# Bayesian nonparametric analysis of stock–recruitment relationships

**Stephan B. Munch, Athanasios Kottas, and Marc Mangel**

**Abstract:** The relationship between current abundance and future recruitment to the stock is fundamental to managing fish populations. However, many different recruitment models are plausible and the data are insufficient to distinguish among them. Although nonparametric methods may be used to circumvent this problem, these are devoid of biological underpinnings. Here, we present a Bayesian nonparametric approach that allows straightforward incorporation of prior biological information and use it to estimate several fishery reference points. We applied this method to artificial data sets generated from a variety of parametric models and compare the results with the fit of Ricker and Beverton–Holt models. We found that the Bayesian nonparametric method fit the data nearly as well as the true parametric model and always performed better than incorrect parametric alternatives. The estimated reference points agree closely with true values calculated for the underlying parametric model. Finally, we apply the method to empirical data for lingcod (*Ophiodon elongatus*) and several salmonids. Since this method is capable of reproducing the behavior of any of the parametric models and provides flexible, data-driven estimates of stock–recruitment relationships, it should be of great value in fisheries applications where the true functional relationship is always unknown.

**Résumé :** La relation entre l'abondance actuelle et le recrutement futur du stock est d'importance capitale dans la gestion des populations de poissons. Plusieurs modèles de recrutement sont cependant plausibles et les données sont insuffisantes pour les distinguer. Bien que des méthodes non paramétriques puissent servir à résoudre le problème, celles-ci ne possèdent pas de fondement biologique. Nous présentons ici une méthode bayésienne non paramétrique qui permet une inclusion directe des renseignements biologiques a priori et nous l'utilisons pour estimer plusieurs points de référence halieutiques. Nous appliquons la méthode à des séries de données artificielles générées par divers modèles paramétriques et comparons les résultats à l'ajustement des modèles de Ricker et de Beverton–Holt. La méthode non paramétrique bayésienne s'ajuste aux données presque aussi bien que le véritable modèle paramétrique et elle fonctionne toujours mieux que les modèles de rechange paramétriques incorrects. Les points de référence estimés correspondent de près aux valeurs réelles calculées par le modèle paramétrique sous-jacent. Nous appliquons enfin la méthode à des données empiriques sur la morue-lingue (*Ophiodon elongatus*) et plusieurs salmonidés. Puisque la méthode peut reproduire le comportement de tous les modèles paramétriques et fournir des estimations des relations stock–recrutement flexibles et basées sur les données, elle peut s'avérer d'une grande utilité dans les applications aux pêches dans lesquelles les véritables relations fonctionnelles restent toujours inconnues.

[Traduit par la Rédaction]

## Introduction

The relationship between stock size and subsequent recruitment is one of the keystone concepts of fishery science, since the parameters for this function translate directly into management reference points and set the ultimate limits on sustainable fishing (Quinn and Deriso 1999; Bravington et al. 2000). However, it is also one of the most problematic: data are difficult to measure and generally noisy, the relationship is surely nonlinear over a range of stock sizes, and a variety of plausible biological mechanisms are consistent with very different functional relationships between spawning stock and recruitment. There is, of course, general agreement about the properties that a stock–recruitment model should possess. First, any extant stock must be able to replace itself so that recruitment should exceed losses owing to mortality over some range of stock sizes. Second, we expect there to be density dependence; recruitment may be nearly proportional to stock close to the origin, but per capita recruitment is expected to decrease at large stock sizes. Third, for closed populations, recruitment should tend to zero as stock goes to zero, although the spatial scale on

Received 1 December 2003. Accepted 13 September 2004. Published on the NRC Research Press Web site at http://cjfas.nrc.ca on 26 August 2005.
J17866

**S.B. Munch**[1,2] **and A. Kottas.** Center for Stock Assessment Research, Jack Baskin School of Engineering, University of California at Santa Cruz, Santa Cruz, CA 95064, USA.
**M. Mangel.** Department of Applied Mathematics and Statistics, Jack Baskin School of Engineering, University of California at Santa Cruz, Santa Cruz, CA 95064, USA.

[1]Corresponding author (e-mail: sbmunch@soe.ucsc.edu).

 doi: 10.1139/F05-073 

which the population is closed is often unclear. Many models may be derived from this set of principles and the data are frequently insufficient to distinguish among them. Consequently, selection of a parametric model is often rather arbitrary.

Over the last decade, parametric Bayesian approaches have been increasingly applied in fisheries (Hilborn and Mangel 1997; McAllister and Kirkwood 1998; Millar 2002). One of the great advantages of the Bayesian approach is that it allows a statement of model probability to be made in cases where there is no a priori biological basis for model selection. However, this approach will not work if the appropriate model is not included in the set of candidates. Moreover, even if the model is appropriately specified, important management parameters such as the slope at the origin may be unduly influenced by points far away (Ludwig 1995). Consequently, several authors have proposed nonparametric approaches. Rothschild and Mullen (1985) and Brodziak et al. (2002) divided the stock–recruitment plane into several regions and estimated transition probabilities among them from the observed time series. Nonparametric density estimators have also been used to construct the distribution of recruitment given stock biomass (e.g., Evans and Rice 1988; Cook 2000). Nonparametric regression and spline methods that fit some locally weighted smoothing function to the stock–recruit data have also been used (e.g., Chapman 1973; Cook 1998).

The chief benefit of these nonparametric approaches is that they allow the data to speak for themselves. This is a highly desirable property, once we recognize that the available biological information is typically insufficient to specify a functional form a priori. However, there are several drawbacks to existing methods. First, they all require the ad hoc specification of a smoothing parameter. Although cross-validation methods may be used to circumvent this problem to some extent, they do not perform well on the relatively small data sets available in fisheries. Second, uncertainty bounds for estimates from these methods rely heavily on asymptotics; given the relatively small samples available in fisheries, these uncertainty bounds will be unreliable. Third, these methods lack biological underpinnings. This makes results hard to interpret and sometimes biologically unreasonable. Bravington et al. (2000) have made some progress in this regard by developing a smoother that is forced to pass through the origin and produce diminishing gains in recruits as spawning stock size increases. Although elegantly incorporated, these biological constraints are by no means certainties and the inclusion of hard constraints is at odds with the nonparametric philosophy of letting the data speak for themselves. We still lack a method that allows prior information about the biology of the stock to be incorporated while simultaneously allowing the data to determine the overall shape of the fitted relationship. Here, we present a method, based on Bayesian nonparametrics (e.g., see Walker et al. 1999; Müller and Quintana 2004), that solves the difficulty. The approach provides a nonparametric fit to the data with several advantages over previous methods. First, the distributions of the parameters that determine the smoothness of the fitted nonparametric function are driven primarily by the data, given limited amounts of prior information. Second, full specification of model uncertainty is possible for the regression function and reference points derived from it without recourse to asymptotic approximations. Third, and most importantly, biological information is consistently incorporated through specification of the prior but does not rigidly constrain the shape of the model within the range of the data. Since this is the first application of Bayesian nonparametric methods to understanding the relationship between stock size and recruitment, we develop the statistical background for our Bayesian nonparametric (BNP) model in some detail. We demonstrate the flexibility of the method on artificial data sets generated from a suite of parametric models and then apply the method to five actual data series.

## Methods

### Model specification

For simplicity, we assume that spawning stock size ($S$) is measured without error and is proportional to total egg production. This assumption may be relaxed, although at the cost of increased complexity (see Discussion). We model recruitment ($R$) as a function of stock size and random processes, although other covariates may be incorporated where appropriate. We assume that stochastic effects are multiplicative and lognormally distributed. Alternative choices for the distribution of stochastic effects are readily incorporated in this framework. We also use log-transformed variables, specifically, $x = \ln S$ and $y = \ln R$. Thus, our stock–recruitment (SR) model becomes

$$y = \Im(x) + \varepsilon \tag{1}$$

where $\varepsilon$ is normally distributed with mean zero and variance $\sigma_\varepsilon^2$. In the standard context, $\Im(x)$ is the log-transformed SR model (e.g., $\Im(x) = \ln a + x - be^x$ for a Ricker model). In the BNP context, however, the function $\Im(x)$ is uncertain and described by a prior probability distribution over random functions. When thought of in this manner, fitting a parametric model, say Ricker or Beverton–Holt, corresponds to assigning positive prior probability to a very narrow region of the biologically plausible space of SR functions. We seek a broader prior specification, encompassing the space of biologically possible regression functions. One very general prior specification for $\Im(x)$ is the Gaussian process (GP). Other prior specifications are possible (Denison et al. 2002) but the GP prior is quite flexible, allows ready interpretation, and lends itself to relatively simple prior specification and posterior inference (e.g., see Neal (1999) and references therein for Bayesian GP regression).

As we are about to describe a GP model with five parameters governing the mean and covariance functions, some discussion of what makes this method nonparametric is in order. This method is referred to as a nonparametric Bayesian approach because we are seeking an SR relationship without specifying parameters that rigidly govern the functional form of the regression model. However, the model clearly makes specific distributional assumptions and in this regard differs from traditional nonparametric methods. Despite this, we retain the nonparametric Bayesian appellation in keeping with the rapidly developing statistical literature

on the subject rather than generate confusion by introducing new terminology.

The GP prior for $\Im(x)$ is fully described by a mean function $\mu(x)$ and covariance function $C(x,x')$, denoted as $\Im(x) \sim \mathrm{GP}(\mu(x), C(x,x'))$. Note that this specification for $\Im(x)$ is equivalent to saying that $\Im(x) = \mu(x) + z(x)$, where $\mu(x)$ can be any of the standard SR models on a log scale (or any other plausible choice for the mean function) and $z(x)$ is a zero mean GP with covariance function $C(x,x')$. Thus, this approach may be thought of as an analysis based on one of the standard SR models in which $C(x,x')$ provides the means by which we can account for systematic deviations from the hypothesized SR function. Consequently, it is mainly through specification of $\mu(x)$ that we may incorporate our prior beliefs about the relationship between stock size and recruitment. In keeping with a long tradition in fisheries biology, we assume that the population is closed, so that there should be no recruits in the absence of any adult stock (i.e., $R = 0$ when $S = 0$) and that recruitment is proportional to stock size near the origin. Because the actual mechanisms of density dependence are unknown for most stocks, we choose to let the data tell us. With this in mind, we choose as our prior mean function

$$\mu(x) = \beta_0 + \beta_1 x \tag{2}$$

which corresponds directly to Cushing's (1973) power function model. Obviously, other choices of the prior mean function are possible. For instance, if we had strong prior evidence that the stock in question exhibited Ricker-type density dependence, we could choose $\mu(x) = \ln(\alpha) + x - \beta e^x$.

Choosing the covariance function is important because it determines how closely related values of $\Im(x)$ are for nearby values of $x$. Thus, prior specification of the covariance function allows us to incorporate our biological beliefs regarding the smoothness of the underlying relationship between stock size and recruitment.

