A NOVEL BAYESIAN METHOD FOR MAKING THE MOST OF SPATIAL FISHERY CATCH AND EFFORT DATA

by

Aaron Springford B.Sc.(Hon), Simon Fraser University, 2004

A RESEARCH PROJECT SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF RESOURCE MANAGEMENT in the School of Resource and Environmental Management

Report No. 458

© Aaron Springford 2008 SIMON FRASER UNIVERSITY Summer 2008

All rights reserved. This work may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

APPROVAL

Name:	Aaron Springford
Degree:	Master of Resource Management
Title of Research Project	t: A Novel Bayesian Method for Making the Most of Spatial Fishery Catch and Effort Data
Report Number:	458
Examining Committee:	Mrs. Ashleen Benson Chair
	Dr. Sean P. Cox Senior Supervisor Assistant Professor School of Resource and Environmental Management Simon Fraser University
	Dr. Randall M. Peterman Supervisor Professor School of Resource and Environmental Management Simon Fraser University

Date Approved:

Abstract

Current fisheries monitoring practices in many regions of the world include precise measures of fishing location. However, spatial information is ignored in most current stock assessments, which assume instead that fishery catch per unit effort (CPUE) observations are independent replicate measurements of average stock density. This aspatial approach misleads estimates of stock status and productivity because: (1) harvesters are not random samplers of stock density, and (2) CPUE observations may be spatially autocorrelated.

This paper introduces a hierarchical Bayesian method describing the spatial distribution of fishery CPUE. The spatial method is applied to British Columbia sablefish (*Anoplopoma fimbria*) and compared to traditional aspatial approaches in a stock assessment context. I show that spatial assessments offer less optimistic estimates of stock status and productivity compared to traditional aspatial assessments, and that the area occupied by the commercially exploitable stock is estimated to have declined by 62 percent from 1990 to 2005.

To my family

"If we knew what it was we were doing, it would not be called research, would it?" — Albert Einstein (1879-1955)

Acknowledgments

I thank my senior supervisor Sean Cox, who gave me just enough room to be creative, but not enough to run amok. Thanks to Randall Peterman for all of the focused help with this manuscript. Thanks to Bill de la Mare for giving me a unique perspective on fisheries management.

I also thank the Fisheries Science and Management Research Group, especially those students in my cohort. Discussions with Pier van Dishoeck and Ashleen Benson have been particularly helpful.

Thanks to the people at the Canadian Sablefish Association, Rob Kronlund, and Malcolm Wyeth for your support and cooperation. Funding for this project was provided by the National Science and Engineering Research Council of Canada, the Canadian Sablefish Association, and the Faculty of Applied Science at Simon Fraser University.

Finally, thanks to my family for all that you have done for me. And special thanks to Gwen for living with "ha ha fish!"

Contents

$\mathbf{A}_{]}$	ppro	val		ii
A	bstra	ıct		iii
D	edica	tion		iv
\mathbf{Q}_1	uotat	tion		v
A	ckno	wledgr	nents	vi
Co	onter	nts		vii
Li	st of	Table	s	ix
Li	st of	Figur	es	x
1	Spa	tial hi	erarchical Bayesian models for analyzing fishery data	1
	1.1	Introd	luction	1
		1.1.1	CPUE-based assessment of fish stocks	3
	1.2	Metho	ods and Computation	7
		1.2.1	Hierarchical Models	7
		1.2.2	Advantages of a Bayesian approach	10
		1.2.3	Statistical models for fishery CPUE	11
		1.2.4	The big n problem \ldots	16
		1.2.5	Conditionally autoregressive hierarchical spatial model	19

		1.2.6	Summaries of model results	21
	1.3	Discus	ssion	25
2	Ac	compai	rison of spatial and aspatial assessments of BC sablefish	h
	usir	ng com	nmercial data	31
	2.1	Introd	luction	31
		2.1.1	British Columbia sablefish fishery	33
	2.2	Metho	ds	35
		2.2.1	Data characteristics and exploratory analysis	35
		2.2.2	Data analysis	42
	2.3	Result	ts	52
		2.3.1	Model selection and resulting indices of abundance $\ . \ . \ .$.	52
		2.3.2	Consequences for stock assessment	53
		2.3.3	Trend in the area occupied	58
	2.4	Discus	ssion	61
\mathbf{A}	MC	CMC p	rocedure for spatial HBM	71
Re	efere	nces		74

List of Tables

2.1	Data extracted from the commercial sablefish database	35
2.2	Symbols used in BC sablefish trap fishery models	43
2.3	Generalized linear model selection statistics	44
2.4	BC sable fish spatial hierarchical Bayesian model selection statistics $% \operatorname{BC}$.	47
2.5	Pairwise correlations of relative abundance indices	56
2.6	Index comparison: Schaefer model fits	59

List of Figures

1.1	Effect of spatial aggregation and disaggregation on fishery CPUE	5
1.2	Hierarchical model structure for spatial fisheries data \ldots	9
1.3	Aspatial GLM for fisheries CPUE data	13
1.4	Hierarchical model for normally distributed fisheries data	17
1.5	CAR Hierarchical model for normally distributed fisheries data	20
1.6	Hierarchical model for presence-absence fisheries data	28
2.1	BC sablefish random stratified survey areas	37
2.2	Spatial concentration of effort	39
2.3	Number of skippers by grid cell	40
2.4	Number of grid cells fished by skipper	41
2.5	Predicted CPUE from spatial HBM for BC sable fish trap data $\ . \ . \ .$	54
2.6	Spatial effects from spatial HBM for BC sable fish trap data	55
2.7	Indices of abundance for BC sablefish	57
2.8	Area occupied by the economically exploitable BC sable fish stock $\ . \ .$	60
2.9	Time series of scdf and heterogeneity index	61
2.10	Variance in HBM prediction of log-CPUE for BC sable fish trap data .	67

Chapter 1

Spatial hierarchical Bayesian models for analyzing fishery data

1.1 Introduction

The use of average fishery catch per unit of effort (CPUE) as an index of fish stock abundance has been the cornerstone of fisheries stock assessment for nearly a century. The use of fishery CPUE has been popular because catch records are easily and cheaply obtained compared to alternative methods of abundance estimation such as tagging experiments. The use of fishery catch statistics to deduce stock abundance goes back to at least the work of Baranov (1918). Baranov envisioned a sedentary fish stock of uniform density over its occupied area, and a fishery that randomly samples this homogeneous stock with independent and identical fishing sets, so that the set-by-set catch is some constant fraction of the total remaining stock. Ricker (1940) showed that a similar relationship would arise for highly mobile fish stocks that re-occupy fished areas instantaneously. However, deviations from the assumptions of Baranov (1918) or Ricker (1940), namely heterogeneous densities of fish or targeted fishing activities, have the potential to affect the predicted proportional relationship between CPUE and stock abundance (Paloheimo and Dickie, 1964).

A famous example of disproportionality between CPUE and stock biomass is the collapse of the northern cod (*Gadus Morhua*) fishery off the coast of Newfoundland,

Canada. The collapse has been characterized as "one of the worst social and economic nightmares in the nations history" (Walters and Maguire, 1996). The ensuing moratorium on fishing led to what was described in the popular media as the single biggest layoff of Canadian workers ever. In 1505, before the large scale deployment of European fishing fleets in North America, the estimated biomass of cod in Newfoundland waters was approximately 7,000,000 tonnes (Rose, 2004). The same analysis estimates current biomass near 400,000 tonnes, or about six percent of the pre-settlement biomass.

Although a combination of factors led to collapse of the northern cod stock, scientists played a role by persistently overestimating the abundance and productivity of the stock (Walters and Maguire, 1996). Several studies have suggested that the stock assessment failed to track the declining abundance because of an assumed linear relationship between fishery CPUE and stock abundance (Hutchings, 1996; Hutchings and Myers, 1994; Walters and Maguire, 1996). Observations of stock decline in some areas hinted that the overall biomass may have been decreasing significantly even though the fishery CPUE remained high (Rose and Kulka, 1999). Despite knowledge on the part of scientists and cod harvesters that the fishery was targeting smaller and smaller areas of high stock density (Rose and Kulka, 1999), the assessment was unable to show a decline in total stock biomass. Some authors (Hutchings, 1996; Hutchings and Myers, 1994; Walters and Maguire, 1996) have recommended in hindsight that a power relationship relating stock biomass to fishery CPUE (i.e. $CPUE = aB^c$), where B is the unknown stock biomass, and a > 0 and 0 < c < 1 are unknown parameters) might be a better way to represent the link between fishery CPUE and stock abundance when the spatial characteristics of the fishery and stock change over time. They argue that the power relationship with 0 < c < 1 is a more appropriate assumption that reflects the ability of the fishery to maintain high catch rates even as the stock declines. However, Shelton (2005) showed that use of a power relationship would not have changed scientific advice for the northern cod appreciably, which casts doubt on the use of the power relationship to reliably avoid such fisheries assessment failures in the future. A better approach to deal with the problem of changing spatial distribution of stocks and fisheries is needed.

1.1.1 CPUE-based assessment of fish stocks

In stock assessment, fishery CPUE is usually assumed proportional to underlying abundance, even though it has long been recognized that the distribution of fish and the characteristics of the fishery can cause serious departures from such a relationship (Ellis and Wang, 2007; Harley et al., 2001; Paloheimo and Dickie, 1964; Peterman and Steer, 1981; Swain and Sinclair, 1994). To account for these departures, analysts tend to focus on standardizing the measure of fishery effort, so that each fishing set can be treated as an identical replicate measurement of stock density. Fishery CPUE standardization attempts to isolate the effects of auxiliary characteristics such as the vessel, the skipper, the type or size of gear used, or the season in which the effort was expended. Recent approaches to standardizing fishery CPUE involve the use of generalized linear models (GLMs) to account for variation arising from these auxiliary characteristics (covariates) (Maunder and Punt, 2004). Treating CPUE as a response and including *Year* as a covariate factor in the GLM allows year effects to be interpreted as an index of relative abundance. In essence, the GLM is:

f(CPUE) = Intercept + Year + Other Covariates + Residual Error

where f is a one-to-one link function. Once the GLM is fit, the effect of changes in stock biomass is assumed to be represented by the Year effects, so that Year can be used as an index of relative stock biomass.

The GLM standardization approach tends to result in a large number of statistically significant model covariates. However, the estimated effects are often small when compared to the scale of the total variation in CPUE, so that little of the observed variation is actually explained by the model. In such cases, the statistical significance of model effects is the result of the large sample size created by assuming that each of the large number of fishery observations (e.g. hundreds or thousands, typically) is independent, rather than the result of important features in the data.

The continued treatment of fishing observations as independent is surprising because there are several mechanisms that may cause spatial correlation among fishery observations (Nishida and Chen, 2004). I term this inattention to spatial autocorrelation the **non-independence oversight**. The spatial distribution of fish stocks is often not random; rather, CPUE observations that are made geographically close together will be on average more alike than observations made far apart. Conspecific interactions or habitat features that are not easily deduced or measured can cause heterogeneity in the density of fish stocks and fishery CPUE. In theory, some of these features might be measured and included as covariates in the GLM, but there are countless others such as water quality or predator-prey interactions which are too costly or impossible to measure. In practise, it will be exceedingly difficult to overcome the non-independence oversight by including enough covariates to explain spatial variation in CPUE and make GLM residuals statistically independent. A more direct and feasible approach to modeling spatial CPUE data is to estimate spatial autocorrelation. There is a paucity of fishery CPUE modeling exercises in the primary literature that attempt such an approach, but Nishida and Chen (2004) is one exception.

In addition to the non-independence oversight, systematic changes in the distribution of fishery effort and stock density over time can also have an effect on GLMderived indices (Campbell, 1998, 2004; Walters, 2003). Stock abundance indices, by design, should track changes in the abundance of a stock. Baranov (1918) and Ricker (1940) viewed CPUE as a relative measure of stock density, defined as the biomass in the area "swept" be the fishing gear. Assuming that harvesters fish randomly over the area occupied by the stock, average fishery CPUE is an estimate of average stock density. If the area occupied by the stock is constant, then the average stock density is proportional to the total biomass of the stock. However, fish harvesters most assuredly do not fish randomly. Rather, areas with higher densities end up being targeted preferentially because such areas are more profitable. Fish harvesters are able to target areas with higher densities based on previous experience if such areas provide consistently high catch rates. Alternatively, harvesters can use modern electronic fish-finding equipment such as sonar to search for and target high density areas.

Given that harvesters do not fish randomly, changes in the spatial distribution of the fish stock are expected to affect time trends in aggregate fishery statistics such as mean fishery CPUE. Consider two scenarios in which a fish stock is declining in biomass (Figure 1.1): in the first scenario - aggregation - the fish stock is decreasing in both biomass and spatial extent simultaneously; in the second scenario - disaggregation - the fish stock is decreasing in biomass, but increasing in spatial extent. Spatial



Figure 1.1: Effect of spatial aggregation and disaggregation on fishery CPUE. Curves represent the simplified distribution of catch per unit effort (CPUE) over an imagined space. Solid curves represent aggregation (left) and disaggregation (right) scenarios. Dashed curves represent depletions proportional to density, and are provided for reference. Harvesters are assumed to be able to target those areas with the highest CPUE; shaded areas represent the region over which most fishery activity is expected to occur.

aggregation of the underlying stock leads to hyperstability (i.e. positive bias - abundance declines faster than average fishery CPUE) because the stock density where the fishery operates remains high even as the average density declines (Hilborn and Walters, 1992). In the second case, spatial disaggregation will lead to hyperdepletion (negative bias) in relative biomass trends because the stock density where the fishery operates declines faster than the average stock density (Hilborn and Walters, 1992). These biases are the result of two effects (Campbell, 1998, 2004; Walters, 2003): (1) disproportionate changes in the local density in those areas where the fishery operates; (2) changes in the area in which the fishery operates leads to changes in the area over which stock density is measured. To address the first bias, measures of density should be averaged by location (i.e. spatially) rather than by fishing set. To address the second bias, the total stock area (even areas without fishing effort) that is considered in the analysis should remain constant throughout a CPUE time series. In cases where the exploited area is changing over time, the CPUE at locations within the total stock area that were not fished needs to be predicted using a model (Campbell, 2004). Ideally, survey measurements could be used to ensure that the exploitable stock is contained within the defined total stock area.

Because harvesters are not random samplers of stock density, there is a disconnect between aspatial GLM year effects and stock biomass from year to year. I term this disconnect the **dimensional oversight**, because GLM year effects measure changes in fishery performance as measured by average set-by-set CPUE, not changes in average stock density. Other studies have addressed the dimensional oversight, but proposed solutions have been somewhat ad hoc. For example, Walters (2003) suggested a method of analysis for stocks in which the area fished was initially a subset of the total historical stock boundary, and effort subsequently moved on to other areas as the initially fished areas declined in abundance. His method used areas defined on a grid and assumed: (1) that prior to the first instance of fishing in an area, the catch rate would have been the same as the average in the first three years that the area was fished (stable biomass before fishing, small area depletion in the first three years fishing); (2) once an area was abandoned, future catch rates in the area would have been the same as the last fishing year (no stock recovery after fishing ceases). Walters (2003) suggested that the dimensional oversight is better recognized than ignored, and that spatial gaps in fishery data might be filled in by borrowing information from neighbouring areas. Campbell (2004), who also suggests analyzing fishery CPUE on

a pre-defined grid, compared a series of assumptions that might be made in order to fill in grid areas that did not receive any fishing effort. However, the assumptions proposed by Campbell (2004) rely on the judgement of the analyst to determine how the data gaps should be filled in, and are thus subjective.

To help deal with the challenge of using fishery CPUE data to infer year-toyear changes in total stock biomass, I present a class of hierarchical Bayesian models (HBMs) that predict fishery CPUE at every point in the fixed total stock area. The HBMs use the spatial autocorrelation of the fishery CPUE data themselves when predicting CPUE at unfished locations, which is more objective than the approach proposed by Walters (2003) or Campbell (2004), and thereby has the potential to better address the non-independence oversight. The HBM approach also addresses the dimensional oversight by averaging CPUE density across space rather than averaging fishery performance. Assuming that the predicted CPUE at a point in space is proportional to the stock biomass density at that same point, the relative total stock biomass is the average predicted CPUE over the total stock area. Spatial fisheries data are treated continuously and on a set-by-set basis, which produces a level of spatial resolution that is limited only by the resolution of the data. This high level of resolution is desirable because changes in fish distribution may occur on a scale that is fine relative to management areas (Kulka et al., 1996). However, not discretizing the data increases the computational burden, especially for large fishery catch and effort datasets. I develop a novel approach that allows large amounts of data to be analyzed by sacrificing some resolution in spatial variation.

