	90%	27%	42%	100%
	0.5	91%	39%	61%	100%
	0.03	90%	57%	75%	100%
	0.01	87%	61%	71%	99%

5.1.6 Systematic Time Shift and Random Amplitude Change

Since in most marine biological data-sets the aforementioned changes rarely occur alone we will now examine the most interesting case, for the purpose of this thesis, of simultaneously occurring changes. We care mostly about systematic shifts in time and thus, we would like to evaluate how our methods respond in the presence of a systematic time shift and a random amplitude change. This stems from our concern that a random change may mask a systematic one and thus we are not that concerned with the cases of two random, or two systematic changes occurring in the same pattern.

Systematic Time Shift 1/10 of a week and Random Amplitude Change

We shift the pattern by 1/10 of a week per year or 1/40 of a month per year and allow the amplitude to change randomly by a 1/10 of its size per year.

The Resampling results for the S_{env} series are given in Table 5.33. We see that the Resampling almost always fails to detect the systematic shift in the presence of a random amplitude change, when the noise has s.d. ≥ 0.5 . This not surprising as the time shift is very small and was not easily detected in the presence of noise when no amplitude change was occurring.

Table 5.33: Systematic Time Shift 1/10 and Random Amplitude Change - S_{env} Resampling Results

Seasonals	Noise s.d.			
	1	0.5	0.25	0.1
S_{env}	8%	17%	69%	98%

The Resampling results for the S_{dens} series can be seen in Table 5.34. The S_{dens} series is hardly ever found trending. The systematic shift of 1/10 is too small to be detected in the presence of a random amplitude change.

Table 5.34: Systematic Time Shift 1/10 and Random Amplitude Change - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	11%	4%	1%

The percentage of correct GAM results for the S_{env} are shown in Table 5.35. The results improve as the noise decreases in size. All the criteria point to a TS pattern when the added noise has $s.d. \leq 1$.

Table 5.35: Systematic Time Shift 1/40 and Random Amplitude Change - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	89%	100%	100%	32%
	1	100%	100%	100%	100%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

The S_{dens} is rarely recognised to have a TS pattern when $\theta = 0.03$. Decreasing the noise in these series, using $\theta = 1$, all criteria, except the BIC, more often than not recognise a TS pattern when $\theta > 0.5$.

Table 5.36: Systematic Time Shift 1/40 and Random Amplitude Change - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	58%	85%	73%	41%
	0.5	35%	65%	51%	27%
	0.03	25%	34%	28%	9%

Systematic Time Shift 2/10 of a week and Random Amplitude Change

Since a shift of 1/10 of a week per years is quite small we double the shift but keep the amplitude change the same. The Resampling results for the S_{env} , Table 5.37 have improved with TS being identified 100% of the time when the noise is small, s.d.= 0.1.

Table 5.37: Systematic Time Shift 2/10 and Random Amplitude Change - S_{env} Resampling Results

Seasonals	Noise s.d.		
	1	0.5	0.25
S_{env}	16%	68%	100%

In Table 5.38 the Resampling results have significantly improved for the S_{dens} series. Nonetheless, the trending pattern is still only recognised 34% of the time when $\theta = 1$.

Table 5.38: Systematic Time Shift 2/10 and Random Amplitude Change - S_{dens} Resampling Results

Seasonals	Theta		
	1	0.5	0.03
S_{dens}	34%	18%	3%

The GAM results for the S_{env} series are presented in Table 5.39. The series is now recognised to have a trending pattern by all criteria even with s.d.= 2.

The results for the S_{dens} series, also have improved. The ANOVA and the AIC continue to point to a TS pattern when $\theta \geq 0.5$ but the GCV score and the BIC still identify a seasonal pattern stable when $\theta = 0.03$.

Table 5.39: Systematic Time Shift 2/40 and Random Amplitude Change - S_{env} GAMs Results

		Criterion			
Seasonals	Noise s.d.	GCV	ANOVA	AIC	BIC
S_{env}	2	100%	100%	100%	100%
	1	100%	100%	100%	100%
	0.5	100%	100%	100%	100%
	0.25	100%	100%	100%	100%

Table 5.40: Systematic Time Shift 2/40 and Random Amplitude Change - S_{dens} GAMs Results

		Criterion			
Seasonals	Theta	GCV	ANOVA	AIC	BIC
S_{dens}	1	80%	100%	100%	53%
	0.5	66%	84%	79%	40%
	0.03	31%	51%	48%	23%

5.1.7 Discussion of Results

Though the above sensitivity analysis is far from complete it gives us an idea of how our methodologies behave, indicating their strengths and weaknesses. We observe the different behaviours of the methods towards the two distinct patterns employed above. Overall, the Resampling seems more reliable when testing the S_{env} pattern while the methods do overall equally well for the S_{dens} .

Both methods identify a systematic change as TS, with the GAMs being the more sensitive of the two in the case of the S_{env} . Similarly, both methods consider a random shift in time mostly as SS. A random amplitude change, however, is considered stable by the Resampling for the S_{env} and S_{dens} while found trending by the ANOVA and AIC criteria for both patterns. In addition, the GAMs appear to be more sensitive towards systematic shift in time in the presence of a random amplitude change than the Resampling. The difference between the results from the two methods is more pronounced for the S_{dens} pattern. This, however, could be attributed to the fact that the GAMs identify a random amplitude change as trending seasonality.

Generally, the ANOVA and AIC seem to identify small random changes in the noise added to the pattern as trends in the seasonals. From all criteria used for model selection the ANOVA is the least conservative, often favouring the TS model while the BIC is the most conservative one, pointing to a stable pattern even when all other criteria disagree. Nonetheless, under extreme circumstances (when there is a great trend in the seasonals or no trend at all) they both succeed in identifying the appropriate type of model, TS or SS. The ANOVA identifies correctly as stable seasonals the random time shift 80% of the time for the S_{env} when the noise s.d. is equal to 0.5 and the BIC, even though conservative, does point to TS 100% of the time for the S_{env} when exhibiting a systematic time shift of 4/40 of a month per year or a shift of 3/40 but with small noise.

Another interesting point is that if there is a systematic shift in time along with a random change in the amplitude of the pattern the methods' successful detection of

a trending seasonality depends on the size of the shift relatively to the size of the random amplitude change. In view of this, it is recommended that once a change of the seasonals has been observed and possibly even verified with some tests the series should be scaled to accommodate this and then the above tests should be conducted.

The above means that if the Resampling algorithm indicates a trending seasonality for a series similar to the S_{env} , this is always of the systematic type. In addition, however, it implies that when the Resampling finds a stable seasonality, there is still a probability that a systematic shift in time exists but in the presence of a random change in the size of the seasonals or that the noise in the series is too big. To summarise, the Resampling does not appear to give false alarms of trending seasonality but may fail to detect one under specific conditions.

For the GAMs, however, the interpretation of the results is not as straightforward. It seems that small changes, which could in fact be simple noise, can bias the results towards trending seasonality. Therefore, the GAMs are considered more sensitive to any kind of change than the Resampling. Looking closely at the output available from the GAMs, one can, nonetheless, see whether the trend identified in the seasonals is of a systematic nature. The plots of the smooth functions fitted to the trending seasonals can shed light to the nature of the detected change. A systematic change (time or amplitude) is manifested as a monotonic line or curve while anything else points to a random change.

Neither approach performs as well on the S_{dens} as on the S_{env} series. The sudden changes in the pattern itself make it hard to detect changes of a smooth nature. In addition, the Resampling is very conservative towards the Negative Binomial created data as systematic changes are often found stable. The GAMs do not perform very well for the Negative Binomial data either but outperform the Resampling in most situations. The S_{dens} pattern corresponds to densities and thus, behaves very differently from the S_{env} pattern.

As a final note the reader is reminded that in nature a change in the seasonal pattern

hardly ever is of only one type. Issues arise when testing for a specific type of change in the presence of other types, “unattended”, changes. This is similar to conducting seasonal unit root tests for specific roots in the presence of other “unattended” unit roots. Therefore, the above methodologies should be applied with care and if possible together as they have different strengths and weaknesses and thus, compliment each other.

Chapter 6

Stonehaven Results

6.1 The Data

In the following two chapters we present the results from our analysis of the data from Stonehaven. The data consists of weekly samples since January 1997 from the Stonehaven long-term monitoring station, western North Sea ($56^{\circ}57.8'N$ $02^{\circ}06.2'W$), in water depth of around 50m. The data can be broken down into three main categories: the environmental series, such as temperature, salinity and nutrients in the water, the phytoplankton, diatoms and dinoflagellates, and the zooplankton series. The length of the examined series varies. The environmental and zooplankton series run from the start of 1997 to the end of 2007 while the phytoplankton series from the start of 2000 to the end of 2008. The phytoplankton series is shorter as, in the beginning of 2000, a new, more reliable, way to count them was introduced. We examine 13 environmental series, 5 phytoplankton and 37 zooplankton series.

Due to unforeseen weather conditions, the samples are not taken on pre-specified dates and so they are not equi-spaced in time. For our analysis we divide the year in 52 weeks and then allocate the observations to weeks according to their sampling dates. The 26th of December is not included in any week but no problems should

occur since there is no sampling taking place on that date. We made no adjustments for the leap years, 2000, 2004 and 2008, as there are no observations taken on the 29th of February in any leap year.

Another option we considered, was to compute monthly averages of the series. Since, however, we are investigating shifts of seasonality and we would expect them to be usually small, hardly by a month, we would risk not capturing them. Looking at the data as daily observations, i.e. Julian dates¹, was also considered as an option but the number of missing values is too great to lead to accurate interpretation of the results. Therefore, we choose to treat the measurements as weekly observations (taking the average on rare occasions when there was more than one observation in a week), since there is no consistency of the sampling dates through the years.

There is, however, a notable trend in the sampling dates. In most weeks the dates of sampling are decreasing, i.e. 25th, 24th, 23rd of January, in consecutive years. A possible reason for this is that they are actually sampling on the same day, e.g. Mondays. Under our assumption of weekly aggregates this causes no problem but if the analysis shows any effects then we can run it through actual dates as mentioned above.

Nonetheless, there are weeks, as defined by us, with no samples and some with more than one sample. Having weeks with no data, introduces missing values in the series that have to be dealt with. Substitution of the missing values will introduce more uncertainty in our analysis, especially since there are many such runs, hence it is avoided. Instead of using interpolation and extrapolation methods to substitute the missing values we settle on using decomposition techniques that will allow for the existence of internal and external missing values in the series.

In addition, one should note that there are no samples taken in very bad weather. Since the collection of samples involves going to sea, no samples are collected when that is thought unsafe. Therefore, the months in winter time (December, January)

¹The sampling dates include leap years so a total of 4017 days for environmental and zooplankton series and 3288 days for the phytoplankton series.

are more sparsely sampled than the rest. That is to say that for some weeks in these months only few samples exist, e.g. only four samples are collected in week 1 as defined by us, between the 27th of December of one year and the 2nd of January of the next year.

6.2 Methodologies

An outline of the general adjustments made to the methodologies presented in Chapter 4, in order to be used for our analysis of the data, follows. Further modifications are needed to suit the individual characteristics of each type of data- environmental, phytoplankton or zooplankton. A section is devoted to each type of data, with details of the data set's collection method and the exact specifications of the methods used discussed.

6.2.1 Resampling

As mentioned before, when allocating observations to weeks within a year, missing values are produced. These missing values present an issue with the Resampling techniques. Resampling, as presented in Chapter 4, cannot be employed on a series with missing values so a modification is devised. The MAs used in the Resampling algorithms to extract the trends in the series cannot be computed in the presence of missing values. Thus, they are substituted by Friedman's super-smoother. The super-smoother can deal with missing values allowing the rest of the Resampling algorithms to remain unchanged. Furthermore, using the super-smoother instead of the MA to compute the trend helps ensure that no seasonality is induced into the series during the extraction of the trend.

Furthermore, the Resampling requires the data to be in the form of a time series so averages of all observations within one defined week are taken. Missing values are allocated to the weeks with no observations. The p-values corresponding to the

appropriate tests, i.e. NS v.s. SS and then either SS v.s. TS or NS v.s. TS, for each series are reported and interpreted.

6.2.2 GAMs

The construction of GAMs does not require the data in the form of a time series. Furthermore, it allows for more than one value per week thus utilizing all collected samples. We construct the three GAMs - NS, SS and TS- and using a number of criteria we then choose the model that best describes the data.

Thin plate regression splines are used to model the trend and trending seasonals. The dimension of the basis for the trend is the default, 10, while higher values are attempted when it is needed. The basis dimension for the trending seasonals, however, has to be specified for the model to be computed, due to convergence issues. We use three knots per season for the trending seasonals since we have at most eleven years of data, meaning at most 5 values to smooth over.

The nature of the data is such that one expects a smooth stable pattern without spikes. Using seasonal factors to model the stable seasonals provides a better fit to the data while a cyclic smoother may be closer to the Natural seasonal patterns and the interpretation of the series. Some seasonal patterns in the data are sharply defined, for particular species for example, and thus it can be difficult for a smoother to capture them but, nonetheless, any systematic changes are unlikely to be spiky. Seasonal factors allow big, abrupt changes to happen while a systematic change takes place more gradually. Furthermore, a seasonal factor corresponds to a whole week making the detection of small changes more difficult. A smooth curve produces a shape closer to the truth and makes the detection of systematic changes easier. Therefore, using seasonal factors for the stable seasonals does not always seem appropriate as the results would not be smooth for the species time series.

Nonetheless, for the environmental series, the factors are deemed appropriate as the patterns to be estimated are quite smooth and thus the seasonalities produced using

the factors will not be highly variable. Thus, for demonstration purposes, factors will be used to estimate the stable seasonals of all environmental series while a cyclic smoother will be employed for all the species series. This way we can demonstrate both methods.

In addition, in order, again, to favour smoother fits we are using $\gamma = 1.4$ instead of the default value of 1 in the construction of the GCV score as per Kim and Gu (2004).

6.3 Environmental Series Results

The environmental series consist of two measurements, one taken at surface level and one at approximately 45m depth. Thus, for example, we have surface Temperature (at 1m) and Temperature at 45m measurements. The one exception is *Chlorophyll a* for which we have only one series. The samples for *Chlorophyll a* detection are taken from a 10m integrating tube sampler (Lund sampler). Fluorometry is the method used to measure *Chlorophyll a*. A one litre sub-sample from the contents of the tube sampler is filtered to a glass fibre filter paper (GF/F, nominal 0.7 micron pore size). The pigments are extracted from the filter paper by soaking in 90% buffered acetone for 16 to 30 hours and measured on a fluorometer, as described in Smith *et al.* (2007).

For most environmental series (exceptions are Temperature, Salinity and surface Silicate) a second sub-sample is taken from the sample to create ‘duplicate’ measurements. When there is a discrepancy between the two measurements the average is used as the appropriate measurement for that collection date. In Resampling the data needs to be in the form of a time series and thus, the averages of the duplicate measurements are taken. Furthermore, in Resampling we are also averaging measurements taken on the same week, using weeks as defined by us. Nevertheless, as mentioned before, the GAMs allow for multiple entries per week and thus no averaging is done, neither for duplicates nor for weekly measurements.

The nutrients are measured by taking a sample of sea-water using Niskin bottles and then a sub-sample is measured in an automated nutrient analyser. The equipment used to measure nutrients in the water has a detection limit of around 0.01, below which zero values are reported. Since the equipment used to collect nutrient information has a detection limit, the zeros in the database most probably correspond to very small, undetectable traces of nutrients rather than their absence. Hence, for our analysis we substitute these zero values by very small but non-zero ones. We create random values in the range of 0 – 0.001 using a Uniform distribution. This ensures that we have non-zero values for our analysis but also that the values entered are random and a lot smaller than the detection limit of the equipment used. The series that have zero values that are substituted in this way are: Nitrate at 1m (surface), Ammonia at 1m, Phosphate at 1m, *Chlorophyll a* and Silicate at 45m.

We also have to consider the extreme values in the database. Extreme values could be caused by malfunctioning equipment or contamination of the sample and thus when one is observed we evaluate whether it is a true value. Examining the original values instead of the averaged ones and reading the various notes on that collection, reporting on malfunctioning equipment etc., we decide what the case is. Because we are interested in the seasonal pattern of these series, it is thought that extreme values, even when correctly recorded, will not contribute to a better estimation of the seasonal effects. Nonetheless, we only remove extreme values that are verified as erroneous. Where an extreme value is removed a missing value is introduced.

A special case of extreme values is the contamination of the Phosphate samples in 2006. All measurements from the 13th of June to the 2nd of October 2006 for Phosphate 45m and from the 13th of June 2006 to the 10th of January 2007 for Phosphate 1m are removed from the database. Missing values are placed in their space.

For the majority of the series presented in this section the Gaussian family is assumed with the identity link function for the construction of GAMs while for a few a log link is used. The Temperature, Salinity, Ammonia, Phosphate and Nitrate series are

modelled using the identity link, thus assuming a decomposition structure:

$$\hat{X}_t = \mu_t = T_t + S_t + R_t, \quad (6.1)$$

as the seasonals and the residuals are unaffected by the trend. On the other hand, the two Silicate series are modelled using a log link which implies a different structure to the above:

$$\log(\mu_t) = T_t + S_t \quad (6.2)$$

$$\hat{X}_t = \mu_t = e^{S_t + T_t} + R_t \quad (6.3)$$

and thus,

$$\hat{X}_t = e^{S_t} * e^{T_t} + R_t. \quad (6.4)$$

For these series the seasonals increase as the trend increases and decrease as the trend decreases, hence we assume a multiplicative relationship between these two components. The residuals are not affected by the changes in the level of the trend and therefore, they are added and not multiplied to the other two components. Durbin & Murphy (1975) discuss seasonal decomposition using mixed additive-multiplicative models similar to the one above. When the structure in (6.4) is assumed appropriate plots verifying this assumption will be presented.

The *Chlorophyll a* series is modelled using the Gamma family and the log link as the Gaussian models provided a very bad fit.

Furthermore, the seasonal pattern for each series is extracted using factors. The *Chlorophyll a* series, however, is an indicator of the phytoplankton biomass and is highly variable. Thus it will be modelled using the aforementioned cyclic smoothers. We believe that since we will use the cyclic components to model the stable seasonality in the individual species series it is appropriate to use it also for the *Chlorophyll*

a series. For the majority of the environmental series presented in this chapter, however, the two ways to model the stable seasonals differ only slightly but when the seasonals are very smooth or the series is noisy the cyclic smoothers or factors, respectively, could be used to estimate a smooth seasonal pattern.

All the environmental series are found seasonal, as it was expected. Therefore, in the following discussion of the results of our analysis we focus on whether the seasonality is stable or trending (i.e. exhibits systematic change). Resampling results are only presented for the stable seasonal (SS) vs trending seasonal (TS) test while all constructed GAMs are reported. The GCV score of GAMs can be used to choose between nested models with the lowest score being preferable. Therefore, for model selection we note the GCV score of the three GAMs, conduct ANOVA between the three models and report AIC and BIC scores. We also report the percentage deviance explained by each GAM as a measure of fit. Furthermore, when we suspect that the series has a trending seasonal pattern we look at the fitted smooth functions corresponding to each week to further verify this. Once we decide that a series is shifting in time we attempt to quantify the shift by examining the occurrence of some easily identifiable peaks in the de-trended fitted values of the TS GAM. This, nonetheless, is not always possible as there may not be identifiable peaks and/or the shape of the seasonal pattern may change greatly over time.

To check our GAMs for violations of any assumptions we create diagnostic plots. Using standardized deviance residuals, \hat{d}_{std}^2 , we create a QQ plot for normality, a plot of the residuals against the response, the residual histogram and a residual plot. When referring to ‘residuals’ below in the context of diagnostic plots we mean the previously defined standardized deviance residuals.

Table 6.1 summarises the results for all environmental series analysed. However, for most series only the surface measurement is presented in the main section of this thesis as similarities occur between the surface and the 45m measurements. Both

${}^2\hat{d}_{std,i} = (\hat{d}_i - \bar{\hat{d}}) / \left(\hat{\sigma} \sqrt{(1 - H_{ii})} \right)$, where H_{ii} are the elements on the leading diagonal of the hat matrix, H , and $\bar{\hat{d}}$ the average of the deviance residuals \hat{d}_i .

series of Salinity and Ammonia are discussed as their behaviours differed slightly. A full description of all results can be found in the electronic Appendix located at the back of the thesis. Plots and tables of results are presented only when they appear noteworthy for the particular series.

Table 6.1: Environmental Series Results

Series		Pattern
Temperature	1m	Pseudo-trending in Amplitude
	45m	Pseudo-trending in Amplitude
Salinity	1m	Trending (Time Shift Earlier)
	45m	Trending (Time Shift Earlier)
Silicate	1m	Pseudo-trending in Amplitude
	45m	Pseudo-trending in Amplitude
Phosphate	1m	Pseudo-trending in Amplitude
	45m	Pseudo-trending in Amplitude
Ammonia	1m	Trending (Time Shift Earlier)
	45m	Trending (Time Shift Earlier)
Nitrate	1m	Trending (Time Shift Earlier)
	45m	Trending (Time Shift Earlier)
<i>Chlorophyll a</i>		Trending (Amplitude Decreasing & Time Shift Earlier)

6.3.1 Temperature

Temperature data usually have a stable seasonal pattern, much like the Kola data discussed in Chapter 4. Sea water temperatures are affected by various factors. IACMST (2001) reports that changes in the local surface heat exchange, local wind field, influx of oceanic water and freshwater run-off are all contributors towards the changes in temperature. For example, uncommonly cold or warm years cause changes in the size of this pattern. If the years are getting colder or warmer gradually, the change will be systematic while if random climatological effects occur the pattern will fluctuate in size and/or timing of occurrence. The two Temperature series are modelled using the Gaussian family with an identity link.

Temperature - 1m

We start by examining the surface temperature (Temperature 1m). This series will be examined first using factors and then using a cyclic component for the seasonals. Thus, we will demonstrate that both methods of modelling the seasonal component produce the same results. This is especially true in the case of the following environmental series while for the species time series the cyclic component will be preferred due to the noise in those series.

The Resampling test between stable and trending seasonality indicates a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.649 \rightarrow$ Stable

This is further verified by the GCV score of the GAMs:

GCV Score: Non-Seas: 7.328, **Stable:** 0.619 & Trending: 0.759

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 1.01, Stable: 94 & Trending: 95.1

The ANOVA, in Table 6.2, however, points to the TS GAM.

Table 6.2: Temperature 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	520	3784.289				
SS	462.633	228.53	57.367	3555.759	125.477	< 2.2e-16
TS	406.649	188.96	55.984	39.57	1.521	0.013




Figure 6.1: The upper left plot depicts the Temperature at 1m with the trend from the SS GAM as a red line. The stable seasonals and the Temperature 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.

Figure 6.1 shows the trend of the SS GAM as a red line over the Temperature at 1m series, marked by a black line. In the same Figure the seasonals from the SS model are also plotted and the two fits for the SS (red) and TS (blue) models.




Figure 6.2: The upper shows the stable seasonals estimated by the SS GAM when using factors while the bottom one when using a cyclic component.

As mentioned before, here the stable seasonals in the GAM are estimated using factors. In Figure 6.2 we see the stable seasonals estimated using factors (above) and using a cyclic component (below). To aid comparison we plot only the estimated seasonals without adding the average trends estimated by the two SS models, with factors and with a cyclic component. The estimated patterns are very similar with the lower being a smoother version of the upper.

The smooth functions corresponding to the trending seasonals as estimated by the TS GAM are plotted in Figures 6.3 and 6.4. Weeks 31 and 50 deviate greatly from a straight line with their C.I. not always including the zero value. Nevertheless, a curve like the one for week 50, in Figure 6.4, is not of much interest to us as it simply implies a random and not systematic change. What is of great importance to our study is a significant monotonic upwards or downwards slope since that would be an




Figure 6.3: The smooth trending seasonals for weeks 1-36 estimated in the TS model for Temperature 1m.

indication of a systematic change, time shift or amplitude change. This leads us to believe that there is not a systematic change occurring in these weeks.

In addition, we would hardly expect a change in the seasonal pattern to occur in only two specific weeks (i.e. weeks 31 and 50) every year. We expect a sequence of trending weeks that would then indicate a change in the seasonals. Therefore, our decision is that this series has a stable seasonal pattern in terms of systematic changes but may exhibit a random change.

Table 6.3: Temperature 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	60.367	116.351
AIC	2521.426	1170.935	1183.652
BIC	2534.221	1428.418	1679.923

Furthermore, the AIC and BIC scores in Table 6.3 verify the choice of stable seasonality.

Nevertheless, to further examine the above identified as potential trending weeks, 31




Figure 6.4: The smooth functions fitted by the TS GAM, corresponding to the trending seasonals for weeks 37-52 for the Temperature 1m.




Figure 6.5: The stable seasonal pattern (red) and the trending seasonal pattern (blue) from the TS GAM fitted to the Temperature at 1m series.

& 50, we look at a graph the stable (red) and trending (blue) seasonals from the TS GAM, Figure 6.5. Using the stable seasonals as a point of reference we can note any changes in the size or the timing of peak of the trending seasonals. In addition, we can examine whether the change is occurring in one direction, systematic change, or both, random change. We thus verify that no systematic change occurs in the pattern of the Temperature at 1m series. There is, however, a small fluctuation in the size of the seasonals which points to a pseudo-trending in amplitude pattern.

We believe that the Temperature 1m series has a stable seasonal pattern and the ANOVA results are influenced by a very small random change in the size of the seasonals.

In Figures 6.6 and 6.7 we see some residual plots for the SS and TS GAMs. The plot on the top left corner is a QQ plot of the residuals. For both models this plot shows that the upper end of the series deviates from normality. The top right plot shows the residuals plotted against the response and is satisfactory. Similarly, the histograms are acceptable. The plots of the residuals from both models show some remaining pattern. This occurs because the fits of the SS and TS models do not go high enough in all the peaks and low enough in all the troughs and thus, some pattern is still left in the residuals.

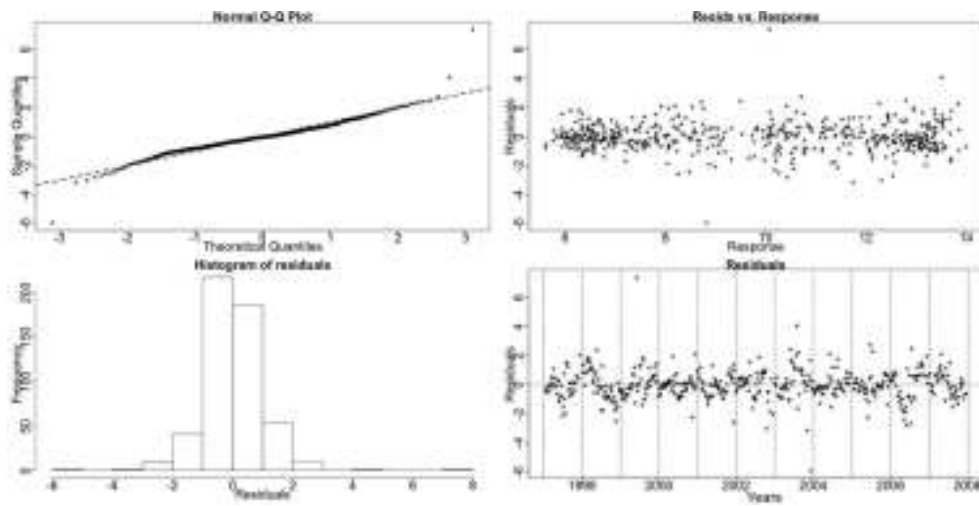


Figure 6.6: The residual checking plots for the SS GAM of the Temperature 1m data. No problems are highlighted.

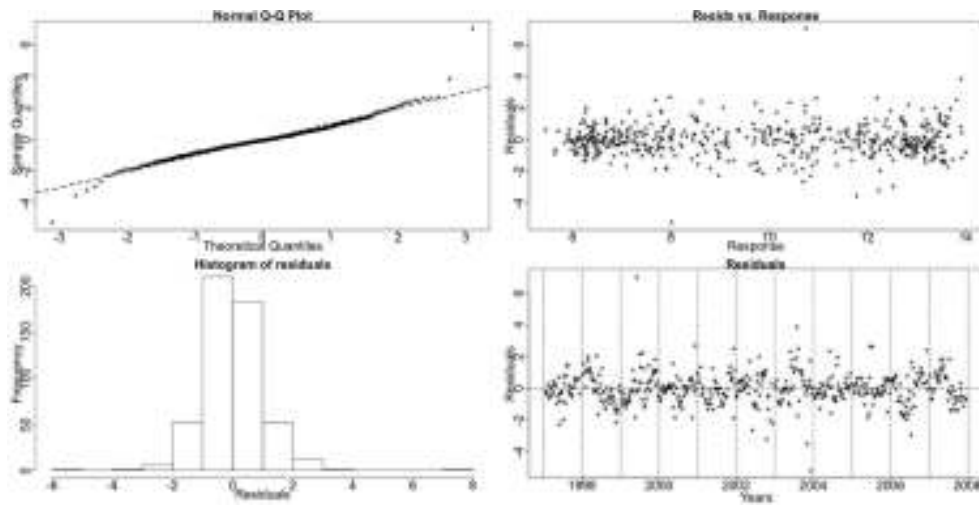


Figure 6.7: Residual checking plots for the TS GAM of the of the Temperature 1m data. No violations are noted.

