Turid Follestad

Stochastic Modelling and Simulation Based Inference of Fish Population Dynamics and Spatial Variation in Disease Risk

Dr. Ing. Thesis

Department of Mathematical Sciences
Norwegian University of Science and Technology
2003

URN:NBN:no-3402

URN:NBN:no-3402

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree "Doktor ingeniør" (Dr. Ing.) at the Norwegian University of Science and Technology (NTNU). The research was financed by the Faculty of Information Technology, Mathematics and Electrical Engineering, NTNU.

First of all I would like to thank my supervisor Professor Håvard Rue for his inspiring guidance and continuous support. I also thank my co-supervisor Professor Arnoldo Frigessi, Norwegian Computing Center, for stimulating discussions.

During the academic year 2000/2001 I spent six months at the Department of Mathematics and Statistics, Lancaster University in England. I would like to thank Professor Peter Diggle for inviting me and for introducing me to the interesting field of spatial epidemiology, and the EU TMR network ERB-FMRX-CT96-0095 on "Statistical and computational methods for the analysis of spatial data" for financial support.

Finally, I thank all my colleagues at the Department of Mathematical Sciences, in particular the members of the statistics group, the administrative staff and those in charge of the computers, for providing a stimulating and supportive working environment.

Trondheim, February 2003

Turid Follestad

URN:NBN:no-3402

Thesis outline

The thesis consists of the following two self-contained parts:

PART I: **A Bayesian hierarchical model for the population dynamics of coastal cod**. Report.

PART II: **Modelling spatial variation in disease risk using Gaussian Markov random field proxies for Gaussian random fields**. *With Håvard Rue*. Report.

Both reports address the problem of sampling based inference in Bayesian hierarchical models, but within two different areas of application. While the study in Part I focus on modelling and inference for a problem in marine biology for which the data are temporally structured, the study in Part II is concerned with the analysis of spatially structured epidemiological data. The parameters of the models are estimated using Markov chain Monte Carlo methods (Gilks, Richardson and Spiegelhalter, 1996). The properties of the models in terms of mixing and convergence of the MCMC algorithms are studied empirically based on simulated data sets, utilising re-parameterisation (Part I) and blocking (Part II) to improve the performance of the parameter estimation procedures.

The major part of the thesis is Part I, where a Bayesian hierarchical model for the population dynamics of cod along the Norwegian Skagerrak coast is presented. The data on which the study is based have been collected as part of monitoring programme that was initiated by a controversy between a fisherman and a scientist during the first few years of the 20th century (Dahl and Dannevig, 1906). They could not agree on whether releasing cod larvae would enhance the population of cod within separate fjords along the coast, and to resolve the controversy, a release experiment followed by beach seine catches was conducted. The experiments and the related discussions motivated the establishment of a long-term monitoring programme that is still conducted (Smith, Gjøsæter, Stenseth, Kittilsen, Danielssen, Solemdal and Tveite, 2003). The data used in the study are samples of the two juvenile age-groups of cod at a set of locations in fjords along the coast, taken annually over a period of 43 years. The model is an extension of the models defined in Bjørnstad, Fromentin, Stenseth and Gjøsæter (1999) and Stenseth, Bjørnstad, Falck, Fromentin, Gjøsæter and Grey (1999), and includes an age-structured model for the dynamics of the population as well as a model for the sampling process generating the data. No data is available for the mature cod, but the mature abundance is represented by a latent variable that is linked to the juvenile cod abundances by a stochastic recruitment process. The parameters and abundances are estimated using a Metropolis-Hastings sampler. Applying the estimation procedure to a simulated data set, the mixing properties are shown to be improved by re-parameterising the model, although mixing problems persists for the adult cod, for which we have no data, and related parameters. Since the model parameters and not the abundances are the quantities of main interest in the study, an alternative approach based on sampling from the

marginal posterior distribution of the parameters, integrating over the abundances, is proposed. Based on the estimated model for the population dynamics, the release experiment from the early 20th century is imitated by conducting a simulation study of the effect of releasing juvenile cod.

In Part II, a spatially structured model for the geographical variation in disease risk is presented. Spatial variation in the risk of a disease is often represented by disease maps, displaying disease incidence of mortality rates in a set of disjoint regions. Since the populations in different regions often vary substantially in size, the variability of these raw estimates of disease risk differs between regions, such that spurious estimates of high risk in regions with low populations can mask the true spatial pattern of the risk. Starting with the works of Clayton and Kaldor (1987) and Besag, York and Mollié (1991) a number of authors have presented Bayesian approaches to the problem of smoothing the crude maps. Typically, the data underlying the maps are aggregated counts in administratively specified regions, and the resulting maps are based on estimates that are disjoint across region boundaries. However, in many cases it seems reasonable to assume that the risk of a disease vary smoothly across the region of interest. In the case of point data, for which the exact locations are known, the model-based geostatistical approach by Diggle, Tawn and Moyeed (1998) or the Poisson-Gamma random field model approach by Best, Ickstadt and Wolpert (2000) can be taken. Recently, Kelsall and Wakefield (2002) presented a geostatistical approach to estimating a smoothly varying surface based on aggregated count data, using an underlying Gaussian random field (GRF) to derive a log-Normal approximation to the joint distribution of relative risks at the regional level. In Part II of this thesis, an alternative approach is presented, based on specifying a GRF on a lattice, and fitting a Gaussian Markov random field (GMRF) to this GRF (Rue and Tjelmeland, 2002). Inference is based directly on the risk surface model as defined on the lattice, avoiding an approximation to the distribution of the regional risks. Utilising the computational convenient structure of the Markov random field, efficient algorithms for estimating the risk surface and the hyper-parameters of the model, based on the routines implemented in Rue and Follestad (2002), are developed. The properties of the sampling based estimation approach is studied empirically using two simulated data sets, and the method is applied to a set of data on mortality from oral cavity cancer in Germany.

References

Besag, J., York, J. and Mollié, A. (1991). Bayesian image restoration with two applications in spatial statistics (with discussion), *Annals of the Institute of Statistical Mathematics* **43**: 1–59.

Best, N. G., Ickstadt, K. and Wolpert, R. L. (2000). Spatial Poisson regression for health and exposure data measured at disparate resolutions, *Journal of the American Statistical Association* **95**: 1076–1088.

Bjørnstad, O. N., Fromentin, J.-M., Stenseth, N. C. and Gjøsæter, J. (1999). Cycles and trends in cod populations, *Proceedings of the National Academy of Sciences of the United States of America* **96**: 5066–5071.

Clayton, D. and Kaldor, J. (1987). Empirical Bayes estimates of age-standardized relative risks for use in disease mapping, *Biometrics* **43**: 671–681.

Dahl, K. and Dannevig, G. M. (1906). Undersøgelser over nytten af torskeudklækning i østlandske fjorde, *Aarsberetning vedkommende Norges fiskerier* pp. 1–121. (In Norwegian).

Diggle, P. J., Tawn, J. A. and Moyeed, R. A. (1998). Model-based geostatistics (with discussion), *Applied Statistics* **47**: 299–350.

Gilks, W. R., Richardson, S. and Spiegelhalter, D. J. (1996). *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, UK.

Kelsall, J. E. and Wakefield, J. C. (2002). Modeling spatial variation in disease risk: A geostatistical approach, *Journal of the American Statistical Association* **97**: 692–701.

Rue, H. and Follestad, T. (2002). GMRFLib: a C-library for fast and exact simulation of Gaussian Markov random fields, *Preprint series in statistics no. 1/2002*, Dept. of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

Rue, H. and Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields, *Scandinavian Journal of Statistics* **29**: 31–49.

Smith, T. D., Gjøsæter, J., Stenseth, N. C., Kittilsen, M. O., Danielssen, D. S., Solemdal, P. and Tveite, S. (2003). A century of manipulating recruitment in coastal cod populations: the Flødevigen experience, *ICES Journal of Marine Science*. In press.

Stenseth, N. C., Bjørnstad, O. N., Falck, W., Fromentin, J.-M., Gjøsæter, J. and Grey, J. S. (1999). Dynamics of coastal cod populations: intra- and intercohort density dependence and stochastic processes, *Proceedings of the Royal Society of London Series B* **266**: 1645–1654.

Part I

A Bayesian hierarchical model for the population dynamics of coastal cod

A Bayesian hierarchical model for the population dynamics of coastal cod

Turid Follestad

Department of Mathematical Sciences
Norwegian University of Science and Technology

URN:NBN:no-3402

Abstract

We present a non-Gaussian and non-linear state-space model for the population dynamics of cod along the Norwegian Skagerrak coast, embedded in the framework of a Bayesian hierarchical model. The model takes into account both process error, representing natural variability in the dynamics of a population, and observational error, reflecting the sampling process relating the observed data to true abundances. The data set on which our study is based, consists of samples of two juvenile age-groups of cod taken by beach seine hauls at a set of sampling stations within several fjords along the coast. The age-structure population dynamics model, constituting the prior of the Bayesian model, is specified in terms of the recruitment process and the processes of survival for these two juvenile age-groups and for the mature population, for which we have no data. The population dynamics is specified on abundances at the fjord level, and an explicit down-scaling from the fjord level to the level of the monitored stations is included in the likelihood, modelling the sampling process relating the observed counts to the underlying fjord abundances.

We take a sampling based approach to parameter estimation using Markov chain Monte Carlo methods. The properties of the model in terms of mixing and convergence of the MCMC algorithm are explored empirically on the basis of a simulated data set, and we show how the mixing properties can be improved by re-parameterisation. Estimation of the model parameters, and not the abundances, is the primary aim of the study, and we also propose an alternative approach to the estimation of the model parameters based on the marginal posterior distribution integrating over the abundances.

Based on the estimated model we illustrate how we can simulate the release of juvenile cod, imitating an experiment conducted in the early 20th century to resolve a controversy between a fisherman and a scientist who could not agree on the effect of releasing cod larvae on the mature abundance of cod. This controversy initiated the monitoring programme generating the data used in our study.

URN:NBN:no-3402

Contents

Chapter 1

Introduction

In this report we describe and explore a Bayesian hierarchical model for the population dynamics of cod along the Norwegian Skagerrak coast. The term population dynamics refers to the dynamic processes governing the life cycle of cod, including the recruitment of new individuals to the population, as well as the growth and survival of the fish at different stages of the life cycle. The main focus is on approaches to the estimation of model parameters, but the model can also be used to estimate the abundances of different age-groups of the population.

The data that are analysed have been collected as a part of a long-term and ongoing monitoring programme with its origins in the early 20th century. The monitoring programme was established following a controversy between a fisherman and a scientist who could not agree on the effect of releasing cod larvae on the abundance of mature cod within a fjord (Dahl and Dannevig, 1906). The investigations that were carried out to try to settle the question did not lead to any conclusion that could be agreed on by both parties, but the discussions following these initial investigations lead to the establishment of annual beach seine catches in a number of fjords along the southern coast of Norway. Similar catches are still being conducted, and they constitute the basis for our modelling.

When building fisheries models, there are two main sources of stochasticity to take into account, giving rise to process and observational error respectively (Polacheck, Hilborn and Punt, 1993). The process error represents natural variability in the dynamics of the population, while the observational error reflects the data collection process, relating the observed data to the actual populations. In general, state-space models (e.g. West and Harrison, 1997) represent a coherent approach to the specification of models that include both sources of error, an approach that has also been taken in the fisheries literature. A review of approaches to statistical modelling of fish population dynamics is given in Section 3.1.

We develop a non-Gaussian state-space model for the population dynamics of Skagerrak cod, embedded in the framework of a Bayesian hierarchical model. The model is an extension of the age-structured model of Bjørnstad, Fromentin, Stenseth and Gjøsæter (1999a), that will be described in Section 3.2. The prior model for the population dynamics, defining the system equation of the state-space model, should reflect the life cycle of cod, that can be described in terms of the rate of recruitment of new individuals to the population and rates of survival between different age groups. Of particular interest is the detection of any between-group and within-group interactions between juvenile stages of the cod. Another important issue is the specification of a model for the sampling process relating the observed data to the true underlying population abundances, and this model constitutes the observation equation of the state-space model. After establishing the model, we illustrate how we can use the model with a set of estimated parameters to simulate the effect of releasing juvenile cod on the mature cod abundance, imitating the release experiment that initiated the monitoring programme producing the data used in the study.

The parameters of the model are estimated by Markov chain Monte Carlo methods, based on annual data on two age groups of juvenile cod from a set of sampling sites in a selection of fjords along the coast. To assess the performance of our sampling algorithm, we apply the approach to a simulated data set. Using single-site Metropolis-Hastings sampling leads to slow convergence and mixing for several of the unknown parameters. This can be related to the facts that the number of parameters, including the annual abundances at each sampling site, is large, and that there are high correlations between some of the parameters. This often leads to slow convergence (see e.g. Gilks and Roberts, 1996). We illustrate how re-parameterising the model leads to improved mixing and convergence, but mixing problems persist for some parameters. However, our main interest is to study the processes governing the population dynamics, and stock assessment is only a secondary goal. Therefore, we suggest an alternative approach that can be taken to estimate the parameters of the model, using a Metropolis-Hastings algorithm sampling from the marginal posterior distribution of the model parameters given the observed counts, marginalising over the unknown populations.

The report is organised as follows. In Chapter 2 we give some biological and historical background on the Flødevigen data set and present the data. A review of general approaches to the modelling of fish population dynamics as well as of previous work on the Flødevigen data is given in Chapter 3. In Chapter 4 the models for the population dynamics and the sampling process are defined, and in Chapter 5 we present the sampling based approach to estimation of parameters and abundances, including results for a set of simulated data as well as the Flødevigen data. In Chapter 7 we describe how the effect of releasing juvenile cod can be simulated using the model. An alternative approach to parameter estimation based on sampling form the marginal posterior distribution, integrating over the abundances, is described in Chapter 8. A summarising discussion is given in Chapter 9.

Chapter 2

Biological and historical background on the Flødevigen data

In this chapter, we give a brief overview of the biology of the cod and some historical background on the data on which our study is based.

2.1 The cod life cycle

The life cycle of the Norwegian Skagerrak cod (*Gadus morhua L.*) can be divided into four main stages (Gjøsæter, Stenseth, Ottersen, Lekve, Dahl, Danielssen, Torstensen and Christie, 2003a; Stenseth, Bjørnstad, Falck, Fromentin, Gjøsæter and Grey, 1999), as illustrated in Figure 2.1. The cod along the Skagerrak coast normally spawn in late winter or early spring (February, March or April), and the eggs are hatched one or two weeks after spawning. The larvae stage lasts until May-June, when the larvae metamorphose into juveniles (small fish), constituting the *0-group* cod of the autumn that year. One year later, at the age of about 1.5 years, the juveniles are referred to as the *1-group* cod. The Skagerrak cod is assumed to mature at the age of two to three years (Gjøsæter, Enersen and Enersen, 1996). We will assume that all cod at the age of two years and older belong to the spawning population.

Based on results from a tagging experiment in the late 1980's, Danielssen and Gjøsæter (1994) conclude that the Norwegian Skagerrak cod populations seem to be relatively isolated and non-migratory. The 0-group and parts of the 1-group cod are generally found closer to the shore in more shallow water than the adult cod, but they tend to move to deeper water during summer and winter, probably due to variations in temperature (Danielssen and Gjøsæter, 1994; Gjøsæter and Danielssen, 1990, and references therein). The juvenile cod tend to settle where they find food and places to hide from

predators (Johannessen and Sollie, 1994). In a study of Atlantic cod by Tupper and Boutilier (1995), the results suggest that cod settle in any habitat, but that growth and survival is best in areas where they find shelter.

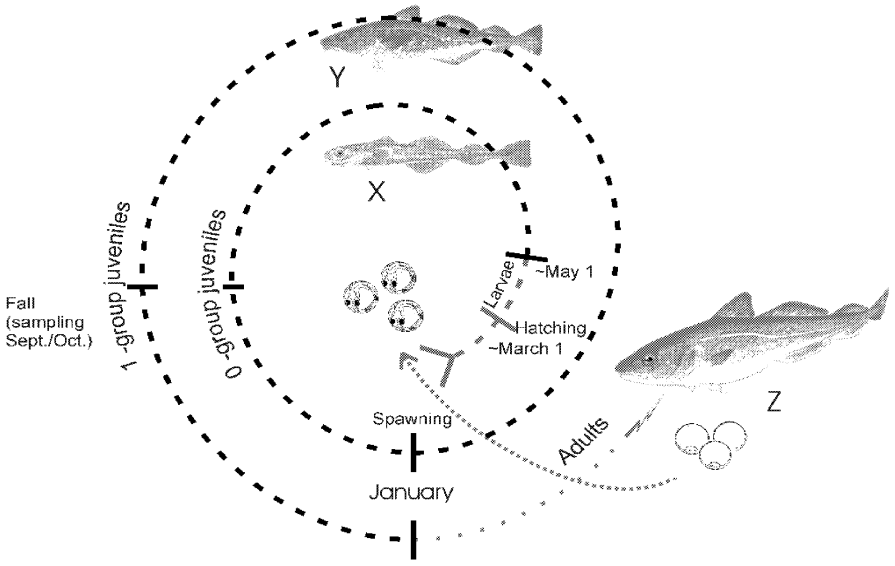

Figure 2.1: The life cycle of cod (Ill.: Ottar N. Bjørnstad, Department of Entomology, Pennsylvania State University).

2.2 Historical background

The data that make up the basis of the study described in this report, is a subset of the data from a unique monitoring programme that has been run at the Norwegian Skagerrak coast in its present form since 1919; see Smith, Gjøsæter, Stenseth, Kittilsen, Danielssen, Solemdal and Tveite (2003) for a recent review. The monitoring programme was initiated following a controversy between Gunder Mathiesen Dannevig (1841-1911), a fisherman and a captain of sailing vessels, and Johan Hjort (1869-1948), a pioneer in marine research, about the effect of releasing hatched cod larvae on the size of the population of adult cod in a restricted area like a fjord.

Dannevig (Figure 2.2, left) was the founder of the Flødevigen Hatchery (later the Flødevigen Marine Research Station), near Arendal in Southern Norway, which was completed in 1884 (Solemdal, Dahl, Danielssen and Moksness, 1984). He argued that there was a deterministic relationship between the number of yolk-sac cod larvae and the

Figure 2.2: Gunder Mathiesen Dannevig at his desk at Flødevigen (left), and Johan Hjort (right) (Institute of Marine Research, Flødevigen Marine Research Station).

number of recruits to the population, and as a consequence, releasing larvae would supplement the free living fish populations (Dahl and Dannevig, 1906). In particular, releasing larvae could thus counteract the effect of over-fishing. His theory was based on the opinion that the cod appeared in local fjord populations with its own egg and larval production, a point of view that was in agreement with observations by fishermen who claimed that the cod in different fjords were of different types. Thus, the released larvae would be relatively abundant enough to enhance the local cod population.

Hjort (Figure 2.2, right) disagreed with Dannevig about the existence of local fjord populations, and was of the opinion that the coastal cod belonged to one large oceanic population that migrated over large areas. He argued that the larvae would be carried by the stream between the fjords and also into the open Skagerrak sea. As a consequence, he found it useless to try to enhance the cod population by releasing cod larvae in a fjord, since the number of released larvae would be relatively small compared to the number of naturally hatched larvae (Dahl and Dannevig, 1906). He was also of the opinion that annual fluctuations in the stock was mainly due to natural variation. Some time during the first few years of the 20th century however, Hjort gave up his theory that there is one large migrating cod population, and instead opposed Dannevigs opinion on the effect of hatching by arguing that artificially hatched larvae was no more capable of living than larvae naturally hatched in the fjords (Schwach, 2000).

As a side comment, it is worth mentioning that the controversy between the scientist Hjort and the practitioner Dannevig is also believed to be motivated by political reasons, related to different views on the future of the Norwegian fisheries. Dannevig wanted to preserve the traditional coastal fisheries, while Hjort's vision was to modernise the fishing industry, and he was of the opinion that artificial hatching would hamper this process of modernisation (Schwach, 2000).

To be able to solve the controversy, and to conclude whether or not releasing hatched larvae was useful, Dannevig proposed to make a series of beach seine hauls in some selected fjords in Southern Norway. After several proposals during the 1890's and early 1900's were rejected, the Norwegian Government eventually decided to fund the hauls. In the years 1903-1905, Dannevig, assisted by a member of Hjort's staff, Knut Dahl, made hauls in two selected fjords, before and after releasing hatched larvae. Dahl also made hauls in nearby fjords with no release. However, the experiment did not help to settle the disagreement between the two parties, since Dannevig and Dahl arrived at contradictory conclusions based on the collected data. Dannevig had, although reluctantly, agreed that Dahl could conduct additional investigations in the fjords where the larvae were released. As described in Dahl and Dannevig (1906), these investigations included studies of the hydrographical conditions of the fjords as well as biological examinations of the cod. Dahl wanted to take into account the results of these additional investigation as well as the results from the fjords with no release in their final report on the study, but Dannevig disagreed, since he was of the opinion that they were not relevant to the problem. As a consequence, the final report (Dahl and Dannevig, 1906) is given as separate contributions from the two. The translated extracts from the conclusions quoted below (Solemdal et al., 1984) should clearly illustrate their lack of agreement. Based on the results from the release experiment, as well as from his additional investigations, Dahl argued that the observed increase in the number of juvenile cod after releasing larvae could not be distinguished from natural fluctuations as observed by an increasing number of juveniles in nearby fjords as well. Some of his results were summarised as illustrated in Figure 2.3. Thus, he concluded that

> *"In my opinion, it has already been proved that such a release (of yolk sac larvae) is not capable, to any provable extent, of affecting the natural stock of yolk sac larvae of even a small and limited area, and even less able of increasing the fishery in such an area."*

On the other hand, Dannevig concluded that

> *"The result of the now completed investigations is new evidence of the accuracy of our theory so often emphasized: that artificial hatching is not merely the most important, but also the only means for maintenance and improvement of the fish stocks in our fjords, and also that the money is more profitably used in order to increase the hatching production than spent on investigations for which the result is known in advance."*

The investigations did not resolve the controversy, but the study, and the general discussion on the effect of hatching, is believed to be one of the factors leading to the further research by Johan Hjort on the causes of the fluctuations in the fisheries (Schwach, 2000). The research resulted in his theory that these fluctuations were due to natural variation

in the strength of different year-classes of juveniles, where a year-class of fish is the group of fish that were spawned in the same year. Further, the environmental conditions during the critical early phases of the development of the larvae were believed to be essential in determining the strength of a year-class (Hjort, 1914; Hjort, 1926). Today, this theory is considered fundamental to fisheries research (Schwach, 2000; Myers and Cadigan, 1993b).

The 1903-1905 beach seine hauls, and the failure to resolve the controversy based on the results from these hauls, motivated establishing a long-term monitoring programme along the same lines, but extending the number of fjords. From 1919 the monitoring programme was established in it's present form (Solemdal et al., 1984). The Flødevigen hatchery was taken over by the Norwegian government in 1918, and is from 1973 a laboratory of the Norwegian Institute of Marine Research. A brief description of the present beach seine sampling procedure is given in Section 2.3.

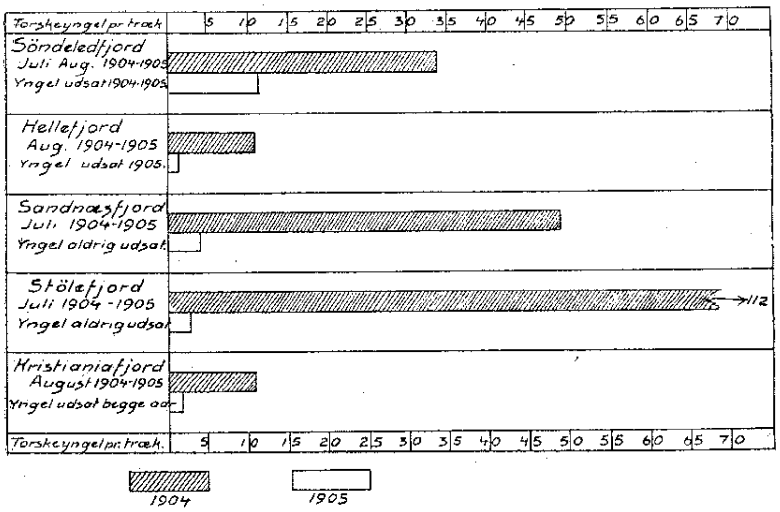

Fig. 8. Grafisk fremstilling af mængden af torskeyngel pr. vadtræk i de undersøgte fjorde juli—august 1904—1905.

Figure 2.3: Illustration of some of the results from the release experiment, according to K. Dahl (Dahl and Dannevig, 1906, p.76). The horizontal bars represent the number of juveniles per haul in five fjords in the autumn of 1904 (shaded) and 1905. Larvae were released in the spring of 1904 in the first two fjords and both years in the last fjord.

Age-group	Mean	Med	Min	Max	Proportion of 0 counts (%)
0-group	13.3	5	0	355	23
1-group	1.98	1	0	34	49

Table 2.1: A summary of the data in the two Risør fjords.

2.3 The data

At present, about 120 stations are sampled, among them 33 have been sampled since the beginning in 1919. The sampling is performed by beach seine hauls between the beginning of September and mid October each year, and follows a strict procedure (see e.g. Gjøsæter, Stenseth, Sollie and Lekve, 2003b). The sampling procedure is illustrated in Figure 2.4, and pictures from the beach seine hauls are given in Figures 2.5 and Figures 2.6. The beach seine catches are operated from a specially designed boat. One person holds one of the ropes of the beach seine while the boat with the seine is rowed along a line as shown in Figure 2.4 (*A*). Another person takes the other rope, and goes ashore at a distance of around 5-15m from the first person (*B*). Then the beach seine is hauled, and the number of individuals of each species are counted (*C*). The juvenile cod is categorised as 0-group or 1-group cod according to their length. The beach seine is $40m$ long and $3.7m$ deep, each haul covers an area of up to $700m^2$, and the depth sampled varies from 3 to $15m$.

In our study we use a subset of the data consisting of catches taken over the period from 1956-1998 in 49 stations within 13 fjords. A map of the sampling region is given in Figure 2.7, and a table of the fjords and stations used in the study as well as plots of the 49 time series of abundances of 0-group and 1-group cod are given in Appendix A. We also select a smaller subset of data to be used as a basis for studying the properties of the statistical model. This dataset, which we will denote the Risør data, is made up from samples from 12 stations in two fjords in the Risør area, Sandnesfjord and Søndeledfjord, shown in the bottom panel of Figure 2.7. We also restrict ourselves to the 38 years between 1957 and 1994. Consequently, this subset consists of 456 observed counts of 0-group and 1-group cod. We will treat the 12 stations as belonging to one single fjord by modelling the total abundances of the different age-groups within the two nearby fjords. Plots of the data, aggregated within fjords and over the fjords, are given in Figure 2.8, while the 12 separate time series of 0-group and 1-group data can be found in Figures A.4, A.5 and A.6 in Appendix A. A summary of the station specific counts is given in Table 2.1. Observe that a substantial fraction of the observed counts are zero.

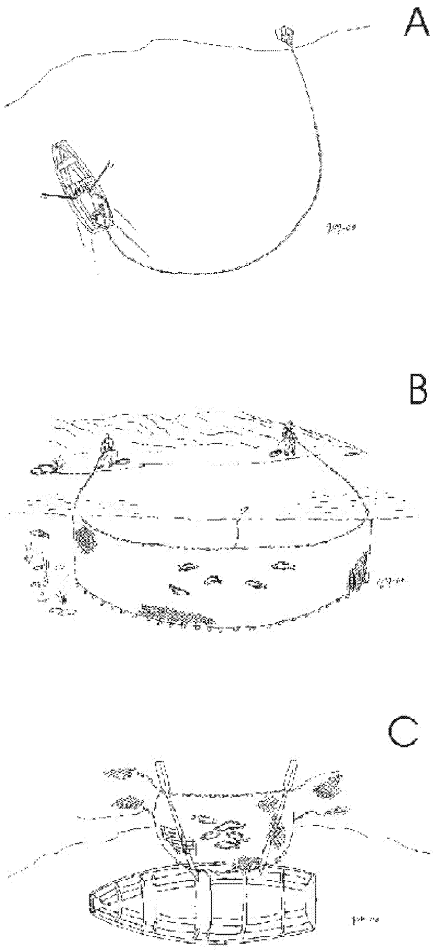

Figure 2.4: Illustration of the sampling procedure for a sampling station (Institute of Marine Research, ill.: Vetle Madsen).

Figure 2.5: The sampling procedure in early years of the study, illustrated by a picture from a beach seine haul in 1907 (Institute of Marine Research). Captain Dannevig can be seen in the middle of the picture, inspecting the beach seine hauls.

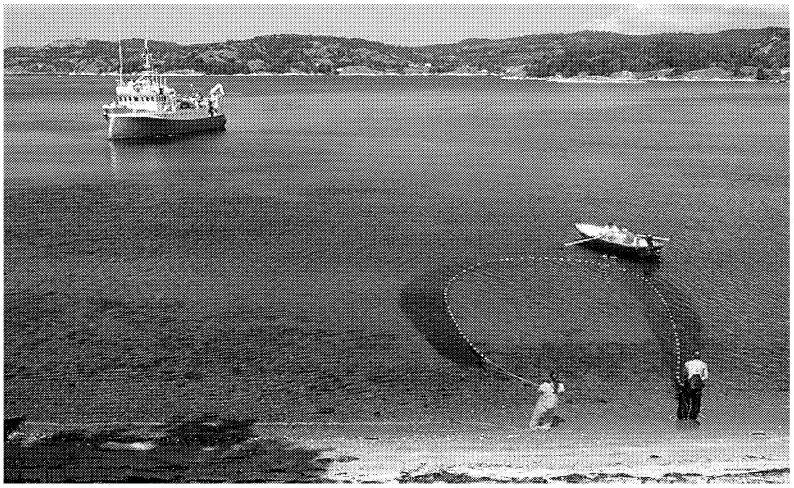

Figure 2.6: The sampling procedure of today. (Institute of Marine Research. Photo: Øystein Paulsen).

The Skagerrak coast

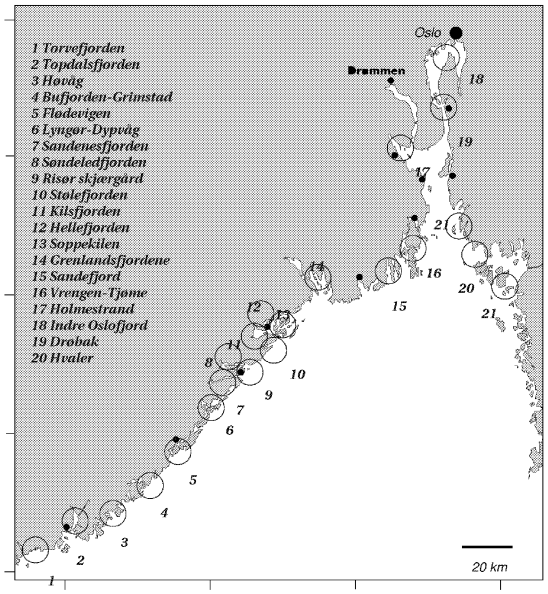

The Risør fjords

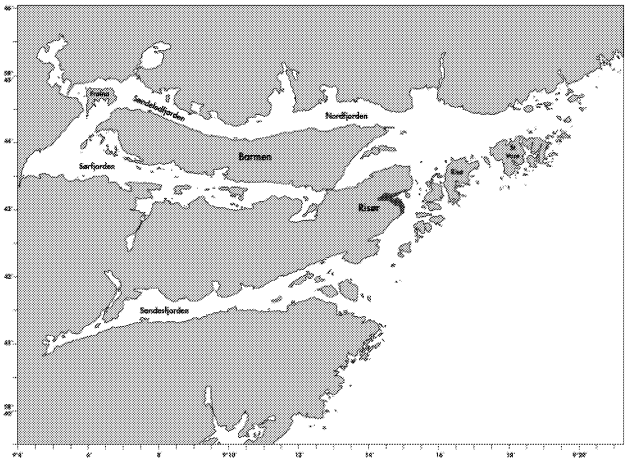

Figure 2.7: The study region, the whole area (top) and the sub-area made up from the Risør fjords (bottom), numbered 7 and 8 in the top map. (The maps are provided by the Institute of Marine Research).

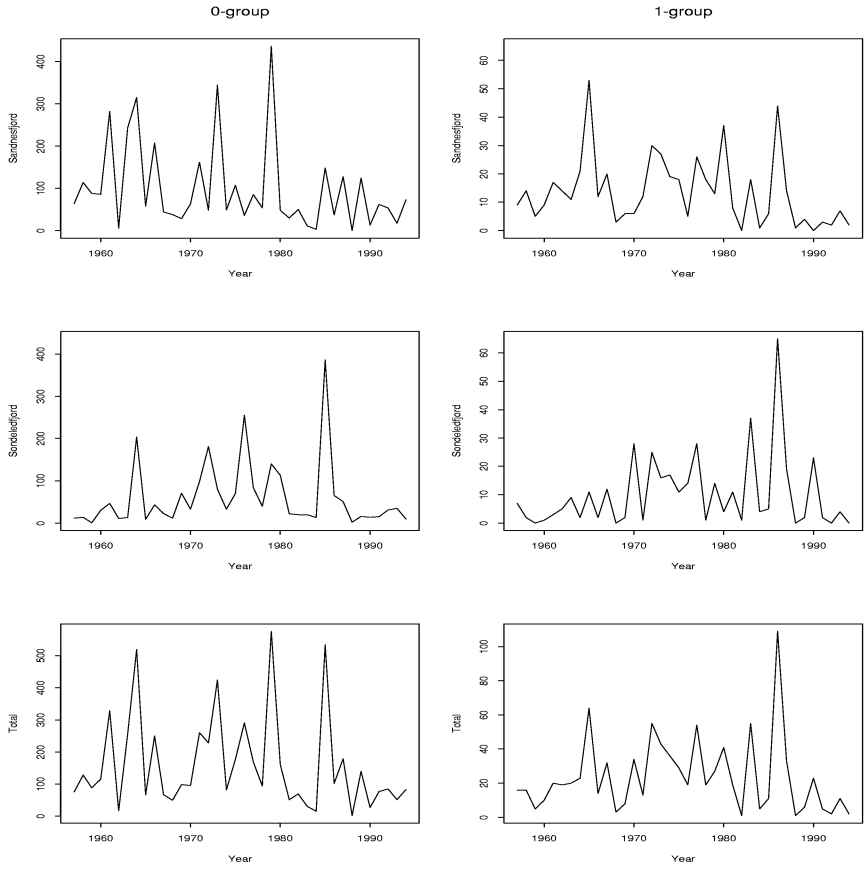

Figure 2.8: Aggregated data from the two fjords in the Risør area: Sandnesfjord (top), Søndeledfjord (middle) and total (bottom).

Chapter 3

Literature review

3.1 Approaches to fish population dynamics modelling

Most fisheries models include structural time series components, relating current and past population abundances (Schnute, 1994). In general, fisheries models can broadly be categorised into one of two main groups: biomass dynamics models and age-structured models, with delay-difference models as an intermediate type (Hilborn and Walters, 1992; Quinn and Deriso, 1999). In biomass dynamics models, current stock abundance, typically in terms of biomass B_t, is related to previous biomass B_{t-1} by adding a surplus production function $q(B_t)$, quantifying the effect of recruitment, growth and natural mortality, and subtracting the catch. The surplus production represents the amount of biomass that can be taken while maintaining the biomass at a constant rate. Parameters are estimated based on catch data, measured in total biomass, and relative abundance data from catches (catch per unit effort, CPUE) or from surveys. If catch-at age data are available, the population dynamics might be specified in terms of abundances at different age-cohorts. Such age-structured models include relations representing survival between the different cohorts as well as recruitment of new individuals to the population. Delay-difference models explicitly model an age-structured dynamics of a population, including biologically meaningful parameters for recruitment, growth and mortality, but the equations describing the age-structured processes are collapsed into one equation, such that in contrast to the age-structured models, age-specific data are not needed to fit the model. Using this type of model, it is implicitly assumed that the age-structured equations modelling the population dynamics hold exactly.

When building and analysing fisheries models, there are two main sources of stochasticity to take into account, giving rise to process and observational error respectively (Polacheck et al., 1993). The process error represents natural variability in the dynamics of the population, and emphasis on the stochasticity in early stages of the life cycle

of the cod dates back to the pioneering work of Hjort (1914). The observational error reflects stochasticity in the data collection process, relating the observed data to the actual populations. The models can be divided into three categories according to which sources of variability that are taken into account. (see e.g. Schnute (1994) and references therein). The first category includes models in which the population dynamics is assumed to be deterministic, such that given the model parameters, the dynamics is fully determined. In this category, observational error is the only source of variability included in the model. In models belonging to the second category, the observations are assumed to be exact, while process error is included to reflect natural variability in the model for the population dynamics. Models in the third category combine the two other types of models, treating both the population dynamics and the sampling process as stochastic. To enable the application of traditional parameter estimation methods, only one source of error has often been taken into account when fitting fisheries models, but it has been pointed out that this simplifying assumption might cause biases in the parameter estimates (De Valpine and Hastings, 2002, and references therein).

A coherent way to include both process and observational errors in fisheries models is to define the models in a state-space modelling framework (West and Harrison, 1997). Schnute (1994) describes how state-space modelling constitutes a general framework for developing sequential age-structured fisheries models, and discusses frequentist and Bayesian approaches to estimation of the model parameters using the Kalman filter an errors-in-variables paradigms. Depending on the type of data, the quantity to be represented by the state vector of the state-space model will typically be the total count or biomass of the fish or the number or biomass of fish categorised in age or length groups. The population dynamics will be represented by the state or system equation, and the sampling process by the observation equation of the state-space model.

For linear Gaussian models, parameters can be estimated using likelihood based estimation within a Kalman filter framework. This approach has been taken by e.g. Sullivan (1992) modelling catch-at-length data, and Freeman and Kirkwood (1995) and Reed and Simons (1996) who use total catch and catch per unit effort data. In the first two of these studies the system equation is linear in the population abundances, while Reed and Simons (1996) use a multiplicative model, such that log populations are linear in the log populations at the previous time step. Kimura, Balsiger and Ito (1996) compare Kalman filter and non-linear least squares estimates of the parameters of a delay-difference model based on simulated relative abundance and catch data, and show that the Kalman filter approach appears to be superior in the presence of process error. Newman (1998) formulates a state-space model for the movement and survival of salmon, defining the components of the state vector as the number of individuals in different regions. A spatially structured dynamic model is formulated based on specifying individual movement and survival, and by integrating over areas he arrives at a corresponding model for aggregated counts in different regions. Parameters are estimated by maximum likelihood, evaluating the likelihood using the Kalman filter.

In general, specifying the structural time series model in terms of a state-space model allows for the incorporation of non-Gaussian distributions and non-linearities in the specification of the population dynamics and the sampling process. However, introducing non-linearities and non-normality imply that classical methods for estimation, like the Kalman filter, is not generally applicable. The extended Kalman filter, relying on linear approximations based on Taylor series expansions, was used by Gudmundsson (1994) to fit parameters of a non-linear state-space model of catch-at-age data. In a recent paper, Schnute and Kronlund (2002) apply the extended Kalman filter to estimate relationships between stock and recruitment based on catch and escapement data for salmon. De Valpine and Hastings (2002) present an alternative method, which they call the "numerically integrated state-space method", for estimating states and parameters of general population models. They extend the Kalman filter approach to non-linear and non-Gaussian state-space models following Kitagawa (1987), by numerically estimating all probability distributions and calculations involved.

Advances in simulation based estimation using Markov chain Monte Carlo (MCMC) methods have increased the applicability of Bayesian methods to fit non-linear and non-Gaussian state-space models (e.g. Carlin, Polson and Stoffer (1992), Geweke and Tanizaki (2001)). These methods have also been introduced to fisheries modelling, and the potential for applying the Bayesian approach in conjunction with MCMC methods within population dynamics modelling is pointed out by several authors (e.g. Punt and Hilborn, 1997; Buckland, Goudie and Borchers, 2000; Myers, 2001). In Schnute and Kronlund (2002), a frequentist extended Kalman filter approach and a Bayesian errors-in-variables approach to estimation of the parameters of a state-space model for stock-recruitment relationships for salmon are contrasted, using MCMC methods to assess the uncertainty of modal estimates of the parameters from the two approaches. Parameters of state-space models representing the population dynamics for tuna are estimated using Gibbs sampling in combination with the adaptive rejection Metropolis algorithm in Meyer and Millar (1999) and Metropolis-Hastings within-Gibbs sampling in Millar and Meyer (2000b). The models are also extended to incorporate age-structured data (Millar and Meyer, 2000a). Knight, Weir and Pettitt (2003) adopt a fully Bayesian hierarchical model for summer whiting, explicitly modelling the sampling process using a Binomial distribution. They include a spatially autoregressive model for the densities of counts, and estimate parameters using MCMC. A Bayesian approach to parameter estimation has also recently been taken in an analysis of the data on which our study is based. Bjørnstad et al. (1999a) use Gibbs sampling to estimate parameters in a dynamic model for the coastal populations of cod along the Norwegian Skagerrak coast. In the next section we give an overview of previous analyses of the Skagerrak cod data.

3.2 A review of previous analyses of the Norwegian Skagerrak cod

As described in Section 2.2, the data from the beach seine hauls along the Skagerrak coast were originally collected to study the effect of releasing artificially hatched cod larvae, and many early publications on the data were related to this problem (see e.g. Gjøsæter, Stenseth, Sollie and Lekve (2003b), Tveite (1971) and references therein). The larval release programme was continued until 1971. Based on analysis of variance on the 0-group data from the period 1920-1969, Tveite (1971) concluded that the release of cod larvae could not be significantly separated from other natural sources of variability as a source of variation in year-class strength.

In recent years, more attention has been drawn to studies of the effect of the juvenile post-settlement stages. Julliard, Stenseth, Gjøsæter, Lekve, Fromentin and Danielssen (2001) study the mortality of juvenile cod using a different set of data resulting from a release-recapture experiment within one fjord. The study is based on capture-mark-recapture modelling, assuming that tagged individuals might be recaptured at a series of different capture occasions. Based on data obtained from fishermen returning tags from recaptured individuals, they estimate parameters representing natural and fishing mortality by maximum likelihood. The results indicate a very hight mortality the first few months of the 0-group juvenile stage.

Spatial and temporal patterns of the abundance of cod, as well as other coexisting species, have been studied by Fromentin and co-workers (Fromentin, Stenseth, Gjøsæter, Bjørnstad, Falck and Johannessen, 1997; Fromentin, Stenseth, Gjøsæter, Johannessen and Planque, 1998; Fromentin, Gjøsæter, Bjørnstad and Stenseth, 2000). Based on descriptive analysis and spectral analysis of the raw 0-group and 1-group data, they identify 2-2.5 year cycles as well as long-term trends in abundance of both age-groups, and report substantial spatial variation both between fjords and between stations within a fjord (Fromentin et al., 1997; Fromentin et al., 2000). Similar results on the temporal pattern is found by Bjørnstad et al. (1999a) using spectral analysis. Fromentin et al. (1998) examine possible relationships between climatic conditions, represented by NAO (the North Atlantic Oscillation index) and long-term fluctuations in abundance of juvenile cod. The results do not give any evidence of the existence of such a correspondence. They also consider the effect of bottom flora coverage on the abundances, and give reasons for a possible relation between the variation in bottom flora coverage and cod abundance.

An important aspect of study within population dynamics of fish is the study of possible density-dependencies between and within age-groups. (Myers, 2001; Myers and Cadigan, 1993a). Density-dependencies can be caused by competition for food, or predation of older individuals on younger. The effect of density-dependencies in survival of juvenile cod has also become a focus of study on the Flødevigen data. Bjørnstad

and co-authors (Bjørnstad, Fromentin, Stenseth and Gjøsæter, 1999b) introduce a test for density-dependent survival by testing for convexity in the survival function relating the 1-group at one year to the 0-group the previous year. They show that the method, that does not require any parametric form of the density function, is relatively robust to the presence of measurement error, and using the method, they conclude that there is significant within-group density-dependency in the survival from 0-group to 1-group cod. Lekve, Ottersen, Stenseth and Gjøsæter (2002) identify skewed length distributions of the juvenile cod, and they suggest that the skewness might partly be caused by competition between the juveniles. Their results further indicate that the growth and survival of juvenile cod can be influenced by environmental factors, like directional wind stress setting up currents transporting food in favourable directions.

In Fromentin, Myers, Bjørnstad, Stenseth, Gjøsæter and Christie (2001), direct (within-group) and delayed (between-group) density-dependent mortality is estimated based on the model defined in Myers and Cadigan (1993a). This model incorporates stochastic recruitment, which includes egg production, hatching and survival to the 0-group, but Fromentin and co-workers add a stochastic component to the density-independent juvenile mortality. It is assumed that the observations are proportional to the true abundances, but with a log-Gaussian observational error added. The parameters are estimated by maximum likelihood based on 0-group and 1-group data that are aggregated in fjords and then log-transformed. The estimates indicate a difference between south-western and north-eastern fjords. There is little evidence of delayed density-dependent mortality in the south-west, but the results indicate such an effect for some of the north-eastern fjords. Within-group density-dependency is found in both regions. Post-settlement density-independent survival seems to have some significance, particularly in the south-western fjords. The variability in recruitment of 0-group cod seems to be larger in the north-western than in the south-eastern fjords.

An age-structured dynamic model for the population dynamics is defined in Bjørnstad et al. (1999a). They aggregate the data in two regions, consisting of the south-western and north-eastern stations respectively, and estimate parameters describing the population dynamics for each region separately. It is assumed that two adult age groups, two and three year old cod, contribute to the adult spawning stock. We here give a brief description of this model, since it constitutes the basis for our hierarchical Bayesian model to be presented in Chapter 4. For each region, let X_t and Y_t be the abundance in year t of the 0-group and 1-group cod respectively, and let Z_t and W_t denote the abundance of two and three year old cod. Then, their population dynamics model can be summarised by

$$\begin{aligned} Y_t &= X_{t-1}\exp(c+\beta\log X_{t-1}+\gamma\log Y_{t-1}) && \text{(juvenile survival)} \\ Z_t &= \lambda Y_{t-1} && \text{(survival from 1-group to 2-year old)} \\ W_t &= \lambda Z_{t-1} && \text{(survival from 2- to 3-year-old)} \\ X_t &= \exp(\alpha_t)(Z_t+W_t) && \text{(recruitment of 0-group cod)} \end{aligned} \tag{3.1}$$

The density-dependent juvenile survival rate is assumed to be log-linear in the log-abundances of the 0-group and 1-group cod the previous year, in correspondence with other studies (Bjørnstad et al., 1999b; Myers and Cadigan, 1993a). Stochasticity enters the model by assuming the annual recruitment rate, represented by $\exp(\alpha_t)$, to be log-Normal with constant mean and variance. The other components of the population dynamics as defined by (3.1) are assumed to be deterministic. Further, they assume the aggregated count data to be noisy observations of the true abundances, with measurement error approximated by Poisson variability. So no explicit model for the sampling process relating true abundances to the actual samples is included. The parameters of the model, except for the adult survival rate λ that is kept fixed, are estimated using the Gibbs sampler with adaptive-rejection sampling, using the aggregated data for both 0-group and 1-group cod. They estimate within- and between-cohort effects, but the within-cohort effect is weaker ($|\gamma| < |\beta|$). Also, a non-zero density-independent effect (c) on juvenile survival is found. Changing the fixed value of the adult survival λ influences the recruitment rate, but it appears to have no effect on other parameters.

In Stenseth et al. (1999), a similar age-structured population dynamics model for the Skagerrak cod is defined, including the log-linear density-dependent survival at the juvenile stages, density-independent survival to adult cod and stochastic recruitment. As in Bjørnstad et al. (1999a), it is assumed that the stochasticity in the recruitment process dominates the stochasticity in the survival processes, but in contrast to that study, only the two year old cod is assumed to contribute to the adult spawning population. Further, they assume no measurement error. They combine the components of the age-structured model into one time series model for the 0-group cod, for which the data is most abundant, eliminating the 1-group and adult cod populations from the population dynamics equations. They identify the resulting model as an ARMA(2,1) model for the 0-group cod populations. Based on their estimates, it is concluded that there is evidence of both within-group and between-group density-dependencies for the juvenile cod. Also, they argue that the results are in agreement with a periodic cycle of 2-2.5 years for the abundance of 0-group cod.

Chan, Stenseth, Lekve, Gjøsæter and Ottersen (2003b) extend the age-structured model of Stenseth et al. (1999) to incorporate abiotic (temperature, NAO) and biotic (coexisting species) factors as covariates in the survival and recruitment models, and to take into account both natural and fishing mortality. In contrast to the previous studies described above, all adult stages are included in the model, but they assume them to have the same biological properties. Also, the model is fully specified on a fjord-level, based on the assumption that the cod spawns in deeper waters of the fjord. As in Stenseth et al. (1999), the variables corresponding to the 1-group and adult populations are eliminated from the dynamic model, leaving an ARMAX(2,2)-model in the state variable corresponding to the 0-group cod. The parameters are estimated using the 0-group data only. The data are aggregated to fjord level, using weights representing the chance of an egg settling as juvenile cod in the given sampling sites in the fjord. The results support the conclu-

sions of the earlier studies that there are significant between-and within-group density-dependencies. The recruitment is estimated to increase by increased water-temperature, and they estimate an increase in fishing mortality over the years of study.

Based on the model of Chan et al. (2003b), a simulation study of the effect of releasing larvae as well as juvenile cod has been conducted (Chan, Stenseth, Kittilsen, Gjøsæter, Lekve, Smith, Tveite and Danielssen, 2003a). Evidence of significant augmented mature population sizes in some fjords are found, but the increases are considered to be too small to be of practical significance.

There are two main weaknesses with the majority of dynamic models defined so far on the Skagerrak cod data. First, the observational error represents measurement error only, which implies that the sampled data are assumed to be proportional to the true abundances. No explicit model for the sampling process, relating the observed counts to the true underlying population abundances, is included. As a consequence, the population dynamics is estimated as if defined on the sampled and not the true underlying abundances. Fromentin et al. (2001) explicitly include proportional catchability in the model, but the model parameters are estimated on centred data. Second, except for the model in Fromentin et al. (2001), the relations defining the survival of juvenile and adult cod are assumed to be deterministic. This assumption is based on the results of earlier studies, indicating that stochastic fluctuations in juvenile and adult survival are small compared to the stochasticity of the recruitment process and survival at earlier stages (Hjort, 1914; Hjort, 1926; Cushing, 1990). However, recent studies have indicated that variation in juvenile stages are also important and should be taken into account (e.g. Fromentin et al., 2001 and references therein). Our aim is to generate and study a model for the population dynamics of the Skagerrak cod by extending and modifying the population dynamics model (3.1), including stochastic variation in all components of the population dynamics as well as an explicit model for the sampling process generating the data.

Chapter 4

The statistical model

We define the statistical model as a state-space model within a Bayesian hierarchical modelling framework. Within this framework, the stochastic model is naturally divided into two components, one specifying the dynamics between the different age groups of cod, and the other representing the sampling process relating the true underlying population abundances to the sampled cod at the monitored locations.

Our statistical model is based on the assumption that each fjord represents a reproductive unit, and that there is no migration between fjords (Danielssen and Gjøsæter, 1994). The total spawning population of cod within a fjord is assumed to contribute to the recruitment of young cod at all monitored locations along the shore. Further, we will define the population dynamics on total fjord populations of all age groups involved. Since the data are samples taken at a limited number of locations along the shore, our model should include a down-scaling between the population abundances at fjord level and the sampled 0-group and 1-group cod at the monitoring stations along the shore. This is incorporated in the model for the sampling process.

For each fjord f, the state vector for each year t is three-dimensional, and made up by the components $X_{f,t}$, $Y_{f,t}$ and $Z_{f,t}$, representing the abundance of 0-group (X), 1-group (Y) and mature (Z) cod respectively. The corresponding observation vector consists of catches of 0-group cod, denoted $X^o_{i,t}$ and 1-group cod $Y^o_{i,t}$, for each year t and at a set of monitored locations i. Thus, for each fjord the observation vector is $2 \times n_f \times n_t$-dimensional, where n_f is the number of stations in fjord f and n_t the number of years of study. For the mature cod, no data is available.

We specify our hierarchical model by first defining the population dynamics model as a prior model for the population abundances, and then the sampling process, constituting the likelihood of our model.

4.1 The population dynamics model as a Bayesian prior

Our model is a modified and extended version of the population dynamics model defined in Bjørnstad et al. (1999a) and Stenseth et al. (1999). The dynamics regulating the population can be characterised by two main processes, the recruitment of new individuals, including spawning and survival to the first juvenile stage, and survival from one year to the next. We specify two relations for survival, one for the survival from 0-group to 1-group cod, and one for the survival between the 1-group and adult cod as well as for the adult cod. In contrast to the models in Bjørnstad et al. (1999a), where two and three year old cod are considered as separate age groups, and Stenseth et al. (1999), where only the two year olds are assumed to contribute to the spawning stock, we do not distinguish between different age groups of mature cod. We assume that the mature population, consisting of all cod at the age of two years or older, experience the same rate of survival. This is in analogy to the approach taken by Chan et al. (2003b).

While Bjørnstad et al. (1999a) and Stenseth et al. (1999) consider the population dynamics, except for the recruitment process, to be deterministic, we impose stochasticity in all components of the state-space model. We specify the prior distributions for recruitment, survival between 0-group and 1-group cod and adult survival in terms of conditional distributions, given past values of all population abundances and other time dependent parameters. To simplify the notation, we will denote by past(t) the collection of all abundances and other time dependent variables up to, but not including, year t. In Figure 4.1 a schematic representation of the dynamics of the life cycle of the cod, as described in Section 2.1, is given in terms of the variables Z_t, X_t and Y_t for one single fjord, suppressing the fjord subscript f. Graphically, our prior model can be illustrated as in Figure 4.2, discarding all parameters. The arrows indicate conditional dependencies. For example, given past abundances, the 1-group population abundance in year t depends on the 0-group and 1-group abundances in year $t-1$, while the 0-group abundance in year t, given past abundances as well as the current abundance of mature cod, depends on the current mature abundance only.

4.1.1 Recruitment of new individuals

Conditional on the past and on the current adult population abundances $Z_{f,t}$, we define the prior mean of the 0-group abundance to be

$$\mathrm{E}(X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}) \;=\; Z_{f,t} \exp(\alpha_0 + \alpha_t); \;\; t = 1, 2, \ldots, n_t. \tag{4.1}$$

In our model, the recruitment rate $\exp(\alpha_0 + \alpha_t)$ reflects the spawning process as well as the survival from the larval to the 0-group stage. We assume that the rate of recruitment of new individuals to the population is density-independent, but allow for time dependencies. The component α_t represents possible effects of time-varying biotic and

Year $t-1$			Year t		Year $t+1$	
Spring		Autumn	Spring	Autumn		Spring
Z_{t-1}	$\xrightarrow[\text{survival}]{\text{spawning,}}$	X_{t-1}	$\xrightarrow[\text{survival}]{\text{density dep.}}$	Y_t	$\xrightarrow{\text{survival}}$	Z_{t+1}
		$\{X^o_{i,t-1}\}$ $\{Y^o_{i,t-1}\}$ (Sampling)		$\{X^o_{i,t}\}$ $\{Y^o_{i,t}\}$ (Sampling)		

Figure 4.1: An illustration of the dynamics of the life cycle of the cod within one fjord in terms of the stochastic variables of the prior model. The subscripts i and t refers to location and year, while the fjord subscript of the adult (Z_t), 1-group (Y_t) and 0-group (X_t) fjord population abundances is suppressed.

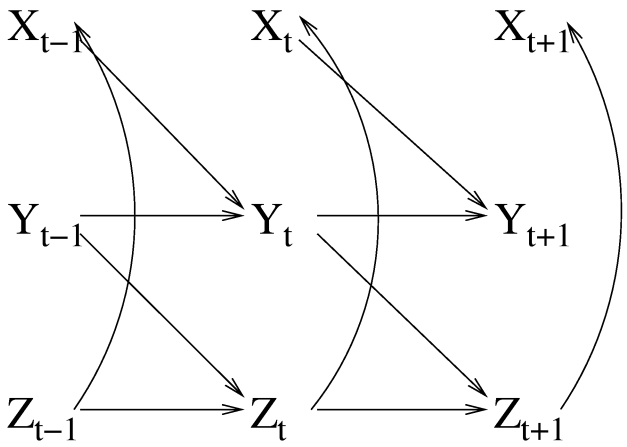

Figure 4.2: Graphical description of the prior model for the population dynamics. The age-groups are denoted by X_t (0-group), Y_t (1-group), and Z_t (mature cod).

abiotic factors, e.g. climatic and hydrographical conditions, on the spawning and early survival of the larvae. We impose the constraint $\sum_t \alpha_t = 0$, such that α_0 represents the mean level of the log recruitment rate, and α_t temporal fluctuations around this mean level.

Since we are dealing with count data, a discrete distribution will in principle be most appropriate. In practice however, the population sizes will be relatively large, since the model for the population dynamics is defined on the fjord level. Therefore, we choose to use a Gaussian assumption on the error in the mean. The variance is assumed to be proportional to the mean number of new recruits to the population, such that the conditional prior distribution of $X_{f,t}$, given the past population history as well as the current adult abundance $Z_{f,t}$, is defined by

$$X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}, \delta_x \sim N(Z_{f,t}\exp(\alpha_0+\alpha_t), Z_{f,t}\exp(\delta_x+\alpha_t)); \quad t = 1,2,\ldots,n_t \tag{4.2}$$

Here, the over-dispersion factor δ_x is given by $\delta_x = \delta_x' + \alpha_0$, where $\exp(\delta_x')$ represents the over-dispersion.

4.1.2 Survival between the juvenile stages

The juvenile survival is specified as survival between the 0-group and 1-group stages from September/October one year to the same time of the next year. We will allow for within- and between-age-group density-dependency in the survival rate, such that the survival rate from year $t-1$ to year t depends on the abundances of 0-group and 1-group cod in year $t-1$. Otherwise, the survival rate is assumed to be independent of time. Earlier studies have indicated that the density-dependent survival rate is approximately log-linear in the 0-group and 1-group-abundances the previous year (Bjørnstad et al., 1999b; Myers and Cadigan, 1993a), and we choose to use a log-linear model in analogue to the model in Bjørnstad et al. (1999b) for the conditional mean of the 1-group cod. Thus, we define the mean number of surviving 0-group cod by

$$\mathrm{E}(Y_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa) = \exp(-\kappa)\, X_{f,t-1}^{1-\beta} Y_{f,t-1}^{-\gamma}; \quad t = 2,3,\ldots,n_t, \tag{4.3}$$

which can also be written as

$$\mathrm{E}(Y_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa) = X_{f,t-1}\exp(-\kappa - \beta\log(X_{f,t-1}) - \gamma\log(Y_{f,t-1})). \tag{4.4}$$

The degree of dependency on the population abundances at time $t-1$ is governed by the parameters β for within-group dependency and γ for between-group dependency. The parameter κ represents the level of the density-independent component of the survival rate. All these parameters are restricted to be non-negative.

Assuming that the survival process for each individual can be approximated by independent Bernoulli trials, we define the conditional prior distribution for $Y_{f,t}$ by a Binomial distribution, given by

$$Y_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa \quad \sim \quad Bin(X_{f,t-1}, \exp(-\kappa)\ X_{f,t-1}^{-\beta} Y_{f,t-1}^{-\gamma}). \tag{4.5}$$

To complete the prior model definition, a non-informative uniform prior $f(y_{f,1})$ is specified for the abundance at year 1, such that

$$f(y_{f,1}) \propto 1;\ y_{f,1} = 0, 1, 2, \ldots \tag{4.6}$$

4.1.3 Adult survival

The number of adult cod in year t will depend on the number of adult cod, $Z_{f,t-1}$, and the 1-year cod abundance, $Y_{f,t-1}$, the previous year. We assume that the Norwegian Skagerrak cod matures at the age of two years, and that all age groups of mature cod are equivalent with respect to reproduction and survival.

We will denote the rate of survival for the adult cod by $\exp(-\theta)$, where θ is the instantaneous mortality. We do not distinguish between natural and fishing mortality, such that the parameter θ will represent the total mortality for the adult and 1-group cod. In analogy to the 1-group cod, we assume that conditional on the past, the number of adults follows a Binomial distribution, given by

$$Z_{f,t} \mid \text{past}(t), \theta \quad \sim \quad Bin((Z_{f,t-1} + Y_{f,t-1}), \exp(-\theta));\quad t = 2, 3, \ldots, n_t, \tag{4.7}$$

and that

$$f(z_{f,1}) \quad \propto \quad 1;\ z_{f,1} = 0, 1, \ldots \tag{4.8}$$

4.1.4 Specification of hyper-priors of the population dynamics

We restrict the parameters α_0, θ, β, γ and κ to be non-negative, and they are all assigned Gamma priors. The overall mean recruitment rate $\exp(\alpha_0)$ could in principle be less than 1 (and thus $\alpha_0 < 0$), but biologically a recruitment rate less than 1 would be meaningless, since this would lead to a continuing decrease in the population. We assign equal and independent Gaussian priors $\alpha_t \sim N(0, \tau_\alpha^{-1});\ t = 1, \ldots, n_t$ to the parameters representing the temporal structure of the log recruitment rate, but impose the constraint $\sum_t \alpha_t = 0$. The common precision τ_α is assumed to be unknown and to be estimated, and is given a vague Gamma prior. The over-dispersion factor δ_x in (4.2) is assigned a Gaussian prior with fixed hyper-parameters, truncated at 0. A summary of the hyper-priors, including specification of the chosen values of the hyper-parameters, will be given in Section 5.1.

4.2 The sampling process

The sampling process is represented by the likelihood of the Bayesian hierarchical model. Since no observations are available for the mature cod, the sampling process is given by the beach seine catches of the juvenile cod, which tend to stay in shallow waters near the shore in the autumn, at the time of the year when the sampling takes place. The sampling processes are assumed to be independent and equivalent for the two juvenile age-groups, so we derive the model for the 0-group cod only.

Defining the population dynamics on the fjord level, we have to incorporate some kind of down-scaling of the population abundances from the fjord level to the level of the monitoring stations. To reflect the sampling procedure as described in Figure 2.4, we find it conceptually convenient to define the model in terms of three steps. First, we have to define the fraction of the 0-group cod population that might happen to be in the area covered by the seine at a location i at the time of sampling. This fraction is the population that has settled near a station i in a year t, using this area as its nursery area. We denote the corresponding stochastic variable by $X^s_{i,t}$. Second, the cod tend to move to and from the shore, such that only a fraction of the population that has settled at a location, will actually be in the area covered by the seine at the time of sampling. We will denote this fraction the catchable population, and it is represented by the variable $X^m_{i,t}$. Finally, some individuals might escape from the seine, leaving the number that is actually caught, which are the data $X^o_{i,t}$. To summarise, the likelihood should reflect the sub-processes

- settling at a location i,
- movement to and from the shore, and
- escaping or not escaping from the beach seine.

The last two processes are defined at the sampling location scale, while the settling process includes down-scaling of the true population abundances to the scale of the data. In Figure 4.3, we give a schematic description of the three sub-steps in our model for the sampling process in terms of the corresponding stochastic variables.

We will assume that the settling process, as well as the actual sampling, depends on the bottom coverage. More specifically, we assume that the juvenile cod tend to settle where there is a large amount of vegetation on the bottom, providing food and shelter, and the probability of escaping the seine is larger if the vegetation is dense. The second sub-process, movement to and from the shore, is assumed to be distributed independently of location and time.

The term "settling" might be a bit misleading the way it is used in the definition of the sampling process. Since the population dynamics apply to the cod at fjord level, there

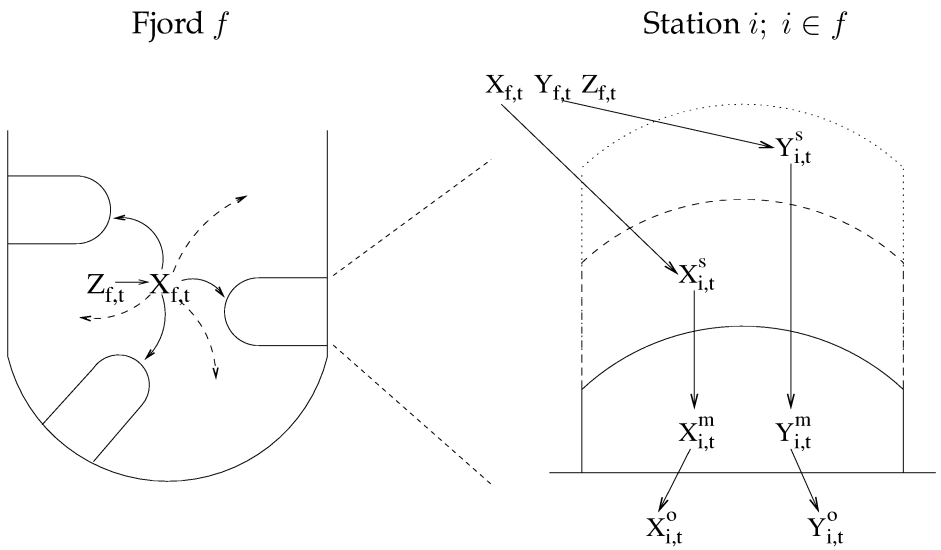

Figure 4.3: A schematic description of the sampling process. The left figure illustrates the settling at different locations in a fjord, and the right figure the sampling process at one sampling station. In the right figure, the full line represents the beach seine, the dashed line the outer boundary of the nursery area for the 0-group cod (X), and the dotted line the outer boundary of the nursery area for the 1-group cod (Y).

is no underlying assumption that the 0-group an 1-group cod settle *permanently* at a specific location. Rather, we assume that the cod spread in the fjord in a way such that on the average, there is a higher population density where the bottom coverage is good. Within this general pattern, we in principle allow each individual to move around.

4.2.1 Specification of the likelihood model

We model the first process, the settling at a location i, using a multinomial distribution. This is motivated by the fact that for each fjord f and year t, the settling process can be considered as an experiment consisting of $X_{f,t}$ independent trials. Each of the $X_{f,t}$ 0-group cod has a certain probability of settling at each of the n_f locations in fjord f, or at another, non-monitored, location along the shore of the fjord. Thus, each experiment has $n_f + 1$ possible outcomes. The independence assumption implies that we assume the cod to move independently of each other, and thus schooling is ruled out.

Let $p^{s,x}_{i,t}$ be the probability of settling at location i in year t for a 0-group individual. If we let $X^s_{0,t}$ denote the number of cod settling in the part of the fjord that is *not* monitored,

the multinomial model for the settling process can be written

$$(X^s_{0,t}, \ldots, X^s_{n_f,t} | p^{s,x}_{0,t}, \ldots, p^{s,x}_{n_f,t}) \sim \text{Multinom}(X_{f,t};\ p^{s,x}_{0,t}, \ldots, p^{s,x}_{n_f,t}). \tag{4.9}$$

Conditionally on the number of cod that has settled at a location, the number of cod that moves to the shore at location i is taken to be Binomial, with parameters $X^s_{i,t}$ and $p^{m,x}_{i,t}$, that is

$$X^m_{i,t} | X^s_{i,t}, p^{m,x}_{i,t} \sim Bin(X^s_{i,t}, p^{m,x}_{i,t}). \tag{4.10}$$

Here, we implicitly assume that the sub-populations at different locations and years behave independently of each other. Similarly, we model the actual sampling of the cod as another Bernoulli trial, such that the observed number of 0-group cod, given the number in the region covered by the beach seine, $X^m_{i,t}$, and the probability of not escaping, $p^{o,x}_{i,t}$, is given by

$$X^o_{i,t} | X^m_{i,t}, p^{o,x}_{i,t} \sim Bin(X^m_{i,t}, p^{o,x}_{i,t}). \tag{4.11}$$

In analogy to $X^s_{0,t}$ in (4.9), let $X^o_{0,t}$ denote the 0-group cod that is not sampled, that is $X^o_{0,t} = X_{f,t} - \sum_{j=1}^{n_f} X^o_{j,t}$. In words, $X^o_{0,t}$ represents all 0-group cod in the part of the fjord that is not monitored, as well as the 0-group cod at the monitored stations that is not caught. Combining (4.9), (4.10) and (4.11), we arrive at a multinomial distribution for the conditional distribution of $(X^o_{0,t}, \ldots, X^o_{n_f,t})$ given the 0-group fjord abundance $X_{f,t}$. The distribution is given by

$$(X^o_{0,t}, \ldots, X^o_{n_f,t} | p^x_{0,t}, \ldots, p^x_{n_f,t}) \sim \text{Multinom}(X_{f,t};\ p^x_{0,t}, \ldots, p^x_{n_f,t}), \tag{4.12}$$

where

$$\begin{aligned} p^x_{i,t} &= p^{s,x}_{i,t} p^{m,x}_{i,t} p^{o,x}_{i,t}; \qquad i = 1, \ldots, n_f, \text{ and} \\ p^x_{0,t} &= 1 - \sum_{j=1}^{n_f} p^{s,x}_{j,t} p^{m,x}_{j,t} p^{o,x}_{j,t}. \end{aligned} \tag{4.13}$$

Observe that since we model the true underlying populations, $X_{f,t}$, at fjord level, the sampling probabilities $p^x_{i,t}$ at the monitored locations will be small. The correlation between the variables in a multinomial model is

$$\text{Cor}(X^o_{i,t}, X^o_{j,t}) = \frac{-p^x_{i,t} p^x_{j,t}}{\sqrt{p^x_{i,t}(1 - p^x_{i,t})}\sqrt{p^x_{j,t}(1 - p^x_{j,t})}}. \tag{4.14}$$

For small sampling probabilities, $\text{Cor}(X^o_{i,t}, X^o_{j,t}) \approx -\sqrt{p^x_{i,t} p^x_{j,t}}$, which is again small for small $p^x_{i,t}$ and $p^x_{j,t}$. Thus, in practice, we do not introduce any large errors by assuming independence and using the marginal Binomial distributions, given by

$$X^o_{i,t} \mid X_{f,t}, p^x_{i,t} \sim Bin(X_{f,t}, p^x_{i,t}). \tag{4.15}$$

To give a rough estimate of the order of magnitude of the sampling probabilities, consider the case where we assume that the probability of settling is equal at all locations. Let $\Omega(f_i)$ be the length of the coastline of the fjord to which location i belongs, and L the distance between the two landing points of the beach seine. Then, the settling probability at location i and year t can be taken to be the fraction of the shore of the fjord covered by the seine, such that $p_{i,t}^{s,x} = L/\Omega(f_i);\ \forall(i,t)$. Assuming further that $p_{i,t}^{o,x} = 1$ and $p_{i,t}^{m,x} = 1$ for all i and t, we get that $p_{i,t}^{x} = p_{i,t}^{s,x} = L/\Omega(f_i);\ \forall t$. For the fjords in the study, we have $L/\Omega(f) < 0.005$, and thus $\text{Cor}(X_{i,t}^o, X_{j,t}^o) < 0.005$, which is negligible.