1.2 Methods and Computation

1.2.1 Hierarchical Models

Hierarchical models are appropriate for data that arise from a combination of processes that operate on different scales, because they specify not only the statistical distribution for the data, but also the distributions of model parameters. Specifying distributions for model parameters creates additional levels or stages in the model

where parameters at lower levels have probability distributions defined at higher levels. For example, consider a model for lengths of fish of a given species in different inlets along a coastline. Data are collected by travelling to each of the inlets and collecting samples of fish. A non-hierarchical model for these data would involve a distribution of fish lengths given an unknown mean fish length for all inlets, pooling all inlets together. The alternative non-hierarchical model would involve a separate distribution of fish lengths for each of the individual inlets. Neither of these nonhierarchical descriptions seem to be suitable descriptions of fish length because inlets will likely differ in their mean fish length, but not to the point where they are statistically independent. A hierarchical model for the same situation would include not only a distribution for fish lengths in each inlet (first stage), but also a (prior) distribution for the mean length across inlets (second stage). Including a distribution for the mean length across inlets specifies how they might be related, which strikes a balance between the assumption that all inlets are the same and the assumption that all inlets are independent. In addition, this balance need not be specified by the analyst, but can be estimated from the data. The Bayesian approach allows for the statistical analysis of such hierarchical structures by allowing population-level parameters (in our case, the mean fish length in an inlet) to be defined probabilistically.

Su et al. (2004) present a Bayesian hierarchical stock-recruitment analysis for pink salmon (*Oncorhynchus gorbuscha*) populations in the Northeast Pacific. The first stage of their model includes parameters that are specific to each population, such as the effect of sea surface temperature on the stock-recruitment relationship. In the second stage, population-specific parameters are correlated in space such that parameters for populations that are geographically close together are more similar than for stocks that are geographically far apart. Correlation based purely on geographic proximity is termed spatial autocorrelation, and in the case of Su et al. (2004) a distribution is specified for the statistical relationship between population parameters. In combination with population-specific parameters (the first hierarchical stage), the spatial autocorrelation parameters (the second hierarchical stage) and their priors (the third hierarchical stage) form a hierarchical model for pink salmon stocks in the Northeast Pacific. The class of models that I consider for spatial fisheries data also have a hierarchical structure (Figure 1.2) consisting of (1) a statistical model for the data (*Response* level); (2) a model for the aspatial and spatial parameters (*Process* level); (3) priors for the aspatial and spatial parameters (*Prior* level); and (4) hyperpriors for the spatial priors (*Hyperprior* level). One could go on specifying higher and higher levels since the probability distribution for a parameter at any one level will entail additional parameters at the next higher level. At some level, the model parameters must be assumed known, and are either chosen so as to be vague with respect to the distribution of the parameter that they describe (i.e. with a diffuse prior), or estimated from the data as a first step in the analysis in what is called an empirical Bayesian approach (Gelman et al., 2004).



Figure 1.2: Hierarchical model structure for spatial fisheries data.

1.2.2 Advantages of a Bayesian approach

The spatial hierarchical model in figure 1.2 is suited to the Bayesian approach to data analysis because the Bayesian approach allows both the observed data and model parameters to be random variables (Banerjee et al., 2004). Bayesian methods are often computationally expensive, but are becoming more straightforward to apply thanks to Markov Chain Monte Carlo (MCMC) techniques, the ready availability of large computer processing power, and automated software to perform the analysis. The challenge of the Bayesian approach lies in the fact that making inferences about model parameters $\boldsymbol{\theta}$ separately requires computing the marginal distribution of each parameter with respect to the other parameters. For example, given data \boldsymbol{y} , the marginal distribution of one model parameter θ_i

$$p(\theta_i | \boldsymbol{y}) = \int_{\boldsymbol{\theta}_{j \neq i}} p(\boldsymbol{\theta} | \boldsymbol{y}) d\boldsymbol{\theta}_{j \neq i}$$

is often impossible to calculate in closed form. As a result, probability distributions for parameters must be evaluated using numerical integration techniques such as MCMC. Markov Chain Monte Carlo techniques enable approximate integrations by generating samples from the distribution of interest, in this case the posterior probability distribution for the model parameters. Software to perform the MCMC sampling, check diagnostics, and compare alternative models is freely available (for a list of alternatives, see Gelman et al., 2004).

The incorporation of prior information in Bayesian analysis is also an attractive feature of the method, because spatial autocorrelative effects may be difficult to distinguish from ordinary aspatial linear effects. This difficulty may arise, for example, if fishing vessels operate largely in non-overlapping areas. Contrast in the number of vessels fishing in a particular location (i.e. several vessels fishing in close proximity to each other) is needed in order for the effect of the fishing vessel to be distinguished from the effect of the location in which the fishing vessel operates. If such contrast is not available, additional prior information about spatial effects or ordinary linear effects (from scientific surveys, for example) may be available to help untangle the model components. The natural incorporation of priors allows Bayesian models to include different sources of information that may otherwise be difficult to consider together, which could ultimately lead to more encompassing and useful results. In cases where prior information is not available, care must be taken to ensure that there is sufficient contrast in the data to allow extrapolation beyond sampled points.

1.2.3 Statistical models for fishery CPUE

My approach in describing statistical models for fishery CPUE is to begin with a familiar aspatial generalized linear model (GLM), which represents the common approach to standardizing fishery CPUE data (Maunder and Punt, 2004). I then modify the aspatial GLM to include spatial effects that take into account spatial autocorrelation in fishery CPUE and heterogeneity in sampling effort over space. This set-by-set approach raises computational problems associated with large numbers of data points typical of fishery catch and effort datasets; therefore, I present a modified spatial model that can be fitted to large amounts of data.

Aspatial Generalized Linear Model

The most commonly used method for modeling CPUE is the GLM. The generalized linear model explains variation in fishery CPUE as a linear combination of specified continuous covariates such as depth, and discrete covariates such as year. The objective of fitting a GLM to fishery CPUE data is to isolate effects on CPUE that are not the result of changes in stock abundance, which are assumed to be represented by year effects.

Suppose that set-by-set fishery CPUE data are available:

$$y_i = \frac{C_i}{E_i} \tag{1.1}$$

where *i* is fishing set, *y* is CPUE, *C* is catch, and *E* is fishing effort expended. The yearly average fishery CPUE could be used as an index of stock abundance, computing the average y_i each year *t*. However, effects other than changes in stock abundance can influence fishery CPUE. Commonly considered fishery-related effects include individual fishing vessels or fishing vessel characteristics. Other factors such as depth

of fishing set, underlying habitat type, and the season in which the fishing event took place, may also influence CPUE, so these factors should also be considered in the model formulation. All of these factors are gathered column by column into a design matrix of covariates

$$X = \begin{bmatrix} 1 & X_{11} & \dots & X_{1p} \\ 1 & X_{21} & \dots & X_{2p} \\ \vdots & \vdots & \vdots \\ 1 & X_{n1} & \dots & X_{np} \end{bmatrix}$$
(1.2)

where n is the total number of fishing records and p is the number of covariates included (i.e. covariates are in columns and fishing sets are in rows). The generalized linear model is then

$$f(\boldsymbol{y}) \stackrel{iid}{\sim} N(X\boldsymbol{\beta}, \tau^2) \tag{1.3}$$

where $\boldsymbol{\beta} = \beta_0, \beta_1, \dots, \beta_p$ is the column vector of linear effects corresponding to the covariates $X_0, X_1, \dots, X_p, \tau^2$ is the residual variance, and f is a link function such that the residual variation in the $f(\boldsymbol{y})$ has a normal distribution. A logarithmic link function is often used for CPUE, in part because it properly constrains CPUE to be positive (Hilborn and Walters, 1992). Because equation 1.3 describes the distribution of the data \boldsymbol{y} , it is often called the data likelihood. Model specification is completed by specifying priors. Unless informative prior knowledge is available, it is suitable to use vague or diffuse priors for model parameters, e.g.

$$\boldsymbol{\beta} \stackrel{iid}{\sim} N(0, 1000) \tag{1.4}$$

$$\tau^2 \sim IG(1, 0.001)$$
 (1.5)

where N is the normal distribution with mean 0 and variance 1000, and IG is the inverse-gamma distribution with mean 1 and 0.001 degrees of freedom. These diffuse priors allow the data to determine parameter estimates. Figure 1.3 is a depiction of the aspatial GLM as a special case of the spatial hierarchical model. The linear model coefficients β estimate the effect that each of the covariates has on the CPUE. The β corresponding to the year effect is interpreted as a relative measure of stock abundance that is linearly independent of other covariate effects in the GLM.



Figure 1.3: Aspatial GLM for fisheries CPUE data.

Point-referenced spatial Hierarchical Bayesian Model

The aspatial GLM analysis is made into a spatial one by adding set-by-set autocorrelated spatial effects to the hierarchical model. Point-referenced spatial models consist of observations made at continuous locations s that occur within a defined spatial domain \mathcal{D} , which is a fixed subset of the two-dimensional space \mathbb{R}^2 . Observations made at locations s are partially defined by their location and are therefore denoted y(s). For fishery CPUE, y(s) is assumed to be a measure of relative stock density at locations s. Point-referenced spatial models have the attractive feature that the spatial domain \mathcal{D} is unchanging, even though the locations s_{obs} will differ over time. In fisheries, this addresses the dimensional oversight by ensuring that inference is made in relation to the domain \mathcal{D} instead of the actual locations s that were observed, which may change over the history of the fishery. For example, if the underlying stock is aggregating over time such that the area occupied by both the stock and the fishery are decreasing, then biases can be avoided when constructing an aggregate index by averaging over the spatial domain \mathcal{D} rather than the subset of \mathcal{D} that is occupied by the fishery (Campbell, 1998, 2004; Walters, 2003). In other words, predicted fishery CPUE density must be averaged over the total stock area, even at locations that receive little or no fishing effort. Thus, averaging over \mathcal{D} requires prediction at unsampled points, which underscores the need to treat CPUE spatially.

Beginning from our aspatial GLM model, spatial effects can be included at the process stage so that

$$f(y(\boldsymbol{s})) \sim N(X\boldsymbol{\beta} + w(\boldsymbol{s}), \tau^2)$$
(1.6)

where w(s) are the spatial random effects at locations s. The spatial random effects do not refer to the effect of any particular covariate on f(y(s)), rather they are random location effects that include patterns of spatial variation. Priors for the w(s) are needed to make the spatial component of the model complete. Priors that specify the correlation structure of the w(s) allow spatial autocorrelation to be included in the model.

Spatial process models commonly assume isotropy: the correlation between two observations made at points s_i and $s_{j\neq i}$ depends solely on the distance v_{ij} between them, and not on other factors such as the direction of separation. Although there are several possible models for isotropic spatial correlation Banerjee et al. (2004), all of them specify a functional relationship between the degree to which points are correlated and the separation between them, i.e. $\operatorname{cov}(s_i, s_{j\neq i}) \equiv \rho(v_{ij}) = g(v_{ij}, \theta_{\text{cov}})$, where g is a decreasing function in v_{ij} , and θ_{cov} are parameters describing the shape of g. It may be worth comparing different models of spatial correlation for a given data set that differ in shape and complexity. However, in practise it is very difficult to determine which spatial correlation model is appropriate by examining the data alone, because many correlation models will fit the data equally well. Nishida and Chen (2004) consider several spatial correlation functions in their analysis of data from the Japanese longline yellowfin tuna fishery, but their results show little sensitivity to the choice of correlation function.

Because of its simplicity, one of the most commonly employed models of isotropic

spatial correlation is the exponential

$$\rho(v_{ij}) = e^{-\phi v_{ij}} \tag{1.7}$$

where v_{ij} is the distance between data points, and $\phi > 0$ is the decay parameter. When i = j, $v_{ij} = 0$, and

$$\rho(v_{ii}) = 1 \tag{1.8}$$

that is, observations made at the same point in space are repeated measures. In this paper, I use the exponential correlation model, but it is also possible to use other correlation models without loss of generality (Banerjee et al., 2004).

A covariance matrix for the locations $\boldsymbol{s} \in \mathcal{D}$ can be constructed from the correlation equation 1.7. For a given distance matrix V, which is the n by n collection of v_{ij} for all of the observed locations \boldsymbol{s} , the correlation matrix $H(\phi)$ is

$$H(\phi) = e^{-\phi V} \tag{1.9}$$

thus, the covariance matrix Σ_w of the spatial random effects is

$$\Sigma_w(\phi, \sigma^2) = \sigma^2 H(\phi) \tag{1.10}$$

where σ^2 is the variance of the spatial effects. The vector of spatial random effects $w(\mathbf{s}) = w(s_1)...w(s_n)$, which are conditionally independent of $X\boldsymbol{\beta}$ (equation 1.6), thus have the prior distribution

$$w(\boldsymbol{s}) \sim N_n(\boldsymbol{0}, \Sigma_w) \tag{1.11}$$

Finally, a vague conjugate prior for the spatial variance is

$$\sigma^2 \sim IG(1, 0.001) \tag{1.12}$$

Note that the model in equation 1.6 has more parameters to estimate than there are data points, because there is a w for each s plus other model parameters (β , τ^2 , σ^2 , ϕ). The number of parameters that need to be estimated can be reduced specifying the marginal distribution of f(y(s)) with respect to w(s), i.e.

$$f(y(\boldsymbol{s})) \sim N_n(X\boldsymbol{\beta}, \Sigma) \tag{1.13}$$

where $\Sigma = \Sigma_w + \tau^2 I$, *I* being the identity matrix. This form of the model includes spatial autocorrelation directly in the covariance matrix Σ .

Year effects on CPUE can be included in the same way as for the aspatial model by including year as a factor in X. However, spatial effects w(s) that vary with time can also be considered. Time-varying spatial effects are appropriate in cases where patterns in spatial variation are changing over time, as is often the case in natural systems that respond to changing environmental conditions. A simple way to incorporate time-varying spatial effects into the model is to estimate a separate set of spatial effects $w(s_r)$ for time periods $r \in 1...R$. The prior distribution for w(s)is then divided into R conditionally independent spatial effects surfaces, each having the prior distribution

$$w(\boldsymbol{s}_r) \sim N_{n_r}(\boldsymbol{0}, \sigma_r^2 H_r(\phi_r)) \tag{1.14}$$

where n_r is the number of observations, σ_r^2 is the spatial variance, ϕ_r is the Exponential correlation parameter, and H_r is the correlation matrix of the observations, all at time r. Figure 1.4 is a depiction of the point-referenced spatial HBM for the case in which there may be a time series of spatial effects surfaces. In this depiction, spatial parameters $\boldsymbol{\sigma}$ and $\boldsymbol{\phi}$ are presented as vectors of length R.

1.2.4 The big n problem

Despite the apparent simplicity of the spatial covariance model in section 1.2.3, such models are not computationally feasible for commercial fishery data analysis in which there are many thousands of observations. The difficulty lies in the matrix inversions required to evaluate the conditional distribution for β (Banerjee et al., 2004). From equation 1.13, this conditional distribution is

$$p(\boldsymbol{\beta}|\tau^2, \sigma^2, \phi, y(\boldsymbol{s})) = N(A\boldsymbol{a}, A)$$
(1.15)

where

$$A^{-1} = X^{T} (\sigma^{2} H(\phi) + \tau^{2} I)^{-1} X + \Sigma_{\beta}^{-1}$$
(1.16)

$$\boldsymbol{a} = X^{T}(\sigma^{2}H(\phi) + \tau^{2}I)^{-1}\boldsymbol{y} + \Sigma_{\boldsymbol{\beta}}^{-1}\mu_{\boldsymbol{\beta}}$$
(1.17)



Figure 1.4: Hierarchical model for normally distributed fisheries data.

and where μ_{β} and Σ_{β} are the mean and variance, respectively, of a normal prior distribution on β . If there are *n* observations, a square matrix of dimension *n* by *n* must be inverted in order to compute *A*. Beyond approximately 500 observations, repeated matrix inversion of the type necessary for MCMC integration techniques becomes infeasible in commonly available computing environments. This is described by Banerjee et al. (2004) as the big *n* problem.