Using a Cyclic Component to Model the Stable Seasonals

We examine now the surface temperature (Temperature 1m) using a cyclic component to model the stable seasonality. The GCV score of the GAMs points to the stable seasonal model:

GCV Score: Non-Seas: 7.328, **Stable:** 0.508 & Trending: 0.759

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 1.01, **Stable:** 93.6 & Trending: 95.1

Table 6.4: Temperature 1m ANOVA - Cyclic Component

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	520	3784.289				
SS	508.406	246.091	11.594	3538.198	630.495	< 2.2e-16
TS	406.653	188.966	101.753	57.125	1.208	0.104

The ANOVA, in Table 6.4, now also points to the SS GAM.




Figure 6.8: The upper left plot depicts the Temperature at 1m with the trend from the SS GAM as a red line. The stable seasonals and the Temperature 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise. Both GAMs use a cyclic component for stable seasonality.

Figure 6.8 shows the trend of the SS GAM as a red line over the Temperature at 1m

series, marked by a black line. In the same Figure the seasonals from the SS model are also plotted and the two fits for the SS (red) and TS (blue) models.

Table 6.5: Temperature 1m AIC & BIC - Cyclic Component

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	14.594	116.347
AIC	2521.426	1118.033	1183.661
BIC	2534.221	1180.167	1679.026

Furthermore, the AIC and BIC scores in Table 6.5 verify the choice of a stable seasonal pattern.

In Figures 6.9 and 6.10 we see the residual plots for the SS and TS GAMs. They are similar to the ones produced when the stable seasonals were fitted by factors.

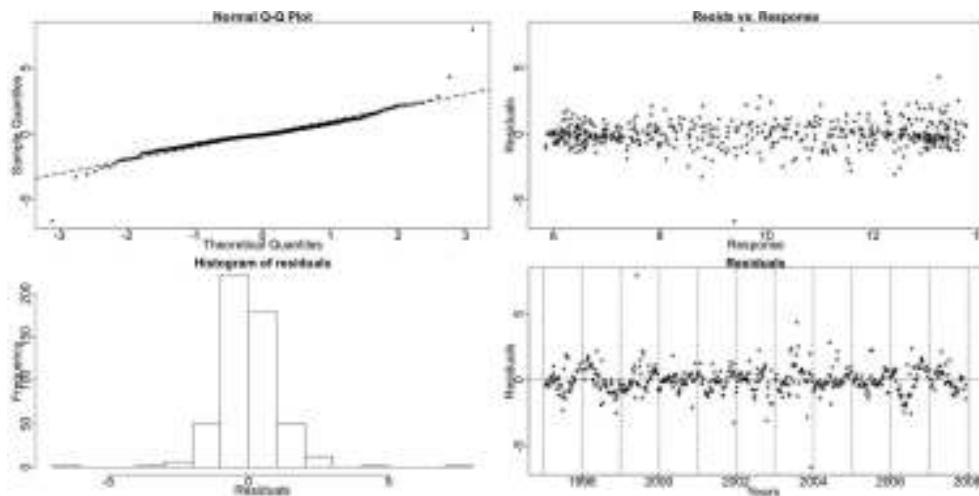


Figure 6.9: The residual checking plots for the SS GAM (using a cyclic component) of the Temperature 1m data. No problems are highlighted.

The two methods for modelling the stable seasonal component in the GAMs produce similar results that lead us to choose the SS GAM.

The results for the Temperature series at both 1m and 45m are similar and have shown that these series have stable seasonal patterns in terms of systematic shifts in

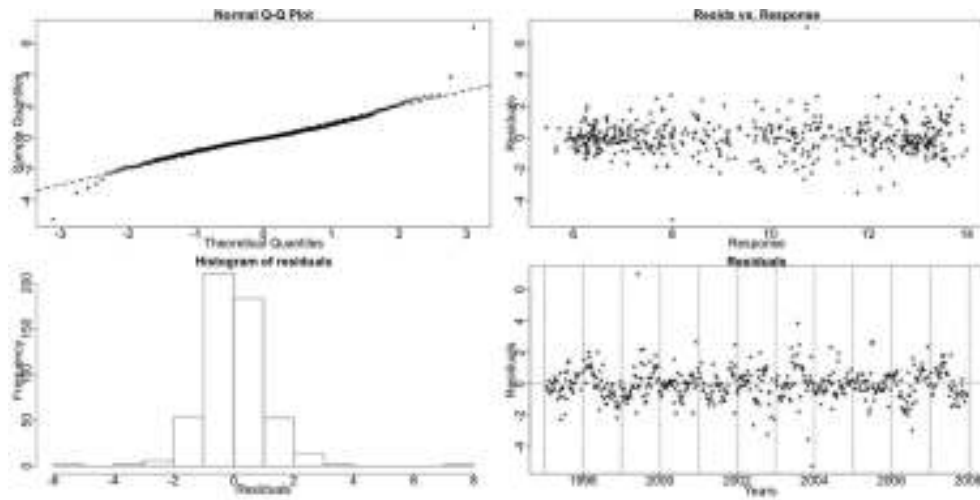


Figure 6.10: Residual checking plots for the TS GAM (using a cyclic component) of the of the Temperature 1m data. No violations are noted.

time. In particular, both patterns are pseudo-trending in amplitude. The existence of an underlying gradual rise in annual average temperatures (increasing trend) does not affect our results and reflects the general patterns seen in the wider northeast Atlantic Ocean of which the North Sea is an attached region (ICES, 2009).

6.3.2 Salinity

Salinity levels are affected by in-flowing oceanic water, freshwater run-off and local evaporation/ precipitation changes according to IACMST (2001). Therefore, salinity data are generally more variable than temperature data. The salinity levels at 1m are also easily affected by rainfall and river inflow since the fresh water is input to the sea surface layers. The salinity levels at 45m, on the contrary, are less affected by rapid fluctuations caused by freshwater influences and are thus expected to be less noisy.

Oceanic water has generally a higher salinity than coastal water, thus water with a high oceanic content that is advected into the sampled area will be indicated in the data by an increase in salinity values. Both salinity levels, however, are generally higher in late summer when the influx of coastal water, often mixed with increased oceanic water, penetrates down along the Scottish east coast (Hay, pers. comm., 2009).

The GAMs constructed for the two Salinity series use an additive structure, as in (6.1), using the Gaussian family with the identity link function. The seasonal component and the residuals are not affected by changes in the level of the trend. Hence, all three components are added to create the series.

Salinity - 1m

The Resampling test for the Salinity at 1m (Salinity 1m) series identifies a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.340 \rightarrow$ Stable

The GCV score of the three GAMs also points to that conclusion:

GCV Score: Non-Seas: 0.063, **Stable:** 0.053 & Trending: 0.061

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 14.6, Stable: 46.9 & Trending: 60.3

This series is noisy, hence the percentage explained deviance by the GAMs is small compared to the percentage deviance explained for the Temperature series.

Table 6.6: Salinity 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	501.482	30.564				
SS	449.529	19.001	51.954	11.563	5.265	< 2.2e-16
TS	392.005	14.199	57.523	4.802	2.305	1.540e-06

The ANOVA results, given in Table 6.6, however, favour the TS GAM.




Figure 6.11: The upper left plot depicts Salinity at 1m with the SS model's trend as a red line. The stable seasonals and the Salinity 1m marked by dots with the stable (red) and trending (blue) fitted GAMs, follow clockwise.

The two fits (SS and TS) of the GAMs for Salinity 1m are shown in the lower graph in Figure 6.11. The TS fit is better than the SS one as it goes higher into the peaks and lowers into the troughs of the series. The trend and stable seasonals from the SS GAM are depicted in the same picture.




Figure 6.12: The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Salinity 1m data.

In Figures 6.12 and 6.13 we can see the fifty two smooth functions fitted by the TS GAM for the trending seasonals. Many weeks deviate from a straight line parallel to the x-axis. The majority of them, however, is not monotonically increasing or decreasing indicating a random change. Nonetheless, there a few weeks, 6 and 19-22 that point to a systematic change.

There is the possibility of a systematic shift in time, that is being initialised towards the end of the series and thus, it cannot easily be detected.

Table 6.7: Salinity 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	8.518	60.471	117.995
AIC	29.895	-108.142	-141.386
BIC	66.227	147.8	358.02

The AIC, Table 6.7, favours the TS GAM while the BIC, in the same table, chooses the NS model. The AIC, similarly to the ANOVA, is more sensitive and points to the TS GAM.




Figure 6.13: The smooth trending seasonals estimated in the TS model for the Salinity 1m series, here weeks 37-52.




Figure 6.14: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Salinity at 1m series.

To further investigate whether there is a systematic change in the timing of the seasonal pattern and identify its direction we plot the stable seasonals and the trending seasonals, as estimated by the TS GAM, on the same plot, Figure 6.14. We note a random amplitude change with the amplitude being bigger towards the start and end of the examined series. There is also a slight shift of the seasonals earlier over the years. This, however, is hard to quantify as there is a change in the size and shape of the seasonals for these years. There appears to be a transition from one seasonal pattern to another one with a slightly different from.

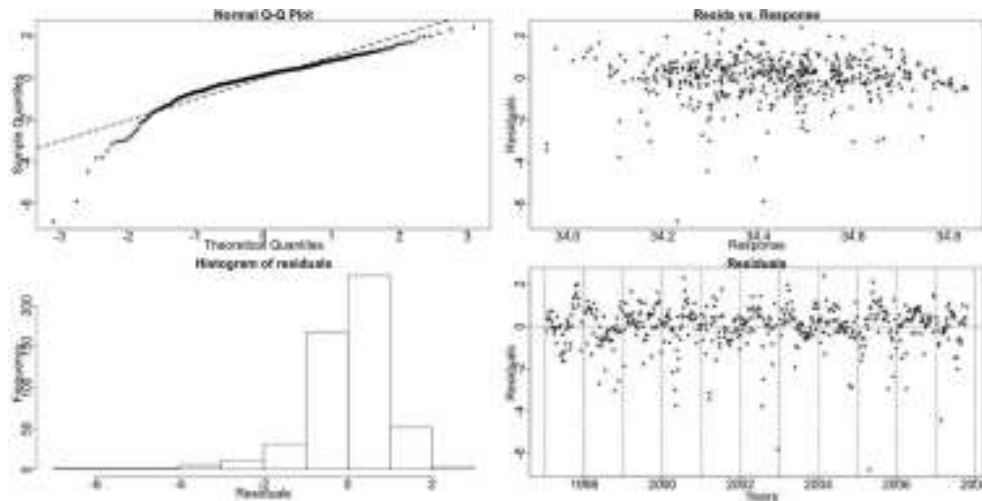


Figure 6.15: The residual checking plots for the SS GAM of Salinity 1m show a deviation from normality.

In Figures 6.15 and 6.16 we see some diagnostic plots for the SS and TS GAMs. The QQ plots and histograms show a deviation from normality near the extremes of the series. These models are not well fitted to the data but were found the most appropriate ones from those attempted.

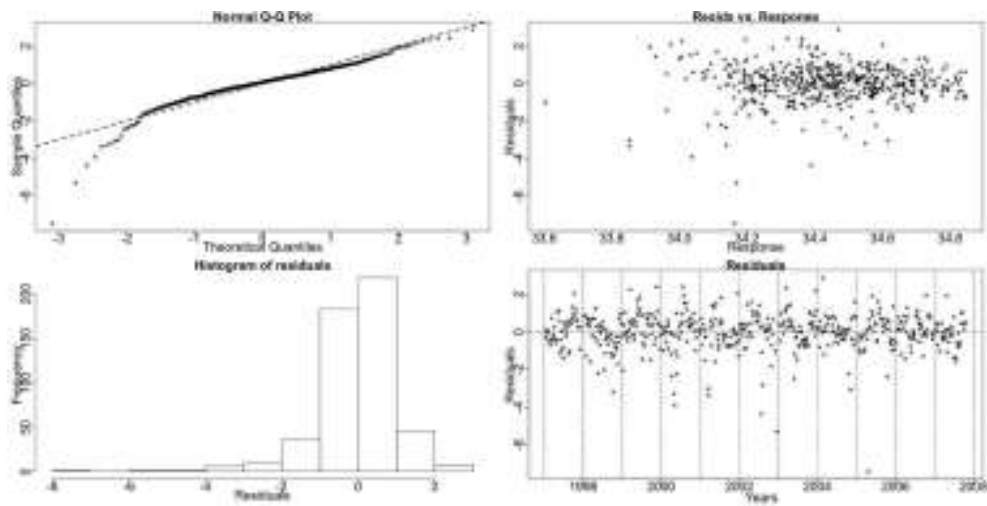


Figure 6.16: Plots of the standardized deviance residuals from the TS GAM of Salinity 1m show a deviation from normality.

Salinity - 45m

As mentioned above, Salinity at 45m is not as easily affected by rainfall and river input as the surface Salinity and thus, the series is more well behaved. The Resampling test finds this series to have a trending seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.018 \rightarrow$ Trending

The GCV score, however, of the SS and TS GAMs is almost the same:

GCV Score: Non-Seas: 0.027, **Stable:** 0.017 & **Trending:** 0.018

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 27, Stable: 66.5 & Trending: 76.1

The GAMs are able to explain a greater portion of the variability in this series than they were for the Salinity at 1m.

Table 6.8: Salinity 45m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	498.66	13.04				
SS	447.537	5.978	51.123	7.062	10.342	< 2.2e-16
TS	390.533	4.265	57.004	1.713	2.752	4.824e-09

In Table 6.8 the ANOVA favours the TS GAM.

The Salinity 45m series along with the fits of the SS (red) and TS (blue) GAMs are shown in Figure 6.17. The terms of the SS GAM are also plotted in the same picture.

In Figures 6.18 and 6.19 we can see the fifty two smooth functions for the trending seasonals in the TS model. A number of weeks (2, 6, 20-26, 36, 39 and 49) deviate from a straight line parallel to the x-axis. Furthermore, a monotonic curve like that for weeks 20-26 indicates a systematic change.

The AIC score, Table 6.9, is lower for the TS GAM, thus preferring it. In the same table, the BIC points to the SS GAM.




Figure 6.17: Clockwise from top left the Salinity 45m series with the trend from the SS GAM (red line), the stable seasonals and the Salinity 45m series (dotted) with the stable (red) and trending (blue) fitted GAMs.

Table 6.9: Salinity 45m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	10.34	13.943	72.595
AIC	-398.205	-692.167	-749.727
BIC	-354.101	-432.148	-248.555

In Figure 6.20 we see the stable (red) and trending (blue) seasonals from the TS GAM. The amplitude of the seasonals fluctuates in the same way as for the seasonals in the Salinity at 1m series. The amplitude in the first three and the last two observed years is greater than for the rest of the years. Furthermore, we notice that the trending seasonal peak at the beginning of the series occurs just after the stable seasonal one and gradually shifts earlier in the year to occur just before the stable peak towards the end of the series. Finally, for week 1 the stable seasonals attribute a low value while the trending ones a high one.

To further observe the changes in the seasonal pattern of Salinity 45m we choose to examine the de-trended fitted values of the TS GAM rather than the original series since the latter is very noisy. In Figure 6.21 we see the de-trended fitted values from




Figure 6.18: The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Salinity 45m data.




Figure 6.19: The smooth trending seasonals estimated in the TS model for the Salinity 45m, weeks 37-52.




Figure 6.20: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Salinity at 45m series.

the TS GAM for Salinity 45m. The trough is marked by dotted lines at weeks 10 and 30. The peak is shifting earlier in the year while the size and shape of the seasonals also seems to be changing.

Figures 6.22 and 6.23 show some diagnostic plots for the SS and TS GAMs. There are no problems with the models.

Salinity at 1m is more easily affected by river input and rain and can thus, change randomly in size and timing. Hence, it is easier to detect a systematic change when looking at a depth of 45m than the surface, as the results have demonstrated. The Resampling results for Salinity 1m can have been affected by the fluctuation in the size of the seasonals and thus, have greater difficulty in detecting a systematic shift. Nonetheless, the GAMs for Salinity at 1m identified the changes in the seasonals. Both Salinity series have patterns with fluctuating amplitude that are also shifting earlier in the year.




Figure 6.21: The de-trended fitted values from TS GAM for Salinity 45m. The vertical lines mark weeks 10 and 30. The peak appears to shift earlier in the year.

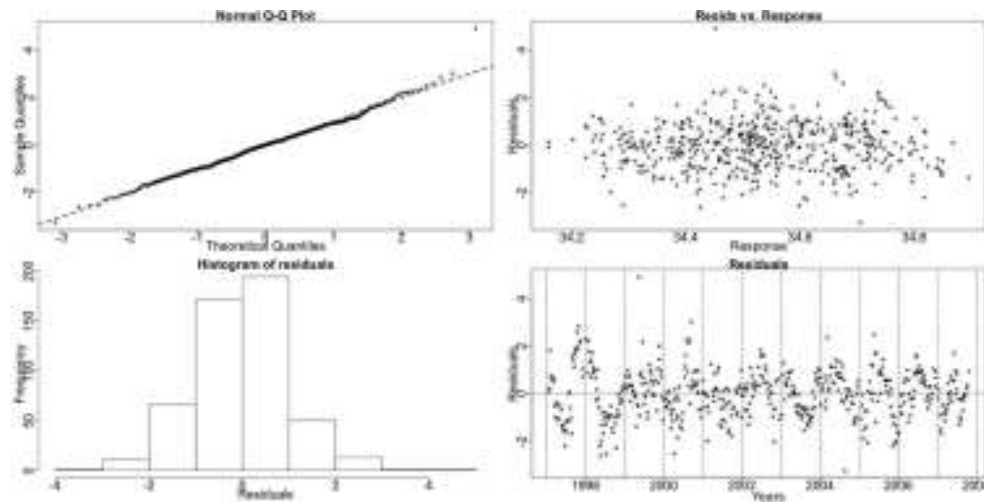


Figure 6.22: The residual checking plots for the SS GAM of Salinity 45m show no violations.

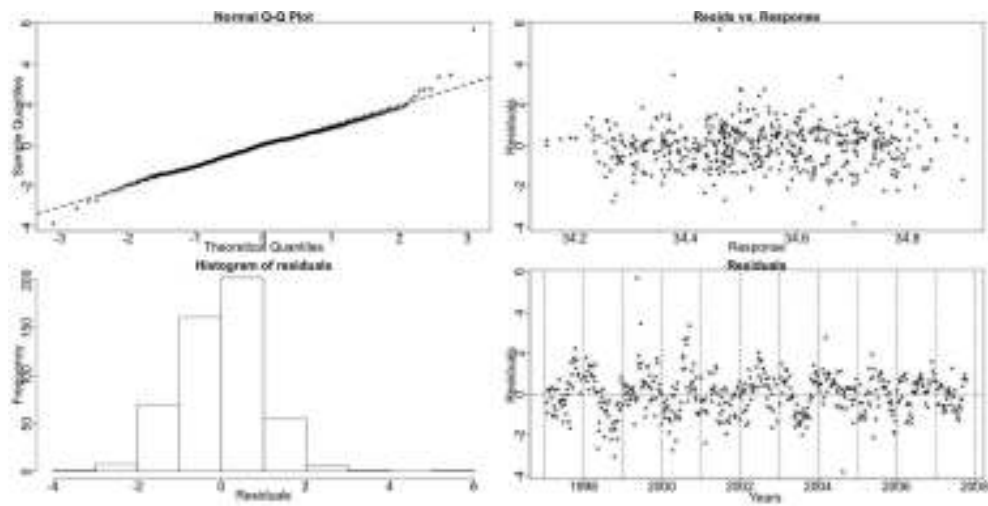


Figure 6.23: The standardized deviance residuals' checking plots from the TS GAM of Salinity 45m show no violations of the assumptions of the model.

6.3.3 Silicate

Silicate, phosphate and nitrate are the three key nutrients used by phytoplankton in their growth that we are examining from the Stonehaven database. Parsons *et al.* (1988) say that silicate is usually higher in river water than in sea water. Thus, silicate builds up during winter time from river inputs and sea sediments. It is consumed mainly by diatoms to create their ‘glass’ exoskeleton. The levels of the surface silicate are more easily affected by extreme events in external conditions such as floods, increased run-off from the land etc.; thus we expect the surface series to be more noisy than the 45m one.

Both Silicate series exhibit a winter peak. Because the weather in winter is poorer than other seasons the sampling during the winter weeks is more often prohibited by the weather conditions than for other weeks in the year. This means that we may not have as much information about the peak of these series as we have for other series that peak in the summer, spring or autumn. From an ecological perspective, however, the spring time when silicate starts being consumed and the autumn when silicate begins to replenish are more critical periods in the annual cycle of these series (Hay, pers. comm., 2009).

Silicate - 1m

The Silicate at 1m (Silicate 1m) data have a stable seasonal pattern according to the Resampling test results.

Resampling Testing Stable vs Trending: $p = 0.191 \rightarrow$ Stable

For this series we use the log link with the Gaussian family to build the GAMs. In Figure 6.24 we see two plots. The top plot shows the mean value of the series for each year plotted against the range of the series for that year. We see a positive connection between them which is further verified by the correlation, 0.442. Thus, the amplitude of the seasonals changes with changes in the level of the trend. It

increases when the trend increases and decreases when the decreases. Hence these two components, trend and seasonal, are multiplicative.

The lower plot in Figure 6.24 shows the fitted values from the SS GAM using a log link and a Gaussian family plotted against the absolute value of the Pearson residuals from that model. We fit a super-smoother, marked by the red line, to check for any relationship between them. The produced line is almost parallel to the x-axis indicating that the residuals are not affected by changes in the trend and thus, should be added to the other two components and not multiplied. Therefore, the log link is appropriate.




Figure 6.24: The top plot shows the mean value of the Silicate at 1m series for each year against the range of the series for that year. The bottom plot shows the fitted values from the SS GAM plotted against the absolute residuals from that model with a smoother running through them, red line.

The GCV score agrees with the Resampling results and also points to the SS GAM:

GCV Score: Non-Seas: 4.126, **Stable:** 1.374 & Trending: 1.410

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 8.67, Stable: 77.5 & Trending: 84

The ANOVA in Table 6.10 favours the TS GAM.

Table 6.10: Silicate 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	558.41	2270.386				
SS	502.627	558.739	55.775	1711.552	27.605	< 2.2e-16
TS	445.331	397.104	57.296	161.635	3.164	8.225e-12




Figure 6.25: Clockwise from top left the Silicate 1m series with the trend from the SS GAM, the stable seasonals and Silicate 1m (dotted) with the stable (red) and trending (blue) fitted GAMs.

The Silicate 1m series along with the fits of the SS (red) and TS (blue) GAMs are shown in Figure 6.25. The trend and seasonals of the SS GAM are also plotted in the same picture.

In Figures 6.26 and 6.27 we can see the fifty two smooth functions for the trending seasonals in the TS model. Weeks 18 and 33 deviate from a straight line. These type of curves though correspond to a random and not a systematic type of change.

Table 6.11 depicts the AIC and BIC scores for the three GAMs. The AIC points to the TS GAM while the BIC the SS model.

This series has a pseudo-trending in amplitude seasonal pattern. This fluctuation in the size is probably the reason that the AIC and the ANOVA point to the TS

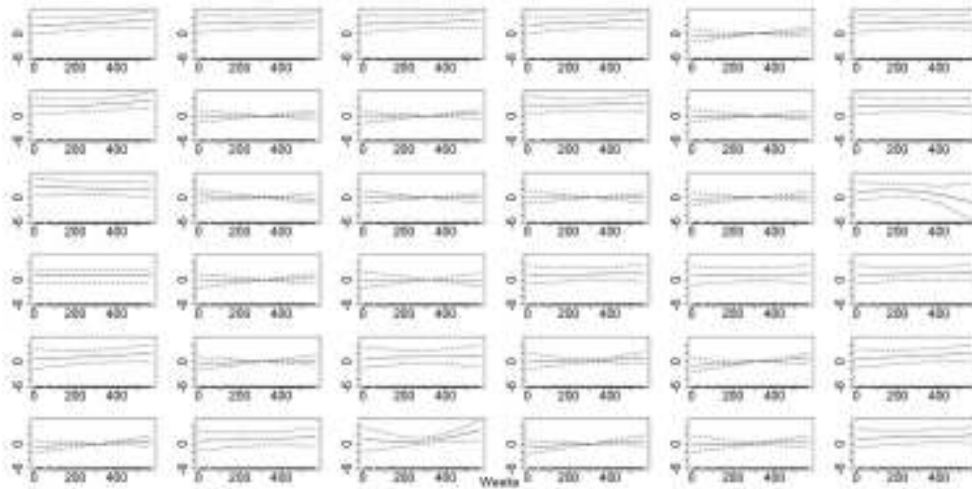


Figure 6.26: The smooth trending seasonals for weeks 1-36, estimated in the TS model for the Silicate 1m data.




Figure 6.27: The smooth trending seasonals estimated in the TS model for the Silicate 1m, weeks 37-52.

Table 6.11: Silicate 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	5.598	61.373	118.669
AIC	2393.966	1716.193	1638.53
BIC	2418.536	1982.139	2152.755

GAM. Nevertheless, this is considered a stable seasonal pattern in the context of this study.

Figures 6.28 and 6.29 show the diagnostic plots for the SS and TS GAMs which seem fine.

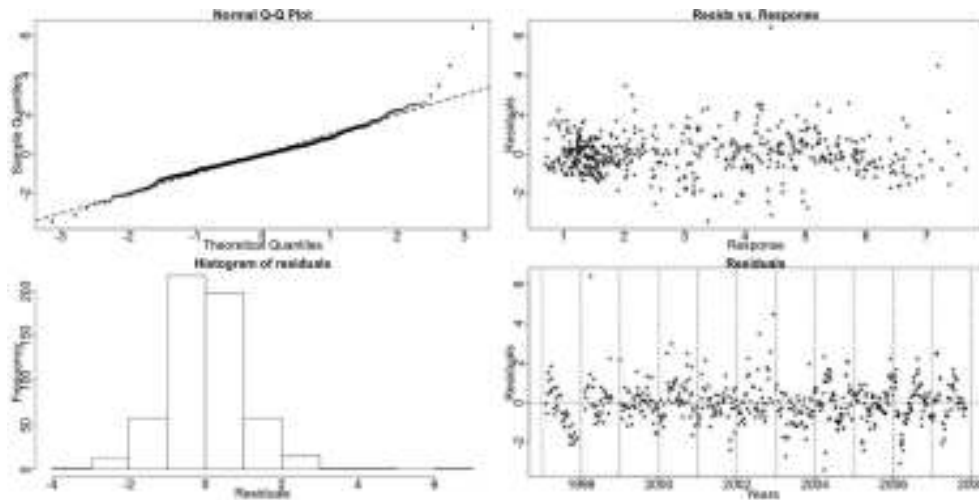


Figure 6.28: The residual checking plots for the SS GAM of Silicate 1m show no violations.

Both Silicate measurements appear to have pseudo-trending in amplitude seasonal patterns. As this chemical is a strongly limiting factor in the growth of diatoms we might expect to see similar patterns with these when we examine the phytoplankton series, as is verified later.

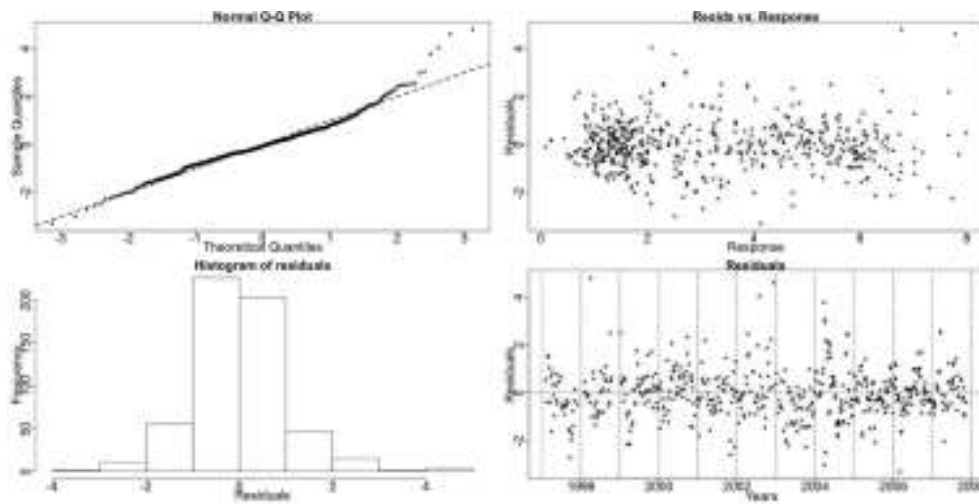


Figure 6.29: The standardized deviance residuals' checking plots from the TS GAM of Silicate 1m show no violations of the assumptions of the model.

6.3.4 Phosphate

As with other nutrients the phosphate concentration is increasing during winter to be consumed by phytoplankton during the rest of the year. Parsons *et al.* (1988) point out that the phosphates concentration may in fact be below the detection limit after a phytoplankton bloom. Hence, changes in phosphate concentration are closely linked to the biological demands of the phytoplankton species. Nonetheless, phosphate is not usually a limiting nutrient for phytoplankton growth in the sea; whereas in freshwater systems it is. In the following series the majority of 2006 data had to be removed due to a contamination of the samples. Missing values were introduced.

Both Phosphate series are modelled using a Gaussian distribution with the identity link function.