So far we have only considered the 0-group population. For the 1-group population, we assume that the situation is equivalent, such that the likelihood of $(Y_{0,t}^o, \ldots, Y_{n_f,t}^o)$ is given by

$$(Y_{0,t}^o, \ldots, Y_{n_f,t}^o | p_{0,t}^y, \ldots, p_{n_f,t}^y) \sim \text{Multinom}(Y_{f,t};\ p_{0,t}^y, \ldots, p_{n_f,t}^y). \tag{4.16}$$

Here,

$$\begin{aligned} p_{i,t}^y &= p_{i,t}^{s,y} p_{i,t}^{m,y} p_{i,t}^{o,y}; \qquad i = 1, \ldots, n_f, \text{ and} \\ p_{0,t}^y &= 1 - \sum_{j=1}^{n_f} p_{j,t}^{s,y} p_{j,t}^{m,y} p_{j,t}^{o,y}. \end{aligned} \tag{4.17}$$

The marginal Binomial distributions for the 1-group cod are given by

$$Y_{i,t}^o \mid Y_{f,t}, p_{i,t}^y \sim Bin(Y_{f,t}, p_{i,t}^y). \tag{4.18}$$

4.2.2 Specification of the sampling probabilities of the likelihood

The sampling probabilities of the 0-group and 1-group cod were defined in terms of the corresponding parameters for the sub-process, as shown in (4.13) and (4.17). We impose prior distributional assumptions on the settling probabilities $p_{i,t}^{s,x}$ and $p_{i,t}^{s,y}$, scaling these probabilities by fixed values of the catchability probabilities $p_{i,t}^{m,x}$ and $p_{i,t}^{m,y}$ and the probabilities $p_{i,t}^{o,x}$ and $p_{i,t}^{o,y}$ of not escaping. These fixed values are specified on the basis of expert biological information.

Based on the fact that the juvenile cod tend to settle where they find food and hiding places (see Section 2.1), it seems reasonable to assume that the settling probability at different stations depends on the amount of vegetation on the bottom. Sampling efficiency might also be related to bottom vegetation. Diving behind the seine during sampling has shown that fish hiding in dense vegetation may sometimes escape under the seine (Flødevigen Marine Research Station, unpublished underwater video and personal communication). Therefore, increase in bottom vegetation probably makes the location a more attractive habitat for young cod, but allows for a less efficient beach seine catch.

Thus, we incorporate a categorical variable $d_{i,t}$, describing the amount of vegetation on the bottom at the sampling site i in year t. The values taken by $d_{i,t}$ are coded as $1, 2, 3, 4, 5$, for "no vegetation", "few plants", "some plants", "many plants" and "bottom totally covered", respectively. We will assume that this variable can be related to the suitability of the habitat as a nursery area for young cod, and also to how easy it is to obtain a good sample of 0-group and 1-group cod individuals.

We first describe the prior distributions for the settling probabilities $p_{i,t}^{s,x}$ and $p_{i,t}^{s,y}$. We will choose equivalent priors for the two age groups, and concentrate on the priors for the 0-group. Using the multinomial likelihood (4.12), the conjugate prior distribution is the Dirichlet distribution. This takes into account that the $p_{i,t}^{s}$'s, including the ones for the non-monitored locations, are dependent, and should sum to 1. We are interested in the probabilities representing the monitored locations, and for these, we have the constraint

$$\sum_{i=1}^{n_f} p_{i,t}^{s,x} < 1; \forall t. \tag{4.19}$$

However, as pointed out in Section 4.2.1, the settling probabilities will be small, such that the constraint (4.19) will be satisfied. Consequently, we define marginal priors for $p_{i,t}^{s,x}$, and the settling processes at different locations are assumed to be independent. We take into account the assumption on habitat preferences due to bottom vegetation by including the bottom coverage $d_{i,t}$ as a covariate, and define the settling probabilities to take one out of 5 different values, depending on the categorical variable $d_{i,t}$. More specifically, we let

$$\text{logit}(p_{i,t}^{s,x}) = \log(L/\Omega(f_i)) + \sum_{k=1}^{5} c_k^{s,x} I_{\{d_{i,t}=k\}}, \tag{4.20}$$

where the parameters $c_k^{s,x};\ k = 1, \ldots, 5$ are random effects that are to be estimated from the data. This choice of prior is motivated as follows. As mentioned in Section 2.1, the juvenile cod tend to live near the shore in the autumn, when the sampling take place. If we assume that the cod settles uniformly along the coastline of fjord, it would seem reasonable to let

$$p_{i,t}^{s,x} = L/\Omega(f_i), \tag{4.21}$$

where $L/\Omega(f_i)$ is the fraction of the coastline of the fjord covered by station i. Recall that L represents the distance between the landing points of the beach seine, and $\Omega(f)$ is a measure of the length of the coastline of fjord f. We adjust this average value by multiplying by a factor $1 + c_k^{'}$ depending on the value k of the covariate $d_{i,t}$, such that

$$p_{i,t}^{x,s} \quad = \quad \frac{L}{\Omega(f)}(1 + \sum_{k=1}^{5} c_k^{'} I_{[d_{i,t}=k]}). \tag{4.22}$$

Since the settling probabilities $p_{i,t}^{s,x}$ will be relatively small, we have that $\text{logit}(p_{i,t}^{s,x}) = \log(p_{i,t}^{s,x}/(1 - p_{i,t}^{s,x})) \approx \log(p_{i,t}^{s,x})$. Defining $c_k^{s,x} = \log(1 + c_k^{'})$, the prior model in (4.20) follows.

We assign equal Gaussian priors to the $c_k^{s,x}$'s, such that

$$p(c_k^{s,x}) \quad \sim \quad N(\mu_{c,x}, \tau_{c,x}) \tag{4.23}$$

where $\mu_{c,x}$ and $\tau_{c,x}$ are both fixed. The parameters will be estimated under the constraint

$$\sum_{i,t} c_{d_{i,t}}^{s,x} \quad = \quad 0, \tag{4.24}$$

such that the $c_k^{s,x}$'s can be interpreted as deviations from an overall mean level of the settling probability for the monitored stations. For the estimated deviations to be valid for non-monitored locations as well, we need to assume that the sampling stations are representative for the whole coastline of the fjord.

The probability of a cod being near the shore, and thus catchable at the time of sampling, is assumed to be constant, but different for the two age groups. We denote these constants by ψ_x and ψ_y for the 0-group and 1-group cod respectively. Finally, in analogy to the settling probabilities $p_{i,t}^{s,x}$ and $p_{i,t}^{s,y}$, we let the probabilities of not escaping, $p_{i,t}^{o,x}$ and $p_{i,t}^{o,y}$, each take one out of 5 different fixed values, $p_k^{o,x};\ k = 1, \ldots, 5$ and $p_k^{o,y};\ k = 1, \ldots, 5$, respectively, depending on the value of the variable $d_{i,t}$. The actual values assigned to these parameters will be specified in Section 5.1.

4.3 Summary of the model

A summary of the components of the model is given in Table 4.1. Prior specifications of the sampling probabilities of the likelihood are given in Table 4.2, and the structure of the prior model, including the parameters, is illustrated in Figure 4.4.

Prior model (population dynamics):

Recruitment: $X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}_t, \delta_x \sim N(Z_{f,t}\exp(\alpha_0+\alpha_t), Z_{f,t}\exp(\delta_x+\alpha_t))$
Juvenile survival: $Y_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa \sim Bin(X_{f,t-1}, \exp(-\kappa)\ X_{f,t-1}^{-\beta} Y_{f,t-1}^{-\gamma})$
Adult survival: $Z_{f,t} \mid \text{past}(t), \theta \sim Bin((Z_{f,t-1}+Y_{f,t-1}), \exp(-\theta))$

An overview of the hyper-priors is given in Table 5.2 in Section 5.1.

Likelihood (sampling process):

0-group: $X_{i,t}^o \mid X_{f,t}, p_{i,t}^x \sim Bin(X_{f,t}, p_{i,t}^x)$
1-group: $Y_{i,t}^o \mid Y_{f,t}, p_{i,t}^y \sim Bin(Y_{f,t}, p_{i,t}^y)$

The sampling probabilities are given by

$$p_{i,t}^x = p_{i,t}^{s,x} p_{i,t}^{m,x} p_{i,t}^{o,x} \text{ and } p_{i,t}^y = p_{i,t}^{s,y} p_{i,t}^{m,y} p_{i,t}^{o,y},$$

and the prior specifications of the components of $p_{i,t}^x$ and $p_{i,t}^y$ are given in Table 4.2.

Table 4.1: A summary of the Bayesian hierarchical model for the population dynamics and the sampling process.

Settling probability	
0-group	$\text{logit}(p_{i,t}^{s,x}) = \log(L/\Omega(f_i)) + \sum_{k=1}^{5} c_k^{s,x} I_{\{d_{i,t}=k\}}$, $p(c_k^{s,x}) \sim N(\mu_{c,x}, \tau_{c,x})$ $\mu_{c,x}$, $\tau_{c,x}$ fixed
1-group	$\text{logit}(p_{i,t}^{s,y}) = \log(L/\Omega(f_i)) + \sum_{k=1}^{5} c_k^{s,y} I_{\{d_{i,t}=k\}}$, $p(c_k^{s,y}) \sim N(\mu_{c,y}, \tau_{c,y})$ $\mu_{c,y}$, $\tau_{c,y}$ fixed
Probability of being catchable	
0-group	$p_{i,t}^{m,x} = \psi_x; \quad \forall(i,t)$ ψ_x fixed
1-group	$p_{i,t}^{m,y} = \psi_y; \quad \forall(i,t)$ ψ_y fixed
Probability of not escaping	
0-group	$p_{i,t}^{o,x} = \sum_{k=1}^{5} p_k^{o,x} I_{[d_{i,t}=k]};\ k = 1, \ldots, 5$ $p_k^{o,x};\ k = 1, \ldots, 5$ fixed
1-group	$p_{i,t}^{o,y} = \sum_{k=1}^{5} p_k^{o,y} I_{[d_{i,t}=k]};\ k = 1, \ldots, 5$ $p_k^{o,y};\ k = 1, \ldots, 5$ fixed

Table 4.2: Prior specifications for the sampling probabilities of the likelihood.

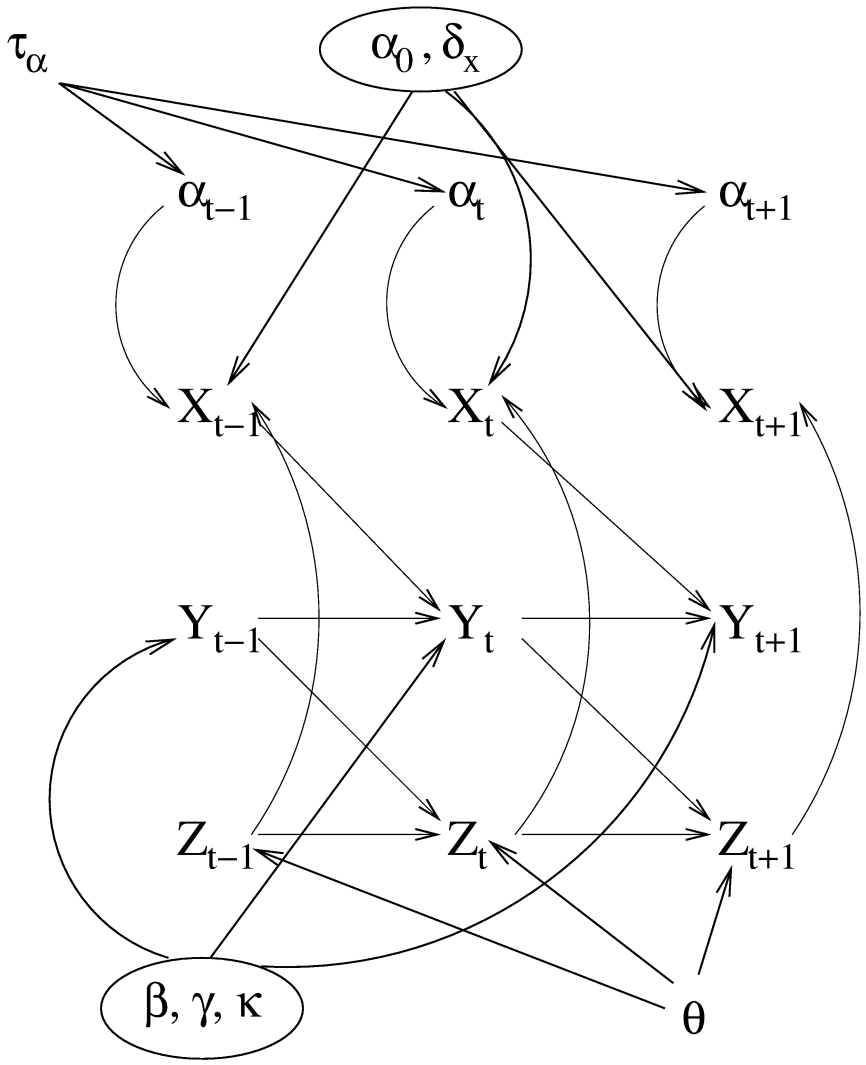

Figure 4.4: Graphical description of the prior model for the population dynamics, including the parameters. The age-groups are denoted by X_t (0-group), Y_t (1-group) and Z_t (mature cod), where the fjord index is suppressed.

Chapter 5

A sampling based approach to parameter estimation

The parameters of the hierarchical Bayesian model are estimated using Markov chain Monte Carlo methods (Gilks, Richardson and Spiegelhalter, 1996; Robert and Casella, 1999). The unknowns to be estimated are the survival and recruitment parameters β, γ, κ, θ, α_0 and $\{\alpha_t\}_{t=1,\ldots,n_t}$, the over-dispersion parameter δ_x, the parameters $\{c_k^{s,x}\}_{k=1,\ldots,5}$ and $\{c_k^{s,y}\}_{k=1,\ldots,5}$ representing the effect of bottom vegetation on the sampling process, and the abundances $\{X_{f,t}\}$, $\{Y_{f,t}\}$ and $\{Z_{f,t}\}$ for all fjords f and years t. In this chapter we study the results from estimating the model parameters as well as the abundances using the Metropolis-Hastings sampler. The parameters $\{\alpha_t\}_t$, $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ are updated in three blocks, one for each group of parameters, while the remaining parameters and the abundances are updated using single-site Metropolis-Hastings steps.

To complete the model, the hyper-parameters of the population dynamics prior and the likelihood defined in Chapter 4 should be specified, and these specifications are given in Section 5.1. Some details on the implementation of the algorithm are given in Section 5.2. To study the convergence and mixing properties of the algorithm, we apply the algorithm to a simulated data set and describe how re-parameterising the model leads to improved mixing. The results of this simulation study are presented in Sections 5.3 and 5.4. The results from applying the re-parametrised model to the Flødevigen data are given in Chapter 6.

5.1 Specification of hyper-parameters

The hyper-parameters are specified on the basis of reasonable choices of prior means and prior ranges of the parameters, partly on the basis of other studies. We first con-

sider the parameters of the sampling process, as defined in Table 4.2. The prior means of the parameters $c_k^{s,x};\ k = 1, \ldots, 5$ of the settling probabilities (4.20), and the similar parameters for the 1-group cod, denoted $c_k^{s,y};\ k = 1, \ldots, 5$, are set equal to zero. The prior variances are specified such that the interval $(-\log(100), \log(100))$ corresponds approximately to four times the standard deviation. As explained in Section 2.1, the 1-group cod tend to stay further away from the shore than the 0-group data, and thus the beach seine catches are less efficient for these data. Therefore, the probability of being catchable is assumed to be larger for the 0-group than the 1-group cod, and the corresponding parameters are fixed to $\psi_x = 0.4$ and $\psi_y = 0.1$ based on biological expert opinion. This means that it is assumed that on the average 40% of the 0-group cod and 10% of the 1-group cod will be in the region covered by the beach seine at the time of sampling. In Section 5.4 we assess the sensitivity of the estimated parameters on the choice of values for ψ_x and ψ_y. The probability of escaping was defined to be a function of the amount of bottom vegetation, represented by the categorical covariate $d_{i,t}$. Assuming that the 0-group cod, which are the smallest, have the highest chances of escaping, these probabilities are fixed according to the values given in Table 5.1.

$d_{i,t}$	$p_{i,t}^{o,x}$	$p_{i,t}^{o,y}$
5 (bottom totally covered)	0.8000	0.8500
4 (many plants)	0.8375	0.8750
3 (some plants)	0.8750	0.9000
2 (few plants)	0.9125	0.9250
1 (no vegetation)	0.9500	0.9500

Table 5.1: Fixed values of the probabilities of not escaping as a function of bottom coverage $d_{i,t}$.

We now turn to the parameters of the population dynamics prior. We need to specify the hyper-parameters for the log recruitment rate parameters α_0 and $\alpha_t; t = 1, \ldots, n_t$ in (4.2), the parameters β, γ and κ in (4.5), and the parameter θ in (4.7). The probability of survival for the juvenile cod is given from (4.5) by $p_y = X_{f,t-1}^{-\beta} Y_{f,t-1}^{-\gamma} \exp(-\kappa)$. Using the estimated monthly survival rates found in Julliard et al. (2001), the probability of survival from 0-group cod in September year t to 1-group cod in September year $t + 1$ is approximately $\exp(-2.5) = 0.08$. If we let $\kappa = 0$ and $\beta = \gamma$, we can obtain a set of fjord specific values of β by using rough estimates of the 0-group and 1-group abundances at fjord level. Assuming that the cod show no habitat preferences and settle uniformly along the coast, a rough estimate of the abundances can be found by upscaling of the observations at station level by the factor $\Omega(f)/(0.9\psi_x L)$ for the 0-group and $\Omega(f)/(0.9\psi_y L)$ for the 1-group using mean levels of the probabilities of settling, being catchable and not escaping given above. Using this strategy, we obtain a set of estimates of β in the range 0.1 to 0.15 corresponding to a survival rate of $\exp(-2.5)$. Further, a value of $\beta = \gamma = 0.5$ leads to fjord specific survival rates which are all less

than 0.001, which seems well below a reasonable limit for the survival probabilities. We take $E(\beta) = E(\gamma) = 0.2$, and fix the hyper-parameters of the Gamma-prior of these to parameters such that $P(\beta > 0.5) \approx 0.05$, leading to Gamma$(2, 10)$-priors.

For κ, we fix the hyper-parameters such that $P(\kappa > 3) \approx 0.05$, and choosing, somewhat arbitrarily, $E(\kappa) = 1$, we arrive at the prior $\kappa \sim$ Gamma$(1.0, 1.0)$. For θ, we use the results from Julliard et al. (2001), where the adult survival rate $\exp(-\theta)$ was estimated to $\exp(-1)$. Further assuming that the adult survival rate is not very likely to be less than $\exp(-3) = 0.05$, we choose $\theta \sim$ Gamma$(1.0, 1.0)$ as for κ.

We now consider the parameters of the recruitment process (4.2). The parameters in the Gamma prior for α_0 is specified such that $E(\alpha_0) = 5$. Assuming a relatively vague prior, we let $\alpha_0 \sim$ Gamma$(0.5, 0.1)$. The parameters modelling the temporal structure in the log recruitment rate, α_t, are assigned equal Normal$(\mu_\alpha, \tau_\alpha^{-1})$ priors with $\mu_\alpha = 0$ and $\tau_\alpha \sim$ Gamma$(0.01, 0.01)$, but they are constrained to sum to zero. Finally, the parameters of the truncated Normal prior distribution of δ_x is $\mu_{\delta_x} = 10$ and $\tau_{\delta_x} = 25^{-1}$.

A summary of the parameters of the population dynamics prior model and the likelihood, their prior distributions and hyper-parameters is given in Table 5.2.

Parameter	Status	Prior distr./value	Hyper-parameters
Population dynamics prior:			
α_0	estimated	Gamma	$\alpha_\alpha = 0.5$, $\beta_\alpha = 0.1$
β	estimated	Gamma	$\alpha_\beta = 2.0$, $\beta_\beta = 10.0$
γ	estimated	Gamma	$\alpha_\gamma = 2.0$, $\beta_\gamma = 10.0$
θ	estimated	Gamma	$\alpha_\theta = 1.0$, $\beta_\theta = 1.0$
κ	estimated	Gamma	$\alpha_\kappa = 1.0$, $\beta_\kappa = 1.0$
α_t	estimated	Normal	$\mu_\alpha = 0.0$, τ_α
τ_α	estimated	Gamma	$\alpha_\tau = 0.01$, $\beta_\tau = 0.01$
δ_x	estimated	Normal (truncated)	$\mu_{\delta_x} = 10$, $\tau_{\delta_x} = 25^{-1}$
Likelihood:			
ψ_x	fixed	$\psi_x = 0.4$	
$p_k^{o,x}; k = 1, \ldots, 5$	fixed	See Table 5.1	
ψ_y	fixed	$\psi_y = 0.1$	
$p_k^{o,y}; k = 1, \ldots, 5$	fixed	See Table 5.1	
$c_k^{s,x}; k = 1, \ldots, 5$	estimated	Normal	$\mu_{c,x} = 0.0, \tau_{c,x} = 5.0^{-1}$
$c_k^{s,y}; k = 1, \ldots, 5$	estimated	Normal	$\mu_{c,y} = 0.0, \tau_{c,y} = 5.0^{-1}$

Table 5.2: Summary of the parameters of the prior model of the abundances and of the likelihood model. The status refers to whether the parameters are to be estimated or are kept fixed.

5.2 Implementation details and preliminary results

The estimation procedure is implemented in C++. The Metropolis-Hastings updates are based on random walk proposals, and the proposal distributions are tuned during a set of initial iterations to give reasonable acceptance rates. The acceptance rates for the model parameters for the final runs using the simulated data or the Flødevigen data are in the range 6-58%, For the simulated data, the acceptance rates for the abundances are between 16 and 47%, while for the Flødevigen data, they are in the range 7-73%.

In the Binomial likelihoods (4.15) and (4.18) for the 0-group and 1-group cod respectively, the abundances $X_{f,t}$ and $Y_{f,t}$ will typically be large, and the sampling probabilities $p^x_{i,t}$ and $p^y_{i,t}$ are small. For computational convenience, we use a Poisson approximation to these Binomial likelihoods for large values of the updates of the abundances.

About 19% of the values of the covariate $d_{i,t}$ are missing for the whole set of monitoring stations, and about 14% for the two fjords in the Risør area. For each iteration of the sampler, the missing values are replaced by a simulated value. In the case of a missing value at time t, we fill in a value sampled in the range $[d_{low}, d_{high}]$, were d_{low} and d_{high} are the smallest and largest values measured at time $t-1$ and $t+1$ at that location. If one of these two is also missing, the corresponding non-missing value closest to t is used.

Initial runs of the sampling algorithm on the Risør data presented in Figure 2.8, indicate that the mixing is poor and the convergence is slow. The convergence and mixing of a single-site Metropolis-Hastings algorithm might be improved by re-parameterising the model, such that the posterior correlations between the parameters are eliminated or reduced (Gilks and Roberts, 1996). To assess the performance of the estimation procedure and the effect of re-parameterisation, we conduct a simulation study running the algorithm on a simulated data set. The simulated data and the results of the study are presented in Section 5.3.

5.3 Simulation study

In this section, we assess the performance of our estimation procedure based on a set of simulated data, described in Section 5.3.1. In Section 5.3.2, results from using the original parameterisation are presented. The effect of re-parameterising the model on the convergence and mixing of the sampling algorithm is studied in Sections 5.3.3 and 5.3.4.

5.3.1 The simulated data set

The simulated data set is generated such as to mimic the subset of data from the monitoring stations in the two Risør fjords. We simulate the Bayesian dynamic model over a period of $n_t = 38$ years using the parameters shown in Table 5.3 and Figure 5.1. The values of the covariate $d_{i,t}$ representing the amount of bottom vegetation are taken to be those of the 12 monitoring stations of the Risør area. The random effects α_t, representing the temporal structure of the log recruitment rate, are set equal to the estimated values from a preliminary run on the Risør data. This time series of estimated values indicate that there is no strong temporal structure in the log recruitment rate, and from the QQ-plot at the bottom right in Figure 5.1, assuming the random effects $\alpha_t; t = 1, \ldots, n_t$ to be Normal with common variance seems appropriate. Their empirical variance is 1.08, and the prior mean of τ_α is set to 1.0. The simulated time series of the 0-group, 1-group and adult populations, together with the data sampled using the likelihood model in (4.15) and (4.18), are given in the middle and bottom panels of Figure 5.1.

	Population dynamics prior:					
Parameter	β	γ	κ	θ	α_0	δ_x
Value	0.1	0.1	0.5	1.0	3.0	10.0
	Sampling model (likelihood):					
	0-group:					
Parameter	$p_1^{o,x}$	$p_2^{o,x}$	$p_3^{o,x}$	$p_4^{o,x}$	$p_5^{o,x}$	ψ_x
Value	0.95	0.9125	0.875	0.8375	0.8	0.4
	1-group:					
Parameter	$p_1^{o,y}$	$p_2^{o,y}$	$p_3^{o,y}$	$p_4^{o,y}$	$p_5^{o,y}$	ψ_y
Value	0.95	0.925	0.9	0.875	0.85	0.1

Table 5.3: The parameters of the simulated data set.

5.3.2 Results from the original parameterisation

Preliminary runs indicate that estimating all parameters based on the simulated data set using the original parameterisation leads to poor mixing for several of the parameters. To get some initial insight into the structure of the parameter estimation problem, the likelihood surface for the posterior model can be studied. In our problem, this surface is high-dimensional and of a complicated form. However, some information can be obtained by studying plots of marginal and pairwise joint full conditional distributions defined for the sampling based estimation algorithm. In Figure 5.2 we have plotted the full conditional distribution for the parameters β, γ, κ, θ, α_0 and δ_x. Similar plots of joint

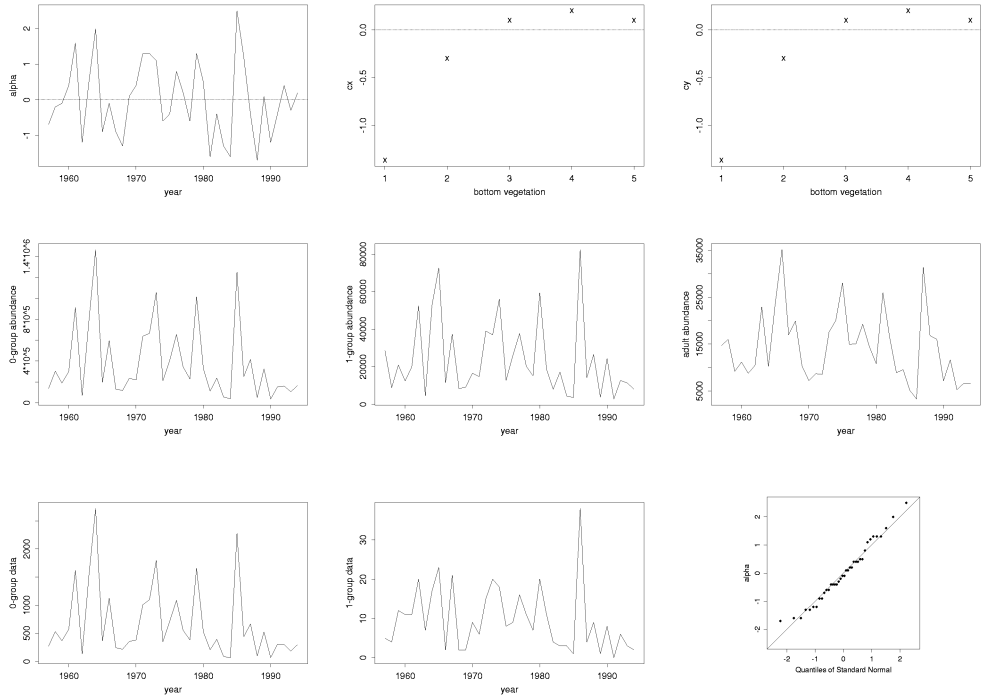

Figure 5.1: The recruitment parameters $\{\alpha_t\}_t$ and the settling probability factors $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ of the simulated data set (top row), the simulated time series of abundances (middle row) and the simulated 0-group and 1-group data aggregated over all locations in the fjords (bottom). The bottom right panel shows a QQ-plot of the recruitment parameter vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_{n_t})^T$.

full conditionals of pairs of parameters are given in Figure 5.3. In both plots, we have conditioned on the true abundances of the simulated data set as well as on true values on the remaining parameters. From Figure 5.2 we observe that the probability mass is highly concentrated for β, γ and θ, and has clear, but less concentrated peaks for κ, α_0 and δ_x. The most striking feature of Figure 5.3 is the strong conditional dependency between the parameters β, γ and κ of the prior model (4.5) of juvenile survival. The other parameters appear to be conditionally independent of each other as well as of β, γ and κ, and the modes are well identified.

To assess the variability of the marginal full conditional distributions, we compute rough estimates of the coefficient of variation (cv) from empirical 2.5% and 97.5% quantiles of these distributions. The resulting values are $cv_\beta = 0.8e^{-4}$, $cv_\gamma = 1.0e^{-4}$, $cv_\kappa =$

$1.0e^{-2}$, $cv_\theta = 1.1e^{-2}$, $cv_{\alpha_0} = 0.8e^{-1}$ and $cv_{\delta_x} = 2.3$. The over-dispersion parameter δ_x in (4.2) has by far the largest coefficient of variation and is thus expected to be the parameter that is most difficult to identify. Based on these results as well as the preliminary runs, we decide to keep δ_x fixed at its true value $\delta_x = 10$ throughout the simulation study.

For the original model, as for all re-parameterisations, the MCMC algorithm was run using 5.1 million iterations, after some initial runs used to tune the algorithm. Trace plots are generated using every 1000th iteration, discarding the initial tuning iterations, leaving a total of 5100 updates. The trace plots for a selection of the parameters of the population dynamics using the original model, which we will denote by Model I, are shown in Figure 5.4. The mixing is poor and the convergence is slow for all parameters, except for the precision τ_α. Further, the trace plots indicate that there is high cross-correlation between the simulated Markov chains for the parameters θ, representing adult mortality and the mean log recruitment rate α_0. This is also clearly seen from the scatter plots in Figure 5.5. There is also some evidence of correlation between β and γ, and between β and κ, but not so strong.

The parameters α_0 and θ are contained in the population priors involving the unknown adult population abundances Z_t, for which we have no data. In Figure 5.6, the MCMC updates for the time-independent parameters are plotted against the updates of the log abundances for year 1976, which is in the centre of the period of study and should be least affected by boundary effects. We observe that the abundance estimate of the adult population Z_{1976} is strongly correlated to α_0 and θ. The levels of β, γ and κ seem to be independent of the adult abundance, but there is some evidence of increasing simulation variance for increasing values of Z_{1976}. The empirical correlation between θ and $Z_{f,t}$ over the last 30 years of the period of study, discarding the initial 8 years, are in the range $[-0.94, -0.88]$. In contrast, all parameters appear to be independent of the 0-group and 1-group abundance estimates. To improve mixing and convergence, we explore different parameterisations of the prior model defined by the conditional distributions in (4.2), (4.5) and (4.7). The re-parametrised models are defined in the next subsection, and the results are given in Section 5.3.4.

5.3.3 Re-parameterising the model

By re-parameterising the model, we aim at reducing the correlation between the parameters to be estimated. Based on the results for the original model, as well as the joint full conditional distributions displayed in Figure 5.3, we now motivate and describe the different re-parameterisations for which we explore the convergence and mixing properties of the corresponding sampling based estimation algorithm. A summary of the different re-parameterisations is given in Table 5.4. The first re-parameterisation, denoted Model II, is based on the evidence of strong correlation between α_0 and θ from

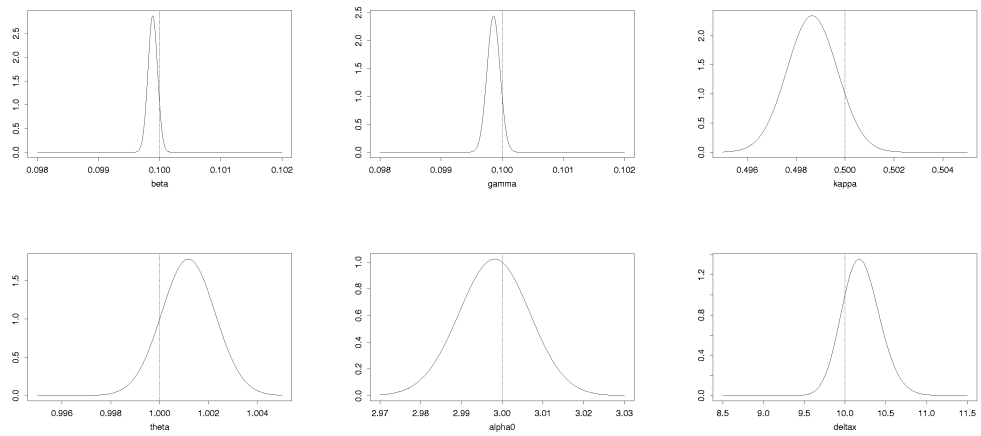

Figure 5.2: Full conditional distributions of the global parameters of the population dynamics. The vertical lines indicate the true values.

the scatter plots of Figure 5.5. This observation contrasts with the results displayed in Figure 5.3, but recall that those results were obtained conditionally on the abundances. In Figure 5.5 we have also included the updates of the parameter given by the difference $\alpha_0 - \theta$. The correlation between this parameters and θ or α_0 is less than the correlation between the two original parameters. Therefore, we re-parameterise the model introducing the parameter $\eta = \alpha_0 - \theta$, and update θ and η instead of θ and α_0. A further motivation of this choice of η is given by the expression obtained by substituting the conditional prior mean of $Z_{f,t}$, $\mathrm{E}(Z_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\})$, for $Z_{f,t}$ in the conditional prior mean of $X_{f,t}$. From (4.2) and (4.7) we have that

$$E(X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}) = Z_{f,t} \exp(\alpha_0 + \alpha_t), \quad \text{and} \tag{5.1}$$

$$E(Z_{f,t} \mid \text{past}(t), \theta) = (Z_{f,t-1} + Y_{f,t-1}) \exp(-\theta). \tag{5.2}$$

Substituting the mean (5.2) for $Z_{f,t}$ in (5.1), we get

$$\begin{aligned} E(X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}) &= (Z_{f,t-1} + Y_{f,t-1}) \exp(-\theta) \exp(\alpha_0 + \alpha_t) \\ &= (Z_{f,t-1} + Y_{f,t-1}) \exp(\alpha_0 - \theta + \alpha_t) \end{aligned} \tag{5.3}$$

Even if the parameter θ is also contained in the conditional prior variance of $Z_{f,t}$, the parameters θ and α_0 might suffer from near non-identifiability, and $\eta = \alpha_0 - \theta$ will possibly be better identified.

Since data are available for the 0-group and 1-group cod but not for the adult cod, the parameters β, γ and κ, governing the survival from the 0-group to the 1-group stage

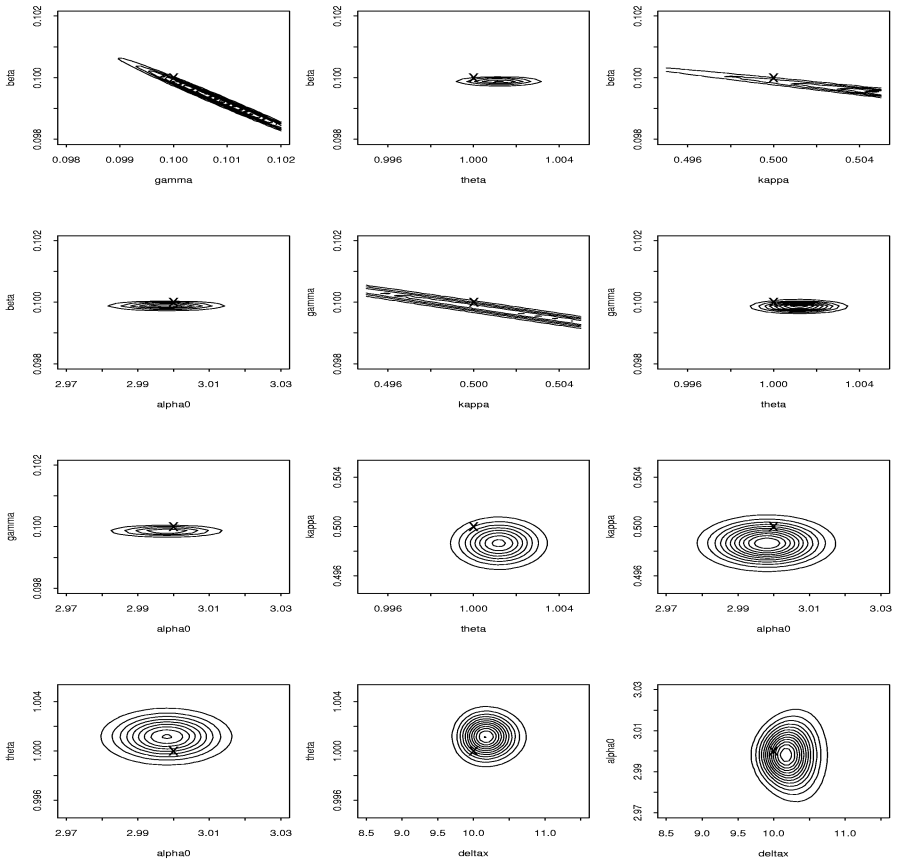

Figure 5.3: Plots of joint full conditional distributions of pairs of parameters. The true values are marked with an 'x'.

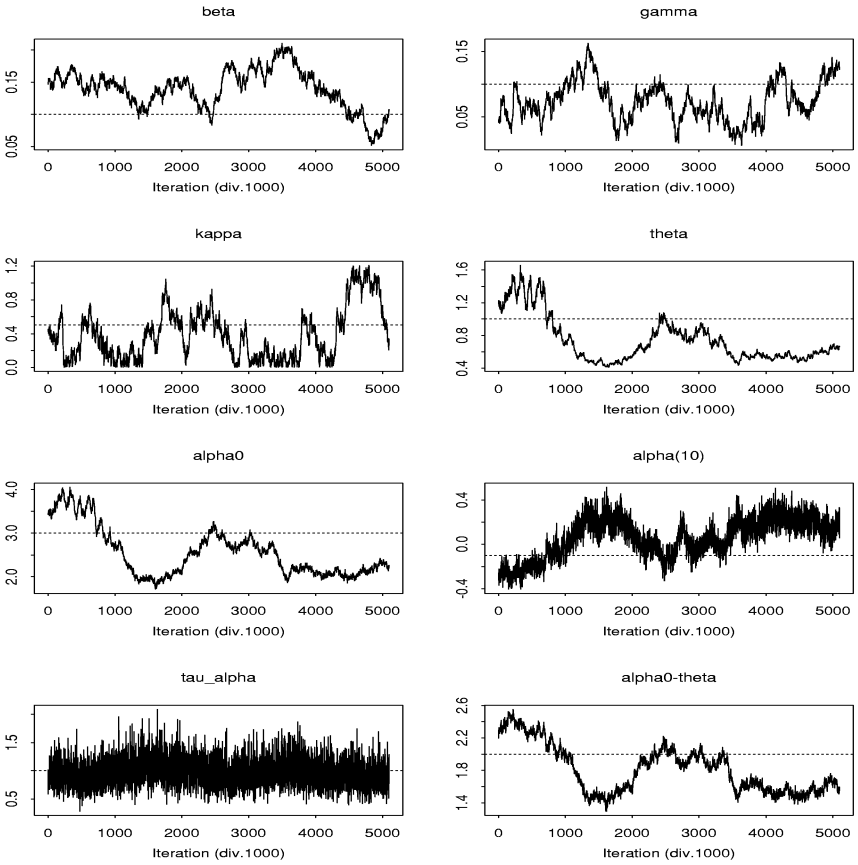

Figure 5.4: Trace plots for a selection of the parameters using the original model. Every 1000th iteration is shown. The horizontal dotted lines indicate the values used to generate the simulated data set.

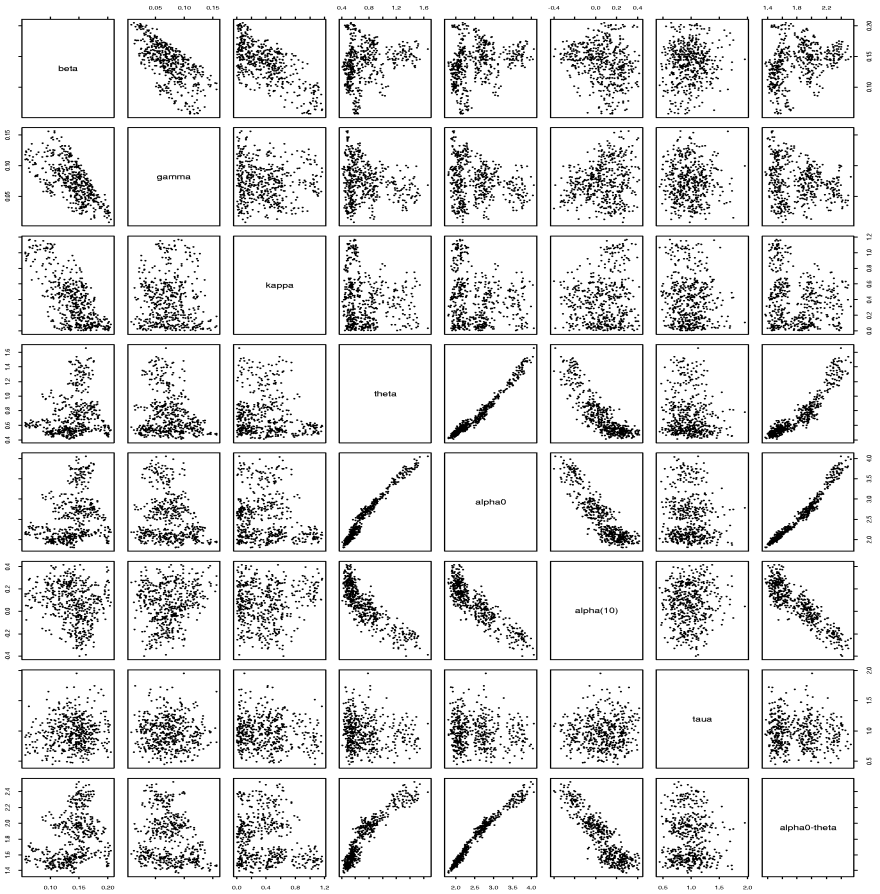

Figure 5.5: Scatter plots based on results from the original model. The Markov chains are thinned by selecting 500 evenly spaced values.

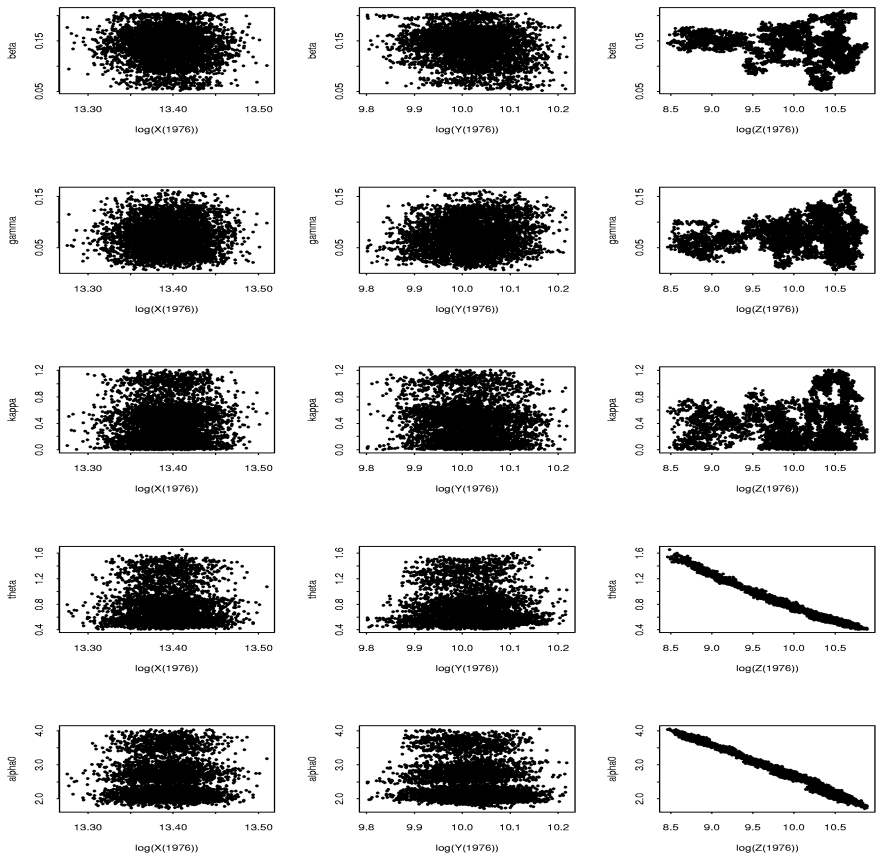

Figure 5.6: Scatter plots of the global parameters of the population dynamics against log estimated abundances for year 1976 and for Model I.

Original model Model I	Model II	Model III
β γ κ θ α_0	β γ κ θ $(\alpha_0 = \eta + \theta)$	β γ $(\kappa = \kappa_1 - \beta\,\overline{\log(X)} - \gamma\,\overline{\log(Y)})$ θ $(\alpha_0 = \eta + \theta)$
	$\eta = \alpha_0 - \theta$	$\eta = \alpha_0 - \theta$ $\kappa_1 = \kappa + \beta\,\overline{\log(X)} + \gamma\,\overline{\log(Y)}$

Model IV	Model V
β γ $(\kappa = \kappa_1 - \beta\,\overline{\log(X)} - \gamma\,\overline{\log(Y)})$ θ $(\alpha_0 = \eta + \theta)$	β γ $(\kappa = \kappa_1 - \beta\,\overline{\log(X)} - \gamma\,\overline{\log(Y)})$ θ $(\alpha_0 = \alpha_0^{re} - \overline{\log(Z)})$
$\eta = \alpha_0 - \theta$ $(\kappa_1 = \eta - \nu)$ $\nu = \eta - \kappa_1$	$\kappa_1 = \kappa + \beta\,\overline{\log(X)} + \gamma\,\overline{\log(Y)}$ $\alpha_0^{re} = \alpha_0 + \overline{\log(Z)})$

Table 5.4: Parameters of the different re-parametrised models. The parameters in parentheses are not estimated, but at each iteration computed from the estimated parameters by the given relation. For further details, see the text.

are expected to be more easily identified that the parameters α_0 and θ. However, the trace plots in Figure 5.4 indicate that the mixing for these three parameters is slow, and for κ, the estimation variance is very large. Observe from Figure 5.5 that β and κ as well as β and γ are negatively correlated, an observation that is consistent with the joint full conditionals shown in Figure 5.3. The mean of the Binomial conditional prior distribution of the survival of juvenile cod, as defined in (4.4), is log-linear in $X_{f,t-1}$ and $Y_{f,t-1}$, with κ representing an intercept in the log-linear model. In analogy to linear regression, the parameters κ_1, β and γ in the re-parametrised model

$$\exp(-\kappa_1 - \beta(\log X_{f,t-1} - \overline{\log X_{\cdot,\cdot}}) - \gamma(\log Y_{f,t-1} - \overline{\log Y_{\cdot,\cdot}})) \tag{5.4}$$

for the prior survival probability of the 0-group cod, as given by (4.5), are expected to be less correlated than κ, β and γ of the original parameterisation. In our model, the covariates $\log X_{f,t-1}$ and $\log Y_{f,t-1}$ are themselves unknowns to be estimated, and using this parametrisation, care should be taken such that all parameters, including the population abundances, are assigned valid values at each update. We will denote by

Model III the re-parameterised model obtained by introducing (5.4) to Model II.

In Model II, we combined the global parameters θ and α_0. If we continue the reasoning based on substitution of means motivating Model II, we can also include the parameter κ_1 in (5.4), representing the average level of juvenile survival, in a linear combination of density-independent parameters. Substituting $E(Y_{f,t-1})$ for $Y_{f,t-1}$ in (5.3), we get (suppressing conditioning on the past and the parameters in the notation)

$$E(X_{f,t}) = (Z_{f,t-1} + \left(\frac{X_{f,t-2}}{\exp(\overline{\log X_{.,.}})}\right)^{1-\beta} \left(\frac{Y_{f,t-2}}{\exp(\overline{\log Y_{.,.}})}\right)^{1-\gamma} \exp(-\kappa_1))\ \exp(-\theta + \alpha_0 + \alpha_t). \tag{5.5}$$

We introduce the parameter

$$\nu = \alpha_0 - \theta - \kappa_1, \tag{5.6}$$

estimate θ, η and ν, and compute α_0 and κ_1 from (5.6) and the relation $\eta = \alpha_0 - \theta$. The resulting re-parameterised model is denoted Model IV.

Finally, we re-parameterise the prior model for the recruitment of new individuals (4.5) by centring the log adult population abundances, $Z_{f,t}$, resulting in Model V. This transformation is motivated by the centring approach to re-parameterisation of random-effects models suggested by Vines, Gilks and Wild (1996), denoted *sweeping* by Gilks and Roberts (1996). The conditional prior mean of $X_{f,t}$ defined in (4.1) can be written as

$$\begin{aligned} \mathrm{E}(X_{f,t} \mid \mathrm{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}) &= \exp(\log(Z_{f,t}) + \alpha_0 + \alpha_t) \\ &= \exp(\log(Z_{f,t}) - \overline{\log(Z_{.,.})} + \overline{\log(Z_{.,.})} + \alpha_0 + \alpha_t) \\ &= \exp(\log(Z_{f,t}) - \overline{\log(Z_{.,.})} + \alpha_0^{re} + \alpha_t), \end{aligned} \tag{5.7}$$

where

$$\alpha_0^{re} = \alpha_0 + \overline{\log(Z_{.,.})}. \tag{5.8}$$

The log mean of the abundances is swept from the abundances to the overall mean of the log recruitment rate, α_0.

5.3.4 Comparison of the results from different re-parameterisations

In this subsection we present the main results from each re-parameterisation and discuss the effect of each model in terms of mixing and convergence of the parameters that are estimated.

In Figure 5.7 trace plots for the parameters involved in the re-parametrisation in Model II are plotted. Comparing the plot of α_0 against θ for the original model to the plot of the new parameter η against θ for the re-parametrised model, also shown in Figure 5.7, we observe that the degree of dependency has not been reduced to any large extent. So not very much is gained using this re-parameterisation.

In contrast, the centring of the log abundances in Model III results in a substantial improvement of the mixing for β, γ and κ, as illustrated in Figure 5.8. We also observe that the correlation between β and the intercept term κ_1 of the re-parameterised model is less than the correlation between β and κ; the absolute value of the empirical correlation between the updates is reduced from 0.69 to 0.36. The true values $\beta = 0.1$, $\gamma = 0.1$ and $\kappa = 0.5$ are all within the credibility intervals generated from empirical 2.5% and 97.5% quantiles. These are $[0.065, 0.194]$ for β, $[0.014, 0.144]$ for γ and $[0.013, 1.417]$ for κ. The convergence remains slow for the parameters α_0 and θ, and these parameters are still strongly correlated to the mature abundance, as illustrated by the scatter plots of Figure 5.9.

From the results for running Model IV (not shown), we find that the convergence and mixing of ν is much slower than that of κ_1, such that nothing is gained by including the parameter ν.

From the trace plot of the updates of α_0^{re} of Model V, shown in Figure 5.10, we can conclude that this parameter is substantially better behaved than α_0. From (5.8) we have that

$$\exp(\alpha_0^{re}) = \exp(\alpha_0)\exp(\overline{\log(Z_{\cdot,\cdot})}), \tag{5.9}$$

and this factor can be regarded as a measure of the overall mean number of new recruits to the population per year, where the mean is taken over all fjords (in our simulated example one fjord) and the whole period of study. Thus, the results indicate that the mean number of recruits is a well behaved quantity in terms of mixing and convergence, but that it is difficult to identify the relative contributions of the abundance of mature cod and the recruitment rate per adult cod to this number.

To summarise the findings regarding the global parameters, trace plots for all parameterisation, except Model IV, are plotted together in Figure 5.11. From the plot it is apparent that the main improvement in convergence and mixing is achieved by applying the centring of re-parameterisation III. Also, we observe that the hyper-parameter τ_α is well-behaved for all parameterisations.

Trace plots for the temporal fluctuations of the log recruitment rate, represented by $\{\alpha_t\}_t$, are shown Figure 5.12 for five of the years. There seems to be no major differences in the mixing properties for the different parameterisations. There is some indication of heterogeneity in the sampling variance, and as illustrated for α_{10} using Model III in the bottom row of Figure 5.9, the updates of the α_t's are correlated to the updates of the mature abundances. The settling probability factors $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ of the sampling probabilities summarised in Table 4.2, are computationally well-behaved for all models, as can be seen from Figures 5.13 and 5.14. The estimated posterior means for $\{\alpha_t\}_t$, $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ are shown in Figure 5.15, together with the true values. The fit is good and the variability is small, except for $c_1^{s,x}$ and $c_1^{s,y}$ which are estimated on the basis of very few observations (< 10), and the first few elements of $\{\alpha_t\}_t$ suffering from initialisation effects. We observe that although the mixing of the α_t's was found

not to be quite satisfactory, and the updates are correlated to the adult abundances, the estimated posterior means indicate that the true temporal structure is well reproduced.

So far, we have focused on the effects on the model parameters. In Figure 5.16 we plot trace plots of the abundances for two years, year 1967 and 1968, for all parameterisations. Convergence and mixing is good for the 0-group and 1-group populations, while for all models the mixing is poor and convergence is very slow for the adult population abundance. Despite the apparent convergence problems for the abundance of adult cod, we proceed by presenting posterior mean estimates to get an indication of the fit to the data, but the poor mixing should be keep in mind when interpreting the results. Estimated and true abundances for the 38 years of study are plotted in Figure 5.17. These plots indicate that the 0-group population abundance is very well estimated. The fit of the 1-group population is quite good, but the abundances seem to be systematically slightly underestimated, particularly for the large abundance estimates. The estimated adult population fluctuates substantially between the re-parameterisations, as expected from the poor mixing.

In Figure 5.18 we further illustrate the behaviour of the abundance estimates, using the estimates based on Model III. Here, we have plotted the estimated posterior means of the abundances for the two juvenile age-groups as a function of time, together with the observed fjord total of sampled individuals of each age-group. The means are taken over every 1000th iteration, discarding the first 100000 iterations after tuning. To simplify visual comparison, the observations of the 0-group and 1-group cod are up-scaled by the constants $\Omega(f)/(n_f\psi_x L)$ and $\Omega(f)/(n_f\psi_y L)$ respectively, where n_f is the number of stations in fjord f. These up-scaling factors are analogous to the values used in Section 5.1 to specify the hyper-priors of the population dynamics priors. Also, the estimated time series for the 0-group and 1-group cod, the latter shifted back one year and up-scaled by a constant, are shown together. We observe that the temporal structure of the abundance estimates follows the data well for both age-groups, with the closest fit for the 0-group data. However, comparing the middle panel to the bottom panel, we are inclined to believe that the estimates for the 1-group cod owe more to the prior model than to the likelihood. Thus, we can conclude that the information inherent in the simulated data seems to be stronger for the 0-group than the 1-group population, and not strong enough to give reliable estimates for the adult populations using the current model.

For Model III we also ran the sampling algorithm for two alternative sets of initial values for the parameters. The results are illustrated by trace plots in Figures 5.19 and 5.20. We observe that different initial values leads to substantially different estimated posterior distributions for θ and α_0 as well as the mature abundances, confirming the lack of convergence for these parameters. The mixing properties of the α_t's are also variable, but although the corresponding posterior means will differ, the resulting temporal structures (not shown) are similar for all three runs. On the other hand, the estimated posterior distributions of β, γ and κ as well as the abundances of the 0-group and 1-

group cod are essentially the same for all runs.

Note that we have simulated our dataset such that the 1-group data contains a substantial number of zeros reflecting the structure of the original dataset for the Risør fjords. Also, the values of the covariates $d_{i,t}$ in the simulation study is chosen to match the values for the stations in this area. Therefore, we will expect the convergence problems to be of the same order or worse for the Risør data.

The parameters showing poorest convergence after re-parameterisation are α_0 and θ, which from Figure 5.9 are found to be highly correlated to the adult populations. However, referring to Figure 5.3, the mode of the joint full conditional distribution of α_0 and θ seems to be well identified. Running the program using the original model and keeping the populations fixed at their true values, this result is confirmed, as can be seen from the trace plots in Figure 5.21. The convergence is fast and mixing is good for all parameters. Observe that since the parameters of the population dynamics as defined in (4.2), (4.5) and (4.7) depend on the data only through the abundances, the data $\{X^o_{i,t}\}$ and $\{Y^o_{i,t}\}$ will not be informative for the population dynamics parameters when we condition on the true population sizes $\{X_{f,t}\}$, $\{Y_{f,t}\}$ and $\{Z_{f,t}\}$. Moreover, fixing the abundances, the estimation problem is equivalent to the problem of estimating the parameters of each of the three components of the prior model separately. Estimating all parameters of the recruitment process (4.2) based on abundances from one fjord only, problems related to near non-identifiability were experienced in the case of fixed abundances as well. Conditioning on the abundances, we observe that for the recruitment prior (4.2), the number of free parameters exceeds the number of pairs ($X_{f,t}$, $Z_{f,t}$). Although in principle no particular attention needs to be payed to this fact in a Bayesian as opposed to in a frequentist setting, near non-identifiability might be experienced if the priors are relatively vague. Therefore, the parameter δ_x was still kept fixed running the model conditionally on the abundances. However, we also simulate another data set, grouping the 12 sampling stations into two fjords, and estimate the model parameters conditionally on the two sets of abundances. In this case, the convergence and mixing are good for all parameters, including δ_x, as illustrated in Figure 5.22.

Since the parameters seem to be easily identified given the true population sizes, we conclude that the convergence and identifiability problems in the simulation based estimation are consequences of high correlations between some of the parameters and the estimates of the population abundances, in particular for the adult cod for which we have no data. As a consequence, a sampling based estimation approach where we sample directly from the marginal posterior of the model parameters, estimating the parameters only and not the abundances, is expected to improve convergence and mixing. However, this requires integration of the joint posterior over the time series of abundances for the 0-group, 1-group and adult populations. In Chapter 8 we describe an approach to this problem using the Metropolis-Hastings algorithm.

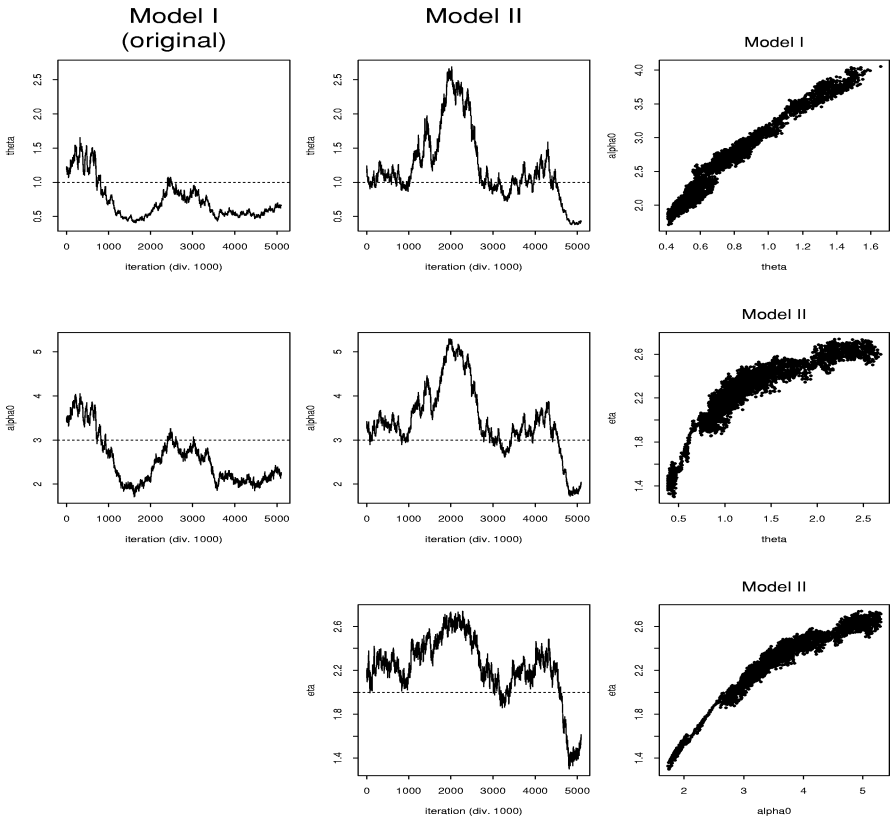

Figure 5.7: Results from Model II, re-parameterising α_0 and θ. Trace plots for α_0 and θ are shown for the original model (left) and the re-parametrised model (middle). The rightmost column shows scatter plots of some pairs of parameters.

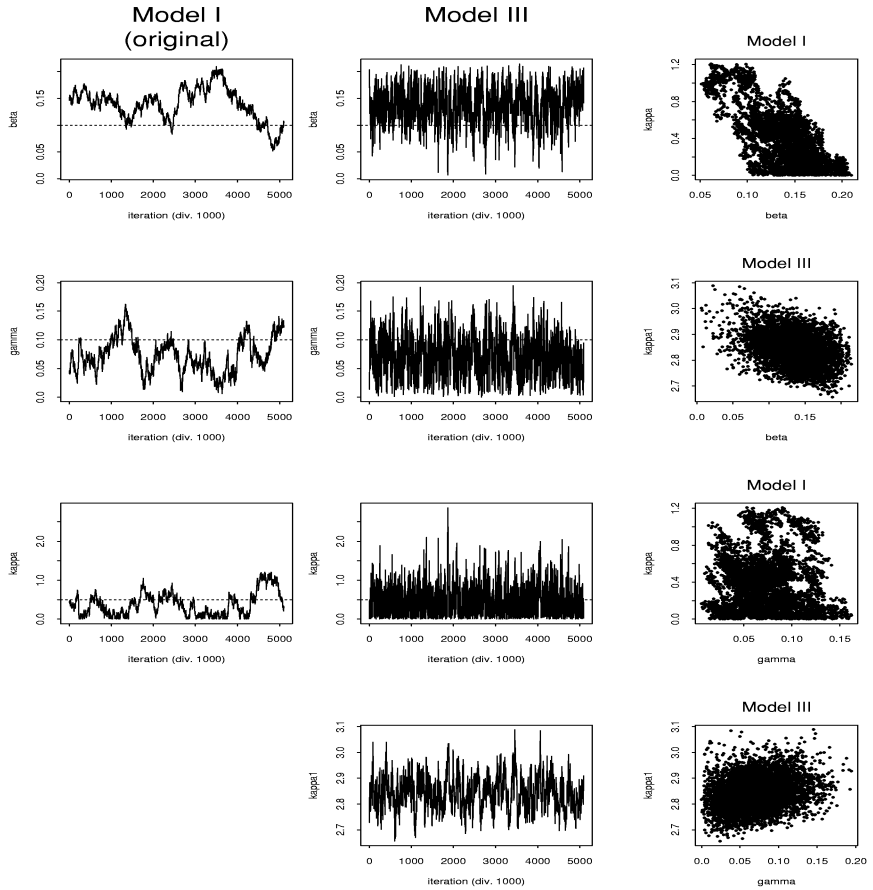

Figure 5.8: Results from Model III, re-parameterising the prior juvenile survival probability. The trace plots are shown for the original model (left) and the re-parametrised model (middle). The rightmost column shows scatter plots of some pairs of parameters.

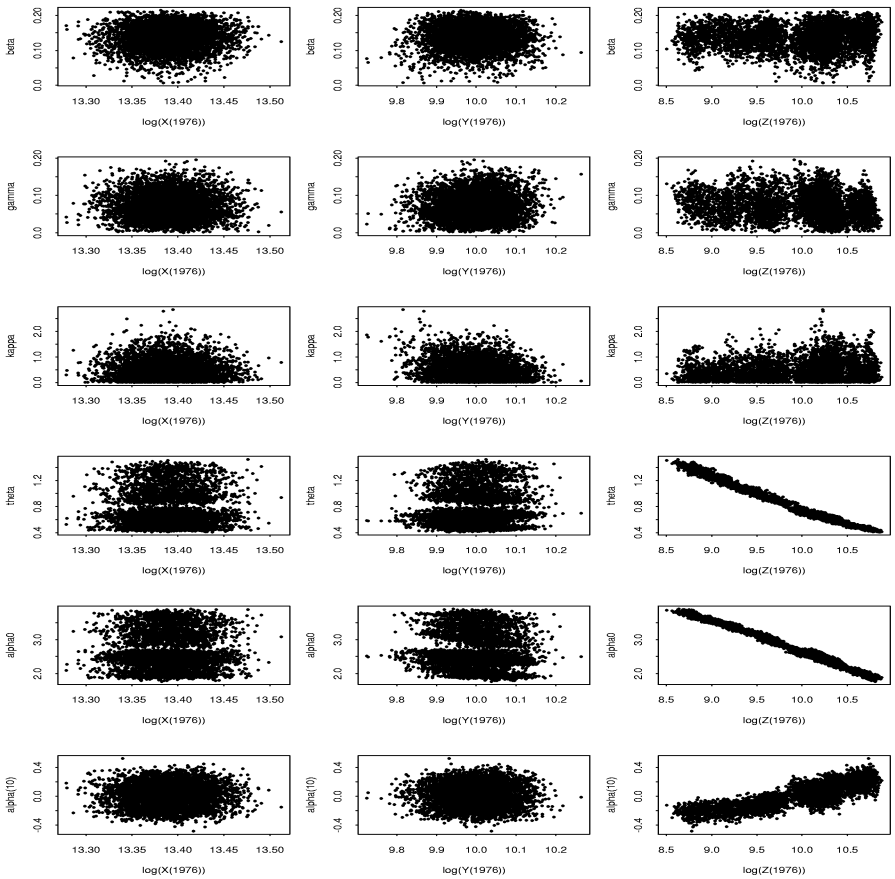

Figure 5.9: Scatter plots of a selection of parameters of the population dynamics against log estimated abundances for year 1976 and for Model III.

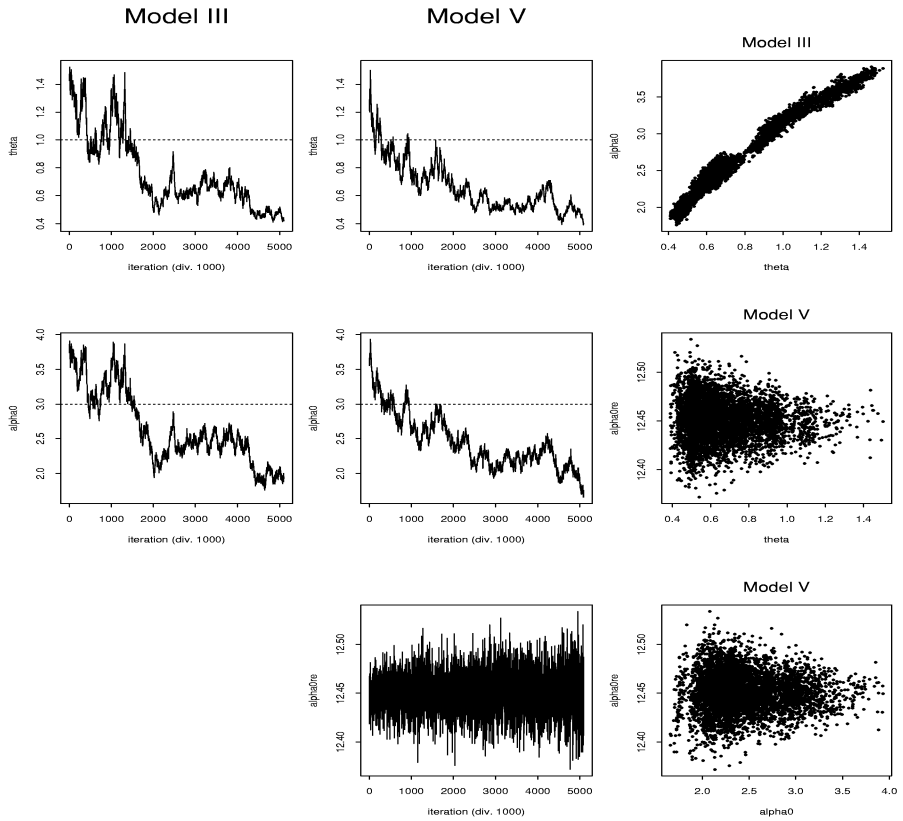

Figure 5.10: Results from Model V, re-parameterising α_0 by centring the log abundance of the adult cod, $\log(Z_{f,t})$. Trace plots for θ and α_0 are shown for the original model (left) and the re-parametrised model (middle). The rightmost column shows scatter plots of some pairs of parameters.

