Bayesian Dynamic Modeling of Stock-Recruitment Relationships

Renate Meyer and Russell B. Millar

Department of Statistics

The University of Auckland

Private Bag 92019

Auckland, New Zealand

Email: meyer@stat.auckland.ac.nz, millar@stat.auckland.ac.nz

ABSTRACT

Assessing the relationship between spawning stock size and the resulting number of adult offspring (recruitment) is one of the most fundamental issues in fisheries stock assessment and an important cornerstone for management decisions on harvest policies in many fisheries. This paper proposes a Bayesian state-space model for fitting stock-recruitment curves. This approach eliminates at least two of the four major problems encountered in traditional stock-recruitment analyses, that of "errors-in-variables bias" and "time-series bias". The state-space model takes the temporal dependencies of the observations into account through a conditional modeling of the observations, given unknown states, and specification of Markovian transition of states. Both process and observation errors are explicitly captured in the state-space model and quantified through posterior distributions of the parameters via the Bayesian paradigm. Beyond bias elimination, this approach is capable of quantifying fundamental uncertainties in parameter estimates and risks of management policies. Problems with posterior computations are overcome using Metropolis-Hastings-within-Gibbs sampling. This novel Bayesian state-space approach to stock-recruitment analysis is illustrated using a dataset on Fraser River pink salmon. The Ricker curve is employed to describe the dependence of recruitment on the spawning stock size. The state-space model is implemented using the readily available software package BUGS.

Keywords: Fisheries stock assessment, Ricker curve, state-space methodology, time-series bias, Markov Chain Monte Carlo, Gibbs sampler, BUGS.

1 Introduction

Estimating the functional form of the dependence of recruitment on the spawning stock size is one of the most fundamental and important issues in fisheries stock assessment (Hilborn and Walters 1992, chapter 7). To suppose constant recruitment no matter how hard the stock is being fished, is a dangerous assumption in fisheries management. Sad historical evidence thereof is given by the stock collapses of blue whales, California sardines, Peruvian anchoveta, and Atlantic cod (Hilborn and Walters 1992 and references therein). Stock-recruitment analysis provides the basis for setting precautionary reference points for spawning stock sizes and utilization rates to prevent recruitment overfishing and an ultimate collapse of the stock.

The analysis of stock-recruitment relationships traditionally consists of two phases: 1) the estimation of the parameters of the functional form relating the biomass/numbers of recruits to the biomass/numbers of spawners, and 2) the calculation of management parameters such as the optimal spawning stock size and harvest rate to maximize the average annual catch, coupled with an evaluation of their associated uncertainties and risks.

A stock-recruitment curve attempts to describe how average recruitment varies with parental stock size. Ricker (1975) postulated four properties that such a curve should have: passing through the origin, positive at all non-zero levels of stock, a non-increasing rate of recruitment (recruits-per-spawner) with increasing stock size, and for semelparous species (those who spawn only once, such as Pacific salmon) recruitment exceeding parental stock over some part of the range of possible stock sizes.

Common functional forms that relate recruitment (R) to spawning stock size (S) and have the above properties include:

the Ricker curve: $R(S) = \alpha S \exp(-\beta S)$ (Ricker 1954), $\alpha, \beta > 0$, where α is the recruits-per-spawner at low stock sizes, β describes how quickly the recruits-per-spawner drop as S increases.

the Beverton-Holt curve: $R(S) = \dfrac{\alpha S}{\beta + S}$, (Beverton and Holt 1957), $\alpha, \beta > 0$, where α is the maximum number of recruits produced, β is the spawning stock needed to produce an average recruitment of $\alpha/2$.

the Deriso curve: $R(X) = \alpha S(1 - \beta\gamma S)^{1/\gamma}$, (Deriso 1980), $\alpha, \beta > 0$, which includes the density-independent model ($R = \alpha S$), the Ricker, and the Beverton-Holt model as special or limiting cases.

Ignoring measurement error in the annual spawning stock size, maximum-likelihood estimates of the parameters are usually obtained by linear or nonlinear regression. For instance, when using the Ricker curve, dividing by S and taking the logarithm allows parameters to be estimated by standard linear regression.

Various pitfalls of such standard estimation procedures are very well known and studied (cf. Hilborn and Walters, 1992, Walters and Korman, 1998, and references therein), the most important of these being known in the fisheries literature as the "errors-in-variables bias", the "time-series bias", nonstationarity, and confounding spawning stock size effects with persistent environmental changes. This paper presents a Bayesian state-space approach that addresses the first two problems.

The "errors-in-variables bias", caused by ignoring measurement error in the independent variable, S, was noted by Walters and Ludwig (1981). They performed simulations of stock-recruitment relationships with errors in measurement of spawning stock size. These demonstrated that very small measurement errors result in a point cloud of spawners and recruits that shows no or very little of the underlying true stock-recruitment relationship. Subsequently, Ludwig and Walters (1981) developed a procedure for estimation of stock-recruitment parameters in the presence of measurement errors based on errors-in-variables models. A serious drawback of errors-in-variables regression, however, is that the large sample asymptotics needed for the estimation of parameter uncertainty does not apply to typical fisheries time series data of 10 to 20 points. Furthermore, to make parameters identifiable, the ratio of the error variances has to be known (Kendall and Stuart, 1973).