Phosphate - 1m

The Resampling results for Phosphate at 1m (Phosphate 1m) indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.316 \rightarrow$ Stable

The GCV scores of the three GAMs, also, point to the SS model:

GCV Score: Non-Seas: 0.037, **Stable:** 0.022 & Trending: 0.024

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 17.7, Stable: 60.8 & Trending: 68.5

Table 6.12: Phosphate 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	642.222	23.201				
SS	591.532	11.055	50.689	12.147	12.823	< 2.2e-16
TS	536.782	8.872	54.751	2.183	2.412	2.978e-07

The ANOVA results, presented in Table 6.12, favour the TS GAM, however.




Figure 6.30: The Phosphate 1m series is plotted in the top left with the trend from the SS GAM (red line). The stable seasonals and the Phosphate 1m (dotted) with the stable (red) and trending (blue) fitted GAMs follow clockwise.

In Figure 6.30 the series of Phosphate 1m with the fits of the SS (red) and TS (blue) GAMs are plotted. The trend and seasonals of the SS GAM are also plotted in the same picture. In addition, in this plot we note that years 2003 and 2004 behave differently to the previous years.

Table 6.13: Phosphate 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	10.778	61.468	116.218
AIC	-303.114	-685.1	-719.032
BIC	-254.826	-409.722	-198.369

The AIC and BIC scores in Table 6.13 point to the TS and SS GAMs, respectively.

Examining further the series, see Figure 6.31, we notice that the size of the seasonals fluctuates while the series has a stable seasonal pattern in terms of systematic




Figure 6.31: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Phosphate at 1m series.

changes. Additionally, from 2002 onwards the main peak is split into two peaks, thus changing the shape of the seasonals.

Figures 6.32 and 6.33 show some diagnostic plots for the SS and TS GAMs. There are no major problems with the models even though the upper tail of the distribution of the residuals deviates from the Gaussian family.

Our analyses do not show any systematic changes in the patterns of both phosphate measurements. This is partly expected as phosphate is seldom a limiting nutrient hence an increase or decrease in supply and demand may be less evident for phosphate than for example Nitrate which is a limiting nutrient in the sea.

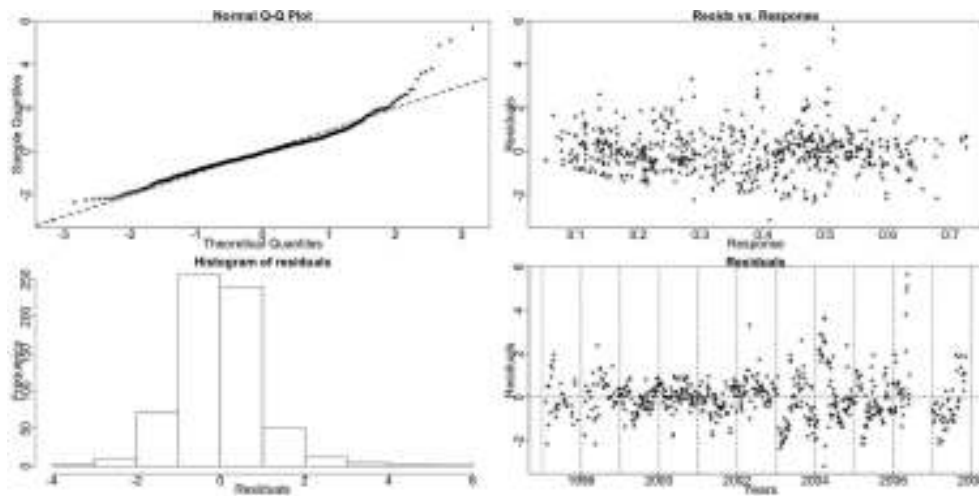


Figure 6.32: The residual checking plots for the SS GAM of Phosphate 1m show no violations.




Figure 6.33: The standardized deviance residuals' checking plots from the TS GAM of Phosphate 1m show no violations of the assumptions of the model.

6.3.5 Ammonia

Ammonia in the water is derived as an excretion product of the animal zooplankton as well as being produced during their death and decay. Ammonia is also rapidly taken up by phytoplankton. The amount produced by the zooplankton in the water may be rapidly consumed by the phytoplankton. Thus a low value of ammonia may reflect a small number of zooplankton present or a rapid consumption rate and should be carefully interpreted. Generally, the level of ammonia increases during the months of high zooplankton abundance. Rapid changes in ammonia levels are hard to interpret. Ammonia samples can easily be contaminated during collection and handling (Hay, pers. comm., 2009).

The GAMs for both Ammonia series are constructed using an additive structure (identity link function) as no evidence to the contrary were found.

Ammonia - 1m

The Resampling results for Ammonia at 1m (Ammonia 1m) indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.285 \rightarrow$ Stable

The GCV scores disagree with the choice of the SS model:

GCV Score: Non-Seas: 0.494, Stable: 0.464 & **Trending:** 0.450

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 4.3, Stable: 27.9 & Trending: 40.9

Table 6.14: Ammonia 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	692.948	336.992				
SS	641.501	253.957	51.448	83.035	4.077	< 2.2e-16
TS	585.16	208.242	56.341	45.715	2.28	1.135e-06

In addition, the ANOVA , in Table 6.14, is in favour of the TS one.




Figure 6.34: Clockwise from top left the Ammonia 1m series with the trend from the SS GAM marked by a red line, the stable seasonals and Ammonia 1m (dotted) with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.34 the series of Ammonia 1m with the fits of the SS (red) and TS (blue) GAMs are plotted. The trend and seasonals of the SS GAM are also plotted.

Figures 6.35 and 6.36 show the smooth functions fitted to each week for the trending seasonals by the TS GAM. Weeks 2, 19, 35, 37, 40, 41 and 42 differ from a straight line parallel to the x-axis. This indicates that the smooths for these weeks are significant and should be included in the model.

Table 6.15: Ammonia 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	7.052	58.499	114.84
AIC	1487.794	1392.946	1366.903
BIC	1519.876	1659.098	1889.385

The AIC and BIC scores in Table 6.15 point to the TS and NS GAMs, respectively.




Figure 6.35: The smooth fitted functions for weeks 1-36, in the TS model for the Ammonia 1m data.




Figure 6.36: The smooth trending seasonals estimated in the TS model for the Ammonia 1m, weeks 37-52.




Figure 6.37: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Ammonia at 1m series.

In Figure 6.37 the stable (red) and trending (blue) seasonals from the TS GAM are plotted over each other. There is a fluctuation in the size of the seasonals of this series. This, however, could hinder the Resampling test's ability to identify a shift in time. There is probably a systematic shift in time that is very small and thus, easily disguised by the random amplitude fluctuations.

In Figures 6.38 and 6.39 the diagnostic plots for the SS and TS GAMs show a deviation from normality.

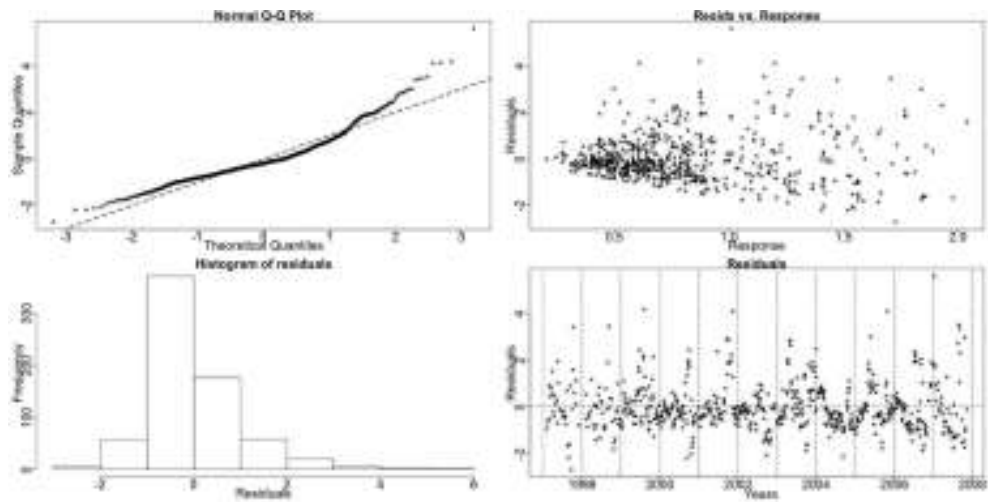


Figure 6.38: The residual checking plots for the SS GAM of Ammonia 1m.

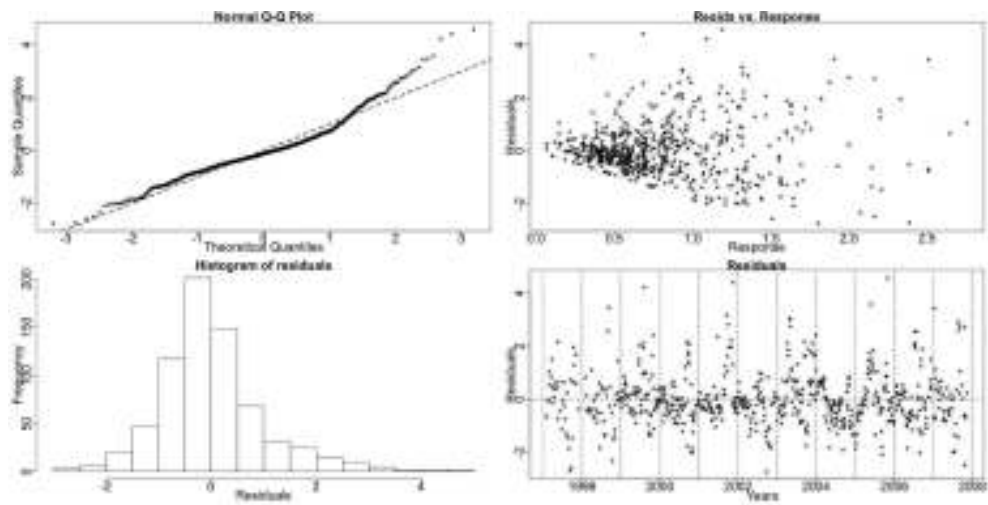


Figure 6.39: The residual checking plots from the TS GAM of Ammonia 1m.

Ammonia - 45m

The Resampling results for Ammonia at 45m (Ammonia 45m) indicate a trending seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.019 \rightarrow$ Trending

The GCV score of the three GAMs, however, points to the SS model:

GCV Score: Non-Seas: 0.486, **Stable:** 0.436 & Trending: 0.452

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 5.04, Stable: 32 & Trending: 46.5

Table 6.16: Ammonia 45m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	687.936	328.899				
SS	635.569	235.522	52.367	93.377	4.812	< 2.2e-16
TS	579.177	185.188	56.392	50.334	2.792	8e-10

The Resampling result is further enforced by the ANOVA, in Table 6.16, which is in favour of the TS GAM, too.

In Figure 6.40 we see the series of Ammonia 45m with the trend from the SS GAM. The stable seasonals and the fitted values of the SS (red) and TS (blue) GAMs are also plotted.

Figures 6.41 and 6.42 show the smooth functions fitted to each week for the trending seasonals by the TS GAM. Weeks 22, 24, 25, 29, 31, 35, 37, 39, 40, 41, 42 and 50 deviate significantly from a straight line parallel to the x-axis.

Table 6.17: Ammonia 45m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	7.064	59.431	115.823
AIC	1465.391	1338.366	1284.29
BIC	1497.78	1608.33	1810.411




Figure 6.40: Clockwise from top left the Ammonia 45m series with the trend from the SS GAM, the stable seasonals and Ammonia 45m (dotted) with the stable (red) and trending (blue) fitted GAMs.

The AIC and BIC scores in Table 6.17 point to the TS and NS GAMs, respectively.

In Figure 6.43 we see the stable (red) and trending (blue) seasonals from the TS GAM. The peak shifts systematically earlier in time while there is also a random amplitude change.

Similarly to Salinity, finding Ammonia at 45m to have a trending seasonal pattern verifies our thoughts about a systematic shift masked by the fluctuation in the size of the pattern for the Ammonia at 1m series. Ammonia at 1m could be trending but the series maybe too noisy for it to be detected. Weeks 40, 41 and 42 were found trending for the Ammonia at 1m as well as for the Ammonia at 45m.

To estimate the shift in time we plot the de-trended fitted values from the TS of this series and then note the occurrence of the main peak, see Figure 6.44. With a black dotted line we mark week 20 and with a red line week 50 in all years. The main peak seems to be shifting earlier in each year.

Figures 6.45 and 6.46 show some diagnostic plots for the SS and TS GAMs. No




Figure 6.41: The smooth fitted functions for weeks 1-36, in the TS model for the Ammonia 45m data.




Figure 6.42: The smooth trending seasonals estimated in the TS model for the Ammonia 45m, weeks 37-52.




Figure 6.43: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Ammonia at 45m series.




Figure 6.44: The de-trended fitted values from the TS GAM fitted to the Ammonia 45m series. The dotted lines mark weeks 20 and 50.

major issues appear with the models even though there is a deviation from normality especially at the upper tail of the series.

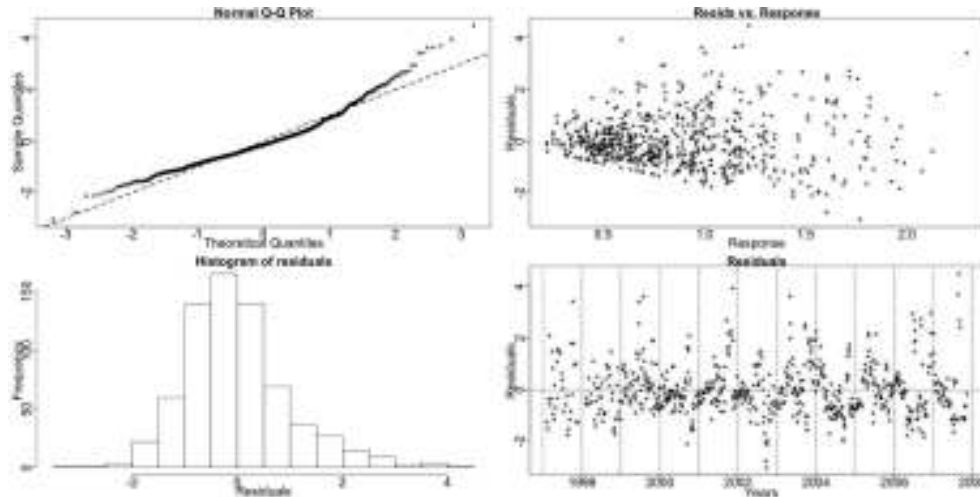


Figure 6.45: The residual checking plots for the SS GAM of Ammonia 45m.

There is evidence of a shift earlier in the year for the seasons of both Ammonia series. This bears a striking resemblance to that seen in the Salinity data. We perceive a shift through some transitory years from one state to another rather than an entirely smooth trend. This may be a reflection of the pattern of zooplankton abundance and decay as they are mainly responsible for the ammonia signal.

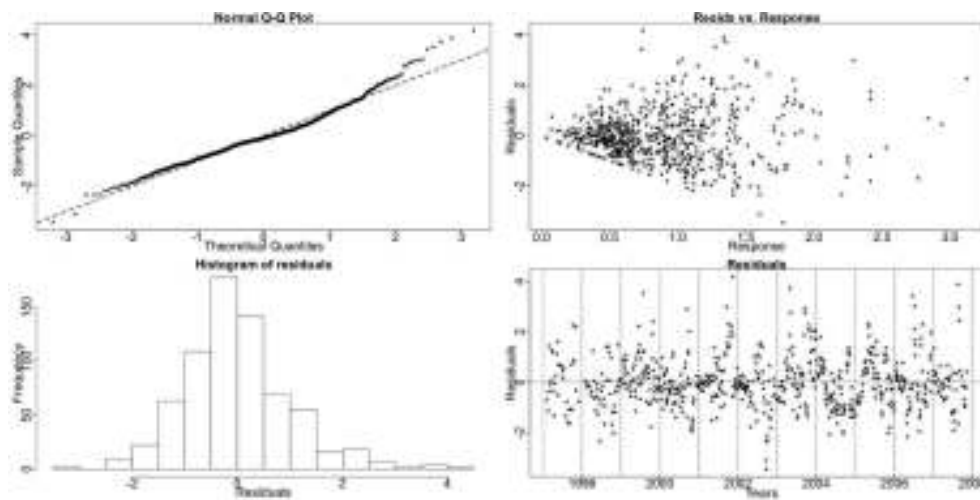


Figure 6.46: The standardized deviance residuals' checking plots from the TS GAM of Ammonia 45m.

6.3.6 Nitrate

Nitrate is the main nutrient in the water that drives the marine phytoplankton production and is quite often present in limiting concentrations. It has a seasonal nature with slow build up in winter months and rapid decline in spring as the light levels increase to allow the phytoplankton community to begin their photosynthesis. When there is sufficient light phytoplankton is reproduced and the levels of nitrate in the water decrease (Hay, pers. comm., 2009).

The GAMs for both Nitrate series examined below are based on an additive structure, using the identity link function, and a Gaussian family.

Nitrate - 1m

The Resampling results for Nitrate at 1m (Nitrate 1m) indicate a stable seasonal pattern. We do note that the p-value is quite small though.

Resampling Testing Stable vs Trending: $p = 0.075 \rightarrow$ Stable

The GCV score of the three GAMs rejects the choice of the SS model in favour of the TS one:

GCV Score: Non-Seas: 13.292, Stable: 3.130 & **Trending:** 3.04

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 2.98, Stable: 82.1 & Trending: 86.7

Table 6.18: Nitrate 1m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	699.902	9229.376				
SS	643.064	1706.362	56.838	7523.014	49.881	< 2.2e-16
TS	587.779	1267.405	55.285	438.956	3.682	1.68e-15

This is further enforced by the ANOVA, in Table 6.18, which is in favour of the TS GAM, too.




Figure 6.47: Clockwise from top left the Nitrate 1m series with the trend from the SS GAM, the stable seasonals and the Nitrate 1m series marked with dots with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.47 the series of Nitrate 1m with the fits of the SS (red) and TS (blue) GAMs are plotted. The trend and residuals of the SS GAM are also plotted. For years 2000 and 2001 neither fit goes low enough in the trough, both overestimating the levels of Nitrate.

Figures 6.48 and 6.49 depict the smooth functions fitted to each week by the TS GAM. The C.I.s of weeks 2, 11-18, 43 and 48-50 differ from a straight line parallel to the x-axis, marking them significant in our model. In particular, for weeks 12-17 the fitted smooth functions exhibit a monotonic behaviour indicating a trending seasonal pattern.

Table 6.19: Nitrate 1m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	4.098	60.936	116.221
AIC	3813.3	2740.293	2641.797
BIC	3831.967	3017.879	3171.226




Figure 6.48: The smooth fitted functions for weeks 1-36, in the TS model for the Nitrate 1m data.




Figure 6.49: The smooth trending seasonals estimated in the TS model for the Nitrate 1m, weeks 37-52.

The AIC and BIC scores in Table 6.19 point to the TS and SS GAMs, respectively.




Figure 6.50: The stable seasonals, red, and the trending seasonals, blue, from the TS GAM fitted to the Nitrate at 1m series.

The two seasonal patterns, stable and trending, estimated by the TS GAM are plotted in Figure 6.50. The peak of the trending seasonals, blue, at the start of the series appears to the right of the stable seasonals, red, and moves gradually earlier in the year. Thus, towards the end of the series the peak of the trending seasonals is located just to the left of the stable one. Furthermore, we can verify that there is a random amplitude change as the overall size of the estimated trending seasonals first decreases and then increases.

To further examine the shift, in Figure 6.51 we see the de-trended fitted values from the TS GAM for Nitrate at 1m plotted by year. Dotted vertical lines are drawn to mark weeks 20 (black) and 50 (red). The summer trough appears to shift gradually earlier in time.

Figures 6.52 and 6.53 show some diagnostic plots for the SS and TS GAMs. The




Figure 6.51: The de-trended fitted values from TS GAM for Nitrate 1m. The vertical lines mark weeks 20 and 50 in black and red, respectively.

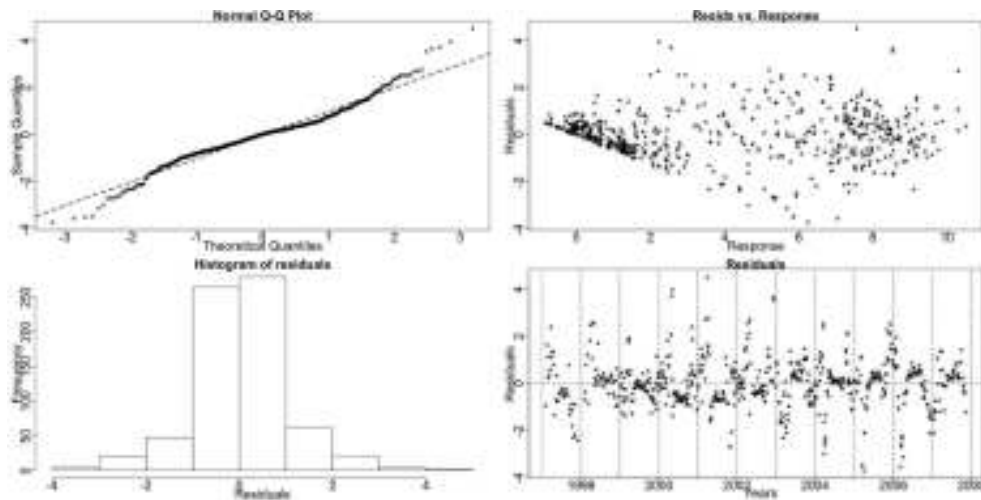


Figure 6.52: The residual checking plots for the SS GAM of Nitrate 1m.

assumption of normality is violated at the extremes of the series. In addition, there are some high residual values due to the model fits over-predicting or under-predicting the Nitrate levels.

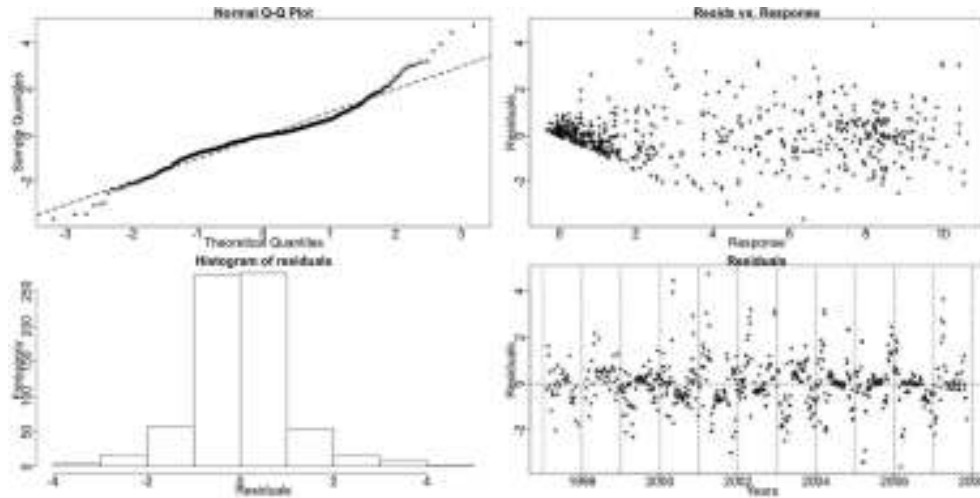


Figure 6.53: The standardized deviance residuals' checking plots from the TS GAM of Nitrate 1m.

The apparent shift earlier in the Nitrate seasonal pattern indicates that there may be evidence of earlier uptake by the phytoplankton and so a shift in the timing of the spring bloom. Nevertheless, the point in time that the Nitrate levels begin to replenish does not appear to be shifting earlier in the year as does the point of its decline. This could possibly translate to a longer period during which blooms take place.

6.3.7 *Chlorophyll a*

Smith *et al.* (2007) state that *Chlorophyll a* is the primary pigment of interest in monitoring programmes as it is the most abundant in photosynthesising plants and is also readily detected by fluorescence detection. Most of the phytoplankton species contain *Chlorophyll a*. The *Chlorophyll a* (*Chlorophyll*) series gives an indication of the phytoplankton biomass and not rate of production as for example there may be a huge population of zooplankton consuming it and thus very little biomass as *Chlorophyll* evident in the water samples (Bresnan, pers.comm., 2008).

The GAMs for this series assume a Gamma family with a log link and a cyclic component for the trending seasonals extraction. The diagnostic plots for the Gaussian family models were very poor indicating that a different family would be more appropriate. Though the diagnostic plots for the Gamma family are not ideal, Figures 6.59 and 6.60, are an improvement from the ones for the Gaussian family.

When using a cyclic component to model the stable seasonality there are two choices for the construction of the TS GAM. One can choose to include or exclude the term corresponding to the stable seasonal (i.e. the cyclic smoother) in the TS model. Including the stable seasonal term ensures that the models are properly nested and thus the ANOVA and the GCV score can be used for comparison. The goodness of fit measures are essentially the same on both occasions.

Nonetheless, the confidence intervals for the smooth terms corresponding to the trending seasonals are wider when the stable term is included. This is an indication that the contribution of the stable cyclic component affects the estimation of the standard error of the trending seasonals. No other differences occur and to better demonstrate the above issue, for this series we will present the results for both ways. For the rest of this thesis, when a series is modelled using a cyclic component for the stable seasonals both models are examined and the plots of the smooth functions are always from the TS GAM without the stable component.

The Resampling results for *Chlorophyll* indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.138 \rightarrow$ Stable

Including the Stable Seasonals' Term in the TS GAM

In agreement with the Resampling results the GCV score, also, prefers the SS model:

GCV Score: Non-Seas: 1.042, **Stable:** 0.424 & Trending: 0.532

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 3.32, Stable: 62.1 & Trending: 70.5

Table 6.20: *Chlorophyll* ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	689.929	709.733				
SS	680.915	278.136	9.015	431.599	117.211	< 2.2e-16
TS	578.751	216.898	102.164	61.239	1.599	4.887e-04

The ANOVA, in Table 6.20, is in favour of the TS GAM, however.




Figure 6.54: Clockwise from top left the *Chlorophyll* series with the trend from the SS GAM, the stable seasonals and *Chlorophyll* (dotted) with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.54 the series of *Chlorophyll* with the fits of the SS (red) and TS (blue) GAMs are plotted. The trend and seasonals of the SS GAM are also plotted.




Figure 6.55: The smooth fitted functions for weeks 1-36, in the TS model for the *Chlorophyll* data.

Figures 6.55 and 6.56 show the smooth functions fitted to each week for the trending seasonals by the TS GAM. Weeks 25, 35 and 36 differ significantly from a straight line parallel to the x-axis. The smooth functions corresponding to these weeks are not monotonic and thus they indicate a random change.

Table 6.21: *Chlorophyll* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	6.07	15.085	117.249
AIC	1810.663	1108.8	1130.208
BIC	1838.246	1177.345	1662.976

The AIC and BIC scores in Table 6.21 both point to the SS GAM.

Figure 6.57 shows the stable (red) and trending (blue) seasonals estimated by the SS and TS GAMs, respectively. It could be considered that there is a systematic amplitude change as the size of the pattern decreases until 2005 and then increases only slightly³. In addition, the shape of the seasonals changes. The sharp peak

³The actual ranges corresponding to years 1997-2008 are: 4.357, 3.921, 3.562, 3.285, 3.070, 2.910, 2.706, 2.631, 2.667, 2.776, 2.912.




Figure 6.56: The smooth trending seasonals estimated in the TS model for *Chlorophyll*, weeks 37-52.

during the first three years is joined by a second one later in the year, almost forming one wider peak. In the last two observed years, however, the shape of the pattern resembles the first years, with a narrower peak.

In Figure 6.58 we see the de-trended fitted values for *Chlorophyll* by the TS GAM. The vertical lines mark weeks 18 and 42 in black and red, respectively. The peak itself changes in size and shape while the timing of its descend appears to shift a little earlier in the year.

Figures 6.59 and 6.60 show some diagnostic plots for the SS and TS GAMs. There is a deviation from normality at the extremes of the series.




Figure 6.57: The stable seasonals, red, and the trending seasonals, blue, from the SS and TS GAM, respectively, fitted to the *Chlorophyll* series.




Figure 6.58: The de-trended fitted values from TS GAM for *Chlorophyll a*. Weeks 18 and 42 are marked by vertical lines in black and red, respectively.

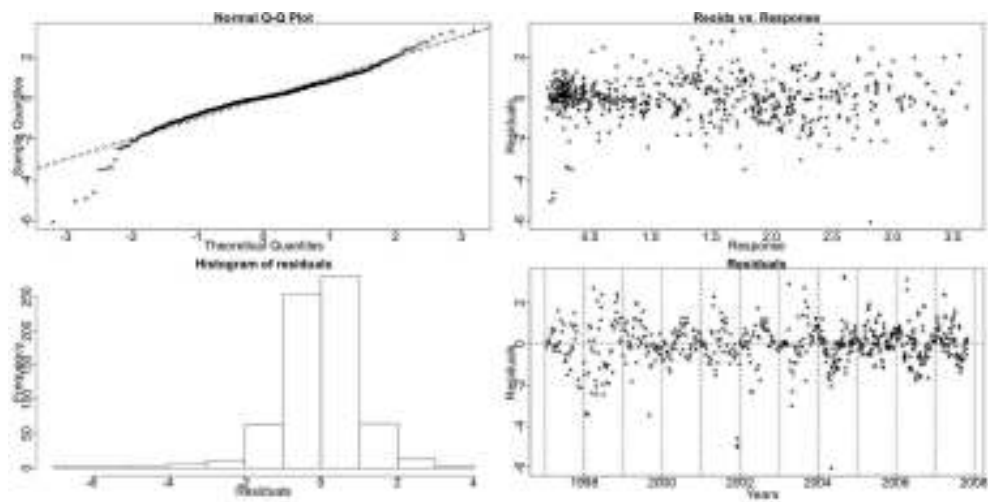


Figure 6.59: The residual checking plots for the SS GAM of *Chlorophyll* show a deviation from normality at the extremes of the series.