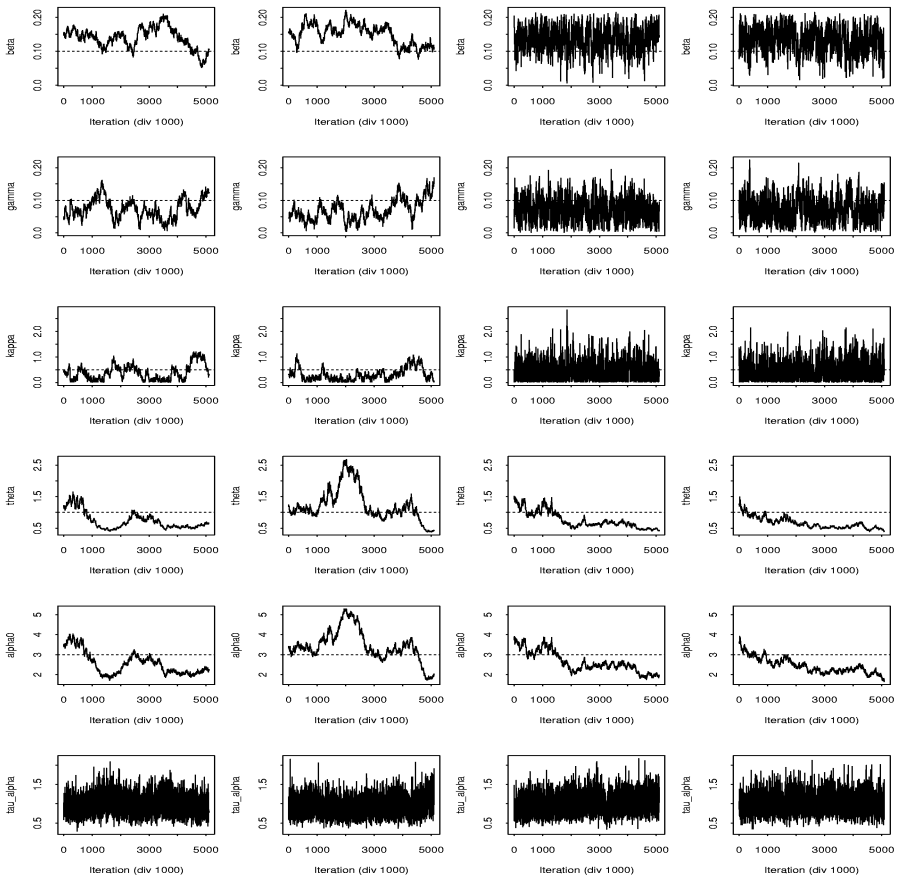

Figure 5.11: Trace plots for the parameters β, γ, κ, θ, α_0 and τ_α for (from left to right) the original model (Model I) and the re-parameterisations Model II, III and V.

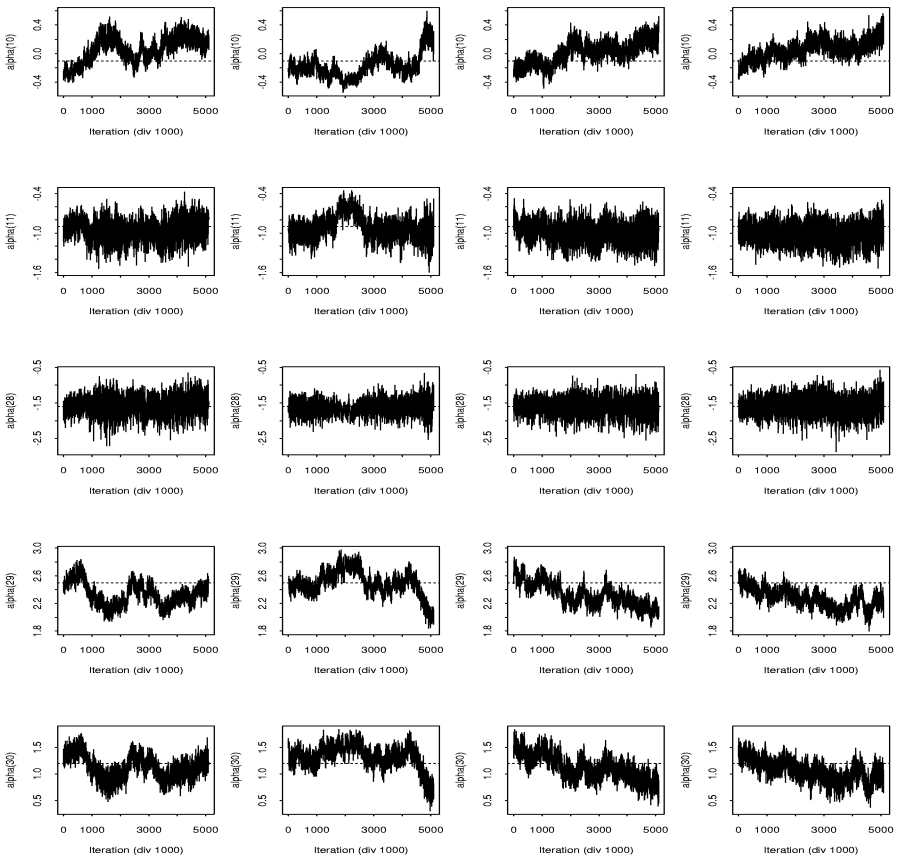

Figure 5.12: Trace plots for $\alpha_{10}, \alpha_{11}, \alpha_{28}, \alpha_{29}$ and α_{30} for (from left to right) the original model (Model I) and the re-parameterisations Model II, III and V.

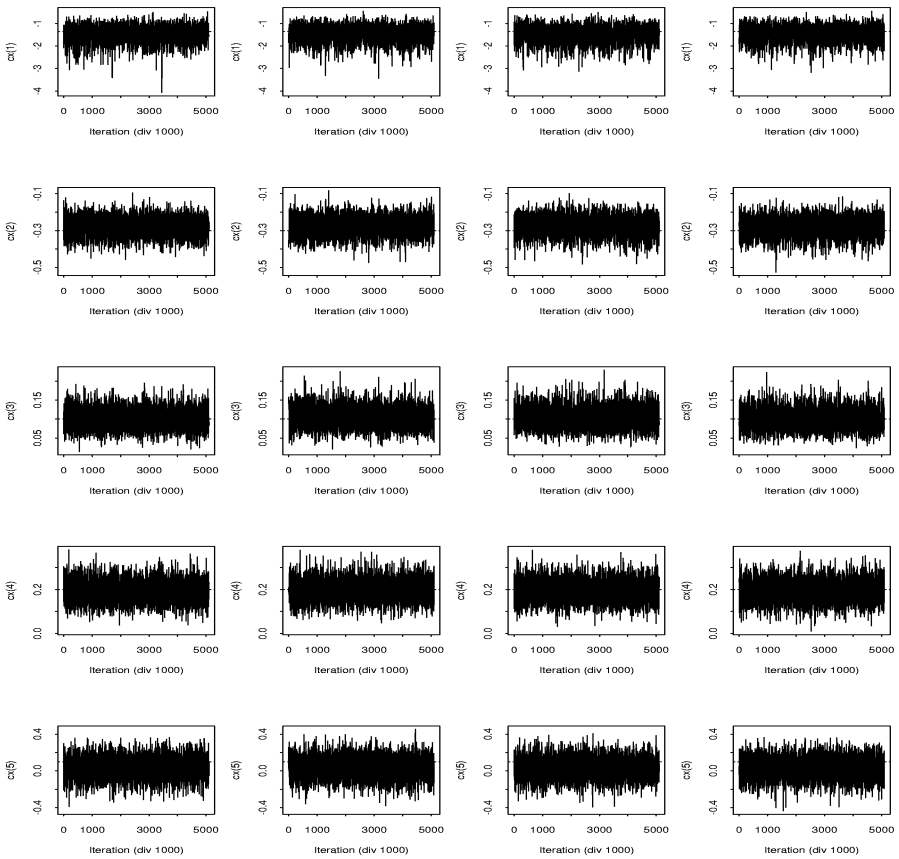

Figure 5.13: Trace plots for $c_1^{s,x}, \ldots, c_5^{s,x}$ for (from left to right) the original model (Model I) and the re-parameterisations Model II, III and V.

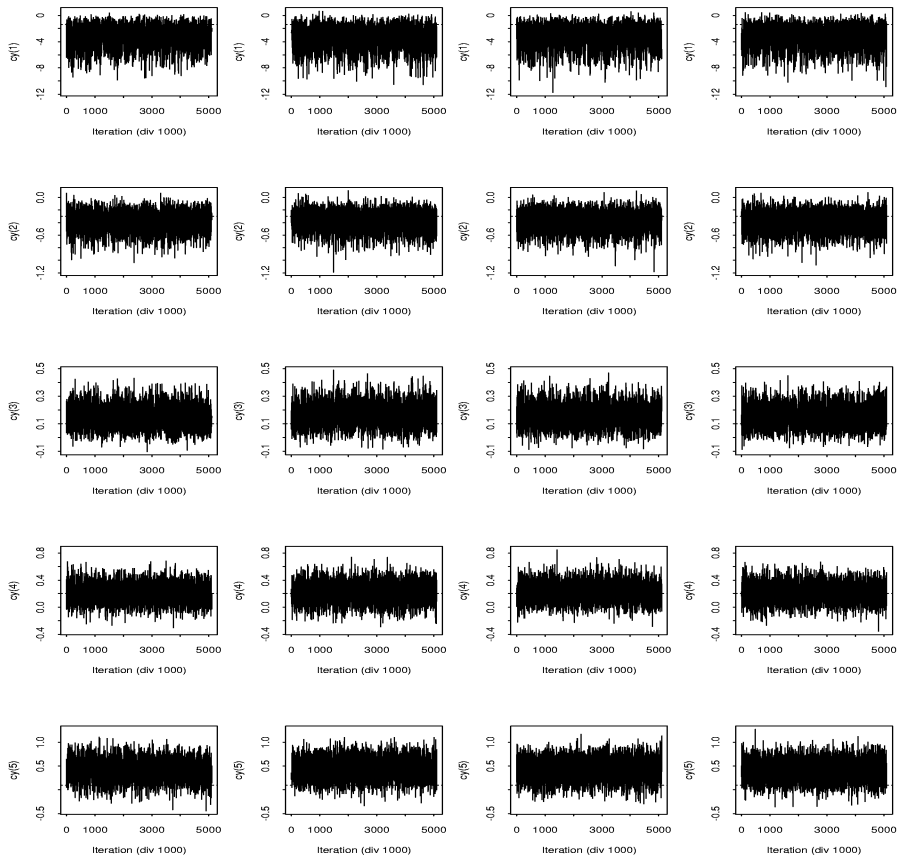

Figure 5.14: Trace plots for $c_1^{s,y}, \ldots, c_5^{s,y}$ for (from left to right) the original model (Model I) and the re-parameterisations Model II, III and V.

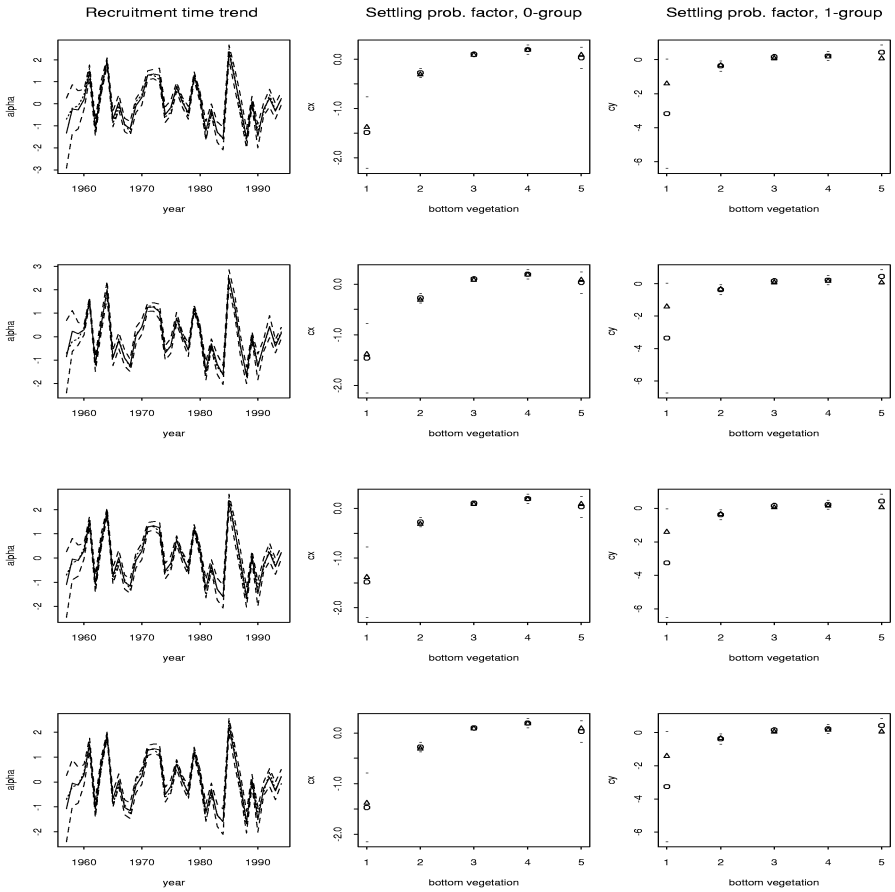

Figure 5.15: Estimated (full line) and true (dotted line) values for the temporal trend in the recruitment rate, $\{\alpha_t\}$, and estimated (circles) and true (triangles) values of the settling probability factors $\{c_k^{s,x}\}$ and $\{c_k^{s,y}\}$. The dashed lines and horizontal bars (-) indicate two times the empirical standard deviation. The plots are for (from the top downward) the original model (Model I) and the re-parameterisations Model II, III and V.

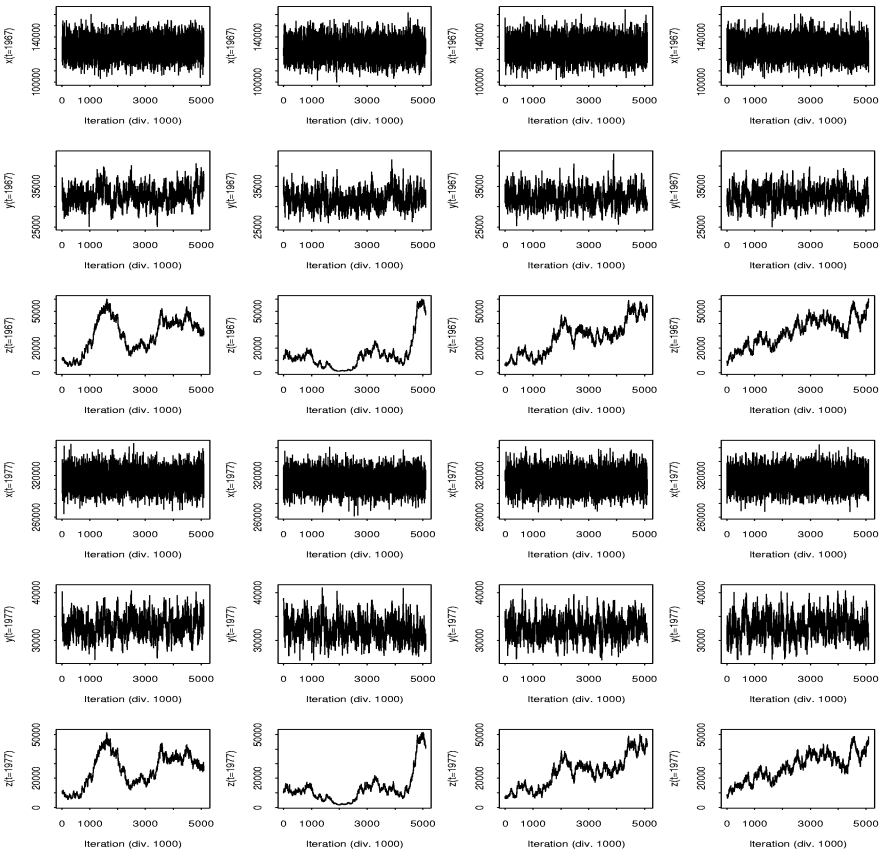

Figure 5.16: Trace plots for two sets of 0-group, 1-group and adult populations, for (from left to right) the original model (Model I) and the re-parameterisations Model II, III and V.

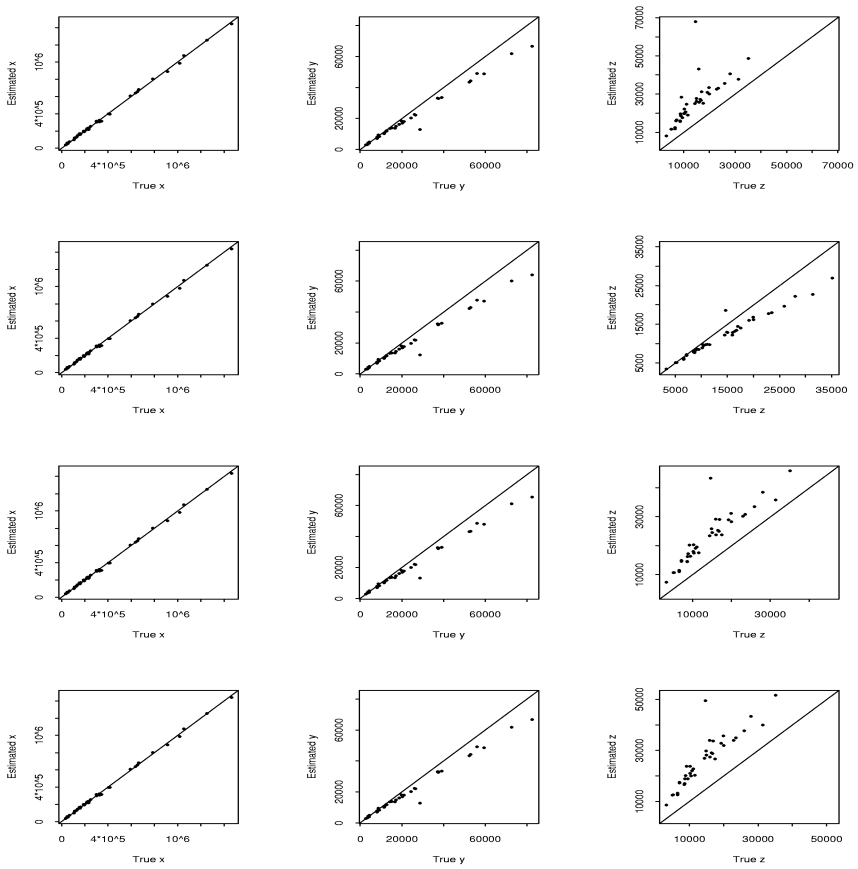

Figure 5.17: Estimated and true populations, for (from the top downward) the original model (Model I) and the re-parameterisations Model II, III and V. (The extreme values for Z are due to initialisation effects).

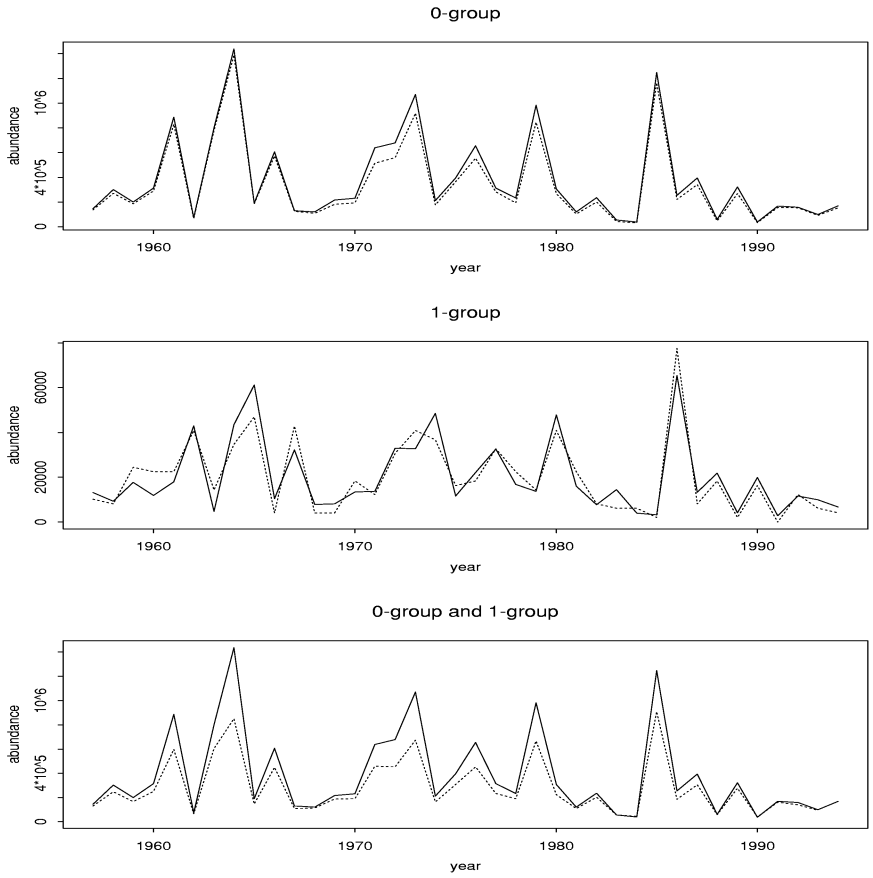

Figure 5.18: Estimated posterior means from Model III, compared to the simulated data. The full lines in the two upper panels show the posterior mean of the 0-group (top panel) and the 1-group (middle panel). The simulated data, in terms of the total number of sampled cod, up-scaled by the factors $\Omega(f)/(Ln_f\psi_x)$ and $\Omega(f)/(Ln_f\psi_y)$ respectively, are added (dotted lines). In the bottom panel, the estimates for the 0-group (full line) and 1-group (dotted line) cod are shown together, the latter being shifted one year back and up-scaled by a constant to simplify visual comparison.

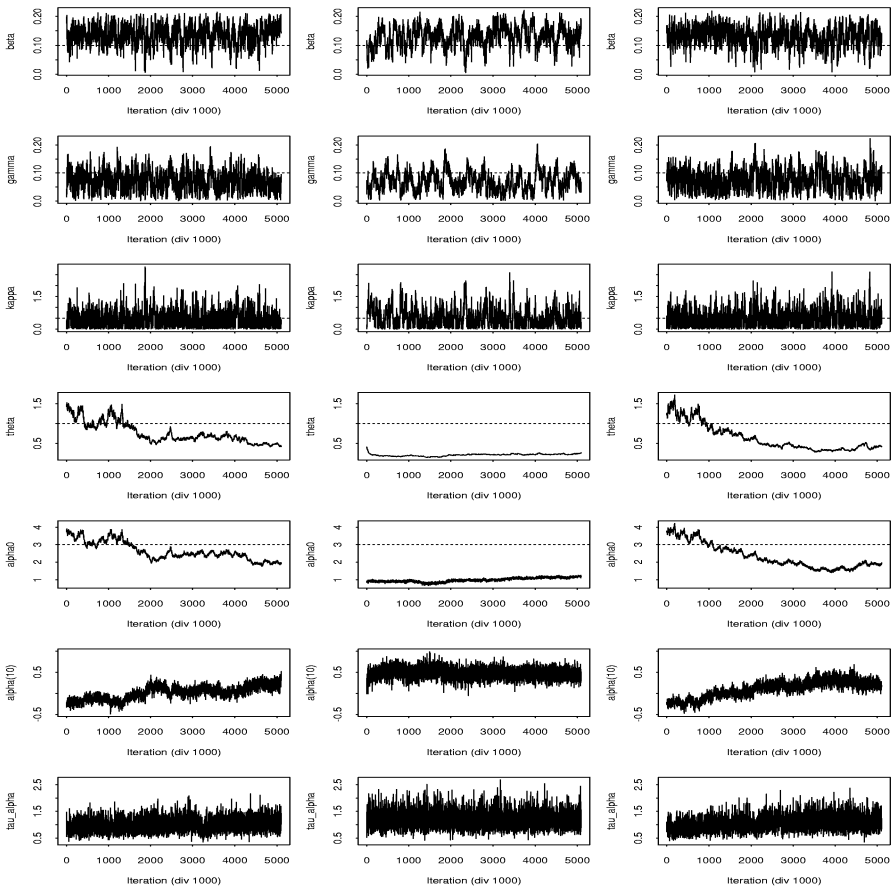

Figure 5.19: Trace plots for the parameters β, γ, κ, θ, α_0, α_{10} and τ_α for Model III, using three different sets of initial values.

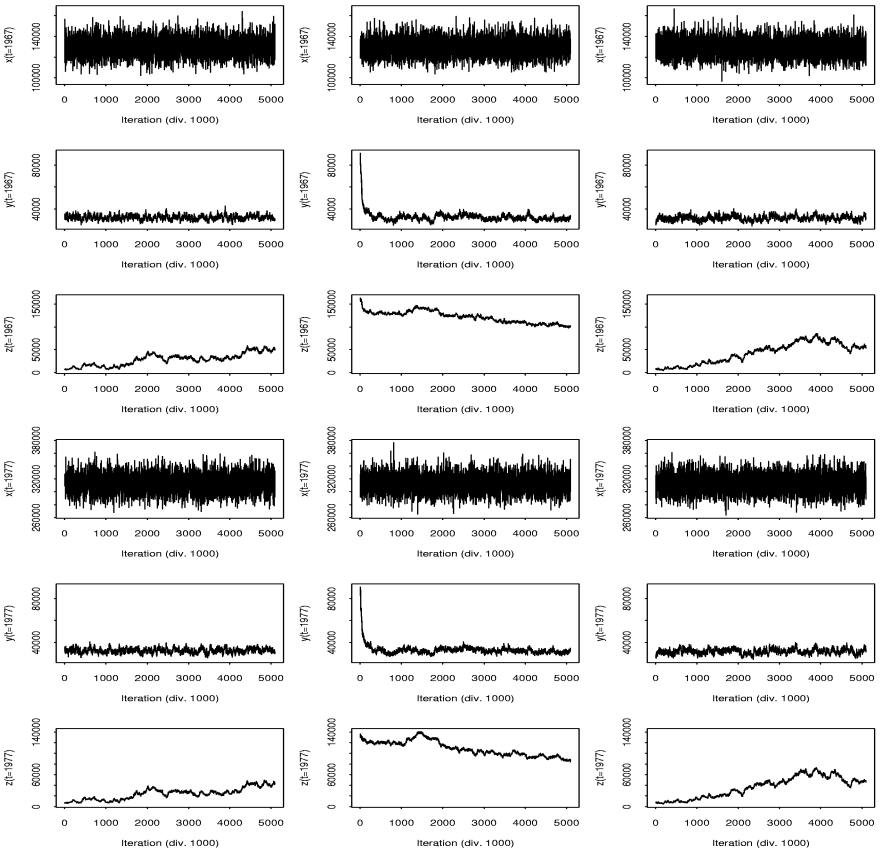

Figure 5.20: Trace plots for two sets of 0-group, 1-group and adult populations for Model III, using three different sets of initial values.

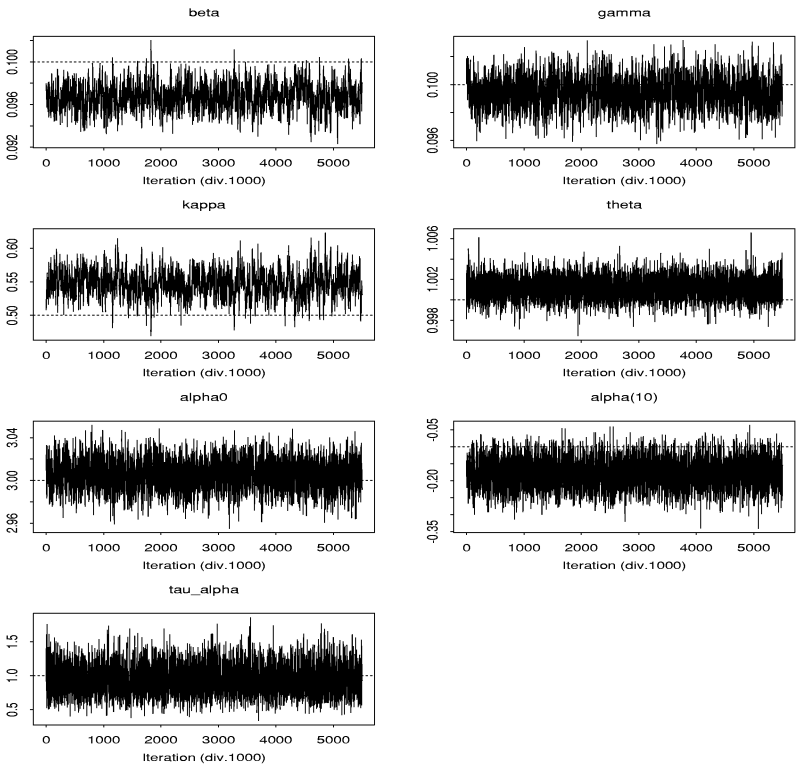

Figure 5.21: Trace plots for a selection of the model parameters, conditional on the true populations using the simulated data for one fjord. The horizontal dotted lines indicate the values used to generate the data.

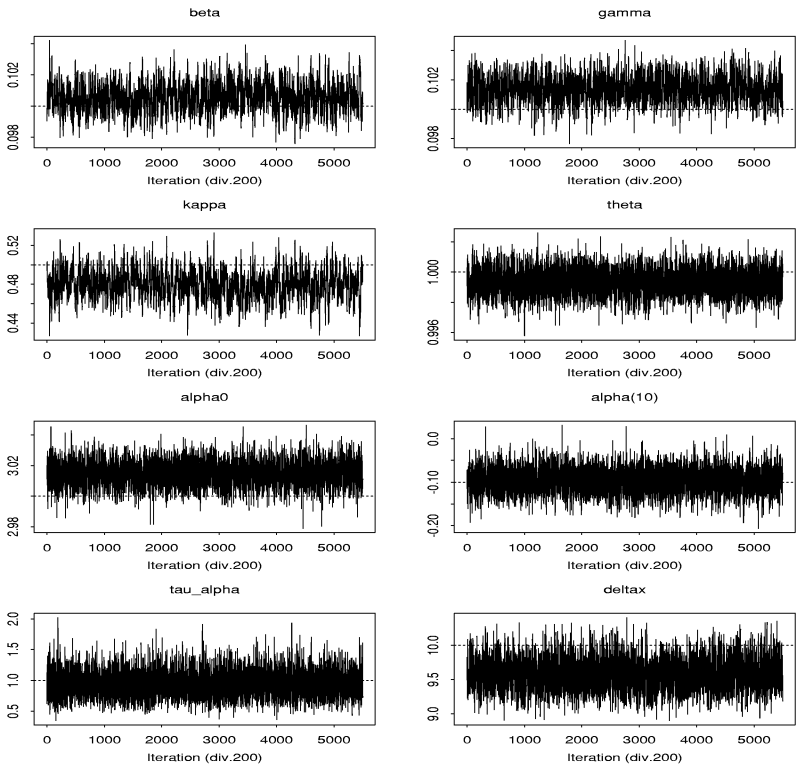

Figure 5.22: Trace plots for a selection of the model parameters, conditional on the true populations, grouping the sampling stations of the simulation study into two fjords. The horizontal dotted lines indicate the values used to generate the data.

5.4 Sensitivity analysis

So far, we have fixed the parameters ψ_x and ψ_y of the likelihood, representing the probability of each individual being in the region covered by the beach seine at the time of sampling. Because the abundances in the likelihood models (4.15) and (4.18) are also unknown, the parameters ψ_x and ψ_y are in practice difficult to identify in the current model without sufficiently informative priors, as we have experienced when attempting to estimate these parameters as well. To check the sensitivity of the parameter estimates to the choice of values for ψ_x and ψ_y, we ran the estimation procedure on the simulated data set for a set of different values of these two parameters. Also, since the convergence of the parameters θ and α_0 is slow due to high cross-correlation, we fixed θ at three different values, but estimated α_0. The sensitivity analysis was run for $\psi_x \in \{0.30, 0.35, 0.40, 0.45, 0.50\}$, $\psi_y \in \{0.050, 0.075, 0.100, 0.125, 0.150\}$ and $\theta \in \{0.8, 1.0, 1.2\}$. The results are presented graphically in Figures 5.23, 5.24, 5.25 and 5.26, and the estimates and corresponding empirical standard deviations for the model parameters are given in Tables B.1, B.2 and B.3 in Appendix B.

Since ψ_x and ψ_y represent the probability of being catchable at the time of sampling, increasing these parameters should lead to decreasing estimates of the posterior mean of the corresponding population abundances. This effect is illustrated in Figure 5.23. Here, the estimated population abundances for all pairs of values of ψ_x and ψ_y are plotted for each of the three selected values of θ. The time series of abundances can be seen to be grouped into five groups. For the 0-group cod (X), the groups correspond to the five selected values of ψ_x, with X decreasing as ψ_x increases. Equivalently, the estimates of the 1-group abundances (Y) are grouped according to the value of ψ_y and decrease with increasing ψ_y. The effect of ψ_x on Y and ψ_y on X are both small as compared to the corresponding within-age-group effect. The estimated abundances of the adult cod follow the same pattern as for the 1-group cod, in accordance with the prior model (4.7), and the effect of ψ_x on the estimated adult population abundance is small. We observe that the juvenile abundances are independent of θ, while the number of adult cod decreases when θ increases such that the corresponding survival probability $\exp(-\theta)$ decreases.

The parameters $\alpha_t;\ t = 1, \ldots, n_t$ in (4.2) are constrained to sum to zero, and they are therefore not expected to be much affected by the change in level of the 0-group and 1-group abundances induced by varying ψ_x and ψ_y. From Figure 5.24 we see that except for the first few years, the values of α_t are essentially unaffected by changing these two parameters, as well as by changes in θ. As shown in the bottom panels of Figure 5.24, the uncertainty of the parameter estimates are relatively large for the first few years, due to initialisation effects.

A similar result is obtained for the settling probability factors $c_k^{s,x};\ k = 1, \ldots, 5$ and $c_k^{s,y};\ k = 1, \ldots, 5$ of the likelihoods for the 0-group and 1-group data respectively, defined in Table 4.2. These parameters represent the relative settlement preferences of the

juvenile cod and as expected, these parameters seem not to be affected by the actual values of ψ_x, ψ_y and θ. Exceptions are $c_1^{s,x}$ and $c_1^{s,y}$, but for these parameters, the estimates are highly variable. This can be explained by the fact that the corresponding value of the bottom vegetation code (1=no vegetation) is recorded for only a few stations and years.

For the remaining population dynamics parameters governing the prior distributions the results are displayed in Figures 5.25 and 5.26. In Figure 5.25 the results are displayed as functions of ψ_x with each line representing a value of ψ_y, and in Figure 5.26 the same results are displayed changing the role of ψ_x and ψ_y. To interpret the results, recall from (4.1), (4.4) and (4.7) in Section 4.1 that the prior mean of the population abundances are given by the expressions

$$\begin{aligned} \mathrm{E}(X_{f,t} \mid \text{past}(t), Z_{f,t}, \alpha_0, \{\alpha_t\}) &= Z_{f,t} \exp(\alpha_0 + \alpha_t), \\ \mathrm{E}(Y_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa) &= X_{f,t-1} \exp(-\kappa - \beta \log(X_{f,t-1}) - \gamma \log(Y_{f,t-1})), \\ \mathrm{E}(Z_{f,t} \mid \text{past}(t), \theta) &= (Z_{f,t-1} + Y_{f,t-1}) \exp(-\theta). \end{aligned}$$

The most noticeable effect of θ is on the recruitment rate α_0 in correspondence with the high positive cross-correlation already experienced between these two parameters and consistent with the results of Bjørnstad et al. (1999a). This result is also in accordance with the prior distributions, since a decrease in the adult survival rate $\exp(-\theta)$ leads to a decrease in the size of the adult population, which needs to be compensated by an increase in the recruitment rate $\exp(\alpha_0)$ to keep the 0-group abundances at the same level. Moreover, we observe that the prior precision τ_α of the α_t's seems to decrease with increasing θ, but the absolute effect is very small.

From Figure 5.25 we observe that increasing ψ_x and thus decreasing the abundance of 0-group cod, has an expected negative effect on α_0. Also, there is in general a slight negative effect on β as well as κ over the selected range of ψ_x, but less pronounced than for α_0. A reduction in β is supported by the expression for the prior mean of $Y_{f,t}$, and we know from Section 5.3.4 that κ and β are negatively correlated. There seems to be little effect on τ_α and γ.

An increase in ψ_y has a positive effect on all parameters except for τ_α, as illustrated in Figure 5.26. The effect seems to be more consistent over the range of ψ_y than for ψ_x, but direct comparison of the size of the effect is difficult because the actual values as well as the relative step-wise increase in ψ_x and ψ_y are different. The positive effect on α_0 can be explained by the fact that increasing ψ_y and thus decreasing the abundance of the 1-group cod will result in a decrease in the adult population, following the expression for the conditional prior mean of $Z_{f,t}$. For given values of $X_{f,t}$, we see from the conditional prior mean that α_0 should increase to compensate for the decrease in the size of the adult population. The effects of ψ_y on β, γ and κ can be explained from the prior model using similar arguments.

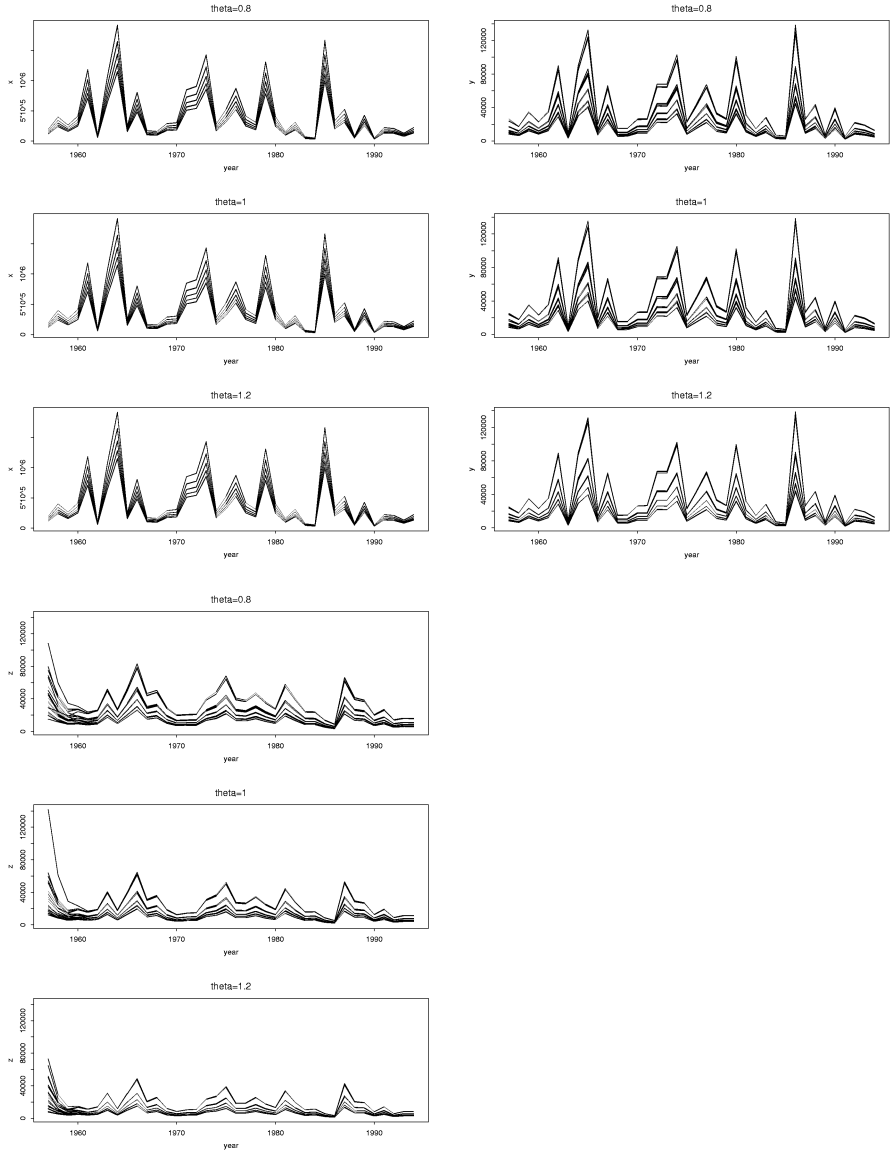

Figure 5.23: Estimated posterior means of the population abundances of the 0-group (three upper left panels),1-group (right panels) and adult (three bottom left panels) cod for different values of ψ_y and ψ_x, plotted for $\theta \in \{0.8, 1.0, 1.2\}$. See the text for further explanation.

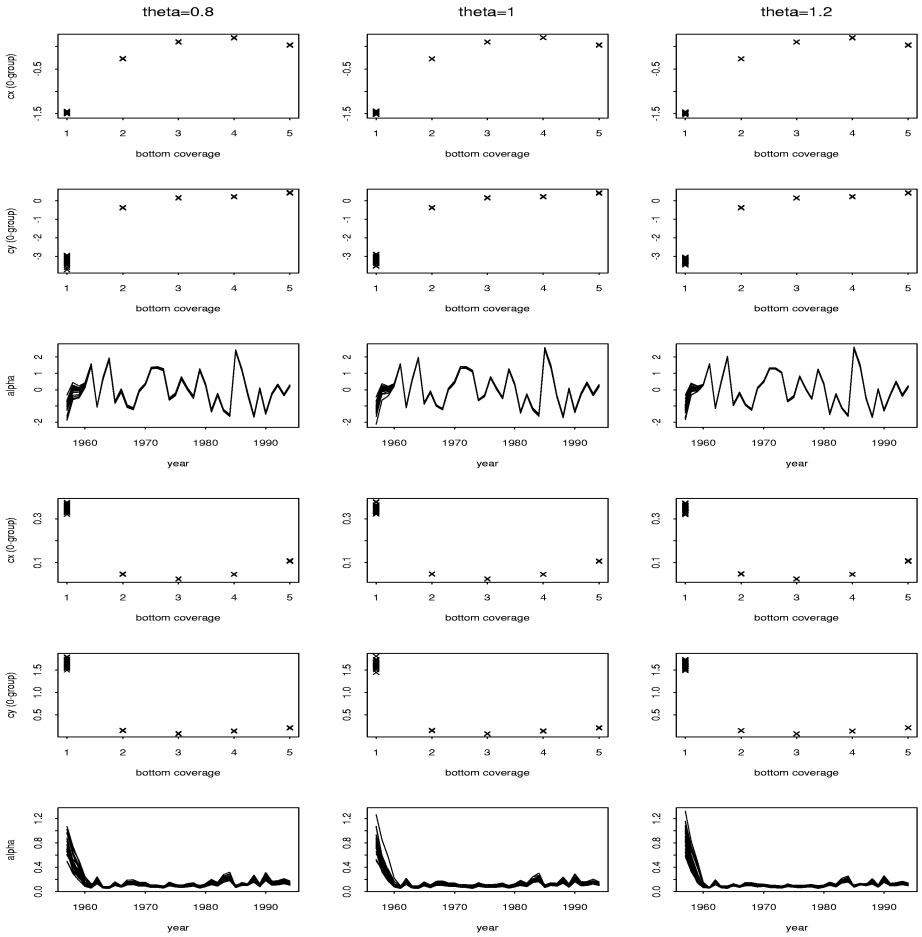

Figure 5.24: Estimated posterior means (three top rows) and empirical standard deviations (three bottom rows) of $\alpha_t;\ t = 1, \ldots, n_t$ and the factors $c_k^{s,x};\ k = 1, \ldots, 5$ and $c_k^{s,y};\ k = 1, \ldots, 5$ for the 25 different combinations of ψ_x and ψ_y and for $\theta \in \{0.8, 1.0, 1.2\}$.

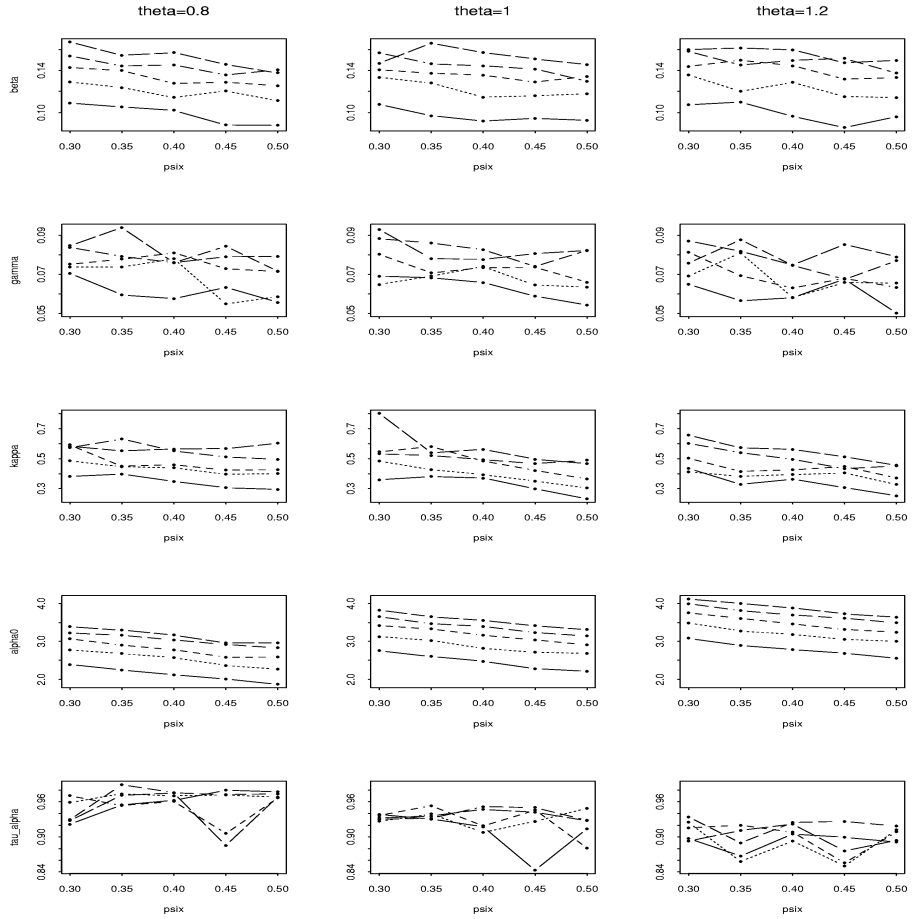

Figure 5.25: Estimated posterior means of the time-independent parameters of the prior model plotted as a function of ψ_x for the five different values of ψ_y and for $\theta \in \{0.8, 1.0, 1.2\}$ (see also Figure 5.26).

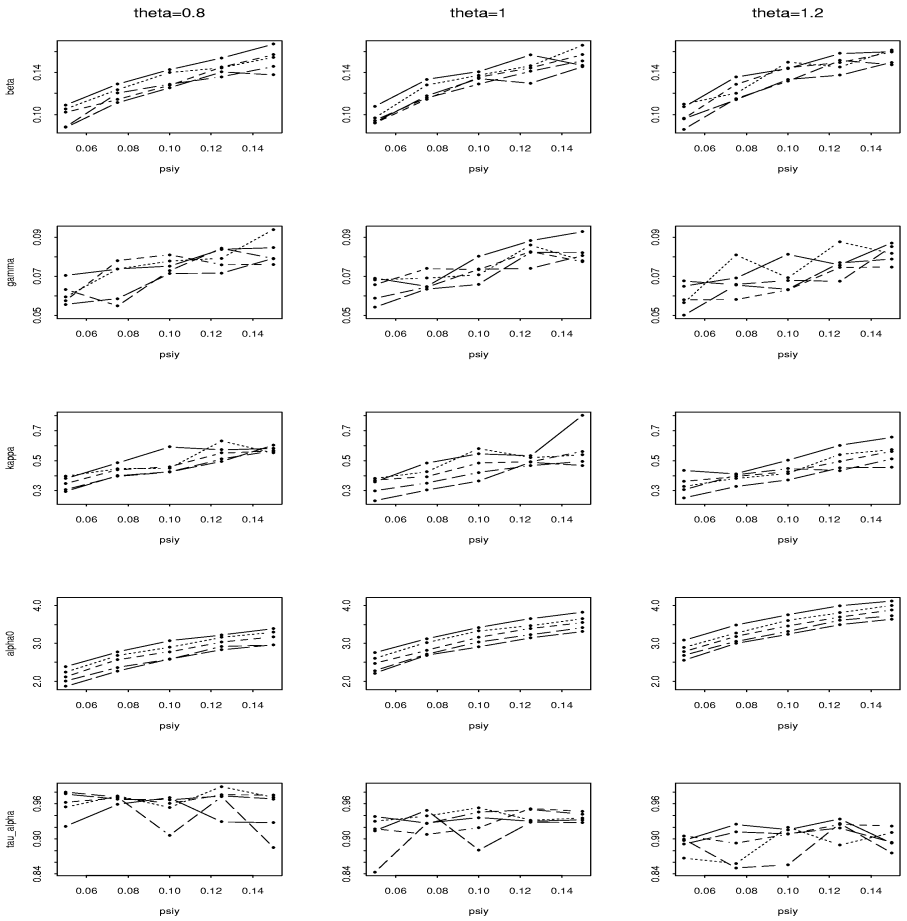

Figure 5.26: Estimated posterior means of the time-independent parameters of the prior model plotted as a function of ψ_y for the five different values of ψ_x and for $\theta \in \{0.8, 1.0, 1.2\}$ (see also Figure 5.25).

5.5 Summary and discussion

Based on a set of simulated data we have explored the properties of the statistical model and a sampling based approach to parameter estimation, focusing on the mixing and convergence properties of the MCMC algorithm. We have shown that mixing is improved by re-parameterisation, but that slow convergence remains a problem for several parameters.

The only unknowns of the likelihood are the settling probability factors $c_k^{s,x};\ k = 1, \ldots, 5$ and $c_k^{s,y};\ k = 1, \ldots, 5$. The estimated posterior means for these parameters are all close to the true values, and the mixing is good and the convergence is fast for all parameterisations.

Re-parameterising the prior model for juvenile survival by centring the log-abundances of the juvenile cod, is shown to improve the convergence and mixing of the parameters β, γ and κ. These parameters have well identified modes with empirical credibility intervals including the true values.

Although the mixing is variable, the estimated posterior means of $\alpha_t;\ t = 1, \ldots, n_t$ reproduce the true temporal structure of the log-recruitment rate well. The same is true for the temporal structure of the abundances of the 0-group and 1-group cod, both for which the mixing is satisfactory. The closest fit is obtained for the 0-group, for which the observed counts are most abundant. The estimated posterior means of the abundances are very close to the true values for the 0-group cod. A close fit to the overall levels of the abundances for the juvenile age-groups is in accordance with the fact that the sampling probabilities of the likelihood model are fixed except for the modulations due to differences in the amount of bottom vegetation. A more general approach taken by Knight et al. (2003) is to define a sufficiently informative prior on the parameters defining the level of the probability, in our model ψ_x and ψ_y. We ran a sensitivity analysis varying these parameters, and the results showed that the parameters representing relative effects are essentially unaffected by changes in the values of ψ_x and ψ_y, but that the choice of values influences the other parameters, in accordance with the prior model.

An unsolved problem is the poor mixing and slow convergence of the adult population abundances, for which we have no data, as well as the parameters α_0 and θ. These parameters were shown to be negatively correlated to the adult abundances, and correlations between population abundance and corresponding model parameters has been identified as a general problem in modelling fish population dynamics (Millar and Meyer, 2000b, and references therein). Moreover, estimating instantaneous mortality, in our model represented by θ, is generally recognised to be a challenging task in fisheries stock assessment (Gavaris and Ianelli, 2002). As illustrated by Millar and Meyer (2000b) in the case of a biomass dynamics model, identifiability might be improved by redefining the dynamics in terms of relative abundances. The results for Model V, presented in Figure 5.10 indicate that sweeping the mean of the log-abundance to the log-recruitment

rate α_0, we arrive at the well-identified quantity $\alpha_0^{re} = \alpha_0 + \overline{\log(Z)})$. Consequently, the mean number of recruits to the population can be identified, but not the relative contributions to this number from the number of adult cod and the recruitment rate.

In our model, we have included an over-dispersion parameter δ_x in the prior model (4.2) for recruitment. In general, the model can be extended to include over-dispersion in the priors for the 1-group and adult cod as well as in the Binomial likelihood models, but this should be weighted against model parsimony, in particular in the light of the convergence and identifiability problems of the current model.

Chapter 6

Results for the Flødevigen data

In this chapter we apply the sampling based estimation approach to the Flødevigen data set. We first present results for the subset of stations in the Risør area, plotted in Figure 2.8. The stations belong in reality to two fjords, eight stations in Sandnesfjord and four in Søndeledfjord, but we treat these twelve stations as belonging to one fjord. In Section 6.2 the results using data for 13 fjords are presented. The results are in both cases based on the re-parameterised model denoted Model III in Table 5.4.

6.1 Results for the Risør fjord

The results from the simulation study, running the algorithm on data from one fjord, indicated that the strong interactions between the abundance of adult cod and the parameters θ and α_0 governing adult survival and recruitment made the identifiability of these parameters difficult. We choose to fix the parameter θ at the value $\theta = 1$ which is in accordance with the results of Julliard et al. (2001). The over-dispersion parameter of the recruitment prior (4.2) is kept fixed at the value $\delta_x = 10$.

Trace plots for a selection of the parameters, based on 5.1 million iterations of the sampler, are shown in Figure 6.1. The convergence and mixing seems to be reasonably good for all parameters, and the mixing properties are similar for the parameters not shown. The posterior marginal means and empirical standard deviations for the global parameters, computed from every 1000th iteration after discarding the first 200000 iterations, are given in Table 6.1. We observe that there seems to be evidence of a positive within-group density-dependency ($\beta > 0$), while the significance of the parameter γ is more doubtful. The density-independent component of the juvenile survival rate, represented by κ, has a large variance. Further evidence of the significance or lack of significance of γ and κ can be obtained by comparing the fit of models including or

Parameter	Posterior mean	Empirical std.dev.
β	0.026	0.009
γ	0.014	0.009
κ	0.079	0.074
α_0	0.84	0.08
τ_α	1.05	0.28

Table 6.1: Results for the main population dynamics parameters from running the Metropolis-Hastings sampler on the Risør data.

not including these two parameters. The fit of a set of Bayesian hierarchical models differing by the definition of the parameter vector $\boldsymbol{\theta}$ of the model can be compared by using the deviance information criterion DIC (Spiegelhalter, Best, Carlin and van der Linde, 2002). This is defined as a function of the estimated posterior mean $\bar{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ by

$$\text{DIC} = D(\bar{\boldsymbol{\theta}}) + 2p_D = \overline{D(\boldsymbol{\theta})} + p_D, \tag{6.1}$$

where $D(\boldsymbol{\theta})$ is the Bayesian deviance

$$D(\boldsymbol{\theta}) = -2\log\{f(\boldsymbol{y} \mid \boldsymbol{\theta})\} + 2\log(f(\boldsymbol{y})) \tag{6.2}$$

and p_D is a measure of the effective number of parameters. The functions $f(\boldsymbol{y} \mid \boldsymbol{\theta})$ and $f(\boldsymbol{y})$ are the likelihood and a standardising functions depending only on the data, respectively. Results from using this criterion on the likelihood as defined by our model indicate that the model discarding γ but keeping κ gives the best fit to the Risør data (Gunnhildur Högnadóttir Steinbakk, Norwegian Computing Center, personal communication).

Plots of the estimated posterior means of $\{\alpha_t\}_t$, the temporal fluctuation of the log-recruitment rate, and the settling probability factors $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ are given in Figure 6.2. For both juvenile age-groups, there seems to be no structure in the dependency of the settling probability on the bottom coverage. For bottom coverage equal to 1, the variance of the estimated parameter is large due to a small number of observations within this category. From the plots of estimated posterior distributions and the corresponding prior distributions for the parameters, shown in Figure 6.3, we observe that the posterior distributions are more concentrated than the priors, confirming that there is information on the parameters in the data.

Trace plots for the abundances, a selection of which are shown in Figure 6.4, indicate that conditionally on θ and δ_x the mixing is good and convergence is achieved for all age-groups. The estimated posterior mean of the abundances of the 0-group and 1-group cod, discarding the first 200000 iterations, are given in the right panels of Figure 6.2. To get an impression of the degree of fit to the data, we have added the observed counts up-scaled by the factors $\Omega(f)/(Ln_f\psi_x)$ and $\Omega(f)/(Ln_f\psi_y)$ for the 0-group and 1-group

respectively, in analogy to Figure 5.18 for the simulated data. The up-scaling represents a guess of an overall mean abundance level based on a homogeneous fjord and assuming no habitat preferences among the individuals of the juvenile age-group. The estimates give a very close fit to the temporal structure of the 0-group cod and a reasonable but not so good fit to the 1-group data. From the bottom right panel of Figure 6.2, where we have plotted 1-group estimates shifted one year back together with the 0-group estimates, we see that the temporal structure of the 0-group and 1-group estimates are essentially identical. This result indicates that the prior acts more strongly than the likelihood on the estimated abundances for the 1-group cod. This was also seen for the simulated data, that were generated to contain a similar number of zero counts as in the 1-group data in the Risør area.

The parameters ψ_x and ψ_y of the sampling process are fixed throughout the study, and the chosen values will influence the level of the abundance estimates. To assess the sensitivity of the estimated parameters on ψ_x, ψ_y and θ, a sensitivity analysis was conducted in a similar way as for the simulated data set, running the algorithm for 1.1 million iterations for each set of values. A summary of the main results is shown in Figures 6.5 and 6.6. The results are similar to those reported for the simulated data set in Section 5.4, but the effects of varying ψ_x and ψ_y on the parameters β and γ are more profound, in particular the effects of ψ_y. Estimated posterior means and corresponding empirical standard deviations from the sensitivity analysis are given in Tables B.4, B.5 and B.6 in Appendix B.

6.2 Results using data from several fjords

Applying the estimation algorithm to the simulation study and to the Risør data, our full model appeared to be near non-identifiable. In both cases, the model was specified for one fjord only. Using data from several fjords, the overall level of mixing and identifiability is expected to be improved compared to the results from the single fjord cases. We ran the algorithm based on Model III using data from a total of 49 stations in 13 fjords, observed in the period 1956-1998. Plots of the 49 time series of counts for the 0-group and 1-group cod are given in Appendix A. In Figures 6.7, 6.8 and 6.9 trace plots resulting from running 5.1 million iterations of the sampling algorithm are shown. The mixing is still not satisfactory for θ and δ_x, but their modes seems to be reasonably well identified. For the other parameters, the mixing properties are similar to the ones observed for the Risør data, fixing δ_x and θ. The convergence for the abundances is now considered to be reasonably good for all age-groups, although the mixing is slightly poorer for the adult cod than for the juvenile age-groups. The parameters θ and δ_x are highly correlated as well as dependent on the adult population abundance, as illustrated by the scatter plots of Figures 6.10 and Figure 6.11. The latter plot is given for one fjord (fjord no. 8), but to different extents the pattern is similar for the other fjords. The

parameters α_0 and θ seem to behave independently when estimating all parameters of the model.

Recall from Section 4.1.1 that we introduced the re-parameterisation $\delta_x = \delta'_x + \alpha_0$, where $\exp(\delta'_x)$ is the over-dispersion of the Gaussian prior of the recruitment process. Studying the scatter plots of Figure 6.10, we observe that the correlation between α_0 and δ_x is much smaller than between α_0 and δ'_x, a fact that supports this re-parameterisation.

Estimated posterior means and empirical standard deviations for the global parameters are given in Table 6.2, and the estimates of $\alpha_t;\ t = 1, \ldots, 43$ and the settling probability factors $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ are given in Figure 6.12. As for the Risør data, the within-group density-dependence is clearly significant ($\beta > 0$), while the estimate of the parameter γ provides no evidence of significant between-group density dependency. The estimated value of the overall mean log-recruitment rate α_0 is relatively small, and the variance is large. Studying the values of the α_t's plotted in Figure 6.12, we observe that the recruitment rate of the year 1988 stands out as particularly low. This coincides with the year of an algae bloom along the Skagerrak coast, believed to have had a profound short term influence on the marine life in the affected areas (Smith et al., 2003; Johannessen and Sollie, 1994). The low values for the first few years are due to initialisation effects. In contrast to what was the case for the Risør data, there now seems to be a structure in the settling probability factors for the 0-group cod. For the 1-group cod the results indicate an increased sampling probability for a totally covered bottom, but there is no similar overall pattern as for the 0-group cod. For these parameters, the mixing is good and the convergence is fast. To check the significance of the differences, we study the estimated posterior means and empirical 95% credibility intervals of the differences $c_k^{s,x} - c_l^{s,x};\ k > l$ and corresponding values for the 1-group cod. The results are given in Figure 6.13, and confirm that the differences for the 0-group cod are significant for all but the differences $c_5^{s,x} - c_3^{s,x}$, $c_5^{s,x} - c_4^{s,x}$ and $c_2^{s,x} - c_1^{s,x}$. Thus, we conclude that the probability of settling seems to increase with increasing amount of bottom vegetation, but with a slight decrease for a totally covered bottom. The indications of the presence of habitat preferences related to bottom vegetation is consistent with results from other studies (e.g. Fromentin et al., 2001; Fromentin et al., 1998; Johannessen and Sollie, 1994).

From the right panels of Figure 6.12, where we have plotted the estimated posterior means of the 0-group and 1-group cod for one of the two fjords in the Risør area, we observe that the general appearance of the estimated times series is as for the Risør data. The fit to the temporal structure of the data is better for the 0-group than the 1-group data, and the similarity between the temporal structure for the two age-groups is apparent from the bottom left panel. Estimated abundances for all age groups and all fjords are shown in Figures C.1 to C.3 in Appendix C.

Plots of the estimated posterior marginal distributions of some parameters are given in Figure 6.14, and plots of the prior distributions are added. As for the Risør data, the results confirm that there is information in the data, since the posteriors are substantially

more concentrated than the priors.

Parameter	Posterior mean	Empirical std.dev.
β	0.059	0.003
γ	0.0014	0.0010
κ	0.023	0.022
θ	2.1	0.1
α_0	0.54	0.31
τ_α	0.29	0.07
δ_x	13.6	0.2
$\delta_x - \alpha_0$	13.1	0.6

Table 6.2: Results for the main population dynamics parameters from running the Metropolis-Hastings sampler on the data from all fjords.

6.3 Summary and discussion

We have estimated the parameters of our Bayesian hierarchical model for two subsets of the Flødevigen data; the Risør data for which only one set of abundances was estimated, and a data set consisting of observed counts from 13 fjords.

In contrast to the model-wise related studies in Stenseth et al. (1999), using the 0-group data only, and Bjørnstad et al. (1999a) using aggregated 0-group and 1-group data, we use the data for both age groups at their original level of aggregation. In this way, the uncertainty inherent in the original data is taken into account, and we have been able to get an impression of the relative degree to which the two sets of data can inform on the parameters of the model. The 0-group data are the most abundant, and was prior to the study believed to be the subset of the data containing the most information on the parameters of interest. Our results, e.g. as represented by the plots of abundances and up-scaled observed counts in Figures 6.2 and 6.12, confirm the assumption. The fact that the prior model seems to dominate the likelihood for the 1-group data can also be a result of an unsatisfactorily specified likelihood model. The variance of the likelihood of the 1-group cod might be increased by including an over-dispersion parameter to give a more realistic representation of the sampling uncertainty.

Comparing Figure 5.18 for the simulated data set and Figures 6.2 and 6.12 for the Flødevigen data, we observe that the magnitude of the estimated 1-group abundance relative to the 0-group abundance is much larger for the real data than for the simulated data set. This indicates that the *relative* values of the sampling probabilities of the 0-group and 1-group cod, as specified by the parameters in Section 5.1 might not reflect the true relations to a satisfactory extent.

The parameters $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ of the likelihood were well identified and the mixing and convergence good for both data sets. The temporal structure of log-recruitment rate, represented by the parameters $\{\alpha_t\}_t$, were reproduced well using both data sets, but the mixing and convergence were improved using data from 13 fjords as compared to the results for the Risør data. The resulting estimated posterior means of $\{c_k^{s,x}\}_k$ and $\{c_k^{s,y}\}_k$ indicated that the settling probability of the 0-group cod increases with the degree of bottom coverage, but with a slight decrease for a totally covered bottom. This latter fact might be due to a less efficient beach seine catch. For the 1-group cod, no significant effect on the sampling probabilities of differences in bottom coverage was found. For the α_t's, the estimated value corresponding to the year 1988 was particularly low, coinciding with the algae bloom at the Skagerrak coast that year.

The mixing and convergence of the parameters β, γ and θ of the prior for juvenile survival turned out to be satisfactory. For the Risør data, we fixed the parameters δ_x and θ to obtain satisfactory mixing and convergence. Running the algorithm on the data set including observations from several fjords, the convergence properties were improved compared to the results from estimating the abundances for one fjord only. However, the mixing was still found to be relatively poor for θ and slightly better but still not satisfactory for δ_x, and these parameters are highly correlated as well as correlated to the adult abundance. These parameters of the model should therefore be interpreted jointly and with care.

Comparing the estimated parameters to the results of other studies, one should take into account that not all parameters of the population dynamics are scale-invariant. Multiplying the population abundances for the three age-groups by a common factor c, the dynamics will change, and thus parameters obtained by models defining the population dynamics on the scale corresponding to the sampled data, will in general not be directly comparable to estimates from a fjord level model as ours. To illustrate, let

$$\begin{aligned} X'_{f,t} &= X_{f,t}/c, \\ Y'_{f,t} &= Y_{f,t}/c, \\ Z'_{f,t} &= Z_{f,t}/c \end{aligned}$$

where c is a positive constant. Substituting for $X_{f,t}$, $Y_{f,t}$ and $Z_{f,t}$ in the prior models (4.2), (4.5) and (4.7), we get

$$cX'_{f,t} \mid \text{past}(t), Z'_t, \alpha_0, \{\alpha_t\}, \delta_x \sim N(cZ'_{f,t}\exp(\alpha_0+\alpha_t), cZ'_{f,t}\exp(\delta_x+\alpha_t)), \quad (6.3)$$

$$cY'_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa \sim Bin(cX'_{f,t-1}, \exp(-\kappa)\ (cX'_{f,t-1})^{-\beta}(cY'_{f,t-1})^{-\gamma}), \quad (6.4)$$

$$cZ'_{f,t} \mid \text{past}(t), \theta \sim Bin(c(Z'_{f,t-1}+Y'_{f,t-1}), \exp(-\theta)). \quad (6.5)$$

Rewriting the conditional prior distribution of $cY'_{f,t}$ as

$$cY'_{f,t} \mid \text{past}(t), \beta, \gamma, \kappa \sim Bin(cX'_{f,t-1}, \exp(-\kappa - c\log(\beta) - c\log(\gamma) - \beta\log(X'_{f,t-1}) - \gamma\log(Y'_{f,t-1}))), \quad (6.6)$$

we see that the parameters β and κ of the survival rate are scale-invariant, while the intercept term $\exp(-\kappa - c\log(\beta) - c\log(\gamma))$ will depend on the scaling factor. Although the parameters α_0 and θ will not be scale-invariant in the conditional prior variances of $X'_{f,t}$ and $Z'_{f,t}$, they are scale-invariant in the conditional prior means. The latter fact indicates that using the results from other studies to fix the parameter θ seems reasonable.

As was shown by the sensitivity analysis varying ψ_x and ψ_y, the global parameters all depend to different extents on the down-scaling from fjord level to station level inherent in the sampling probabilities of the likelihood. Thus, even though β and γ are scale invariant in the prior, differences in models for the sampling process will cause direct comparison of absolute values of these parameters to be difficult. However, the presence or absence of within- or between group density-dependency applying different models can still be compared by studying the significance of β and γ.

Keeping these comments in mind we can make rough comparisons of the estimated posterior means to estimated values obtained from other studies on cod along the Norwegian Skagerrak coast. Our estimated values of θ and α_0, using data from 13 fjords, are $\widehat{\theta} = 2.1$ and $\widehat{\alpha_0} = 0.54$, corresponding an adult survival rate of $\exp(-2.1) = 0.12$ and a recruitment rate of $\exp(0.54) = 1.71$. The estimated value of θ is relatively large and of α_0 relatively small compared to previous results. For example, θ has previously been estimated to 1.0 (Julliard et al., 2001), using a different set of data, and 0.56 (Chan et al., 2003b), both smaller than our estimate of 2.1. Chan et al. (2003b) distinguish between natural mortality and fishing mortality, such that the estimate $\exp(-0.56) = 0.57$ is interpreted as a baseline adult survival rate, with temporal fluctuations around this level due to fishing mortality. In the study of Bjørnstad et al. (1999a), four different estimates of α_0 in the range 1.93 to 2.73 were found, which are all larger than our estimate of 0.54.