The potential problem of "time-series bias" was noted by Hilborn and Starr (1984) and first described by Walters (1985) for the general problem of estimating functional relationships from time-series data. As the cause of this bias, Walters (1985) identified correlation between deviations of the dependent variable and subsequent levels of the independent variable, when these levels are not fixed by an experimental design. When modeling natural populations, observation errors in the dependent variable typically influence subsequent values of the independent variable through the system dynamics. In the case of stock-recruitment relationships, the independent variable, the estimated spawning stock size, is not independent of the error around the mean recruitment relationship. Positive errors in observed recruitment lead to over-estimates of spawning stocks later in time, and negative errors to under-estimates, as will be formalized in the next chapter. For instance in the case of pink salmon, a semelparous species that mature at age 2, the population dynamics have been described by two relations between spawners and recruits (using the Ricker curve):

$$R_y = \alpha S_{y-2} \mathrm{e}^{-\beta S_{y-2} + u_y} \quad \text{and} \tag{1}$$

$$S_y = s_y R_y \mathrm{e}^{v_y} \tag{2}$$

where S_y is the number of recruits R_y that survive to spawn in year y, and u_y and v_y are independent, zero mean random variables. As pointed out by Walters (1985), a combination of these two equations results in a nonlinear autoregressive model. Standard regression estimates for α and β would be biased (Box and Jenkins, 1976). As shown by Walters (1985, 1990) time-series bias results in upward biases of the productivity parameter (α) and downward bias of the density dependence parameter (β) in the Ricker curve. These in turn result in overestimates of the optimum exploitation rate and underestimates of the optimum spawning stock size, and thereby overfishing. Factors affecting time-series bias, such as environmental cycles or trends, have been studied by Caputi (1988), Kope (1988), and Walters and Collie (1988).

A partial bias correction method for the Ricker stock-recruitment model has been developed by Walters (1990). Simulation studies have been carried out to quantify time-series bias (Walters 1981, Korman et al. 1995, Myers and Barrowman 1995). Myers and Barrowman (1995) found that for semelparous species, the time-series bias is important for all populations, and that for iteroparous species, large biases occur if the populations are exploited at close to the maximum that is biologically possible. Korman et al. (1995) used data from 30 sockeye salmon stocks and Monte Carlo simulations to examine the importance of time-series bias on estimates of optimal harvest rate, optimal escapement, and sustainable yield, comparing least squares procedures with the bias correction method. They found that the effect of time-series bias is greatest for low-productivity stocks that exhibit a high degree of autocorrelation among annual recruitment errors. The correction method reduced the bias in optimal escapement estimates under a limited set of conditions but at the price of increased variance.

Apart from partial bias correction formulas and Monte Carlo simulation to estimate the time-series bias, no statistical "remedy" has been found so far. Walters (1985) mentions one possible way to avoid this bias altogether by developing a complete model for the observations (R_y, S_y), $y = 1, \ldots, N$, and applying Bayes estimation procedures. But he deems these too complicated from a computational point of view. This was a valid comment at that point in time. However, with the immense progress made in Bayesian posterior computation through the development of Markov Chain Monte Carlo (MCMC) techniques, this approach is now viable and is the objective of this paper. Moreover, a Bayesian approach is most appropriate for the second part of the stock-recruitment analysis, namely the assessment of management parameters through quantification of uncertainties and risks.

This paper is organized as follows: In the second section we set up a state-space model for analysing stock-recruitment relationships. The third section describes the Bayesian approach to parameter estimation via a specific MCMC technique and the fourth section describes the estimation of management reference points and the assessment of their uncertainty within the Bayesian framework. The fifth section gives the results of fitting a stock-recruitment state-space model to data on Fraser River pink salmon. The paper concludes with a discussion and outlook.

2 State-space model

We will illustrate the Bayesian state-space approach using the stock-recruitment data of Fraser River pink salmon as a running example. Pink salmon (*Oncorhynchus gorbuscha*) of the Pacific Ocean have a 2-year life history almost without exception. In even-numbered years pinks are very scarce in the Fraser River and most adjacent streams of the lower Strait of Georgia and Puget Sound, but are abundant in odd-numbered years – a situation which has existed from the time of the region's first historical records during the 19th century. The data (Ricker 1989) presented in Table 1, consists of odd-year measurements of the numbers of spawners (escapement going upstream to spawn) and the number of pink salmon caught. The third column gives the total run or recruitment, i.e. the sum of spawners and fish caught. During the years of investigation, the Fraser River supported the largest odd-year pink salmon population in British Columbia. Between 1979 and 1987, odd-year pink salmon runs from the Fraser River averaged about 15 million fish (Ricker 1989). During odd-numbered years, Fraser River pink salmon can be the largest stock of Pacific salmon returning to the B.C. coast.

Stock-recruitment analysis is intrinsically uncertain. Two sources of uncertainty are the variability in annual population dynamics ("process errors") caused e.g. by spatial and temporal fluctuations in environmental conditions, and the variation in the observations due to measurement and sampling errors ("observation errors"). For any realistic stock assessment, these uncertainties have to be build into the stock assessment model. To this end, a state-space model is one of the most appropriate and promising tools. A state-space model relates the observations to unknown states through stochastic observation equations. The sequence of unknown states is allowed to follow a Markovian transition given by the stochastic state equations. Thereby, a state-space model can explicitly account for observation and process errors. A general introduction to nonlinear state-space models for fisheries stock assessment is given in Meyer and Millar (1999a).