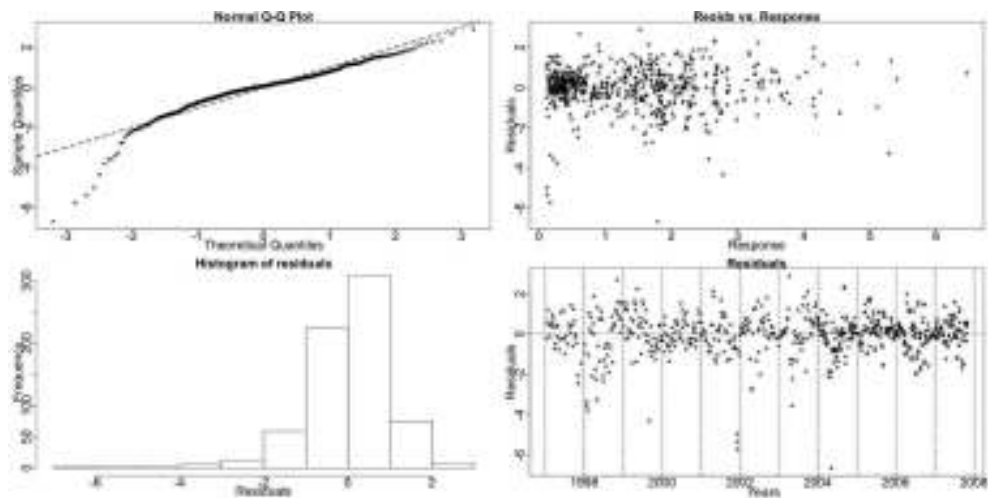


Figure 6.60: The standardized deviance residuals' checking plots from the TS GAM of *Chlorophyll* show a deviation from normality at the lower end of the series.

Without the Stable Seasonals' Term in the TS GAM

The GCV score of the new TS GAM is:

GCV Score: Trending: 0.532

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 3.32, Stable: 62.1 & Trending: 70.5




Figure 6.61: Clockwise from top left the *Chlorophyll* series with the trend from the SS GAM, the stable seasonals and *Chlorophyll* (dotted) with the SS (red) and TS (blue) fitted GAMs. The TS GAM does not include the stable cyclic smoother.

In Figure 6.61 the series of *Chlorophyll* with the fits of the SS (red) and TS (blue) GAMs are plotted. The trend and seasonals of the SS GAM are also plotted.

Figures 6.55 and 6.56 show the smooth functions fitted to each week for the trending seasonals by the TS GAM. Weeks 25, 35 and 36 differ significantly from a straight line parallel to the x-axis. The smooth functions corresponding to these weeks are not monotonic and thus they indicate a random change.

The AIC and BIC scores in Table 6.22 both point to the SS GAM.

The diagnostic plots for the TS GAM are shown in Figure 6.64. There is a deviation from normality at the extremes of the series.

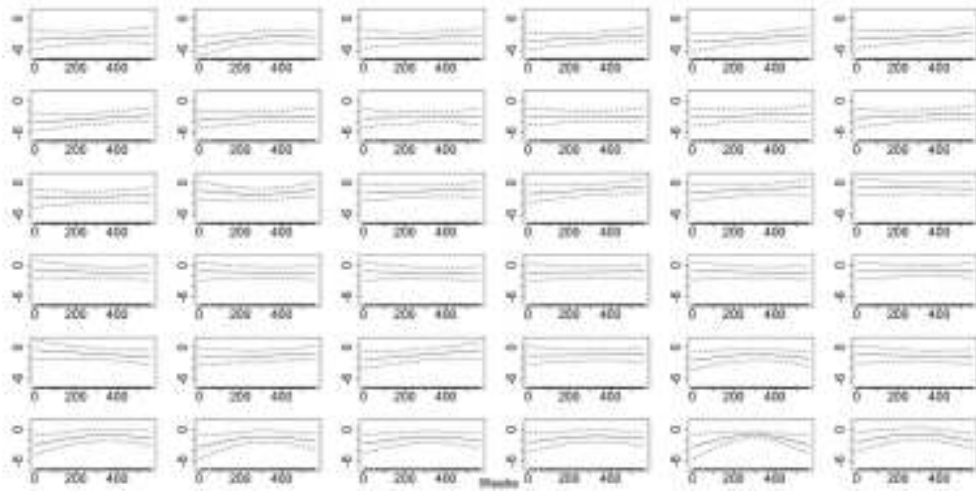


Figure 6.62: The smooth fitted functions for weeks 1-36, in the TS model (without the stable cyclic smoother) for the *Chlorophyll* data.

Table 6.22: *Chlorophyll* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	6.07	15.085	117.317
AIC	1810.663	1108.8	1130.374
BIC	1838.246	1177.345	1663.451

The analyses with either TS GAM points to the existence of trending seasonality. Furthermore, we note that the scores and test results from the goodness of fit measures that are used are very similar in both occasions. We prefer, however, to use the ANOVA and the GCV score for inference and thus we prefer the TS model that includes the stable term. When looking at the smooth functions fitted to each week the other TS GAM provides tighter confidence intervals and will thus be preferred for that particular purpose. For the rest of this thesis both models will be examined. The ANOVA results and the GCV scores comparisons will always correspond to the TS model that includes the stable component while when plots of the smooth functions are presented these will be from the TS GAM without the stable seasonal term. The AIC and BIC values are always very similar and will thus not be reported twice.




Figure 6.63: The smooth trending seasonals estimated in the TS model (without the stable cyclic smoother) for *Chlorophyll* , weeks 37-52.

In the *Chlorophyll* series analysis we note systematic changes in the pattern. There is evidence that the size of the pattern (amplitude) is gradually decreasing with time. In the first three years there seems to be a more sudden decline in spring. Ensuing years show a more stable pattern while in the last three a slight increase in the size is noted. There is also an indication of an earlier decline in the concentration of *Chlorophyll* through autumn.

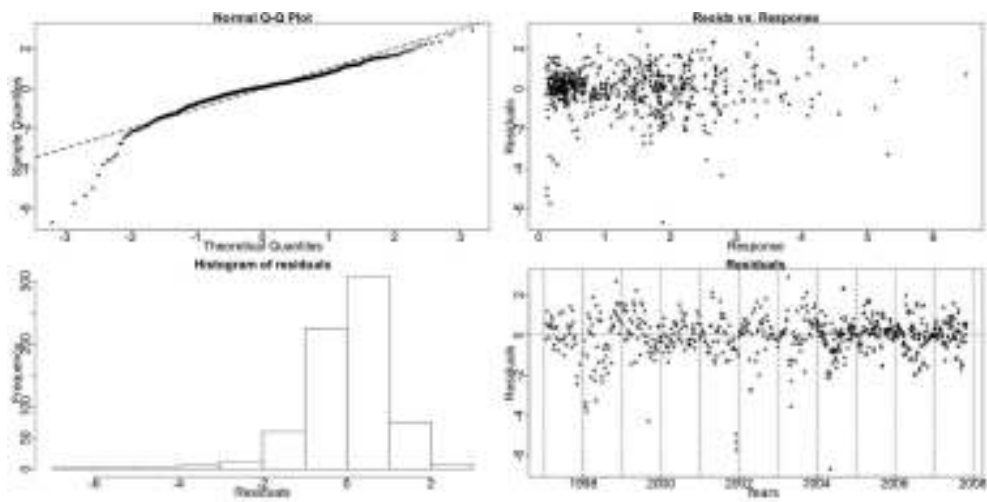


Figure 6.64: The standardized deviance residuals' checking plots from the TS GAM of *Chlorophyll* show a deviation from normality at the lower end of the series.

6.4 Phytoplankton Series Results

The phytoplankton data consist of diatoms and dinoflagellates. Bresnan *et al.* (2009) describe the procedure of collecting and analyzing phytoplankton samples. A 10m Lund tube is used and one litre of sea water is preserved. A 50ml subsample is analysed under an inverted microscope. After the homogeneity of the sample is verified all species present are recorded. The species in ten random fields of view are counted and multiplied up to give densities in one litre of water.

The phytoplankton series are densities derived from counts, and thus no malfunctioning equipment could interfere with the data. Phytoplankton species reproduce by cell division (< 2 days/division), therefore very rapid population growth is common when growth conditions are good. Extreme values are considered correct and are included in the analysis. Each species is generally adapted to grow optimally in the environmental conditions during a part of Nature's seasonal cycle. Because they are based on counts there are many zero values in the data. For some species there are more zero values than non-zero ones. Furthermore, commonly the sharp rises in the population of the phytoplankton create high spikes in the series.

When conducting the Resampling tests we use the logarithm of the examined series after we add 1 to it. We are thus, accommodating for the high variability in the spikes of the seasonal pattern. For Resampling tests we take the average of multiple entries in one pre-defined week to create a time series. This is not necessary for the GAMs.

Since we are analysing densities the Gamma and then the Normal families were assumed for our models. Two series, the Total Diatoms and Total Phytoplankton, included a small number of zeroes and thus we are able to fit the GAMs using a log Normal distribution with the identity link. Hence, we use the logarithm of the series after we add 1 to all values and fit using the Normal family.

Nevertheless, for the majority of the series of phytoplankton species and groups, the number of zero values is restrictive and the models cannot be fitted or when they

are fitted the fit is very poor. Therefore, we use Binomial models for most series, trying in this way to model the presence or absence of each species. All the non-zero values in a series are substituted by ones and then the models are fitted. We create the same three GAMs as before, NS, SS and TS, but now we assume the Binomial family with a logit link function, where:

$$\text{logit}(p) = \log\left(\frac{p}{1-p}\right). \quad (6.5)$$

Considering the UBRE, AIC and BIC score of the three models and the results from the ANOVA between the models we choose the most appropriate GAM. To evaluate the fit of the GAMs we cannot perform the usual checks when dealing with Binomial models as the asymptotic results supporting the usual Chi-Squared approximation involved in the likelihood-ratio tests is unreliable (Dobson, 2002). Instead, we compare the residual deviance of the model under examination, as fitted to the original series, with the residual deviances of the same model when fitted to 99 randomized series that have the same distribution as the original series.

We create 99 random uniform sets, U_t , equal to the length of the data and compare them to the fitted values, Y_t from the GAM in question. Thus, we create 99 series, F_t^i , $i = 1, \dots, 99$, as in (6.6) to which the model is fitted.

$$F_t^i = 0, \quad U_t > Y_t \quad \& \quad F_t^i = 1, \quad U_t < Y_t \quad (6.6)$$

Then we compare the deviance of the GAM fitted to the original series to the 99 deviances obtained by fitting the same GAM to the F_t^i series. If the GAM is a good fit then its deviance should belong to the distribution of the other 99 deviances.

The seasonal patterns of the series fitted by Binomial GAMs have sudden sharp peaks and thus, we believe that a trend in the seasonals would be traceable looking at the presence or absence of the species.

The diagnostic check for the log Normal models is done by examining the standardized deviance residuals. As in the previous section we create a QQ plot, a plot of the residuals against the response, the residual histogram and a residual plot. When referring to ‘residuals’ below in context of diagnostic plots we mean the previously defined standardized deviance residuals.

The Resampling tests are always performed on the densities of the series. All the series in this section were found seasonal when testing for seasonality with Resampling. Therefore, only the Resampling test results between stable and trending seasonality are reported.

All phytoplankton series analysed are found to have seasonal patterns that are stable in terms of systematic shifts. Only selected results are described below as many series exhibit similar behaviour. A full description of the results for all series examined can be found in the electronic Appendix I at the end of this thesis. First the results for the diatoms are presented and then for the dinoflagellates while lastly we include the results for the Total Phytoplankton series which is an aggregate of all observed species.

6.4.1 Diatoms

Diatoms are photosynthesising algae that have a siliceous skeleton. They need sunlight and chemical nutrients for growth thus, are found closer to the surface (University College London, 2009). Furthermore, diatoms need silicate for their silicate outer wall. This nutrient in particular is important for diatom growth and can limit the amount of diatoms present in the water.

Chaetoceros

The *Chaetoceros* series includes the total of all observed *Chaetoceros* species as it is hard to accurately distinguish between them. They are most active in the upper depths that receive sunlight and are prey to many zooplankton species such as copepods (Bresnan, pers.comm., 2008).

This is an aggregate series of several species of different sizes that may exhibit peaks at different periods (i.e. have slightly different seasonal cycles). Additionally, due to the fast reproductive and adaptive abilities of these species we might not be able to perceive a seasonal change without examining individual species time series. It is also possible that existing species may be supplemented or replaced by new ones carried into the region by advection from adjacent regions. The above is true for all series examined in this section, as identifying individual species is not always feasible and thus we look at aggregates. Nonetheless, we believe that these analyses, limited as they are, constitute a first attempt at exploring these species' seasonal patterns.

The Resampling results for *Chaetoceros* indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.144 \rightarrow$ Stable

Additionally, the UBRE score prefers the SS model.:

UBRE Score: Non-Seas: 0.157, **Stable:** -0.067 & Trending: 0.111

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 1.15, Stable: 24.5 & Trending: 42.6

Table 6.23: *Chaetoceros* ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	P(> Chi)
NS	416.193	476.899			
SS	409.491	364.318	6.701	112.590	1.707e-21
TS	349.680	276.858	59.81	87.440	0.01

The ANOVA results, however, presented in Table 6.23, favour the TS GAM.




Figure 6.65: Clockwise from top left *Chaetoceros* with the trend (red line) from the SS model, the stable seasonals and *Chaetoceros* (dotted) with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.65 the series of *Chaetoceros* with the fits of the SS (red) and TS (blue) GAMs, the trend and seasonals from the SS model are plotted.

Table 6.24: *Chaetoceros* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	2.807	9.509	69.319
AIC	482.513	383.324	415.498
BIC	493.851	421.720	695.397

The AIC and BIC scores in Table 6.24 agree with the choice of the SS GAM. We believe that this series has a stable seasonal pattern.

Diagnostic checking was performed for the SS and TS models as described above. The original deviance for the TS model is greater than the 99 other derived deviances. This indicates that the TS model is not a good fit. The SS model, however, is a good fit as the original deviance is the 41st percentile of the 99 deviances.

Total Diatoms

This series includes all diatoms in the Stonehaven database. Since this series has a sufficient number of non zero values we use a log Normal distribution for the GAMs. The series, however, exhibits very sudden high spikes, for example, going from 0 to 3000000 with a seemingly random allocation of zero values, see Figure 6.66 of the logged values. This is very hard to model, therefore, we remove the zero values from the series, inserting missing values in their place, when creating the GAMs. The three GAMs are created for the logged series after the missing values are removed. We thus, find a model for the series conditioned on the presence of diatoms. There are 35 zeroes in a total of 419 samples giving a probability of 0.084. Dividing the obtained conditional model by 0.084 we obtain the unconditional model.

This is an aggregate of many different species that have various patterns. As a result the interpretation of the following analysis is very hard. From a biological point of view, however, the behaviour of the seasonal pattern of all diatoms is interesting as the diatoms share the same characteristics, for example need silicate and are all photosynthesising. Thus, diatoms are more sensitive to a change in the spring weather conditions (Bresnan, pers.comm., 2008).




Figure 6.66: The Logged series of Total Diatoms, after adding 1.

The Resampling results for the Total Diatoms series indicate a stable seasonal pattern. **Resampling Testing Stable vs Trending:** $p = 0.40 \rightarrow$ Stable

Additionally, the UBRE score also prefers the SS model:

UBRE Score: Non-Seas: 3.324, **Stable:** 1.761 & Trending: 3.051

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.064, Stable: 48.9 & Trending: 64.9

Table 6.25: Total Diatoms ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	382.000	1257.715				
SS	377.096	642.522	4.904	615.192	73.623	<2.2e-16
TS	278.194	442.035	98.901	200.487	1.276	0.064

The ANOVA results, presented in Table 6.25, favour the SS GAM.




Figure 6.67: Clockwise from top left the logged Total Diatoms series with the trend (red line) from the SS GAM, the stable seasonals and the series (dotted) with the stable (red) and trending (blue) fitted GAMs.

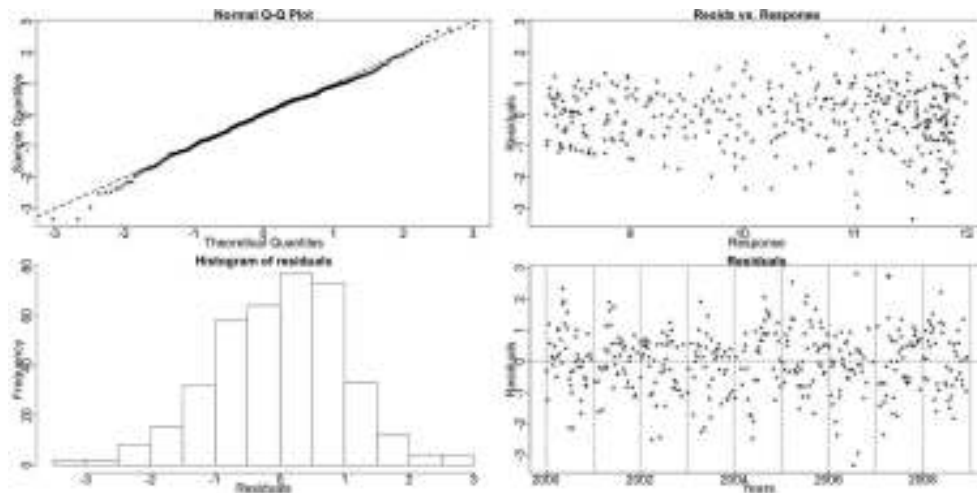
In Figure 6.67 the series of Total Diatoms with the fits of the SS (red) and TS (blue) GAMs are plotted. In the same Figure, we see the trend and seasonals from the SS model.

Table 6.26: Total Diatoms AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3.000	7.904	106.806
AIC	1551.326	1303.220	1357.403
BIC	1563.178	1334.447	1779.353

Both the AIC and BIC scores in Table 6.26 point to the SS GAM. This series' seasonal pattern is stable.

In Figures 6.68 and 6.69 we see the residual plots for the SS and TS models, respectively. The SS model's QQ plot does not show any deviation from normality, while the TS's one is problematic at the lower end of the series.

**Figure 6.68:** Diagnostic plots for the Total Diatoms SS model.

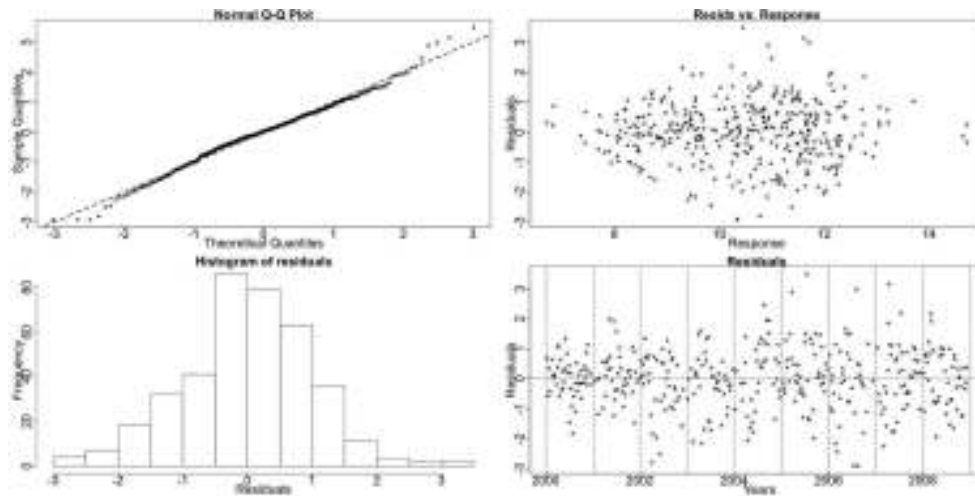


Figure 6.69: Diagnostic check for the Total Diatoms TS model.

6.4.2 Dinoflagellates

Dinoflagellates are usually more abundant during late spring and summer and in some occasions can form massive algal blooms. Unlike the diatoms' silicate exoskeleton, the dinoflagellates' one is made of cellulose. In addition, not all dinoflagellates have *Chlorophyll a*. Other pigments are often present and many species are mixotrophic; i.e. able to photosynthesise or feed on other organisms. Some other dinoflagellates are fully heterotrophic feeding only on other species. As with many diatoms, the dinoflagellate species are able to survive as resting cysts and resistant stages in sediments. The environmental cues that trigger excystment and often blooms are very poorly understood (Amorim, pers. comm., 2009).

Alexandrium

The genus *Alexandrium* includes a few species, some being toxic dinoflagellates that can be responsible for seasonal harmful algal blooms. The series *Alexandrium* is a

total of all recorded *Alexandrium* species. The Resampling results for *Alexandrium* indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.114 \rightarrow$ Stable

The UBRE score, however, prefers the TS model:

UBRE Score: Non-Seas: -0.513 , Stable: -0.615 & **Trending:** -0.662

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 9.05, Stable: 37.1 & Trending: 76.4

Table 6.27: *Alexandrium* ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	P(> Chi)
NS	414.640	191.752			
SS	408.724	132.528	5.900	59.224	5.707e-11
TS	385.169	49.692	23.555	82.836	1.592e-08

The ANOVA results, presented in Table 6.27, favour the TS GAM, too.




Figure 6.70: The *Alexandrium* series with the trend (red line) from the SS GAM, the stable seasonals and the series as dots with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.70 the series of *Alexandrium* with the fits of the SS (red) and TS (blue)

GAMs and the terms of the SS GAM are plotted.

Table 6.28: *Alexandrium* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	4.376	10.276	33.831
AIC	200.504	153.080	117.354
BIC	218.174	194.574	253.960

The AIC and BIC scores, in Table 6.28, point to the TS and SS GAM, respectively. This series has a pseudo-trending seasonal pattern and this random change in the size of the seasonals is the reason that the ANOVA and the AIC prefer the TS GAM.

Diagnostic checking was performed for the SS and TS models as described above. The original deviance for the TS model is the 60th percentile of the other 99 deviances. This indicates that the TS model is a good fit. The SS model is, also, a good fit as the original deviance is the 50th percentile of the 99 deviances.

Total Dinoflagellates

The Total Dinoflagellates series includes all dinoflagellates species in the Stonehaven database. The Resampling results for Total Dinoflagellates indicate a stable seasonal pattern.

As this is an aggregate series of many different species that have various patterns the interpretation of the following analysis is not straightforward. Nevertheless, as with the Total Diatoms series, all dinoflagellates share some characteristics and from a biological point of view the behaviour of the Total Dinoflagellates series is important.

Resampling Testing Stable vs Trending: $p = 0.277 \rightarrow$ Stable

Additionally, the UBRE score prefers the SS model:

UBRE Score: Non-Seas: 0.296, **Stable:** -0.050 & Trending: 0.024

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 3.83, Stable: 32.1 & Trending: 60.6

Table 6.29: Total Dinoflagellates ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	P(> Chi)
NS	415.518	533.046			
SS	411.166	375.779	4.352	157.267	1.12e-32
TS	311.353	184.04	99.813	191.739	9.055e-08

The ANOVA results, presented in Table 6.29, favour the TS GAM, however.




Figure 6.71: The Total Dinoflagellates series with the trend (red line) from the SS GAM, the stable seasonals and the Total Dinoflagellates series marked by dots with the stable (red) and trending (blue) fitted GAMs.

Figure 6.71 shows plots of the series of Total Dinoflagellates, the trend and seasonals of the SS GAM and the series as dots with the fits of the SS (red) and TS (blue) GAMs are plotted.

In Table 6.30 both the AIC and the BIC scores point to the SS one. We believe that this pattern is stable.

Diagnostic checking was performed for the SS and TS models as described above. The original deviance for the TS model is greater than the 99 other derived deviances.


Table 6.30: Total Dinoflagellates AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3.482	7.834	107.647
AIC	540.008	391.447	399.334
BIC	554.082	423.081	833.999

This indicates that the TS model is not a good fit. The SS model, however, is a good fit as the original deviance is the 30th percentile of the 99 deviances.

6.4.3 Total Phytoplankton

The Total Phytoplankton series includes all phytoplankton species in the Stonehaven database. Similarly to the Total Diatoms series, we remove the zero values from this series and use a log Normal distribution for the GAMs. The zero values occur randomly, see Figure 6.72, with a total of 27 in 419 measurements, giving thus a probability of 0.064 for absence.

**Figure 6.72:** The logged series of Total Phytoplankton.

The Resampling results for Total Phytoplankton indicate a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.116 \rightarrow$ Stable

Additionally, the GCV score prefers the SS model:

GCV Score: Non-Seas: 3.216, **Stable:** 1.516 & Trending: 2.538

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.454, **Stable:** 54.9 & Trending: 66.9

Table 6.31: Total Phytoplankton ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	389.488	1238.050				
SS	384.168	561.386	5.319	676.661	87.052	<2.2e-16
TS	284.491	377.569	99.678	183.817	1.39	0.019

The ANOVA results, presented in Table 6.31, favour the TS GAM, however.




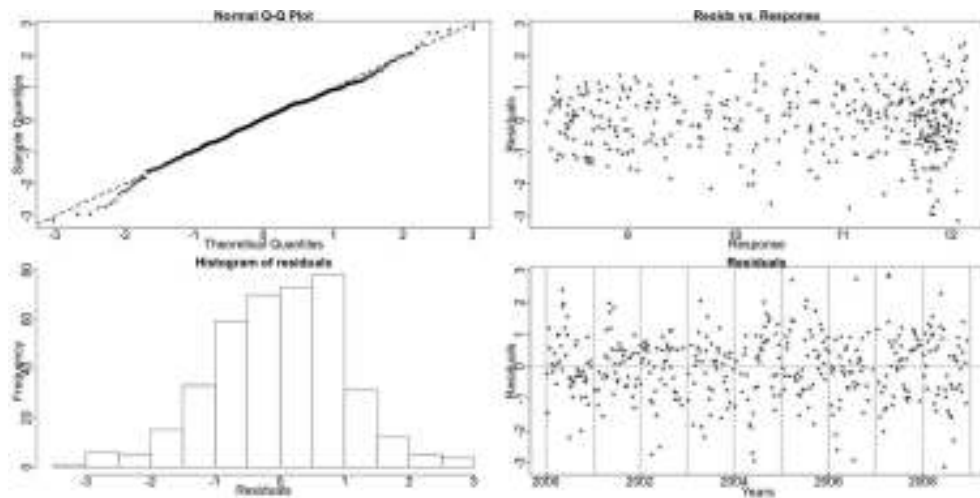
Figure 6.73: Clockwise from top left the Total Phytoplankton series with the trend (red line) from the SS GAM, the stable seasonals and the logged series of Total Phytoplankton marked by dots with the stable (red) and trending (blue) fitted GAMs.

In Figure 6.73 the series of Total Phytoplankton with the fits of the SS (red) and TS (blue) GAMs are plotted. In the same Figure we see the stable seasonals and trend from the SS model.

Table 6.32: Total Phytoplankton AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3.512	8.832	108.509
AIC	1570.284	1270.897	1314.763
BIC	1584.232	1305.970	1745.681

Both the AIC and BIC scores in Table 6.32 point to the SS GAM. There is a fluctuation of the size of the seasonals of this series but the pattern is stable in terms of systematic shifts in time.

**Figure 6.74:** Diagnostic check for the Total Phytoplankton SS model.

In Figures 6.74 and 6.75 residual plots for the SS and TS model are shown. There is a slight violation of normality at the ends of the series.

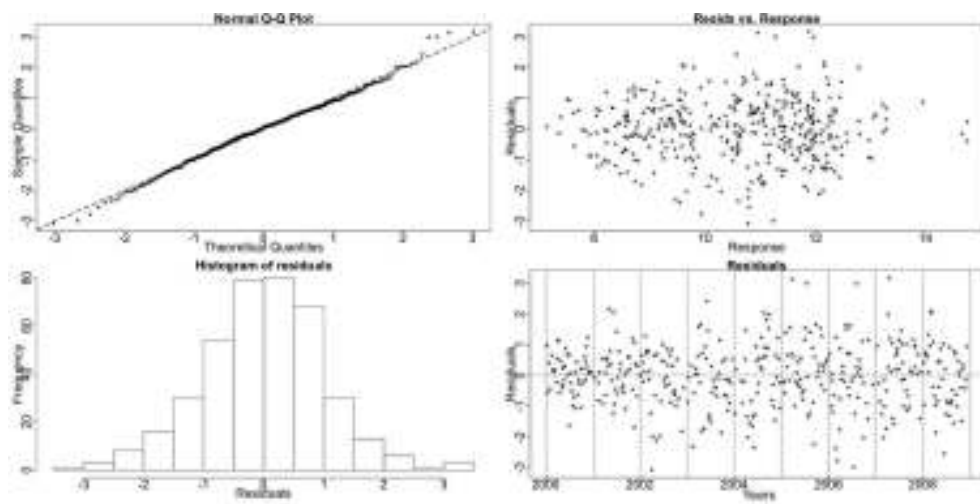


Figure 6.75: Diagnostic check for the Total Phytoplankton TS model.