A generally useful choice for the covariance function is

$$C(x,x') = \tau^2 \exp(-\phi |x - x'|^\alpha), \quad \tau > 0,\ \phi > 0,\ \alpha \in (0,2] \tag{3}$$

Here, $\tau^2$ is the variance and parameters $\phi$ and $\alpha$ control the interdependence of $\Im(x)$ and $\Im(x')$ as a function of the distance between $x$ and $x'$. Specifically, $\alpha$ affects the fine-scale variability or graininess in $\Im(x)$, while $\phi$ (for a fixed $\alpha$) controls the large-scale variability or range of dependence. To build intuition, we illustrate the effects of each of these parameters on a realization of $\Im(x)$ (Fig. 1). Note a key assumption here: the covariance function is more properly $C(\Im(x), \Im(x'))$ but we assume that the interdependence of values in the stock recruitment relationship arise only through the distance between log-stock sizes (i.e., we assume that deviations from the mean function are isotropic). In the limiting case where the prior mean function accurately describes the behavior of the data, deviations from $\mu(x)$ possess no systematic behavior and $C(x,x')$ should tend toward a delta function.

Other choices of $C(x,x')$ are possible, although care must be taken to ensure that $C(x,x')$ is a proper covariance function (i.e., a symmetrical, nonnegative definite function). For example, a simple extension is

$$C(x,x') = \gamma_0 + \gamma_1 xx' + \tau^2 \exp(-\phi |x - x'|^\alpha), \quad \gamma_0 > 0,\ \gamma_1 > 0 \tag{4}$$

In this case, in the limit as $\tau$ goes to zero, the first two terms enforce a linear dependence of $\Im(x)$ (more specifically, deviations from the mean function) on $x$ where the variability in the slope is determined by $\gamma_1$.

To simplify notation, we collect the parameters specifying the prior for $\Im(x)$ in the vector $\boldsymbol{\theta} = (\beta_0, \beta_1, \tau^2, \phi, \alpha)$. Given the specifications in eqs. 1–3, the full Bayesian model becomes

$$\begin{aligned} y_i | \Im(x_i), \sigma_\varepsilon^2 &\sim N(\Im(x_i), \sigma_\varepsilon^2), \quad i = 1, \ldots, n \\ \Im(x) | \boldsymbol{\theta} &\sim \mathrm{GP}(\mu(x), C(x,x')) \\ \boldsymbol{\theta}, \sigma_\varepsilon^2 &\sim p(\boldsymbol{\theta})\, p(\sigma_\varepsilon^2) \end{aligned} \tag{5}$$

where $p(\boldsymbol{\theta})$ and $p(\sigma_\varepsilon^2)$ are independent priors for $\boldsymbol{\theta}$ and $\sigma_\varepsilon^2$. We discuss prior specification for these parameters below.

It is helpful to interpret the general model in the following way: given a specific set of observed recruitment and stock sizes, $(\boldsymbol{x}, \boldsymbol{y}) = ((x_1, \ldots, x_n), (y_1, \ldots, y_n))$, the GP prior for $\Im(x)$ means that $\{\Im(x_1), \ldots, \Im(x_n)\}$ is $n$-variate normal, i.e.,

$$(\Im(x_1), \ldots, \Im(x_n)) | \boldsymbol{\theta} \sim N_n(\boldsymbol{\mu}^n(\boldsymbol{\theta}), \mathbf{C}^n(\boldsymbol{\theta})) \tag{6}$$

**Fig. 1.** Effects of parameters of the covariance function on the sampled Gaussian process. Functions are sampled from a multivariate normal with the given covariance function on a grid of 50 points using the same random seed for each panel. The mean function is zero throughout. (*a*) The covariance function parameters are $\phi = 0.4$, $\alpha = 2$, and $\tau = 0.4$. The remaining panels show the results of individually halving each parameter singly while holding the others constant: (*b*) $\phi = 0.2$, (*c*) $\alpha = 1$, and (*d*) $\tau = 0.2$. Note that reducing $\phi$ (Fig. 1*b*) decreases the number of turning points in the sampled function, reducing $\alpha$ (Fig. 1*c*) sharpens the corners, and reducing $\tau$ (Fig. 1*d*) merely changes the scale.

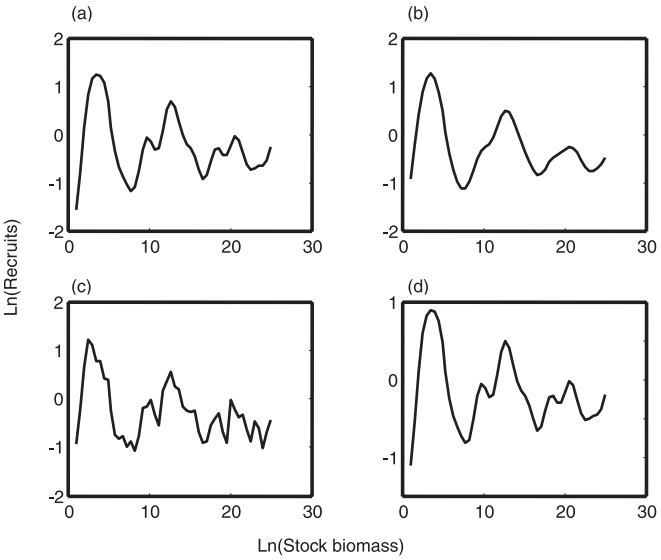

**Table 1.** Model formulations and parameters used to generate each of the simulated data sets.

| Model | Parameterization | $\alpha$ | $\beta_1$ | $\beta_2$ | $C$ | $d$ |
|---|---|---|---|---|---|---|
| Ricker | $R = \alpha S\, e^{-\beta_1 S}$ | 5.437 | 0.004 | na | na | na |
| Beverton–Holt | $R = \frac{\alpha S}{1 + \beta_2 S}$ | 5.000 | na | 0.010 | na | na |
| Shepherd | $R = \frac{\alpha S}{1 + (\beta_2 S)^C}$ | 6.000 | na | 0.0063 | 1.50 | na |
| Saila–Lorda | $R = \alpha S^C e^{-\beta_1 S}$ | 0.006 | 0.010 | na | 2.50 | na |
| Open-mixture | $R = \frac{\alpha S^C e^{-\beta_1 S}}{1000} + \frac{\alpha S}{1 + \beta_2 S} + d$ | 6.200 | 0.0200 | 0.015 | 2.65 | 20 |

**Note:** "na" under a parameter heading indicates that a given model does not include that parameter.

where the mean vector $\boldsymbol{\mu}^n(\boldsymbol{\theta}) = \{\beta_0 + \beta_1 x_1, \ldots, \beta_0 + \beta_1 x_n\}$ and the $(i,j)$th entry of the covariance matrix is $\mathbf{C}^n(\boldsymbol{\theta})_{i,j} = \tau^2 \exp(-\phi|x_i - x_j|^\alpha)$.

## Prior specification

In general, care must be taken with improper priors for the parameters $\theta$ and $\sigma_\varepsilon^2$, as these may lead to improper posteriors (e.g., see Berger et al. 2001). We choose proper priors that can be specified with small amounts of prior information. Specifically, we take

$$(7)\qquad \beta_0 \sim N(a_{\beta_0}, b_{\beta_0}), \quad \beta_1 \sim \Gamma(a_{\beta_1}, b_{\beta_1}),$$
$$\tau^2 \sim \mathrm{IG}(a_\tau, b_\tau), \quad \phi \sim \mathrm{IG}(a_\phi f, b_\phi), \quad \sigma_\varepsilon^2 \sim \mathrm{IG}(a_\sigma, b_\sigma)$$

where $\Gamma(a,b)$ denotes a gamma density with mean $a/b$ and $\mathrm{IG}(a,b)$ denotes an inverse gamma density with mean $b/(a-1)$ for $a > 1$.

From previous studies (Myers 2001), we expect that the slope at the origin should be between 2 and 7 and we therefore set $a_{\beta_0} = \ln(4.5)$. For the exponent in our Cushing prior, we expect recruits to be nearly proportional to stock size at low density and set $a_{\beta_1} = b_{\beta_1}$ so that $E(\beta_1) = 1$, allowing for the possibility of both compensation ($\beta_1 < 1$) and depensation ($\beta_1 > 1$) in the mean function. We also set $a_\tau = a_\phi = a_\sigma = 2$, resulting in infinite prior variances for $\tau^2$, $\phi$, and $\sigma_\varepsilon^2$. To fully specify the priors for $\beta_0$, $\beta_1$, $\tau$, and $\sigma_\varepsilon^2$, we note that the marginal response variance is $\mathrm{Var}(y_i) = b_\sigma + b_\tau + b_{\beta_0} + \mathrm{Var}(\beta_1)x_i^2$ (see eq. 9 below). We base our prior on a prior guess at the range of response variables $r_y$ and use $(r_y/4)^2$ as a guess at $\mathrm{Var}(y_i)$. We then inflate this guess by a factor of 40 in determining the parameter variances. Hence, using a prior guess $\tilde{x}$ as a rough center for the predictor, we divide the variance in $y_i$ evenly among the four components by setting $b_\sigma = b_\tau = b_{\beta_0} = 10(r_y/4)^2$ and $a_{\beta_1} = b_{\beta_1} = 1.6\tilde{x}\, r_y^{-2}$, so that $\mathrm{Var}(\beta_1) = 10(r_y/4)^2\tilde{x}^{-2}$. It is more difficult to imagine appropriate prior information for $\phi$ and $\alpha$. We note that $\phi$ controls the correlation among points and how rapidly the correlation dies away with distance. From the spatial statistics literature where GP priors are common (Cressie 1993), $3/\phi$ is called the range of dependence, i.e., the value of the distance $|x - x'|$ that gives $\mathrm{Corr}(\Im(x), \Im(x')) \approx 0.05$. We use one-half the range of observed $x_i$s, $r_x/2$, as a guess at the range of dependence. Thus, we set the prior mean of $\phi$, $b_\phi$, equal to $6/r_x$.

The remaining parameter for which prior specification is required is $\alpha$. For $C(x,x')$ to remain nonnegative definite, $\alpha$ must lie within the interval $(0, 2)$. As noted previously, $\alpha$ controls the fine-scale variability in $\Im(x)$. More specifically, $\alpha$ determines whether or not realizations of $\Im(x)$ are differentiable. Most commonly employed SR models are differentiable (the chief exception being the "hockey stick" (Barrowman and Myers 2000)), in keeping with the notion that abrupt changes in the biology of the recruitment process at a specific, single stock size are unlikely. From standard theory of random processes (Papoulis 1984), the derivative of $\Im(x)$ exists in a mean-square sense only if $\partial^2 C(x,x')/\partial x \partial x'$ exists. From eq. 3, this is only true if $\alpha = 2$. Thus, for all subsequent analyses, we set $\alpha = 2$ and the parameter vector $\boldsymbol{\theta}$ is reduced to $(\beta_0, \beta_1, \tau_2, \phi)$.