Banerjee et al. (2004) list several big n strategies. For example, simply ignoring some of the data by subsampling is one potential approach. Although it may seem like a bad idea to thin the data, observations aren't expected to be statistically independent anyway. The information lost may be slight if the data points that are actually used show good contrast in X and are representative of the spatial domain \mathcal{D} . To check the effect of subsampling, several subsamples can be taken and model results compared for consistency. Another possible big n strategy is Coarse-Fine Coupling (Higdon et al., 2003). In this method, two Markov chains - fine and coarse - are computed simultaneously during the Bayesian analysis. The fine chain includes all of the untransformed data, whereas the coarse chain is based on some transformation (e.g. local average) of the data from n dimensions to some smaller dimension. The lower dimensionality of the coarse chain means that parameter samples will be generated much more quickly than for the fine chain. After some number of updates, a swap of parameter values between the two chains is allowed according to a swapping rule. This procedure allows the fine chain to achieve the desired amount of posterior resolution while the coarse chain explores the parameter space quickly. The Coarse-fine Coupling option still requires that the fine chain can be calculated, however slowly, within the physical limits of the available computing environment, but doesn't rely on the fine chain to map the entire parameter space. The advantage of coarse-fine coupling is that it carries the original resolution of the data through to the results.

Another option for performing a spatial hierarchical Bayesian analysis of a very large dataset is to specify the priors for the random spatial effects using a conditionally autoregressive (CAR) formulation instead of the continuous isotropic model I described for the point-referenced HBM. With the CAR prior, the spatial domain \mathcal{D} is treated as a grid, with random spatial effects assigned according to the grid cell ς in which an event takes place.

The spatial random effects $w(\varsigma)$ are shared by all observations within an area or grid cell, reducing the number of spatial random effects that must be estimated from n (the dimension of s) to the number of areas or grid cells m (the dimension of ς). The assumption is that, conditional on covariates X, fishing event locations have occurred randomly within each grid cell so that $w(\varsigma)$ represent grid cell average spatial effects. Assuming that all observations within a grid cell share the same spatial effect leads to a loss of spatial effect resolution. However, the sacrifice may be worth making in cases where the number of observations is prohibitively large, and predictions based on grid cell average spatial effects are sufficient. Unfortunately, assessing the appropriateness of the CAR formulation for stock assessment using a simulation approach is not practical since repeated model fitting remains much too computationally expensive.

1.2.5 Conditionally autoregressive hierarchical spatial model

The conditionally autoregressive hierarchical spatial model is nearly identical to the point-referenced HBM; the first stage of the model described in equation 1.6 remains the same, keeping the set-by-set covariate information from X, but the second stage specification is adjusted to reflect the new CAR prior. Spatial effects are pooled according to the grid cell in which the fishing event occurred, i.e. $w(s_i) = w(\varsigma_j)$ for all $s_i \in \varsigma_j$. In conditionally autoregressive priors, a normal specification (Besag, 1974) is often used, i.e.

$$w(\varsigma_k)|w(\varsigma_{l\neq k}) \sim N\left(\frac{\sum_{l\neq k} b_{kl} w(\varsigma_l)}{\sum_{l\neq k} b_{kl}}, \frac{\sigma^2}{\sum_{l\neq k} b_{kl}}\right)$$
(1.18)

where b_{kl} define the objects or cells that are assumed to influence the value of $w(\varsigma_k)$, and $(k, l) \in 1...m$. For spatial effects $w(\varsigma)$ estimated on a grid, the simplest choice is to set $b_{kl} = 1$ for bordering cells and $b_{kl} = 0$ otherwise. The spatial hierarchical model with the CAR prior is depicted in figure 1.5. In the figure, I present spatial hyperparameters as time-dependent vectors (σ^2 is σ_r^2 for $r \in 1...R$), as was the case in the point-referenced HBM. This time-dependent CAR model fits a separate set of spatial effects for each time period, i.e.

$$w(\varsigma_{k,r})|w(\varsigma_{l\neq k,r}) \sim N\left(\frac{\sum_{l\neq k} b_{kl}w(\varsigma_{l,r})}{\sum_{l\neq k} b_{kl}}, \frac{\sigma_r^2}{\sum_{l\neq k} b_{kl}}\right)$$
(1.19)

From equation 1.18, the probability distribution of $w(\boldsymbol{\varsigma})$ is defined conditionally, where each $w(\varsigma_k)$ is a function of its neighbours. Therefore, spatial effects are estimated one at a time rather than simultaneously and no large matrix inversions are required. Computing the joint distribution $p(w(\varsigma_1)...w(\varsigma_m))$ thus becomes a tractable problem for large datasets.

The CAR prior for $w(\varsigma)$ is an example of a pairwise difference specification (Besag et al., 1995). Pairwise difference priors are improper, which means that an arbitrary constant can be added to each of the $w(\varsigma_k)$ and the CAR specification in equation 1.18 would still be satisfied. Improper prior distributions do not guarantee proper posterior distributions, but this is not of practical importance here because posterior



Figure 1.5: CAR Hierarchical model for normally distributed fisheries data.

distributions will be proper so long as there is at least some spatial parameter information available in the data (Banerjee et al., 2004). However, the overall level for $w(\varsigma)$ is not distinguishable from β_0 in equation 1.13, therefore the constraint

$$\sum_{k=1}^{m} w(\varsigma_k) = 0$$
 (1.20)

is often imposed in order to estimate β_0 (Banerjee et al., 2004). When a set of spatial effects is estimated for each of R time periods, as in equation 1.19, the above constraint is applied to each set of spatial effects (i.e., for $r \in 1 \dots R$).

In terms of computation, the conditional dependence of the CAR prior is convenient for Gibbs sampling routines, since the Gibbs sampler draws posterior parameter samples from the conditional distribution of each parameter in turn. The Gibbs sampling distributions for the spatial parameters thus follow directly from the definition of the prior and the data likelihood, conditioned on the values of other model parameters. For example, in the normal likelihood - normal CAR prior case, the conditional Gibbs sampling distribution for the spatial effects is a normal distribution whose mean is a weighted average of the CAR prior mean and $y(s_i) - X_i\beta$ for all $s_i \in \varsigma_j$, with weights inversely proportional to the variance of the prior and the likelihood (Gelman et al., 2004). The use of conjugate prior distributions and the conditional autoregressive prior on spatial effects allows conditional distributions to be specified for each parameter in the HBM, avoiding the need for alternative samplers such as the hybrid Metropolis-Gibbs sampler employed by Su et al. (2004). For details, see the Appendix.

1.2.6 Summaries of model results

In describing possible spatial models for fishery data, my approach has been to begin with the familiar aspatial generalized linear model (GLM), and subsequently describe modifications that account for spatial effects. The resulting spatial hierarchical Bayesian models take into account spatial autocorrelation in fishery CPUE and heterogeneity in sampling effort over space. In response to the big n problem, I present the CAR prior for spatial effects estimated on a grid. I now turn to the subject of what sort of information can be extracted from the results of fitting the spatial HBMs to fishery data. In the simple aspatial GLM case, the information extracted typically consists of the year effects on fishery CPUE. The spatial HBMs, however, can yield much richer results because they do not assume that the fishery and stock are each homogeneous in space. As a result, there is potentially much more information available in the spatial model results that can be used to infer trends in stock characteristics such as stock density profiles and stock extent.

The spatial HBMs allow a much richer posterior analysis because they address CPUE and stock density on a local scale rather than in aggregate. The assumption made for the spatial HBMs is that the CPUE of a fishing set is proportional to the stock density (local biomass) at the location of the set, which addresses the dimensional oversight described earlier by relating fishery CPUE to stock density directly. In addition, if suitable values for X can be obtained for all $s \in \mathcal{D}$, then geographic maps of predicted CPUE, generated from spatial HBM results, can reveal spatial contrasts in stock density. Spatio-temporal shifts in patterns of stock density, be they aggregation, disaggregation, or directional shifts, can be visualized by plotting a time series of predicted CPUE maps. Plotting maps of spatial effects w(s) shows areas where the fishery CPUE is higher or lower than expected given the aspatial covariates. Over time, regions in \mathcal{D} with consistently positive spatial effects are suggestive of areas of high biomass production that can direct further scientific investigations or inform spatial fishery management.

Maps of model-predicted CPUE are fine for visualizing trends in abundance, but many assessments of stock status require a single measure of relative stock abundance from year to year. The spatial HBM predicts CPUE for every point in \mathcal{D} , so the average predicted CPUE over \mathcal{D} is one candidate for a relative index of stock biomass. This spatial index of relative biomass will be more appropriate than covariate year effects in the likely event that regions in \mathcal{D} have gone unfished at some point in the time series (Campbell, 1998, 2004; Walters, 2003). In such cases, extrapolations made to fill in spatial gaps affect the average predicted CPUE in such a way as to depart from trends in the year effects alone.

The average predicted CPUE over \mathcal{D} can be evaluated numerically using Monte Carlo integration by randomly choosing a set of locations $\tilde{s} \in D$, and evaluating the posterior predicted CPUE $y(\tilde{s})$ according to

$$p(y(\tilde{\boldsymbol{s}})|y(\boldsymbol{s})) = \int p(y(\tilde{\boldsymbol{s}})|y(\boldsymbol{s}), \boldsymbol{\beta}, w(\boldsymbol{s}), \tau^2) p(\boldsymbol{\beta}, w(\boldsymbol{s}), \tau^2|y(\boldsymbol{s})) d\boldsymbol{\beta} dw(\boldsymbol{s}) d\tau^2 \quad (1.21)$$

which can be evaluated using the posterior samples from $p(\beta, w(s), \tau^2 | y(s))$ calculated during MCMC sampling. Suitable values must be specified for covariates that are not spatially referenced (i.e., that cannot be determined from the randomly chosen set of locations \tilde{s}). For example, the non-spatial covariate values might be set at some value of interest (e.g. a particular seasonal effect), or set at the average of the observed values. As approximated by the predicted CPUE at locations \tilde{s} , the average predicted CPUE across locations is then

$$\bar{y}(\boldsymbol{s}) = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} y(\boldsymbol{s}) d\boldsymbol{s} \approx \frac{1}{|\mathcal{D}|} \sum_{\mathcal{D}} y(\tilde{\boldsymbol{s}})$$
(1.22)

where $|\mathcal{D}|$ is the area of \mathcal{D} . The mean predicted CPUE over \mathcal{D} is a measure of average stock density. However, several other statistics related to stock abundance and distribution, such as the extent of the stock, can be estimated from the spatial cumulative distribution function. Furthermore, uncertainties in these statistics can be carried through the analysis using the posterior distribution of the spatial HBM.

Spatial cumulative distribution function

Suppose that a spatial hierarchical Bayesian model has been fit to fishery CPUE data. Equations 1.6, and 1.14 or 1.18, define a predictive model for the response surface y(s) for all possible locations in \mathcal{D} , where the surface height is the predicted CPUE at each location. Suppose that this predicted CPUE surface is integrated, beginning from its lowest predicted CPUE point and ending at some higher predicted CPUE point which I denote u. If the spatial domain \mathcal{D} is normalized so that its area is equal to one, then the integration can be seen as a cumulative distribution function for CPUE over space - the spatial cumulative distribution function (scdf). The scdf can be used to evaluate the proportion of the fixed region \mathcal{D} with response y(s) less than or equal to some value u, and is defined as

$$F_{\mathtt{scdf}}(u) = Pr[\mathbf{s} \in \mathcal{D} : y(\mathbf{s}) \le u] = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} z_u(\mathbf{s}) d\mathbf{s}$$
(1.23)

where $|\mathcal{D}|$ is the area of \mathcal{D} , $z_u(s) = 1$ if $y(s) \leq u$ and zero otherwise (Banerjee et al., 2004). Notice that the scdf is itself random because $z_u(s)$ is based on the posterior probability of y(s). In other words, the uncertainty in the posterior predictive distribution for y(s) is carried through to the calculation of the scdf.

The scdf describes the distribution of CPUE over the spatial domain \mathcal{D} , and therefore, provides the basis for a range of summary statistics. The fiftieth percentile of the scdf is a natural choice for an index of stock abundance because it represents the median predicted y(s) over \mathcal{D} . However, there are other indices that may reveal changes in the spatial characteristics of the stock over time. For instance, at some level of CPUE defined by the analyst, the stock density may be considered to be so low that either the stock is absent or not economically exploitable. The area occupied by the stock can be estimated by tracking the proportion of \mathcal{D} with CPUE higher than this critically low level. Alternatively, the range of predicted CPUE values may also be of interest as an indicator of aggregative or dissaggregative processes. An increasing range of CPUE values coincident with decreasing mean spatial CPUE might imply stock aggregation where fish are concentrating in certain areas and vacating others.

Equation 1.23 can be evaluated numerically using Monte Carlo integration in identical fashion to the evaluation of $\bar{y}(s)$ by randomly choosing a set of locations $\tilde{s} \in D$, and evaluating the posterior predicted response $y(\tilde{s})$ according to equation 1.21. The index of abundance corresponding to the fiftieth percentile of the scdf, which I denote y_{50} , is defined as

$$0.5 = F_{\text{scdf}}(y_{0.5}) = \Pr[\mathbf{s} \in \mathcal{D} : y(\mathbf{s}) \le y_{0.5}] = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} z_{y_{0.5}}(\mathbf{s}) d\mathbf{s}$$
(1.24)

The tenth $(y_{0.1})$ and ninetieth $(y_{0.9})$ percentiles are defined analogously. From these percentiles, I define the index of heterogeneity \mathcal{H} as

$$\mathcal{H}(0.9, 0.1) = y_{0.9} - y_{0.1} \tag{1.25}$$

which is meant to capture the disparity, i.e. the range, between those areas with the highest and the lowest expected CPUE values. The index of heterogeneity might be calculated for some other combination of CPUE percentiles since alternative percentiles may be more appropriate depending on the fraction of \mathcal{D} that is actually occupied by the stock and/or the fishery. Trends in $\mathcal{H}(0.9, 0.1)$ will not be obvious, for example, if the stock occupies only a small fraction of the spatial domain. In such a case, an index of heterogeneity such as $\mathcal{H}(0.95, 0.8)$ might be more appropriate because the bulk of the area \mathcal{D} will have a near zero stock density so that small scdf percentiles will be consistently negligible.

Changes over time in the area occupied by the stock can also be tracked using the scdf. If some CPUE value y_{abs} is defined for which the underlying stock is assumed to be effectively absent, or perhaps economically unexploitable, then the scdf value $F_{scdf}(y_{abs})$ is the proportion of \mathcal{D} in which the stock is effectively absent. The area occupied by the stock in year t is then

$$\Psi_t = |\mathcal{D}| \left(1 - F_{\mathsf{scdf}}(y_{\mathsf{abs}})_t \right) \tag{1.26}$$

Determining a value for y_{abs} that is representative of effective absence may be difficult without information relating CPUE to the underlying abundance. There is no standard way to define low abundance or effective absence directly from CPUE, but one approach is to examine the distribution of historical CPUE on a set-by-set basis, and choose y_{abs} such that it defines the lower tail of the distribution. Determining a value for y_{abs} that is representative of economic exploitability is possible by making enquiries into vessel profitability as a function of CPUE.

1.3 Discussion

By using spatial HBMs to analyze fishery data and generate stock abundance indices, the relationship between CPUE and average stock density is treated on a local scale. The dimensional oversight is addressed by treating CPUE spatially, weighting information according to where it was collected. Instead of assuming that aggregate fishery CPUE is related to total stock abundance, the spatial hierarchical approach makes the weaker assumption that the CPUE at a given location is proportional to the stock density at that location. The location-specific CPUEs are then integrated over all possible locations to obtain a relative measure of the total stock biomass. The typical aggregate approach, in which each CPUE observation has equal weight in the calculation of the stock abundance index, is a measure of average fishery performance.