For both sets of data we estimated a significant within-group density-dependent effect on the juvenile survival rate. This is in accordance with other studies (Fromentin et al., 2001; Bjørnstad et al., 1999a; Bjørnstad et al., 1999b; Stenseth et al., 1999). The between-group density dependency is smaller, and not significant in our study. In the light of the discussion above regarding the information inherent in the 1-group data, this is not surprising. Fromentin et al. (2001) found some evidence of between-group density-dependence in north-eastern fjords, estimating a significant effect for a few fjords. Based on aggregated data, Bjørnstad et al. (1999a) estimated a significant density-dependency between groups, but smaller than the within-group dependency. Stenseth et al. (1999) also found evidence of between-group density dependencies, but they did not take the stochasticity of the population dynamics into account. Taken together, the results of the different analyses based on the Flødevigen data are not conclusive when it comes to the between-group density-dependency.

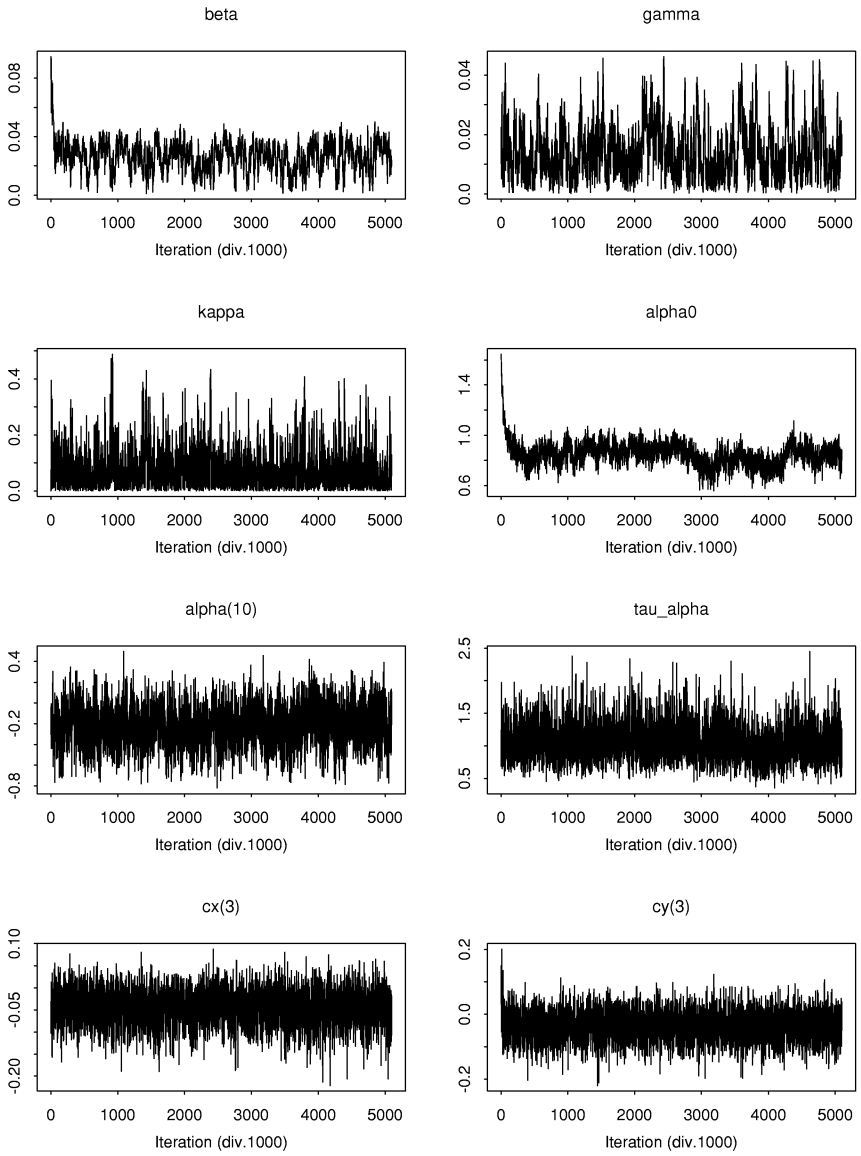

Figure 6.1: Trace plots for the Risør data, using Model III. The plots show every 1000th of a total of 5.1 million iterations after tuning. The parameters denoted by $cx(3)$ and $cy(3)$ are the settling probability factors $c_3^{s,x}$ and $c_3^{s,y}$ respectively.

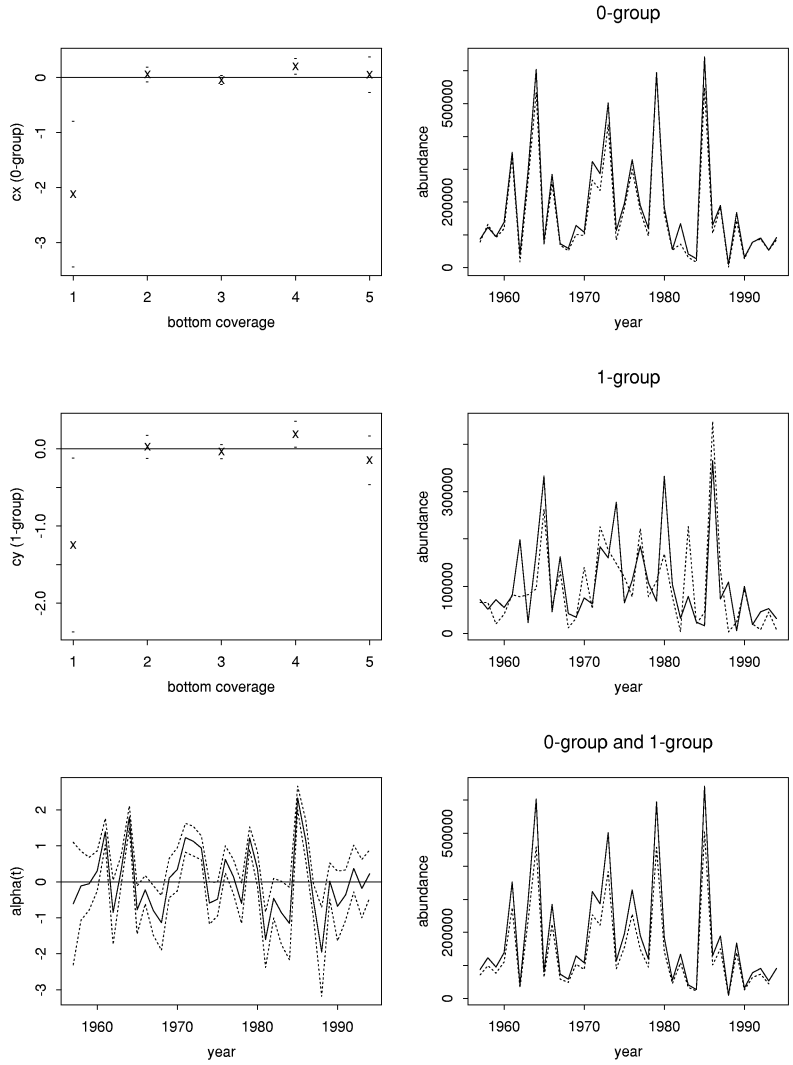

Figure 6.2: Estimated posterior means from Model III based on the Risør data. The three left panels show the settling probability factors $\{c_k^{s,x}\}$ (top) and $\{c_k^{s,y}\}$ (middle) and the temporal structure of the log recruitment rate, $\{\alpha_t\}_t$, including two times the empirical standard deviation of the updates. The three right panels show estimated abundances compared to the real data. The full lines in the two upper panels show the posterior mean of the 0-group (top) and the 1-group (middle). The dotted lines represent the total observed catch of cod, up-scaled by the factors $\Omega(f)/(Ln_f\psi_x)$ and $\Omega(f)/(Ln_f\psi_y)$, respectively. The bottom right panel shows the estimates for the 0-group (full line) and 1-group (dotted line) shifted one year back and up-scaled by a constant. The estimates are means taken over every 1000th iteration.

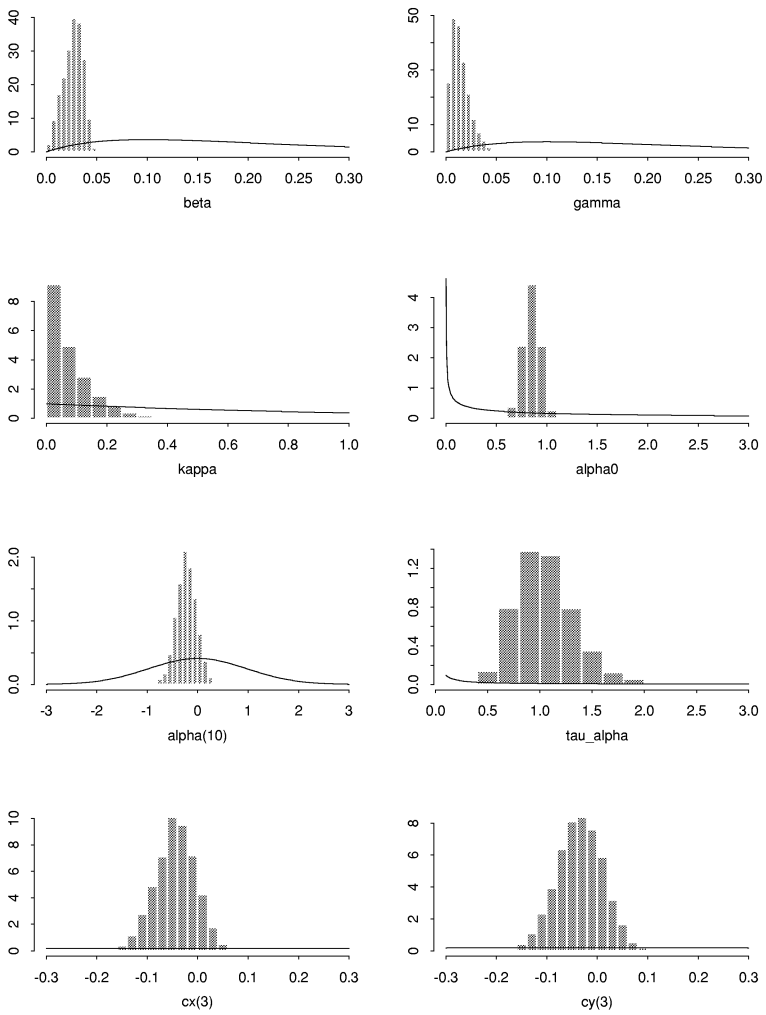

Figure 6.3: Histogram of the estimated marginal posterior distributions for the main population dynamics parameters based on the Metropolis-Hastings sampler, using the Risør data. The prior distributions are added as full lines.

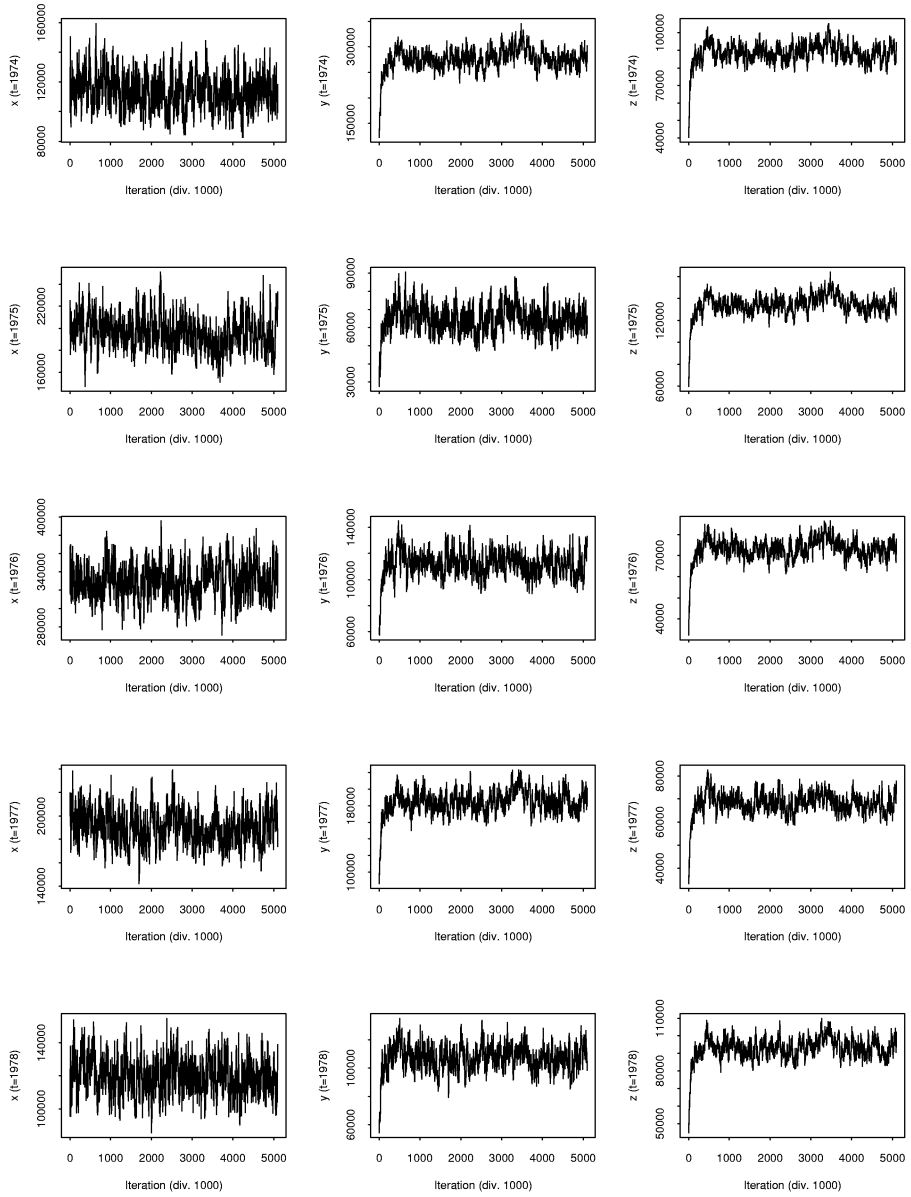

Figure 6.4: Trace plots for the abundances of 0-group (left panels), 1-group (middle panels) and adult (right panels) cod for the years 1974-1978 for the Risør data, using Model III. Every 1000th of 5.1 million iterations after tuning is shown.

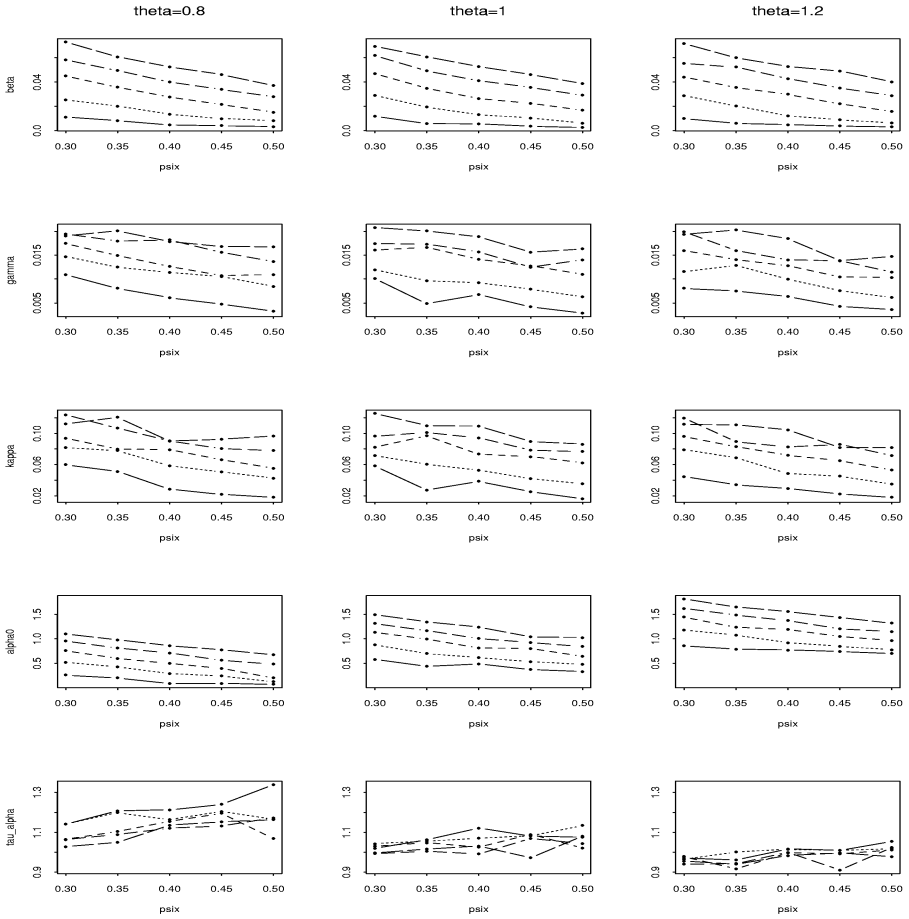

Figure 6.5: Estimated posterior means of the time-independent parameters of the prior model plotted as a function of ψ_x for the five different values of ψ_y and for $\theta \in \{0.8, 1.0, 1.2\}$ using the Risør data (see also Figure 6.6).

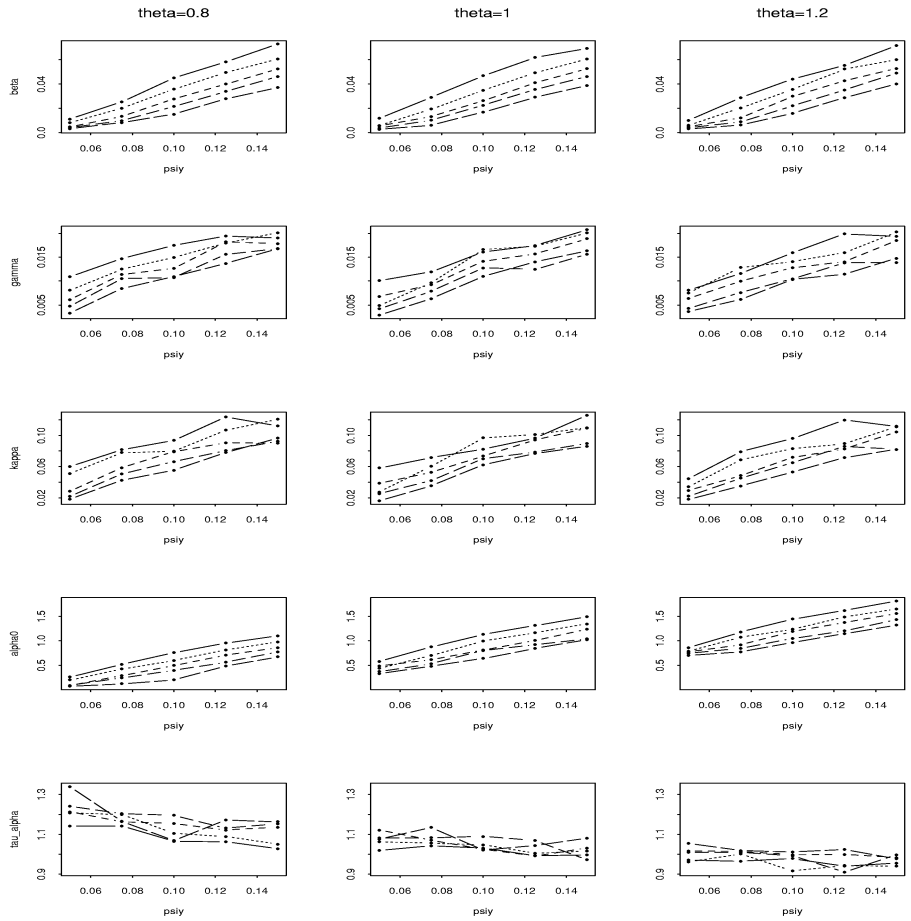

Figure 6.6: Estimated posterior means of the time-independent parameters of the prior model plotted as a function of ψ_y for the five different values of ψ_x and for $\theta \in \{0.8, 1.0, 1.2\}$, using the Risør data (see also Figure 6.5).

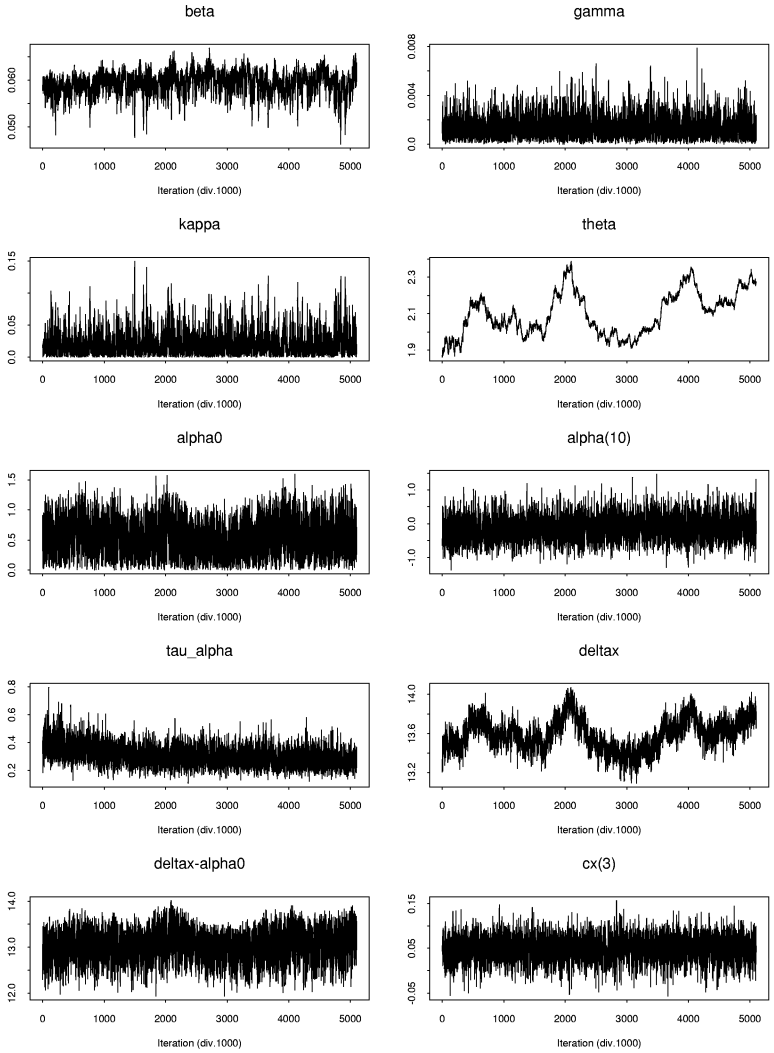

Figure 6.7: Trace plots for a selection of the model parameters based on data from all 13 fjords. Every 1000th of 5.1 million iterations is shown.

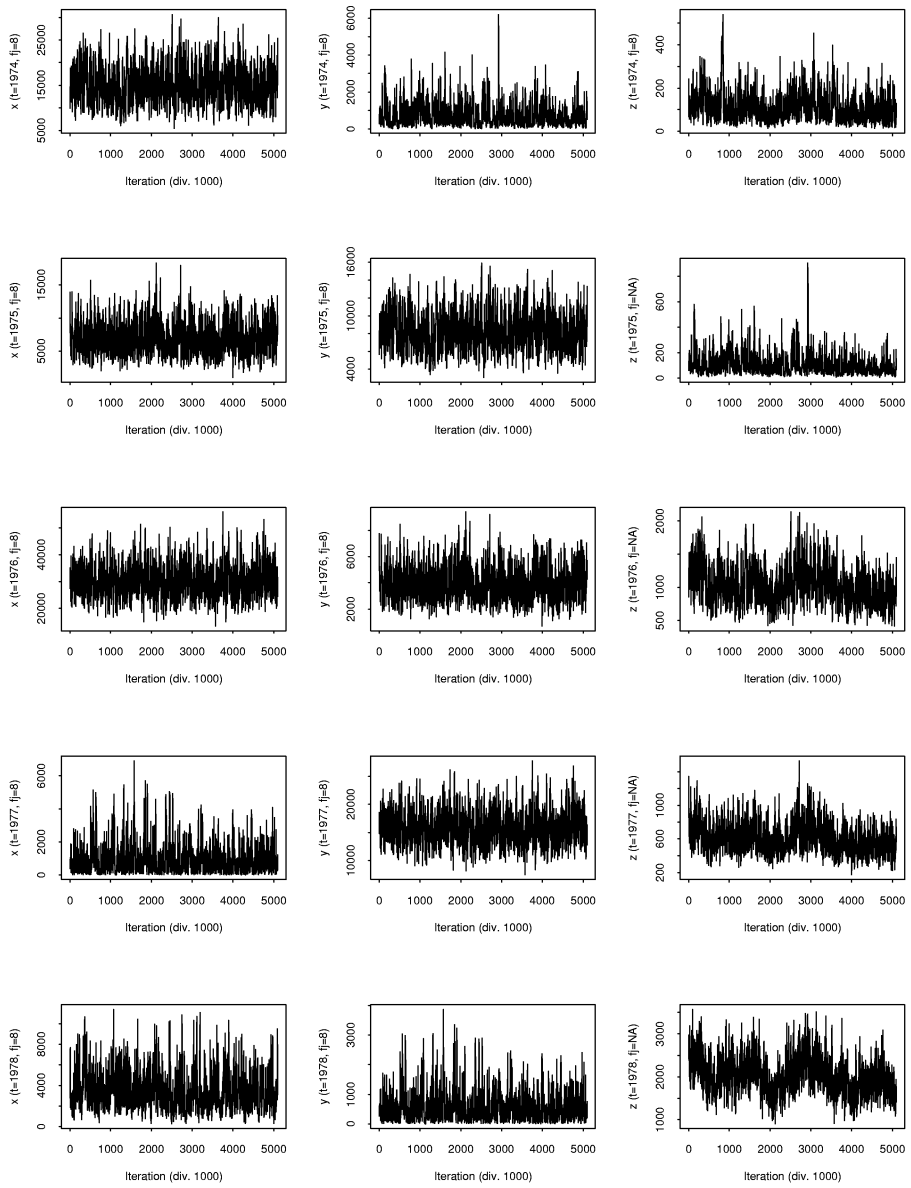

Figure 6.8: Trace plots for the abundances of 0-group (left panels), 1-group (middle panels) and adult (right panels) cod for fjord 8 (one of the Risør fjords) for the years 1974-1978, using Model III. Every 1000th of 5.1 million iterations is shown.

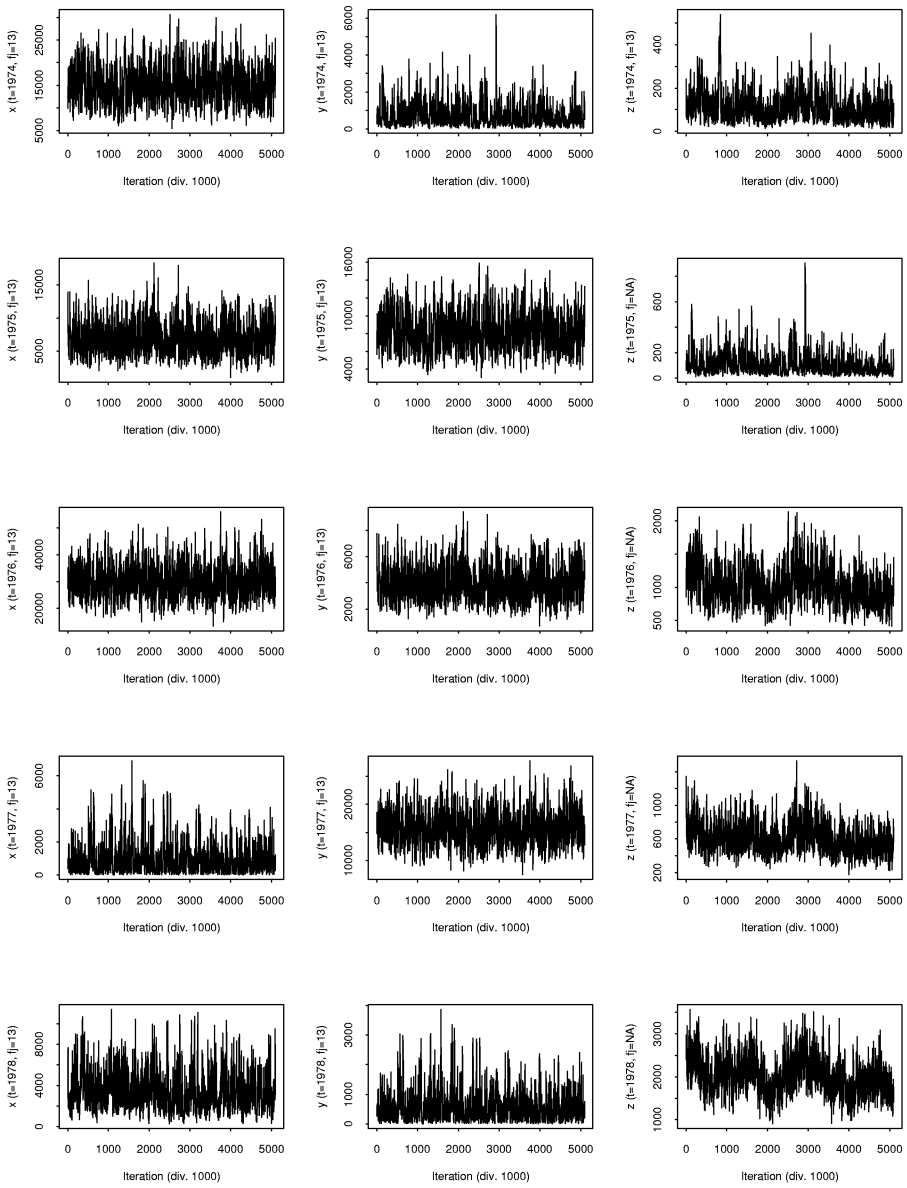

Figure 6.9: Trace plots for the abundances of 0-group (left panels), 1-group (middle panels) and adult (right panels) cod for fjord 13 for the years 1974-1978, using Model III.

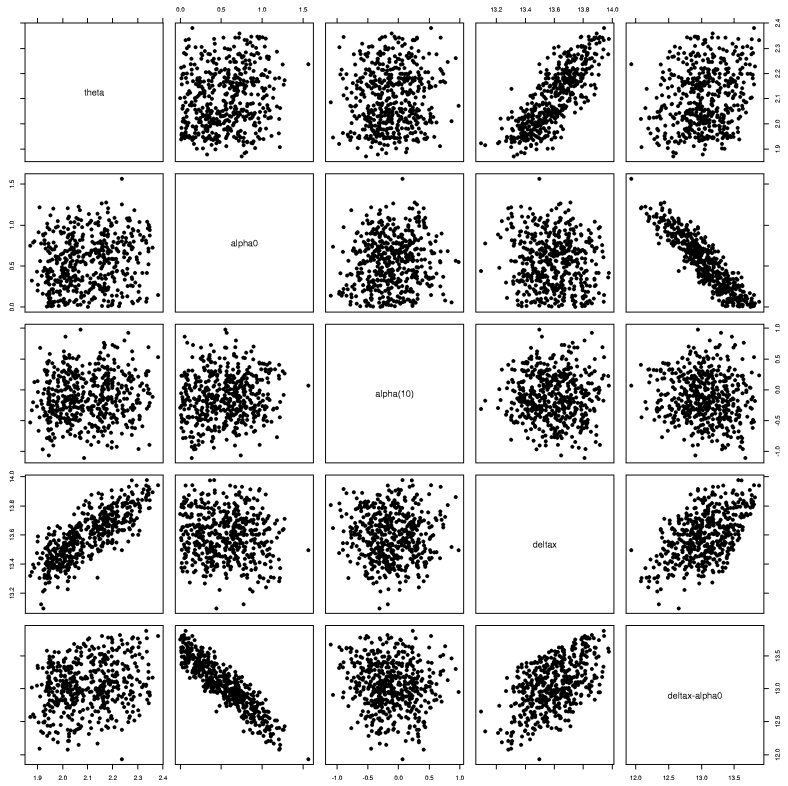

Figure 6.10: Scatter plots for some of the model parameters using data from 13 fjords.

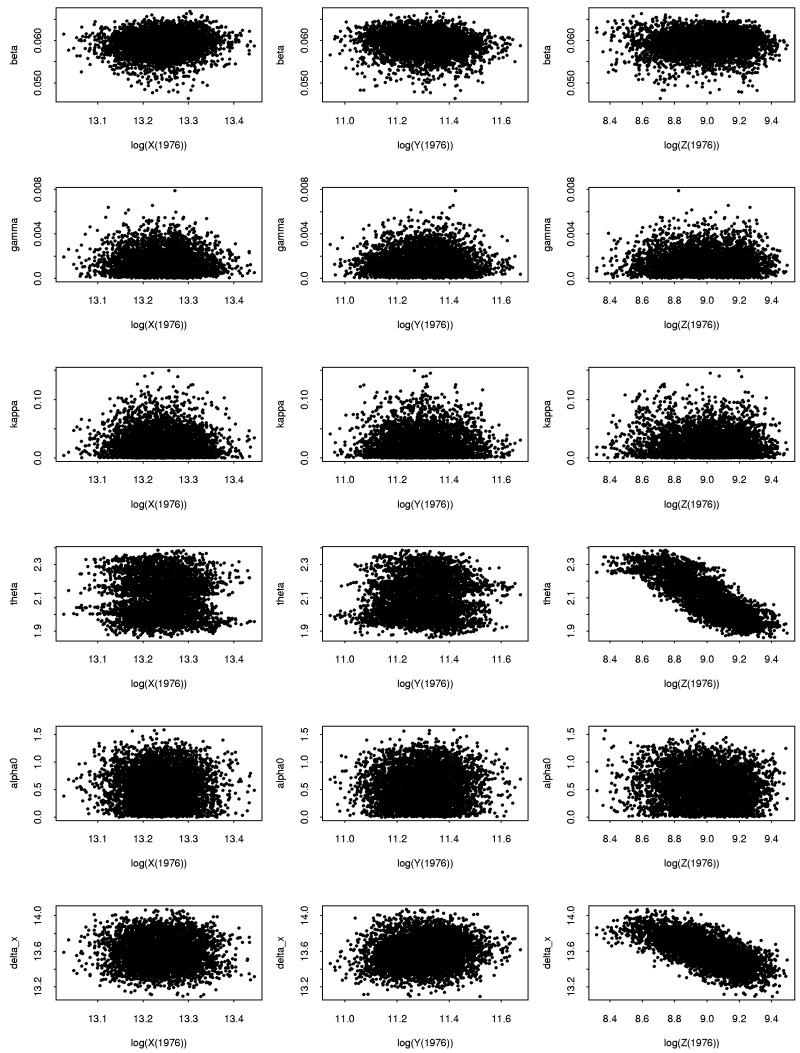

Figure 6.11: Scatter plots of the global parameters of the population dynamics against corresponding values the log estimated abundances for fjord no. 8 in year 1976.

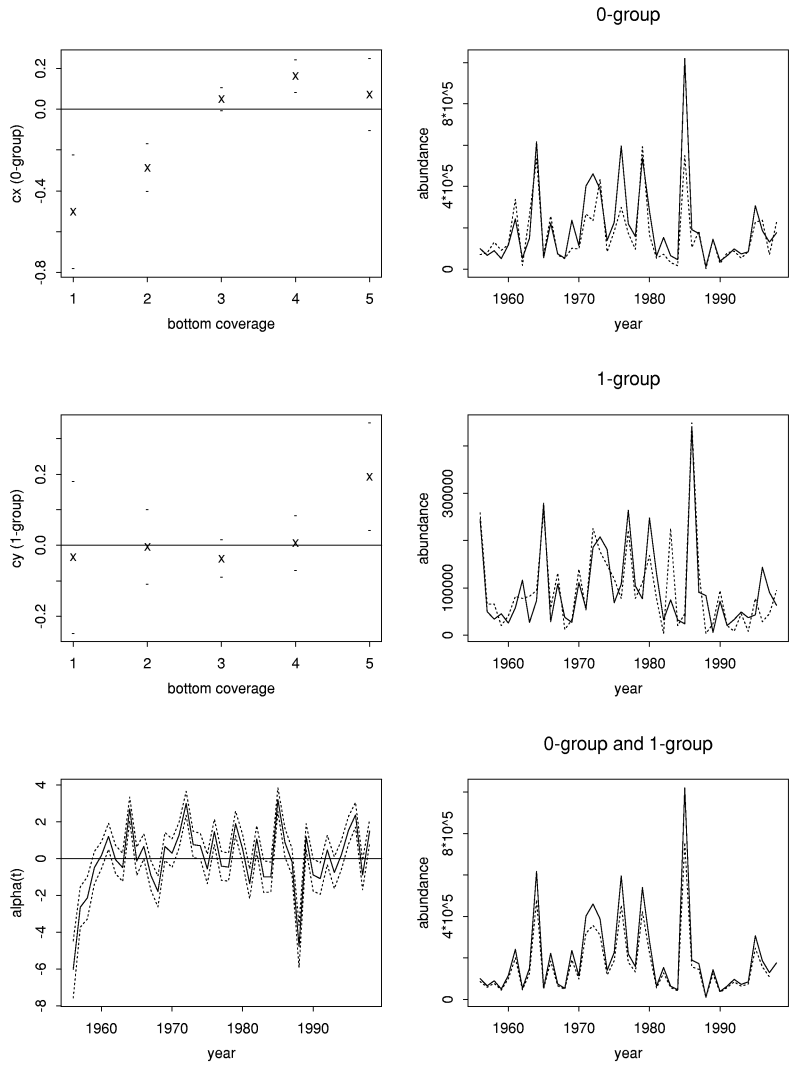

Figure 6.12: Estimated posterior means from applying Model III on data from 13 fjords. See Figure 6.2 for explanation. The three right panels show estimated abundances and up-scaled data for Søndeledfjord, one of the two fjords in the Risør area displayed in Figure 6.2.

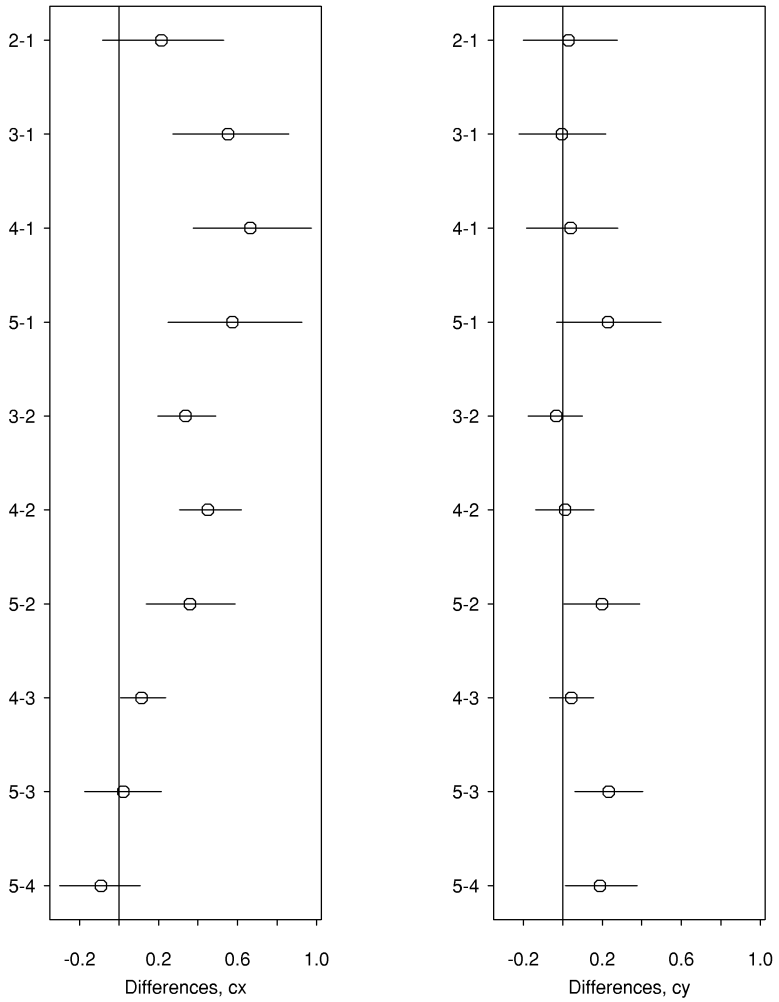

Figure 6.13: Estimated posterior means and empirical 95% credibility intervals for the differences $c_k^{s,x} - c_l^{s,x};\ k < l$ and $c_k^{s,y} - c_l^{s,y};\ k < l$ using data from 13 fjords.

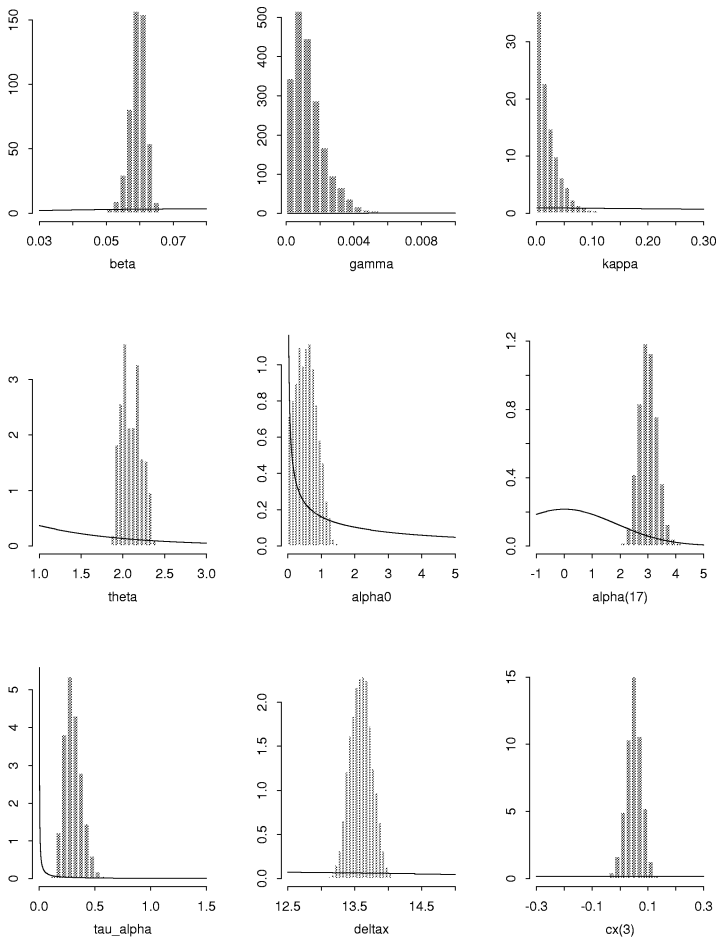

Figure 6.14: Histogram of the estimated marginal posterior distributions for the main population dynamics parameters based on the Metropolis-Hastings sampler, using data from all fjords. The prior distributions are added as full lines.

Chapter 7

A simulation study of the effect of releasing 0-group cod

As described in Section 2.2, the data on which we base our study originated from a discussion on the effect of releasing cod larvae on the population of cod within a fjord. Having established a model for the population dynamics and estimated the model parameters, we are now able to simulate the behaviour of the populations of the different age-groups after releasing young cod. The process of survival from the larval to the 0-group stage is not explicitly included in our model, but is defined as part of the recruitment process. Consequently, we are not able to simulate the original experiment releasing cod larvae, but we can study the effect on the future development of the population of any excess of 0-group juvenile cod resulting from a release of cod at the larval stage.

We perform the study by sampling from the prior model using the parameters of Table 6.2, estimated using the data from 13 fjords. Because of the Markov structure of the model, conditional on current and past abundances the data are not informative on future abundances. Consequently, we predict the future behaviour of the populations by sampling from the prior model. As initial values for the population abundances, we use the estimated posterior means of abundances of the final year for which we have data (1998). The estimated over-dispersion of $\exp(13.1)$ leads to relatively large prior variance for the 0-group abundances, such that sampling from the Gaussian prior might spuriously lead to negative values of $X_{f,t}$. To avoid this problem, we sample from a truncated Normal distribution.

We first simulate the population dynamics using the prior model given by (4.2), (4.5) and (4.7) for 10 years. Then we repeat the simulation procedure, using the same seed, but add a number $N_{x,f}$ to the initial value $\widehat{X}_{f,0}$ of the 0-group population abundance in fjord f. The values of the parameters of the population dynamics are taken to be the estimated posterior mean values except for the temporal structure $\alpha_t;\ t = 1, \ldots, 10$ of

the recruitment rate, which are generated by $\alpha_t \sim N(0, \hat{\tau}_\alpha^{-1})$. We simulate two scenarios, releasing an amount of 0-group cod corresponding to 50% and 100% of the initial abundances respectively. That is, the initial abundances are increased to $1.5\widehat{X}_{f,0}$ and $2\widehat{X}_{f,0}$. The results from the simulated experiment are shown in Figures 7.1 to 7.4, where we have plotted the median population abundances over 100 iterations. There seems to be an immediate effect of the release on the mature abundance for both amounts of released cod, but to assess the significance of this effect as well as of any long-term effects, the variability should be taken into account. To get an indication of the significance of the results, we perform t-tests based on the 100 realisations of the future mature abundances for each fjord and each year, testing for the differences in mean abundance with and without release. For both amounts of released 0-group cod, the differences for the second year after release, corresponding to the first year when the enhanced year-class is assumed to contribute to the mature population, are significant at the 95% level for all fjords. Releasing an amount corresponding to 100% of the initial populations, significant differences are also identified for year four after release in seven of the twelve fjords. Otherwise, a total of three differences for years three and six are the only ones that are found to be significant at the 95% level. The fact that there is evidence of significant effects for years two and four but not for the intermediate year, is in correspondence with the 2-2.5 year cycle for the Skagerrak cod population identified in earlier studies (e.g. Stenseth et al., 1999). We conclude that based on our model and the estimated parameters, there is evidence of a short-term effect of releasing 0-group cod amounting to 50% or 100% of the natural population abundances, but no significant long-term effects are found. These results are in correspondence with the main results from a similar study conducted by Chan et al. (2003a), extending the scenario to include several successive years of release. Although they identify statistically significant long-term effects for some fjords, these effects are considered to be too small to have any practical significance.

Using the estimated posterior means of the model parameters, we do not take into account the variability in the parameter estimates. If we conduct the simulation experiment using different updated values of the model parameters at each iteration, the simulation variance will increase. Therefore, although the results regarding significance of the short-term effect might be altered, taking the variability of the parameters into account should not lead to different conclusions regarding the long-term effects.

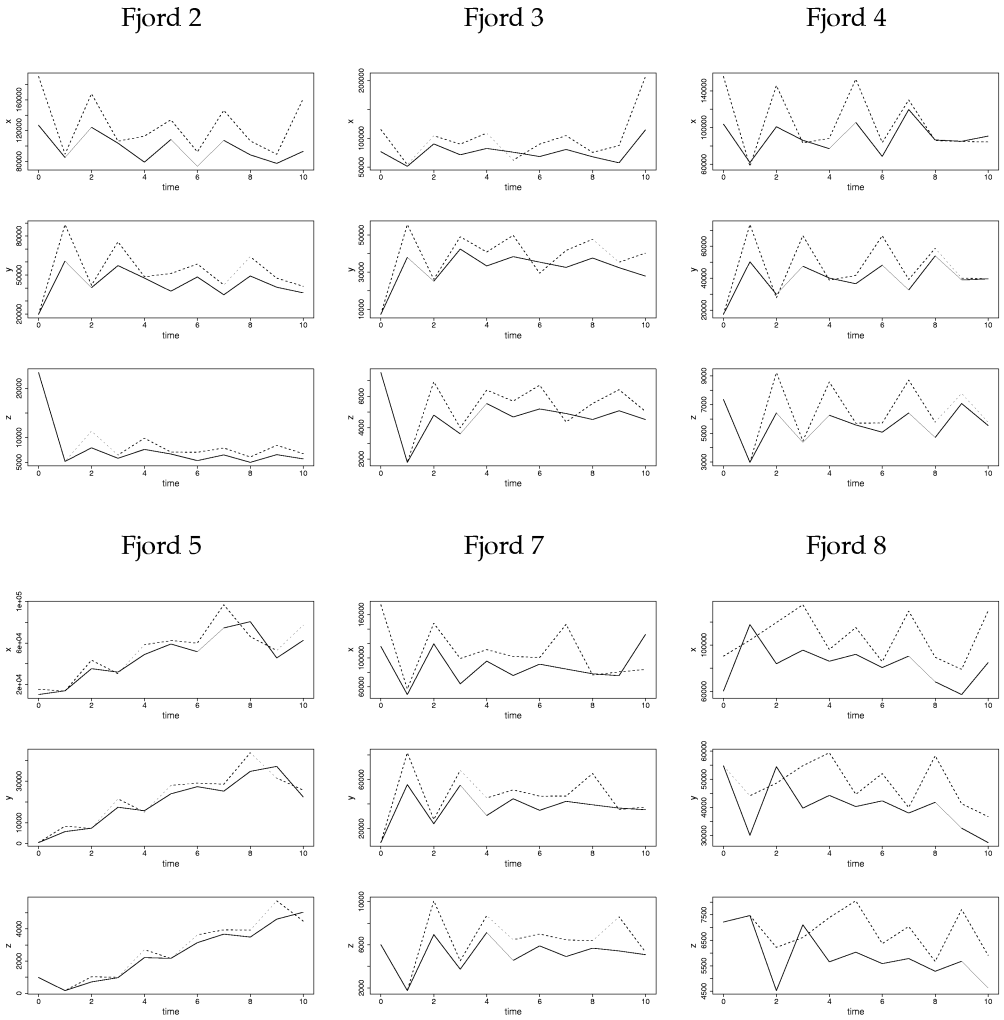

Figure 7.1: Results from the simulation study of the effect of releasing juvenile cod for six of the fjords. Each panel show the median over 100 iterations with (dashed line) and without (full line) releasing an amount of 0-group cod corresponding to 50% of the initial abundance.

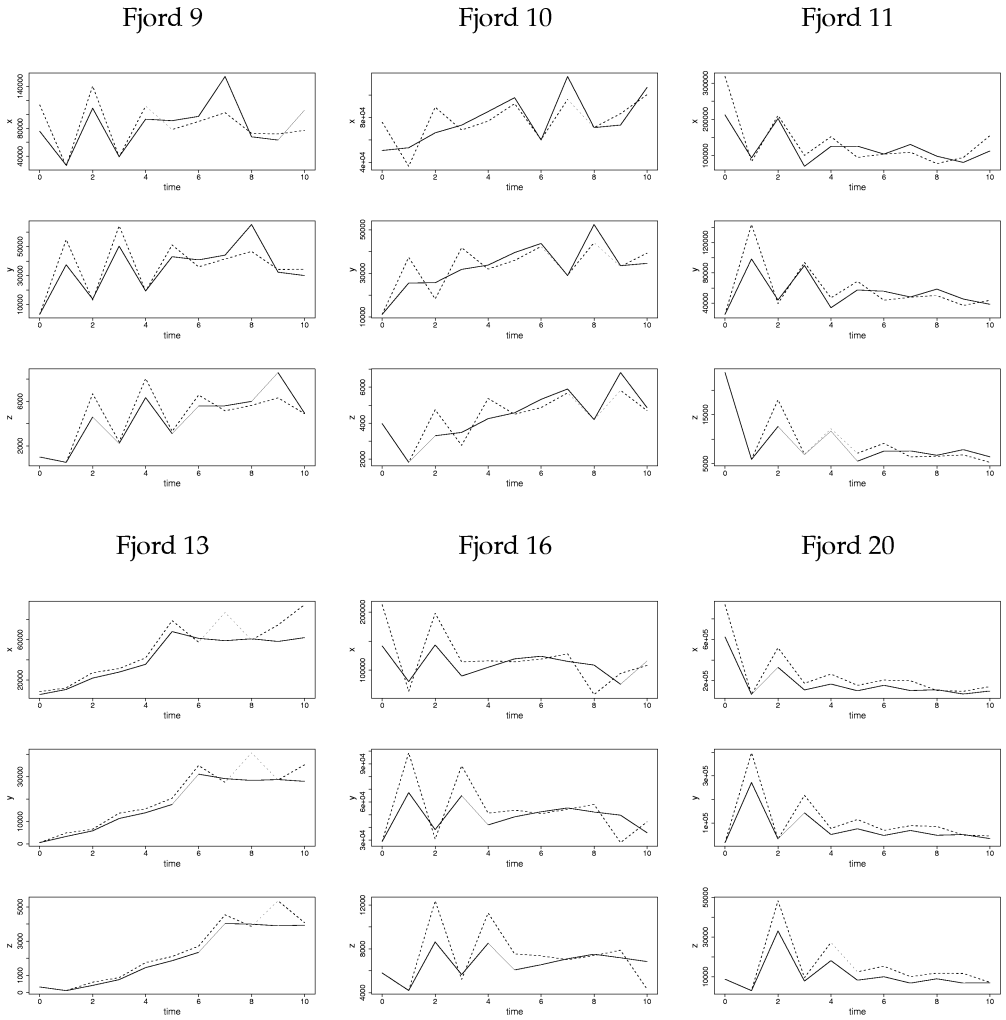

Figure 7.2: Results from the simulation study of the effect of releasing juvenile cod for six of the fjords. Each panel show the median over 100 iterations with (dashed line) and without (full line) releasing an amount of 0-group cod corresponding to 50% of the initial abundance.

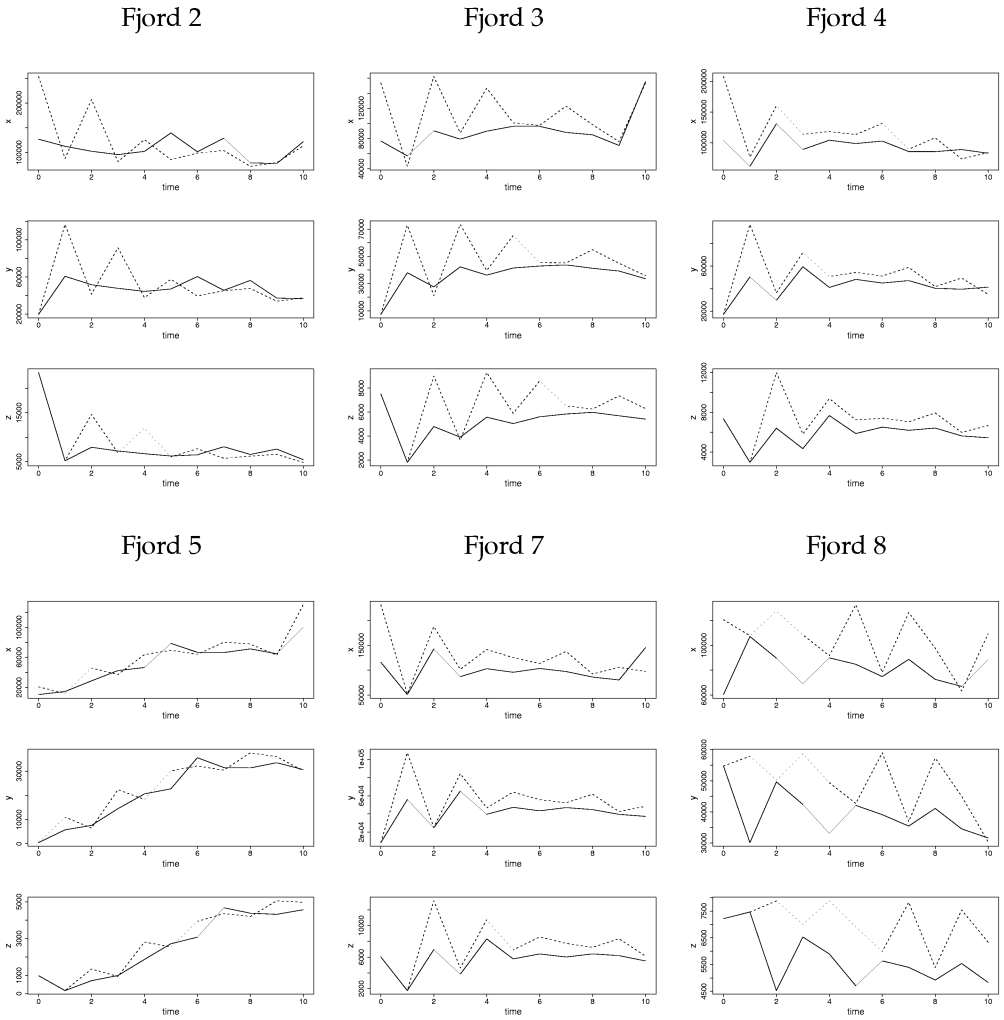

Figure 7.3: Results from the simulation study of the effect of releasing juvenile cod for six of the fjords. Each panel show the median over 100 iterations with (dashed line) and without (full line) releasing an amount of 0-group cod corresponding to 100% of the initial abundance.

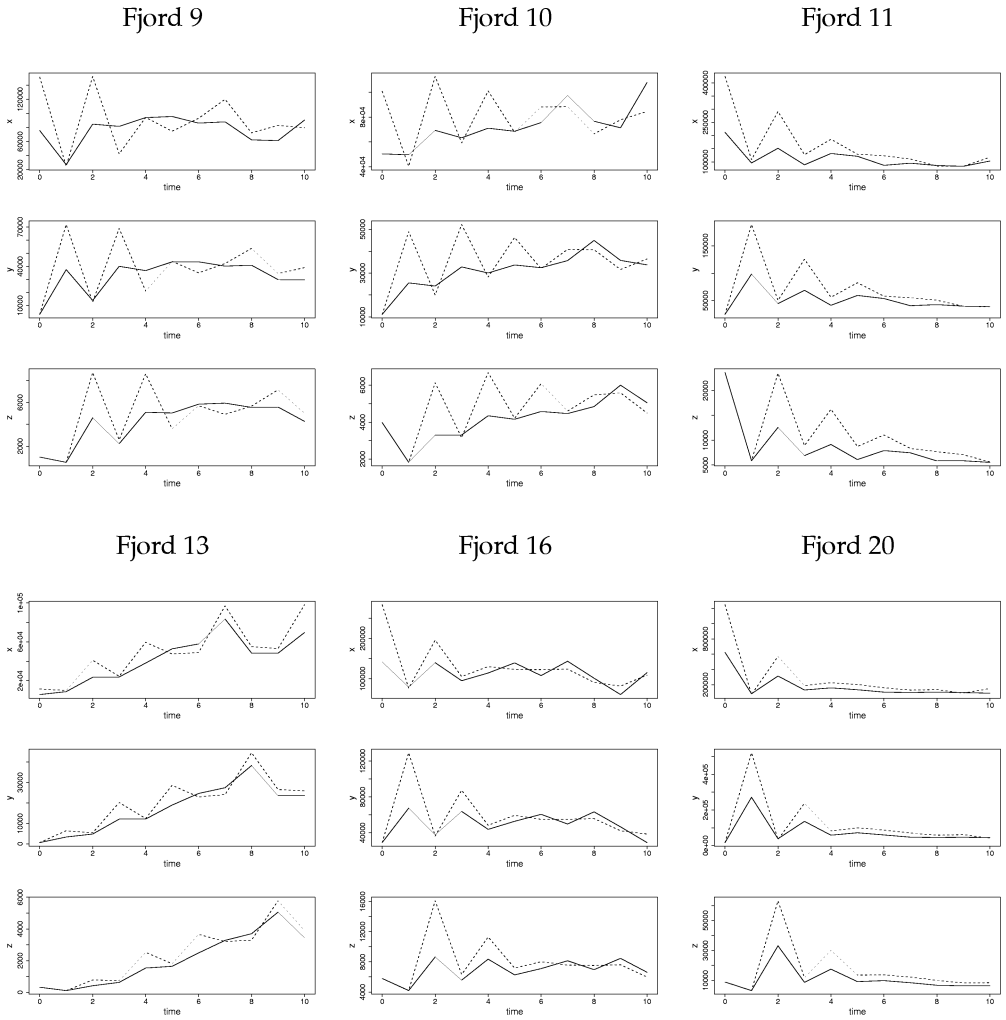

Figure 7.4: Results from the simulation study of the effect of releasing juvenile cod for six of the fjords. Each panel show the median over 100 iterations with (dashed line) and without (full line) releasing an amount of 0-group cod corresponding to 100% of the initial abundance.

Chapter 8

Sampling from the marginal posterior distribution of the parameters

The results from estimating the parameters of the model using the simulated data, indicated that the main estimation problems using single site updating are due to high correlations between the adult abundance and related model parameters. However, in many cases, the main interest is to estimate the model parameters governing the population dynamics and sampling process, while the abundances are of secondary interest. In this chapter we propose an alternative approach, were the parameters of the model are estimated by sampling from the marginal posterior of the model parameters, integrating over the population abundances. In Section 8.1 we give a description of the approach, and in Section 8.2 we study its performance using a set of simulated data.

8.1 The marginal posterior sampling approach

We first describe the approach for a general discrete time state-space model, specified in a Bayesian hierarchical model framework, before specifying the sampling procedure for the population dynamics model for the cod.

Let $\boldsymbol{V}_t$ denote the state vector of the model at time t. Further, let $\boldsymbol{V}$ denote the vector obtained by stacking the state vectors $\boldsymbol{V}_t; t = 1, \ldots, n_t$ on top of each other, that is

$$\boldsymbol{V} = \begin{pmatrix} \boldsymbol{V}_1 \\ \boldsymbol{V}_2 \\ \vdots \\ \boldsymbol{V}_{n_t} \end{pmatrix}. \tag{8.1}$$

Similarly, let $\boldsymbol{D}_t$ denote the observed vector quantity at time t, and $\boldsymbol{D}$ the corresponding

stacked vector, such that

$$\boldsymbol{D} = \begin{pmatrix} \boldsymbol{D}_1 \\ \boldsymbol{D}_2 \\ \vdots \\ \boldsymbol{D}_{n_t} \end{pmatrix}. \tag{8.2}$$

Finally, let

$$\boldsymbol{\theta} = \begin{pmatrix} \boldsymbol{\theta}_V \\ \boldsymbol{\theta}_D \end{pmatrix}, \tag{8.3}$$

denote the vector of the unknown parameters to be estimated. Here, $\boldsymbol{\theta}_V$ and $\boldsymbol{\theta}_D$ are the parameter vector of the system equation and the observation equation respectively, specifying the distributional assumptions of the state-space model. In our population dynamics model, the state vector $\boldsymbol{V}_t$, the data vector $\boldsymbol{D}_t$ and the unknown parameters $\boldsymbol{\theta}$ to be estimated are

$$\begin{aligned} \boldsymbol{V}_t &= (X_t, Y_t, Z_t)^T, && (8.4) \\ \boldsymbol{D}_t &= \left((\{X^o_{i,t}\}_i)^T, (\{Y^o_{i,t}\}_i)^T\right), \text{ and} && (8.5) \\ \boldsymbol{\theta} &= (\beta, \gamma, \kappa, \theta, \alpha_0, (\{\alpha_t\}_t)^T, \delta_x, (\{c^{s,x}_k\}_k)^T, (\{c^{s,y}_k\}_k)^T)^T. && (8.6) \end{aligned}$$

Here, we have suppressed the fjord subscript f of the abundances.

We specify the system and observation equations in the framework of a Bayesian hierarchical model, such that conditionally on the parameters $\boldsymbol{\theta}$, the system equation is represented by the conditional priors

$$\begin{aligned} &f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t-1}, \boldsymbol{\theta}_V); \quad t = 2, \ldots, n_t, \\ &f(\boldsymbol{v}_1 \mid \boldsymbol{\theta}_V), \end{aligned} \tag{8.7}$$

utilising the Markov property of the state-space model, and the observation equation by the likelihood

$$\prod_{t=1}^{n_t} f(\boldsymbol{d}_t \mid \boldsymbol{v}_t, \boldsymbol{\theta}_D). \tag{8.8}$$

The model is completed by defining the prior $f(\boldsymbol{\theta})$ of the parameter vector $\boldsymbol{\theta}$. A graphical description of the general model is given in Figure 8.1.

We will estimate the parameters $\boldsymbol{\theta}$ by sampling from the joint marginal posterior distribution of $\boldsymbol{\theta}$, given observations of the data vector $\boldsymbol{D}$. Referring to Figure 8.1, this corresponds to establishing direct edges between the nodes representing the data and the parameters, eliminating the intermediate nodes representing the abundances. The distribution is given by

$$f(\boldsymbol{\theta} \mid \boldsymbol{d}) = \int_{\boldsymbol{v}} f(\boldsymbol{\theta}, \boldsymbol{v} \mid \boldsymbol{d})\, d\boldsymbol{v}, \tag{8.9}$$

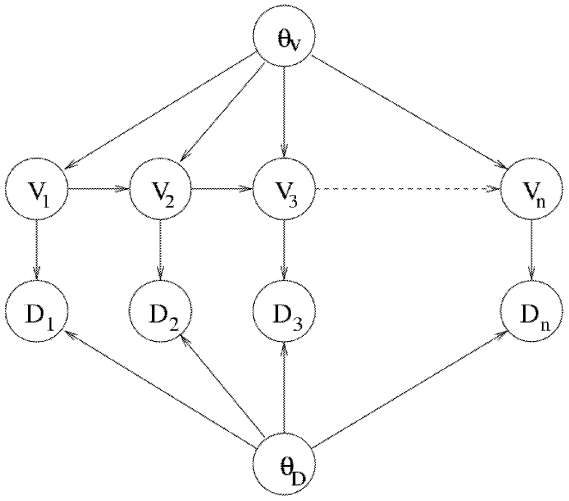

Figure 8.1: A graphical representation of the general state-space model, where $n = n_t$.

integrating over all values of the population dynamics vector $\boldsymbol{V}$. Using Bayes theorem, the marginal posterior distribution can be expressed by

$$f(\boldsymbol{\theta} \mid \boldsymbol{d}) = \frac{f(\boldsymbol{v}, \boldsymbol{\theta} \mid \boldsymbol{d})}{f(\boldsymbol{v} \mid \boldsymbol{\theta}, \boldsymbol{d})} \propto \frac{f(\boldsymbol{d} \mid \boldsymbol{v}, \boldsymbol{\theta}) f(\boldsymbol{v} \mid \boldsymbol{\theta}) f(\boldsymbol{\theta})}{f(\boldsymbol{v} \mid \boldsymbol{\theta}, \boldsymbol{d})}, \tag{8.10}$$

This identity is similar to the "basic marginal likelihood identity" used by Chib (1995) to compute the density of the marginal likelihood from the output from Gibbs sampling, as well as an analogue to the formula for the Bayesian predictive density of a future observation presented in Besag (1989). Sampling based estimates of $\boldsymbol{\theta}$ can be obtained using the Metropolis-Hasting algorithm with (8.10) as the target distribution. Schematically, the algorithm for obtaining m samples using uniform proposal distributions can be described as in Figure 8.2.

For each proposal $\boldsymbol{\theta}^{'}$ of $\boldsymbol{\theta}$, we need to evaluate (8.10) in both the current and proposed value of $\boldsymbol{\theta}$ in order to compute the acceptance rate. Since the normalising constants in general will depend on the parameters to be estimated, we need to evaluate normalised densities, a fact that might increase the computational burden substantially. In the nominator of (8.10) we recognise the likelihood and the joint prior distribution of the population sizes $\boldsymbol{v}$ and the model parameters $\boldsymbol{\theta}$. These can be computed from the model specification, as defined by the conditional prior distributions (4.2), (4.5) and (4.7) and the likelihood given by (4.15) and (4.18), so we focus on the evaluation of the denominator in (8.10).