6.5 Notes on Results

Both seasonal decomposition based testing procedures, Resampling and the GAMs, are used to test in each case a specific hypothesis under the specified models. Thus, the Resampling result is dependent on the estimated trends and stable/trending seasonals. Similarly, the criteria used for GAM selection test not whether the series has stable or trending seasonality but if the specified stable or trending model is a better fit to the data. This implies that a ‘poor’ stable model could give a trending result and the opposite. In addition, the BIC is more conservative and thus, it has been observed that if the series seems to be trending (the other results point us in that direction) then the BIC often indicates a stable model while if the series seems stable then the BIC points to a non-seasonal model. The ANOVA is more sensitive than the rest of the criteria used and is influenced by any type of change, thus preferring trending than stable models. Nevertheless, when there is no change in the seasonals the ANOVA correctly points to the SS GAM while if there is a big systematic change in the seasonals the conservative BIC will identify a TS GAM.

Chapter 7

Stonehaven Zooplankton Results

7.1 The zooplankton data

The results from the analysis of the zooplankton time series from Stonehaven are presented in this chapter. The zooplankton series share some similarities with the environmental and phytoplankton series of the previous section. They also are weekly observations but unlike the environmental series these are counts much like the phytoplankton, and not measurements. Thus, a zero in the series cannot be attributed to the detection limit of the equipment but can mean either a zero in the population or a zero in the sample.

The sampling procedure involves lowering a plankton net¹ to a depth of 45m and hauling vertically to sample the water column from 45m near the seabed at around 50m. The sample is preserved in 4% buffered formaldehyde and returned to the laboratory for species identification and counts. After removing the big species for direct count the sample is made up to a standard volume of 200ml and a sub-sample of suitable volume (2.5, 5 or 10ml) is then extracted and everything in it is identified to a suitable taxonomic level and counted. The raw numbers of count are derived

¹a bongo net with 40 cm diameter and 200 μ m mesh

by multiplying up the counts from the sub-sample to the standard (200ml). These species counts per sample are then corrected to generate average densities of the species in the water column, expressed as numbers per cubic metre or as standing stock (i.e. the number beneath a square meter of water surface). This assumes the net samples a cylindrical volume with 70% efficiency.

The database does not include information on the volume of the sub-sample used which would be very useful so as to consider the weight that each observation carries. In year 1999 (16th of March) the diameter of the net changed from 30 to 40cm (BO200 to BG200) which caused the value of that constant to change. In this section we are examining the densities of some zooplankton species.

7.1.1 Resampling

The Resampling testing techniques presented in chapter 4 are used with the same modifications discussed previously to allow for the many missing values in the series. The observations are allocated weeks in the year (52 in total, not including the 26th of December) according to the collection date. In order to form time series, when two or more observations fall in the same week, averages are used. The zooplankton series often exhibit high spikes, large differences in the counts, so the logarithm of the series is computed and tested using these procedures ².

All the reported testing results refer to tests of whether the seasonal pattern is stable or trending since the nature of the data implies that they are in fact all seasonal; all results between non-seasonal or seasonal verify this, unless stated otherwise.

²More accurately, the logarithm of the series plus one $\log(x_t + 1)$, to account for the zeros in the series.

7.1.2 GAMs

GAMs are used creating the aforementioned three types of models: non-seasonal, NS, (smooth function of time), stable seasonal, SS, (smooth function of time and a cyclic component for seasons) and trending seasonal, TS, (smooth function of time, a cyclic component for seasons and smooth function of time by factors of seasons) and testing for the best suited one.

The data are not used in the form of a time series so multiple entries for one week are allowed, thus utilizing all available data points. The Poisson family was at first assumed, as the data are in the form of counts, with the log link function, but there arose issues caused by overdispersion and zero inflated counts. Furthermore, the variation in the exact volume of the sample actually counted drew us to look at the densities of the zooplankton database instead of the raw counts. For most series a log Normal distribution gives an adequate fit while Gamma distribution was also examined but rejected. The constant relating to the volume of water sampled is used as an offset to account for the change in the width of the net. The stable seasonals are estimated by a cyclic polynomial with 10 knots and the number of knots for the trending seasonality terms are set equal to three as in the previous chapter.

For the Gaussian family GAMs use GCV score, an estimate of the prediction error, to find the appropriate smoothness for each applicable model term. To avoid overfitting and to favour smoother fits we are using $\gamma = 1.4$ instead of the default value 1 in the construction of the model as advocated by Kim & Gu (2004). ANOVA is used between non-seasonal versus seasonal and stable seasonal versus trending seasonal models. Most series were found seasonal so only the stable versus trending p-values are reported for the majority of them while values for non-seasonal versus stable seasonal and non-seasonal versus trending seasonal are reported for the series that were tested non seasonal. Finally AIC and BIC values for all three models are reported with the minimum highlighted in bold. For each model the GCV score is also presented with the lowest in bold. We also report the percentage deviance explained by each GAM as a measure of fit. When the results point to a trending

seasonal pattern we further examine the smooth terms fitted to each week by looking at plots of each smooth function.

When fitting GAMs one has to check for any assumptions violations by checking the residuals through diagnostic plots. As before, we are using standardised deviance residuals, \hat{d}_{std} ³ for our diagnostic plots. We construct a QQ plot and a histogram of the residuals to test the assumption of normality. A plot of the residuals against the response is also examined. The residuals should appear evenly scattered while a trend in the variability of the residuals would flag a violation of the constant variance assumption. Finally, we look at the residual values. There should be no patterns in them as that would indicate that there is still information to be extracted from them and they are not just ‘noise’.

The results are presented by genus and then species. In Table 7.1 the results for all zooplankton series examined are summarised. Only the most commonly known species and series with interesting results are presented in this section. For the first series of each species presented the results from the GAMs are more extensively given, including a variety of plots and tables. When similarities occur, the results for the rest of the series of the same species are not presented. Extensive plots and tables of results are presented only when they appear noteworthy or in need of further interpretation. A full description of the results for all zooplankton series examined can be found in the form of an electronic appendix at the end of this thesis. As the Resampling is used on the logged data and the GAMs assume a log Normal distribution, all plots shown below are with regards to the logged data.

Not all species reported here were recorded since 1997. *Acartia*, *Centropages*, *Temora*, *Metridia lucens*, *Paracalanus parvus* and *Pseudocalanus elongatus* were included in the database in 1999 so these series are shorter by two years.

³ $\hat{d}_{std,i} = (\hat{d}_i - \bar{\hat{d}}) / (\hat{\sigma} \sqrt{1 - H_{ii}})$, where H_{ii} are the elements on the leading diagonal of the hat matrix, H , and $\bar{\hat{d}}$ the average of the deviance residuals \hat{d}_i .

Table 7.1: Zooplankton Series Results

Species & Copepodite Stage		Pattern
<i>Acartia clausi</i>	4, 5, 6f	Pseudo-trending in Amplitude
	6m	Trending (Time Shift earlier)
Total <i>Acartia</i>		Pseudo-trending in Amplitude
<i>Centropages typicus</i>	1-5 and 6m	Pseudo-trending in Amplitude
	6f	Pseudo-trending in Amplitude and Phase
<i>Centropages hamatus</i>	4-6	Pseudo-trending in Amplitude
<i>Calanus finmarchicus</i>	5	Trending (Time Shift earlier)
	6f and 6m	Pseudo-trending in Amplitude
<i>Calanus helgolandicus</i>	5	Trending (Time Shift Earlier)
	6f and 6m	Trending (↑ Amplitude & Time Shift Earlier)
Juvenile <i>Calanus</i>	1-4	Trending (↑ Amplitude & Time Shift Earlier)
Calanoid Copepods		Pseudo-trending in Amplitude
<i>Temora longicornis</i>	1-6	Pseudo-trending in Amplitude
Total <i>Temora</i>		Pseudo-trending in Amplitude
<i>Paracalanus parvus</i>	1-6	Pseudo-trending in Amplitude
<i>Pseudocalanus elongatus</i>	1-6	Pseudo-trending in Amplitude
<i>Metridia lucens</i>	1-5 and 6m	Pseudo-trending in Amplitude
	6f	Non Seasonal
<i>Oithona</i>	1-6	Pseudo-trending in Amplitude
Total <i>Oithona</i>		Pseudo-trending in Amplitude
<i>Lamellibranchiata</i>		Pseudo-trending in Amplitude

7.2 Results by species

Copepods

Copepods are small crustaceans that are so abundant that often represent more than 90% of the biomass in mesozooplankton samples. They are mainly herbivores and as such represent a major pathway for the transfer of energy up the marine food web. Some copepod species exist in the water all year and some produce resting stages that allow them to survive winter. It is a feature of most copepods that the adult males either do not feed or feed less and consequently have a shorter life span than adult females; whose feeding as adults is translated mainly into egg production rather than somatic growth (Hay, pers. comm., 2009).

7.2.1 *Acartia*

Acartia is a small calanoid copepod that is found in shelf seas. The genus *Acartia* is common and abundant all around the world's oceans. It is an important link in the food chain as it is largely a herbivore mediating energy transfer between phytoplankton and developing larvae of commercial fish species. *Acartia clausi* is resident in the North Sea all year, but late in the year the eggs produced are able to delay hatching and so rest in the seabed sediments to hatch in spring through early summer (Hay, pers. comm., 2009). We look at some developmental stages of *Acartia clausi* and the total counts of *Acartia*, including all observed stages of *clausi*, *longeremis*, *discudata* and *bifilosa*, which are, relative to *clausi*, rare species in the samples .

Acartia clausi

Acartia clausi has a life-cycle that consists of six naupliar stages followed by five copepodid larval stages to reach adulthood in stage 6. We examine densities of

Acartia clausi copepodite stages 4 (C4), 5 (C5) and 6 (C6f) females and males (C6m). The 200 micron net mesh of the sampler does not trap the early nauplii or the first three copepodite stages in a quantitative way as they pass through this mesh. *Acartia clausi* exhibits one peak per year for all examined stages.

Acartia clausi copepodite stage 4

Acartia clausi C4 exhibits one major peak per year and appears to have a stable seasonal pattern. The Resampling test yields an insignificant p-value when used on the logged *Acartia clausi* C4 series to test for trending seasonality. The p-value indicates a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.312 \rightarrow$ Stable

This is further verified by the GCV score of the GAMs:

GCV Score: Non-Seas: 4.774, **Stable:** 1.297 & Trending: 2.059

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.762, Stable: 75.2 & Trending: 83.3

Table 7.2: *Acartia clausi* C4 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	429.842	2033.470				
SS	417.326	508.417	12.516	1525.058	100.018	$< 2.2e - 16$
TS	314.726	341.736	102.6	166.681	1.496	0.005

Looking at the ANOVA results, Table 7.2, the trending model is preferred.

In Figure 7.1 we see the logged *Acartia clausi* C4 with the fitted values from the two, SS (red) and TS (blue), models. In the same Figure the trend and seasonals from the SS GAM are also plotted.

The AIC and BIC scores, in Table 7.3, point to the SS GAM. The GAMs, as mentioned before, are sensitive to any kind of change and the change in the amplitude of the seasonals, see Figure 7.1, could be the reason the ANOVA prefers the trending model. This series has a pseudo-trending in amplitude seasonal pattern.




Figure 7.1: The logged *Acartia clausi* C4 series with the trend from the SS model, the stable seasonals and the series as dots with the fits from the stable (red) and trending (blue) GAMs. Dotted lines denote the years.

Nevertheless, this amplitude change could affect the sensitivity of the above Resampling testing procedure. Thus, the series could be trending in time but the above test fails to detect it. One should attempt to accommodate for these changes in the size of the seasonals and then test again the series for trends in time. Testing for one specific type of change (for us systematic time shift) in the presence of other unattended changes can lead to ambiguity in the interpretation of the results.

We, nonetheless, believe that this series' seasonal pattern does not exhibit a systematic change, only a random amplitude change.

Table 7.3: *Acartia clausi* C4 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3.158	15.667	118.274
AIC	1901.480	1327.674	1361.256
BIC	1914.329	1391.443	1842.447

The diagnostic plots for the SS models are shown in Figure 7.2. The QQ plot is satisfactory as the residuals display only little deviation from the theoretical quantiles

at the extremes of the series. Nevertheless, the histogram of the residuals looks good. The plot of residuals against the fitted values seems fine and the lower right plot of the residuals themselves shows a slight remaining pattern, but not anything problematic.

Figure 7.3 shows the diagnostic plots for the TS model. The plots look similar to those for the SS model. The QQ plot shows a slight deviation from normality at the extremes of the series which is also seen in the histogram of the residuals. The plot of the residuals against the fitted values and the residuals themselves show no problems. The appearance of straight lines in this plot is due to the fact that the residuals take a small number of values for the corresponding low values of the response and thus, it is not regarded as a problem.

The *Acartia clausi* C5 and C6f series examined behave similarly to *Acartia clausi* C4, see electronic appendix for details.

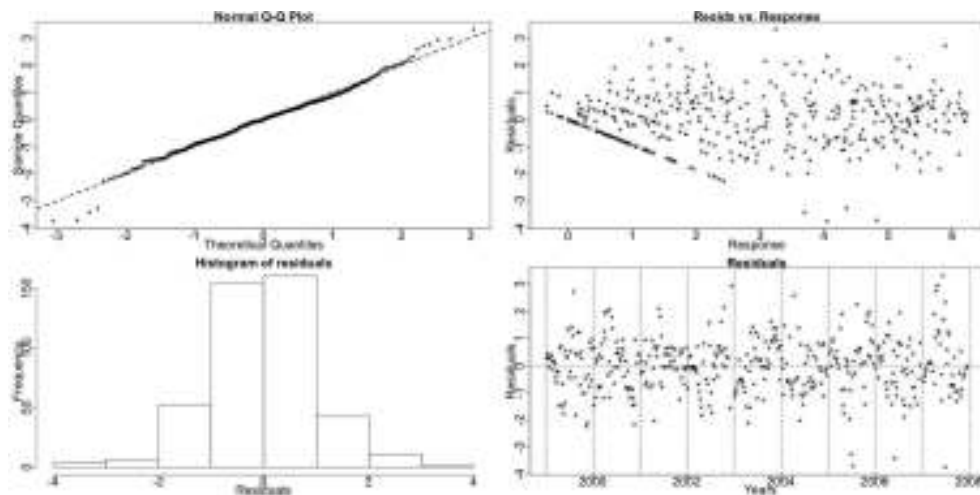


Figure 7.2: Diagnostic plots for the SS GAM for *Acartia clausi* 4. The QQ plot shows a slight problem with the Gaussian assumption at the extremes of the series.

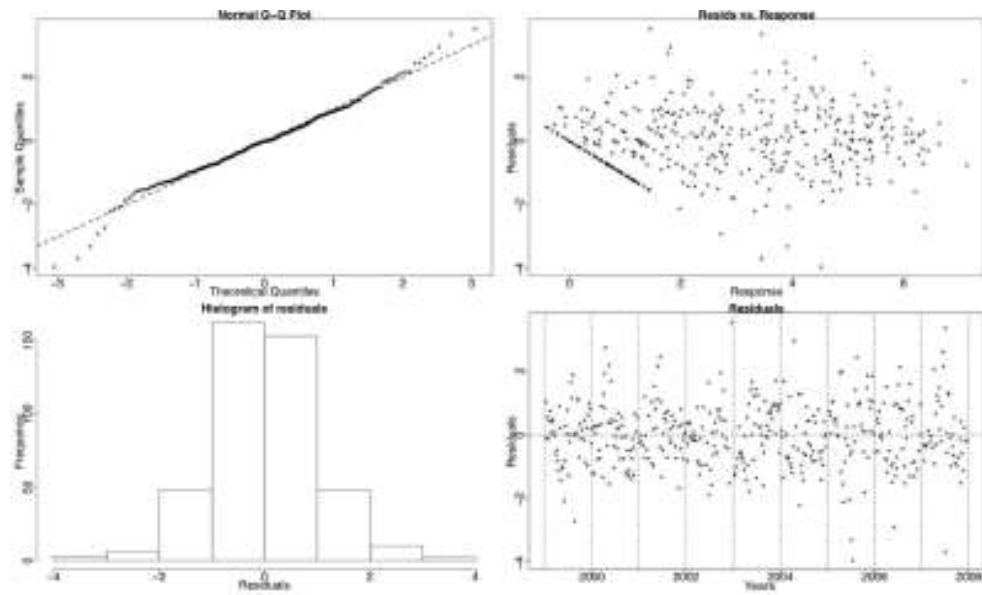


Figure 7.3: Diagnostic plots for the TS GAM for *Acartia clausi* 4. The QQ plot shows a small problem with the Gaussian assumption at the extremes of the series.

Acartia clausi copepodite stage 6 males

The *Acartia clausi* C6m series has a seasonal cycle that changes systematically in time. The Resampling test for trending seasonality points to stable seasonality.

Resampling Testing Stable vs Trending: $p = 0.07 \rightarrow$ Stable

The GCV score of the GAMs also favours stable seasonality:

GCV Score: Non-Seas: 4.935, **Stable:** 1.381 & Trending: 2.076

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.596, Stable: 74.7 & Trending: 83.3

Table 7.4: *Acartia clausi* C6m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	429.94	2103.506				
SS	416	536.255	13.94	1567.251	87.214	$< 2.2e - 16$
TS	316.876	352.461	99.124	183.793	1.667	4.913e-4

The ANOVA, however, supports a trending seasonal pattern, Table 7.4.




Figure 7.4: Clockwise from top left corner, the logged *Acartia clausi* C6m with the trend from the SS model, the stable seasonals and the series as dots with the fitted values from the stable (red) and trending (blue) GAMs.

Figure 7.4 show the *Acartia clausi* C6m logged densities with the prospective fits from the SS (red) and TS (blue) models. In the same Figure the trend and seasonals from the SS GAM are plotted.




Figure 7.5: The smooth fitted functions for weeks 1-36, in the TS model for the *Acartia clausi* C6m data.

In Figures 7.5 and 7.6 the smooth functions fitted to each week by the TS GAM are shown. The smooth functions corresponding to a number of weeks differ from a straight line parallel to the x-axis. Specifically, the smooth functions of some weeks like weeks 16-18, 37 and 38 are monotonic indicating a possible systematic change.

Table 7.5: *Acartia clausi* C6m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3.06	17	116.124
AIC	1915.91	1353.354	1370.306
BIC	1928.359	1422.518	1842

The AIC and BIC scores point to the SS GAM, however.

To further examine the type of changes the seasonal pattern exhibits, we plot the




Figure 7.6: The smooth fitted functions for weeks 37-52, in the TS model for the *Acartia clausi* C6m data.

stable (red) and trending (blue) patterns estimated by the SS and TS GAMs, respectively, Figure 7.7. There is a random change in the size of the estimated trending seasonals but also there appears to be a slight systematic shift earlier in the year. In particular it is the descend from the peak that appears to be shifting earlier in time.

Our decision is that the pattern exhibits a random amplitude change but also starts shifting slightly earlier in the year.

The diagnostic plots for the SS and TS GAMs are similar to the ones presented for *Acartia clausi* C4.




Figure 7.7: The stable seasonals, red, and the trending seasonals, blue, from the SS and TS GAMs, respectively, fitted to the *Acartia clausi* C6m series.

Total *Acartia*

The total of all *Acartia* series available in the database includes all observed stages of *clausi*, *longeremis*, *discudata* and *bifilosa*. It has a stable seasonal pattern that does not shift with time. The Resampling results point to a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.784 \rightarrow$ Stable

The GCV score of the GAMs coincides with the Resampling results:

GCV Score: Non-Seas: 4.736, **Stable:** 1.191 & Trending: 2.043

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 6.88, Stable: 77.9 & Trending: 83.7

Table 7.6: Total *Acartia* ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	424.141	1943.114				
SS	415.537	461.211	8.604	1481.903	155.174	$< 2.2e - 16$
TS	314.972	340.008	100.565	121.203	1.116	0.238

All the results, Tables 7.6 and 7.7, point to a stable seasonal pattern.

The terms of the fitted GAM with stable seasonality are plotted in Figure 7.8. In the same Figure we see the series with the fitted values from the stable (red) and trending (blue) GAMs.

This series has a stable seasonal pattern in terms of systematic changes. It does exhibit though a random amplitude change.




Figure 7.8: Total *Acartia* logged densities with the trend, as a red line, and seasonals from the SS GAM. The lower plot shows the series as dots with the fits from the SS (red) and TS (blue) GAMs.

Table 7.7: Total *Acartia* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	8.859	17.463	118.028
AIC	1893.245	1289.154	1358.574
BIC	1929.286	1360.201	1838.763

7.2.2 *Centropages*

Centropages is a common genus of calanoid copepods. We will be looking at *Centropages typicus* (*C. typicus*) and *Centropages hamatus* (*C. hamatus*). At Helgoland Roads, these two *Centropages* species coexist and a seasonal succession is noticed between *C. hamatus* and *C. typicus*. *C. typicus* prefers the saline and relatively warm waters of Atlantic origin and is more abundant there in the second half of the year while maximum abundances of *C. hamatus* are usually found, before the *C. typicus* seasonal peak (Bonnet, 2007).

C. typicus is a larger species with a more carnivorous diet and being optimally adapted to warmer waters it is found abundantly only through the warmer months. *C. hamatus* is a North Sea resident and more abundant species that is more omnivorous in its diet. Like *Acartia* it is known to produce resting eggs later in the year that overwinter in the seabed sediments and hatch in Spring to develop a new population, which reproduces with subitaneous eggs that develop normally to hatch into naupliar stages (Hay, pers. comm., 2009).

Centropages typicus

C. typicus is a calanoid copepod with a wide range of distribution in the North Atlantic and adjacent shelf seas. It feeds on a wide range of prey both phytoplankton and animal prey (e.g. nauplii of copepods). Stocks increase in temperatures between 13°C and 20°C causing a seasonal peak in summer and up to early winter in the North Sea. Its over-wintering behaviour is not understood but since it cannot survive in the North Sea it is associated with influx of water from the Atlantic ocean (Carlotti & Harris, 2007; Bonnet, 2007).

C. typicus copepodite stages 1 to 4

C. typicus copepodite stages 1 to 4 (C1-4) exhibits a stable seasonal pattern with one peak per year. The Resampling test points to a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.809 \rightarrow$ Stable

The GCV score of the GAMs agrees with the Resampling results:

GCV Score: Non-Seas: 0.839, **Stable:** 0.572 & Trending: 0.916

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 15.1, Stable: 44.5 & Trending: 64.6

Table 7.8: *C. typicus* C1-4 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	422.475	340.473				
SS	416.198	222.515	6.278	117.958	35.145	$< 2.2e - 16$
TS	308.248	141.977	107.949	80.538	1.62	0.001

The ANOVA , Table 7.8, is in favour of a trending seasonal model while the Resampling p-value, the GCV score, the AIC and BIC, Table 7.9, a stable one. There is a random change in the size of the seasonals and, as explained before, the GAMs are very sensitive to any kind of change. We believe that there is a fluctuation in the size of the seasonals but not a systematic shift in time, i.e. pseudo-trending seasonality.

In Figure 7.9 we see the logged series of *C. typicus* C1-4 with the terms from the SS GAM and the SS (red) and TS (blue) fitted GAMs.

Table 7.9: *C. typicus* C1-4 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	10.525	16.802	124.752
AIC	1144.157	972.965	994.753
BIC	1186.975	1041.324	1502.296

Figures 7.10 and 7.11 show the diagnostic plots for the two GAMs. Both QQ plots show a deviation from normality.

Finally, as a note we mention that the somewhat intriguing pattern that appears in the residuals is just a product of the smooth fit. The models fit smooth and gradual




Figure 7.9: *C. typicus* C1-4 with the trend from the SS model, the stable seasonals and the series (marked as dots) with the two fits from the stable (red) and trending (blue) GAMs.

increases and decreases in the numbers of *C. typicus* C1-4 but the real data exhibits very sharp rises and falls which causes a pattern to appear in the residuals. This could be seen as an indication that these type of time series (rapid increase and decrease) are not suitably modelled by GAMs. Another possibility is to model them as binary data with the Binomial family as we did for most phytoplankton series in the previous chapter.

The results for the other *C. typicus* series are very similar to those for *C. typicus* C1-4. The series of *C. typicus* C6m, however, has only 75 non zero values and the diagnostic plots for the GAMs show a poor fit. Modelling it as a binary variable with a Binomial GAM would be better suited. For detailed results for these series please see the electronic appendix.




Figure 7.10: *C. typicus* C1-4 diagnostic plots for the SS GAM.

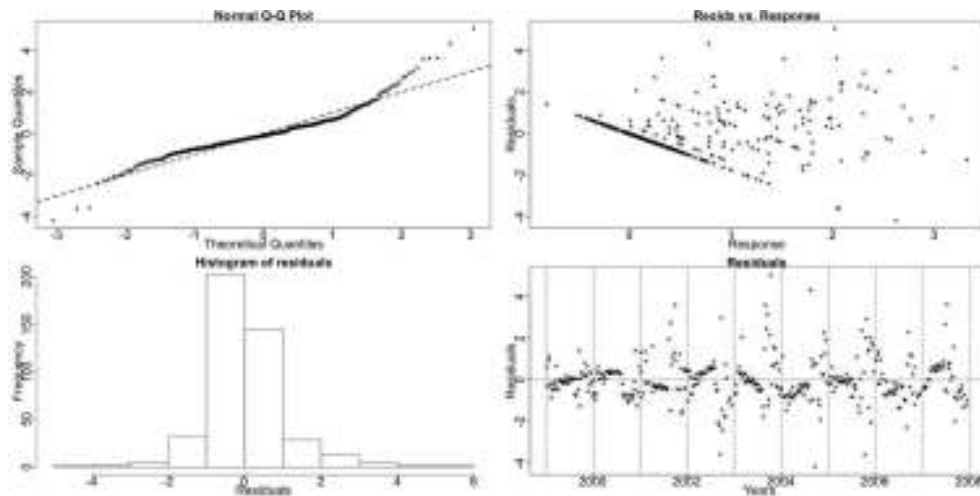


Figure 7.11: *C. typicus* C1-4 diagnostic plots for the TS GAM.

Centropages hamatus

Centropages hamatus (*C. hamatus*) is like *C. typicus* a very abundant calanoid copepod. However, contrary to *C. typicus*, resting eggs of *C. hamatus* have been identified in the Southern North Sea. This is commonly interpreted to mean that these eggs initiate a first generation which hatches and then develops to adult-hood during March, when the first females are recorded in the plankton (Bonnet, 2007). All analysed hamatus series that follow are found to have a seasonal pattern which is stable in time.

Centropages hamatus copepodite stage 4

C. hamatus C4 exhibits one major peak per year. It starts in spring and lasts till autumn. The Resampling results identify a stable seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.531 \rightarrow$ Stable

This is further verified by the GCV score of the GAMs:

GCV Score: Non-Seas: 1.475, **Stable:** 0.791 & Trending: 1.181

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 2.51, Stable: 49.5 & Trending: 66.5

Table 7.10: *C. hamatus* C4 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	430	628.765				
SS	424.657	325.462	5.343	303.303	74.062	$< 2.2e - 16$
TS	324.255	216.137	100.402	109.326	1.634	0.001

The seasonality of the series changes in terms of size in a random way, see Figure 7.12. The ANOVA, Table 7.10, probably picks up this variation in size and proposes the trending GAM. However, the GCV score, the p-value from the Resampling and the AIC and BIC scores, Table 7.11, all indicate a stable model. This pattern is pseudo-trending in amplitude.




Figure 7.12: Clockwise from top left: *C. hamatus* C4 with the trend from the SS model, the stable seasonals and the series, marked as dots, with the fit from the stable (red) and trending (blue) GAMs.

In Figure 7.12 the logged *C. hamatus* C4 series with the two prospective fits from the SS (red) and the TS (blue) models are shown. In the same Figure the trend and seasonals from the SS GAM are also plotted.

Table 7.11: *C. hamatus* C4 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	8.343	108.745
AIC	1394.106	1120.316	1144.287
BIC	1406.312	1154.261	1586.708

Figures 7.13 and 7.14 show the diagnostic plots for the two GAMs. Both the SS and the TS GAMs give an adequate fit. They both exhibit a small deviation from the normality assumption in the QQ plots and histograms but that does not look significant enough to prove them inappropriate.

The other *C. hamatus* series (C5, C6f and C6m) behave similarly to this one and thus their results will not be presented here.

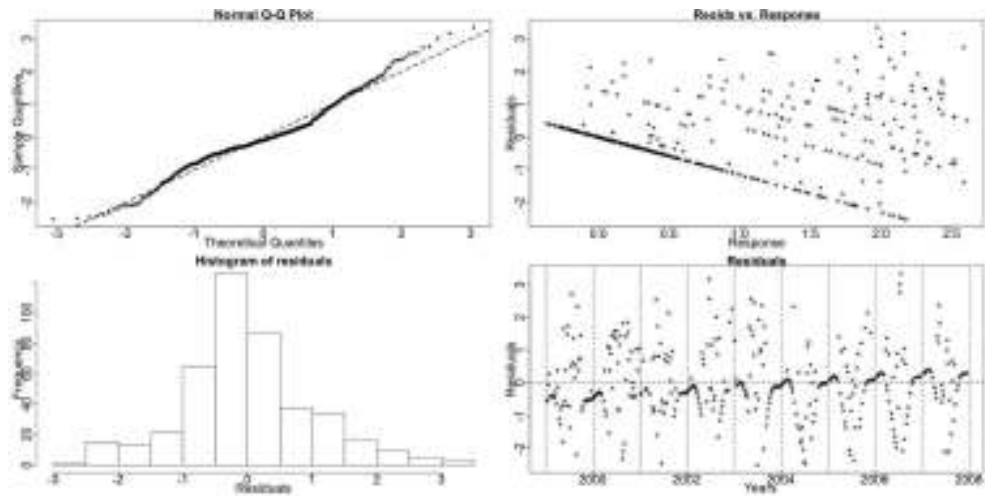


Figure 7.13: *C. hamatus* C4 diagnostic plots for the SS GAM.