## Posterior inference for $\theta$ and $\sigma_\varepsilon^2$

Although it is possible to sample from the joint posterior distribution for $\Im(x)$, $\theta$, and $\sigma_\varepsilon^2$ given $\boldsymbol{x}, \boldsymbol{y}$ under the model specification, posterior sampling is greatly simplified by marginalizing over $\Im(x)$ to obtain a model with one less layer in the hierarchy. To do so, first note that, as in eq. 5, $\boldsymbol{y}$ is conditionally $n$-variate normal. The expected value of $\boldsymbol{y}$ taken over $\Im(x)$ is then $\boldsymbol{\mu}^n(\boldsymbol{\theta})$ (defined below eq. 6). From eq. 1, the covariances for $\boldsymbol{y}$ are given by

$$(8)\qquad \mathrm{Cov}[\Im(x_i) + \varepsilon_i), (\Im(x_j) + \varepsilon_j)] = \mathrm{Cov}[\Im(x_i), \Im(x_j)] + \mathrm{Cov}[\varepsilon_i, \varepsilon_j] = C(x_i, x_j) + \sigma_\varepsilon^2 I_n(i,j)$$

with $I_n(i,j) = 1$ for $i = j$ and 0 otherwise. Thus, a two-level model in which the explicit conditioning of $\boldsymbol{y}$ on $\Im(x)$ has been integrated out is

$$(9)\qquad \boldsymbol{y}|\boldsymbol{\theta}, \sigma_\varepsilon^2 \sim N_n(\boldsymbol{\mu}^n(\boldsymbol{\theta}), C^n(\boldsymbol{\theta}) + \sigma_\varepsilon^2 I_n)$$
$$\boldsymbol{\theta}, \sigma_\varepsilon^2 \sim p(\boldsymbol{\theta})\, p(\sigma_\varepsilon^2)$$

We use the Metropolis–Hastings algorithm (Hastings 1970) to obtain draws from the joint posterior of $(\boldsymbol{\theta}, \sigma_\varepsilon^2, p(\boldsymbol{\theta}), \sigma_\varepsilon^2|\,\text{data})$. Details on the specific implementation of the algorithm are given in Appendix A.

## Posterior prediction for $\Im(x)$

Based on the observed data $(\boldsymbol{x}, \boldsymbol{y})$, our goal is to predict values for $y$ and $\Im(x)$ over a range of values of $x$. That is, we estimate the SR function over a grid of $k$ points that we will

**Fig. 2.** Fits to simulated data: low-variance simulations ($V_p$ = 0.01). The figures illustrate the fit of parametric (Ricker and Beverton–Holt) and Bayesian nonparametric estimates of stock–recruitment relationships when the data are generated by a known model. True models are (*a*) Ricker, (*b*) Beverton–Holt, (*c*) Shepherd, (*d*) Saila–Lorda, and (*e*) open-mixture. In each panel, the true model is depicted by the blue line. Blue points are data sampled from the true model. Under the Bayesian nonparametric method, the solid black line gives the posterior mean and the broken lines are 95% pointwise uncertainty bands. Fitted Ricker and Beverton–Holt models are shown by the green and magenta lines. A LOESS smoother is indicated by the red line.

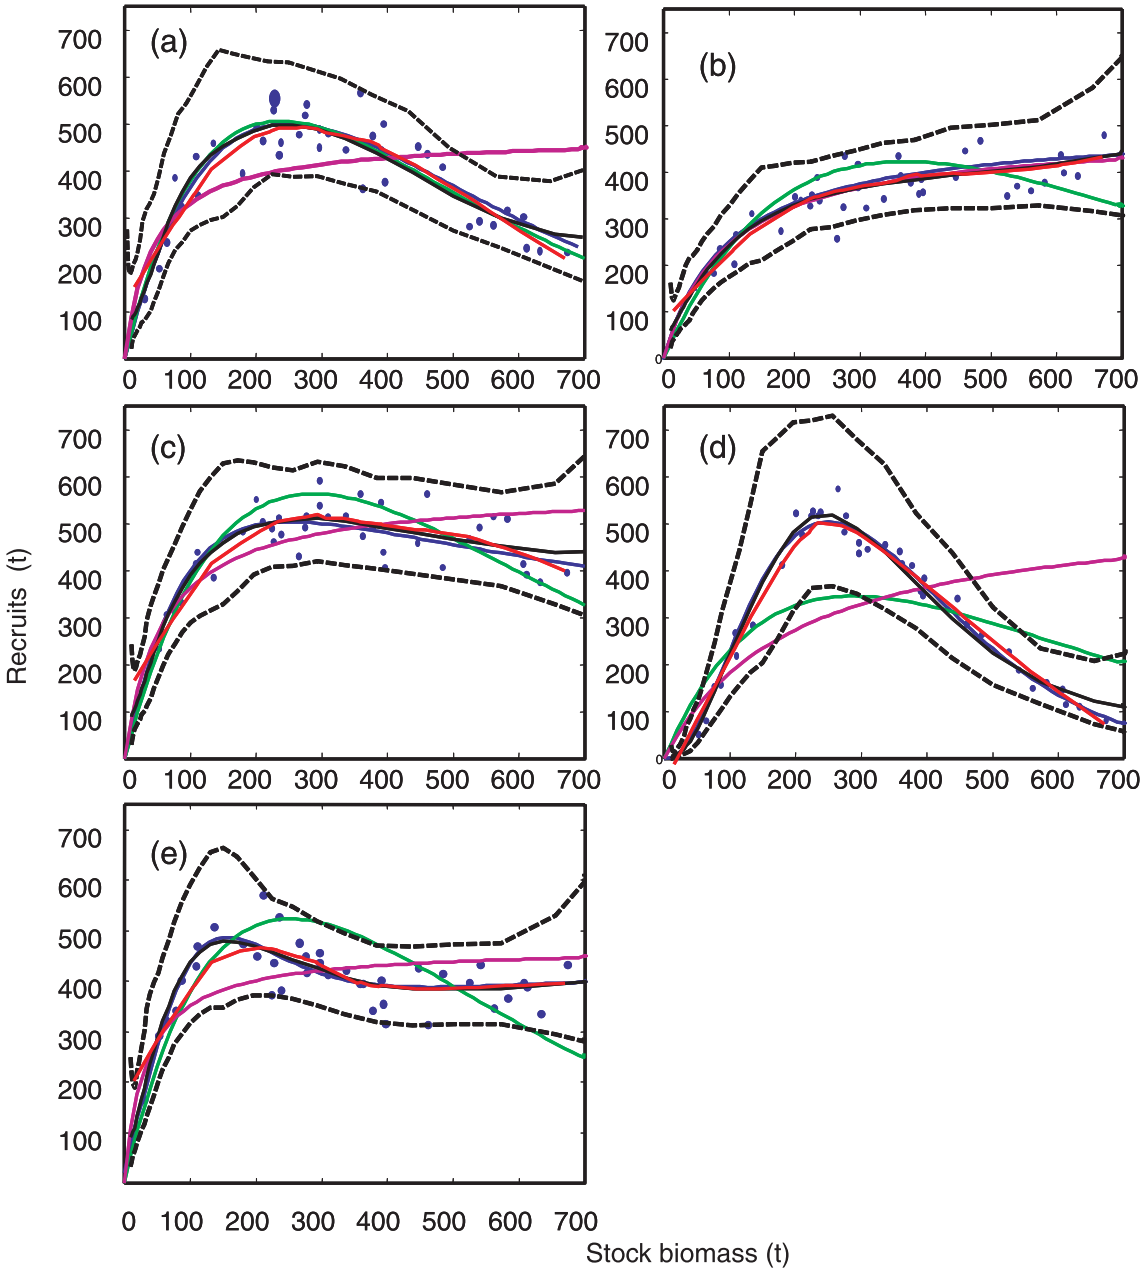


denote by $\boldsymbol{x}_{\text{new}} = (x_{\text{new }1},..., x_{\text{new }k})$. As derived in Appendix B, $\boldsymbol{y}_{\text{new}}$ and $\Im(\boldsymbol{x}_{\text{new}})$ are multivariate normal conditional on particular choices of $\boldsymbol{\theta}$, $\sigma_\varepsilon^2$. Specifically

$$(10)\quad p(\boldsymbol{y}_{\text{new}}|\boldsymbol{\theta}, \sigma_\varepsilon^2, \text{data}, \boldsymbol{x}_{\text{new}}) = N^k(M_{\text{new}}, S_{\text{new}})$$

$$(11)\quad p(\Im(\boldsymbol{x}_{\text{new}})|\boldsymbol{\theta}, \sigma_\varepsilon^2, \text{data}, \boldsymbol{x}_{\text{new}}) = N^k(M_{\text{new}}, S_{\text{new}} - \sigma_\varepsilon^2 I_k)$$

where $M_{\text{new}}$ and $S_{\text{new}}$ are the conditional mean vector and conditional covariance matrix described in Appendix B. Of course, we want the posterior distributions for $\boldsymbol{y}_{\text{new}}$ and $\Im(\boldsymbol{x}_{\text{new}})$ conditioned only on the data, not a particular choice of the parameter vector. Although marginalizing over $\boldsymbol{\theta}$ cannot be done analytically, the posterior distributions $p(\boldsymbol{y}_{\text{new}}|\text{data}, \boldsymbol{x}_{\text{new}})$ and $p(\Im(\boldsymbol{x}_{\text{new}})|\text{data}, \boldsymbol{x}_{\text{new}})$ may be obtained through Monte Carlo integration by sampling from the conditional densities 10 and 11 at each step of the Metropolis–Hastings algorithm. All calculations were implemented in Matlab v. 6.5 (The MathWorks, Inc., Natick, Massachusetts).

## Reference points

For any method of estimating the SR relationship to be of

**Fig. 3.** Fits to simulated data: high-variance simulations ($V_p$ = 0.1). The figures illustrate the fit of parametric (Ricker and Beverton–Holt) and Bayesian nonparametric estimates of stock–recruitment relationships when the data are generated by a known model. True models are (*a*) Ricker, (*b*) Beverton–Holt, (*c*) Shepherd, (*d*) Saila–Lorda, and (*e*) open-mixture. In each panel, the true model is depicted by the blue line. Blue points are data sampled from the true model. Under the Bayesian nonparametric method, the solid black line gives the posterior mean and the broken lines are 95% pointwise uncertainty bands. Fitted Ricker and Beverton–Holt models are shown by the green and magenta lines. A LOESS smoother is indicated by the red line.