Let $\boldsymbol{d}_1^t$ denote the vector obtained by stacking the data vectors for all points in time up to

$\boldsymbol{\theta}^{(0)} = \boldsymbol{\theta}_0$
for $(i = 1$ to $m)$
{
- Propose $\boldsymbol{\theta}^{'}$ from $\boldsymbol{\theta}^{'} \sim$ Unif$(\boldsymbol{\theta}^{(i-1)} \pm \frac{\boldsymbol{h}}{2})$
- Compute the acceptance probability
$\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}^{'}) = \frac{f(\boldsymbol{\theta}^{'}|\boldsymbol{d})}{f(\boldsymbol{\theta}^{(i-1)}|\boldsymbol{d})}$
- $\boldsymbol{\theta}^{(i)} = \boldsymbol{\theta}^{'}$ with prob. $\alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}^{'})$,
$\boldsymbol{\theta}^{(i)} = \boldsymbol{\theta}^{(i-1)}$ with prob. $1 - \alpha(\boldsymbol{\theta}^{(i-1)}, \boldsymbol{\theta}^{'})$.
}

Figure 8.2: A summary of the Metropolis-Hastings algorithm for sampling from the marginal posterior for the parameters $\boldsymbol{\theta}$, using uniform proposal distributions.

and including time t, and let $\boldsymbol{v}_1^t$ represent the corresponding quantity for given values of the state vector $\boldsymbol{V}$. The denominator can be evaluated using the forward-filtering-backward-sampling method (Frühwirth-Schnatter, 1994; Carter and Kohn, 1994), expressing the denominator in terms of the backward-sampling distribution

$$f(\boldsymbol{v}_1^{n_t} \mid \boldsymbol{d}_1^{n_t}, \boldsymbol{\theta}) = f(\boldsymbol{v}_{n_t} \mid \boldsymbol{d}_1^{n_t}, \boldsymbol{\theta}) \prod_{t=n_t-1}^{1} f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t+1}, \boldsymbol{d}_1^t, \boldsymbol{\theta}) \tag{8.11}$$

where

$$f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t+1}, \boldsymbol{d}_1^t, \boldsymbol{\theta}) \propto f(\boldsymbol{v}_t \mid \boldsymbol{d}_1^t, \boldsymbol{\theta}) f(\boldsymbol{v}_{t+1} \mid \boldsymbol{v}_t, \boldsymbol{\theta})). \tag{8.12}$$

The forward-filtering distributions $f(\boldsymbol{v}_t \mid \boldsymbol{d}_1^t, \boldsymbol{\theta})$ are given by

$$f(\boldsymbol{v}_t \mid \boldsymbol{d}_1^t, \boldsymbol{\theta}) \propto \int_{\boldsymbol{v}_{t-1}} f(\boldsymbol{v}_{t-1} \mid \boldsymbol{d}_1^{t-1}, \boldsymbol{\theta}) f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t-1}, \boldsymbol{\theta}) f(\boldsymbol{d}_t \mid \boldsymbol{v}_t, \boldsymbol{\theta}) d\boldsymbol{v}_{t-1}. \tag{8.13}$$

This integral is not available in closed form, and some form of numerical approximation is needed. We proceed by defining a discretisation of the state-space, and consequently replacing the integral in (8.13) by a sum over all values $\boldsymbol{v}_{t-1}^d$, where we use superscript d to denote discretised abundances. Thus, the integral (8.13) is approximated by

$$\begin{aligned} \int_{\boldsymbol{v}_{t-1}} & f(\boldsymbol{v}_{t-1} \mid \boldsymbol{d}_1^{t-1}, \boldsymbol{\theta}) f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t-1}, \boldsymbol{\theta}) f(\boldsymbol{d}_t \mid \boldsymbol{v}_t, \boldsymbol{\theta}) d\boldsymbol{v}_{t-1} \\ & \approx \sum_{\boldsymbol{v}_{t-1}^d} f(\boldsymbol{v}_{t-1}^d \mid \boldsymbol{d}_1^{t-1}, \boldsymbol{\theta}) f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t-1}^d, \boldsymbol{\theta}) f(\boldsymbol{d}_t \mid \boldsymbol{v}_t, \boldsymbol{\theta}) \Delta \boldsymbol{v}_{t-1}^d. \end{aligned} \tag{8.14}$$

The sum (8.14) can in principle be evaluated for any discretisation of the state-vector. However, even for small dimensions q of the state vector $\boldsymbol{V}_t$, the computational cost increases rapidly with the resolution of the discretisation. For simplicity, we assume that

the grid size k is equal for all elements of the state vector, but in general, the discretisation can be of different resolutions for the different elements. As shown in Table 8.1, evaluation of the sum (8.14) requires an $\mathcal{O}(n_t k^{2q})$ - algorithm, where n_t is the number of observations in time. The normalising constant has to be evaluated since it might depend on the unknown parameter vector $\boldsymbol{\theta}$.

Source	Original scale	Virtual scale
Un-normalised FFBS distribution	$\mathcal{O}(k^q)$	$\mathcal{O}(k^{q-(r-1)})$
Normalisation	$\mathcal{O}(k^q)$	$\mathcal{O}(k^q)$
Time (no. of observations)	$\mathcal{O}(n_t)$	$\mathcal{O}(rn_t)$
Total	$\mathcal{O}(n_t k^{2q})$	$\mathcal{O}(rn_t k^{2q-(r-1)})$

Table 8.1: Computational cost of evaluating the sum (8.14) approximating the unnormalised forward-filtering distributions (8.13) and of evaluating the normalising constant. The cost is computed for the original and the virtual time scale.

To reduce the computational cost, we introduce a new time-scale, splitting the original time step into sub-steps. The motivation is to define a sub-time scale utilising the structure of the conditional prior specification, such that only one or a few components of the state vector are updated at each sub-step. The resulting virtual time scale using r sub-steps is illustrated in Figure 8.3. Using this virtual time scale, the computational cost is reduced from $\mathcal{O}(n_t k^{2q})$ to $\mathcal{O}(rn_t k^{2q-(r-1)})$, as shown in Table 8.1.

$$\begin{array}{ccccccccc}
\boldsymbol{v}_{t-1} & \rightarrow & \boldsymbol{v}_{t-(r-1)/r} & \cdots & \rightarrow & \boldsymbol{v}_{t-1/r} & \rightarrow & \boldsymbol{v}_t & \\
t-1 & & & & & & & t & \text{(real)} \\
r(t-1) & & rt-(r-1) & \cdots & & rt-1 & & rt & \text{(virtual)}
\end{array}$$

Figure 8.3: A virtual time scale using r sub-steps.

The effect on the summation (8.14) of introducing the virtual time scale is best illustrated using an example, and we apply the above procedure to our state-space model for the cod population dynamics. The original time step is split into three sub-steps, one for each element of the state vector. The resulting updating scheme on the new virtual time scale is illustrated in Figure 8.4. Depending on the sub-time-step t, the sum (8.14) is taken over all values of the discretisation of Z_{t-1} (the z-step), X_{t-1} (the x-step) or Y_{t-1} (the y-step). The computational cost is reduced from $\mathcal{O}(n_t k^6)$ to $\mathcal{O}(3n_t k^4)$, as described in Table 8.2.

Defining the abundances on the extended virtual time scale and denoting the virtual time by s, the conditional priors $f(\boldsymbol{v}_s \mid \boldsymbol{v}_{s-1}, \boldsymbol{\theta})$, entering into the summations (8.14), are

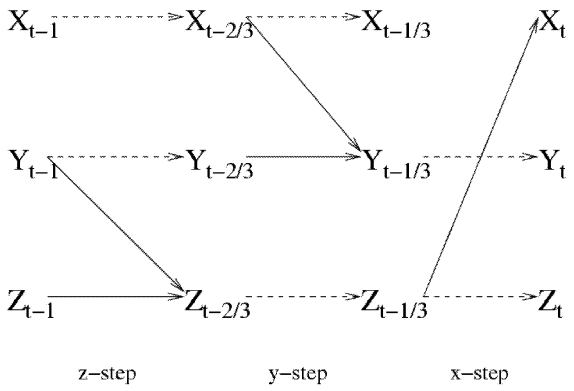

Figure 8.4: Structure of sub-model, introducing a virtual time scale. The figure shows one time-step in the real time scale, corresponding to three time steps in the virtual scale. The virtual intermediate sub-steps are denoted $t - 2/3$ and $t - 1/3$. The full lines represent a change of value, while the dotted lines indicate that the corresponding variable is kept fixed.

given from (4.2), (4.5) and (4.7) by

$$\begin{array}{ll} \text{z-step:} & f(\boldsymbol{v}_s \mid \boldsymbol{v}_{s-1}, \boldsymbol{\theta}) = f(z_s \mid y_{s-1}, z_{s-1}, \boldsymbol{\theta}) I_{[x_s = x_{s-1}]} I_{[y_s = y_{s-1}]}, \\ \text{y-step:} & f(\boldsymbol{v}_s \mid \boldsymbol{v}_{s-1}, \boldsymbol{\theta}) = f(y_s \mid y_{s-1}, x_{s-1}, \boldsymbol{\theta}) I_{[x_s = x_{s-1}]} I_{[z_s = z_{s-1}]}, \\ \text{x-step:} & f(\boldsymbol{v}_s \mid \boldsymbol{v}_{s-1}, \boldsymbol{\theta}) = f(x_s \mid z_{s-1}, \boldsymbol{\theta}) I_{[z_s = z_{s-1}]} I_{[y_s = y_{s-1}]}. \end{array} \tag{8.15}$$

The joint prior $f(\boldsymbol{v} \mid \boldsymbol{\theta})$ needed in the nominator of (8.10) is

$$\begin{aligned} f(\boldsymbol{v} \mid \boldsymbol{\theta}) &= \prod_{s=1}^{3n_t} f(\boldsymbol{v}_s \mid \boldsymbol{v}_{s-1}, \boldsymbol{\theta}) \\ &= \prod_{s;s\%3=1} f(z_s \mid z_{s-1}, y_{s-1}, \boldsymbol{\theta}) \prod_{s;s\%3=2} f(y_s \mid x_{s-1}, y_{s-1}, \boldsymbol{\theta}) \prod_{s;s\%3=0} f(x_s \mid z_s, \boldsymbol{\theta}) \\ &= \prod_{t=1}^{n_t} f(\boldsymbol{v}_t \mid \boldsymbol{v}_{t-1}, \boldsymbol{\theta}), \end{aligned} \tag{8.16}$$

where t denotes the time steps of the real time scale. The likelihood is evaluated by

$$f(\boldsymbol{d} \mid \boldsymbol{v}, \boldsymbol{\theta}) = \prod_{i=1}^{n_s} \prod_{t=1}^{n_t} f(x^o_{i,t} \mid x_{i,t}, \{c^{s,x}_k\}) f(y^o_{i,t} \mid y_{i,t}, \{c^{s,y}_k\}), \tag{8.17}$$

where the product is taken over all n_s sampling locations i in the fjord, and all n_t years t. We assume independent priors for the parameters $\boldsymbol{\theta}$, so $f(\boldsymbol{\theta}) = \prod_j f(\theta_j)$.

Source	Original scale	Virtual scale
Un-normalised FFBS distribution	$\mathcal{O}(k^3)$	$\mathcal{O}(k)$
Normalisation	$\mathcal{O}(k^3)$	$\mathcal{O}(k^3)$
Time (no. of observations)	$\mathcal{O}(n_t)$	$\mathcal{O}(3n_t)$
Total	$\mathcal{O}(n_t k^6)$	$\mathcal{O}(3n_t k^4)$

Table 8.2: Computational cost of evaluating the sum (8.14) for the population dynamics model. The cost is computed for the original and the virtual time scale.

8.2 An application of the approach to simulated data

To explore the properties of the sampling algorithm, we apply the method to a set of simulated data. To reduce the computational burden, we use a shorter time series than in Section 5.3. We simulate 10 years of data, but generate data for 20 stations, to compensate in part for the reduced number of observations in time. Also, we do not use observed values of the covariate $d_{i,t}$, but sample $d_{i,t}$ at random, assuming a uniform distribution over all possible values. This is done in order to reduce the variability in precision of the estimates of the settling probability factor due to the sparsity in the number of data points for extreme values of $d_{i,t}$. The parameters determining the survival rates as well as the log recruitment rate α_0 are equal to the parameters for the full simulated data set, given in Table 5.3, while the over-dispersion parameter is increased to $\delta_x = 11.0$. The temporal fluctuations in the log recruitment rate are specified by sampling from the distribution given by $\alpha_t \sim N(0, \tau_\alpha^{-1})$, where $\tau_\alpha = 1.0$, and adjusted such that $\sum_{t=1}^{10} \alpha_t = 0$. The resulting data set is illustrated in Figure 8.5.

A crucial step of the proposed algorithm is to choose a representative discretisation. The discretisation should cover the range of possible sizes of the populations, and at the same time, the resolution k should be kept at a minimum so as to reduce the computational cost, which was shown to be proportional to k^4. The abundances are unknown, but some guidance to the range of values can be given from the a rough up-scaling of the data, and this can be combined with biological information on the order of magnitude of the abundances. Referring to the considerations on sampling probabilities in Section 4.2.2 and the corresponding prior specifications in Section 5.1, rough estimates of the number of individuals of the juvenile age groups, relying on the data as well as prior information on cod movement, can be obtained by up-scaling the mean number of counts at each point in time by a factor $\Omega(f)/(0.9\psi_x L)$ for the 0-group and $\Omega(f)/(0.9\psi_y L)$ for the 1-group. For the adult population no data are available, so we estimate the abundances sequentially in time by $\hat{Z}_t = (\hat{Z}_{t-1} + \hat{Y}_{t-1}) \exp(-\theta)$. The survival rate $\exp(-\theta)$ is unknown, and this quantity is sampled for each t. The values are sampled uniformly over the interval $[\exp(-1) * 0.8, \exp(-1) * 1.2]$, centred at the value

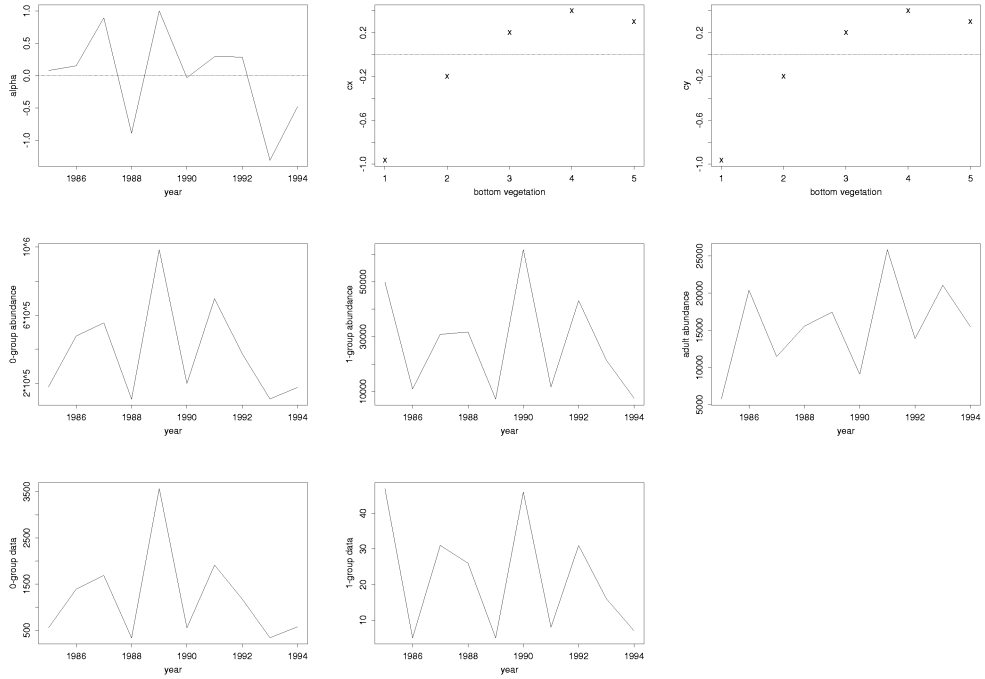

Figure 8.5: Parameters (upper panels), simulated population abundances on fjord level (middle panels) and sampled data as a total for all 20 stations (bottom panels) for the simulated data set.

$\exp(-\theta) = \exp(1)$. For the first year, we use equilibrium considerations, estimating Z_1 by the value $\hat{Z}_1 = \overline{\hat{Y}_t}\exp(-\theta)/(1-\exp(-\theta))$, assuming the mature population to be stable, and using the average rough abundance estimate of the 1-group cod as an estimate of Y_0. Formally, we generate rough estimates of the abundances for each of the 10 years by

$$\begin{aligned}
\hat{X}_t &= \frac{1}{n_s}\sum_{i=1}^{n_s} X_{i,t}\frac{\Omega(f)}{0.9L\psi_x}, \\
\hat{Y}_t &= \frac{1}{n_s}\sum_{i=1}^{n_s} Y_{i,t}\frac{\Omega(f)}{0.9L\psi_y}, \\
\hat{Z}_1 &= \overline{\hat{Y}_t}\frac{\exp(-\theta)}{1-\exp(-\theta)}, \\
\hat{Z}_t &= (\hat{Y}_{t-1} + \hat{Z}_{t-1})\exp(-\theta).
\end{aligned} \tag{8.18}$$

The variability of the estimates (8.18) of the abundances is assessed in Appendix D.

The choice of discretisation based on the abundance estimates from the up-scaling as given in (8.18) is illustrated in Figure 8.6. The spread of the abundance estimates for the 10 years gives an estimate of the range of probable values for the abundances. Studying the empirical distribution of the estimates, we get an indication of how to choose the discretisation such as to obtain highest resolution in the regions of the state space with highest density. As illustrated in Figure 8.6, by plotting the sorted abundance estimates for the 10 years we find that equal spacing in log scale for the 0-group and 1-group abundances and in the original scale for the adult abundance seems appropriate. In the bottom panels, discretisations in 10 intervals are shown, and for comparison, the true simulated abundances are added as dotted lines.

To study the effect of varying the resolution, we evaluate the marginal posterior distributions, ignoring priors, for each parameter in turn, keeping the remaining parameters fixed at their true values. That is, for each parameter θ_j we evaluate

$$f(\theta_j \mid \boldsymbol{d}, \boldsymbol{\theta}_{-j}) \propto \frac{f(\boldsymbol{d} \mid \boldsymbol{v_e}, \boldsymbol{\theta}) f(\boldsymbol{v_e} \mid \boldsymbol{\theta})}{f(\boldsymbol{v_e} \mid \boldsymbol{\theta}, \boldsymbol{d})}, \tag{8.19}$$

in a selected value $\boldsymbol{v}_e$. Defining $\boldsymbol{x}^d$, $\boldsymbol{y}^d$ and $\boldsymbol{z}^d$ to be the length k vectors of discretised values of the 0-group, 1-group and adult abundances, we choose $\boldsymbol{v}_e$ to be the value $(x^d_{k/2}, y^d_{k/2}, z^d_{k/2})$. The results are shown in Figures 8.7 and 8.8. For α_0 and δ_x there is hardly any effect of increasing the resolution of the discretisation, except for a small shift in the distribution for α_0 going from resolution 10 to resolution 20. This might indicate that the data have limited information on these parameters, in accordance with the near-identifiability problems experienced when estimating abundances as well as model parameters based on data from one fjord only. The effect of increasing the resolution is much more pronounced for the parameters β and γ, the density-dependent effects of the juvenile survival, as well as for the density-independent survival parameters θ and κ. From the plots we observe that the effect is large going from resolution 10 to resolution 20 and still substantial increasing the resolution to 30. By further increasing the resolution the posteriors grow smoother, but the range of values with significant posterior probability does not change much. Based on these results we conclude that using a discretisation of resolution 30 seems sufficient to estimate the posterior range of these parameters, but the resolution should be further increased if estimating the full posterior distribution is an important issue.

We ran the sampling algorithm using $K = 30$, basing the discretisation on the results presented in Figure 8.6. The computational cost is still relatively large; each iteration took 6-7 seconds on a 2.8 GHz Pentium 4 computer. Estimating all parameters, the convergence is slow, so in order to improve the convergence properties, we fix δ_x and θ to the true values. Trace plots from running the algorithm for 520000 iterations using the simulated data set are given in Figure 8.9. For comparison, in Figures 8.10 and 8.11 the trace plots are shown together with similar plots from applying the Metropolis-

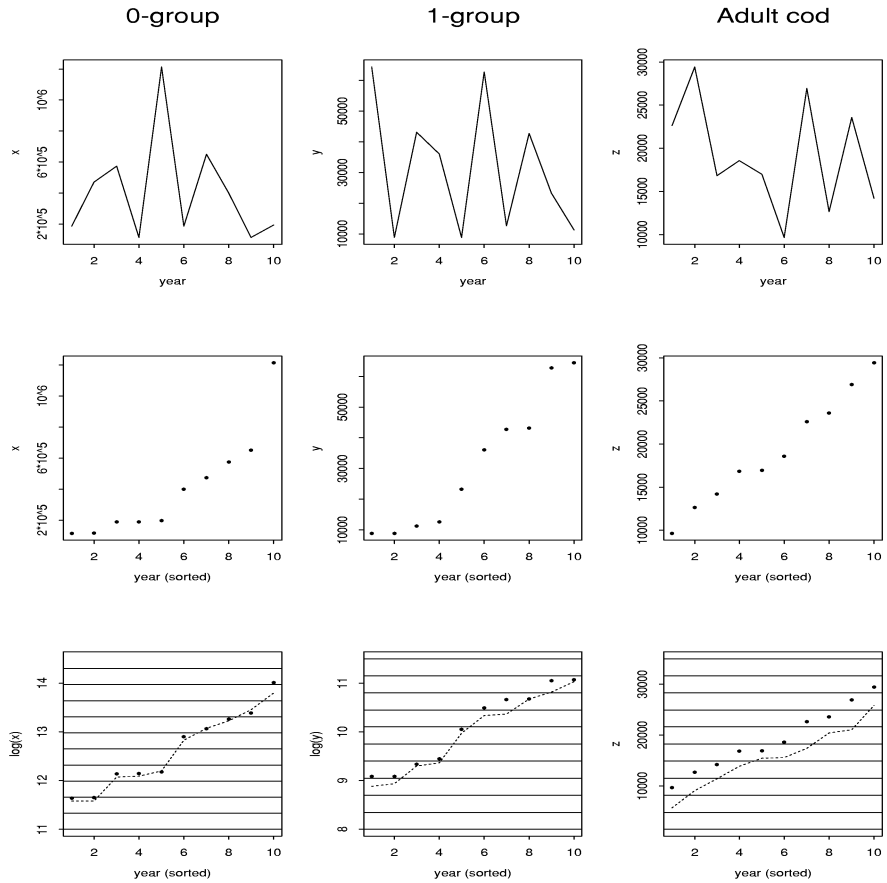

Figure 8.6: Discretisation of the abundances for (from left to right) the 0-group, 1-group and adult population abundances. The top panels show the estimated abundances from up-scaling, the middle panel the sorted estimates, and the bottom panel the resolution 10 discretisation based on log abundances for the juvenile cod. The dotted lines are the true simulated abundances.

Hastings algorithm of Chapter 5 to the same data set. For α_0 and the hyper-parameter τ_α the algorithm is still not considered to have converged after the 520000 iterations, and from Figures 8.10 and 8.11 the mixing is extremely poor as compared to the result obtained when estimating the abundances as well. From Figure 8.8, the parameter θ seems to be better identified than α_0 such that fixing α_0 and estimating θ might be a more reasonable choice using the marginal posterior sampling approach. The convergence properties for the temporal random effects α_t, represented in the plots by the updates of α_5, is better, and the levels are comparable to the ones obtained by the algorithm of Chapter 5. However, we observe that the simulation variance is smaller than for the other approach, and this might be related to the fact that the updates of the prior precision τ_α stay near the initial value 1. The range of the posterior distributions for β, γ and κ as well as the parameters $\{c_l^{s,x}\}_l$ and $\{c_l^{s,y}\}_l$ are all well identified. These parameters were also the best behaved parameters in our study in Chapter 6, and the results confirm the initial findings presented in Figure 8.7.

8.3 Summary and discussion

We have presented an approach to sampling based estimation of the model parameters, sampling from the marginal posterior distribution of the parameters. Using this approach, we eliminate the need for updating the population abundances, which for the mature cod were shown in Chapters 5 and 6 to be highly correlated with several of the model parameters. The approach is based on specifying the marginal posterior distribution in terms of the prior and likelihood models and the joint full conditional distribution of the abundances, evaluating this full conditional distribution using the forward-filter-backward-sampling method. The requirement of evaluation of the normalising constants of the forward-filtering distributions leads to a substantial increase in computational cost. By introducing a virtual time-scale, we obtained a reduction of the cost of the sampling algorithm from $\mathcal{O}(n_t k^6)$ to $\mathcal{O}(3n_t k^4)$ for our hierarchical model for the Skagerrak cod.

Applying the algorithm to a simulated data set, the ranges of the parameters related to juvenile survival, as well as the parameters measuring relative effects, were well identified, and comparable to results from using the Metropolis-Hastings algorithm estimating the abundances as well. However, slow convergence still remains a problem. As we observed in Section 6.2, increasing the amount of data and the number of fjords improved on the convergence and mixing of the Metropolis-Hastings sampler. Since the fjord populations are assumed to be separate, the marginal posterior sampling approach can in principle easily be extended to the case of several fjords of data by replacing the

marginal posterior (8.10) by

$$\prod_f f(\boldsymbol{\theta} \mid \boldsymbol{d}_f) = \prod_f \frac{f(\boldsymbol{v}_f, \boldsymbol{\theta} \mid \boldsymbol{d}_f)}{f(\boldsymbol{v}_f \mid \boldsymbol{\theta}, \boldsymbol{d}_f)} \propto \prod_f \frac{f(\boldsymbol{d}_f \mid \boldsymbol{v}_f, \boldsymbol{\theta}) f(\boldsymbol{v}_f \mid \boldsymbol{\theta}) f(\boldsymbol{\theta})}{f(\boldsymbol{v}_f \mid \boldsymbol{\theta}, \boldsymbol{d}_f)}. \tag{8.20}$$

However, the CPU time will be roughly proportional to the number of fjords, and although we have made efforts to optimise the code such as to reduce the computational cost, we are still left with a substantial number of evaluations for each iteration of the sampler.

The choice of discretisation is obviously a crucial point. The resolution should be high enough to avoid point masses in the prior distributions as defined on the discretised abundances, at the same time as we need to focus on parsimony in the number of levels to keep the computational cost down. We have used the same resolution for all age-groups, but alternatively the resolution could be allowed to vary between the age-groups. In light of the poor mixing of α_0 and τ_α, a more dense discretisation for the adult cod might be needed.

The results of the previous chapters indicated that the main convergence problems were related to the mature population abundances. To reduce the computational burden of the marginal posterior sampling algorithm, an alternative would be to marginalise over the mature abundances only, and update the 0-group and 1-group abundances with the model parameters. The cost of updating the mature abundance will be $\mathcal{O}(n_t k^2)$. The prior model for juvenile survival (4.5) does not depend on the $Z_{f,t}$'s such that MCMC updates for the 1-group cod can be generated by a Metropolis-Hastings step with the full conditional as target distribution. However, the prior (4.2) of the 0-group cod does depend on the mature cod. Marginalising the posterior distribution for $X_{f,t}$ using an analogue to expression (8.10) for integrating out the $Z_{f,t}$'s requires a $\mathcal{O}(n_t^2 k^2)$-algorithm for each fjord, such that the cost of the algorithm will be proportional to the square of the number of years of data.

Although poor mixing still remains a problem for several parameters, we have illustrated that the marginal posterior sampling approach is capable of capturing information in the data on the parameters. We therefore consider the approach to have a potential in the estimation of the parameters of the population dynamics model in cases were the abundances are of secondary interest, but at present the approach has practical limitations due to the large amount of CPU time required.

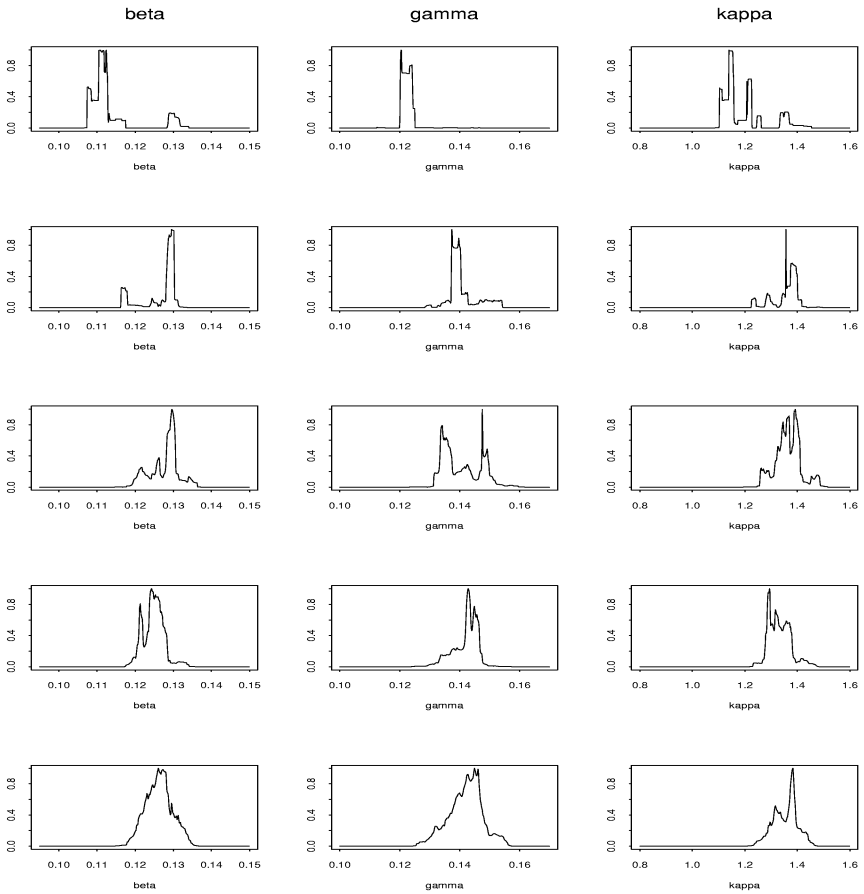

Figure 8.7: The effect of changing the resolution of the discretisation. The plots show the marginal posterior distribution for each parameter (ignoring the prior), fixing the rest of the parameters at their true values. The resolutions are (from the top downward) 10^3, 20^3, 30^3,40^3 and 50^3.

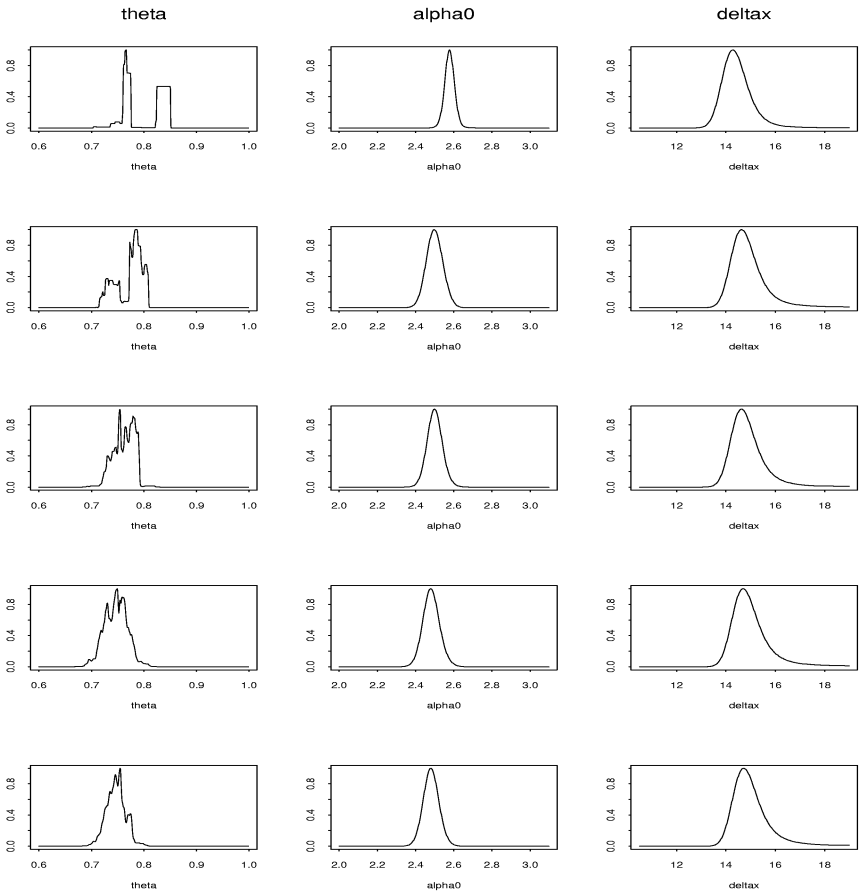

Figure 8.8: The effect of changing the resolution of the discretisation. The plots show the marginal posterior distribution for each parameter (ignoring the prior), fixing the rest of the parameters at their true values. The resolutions are (from the top downward) 10^3, 20^3, 30^3,40^3 and 50^3.

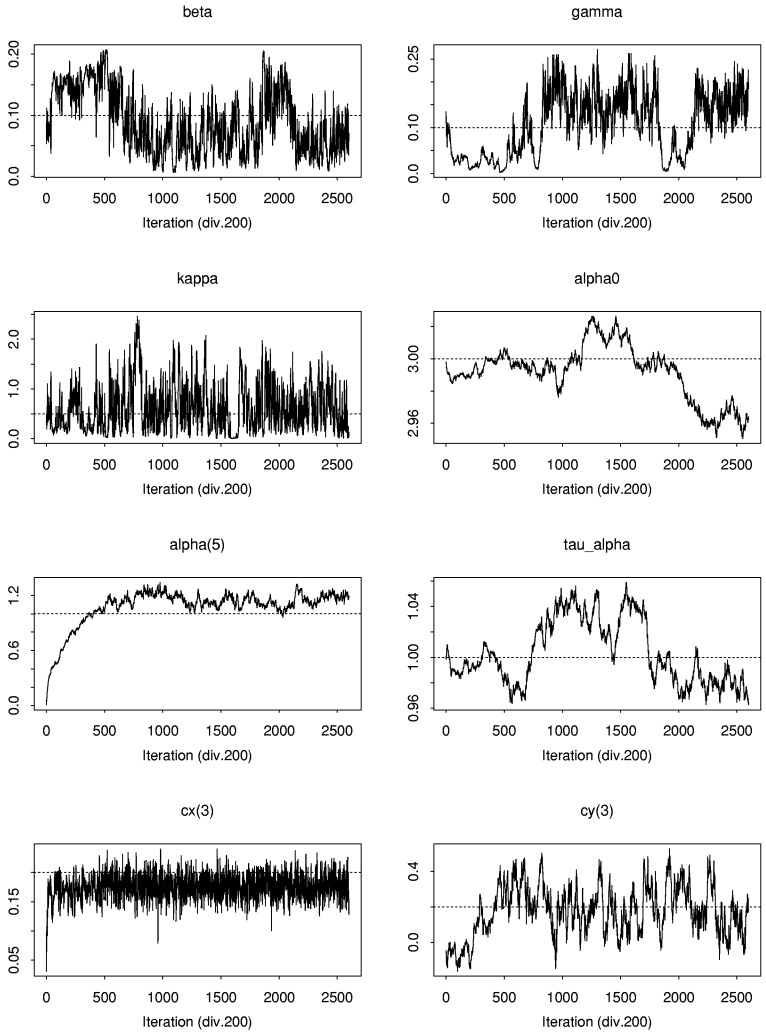

Figure 8.9: Trace plots from applying the marginal posterior sampling approach to the simulated data set of 10 years and 20 stations. The horizontal dotted lines indicate the values used to generate the data.

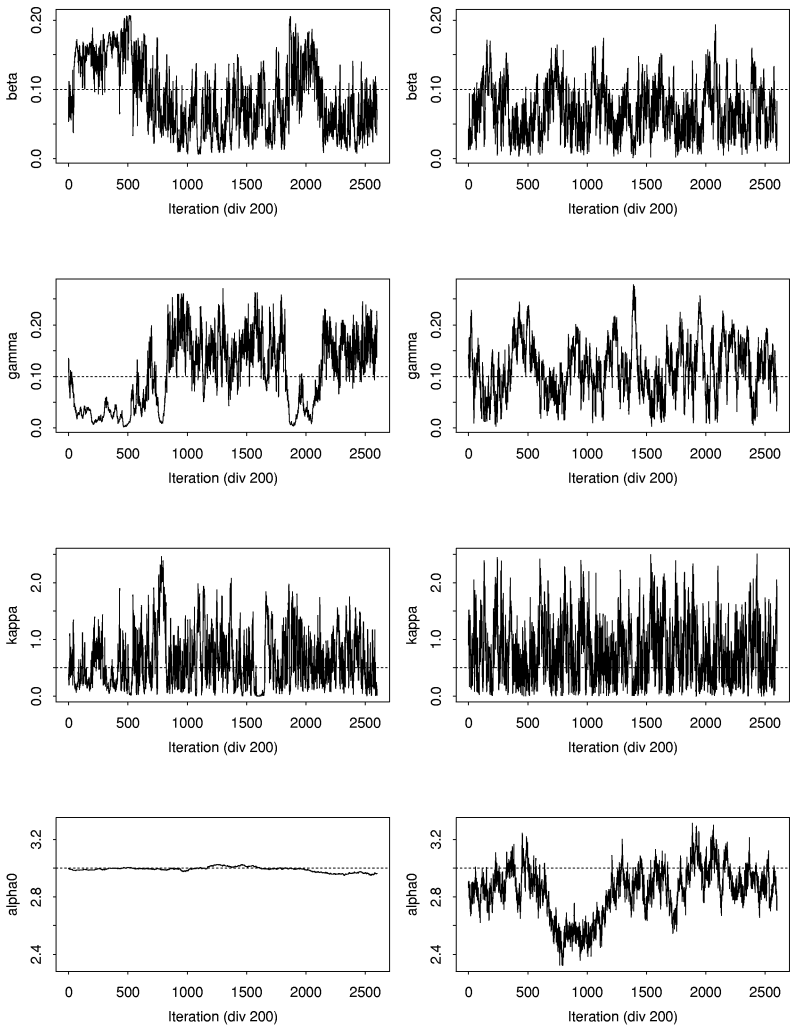

Figure 8.10: Trace plots from the marginal posterior sampling approach (left panels) and corresponding results from running the single-site Metropolis-Hastings algorithm (right panels) for the simulated data set of 10 years and 20 stations. The horizontal dotted lines indicate the values used to generate the data.

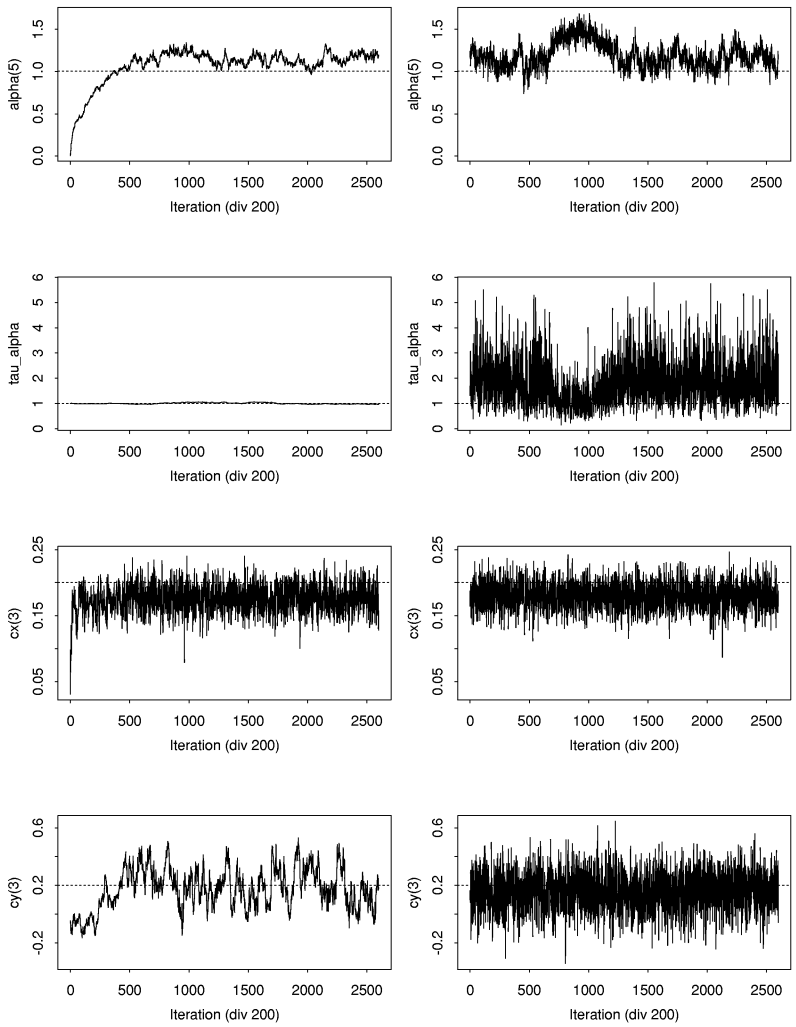

Figure 8.11: Trace plots from the marginal posterior sampling approach (left panels) and corresponding results from running the single-site Metropolis-Hastings algorithm (right panels) for the simulated data set of 10 years and 20 stations. The horizontal dotted lines indicate the values used to generate the data.

Chapter 9

Discussion

We have illustrated how a model for the population dynamics of coastal cod can be specified within the framework of a Bayesian hierarchical model, estimating the parameters using a fully Bayesian approach. The model represents an improvement compared to other models applied to these data since stochastic variation is accounted for at all levels of the model, including an explicit stochastic model for the sampling process. Consequently, the population dynamics is specified for the fjord level populations and not for populations proportional to the sampled counts of individuals.

Using a set of simulated data, we have studied the properties of the model in terms of convergence and mixing of a Metropolis-Hastings approach to the estimation of the model parameters and the abundances. Re-parameterising the model was shown to improve mixing compared to the results for the model on its original form. However, convergence problems persist due to high correlations between the adult abundance, for which we have no data, and related model parameters, in particular the instantaneous mortality θ and the level of the log-recruitment rate α_0. Increasing the number of fjords, the problem is reduced, but slow convergence and unsatisfactory mixing remains a problem for the adult abundance as well as the instantaneous adult mortality θ and the over-dispersion parameter δ_x of the prior model for recruitment. For the Risør data, for which abundances for only one fjord were estimated, the parameters θ and δ_x were kept fixed.

The estimated marginal posterior distributions of the model parameters were shown to be more concentrated than the priors, confirming that there is information in the data on the parameters of the model. Moreover, the estimated temporal structure of the rate of recruitment of new individuals to the population indicated a particularly low rate in 1988, recognising the algae bloom of that year. The results for the juvenile abundances showed that the temporal structure of the estimated posterior means was reproduced. The fit was better for the 0-group than of the 1-group abundances, confirming prior belief that there is more information in the 0-group data than the 1-group data, a result

that can be related to the larger fraction of zero counts in the 1-group data.

Of particular interest to biologists is to identify the presence or absence of density-dependency in survival between the juvenile stages. The results from our study gave clear evidence of a within-group density-dependency, in correspondence with other studies. On the other hand, no evidence was found for a corresponding significant density-dependency between the 0-group and 1-group cod.

The unknown parameters of the likelihood, relating the sampling probability to a categorical covariate representing the amount of vegetation on the bottom, were computationally well-behaved and well identified for all parameterisations of the model. The results from using the data for 13 fjords indicated that the settling probability increases with the degree of bottom coverage, but with a slight decrease for a totally covered bottom, possibly due to a less efficient beach seine catch.

Running a sensitivity analysis varying the factors ψ_x and ψ_y of the sampling probabilities, that were fixed based on biological information, posterior mean estimates of the abundances, as well as for the population dynamics parameters that were not estimated relative to a reference level, were shown to depend on the chosen values of these factors. Moreover, not all parameters are scale-invariant. Therefore, care should be taken when interpreting estimated values as well as comparing these values to the results from other studies.

In our problem, estimating the population abundances is of secondary interest compared to the estimation of the parameters of the population dynamics. We proposed an approach to parameter estimation based on sampling from the marginal posterior distribution of the model parameters, integrating over the population abundances. Although the mixing turned out to be poor for several parameters, applying the approach to a simulated data set it was shown that the method was able to capture information on the parameters available in the data. Introducing a virtual time scale we obtained a significant reduction of the computational cost, but the required CPU time is still large, such that at present the approach has practical limitations.

An interesting extension of the model would be to estimate temporal trends in the parameters. Since the region of study is a popular region for leisure fishing, it would be of particular interest to allow for temporal variability in the parameter θ representing the adult mortality, such that trends in the number of tourists visiting the area might be reflected in temporal fluctuations of the fishing mortality. In our model we consider only total mortality, but as suggested by Chan et al. (2003b), we could distinguish between natural and fishing mortality.

Another extension could be to allow the parameters to be fjord specific, in light of the results of e.g. Fromentin et al. (2001), that indicated that the population dynamics is somewhat different in the north-eastern and the south-western fjords. It is suggested that the regional differences in the population dynamics, e.g. in the extent of predation

and cannibalism, relate to differences in the amount of bottom vegetation. In Section 3.2 we pointed out that several studies indicate that abiotic factors like hydrography and climate influence the population dynamics, and in general, incorporating more biotic and abiotic covariates in the parameters of the Bayesian hierarchical model would be an interesting direction for further work.

Finally, we return to the question initiating the monitoring programme that has generated the data on which our study is based. To imitate the experiment that was conducted in the early 20th century to settle the dispute between G.M. Dannevig and J.Hjort on the effect of releasing cod larvae, we performed a simulation study of the effect of releasing juvenile cod on the abundance of the mature population. In the real word experiment, the cod was released at the larval stage, but since our model does not distinguish between the rates of recruitment and of survival to the 0-group stage, we simulated the release of juvenile 0-group cod. Consequently, in our simulated experiment we focused on the effect on the mature population conditionally on the survival to the 0-group stage. The results indicated that there seems to be an immediate effect on the size of the adult population releasing amounts of 0-group cod corresponding to 50% and 100% of the initial populations, but no significant long-term effect of releasing juvenile cod was found.

Acknowledgements

The author would like to thank the Institute of Marine Research, Flødevigen Marine Research Station in Arendal for providing the data, and in particular Jakob Gjøsæter for providing information on the biological and historical background on the data. The author would also like to thank Nils Christian Stenseth and Kyrre Lekve at the Department of Biology, University of Oslo as well as Jakob Gjøsæter for interesting discussions of biological issues and interpretation of results. Finally, the author is grateful to Håvard Rue, Norwegian University of Science and Technology, and Arnoldo Frigessi, Norwegian Computing Center, for advice and stimulating discussions concerning the statistical modelling and inference.

Bibliography

Besag, J. (1989). A candidate's formula: A curious result in Bayesian prediction, *Biometrika* **76**: 183–183.

Bjørnstad, O. N., Fromentin, J.-M., Stenseth, N. C. and Gjøsæter, J. (1999a). Cycles and trends in cod populations, *Proceedings of the National Academy of Sciences of the United States of America* **96**: 5066–5071.

Bjørnstad, O. N., Fromentin, J.-M., Stenseth, N. C. and Gjøsæter, J. (1999b). A new test for density-dependent survival: the case of coastal cod populations, *Ecology* **80**: 1278–1288.

Buckland, S. T., Goudie, I. B. J. and Borchers, D. L. (2000). Wildlife population assessment: past developments and future directions, *Biometrics* **56**: 1–12.

Carlin, B. P., Polson, N. G. and Stoffer, D. S. (1992). A Monte Carlo approach to nonnormal and nonlinear state-space modeling, *Journal of the American Statistical Association* **87**: 493–500.

Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models, *Biometrika* **81**: 541–553.

Chan, K.-S., Stenseth, N. C., Kittilsen, M. O., Gjøsæter, J., Lekve, K., Smith, T., Tveite, S. and Danielssen, D. (2003a). Assessing the effectiveness of releasing cod larvae for stock improvement with monitoring data, *Ecological Applications* **13**: 3–22.

Chan, K.-S., Stenseth, N. C., Lekve, K., Gjøsæter, J. and Ottersen, G. (2003b). Population dynamics of cod along the Norwegian Skagerrak coast (I): a spatio-temporal model for cod incorporating both biotic and abiotic effects, Unpublished.

Chib, S. (1995). Marginal likelihood from the Gibbs output, *Journal of the American Statistical Association* **90**: 1313–1321.

Cushing, D. H. (1990). Plankton production and year-class strength in fish populations: an update of the match/mismatch hypothesis, *Advances in Marine Biology* **26**: 249–293.

Dahl, K. and Dannevig, G. M. (1906). Undersøgelser over nytten af torskeudklækning i østlandske fjorde, *Aarsberetning vedkommende Norges fiskerier* pp. 1–121. (In Norwegian).

Danielssen, D. S. and Gjøsæter, J. (1994). Release of 0-group cod, Gadus morhua L., on the southern coast of Norway in the years 1986-1989, *Aquaculture and Fisheries Management* **25**: 129–142.

De Valpine, P. and Hastings, A. (2002). Fitting population models incorporating process noise and observation error, *Ecological Monographs* **72**: 57–76.

Freeman, S. N. and Kirkwood, G. P. (1995). On a structural time-series method for estimating stock biomass and recruitment from catch and effort data, *Fisheries Research* **22**: 77–98.

Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models, *Journal of Time Series Analysis* **15**: 183–202.

Fromentin, J.-M., Gjøsæter, J., Bjørnstad, O. N. and Stenseth, N. C. (2000). Biological processes and environmental factors regulating the dynamics of the Norwegian Skagerrak cod populations since 1919, *ICES Journal of Marine Science* **57**: 330–338.

Fromentin, J.-M., Myers, R. A., Bjørnstad, O. N., Stenseth, N. C., Gjøsæter, J. and Christie, H. (2001). Effects of density-dependent and stochastic processes on the regulation of cod populations, *Ecology* **82**: 567–579.

Fromentin, J.-M., Stenseth, N. C., Gjøsæter, J., Bjørnstad, O. N., Falck, W. and Johannessen, T. (1997). Spatial patterns of the temporal dynamics of three gadoid species along the Norwegian Skagerrak coast, *Marine Ecology Progress Series* **155**: 209–222.

Fromentin, J.-M., Stenseth, N. C., Gjøsæter, J., Johannessen, T. and Planque, B. (1998). Long-term fluctuations in cod and pollack along the Norwegian Skagerrak coast, *Marine Ecology Progress Series* **162**: 265–278.

Gavaris, S. and Ianelli, J. N. (2002). Statistical issues in fisheries' stock assessments (with discussion), *Scandinavian Journal of Statistics* **29**: 245–271.

Geweke, J. and Tanizaki, H. (2001). Bayesian estimation of state-space models using the Metropolis-Hastings within Gibbs sampling, *Computational Statistis and Data Analysis* **37**: 151–170.

Gilks, W. R. and Roberts, G. O. (1996). Strategies for improving MCMC, *in* W. R. Gilks, S. Richardson and D. J. Spiegelhalter (eds), *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, UK, pp. 89–114.

Gilks, W. R., Richardson, S. and Spiegelhalter, D. J. (1996). *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, UK.

Gjøsæter, J. and Danielssen, D. S. (1990). Recruitment of cod (Gadus morhua), whitting (Merlangius merlangus) and pollack (Pollachius pollachius) in the Risør area on the Norwegian Skagerrak coast 1945 to 1985, *Flødevigen rapportserie* **1**: 11–31.

Gjøsæter, J., Enersen, K. and Enersen, S. E. (1996). Ressurser av torsk og andre fisk i fjorder på den Norske Skagerrak-kysten, *Project report: Fisken og havet 23*, Institute of Marine Research. (In Norwegian).

Gjøsæter, J., Stenseth, N. C., Ottersen, G., Lekve, K., Dahl, E., Danielssen, D. S., Torstensen, E. and Christie, H. (2003a). The key fish species along the Norwegian Skagerrak coast with an emphasis on immature cod: their environment and their interactions, Unpublished.

Gjøsæter, J., Stenseth, N. C., Sollie, A. and Lekve, K. (2003b). The Flødevigen beach seine surveys: a unique long-term monitoring program, Unpublished.

Gudmundsson, G. (1994). Time series analysis of catch-at-age observations, *Journal of the Royal Statistical Society, Series C - Applied Statistics* **43**: 117–126.

Hilborn, R. and Walters, C. J. (1992). *Quantitative Fisheries Stock Assessment. Choice, Dynamics and Uncertainty*, Chapman and Hall, New York.

Hjort, J. (1914). Fluctuations in the great fisheries of northern Europe viewed in the light of biological research, *Rapports et Procès-Verbaux du Conseil Permanent International pour l'Exploration de la Mer* **20**: 1–228.

Hjort, J. (1926). Fluctuations in the year classes of important food fishes, *Journal du Counceil Permanent International pour l'Exploration de la Mer* **1**: 5–38.

Johannessen, T. and Sollie, A. (1994). Overvåking av gruntvannsfauna på Skagerrakkysten - historiske forandringer i fiskefauna 1919-1993, og ettervirkninger av den giftige algeoppblomstringen i mai 1988, *Project report: Fisken og havet 10*, Institute of Marine Research. (In Norwegian).

Julliard, R., Stenseth, N. C., Gjøsæter, J., Lekve, K., Fromentin, J.-M. and Danielssen, D. S. (2001). Natural mortality and fishing mortality in a coastal cod population: a release-recapture experiment, *Ecological Applications* **11**: 540–558.

Kimura, D. K., Balsiger, J. W. and Ito, D. H. (1996). Kalman filtering the delay-difference equation: practical approaches and simulations, *Fishery Bulletin* **94**: 678–691.

Kitagawa, G. (1987). Non-Gaussian state-space modeling of nonstationary time series (with discussion), *Journal of the American Statistical Association* **82**: 1032–1063.

Knight, S. M., Weir, I. S. and Pettitt, A. N. (2003). Estimating fish numbers: a case study involving the Maroochy River in south east Queensland, *Australian and New Zealand Journal of Statistics*. Submitted.

Lekve, K., Ottersen, G., Stenseth, N. C. and Gjøsæter, J. (2002). Length dynamics in juvenile coastal Skagerrak cod: Effects of biotic and abiotic processes, *Ecology* **83**: 1676–1688.

Meyer, R. and Millar, R. B. (1999). Bayesian stock assessment using a state-space implementation of the delay difference model, *Canadian Journal of Fisheries and Aquatic Sciences* **56**: 37–52.

Millar, R. B. and Meyer, R. (2000a). Bayesian state-space modeling of age-structured data: fitting a model is just the beginning, *Canadian Journal of Fisheries and Aquatic Sciences* **57**: 43–50.

Millar, R. B. and Meyer, R. (2000b). Non-linear state space modelling of fisheries biomass dynamics by using Metropolis-Hastings within-Gibbs sampling, *Journal of the Royal Statistical Society, Series C - Applied Statistics* **49**: 327–342.

Myers, R. A. (2001). Stock and recruitment: generalizations about maximum reproductive rate, density dependence, and variability using meta-analytic approaches, *ICES Journal of Marine Science* **58**: 937–951.

Myers, R. A. and Cadigan, N. G. (1993a). Density dependent juvenile mortality in marine demersal fish, *Canadian Journal of Fisheries and Aquatic Sciences* **50**: 1576–1590.

Myers, R. A. and Cadigan, N. G. (1993b). Is juvenile natural mortality in marine demersal fish variable?, *Canadian Journal of Fisheries and Aquatic Sciences* **50**: 1591–1598.

Newman, K. B. (1998). State-space modeling of animal movement and mortality with application to salmon, *Biometrics* **54**: 1290–1314.

Polacheck, T., Hilborn, R. and Punt, A. E. (1993). Fitting surplus production models: comparing methods and measuring uncertainty, *Canadian Journal of Fisheries and Aquatic Sciences* **50**: 2597–2607.

Punt, A. E. and Hilborn, R. (1997). Fisheries stock assessment and decision analysis: the Bayesian approach, *Reviews in Fish Biology and Fisheries* **7**: 35–63.

Quinn, T. J. and Deriso, R. B. (1999). *Quantitative Fish Dynamics*, Oxford University Press, New York.

Reed, W. J. and Simons, C. M. (1996). Analyzing catch-effort data by means of the Kalman filter, *Canadian Journal of Fisheries and Aquatic Sciences* **53**: 2157–2166.

Robert, C. P. and Casella, G. (1999). *Monte Carlo statistical methods*, Springer-Verlag, New York.

Schnute, J. T. (1994). A general framework for developing sequential fisheries models, *Canadian Journal of Fisheries and Aquatic Sciences* **51**: 1676–1688.

Schnute, J. T. and Kronlund, A. R. (2002). Estimating salmon stock-recruitment relationships from catch and escapement data, *Canadian Journal of Fisheries and Aquatic Sciences* **59**: 433–449.

Schwach, V. (2000). *Havet, fisken og vitenskapen. Fra fiskeriundersøkelser til havforskningsinstitutt 1860-2000*, Havforskningsinstituttet. (In Norwegian).

Smith, T. D., Gjøsæter, J., Stenseth, N. C., Kittilsen, M. O., Danielssen, D. S., Solemdal, P. and Tveite, S. (2003). A century of manipulating recruitment in coastal cod populations: the Flødevigen experience, *ICES Journal of Marine Science*. In press.

Solemdal, P., Dahl, E., Danielssen, D. S. and Moksness, E. (1984). The cod hatchery in Flødevigen - background and realities, *in* E. Dahl, D. S. Danielssen, E. Moksness and P. Solemdal (eds), *The Propagation of Cod, Gadus morhua L., Part I*, Institute of Marine Research, Flødevigen Biological Station, Arendal, Norway, pp. 11–45.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and van der Linde, A. (2002). Bayesian measures of model complexity and fit (with discussion), *Journal of the Royal Statistical Society, Series B* **64**: 583–639.

Stenseth, N. C., Bjørnstad, O. N., Falck, W., Fromentin, J.-M., Gjøsæter, J. and Grey, J. S. (1999). Dynamics of coastal cod populations: intra- and intercohort density dependence and stochastic processes, *Proceedings of the Royal Society of London Series B* **266**: 1645–1654.

Sullivan, P. J. (1992). A Kalman filter approach to catch-at-length analysis, *Biometrics* **48**: 237–257.

Tupper, M. and Boutilier, R. G. (1995). Effects of habitat on settlement, growth, and postsettlement survival of Atlantic cod (Gadus morhua), *Canadian Journal of Fisheries and Aquatic Sciences* **52**: 1834–1841.

Tveite, S. (1971). Fluctuations in year-class strength of cod and pollack in southeastern Norwegian coastal waters during 1920-1969, *Fiskeridirektoratets Skrifter Serie Havundersøkelser* **16**: 65–76.

Vines, S. K., Gilks, W. R. and Wild, P. (1996). Fitting Bayesian multiple random effects models, *Statistics and Computing* **6**: 337–346.

West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models, 2nd ed.*, Springer-Verlag, New York.

Appendix A

The data

Fjord number	Fjord name	Station number	Station name
1	Torvefjord	3	Rauskjær
2	Topodalsfjord	12	Justnes
		15	Børresholmen
		16	Solbustad
		17	Vigvold
		18	Søm - øst
3	Høvåg	31	Jakteviga
		33	Steindalsfjord, Kvarsnes
		36	Lusekilen indre
		37	Lusekilen ytre
		39	Østervik - odde
		41	Fjelldalsøya - ytre
4	Bufjord - Grimstad	56	Bufjorden - ytre
		57	Bufjorden - indre
5	Flødevigen	65	Lillehavn
		66	Kumkrogen

Table A.1: Fjords and stations used in the study.

Fjord number	Fjord name	Station number	Station name
(**R**) 7	Sandnesfjord, Risør	91	Laget - indre
		92	Laget - ytre
		93	Håholmen, odden
		94	Håholmen, bukta
		95	Åmland - ytre
		96	Åmland - indre
		97	Løkvik - indre
		98	Løkvik - ytre
(**R**) 8	Søndeledfjord, Risør	111	Sundet, nord
		112	Sundet, syd
		121	Øymoen - ytre
		122	Øymoen - indre
9	Risør skerries	141	Varøybukta
		142	Varøysund
10	Stølefjord, Kragerø	151	Stølefjorden - indre
		152	Stølefjorden - ytre
11	Kilsfjord, Kragerø	161	Langvarp - Eidskilen
		163	Lyngdalen
		164	Blankenberg
13	Soppekilen, Kragerø	182	Soppekilen - ytre
		185	Soppekilen - indre
16	Nøtterø - Tjøme	231	Saltbu
		232	Ordal
		233	Sevik
		235	Breivik - øst
		238	Årøysundet - indre
20	Hvaler	341	Dypeklo, holmen
		342	Dypeklo, fastland
		343	Hellesvikskilen
		344	Dragesund - vest
		345	Dragesund - øst
		348	Papperhavn, indre
		349	Papperhavn, smalt sund

Table A.1 (cont.): Fjords and stations used in the study. The two fjords making up the subset denoted the "Risør data" are marked with an (**R**).

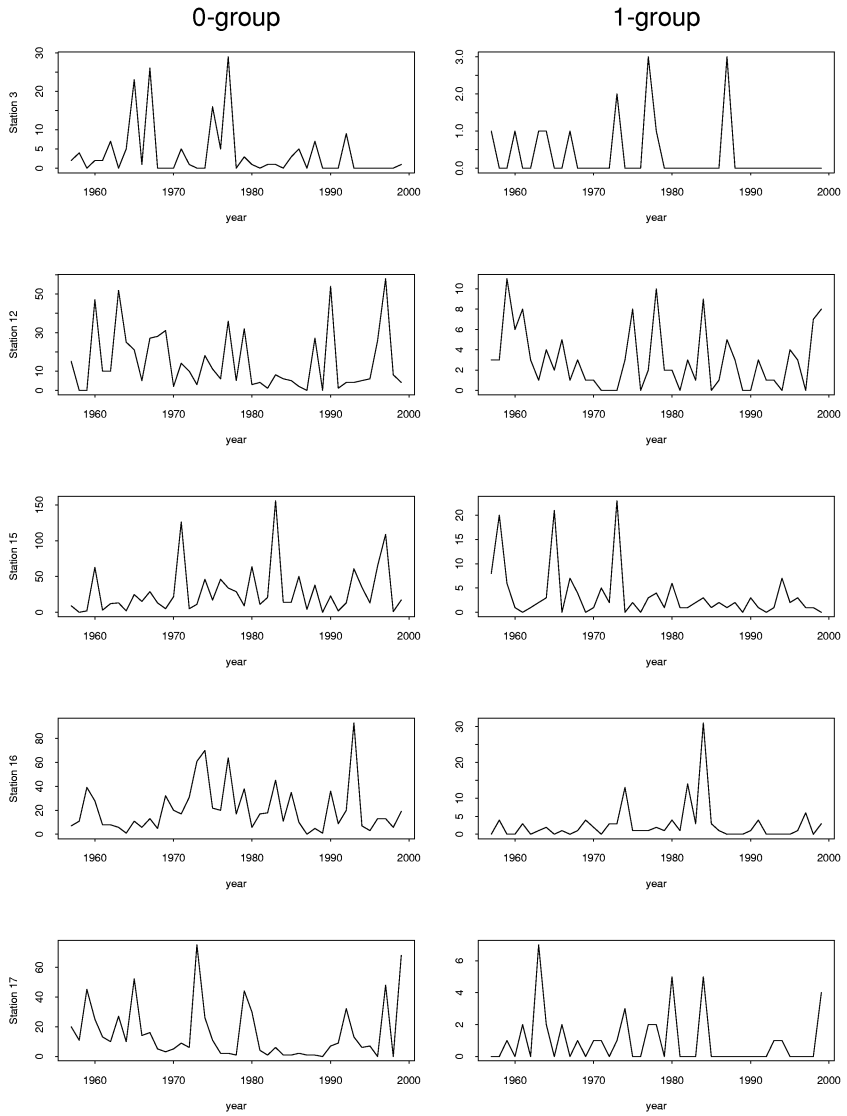

Figure A.1: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

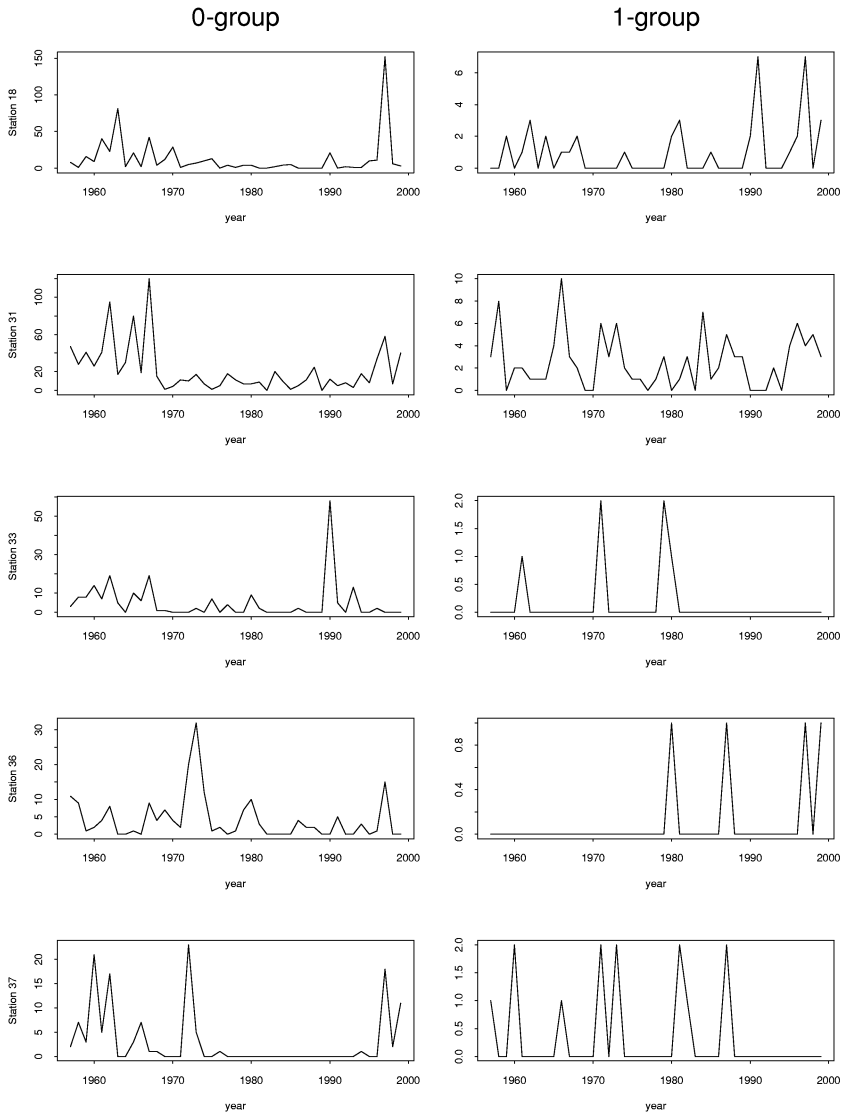

Figure A.2: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

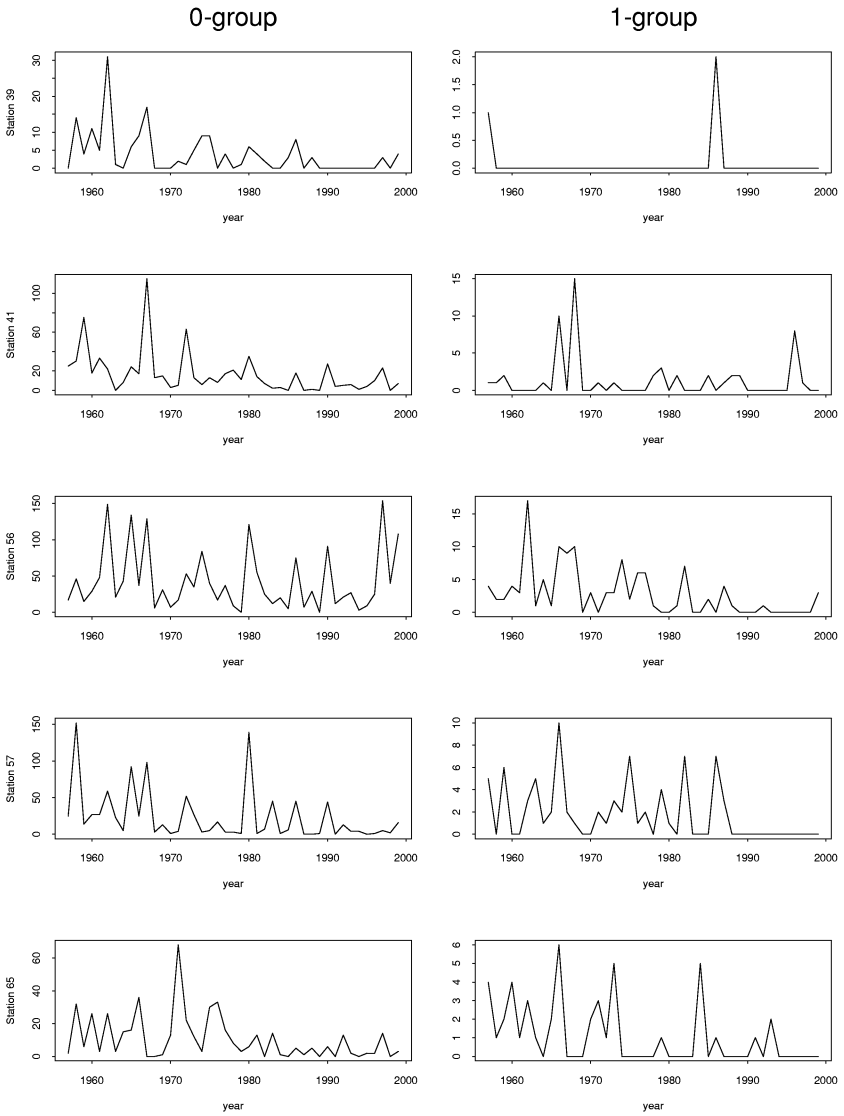

Figure A.3: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

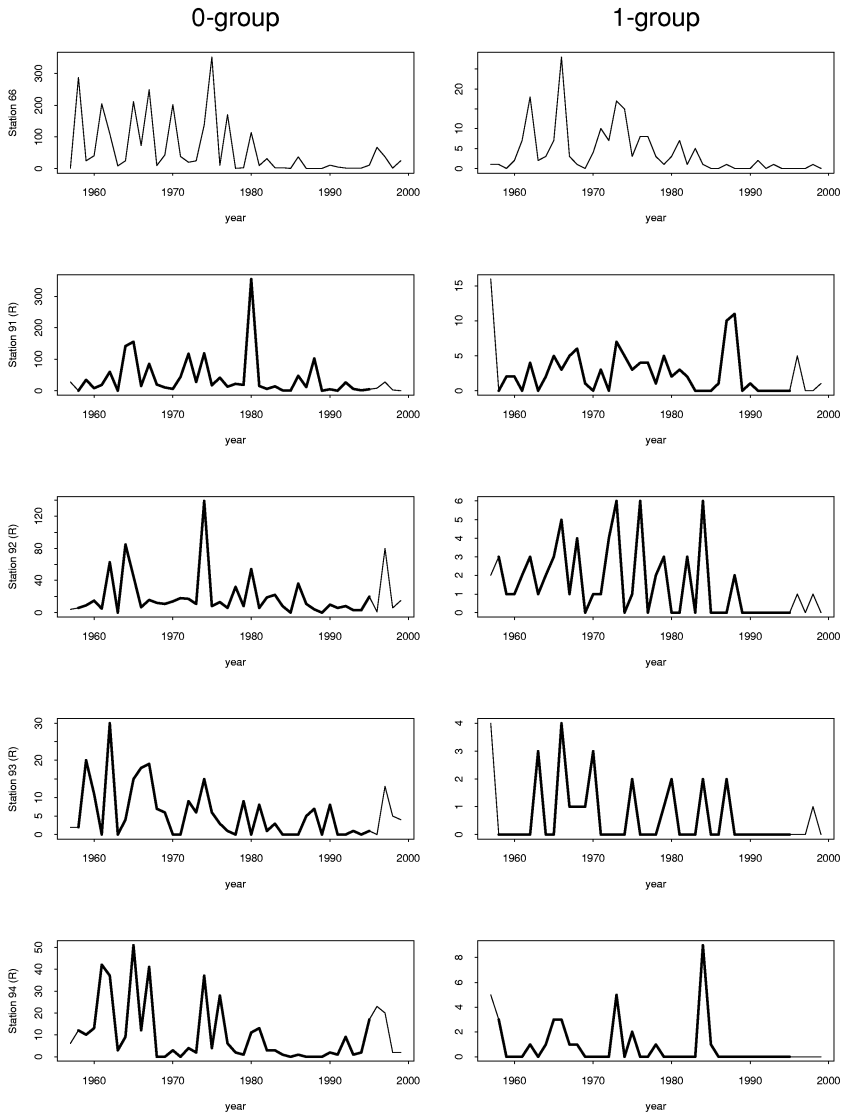

Figure A.4: Observed station specific counts of 0-group cod (left) and 1-group cod (right). The stations in the Risør area are marked with an (R), and the subset of observations for the period 1957-1994 are enhanced.