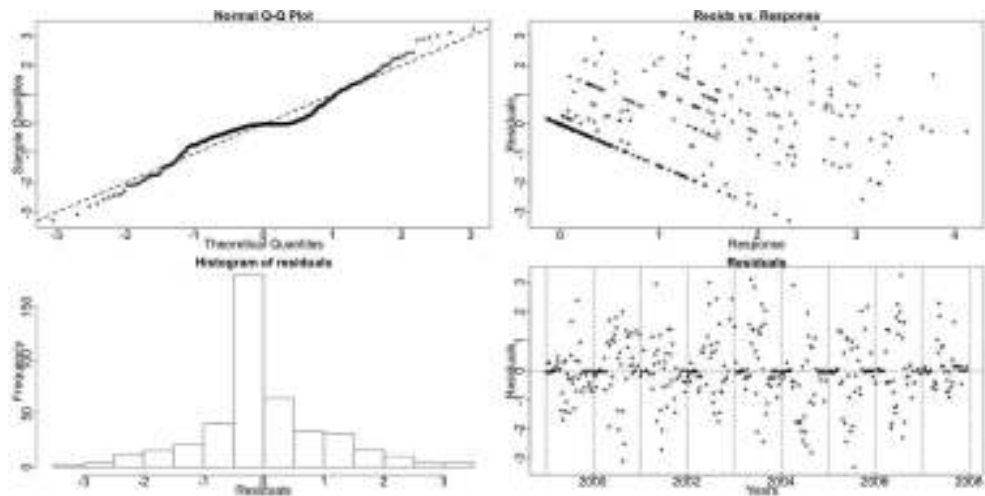


Figure 7.14: *C. hamatus* C4 diagnostic plots for the TS GAM.

7.2.3 *Pseudocalanus elongatus* copepodite stages 1 to 6

Pseudocalanus is considered one of the most common and typical coastal copepods in the North Sea. It is a largely herbivorous species with no resting/overwintering developmental stage (Hay, pers. comm., 2009). *Pseudocalanus elongatus* has six naupliar stages and five copepodite stages before it reaches adulthood in copepodite stage 6. It favours temperature of $5^{\circ} - 15^{\circ}\text{C}$ while reduced growth has been verified at 20°C (Stegert *et al.*, 2007). *Pseudocalanus elongatus* C1-6 has seasonals that peak twice within a year. The first peak is in spring while the second in autumn. The pattern appears to be stable in terms of systematic shifts but the size of the seasonals changes randomly over the years.




Figure 7.15: Clockwise from top left: the logged *Pseudocalanus elongatus* C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs.

In Figure 7.15 we see the logged *Pseudocalanus elongatus* C1-6 series, the terms and fitted values from the SS GAM and the fitted values from the TS GAM.

7.2.4 *Temora longicornis*

Temora longicornis (*T. longicornis*) is a small, common and often abundant copepod species endemic to the region and an important component of the plankton community. This small copepod is omnivorous, also able to produce resting eggs when conditions are poor, enabling it to overwinter in a resting state in the sediments; to reappear when conditions improve in spring (Hay, pers. comm., 2009). All the following *Temora* series are *T. longicornis* and are found to have a stable seasonal pattern in time with two seasonal peaks. It should be noted that the 200 micron mesh of the sampler does not catch the small C1- C3 copepodites of this small species very well.

Temora longicornis copepodite stage 1

The series of *T. longicornis* C1 has a stable seasonal pattern according to the Resampling results.

Resampling Testing Stable vs Trending: $p = 0.241 \rightarrow$ Stable

The GCV score of the GAMs agrees with the Resampling conclusion:

GCV Score: Non-Seas: 0.851, **Stable:** 0.737 & Trending: 1.136

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.308, Stable: 20.6 & Trending: 47.3

Table 7.12: *T. longicornis* C1 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	430	362.798				
SS	417.485	289.028	12.515	73.77	8.514	5.696e-15
TS	316.323	191.856	101.162	97.173	1.584	0.001

All the tests and criteria, except the ANOVA point to the selection of the stable seasonal GAM, Tables 7.12 and 7.13.




Figure 7.16: From top left, clockwise we see the logged series of *T. longicornis* C1 with the trend from the SS model, the stable seasonals and the logged series marked as dots with the fitted values from the SS (red) and TS (blue) GAMs.

In Figure 7.16 we see the fitted values from the SS (red) and TS (blue) models and the terms of the SS one. The size of the seasonal pattern changes randomly in time. The ANOVA probably prefers the TS GAM because it picks up this random change.

Table 7.13: *Temora* C1 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	15.515	116.677
AIC	1156.545	1083.372	1108.67
BIC	1168.75	1146.493	1583.36

Figures 7.17 and 7.18 show the diagnostic plots for the two GAMs. The normality assumption is violated in both models.

The other *Temora longicornis* series behave similarly to this one but have better diagnostic plots for both GAMs, see electronic appendix.




Figure 7.17: The diagnostic plots for the SS model of *T. longicornis* C1.

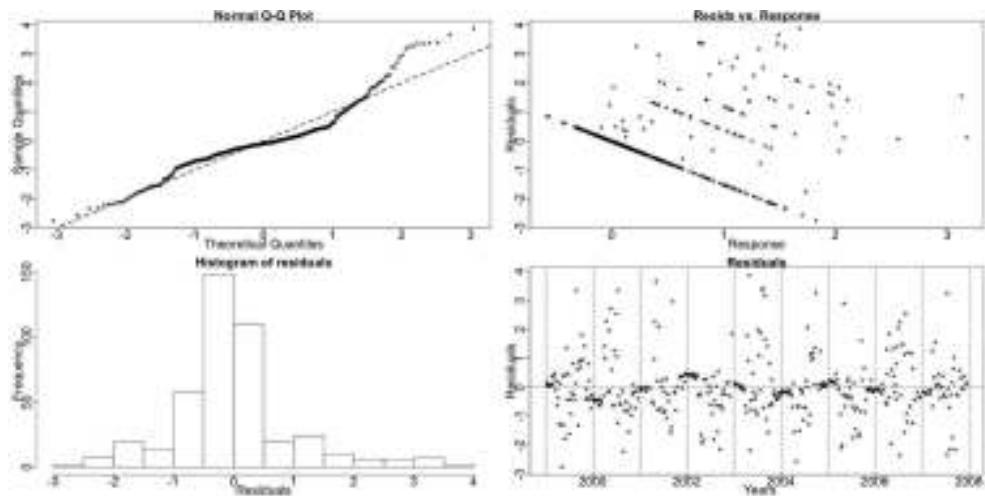


Figure 7.18: The diagnostic plots for the TS model of *T. longicornis* C1.

Total *Temora longicornis*

The Total *T. longicornis* series includes all the observed stages of *Temora longicornis*. It has a small peak in spring and a bigger one in autumn. This is an aggregate series of the developmental stages and thus, the interpretation of the results is more difficult.




Figure 7.19: Clockwise from top left: the logged Total *T. longicornis* with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS(blue) GAMs.

In Figure 7.19 we see the logged Total *T. longicornis* series, the terms and fitted values from the SS GAM and the fitted values from the TS GAM. The pattern changes randomly in size but is stable in terms of systematic changes. This series has a pattern whose amplitude fluctuates but does not exhibit any systematic changes.

7.2.5 *Paracalanus parvus* copepodite stages 1 to 6

Paracalanus parvus is a small, mainly herbivorous coastal copepod found in temperate waters. Usually in the northern North Sea *Paracalanus* appears less abundant than other small copepods such as *Pseudocalanus*, *Acartia* and *Temora*. Like *Pseudocalanus* it has no resting stage for overwintering; instead it survives as best it can on the food available or stored as lipid in its body (Hay, pers. comm., 2009). *Paracalanus parvus* C1-6 series has a seasonal pattern with two peaks, one in spring and one in autumn. We believe that the seasonal pattern does not exhibit systematic shifts but a random amplitude change.




Figure 7.20: Clockwise from top left: the logged *Paracalanus parvus* C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS(blue) GAMs.

In Figure 7.20 we see the logged *Paracalanus parvus* C1-6 series, the terms and fitted values from the SS GAM and the fitted values from the TS GAM.

7.2.6 Small Calanoid Copepods copepodite stages 1 to 6

The small Calanoid Copepods group is compounded of a number of species mainly *Acartia*, *Temora*, *Pseudocalanus*, *Paracalanus* and *Centropages*. It is of interest to look at the whole group as any trends would be very significant for overall food web dynamics in the region. It should be noted that the 200micron mesh size of the sampling net does not catch the C1 -C3 very well as they are too small. The series of Small Calanoid Copepods copepodite stages 1 to 6 (C1-6) has a stable seasonal pattern in terms of systematic changes according to the Resampling results.

Resampling Testing Stable vs Trending: $p = 0.271 \rightarrow$ Stable

The GCV score of the GAMs also points in the same direction:

GCV Score: Non-Seas: 1.777, **Stable:** 0.919 & Trending: 1.301

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 14.6, Stable: 57.6 & Trending: 68.2

Table 7.14: Small Calanoid Copepods C1-6 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	517.359	892.112				
SS	509.714	442.351	7.644	449.761	67.797	$< 2.2e - 16$
TS	411.838	331.725	97.877	110.626	1.403	0.013

The ANOVA, Table 7.14, points to a trending seasonal pattern while all other test results and criteria, Table 7.15, employed point to the stable seasonal model.

In Figure 7.21 the fitted values of the SS (red) and TS (blue) GAMs and the estimated terms of the SS GAM are shown. There are two main peaks in the seasonal pattern.

There is a fluctuation in the size of the seasonals, see Figure 7.21, but no change of a systematic nature occurs. This fluctuation could be the reason the ANOVA prefers the TS GAM.




Figure 7.21: The logged series of Small Calanoid Copepods C1-6 with the trend from the SS model, the stable seasonals and the logged series plotted as dots with the fitted values from the SS (red) and TS (blue) GAMs.

Table 7.15: Small Calanoid Copepods C1-6 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	9.641	17.286	115.162
AIC	1789.887	1436.193	1480.565
BIC	1831.01	1509.921	1971.766

Figures 7.22 and 7.23 show the diagnostic plots for the two models. The QQ plots and histograms betray a deviation from the assumption of normality towards the lower values of the series.

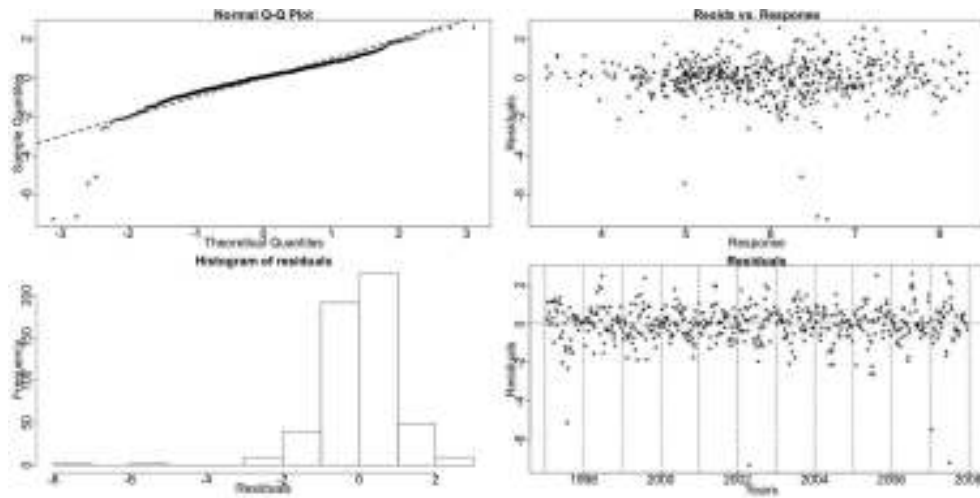


Figure 7.22: The diagnostic plots for the SS model of Small Calanoid CopepodsC1-6.

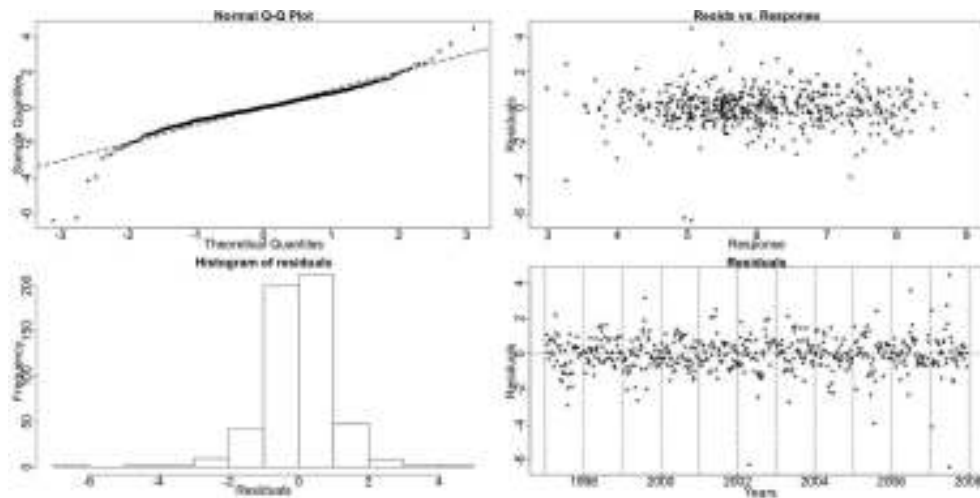


Figure 7.23: The diagnostic plots for the TS model of Small Calanoid CopepodsC1-6.

7.2.7 *Calanus finmarchicus*

Calanus finmarchicus (*C. finmarchicus*) is one of the most common copepods in the North Sea. It is an important component of the North Sea food web as its juveniles are food for a lot of commercial fish in spring and early summer. This species is largely herbivorous although able to feed on some microzooplankton and even its own eggs. Its life-cycle consists of egg, six naupliar stages and five copepodid stages (Hay, pers. comm., 2009). In stage 5 *finmarchicus* goes into diapause through the winter, descending to depths over 500m off the edge of the continental shelf, to emerge in early spring as a stage six adult.

The exact environmental conditions that trigger and stop the diapausing stage are not known. *C. finmarchicus* exhibits a reproductive peak in April till June while later the diapausing stages (C5 and some C6 females) sink to the bottom to over-winter. It favours cool waters 0° - 15° and it is believed to have no resident population in the North Sea (Bonnet *et al.*, 2005). The North Sea is shallow and thus, animals enter diapause partially which leads to low survival rates (Hirche, 1983). General belief is that the population is re-initiallized every year by inflow of adults from the North Atlantic.

In recent years research shows that changes of the climatological conditions impact on the number and distribution of *C. finmarchicus*. Work on the CPR⁴ data shows that in the last 40 years it has shifted progressively northwards with its numbers decreasing (Reid *et al.*, 2003).

Calanus finmarchicus copepodite stage 5

The *C. finmarchicus* C5 series exhibits two (major) peaks in the year. The first one can be attributed to population development within the influx of water from the North Atlantic while the next one is a first reproduced generation. Although as

⁴Continuous Plankton Recorder

with most copepods reproduction is continuous while food resources are sufficient, climatological factors influence the reproduction stages of *C. finmarchicus* so more generations can be produced under appropriate circumstances. These could manifest as other peaks within a year (Hay, pers. comm., 2009).

The results from the Resampling tests for trending seasonal pattern on the logged series suggest a systematic time shift in the occurrence of the peaks and troughs of the series, i.e. of the periodic component.

Resampling Testing Stable vs Trending: $p = 0.001 \rightarrow$ Trending

The GCV score of the GAMs, however, points to a stable seasonal pattern:

GCV Score: Non-Seas: 5.35, **Stable:** 0.439 & Trending: 0.550

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 0.762, Stable: 29.2 & Trending: 51.8

Table 7.16: *C. finmarchicus* C5 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	524	282.883				
SS	509.998	211.56	14.002	71.322	12.279	$< 2.2e - 16$
TS	415.402	144.086	94.596	67.474	2.056	$6.844e - 07$

The existence of trend in the seasonals is further verified by the GAMs' results. The ANOVA, Table 7.16, and AIC, Table 7.17, agree with the Resampling outcome while the more conservative BIC does not recognise the benefit of including trending seasonals in the model.

In Figure 7.24 the logged series of *C. finmarchicus* C5 with the SS terms and the fit from both the SS and the TS GAMs. The fit from the TS GAM seems better. Looking at the first plot in Figure 7.24 we notice the change in the size of the seasonals but in this case it is not big enough relatively to the rate of shift to mask it from the Resampling test.

The smooth trending seasonals as fitted to each week are shown in Figures 7.25




Figure 7.24: The logged series of *C. finmarchicus* C5 with the trend from the SS model, the stable seasonals and the series as dots with the fit from the SS (red) and TS (blue) GAMs.

and 7.26. Weeks 29-31 differ from a straight line parallel to the x-axis and are monotonic.

In Figure 7.27 we see the two seasonal patterns, stable (red) and trending (blue), estimated by the SS and TS GAMs, respectively, for the *C. finmarchicus* C5 series. There is a fluctuation in the size of the seasonals and a shift earlier in the year. We note that the pattern changes throughout the years and the second peak slowly becomes the dominant one. Specifically, the declining slope of the major peak and the small winter peak appear to be shifting earlier in time.

In Figure 7.28 we see a plot of the de-trended fitted values from TS GAM for *C. finmarchicus* C5 with weeks 17, 22, 34 and 39 marked by vertical lines in black, red, green and blue colour, respectively. We are using the vertical lines to highlight some peaks and troughs in the seasonals in order to observe the shift earlier in time. This means that we use the peaks and troughs to quantify the rate of change while marine biologists would probably use the time of first rise to the peak, however, this is hard to identify.




Figure 7.25: The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of *C. finmarchicus* C5.




Figure 7.26: The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of *C. finmarchicus* C5.




Figure 7.27: The stable (red) and trending (blue) seasonal patterns from the SS and TS GAMs, respectively, for *C. finmarchicus* C5.


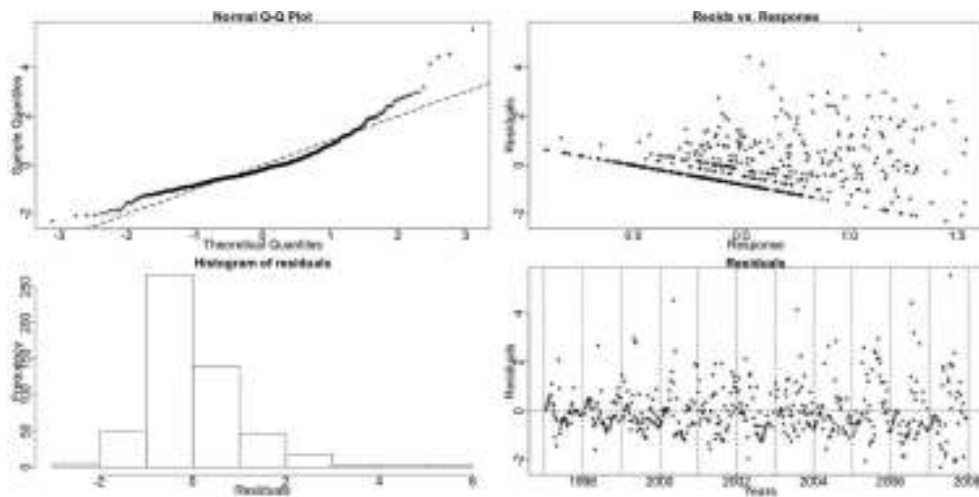


Figure 7.28: The de-trended fitted values from the TS GAM for *C. finmarchicus* C5. Weeks 17, 22, 34 and 39 are marked by dotted lines in black, red, green and blue, respectively.

Table 7.17: *C. finmarchicus* C5 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	17.002	111.598
AIC	1172.462	1047.651	1034.808
BIC	1185.258	1120.170	1510.809

Figures 7.29 and 7.30 depict the diagnostic plots for the SS and TS fitted GAMs. The QQ plots and histograms of the residuals show a right skewed distribution of residuals for both models. Nonetheless, there is an apparent gradual improvement as we move from the SS to the TS GAM.

**Figure 7.29:** The diagnostic plots for the SS GAM of *C. finmarchicus* C5.

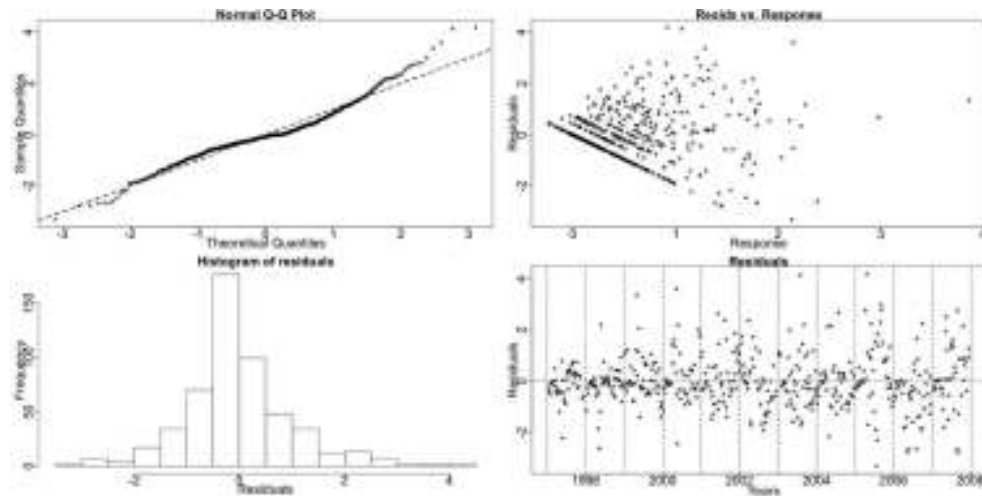


Figure 7.30: The diagnostic plots for the TS GAM of *C. finmarchicus* C5.

Calanus finmarchicus copepodite stage 6 females

The Resampling test results for trending seasonality for the *C. finmarchicus* C6f logged series show that there is no systematic time shift in its pattern.

Resampling Testing Stable vs Trending: $p = 0.087 \rightarrow$ Stable

The GCV score of the GAMs reinforces the belief of a stable seasonal pattern:

GCV Score: Non-Seas: 0.183, **Stable:** 0.162 & Trending: 0.226

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 2.88, Stable: 17.5 & Trending: 40.1

Table 7.18: *C. finmarchicus* C6f ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	524	95.327				
SS	516.436	80.943	7.564	14.384	12.133	2.851e-15
TS	414.763	58.834	101.673	22.109	1.533	0.002

Figure 7.31 shows the fit of the stable (red) and trending (blue) GAMs to the logged *C. finmarchicus* C6f series and the terms of the SS model.




Figure 7.31: The series of *C. finmarchicus* C6f with the trend from the SS model, the stable seasonals and the series marked by dots with the fit from the SS (red) and TS (blue) GAMs, appear clockwise from top left.

The ANOVA indicates the existence of a trend in the seasonals but it is of a random change in the amplitude of the pattern and not a systematic change. There is also the possibility that there is a small shift in time but it is only now starting and hence, it is not detected by the Resampling. In this occasion re-scaling the series to accommodate for the rapid changes in the amplitude of the seasonals and then testing again could help detect it.

Table 7.19: *C. finmarchicus* C6f AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	10.564	112.237
AIC	600.323	529.412	564.949
BIC	613.119	574.472	1043.672

The AIC and BIC scores, Table 7.19, verify our decision of stable seasonality.

The diagnostic plots for the GAMs are similar to the ones for *C. finmarchicus* C5.

Calanus finmarchicus copepodite stage 6 males

The results for *C. finmarchicus* C6m series are similar to the ones for C6f. The series has a seasonal pattern that does not exhibit a systematic shift in time.




Figure 7.32: Clockwise from top left: logged *C. finmarchicus* C6m with the trend from the SS model, the stable seasonals and the series as dots with the fitted values from the SS (red) and TS (blue) GAMs.

Figure 7.32 shows the SS (red) and TS (blue) GAMs' fitted values and the terms of the SS one. This series has a pseudo-trending in amplitude seasonal pattern.

7.2.8 *Calanus helgolandicus*

Calanus helgolandicus (*C. helgolandicus*) is a southern temperate species that has not been studied as extensively as *C. finmarchicus* and thus less is known about its life-cycle. It is not thought to have a diapause stage like *C. finmarchicus* but like its congener it is largely herbivorous (Hay, pers. comm., 2009). It favours warmer waters than *C. finmarchicus* but as the water temperature increases in the North Sea greater numbers of *C. helgolandicus* are observed. It has a resident population and at cold temperatures -around 6°C or less- it cannot reproduce and barely survives. It exhibits two peaks, one in spring and one in autumn but there is a variation in occurrence (Bonnet *et al.*, 2005).

Calanus helgolandicus copepodite stage 5

When the logged *C. helgolandicus* C5 series is tested there is evidence of a trending seasonal pattern. The Resampling p-value indicates that the pattern is not stable.

Resampling Testing Stable vs Trending: $p = 0 \rightarrow$ Trending

The GCV score of the GAMs, however, favours the SS GAM:

GCV Score: Non-Seas: 1.444, **Stable:** 0.931 & Trending: 1.144

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 11.2, Stable: 46.2 & Trending: 65.4

Table 7.20: *C. helgolandicus* C5 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	521.451	741.384				
SS	510.049	449.106	11.402	292.279	29.113	$< 2.2e - 16$
TS	410.449	288.613	99.601	160.492	2.292	6.247e-09

The ANOVA, Table 7.20, agrees with the Resampling result and points to the TS

model.

The BIC, Table 7.21, and the GCV score prefer the stable seasonal GAM while all other criteria used for GAM selection coincide with the above verdict of trending seasonality.




Figure 7.33: Clockwise from top left we see the logged *C. helgolandicus* C5 series with the trend from the SS model, the stable seasonals and the logged series as dots with the fitted values of the SS (red) and TS (blue) GAMs.

In Figure 7.33 we see the fits of the SS (red) and TS (blue) models and the terms of the SS one. The TS one provides a better fit.

The series has two main peaks, the first one (spring) can be attributed to influx of mixed coastal and oceanic water from the warmer south and west, since it overwinters with very poor survival while the second peak (early autumn) is the main generation population. There is a third peak in late autumn that could be another generation. All three peaks are shifting slightly earlier every year, Figure 7.33.

The smooth functions fitted to the trending seasonals by the TS GAM are shown in Figures 7.34 and 7.35. Weeks 29, 30, 32-37, 39, 41, 46 and 47 differ greatly from a straight line parallel to the x-axis. From these Figures we can also see the nature of the change. Weeks 29, 30, 32-37 and 47 exhibit a monotonic increase of value




Figure 7.34: The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of *C. helgolandicus* C5.

which points to a systematic change while weeks 39, 41 and 46 exhibit an increase and decrease in value which points to a fluctuation in either phase or amplitude of the seasonals.

Table 7.21: *C. helgolandicus* C5 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	5.549	16.951	116.551
AIC	1684.354	1443.494	1410.114
BIC	1708.022	1515.795	1907.240

In Figure 7.36 we see the two seasonals estimated by the SS and TS GAMs, respectively, fitted to the *C. helgolandicus* C5 series. The stable one is marked by a red line and the trending by a blue line. The pattern is shifting earlier in the year. The amplitude of the individual peaks changes randomly with time while overall the size of the pattern is increasing systematically.

The de-trended fitted values from the TS model are plotted in Figure 7.37. Weeks 16, 22, 37 and 46 are marked by vertical lines in black, red, green and blue colours, respectively. These vertical lines are marking the most prevalent peaks in the series.




Figure 7.35: The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of *C. helgolandicus* C5.

We can further verify with this plot that the peaks shift earlier in the year.

Figures 7.38 and 7.39 show the diagnostic plots for the two models. The models' QQ plots and histograms betray a slight deviation from the assumption of normality.




Figure 7.36: The stable, red, and trending, blue, seasonals estimated by the SS and TS GAMs, respectively, for the *C. helgolandicus* C5 series.




Figure 7.37: The de-trended fitted values from the TS model for the *C. helgolandicus* C5 series. Marked by the black, red, green and blue line are, respectively, weeks 16, 22, 37 and 46.

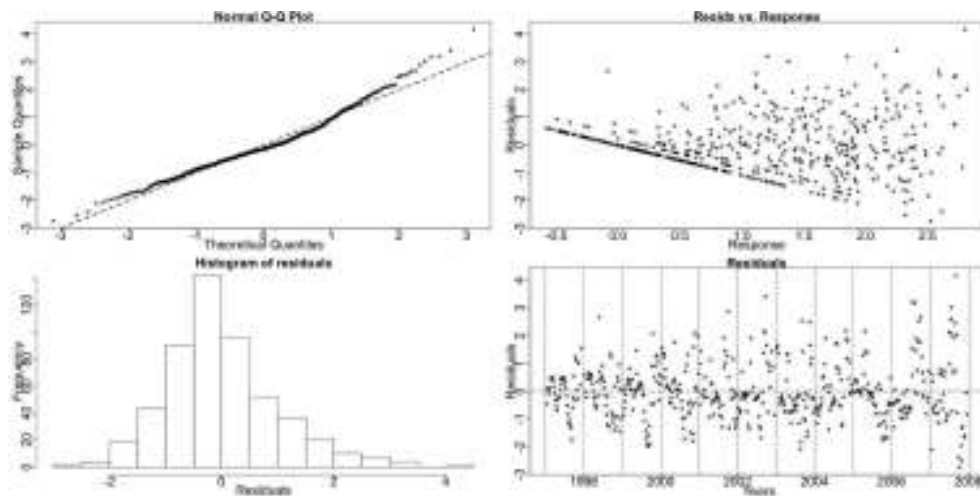


Figure 7.38: The diagnostic plots for the SS model of *C. helgolandicus* C5.