The case of the northern cod stock is an excellent example of the shortcomings of the typical aggregate approach to addressing the dimensional oversight (Rose and Kulka, 1999). It was recognized that the assumption of CPUE proportional to total stock biomass did not hold, and that some correction had to be made. However, when the stock collapsed and criticisms were targeted at the stock assessment methodology, the suggested remedy was not to find and use an index that was proportional to average stock density. Rather, the suggestion was to fit a power function to the non-proportional relationship between average fishery CPUE and total stock biomass (Hutchings, 1996; Hutchings and Myers, 1994; Walters and Maguire, 1996). This solution relies not only on the power function being a suitable model for the relationship, but also on the assumption that it is stationary over time. The assumption of stationarity is particularly problematic because the power function is an emergent relationship based on factors ranging from the distribution of the stock to the management of the fishery, any or all of which may be changing. Even in the unlikely event that the unseen factors contributing to the CPUE-biomass relationship remain constant, the relationship will only be estimable if there is enough contrast in stock biomass. In the northern cod case, the chosen stock abundance index was unable to track an overall decline because average fishery CPUE remained stable even as local abundance declined dramatically over most of the fishing grounds (Rose and Kulka, 1999). Such a situation undoubtedly reduces managers' confidence in stock assessment estimates of stock status.

Another advantage of the spatial HBM is that the Bayesian approach ensures that information and uncertainty are properly handled throughout. The Bayes posterior distributions contain all of the information that is available for model parameters, assuming that the model is correctly specified. As such, spatial HBM derived indices of abundance are amenable to inclusion in integrated stock assessment models, which are considered to be important tools in modern stock assessment (Maunder et al., 2006). The posterior probability distributions can also be used to guide future research activities (e.g. survey designs) for resolving model uncertainties. Maps of variance in the predicted CPUE can be useful in deciding where to focus survey effort.

Extension: Spatial HBM for presence-absence data

The hierarchical Bayesian framework can be used for fishery data other than CPUE. For example, many species caught in commercial fisheries occur very infrequently. Although such rare species are probably not targeted quota species, their status may be of conservation concern. In establishing benchmarks for listing endangered or threatened species, the Word Conservation Union (IUCN) has developed a set of criteria (IUCN, 2001). When absolute abundance information is not available, trends in either relative abundance or in the spatial distribution of a species based on presence/absence can be used to list that species as threatened or endangered.
Assuming that the probability of catching a rare species is related to its presence, the spatial HBM can be modified to model the set-by-set probability of capture rather than set-by-set CPUE by using a Bernouilli distribution rather than a normal distribution in the response stage. The Bernouilli distribution represents a stochastic process that can have either a success or failure outcome. For the presence-absence version of the point-referenced spatial HBM, the response is

$$y(s) \sim \text{Bern}(p(s))$$
 (1.27)

where y(s) is the vector of observations and p(s) is the probability that the species is observed at s. If the rare species is caught in set i, then $y(s_i)$ is one, otherwise $y(s_i)$ is zero. Similar to the CPUE analysis, the model for the probability p(s) is a GLM formulation with spatial effects added, i.e.

$$logit(\boldsymbol{p}) \sim N_n(X\boldsymbol{\beta} + \boldsymbol{w}, \tau^2 I)$$
(1.28)

Here, the link function is the logistic, but otherwise the model is very similar to the spatial CPUE model. Spatial effects w(s) are defined as in equation 1.14 or equation 1.19, so the probability of observing a species in a given fishing event is correlated in space. Figure 1.6 is a depiction of the spatial presence-absence model. Note that compared to the spatial CPUE model, only the response portion of the model has been modified.

As in the CPUE model, a measure of the area occupied can be obtained from the scdf by defining a value for y_{abs} , the probability of catch associated with effective absence. The probability of catch is modeled on the logit scale (equation 1.28), so a strictly zero probability of catch is never predicted. As a result, a non-zero y_{abs} must be chosen when defining area occupied. Unfortunately, the scdf value corresponding to very small values of y_{abs} may be very uncertain. In cases where the uncertainty in the area occupied calculated from the scdf is large, an alternative approach that is compatible with IUCN criteria is to assume that the probability of catch is proportional to the stock density. This assumption is most appropriate in cases where the probability p of observing the species is very small, and there are very few individuals caught in a set. Under this assumption, a time series of the mean probability of catch



Figure 1.6: Hierarchical model for presence-absence fisheries data.

over space $(\bar{y}(s))$ will track trends in relative abundance. This alternative approach may be preferred for species that are caught extremely rarely because of the difficulty associated with specifying y_{abs} and the uncertainty associated with estimating the area occupied based on relatively few sightings. Although absolute abundance information is preferred when using IUCN criteria, the use of time series of either the area occupied or relative abundance is acceptable for assessing population status. Thus, spatial HBMs can be used for listing stocks even in cases where absolute abundance information is not available or too costly to collect, and only incidental catch records are available.

Shifting stock distributions

There is sometimes a need to model directional changes in the distribution of fish stocks. For example, fish stocks may shift their geographic distribution in response to long-term changes in climate (Perry et al., 2005). It is imperative that the spatial

domain \mathcal{D} encompass the entire stock for the entire time series in order to properly address the dimensional oversight (Walters, 2003). However, fishery-dependent data may not be available for large portions of \mathcal{D} where the stock is absent, with the possible exception of multi-species fisheries in which effort occurs throughout \mathcal{D} . In cases where there is strong evidence of a regime shift, it may be appropriate to perform a separate independent analysis for the period before and the period after the shift, each with a separately defined \mathcal{D} . Alternatively, strong priors on spatial effects could be used to constrain stock density to be low for regions in which the stock is thought to be vacant, either because the stock has yet to colonize the region, or because the stock has vacated the region to colonize elsewhere.

The example of a fish stock exhibiting directional change in its distribution is strong motivation to ensure that the spatial domain \mathcal{D} is defined to be large enough to encompass the entire stock. Too often, survey effort is targeted at the areas thought to represent the *core* of the stock. This is due to survey designs that attempt to estimate stock abundance to maximum precision by sampling in proportion to presumed stock density. However, this objective may be short-sighted given that the fishery will already be targeting what it thinks are the core areas of the stock, collecting stock density information in the process. Rather, it may be of value to target some of the survey effort at peripheral areas to ensure that \mathcal{D} is suitably defined.

MacCall (1990) derives expected spatial distributions of fish stocks given densitydependent habitat selection and logistic population growth. His results predict that fish will distribute themselves according to habitat suitability, which leads to continuous patterns of spatial stock density. If this is the case, then the spatial HBM approach outlined here is very well suited to describing such patterns. Environmental covariates that influence CPUE (and inferred density) can be included as linear effects, whereas local processes either for which there is no available data, or which are not measurable, are included as random spatial effects. Maps of predicted CPUE show patterns in assumed stock density, which MacCall (1990) links to areas of high habitat suitability. Such areas are priorities for conservation, and could be identified using the results of the spatial HBM. The theoretical treatment of MacCall (1990) also underscores the need to sample marginal areas, which become desirable habitat when increases in stock abundance lead to increases in core area density, or when changing environmental conditions make current core areas less suitable.

Summary

Inferring fish stock abundance and various other parameters from commercial fisheries data can be problematic. Heterogeneities in the stock and the fishery mean that aggregate commercial CPUE is often not proportional to total stock biomass. In addition, spatial autocorrelation can induce partial pseudoreplication if observations are assumed to be conditionally independent, as is the case in typical GLM analyses used to standardize commercial CPUE. Treatment of the problem with spatial HBMs can recognize and estimate spatial autocorrelation in commercial fishery data, allowing for prediction at unsampled points by borrowing information across space. Unfortunately, the size of typical fishery datasets makes the calculation of spatial HBM posterior distributions challenging. One solution to the computational big n problem is to use a conditionally autoregressive areal model in which observations occurring in a small area share the spatial effect for that area. The Bayesian framework allows the incorporation of prior information, and results in a predictive model for CPUE. Once posterior distributions are evaluated, a series of statistics that are spatially referenced and proportional to area, rather than fishing intensity, can be calculated directly using the spatial cumulative distribution function. These scdf-derived statistics include indices of abundance, measures of spatial range, and measures of stock aggregation. The method is useful not only for targeted fishery CPUE analysis, but also for rare species of conservation concern where presence-absence records are available.

Chapter 2

A comparison of spatial and aspatial assessments of BC sablefish using commercial data

2.1 Introduction

Using fishery-dependent data such as commercial catch records to ascertain trends in the state of fish stocks is appealing for at least two reasons. First, much of the effort to collect the data is already being made in the course of fishing activities. Second, catch records may be the only (or the most complete) data source available. However, the information collected from fishery observations must be properly interpreted in order to be useful in determining stock trends. Proper interpretation means recognizing that common statistical model assumptions may not hold in all cases. In particular, the assumption of independence of observations made in standard generalized linear model analyses will almost certainly not hold for fishery data because some amount of spatial autocorrelation is expected when dealing with observations made close together in space (Nishida and Chen, 2004). In addition, because the primary goal of fish harvesters is to catch fish, not to survey the population, the volume of observations will be weighted towards areas where fish tend to occur at the expense of areas where fish do not tend to occur (Campbell, 1998, 2004; Walters, 2003). In other words, aggregated commercial catch rate data will tend not to be representative of the average underlying fish population density.

Despite some strong assumptions, commercial fishery data is commonly used to assess fish stocks, whether due to necessity or the lure of abundant data. In typical aspatial generalized linear model (GLM) analysis of fishery catch rate data, the assumption is that, conditional on specified covariates, fishery observations are independent replicate measurements of a constant but unknown mean stock density. However, fishery data with unbalanced sampling effort (i.e. fishing that is not random with respect the area over which the stock is distributed) and spatial autocorrelation will violate the aspatial GLM assumptions.

In this chapter, I analyze commercial catch and effort data form the British Columbia (BC) sablefish fishery using a spatial hierarchical Bayesian model (HBM) that accounts for unbalanced sampling and spatial autocorrelation. The approach taken makes the weaker assumption that the stock density *at the location of a fishing set*, not the *total* stock biomass, is proportional to the set-by-set catch per unit effort (CPUE).

The first objective of this analysis is to determine the consequences of using the spatial HBM for stock assessment of BC sablefish. To accomplish this objective, I fit a Schaefer biomass dynamics model using each of four relative stock biomass indices. The four indices are: (1) the nominal fishery CPUE; (2) the aspatial GLM year effects; (3) the mean spatial CPUE, predicted using the HBM; and (4) the spatial median CPUE, predicted using the HBM. The second objective of this analysis is to determine whether the area occupied by the commercially exploitable stock has declined over the fishery time series, and whether aggregative or disaggregative processes appear to be occurring for BC sablefish. Beyond calculating an index of abundance, the results of fitting the spatial HBM can be used to map predicted CPUE over space. These predicted CPUE maps can then be used to calculate area occupied and spatial distribution statistics.

2.1.1 British Columbia sablefish fishery

Sablefish (Anoplopoma fimbria) is a deepwater fish species found throughout the North Pacific Ocean, from Mexico to Japan. The bathymetric range of adult sablefish extends from 200 metres to greater than 1900 metres depth, with the highest densities occurring between 350 and 900 metres (McFarlane and Beamish, 1983). Juvenile sablefish inhabit inshore waters shallower than 200 metres and move to deeper waters as they mature. The analysis of tagging studies by Kimura et al. (1998) suggest that there are two Northeast Pacific sablefish stocks that mix off the BC coast near southwest Vancouver Island. Tagging studies suggest that most sablefish off the coast of British Columbia (BC), Canada, are relatively sedentary, and that within stock areas, movement is low. Kimura et al. (1998) observed more movement of tagged sablefish in the northern stock than in the southern stock for sablefish tagged in US waters. For sablefish tagged in Canadian waters, Beamish and McFarlane (1983) found that more than four in five tag recaptures occur less than 200 kilometres from the point of release.

Sablefish have been harvested off the BC coast since the late nineteenth century primarily by longline hook or longline trap gear (McFarlane and Beamish, 1983). During the sixteen year period from 1990 to 2005, more than seventy-five percent of the sablefish quota was taken using longline trap gear (Haist et al., 2005), which consists of modified Korean conical traps that are baited with a combination of squid and hake (Wyeth et al., 2004). A typical longline trap set deploys about sixty traps for a soak time of about 48 hours. I focus my analysis on the BC sablefish long-line trap catches because the majority of the catch is taken using traps. Prior to the enforcement of Canada's Exclusive Economic Zone in 1977, Japanese vessels accounted for the majority of catch. From the time that Japanese vessels were excluded in 1977 until 1989, the Canadian fishery was regulated by various mechanisms of effort limitation designed to achieve target total allowable catches (Haist et al., 2004). However, beginning in 1990, the Canadian sablefish fishery changed to an Individual Transferable Quota (ITQ) management system. My analysis uses data from this later period (1990 - 2005) exclusively because transitions between management systems

may affect fishery performance (i.e. CPUE) independent of actual changes in stock biomass. However, the period of study was not without changes in management, the most notable changes being the introduction of trap escape rings in 1998, and the use of a two-year quota for 2002 and 2003 (Haist et al., 2004). Most of the two year quota was taken in 2002, resulting in very little quota remaining for the 2003 season. Thus, the noticeable increase in fishery CPUE from 2002 to 2003 may be the result of changes in targeting by the fleet due to the shortened 2003 season rather than an increase in stock biomass. This could cause stock production to be overestimated for 2003 if fleet-targeting effects are not accounted for.

Fisherv-independent surveys for BC sablefish include (Kronlund et al., 2002): (1) standardized survey, which began in 1988, and has remained largely unchanged since 1990; (2) stratified random survey, which began in 2003; and (3) tagging program that began in the offshore in 1991 and was modified to include BC mainland inlets in 1994. The standardized survey was designed to index historic offshore trap fishing areas (Wyeth and Kronlund, 2003). The survey began in 1988 with eight areas, and was modified to include a ninth area in 1994. Not all areas have been surveyed in each year. Longline trap sets are targeted at several depth ranges, with the exact location of the set left to the discretion of the survey skipper. The more recently implemented stratified random survey, described in Wyeth et al. (2004), divides the BC offshore into five bordering areas that cover most of the area targeted by the commercial trap fishery (Figure 2.1). The stratified survey further divides these areas into three depth strata, and survey set locations are randomly chosen from each strata. Five sets are made at random locations in each depth strata and area combination for a total of seventy-five stratified survey sets per year. The design of the stratified random survey is on the whole more representative of the entire BC sablefish stock than the standardized survey, which focuses primarily on the core fishing areas (Kronlund et al., 2002).

2.2 Methods

2.2.1 Data characteristics and exploratory analysis

Spatial catch and effort data

The data used in this study were obtained from a database of sablefish commercial catch and effort compiled by the Department of Fisheries and Oceans Canada (DFO). The DFO database contains information on commercial fishing events as recorded in fishing logbooks and corroborated by a dockside monitoring program. Data from the commercial sablefish fishery deploying longline trap gear were used in this analysis; the records of interest for the individual fishing sets are summarized in Table 2.1. Sets with incomplete or missing records were assumed to be missing at random and were excluded from the analysis. The resulting dataset included 45,034 fishing sets.

Name	Description	Values or Units
Year	The year of the fishing event	1990 to 2005
Quarter	The quarter (season)	November-January,
		February-April,
		May-July, or August-October
Skipper	Numerical skipper code	1 to 50
Vessel	Numerical vessel code	1 to 29
Catch	Retained round weight of sablefish	Kilograms
Traps	Traps deployed in the set	Number of traps
Lat	Mean latitude of fishing set	Degrees
Lon	Mean longitude of fishing set	Degrees
Depth	Mean fishing depth	Depth in metres

Table 2.1: Data extracted from the commercial sablefish database.

Beginning in 1998, some harvesters began using traps with escape rings to allow undersized fish to escape. Escape rings were made mandatory beginning in 1999. Variation in trap effectiveness may include factors such as the use of escape rings, and differences in fishing practise such as trap baiting, length of string, and soak time. Unfortunately, nearly all vessels began using escape rings only once they became mandatory in 1999, so the effect of escape ring introduction on CPUE is confounded with changes in abundance from 1998 to 1999. I adopt the convention of previous assessments of BC sablefish, performed by sablefish biologists who have attributed other sources of variation in trap effectiveness to the practises of the vessel or skipper associated with the fishing set, by using "trap" as the unit of effort (e.g. Kronlund et al., 2003).