As the actual measurements (estimated by independent sampling procedures) each year are the number of spawners (S_y) when they go upstream and the number of pink salmon caught (C_y) before their upstream migration, it seems more reasonable to treat their observation errors as independent rather than those of S_y and $R_y = S_y + C_y$ which are obviously correlated. Assuming a lognormal distribution for the estimated number of spawners and fish caught, the observation equations that relate the observations to unknown true states are given by:

$$S_y = S_y^{\text{true}} \mathrm{e}^{v_y^S} \tag{3}$$

$$C_y = C_y^{\text{true}} \mathrm{e}^{v_y^C} \tag{4}$$

where v_y^S and v_y^C are iid $N(0, \tau_S^2)$ and $N(0, \tau_C^2)$, respectively, for $y = 1, 3, \ldots, N$. Peterman (1981) examined a number of data sets and concluded that the observed pattern of variation conformed to a lognormal distribution. One feature of the lognormal distribution is that the amount of variation is proportional to the mean.

Using the symbols R_y for the true unknown recruitment and F_y for the true unknown instantaneous fishing mortality in year y,

$$S_y^{\text{true}} = \mathrm{e}^{-F_y} R_y \quad \text{and} \tag{5}$$

$$C_y^{\text{true}} = (1 - \mathrm{e}^{-F_y}) R_y. \tag{6}$$

We treat the R_y's and F_y's as the unknown states. For the Ricker curve, which is often used for Pacific salmon stocks, recruitment declines at higher stock sizes. The parameter α gives the recruits-per-spawner at low stock sizes, and β describes how quickly the recruits-per-spawner decrease as S increases. Using the Ricker curve and assuming that the fishing effort changes only gradually from year to year the state equations modeling the biannual population dynamics are given by:

$$R_y | R_{y-2}, F_{y-2}, \sigma_R^2 = \alpha \mathrm{e}^{-F_{y-2}} R_{y-2} \mathrm{e}^{-\beta \mathrm{e}^{-F_{y-2}} R_{y-2} + u_y^R}, \tag{7}$$

$$F_y | F_{y-2}, \sigma_F^2 = F_{y-2} \mathrm{e}^{u_y^F}. \tag{8}$$

Equation (7) is obtained from substituting (2) into (1) with $s_y = \mathrm{e}^{-F_y}$ and u_y^R and u_y^F are conditionally independent $N(0, \sigma_R^2)$ and $N(0, \sigma_F^2)$, respectively. (The state equations have to be supplemented by the distribution of initial states (R_{-1}, F_{-1}).) For a theoretical justification of the lognormal error assumption, based on assuming that the total survival from egg to recruit is a product of survival of individuals in various life history stages and the central limit theorem, see Hilborn and Walters (1992), p. 264, and Walters and Korman (1998).

Using equations (5) and (6), the observation equations (3) and (4) which specify the conditional distribution of the observations given the unknown states and variance parameters are rewritten as:

$$S_y|R_y, F_y, \tau_S^2 \quad = \quad \mathrm{e}^{-F_y} R_y \mathrm{e}^{v_y^S}, \tag{9}$$

$$C_y|R_y, F_y, \tau_C^2 \quad = \quad (1 - \mathrm{e}^{-F_y}) R_y \mathrm{e}^{v_y^C}. \tag{10}$$

3 Bayesian approach via MH within Gibbs

As explained in more detail in Meyer and Millar (1999a), the Bayesian approach to statistical inference in general consists of two steps:

i) Specification of a joint probability model, $p(\mathbf{x}, \boldsymbol{\theta})$, for all observables, $\mathbf{x}$, and unobservables, $\boldsymbol{\theta}$. (We use the notation $p(x)$ for the density function of a random variable x.) As $p(\mathbf{x}, \boldsymbol{\theta}) = \mathbf{p}(\boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta})$ this amounts to specifying a prior distribution $p(\boldsymbol{\theta})$ for the unobservables and a sampling distribution or likelihood $p(\mathbf{x}|\boldsymbol{\theta})$.

ii) Conditioning on the data to obtain the posterior distribution of the unobservables given the observations by applying Bayes theorem:

$$p(\boldsymbol{\theta}|\mathbf{x}) = \frac{p(\boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta})}{p(\mathbf{x})} = \frac{p(\boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta})}{\int p(\boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta}) d\boldsymbol{\theta}}.$$

In this specific stock-recruitment state-space model, the vector of observations is

$$\mathbf{x} = ((S_y, C_y), y = 1, 3, \ldots, N)$$

and the unobservables are the parameters of the state-space model and the unknown states:

$$\boldsymbol{\theta} = (\alpha, \beta, \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2, R_y, F_y, y = 1, 3, \ldots, N).$$

Making use of the conditional independence assumption in (9) and (10), the likelihood is given by a product of lognormal distributions:

$$\begin{aligned} p\left(S_1, C_1, \ldots, S_N, C_N \right. & \mid \left. \alpha, \beta, \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2, R_1, F_1, \ldots, R_N, F_N\right) \\ & = \textstyle\prod_{y=1,3}^{N} p(S_y|R_y, F_y, \tau_S^2) p(C_y|R_y, F_y, \tau_C^2). \end{aligned} \tag{11}$$