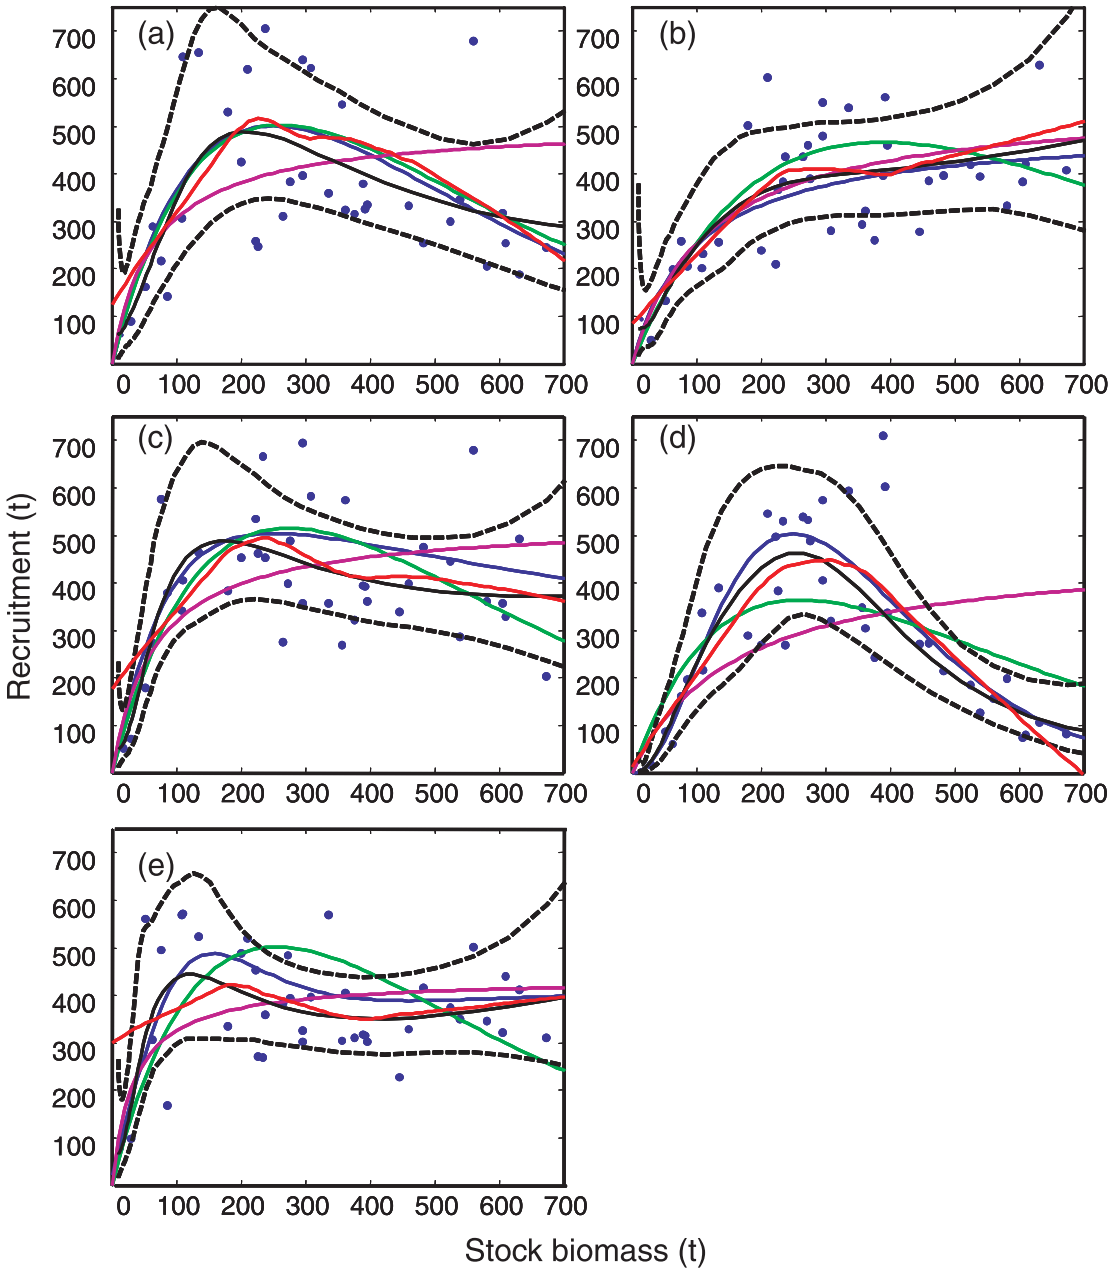


practical value, it is important that certain biological reference points may be calculated in a straightforward manner. Here, we outline how posterior predictive distributions for unfished biomass ($B_0$), the steepness ($h$) of the SR curve (Mace and Doonan 1988), and the stock biomass at maximum sustainable yield ($B_{MSY}$) can be obtained in our BNP framework.

Although more realistic models are certainly possible, we assume for simplicity that the stock dynamics on an annual time step are given by

$$S_{t+1} = S_t - (M + F)S_t + R(S_t) \quad (12)$$

where $M$ and $F$ are the fractions of the population removed by natural and fishing mortality over the course of a year and $R(S_t)$ is recruitment written explicitly as a function of stock at time $t$. Thus, at equilibrium in the absence of fishing, $B_0$ is the solution of $R(S) = MS$. Correspondingly, $x_0 = \ln(B_0)$ is the value of $x$ at which $\Im(x)$ crosses the line $y = x + \ln(M)$. For any posterior predictive realization of $\Im(\boldsymbol{x}_{new})$, there may be several such points, the number of which de-

**Table 2.** Goodness of fit to the simulated data sets with low and high variance.

| | | Residual sum of squares | | | AIC | | |
|---|---|---|---|---|---|---|---|
| Error variance | True model | Ricker | Beverton–Holt | BNP | Ricker | Beverton–Holt | BNP |
| Low | Ricker | 0.48 | 3.18 | 0.49 | –171.22 | –95.32 | –163.79 |
| | Beverton–Holt | 1.07 | 0.46 | 0.44 | –138.82 | –172.80 | –168.74 |
| | Shepherd | 0.65 | 0.79 | 0.33 | –159.00 | –151.06 | –180.48 |
| | Saila–Lorda | 13.35 | 19.62 | 0.40 | –37.89 | –22.50 | –172.48 |
| | Open-mixture | 1.69 | 1.24 | 0.36 | –120.53 | –133.11 | –176.53 |
| High | Ricker | 6.21 | 8.65 | 5.94 | –68.49 | –55.27 | –64.31 |
| | Beverton–Holt | 3.87 | 3.42 | 2.98 | –87.40 | –92.40 | –91.83 |
| | Shepherd | 3.54 | 4.83 | 2.75 | –90.95 | –78.54 | –95.09 |
| | Saila–Lorda | 16.35 | 22.90 | 3.32 | –29.78 | –16.32 | –87.53 |
| | Open-mixture | 5.27 | 4.27 | 2.82 | –75.07 | –83.52 | –94.10 |

**Note:** True models used to simulate the data are as described in the text and parameterized in Table 1. Residual sums of squares were calculated from the deviation between the observed data and the posterior mean model. AIC, Akaike's Information Criterion; BNP, Bayesian nonparametric model.

pends strongly on the amount of variability in the data. Here, we focus attention solely on those points that correspond to locally stable equilibria, i.e., where $\Im(x)$ is greater than $x + \ln(M)$ for $x$ less than $x_0$ and less than $x + \ln(M)$ for $x$ greater than $x_0$. Specific values of $x_0$ and $\Im(x_0)$ may be obtained by interpolating between the grid points of a realization from $\Im(\boldsymbol{x}_{\text{new}})$. For this purpose, we found that a grid of $k = 150$ points provided sufficient resolution. The steepness of the SR curve is defined as recruitment at 20% of $B_0$ expressed as a fraction of recruitment at $B_0$ (Mace and Doonan 1988). Therefore, after interpolating $\Im(x_0 + \ln(0.2))$ from the realization of $\Im(\boldsymbol{x}_{\text{new}})$, steepness can be estimated by calculating $\exp[\Im(x_0 + \ln(0.2)) - \Im(x_0)]$. Since yield in this model is given by $FS$, $B_{\text{MSY}}$ is simply the value of $S$ for which $R(S) - MS$ is maximized and can be found directly from $\exp[\Im(\boldsymbol{x}_{\text{new}})]$ and $\exp(\boldsymbol{x}_{\text{new}})$. Posterior predictive distributions for each of these reference points were estimated from the collection of values obtained from realizations of $\Im(\boldsymbol{x}_{\text{new}})$ calculated at each step of the Metropolis–Hastings algorithm.

## Testing the method

We demonstrate application of the method on two sets of data. The first are data simulated from a suite of five parametric models (Table 1): Ricker (Ricker 1954), Beverton–Holt (Beverton and Holt 1957), Shepherd (Shepherd 1982), Saila–Lorda (Iles 1994), and an atypical model for a potentially open population. Model formulations and parameters used in the simulations are provided in Table 1. The fifth parametric model, open-mixture, is a linear combination of the Beverton–Holt and Saila–Lorda models with an additional supply of recruits ($d$) from outside the specified stock area.

We chose parameters for all models so that the maximum and equilibrium recruitment levels were nearly the same. For each of the simulations, we sampled 40 stock sizes from a triangular distribution with support on [0,750] and mode 250. Given the chosen parameters, this sampling tends to emphasize points near the stock size of maximum recruitment. For each sampled stock size, observed recruitment was sampled as $R = f(S)e^{\bar{\omega}}$, where $\bar{\omega} \sim N(-1/2V_p, V_p)$. We conducted simulations using a small residual error $V_p = 0.01$ and a larger observation error $V_p = 0.1$. For each data set so constructed, we fit the BNP model with priors as described above using a burn-in period of greater than 10 000 draws (see Appendix A) followed by a posterior sampling period of another 10 000 draws.

**Fig. 4.** True $\sigma_\varepsilon^2$ and posterior expectations for data simulated from Ricker (open circles) and Beverton–Holt models (solid circles).