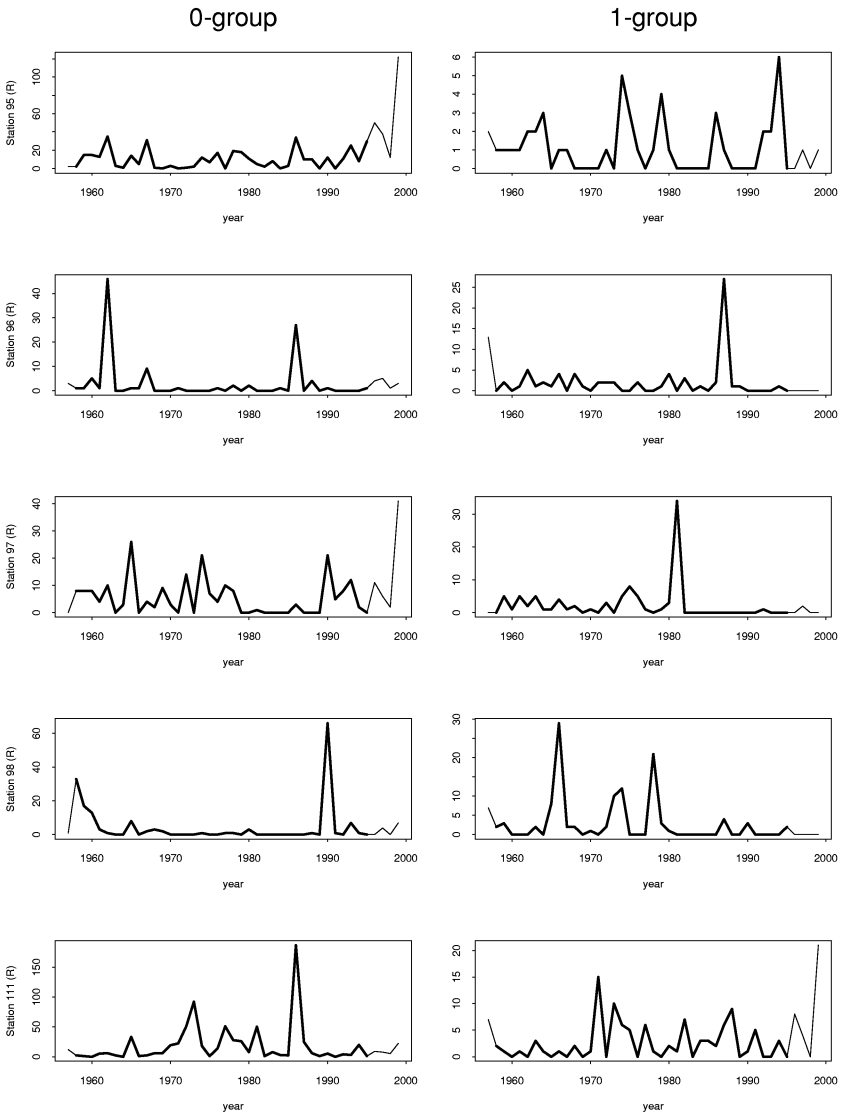

Figure A.5: Observed station specific counts of 0-group cod (left) and 1-group cod (right). The stations in the Risør area are marked with an (R), and the subset of observations for the period 1957-1994 are enhanced.

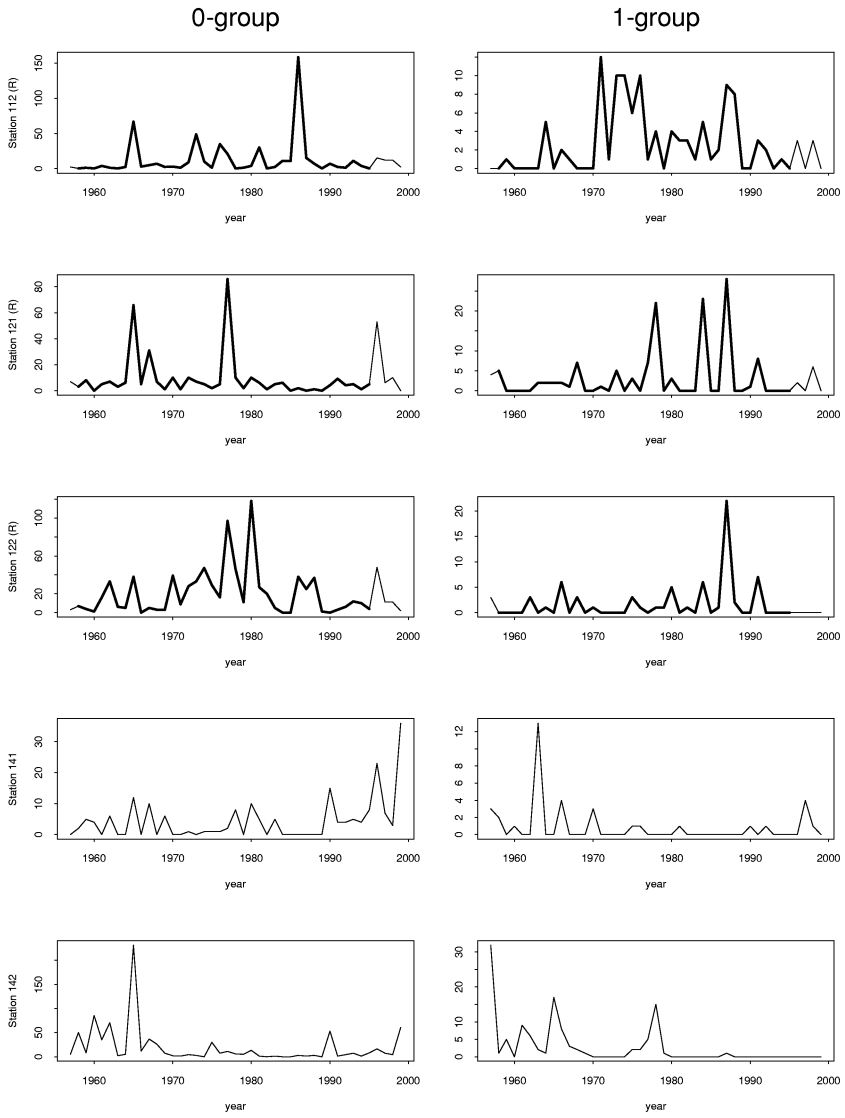

Figure A.6: Observed station specific counts of 0-group cod (left) and 1-group cod (right). The stations in the Risør area are marked with an (R), and the subset of observations for the period 1957-1994 are enhanced.

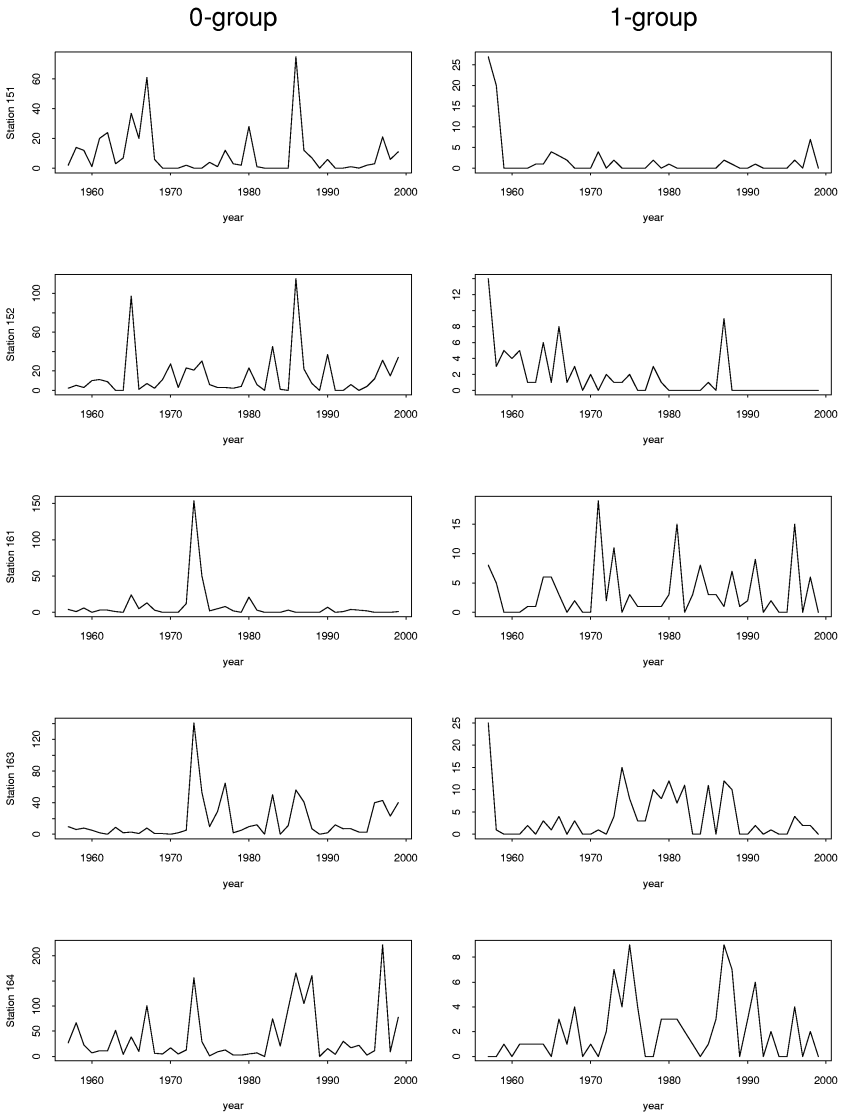

Figure A.7: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

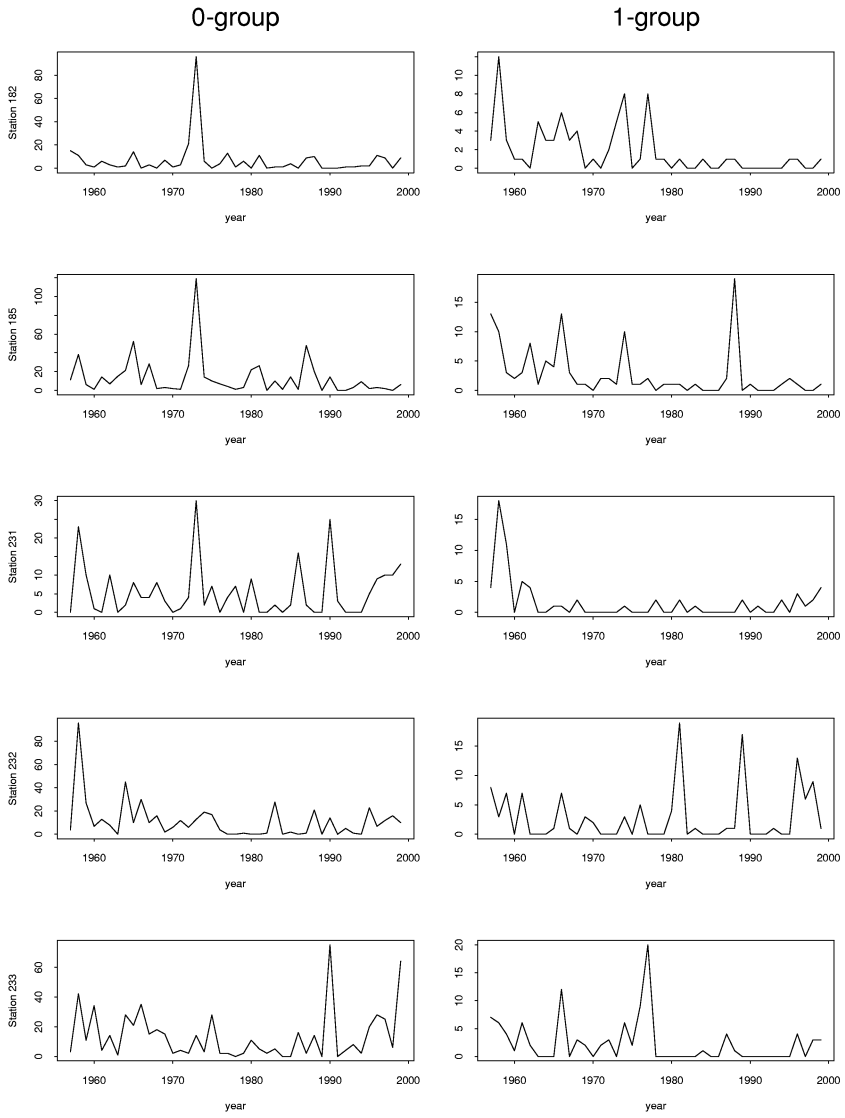

Figure A.8: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

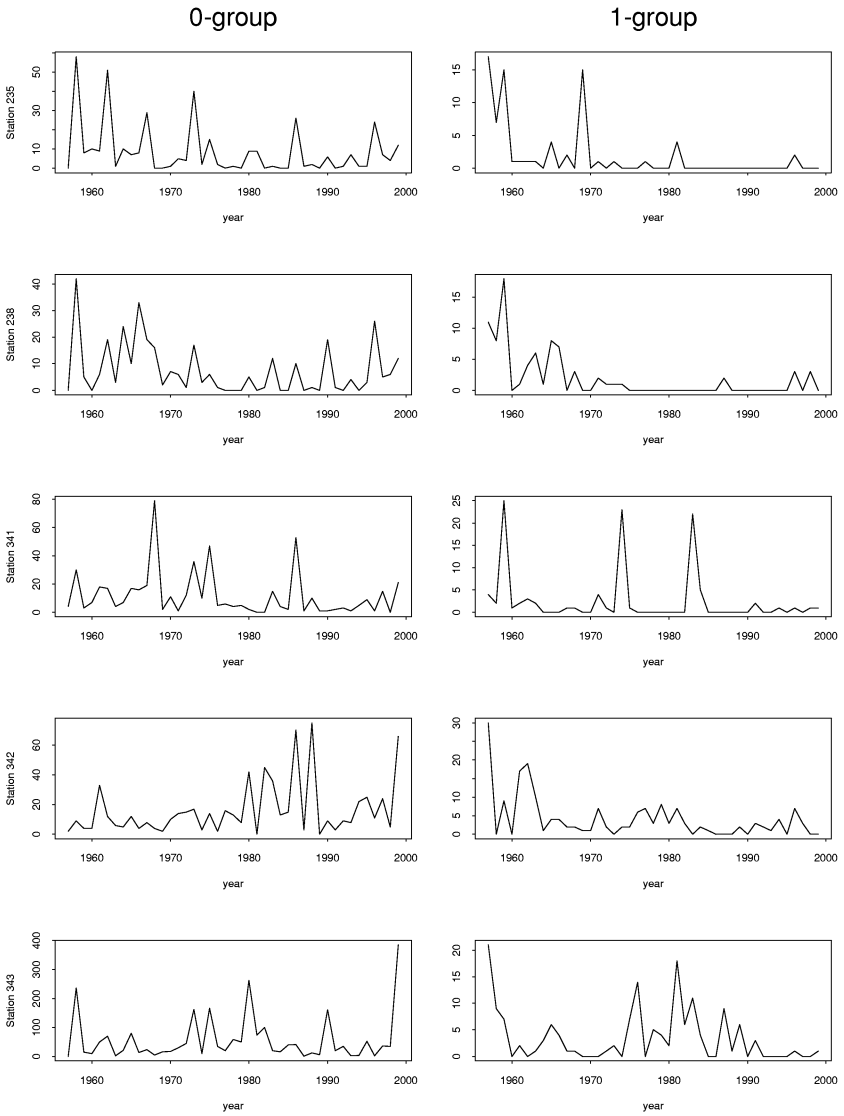

Figure A.9: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

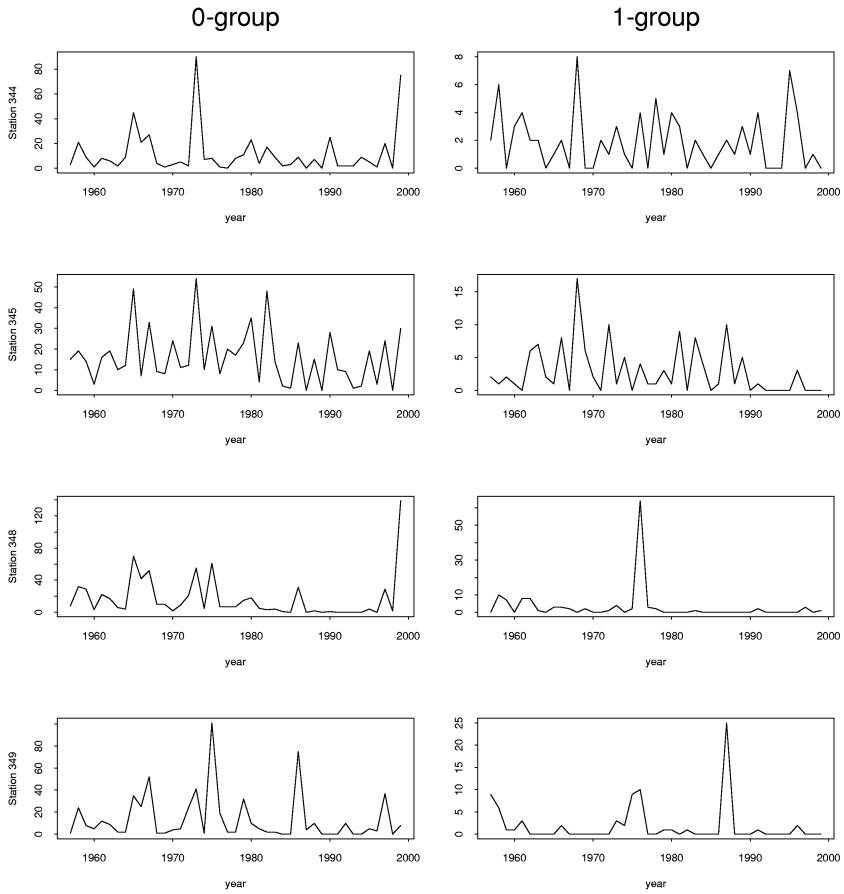

Figure A.10: Observed station specific counts of 0-group cod (left) and 1-group cod (right).

Appendix B

Results from the sensitivity analysis

ψ_x	ψ_y	$\hat{\beta}$	$\widehat{\mathrm{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\mathrm{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\mathrm{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\mathrm{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\mathrm{sd}(\tau_\alpha)}$
0.30	0.050	0.109	0.034	0.071	0.036	0.381	0.301	2.39	0.11	0.92	0.23
0.30	0.075	0.129	0.033	0.074	0.035	0.485	0.389	2.78	0.07	0.96	0.22
0.30	0.100	0.143	0.038	0.075	0.035	0.591	0.451	3.07	0.08	0.97	0.23
0.30	0.125	0.154	0.035	0.084	0.039	0.572	0.484	3.22	0.07	0.93	0.22
0.30	0.150	0.167	0.039	0.085	0.042	0.583	0.475	3.39	0.07	0.93	0.23
0.35	0.050	0.105	0.026	0.060	0.025	0.397	0.322	2.25	0.09	0.95	0.23
0.35	0.075	0.124	0.032	0.074	0.034	0.446	0.355	2.68	0.07	0.97	0.23
0.35	0.100	0.140	0.032	0.078	0.030	0.449	0.384	2.90	0.08	0.95	0.23
0.35	0.125	0.144	0.038	0.079	0.037	0.631	0.489	3.16	0.05	0.99	0.23
0.35	0.150	0.155	0.034	0.094	0.040	0.551	0.415	3.30	0.09	0.97	0.23
0.40	0.050	0.102	0.027	0.058	0.031	0.348	0.267	2.12	0.07	0.96	0.23
0.40	0.075	0.114	0.035	0.078	0.040	0.439	0.374	2.57	0.08	0.97	0.23
0.40	0.100	0.128	0.034	0.081	0.040	0.460	0.337	2.78	0.07	0.96	0.23
0.40	0.125	0.145	0.031	0.076	0.030	0.550	0.417	3.04	0.08	0.98	0.23
0.40	0.150	0.157	0.035	0.076	0.036	0.564	0.419	3.17	0.07	0.98	0.23
0.45	0.050	0.088	0.027	0.063	0.028	0.307	0.249	2.01	0.09	0.98	0.23
0.45	0.075	0.121	0.031	0.055	0.028	0.396	0.341	2.37	0.08	0.97	0.23
0.45	0.100	0.129	0.032	0.073	0.034	0.424	0.353	2.58	0.07	0.91	0.22
0.45	0.125	0.136	0.034	0.084	0.037	0.511	0.406	2.92	0.07	0.97	0.23
0.45	0.150	0.146	0.035	0.079	0.037	0.566	0.447	2.96	0.11	0.89	0.22
0.50	0.050	0.088	0.022	0.056	0.025	0.296	0.239	1.87	0.08	0.98	0.23
0.50	0.075	0.111	0.032	0.059	0.029	0.400	0.349	2.27	0.09	0.97	0.23
0.50	0.100	0.125	0.037	0.072	0.037	0.426	0.370	2.59	0.09	0.97	0.24
0.50	0.125	0.141	0.036	0.072	0.035	0.495	0.403	2.83	0.08	0.97	0.23
0.50	0.150	0.138	0.036	0.079	0.035	0.602	0.493	2.96	0.07	0.97	0.23

Table B.1: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 0.8$, using the simulated data set.

ψ_x	ψ_y	$\hat{\beta}$	$\widehat{\text{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\text{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\text{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\text{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\text{sd}(\tau_\alpha)}$
0.30	0.050	0.108	0.029	0.069	0.031	0.359	0.327	2.76	0.06	0.94	0.22
0.30	0.075	0.134	0.032	0.065	0.031	0.484	0.374	3.13	0.06	0.93	0.22
0.30	0.100	0.141	0.041	0.080	0.038	0.545	0.433	3.42	0.06	0.94	0.22
0.30	0.125	0.157	0.033	0.088	0.038	0.532	0.410	3.66	0.09	0.93	0.22
0.30	0.150	0.147	0.045	0.093	0.039	0.801	0.602	3.83	0.08	0.93	0.22
0.35	0.050	0.097	0.032	0.068	0.033	0.382	0.320	2.60	0.10	0.93	0.22
0.35	0.075	0.128	0.032	0.069	0.031	0.426	0.393	3.03	0.06	0.94	0.22
0.35	0.100	0.137	0.037	0.071	0.036	0.581	0.461	3.33	0.07	0.95	0.22
0.35	0.125	0.146	0.038	0.086	0.040	0.520	0.428	3.47	0.10	0.93	0.22
0.35	0.150	0.166	0.031	0.078	0.032	0.539	0.394	3.65	0.07	0.94	0.22
0.40	0.050	0.092	0.029	0.066	0.033	0.371	0.320	2.48	0.08	0.92	0.22
0.40	0.075	0.115	0.037	0.074	0.039	0.394	0.336	2.82	0.06	0.91	0.22
0.40	0.100	0.136	0.033	0.073	0.036	0.484	0.389	3.16	0.10	0.92	0.22
0.40	0.125	0.145	0.034	0.083	0.035	0.493	0.399	3.39	0.07	0.95	0.22
0.40	0.150	0.157	0.035	0.078	0.036	0.561	0.431	3.55	0.06	0.95	0.22
0.45	0.050	0.094	0.026	0.059	0.028	0.299	0.252	2.28	0.09	0.84	0.21
0.45	0.075	0.116	0.031	0.065	0.033	0.352	0.314	2.71	0.08	0.93	0.23
0.45	0.100	0.129	0.036	0.074	0.037	0.421	0.350	3.04	0.08	0.95	0.22
0.45	0.125	0.141	0.036	0.074	0.037	0.466	0.392	3.23	0.07	0.95	0.23
0.45	0.150	0.151	0.036	0.081	0.039	0.496	0.399	3.42	0.07	0.94	0.22
0.50	0.050	0.092	0.023	0.054	0.025	0.233	0.184	2.21	0.09	0.91	0.22
0.50	0.075	0.118	0.029	0.064	0.033	0.304	0.262	2.69	0.07	0.95	0.23
0.50	0.100	0.134	0.029	0.066	0.030	0.364	0.301	2.91	0.10	0.88	0.23
0.50	0.125	0.130	0.033	0.082	0.034	0.490	0.411	3.15	0.08	0.93	0.22
0.50	0.150	0.146	0.034	0.082	0.034	0.467	0.388	3.31	0.09	0.93	0.22

Table B.2: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 1.0$, using the simulated data set.

ψ_x	ψ_y	$\widehat{\beta}$	$\widehat{\text{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\text{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\text{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\text{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\text{sd}(\tau_\alpha)}$
0.30	0.050	0.107	0.031	0.065	0.031	0.435	0.355	3.08	0.07	0.90	0.21
0.30	0.075	0.136	0.034	0.069	0.033	0.412	0.342	3.49	0.06	0.92	0.22
0.30	0.100	0.144	0.032	0.081	0.035	0.503	0.390	3.76	0.07	0.92	0.22
0.30	0.125	0.158	0.039	0.076	0.036	0.602	0.496	4.00	0.06	0.93	0.22
0.30	0.150	0.160	0.040	0.087	0.039	0.656	0.484	4.12	0.09	0.89	0.22
0.35	0.050	0.110	0.027	0.057	0.028	0.329	0.292	2.89	0.07	0.87	0.22
0.35	0.075	0.120	0.033	0.081	0.036	0.382	0.321	3.28	0.08	0.86	0.21
0.35	0.100	0.150	0.031	0.069	0.032	0.415	0.337	3.61	0.08	0.92	0.21
0.35	0.125	0.145	0.038	0.088	0.041	0.540	0.431	3.81	0.10	0.89	0.22
0.35	0.150	0.162	0.038	0.082	0.036	0.573	0.455	4.00	0.08	0.91	0.22
0.40	0.050	0.096	0.030	0.058	0.029	0.363	0.309	2.79	0.06	0.90	0.21
0.40	0.075	0.129	0.031	0.058	0.031	0.394	0.334	3.19	0.07	0.89	0.21
0.40	0.100	0.145	0.030	0.063	0.032	0.426	0.354	3.46	0.08	0.91	0.22
0.40	0.125	0.150	0.035	0.075	0.034	0.495	0.404	3.70	0.07	0.92	0.22
0.40	0.150	0.160	0.035	0.075	0.033	0.561	0.453	3.88	0.06	0.92	0.22
0.45	0.050	0.086	0.028	0.068	0.031	0.307	0.259	2.69	0.07	0.90	0.21
0.45	0.075	0.115	0.032	0.066	0.034	0.404	0.346	3.05	0.10	0.85	0.22
0.45	0.100	0.132	0.031	0.068	0.031	0.448	0.354	3.31	0.12	0.86	0.22
0.45	0.125	0.152	0.031	0.067	0.033	0.434	0.357	3.61	0.08	0.93	0.22
0.45	0.150	0.148	0.038	0.085	0.038	0.511	0.430	3.73	0.10	0.88	0.23
0.50	0.050	0.096	0.022	0.050	0.024	0.253	0.218	2.56	0.08	0.89	0.21
0.50	0.075	0.114	0.027	0.066	0.029	0.327	0.284	3.00	0.07	0.91	0.22
0.50	0.100	0.133	0.033	0.063	0.034	0.371	0.322	3.24	0.09	0.91	0.22
0.50	0.125	0.137	0.034	0.077	0.038	0.452	0.356	3.49	0.07	0.92	0.22
0.50	0.150	0.150	0.032	0.079	0.037	0.456	0.375	3.64	0.09	0.89	0.23

Table B.3: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 1.2$, using the simulated data set.

ψ_x	ψ_y	$\hat{\beta}$	$\widehat{\text{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\text{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\text{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\text{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\text{sd}(\tau_\alpha)}$
0.30	0.050	0.011	0.006	0.011	0.005	0.060	0.055	0.26	0.10	1.14	0.31
0.30	0.075	0.025	0.008	0.015	0.009	0.082	0.075	0.52	0.09	1.14	0.30
0.30	0.100	0.045	0.010	0.017	0.010	0.094	0.087	0.76	0.11	1.06	0.29
0.30	0.125	0.058	0.013	0.019	0.011	0.124	0.119	0.96	0.07	1.06	0.28
0.30	0.150	0.073	0.012	0.019	0.011	0.112	0.103	1.10	0.07	1.03	0.27
0.35	0.050	0.008	0.006	0.008	0.006	0.051	0.063	0.20	0.13	1.21	0.34
0.35	0.075	0.020	0.008	0.012	0.008	0.078	0.076	0.43	0.09	1.20	0.33
0.35	0.100	0.036	0.009	0.015	0.009	0.080	0.071	0.60	0.08	1.11	0.30
0.35	0.125	0.049	0.012	0.018	0.011	0.107	0.098	0.82	0.09	1.09	0.29
0.35	0.150	0.061	0.014	0.020	0.012	0.120	0.118	0.98	0.08	1.05	0.28
0.40	0.050	0.005	0.004	0.006	0.005	0.029	0.029	0.09	0.08	1.21	0.34
0.40	0.075	0.013	0.005	0.011	0.007	0.059	0.054	0.29	0.12	1.16	0.33
0.40	0.100	0.028	0.009	0.013	0.009	0.079	0.077	0.50	0.10	1.16	0.32
0.40	0.125	0.040	0.012	0.018	0.012	0.091	0.084	0.71	0.09	1.12	0.30
0.40	0.150	0.052	0.011	0.018	0.011	0.090	0.083	0.86	0.07	1.14	0.30
0.45	0.050	0.004	0.004	0.005	0.004	0.022	0.026	0.10	0.08	1.24	0.36
0.45	0.075	0.010	0.005	0.011	0.007	0.051	0.048	0.25	0.12	1.20	0.34
0.45	0.100	0.021	0.007	0.011	0.006	0.066	0.063	0.40	0.08	1.20	0.33
0.45	0.125	0.034	0.011	0.016	0.010	0.080	0.079	0.57	0.10	1.13	0.31
0.45	0.150	0.046	0.011	0.017	0.011	0.092	0.092	0.78	0.08	1.15	0.31
0.50	0.050	0.003	0.002	0.003	0.003	0.018	0.020	0.08	0.07	1.34	0.39
0.50	0.075	0.008	0.004	0.008	0.005	0.043	0.039	0.13	0.10	1.17	0.34
0.50	0.100	0.015	0.006	0.011	0.006	0.055	0.047	0.20	0.13	1.07	0.31
0.50	0.125	0.028	0.009	0.014	0.008	0.078	0.072	0.49	0.09	1.17	0.33
0.50	0.150	0.037	0.011	0.017	0.010	0.097	0.097	0.68	0.08	1.16	0.32

Table B.4: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 0.8$, using the Risør data.

ψ_x	ψ_y	$\hat{\beta}$	$\widehat{\mathrm{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\mathrm{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\mathrm{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\mathrm{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\mathrm{sd}(\tau_\alpha)}$
0.30	0.050	0.012	0.006	0.010	0.006	0.058	0.055	0.58	0.14	1.02	0.27
0.30	0.075	0.029	0.008	0.012	0.008	0.072	0.076	0.88	0.09	1.04	0.27
0.30	0.100	0.047	0.011	0.016	0.010	0.082	0.073	1.13	0.06	1.03	0.26
0.30	0.125	0.062	0.011	0.017	0.011	0.096	0.088	1.31	0.06	0.99	0.26
0.30	0.150	0.070	0.013	0.021	0.011	0.126	0.120	1.49	0.05	1.00	0.25
0.35	0.050	0.006	0.005	0.005	0.004	0.027	0.030	0.44	0.11	1.06	0.28
0.35	0.075	0.019	0.006	0.010	0.006	0.060	0.053	0.70	0.08	1.06	0.28
0.35	0.100	0.035	0.010	0.017	0.009	0.097	0.085	1.00	0.07	1.05	0.28
0.35	0.125	0.049	0.012	0.017	0.011	0.101	0.091	1.17	0.07	1.01	0.26
0.35	0.150	0.061	0.013	0.020	0.012	0.110	0.099	1.35	0.07	1.02	0.26
0.40	0.050	0.005	0.004	0.007	0.006	0.038	0.037	0.49	0.12	1.12	0.31
0.40	0.075	0.013	0.006	0.009	0.006	0.053	0.046	0.62	0.11	1.07	0.29
0.40	0.100	0.026	0.009	0.014	0.008	0.073	0.065	0.82	0.08	1.03	0.27
0.40	0.125	0.041	0.010	0.016	0.009	0.094	0.088	1.01	0.07	0.99	0.26
0.40	0.150	0.053	0.013	0.019	0.011	0.109	0.109	1.24	0.07	1.03	0.27
0.45	0.050	0.004	0.002	0.004	0.003	0.025	0.027	0.38	0.11	1.08	0.30
0.45	0.075	0.010	0.005	0.008	0.004	0.042	0.036	0.54	0.08	1.08	0.30
0.45	0.100	0.023	0.009	0.013	0.008	0.070	0.064	0.81	0.08	1.09	0.29
0.45	0.125	0.036	0.009	0.012	0.007	0.079	0.076	0.93	0.06	1.07	0.28
0.45	0.150	0.046	0.012	0.016	0.010	0.090	0.086	1.04	0.08	0.97	0.26
0.50	0.050	0.003	0.002	0.003	0.003	0.016	0.022	0.34	0.10	1.08	0.30
0.50	0.075	0.006	0.004	0.006	0.004	0.035	0.033	0.49	0.09	1.13	0.32
0.50	0.100	0.017	0.007	0.011	0.006	0.062	0.061	0.64	0.11	1.02	0.28
0.50	0.125	0.029	0.009	0.014	0.008	0.077	0.069	0.85	0.09	1.04	0.29
0.50	0.150	0.039	0.011	0.016	0.010	0.086	0.082	1.03	0.08	1.08	0.29

Table B.5: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 1.0$, using the Risør data.

ψ_x	ψ_y	$\widehat{\beta}$	$\widehat{\text{sd}(\beta)}$	$\hat{\gamma}$	$\widehat{\text{sd}(\gamma)}$	$\hat{\kappa}$	$\widehat{\text{sd}(\kappa)}$	$\hat{\alpha}$	$\widehat{\text{sd}(\alpha)}$	$\hat{\tau}_\alpha$	$\widehat{\text{sd}(\tau_\alpha)}$
0.30	0.050	0.010	0.005	0.008	0.006	0.045	0.040	0.86	0.08	0.97	0.25
0.30	0.075	0.029	0.008	0.012	0.008	0.079	0.079	1.18	0.07	0.97	0.25
0.30	0.100	0.044	0.011	0.016	0.010	0.096	0.098	1.44	0.07	0.98	0.24
0.30	0.125	0.056	0.012	0.020	0.011	0.120	0.106	1.62	0.07	0.94	0.24
0.30	0.150	0.072	0.013	0.019	0.012	0.112	0.109	1.81	0.06	0.96	0.23
0.35	0.050	0.006	0.003	0.007	0.006	0.034	0.032	0.79	0.11	0.96	0.25
0.35	0.075	0.020	0.008	0.013	0.008	0.069	0.060	1.08	0.07	1.00	0.26
0.35	0.100	0.036	0.009	0.014	0.008	0.083	0.079	1.24	0.10	0.92	0.25
0.35	0.125	0.052	0.012	0.016	0.012	0.089	0.083	1.49	0.06	0.94	0.24
0.35	0.150	0.060	0.014	0.020	0.013	0.111	0.104	1.65	0.07	0.94	0.24
0.40	0.050	0.005	0.003	0.006	0.005	0.029	0.029	0.78	0.08	1.02	0.27
0.40	0.075	0.012	0.006	0.010	0.006	0.048	0.044	0.92	0.07	1.01	0.27
0.40	0.100	0.030	0.009	0.013	0.007	0.072	0.067	1.19	0.08	1.00	0.25
0.40	0.125	0.043	0.009	0.014	0.008	0.083	0.078	1.37	0.06	1.00	0.26
0.40	0.150	0.053	0.013	0.018	0.011	0.104	0.100	1.56	0.07	0.98	0.25
0.45	0.050	0.004	0.003	0.004	0.003	0.022	0.025	0.75	0.10	1.01	0.27
0.45	0.075	0.009	0.005	0.008	0.004	0.045	0.042	0.85	0.08	1.01	0.26
0.45	0.100	0.022	0.007	0.010	0.007	0.065	0.057	1.05	0.09	0.99	0.26
0.45	0.125	0.035	0.010	0.014	0.008	0.086	0.080	1.20	0.10	0.91	0.26
0.45	0.150	0.049	0.009	0.014	0.008	0.082	0.076	1.44	0.07	1.00	0.25
0.50	0.050	0.003	0.002	0.004	0.003	0.018	0.018	0.70	0.08	1.05	0.29
0.50	0.075	0.006	0.004	0.006	0.004	0.035	0.031	0.78	0.10	1.02	0.28
0.50	0.100	0.016	0.006	0.010	0.006	0.053	0.047	0.97	0.09	1.01	0.27
0.50	0.125	0.029	0.007	0.011	0.006	0.071	0.065	1.15	0.08	1.02	0.27
0.50	0.150	0.040	0.011	0.015	0.009	0.082	0.080	1.32	0.08	0.98	0.26

Table B.6: Sensitivity of parameters to different values of ψ_x and ψ_y for $\theta = 1.2$, using the Risør data.

Appendix C

Estimated abundances for all fjords

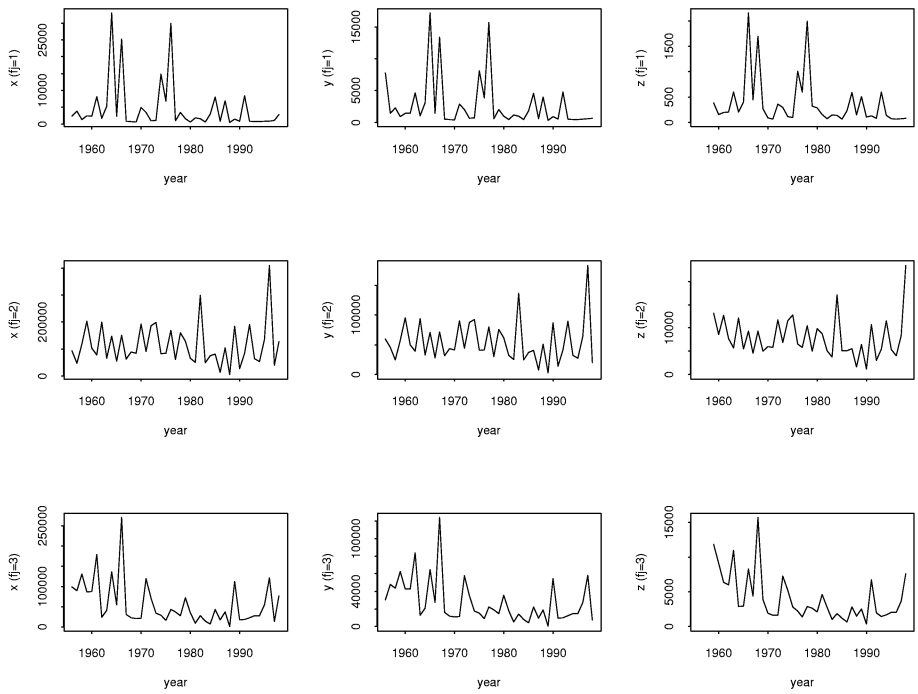

Figure C.1: Estimated abundances for the 0-group (left), 1-group (middle) and adult (right) cod based on data from 13 fjords. The first 3 years are discarded for the adult cod because of boundary effects.

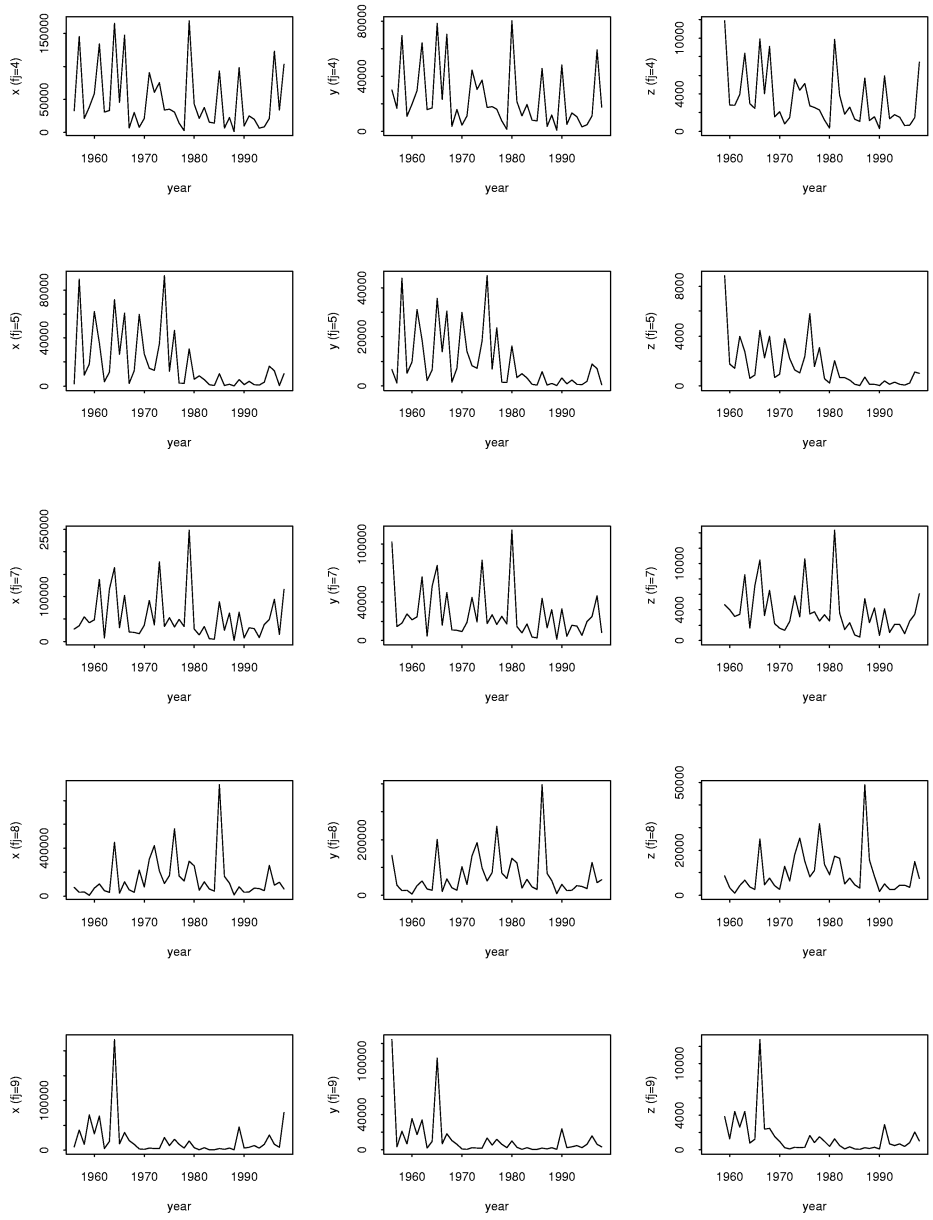

Figure C.2: Estimated abundances for the 0-group (left), 1-group (middle) and adult (right) cod based on data from 13 fjords. The first 3 years are discarded for the adult cod because of boundary effects.

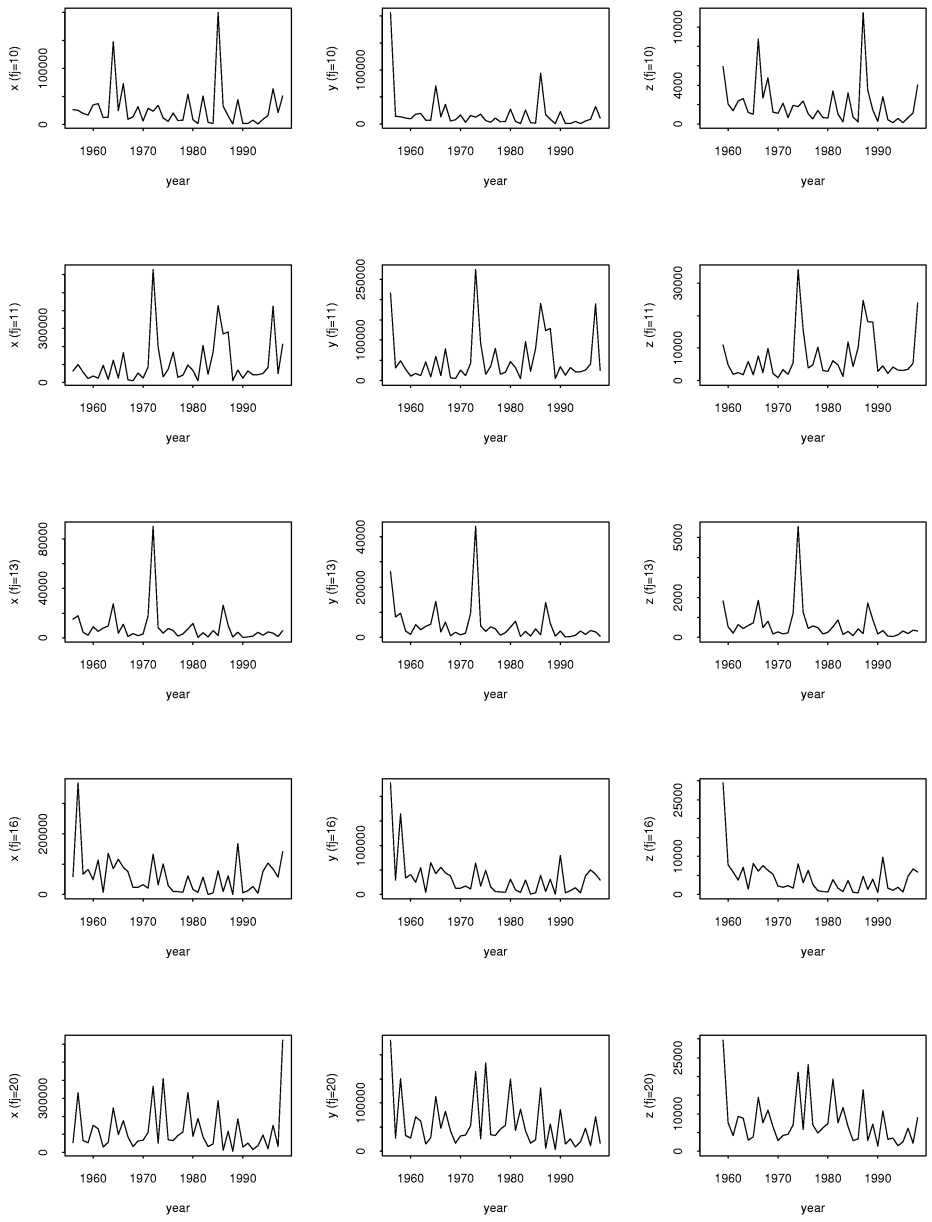

Figure C.3: Estimated abundances for the 0-group (left), 1-group (middle) and adult (right) cod based on data from 13 fjords. The first 3 years are discarded for the adult cod because of boundary effects.

Appendix D

Assessment of the variability of the discretisation for the marginal posterior sampling algorithm

Here, we describe how the variability of the estimates (8.18), used to specify the discretisation of the abundances for the marginal sampling algorithm, can be assessed. We draw 100 samples of data from the likelihood models (4.15) and (4.18) given the population abundances of the simulated data set, and the time series of up-scaled estimated abundances (8.18) are computed for each sample. In Figure D.1 the results are summarised. The left panels show the up-scaled abundances and the middle panels the true values for each of the three age groups. In the right three panels relative differences between estimated and true abundances are shown. The differences are smaller for the 0-group than for the 1-group cod. This is as expected, since the data are more abundant for the 0-group than the 1-group cod, for which the proportion of zeros in the station specific data is large (46%). For the adult cod, for which we have no data, the deviations are of about the same order of magnitude as for the 1-group. The large deviations of the first year for the adult cod are due to the uncertainty in the equilibrium estimate of Z_1, and could be discarded when choosing the discretisation.

The plots in Figure D.1 represent the variability of the sampled data generated from one realisation of population abundances from the population dynamics prior model. To assess the effect of the uncertainty of the simulated abundances, we repeat the procedure described above, but generate the sampled data sets from 100 separate realisation from the prior model. Similar results as presented in Figure D.1 for the one sample case are given in Figure D.2. The variability of the population estimates increases with abundance, which is in accordance with the prior model specification, and the pattern of the deviances between up-scaled and true values are as for the single population sample case.

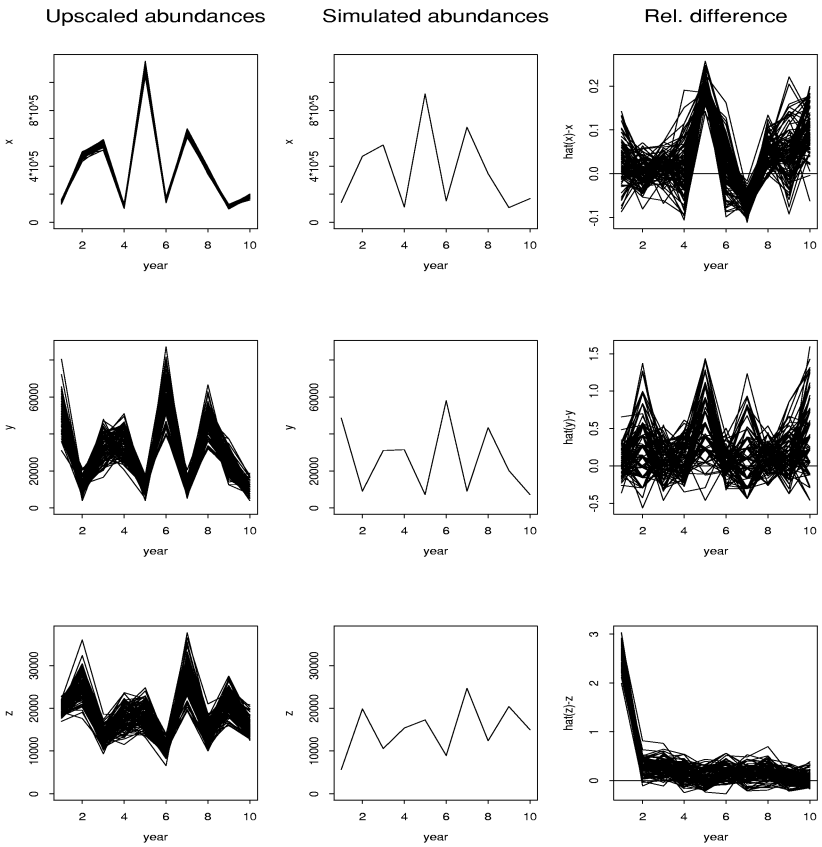

Figure D.1: Simulated abundances and abundances estimated by up-scaling. The estimated abundances are generated from repeated samples of data generated from the same realisation from the prior model.

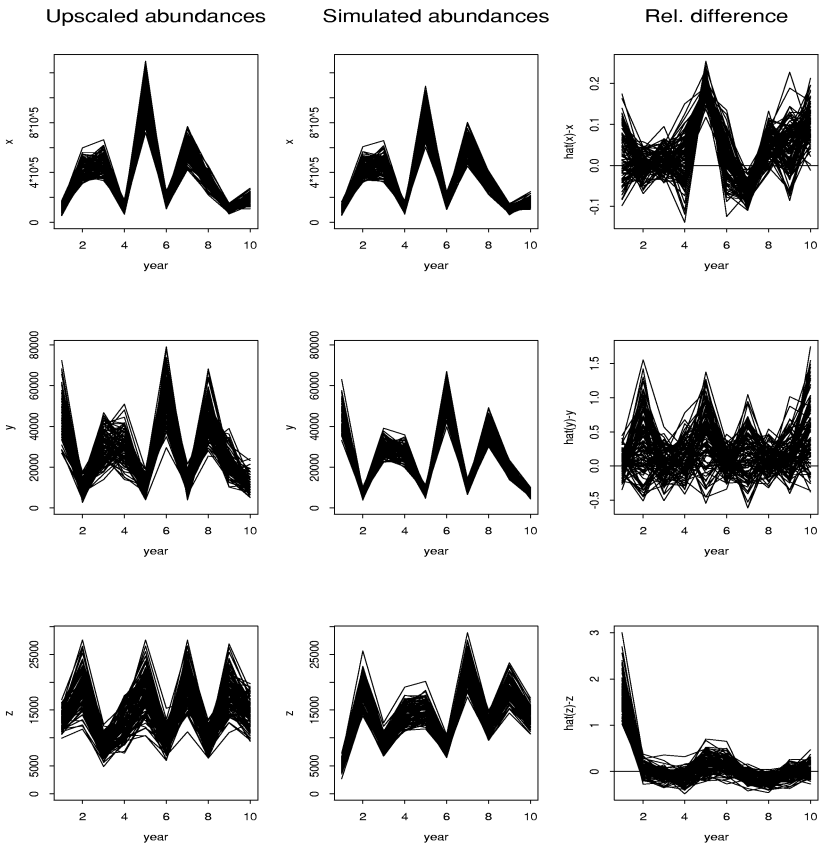

Figure D.2: Simulated abundances and abundances estimated by up-scaling. The estimated abundances are generated from samples of data generated from different realisations from the prior model.

URN:NBN:no-3402

Part II

Modelling spatial variation in disease risk using Gaussian Markov random field proxies for Gaussian random fields

Modelling spatial variation in disease risk using Gaussian Markov random field proxies for Gaussian random fields

Turid Follestad and Håvard Rue

Department of Mathematical Sciences
Norwegian University of Science and Technology

Abstract

Analyses of spatial variation in disease risk based on area-level summaries of disease counts are most often based on the assumption that the relative risk is uniform across each region. Such approaches introduce an artificial piecewise-constant relative risk-surface with discontinuities at regional boundaries. A more natural approach is to assume that the spatial variation in risk can be represented by an underlying smooth relative risk-surface over the area of interest. This approach was taken by Kelsall and Wakefield (2002), who used an underlying Gaussian random field (GRF) to derive a multivariate log-Normal distribution of the risk at the regional level. The derivation rely on the approximation $\sum_i \exp(x_i) \approx \exp(\sum_i x_i)$, which is frequently used in similar contexts in the geostatistics literature, but the different sizes and shapes of the regions typically found in disease mapping applications indicate that the validity of the approximation is questionable.

We propose an approach to the modelling of a smoothly varying risk surface based on aggregated data avoiding this approximation. We also derive computationally efficient block MCMC-algorithms using a re-formulation of the geostatistical GRF model using Gaussian Markov random fields (GMRFs). We make extensive use of recent developments for GMRFs, including a method for fitting GMRFs to Gaussian random fields and computationally efficient algorithms for GMRFs based on numerical methods for sparse matrices. We demonstrate our approach on simulated data as well as a set of German oral cavity cancer mortality data from the period 1986–90, which have been previously analysed in the literature.

URN:NBN:no-3402

1 Introduction

Disease maps displaying the geographical variability of disease incidence or mortality rates across a region of interest, are valuable tools in spatial epidemiology. By studying a disease map, regions with particularly high or low rates can be identified, and this information can be used as input to ethological studies as a guideline in defining and validating hypotheses about a disease. For an overview of the history of disease mapping, see e.g. Walter (2000). Disease incidence or mortality data can be available as point data for which the exact location of each case is known, or more commonly as aggregated or areal summary data, often due to confidentiality reasons. For rare and non-infectious diseases, the aggregated incidence or mortality counts $y_i;\ i = 1, \ldots, m$ in a set of m regions are commonly assumed to be conditionally independent given the stratum-weighted relative risks R_i of the regions, and to follow Poisson distributions with mean given by $E_i R_i$. The value E_i represents the expected number of cases in region i, typically given as a population-weighted sum of stratum-specific probabilities of disease, computed from the data assuming uniform risk across the study area. The maximum likelihood estimate of the relative risk in region i is the standardised mortality (or incidence) ratio SMR y_i/E_i. Figure 1 shows the observed aggregated counts and SMR for a set of data on mortality from oral cavity cancer in Germany, that will be analysed in Section 6 (Knorr-Held and Raßer, 2000). From Figure 1 we observe that there is a tendency toward high risk in the north-east and in the south-west and low risk in the east. However, for small populations at risk and for rare diseases, the SMR as an esti-

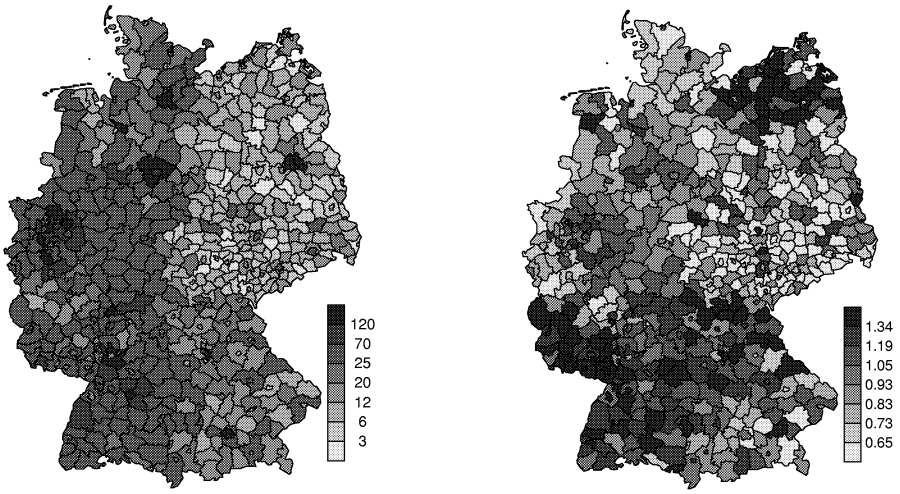

Figure 1: The observed counts (left) and the standardised mortality ratio (SMR) (right) for the German oral cavity cancer data.

mator of the relative risk can be highly variable. It can give rise to spurious estimates of high risk in regions with low populations, masking the true spatial pattern of the risk over the area of interest. Therefore, conclusions drawn from maps of the SMR can be misleading. To overcome this problem, a number of authors have developed statistical approaches to improve on raw estimates of disease risk. Reviews of statistical methods for mapping disease risk are provided by e.g. Diggle (2000), Wakefield, Best and Waller (2000) and Mollié (1996), the latter two focusing on Bayesian approaches.

Taking a Bayesian approach, the risk estimates of sparsely populated or low frequency regions are smoothed toward an overall prior mean. Since it is often the case that the relative risk tend to be similar in neighbouring regions, disease maps can also be improved on by adding spatial correlation to the prior model. The estimates of the risk in each region can then "borrow strength" from neighbouring regions. This can be accommodated by including a spatially structured component within a random effects model for the disease risk, an approach first taken by Clayton and Kaldor (1987). A commonly used approach, proposed by Besag, York and Mollié (1991), is to model the log relative risk as

$$\log(R_i) = \boldsymbol{\beta}^T \boldsymbol{z}_i + u_i + v_i, \tag{1}$$

where $\boldsymbol{z}_i$ is a vector of covariates, including an intercept term, and u_i and v_i are spatially structured and unstructured random effects, respectively. The spatially structured random effect is assigned a Gaussian Markov random field prior, such that $f(u_i|\boldsymbol{u}_{-i}) = f(u_i \mid \boldsymbol{u}_{\delta i})$, where $\boldsymbol{u}_{-i}$ denotes all elements of the vector $\boldsymbol{u}$ except element i, and $\delta(i)$ is the set of neighbouring regions of region i. To specify the Markov random field prior, we need to define which regions are neighbours. The level of aggregation of areal summary data is often defined by administratively specified regions, and therefore alternative definitions to the square neighbourhoods often used in the case of lattice data are needed. An approach taken by many authors, e.g. Clayton and Kaldor (1987), Bernardinelli, Pascutto, Best and Gilks (1997), Knorr-Held and Besag (1998) and Waller, Carlin, Xia and Gelfand (1997), is to define two regions as neighbours if they share a common boundary. This will work well if the regions do not differ much in size and shape, but this is often not the case. An alternative to the adjacency approach is to specify the joint distribution of the heterogeneity effects u_i, defining spatial structure of the covariance matrix as a function of differences between the region centres (e.g. Wakefield and Morris, 2001; Wakefield and Morris, 1999). However, similar objections apply to this method as to the adjacency based methods, as the size and shape of the regions are still not taken into account, and in both cases the inference will depend on the level of aggregation of the data.

The method of Besag et al. (1991) does not naturally allow for discontinuities in the spatial structure of the risk. In a recent paper, Fernández and Green (2002) present an alternative approach, developing a spatially structured mixture model where GMRF priors are specified for the weights in the mixture. Using a mixture of Poisson distributions, the method is applied in a disease mapping context, and it is illustrated how

the approach represents an improvement over the method of Besag et al. (1991) in cases where the spatial pattern has step-like discontinuities. The approach is related to that of Knorr-Held and Raßer (2000) identifying clusters of constant risk.

In general, spatial heterogeneity of the disease risk will be a confounder for unmeasured spatially structured factors influencing the disease risk. In most cases, there is no reason that these risk factors are region specific and discontinuous at region boundaries. Thus, the relative risk is not expected to be constant within regions and disjoint across regions. On the contrary, it seems reasonable to believe that the underlying risk surface is varying continuously over the region of study. In cases where the observations can be regarded as point data, a smoothly varying risk surface and the corresponding hyper-parameters can be estimated using extensions of classical geostatistical or point process approaches. Using data for which the exact locations are known, Diggle, Tawn and Moyeed (1998) propose a model-based geostatistical approach embedding the classical linear geostatistical methods for Gaussian data within a framework analogous to the generalised linear models (McCullagh and Nelder, 1989) for mutually independent data. Consequently, they allow for data for which the stochastic variation is assumed to be non-Gaussian. Another approach is taken by Best, Ickstadt and Wolpert (2000), who specify a Poisson-Gamma random field model for the disease risk. The approach is based on the methodological framework presented in Wolpert and Ickstadt (1998) and extended in Ickstadt and Wolpert (1999) to include location-specific covariates measured at different levels of spatial aggregation and individual attributes like age and gender. The point locations of individual cases and the corresponding attributes are regarded as a marked point process, and the spatial structure of the residual risk surface is represented by a kernel smoothed Gamma random field. The risk surface can be estimated at any level of spatial aggregation.

When the disease incidence or mortality data are only available as aggregated counts, the approaches to risk-surface estimation described above are not directly applicable. Kelsall and Wakefield (2002) propose a geostatistical approach to modelling the joint distribution of the area-level relative risks in such situations. They specify a model for an underlying continuously varying risk surface $R(s);\ s \in \mathcal{A}$, assuming the log risk surface $S(s) = \log(R(s))$ to be a realisation of a Gaussian random field (GRF). Based on this GRF model, area-level relative risks R_i in a set of regions $\mathcal{A}_i;\ i = 1, \ldots, m$, forming a partition of the study region $\mathcal{A}$, are defined by

$$R_i = \int_{\mathcal{A}_i} R(s) f_i(s) ds, \tag{2}$$

where $f_i(s);\ i = 1, \ldots, m$ are weight functions depending on the stratum-specific population density distribution in region $\mathcal{A}_i$. Conditionally on these relative risks, the data are assumed to be independent realisations from a Poisson distribution with mean $R_i E_i$. To allow for computational feasibility, they approximate the joint distribution of the region-level risks $R_i;\ i = 1, \ldots, m$ by a multivariate log-Normal distribution with moments that are derived from the moments of the Normal distribution of $S(s)$, using

numerical methods to evaluate the integrals involved. The approximation is essentially equivalent to approximating the distribution of $S_i = \log(R_i)$ by the distribution of $\int_{\mathcal{A}_i} \log R(s) f_i(s) ds$. As pointed out by the authors, the approximation is best when the size of the regions are relatively small and the regions are of about the same shape, and the log-Gaussian assumption is exact only in the limit when the regions are of the same shape and size, and the size tends to zero. The parameters of the model, including the log-risk surface at a set of locations s_k, are estimated by Markov chain Monte Carlo methods (Gilks, Richardson and Spiegelhalter, 1996), using Gibbs sampling in combination with adaptive rejection sampling.

We propose an alternative approach to the estimation of a smooth risk surface based on aggregated count data, representing the Gaussian random field defining the prior for the log-risk surface by a Gaussian Markov random field defined on a lattice. The basis of the model formulation is as in Kelsall and Wakefield (2002), but while they use the geostatistical model to derive an approximation to the joint distribution of the regional-level log-risk, and base the inference on the resulting regional-level stochastic model, we avoid the approximation by working directly on a lattice representation of the model. We replace the integral expression (2) for the regional level relative risk R_i by a sum over the exponentiated values of the GMRF for the lattice nodes falling within $\mathcal{A}_i$. Due to the conditional independence structure of the GMRF, using a GMRF proxy to the GRF allows for the use of computationally efficient algorithms for sampling based inference. However, the spatial structure is often intuitively easier to specify and interpret using a geostatistical GRF formulation than the conditional formulation represented by the GMRF. Therefore, we specify the spatial structure of the random field in terms of the correlation function for the GRF, using the procedure in Rue and Tjelmeland (2002) to fit the GMRF to the GRF. Thus, our approach relies on the assumptions that the smooth relative risk surface can be represented on a lattice and that the GRF as defined on this lattice can be well estimated by a GMRF.

Drawing on the routines for fast and exact simulation of GMRF implemented in Rue and Follestad (2002), we develop an efficient block-sampling algorithm for estimating the log-risk surface and the parameters of the model. For each block, the elements of the lattice based log-risk surface are updated using a Metropolis-Hastings step, generating a proposal from a Gaussian approximation to the full conditional distribution. This can be done efficiently after re-formulating the problem of sampling from the proposal distribution to a computationally convenient conditional sampling problem.

The report is organised as follows. In Section 2 we present the statistical model, and an overview of our approach to estimating the log risk surface and the hyper-parameters is presented in Section 3. More details on the estimation algorithm are given in Section 4. In Section 5 we present results for a simulated data set, and results for the German oral cavity cancer data are given in Section 6. The method and the results are summarised and discussed in Section 7.

2 The statistical model

The statistical model is based on disease incidence or mortality data available as aggregated counts y_i in a set of m disjoint regions denoted $\mathcal{A}_i;\ i = 1, \ldots, m$, such that $\mathcal{A} = \cup_i \mathcal{A}_i$ is the overall region of study. Following the approaches of Kelsall and Wakefield (2002), Best et al. (2000) and Diggle et al. (1998) we assume that the geographical variation in the risk of the disease can be represented by a smoothly varying surface $R(s);\ s \in \mathcal{A}$. In this section we specify a lattice based model for the risk surface, first presenting a Gaussian random field model, and then a Gaussian Markov random field proxy to this model.

2.1 A Gaussian random field model on a lattice

The log-risk surface $S(s) = \log R(s)$ is assumed to be a realisation of a Gaussian random field. The basis of our modelling approach is as in Kelsall and Wakefield (2002), but we explicitly define the Gaussian random field model on a lattice overlaying the study region $\mathcal{A}$. Throughout our study, using simulated data as well as the real dataset, we use the 544 districts of Germany for which the German oral cavity data are defined as our region of interest. A map of the study region with an overlaying lattice consisting of $n_{GRF} = 16824$ nodes is given in Figure 2. For a better visual impression of the resolution of the lattice, see the top right panel of Figure 6. The number of lattice nodes within each region is in the range 1 to 136, with a median number of 29.