The joint prior distribution can be easily specified by repeatedly applying the definition of conditional probability densities and using conditional independence of R_y and F_y in (7) and (8), respectively, as:

$$\begin{aligned} p(\alpha, \beta, \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2, R_1, F_1, \ldots, R_N, F_N) &= p(\alpha, \beta, \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2) p(R_{-1}) p(F_{-1}) \\ & \prod_{y=1}^{N} p(R_y|R_{y-2}, F_{y-2}, \sigma_R^2) p(F_y|F_{y-2}, \sigma_F^2). \end{aligned} \tag{12}$$

Thus, the joint prior will be completely determined once we specify $p(\alpha, \beta, \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2)$, $p(R_{-1})$ and $p(F_{-1})$.

Recently, Punt and Hilborn (1997) have given some general guidelines for selecting prior distributions for parameters of stock assessment models. They advocate the elicitation of informative prior distribution from historical information, basic biological knowledge, and expert judgement. Along similar lines, Hilborn and Liermann (1998) recommend to "stand on the shoulders of giants" and make use of previous experience, such as information from stock assessments of related species, in deriving formal priors via a meta-analysis. Although we do not agree with Walters and Ludwig (1994) that noninformative priors should be the default choice when performing Bayesian stock assessments, we want to avoid the usual criticism of contaminating objective statistical inference with personal, subjective beliefs. Thus, we follow Walters and Korman (1998) in choosing noninformative priors for the model parameters in this particular application.

We assume a uniform distribution for the regression parameter β. As α acts like a scale parameter, the noninformative distribution is uniform for $\log(\alpha)$, or equivalently $f(\alpha) \propto 1/\alpha$. Similarly, we assume the standard noninformative density for a scale parameter, i.e. $p(\theta) \propto 1/\theta$, for the variances $\theta = \sigma_R^2, \sigma_F^2, \tau_S^2, \tau_C^2$. To avoid problems with improper distributions (and the software that we use for implementation actually requires specification of proper distributions), we use the standard approximation of this improper prior by a gamma(0.001,0.001) distribution for the inverse of θ (cf. Meyer and Millar (1999b)). We assume reasonably diffuse lognormal distributions for $p(R_{-1})$ and $p(F_{-1})$.

The main computational problem that has hindered the application of Bayesian inference in the past is the high-dimensional integration to derive the normalization constant of the posterior density. With the development of MCMC methods over the last decade, this has been overcome (Gilks et al. 1996). MCMC methods are simulation techniques that draw a sample from the posterior distribution. Instead of generating a sequence of independent samples from the posterior, a Markov chain is constructed such that its equilibrium distribution is the joint posterior. Thus, after running the Markov chain for a certain "burn-in" period, correlated samples from the limiting distribution are obtained. Using this sample, any characteristic of the posterior, e.g. its mean and variance, can be consistently estimated by its empirical counterparts, e.g. the sample average and the sample variance, respectively.

One of the most well-known MCMC methods is the Gibbs sampler, a special case of the Metropolis-Hastings (MH) algorithm (Gilks et al. 1996). It generates a sample from the joint posterior by sampling in turn from each of the univariate full conditional posterior distributions, as explained in detail in Meyer and Millar (1999a). We advocate the use of the Gibbs sampler in practical stock assessments since all-purpose, non-commercial software, BUGS (Spiegelhalter et al. 1996), is available that can handle the two main tasks necessary for implementation of the Gibbs sampler. These tasks are to i) construct and ii) to sample from the full conditional posterior densities. (These can have a non-standard form in general, so that special Monte Carlo techniques are required to sample from them.) BUGS can perform these tasks for a variety of complex models such as random effects, generalized linear, proportional hazards, latent variable, and frailty models. As shown by Meyer and Millar (1999b), state-space models are also amenable to a Bayesian analysis via BUGS. This has the major advantage that no one-off program in a low-level programming language such as C or FORTRAN has to be written for every new stock-recruitment analysis. Only the prior and sampling distributions for unobservables and observables, respectively, have to be specified in a BUGS program. The tedious task of constructing the full conditionals (Meyer and Millar, 1999a) is automated by BUGS using directed acyclic graphs (Meyer and Millar, 1999b). Sophisticated routines such as adaptive rejection sampling to sample from log-concave full conditionals and MH algorithms based on slice sampling to sample from non-logconcave full conditional densities have been implemented in BUGS and are continuously being refined. The implementation of a state-space model in BUGS is explained in depth in Meyer and Millar (1999b) using a Schaefer surplus production model as a running example. The BUGS code for the stock-recruitment state-space model here-in is given in the Appendix.