Bathymetry data were obtained from the Scripps Institution of Oceanography (http://topex.ucsd.edu/WWW_html/mar_topo.html) via the R (R Development Core Team, 2007) package PBSmapping (Schnute et al., 2006). The bathymetry data have a longitudinal resolution of approximately 0.033 degrees (about 2.3 kilometres) and a latitudinal resolution of approximately 0.019 degrees (about 2.1 kilometres). These bathymetry data, although not necessary for fitting the model because depth is recorded for each fishing set, were used to generate the posterior predicted CPUE surfaces for mapping and subsequent analysis.

Study area

The spatial domain \mathcal{D} was defined by assuming that the spatial and depth distribution of the exploited stock overlaps with past fishing activity. The stratified random survey divides the BC offshore into five bordering areas that cover most of the area targeted by the commercial trap fishery, so I defined the spatial domain by dividing the union of the five areas into 10-km square grid cells, and eliminating those cells that weren't fished by the commercial trap fishery during the period from 1990 to 2005 (Figure 2.1). Thus, the spatial domain included all of the 10-km grid cells in the five stratified survey areas that received some fishing effort during the period. The domain \mathcal{D} was further refined by limiting the depth range of the exploited stock to past (1990 to 2005) fishing activity, which was conducted between depths of 176 m and 1,312 m.



Figure 2.1: Map of study area. Shaded areas are the British Columbia sablefish random stratified survey areas. Overlain are the 10km grid cells with at least one commercial trap set during the period 1990-2005.

Exploratory analysis of spatial catch and effort

An exploratory analysis of spatial catch and effort data reveals two concerns for indices of abundance based on nominal fishery CPUE or GLM year effects. First, although there is no obvious trend in the number of cells fished by year (Figure 2.2a), the area of the grid cells used by the fishery has varied over the period from approximately 12,100 to 28,200 square kilometres. As a result, nominal fishery CPUE may not be measuring the stock in a representative way, with some years focused on a smaller region of the BC coast, and other years focused a larger region. Second, the number of cells which received at least one percent of the total effort shows a trend over the period examined, suggesting a spatial concentration of effort over time (Figure 2.2b). These two aspects of the fishery effort underscore the need for a spatial approach to address the dimensional oversight.

Differences in the number of fishing sets made in each grid cell can also present analytical challenges due to an unbalanced statistical design. For example, the skipper effect on CPUE may be confounded with location of the fishing event if the skipper has only fished in a small number of locations. However, the concern over such confounded effects is mitigated when:

- 1. There are multiple skippers fishing in each cell, resulting in the location effect having contrast in skipper effects. Although this is generally true for the BC sablefish trap fishery data I examined, there are some spatial cells that were only fished by one or two skippers (Figure 2.3).
- 2. There are multiple cells fished by each skipper, resulting in the skipper effect having contrast in spatial effects. Although a large number of skippers fished in relatively few cells, there are many skippers that have fished a relatively large number of cells (Figure 2.4).

The story is similar for vessels. Although a few vessels have only fished a small number of cells, and some cells have been fished by only a few vessels, there are a number of vessels that have fished a large majority of the spatial domain \mathcal{D} . These vessels have covered many spatial cells and provide the contrast necessary to differentiate spatial





Figure 2.2: Spatial concentration of effort. (a) Annual proportion of 10-km grid cells that were not fished by the commercial trap fishery. The high point in 2003 is a consequence of the reduced fishery quota for that year. (b) Annual proportion of spatial cells receiving at least one percent of the total trap fishery sets.



Figure 2.3: Number of skippers by grid cell. The total number of skippers fishing in each 10-km grid cell for the BC sablefish trap fishery 1990-2005.



Figure 2.4: Number of grid cells fished by skipper. The total number of 10-km grid cells fished for each of the skippers in the BC sablefish trap fishery 1990-2005.

2.2.2 Data analysis

The first objective of this study was to determine the consequences of using the spatial HBM for stock assessment of BC sablefish. The first step in addressing this objective was to calculate the four indices of abundance for comparison. The first index the nominal CPUE index, which involved calculating the mean CPUE directly from the set-by-set commercial catch and effort. The second index was an aspatial GLM index, which involved fitting a GLM to the commercial CPUE data and extracting year effects. The third and fourth indices were based on a spatial hierarchical Bayesian model (HBM), beginning with the GLM structure and adding spatial effects. Notation for the various quantities in the models (Table 2.2) follows the notation developed in Chapter 1.

Nominal CPUE index

The nominal CPUE index was calculated by averaging set-by-set fishery CPUE for each year from 1990 to 2005, i.e.

$$Q_t^{\text{CPUE}} = \frac{1}{n_t} \sum_i y_{i,t} \tag{2.1}$$

where i indexes fishing set and n_t is the total number of fishing sets that took place in year t.

Aspatial GLM index

Using the nominal CPUE index as a measure of relative stock abundance assumes that, within a year, each fishing set is a replicate measurement of the average stock density. However, factors other than which year the fishing set was made may influence CPUE, such as the effect of individual fishing vessels or of fishing vessel characteristics, the depth of fishing set, the underlying habitat type, or the season in which the fishing event took place. The generalized linear model allows the effects of these other factors on CPUE to be included and estimated. Once the GLM has been fit to the data, the *Year* effect can be used as the relative stock abundance index.

Symbol	Description
n	Number of fishing sets (45,034)
C	Set-by-set Catch
$oldsymbol{E}$	Set-by-set number of <i>Traps</i>
y	CPUE - see equation 1.1
$f(oldsymbol{y})$	Natural logarithm of CPUE
p	Number of linear covariates
X	Design matrix of p covariates - Examined covariates include:
	Year, Quarter, Vessel, Skipper, Depth, Squared Depth
$oldsymbol{eta}$	Vector of p linear effects
$ au^2$	Variance of Response
s	locations of the n fishing sets
$w(oldsymbol{s})$	set-specific spatial random effects
$w(\boldsymbol{\varsigma})$	grid cell-specific spatial random effects
R	Number of time blocks for spatial effects
σ^2	Variance of the R CAR prior(s)
t	Index of year
i	Index of fishing set
j	Index of linear covariates $\boldsymbol{\beta}$
k	Index of grid cell
r	Index of time block of spatial effects

Table 2.2: Symbols used in specifying the BC sablefish commercial trap fishery models.

I fit fifteen candidate generalized linear models (Table 2.3) specified according to equations 1.1, 1.2, and 1.3 using the glm function in the statistical computing environment R (R Development Core Team, 2007). The design matrix in equation 1.2 included an intercept and covariate values for *Year*, *Quarter*, *Depth*, and squared *Depth*. Some models also included *Skipper* and others included *Vessel*. Covariates *Year*, *Quarter*, *Skipper*, and *Vessel* were all treated as factors, whereas *Depth* and squared *Depth* were continuous. Previous analysis of the available survey data showed that the highest CPUEs are observed at depths ranging from 457 to 824 m (250 to 450 fm), which suggested a quadratic relationship between the logarithm of CPUE and depth (Wyeth et al., 2004). The quadratic relationship was also allowed to vary over time by estimating separate *Depth* and squared *Depth* effects for time blocks lasting one, two, four, eight, or all sixteen years.

Table 2.3: Generalized linear model selection statistics. Covariates Year, Quarter, Depth, and squared Depth were included in all models examined. Some models also included Skipper or Vessel effects. Except for models 1 - 3, Depth and squared Depth effects were aggregated into multi-year time blocks. The chosen model is shown in **bold**.

	Covar	riate		Statistic	
Model	Skipper	Vessel	Depth Blocks	AIC	Deviance
1	Yes	No	16	321598	3315253
2	No	Yes	16	322125	3357362
3	No	No	16	327034	3748677
4	Yes	No	8	321984	3346125
5	No	Yes	8	322445	3383706
6	No	No	8	327472	3788000
7	Yes	No	4	322228	3365479
8	No	Yes	4	322681	3402676
9	No	No	4	327716	3809967
10	Yes	No	2	322389	3378169
11	No	Yes	2	322859	3416825
12	No	No	2	327980	3833055
13	Yes	No	1	322395	3378884
14	No	Yes	1	322877	3418431
15	No	No	1	328039	3838395

Akaike's Information Criterion (AIC) and the residual deviance were calculated for each of the candidate model fits. The AIC (Akaike, 1973) is an information theoretic model selection criterion that balances model parsimony with model fit. Models with lower AIC values are considered to be better. In addition to the AIC, the *Depth* and squared *Depth* parameters were checked to ensure that a realistic relationship between CPUE and *Depth* was being estimated. Some time blocks lacked fishing sets made at relatively deep or shallow depths, resulting in poor contrast in *Depth*. Those models that estimated a positive *Depth* effect and a negative squared *Depth* effect in some time blocks were removed from consideration because they predict relatively higher CPUE for deep and shallow depths compared to intermediate depths. Such estimates of CPUE versus *Depth* have no biological foundation, and are purely the result of poor contrast in *Depth*. If CPUE increased as the square of depth, then harvesters would target locations that are as deep as possible, which is clearly not the case.

I extracted the *Year* effects from the best model as the GLM index of relative abundance. Because the raw *Year* effects are for CPUE on the natural log scale, I transformed the effects to the arithmetic scale according to

$$Q_t^{\mathsf{GLM}} = e^{\beta_0} e^{\beta_{j_t}} \tag{2.2}$$

where β_{j_t} is the linear model effect corresponding to year t. Only differences between the base year (1990) and subsequent years are estimated by the *Year* effects, so in the base year, β_{j_1} is zero and $Q_1^{\text{GLM}} = e^{\beta_0}$.

Spatial Hierarchical Bayesian Model

The first step in estimating the spatial stock biomass indices was to formulate and fit a series of spatial HBMs. The BC sablefish data set contains 45,034 entries, which precludes a continuous spatial effects model such as the one depicted in Figure 1.4 due to the big *n* problem. Therefore, I used the CAR prior (equation 1.18) for spatial effects, defined on the 10-km grid as shown in Figure 2.1. This is equivalent to the CAR spatial hierarchical model (figure 1.5). I use the same notation as in the model definition (equations 1.1, 1.2, 1.4, 1.5, 1.6, 1.12, 1.19, and 1.20); Table 2.2 contains a brief description of specific BC sablefish symbol definitions used in the CAR spatial HBM.

Set-specific CPUE was computed as in equation 1.1, and was subsequently logtransformed (as in equation 1.6) according to

$$f(y_i) = \log_e(y_i + 0.001) \tag{2.3}$$

to allow \boldsymbol{y} to be modeled using a normal error distribution. The value 0.001 was added to each observation to allow zero catches in the analysis.

The design matrix included the same covariate values as the aspatial GLM analysis except that, in similar fashion to the aspatial GLM analysis, *Depth* and squared *Depth* effects were estimated for separate time blocks lasting one, two, or four years. Aggregating *Depth* effects across years was necessary due to insufficient contrast on a finer time scale. The spatial HBMs also included either four time blocks of spatial effects $w(\varsigma_r)$, r = 1...4, defined according to equation 1.19, or one set of spatial effects covering all sixteen years. This allowed patterns in spatial variation to vary over time for some models. The number of *Depth* blocks and spatial effects were chosen so that the candidate HBMs covered a range of temporal complexity from sixteen *Depth* blocks and four spatial blocks to four *Depth* blocks and one spatial block (Table 2.4).

The spatial HBMs were fit using a Gibbs sampler (Geman and Geman, 1984) that was custom-coded in R. The Gibbs sampler is a Markov Chain Monte Carlo (MCMC) method for estimating the joint posterior distribution of the model parameters by drawing from the posterior distribution of each model parameter, conditional on the values of the other parameters. Compared to other MCMC methods such as the Metropolis Hastings algorithm (Hastings, 1970; Metropolis et al., 1953), the Gibbs sampler has the advantage that it can sample from the joint posterior without the need for any *tuning*. Since hierarchical models can become rather complex, it is also generally easier to think about and describe hierarchical model parameters conditionally.

Vague priors were used for all model parameters that were not specified conditionally. Each linear coefficient in the column vector β (equation 1.4) was given a normal prior with mean zero and large variance of 1000. The variance of the model residuals, τ^2 (equation 1.5), and the variance of the CAR priors σ^2 (equation 1.12), were each given Inverse-Gamma priors with mean 1 and 0.001 degrees of freedom. Combined with distributions for other model parameters, the joint posterior probability

Table 2.4: BC sablefish spatial hierarchical Bayesian model selection statistics. Covariates Year, Quarter, escape rings, Depth, and squared Depth were included in all models examined. Models with aggregated Depth (and squared Depth) blocks were considered due to apparent data coverage issues (see text). For each model, the average deviance (D_{avg}) , the best-fit deviance $(D_{\hat{\theta}})$, the effective number of parameters ν , and the deviance information criterion (DIC) were calculated. Model F, in **bold**, was chosen based on its low DIC statistic.

	Covar	riate	Blo	\mathbf{cks}	Statistic			
Model	Skipper	Vessel	Spatial	Depth	D_{avg}	$D_{\hat{\theta}}$	ν	DIC
A	Yes	No	4	16	45037	44940	97	45135
В	No	Yes	4	16	45039	44961	78	45116
\mathbf{C}	No	No	4	16	45038	44984	55	45093
D	Yes	No	4	8	45038	44956	82	45120
Ε	No	Yes	4	8	45042	44980	62	45104
\mathbf{F}	No	No	4	8	45026	44986	40	45065
G	Yes	No	4	4	45036	44961	75	45110
Η	No	Yes	4	4	45038	44984	55	45093
Ι	No	No	4	4	45040	45008	32	45072
J	Yes	No	1	4	45028	44949	79	45107
Κ	No	Yes	1	4	45033	44975	58	45092
L	No	No	1	4	45038	45005	33	45071

distribution for the spatial HBM is

$$p(\tau^{2}) \prod_{r=1}^{R} CAR(w(\boldsymbol{s})|\sigma_{r}^{2}) p(\sigma_{r}^{2}) \prod_{j=1}^{p} p(\beta_{j}) \prod_{i=1}^{n} p(y_{i}|\boldsymbol{\beta}, w(s_{i}), \boldsymbol{\sigma}^{2}, \tau^{2}) \quad (2.4)$$

$$\propto IG(\tau^{2}|1, 0.001) \prod_{r=1}^{R} CAR(w(\boldsymbol{s})|\sigma_{r}^{2}) IG(\sigma_{r}^{2}|1, 0.001)$$

$$\times \prod_{j=1}^{p} N(\beta_{j}|0, 1000) \prod_{i=1}^{n} N(f(y_{i})|X_{i}.\boldsymbol{\beta} + w(s_{i}), \tau^{2})$$

Because of the conjugacy of model priors, the conditional distribution for each model parameter is available in closed form, which enables a Gibbs sampling approach to generating the approximate joint posterior. I developed a Gibbs sampler to draw samples from the conditional posterior distribution of each model parameter in turn. The conditional distribution for each model parameter can be extracted from the joint posterior by treating all other model parameters as constants (e.g., Su et al., 2004).

When choosing an appropriate model for the data, models with a high predictive power for future data are desirable. The deviance information criterion (DIC) is an estimate of the expected predictive deviance of a given model (Spiegelhalter et al., 2002). Models having a lower DIC are expected to be better out-of-sample predictors (Gelman et al., 2004). The expected predictive deviance is a function of the expected deviance of a point estimate based on a set of model parameters and the expected deviance averaged over representative posterior parameter samples.

The DIC is composed of two components, each of which can be estimated from posterior parameter samples according to

$$DIC = 2\hat{D}_{avg}(\zeta) - D_{\hat{\theta}}(\zeta) \tag{2.5}$$

where the generic symbols ζ and θ represent data and parameters in the model, respectively. The average expected deviance is

$$\hat{D}_{avg}(\zeta) = \frac{1}{L} \sum_{l=1}^{L} D(\zeta, \theta_l)$$
(2.6)

for posterior samples $l = 1, \dots, L$, and the expected deviance of a point estimate is

$$D_{\hat{\theta}}(\zeta) = D(\zeta, \hat{\theta}(\zeta)) \tag{2.7}$$

An appropriate choice for the point estimate $\hat{\theta}$ is the mean of the posterior simulations. Calculating both the average expected deviance and the expected deviance of a point estimate requires estimates of the deviance, which is a measure of the discrepancy between the data and model. It is defined as twice the negative log likelihood

$$D(\zeta, \theta) = -2\log_e p(\zeta|\theta) \tag{2.8}$$

The expected deviance of a point estimate and average expected deviance can also be used to estimate the effective number of parameters in the model. Hierarchical model structures, the value of priors and hyperpriors, and observed data, all constrain

the value of parameters such that the effective number of parameters is less than the actual number of parameters in the model. The results of model fits reflect these constraints. The effective number of parameters

$$\nu = \hat{D}_{avg}(\zeta) - D_{\hat{\theta}}(\zeta) \tag{2.9}$$

is a measure of how much actual fitting of the model to the data occurs. It is the number of unconstrained parameters in the model, and represents the effect on the deviance of fitting the model.