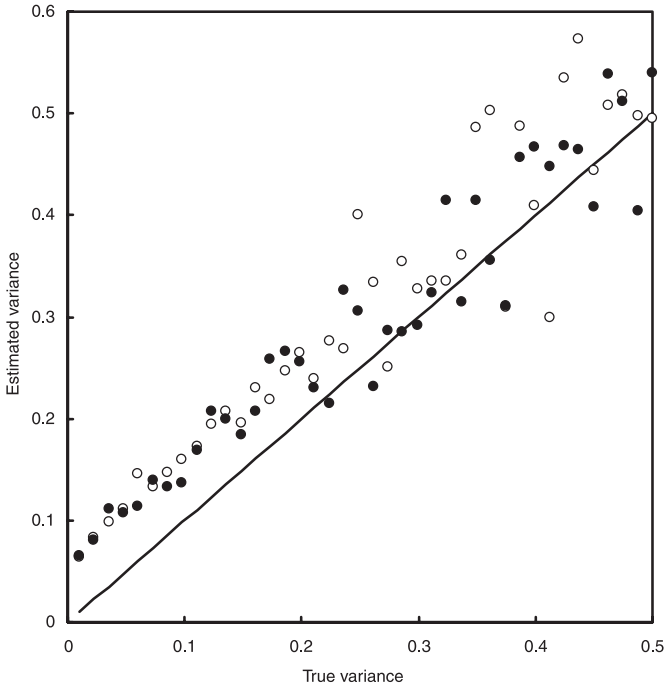


For quantitative measures of fit, we calculated the residual sum of squares ($SS = \Sigma(y - E[\Im(x)])^2$) and Akaike's Information Criterion (AIC) (Burnham and Anderson 1998). For comparison, we calculated sum of squares and AIC for Ricker and Beverton–Holt models fit to each data set as well using a Bayesian approach with uninformative priors for all of the parameters. We did not compare the fit of Shepherd and Sailia–Lorda models to our simulated data because our intention is merely to demonstrate the unifying flexibility of the nonparametric approach. We also fit a LOESS (Cleveland 1981) smoother to each. Since the stiffness of the

**Fig. 5.** Posterior predictive distributions of (*a* and *b*) $B_0$, (*c* and *d*) $B_{MSY}$, and (*e* and *f*) and steepness for Ricker model (left panels) and Beverton–Holt model (right panels). Solid triangles indicate the true value and open triangles indicate the posterior mean. Lines are kernel density estimates for the posteriors.

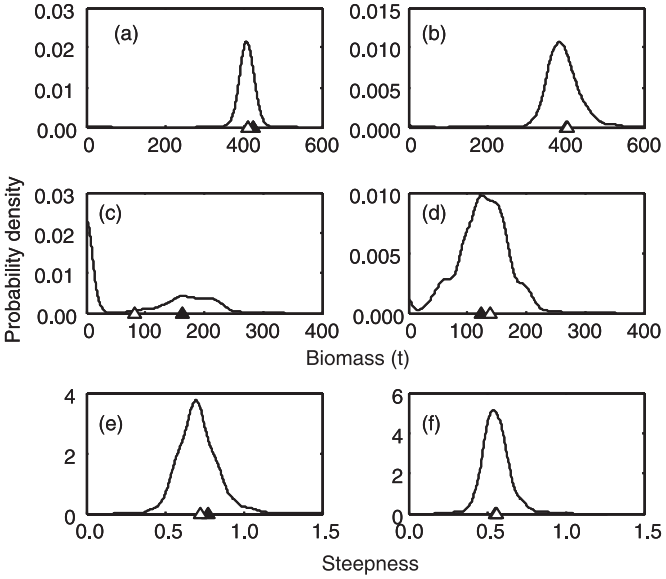


LOESS is arbitrary and the minimum sum of squares is zero, we do not calculate a measure of fit but include LOESS with stiffness set to 0.6 in the figures for visual comparison.

One of the chief parameters of interest that come directly from this model specification is the error variance $\sigma_\varepsilon^2$. To address how reliably our nonparametric approach could recover estimates of error variance, we simulated data from the Ricker and Beverton–Holt models in the manner described above for 100 values of $V_p$ ranging from 0.01 to 0.5. We also calculate posterior distributions for $B_0$, $h$, and $B_{MSY}$ based on 5000 realizations of $\Im(\boldsymbol{x}_{new})$ on a grid of 150 points.

We demonstrate the method on actual observations of stock and recruits for sockeye (*Oncorhynchus nerka*), chum (*Oncorhynchus keta*), and pink salmon (*Oncorhynchus gorbuscha*) from the Weaver Creek spawning channel, British Columbia (Essington et al. 2000), and northern and southern lingcod (*Ophiodon elongatus*) stocks from central California through Alaska (Jagielo et el. 2000). For each of these series, we fit the BNP model as well as the Ricker Beverton–Holt models and a LOESS smoother (stiffness = 0.6). To explore how much could be learned about the parameters of our GP model from typical fisheries data, posterior distributions were generated for the most and least informative of these data sets (sockeye salmon and southern lingcod, respectively).

## Results

The fit of the BNP model to each of the simulated data sets was quite good (Figs. 2 and 3). In each of the low-variance cases, the shape of the BNP model closely follows that of the simulated model over most of the range covered by the data. Departures from the true model are somewhat greater in the high-variance cases. Posterior predictive uncertainty bands for $\Im(x)$ obtained from the Monte Carlo samples tend to be fairly narrow in regions with sufficient data and wider in regions where data are scarce. The most substantial departures from the model occur when overcompensation is present (i.e., Ricker, Shepherd, and Saila–Lorda) primarily at the largest stock sizes. This deviation is a reflection of the choice of prior mean function and the scarcity of

**Fig. 6.** Stock–recruitment functions for three salmon species: (*a*) chum (*Oncorhynchus keta*), (*b*) pink (*Oncorhynchus gorbuscha*), and (*c*) sockeye (*Oncorhynchus nerka*). Blue points are observed data. Under the Bayesian nonparametric method, the solid black line gives the posterior mean and the broken lines denote 95% pointwise uncertainty bands. Fitted Ricker and Beverton–Holt models are shown by the green and magenta lines. A LOESS smoother is indicated by the red line. Data are from Essington et al. (2000).

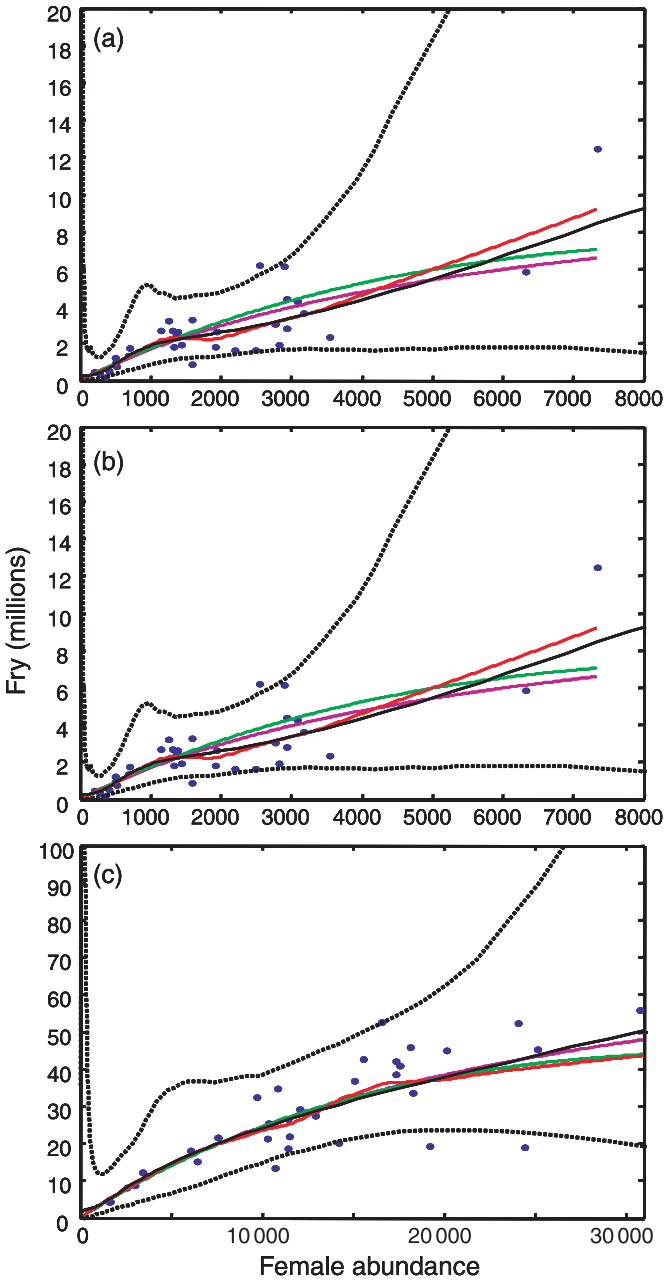

**Fig. 7.** Lingcod (*Ophiodon elongatus*) stock–recruitment models for (*a*) northern and (*b*) southern stocks. Blue points are observed data. Under the Bayesian nonparametric method, the solid black line gives the posterior mean and the broken lines denote 95% pointwise uncertainty bands. Fitted Ricker and Beverton–Holt models are shown by the green and magenta lines. A LOESS smoother is indicated by the red line. Data are from Jagielo et al. (2000).

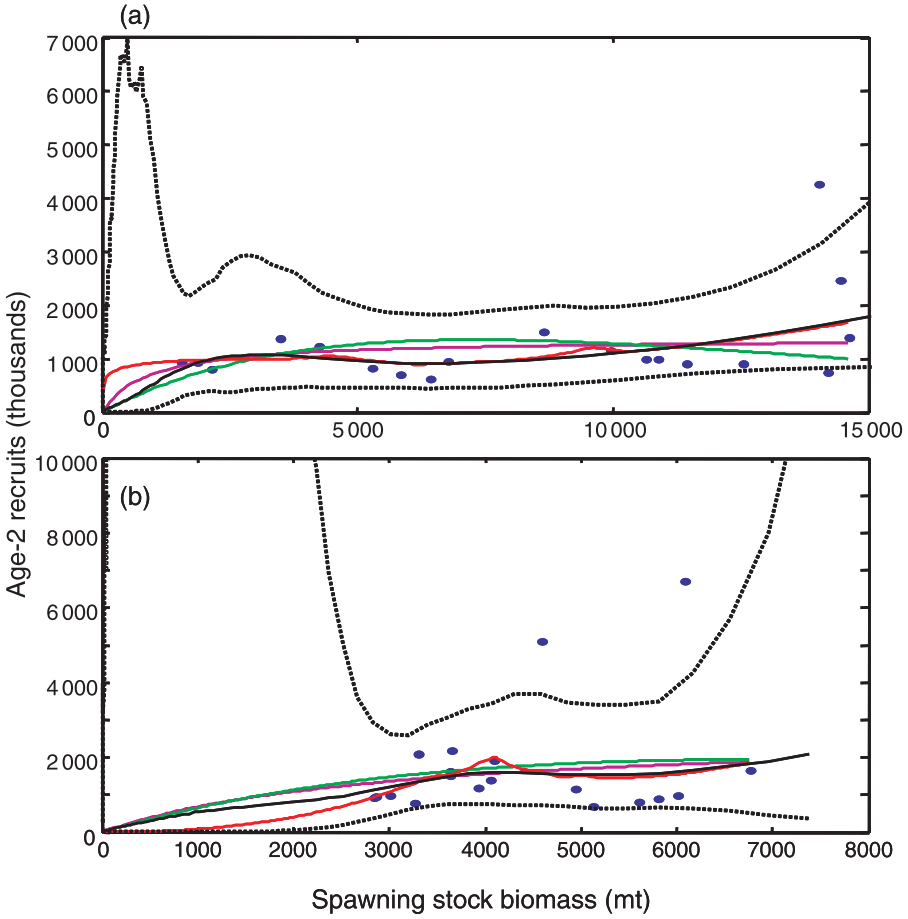


data at the largest stock sizes; had we assumed an overcompensating prior mean function, this deviation would not occur. Given that we have explicitly used the incorrect prior mean function for all of the simulated data sets, these results indicate that the BNP method is robust to misspecification of the mean and can be expected to work quite well when sufficient shape information is present in the data.