4 Management reference points and uncertainties

Various important management parameters can be derived from the productivity parameter, α, and the density dependence parameter, β. We use the formulae that were first presented by Ricker (1975) and also given in Table 7.2 of Hilborn and Walters (1992). These include the stock size, $S_{max} = \frac{1}{\beta}$, which produces maximum recruitment, $R_{max} = \frac{\alpha}{\beta}\mathrm{e}^{-1}$, the more important optimum stock size, S_{MSY}, that produces the so-called "maximum sustainable yield", MSY, and the optimum harvest rate for MSY, u_{MSY}. MSY is the largest average catch that can be captured from the stock under existing environmental conditions on a sustainable basis (in the absence of process error). S_{MSY} is the stock size at which the annual median catch $C = R - S = \alpha S\mathrm{e}^{-\beta S} - S$ is maximized. This point is usually obtained by iterative numerical routines, but Hilborn (1985) gave very useful approximate formulae that bypass the need for numerical optimization techniques:

$$\begin{aligned} S_{MSY} &\approx \frac{\log\alpha}{\beta}(0.5 - 0.07 * \log\alpha), \\ u_{MSY} &\approx 0.5 * \log\alpha - 0.07 * (\log\alpha)^2, \end{aligned}$$

$$MSY \approx \alpha S_{MSY}\, \mathrm{e}^{-\beta S_{MSY}} - S_{MSY}.$$

More important than simply calculating a point estimate of these management parameters is a realistic assessment of their uncertainties. The Bayesian approach in conjunction with the state-space model is particularly suited, as the state-space methodology accounts for both measurement and process errors and eliminates time series bias by explicitly incorporating the temporal structure of the observations through conditional Markovian modeling. Bayesian inference quantifies uncertainties and risks through giving a posterior probability distribution of all derived parameters. Using a sampling-based technique for posterior calculation such as the Gibbs sampler has the additional advantage that the posterior distribution of any parameter transformation $g(\alpha, \beta)$, such as S_{MSY} can be easily derived from the joint posterior distribution of α and β. If $\{(\alpha^{(m)}, \beta^{(m)}),\, m = 1, \ldots, k\}$ constitutes a sample from the posterior distribution of α and β then $\{g(\alpha^{(m)}, \beta^{(m)}),\, m = 1, \ldots, k\}$ is a sample from the posterior distribution of $g(\alpha, \beta)$. Using a large enough sample size k, any summary statistic of the posterior distribution can be computed to the required degree of accuracy.

5 Results

We used BUGS for implementation of the Gibbs sampler to sample from the joint posterior of the parameters and the states in the stock-recruitment state-space model applied to the data on Fraser River pink salmon. After a burn-in period of 10000, we performed 500,000 cycles of the Gibbs sampler and thinned the chain by taking every

100th sample yielding a final chain of length 5000.

Extensive convergence diagnostics for all unobservables were performed using the CODA software of Best et al. (1995). The Raftery and Lewis convergence diagnostics (estimating the 2.5% quantile of each parameter to an accuracy of 0.005 with probability 0.95) confirmed that the thinning, burn-in period, and sample size were sufficient. All chains passed the Heidelberger and Welch stationarity and halfwidth test. Lags and autocorrelation within each chain were reasonably low. For most parameters, Geweke's Z-scores did not fall within the extreme tails of a standard normal distribution suggesting that the chain fully converged.

Kernel density plots for the marginal posterior densities of the parameters α, β, and the management parameters $S_{max}, R_{max}, S_{MSY}, u_{MSY}$, and MSY are given in Figure 1. Summary statistics including the mean, standard deviation, and the 25%, 50% and 75% quantiles for the regression, variance, and management parameters are given in Table 2.

For this example, the ML estimates of the stock-recruitment parameters (obtained by standard linear regression of $\log(R/S) = \log(\alpha) - \beta S$, and thereby ignoring process error and measurement error in S),

$$\begin{aligned}\alpha_{ML} &= 5.034 \text{ with standard deviation of } 1.355,\\ \beta_{ML} &= 0.136 \text{ with standard deviation of } 0.1032,\end{aligned}$$

are very close to the Bayes estimators, i.e. the posterior medians,

$$\begin{aligned}\alpha_B &= 4.734 \text{ with } 95\% \text{ probability interval of } [3.12, 7.83],\\ \beta_B &= 0.135 \text{ with } 95\% \text{ probability interval of } [0.017, 0.30].\end{aligned}$$

The trajectories of the two sets of states, i.e. the posterior medians of $\{R_y\}$ and $\{F_y\}$ are given in Figure 2 and 3, respectively. These are overlaid by the corresponding observed quantities $S_y + C_y$ and $-\log(1 - C_y/(S_y + C_y))$, respectively.

Figure 4 shows a plot of the observed pairs of numbers of recruits and spawners (squares) and the posterior medians of the numbers of recruits and spawners (dots) as well as the average recruitment curves using the ML estimates, $\alpha_M L, \beta_M L$ and the Bayes estimates (posterior medians), α_b, β_B, for the parameters of the Ricker model.

6 Discussion and Outlook

We proposed a state-space approach to fitting stock-recruitment curves that eliminates at least two of the four major problems encountered in traditional stock-recruitment analysis, that of "errors-in-variables bias" and "time series bias". No fully satisfactory correction methods for these biases have been known before. This is accomplished by developing a complete probability model for the observations, i.e. the annual numbers of recruits and spawners, and proceeding with Bayesian inference. The state-space model takes the temporal dependencies of the observations into account through a conditional modeling of the observations, given unknown states, and specification of

Markovian transitions of states. Thus it avoids time-series bias which occured in standard estimation procedures from ignoring the time series structure of the observations. Furthermore, both process and measurement errors are explicitly captured in the state-space model and quantified through posterior distributions of the parameters via the Bayesian paradigm.