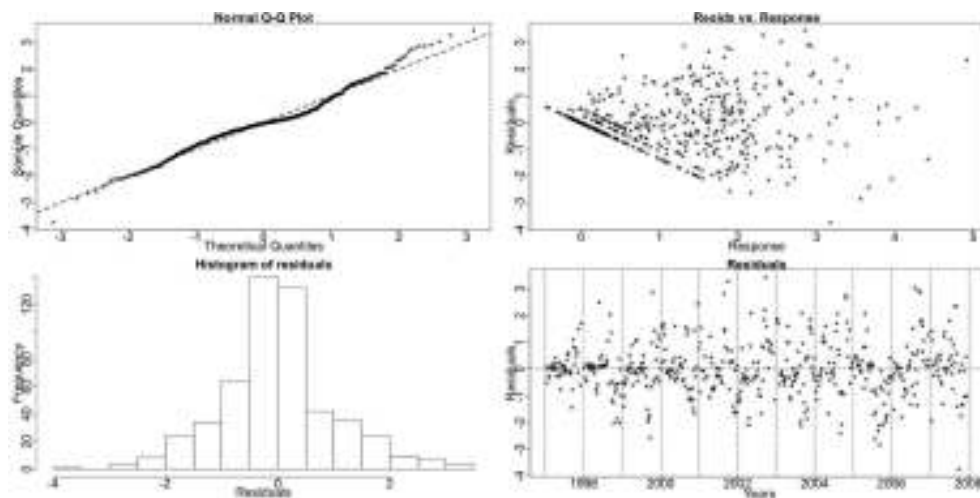


Figure 7.39: The diagnostic plots for the TS model of *C. helgolandicus* C5.

Calanus helgolandicus copepodite stage 6 females

The *C. helgolandicus* C6f series appears to have a trending seasonal pattern according to the Resampling results.

Resampling Testing Stable vs. Trending: $p = 0 \rightarrow$ Trending

The GCV score of the GAMs points to the SS GAM but the score of the TS one is only slightly greater:

GCV Score: Non-Seas: 1.228, **Stable:** 0.817 & Trending: 0.851

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 11.3, Stable: 43.2 & Trending: 69.9

Table 7.22: *C. helgolandicus* C6f ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	524	639.012				
SS	516.917	409.27	7.083	229.742	40.966	$< 2.2e - 16$
TS	411.811	216.791	105.105	192.479	3.479	$< 2.2e - 16$

The Resampling p-value, the ANOVA, Table 7.22 and the AIC, Table 7.23, all point to the existence of a trending seasonal pattern. Hence, the GAM that best describes the data is the trending seasonal model.

Figure 7.40 shows the fitted values of the SS (red) and TS (blue) models and the terms of the SS one. *C. helgolandicus* C6f exhibits a variety of peaks per year.

The smooth functions fitted by the TS GAM for the trending seasonals are plotted in Figures 7.41 and 7.42. Many weeks deviate from a straight line parallel to the x-axis, for example weeks 29-42. In particular, weeks 29-37 are monotonically increasing indicating a systematic change.

Figure 7.43 shows the stable (red) and the trending (blue) seasonals estimated by the SS and TS GAMs, respectively, for the *C. helgolandicus* C6f series. The pattern exhibits a systematic increase in amplitude and a systematic phase change.




Figure 7.40: The logged *C. helgolandicus* C6f series with the trend from the SS model, the stable seasonals and the series, this time marked by dots, with the fitted values from the SS (red) and TS (blue) GAMs.

Table 7.23: *C. helgolandicus* C6f AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	10.083	115.189
AIC	1601.094	1380.902	1256.871
BIC	1613.890	1423.910	1748.185

In Figure 7.44 we see a plot of the de-trended fitted values of the TS GAM fitted to the logged data. Weeks 30 and 46 are marked by vertical lines in black and red, respectively. These lines mark the autumn peak of *C. helgolandicus* C6f. The peak shifts earlier in the year while it increases with time.

Figures 7.45 and 7.46 show the diagnostic plots for the SS and TS models. Both models deviate slightly from normality but there is nothing alarming about these plots.




Figure 7.41: The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of *C. helgolandicus* C6f.




Figure 7.42: The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of *C. helgolandicus* C6f.




Figure 7.43: The stable (red) and trending (blue) seasonals from the SS and TS models, respectively, for the *C. helgolandicus* C6f series.




Figure 7.44: The de-trended fitted values from the TS model for the *C. helgolandicus* C6f series. Marked by the black and red lines are, respectively, weeks 30 and 46. The pattern shifts earlier in the year.

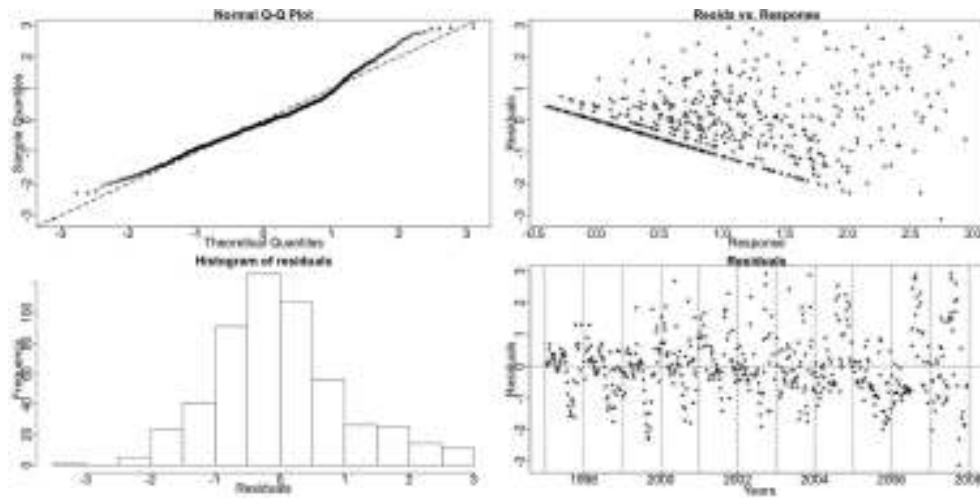


Figure 7.45: The diagnostic plots for the SS model of *C. helgolandicus* C6f.

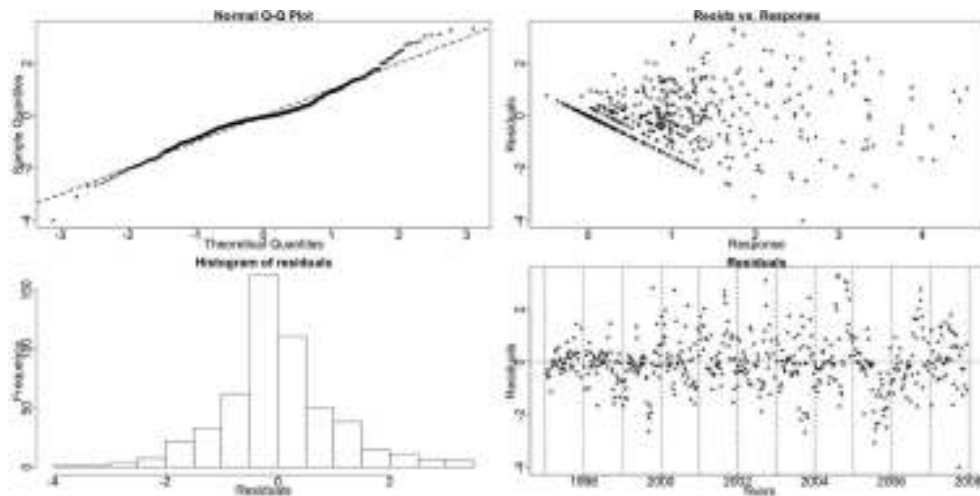


Figure 7.46: The diagnostic plots for the TS model of *C. helgolandicus* C6f.

Calanus helgolandicus copepodite stage 6 males

The *C. helgolandicus* copepodite stage 6 males exhibit a behaviour very similar to the females of that stage. Thus, the Resampling test points to a trending seasonal pattern.

Resampling Testing Stable vs. Trending: $p = 0 \rightarrow$ Trending

The GCV score of the GAMs reinforces the belief of a trending seasonal pattern:
GCV Score: Non-Seas: 0.658, Stable: 0.422 & **Trending:** 0.381

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 9.31, Stable: 44.1 & Trending: 73.5

Table 7.24: *C. helgolandicus* C6m ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	524	342.428				
SS	516.678	211.141	7.322	131.287	43.877	$< 2.2e - 16$
TS	415.934	100.196	100.744	110.945	4.572	$< 2.2e - 16$

Similar to the females of *C. helgolandicus* all testing methods, Tables 7.24 and 7.25, apart from BIC, point to a trending seasonal pattern.

In Figure 7.47 we can see the fitted values if the SS (red) and TS (blue) GAMs and the terms of the SS one. The TS model provides a better fit to the data.

Figures 7.48 and 7.49 show the fifty-two smooth functions fitted for the trending seasonals in the TS GAM. Weeks 28-39 deviate from a straight line parallel to the x-axis. The monotonic nature of weeks 29-37 indicates a systematic change.

To examine further the types of change in the pattern of the *C. helgolandicus* C6m series we plot the two patterns estimated by the SS and TS GAMs together, see Figure 7.50. The pattern changes systematically in amplitude and phase. The size




Figure 7.47: The logged *C. helgolandicus* C6m series with the trend from the SS model, the stable seasonals and the series, this time marked by dots, with the fitted values from the SS (red) and TS (blue) GAMs.

Table 7.25: *C. helgolandicus* C6m AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	3	10.322	111.066
AIC	1272.944	1033.247	842.657
BIC	1285.739	1077.274	1316.387

of the pattern increases over the years and the pattern itself is shifting earlier in the year.

Figure 7.51 shows a plot of the de-trended fitted values of the TS GAM fitted to the logged data. Week 31 is marked by a black vertical line while week 41 by a red one. The peak that develops between these two vertical line slowly shifts forward in the year while it increases in size.

Figures 7.52 and 7.53 show the diagnostic plots for the two GAMs. The QQ plots and histograms show a slight deviations from the normality assumption.




Figure 7.48: The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of *C. helgolandicus* C6m.




Figure 7.49: The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of *C. helgolandicus* C6m.




Figure 7.50: The stable (red) and trending (blue) seasonal patterns estimated by the SS and TS GAMs, respectively, for the *C. helgolandicus* C6m series.




Figure 7.51: The de-trended fitted values from the TS model for the *C. helgolandicus* C6m series. Marked by the black and the red lines are, respectively, weeks 31 and 41.

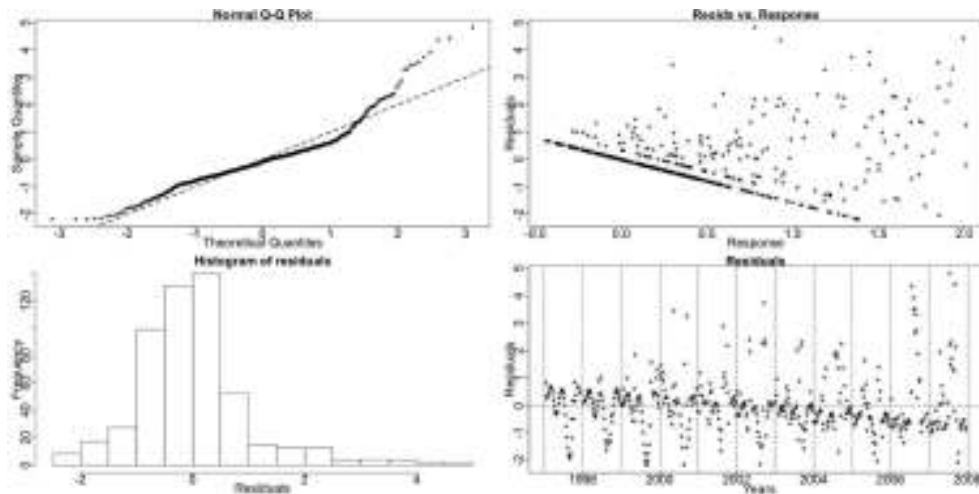


Figure 7.52: The diagnostic plots for the SS model of *C. helgolandicus* C6m.

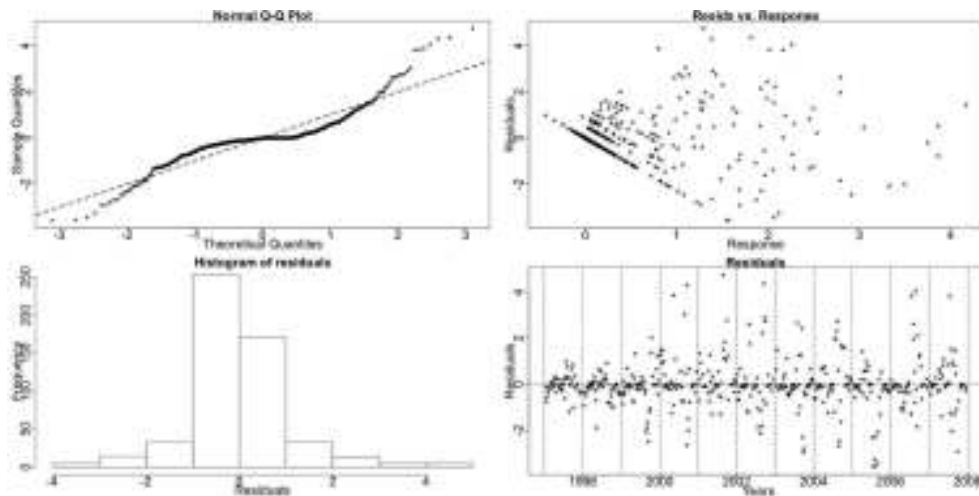


Figure 7.53: The diagnostic plots for the TS model of *C. helgolandicus* C6m.

7.2.9 Juvenile *Calanus* copepodite stages 1 to 4

Calanus in the early copepodite stages 1-4 (C1-4) cannot be accurately distinguished between the two dominant species, *finmarchicus* and *helgolandicus*. Therefore, the Juvenile *Calanus* stages 1-4 series includes juveniles from both species.

The Resampling test indicates that this series has a trending seasonal pattern.

Resampling Testing Stable vs Trending: $p = 0.002 \rightarrow$ Trending

The GCV score of the GAMs, however, is lower for the SS model:

GCV Score: Non-Seas: 2.471, **Stable:** 1.338 & Trending: 1.594

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 9.84, Stable: 53.4 & Trending: 70.8

Table 7.26: Juvenile *Calanus* C1-4 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	522.078	1272.507				
SS	513.365	657.018	8.713	615.49	55.193	$< 2.2e - 16$
TS	413.791	412.421	99.574	244.596	2.465	2.291e-10

The BIC, Table 7.27, and the GCV score select the stable GAM but the AIC, the ANOVA, Table 7.26, and the Resampling results point to a trending seasonality.

In Figure 7.54 we see the terms of the SS GAM and the fitted values of the the SS (red) and TS (blue) GAMs. The seasonal pattern appears to change in size.

Figures 7.55 and 7.56 show the fifty-two smooth functions fitted for the trending seasonals in the TS GAM. Weeks 26-36, 37 and 38 deviate from a straight line parallel to the x-axis. Weeks 27-36 are monotonic thus indicating a systematic change.

Figure 7.57 depicts the stable (red) and trending (blue) seasonals as estimated by the SS and TS GAMs, respectively, fitted to the Juvenile *Calanus* C1-4 series. The amplitude of the seasonals fluctuates. Both peaks appear to systematically shift




Figure 7.54: Clockwise from top left: the logged Juvenile *Calanus* C1-4 with the trend from the SS model, the stable seasonals and the logged series plotted as dots with the fitted values of the SS (red) and TS (blue) models.

Table 7.27: Juvenile *Calanus* C1-4 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	4.922	13.635	113.209
AIC	1967.258	1636.981	1591.188
BIC	1988.251	1695.139	2074.058

earlier in the year but it is more pronounced for the second peak. Furthermore, the amplitude of the second peak is gradually increasing while the amplitude of the first one fluctuates.

In Figure 7.58 we see a plot of the de-trended fitted values from the TS model. Weeks 14, 23, 32 and 41 are marked by dotted vertical lines in black, red, green and blue, respectively. The Juvenile *Calanus* C1-4 series has two main peaks, one in spring and one in autumn. The second peak, marked by the green and blue lines, moves earlier in the year. The same is true for the first peak, marked by the black and red lines, but the shift is not of the same rate, i.e. shifts slower.

It is likely, from what we know of the differences in C5 and adult seasonal abundances




Figure 7.55: The smooth trending seasonals for weeks 1-36 estimated in the TS model for the log of Juvenile *calanus* C1-4.




Figure 7.56: The smooth trending seasonals for weeks 37-52 estimated in the TS model for the log of Juvenile *calanus* C1-4.




Figure 7.57: The stable (red) and trending (blue) seasonal patterns from the SS and TS models, respectively, for the Juvenile *Calanus* C1-4.

and biology, that a significant number of juvenile *Calanus* in the first peak are *C. finmarchicus*, whereas the majority in the second peak are *C. helgolandicus*. This fits with the analysis results for the *C. finmarchicus* C5 trend and with the trends seen in all the *C. helgolandicus* data series. i.e. the *C. helgolandicus* trend indicates a more pronounced shift than is evident for *C. finmarchicus*.

Figures 7.59 and 7.60 show the diagnostic plots for the two models. Both models are satisfactory.




Figure 7.58: The de-trended fitted values from the TS model for the Juvenile *Calanus* C1-4. Vertical lines mark weeks 14, 23, 32 and 41 in colours black, red, green and blue, respectively.

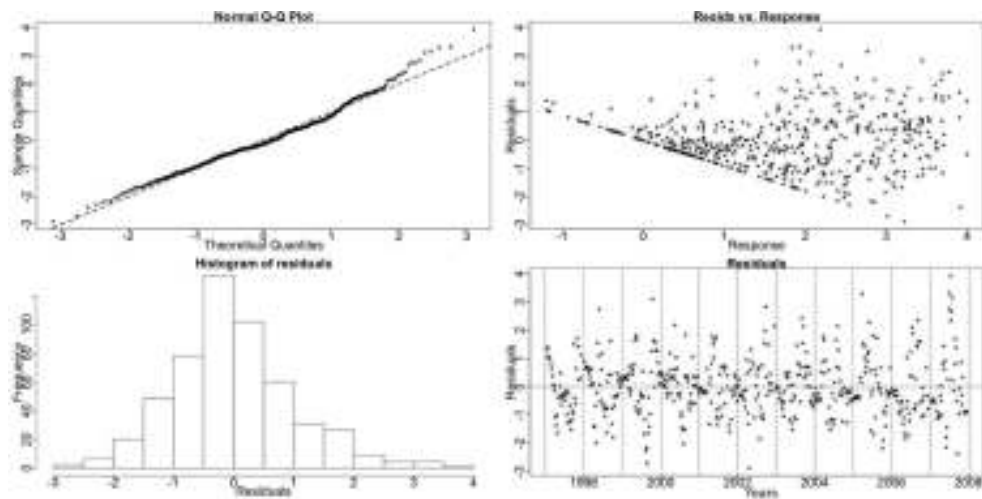


Figure 7.59: The diagnostic plots for the SS model of Juvenile *Calanus* C1-4.

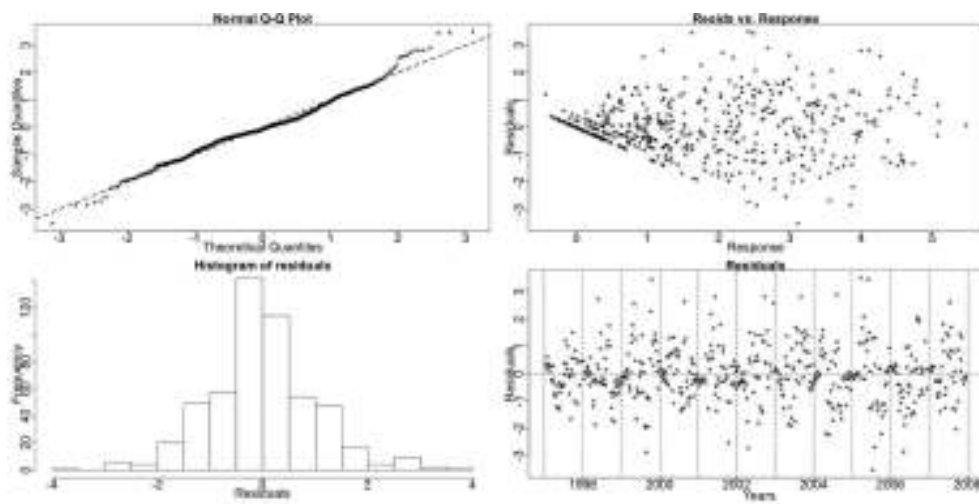


Figure 7.60: The diagnostic plots for the TS model of Juvenile *Calanus* C1-4.

7.2.10 *Oithona*

Oithona is a very small non-calanoid copepod belonging to the order Cyclopoida that is an omnivorous ambush feeder, feeding on large phytoplankton cells and microzooplankton. The *Oithona* genus is found very commonly in the North Sea and mainly as the very abundant species *Oithona similis*. Two other species of *Oithona* are found in smaller abundances. These are *O. nana*, a small inshore species, occasionally abundant, and *O. plumifera* that is an oceanic species and occasional visitor to the samples off Stonehaven. The adult copepods can be distinguished but juvenile stages are difficult to separate into species. As with the other small copepods the juvenile C1 - C3 stages are not sampled by the 200 micron plankton net with the same efficiency as the larger, later developmental stages (C4 - C6).

Nielsen & Sabatini (1996) state that the *Oithona* species do not exhibit as much variation as the calanoid copepods in terms of their biomass and production. We examine an aggregate series of *Oithona* stages C1-6.

Oithona copepodite stages 1 to 6

The *Oithona* C1-6 has a seasonal pattern that peaks in spring and summer. The Resampling recognises a stable seasonal pattern.

Resampling Testing Stable vs. Trending: $p = 0.083 \rightarrow$ Stable

The GCV score of the GAMs also points to a stable seasonality:

GCV Score: Non-Seas: 1.453, **Stable:** 1.295 & Trending: 1.687

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 45.9, Stable: 53.5 & Trending: 68.3

The Resampling p-value, the GCV score and the BIC suggest a stable seasonal pattern. However, the ANOVA, Table 7.28, and AIC, Table 7.29, indicate a trending model as appropriate.

Table 7.28: *Oithona* C1-6 ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	517.528	730.096				
SS	511.054	627.93	6.474	102.167	12.844	2.368e-14
TS	411.132	427.698	99.922	200.232	1.926	4.315e-06




Figure 7.61: Clockwise from top left: the logged *Oithona* C1-6 with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs.

In Figure 7.61 we see the logged *Oithona* C1-6 series, the terms and fitted values from the SS GAM and the fitted values from the TS GAM.

The seasonal pattern of the series changes randomly in size. This is probably the reason that ANOVA and AIC favour the TS GAM.

In Figures 7.62 and 7.63 the diagnostic plots for the two GAMs are shown. In both QQ plots there is a deviation from normality at the lower end of the series. This is seen as a long left tail in the histograms. The plots are nonetheless satisfactory.

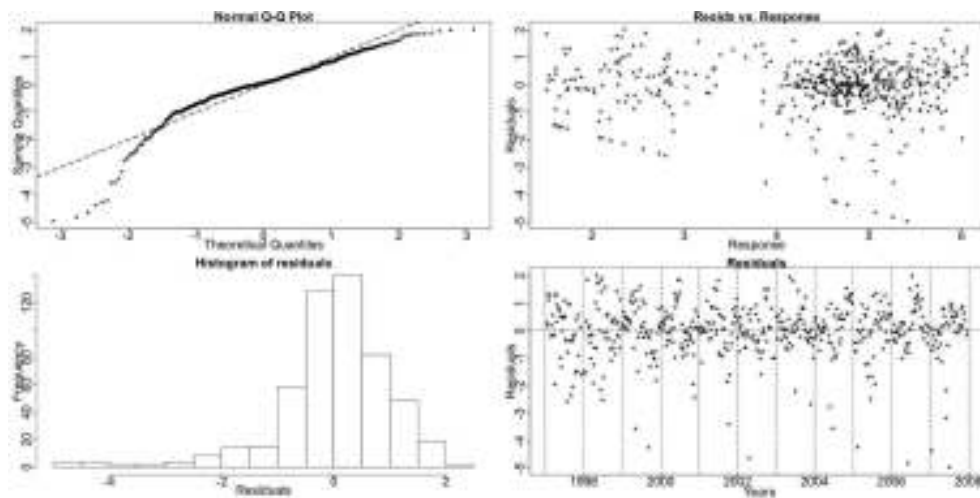


Figure 7.62: Diagnostic plots for the stable seasonal GAM for *Oithona* C1-6.

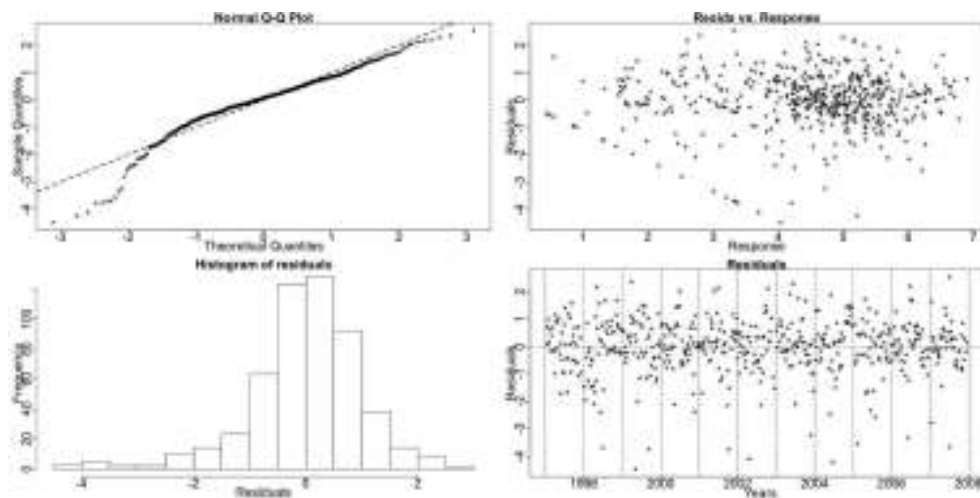


Figure 7.63: Diagnostic plots for the trending seasonal GAM for *Oithona* C1-6.

Table 7.29: *Oithona* C1-6 AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	9.472	15.946	115.868
AIC	1684.13	11617.784	1615.639
BIC	1724.53	1685.797	2109.85

7.2.11 *Lamellibranchiata*

Lamellibranchiata is a class of *Phylum Mollusca* and in the plankton samples it is the larvae of these bottom dwelling adult molluscs with opposing shells, that includes species such as scallops and cockles, that are caught in the sampler. There are a number of species of *lamellibranch mollusc* in the region and so, since their larvae cannot be distinguished easily even under the light microscope, they are aggregated into the general category here. It is a feature of the plankton that many bottom dwelling organisms produce planktonic larvae. Other common examples would be *echinoderms* (starfish), *polychaetes* (worms), *decapods* (crabs) and the *gastropod* (spiral shelled) molluscs. It exhibits one major peak in its seasonal pattern in the late summer. The Resampling find this series to have a stable seasonal pattern.

Resampling Testing Stable vs. Trending: $p = 0.568 \rightarrow$ Stable

The GCV score of the GAMs, also, chooses stable seasonality:

GCV Score: Non-Seas: 3.563, **Stable:** 1.521 & Trending: 1.806

The percentage deviance explained by each GAM is presented below:

% Deviance Explained: Non-Seas: 11.6, Stable: 63.7 & Trending: 77.7

Table 7.30: *Lamellibranchiata* ANOVA

	Resid. Df	Resid. Dev	Df	Deviance	F	Pr(>F)
NS	521.601	1830.72				
SS	514.772	752.702	6.83	1078.018	107.947	$< 2.2e - 16$
TS	412.525	462.851	102.247	289.851	2.527	4.891e-11

The series has a stable seasonal pattern in terms of systematic shifts. There is a fluctuation in the size of the seasonals for some years and that is probably the cause for the ANOVA, Table 7.30, and AIC, Table 7.31, choosing the trending model. Nevertheless, the Resampling, the BIC and GCV score indicate that a stable model is preferable.

In Figure 7.64 we see the logged *Lamellibranchiata* series, the terms and fitted values




Figure 7.64: Clockwise from top left: the logged *Lamellibranchiata* with the trend from the SS model, the stable seasonals and the logged series (dotted) with the fitted values of the SS (red) and TS (blue) GAMs.

from the SS GAM and the fitted values from the TS GAM.

Table 7.31: *Lamellibranchiata* AIC & BIC

	Non-Seasonal	Stable Seasonal	Trending Seasonal
df	5.399	12.228	114.475
AIC	2159.528	1705.682	1654.400
BIC	2182.555	1757.839	2142.670

In Figures 7.65 and 7.66 the diagnostic plots for the two GAMs are shown. Both models' QQ plots betray a deviation from normality at the extremes of the series.