Comparison with the fit of the parametric models is generally quite favorable. The BNP fit (sum of squares or AIC) (Table 2) was usually close to the fit of the correct parametric model and was always better than the fit of the incorrect model. While it is not surprising that the Ricker and Beverton–Holt models were not adequate to describe data generated from more complicated models, we argue that the ecology underlying the recruitment process is likely to be more complicated than all of the standard models. The salient point is that, having demonstrated that the BNP approach is adequate to describe data from any of the parametric models, it provides a unifying approach that should be preferable when none of the standard models can be assumed to be correct. In real fisheries, the relationship between stock and recruitment is likely to involve several different sources of density dependence and, potentially, sources of recruits from other stocks. The open-mixture model is a simple example. The BNP method was the only one to accurately describe data from this model, highlighting its utility in situations where the underlying dynamics are complex and unknown.

A common concern when fitting nonparametric models to noisy data is that the resulting model may be over-fit. However, estimates of the error variance with the BNP model were quite good (Fig. 4). The expected value of the posterior $\sigma_\varepsilon^2$ was, on average, within 15% of the true value for $\sigma_\varepsilon^2 > 0.2$. For $\sigma_\varepsilon^2 < 0.2$, the estimated variance was substantially greater than the true value, although confidence regions from the posterior $\sigma_\varepsilon^e$ typically contained the true value for $\sigma_\varepsilon^2$ as small as 0.1. Overall, the estimated variance was more accurate for data generated using the Beverton–Holt model than for the Ricker model. This, again, is a consequence of our choice of the prior mean function. Overcompensatory choices of prior mean function would produce substantially better estimates of the variance under the Ricker model. In contrast with smoother-based methods, overfitting rarely occurred. Although the error variance estimated using a parametric model is strongly influenced by proper model choice, the BNP method allows reasonably accurate estimation of the error variance in the absence of model specification. Since recruitment variability is probably the single most important variable driving uncertainty in future stock sizes, this

**Fig. 8.** Prior and posterior probability densities for the parameters of the Bayesian nonparametric model fit to the sockeye salmon (*Oncorhynchus nerka*) data. The broken lines show the prior distributions for each parameter. The solid lines are kernel density estimates for the posteriors. (*a*) $\beta_0$; (*b*) $\beta_1$; (*c*) $\phi$; (*d*) $\tau$; (*e*) $\sigma^2_\varepsilon$.

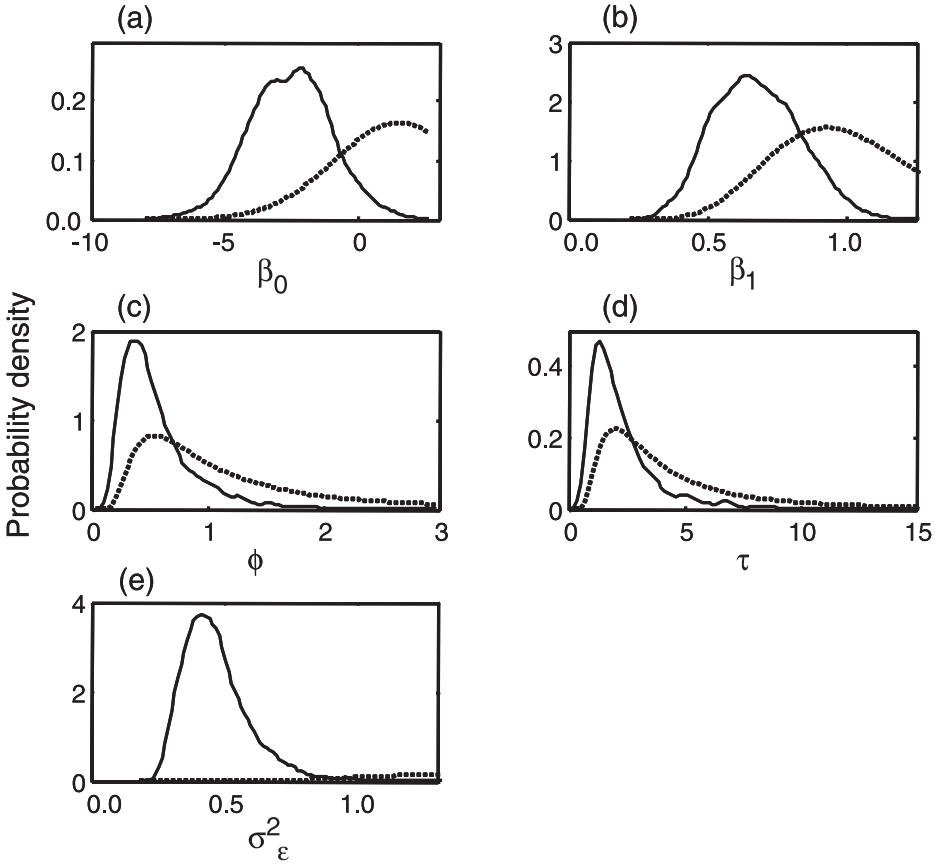


feature of the BNP approach will be of great utility in risk assessments.

### Reference points

The ability of the BNP method to estimate reference points was generally good (Fig. 5). Posterior means for $B_0$ and steepness were within a few percent of the true values for data sampled from both the Ricker and Beverton–Holt models. For Beverton–Holt data, the posterior mean estimated $B_{MSY}$ to within <1% of the true value. When there are little or no data near the origin, estimates of $\Im(\boldsymbol{x}_{new})$ from the Ricker model tend to show large variability close to zero. This is a consequence of the isotropic specification for $C(x,x')$ and the overcompensatory departures from the prior mean at large stock sizes. As a result, the posterior mean for $B_{MSY}$ for Ricker data was substantially biased owing to a large number of samples with solutions near the origin.

### Fit to actual data

For each of the salmon stocks (Fig. 6), the Ricker, Beverton–Holt, and BNP models all indicate recruitment increasing continuously with stock size. In the case of both chum and pink salmon, the BNP fit increases from the origin, flattens slightly at stock sizes around 25% of the observed maximum, and then increases again. Although it is impossible to say for certain, this pattern of recruitment may indicate alternative regimes, each characterized by different relationships. In contrast, the fit to the lingcod data (Fig. 7) is essentially flat for both northern and southern stocks. Note that the BNP extrapolation to the origin appears considerably more reasonable than that obtained with the LOESS smoother and more closely approximates the extrapolation obtained using the parametric models. However, uncertainty bands near the origin widen dramatically, indicating that there is no SR information in this region and that the extrapolations are based primarily on the prior mean function. In contrast, extrapolations from the LOESS are simply regression lines based on the 60% of the data closest to the origin.

Each of the posterior distributions for the parameters of the BNP model fit to the sockeye salmon data (Fig. 8) showed substantial reductions in variance and shifts of the mean. This indicates that when the data are reasonably informative, it is possible to learn about all of the parameters of the BNP model. Similarly, posterior distributions based on the southern lingcod stock (Fig. 9) indicate learning for the mean function and residual variance. However, the posterior distributions for the parameters of the covariance function showed little change from the priors.

## Discussion

The BNP approach that we have described offers a number of advantages over previous methods of defining the relationship between stock size and subsequent recruitment.

**Fig. 9.** Prior and posterior probability densities for the parameters of the Bayesian nonparametric model fit to the southern lingcod (*Ophiodon elongatus*) data. The broken lines show the prior distributions for each parameter. The solid lines are kernel density estimates for the posteriors. (*a*) $\beta_0$; (*b*) $\beta_1$; (*c*) $\phi$; (*d*) $\tau$; (*e*) $\sigma^2_\varepsilon$.

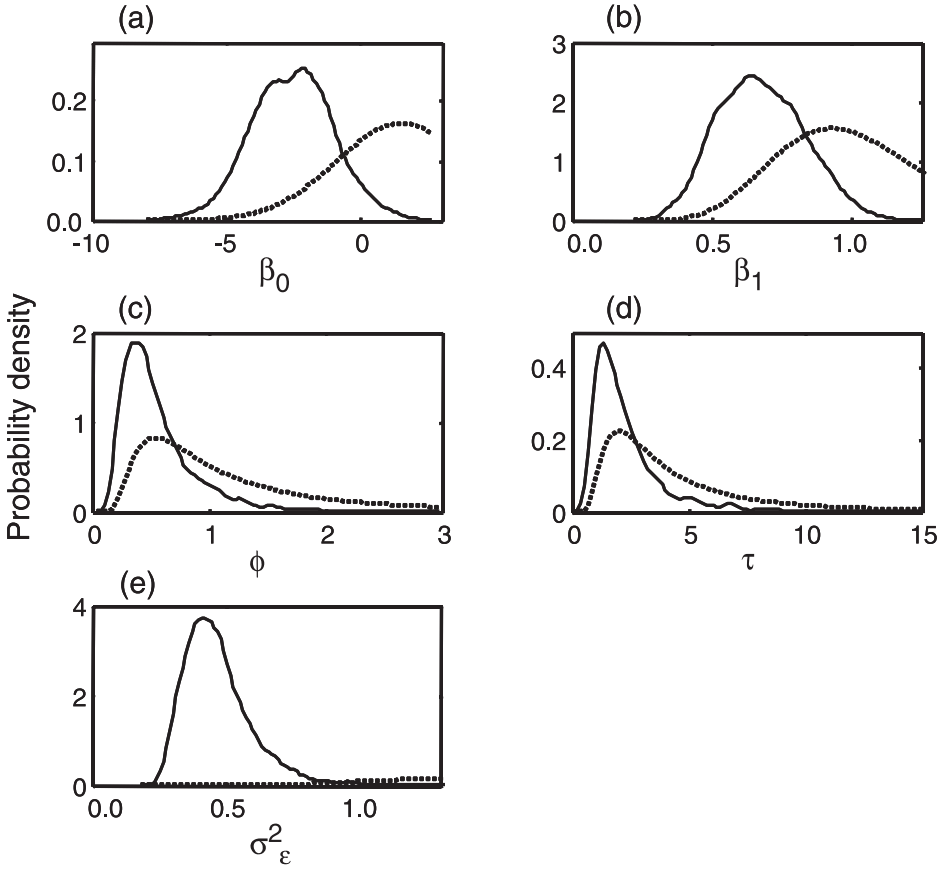


**Fig. 10.** Comparison of confidence bounds. The blue dots are the same data sampled from a Ricker model shown in Fig. 3*a*. The solid and broken black lines are the Bayesian nonparametric mean and confidence interval. The solid and broken green lines are the fit of the Ricker model including a 95% posterior uncertainty region for the Ricker fit.