We agree with Walters and Korman (1998) that one of the major problems is a "fundamental destruction of information about the stock-recruitment relationship due to problems with the data" and that it is absolutely necessary to get more informative data via improved monitoring and experimental management programs. However, we do not think that this can be used as an excuse for not using more elaborate statistical models for parameter estimation. As demonstrated using the Fraser River pink salmon data, a state-space modeling approach is capable of eliminating errors-in-variables and time-series bias. This is preferable to the four possible remedies for "time-series bias" suggested by Walters and Korman (1998), namely (1) use of bias correction formulae like Walters (1990), (2) waiting for more data while avoiding policies like constant exploitation rate that tend to exaggerate the bias, (3) deliberate management experiments aimed at systematically obtaining more observations at extreme stock sizes, and (4) using habitat and other auxiliary data to provide an independent, conservative estimate of the β parameter. All these may incur considerable management risk.

Walters and Korman (1998) state that existing estimation procedures most often cannot provide the answers needed for management, such as good spawning stock reference points and an assessment of their uncertainty necessary to design robust and adaptive management policies. They point out the main difficulties with the frequentist approach via bootstrap sampling, especially with assessing time series bias in the parameters. On the other hand, the Bayesian state-space approach enables the fisheries scientist to measure the reliability of parameter estimates and admit realistic uncertainties in important management reference points. Apart from their arguments in favour of the Bayesian approach (easy implementation via readily available software, easier justification of Bayesian than frequentist measure of uncertainty from a management risk evaluation perspective) we want to emphasize the coherence of the axioms of Bayesian statistical inference. By giving the unknown parameters of the sampling distribution a probability structure, the Bayesian paradigm authorizes a quantitative discourse on these parameters. It is the only coherent approach that allows for conditioning on the observations and thus for an effective implementation of the unassailable "Likelihood Principle" (Robert, 1994), a principle which is widely accepted by both Bayesians and frequentists. Furthermore, the Bayesian paradigm enables the stock assessment scientist to make use of substantive knowledge, expert judgment, and information from inferences for related species and stocks through the construction of formal prior distributions.

Certainly, for a small dataset like the one on Fraser River pink salmon, there might not be a very noticable difference between the parameter estimates derived from standard estimation procedures such as linear regression and that of Bayesian estimates such as posterior medians derived from fitting a highly sophisticated state-space model using noninformative priors. However, using models that realistically account for measurement and process errors and avoid time-series bias will become more and more important when several datasets are combined in a meta-analysis as e.g. in Liermann and Hilborn (1997). Using faulty models for each individual dataset carries the danger of propagating the errors in the overall assessment. Hence, we see that Bayesian hierarchical state-space models provide an enormous potential for combining information from the collection of historical datasets.

The state-space approach has the further potential of incorporating nonstationarity in the stock-recruitment relationship, e.g. caused by a change in habitat conditions or human impacts. This can be easily achieved by allowing the slope and intercept parameters, just like the F_y's, to dynamically vary over time according to a first- or second-order random walk (West and Harrison, 1989). From a Bayesian point of view, this defines a so-called "smoothness prior" for $(\alpha_1, \beta_1, \ldots, \alpha_N, \beta_N)$. Such "dynamic" models are more flexible than fixed-parameter approaches.

The Bayesian approach using "Metropolis-Hastings-within-Gibbs" sampling as in BUGS offers the opportunity of robustifying the stock-recruitment state-space model by using non-Gaussian error distributions with longer than normal tails. Instead of the normal distribution for the logarithm of states, one could assume Student-t-distributions with low degrees of freedom, i.e. $u_y^R \sim t(0, \sigma_R^2, df)$ and $u_y^F \sim t(0, \sigma_F^2, df)$. Similarly, a heavy-tailed distribution might be more appropriate for the logarithm of observations to allow for crude measurement errors. Following the terminology of Fahrmeir and Kunstler (1998), we call these *robust* state-space models as they allow for additive and innovation outliers by choosing a heavy-tailed distribution for the observation and process errors, respectively. For models with heavy-tailed observation error distribution, the filters and smoothers are robust against additive outliers. State-space models that are robust against innovation outliers have proven useful for fitting time series with change points (Fahrmeir and Kunstler, 1998). In the fisheries context, these change points could be sudden shifts of the level of fishing mortality and/or changes of recruitment patterns over time. Resulting filters and smoothers are "edge-preserving", they provide smooth fits for periods with only small annual variation but do not blur edges or change points. Consequently, using robust state-space models has the potential of checking for environmental effects. The benefits of fitting such robust state-space models to fisheries data still have to be explored and this remains an important topic for further research.

References

Best, N.G., Cowles, M.K., and Vines, S.K. 1995. CODA manual version 0.30. MRC Biostatistics Unit, Cambridge, U.K.

Beverton, R.J.H., and Holt, S.J. 1957. On the dynamics of exploited fish populations. Fisheries Investment Series 2, Vol. 19, U.K. Ministry of Agriculture and Fisheries, London.

Box, D.E.P., and Jenkins, G.M. 1976. Time series analysis: forecasting and control. Holden-Day, San Francisco, 575 p. (see p. 208-284).

Caputi, N. 1988. Factors affecting the time series bias in stock-recruitment relationships and the interaction between time series and measurement error. Can. J. Fish. Aquat. Sci. 45:178-184.