Having selected the best candidate spatial HBM, I estimated the spatial CPUE indices by stratified Monte-Carlo integration of the spatial HBM CPUE surface. In the Monte-Carlo integration scheme, I randomly chose twenty-five points in each of the spatial grid cells in each year, and predicted CPUE at each point for each of 3,000 posterior parameter samples. To calculate a predicted CPUE at each point, covariate values and levels had to be specified. The Year covariate was determined by the integration scheme (twenty-five points in each cell in each year). I set the Quarter effect to the fall (August-October) to allow future comparisons to indices derived from surveys, which are done in the fall. Finally, I calculated *Depth* and squared *Depth* by interpolation of bathymetry data using multilevel B-splines in the R package MBA (Finley and Banerjee, 2006). The resulting model-predicted CPUE surface consisted of 113,200 points $y(\tilde{s})$ over sixteen years, each point estimate including uncertainty in the form of 3,000 predictions, one for each posterior sample. Extrapolations beyond the depth range historically fished were omitted when calculating statistics. Mean predicted CPUE and eighty percent credibility intervals were calculated at each point to generate final results. The mean spatial CPUE index was calculated as the average of the Monte-Carlo samples in each year using

$$Q_t^{\text{HBM1}} \equiv \bar{y}(\boldsymbol{s}_t) \approx \frac{1}{n_{\mathcal{D}t}} \sum_{\mathcal{D}} y(\tilde{\boldsymbol{s}_t})$$
(2.10)

where $n_{\mathcal{D}t}$ was the number of locations \tilde{s}_t that were within the depth range 176 m - 1,312 m in year t.

I used the spatial cumulative distribution function to generate estimates of the median spatial CPUE index $Q^{\text{HBM2}} \equiv y_{50}$ (equations 1.23 and 1.24) and the 10th, 30th,

70th, and 90th percentiles for each year. The index of heterogeneity $\mathcal{H}(0.9, 0.1)$, which is the difference between the 90th and the 10th scdf percentiles, was also calculated for each year (equation 1.25). The area occupied by the stock Ψ was also calculated for each year (equation 1.26) using a low-abundance threshold of six kilograms per trap, which corresponds to the lowest tenth percentile of observed set-by-set commercial CPUE over the period 1990 - 2005. A catch of six kilograms is roughly two sablefish (Wyeth et al., 2004).

Consequences for Stock Assessment

I evaluated the consequences for stock assessment of choosing each of the four relative abundance indices derived from spatial catch and effort data. The intent was to compare the different indices of relative abundance rather than to accurately estimate actual BC sablefish population parameters. In this spirit, I chose a Schaefer surplus production model (Schaefer, 1954) for stock assessment to make comparisons because a more complicated model would have made interpreting differences in results difficult. The Bayesian production model was parametrized in terms of the theoretical maximum sustainable yield (MSY), and the optimal exploitation rate (h_{MSY}) corresponding to MSY. Posterior parameter samples were generated using a Metropolis Hastings algorithm (Hastings, 1970; Metropolis et al., 1953).

The surplus production model consists of a deterministic process model and random observation model. The deterministic process model tracks the biomass B of the stock according to

$$B_{t+1} = B_t + \varrho B_t \left(1 - \frac{B_t}{K}\right) - \mathcal{C}_t \tag{2.11}$$

where K is the unfished biomass, ρ is the rate of growth when B is much less than K, and C is the total catch of BC sablefish in all fisheries. This biomass model was parametrized in terms of the optimum harvest rate $h_{MSY} = \frac{\rho}{2}$ and the maximum sustainable yield $MSY = \frac{\rho K}{4} = \frac{h_{MSY}K}{2}$, and the state variable was redefined to be the stock biomass depletion $P = \frac{B}{K}$ (Hilborn and Walters, 1992). Using this parametrization, the Schaefer model becomes

$$P_{t+1} = P_t + 2h_{MSY}P_t(1 - P_t) - \frac{h_{MSY}C_t}{2MSY}$$
(2.12)

The full time series of catches (1965-2005) available for the BC sablefish fishery were included. The initial depletion P_1 (in 1965) was assumed to be 1 (i.e. unfished biomass level).

The observation model specified the depletion in year t to be proportional to the relative abundance index, with lognormal observation error, so that

$$Q_t = q P_t e^{\epsilon_t} \tag{2.13}$$

where q is the catchability (the fraction of P_t caught per unit effort), and $\epsilon_t \sim N(0, \tau_Q^2)$, where τ_Q^2 is the variance of ϵ . The negative log-likelihood function for this observation model is

$$-\log_e p(\boldsymbol{Q}|h_{MSY}, MSY) \propto \frac{T}{2}\log_e \left[\sum_{t=1}^T \left(\log(Q_t/P_t) - \overline{\log q}\right)^2\right]$$
(2.14)

where $\overline{\log q}$ is the conditional maximum likelihood estimate of $\log q$:

$$\overline{\log q} = \frac{1}{T} \sum_{t=1}^{T} \log(Q_t / P_t)$$
(2.15)

I specified priors for the model parameters MSY and h_{MSY} . The short time series of catch and the declining trend in the four relative abundance indices (known as a one-way trip (Hilborn and Walters, 1992)) contain limited information on either the size of the stock or MSY. Due to the declining index trends, I used an informative prior for the harvest rate parameter: $p(h_{MSY}) \sim \text{Beta}(8,64)$. This prior was based on the most recent assessment of neighbouring Alaska sablefish stocks (Hanselman et al., 2007), which estimates $h_{MSY} = 0.101$. The prior for h_{MSY} was designed on the premise that BC sablefish will have similar but not identical characteristics to Alaska sablefish, which is not unreasonable given that they may be part of the same biological population (Kimura et al., 1998). The prior for MSY was a normal distribution centred on the average historical catch, with a coefficient of variation set to one, i.e. $p(MSY) \sim N(4200, 4200^2)$. This vague prior allows the indices to determine the posterior probability distribution for MSY in tonnes.

2.3 Results

2.3.1 Model selection and resulting indices of abundance

Aspatial GLM

Comparing the candidate aspatial generalized linear models using the AIC and residual deviance suggested that the most complex model (model 1) was the most suitable GLM for describing the data (Table 2.3). However, all of the models that estimated *Depth* effects in more than two blocks (eight years per block) reversed the biologically realistic CPUE-*Depth* relationship, predicting lower CPUE in the intermediate depths compared to the deeper or shallower depths. This reversal was likely due to limited contrast in the depths fished for finer timescales. On this basis, model 1 was not chosen. Of those GLM candidates that estimated appropriate CPUE vs *Depth* relationships (models 10-15), model 10 had the lowest AIC and residual deviance (Table 2.3). The chosen model included skipper effects and two eight-year depth effect blocks.

Spatial HBM

Statistics of fits were used to compare candidate spatial hierarchical Bayesian models (Table 2.4). Differences in the average residual deviance (D_{avg}) among candidate models were small. As a result, the simpler models that did not include either *Skipper* or *Vessel* effects had lower DIC values due to their relative parsimony (fewer parameters). More complex models with *Skipper* or *Vessel* effects (e.g. A, B, D, E, etc.) did not fit the data any better than the simple models, suggesting that variability in CPUE is better explained by spatial effects than by the effect of particular skippers or fishing vessels. The model with the lowest DIC was model F, which had no skipper or vessel effects, four spatial effects blocks, and eight depth effect blocks (Table 2.4). Despite estimating 1172 individual parameters, model F only explained the equivalent of forty independent parameters, reflecting a high degree of correlation among effects. Spatial heterogeneities in the inferred stock density appear to violate the assumptions of the nominal CPUE and aspatial GLM analyses. Examination of HBMpredicted CPUE maps (Figure 2.5) shows that the CPUE-*Depth* relationship is not uniform along the coast. The spatial effects w(s), which represent the effect of processes specific to location that are not otherwise specified in the model, capture departures from the mean CPUE-*Depth* relationship. As a result, the spatial effects show areas with higher or lower CPUE than expected from aspatial covariates alone (Figure 2.6). From north to south, the time series of spatial effects maps shows three large areas with positive spatial effects: (1) the north end of the Queen Charlotte Islands (stratified survey area 5); (2) the western edge of Queen Charlotte Sound (stratified survey area 3); (3) and west of Nootka Island (central part of stratified survey area 2). The area off the west coast of Vancouver Island south of Nootka Island (stratified area 1) had consistently below average CPUE.

2.3.2 Consequences for stock assessment

All four indices show a decline over the period 1990 - 2005: nominal CPUE declined by about thirty-two percent from 1990 to 2005; aspatial GLM year effects declined by about thirty-one percent; mean spatial CPUE declined by about sixty percent; and spatial median CPUE declined by about sixty-one percent. Calculation of pairwise correlations between indices shows an apparent difference between the aspatial indices and the spatial indices (Table 2.5). Comparing the indices directly requires them to be put on a similar scale by dividing each series by its mean (Figure 2.7). When compared in this way, the nominal CPUE index and the aspatial GLM index show a sharp decline from the beginning of the time series to the late 1990s, after which they vary but do not decline further. Both of the aspatial indices also show a one-year peak in 2003, coincident with the second year of a two-year quota in which most of the allowable catch was taken in the first year. The mean spatial CPUE and the spatial median CPUE indices, on the other hand, show a more steady decline over the entire period 1990 - 2005. The 2003 peak evident in the aspatial indices is not as pronounced in the spatial indices, presumably because the aspatial indices measure



Figure 2.5: Predicted CPUE from spatial HBM for BC sablefish trap data. Predicted values of the spatial hierarchical Bayesian model (HBM) for the BC sablefish commercial trap fishery, averaged by ten kilometre grid cell. The CPUE values plotted are in units of kilograms per trap.



Figure 2.6: Spatial effects from spatial HBM for BC sablefish trap data. Spatial effects from the spatial HBM for the BC sablefish commercial trap fishery, averaged by ten kilometre grid cell. Values are plotted on the arithmetic scale. For example, all other covariates being equal, a spatial effect of 1.2 would correspond to twenty percent higher predicted CPUE than a spatial effect of 1.

fishery performance rather than average stock density.

Table 2.5: Pairwise correlations of relative abundance indices. Q^{CPUE} is the nominal CPUE index, Q^{GLM} is the aspatial GLM index, Q^{HBM1} is the mean spatial CPUE index, and Q^{HBM2} is the spatial median CPUE index. There appears to be a difference between the aspatial indices and the spatial indices.

	$Q^{ t cpue}$	$oldsymbol{Q}^{ extsf{glm}}$	$Q^{ t HBM1}$	$Q^{ t ext{HBM2}}$
$oldsymbol{Q}^{ extsf{CPUE}}$	1	0.946	0.776	0.742
$oldsymbol{Q}^{ extsf{glm}}$		1	0.630	0.589
$Q^{ t ext{HBM1}}$			1	0.991
$Q^{ ext{HBM2}}$				1

In the context of stock assessment, the mean spatial CPUE index and the spatial median CPUE index result in less optimistic estimates of maximum yield than do the nominal CPUE index and the aspatial GLM index (Table 2.6). However, stock assessments based on all four indices estimate average historic catches to be in excess of MSY. What differs among the stock assessments is the degree to which the stock has been overfished. I calculated the average ratio of estimated MSY to average catch from 1965-2005 (4200 tonnes). According to the aspatial GLM assessment, which is the most optimistic of the four assessments, catches are estimated to have been on average eleven percent higher than estimates of the theoretical maximum yield. However, the spatial assessments are more pessimistic; they estimate catches to have been on average thirty-six to thirty-nine percent higher than estimates of MSY.

The spatial assessments based on the HBM also estimate the stock to be currently more depleted than either the nominal CPUE or the GLM-based assessments (Table 2.6), which has implications for management. For example, the aspatial GLM stock assessment estimates the current (2005) depletion (B_{final}/K) to be 0.445, near the assumed stock depletion at MSY (i.e. P = 0.5). Management advice based on the aspatial GLM analysis would thus be along the lines of a slight reduction in catch from average historic levels to be near the MSY estimate of 4,050 t. The



Figure 2.7: Indices of abundance for BC sablefish. The top panels show the four time series of relative abundance: nominal CPUE index (open circles), aspatial GLM year index (solid circles), mean spatial CPUE index (open squares) and spatial median CPUE index (open triangles). The bottom panel compares the four series by scaling each series to its mean, and plotting the difference between each scaled series and the scaled nominal CPUE (e.g. $Q^{\text{GLM}}/\bar{Q}^{\text{GLM}} - Q^{\text{CPUE}}/\bar{Q}^{\text{CPUE}}$ for the scaled aspatial GLM year index, where \bar{Q} is the mean of an index time series).

spatial stock assessments, on the other hand, estimate the current depletion to be about 0.148, which corresponds to about one-third of the assumed depletion of stock biomass at MSY (P = 0.5). Management advice based on the spatial assessment would thus call for catches smaller than MSY that would allow the stock to rebuild towards MSY biomass levels. Although the stock assessment estimates of population parameters could be improved by the inclusion of survey, tagging, or size composition data, which are all available for BC sablefish, the current method allows at least a comparison of the information provided by the alternative fishery-dependent indices. In this context, if the stock assessment based on the spatial HBM is closer to the truth than the stock assessment based on the aspatial GLM index, then there is the potential for not achieving management objectives as a result of the dimensional and non-independence oversights of the aspatial fishery CPUE analysis.

2.3.3 Trend in the area occupied

The second objective of this analysis was to determine if the area occupied by the economically exploitable stock has declined over the fishery time series, and whether aggregative or disaggregative processes appear to be occurring for BC sablefish. Using an economically exploitable cutoff of six kilograms per trap, which is roughly equivalent to two sablefish, the area occupied by the BC sablefish stock is estimated to have decreased by about 62 percent from 1990 - 2005 (Figure 2.8). The area occupied by the economically exploitable stock was near the area $|\mathcal{D}|$ from the beginning of the time series in 1990 to 1999; most of the decline in the area occupied occurred late in the time series from 2000 to 2005. This contrasts with the nominal CPUE history of the fishery, which showed a rapid decline early in the time series followed by a levelling off in CPUE.

Although the area occupied by the economically exploitable stock has declined, trends in percentiles of the scdf and the heterogeneity index \mathcal{H} do not suggest aggregative processes in the BC sablefish stock. The percentiles of the scdf, which quantify the spatial distribution of CPUE over the domain \mathcal{D} , show that some relatively high CPUE areas existed early in the time series, but that since the mid-1990s, these high

59

Table 2.6: Comparison of indices of abundance using five estimated quantities. A Schaefer surplus production model was fit to each of four commercial fishery indices using a Metropolis-Hastings algorithm. Listed are the posterior mean, posterior median, and the 2.5 and 97.5 percentiles of the posterior for each of five quantities: (1) h_{MSY} is the theoretical optimum harvest rate; (2) MSY is the maximum sustainable yield (in tonnes); (3) K is the unfished biomass (in tonnes); (4) P_{final} is the stock depletion in 2005; (5) \bar{C}/MSY is the ratio of the average historical catch (1965-2005) to MSY.

	h_{MSY}	MSY	K	P_{final}	\bar{C}/MSY
Nominal CPUE					
mean	0.088	3,570	88,600	0.329	1.22
2.5%	0.039	2,590	52,500	0.211	0.98
50.0%	0.084	3,540	83,500	0.306	1.19
97.5%	0.155	4,290	153,000	0.651	1.62
Aspatial GLM					
mean	0.092	4,050	98,700	0.445	1.11
2.5%	0.038	2,750	$50,\!600$	0.264	0.51
50.0%	0.088	3,770	86,000	0.396	1.12
97.5%	0.164	8,290	242,000	0.868	1.52
Mean Spatial CPUE					
mean	0.080	$3,\!180$	$86,\!100$	0.148	1.36
2.5%	0.033	2,100	$54,\!100$	0.102	1.07
50.0%	0.077	$3,\!240$	83,900	0.144	1.30
97.5%	0.146	$3,\!950$	125,700	0.214	2.00
Spatial Median CPUE					
mean	0.074	$3,\!100$	89,500	0.147	1.39
2.5%	0.040	$2,\!350$	$56,\!800$	0.101	1.08
50.0%	0.068	$3,\!090$	90,100	0.144	1.37
97.5%	0.137	$3,\!900$	$117,\!600$	0.210	1.79



Figure 2.8: Area occupied by the economically exploitable BC sablefish stock (1990-2005). The area (km^2) with HBM-predicted CPUE greater than six kilograms (approximately two sablefish) per trap, estimated from the scdf. Error bars correspond to 80% credibility intervals.