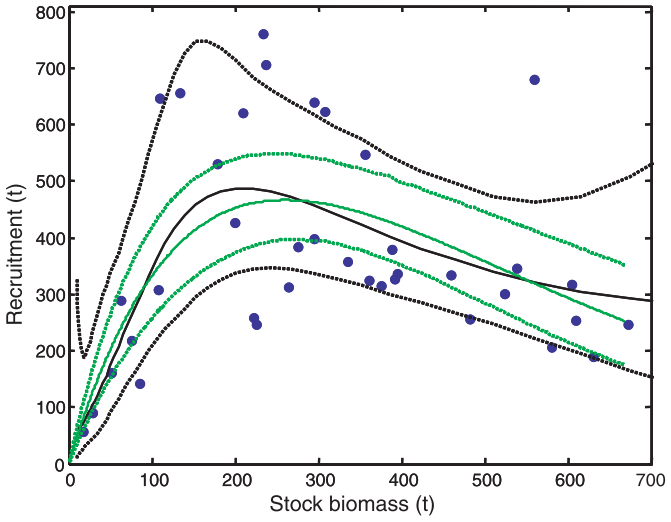


Regarding previous nonparametric approaches, the advantages are clear: no ad hoc smoothing parameter need be chosen, and biological information is easily incorporated through specification of the prior mean function. Moreover, choice of the prior mean function does not strictly dictate the shape of $\Im(x)$ within the range of the data.

As with other nonparametric methods (and parametric models as well; Ludwig 1995), biologically unreasonable fits are possible outside the range of the data. However, outside the range of the data, confidence limits grow rapidly and the posterior returns to the prior mean function, indicating that there is little information added by the data to points outside the range of the data (Ludwig 1995). Consequently, choosing a prior mean function close to the true underlying function will produce better extrapolations. Similarly, uncertainty bands near the origin tend to be large when there are no data near the origin, as in the lingcod example. Although this may at first glance appear to be a drawback, we argue that this is precisely what one wants from an inferential technique; regardless of their importance as management tools, point estimates of recruitment from parametric models far outside the observed range of the data are little more than elegantly concocted fictions.

The chief advantage of BNP over parametric approaches to inference is that only one model is required to fit the data regardless of the true underlying dynamics; dealing with model uncertainty is explicitly part of the BNP approach. Thus, one could envision a scenario in which Ricker density dependence is postulated in the choice of prior mean function, but the posterior mean behaves as a Beverton–Holt. Moreover, it is certainly possible to use any of the three pa-

rameter models as the prior mean function and retain the biological interpretations of the parameters in cases where the deviations from the prior mean function are not significant (as would be evidenced from a nearly diagonal structure in the posterior covariance function). With previous methods, multiple models (parametric or nonparametric) would have to be fit. Subsequently, researchers could either conduct a probabilistic assessment (e.g., via AIC) of which model was better or use the posterior probabilities associated with each model in Bayesian model averaging. Model averaging is an effective means of dealing with uncertainty in model formulation, provided that linear combinations of the candidate models span all of the SR relationships that are plausible for a given stock. If this condition cannot be met, then more flexible methods are preferable.

The BNP model allows good description of data sampled from parametric models, coming nearly as close to the underlying true model as fitted parametric curves. Moreover, with an unknown true model, BNP is the only single method that works effectively in all cases. It is, for example, the only single model that came close to the true model in the case of the open-mixture SR relationship. It should be noted, however, that the added flexibility of the BNP method does not come without costs. First, as a consequence of relaxing the prior specification of the functional form of the model, posterior confidence bounds for the fitted function are substantially wider than those of a correct parametric model (Fig. 10). However, it should be noted that the uncertainty bounds for the Ricker model are somewhat pathological in that they are insensitive to the uncertainty in the observed data; the bounds are of a uniform width in data-rich and data-poor regions alike. Second, when the correct model is known, parametric models allow easy analytical calculation of management reference points. Calculation of comparable reference points for any nonparametric method must be done numerically. However, given the large number of complex ecological and abiotic factors contributing to recruitment, it seems likely that there is no single SR relationship and certainly no analytical formula for one, so that there are no formulae (simple or complex) relating parameters of the SR relationship to reference points. As begun in this manuscript, it is possible to develop probabilistic definitions for important reference points. Further efforts are needed, particularly in developing an index of depensation free from rigid parametric specifications.

While the BNP approach represents a step forward in modeling the SR relationship, there are a number of issues that we have not explicitly addressed. These problems in the modeling of SR data have been raised by other authors (reviewed in Needle 2002) and are raised here to point out directions in which the BNP approach may be expanded. The first is that we have assumed measurement errors in stock biomass are small compared with estimates of recruitment. This may be explicitly addressed in the BNP framework by including an additional layer in the model hierarchy representing the unobserved, true state of the stock. Second, we have explicitly ignored the time series nature of SR data and have also explicitly assumed a stationary covariance function. However, it is possible to incorporate a time series structure in the BNP framework and it is straightforward to allow nonstationarity in the GP prior over stock sizes. Moreover, it is possible to incorporate temporal nonstationarity in the BNP framework by adopting approaches commonly used in dynamic linear models (West and Harrison 1997), which would allow us to account for the possibility that the nature of the relationship between stock and recruits is not constant over time. Third, stock biomass may not be an appropriate indicator of reproductive output (Rochet 2000). This is a valid concern, but it cannot be addressed in a statistical framework. Rather, the realities of reproductive biology need to be addressed by the regular collection of relevant data. A future paper will address these and other extensions currently under development.

Stock assessment using BNP methods provides a new tool for fishery management that allows one to use the data to dictate the shape of the SR relationship. Given the great objective flexibility of this method and the certainty that for real fish populations the true underlying SR relationship, if one exists, is unknown, the potential contribution to fishery management of this new approach is great.

## Acknowledgements

This work was partially supported by the Center for Stock Assessment Research, University of California at Santa Cruz.

## References

Barrowman, N.J., and Myers, R.A. 2000. Still more spawner recruitment curves: the hockey stick and its generalizations. Can. J. Fish. Aquat. Sci. **57**: 665–676.

Berger, J.O., De Oliveira, V., and Sanso, B. 2001. Objective Bayesian analysis of spatially correlated data. J. Am. Stat. Assoc. **96**: 1361–1374.

Beverton, R.J.H., and Holt, S.J. 1957. On the dynamics of exploited fish populations. G.B. Minist. Agric. Fish. Food Fish. Invest. Ser. II, **19**: 533.

Bravington, M.V., Stokes, T.K., and O Brien, C.M. 2000. Sustainable fishing: the bottom line. Mar. Freshw. Res. **51**: 465–475.

Brodziak, J.K.T., Overholtz, W.J., and Rago, P.J. 2002. Reply: Does spawning stock affect recruitment of New England groundfish? Interpreting spawning stock and recruitment data in New England groundfish. Can. J. Fish. Aquat. Sci. **59**: 193–195.

Burnham, K.P., and Anderson, D.R. 1998. Model selection and inference. A practical information-theoretic approach. Springer-Verlag, New York.

Chapman, D.G. 1973. Spawner–recruit models and estimation of the level of maximum sustainable catch. Rapp. P.-V. Réun. Cons. Int. Explor. Mer, **164**: 325–332.

Cleveland, W.S. 1981. LOWESS — a program for smoothing scatterplots by robust locally weighted regression. Am. Stat. **35**: 54.

Cook, R.M. 1998. A sustainability criterion for the exploitation of North Sea cod. ICES J. Mar. Sci. **55**: 1061–1070.

Cook, R.M. 2000. A rough guide to population change in exploited fish stocks. Ecol. Lett. **3**: 394–398.

Cressie, N. 1993. Statistics for SPATIAL DATA. Revised edition. Wiley, New York.

Cushing, D.H. 1973. Dependence of recruitment on parent stock. J. Fish. Res. Board Can. **30**: 1965–1976.

Denison, D.G.T., Holmes, C.C., Mallick, B.K., and Smith, A.F.M. 2002. Bayesian methods for nonlinear classification and regression. John Wiley & Sons, New York.

Essington, T.E., Quinn, T.P., and Ewert, V.E. 2000. Intra- and interspecific competition and the reproductive success of sympatric Pacific salmon. Can. J. Fish. Aquat. Sci. **57**: 205–213.

Evans, G.T., and Rice, J.C. 1988. Predicting recruitment from stock size without mediation of a functional relation. J. Cons. Int. Explor. Mer, **45**: 73–81.

Hastings, W.K. 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika, **57**: 97–109.

Hilborn, R., and Mangel, M. 1997. The ecological detective. Confronting models with data. Princeton University Press, Princeton, N.J.

Iles, T.C. 1994. A review of stock–recruitment relationships with reference to flatfish populations. Neth. J. Sea Res. **32**: 399–420.

Jagielo, T.H., Wilson-Vandenberg, D., Sneva, J., Rosenfield, S., and Wallace, F. 2000. Assessment of lingcod (*Ophiodon elongates*). *In* Status of the Pacific Coast groundfish fishery through 2000 and recommended biological catches for 2001: stock assessment and fishery evaluation. Pacific Fishery Management Council, Portland, Oreg.

Ludwig, D. 1995. Uncertainty and fisheries management. Lect. Notes Biomath. **100**: 516–528.

Mace, P.M., and Doonan, I.J. 1988. A generalized bioeconomic simulation model for fish population dynamics. N.Z. Fish. Assess. Res. Doc. 88/4.

McAllister, M.K., and Kirkwood, G.P. 1998. Bayesian stock assessment: a review and example application using the logistic model. ICES J. Mar. Sci. **55**: 1031–1060.

Millar, R.B. 2002. Reference priors for Bayesian fisheries models. Can. J. Fish. Aquat. Sci. **59**: 1492 1502.

Müller, P., and Quintana, F.A. 2004. Nonparametric Bayesian data analysis. Stat. Sci. **19**: 95–110.

Myers, R.A. 2001. Stock and recruitment: generalizations about maximum reproductive rate, density dependence, and variability using meta-analytic approaches. ICES J. Mar. Sci. **58**: 937–951.

Neal, R.M. 1999. Regression and classification using Gaussian process priors. *In* Bayesian statistics 6. *Edited by* J.M. Bernardo, J.O. Berger, A.P. David, and A.F.M. Smith. Clarendon Press, Oxford, UK.

Needle, C.L. 2002. Recruitment models: diagnosis and prognosis. Rev. Fish Biol. Fish. **11**: 95–111.

Papoulis, A. 1984. Probability, random variables and stochastic process. McGraw-Hill Book Company, New York.