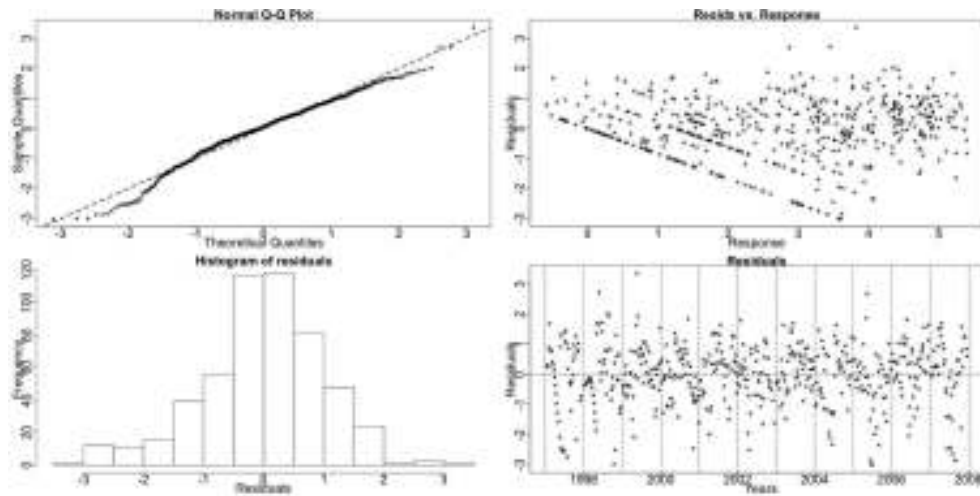


Figure 7.65: Diagnostic plots for the SS GAM for *Lamellibranchiata*.

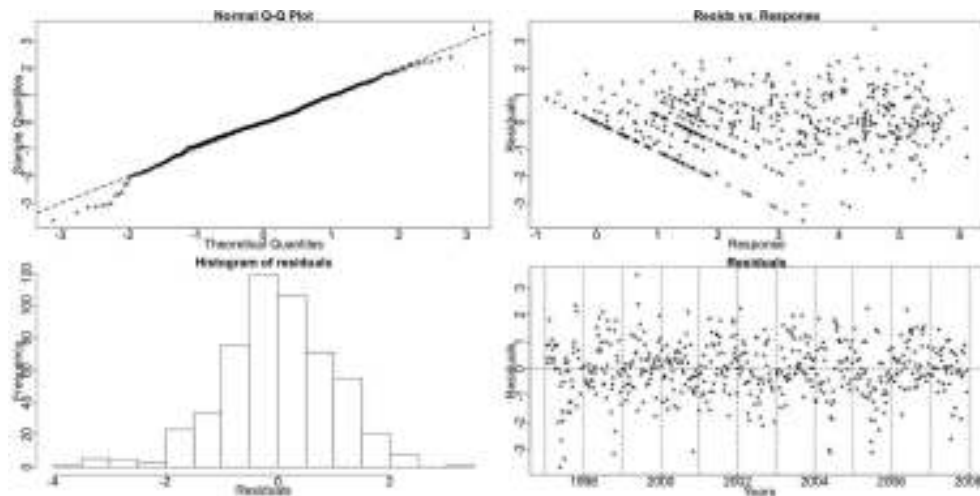


Figure 7.66: Diagnostic plots for the TS GAM for *Lamellibranchiata*.

7.3 Discussion

In this section we presented our analysis of some zooplankton species from the Stonehaven database. The series of *Acartia clausi* copepodite stage 6 males (C6m), *Calanus finmarchicus* C5 and *Calanus helgolandicus* C5 were found to have patterns systematically shifting earlier in time. *Calanus helgolandicus* C6f and C6m and Juvenile *Calanus* copepods C1 - 4 were recognised to have seasonal patterns whose amplitude is systematically increasing as well as shifting earlier in the year. These systematic changes may be linked to climate change. Especially the identified shifts in the patterns of the two main *Calanus* species are important because of their role as prey in the life cycle of commercially important fish such as cod and haddock.

There are a number of issues that arise from the analyses of the above zooplankton series. We are using the GAMs to model smooth functions that in some cases, for example the *Lamellibranchiata* series, do not seem to capture all the signal in the series. This can be seen in Figure 7.64, where the fitted values of both the SS and the TS GAMs do not go high enough into the peaks and low enough into the troughs of the series. The results of this ‘imperfect’ fit is seen as remaining pattern in the residual plots in Figures 7.65 and 7.66.

In addition, the GAMs do not fit well series that have high variability such as the *T. longicornis* C1 series. In Figure 7.16 the series exhibits high spikes. These large differences in the densities are hard to model, especially using smooth functions as we are.

Fitting the GAMs we are trying to model the periodicity in the series as a smooth cyclic component that exhibits smooth changes either systematic or fluctuations. We are interested in these smooth gradual changes because the effect of climate change and general long term environmental changes is expected to be slow and progressive. Nonetheless, for some of the examined series it can be seen that the periodic component is not very smooth as, especially in terms of the species series, the populations are sporadically distributed and affected by the currents. This means

that there is great variability within the samples.

This may affect the comparison of the GAMs when choosing the appropriate type of seasonality. As mentioned before, the main assumption underlying our testing techniques is that we are testing between the constructed models to infer about seasonality and thus if the models are not modelling each term, trend, stable and trending seasonality, correctly, our conclusions may be biased.

Finally, when a series has more zero values than non-zero ones, for example the *Acartia clausi* series, the GAMs with the log Normal distribution do not provide a good fit. As previously discussed, one could convert the series to a binary one (substituting all non-zero values with '1's) and fit GAMs with the Binomial distribution. This would be the same as what was done with some of the phytoplankton series in Chapter 6. It is our belief that a systematic shift in time would still be recognized by examining seasonal trends in the presence and absence of the series but nonetheless, some information would be lost.

Chapter 8

Discussion

8.1 Discussion on Results

In this thesis we have considered ways to examine the nature of seasonality in some marine biological time series and note the existence of systematic changes in them. Systematic changes would relate to climate change and its effects on marine life. Some commonly used testing methodologies, mostly from econometrics, are considered but are, however, not ideally suited to solve our problem. Thus, we introduced Resampling tests which are robust and simple to use. Additionally, we present a way to use GAMs to test for changes in the seasonals. The construction of the GAMs has to be modified appropriately depending on the nature and behaviour of the series but nonetheless, they have proved a useful tool.

A brief sensitivity analysis of both methodologies in order to appreciate their strengths and weaknesses was conducted. We have noted that the Resampling almost always detects a systematic change, never detects a random change, in time or in the size of the seasonals, and can fail to detect a systematic one if it is masked by a random change in the size of the seasonals or by noise. Furthermore, in the context of the

GAMs, we verified that the BIC is a more conservative criterion than the AIC, pointing more often to stable seasonality, and that the ANOVA is highly sensitive to any kind of change and favours a verdict of trending seasonality. In addition, one great difference between these two methods, Resampling and GAMs, is that the GAMs recognise a random change in the seasonals as trending seasonality while the Resampling does not. Overall the GAMs are more sensitive to any kind of change.

Using the above methodologies, we analysed time series from the Stonehaven long-term monitoring station and presented our results. We have established some environmental series that are shifting systematically earlier in time, Salinity, Ammonia and Nitrate while the *Chlorophyll a* series is shifting earlier in time and exhibits a systematic change in the size, decreasing. Furthermore, we examined some phytoplankton series from which none were found to exhibit systematic change. Finally, we found some zooplankton series, *Acartia clausi* C6m, *C. finmarchicus* C5, *C. helgolandicus* C5, C6f and C6m and *Juvenile Calanus* C1-4, that shift earlier in time. In particular, *C. helgolandicus* C6f and C6m and *Juvenile Calanus* C1-4, have patterns whose amplitude increases systematically in time.

8.2 Areas for Further Work

8.2.1 Statistical Interest

There are many ways in which the methodologies used can be improved and further developed. The Resampling methodology handles well smooth patterns, as for example the one from the Temperature data, but is not very accurate when used on densities or count data. Modifying the way that the seasonal decomposition is performed within this algorithm could greatly improve this issue. The Resampling algorithm as is explained here aims to be a simple yet robust testing framework. Nonetheless, it is possible to create different versions of it according to the data in question.

From our analysis of the phytoplankton and zooplankton series we know that the GAMs are not very appropriate when the pattern in a series is not smooth. Most series change smoothly in time, meaning that even when they include many zeroes and high peaks the mean changes gradually. Nonetheless, for some the peaks are very sudden. This kind of behaviour cannot be successfully modelled in terms of smooth functions, as in GAMs. Our suggestion was to convert the series into binary and still use GAMs but assuming the Binomial family. We believe this to have proven valid but still believe that a different method altogether could be devised for this type of series. Reducing the patterns of species to presence-absence causes a significant loss of information. In population studies it is often the rates of population growth or decline, implied or demonstrated by the abundance data, that may yield the best indicator of change or effect.

Another suggestion for future work is to attempt to combine the two methodologies together. Meaning that the residuals from the appropriate GAM could be used in the Resampling algorithm, treating GAM as a means to seasonally decompose the different types of series. Those residuals would then be repeatedly randomized and the above GAM fitted many times so that the residuals could be used for inference. The GAMs are computationally intensive and this was the main reason we did not try to combine the two methods. Nonetheless, this could prove to be a successful testing framework.

In this thesis we attempt to find ways to test for systematic changes in the seasonal pattern with specific interest in systematic time shifts. Nonetheless, since patterns in nature exhibit more than one change simultaneously we believe that it would be worth while trying to find ways to isolate the different types of change. This should be done after a specific type of change is verified. Isolating the type of change we are interested in by controlling for all others present would increase the reliability of the testing results as no interference would be present. For example, we know that systematic time shift cannot always be detected in the presence of random amplitude change and have thus suggested re-scaling the series in order to remove the effects of this random change.

A more in-depth sensitivity analysis of the methods would provide further insight into the ways they need to be improved. We have only examined a few interesting cases in Chapter 5 and feel that many more questions regarding the performance of the two methods could be answered by a more thorough sensitivity analysis. In particular, a sensitivity analysis combined with different types of controls used to isolate the different types of changes present in a series would be of great benefit. This would provide us with information on how the methods perform but also how successful our attempts to isolate the changes are.

The model selection for the GAMs is based on a number of criteria which, valid as they are, prove to be too sensitive or too conservative for our purposes. A more suitable criterion could be constructed to be used for model selection for similar to the above analyses. The effectiveness of this new criterion could be benchmarked against the other criteria with a sensitivity analysis.

Furthermore, when using GAMs we are modelling the mean of the series, for example of a Normal distribution. We could, however, attempt to model in a similar fashion the variability present in the series. Thus, for example, in addition to modelling the μ for a Gaussian GAM we could try to model the σ^2 . Creating again three models, NS, SS and TS, we could then choose the best model that describes the variance of the data. In the context of testing for seasonality both results would be taken into consideration for the final verdict of whether the series exhibits systematic changes or not. However, if one wished to use the GAMs as a means to forecast the future, the two chosen models, one for the mean and one for the variability could be used together to forecast future values. This, however, would not be straightforward as the way the results will be combined is not clear.

In our analysis, when a systematic shift in time is successfully detected we attempt to estimate the rate of change. Our method is an *ad hoc* one, examining major peaks, that does not correspond to the one marine biologists would often find more suitable in studies of species phenology, examining the timing of first occurrence. We use the occurrence of easily identifiable peaks to measure the shift while biologists might be

more interested in a shift in the initial rise of the peaks. As mentioned before, this is hard to identify in highly variable series and thus, was not used by us. Nonetheless, a more scientific way to measure the rate of shift, possibly producing C.I.s for it could be developed.

8.2.2 Biological Interest

We have examined the types of seasonality in a number of series from Stonehaven but have made no attempts to consolidate our results. Further investigation of the existing results would be useful to marine biologists. For example, when two series are shifting earlier in time it is possible that one is influenced by the other or that both are influenced by a third unobserved source.

Especially, since we have analysed environmental series which provide the condition for the blooms in marine species, as well as phytoplankton and zooplankton series, which are thought to be all part of the food chain, connections between the behaviours of some of these series are inevitable. A crude example would be that more Nitrate/ Nutrients in the water encourages the production of more phytoplankton which in turn could yield higher zooplankton levels. Else, as may become evident from fuller analysis of these results, those species which are at the edges of their adaptive ranges, such as the two *Calanus* species, may show greater sensitivity to changes in their environment; than species more fitted to the range of environmental factors measured. It is also the case that there may be other factors not yet considered that might affect species population behaviours and sensitivities.

In marine science many types of complex models connecting different series exist. Similar models could be developed to model series for which a connection is identified from extending our analysis. The purpose of these models would be to forecast future populations.

In this study we have focused on the most important and common species of zooplankton and phytoplankton. The Stonehaven long-term monitoring station has a

wealth of time series covering many other species and groups of organisms, that we, due to time constraints, did not analyze. Analyzing and comparing more series could help shed light on their behaviours and help marine biologists in their work.

In addition, there is a second Scottish monitoring station, in Loch Ewe, whose series do not run as far back as the Stonehaven ones but nonetheless, include almost the same variety of species. There are also similar time series in other regions (e.g. Helgoland station in the German Bight of the North Sea or station L4 in the English channel). These stations are set, so that apart from individual monitoring, their results can also be compared to draw conclusions based on similarities and differences in the behaviour of the various series in the different environments.

Finally, we have mentioned the fact that most of the series analyzed here are affected by many phenomena that are not examined. Series like river inflow and rainfall data are available and could be examined to enhance our study and look for correlations between them. Equally interesting results are anticipated if series like the North Atlantic Oscillation (NAO) and the Joint North Sea Information System (JONSIS, line east of Orkney Isles) are examined as the influences of the climatic and hydrographic factors they represent, on the distribution and population of mainly the zooplankton species, is currently much debated.

The time series analyses presented here consist a step towards moving from finding or expecting correlations between natural phenomena and understanding the causality and mechanisms involved. Such insights and models will be critical to developing understanding and management of ecosystems in the future.

Bibliography

- Abraham, B. & Ledolter, J. (1983). *Statistical Methods for Forecasting*. John Wiley and Sons, Inc.
- Amorim, A. L. (pers. comm., 2009). Marine Scotland.
- Anderson, J. & Bausch, C. (2006). Climate Change and Natural Disasters: Scientific evidence of a possible relation between recent natural disasters and climate change. In *DG Internal Policies of the UNION*, European Parliament. URL http://www.europarl.europa.eu/comparl/envi/pdf/externalexpertise/ieep_6leg/naturaldisasters.pdf. Retrieved: August, 5, 2009.
- Averkiev, A., Gustoev, D., Karpova, I., & Seryakov, E. (2005). Investigations and Forecasting of Long-Term Temperature Variations in the Kola Section. International symposium, Murmansk. PINRO.
- Barry, S. C. & Welsh, A. (2002). Generalized additive modelling and zero inflated count data. *Ecological Modelling*, 157, 179–188.
- Beaugrand, G. & Reid, P. C. (2003). Long-term changes in phytoplankton, zooplankton and salmon linked to climate. *Global Change Biology*, 9, 801–817.
- Beaulieu, J. J. & Miron, J. A. (1993). Seasonal unit roots in aggregate U. S. data. *Journal of Econometrics*, 55, 305–328.

- Bell, S. (2009). When the time is right: phenology and flowerings. *Significance*, 6(3), 113–116.
- Bochkov, Y. (2005). Large-Scale Variations in Water Temperature Along the ‘Kola Meridian’ Section and Their Forecasting. International symposium, Murmansk. PINRO.
- Bonnet, D. (2007). Comparative seasonal dynamics of *Centropages typicus* at seven coastal monitoring stations in the north sea, english channel and bay of biscay. *Progress In Oceanography*, 72(2-3), 233 – 248. URL <http://www.sciencedirect.com/science/article/B6V7B-4MVN0WW-1/2/e26c693f3dd54318fe354bde46f7dea7>. The Biology and Ecology of *Centropages typicus*.
- Bonnet, D., Richardson, A., Harris, R., Hirst, A., Beaugrand, G., Edwards, M., Ceballos, S., Diekman, R., *et al.* (2005). An overview of *Calanus helgolandicus* ecology in European waters. *Progress In Oceanography*, (65), 1–53.
- Box, G. & Jenkins, G. (1976). *Time series analysis: Forecasting and control*. San Francisco: Holden-Day, revised ed.
- Bresnan, E. (pers.comm., 2008). Marine Scotland.
- Bresnan, E., Hay, S., Hughes, S. L., Fraser, S., Rasmussen, J., Webster, L., Slesser, G., Dunn, J., & Heath, R. M. (2009). Seasonal and interannual variation in the phytoplankton community in the north east of Scotland. *Journal of Sea Research*, 61(1-2), 17–25.
- Burridge, P. & Taylor, A. M. R. (2001a). On regression-based tests for seasonal unit roots in the presence of periodic heteroscedasticity. *Journal of Econometrics*, 104, 91–117.
- Burridge, P. & Taylor, A. M. R. (2001b). On the properties of regression-based tests for seasonal unit roots in the presence of higher-order serial correlation. *Journal of Business & Economic Statistics*, 19(3), 374–379.

- Burridge, P. & Taylor, A. M. R. (2003). Bootstrapping the HEGY seasonal unit root tests. *Journal of Econometrics*, *123*, 67–87.
- Busetti, F. & Harvey, A. (2003). Seasonality tests. *Journal of Business & Economic Statistics*, *21*, 420–437.
- Busetti, F. & Taylor, A. M. R. (2003). Testing against stochastic trend and seasonality in the presence of unattended breaks and unit roots. *Journal of Econometrics*, *117*, 21–53.
- Canova, F. & Hansen, B. E. (1995). Are seasonal patterns constant over time? A test for seasonal stability. *Journal of Business & Economic Statistics*, *13*, 237–252.
- Carlotti, F. & Harris, R. (2007). The biology and ecology of *Centropages typicus*: An introduction. *Progress In Oceanography*, *72*(2-3), 117 – 120. URL <http://www.sciencedirect.com/science/article/B6V7B-4MX56CG-2/2/7bb50f019aafe283b94cd0e19b307eff>. The Biology and Ecology of *Centropages typicus*.
- Carpenter, S. R., Caraco, N. F., Correll, D. L., Howarth, R. W., Sharpley, A. N., & Smith, V. H. (1998). Nonpoint pollution of surface waters with phosphorus and nitrogen. *Ecological Applications*, *8*(3), 559–568.
- Chambers, M. J. & McGarry, J. S. (2002). Testing for seasonal unit roots by frequency domain regression. Working papers (Economics) 134, Loughborough University. URL <http://hdl.handle.net/2134/360>. Loughborough University, Working papers (Economics) : [134].
- Chatfield, C. (2000). *Time-Series Forecasting*. Chapman & Hall/CRC, first ed.
- Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). STL: A Seasonal-Trend Decomposition Based on Loess. *Journal of Official Statistics*, *6*(1), 3–73.

- Cleveland, W. S. (1981). LOWESS: A Program for Smoothing Scatterplots by Robust Locally Weighted Regression. *The American Statistician*, 35(No. 1), 54.
- Cochrane, J. H. (1988). How big is the random walk in GNP? *The Journal of Political Economy*, 96(5), 893–920.
- Cook, I. B., Smith, M. T., & Mann, E. M. (2005). The North Atlantic Oscillation and regional phenology prediction over Europe. *Global Change Biology*, 11, 919–926.
- da Silva Lopes, A. C. B. (2004). Deterministic seasonality in Dickey-Fuller tests: Should we care? *Econometrics* 0402007, EconWPA, RePEc:wpa:wuwpem:0402007. URL <http://ideas.repec.org/p/wpa/wuwpem/0402007.html>.
- Department for Environment, Food & Rural Affairs, DEFRA (2008). Climate change and energy. URL <http://www.defra.gov.uk/environment/climatechange/about/index.htm>. Retrieved: April, 13, 2008.
- Dickey, D. A. & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74, 427–431.
- Dickey, D. A., Hasza, H. P., & Fuller, W. A. (1984). Testing for unit roots in seasonal time series. *Journal of the American Statistical Association*, 79, 355–367.
- Dobson, A. J. (2002). *An Introduction to Generalized Linear Models*. Chapman & Hall/CRC, second ed.
- Durbin, J. & Murphy, M. J. (1975). "seasonal adjustment based on a mixed additive-multiplicative model". *Journal of the Royal Statistical Society: Series A (General)*, 138(3), 385–410.
- Faraway, J. J. (2006). *Extending the Linear Model with R*. Texts in Statistical Science Series. Chapman & Hall/CRC, Taylor & Francis Group.
- Findley, D. F., McElroy, T. S., & Wills, K. C. (2005). Modifications of SEATS'

- Diagnostic for Detecting Over- or Underestimation of Seasonal Adjustment Decomposition Components. U.S. Census Bureau.
- Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., & Chen, B. C. (1998). New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. *Journal of Business & Economic Statistics*, *16*(2), 127–152.
- Francis, J. G. F. (1961). The QR Transformation A Unitary Analogue to the LR Transformation - Part 1. *The Computer Journal*, *4*(3), 265–271.
- Francis, J. G. F. (1962). The QR Transformation - Part 2. *The Computer Journal*, *4*(4), 332–345.
- Friedman, J. H. (1984). A variable span scatterplot smoother. Technical Report No. 5, Laboratory for Computational Statistics, Stanford University.
- Fuller, W. A. (1976). *Introduction to statistical time series*. New York: Wiley.
- Guisan, A., Edwards, T. C. J., & Hastie, T. (2002). Generalized linear and generalized additive models in studies of species distributions: setting the scene. *Ecological Modelling*, *157*, 89–100.
- Harvey, A. C. (1990). *Forecasting, structural time series models and the Kalman filter*. Cambridge, U.K.:Cambridge University Press.
- Harvey, A. C. (1993). *Time Series Models*. MIT Press, second ed.
- Harvey, A. C. (2001). Testing in unobserved components models. *Journal of Forecasting*, *20*, 1–19.
- Hastie, T. & Tibshirani, R. (1990). *Generalized Additive Models*. Chapman & Hall.
- Hay, S. (pers. comm., 2009). Marine Scotland.
- Heath, R. M., Backhaus, O. J., Richardson, K., McKenzie, E., Slagstad, D., Beare, D., Dunn, J., Fraser, G. J., Gallego, A., Hainbucher, D., Hay, S., Jónasdóttir, S., Madden, H., Mardaljevic, J., & Schacht, A. (1999). Climate fluctuations and the

- spring invasion of the North Sea by *Calanus finmarchicus*. *Fisheries Oceanography*, 8(1), 163–176.
- Hirche, H. J. (1983). Overwintering of *Calanus finmarchicus* and *Calanus helgolandicus*. *Marine Ecology Progress Series*, 11, 281–290.
- Hood, C. C. (2002). Comparison of Time Series Characteristics for Seasonal Adjustments from SEATS and X-12-ARIMA. ASA proceedings, U.S. Census Bureau.
- Hylleberg, S. (1995). Tests for seasonal unit roots. General to specific or specific to general? *Journal of Econometrics*, 69, 5–25.
- Hylleberg, S., Engle, R. F., Granger, C. W. J., & Yoo, B. (1990). Seasonal integration and cointegration. *Journal of Econometrics*, 44, 215–238.
- IACMST (2001). Climate of UK Waters at the Millenium. Status and Trends. Tech. rep.
- IACMST (2005). Charting Progress - an Integrated Assessment of the State of UK Seas. Tech. Rep. 2 - Marine Processes and Climate.
- ICES (2006). Report of the ICES Advisory Committee on Fishery Mangement, Advisory Committee on the Marine Environment and Advisory Committe on Ecosystems. Book 6 North Sea, International Council for the Exploration of the Sea.
- ICES (2009). Report on Ocean Climate 2008. (298).
- Kim, Y. J. & Gu, C. (2004). Smoothing spline Gaussian regression: more scalable computation via efficient approximation. *Journal of The Royal Statistical Society, B(66)*, 337–356.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54, 159–178.

- Maddala, G. S. & Kim, I.-M. (1998). *Unit Roots, Cointegration, and Structural Change*. Cambridge University Press.
- Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and Applications*. John Wiley and Sons, Inc, 3rd ed.
- Maravall, A. (2002). Brief Description of the TRAMO-SEATS Methodology. *In Modeling Seasonality and Periodicity, Proceedings of the 3rd International Symposium on Frontiers of Time Series Modelling, Tokyo*.
- Marshall, T. (pers. comm., 2004). Kola data.
- McFedries, P. (2006). Changing Climate, Changing Language. URL <http://www.spectrum.ieee.org/energy/environment/changing-climate-changing-language>. Retrieved: June, 15, 2009.
- Menzel, A., Sparks, T. H., Estrella, N., & Eckhardt, S. (2005). ‘SSW to NNE’ – North Atlantic Oscillation affects the progress of seasons across Europe. *Global Change Biology*, (11), 909–918.
- National Research Council (2000). *Clean Coastal Waters: Understanding and Reducing the Effects of Nutrient Pollution*. National Academy Press, Washington, D. C.
- National Statistics Online (1997). Unemployment data. URL <http://www.statistics.gov.uk>. Retrieved: 1997.
- Nielsen, T. G. & Sabatini, M. (1996). Role of cyclopid copepods *Oithona* spp. in North Sea plankton communities. *Marine Ecology Progress Series*, 139, 79–93.
- Nixon, S. W. (1995). Coastal Marine Eutrophication: a definition, social causes and future. *Ophelia*, (41), 199–219.
- Ottersen, G., Edlandsvik, B., & Loeng, H. (2005). The Kola Section – An Indicator of Barents Sea Climate. Tech. rep., International symposium, Murmansk. PINRO.

- Parsons, R., T., Takahashi, M., & Hargrave, B. (1988). *Biological Oceanographic Processes*. Pergamon Press, third ed.
- Phillips, P. C. B. (1987). Time series regression with a unit root. *Econometrica*, *55*, 277–301.
- Phillips, P. C. B. & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, *75*, 335–346.
- Plymouth Marine Sciences Partnership (2009). Excellence in Marine Sciences, Education, Engineering and Technology. URL web.pml.ac.uk/pmsp/documents/P MSP-Brochure-new.pdf. Retrieved: July, 08, 2009.
- Pons, G. (2004). Testing monthly seasonal unit roots with monthly and quarterly information. *Journal of Time Series Analysis*, *27*, 191–209.
- Reid, P. C., Edwards, M., Beaugrand, G., Skogen, M., & Stevens, D. (2003). Periodic changes in the zooplankton of the North Sea during the twentieth century linked to oceanic inflow. *Fisheries Oceanography*, (12:4/5), 260–269.
- Said, E. & Dickey, D. A. (1984). Testing for unit roots in autoregressive moving average models of unknown order. *Biometrika*, *71*, 599–607.
- Sherry, A. R., Zhou, X., Gu, S., Ill, A., Schimel, S. D., Verburg, S. P., Wallace, L. L., & Yiqi, L. (2007). Divergence of reproductive phenology under climate warming. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(1), 198–202.
- Silverman, B. W. & Green, P. J. (1993). *Nonparametric regression and Generalized Linear Models. A roughness penalty approach..* No. 58 in Monographs on Statistics and Applied probability. Chapman & Hall/CRC, 1st ed.
- Smith, K., Webster, L., Bresnan, E., Fraser, S., Hay, S., & Moffat, C. (2007). A review of analytical methodology used to determine phytoplankton pigments in the marine environment and the development of an analytical method to determine

uncorrected chlorophyll a, corrected chlorophyll a and phaeophytin a in marine phytoplankton. Internal Report 03/07, Fisheries Research Services.

Stegert, C., Kreuz, M., Carlotti, F., & Moll, A. (2007). Parameterisation of a zooplankton population model for *Pseudocalanus elongatus* using stage durations from laboratory experiments. *Ecological Modelling*, 206(3-4), 213 – 230. URL <http://www.sciencedirect.com/science/article/B6VBS-4NYSHFX-1/2/2f2eb9cb73895181c07f880a0c2f1c14>.

Taylor, A. M. R. (2003). Robust stationarity tests in seasonal time series processes. *Journal of Business & Economic Statistics*, 21, 156–163.

UK Phenology Network (2009). Nature's Calendar. URL <http://www.naturescalendar.org.uk/>. Retrieved: June, 02, 2009.

University College London (2009). Microfossil Image Recovery And Circulation for Learning and Education. URL <http://www.ucl.ac.uk/GeolSci/micropal/diatom.html>.

Venables, W. N. & Dichmont, C. (2004). GLMs, GAMs and GLMMs: An overview of theory for applications in fisheries research. *Fisheries Research*, 70, 319–337.

White, H. (1984). *Asymptotic Theory for Econometricians*. New York : Academic Press.

Wood, S. N. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1), 95–114.

Wood, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Texts in Statistical Science Series. Chapman & Hall/CRC.

Wood, S. N. & Augustin, N. H. (2002). GAMs with integrated model selection using penalized regression splines and applications to environmental modelling. *Ecological Modelling*, 157(2-3), 157 – 177. URL <http://www.sciencedirect.com/science/article/B6VBS-470V3VF-5/2/4cc09a9ccbba184f3e295d456704e298>.

Yndestad, H., Turrell, W. R., & Ozhigin, V. (2008). Lunar nodal tide effects on variability of sea level, temperature, and salinity in the Faroe-Shetland Channel and the Barents Sea. *Deep-Sea Research, I*(55), 1201–1217.