CPUE areas have been depleted and the stock has been declining steadily throughout the spatial domain (Figure 2.9). If stock aggregation were occurring in BC sablefish, the heterogeneity index would be expected to increase as the average stock density declined, because fish would be vacating areas of relatively low density in favour of areas of higher density. The heterogeneity index \mathcal{H} estimates that the gap between the highest CPUE grid cells and the lowest CPUE grid cells narrowed rapidly from about eighteen kilograms per trap in 1994 to about six kilograms per trap in 1996, after which the gap has shown little trend. The trends in average stock density and the heterogeneity index suggest that the highest density areas have been fished down early in the time series, and that the current stock is being depleted more or less uniformly.



Figure 2.9: Time series of scdf and heterogeneity index. Time series of scdf percentiles for BC sablefish, smoothed to show trends, and time series of the heterogeneity index. Plotted in the left panel are the smoothed median (thick line), and the smoothed 10^{th} , 30^{th} , 70^{th} , and 90^{th} scdf percentiles (thin lines). Plotted in the right panel is the heterogeneity index $\mathcal{H}(0.1, 0.9)$, which is the difference between the 90^{th} and 10^{th} scdf percentiles. Error bars correspond to 80% credibility intervals.

2.4 Discussion

Case of BC sablefish

This study showed consequences for a BC sablefish stock assessment of using spatial indices of stock biomass rather than typically used aspatial indices. The estimated decline in the spatial indices resulted in lower estimates of MSY, h_{MSY} , and stock abundance when compared to estimates derived from the pattern of decline in aspatial fishery indices. Although all four assessments estimate historical catch to be in excess of MSY, differences in stock assessment have implications for management. Due to the tenuous assumptions of the aspatial indices, and given evidence for spatial heterogeneity in stock density above and beyond aspatial effects such as depth, more confidence ought to be placed on the spatial assessments. If managers followed aspatial management advice, when in fact the spatial assessments are accurate, they would

61

tend to allow overfishing of the stock due to the difference in the estimates of current biomass - the aspatial GLM assessment estimated the current stock biomass to be forty-five percent of the unfished level whereas the mean spatial HBM assessment estimated the current stock biomass to be only fifteen percent of unfished biomass.

Estimates of the area occupied by the economically exploitable stock showed a marked reduction, especially late in the period 1990 - 2005. Decrease in the area occupied is a management concern because internationally supported criteria for listing populations as threatened or endangered include trends in area occupied. The possibility that decreases in area occupied are coincident with aggregation are of special concern, especially given the prevalence of aspatial stock assessments. The BC sable-fish stock does not display evidence for aggregation; rather, the results of the spatial HBM suggest that current decreases in stock biomass are occurring evenly throughout the spatial domain. Even so, the relatively low spatial HBM-based biomass estimate for BC sablefish, combined with the sharp reduction in economically exploitable area, should be of concern for BC sablefish managers, whose primary mandate is conservation of the stock.

Challenges of using spatial fishery CPUE data

According to Hilborn and Walters (1992), the first rule of using fishery catch and effort data for estimating stock abundance is to spatially stratify the analysis. This rule is based on the assertion that fishermen are not random samplers of the stock density, thus the average rate of catch in a fishery is not representative of the average density of the stock. This advice is revisited in the *Folly and Fantasy* arguments of Walters (2003): the *Folly* is that nominal fishery CPUE represents average success of harvesters, not average density of the stock; the *Fantasy* is that stock density in fished areas is representative of density in unfished areas. In the case of GLM standardization of CPUE, the arguments of Walters (2003) are what I call the dimensional oversight.

As pointed out by Walters (2003), stratification of fishery CPUE into small areas or grid cells does not guarantee a proper spatial treatment. Once stratified, CPUE values must be generated for unfished cells so that changes in effort distribution do
not influence the average density index *per se.* Assuming that there are no fish when in fact there are fish in the unfished cells will introduce bias (Hilborn and Walters, 1992), so a better method is required. Walters (2003) chose to fill in missing cell values by extrapolating through time using a three-year average for backcasting and the last observed CPUE value for forecasting. Campbell (1998) presented an alternative way to fill in missing area values by assuming that there is no stock aggregation or disaggregation occurring, i.e. that the distribution of area CPUEs in a year, scaled so that the highest area CPUE is one, is constant across all years. In order to fill in areas without CPUE observations, these approaches make specific assumptions that are generally not strictly justified. The spatial cells are also assumed to be statistically independent even though in practise they will tend to be spatially autocorrelated (Nishida and Chen, 2004). In addition, neither of these approaches is truly spatial because they do not make use of location information; the areas or cells in these analyses could be permuted randomly, and the analysis would be unchanged.

The spatial HBM addresses the Folly and Fantasy arguments while incorporating spatial autocorrelation, and treats inferred stock density spatially so that quantities such as area occupied can be estimated directly. However, the spatial HBM analysis of fishery catch and effort data has its limitations, some of which are related to unbalanced sampling by the fishery. In the case of BC sablefish, unbalanced sampling may have been the reason that the HBM without skipper or vessel effects had the highest predictive power, i.e. skippers and vessels did not visit enough spatial cells to tell whether variations in CPUE were due to location or skipper/vessel effects. This lack of contrast due to targeting also meant that spatial effects and depth effects had to be estimated in blocks to ensure adequate spatial coverage and contrast. The use of grid-cell aggregate spatial effects also limits the resolution at which reliable predictions of CPUE can be made. For example, the spatial HBM with CAR prior provides reliable estimates of mean CPUE for 10-km² areas and larger, but prediction within a grid cell will not capture the effects of finer spatial features that are deviations from aspatial covariate effects. Finally, spatial effects were assumed to be temporally independent even though the spatial effects show some evidence of consistent patterns over time (figure 2.6).

The shortcomings of the spatial HBM analysis may have implications for stock assessment and subsequent management advice. Although the spatial HBM addresses unbalanced fishery CPUE data, areas that do not receive much fishing effort increase the amount of prediction uncertainty. These low-effort areas tend to occur at relatively deep and shallow depths because catch rates are lower at these depths, which increases prediction uncertainty at these depths in particular, and uncertainty in estimates of average spatial CPUE overall. Decreasing uncertainty in the average spatial CPUE should decrease uncertainty in stock assessment parameter estimates and improve management advice and performance. The assumption of temporal independence also contributed to uncertainty, especially in the low-effort areas. In similar fashion to the lack of contrast, temporal independence increases uncertainty in the spatial CPUE indices, and reduce certainty in stock assessment parameter estimates.

The reason that including spatial or temporal autocorrelation reduces uncertainty in parameter estimates is that information in the form of a prior is borrowed across grid cells or years. In the spatial HBM, the conditionally autoregressive prior constrains spatial variation according to the characteristics of the data, resulting in *shrinkage* of the spatial effects towards the prior mean (Greenland, 2000). Although this shrinkage has the effect of introducing bias, because the prior mean will be different from the within-cell mean, it can greatly increase the precision of estimates, especially for parameters that have very little data available to estimate them (e.g. spatial effects in grid cells with few or no fishing sets). In fact, shrinkage (or partial *pooling* of effects) typically results in superior parameter estimates in terms of mean squared error, when compared to maximum likelihood estimates (Carlin and Louis, 2000). However, care should be taken when making one-off spatial HBM-based predictions at individual locations due to the possibility that sizeable bias has been introduced by the prior.

Even though the spatial HBM approach that I have outlined addresses to the extent possible the non-independence oversight and the dimensional oversight, the practise of using fishery CPUE to index stock density will be forever suspect. Changes in the management regime of a fishery, for example, can have an impact on CPUE, as Sigler and Lunsford (2001) demonstrated for the Alaskan sablefish fishery. In the BC

sablefish fishery, the fleetwide introduction of escape rings in 1999 may have affected catch rates, but the effect is completely confounded with year-to-year changes in stock density. Whether or not escape rings had any effect on catch rates is unknown, but escape rings are designed to let sub-legal fish avoid capture. Therefore, the effect of escape rings on the retained weight in catch, which is the measure of catch used in this analysis, should be nil because fish escaping through rings would have had to be released had they been brought on-board.

Another possible effect of management on fishery CPUE was the use of a twoyear quota for 2002-2003. Most of the quota was taken in 2002, which resulted in targeted fishing effort during the winter in the northern region of BC. Winter catch rates tend to be high, especially in the north. As a result, the two-year quota had an apparent effect on both nominal CPUE and GLM year indices because of a particularly unbalanced dataset for 2003. The spike in catch rates was less pronounced in the spatial HBM indices presumably because the use of four-year spatial effects blocks resulted in borrowed information for southern areas in 2003. Unfortunately, the twoyear quota was coincident with very high catch rates not only in the fishery, but also in the survey (Wyeth et al., 2004). It is therefore possible that at least some portion of the CPUE spike in 2003 is due to recruitment to the stock.

High-grading is another issue with fishery data that can have insidious effects on the relationship between CPUE and density (Gillis et al., 1995). High-grading occurs when smaller fish are released to conserve room in the vessel's hold in anticipation of larger fish caught later in a trip. Using retained weight in catch in the presence of high-grading is suspect because fish that would otherwise be considered part of the stock are not counted in CPUE statistics (Gillis et al., 1995). It is unclear whether high-grading is an issue for BC sablefish; there is a price premium paid for larger fish, but the reported rate of discards is low (around three percent).

Issues of management and harvester incentives aside, there is still no guarantee that set-by-set CPUE will be proportional to local stock density (Hilborn and Walters, 1992). In fact, if catchability is not constant (equation 2.13) but rather random over time, then periods of low catchability will tend to be sampled more than periods of high catchability, resulting in a hyperstable relationship between CPUE and density (Cooke and Beddington, 1984). High catchability periods will be undersampled simply because high catchability results in large catches, which take more time to handle than small catches. In the case of the sablefish trap fishery, more time will be spent travelling to and from fishing grounds to offload larger catches during periods of high catchability. As a result, it may be prudent to allow hyperstability in the relationship between CPUE indices and stock biomass when fitting stock assessment models, even for spatial HBM-derived indices such as the mean spatial CPUE.

Future improvements and extensions

Improving estimates in areas with high prediction uncertainty should be the first step towards improving management advice for BC sablefish using the spatial HBM method. In order to improve contrast in depth, the fishery-independent scientific survey data could be incorporated into the analysis, either by assuming that survey and commercial sets are equivalent, by estimating an effect for the difference between survey and commercial CPUE, or by modeling the stock density - commercial CPUE and stock density - survey CPUE relationships explicitly using an additional hierarchical level. This approach could be coupled with improved survey design that targets areas where there are few commercial sets, to improve estimation over the entire area \mathcal{D} . For example, survey design in a given year could be a function of CPUE prediction uncertainty, which is tied to the pattern of fishing effort in the current or previous year. A prediction variance map such as the ones in Figure 2.10 could be used to prioritize survey effort, keeping in mind that survey measurements that overlap with commercial effort may be necessary for estimating the relationship between survey and commercial CPUE. Such an approach is often used to site air-quality monitoring stations (e.g. Zidek et al., 2000).

More complex interactions between space and time effects may also be in order. The current HBM treats autocorrelation in the spatial effects over four-year blocks of time. Presumably, the CPUE in a spatial cell is not only autocorrelated with neighbouring cells (spatial autocorrelation), but also with CPUE in the same spatial



Figure 2.10: Variance in spatial hierarchical Bayesian model predictions of log-CPUE for the BC sablefish commercial trap fishery, averaged by ten kilometre grid cell.

cell during neighbouring years (temporal autocorrelation). Modeling spatial autocorrelation allowed the spatial HBM to borrow information across cells that had few observations, and improve CPUE prediction where data were sparse. Adding temporal autocorrelation to the mix would allow information to also be borrowed across time, further improving prediction. This idea was presented in a stock assessment context by Walters (2003). Such models have been used in epidemiological studies with some success (e.g. Waller et al., 1997), and should be an improvement for stocks such as BC sablefish that show consistent spatial effects patterns over time.

Quantities other than CPUE could also be mapped, if available. In particular, mapping the average body size or age of sablefish would complement the information provided by CPUE. Size information is available from survey data. Spatial patterns of growth might be estimated either by assuming that fish are mostly sedentary, or estimating rates of movement from tag releases and recoveries. Such demographic information might be used to inform ecological models such as the *basin model* of MacCall (1990), which uses habitat suitability to form an integrated view capable of predicting patterns of abundance, productivity, movement, and distribution. In addition, the patterns in body size could provide evidence of patterns of recruitment, and would be very instructive for management of stocks such as BC sablefish, whose recruitment patterns and dynamics are still to some degree uncertain.

Finally, the stock assessment model could be improved. The use of the Schaefer biomass dynamics model was motivated by a desire not to muddle the comparison of the aspatial and spatial indices with complicated stock assessment. However, the Schaefer model does make strong assumptions about the relationship between stock biomass and production, namely that it follows a deterministic logistic growth curve with a one-year time lag, regardless of the composition of the stock or of prevailing environmental conditions. Although production in the form of individual growth and mortality occurs on a yearly timestep, recruitment to either the fishery or the spawning stock takes several years (Hanselman et al., 2007). Variation in spawning success from year to year is also evident. A more flexible stock assessment, incorporating ageing information, and allowing for random variation in production, would represent an improvement if the goal was to actually assess stock status rather than to simply compare alternative indices.

Making the most of spatial fishery CPUE data

Fisheries management is increasingly moving away from input-control systems such as effort limitation to output-control systems such as quotas. The success of these output-control systems depends strongly on the accuracy of stock assessments to provide information on both stock abundance and production (Walters and Pearse, 1996). Accurate stock assessment in turn requires a time series of data that is related to stock abundance. Given a time series of such data, a relationship to stock abundance must be either specified or estimated. Scientific research surveys allow for simple relationships between the data series and abundance because the survey can be designed specifically to measure stock density. However, many fisheries do not have enough research survey information to estimate abundance with a reasonable degree of certainty, and must instead use fishing records. Stock assessment scientists are left trying to estimate the relationship between fishing data and abundance, which is not a simple task because of noise in the data and insufficient contrast in abundance (Hilborn and Walters, 1992). The relationship may also be changing over time, making estimation even more difficult. The resulting estimates of abundance can thus mislead output-control management systems, as was the case in the northern cod fishery (Hutchings, 1996; Hutchings and Myers, 1994; Rose and Kulka, 1999; Walters and Maguire, 1996). Spatial information, which can be used to improve estimates, is typically not used.

Even in the earliest days of fishery stock assessment, Baranov (1918) acknowledged that fish stocks are not distributed uniformly over space (Paloheimo and Dickie, 1964). However, these early assessments were limited by the coarse spatial resolution of fishery data, and simple analyses that could be performed without the aid of modern computers. Today, many fisheries have time series of spatially-referenced catch and effort data that have been collected and catalogued at some expense. Despite the volume of spatial data that exists for many stocks, modern stock assessment largely ignores this spatial information during analysis. Modern stock assessments that do include spatial information do not use it to its fullest potential. Typically, "spatial" analysis is limited to calculating areal averages, for example by including *area* effects in a GLM analysis, but these effects do not actually use location information because areas are assumed to be independent. I have been unable to find an analysis that makes use of spatial information on the scale of a fishing-gear set to construct an index of stock abundance. The time is now to make the most of spatial fishery information that has been so expensively collected to improve fish stock assessment and address the shortcomings of aspatial assessments that were identified nearly a century ago.