Quinn, T.J.I., and Deriso, R.B. 1999. Quantitative fish dynamics. Oxford University Press, New York.

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

Rochet, M.J. 2000. Does the concept of spawner per recruit make sense? ICES J. Mar. Sci. **57**: 1160–1174.

Rothschild, B.J., and Mullen, A.J. 1985. The information-content of stock-and-recruitment data and its nonparametric classification. J. Cons. Int. Explor. Mer, **42**: 116–124.

Shepherd, J.G. 1982. A versatile new stock–recruitment relationship for fisheries and the construction of sustainable yield curves. J. Cons. Int. Explor. Mer, **40**: 67–75.

Walker, S.G., Damien, P., Laud, P.W., and Smith, A.F.M. 1999. Bayesian nonparametric inference for random distributions and related functions (with discussion). J. R. Stat. Soc. Ser. B (Methodol.), **61**: 485–527.

West, M., and Harrison, J. 1997. Bayesian forecasting and dynamic models. Springer Series in Statistics. Springer-Verlag, New York.

## Appendix A. Markov chain Monte Carlo (MCMC) sampling for the parameters governing $\Im(x)$.

This appendix describes an MCMC algorithm (Gamerman 1997) for sampling $\boldsymbol{\theta}$ in the more general setting where $\alpha$ is not fixed at 2. Because the elements of $\boldsymbol{\theta}$ and $\sigma_\varepsilon^2$ have restricted support, sampling is facilitated by defining a new set of parameters ($\boldsymbol{\eta}$) with support on $R^6$. This new parameter vector is given by $\boldsymbol{\eta} = \{\beta_0, \log \beta_1, \log \tau^2, \log \phi, \log[\alpha/(2-\alpha)], \log \sigma_\varepsilon^2\}$ and is sampled from a six-variate Gaussian ($N_6$) proposal distribution with mean given by the current value of $\boldsymbol{\eta}$ and covariance matrix $D$. These proposed values were accepted or rejected using the standard Metropolis–Hastings algorithm (Hastings 1970). Since we sample the vector of proposed values, $\boldsymbol{\eta}$, from a multivariate normal, the corresponding density for proposed new parameter values $\{\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}}\}$ is given by

$$\text{(A1)} \quad q(\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}} | \boldsymbol{\theta}, \sigma_\varepsilon^2, D) = \frac{2}{\beta_1^{\text{new}} \tau^{2\,\text{new}} \phi^{\text{new}} \alpha^{\text{new}} (2-\alpha^{\text{new}}) \sigma_\varepsilon^{2\,\text{new}}} N_6(\boldsymbol{\eta}(\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}}) | \boldsymbol{\theta}, \sigma_\varepsilon^2, D)$$

where the fraction following the equal sign is the Jacobian of the transformation relating $\boldsymbol{\eta}$ to $(\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}})$.

Because of the symmetry of the Gaussian proposal density, the acceptance probability ($p$) for a particular draw $(\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}})$ simplifies to

$$\text{(A2)} \quad p = \min\left\{1, \frac{\bar{p}(\boldsymbol{\theta}^{\text{new}}, \sigma_\varepsilon^{2\,\text{new}})}{\beta_1 \tau^2 \phi \alpha (2-\alpha) \sigma_\varepsilon^2} \cdot \frac{\beta_1^{\text{new}} \tau^{2\,\text{new}} \phi^{\text{new}} \alpha^{\text{new}} (2-\alpha^{\text{new}}) \sigma_\varepsilon^{2\,\text{new}}}{\bar{p}(\boldsymbol{\theta}, \sigma_\varepsilon^2)}\right\}$$

where $\bar{p}(\boldsymbol{\theta}, \sigma_\varepsilon^2)$ is the unnormalized posterior from the model in eq. 9.

During the initial sampling, the elements of $\boldsymbol{\eta}$ are sampled independently (i.e., $D$ is initially diagonal), but substantial improvements in sampling may be had with a better choice of $D$. We improved sampling by iteratively updating $D$ based on the covariance of the sampled parameters after each 5000 draws. After four such iterations, we used a burn-in period of another 5000 samples, which was sufficient to ensure convergence.

### References

Gamerman, D. 1997. Markov chain Monte Carlo: stochastic simulation for Bayesian inference. CRC Press, London, UK.

Hastings, W.K. 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika, **57**: 97–109.

## Appendix B. Posterior sampling for $y_{new}$ and $\Im(x_{new})$.

The quantities of interest are $p(\Im(\boldsymbol{x}_{new})|\,\text{data}, \boldsymbol{x}_{new})$ and $p(\boldsymbol{y}_{new}|\,\text{data}, \boldsymbol{x}_{new})$. Note that

$$\text{(B1)}\quad p(\Im(\boldsymbol{x}_{new}|\,\text{data}, \boldsymbol{x}_{new}) = \int p(\Im(\boldsymbol{x}_{new})|\,\text{data}, \boldsymbol{\theta}, \sigma_\varepsilon^2, \boldsymbol{x}_{new}) p(\boldsymbol{\theta}, \sigma_\varepsilon^2|\,\text{data})\mathrm{d}\boldsymbol{\theta}\sigma_\varepsilon^2$$

and

$$\text{(B2)}\quad p(\boldsymbol{y}_{new}|\,\text{data}, \boldsymbol{x}_{new}) = \int p(\boldsymbol{y}_{new}|\,\text{data}, \boldsymbol{\theta}, \sigma_\varepsilon^2, \boldsymbol{x}_{new}) p(\boldsymbol{\theta}, \sigma_\varepsilon^2|\,\text{data})\mathrm{d}\boldsymbol{\theta}\sigma_\varepsilon^2$$

Thus, the posterior predictive distributions B1 and B2 can be obtained by Monte Carlo integration of the conditional distributions for $\Im(\boldsymbol{x}_{new})$ and $\boldsymbol{y}_{new}$. To do so, we require the distribution for $\boldsymbol{y}_{new} = \{y_{new\,1}, ..., y_{new\,k}\}$ given the data, parameters, and $\boldsymbol{x}_{new}$. Based on eq. 1, the analogous distribution for $\Im(x)$ follows directly.

The distribution for $\boldsymbol{y}_{new}$ given the data, parameters, and $\boldsymbol{x}_{new}$ is obtained from the joint distibution of $\boldsymbol{y}$ and $\boldsymbol{y}_{new}$:

$$\text{(B3)}\quad p(\boldsymbol{y}_{new}|\,\boldsymbol{\theta}, \sigma_\varepsilon^2, \text{data}, \boldsymbol{x}_{new}) = \frac{p(\boldsymbol{y}_{new}, \boldsymbol{y}|\,\boldsymbol{\theta}, \sigma_\varepsilon^2, \boldsymbol{x}, \boldsymbol{x}_{new})}{p(\boldsymbol{y}|\,\boldsymbol{\theta}, \sigma_\varepsilon^2, \boldsymbol{x})}$$

The term in the denominator is the density $N_n(\boldsymbol{\mu}^n(\boldsymbol{\theta}), C^n(\boldsymbol{\theta}) + \sigma_\varepsilon^2 \boldsymbol{I}_n)$ given in eq. 9, while the term in the numerator is an $(n + k)$-variate normal distribution. The mean vector for this distribution is given by $(\boldsymbol{\mu}^n(\boldsymbol{\theta})', \boldsymbol{\mu}^k(\boldsymbol{\theta})')'$, where $\boldsymbol{\mu}^k(\boldsymbol{\theta}) = (\beta_0 + \beta_1 x_{new\,1}, ..., \beta_0 + \beta_1 x_{new\,k})'$. The covariance matrix is given by $\mathbf{C}^{n+k}(\boldsymbol{\theta}) + \sigma_\varepsilon^2 I_{n+k}$ where $\mathbf{C}^{n+k}(\boldsymbol{\theta})$ is

$$\text{(B4)}\quad \mathbf{C}^{n+k}(\boldsymbol{\theta}) = \begin{pmatrix} C^n(\boldsymbol{\theta}) & C^{n,\text{new}}(\boldsymbol{\theta}) \\ {\scriptstyle (n \times n)} & {\scriptstyle (k \times k)} \\ C^{n,\text{new}}(\boldsymbol{\theta})^T & C^{\text{new}}(\boldsymbol{\theta}) \\ {\scriptstyle (k \times n)} & {\scriptstyle (k \times k)} \end{pmatrix}$$

where $C^{\text{new}}(\boldsymbol{\theta}_{j,i}) = \tau^2 \exp(-\phi|x_{new\,i} - x_{new\,j}|^\alpha)$ and $C^{n,\text{new}}(\boldsymbol{\theta})_{i,j}^T = \tau^2 \exp(-\phi|x_{new\,i} - x_j|^\alpha)$.

Based on standard theory for multivariate normal variables, the conditional distribution for $\boldsymbol{y}_{new}$ in eq. B3 is also multivariate normal with mean vector $(\boldsymbol{M}_{new})$ and covariance matrix $(\mathbf{S}_{new})$ given by

$$\text{(B5)}\quad \boldsymbol{M}_{new} = \boldsymbol{\mu}^k(\boldsymbol{\theta}) + C^{n,\text{new}}(\boldsymbol{\theta})^T (C^n(\boldsymbol{\theta}) + \sigma_\varepsilon^2 I_n)^{-1} (\boldsymbol{y} - \boldsymbol{\mu}^n(\boldsymbol{\theta}))$$

and

$$\text{(B6)}\quad \mathbf{S}_{new} = C^{\text{new}}(\boldsymbol{\theta}) + \sigma_\varepsilon^2 I_k - C^{n,\text{new}}(\boldsymbol{\theta})^T (C^n(\boldsymbol{\theta}) + \sigma_\varepsilon^2 I_n)^{-1} C^{n,\text{new}}(\boldsymbol{\theta})$$

It can be readily seen from eq. B1 that the corresponding conditional distribution for $\Im(\boldsymbol{x}_{new})$ is also multivariate normal with mean vector $\boldsymbol{M}_{new}$ and covariance matrix given by $\mathbf{S}_{new} - \sigma_\varepsilon^2 I_k$. Having established $p(\boldsymbol{y}_{new}|\,\boldsymbol{\theta}, \sigma_\varepsilon^2, \text{data}, \boldsymbol{x}_{new})$ and $p(\Im(\boldsymbol{x}_{new})|\,\boldsymbol{\theta}, \sigma_\varepsilon^2, \text{data}, \boldsymbol{x}_{new})$ through eqs. B3–B6, draws from the posterior predictive distributions for $\boldsymbol{y}_{new}$ and $\Im(\boldsymbol{x}_{new})$ are readily obtained using eqs. B1 and B2 and the posterior draws for $\boldsymbol{\theta}$ and $\sigma_\varepsilon^2$. Hence, full inference is available for any feature of the posterior predictive distribution for the stock–recruitment relationship.