The multivariate Normal joint prior distribution of the log-risks $S(\boldsymbol{s}_j);\ j = 1, \ldots, n_{GRF}$ is given by the moments

$$\begin{aligned} \mathrm{E}(S(\boldsymbol{s}_j)) &= \mu_j, \\ \mathrm{Var}(S(\boldsymbol{s}_j)) &= \sigma^2, \\ \mathrm{Corr}(S(\boldsymbol{s}_j), S(\boldsymbol{s}_k)) &= \rho(|\boldsymbol{s}_j - \boldsymbol{s}_k|; \boldsymbol{\theta}_\rho). \end{aligned} \tag{3}$$

The correlation function is assumed to be isotropic. The mean vector $\boldsymbol{\mu}$, the marginal variance σ^2 and the parameters $\boldsymbol{\theta}_\rho$ of the correlation function ρ are taken to be unknown and are assigned prior distributions as described in Section 3. Thus, our prior model of the risk surface and the corresponding hyper-parameters is the lattice analogue of the geostatistical model of Kelsall and Wakefield (2002). Given the log-risk surface, the data are assumed to be conditionally independent realisations from Poisson distributions given by

$$y_i \mid R_i \sim \mathrm{Pois}(E_i R_i), \tag{4}$$

where the regional level relative risks R_i are computed from the underlying lattice-based risk surface by

$$R_i = \sum_{j:\ \boldsymbol{s}_j \in \mathcal{A}_i} R(\boldsymbol{s}_j) w(\boldsymbol{s}_j). \tag{5}$$

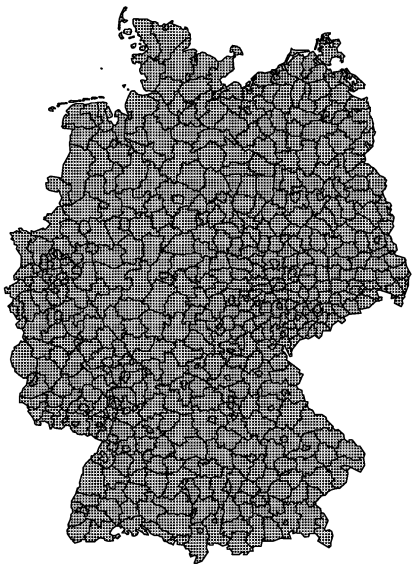

Figure 2: The map of Germany with its 544 districts, overlaid by the lattice for the GRF.

Here, the population density distributions $f_i(s)$ of the continuous surface analogue (2) are replaced by a set of weights $w(\boldsymbol{s}_j)$ which should satisfy the constraint

$$\sum_{j:\, \boldsymbol{s}_j \in \mathcal{A}_i} w(\boldsymbol{s}_j) = 1; \forall i. \tag{6}$$

In the approach by Kelsall and Wakefield (2002), the GRF prior model for the log-risk surface is used to generate a multivariate log-Normal approximation to the regional-level relative risks. To avoid computing such an approximation, we base inference directly on the risk surface model as defined on the lattice. However, for the lattice model to be a reasonable approximation to the smooth surface, the resolution should be relatively high, and consequently the number of nodes of the lattice will typically be large. Using the Gaussian random field representation of the prior model, estimation routines will be computationally expensive since we need to perform matrix operations on the $n_{GRF} \times n_{GRF}$ covariance matrix $\boldsymbol{\Sigma}$, which in general is a full matrix. Moreover, using MCMC methods with single-site updating, the convergence will be slow, due to the high correlations inherent in the prior model. Similar problems arise when estimating the hyper-parameters, because of the strong interaction with the elements of the risk surface. On the other hand, due to the high dimensionality and the full structure of the covariance matrix, updating all elements of the surface in one block will be prohibitive. In the next subsection we describe how the GRF can be represented by a more

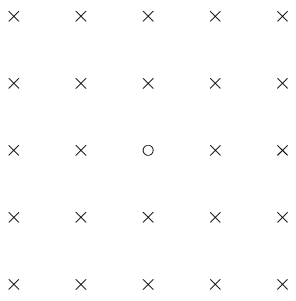

Figure 3: The neighbours ($\times$) of an arbitrary lattice node ($\circ$) using a 5×5 neighbourhood scheme for the GMRF.

computationally convenient Gaussian Markov random field on the lattice.

2.2 A Gaussian Markov random field proxy to the Gaussian random field model

To reduce the computational cost of the inference, we propose to represent the log-risk surface by a vector variable $\boldsymbol{x} = \{x_j\}_{j=1,\ldots,n_{GRF}}$, which is assumed to be a realisation of a Gaussian Markov random field (GMRF). A GMRF is a GRF with the additional property that the conditional distribution of the GMRF at lattice node j, given the values at all other lattice nodes, only depends on the values at the nodes within a neighbourhood $\delta(j)$ of j. Different definitions of neighbourhoods are possible, but we choose to define the neighbourhood $\delta(j)$ of node j to be an 5×5 square neighbourhood, as illustrated in Figure 3. Since we are dealing with a finite lattice, the number of neighbours of the lattice nodes along the boundary of the lattice will be different from the number given by the 5×5 neighbourhood scheme, see Figure 4. To reduce the impact of any boundary effects induced by using a finite lattice, we extend the support of the GMRF to include a set of nodes outside the region of interest. We will denote the sub-vector corresponding to the nodes falling within the region of interest by $\boldsymbol{x}_{\mathcal{A}}$, and the nodes external to this region by $\boldsymbol{x}_{-\mathcal{A}}$, such that $\boldsymbol{x} = (\boldsymbol{x}_{\mathcal{A}}^T, \boldsymbol{x}_{-\mathcal{A}}^T)^T$. The extended lattice, consisting of $n = 31089$ nodes, is shown in Figure 5.

For a general GMRF $\boldsymbol{x}$, the joint distribution is given by

$$\boldsymbol{x} \sim N(\boldsymbol{\mu}, \boldsymbol{Q}^{-1}), \tag{7}$$

where the mean vector $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{\theta})$ and the precision matrix $\boldsymbol{Q} = \boldsymbol{Q}(\boldsymbol{\theta})$ both may depend on a set of unknown parameters $\boldsymbol{\theta}$. Because of the conditional independence structure of the GMRF, only the elements Q_{ij} of the precision matrix for which i and j are neighbours are non-zero. The nodes of the lattice can be re-ordered such as to minimise the

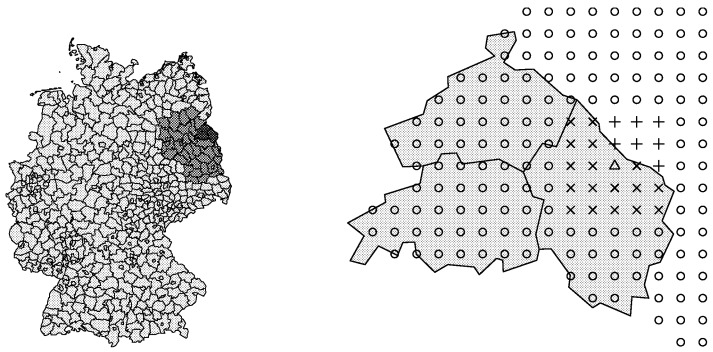

Figure 4: The neighbouring scheme along the boundary of the study region. The neighbours of a node (triangle) at the boundary are partly within the study region (x) and partly outside the region (+).

bandwidth of the corresponding precision matrix (Knorr-Held and Rue, 2002), and due to the band-structure of the matrix, working with a GMRF instead of a GRF can lead to significant reductions in computational cost. This fact is utilised in our sampling based estimation approach described in Sections 3 and 4. There, we make extensive use of efficient algorithms for generating samples from joint and conditional distributions of a GMRF as well as for generating samples conditionally on linear constraints. A sample from the joint distribution of $\boldsymbol{x}$ can be generated by

$$\boldsymbol{x} = \boldsymbol{L}^{-1}\boldsymbol{z} + \boldsymbol{\mu}, \tag{8}$$

where $\boldsymbol{z}$ is a vector of n independent realisation from the standard Normal distribution, and $\boldsymbol{L}$ is the Cholesky factor of the precision matrix $\boldsymbol{Q}$. For a banded symmetric positive definite matrix $\boldsymbol{Q}$ with bandwidth b_w, the Cholesky factorisation

$$\boldsymbol{Q} = \boldsymbol{L}\boldsymbol{L}^T \tag{9}$$

can be computed in $\mathrm{O}(nb_w^2)$ flops (Rue, 2001), such that as long as the bandwidth is kept small, efficient samples can be generated from the joint distribution. In our application, we need to generate conditional samples for a subset of the lattice nodes, given the realisation of the GMRF for the remaining nodes. The conditional distribution $\pi(\boldsymbol{x}_{\mathcal{S}}|\boldsymbol{x}_{-\mathcal{S}})$, where $\mathcal{S}$ is a subset of $\mathcal{A}$, is Normal with moments

$$\begin{aligned} \mathrm{E}(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}}) &= \boldsymbol{\mu}_{\mathcal{S}} - \boldsymbol{Q}_{\mathcal{S},\mathcal{S}}^{-1}\boldsymbol{Q}_{\mathcal{S},-\mathcal{S}}\,(\boldsymbol{x}_{-\mathcal{S}} - \boldsymbol{\mu}_{-\mathcal{S}}), & (10)\\ \mathrm{Var}(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}}) &= \boldsymbol{Q}_{\mathcal{S},\mathcal{S}}^{-1}. & (11) \end{aligned}$$

Here, $\boldsymbol{Q}_{\mathcal{S},\mathcal{S}}$ is the $n_{\mathcal{S}} \times n_{\mathcal{S}}$ diagonal block of $\boldsymbol{Q}$ corresponding to the subset $\mathcal{S}$, with bandwidth less than or equal to the bandwidth of $\boldsymbol{Q}$. Each element l of the mean vector

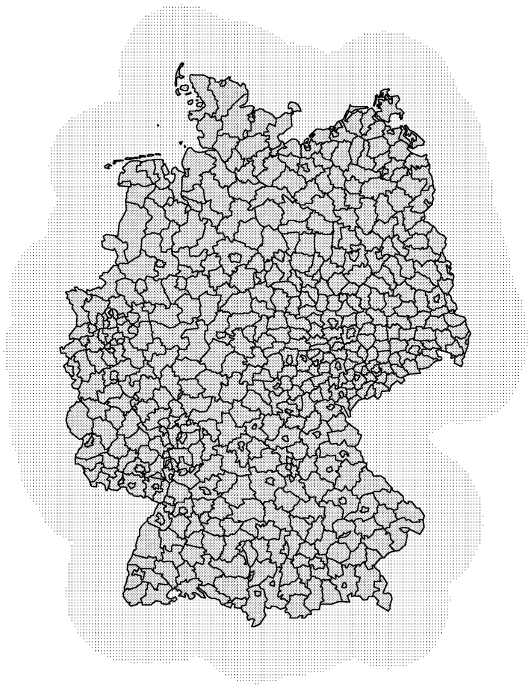

Figure 5: The map of Germany with its 544 districts, overlaid by the lattice GMRF model including a boundary region.

$\boldsymbol{\mu}_{\mathcal{S}|-\mathcal{S}} = \mathrm{E}(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}})$ will only depend on elements of $\boldsymbol{x}_{-\mathcal{S}} - \boldsymbol{\mu}_{-\mathcal{S}}$ at the nodes within the neighbourhood of node l.

We also need to generate samples from a GMRF $\boldsymbol{x}$ conditionally on a linear *soft* constraint $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} + \boldsymbol{\epsilon}$ for a $p \times n$ matrix $\boldsymbol{A}$, a p-vector $\boldsymbol{b}$ and $\boldsymbol{\epsilon} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma})$. This constraint can be interpreted as a generalisation of the *hard* constraint $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, where the quantities representing the linear combinations defining the constraint are observed with noise. In general, a sample conditionally on a soft constraint can be generated by first generating an unconditional sample $\boldsymbol{x}_u$ for $\boldsymbol{x}$ from (7) and an $\boldsymbol{\epsilon} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma})$, and then computing the conditional sample $\boldsymbol{x}_c$ from

$$\boldsymbol{x}_c = \boldsymbol{x}_u - \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{\epsilon} - \boldsymbol{b}). \tag{12}$$

In geostatistics, this result is referred to as conditional simulation using kriging (Cressie, 1993, Section 3.6.2) and the validity of (12) as a sample from $\pi(\boldsymbol{x}|\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} + \boldsymbol{\epsilon})$ follows directly from Normal distribution theory, as shown in Appendix A.3. As long as the number of constraints p is relatively small compared to the number of nodes in the

lattice, all computations involved in evaluating (12) can be done efficiently using the Cholesky factorisation (9).

To fully specify the joint distribution (7) we need to specify the non-zero elements of the precision matrix $\boldsymbol{Q}$. However, based on prior information it is often intuitively easier to specify a model for the correlation structure for a Gaussian random field than to specify the elements of the corresponding precision matrix for the GMRF. Rue and Tjelmeland (2002) show how the elements of the precision matrix $\boldsymbol{Q}$ of a GMRF can be estimated from the covariance function that defines the elements of the covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{Q}^{-1}$. Let $\rho(h; \boldsymbol{\theta}_\rho)$ be the correlation function specifying the correlation between two points of distance h, where h is measured in lattice coordinates. Further, let

$$\boldsymbol{Q} = \tau \boldsymbol{Q}' = \tau \boldsymbol{C}^{-1} \tag{13}$$

where $\boldsymbol{C}$ is the correlation matrix of the GRF and $\tau = 1/\sigma^2$ is the marginal precision. For a given value of the parameter vector $\boldsymbol{\theta}_\rho$, Rue and Tjelmeland (2002) estimate the non-zero elements of the standardised precision matrix $\boldsymbol{Q}'$ by matching the correlation function as defined by these elements to the correlation function ρ of a Gaussian random field. For the exponential, Gaussian, spherical and Matérn classes of correlation functions, they conclude that using a 5×5 neighbourhood the approach gives a good fit to the target correlation function. Among these four classes of functions, the exponential is the one with the best fit.

Using the GMRF prior model for the log relative risk surface, expression (5) for the relative risk R_i at the regional level is replaced by

$$R_i = \sum_{j \in \mathcal{A}_i} \exp(x_j)\ w(\boldsymbol{s}_j), \tag{14}$$

where the sum is taken over the n_i nodes of the lattice falling within region $\mathcal{A}_i$. In what follows we will assume that the weights are constants given by $w(\boldsymbol{s}_j) = 1/n_i;\ j \in \mathcal{A}_i$. This corresponds to an assumption of uniform population density which is often made in disease mapping applications. This does not represent any loss of generality, since the method can easily be modified to allow for non-uniform population distributions by replacing the GMRF $\boldsymbol{x}$ by another GMRF $\boldsymbol{x}'$ with elements $x'_j = x_j + \log(w(\boldsymbol{s}_j))$. In terms of the log-risk surface $\boldsymbol{x}$, the Poisson likelihood model for the incidence counts becomes

$$y_i \mid \boldsymbol{x} \sim \text{Pois}\left(\frac{E_i}{n_i} \sum_{j \in \mathcal{A}_i} \exp(x_j)\right). \tag{15}$$

For notational convenience we define $E'_i = E_i/n_i$, such that in what follows, E'_i is to be interpreted as the expected number of cases per lattice node falling within region $\mathcal{A}_i$.

Let $\boldsymbol{\theta} = (\tau, \boldsymbol{\theta}_\rho^T, \boldsymbol{\theta}_\mu^T)^T$ denote all unknown hyper-parameters of the model, including the precision τ, the parameters $\boldsymbol{\theta}_\rho$ of the correlation function and any parameters $\boldsymbol{\theta}_\mu$ defining the mean vector $\boldsymbol{\mu}$. We take a fully Bayesian approach to parameter estimation,

and the Bayesian hierarchical model for the disease mapping problem is summarised in Table 1. The prior distribution $\pi(\boldsymbol{\theta})$ of the hyper-parameters is specified in Section 3.

Likelihood:	$y_i \mid \boldsymbol{x}$	$\sim$	$\text{Pois}(E_i' \sum_{j \in \mathcal{A}_i} \exp(x_j))$
GMRF prior:	$\boldsymbol{x} \mid \boldsymbol{\theta}$	$\sim$	$N(\boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{Q}(\boldsymbol{\theta})^{-1})$
Hyper-prior:	$\boldsymbol{\theta}$	$\sim$	$\pi(\boldsymbol{\theta})$

Table 1: A summary of the Bayesian hierarchical model.

We end this section by pointing out some computational pitfalls that are still present using the GMRF representation of the model in combination with an MCMC based approach to parameter estimation. In our approach to estimation of the risk surface and the corresponding hyper-parameters we need to sample from the posterior distribution of the log-risk surface $\boldsymbol{x}$ given count data $\boldsymbol{y}$. In Knorr-Held and Rue (2002) it is illustrated how the use of block-sampling leads to substantial improvement in mixing for MCMC updating schemes for a similar model, but where the GMRF prior for the log-risk is defined on the same level of aggregation as the data, using a common boundary neighbourhood specification. The observations are conditionally independent given the regional risks and the hyper-parameters, and thus the conditional independence structure for the posterior is the same as for the GMRF prior. Since the data in our case are aggregated in regions that in general extends over the size of the local neighbourhoods of the GMRF model for $\boldsymbol{x}$, conditioning on the data will destroy the computationally convenient local neighbourhood structure inherent in the prior. For an illustration, consider the plots in the bottom panels of Figure 6. There, we have visualised the conditional independence structure of the prior model and of the posterior model conditioning on the data for a subset of the study region. The subset is given by the lattice nodes within the shaded region in the top left panel, plotted in larger scale and overlaid by the GMRF lattice in the top right panel. Consequently, to preserve computational efficiency, alternative methods are needed.

We have pointed out the potential problem of slow convergence of the hyper-parameters using MCMC methods with the Gaussian random field. Inference for the hyper-parameters will still be problematic using the lattice based GMRF modelling approach, unless we are able to update the hyper-parameters jointly with a subset of the GMRF. Finally, introducing the additional boundary region nodes to reduce the impact of boundary effects, might slow down the convergence and mixing for the hyper-parameters.

In the next sections we present our sampling based approach to parameter estimation, and discuss how the potential problems listed above can be handled within our sampling algorithm. We first give an overview of the method, and then present the approach in more detail in Section 4.

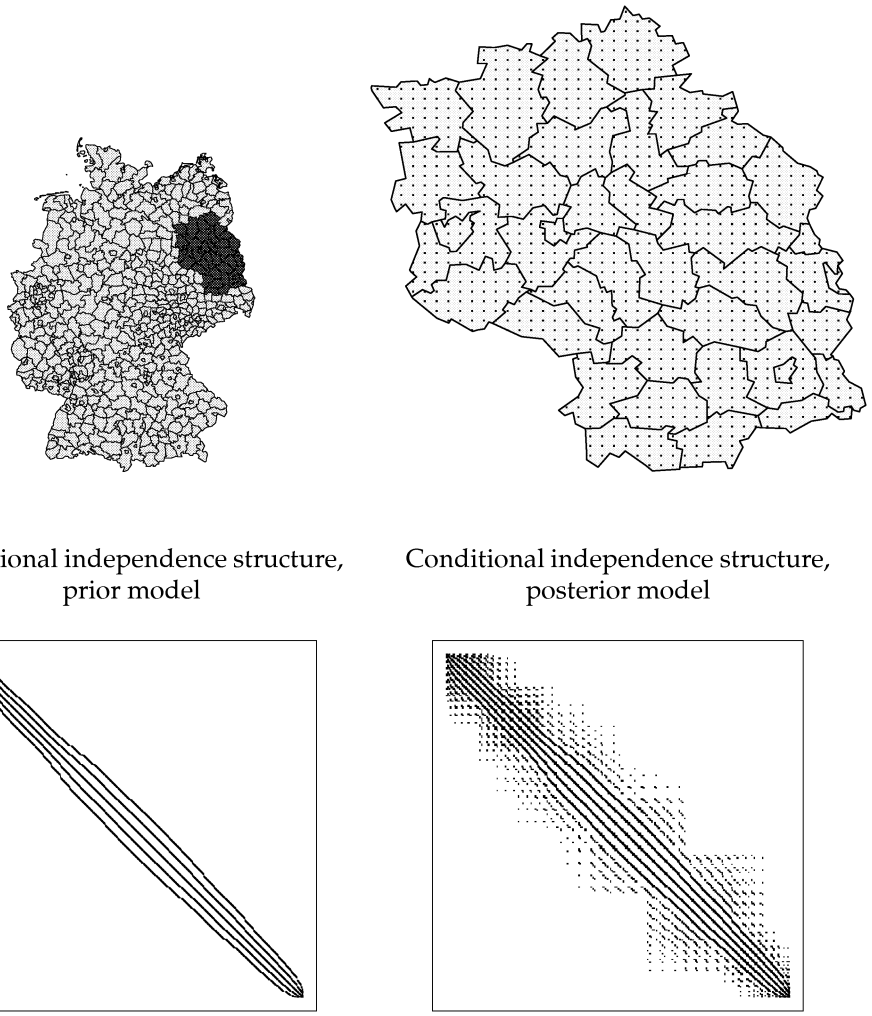

Figure 6: A subset of the region of study (shaded region, top left), with the lattice nodes added (top right). The bottom panels illustrate the conditional independence structure of the prior model (bottom left) and when conditioning on the data (bottom right), for the subset of lattice nodes corresponding to the shaded region.

3 A sampling based estimation approach

In this section we present a sampling based approach to the estimation of the unknown quantities of the model. These are the log risk surface represented by $\boldsymbol{x}$ and the hyper-parameters $\boldsymbol{\theta}$.

Given the Poisson likelihood, the GMRF prior (7) of $\boldsymbol{x}$ and the joint prior distribution $\pi(\boldsymbol{\theta})$ of the parameters $\boldsymbol{\theta}$, the joint posterior distribution of $\boldsymbol{x}$ and $\boldsymbol{\theta}$ given the data $\boldsymbol{y}$ is given by

$$\pi(\boldsymbol{x}, \boldsymbol{\theta} \mid \boldsymbol{y}) \propto \prod_i \{\pi(y_i \mid \boldsymbol{x}, \boldsymbol{\theta})\} \pi(\boldsymbol{x} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}). \tag{16}$$

We will estimate the log risk surface and the hyper-parameters using Markov chain Monte Carlo methods. As pointed out in Section 2, using single-site updating will typically lead to poor mixing and slow convergence due to strong correlations in the prior model. On the other hand, updating $\boldsymbol{x}$ or $\boldsymbol{x}$ and $\boldsymbol{\theta}$ in one block requires careful choice of the proposal distribution to obtain a reasonably high acceptance rate. Knorr-Held and Rue (2002) illustrate the use of block-updating in Markov random field models for disease mapping applications where the relative risk is defined at the regional level. In their case, the Poisson likelihood is given by (4) with $R_i = \exp(x_i)$, where x_i is the log-risk of region i. The joint proposal of the log relative risks $\boldsymbol{x}$ conditional on a proposed value of $\boldsymbol{\theta}$ is generated by using a local quadratic approximation to the posterior. Utilising the band structure of the prior precision matrix $\boldsymbol{Q}$ they obtain computationally efficient samples from the proposal distribution. Applying the approach to a dataset on Insulin dependent Diabetes Mellitus in 366 districts of Sardinia, it is shown that the convergence and mixing of the hyper-parameters are greatly improved by blocking. However, since the spatial model is specified on the regional level, the total number of parameters to be updated are much smaller than in our application, where the Markov random field is specified on a lattice of $n = 31089$ nodes. As a consequence, the acceptance rate of a joint proposal is likely to be reduced compared to the ones reported in their study, and the computational cost of generating the sample is increased. Also, the efficiency of the approach as applied to our problem is reduced because the level of aggregation of the data typically extends the size of the neighbourhood of the GMRF model, as discussed at the end of Section 2.

As a compromise between a full block sampler and the single site Gibbs sampler, we update the hyper-parameters and a *subset* of the elements of $\boldsymbol{x}$ in one block. We split the vector $\boldsymbol{x}$ in the two sub-vectors $\boldsymbol{x}_{\mathcal{A}}$ and $\boldsymbol{x}_{-\mathcal{A}}$, representing the lattice nodes within the region of study $\mathcal{A}$ and in the boundary region respectively, and choose to block-update the hyper-parameters jointly with the subset $\boldsymbol{x}_{-\mathcal{A}}$. As we show in Section 3.3, this is equivalent to updating the hyper-parameters by sampling from the marginal posterior of $\boldsymbol{\theta}$ given $\boldsymbol{x}_{\mathcal{A}}$. Conditionally on $\boldsymbol{x}_{-\mathcal{A}}$ and $\boldsymbol{\theta}$, the elements of $\boldsymbol{x}_{\mathcal{A}}$ are updated in sub-blocks corresponding to the lattice nodes within one or more regions.

In the following subsections we describe our approach to estimation of the log-risk surface as well as the hyper-parameters, and discuss how the potential computational pitfalls pointed out at the end of Section 2 can be handled.

3.1 Step 1: The log risk surface

The sub-vector $\boldsymbol{x}_{\mathcal{A}}$ corresponding the elements of $\boldsymbol{x}$ falling within the study region for which data are available, are updated conditionally on $\boldsymbol{x}_{-\mathcal{A}}$ and the hyper-parameters $\boldsymbol{\theta}$. We update $\boldsymbol{x}_{\mathcal{A}}$ in blocks defined in terms of the regions corresponding to the level of aggregation of the data, using a Metropolis-Hastings step for each block. In this section we give an overview of the approach for a general likelihood, deferring a more detailed description of the sampling procedure for the Poisson likelihood case to Section 4. Also, we first describe the method in terms of m blocks, where each block is made up from the elements of $\boldsymbol{x}_{\mathcal{A}}$ falling within a single region. In Section 3.2 we describe the straightforward extension of the method to blocks made up of from several regions, and discuss briefly considerations to be made by choosing the size of the blocks.

The full conditional distribution for the n_i elements $x_{j_i};\ j_i = 1, \ldots, n_i$ within region $\mathcal{A}_i$ is given by

$$\pi(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta}, \boldsymbol{y}) \propto \pi(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta}) \pi(y_i \mid \boldsymbol{x}_{\mathcal{A}_i}, \boldsymbol{\theta}). \tag{17}$$

Here we denote by $\boldsymbol{x}_{-\mathcal{A}_i}$ all elements of the vector $\boldsymbol{x}$ except for the elements within region $\mathcal{A}_i$. The posterior distribution is in general non-standard, and we use a Metropolis-Hastings step to generate an update of $\boldsymbol{x}_{\mathcal{A}_i}$. As a proposal distribution for $\boldsymbol{x}_{\mathcal{A}_i}$ we use a quadratic approximation to (17), and we illustrate below that by re-formulating this quadratic approximation to the distribution of a conditional sampling problem, sampling can be done efficiently.

The conditional prior distribution $\pi(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta})$ of $\boldsymbol{x}_{\mathcal{A}_i}$ is Gaussian with mean $\boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}}$ depending on the values of $\boldsymbol{x}$ in the set of nodes given by $\delta_{\mathcal{A}_i} = \cup_{j \in \mathcal{A}_i} \delta(j)$, and precision matrix $\boldsymbol{Q}_{\mathcal{A}_i}$ given by the $n_i \times n_i$ diagonal block of $\boldsymbol{Q}$ corresponding to the nodes within region $\mathcal{A}_i$. The matrix $\boldsymbol{Q}_{\mathcal{A}_i} = \boldsymbol{Q}_{\mathcal{A}_i}(\boldsymbol{\theta})$ and the vector $\boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}} = \boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}}(\boldsymbol{\theta})$ will both in general depend on $\boldsymbol{\theta}$, but for notational convenience we suppress explicit reference to the dependency on $\boldsymbol{\theta}$ in what follows. Thus, the log-posterior distribution corresponding to (17) becomes

$$\log(\pi(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta}, \boldsymbol{y})) = \\ -\frac{1}{2}(\boldsymbol{x}_{\mathcal{A}_i} - \boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}})^T \boldsymbol{Q}_{\mathcal{A}_i}(\boldsymbol{x}_{\mathcal{A}_i} - \boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}}) + h_i(\boldsymbol{x}) + \text{const}, \tag{18}$$

where $h_i(\boldsymbol{x})$ is the log-likelihood of the observed count for region $\mathcal{A}_i$. Introducing the vector $\boldsymbol{d}_i$ given by $\boldsymbol{d}_i = \boldsymbol{Q}_{\mathcal{A}_i}\boldsymbol{\mu}_{\mathcal{A}_i|\delta_{\mathcal{A}_i}}$ and re-arranging terms, the log-posterior distribution

(18) can be written in the form

$$\log(\pi(\boldsymbol{x}_{\mathcal{A}_i}|\boldsymbol{x}_{-\mathcal{A}_i},\boldsymbol{\theta},\boldsymbol{y})) \;=\; -\frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T\boldsymbol{Q}_{\mathcal{A}_i}\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{d}_i^T\boldsymbol{x}_{\mathcal{A}_i} + h_i(\boldsymbol{x}) + \text{const.} \tag{19}$$

A Gaussian approximation to (19) can be found by replacing the term $h_i(\boldsymbol{x})$ by a quadratic approximation

$$h_i(\boldsymbol{x}) \approx -\frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T\boldsymbol{B}_i\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{b}_i^T\boldsymbol{x}_{\mathcal{A}_i}, \tag{20}$$

where $\boldsymbol{B}_i$ and $\boldsymbol{b}_i$ in general depend on the observation y_i and the parameters $\boldsymbol{\theta}$. We use a second order Taylor expansion of $h_i(\boldsymbol{x})$ to define the quadratic approximation, as will be described in Section 4 for the Poisson likelihood case. Substituting (20) for $h_i(\boldsymbol{x})$ in (19) and collecting terms that are linear and quadratic in $\boldsymbol{x}_{\mathcal{A}_i}$, a quadratic approximation to the full conditional density (19), which we denote by $\pi_N(\boldsymbol{x}_{\mathcal{A}_i}|\boldsymbol{x}_{-\mathcal{A}_i},\boldsymbol{y})$, is

$$\begin{aligned}\log(\pi_N(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i},\boldsymbol{\theta},\boldsymbol{y})) &= -\frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T(\boldsymbol{Q}_{\mathcal{A}_i}+\boldsymbol{B}_i)\boldsymbol{x}_{\mathcal{A}_i} + (\boldsymbol{d}_i+\boldsymbol{b}_i)^T\boldsymbol{x}_{\mathcal{A}_i} + \text{const} \\ &= -\frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T(\boldsymbol{Q}_{\mathcal{A}_i}+\boldsymbol{B}_i)\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{c}_i^T\boldsymbol{x}_{\mathcal{A}_i} + \text{const},\end{aligned} \tag{21}$$

where we have defined $\boldsymbol{c}_i = \boldsymbol{d}_i + \boldsymbol{b}_i$. This Gaussian approximation is to be used as a proposal distribution in a Metropolis-Hastings step for updating $\boldsymbol{x}_{\mathcal{A}_i}$. However, the precision matrix $\boldsymbol{Q}_{\mathcal{A}_i}+\boldsymbol{B}_i$ of the Gaussian distribution defined by (21) is in general a full matrix, such that the computationally convenient band structure of the prior precision matrix $\boldsymbol{Q}_{\mathcal{A}_i}$ is lost. This effect of conditioning on the data was illustrated in Figure 6. If the elements of $\boldsymbol{x}$ are updated for each region $\mathcal{A}_i$ in turn, this does not necessary imply any significant loss of efficiency, since the number of lattice nodes within each region is typically relatively small and not very much larger than the number of neighbours of a lattice node. But in the general case when elements are updated in larger blocks, preserving the band structure might lead to substantial computational savings.

The general idea of our sampling approach is to re-formulate the problem of sampling directly from the Gaussian proposal distribution (21) to a conditional sampling problem for which the band structure of the precision matrix $\boldsymbol{Q}_{\mathcal{A}_i}$ is preserved. The symmetric matrix $\boldsymbol{B}_i$ can be expressed by

$$\boldsymbol{B}_i = \boldsymbol{D}_i + \boldsymbol{A}_i^T\boldsymbol{A}_i, \tag{22}$$

where $\boldsymbol{D}_i$ is a $n_i \times n_i$ diagonal matrix, possibly with zeros on the diagonal, and $\boldsymbol{A}_i$ is a $1 \times n_i$ matrix. Substituting (22) for $\boldsymbol{B}_i$ in (21) and re-arranging terms, we arrive at the expression

$$\log(\pi_N(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i},\boldsymbol{\theta},\boldsymbol{y})) = -\frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T(\boldsymbol{Q}_{\mathcal{A}_i}+\boldsymbol{D}_i)\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{c}_i^T\boldsymbol{x}_{\mathcal{A}_i} - \frac{1}{2}\boldsymbol{x}_{\mathcal{A}_i}^T\boldsymbol{A}_i^T\boldsymbol{A}_i\boldsymbol{x}_{\mathcal{A}_i} + \text{const.} \tag{23}$$

As long as the matrix $\boldsymbol{Q}_{\mathcal{A}_i} + \boldsymbol{D}_i$ is positive definite, a requirement that is discussed in Section 4.2, the first two terms on the right define the log-density of a Gaussian variable

for which the precision matrix $\boldsymbol{Q}_{\mathcal{A}_i} + \boldsymbol{D}_i$ has the same bandwidth as $\boldsymbol{Q}_{\mathcal{A}_i}$. The last term can be recognised as the log density, up to a constant, of a Gaussian variable with mean $\boldsymbol{A}_i \boldsymbol{x}_{\mathcal{A}_i}$ and covariance matrix $\boldsymbol{I}$, evaluated in $\boldsymbol{0}$. Sampling from (23) is shown in Section 4.2 to be equivalent to sampling from the conditional distribution

$$\pi(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{y}, \boldsymbol{z}^* = \boldsymbol{0}), \tag{24}$$

where $\boldsymbol{z}^* \mid \boldsymbol{x}_{\mathcal{A}_i} \sim N(\boldsymbol{A}_i \boldsymbol{x}_{\mathcal{A}_i}, \boldsymbol{I})$. In Section 4.2 we further illustrate that because of the band structure of $\boldsymbol{Q}_{\mathcal{A}_i} + \boldsymbol{D}_i$, sampling from (24) is computationally much more efficient than sampling directly from the Gaussian approximation as defined by (21), for which the precision matrix is in general a full matrix.

The proposed method for sampling from the full conditional distribution for $\boldsymbol{x}_{\mathcal{A}_i}$, given by $\pi(\boldsymbol{x}_{\mathcal{A}_i} | \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta}, \boldsymbol{y})$ in (19), can be summarised in the following steps.

1. Approximate the likelihood part of the full conditional distribution by a quadratic function in $\boldsymbol{x}_{\mathcal{A}_i}$, obtaining a Normal approximation to the full conditional distribution. The quadratic approximation is computed by Taylor expansion around the conditional mode.
2. Re-formulate the problem of sampling from this Normal approximation as a conditional simulation problem, where the band structure of the precision matrix is preserved.
3. Generate a sample from the Normal approximation based on the re-formulated problem. This can be done using efficient algorithms utilising the band structure of the precision matrix, described in Section 2.2.
4. Use the sample from 3. as a proposed value for $\boldsymbol{x}_{\mathcal{A}_i}$ in a Metropolis-Hastings step, compute the acceptance probability and accept or reject this value.

3.2 Updating blocks of general subsets of $\boldsymbol{x}$

In Section 3.1 we described our approach to updating the elements of $\boldsymbol{x}$ for each region separately. An equivalent approach can be taken to update larger subsets of $\boldsymbol{x}_{\mathcal{A}}$ or all elements of $\boldsymbol{x}_{\mathcal{A}}$ jointly, given the parameters $\boldsymbol{\theta}$ and the boundary elements $\boldsymbol{x}_{-\mathcal{A}}$ of the random field. The full conditional distribution of $\boldsymbol{x}_{\mathcal{S}}$ for a general subset $\mathcal{S}$ of $\mathcal{A}$ given $\boldsymbol{x}_{-\mathcal{S}}$, the data $\boldsymbol{y}$ and the hyper-parameters $\boldsymbol{\theta}$, is given by

$$\pi(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}}, \boldsymbol{\theta}, \boldsymbol{y}) \propto \pi(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}}, \boldsymbol{\theta}) \prod_{i:\, \mathcal{A}_i \in \mathcal{S}} \pi(y_i \mid \boldsymbol{x}_{\mathcal{S}}, \boldsymbol{x}_{-\mathcal{S}}, \boldsymbol{\theta}). \tag{25}$$

In analogy to (19), the corresponding log-density can be written as

$$\log(\pi(\boldsymbol{x}_{\mathcal{S}} | \boldsymbol{x}_{-\mathcal{S}}, \boldsymbol{y}, \boldsymbol{\theta})) \;=\; -\frac{1}{2} \boldsymbol{x}_{\mathcal{S}}^T \boldsymbol{Q}_{\mathcal{S}} \boldsymbol{x}_{\mathcal{S}} + \boldsymbol{d}_{\mathcal{S}}^T \boldsymbol{x}_{\mathcal{S}} + \sum_{i:\, \mathcal{A}_i \in \mathcal{S}} h_i(\boldsymbol{x}) + \text{const}, \tag{26}$$

where $h_i(\boldsymbol{x})$ is the log-likelihood of the observed count for region $\mathcal{A}_i$, and the sum is taken over the $m_{\mathcal{S}}$ regions corresponding to the subset $\mathcal{S}$. By substituting a quadratic approximation computed by a second order Taylor expansion for $\sum_i h_i(\boldsymbol{x})$ in (26) and re-arranging terms, it is shown in Appendix A.2, using a Poisson likelihood, that the corresponding Gaussian approximation to (26) can be expressed by

$$\log(\pi_N(\boldsymbol{x}_{\mathcal{S}} \mid \boldsymbol{x}_{-\mathcal{S}}, \boldsymbol{\theta}, \boldsymbol{y})) = \\ -\frac{1}{2}\boldsymbol{x}_{\mathcal{S}}^T(\boldsymbol{Q}_{\mathcal{S}} + \boldsymbol{D}_{\mathcal{S}})\boldsymbol{x}_{\mathcal{S}} + \boldsymbol{c}_{\mathcal{S}}^T\boldsymbol{x}_{\mathcal{S}} - \frac{1}{2}\boldsymbol{x}_{\mathcal{S}}^T\boldsymbol{A}_{\mathcal{S}}^T\boldsymbol{A}_{\mathcal{S}}\boldsymbol{x}_{\mathcal{S}} + \text{const}, \quad (27)$$

where $\boldsymbol{D}_{\mathcal{S}}$ is a diagonal matrix and $\boldsymbol{A}_{\mathcal{S}}$ a $m_{\mathcal{S}} \times n_{\mathcal{S}}$ matrix, with $n_{\mathcal{S}}$ equal to the number of lattice nodes within the subset $\mathcal{S}$ of regions. This Gaussian distribution is of the same form as (23) for the single region case, and the approach for generating samples from (27) to be described in Section 4.2 can be applied in the case of general subsets as well.

An extension of single region blocks to larger blocks can be generated by including the lattice nodes corresponding to the neighbours of the region, where we define two regions as neighbours if they share a common boundary, and further extensions can be made by adding the neighbours of the neighbours and so on. In Figure 7 we illustrate the size of the blocks corresponding to different choices of the number of neighbours to include in each block. We define the term 1. order neighbourhood to mean all neighbours of a region, 2. order neighbourhood to mean all neighbours as well as all neighbours of the neighbours and so on. The choice of the number of regions to be updated in each sub-block is a trade-off between computational cost and the acceptance probabilities of the Metropolis-Hastings steps. Using the sampling approach outlined in Section 3.1, the problem of reduced computational efficiency due to the fact that the band structure of the precision matrix was not preserved in the Gaussian approximation to the posterior, has been handled. Therefore, the computational cost is expected to be reduced by increasing the size of the blocks and thus reduce the number of blocks. On the other hand, although increased block size might improve mixing due to larger differences between proposed and current values, increasing the number of elements of each block will reduce the quality of the Gaussian approximation, such that the acceptance rate is typically reduced. This should be kept at a reasonable level to ensure proper mixing of the MCMC algorithm.

3.3 Step 2: The hyper-parameters

The hyper-parameters $\boldsymbol{\theta}$ are updated jointly with the remaining elements $\boldsymbol{x}_{-\mathcal{A}}$ of the GMRF. As pointed out in the beginning of this section, updating all element of the GMRF in one block will most likely lead to low acceptance rates, and therefore we block the hyper-parameters with a subset of the elements of $\boldsymbol{x}$.

Using the fact that $\boldsymbol{x}_{-\mathcal{A}}$ is conditionally independent of the data $\boldsymbol{y}$ given $\boldsymbol{x}_{\mathcal{A}}$ and $\boldsymbol{\theta}$, the

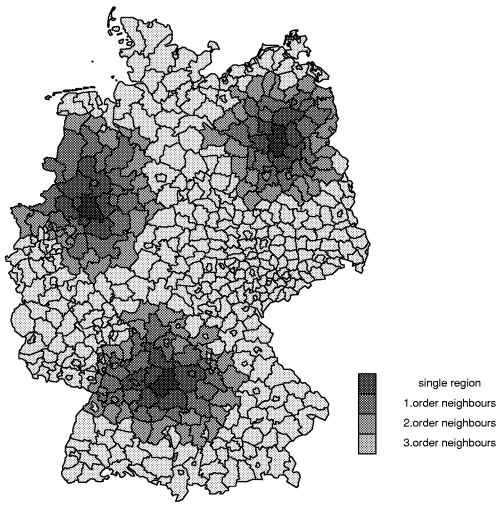

Figure 7: Different possible structures of the blocks for block-updating the log-risk surface. The 1. order neighbours are additional nodes from the 1. order neighbourhood of the single regions, 2. order neighbours are additional nodes in the 2. order neighbourhood and 3. order neighbours the nodes added to the 2. order neighbourhood from the 3. order neighbourhood.

full conditional distribution of $(\boldsymbol{\theta}, \boldsymbol{x}_{-\mathcal{A}})$ is

$$\pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y}) = \pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}}) \propto \pi(\boldsymbol{x}_{\mathcal{A}}, \boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta}). \tag{28}$$

This distribution is in general a non-standard distribution, depending on the form of the hyper-prior $\pi(\boldsymbol{\theta})$, and updates are generated using a Metropolis-Hastings step. First, given the current value $\boldsymbol{\theta}$ of the hyper-parameters, a new parameter vector $\boldsymbol{\theta}'$ is sampled from a proposal distribution $q(\boldsymbol{\theta} \to \boldsymbol{\theta}')$, and then a proposed value $\boldsymbol{x}'_{-\mathcal{A}}$ is generated by sampling from the conditional distribution $\pi(\boldsymbol{x}'_{-\mathcal{A}}|\boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y}) = \pi(\boldsymbol{x}'_{-\mathcal{A}}|\boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})$. Since the joint distribution of $\boldsymbol{x}$ given $\boldsymbol{\theta}$ is Gaussian, it follows that the conditional distribution $\pi(\boldsymbol{x}'_{-\mathcal{A}}|\boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})$ is also Gaussian, with a precision matrix that has the same bandwidth as $\boldsymbol{Q}$. A proposed value can therefore be generated efficiently by sampling directly from this distribution. The proposed value is accepted or rejected according to the acceptance probability

$$\begin{aligned}
\alpha((\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta}), (\boldsymbol{x}'_{-\mathcal{A}}, \boldsymbol{\theta}')) &= \min\left(1, \frac{\pi(\boldsymbol{x}'_{-\mathcal{A}}, \boldsymbol{\theta}' \mid \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y})\ \pi(\boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta}, \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y})\ q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}{\pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y})\ \pi(\boldsymbol{x}'_{-\mathcal{A}} \mid \boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y})\ q(\boldsymbol{\theta} \to \boldsymbol{\theta}')}\right) \\
&= \min\left(1, \frac{\pi(\boldsymbol{x}'_{-\mathcal{A}}, \boldsymbol{\theta}' \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta}, \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}{\pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}'_{-\mathcal{A}} \mid \boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta} \to \boldsymbol{\theta}')}\right),
\end{aligned} \tag{29}$$

again utilising the conditional independence between $\boldsymbol{x}_{-\mathcal{A}}$ and $\boldsymbol{y}$ given $\boldsymbol{x}_{\mathcal{A}}$.

Recall the potential pitfall listed at the end of Section 2, pointing at the fact that the inclusion of the boundary region will in general be expected to slow down the convergence of the hyper-parameters. However, writing out the expression for the acceptance probability, we can show that blocking the hyper-parameters and the boundary nodes essentially eliminates this problem. There is still an effect of the boundary nodes on the subset of $\boldsymbol{x}_{\mathcal{A}}$ corresponding to the inner nodes close to the outer boundary of the study region, but this effect is supposed to be minor. Expanding the distribution $\pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta}|\boldsymbol{x}_{\mathcal{A}})$, (29) can be expressed by

$$\begin{aligned}
\alpha((\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta}), (\boldsymbol{x}'_{-\mathcal{A}}, \boldsymbol{\theta}')) &= \min\left(1, \frac{\pi(\boldsymbol{x}'_{-\mathcal{A}}, \boldsymbol{\theta}' \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta}, \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}{\pi(\boldsymbol{x}_{-\mathcal{A}}, \boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}'_{-\mathcal{A}} \mid \boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta} \to \boldsymbol{\theta}')}\right) \\
&= \min\left(1, \frac{\pi(\boldsymbol{x}'_{-\mathcal{A}} \mid \boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{\theta}' \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta}, \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}{\pi(\boldsymbol{x}_{-\mathcal{A}} \mid \boldsymbol{\theta}, \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}})\ \pi(\boldsymbol{x}'_{-\mathcal{A}} \mid \boldsymbol{\theta}', \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta} \to \boldsymbol{\theta}')}\right) \\
&= \min\left(1, \frac{\pi(\boldsymbol{\theta}' \mid \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta}' \to \boldsymbol{\theta})}{\pi(\boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}})\ q(\boldsymbol{\theta} \to \boldsymbol{\theta}')}\right). \qquad (30)
\end{aligned}$$

Consequently, sampling the hyper-parameters jointly with the elements of $\boldsymbol{x}$ outside the region of interest is equivalent to sampling the hyper-parameters from the marginal posterior distribution $\pi(\boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}}, \boldsymbol{y}) = \pi(\boldsymbol{\theta} \mid \boldsymbol{x}_{\mathcal{A}})$, integrated over the outer elements $\boldsymbol{x}_{-\mathcal{A}}$. In effect, using this approach the influence of the boundary nodes on the convergence of the hyper-parameters should be insignificant.

4 Efficient sampling from the full posterior of the log risk surface

In this section we describe in more detail our approach to the generation of samples from the full conditional distribution of $\boldsymbol{x}_{\mathcal{A}_i}$ given by (19), using the Poisson likelihood (15). We describe the sampling routine in terms of blocks made up from sets of lattice nodes corresponding to single regions, but as we pointed out in Section 3.2, the sampling problem for the general case has the same structure. In Section 4.1 we compute the quadratic approximation to (19) using Taylor expansion, and in Section 4.2 we describe the method for sampling from the resulting Gaussian approximation.

4.1 A Taylor expansion based Gaussian approximation to the posterior

Here, we establish the quadratic approximation to the full conditional distribution of $\boldsymbol{x}_{\mathcal{A}_i}$ analytically by computing the second order Taylor expansion of the log-likelihood

part $h_i(\boldsymbol{x})$ of (19). For the Poisson likelihood (15), $h_i(\boldsymbol{x})$ becomes

$$h_i(\boldsymbol{x}) = y_i \log(E_i' \sum_{j \in \mathcal{A}_i} \exp x_j) - E_i' \sum_{j \in \mathcal{A}_i} \exp x_j. \tag{31}$$

The expansion is computed around a point $\boldsymbol{x}^0_{\mathcal{A}_i}$ taken to be the mode of the full conditional distribution (19), found numerically given the current value of $\boldsymbol{\theta}$. Expressing $h_i(\boldsymbol{x})$ in terms of the gradient (first order derivative) $\boldsymbol{g}_i(\boldsymbol{x}_{\mathcal{A}_i})$ and the Hessian (second order derivative) $\boldsymbol{G}_i(\boldsymbol{x}_{\mathcal{A}_i})$ of the Taylor expansion, we get

$$\begin{aligned} h_i(\boldsymbol{x}) &\approx g(\boldsymbol{x}^0_{\mathcal{A}_i}) + \boldsymbol{g}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i})(\boldsymbol{x}_{\mathcal{A}_i} - \boldsymbol{x}^0_{\mathcal{A}_i}) - \frac{1}{2}(\boldsymbol{x}_{\mathcal{A}_i} - \boldsymbol{x}^0_{\mathcal{A}_i})^T(-\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}))(\boldsymbol{x}_{\mathcal{A}_i} - \boldsymbol{x}^0_{\mathcal{A}_i}) & (32)\\ &= -\frac{1}{2}\boldsymbol{x}^T_{\mathcal{A}_i}(-\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}))\boldsymbol{x}_{\mathcal{A}_i} + (\boldsymbol{g}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i}) - (\boldsymbol{x}^0_{\mathcal{A}_i})^T\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}))\boldsymbol{x}_{\mathcal{A}_i}, & (33)\end{aligned}$$

discarding terms not depending on $\boldsymbol{x}_{\mathcal{A}_i}$. In terms of $\boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$ and $\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$, the matrix $\boldsymbol{B}_i$ and the vector $\boldsymbol{b}_i$ in the quadratic approximation $h_i(\boldsymbol{x}) \approx -\frac{1}{2}\boldsymbol{x}^T_{\mathcal{A}_i}\boldsymbol{B}_i\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{b}_i^T\boldsymbol{x}_{\mathcal{A}_i}$ are given by $\boldsymbol{B}_i = -\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$ and $\boldsymbol{b}_i = \boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) - \boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{x}^0_{\mathcal{A}_i}$, such that the Gaussian approximation (21) is

$$\log(\pi_N(\boldsymbol{x}_{\mathcal{A}_i} \mid \boldsymbol{x}_{-\mathcal{A}_i}, \boldsymbol{\theta}, \boldsymbol{y})) = -\frac{1}{2}\boldsymbol{x}^T_{\mathcal{A}_i}(\boldsymbol{Q}_{\mathcal{A}_i} + (-\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}))\boldsymbol{x}_{\mathcal{A}_i} + \boldsymbol{c}_i^T\boldsymbol{x}_{\mathcal{A}_i} + \text{const}, \tag{34}$$

with $\boldsymbol{c}_i = \boldsymbol{d}_i + \boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) - \boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{x_0}$. Thus, the full conditional distribution for $\boldsymbol{x}_{\mathcal{A}_i}$ is approximated by a Gaussian distribution with mean $\widetilde{\boldsymbol{\mu}}_i$ and precision matrix $\widetilde{\boldsymbol{Q}}_i$, given by

$$\begin{aligned} \widetilde{\boldsymbol{Q}}_i &= \boldsymbol{Q}_{\mathcal{A}_i} - \boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) & (35)\\ \widetilde{\boldsymbol{\mu}}_i &= \widetilde{\boldsymbol{Q}}_i^{-1}(\boldsymbol{d}_i + \boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) - \boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{x}^0_{\mathcal{A}_i}). & (36)\end{aligned}$$

As shown in Appendix A.1, the gradient $\boldsymbol{g}_i(\boldsymbol{x})$ and the Hessian $\boldsymbol{G}_i(\boldsymbol{x})$, evaluated in the mode $\boldsymbol{x}^0_{\mathcal{A}_i}$, are given by

$$\begin{aligned} \boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) &= (\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i')\ \boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) & (37)\\ \boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) &= (\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i')\ \text{diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})) - \frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})^2}\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i}), & (38)\end{aligned}$$

where we have used the definitions

$$\begin{aligned} \boldsymbol{a}_i(\boldsymbol{x}_{\mathcal{A}_i}) &= (\exp(x_j))^T_{j \in \mathcal{A}_i}, & (39)\\ S_i(\boldsymbol{x}_{\mathcal{A}_i}) &= \sum_{j \in \mathcal{A}_i} \exp(x_j) = \sum_{k=1}^{n_i} a_{ik}. & (40)\end{aligned}$$

The mean $S_i(\boldsymbol{x}_{\mathcal{A}_i})/n_i = \frac{1}{n_i}\sum_{j \in \mathcal{A}_i} \exp(x_j)$ is equal to the relative risk in region $\mathcal{A}_i$ conditionally on $\boldsymbol{x}$.

Observe that $\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$ is of the form

$$\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) = -(\boldsymbol{D}_i + \boldsymbol{H}_i) \tag{41}$$

where $\boldsymbol{D}_i$ is a diagonal matrix and $\boldsymbol{H}_i$ a rank one matrix defined by

$$\boldsymbol{D}_i = -(\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i') \text{ diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})) \tag{42}$$

$$\boldsymbol{H}_i = \frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})^2}\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i}) = \boldsymbol{A}_i^T\boldsymbol{A}_i. \tag{43}$$

Here, we have introduced the $1 \times n_i$ matrix $\boldsymbol{A}_i$ given by

$$\boldsymbol{A}_i = \frac{\sqrt{y_i}}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})}\boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i}). \tag{44}$$

Both $\boldsymbol{D}_i$ and $\boldsymbol{H}_i$ depend on the observed count y_i and the point $\boldsymbol{x}^0_{\mathcal{A}_i}$. Substituting the sum $-(\boldsymbol{D_i} + \boldsymbol{H}_i)$ for $\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{A}_i})$ in (34) using the expression in (43) for $\boldsymbol{H}_i$ and re-arranging terms, we arrive at the expression (23) for the proposal distribution. As we will describe in Section 4.2, this re-formulation can be utilised to reduce the computational cost of sampling from (34).

4.2 Sampling algorithm

In Section 4.1 we re-formulated the problem of sampling from the Normal approximation (34) to the full conditional distribution (19) in terms of the general problem of sampling from a distribution on the form

$$\log(\pi(\boldsymbol{x})) = -\frac{1}{2}\boldsymbol{x}^T(\boldsymbol{Q} + \boldsymbol{D})\boldsymbol{x} + \boldsymbol{c}^T\boldsymbol{x} - \frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} + \text{const}, \tag{45}$$

where $\boldsymbol{Q} + \boldsymbol{D}$ is a band matrix. Here, we have suppressed the subscripts $\mathcal{A}_i$ and i, the dependency of $\boldsymbol{x}^0_{\mathcal{A}_i}$ and the conditioning on $\boldsymbol{x}_{\mathcal{A}_i}$, $\boldsymbol{y}$ and $\boldsymbol{\theta}$ for notational convenience.

The first two terms of (45) is the log-density function, up to a constant, of a Gaussian vector variable $\boldsymbol{x}^*$ with mean an precision matrix given by

$$\boldsymbol{\mu}^* = (\boldsymbol{Q}^*)^{-1}\boldsymbol{c}, \tag{46}$$

$$\boldsymbol{Q}^* = \boldsymbol{Q} + \boldsymbol{D}. \tag{47}$$

Rewriting the last term of (45) as

$$-\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = -\frac{1}{2}(\boldsymbol{0} - \boldsymbol{A}\boldsymbol{x})^T\boldsymbol{I}(\boldsymbol{0} - \boldsymbol{A}\boldsymbol{x}), \tag{48}$$

we observe that this term, up to a constant, is equivalent to the multivariate Normal log-density of another vector variable $\boldsymbol{z}^*$ with mean $\boldsymbol{Ax}$ and covariance matrix $\boldsymbol{I}$, that is evaluated in the value $\boldsymbol{z} = \boldsymbol{0}$. Thus, we have introduce two new variables $\boldsymbol{x}^*$ and $\boldsymbol{z}^*$, with distributions given by

$$\begin{aligned} \boldsymbol{x}^* &\sim N(\boldsymbol{\mu}^*, (\boldsymbol{Q}^*)^{-1}) \quad (49) \\ \boldsymbol{z}^*|\boldsymbol{x}^* = \boldsymbol{x} &\sim N(\boldsymbol{Ax}, \boldsymbol{I}), \quad (50) \end{aligned}$$

where $\boldsymbol{\mu}^*$ and $\boldsymbol{Q}^*$ are defined by (46) and (47). In terms of the variables $\boldsymbol{x}^*$ and $\boldsymbol{z}^*$, the distribution (45) is equivalent to the conditional distribution of $\boldsymbol{x}^*$ given $\boldsymbol{z}^* = \boldsymbol{0}$, evaluated in $\boldsymbol{x}^* = \boldsymbol{x}$. We denote this distribution by $\pi_{\boldsymbol{x}^*|\boldsymbol{z}^*}(\boldsymbol{x}^*|\boldsymbol{z}^* = \boldsymbol{0})$, and it can be expressed by

$$\pi_{\boldsymbol{x}^*|\boldsymbol{z}^*}(\boldsymbol{x}|\boldsymbol{z}^* = \boldsymbol{0}) \propto \pi_{\boldsymbol{x}^*}(\boldsymbol{x})\ \pi_{\boldsymbol{z}^*|\boldsymbol{x}^*}(\boldsymbol{0}|\boldsymbol{x}^* = \boldsymbol{x}), \quad (51)$$

using a compact notation. Consequently, sampling from (45), and thus the Gaussian proposal distribution (21), is equivalent to sampling from the conditional distribution given by (51). Since

$$\boldsymbol{z}^* = \boldsymbol{Ax}^* - \boldsymbol{\epsilon}^*; \ \ \boldsymbol{\epsilon}^* \sim N(\boldsymbol{0}, \boldsymbol{I}), \quad (52)$$

conditioning on $\boldsymbol{z}^* = \boldsymbol{0}$ is equivalent to conditioning on $\boldsymbol{Ax}^* = \boldsymbol{\epsilon}^*$. In our application, we have that

$$\boldsymbol{Ax} = (S_i(\boldsymbol{x}^0_{\mathcal{A}_i})^2/y_i) \sum_{j \in \mathcal{A}_i} \exp(x_{0,j})\ x_j \quad (53)$$

is a weighted sum of the lattice specific log relative risks x_j within each region. Conditioning on $\boldsymbol{Ax}^* = \boldsymbol{\epsilon}^*$ can be interpreted as generating samples for which $\mathrm{E}(\boldsymbol{Ax}^*) = \boldsymbol{0}$, and where the elements of $\boldsymbol{Ax}^*$ should be independent Gaussian variables with common variance 1.

To generate a sample $\boldsymbol{x}_c$ from the conditional distribution $\pi(\boldsymbol{x}^*|\boldsymbol{Ax}^* = \boldsymbol{\epsilon}^*)$, we use the approach given by equation (12) in Section 2.2. We first generate an *unconditional* sample $\boldsymbol{x}_u$ from $\pi(\boldsymbol{x}^*)$ and an $\epsilon^* \sim N(\boldsymbol{0}, \boldsymbol{I})$, and then compute $\boldsymbol{x}_c$ by adjusting for the constraint $\boldsymbol{Ax}^* = \boldsymbol{\epsilon}^*$ using the expression

$$\boldsymbol{x}_c = \boldsymbol{x}_u - (\boldsymbol{Q}^*)^{-1}\boldsymbol{A}^T(\boldsymbol{A}(\boldsymbol{Q}^*)^{-1}\boldsymbol{A}^T + \boldsymbol{I})^{-1}(\boldsymbol{Ax}_u - \boldsymbol{\epsilon}^*) \quad (54)$$

The precision matrix $\boldsymbol{Q}^*$ has the same bandwidth as the prior precision matrix $\boldsymbol{Q}$ of $\boldsymbol{x}$. Utilising the band structure of $\boldsymbol{Q}^*$, samples from (51) can be generated efficiently using the methods described and implemented in Rue and Follestad (2002).

To compute (54) we need to evaluate the matrix expression $(\boldsymbol{A}(\boldsymbol{Q}^*)^{-1}\boldsymbol{A}^T + \boldsymbol{I})^{-1}$. The matrix $\boldsymbol{A}(\boldsymbol{Q}^*)^{-1}\boldsymbol{A}^T + \boldsymbol{I}$ is in general a full matrix, but it is of dimension m_c by m_c, where m_c is the number of rows in $\boldsymbol{A}$. So far we have considered the sampling problem updating one region at a time, for which $m_c = 1$. But even for generalisations to larger subsets of regions, we usually have that $m_c \ll n$.

As pointed out in Section 3, the matrix $\boldsymbol{Q} + \boldsymbol{D}$ should be positive definite for the Gaussian distribution defined by (46) and (47) to be proper. The matrix $\boldsymbol{Q}^*$ is positive definite iff $\boldsymbol{x}^T\boldsymbol{Q}^*\boldsymbol{x} > 0;\ \forall \boldsymbol{x} > 0$. Substituting $\boldsymbol{Q} + \boldsymbol{D}$ for $\boldsymbol{Q}^*$ we get

$$\begin{aligned} \boldsymbol{x}^T\boldsymbol{Q}^*\boldsymbol{x} &= \boldsymbol{x}^T\boldsymbol{Q}\boldsymbol{x} + \boldsymbol{x}^T\boldsymbol{D}\boldsymbol{x} \\ &= \boldsymbol{x}^T\boldsymbol{Q}\boldsymbol{x} - (\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i')\boldsymbol{x}^T\text{diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i}))\boldsymbol{x}. \end{aligned} \tag{55}$$

Since all elements of $\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$ are strictly positive, the sign of the last term on the right is determined by the sign of the factor

$$\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i' = (y_i - S_i(\boldsymbol{x}^0_{\mathcal{A}_i})E_i')/S_i(\boldsymbol{x}^0_{\mathcal{A}_i}) = (y_i - \text{E}(y_i \mid \boldsymbol{x}^0_{\mathcal{A}_i}))/S_i(\boldsymbol{x}^0_{\mathcal{A}_i}). \tag{56}$$

This term can in general be of either sign, such that the precision matrix $\boldsymbol{Q}^*$ is not guaranteed to satisfy the positive definiteness requirement. In the case where the prior variance $1/\tau$ is large, such that $\boldsymbol{x}^T\boldsymbol{Q}^*\boldsymbol{x}$ is small, the second term of (55) will dominate, and the chance is higher that the positive definiteness requirement is not met. However, this problem can in general be dealt with by a slight modification of our sampling algorithm, replacing the diagonal elements $d_{j,j}$ of the matrix $\boldsymbol{D}$ by $\max(d_{j,j}, 0.0)$. The corresponding change in proposal distribution is corrected for by the acceptance probability of the MCMC algorithm.

To summarise the sampling approach, a sample from the conditional distribution (45) can be obtained by the following steps:

1. Sample a value $\boldsymbol{x}_u$ from the unconditional distribution (49).
2. Sample an $\boldsymbol{\epsilon}^* \sim N(\boldsymbol{0}, \boldsymbol{I})$.
3. Compute $\boldsymbol{x}_c$ using (54). Then $\boldsymbol{x}_c$ will be a sample from the posterior distribution given by (51), and consequently from (45).

4.3 Some computational details

The sampling algorithm is implemented in `C`, and is based on the routines for fast and exact simulation of Gaussian Markov random fields implemented in the library `GMRFLib` (Rue and Follestad, 2002). The library provides general algorithm for generating samples from a GMRF, including conditional samples for hard and soft linear constraints, and the algorithms are based on the Cholesky factorisation (9) of the precision matrix $\boldsymbol{Q}$ (Rue, 2001).

When applying the `C`-routines to our problem, we have utilised the structure of the specific problem to further reduce the computational cost. Using the fact that the range of

the correlation function is given for a set of discrete values, the normalising constant of the posterior distribution of $\boldsymbol{x}$ can be computed once at the beginning of a run, and stored for later use. Also, the sub-graphs representing the subset of nodes in each block in the block-updates of the log-risk surface $\boldsymbol{x}$, as well as the sub-graph representing the boundary region nodes, are computed only once. A timing of the computer program, running the program for 1000 iterations, reveals that 32% of the time is spent evaluating the elements of the precision matrix $\boldsymbol{Q}$, and 12% in evaluating the log-likelihood. Further, about 21% of the CPU time is spent on setting up the algorithm, including specifying the neighbourhood structure and computing sub-graphs. Consequently, more CPU time is spent on setting up the problem than actually performing the computations generating the samples.

5 Simulation study

In this section we illustrate the performance of the method by applying the sampling algorithm to two simulated data sets, both generated using the study region of the real data. A plot of the study region made up from the 544 districts of Germany, overlaid by a lattice consisting of $n = 31089$ nodes including boundary nodes, were shown in Figure 5.

The simulated data as well as the real data set are standardised such that the overall risk for the region of study is 1. Therefore, no intercept term is included in the model for the log-risk surface, and the prior mean $\boldsymbol{\mu}$ of $\boldsymbol{x}$ is taken to be $\mathbf{0}$. The remaining hyper-parameters of the GMRF prior are the precision τ and the parameters specifying the spatial structure of the random field. We model the spatial dependency using an isotropic one-parameter exponential correlation function given by

$$\rho(h; r) = \exp(-3h/r). \tag{57}$$

Here h is the distance between two nodes of the lattice, and r is the distance for which the correlation is reduced to 0.05. In the simulated data sets, the range parameter of the exponential correlation function (57) is set equal to $r = 40$ measured in lattice coordinates for both data sets. The precisions are taken to be different, and the chosen values are $\tau = 24$ and $\tau = 8$, corresponding to standard deviations of 0.20 and 0.35 respectively. For each data set, we first generate a realisation of the log-risk surface $\boldsymbol{x}$ from the GMRF prior (7), and conditionally on $\boldsymbol{x}$ a set of regional count data is sampled from the Poisson distribution given by the likelihood (15). We define the expected number of cases E_i to be the ones given in the data set used in Section 6, ranging from 3.0 to 393.1 and with a median of 19. A summary of the two simulated data sets used in the study is given in Table 2, and the realisations of $(\exp(x_j))_{j=1,\ldots,n}$ and the corresponding regional relative risks, given by the mean $(\sum_{j \in \mathcal{A}_i} \exp(x_j))/n_i$ over the n_i lattice nodes within region i, are shown in Figure 8.

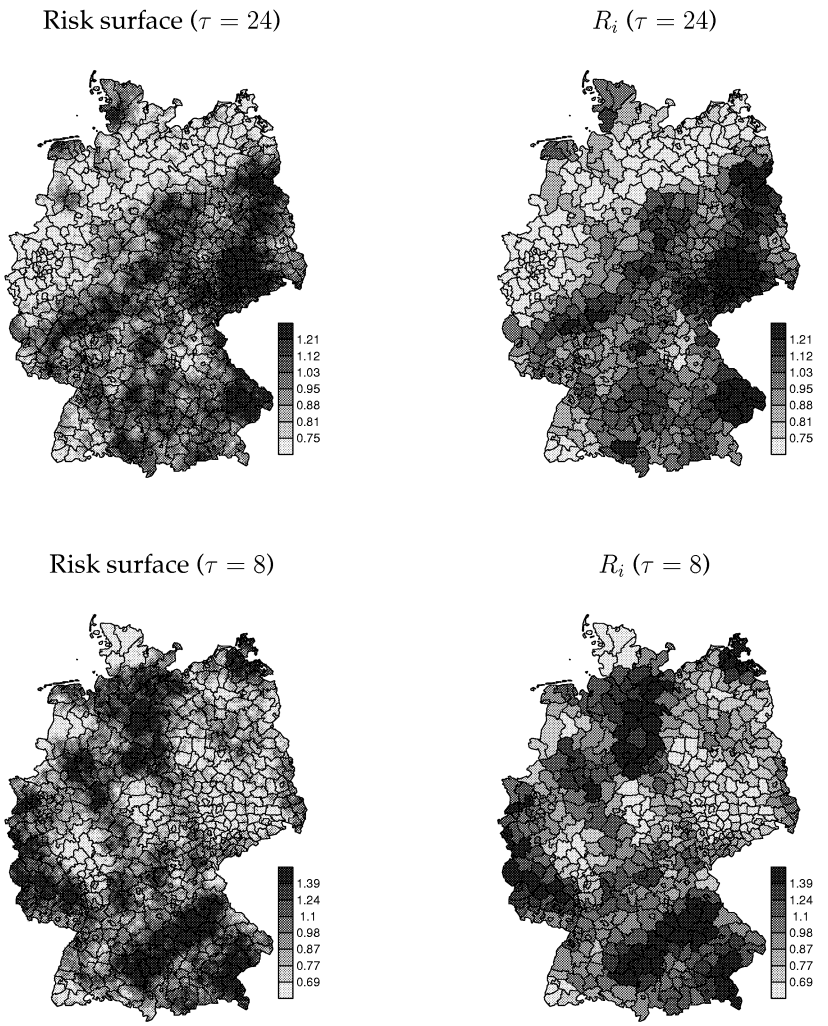

Figure 8: The true risk surfaces and corresponding regional level relative risks for the two simulated data sets.

			Aggregated counts (y_i)				
Data set	τ	r	Min.	2.5% quantile	Median	97.5% quantile	Max.
I	24	40	0	3	19	87	403
II	8	40	1	2	18	111	461
			Relative risk (R_i)				
Data set	τ	r	Min.	2.5% quantile	Median	97.5% quantile	Max.
I	24	40	0.50	0.61	0.98	1.43	1.85
II	8	40	0.45	0.49	095	1.71	2.60

Table 2: The aggregated counts and the true relative risks of the simulated data sets. The quantiles are given as the empirical quantiles of the simulated values.

The prior distribution for the precision τ and the range parameter r are assumed to be independent. To reduce the computational burden, we use a discrete prior distribution for the range parameter r, such that the determinant of the precision matrix, needed for the evaluation of normalising constants, can be computed once at the beginning of the sampling procedure. The discretisation is done in $n_r = 2001$ steps $r_k;\ k = 1, \ldots, n_r$, where the range at step k is equal to $r_k = (k-1)0.05$ measured in lattice coordinates. The discrete prior distribution is defined on the indexes k, such that

$$\pi(k) \quad \propto \quad \frac{1}{k};\ k = 1, \ldots, n_r. \tag{58}$$

The precision τ, which is constrained to be positive, is assigned a Gamma prior, $\tau \sim$ Gamma$(\alpha_\tau, \beta_\tau)$. Based on the recommendations in Kelsall and Wakefield (1999) and the discussion of prior sensitivity in Pascutto, Wakefield, Best, Richardson, Bernardinelli, Staines and Elliott (2000), we choose the parameters of the Gamma$(\alpha_\tau, \beta_\tau)$-prior for τ such that more weight is given to small variances than the Gamma(ϵ, ϵ)-prior for small ϵ frequently used in this type of applications. Specifically, we choose $\alpha_\tau = 0.2$ and $\beta_\tau = 0.0002$. For the range parameter r we use the discrete prior (58) on the range indexes $k = 1, 2, \ldots, 2001$ corresponding to the values $0.0, 0.05, 0.1, 0.15, \ldots, 100.0$ of the range, as measured in lattice coordinates.