Appendix A

MCMC procedure for spatial HBM

Posterior probability distributions for the HBMs were estimated using a Gibbs sampler. Three Gibbs MCMC chains were calculated, starting from different points in the parameter space. This was done to check the MCMC chains for convergence to the posterior distribution. In addition to examining trace plots of the MCMC chains, I monitored the values of the potential scale reduction factor to assess convergence. The potential scale reduction factor (Gelman and Rubin, 1992) calculates the within-chain variance and the between-chain variance in order to estimate the factor by which the scale of the distribution being sampled might be reduced if the chain were calculated indefinitely. Values of the potential scale reduction factor near one for all model parameters suggests convergence of the MCMC chains on the posterior. Once the chains had appeared to converge, I discarded the samples up to that point as *burn-in*. I took 500 additional draws from the posterior, which were also subsequently discarded, to calculate the within-chain correlation. Finally, 3000 samples were drawn to represent the joint posterior distribution. These final samples were thinned based on the previously calculated within-chain correlation, to obtain largely independent samples from the posterior.

Sampling distribution for CAR spatial effects

Because the CAR prior for the spatial parameters $w(\varsigma)$ is Normal and conjugate with the Normal likelihood of the transformed data f(y(s)), the conditional posterior distribution for the spatial parameters is Normal (Banerjee et al., 2004)

$$p\left(w(\varsigma_{k,r})|y(\boldsymbol{s}),\boldsymbol{\beta},w(\varsigma_{l\neq k,r}),\sigma_r^2,\tau^2\right) = N(\mu_{k,r}^*,\tau_{k,r}^{2*})$$
(A.1)

where

$$\mu_{k,r}^* = \frac{\tau^2}{\tau^2 + n_{k,r}\sigma_{w(r)}^2} \mu_{w(k,r)} + \frac{n_{k,r}\sigma_{w(r)}^2}{\tau^2 + n_{k,r}\sigma_{w(r)}^2} \bar{w}_{k,r}$$
(A.2)

and

$$\tau_{k,r}^{2*} = \frac{\tau^2 \sigma_{w(r)}^2}{\tau^2 + n_{k,r} \sigma_{w(r)}^2}$$
(A.3)

Here, $n_{k,r}$ is the number of sets in cell k, time block r. The prior mean $\mu_{w(k,r)}$ from the CAR prior definition (equation 1.19) is

$$\mu_{w(k,r)} = \frac{\sum_{l \neq k} b_{kl} w(\varsigma_{l,r})}{\sum_{l \neq k} b_{kl}}$$
(A.4)

and $\sigma_{w(r)}^2$ is the prior variance from the CAR prior definition (equation 1.19)

$$\sigma_{w(r)}^2 = \frac{\sigma_r^2}{\sum_{l \neq k} b_{kl}} \tag{A.5}$$

Finally, the mean $\bar{w}_{k,r}$ for the spatial effects in grid cell k, time block r is

$$\bar{w}_{k,r} = \frac{\sum_{i:s_i \in \varsigma_{k,r}} \left(f(y(s_i)) - X_i \boldsymbol{\beta} \right)}{n_{k,r}} \tag{A.6}$$

Sampling distribution for linear effects

The sampling distribution for model linear effects using the continuous spatial effects model is given in equations 1.15 and 1.16. This sampling distribution is slightly different in the case of the CAR spatial effects model because the spatial effects are sampled directly. Given the CAR spatial effects, the conditional posterior distribution for β is

$$p(\boldsymbol{\beta}|\tau^2, w(\boldsymbol{s}), y(\boldsymbol{s})) = N_p(A_{CAR}\boldsymbol{a}_{CAR}, A_{CAR})$$
(A.7)

where

$$A_{CAR}^{-1} = \frac{1}{\tau^2} X^T X + \Sigma_{\beta}^{-1}$$
(A.8)

$$\boldsymbol{a}_{CAR} = \frac{1}{\tau^2} X^T(f(\boldsymbol{y}(\boldsymbol{s})) - \boldsymbol{w}(\boldsymbol{s})) + \Sigma_{\boldsymbol{\beta}}^{-1} \mu_{\boldsymbol{\beta}}$$
(A.9)

and where μ_{β} and Σ_{β} are the mean and variance, respectively, of the normal prior distribution on β .

Sampling distributions for variance parameters

The conditional posterior distributions for the variance parameters τ^2 and σ^2 can be parametrized as a scaled inverse- χ^2 distribution (Gelman et al., 2004). Using this parametrization, I used a vague inverse- χ^2 prior (degrees of freedom $\nu_0 = 0.001$, scale $\tau_0^2 = 1$) for all variance parameters. Using this prior, the conditional posterior for the residual error variance τ^2 was

$$p(\tau^2|y(\boldsymbol{s}), w(\boldsymbol{s}), \boldsymbol{\beta}) = \text{Inv} - \chi^2 \left(n + 0.001, \frac{\sum_{i=1}^n \left(f(y(s_i)) - X_i \boldsymbol{\beta} - w(s_i) \right)^2 + 0.001}{n + 0.001} \right)$$
(A.10)

and the conditional posterior for the CAR spatial effects in time block r was

$$p(\sigma_r^2 | w(\boldsymbol{s})) = \text{Inv} - \chi^2 \left(n_r + 0.001, \frac{\sum_k w(\varsigma_{k,r})^2 + 0.001}{n_r + 0.001} \right)$$
(A.11)

References

- Akaike, H. (1973). Information theory as an extension of the maximum likelihood principle. In Petrov, B. and Csaksi, F., editors, *Proceedings of the 2nd International* Symposium on Information Theory, pages 267 – 281.
- Banerjee, S., Carlin, B., and Gelfand, A. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall / CRC Press, Boca Raton.
- Baranov, F. (1918). [On the question of the biological basis of fisheries]. Nauchnyi issledovatelskii iktiologicheskii Institut Izvestiia, 1(1):81 128.
- Beamish, R. and McFarlane, G. (1983). Summary of results of the Canadian sablefish tagging program. In Melteff, B., editor, *Proceedings of the international sablefish* symposium, volume AK-SG-83-8 of Lowell Wakefield Series, pages 59 – 80.
- Besag, J. (1974). Spatial interaction and statistical-analysis of lattice systems. Journal of the Royal Statistical Society Series B Methodological, 36(2):192 236.
- Besag, J., Green, P., Higdon, D., and Mengersen, K. (1995). Bayesian computation and stochastic-systems. *Statistical Science*, 10(1):3 41.
- Campbell, R. (1998). Catch and effort analysis with uncertain stock and effort dynamics: Southern bluefin tuna longline. In Funk, F., editor, Fishery stock assessment models : proceedings of the International Symposium on Fishery Stock Assessment Models for the 21st Century, volume AK-SG-98-01 of Alaska Sea Grant College Program, pages 75 – 95.
- Campbell, R. (2004). CPUE standardisation and the construction of indices of stock abundance in a spatially varying fishery using general linear models. *Fisheries Research*, 70(2-3):209 – 227.
- Carlin, B. and Louis, T. (2000). *Bayes and empirical Bayes methods for data analysis*. Chapman and Hall / CRC Press, Boca Raton.

- Cooke, J. and Beddington, J. (1984). The Relationship between Catch Rates and Abundance in Fisheries. IMA Journal of Mathematics Applied in Medicine & Biology, 1:391 – 405.
- Ellis, N. and Wang, Y. (2007). Effects of fish density distribution and effort distribution on catchability. *ICES Journal of Marine Science*, 64(1):178–191.
- Finley, A. O. and Banerjee, S. (2006). MBA: Multilevel B-spline Approximation. R package version 0.0-3.
- Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2004). *Bayesian Data Analysis*. Chapman and Hall / CRC Press.
- Gelman, A. and Rubin, D. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7:457 511.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6:721 – 741.
- Gillis, D., Peterman, R., and Pikitch, E. (1995). Implications of trip regulations for high grading - a model of the behavior of fishermen. *Canadian Journal of Fisheries* and Aquatic Sciences, 52(2):402–415.
- Greenland, S. (2000). Principles of multilevel modelling. International Journal Of Epidemiology, 29(1):158–167.
- Haist, V., Kronlund, A., and Wyeth, M. (2004). Sablefish (Anoplopoma fimbria) in British Columbia, canada: Stock assessment update for 2003 and advice to managers for 2004. *Can. Stock Assess. Sec. Res. Doc.*, 2004-055.
- Haist, V., Kronlund, A., and Wyeth, M. (2005). Sablefish (Anoplopoma fimbria) in British Columbia, canada: Stock assessment update for 2004 and advice to managers for 2005. *Can. Stock Assess. Sec. Res. Doc.*, 2005-031.
- Hanselman, D., Lunsford, C., Fujioka, J., and Rodgveller, C. (2007). Alaska sablefish assessment for 2008. In *Bering Sea, Aleutian Islands, and Gulf of Alaska SAFE*, pages 195 – 312. North Pacific Fishery Management Council, National Marine Fisheries Service.
- Harley, S., Myers, R., and Dunn, A. (2001). Is catch-per-unit-effort proportional to abundance? *Canadian Journal of Fisheries and Aquatic Sciences*, 58(9):1760–1772.

- Hastings, W. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97 109.
- Higdon, D., Lee, H., and Holloman, C. (2003). Markov chain Monte Carlo-based approaches for inference in computationally intensive inverse problems. In Bernardo, J., Bayarri, M., Berger, J., Dawid, A., Heckerman, D., Smith, A., and West, M., editors, *Bayesian Statistics* 7. Oxford University Press.
- Hilborn, R. and Walters, C. (1992). *Quantitative Fisheries Stock Assessment: Choice, Dynamics, and Uncertainty.* Kluwer Academic Publishers, Norwell, Massachusetts.
- Hutchings, J. (1996). Spatial and temporal variation in the density of northern cod and a review of hypotheses for the stock's collapse. *Canadian Journal of Fisheries* and Aquatic Sciences, 53(5):943 – 962.
- Hutchings, J. and Myers, R. (1994). What can be learned from the collapse of a renewable resource - Atlantic cod, (Gadus morhua), of Newfoundland and Labrador. *Canadian Journal of Fisheries and Aquatic Sciences*, 51(9):2126 – 2146.
- IUCN (2001). *IUCN Red List Categories and Criteria: Version 3.1.* IUCN Species Survival Commission, Gland, Switzerland and Cambridge, UK.
- Kimura, D., Shimada, A., and Shaw, F. (1998). Stock structure and movement of tagged sablefish, Anoplopoma fimbria, in offshore northeast Pacific waters and the effects of El Nino Southern Oscillation on migration and growth. *Fishery Bulletin*, 96(3):462 – 481.
- Kronlund, A., Haist, V., Wyeth, M., and Hilborn, R. (2003). Sablefish (Anoplopoma fimbria) in British Columbia, canada: Stock assessment for 2002 and advice to managers for 2003. *Can. Stock Assess. Sec. Res. Doc.*, 2003-071.
- Kronlund, A., Wyeth, M., and Hilborn, R. (2002). Review of survey, commercial fishery and tagging data for sablefish (Anoplopoma fimbria) in British Columbia (supplement to the November 2001 sablefish stock assessment). Can. Stock Assess. Sec. Res. Doc., 2002-74.
- Kulka, D., Pinhorn, A., Halliday, R., Pitcher, D., and Stansbury, D. (1996). Accounting for changes in spatial distribution of groundfish when estimating abundance from commercial fishing data. *Fisheries Research*, 28(3):321 – 342.
- MacCall, A. (1990). Dynamic Geography of Marine Fish Populations. University of Washington Press, Seattle.

- Maunder, M. and Punt, A. (2004). Standardizing catch and effort data: a review of recent approaches. *Fisheries Research*, 70(2-3):141 159.
- Maunder, M., Sibert, J., Fonteneau, A., Hampton, J., Kleiber, P., and Harley, S. (2006). Interpreting catch per unit effort data to assess the status of individual stocks and communities. *ICES Journal of Marine Science*, 63(8):1373 – 1385.
- McFarlane, G. and Beamish, R. (1983). Biology of adult sablefish (Anoplopoma fimbria). In Melteff, B., editor, *Proceedings of the international sablefish symposium*, volume AK-SG-83-8 of *Lowell Wakefield Series*, pages 59 – 80.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, E. (1953). Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087 – 1091.
- Nishida, T. and Chen, D. (2004). Incorporating spatial autocorrelation into the general linear model with an application to the yellowfin tuna (Thunnus albacares) longline CPUE data. *Fisheries Research*, 70(2-3):265 274.
- Paloheimo, J. and Dickie, L. (1964). Abundance and fishing success. Rapp. Proc. Verb. Reun. Cons. Int. Explor. Mer, 155:152 – 163.
- Perry, A., Low, P., Ellis, J., and Reynolds, J. (2005). Climate change and distribution shifts in marine fishes. *Science*, 308(5730):1912–1915.
- Peterman, R. and Steer, G. (1981). Relation between sport-fishing catchability coefficients and salmon abundance. *Transactions of the American Fisheries Society*, 110(5):585–593.
- R Development Core Team (2007). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- Ricker, W. (1940). Relation of 'catch per unit effort' to abundance and rate of exploitation. Journal of the Fisheries Research Board of Canada, 5:43 70.
- Rose, G. (2004). Reconciling overfishing and climate change with stock dynamics of Atlantic cod (Gadus morhua) over 500 years. *Canadian Journal of Fisheries and* Aquatic Sciences, 61(9):1553 – 1557.
- Rose, G. and Kulka, D. (1999). Hyperaggregation of fish and fisheries: how catchper-unit-effort increased as the northern cod (Gadus morhua) declined. *Canadian Journal of Fisheries and Aquatic Sciences*, 56:118 – 127.

- Schaefer, M. (1954). Some aspects of the dynamics of populations important to the management of commercial marine fisheries. Bulletin of the Inter-American Tropical Tuna Commission, (1):27 – 56.
- Schnute, J., Boers, N., Haigh, R., and others. (2006). PBSmapping: PBS Mapping 2. R package version 2.09.
- Shelton, P. (2005). Did over-reliance on commercial catch rate data precipitate the collapse of northern cod? *ICES Journal of Marine Science*, 62(6):1139 1149.
- Sigler, M. and Lunsford, C. (2001). Effects of individual quotas on catching efficiency and spawning potential in the Alaska sablefish fishery. *Canadian Journal of Fisheries and Aquatic Sciences*, 58(7):1300–1312.
- Spiegelhalter, D., Best, N., Carlin, B., and van der linde, A. (2002). Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society Series B -Statistical Methodology, 64:583 – 616.
- Su, Z., Peterman, R., and Haeseker, S. (2004). Spatial hierarchical Bayesian models for stock-recruitment analysis of pink salmon (Oncorhynchus gorbuscha). *Canadian Journal of Fisheries and Aquatic Sciences*, 61(12):2471 – 2486.
- Swain, D. and Sinclair, A. (1994). Fish distribution and catchability what is the appropriate measure of distribution? *Canadian Journal of Fisheries and Aquatic Sciences*, 51(5):1046–1054.
- Waller, L., Carlin, B., Xia, H., and Gelfand, A. (1997). Hierarchical spatiotemporal mapping of disease rates. *Journal of the American Statistical Association*, 92(438):607–617.
- Walters, C. (2003). Folly and fantasy in the analysis of spatial catch rate data. Canadian Journal of Fisheries and Aquatic Sciences, 60(12):1433 – 1436.
- Walters, C. and Maguire, J. (1996). Lessons for stock assessment from the northern cod collapse. *Reviews in Fish Biology and Fisheries*, 6(2):125 137.
- Walters, C. and Pearse, P. (1996). Stock information requirements for quota management systems in commercial fisheries. *Reviews in Fish Biology and Fisheries*, 6(1):21 – 42.
- Wyeth, M. and Kronlund, A. (2003). Sablefish (Anoplopoma fimbria) research and assessment surveys conducted in British Columbia waters from 1996 through 2000. *Can. Data. Rep. Fish. Aquat. Sci.*, 1116.

- Wyeth, M., Kronlund, A., and Elfert, M. (2004). Summary of the 2003 British Columbia sablefish (Anoplopoma fimbria) research and assessment survey. *Can. Data. Rep. Fish. Aquat. Sci.*, 1148.
- Zidek, J., Sun, W., and Le, N. (2000). Designing and integrating composite networks for monitoring multivariate Gaussian pollution fields. *Journal of the Royal Statistical Society Series C - Applied Statistics*, 49(Part 1):63–79.