Cheng, P., Levy, D.A., and Nealson, P.A. 1991. Hydroacoustic estimation of Fraser River pink salmon abundance and distribution at Mission, B.C., in 1987. Pacific Salmon Commission Technical Report No. 3.

Deriso, R.B. 1980. Harvesting strategies and parameter estimation for an age-structured model. Can. J. Fish. Aquat. Sci. 37:268-282.

Fahrmeir, L., and Kunstler, R. 1998. Penalized likelihood smoothing in robust state space models. preprint

Gilks, W.R., Richardson, S., and Spiegelhalter, D.J. 1996. Markov chain Monte Carlo in practice. Chapman and Hall, London, U.K.

Hilborn, R. 1985. Simplified calculation of optimum spawning stock size from Ricker's stock recruitment curve. Can. J. Fish. Aquat. Sci. 42: 1833-1834.

Hilborn, R., and Liermann, M. 1998: Standing on the shoulders of giants: learning from experience in fisheries. Rev. Fish. Biol. Fish. 8:1-11.

Hilborn, R., and Starr, P.J. 1984. Making stock-recruitment work. *In:* P.E.K. Symons and M. Waldichuk (Eds.) *Proceedings of the Workshop on Stream Indexing for Salmon Escapement Estimation.* Canadian Technical Report on Fisheries and Aquatic Science, No. 1326, pp. 227-244.

Hilborn, R., and Walters, C.J. 1992. Quantitative fisheries stock assessment: choice, dynamics and uncertainty. Chapman & Hall, New York.

Kendall, M.G., and Stuart, A. 1973. The advanced theory of statistics, Vol. II, Hafner Pub. Co., New York.

Kope, R.G. 1988. Effects of the initial conditions on Monte Carlo estimates of bias in estimating functional relationships. Can. J. Fish. Aquat. Sci. 45:185-187.

Korman, J., Peterman, R.M., and Walters, C.J. 1995. Empirical and theoretical analyses of correction of time-series bias in stock-recruitment relationships of sockeye salmon (*Onchorhunchus nerka*). Can. J. Fish. Aquat. Sci. 52:2174-2189.

Liermann, M., and Hilborn, R. 1997. Depensation in fish stocks: a hierarchic Bayesian meta-analysis. Can. J. Fish. Aquat. Sci. 54:1976-1984.

Ludwig, D., and Walters, C.J. 1981. Measurement errors and uncertainty in parameter estimates for stock and recruitment. Can. J. Fish. Aquat. Sci. 38:711-720.

Meyer, R., and Millar, R.B. 1999a. Bayesian stock assessment using a state-space implementation of the delay difference model. Can. J. Fish. Aquat. Sci. **56**:37-52.

Meyer, R., and Millar, R.B. 1999b. BUGS in Bayesian Stock Assessments.Can. J. Fish. Aquat. Sci. **56**: In press.

Myers, R.A.M., and Barrowman, N.J. 1995. Time series bias in the estimation of density-dependent mortality in stock-recruitment models. Can. J. Fish. Aquat. Sci. 52:223-232.

Punt, A.E., and Hilborn, R. 1997. Fisheries stock assessment and decision analysis: a Bayesian approach. Rev. Fish. Biol. Fish. 7:35-63.

Ricker, W.E. 1954. Stock and recruitment. J. Fish. Res. Board Can. 11:559-623.

Ricker, W.E. 1975. Computation and Interpretation of Biological Statistics of Fish Populations. Fisheries Research Board of Canada, Bulletin No. 191.

Ricker, W.E. 1989. History and Present State of the Odd-Year Pink Salmon Runs of the Fraser River Region. Can. Tech. Rep. Fish. Aquat. Sci. 1702, 37 p.

Robert, C.P. 1994. The Bayesian Choice. A Decision-Theoretic Motivation. Springer-Verlag, New York.

Spiegelhalter, D.J., Thomas, A., Best, N., and Gilks, W.R. 1996. BUGS 0.5, Bayesian inference using Gibbs sampling. Manual (version ii) Cambridge, UK: MRC Biostatistics Unit.

Walters, C.J. 1985. Bias in the estimation of functional relationships from time series data. Can. J. Fish. Aquat. Sci. 42:147-149.

Walters, C.J. 1990. A partial bias correction factor for stock-recruitment parameter estimation in the presence of autocorrelated environmental effects. Can. J. Fish. Aquat. Sci. 47:516-519.

Walters, C.J., and Ludwig, D. 1981. Effects of measurement error on the assessment of stock-recruitment relationships. Can. J. Fish. Aquat. Sci. 38:704-710.

Walters, C.J., and Ludwig, D. 1994. Calculation of Bayes posterior probability distribution for key population parameters. Can. J. Fish. Aquat. Sci. 51: 713-722.

Walters, C.J., and Collie, J.S. 1988. Is research on environmental factors useful for fisheries management? Can. J. Fish. Aquat. Sci. 45:1848-1854.

Walters, C.J., and Korman, J. 1998. Analysis of stock and recruitment data for deriving escapement reference levels: models, fitting, precision, diagnostics, pitfalls. Technical Report.

West, M., and Harrison, J. 1989. Bayesian forecasting and dynamic models. Springer-Verlag, New York.