The elements of $\boldsymbol{x}_{\mathcal{A}}$, representing the log relative risk surface within the 544 regions, are updated using the block-sampling approach described in Section 3.2. As pointed out in that section, the optimal choice of block-size can be considered to be a trade-off between computational cost and the acceptance probabilities of the Metropolis-Hastings steps. To study the effect of changing the block-size on the acceptance probabilities, we ran 11000 iterations of the sampler on data set I of Table 2 for four different choices of blocks, keeping the hyper-parameters fixed at their true values. The blocks are made up from single regions, 1. order neighbourhoods, 2. order neighbourhoods and 3. order neighbourhoods respectively, using the neighbourhood definitions given in Section 3.2

and Figure 7. The four different block sizes are also illustrated in Figure 9. In our sampling algorithm, the blocks are slightly modified such that the different blocks of one run of the sampler are disjoint, and such that regions with only one neighbour are added to one of the adjacent blocks. This last modification applies to city regions, like two regions within the 2. order neighbourhood block of Figure 9, as well as some of the regions at the boundary. To avoid boundary effects between blocks, we generate the blocks randomly, updating the partition into blocks at every 10th step of the sampler.

The resulting acceptance rates for the four different choices of the block structure are displayed in Figure 10. We observe that the acceptance rates for the single region blocks are very large, with a median acceptance probability of 0.95, indicating that the Gaussian approximation is a good approximation to the posterior distribution (19). For the blocks based on 1., 2. and 3. order neighbourhoods, the median acceptance probabilities are gradually decreased, taking the values 0.66, 0.35 and 0.16 respectively. The acceptance probabilities seem to be independent of the size of the regions, represented by the number of lattice nodes within the region, but they increase as the mean of the regional level risk approaches 1.0. This result is as expected, since the Gaussian approximation (27) to the posterior distribution (25) of $\boldsymbol{x}_S$ is expected to be better when the values of $\boldsymbol{x}_S$ are small, and thus the corresponding regional level relative risk close to 1.0. Based on these results, we choose to use blocks made up from a region and its 1. order neighbourhood.

The convergence of the MCMC algorithm is assessed by visual inspection of trace plots. The total number of parameters of the risk surface is too large for an inspection of all trace plots to be feasible. So in addition to the hyper-parameters τ and r, we study trace plots of the relative risk R_i of a selected number of regions, and for a subset of the corresponding elements of the log-risk vector $\boldsymbol{x}$. The values of the regional relative risks at iteration k, denoted $R_i^{(k)};\ i = 1, \ldots, m$, are generated from the current values of $\boldsymbol{x}$ by

$$R_i^{(k)} = \frac{1}{n_i} \sum_{j \in \mathcal{A}_i} \exp(x_j^{(k)}), \tag{59}$$

where $x_j^{(k)}$ is the k'th update of x_j. To get an impression of the behaviour of the algorithm for the remaining regions, we compute the mean acceptance probability of the Metropolis-Hastings steps for all regions.

In Figure 11, we show a subset of trace plots for data set I, after running the MCMC algorithm for 101000 iterations. The convergence is fast and the algorithm mixes well for the majority of the relative risk estimates, but the trace plots for region 16 indicate that although the convergence seems to be fast, the mixing is relatively poor for this region. The mean acceptance probability for the corresponding elements of $\boldsymbol{x}$ is 1.9%. The mean acceptance probabilities for all regions are plotted in the top panels of Figure 12, and the acceptance rate for region 16 is seen to be the lowest among the 544 regions, for which the second smallest value is 8.9%. Region 16 corresponds to a region with a small true

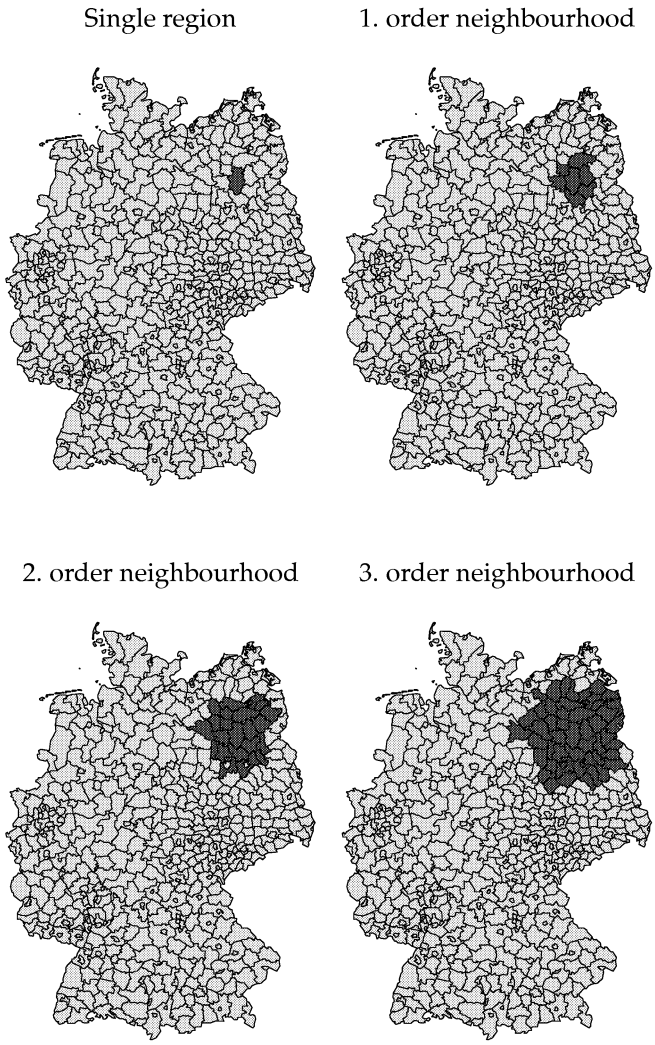

Figure 9: An illustration of the blocks used in the block-sampling of the log-risk surface. The risk surface for the lattice nodes within the dark shaded regions are updated conditionally on the nodes of all remaining regions as well as the boundary nodes.

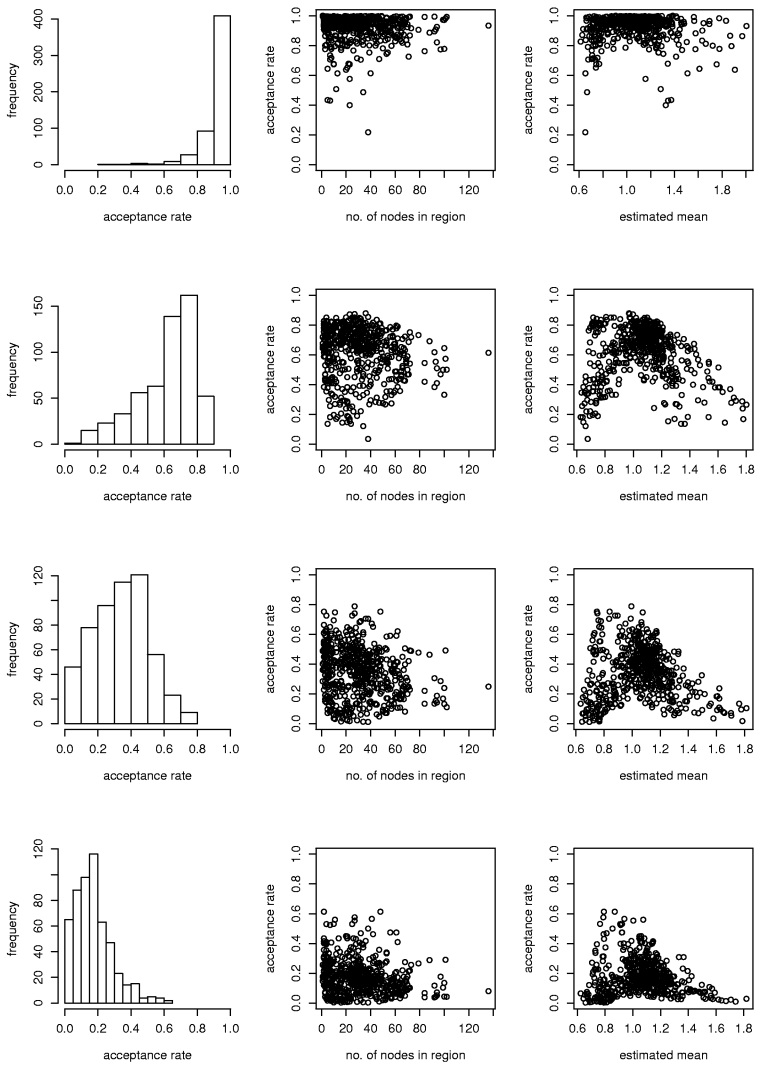

Figure 10: Histograms of the mean acceptance probabilities of the log-risk of the 544 regions using the block-MCMC algorithm (left), the mean acceptance probabilities plotted against the number of nodes within the region (middle), and the same values plotted against the estimated log-risk (left). The plots are given for blocks of single regions and for 1. order, 2. order and 3. order neighbourhoods (from the top and downward), and are based on results from using data set I.

($R_{16} = 0.61$) as well as estimated ($\widehat{R}_{16} = 0.65$) relative risk and a large aggregated count ($y_{16} = 202$). From Table 3 we observe that the true risk is relatively similar for region 16 and its neighbours, but the expected and observed aggregated counts are an order of magnitude larger. From Table 2 it is clear that the observed count of region 16 is in the tail of the empirical distribution of the observed counts, and this might explain why the Gaussian approximation is relatively poor for the elements of the log-risk within this region. (Region 16 includes Hamburg, and since we use the expected counts of the German oral cavity cancer data to generate our simulated data set, this explains the high count of this region).

Region no. (i)	y_i	E_i	R_i	SMR
16	**202**	**314.9**	**0.61**	**0.64**
6	17	30.7	0.67	0.55
9	45	50.9	0.62	0.88
13	30	39.9	0.72	0.75
15	22	38.5	0.70	0.57
38	18	38.1	0.63	0.47
44	28	30.6	0.73	0.91

Table 3: The expected (E_i) and observed (y_i) aggregated counts, true relative risks (R_i) and SMR for region 16 and its neighbours for data set I.

The mixing for τ and r is poorer, but the trace plot of τ indicates that the algorithm has converged for this parameter. For the range parameter r, the mixing is not uniform over the range of possible values. The poor mixing for some neighbouring values for r is due to larger differences between the corresponding neighbouring prior models than the typical differences between neighbouring models over the range of values of r. This is a result of the procedure used for fitting GMRFs to GRFs (Rue and Tjelmeland, 2002). The additional constraint that the coefficients of the precision matrix of the GMRF, computed for each value of the range, should also be near continuous with respect to the range, is not accounted for in the fitting procedure. The effect could be reduced by increasing the resolution for the range r in (58), but probably we need to add explicit smoothing constraints of the parameters with respect to range in the fitting procedure.

To illustrate the effect of slow convergence and mixing of the hyper-parameters on the estimates of the log-risk surface, we have plotted the updated values of some elements of $\boldsymbol{x}$ against corresponding values for τ and r. From the resulting scatter plots shown in the top two rows of Figure 13, we observe that the posterior variance decreases for increasing values of the precision τ, but the posterior means of the elements of $\boldsymbol{x}$ appear to be stable despite the poor mixing of the individual parameters τ and r. Therefore, we proceed by presenting results for the relative risk surface based on estimated posterior means, but the poor mixing of the hyper-parameters should be kept in mind.

We discard the first 1000 iterations and use the remaining 100000 iterations to compute estimated posterior mean values of the relative risk surface and the relative risks within each region. The results for data set I are reported in Figure 14. From the top and middle left panels we observe that the simulation algorithm reproduces the structure of the simulated risk surface well. The corresponding true and estimated values of the regional level risks are plotted in the top and middle right panels. Comparing the estimated values to the standardised mortality ratio (SMR) added in the bottom right panel, the algorithm is seen to smooth the disease map based on the SMR toward the true risk surface. In the bottom left panel we have plotted the estimated probability that the risk R_i exceeds 1.

A selection of trace plots and the estimated posterior mean values for data set II are given in Figures 15 and 16. The general pattern is similar to the results from data set I. The mixing is relatively good and the convergence is fast for the log-risk surface for most regions, but for the hyper-parameters, convergence is not achieved after the 101000 iterations. However, there are more regions for which the mixing for the corresponding elements of $\boldsymbol{x}$ is relatively poor. In Figures 15, we have included trace plots for the regional level relative risk and two corresponding elements of $\boldsymbol{x}$ for a region for which the acceptance rate is extremely low (0.09%). This region is the same as the one with lowest acceptance rate for data set I. In the sampling algorithm we have used the same block-size as for data set I, producing the mean acceptance probabilities illustrated in the middle panels of Figure 12. We observe that the acceptance rates are in general lower than for data set I, and the lowest for the regions with the most extreme values of the risk. This overall decrease in acceptance probabilities corresponds to the fact that the Gaussian approximation is a better fit to the posterior distribution for the data set with the smaller variance. To increase the average acceptance probability, the block-sizes should be reduced to include single regions only for this simulated data set, and in the further discussion of the results, the low acceptance rates for some of the regions should be kept in mind.

As for data set I the posterior mean level of the elements of $\boldsymbol{x}$ seems to be stable despite the convergence problems apparent for the hyper-parameters, as illustrated by the scatter plots in the bottom two rows of Figure 13. For region number 16, there is some indications of negative association between τ and the estimated level of the log-risk, but as we pointed out above, the acceptance rate is very low, and we know from Figure 15 that the mixing is poor for the elements of $\boldsymbol{x}$ within this region. Therefore, the results presented below, based on estimated posterior means, are hoped to be representative despite the poor mixing of the hyper-parameters.

Comparing the differences between the SMR and the estimated risk surface for the two data sets, we observe from Figures 14 and 16 that degree of smoothing is less pronounced for the case $\tau = 8$ than for the data set with $\tau = 24$. Thus, increasing the prior variance of the underlying risk surface seems to reduce the degree of smoothing. This effect is also illustrated in Figure 17, where we have plotted the differences between

the estimated and true relative risks together with corresponding differences between the estimated risks and the observed SMR. Some of the larger differences for data set II correspond to regions for which the acceptance probabilities are small and the mixing relatively poor, but comparing similar differences discarding risk estimates for which the acceptance probabilities are all larger than 20%, the tendency is similar.

We end this section by an illustration of how we can assess the validity of the approximation

$$\frac{1}{n_i}\sum_{j\in\mathcal{A}_i}\exp(x_j) \approx \exp(\frac{1}{n_i}\sum_{j\in\mathcal{A}_i}x_j), \tag{60}$$

which is an analogue to the approximation $\log(R_i) = \int_{\mathcal{A}_i}\log R(s)f_i(s)ds$, underlying the geostatistical approach of Kelsall and Wakefield (2002). Let

$$\tilde{R}_i \;\; = \;\; \exp(\frac{1}{n_i}\sum_{j\in\mathcal{A}_i}x_j), \tag{61}$$

be the approximation to the relative risk $R_i = \frac{1}{n_i}\sum_{j\in\mathcal{A}_i}\exp(x_j)$ of region $\mathcal{A}_i$, and let further $\widehat{\widetilde{R}}_i = \overline{\widetilde{R}_i^{(\cdot)}}$ and $\widehat{R}_i = \overline{R_i^{(\cdot)}}$ be the corresponding posterior mean estimates based on the updates $\widetilde{R}_i^{(k)}$ and $R_i^{(k)}; k = 1,\ldots,100000$ from the Metropolis-Hastings algorithm. By Jensens inequality

$$\exp(\frac{1}{n_i}\sum_{j\in\mathcal{A}_i}x_j) \leq \frac{1}{n_i}\sum_{j\in\mathcal{A}_i}\exp(x_j), \tag{62}$$

such that the estimated posterior means of $\widetilde{R}_i$ should be smaller than or equal to the corresponding values for R_i. In Figure 18 we have plotted $\widehat{\widetilde{R}}_i$ against $\widehat{R}_i$ and $\widehat{\widetilde{R}}_i/\widehat{R}_i$ as a function of the number of lattice nodes within each region, for data sets I and II as well as for the oral cavity cancer data analysed in the next section. We observe that (61) is a good approximation to R_i. As expected, the approximation is better the smaller the number of lattice nodes within the region, and in accordance with Jensens inequality, $\widehat{\widetilde{R}}_i \leq \widehat{R}_i; \forall i$. In the bottom panels of Figure 18, the variability of the updates of the fraction $\widetilde{R}_i^{(k)}/R_i^{(k)}$ is illustrated by plotting histograms of updated values for a region with a relatively large number (53) of lattice nodes. We observe that the variability is largest for data set II, with a minimum value of about 0.92. The mean acceptance probabilities for the risk updates for this region are 0.56, 0.23 and 0.42 for data set I, data set II and the oral cavity cancer data respectively.

6 Oral cavity cancer data

We apply our estimation approach to a set of data on mortality from oral cavity cancer for males in Germany, over the period 1986-1990. We do not intend to do a thorough

analysis of these data, but include the analysis to illustrate the method as applied to a set of real data. The data are given as counts for each of the 544 districts of Germany. The counts range from 1 to 501, with a median count of 19, and the empirical 2.5% and 97.5% quantiles of the observed counts are 3 and 124. The standardised mortality ratios (SMR) for the data were shown in the right panel of Figure 1. The data were analysed by Knorr-Held and Raßer (2000) who identified clusters of elevated or lowered risk using a Bayesian approach based on reversible jump MCMC.

From the bottom panels of Figure 19, we observe that as for the simulated data, the mixing of the hyper-parameters is relatively poor. There is some evidence that the algorithm has converged after about 40000 iterations, but more effort is needed to get reliable estimates of the hyper-parameters using a reasonable amount of computational effort. From trace plots of a selected number of elements of $\boldsymbol{x}$, some of which are plotted in Figure 19, we observe that the mixing is good despite the poor mixing of the hyper-parameters, and the convergence is fast. The acceptance rates are reasonably high for all but a few regions, as illustrated in the bottom panels of Figure 12. The data for the regions for which the mean acceptance probabilities of the log-risk updates are less than 10% are listed in Table 4, and we observe that they all have a relatively large or small SMR or a high observed count, one of which is the maximum observed count (501).

Region no. (i)	y_i	E_i	SMR
197	111	73.0	1.52
322	117	72.9	1.60
324	53	30.8	1.72
328	501	393.1	1.27
414	52	98.5	0.53
443	15	28.1	0.53

Table 4: The expected (E_i) and observed (y_i) aggregated counts SMR for the regions for which the acceptance rates of the log-risk updates for the oral cavity cancer data are less than 10%.

The results from applying our GMRF approach to the data, using blocks made up from regions and their 1. order neighbours, are summarised in Figure 20. The estimated log-risk surface and the corresponding estimated posterior means of the regional relative risks are shown in the upper two panels. The results can be compared to the standardised mortality ratios (SMR) shown in the bottom right panel. We observe that the overall spatial pattern of the estimated relative risk and the SMR are similar, with elevated risk in the north-eastern and south-western parts, but that the estimated spatial risk surface is smoother. The estimated posterior mean relative risks at the regional level vary between 0.57 and 1.54. The results are similar to the ones obtained by Knorr-Held and Raßer (2000). They reported estimated posterior median relative risks in the range 0.65

and 1.42 using their Bayesian cluster detection approach, and between 0.56 and 1.56 using the method of Besag et al. (1991). The estimated spatial pattern is similar to theirs, but their Bayesian clustering approach leads to a somewhat smoother map. However, the smoothness of the map using our approach will depend on the range parameter r, and since the convergence can be questioned, the result should be interpreted with care.

7 Discussion

We have presented an approach to estimation of a spatially varying risk surface based on aggregated count data, using a Gaussian Markov random field prior defined on a lattice. The method is exact in the sense that the posterior mean estimates are generated on the basis of samples from a Markov chain that converges to the correct posterior distribution. This represents an improvement over the geostatistical approach of Kelsall and Wakefield (2002) using a log-Normal approximation to the regional relative risk, in particular in applications for which the regions representing the level of aggregation of the data vary substantially in size and shape. As illustrated in Section 5, for the regions of our study the approximation gives very similar results, but in general the approximation should be justified for the actual set of regions at hand.

We are still left with the problem of convergence of the MCMC sampling algorithm. For the simulated examples and the data set analysed in Sections 5 and 6 the convergence is fast for the elements of the log-relative risk surface $\boldsymbol{x}$, and the mixing is good except for elements of $\boldsymbol{x}$ corresponding to extremes within the range of the relative risks. The acceptance rates for the Metropolis-Hastings sampler are increased by reducing the size of the blocks in the block-MCMC algorithm, at the expense of increased computational cost. For the hyper-parameters, the mixing turned out to be relatively poor, but the estimated posterior means of the elements of the log-risk surface seemed to be stable despite the poor mixing of the individual hyper-parameters. Using a single site Metropolis-Hastings sampling approach, convergence and mixing is often improved by re-parameterisation, but this will have less effect in our case, since we already accept or reject the proposed values of the range parameter r and the precision τ jointly. We chose to block the hyper-parameters with the boundary nodes, an approach that was shown to be equivalent to sampling τ and r from the marginal posterior distribution of (τ, r), integrating over the boundary nodes. To study the effect of blocking on the mixing of the hyper-parameters, other blocking strategies, like including the nodes corresponding to a random sample of inner regions in the block, could be explored. We ran the sampling algorithm including a randomly chosen inner region and it's 1. order neighbourhood in the block, but no improvement in mixing of r or τ was gained.

We have illustrated our approach using the exponential correlation function to specify the spatial correlation structure. This could be replaced by alternative, more flexible

classes of models, like the Matérn class, based on Bessel functions. In Hrafnkelsson and Cressie (2003) a simpler alternative approach to that of Rue and Tjelmeland (2002) is proposed to fit a GMRF to a geostatistical GRF model using the Matérn class of correlation functions.

The commonly used log-Gaussian random effects model for the regional level relative risk, as given by (1), includes a spatially unstructured as well as spatially structured effect, such that the degree of spatial dependency can be assessed by studying the relative values of the estimated precisions of the two effects. A spatially un-structured effect can also be introduced our model, and the proposed sampling based approach to parameter estimation can be applied to the resulting model after a re-parameterisation adding another level to the hierarchical model (see Knorr-Held and Rue, 2002).

As the methods of Best et al. (2000), using a Poisson-Gamma model with identity link, and Kelsall and Wakefield (2002) using similar distributional assumptions as in our model, our method is aggregation consistent, such that the estimated spatial structure is independent of the level of aggregation of the data. Also, the method can be extended to include covariates observed at different non-nested levels of aggregation, using all covariates at their original level of aggregation. This is an appealing feature, since ethological studies often involves data observed at the individual level, as point observations and as aggregated data.

The results from applying our approach to the German oral cavity cancer data turned out to be very similar to those reported by Knorr-Held and Raßer (2000). A closer look at the resulting risk surface displayed in the top left panel of Figure 20, reveals an apparent difference between the general level of the risk in the former German Democratic Republic (GDR), including Eastern Berlin, and Western Germany (BRD). This could be due to different routines for reporting cases, and the effect could be taken into account by including an indicator variable representing former country (GDR or BRD) as a covariate of the model.

We conclude that using GMRFs as proxies for GRFs on a lattice allows for the development of an aggregation consistent approach to estimating a smoothly varying risk surface based on aggregated count data. Applying the approach to simulated data as well as a set of real data using computationally efficient block-MCMC algorithms for parameter estimation, we have shown that the method reproduces the risk surface well. Despite blocking the hyper-parameters with the boundary nodes, further work seems to be needed to improve mixing and convergence of the hyper-parameters.

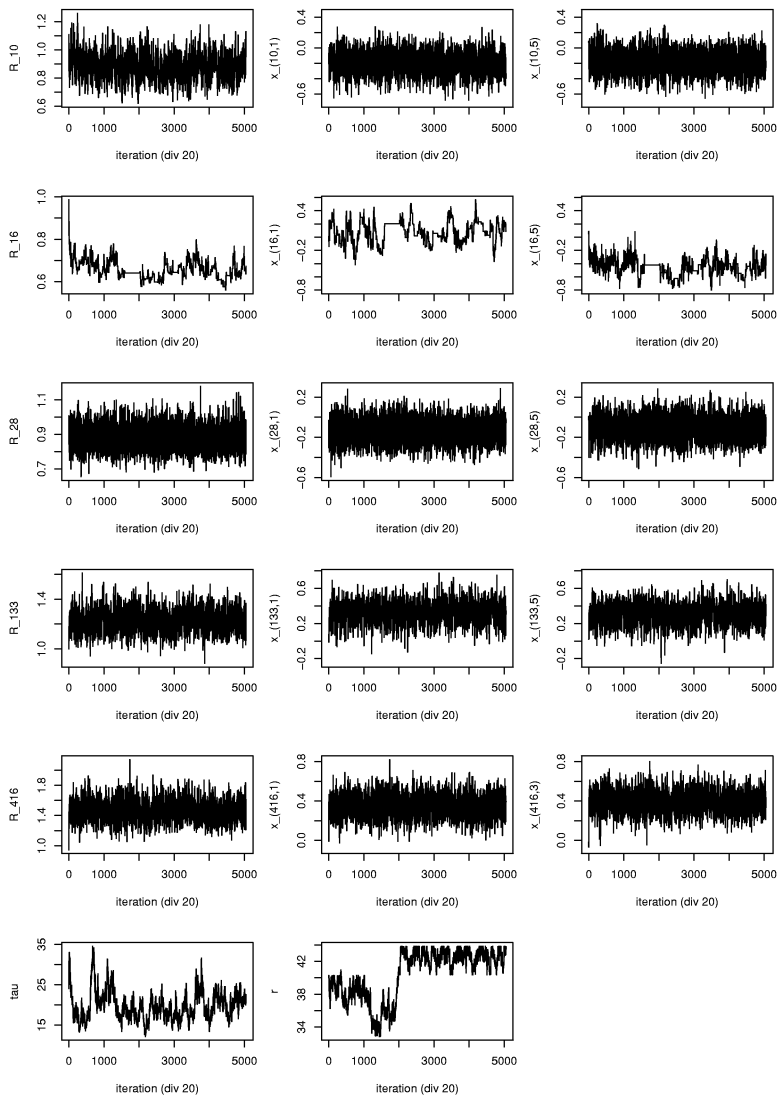

Figure 11: Selected trace plots for the simulated data set I, with $\tau = 24$. The five top rows show trace plots of the regional risk R_i for five regions (left) and of two elements of $\boldsymbol{x}$ falling within each region (middle and right). Every 20th iteration is shown.

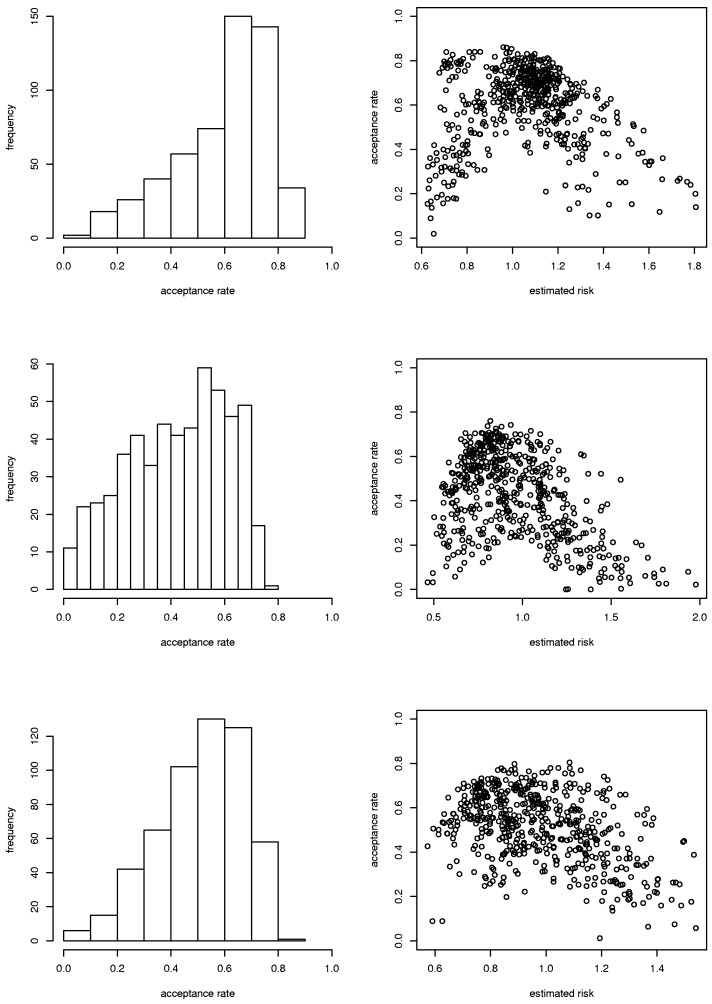

Figure 12: The acceptance rates for the block-MCMC algorithm for (from the top downward) data set I, data set II and the oral cavity cancer data. The left panels show histograms of the acceptance rates of the log-risk $\boldsymbol{x}$ within each region, and in the right panels the acceptance rates are plotted against the estimated risk.

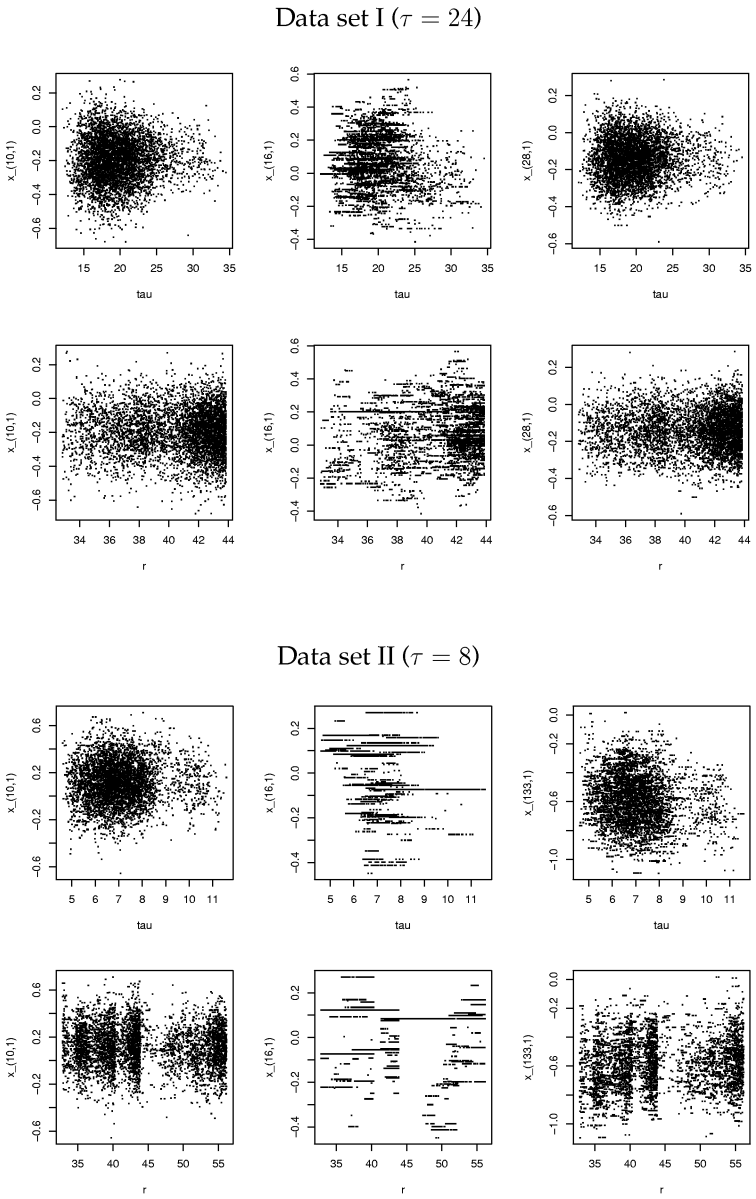

Figure 13: Scatter plots of the hyper-parameters and elements of $\boldsymbol{x}$ for data set I (top rows) and data set II (bottom rows). The numbering of the selected elements of $\boldsymbol{x}$ corresponds to the numbering in Figure 11 and Figure 15, respectively.

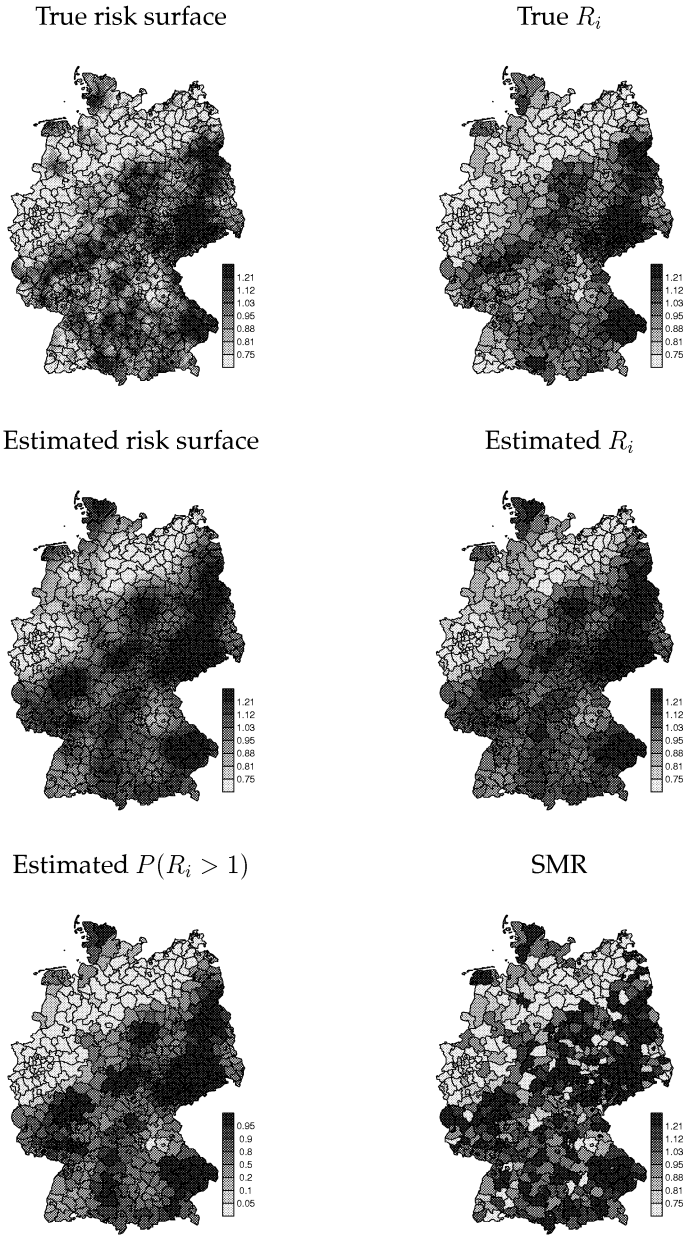

Figure 14: Results for the simulated data set I, with $\tau = 24$ and $r = 40$. The true values of the risk surface and regional risks, also shown in Figure 8, are added for reference.

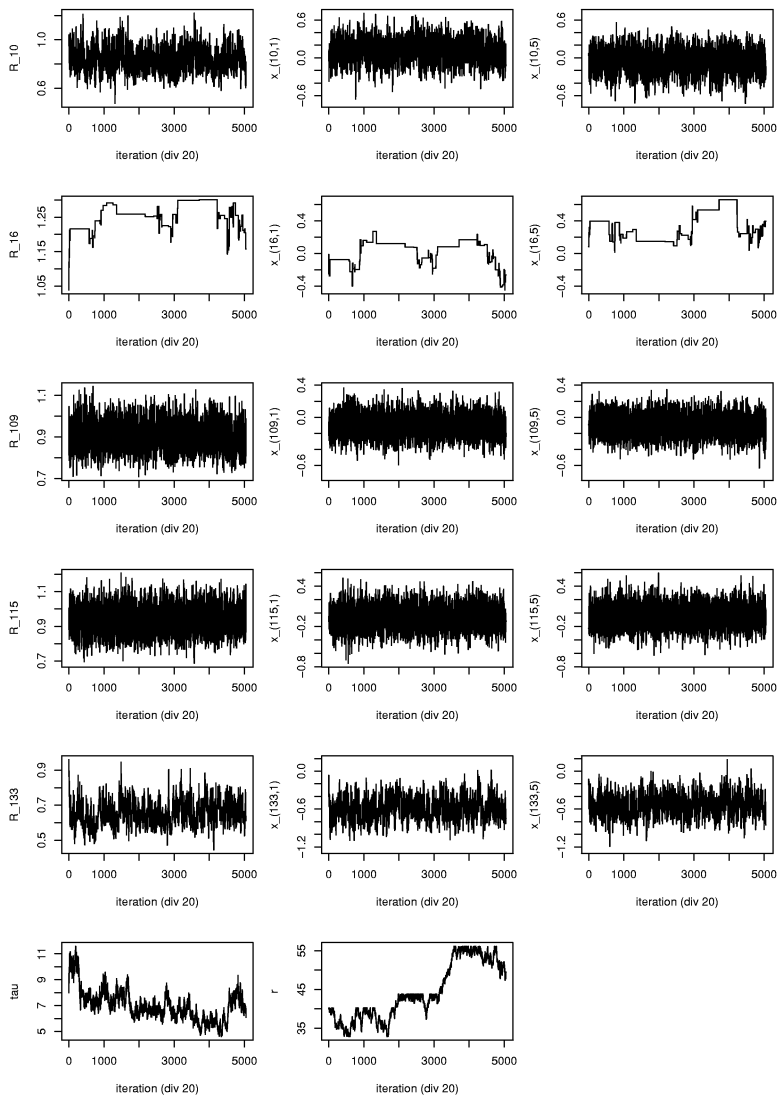

Figure 15: Selected trace plots for the simulated data set II, with $\tau = 8$. The five top rows show trace plots of the regional risk R_i for five regions (left) and of two elements of $\boldsymbol{x}$ falling within each region (middle and right). Every 20th iteration is shown.

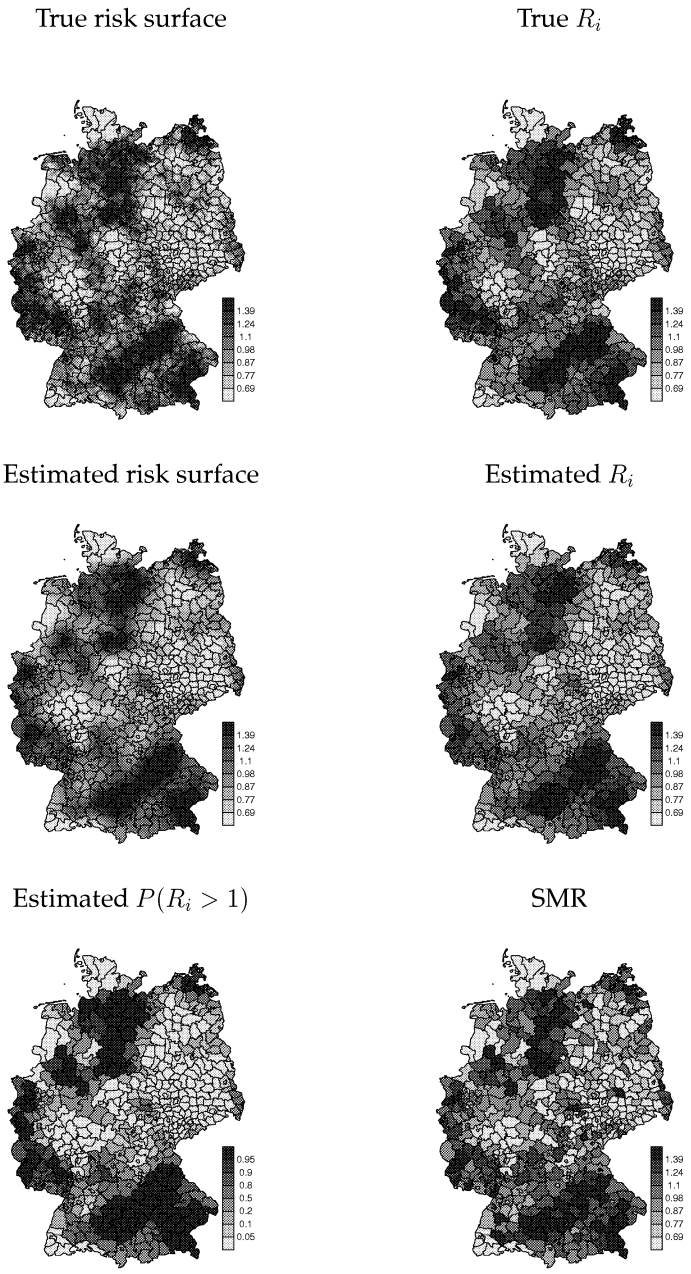

Figure 16: Results for the simulated data set II, with $\tau = 8$ and $r = 40$. The true values of the risk surface and regional risks, also shown in Figure 8, are added for reference.

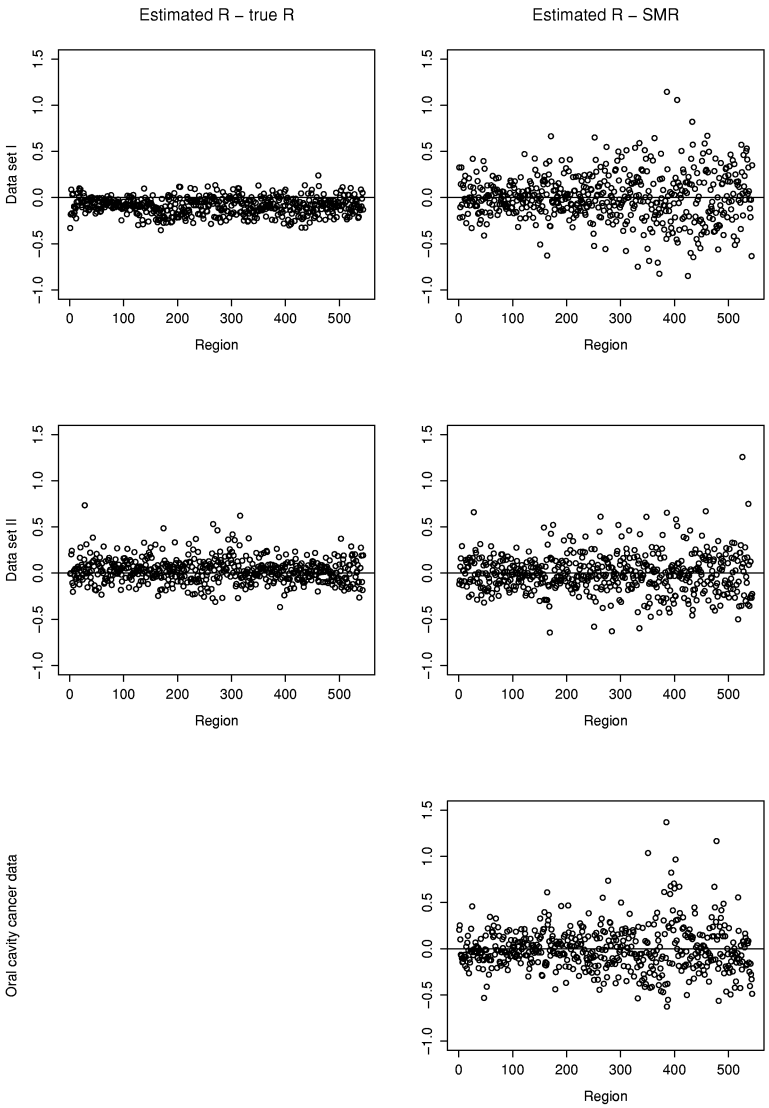

Figure 17: Differences between the estimated and true relative risk (left) and between estimated risk and SMR (right) for the simulated data sets I (top) and II (middle), and differences between estimated risk and SMR for the oral cavity cancer data (bottom).

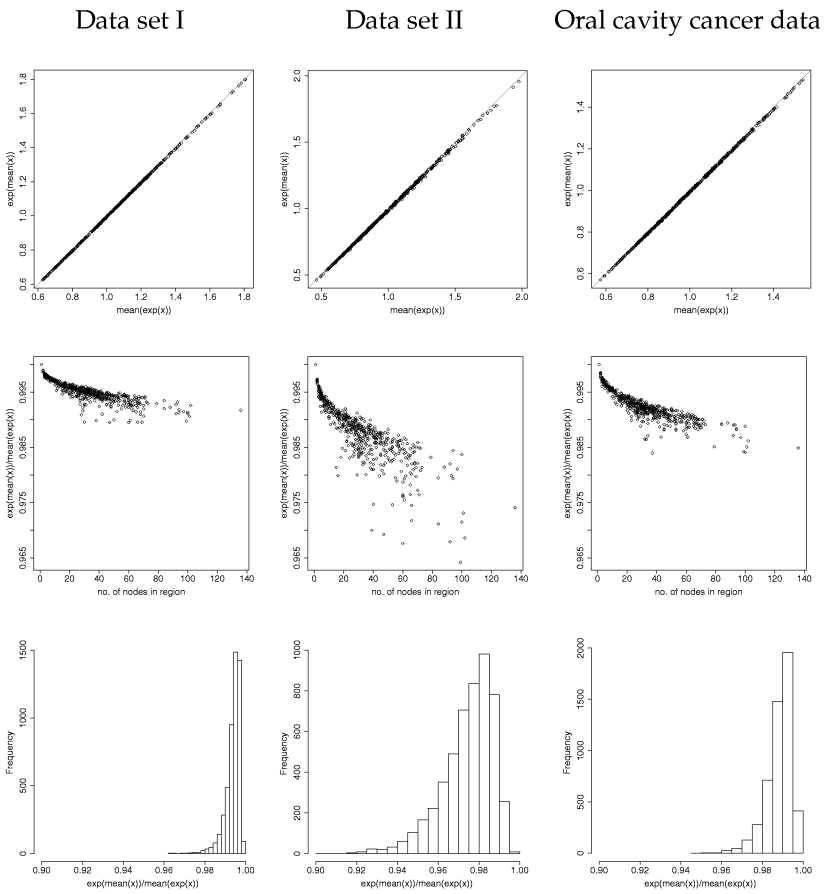

Figure 18: Plots of the posterior mean estimates $\widehat{\widetilde{R}}_i$ against $\widehat{R}_i$ (top panels) and $\widehat{\widetilde{R}}_i/\widehat{R}_i$ as a function of the number of lattice nodes within the region for the two simulated data sets and for the oral cavity cancer data. The bottom panels show histograms of the fractions $\widetilde{R}_{10}^{(k)}/R_{10}^{(k)}$ for region 10, which has 53 lattice nodes.

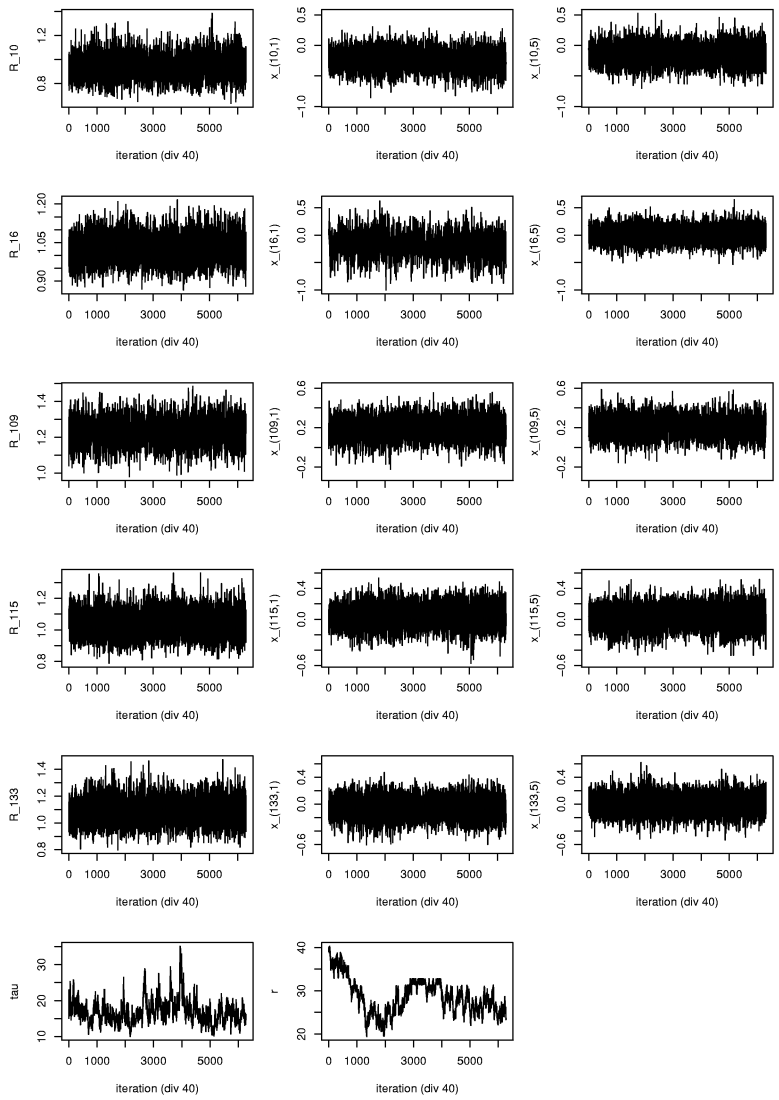

Figure 19: Selected trace plots for the oral cavity cancer data. The five top rows show trace plots of the regional risk R_i for five regions (left) and of two elements of $\boldsymbol{x}$ falling within each region (middle and right). Every 40th iteration is shown.

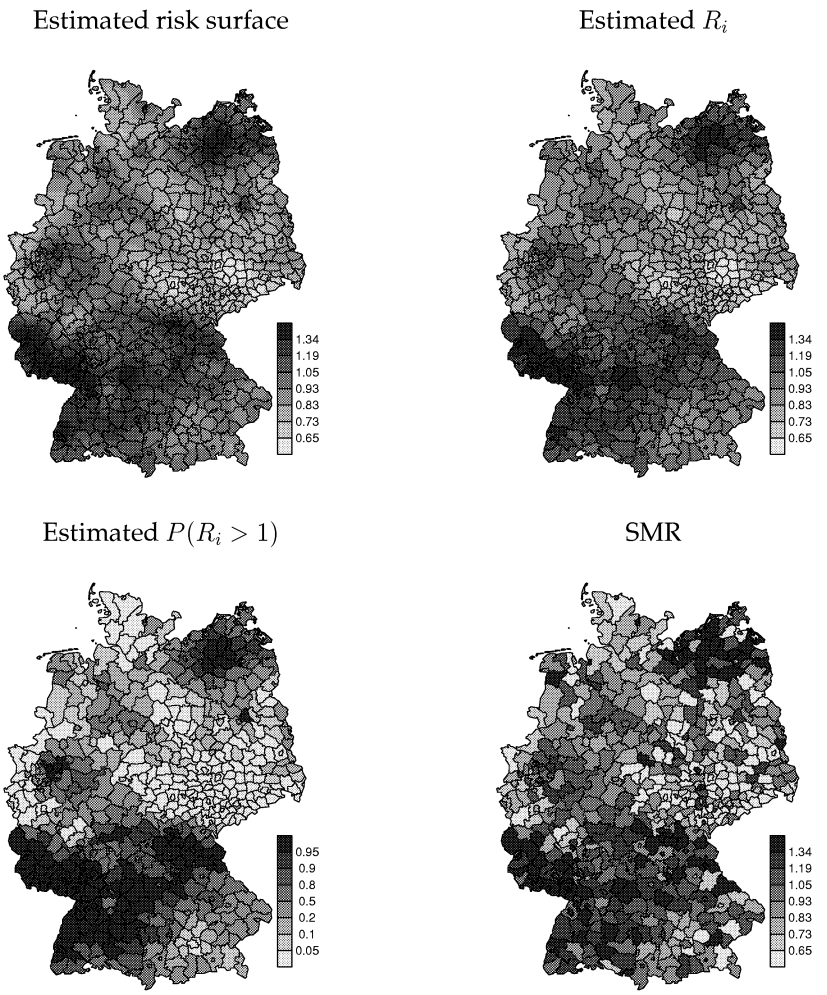

Figure 20: Results for the German oral cavity cancer data.

Acknowledgements

The authors would like to thank Leonhard Knorr-Held for making the German oral cavity cancer data available and for providing the data routines for generating the map of Germany, and Håkon Tjelmeland for stimulating discussions.

References

Bernardinelli, L., Pascutto, C., Best, N. G. and Gilks, W. R. (1997). Disease mapping with errors in covariates, *Statistics in Medicine* **16**: 741–752.

Besag, J., York, J. and Mollié, A. (1991). Bayesian image restoration with two applications in spatial statistics (with discussion), *Annals of the Institute of Statistical Mathematics* **43**: 1–59.

Best, N. G., Ickstadt, K. and Wolpert, R. L. (2000). Spatial Poisson regression for health and exposure data measured at disparate resolutions, *Journal of the American Statistical Association* **95**: 1076–1088.

Clayton, D. and Kaldor, J. (1987). Empirical Bayes estimates of age-standardized relative risks for use in disease mapping, *Biometrics* **43**: 671–681.

Cressie, N. A. C. (1993). *Statistics for Spatial data*, Second edn, John Wiley & Sons, New York.

Diggle, P. J. (2000). Overview of statistical methods for disease mapping and its relationship to cluster detection, *in* P. Elliott, J. C. Wakefield, N. G. Best and D. J. Briggs (eds), *Spatial Epidemiology. Methods and Applications*, Oxford University Press, New York, pp. 87–103.

Diggle, P. J., Tawn, J. A. and Moyeed, R. A. (1998). Model-based geostatistics (with discussion), *Applied Statistics* **47**: 299–350.

Fernández, C. and Green, P. J. (2002). Modelling spatially correlated data via mixtures: a Bayesian approach, *Journal of the Royal Statistical Society, Series B* **64**: 805–826.

Gilks, W. R., Richardson, S. and Spiegelhalter, D. J. (1996). *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, UK.

Hrafnkelsson, B. and Cressie, N. (2003). Hierarchical modeling of count data with application to nuclear fall-out, *Journal of Environmental and Ecological Statistics*. To appear.

Ickstadt, K. and Wolpert, R. L. (1999). Spatial regression for marked point processes (with discussion), *in* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 6*, Oxford University Press, Oxford, UK, pp. 323–341.

Kelsall, J. E. and Wakefield, J. C. (1999). Contribution to the discussion of Best, N. G., Arnold, R. A., Thomas, A.,Waller, L. A. and Conlon, E. M.: "Bayesian models for spatially correlated disease and exposure data", *in* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 6*, Oxford University Press, Oxford, UK, pp. 131–156.

Kelsall, J. E. and Wakefield, J. C. (2002). Modeling spatial variation in disease risk: A geostatistical approach, *Journal of the American Statistical Association* **97**: 692–701.

Knorr-Held, L. and Besag, J. (1998). Modelling risk from a disease in time and space, *Statistics in Medicine* **17**: 2045–2060.

Knorr-Held, L. and Raßer, G. (2000). Bayesian detection of clusters and discontinuities in disease maps, *Biometrics* **56**: 13–21.

Knorr-Held, L. and Rue, H. (2002). On block updating in Markov random field models for disease mapping, *Scandinavian Journal of Statistics* **29**: 597–614.

McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*, Second edn, Chapman & Hall, London, UK.

Mollié, A. (1996). Bayesian mapping of disease, *in* W. R. Gilks, S. Richardson and D. J. Spiegelhalter (eds), *Markov Chain Monte Carlo in Practice*, Chapman & Hall, London, UK, pp. 359–379.

Pascutto, C., Wakefield, J. C., Best, N. G., Richardson, S., Bernardinelli, L., Staines, A. and Elliott, P. (2000). Statistical issues in the analysis of disease mapping data, *Statistics in Medicine* **19**: 2493–2519.

Rue, H. (2001). Fast sampling of Gaussian Markov random fields, *Journal of the Royal Statistical Society, Series B* **63**: 325–338.

Rue, H. and Follestad, T. (2002). GMRFLib: a C-library for fast and exact simulation of Gaussian Markov random fields, *Preprint series in statistics no. 1/2002*, Dept. of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

Rue, H. and Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields, *Scandinavian Journal of Statistics* **29**: 31–49.

Wakefield, J. C. and Morris, S. E. (1999). Spatial dependence and errors-in-variables in environmental epidemiology (with discussion), *in* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith (eds), *Bayesian Statistics 6*, Oxford University Press, Oxford, UK, pp. 657–684.

Wakefield, J. C. and Morris, S. E. (2001). The Bayesian modeling of disease risk in relation to a point source, *Journal of the American Statistical Association* **96**: 77–91.

Wakefield, J. C., Best, N. G. and Waller, L. (2000). Bayesian approaches to disease mapping, *in* P. Elliott, J. C. Wakefield, N. G. Best and D. J. Briggs (eds), *Spatial Epidemiology. Methods and Applications*, Oxford University Press, New York, pp. 104–127.

Waller, L. A., Carlin, B. P., Xia, H. and Gelfand, A. E. (1997). Hierarchical spatio-temporal mapping of disease rates, *Journal of the American Statistical Association* **92**: 607–617.

Walter, S. D. (2000). Disease mapping: a historical perspective, *in* P. Elliott, J. C. Wakefield, N. G. Best and D. J. Briggs (eds), *Spatial Epidemiology. Methods and Applications*, Oxford University Press, New York, pp. 223–239.

Wolpert, R. L. and Ickstadt, K. (1998). Poisson/gamma random field models for spatial statistics, *Biometrica* **85**: 251–267.

A Computational details

A.1 The gradient and the Hessian of the Poisson log-likelihood

Define

$$\boldsymbol{a}_i(\boldsymbol{x}_{\mathcal{A}_i}) = (\exp(x_j))_{j\in\mathcal{A}_i}^T \tag{63}$$

$$S_i(\boldsymbol{x}_{\mathcal{A}_i}) = \sum_{j\in\mathcal{A}_i}\exp(x_j) = \sum_{k=1}^{n_i} a_{ik} \tag{64}$$

where x_j denotes element j of the log-risk surface $\boldsymbol{x}$. We compute the gradient vector $\boldsymbol{g}_i(\boldsymbol{x})$ and the Hessian matrix $\boldsymbol{G}_i(\boldsymbol{x})$ of the Poisson log-likelihood function given by $h_i(\boldsymbol{x}) = y_i \log(E_i' \sum_{j\in\mathcal{A}_i}\exp x_j) - E_i' \sum_{j\in\mathcal{A}_i}\exp x_j$, which define the second order Taylor approximation (33) of the conditional posterior distribution $\pi(\boldsymbol{x}_{\mathcal{A}_i}|\boldsymbol{x}_{-\mathcal{A}_i},\boldsymbol{\theta},\boldsymbol{y})$. The elements of $\boldsymbol{g}_i(\boldsymbol{x})$ and $\boldsymbol{G}_i(\boldsymbol{x})$ are given by

$$\frac{\partial h_i(\boldsymbol{x})}{\partial x_k} = (\frac{y_i}{S_i(\boldsymbol{x}_{\mathcal{A}_i})} - E_i')I_{[k\in\mathcal{A}_i]}\exp(x_k) \tag{65}$$

$$\frac{\partial^2 h_i(\boldsymbol{x})}{\partial x_k \partial x_l} = \begin{cases} \big((\frac{y_i}{S_i(\boldsymbol{x}_{\mathcal{A}_i})} - E_i')I_{[k\in\mathcal{A}_i]} - \frac{y_i}{S_i(\boldsymbol{x}_{\mathcal{A}_i})^2}I_{[k\in\mathcal{A}_i]}\exp(x_l)\big)\exp(x_k) & \text{if } l = k \\ -\frac{y_i}{S_i(\boldsymbol{x}_{\mathcal{A}_i})^2}I_{[k\in\mathcal{A}_i\cap l\in\mathcal{A}_i]}\exp(x_l)\exp(x_k) & \text{if } l \neq k. \end{cases} \tag{66}$$

Consequently, the vector $\boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$ and the matrix $\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i})$, evaluated in the mode $\boldsymbol{x}^0_{\mathcal{A}_i}$ of the posterior distribution of $\boldsymbol{x}_{\mathcal{A}_i}$, can be expressed by

$$\boldsymbol{g}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) = (\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i')\,\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) \tag{67}$$

$$\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{A}_i}) = (\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})} - E_i')\,\mathrm{diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})) - \frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{A}_i})^2}\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{A}_i})\boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{A}_i}), \tag{68}$$

establishing the expressions (37) and (38) of Section 4.1.

A.2 A Gaussian approximation to the posterior of $\boldsymbol{x}_{\mathcal{S}}$ for general sets of regions $\mathcal{S}$

Here, we establish the Gaussian approximation to the conditional posterior distribution $\pi(\boldsymbol{x}_{\mathcal{S}}|\boldsymbol{x}_{-\mathcal{S}},\boldsymbol{\theta},\boldsymbol{y})$ of the log-risk $\boldsymbol{x}_{\mathcal{S}}$ for blocks of lattice nodes corresponding to a set $\mathcal{S}$ of several regions, given by equation (26) in Section 3.2. In analogy to expression (33) for the single region block case, a Taylor expansion based quadratic approximation to the log-likelihood part $\sum_i h_i(\boldsymbol{x})$ of (26) is given by

$$\sum_i h_i(\boldsymbol{x}) \approx -\frac{1}{2}\boldsymbol{x}_{\mathcal{S}}^T(-\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}}))\boldsymbol{x}_{\mathcal{S}} + (\boldsymbol{g}^T(\boldsymbol{x}^0_{\mathcal{S}}) - (\boldsymbol{x}^0_{\mathcal{S}})^T\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}}))\boldsymbol{x}_{\mathcal{S}}, \tag{69}$$

discarding terms not depending on $\boldsymbol{x}_{\mathcal{S}}$. The vector $\boldsymbol{g}(\boldsymbol{x}^0_{\mathcal{S}})$ and the matrix $\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}})$ are the gradient and the Hessian of $\sum_i h_i(\boldsymbol{x})$ evaluated in the mode $\boldsymbol{x}^0_{\mathcal{S}}$ of the posterior distribution of $\boldsymbol{x}_{\mathcal{S}}$.

We derive the Gaussian approximation using the Poisson likelihood (15). As for the single region case in Appendix A.1, define

$$\boldsymbol{a}_i(\boldsymbol{x}_{\mathcal{S}}) = \left(\exp(x_j) I_{[j\in\mathcal{A}_i]}\right)_{j\in\mathcal{S}}, \text{ and} \tag{70}$$

$$S_i(\boldsymbol{x}_{\mathcal{S}}) = \sum_{j\in\mathcal{S}} (\exp(x_j) I_{[j\in\mathcal{A}_i]}) = \sum_{k=1}^{n_{\mathcal{S}}} a_{ik}. \tag{71}$$

Because of the conditional independence structure of the likelihood, the gradient and the Hessian defining the Taylor expansion of $\sum_i h_i(\boldsymbol{x})$ are given from (67) and (68) by the sums

$$\boldsymbol{g}(\boldsymbol{x}^0_{\mathcal{S}}) = \sum_i \left(\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{S}})} - E_i'\right) \boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{S}}) \tag{72}$$

$$\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}}) = \sum_i \left(\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{S}})} - E_i'\right) \operatorname{diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{S}})) - \sum_i \frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{S}})^2} \boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{S}}) \boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{S}}), \tag{73}$$

such that the precision matrix $-\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}})$ of the quadratic approximation to $\sum_i h_i(\boldsymbol{x})$ is block diagonal with blocks $-\boldsymbol{G}_i(\boldsymbol{x}^0_{\mathcal{S}})$ corresponding to (68) for each block $\mathcal{S}$. Define the $m_{\mathcal{S}} \times n_{\mathcal{S}}$ matrix $\boldsymbol{A}_{\mathcal{S}}$ by

$$\boldsymbol{A}_{\mathcal{S}} = \left(\frac{\sqrt{y_i}}{S_i(\boldsymbol{x}^0_{\mathcal{S}})} \boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{S}}) \right)_{i:\, \mathcal{A}_i \in \mathcal{S}}, \tag{74}$$

where $m_{\mathcal{S}}$ and $n_{\mathcal{S}}$ are the number of regions and number of lattice nodes within the block $\mathcal{S}$, respectively. In correspondence with the quadratic approximation for the single region case, the matrix $\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}})$ is of the form

$$\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}}) = -(\boldsymbol{D}_{\mathcal{S}} + \boldsymbol{H}_{\mathcal{S}}) \tag{75}$$

for a diagonal matrix $\boldsymbol{D}_{\mathcal{S}}$ and rank one matrix $\boldsymbol{H}_{\mathcal{S}}$ given by

$$\boldsymbol{D}_{\mathcal{S}} = -\sum_i \left(\frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{S}})} - E_i'\right) \operatorname{diag}(\boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{S}})) \tag{76}$$

$$\boldsymbol{H}_{\mathcal{S}} = \sum_i \frac{y_i}{S_i(\boldsymbol{x}^0_{\mathcal{S}})^2} \boldsymbol{a}_i(\boldsymbol{x}^0_{\mathcal{S}_i}) \boldsymbol{a}_i^T(\boldsymbol{x}^0_{\mathcal{S}}) = \boldsymbol{A}_{\mathcal{S}}^T \boldsymbol{A}_{\mathcal{S}}. \tag{77}$$

Substituting (69) for $\sum_i h_i(\boldsymbol{x})$ in the posterior distribution (26), using the expressions for $\boldsymbol{g}(\boldsymbol{x}^0_{\mathcal{S}})$ and $\boldsymbol{G}(\boldsymbol{x}^0_{\mathcal{S}})$ derived above and collecting terms that are linear and quadratic in $\boldsymbol{x}_{\mathcal{S}}$, we arrive at the Gaussian approximation given by (27).

A.3 Conditioning on a soft linear constraint

Here, we use Normal distribution theory to check the validity of equation (54) as a sample from the GMRF $\boldsymbol{x}$ conditionally on a soft linear constraint. Let

$$\boldsymbol{x} \sim N(\boldsymbol{\mu}, \boldsymbol{Q}) \tag{78}$$

and consider the general problem of sampling from the conditional distribution

$$\boldsymbol{x} \mid \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} + \boldsymbol{\epsilon}, \tag{79}$$

where $\boldsymbol{\epsilon} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma})$. This is equivalent to sampling from the distribution

$$\boldsymbol{x} \mid \boldsymbol{z} = \boldsymbol{b}, \tag{80}$$

where $\boldsymbol{z} = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{\epsilon}$. The sampling problem of Section 4 is a special case of (79) for which $\boldsymbol{\mu} = \boldsymbol{\mu}^*$ and $\boldsymbol{Q} = \boldsymbol{Q}^*$ as defined by (46) and (47), and where $\boldsymbol{\Sigma} = \boldsymbol{I}$ and $\boldsymbol{b} = \boldsymbol{0}$.

Let $\boldsymbol{x}_u$ be an *unconditional* sample for $\boldsymbol{x}$ from (78), and let

$$\boldsymbol{x}_c = \boldsymbol{x}_u - \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}(\boldsymbol{z} - \boldsymbol{b}), \tag{81}$$

where $\boldsymbol{z} = \boldsymbol{A}\boldsymbol{x}_u - \boldsymbol{\epsilon}$. We will show that $\boldsymbol{x}_c$ has the same distribution as a sample from $\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}$, and thus from (79) by comparing the mean and variance of $\boldsymbol{x}_c$ computed by (81) to the moments of (80).

Using multivariate Normal distribution theory, the mean vector and covariance matrix of the distribution $\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}$ can be shown to be

$$\begin{aligned}
\mathrm{E}(\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}) &= \boldsymbol{\mu} + \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{\mu}) \\
\mathrm{Cov}(\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}) &= \boldsymbol{Q}^{-1} + \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}\boldsymbol{A}\boldsymbol{Q}^{-1} \\
&= (\boldsymbol{Q} + \boldsymbol{A}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{A})^{-1}.
\end{aligned} \tag{82}$$

The mean vector and covariance matrix of $\boldsymbol{z} = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{\epsilon}$ is

$$\mathrm{E}(\boldsymbol{z}) = \boldsymbol{A}\boldsymbol{\mu}, \text{ and} \tag{83}$$

$$\mathrm{Cov}(\boldsymbol{z}) = \boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma}, \tag{84}$$

such that

$$\begin{aligned}
\mathrm{E}(\boldsymbol{x}_c) &= \boldsymbol{\mu} - \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}(\boldsymbol{A}\boldsymbol{\mu} - \boldsymbol{b}) \\
&= \mathrm{E}(\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}),
\end{aligned} \tag{85}$$

and

$$\begin{aligned}
\mathrm{Cov}(\boldsymbol{x}_c) &= \boldsymbol{Q}^{-1} + \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}\boldsymbol{A}\boldsymbol{Q}^{-1} \\
&= \boldsymbol{Q}^{-1} + \boldsymbol{Q}^{-1}\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^T + \boldsymbol{\Sigma})^{-1}\boldsymbol{A}\boldsymbol{Q}^{-1} \\
&= \mathrm{Var}(\boldsymbol{x}|\boldsymbol{z} = \boldsymbol{b}).
\end{aligned} \tag{86}$$

Consequently, $\boldsymbol{x}_c$ computed by (81) has the same first and second order moments as a sample from (79), and since the corresponding distributions are both Gaussian, the validity of (54) as an update of (45) follows.