Appendix

Table 1: Size of spawning stock, catch, and total run (in millions of fish) of Fraser River pink salmon from Ricker (1989)

Year	Spawners	Catch	Total Run	(in millions of fish)
1957	2.24	6.15	8.39	
1959	1.07	5.39	6.46	
1961	1.09	0.80	1.89	
1963	1.95	3.53	5.48	
1965	1.19	1.13	2.32	
1967	1.83	11.14	12.97	
1969	1.53	2.40	3.93	
1971	1.81	7.95	9.77	
1973	1.75	5.04	6.79	
1975	1.37	3.52	4.89	
1977	2.39	5.85	8.24	
1979	3.56	10.84	14.40	
1981	4.49	14.19	18.68	
1983	4.63	10.72	15.35	
1985	6.46	12.40	18.86	
1987	3.22	3.84	7.06	

Table 2: Summary for sample size of 5000 from posterior density from the stock-recruitment state-space model.

Parameter	Mean	SD	25%	Median	75%
α	4.92	1.24	4.06	4.73	5.53
β	0.141	0.0735	0.0883	0.135	0.187
σ_R^2	0.142	0.0920	0.0844	0.127	0.183
σ_F^2	0.00964	0.0139	0.00215	0.00512	0.0116
τ_S^2	0.0212	0.0309	0.0029	0.0093	0.0262
τ_C^2	0.428	0.207	0.29	0.384	0.509
R_{max}	18.6	20.7	10.5	13.0	18.2
S_{max}	11.9	16.5	5.34	7.41	11.3
S_{MSY}	6.72	8.31	3.42	4.47	6.53
u_{MSY}	0.607	0.0658	0.563	0.608	0.650
MSY	9.59	9.24	5.99	7.35	9.63

Figure 1 : The graphs show kernel estimates of posterior marginal densities obtained from MCMC samples using the Bayesian state-space approach to parameter estimation in the stock-recruitment model for: the productivity paramter α, the density dependence parameter β, and the management parameters S_{max}, R_{max}, S_{MSY}, u_{MSY}, and MSY.

Figure 2 : Trajectory of the posterior medians of the unobserved states R_y and the corresponding biannual observed quantities $S_y + C_y$.

Figure 3 : Trajectory of the posterior medians of the unobserved states F_y and the corresponding biannual observed quantities $-\log(1 - C_y/(S_y + Cy))$.

Figure 4 : Observed biannual numbers of recruits versus observed annual numbers of spawners (squares) and posterior medians of recruits versus spawners (dots) overlaid by the average recruitment curves using ML estimates (dashed line) and the Bayes estimates (solid line).

BUGS Code

```
model stock-recruitment;
const
     N=16;

var
     YEAR[N], Cobs[N], Cobsmean[N], Sobs[N], Sobsmean[N],
     R0, R[N], Rmean[N], F0, F[N], Fmean[N],
     alpha, ialpha, beta,
     sigmaR2, isigmaR2, sigmaF2, isigmaF2,
     tauC2, itauC2, tauS2, itauS2,
     Smax, Rmax, SMSY, uMSY, MSY;

data YEAR, Sobs, Cobs  in "fraserpink.dat";
inits in "fraserpink.in";

{
#  sampling distribution (likelihood):

for (i in 1:N) {
     Cobsmean[i] <- log(R[i]*(1-exp(-F[i])));
     Sobsmean[i] <- log(R[i]*exp(-F[i]));
     Cobs[i] ~ dlnorm(Cobsmean[i],itauC2);
     Sobs[i] ~ dlnorm(Sobsmean[i],itauS2);
     }

# prior distribution of states:

F0 ~ dlnorm(0.26,8.3)I(0.1,2);
R0  ~dlnorm(2.2,2.1)I(1,15);

Rmean[1] <- log(alpha*exp(-F0)*R0*exp(-beta*exp(-F0)*R0));
Fmean[1] <- log(F0);
for (i in 2:N) {
    Rmean[i] <- log(alpha*exp(-F[i-1])*R[i-1]*exp(-beta*exp(-F[i-1])*R[i-1]));
    Fmean[i] <- log(F[i-1]);
    R[i] ~ dlnorm(Rmean[i],isigmaR2)I(0.1,25);
    F[i] ~ dlnorm(Fmean[i],isigmaF2)I(0.1,3);
    }

# noninformative prior for alpha:
```

```
ialpha ~ dgamma(0.001,0.001);
alpha <- 1/ialpha;

# noninformative prior for beta:

beta ~ dunif(0,1);

# noninformative priors for variances:

isigmaR2 ~ dgamma(0.001,0.001);
isigmaF2 ~ dgamma(0.001,0.001);
itauC2   ~ dgamma(0.001,0.001);
itauS2   ~ dgamma(0.001,0.001);

sigmaR2 <- 1/isigmaR2;
sigmaF2 <- 1/isigmaF2;
tauC2   <- 1/itauC2;
tauS2   <- 1/itauS2;

# calculation of management parameters:

Smax <- 1/beta;
Rmax <- alpha/beta*exp(-1.0);
SMSY <- log(alpha)/beta*(0.5-0.07*log(alpha));
uMSY <- 0.5*log(alpha)-0.07*log(alpha)*log(alpha);
MSY  <- alpha*SMSY*exp(-beta*SMSY)-SMSY;
}
```