CCSBT-ESC/1009/10

CSIRO

UPDATED TECHNICAL SPECIFICATIONS AND PERFORMANCE ANALYSES FOR THE BREM (BIOMASS RANDOM EFFECT MODEL) SUITE OF CANDIDATE MANAGEMENT PROCEDURES

Richard Hillary
J. Paige Eveson
Marinelle Basson
Campbell Davies

Prepared for the 15th Meeting of the Extended Scientific Committee,
4th – 10th September 2010, Taipei, Taiwan.

Contents

1 Introduction 2

2 The BREM model 2

2.1 The population and estimation model 3

2.2 Performance on historical data 5

2.3 Including catch data in the model 7

3 Updated BREM harvest control rule 8

4 Tuning and fixed MP parameters 10

5 Figures 13

Abstract

This paper details (i) a brief explanation of the relative abundance population model at the core of the BREM (biomass random effect model) suite of candidate management procedures (CMPs), (ii) an updated performance assessment of this model on the latest historical data, and (iii) the changes made to the harvest control rule used within the BREM CMP as per the recommendations of the OMMP meeting in Seattle. The estimation framework and core population model was found to explain the historical data well, with the new changes in the underlying harvest controle rule more than able to cover the recommendations of the OMMP group.

1 Introduction

A set of candidate management procedures were defined and, for the model-based ones, tested on historical data in [1]. This document serves as an update to the sections of [1] concerned with the BREM suite of model-based CMPs. It also serves as the technical accompaniment to [2] which details the evaluation of the suite of BREM CMPs defined in this paper. With a model-based MP it is important to establish some basic performance statistics and criteria prior to full MP testing - the most fundamental of which would be how the model performs on the historical set of data available. The next phase is to clearly outline the harvest control rule (HCR) of the MP - the functional combination of the model parameters and process variables and the historical harvest levels that yield the future harvest level.

2 The BREM model

The BREM (**B**iomass **R**andom **E**ffects **M**odel) model-based approach is a relative abundance one where the dynamics of the adult biomass are decomposed into random recruitment and growth effects. This approach uses a variant on a model advocated in [3] which looked to estimate trends in recruitment and adult biomass as well as adult biomass net growth using random-effect methods.

2.1 The population and estimation model

The core population model is itself very simple: recruitment (R_y) and adult (B_y) biomass are related as follows:

$$B_{y+1} = R_y + g_y B_y, \tag{1}$$

where g_y is the adult biomass net growth effect (encompassing natural mortality, surplus production and exploitation effects). For the recruitment process the following model is assumed:

$$R_y = \exp\left(\mu_R + \epsilon_y^R\right), \tag{2}$$

with $\epsilon_y^R \sim N\left(-\sigma_R^2/2, \sigma_R^2\right)$. For the g_y a conceptually similar model is assumed and

$$g_y = \exp\left(\mu_g + \epsilon_y^g\right), \tag{3}$$

with $\epsilon_y^g \sim N\left(-\sigma_g^2/2, \sigma_g^2\right)$. For the aerial survey data I_y^{AS} a lognormal relationship to the recruiting biomass is assumed but with a one-year delay: $I_y^{AS} \sim LN\left(q^R R_{y+1}, \sigma_{AS}^2\right)$. The reason for this delay is because we assume that the aerial survey covers ages 2 to 4 and that the CPUE covers ages 4 to 12/18. To make sure that we are more likely to detect the movement of a signal in the aerial survey appearing in the CPUE data this delay is assumed as R_y represents the recruitment biomass contribution to the adult biomass (assumed covered by the CPUE). The situation is simpler for the CPUE likelihood and these data are assumed log-normally distributed about the adult biomass: $I_y^B \sim LN\left(q^B B_y, \sigma_B^2\right)$.

The model as it stands is non-identifiable which was explored at length in [3]. Without at least some information as to the ratio of the recruit and adult catchability parameters q^R/q^B then it will be impossible to identify how much recruitment affects biomass trends and how much the net growth of the biomass affects the biomass trends. To solve this problem we look to the output from the OM runs. From the grid runs we can extract the log catchability parameters for both the aerial survey and the CPUE data. Given the grid samples over parameters that will clearly alter this ratio (natural mortality, steepness, age range covered by the CPUE) we bootstrapped the mean difference in the log-catchabilities to obtain a best estimate of this ratio. The bootstrapped mean ratio was very precise (around a 3% CV) with an expected value of $q^{AS}/q^{CPUE} = 50254.36$. However, we need to account for the fact that the CPUE in the OM is in *numbers* but here we are trying to relate biomass to biomass. To take account of this in our catchability ratio consider the following ratio:

$$\psi_q = \frac{\sum\limits_{i=a_l}^{a_u} \pi_i^s w_a}{\sum\limits_{i=a_l}^{a_u} \pi_i^s} \tag{4}$$

where π_i^s is the survival probability from age 0 to age i and a_l and a_u are the minimum and maximum ages observed in the CPUE, respectively. This ratio is readily calculable from the grid files outputted from the OM. For each sampled grid cell this ratio was calculated and then a bootstrapped mean and CV were calculated, to robustify the estimates given the banding by M grid option. As with the q ratio estimates the numbers were very precise: a mean and CV of 0.0616 and 0.026, respectively. Assuming $q^B = 1$ this lead to a value of $q^R = q^{AS}/q^{CPUE} \times \psi_q = 3100.912$. In terms of the recruitment biomass variance term a value of $\sigma_R = 0.376$ is employed, as this corresponds to the amount of variation one would expect to see in the aerial survey index (covering ages 2, 3 and 4 for a selectivity of 0.5, 1 and 1, respectively) *only* due to variation in recruitment at age 0 with an assumed SD of 0.6 as per the OM. This was calculated by running a stochastic per-recruit unexploited population for 100 years (with the mean M-vector from the OM) and calculating the SD in the population covered by the aerial survey. The reason for choosing a value of $\sigma_g = 0.246$ was based on a CV of 0.25 which is the mid-point of the process error applied when simulating the CPUE data. In terms of the observation error assumed in the estimation scheme CVs of 0.15 and 0.2 were assumed for the aerial survey and the CPUE data, respectively, given the recent estimates from the aerial survey and the minimum value assumed in the OM conditioning.

The actual parameters to be estimated are μ_R, μ_g, ϵ_y^R and ϵ_y^g. To avoid identification issues with the recruitment in the first year and net growth year effects in the last year, respectively, they were penalised to have mean zero across years (with $-100 \times |\mathbb{E}[\epsilon_y^\bullet]|$ extracted from the log-likelihood). Although maximum posterior density estimates were used in the MP testing full MCMC routines were developed to explore the parametric and process variable uncertainty in the underlying models in this phase - the chief reason being we can obtain more detailed information about the variation in the derived trends such as stock growth, recruitment and biomass not retrievable from the ADMB runs. While using the term random effect, to be clear this model is more of a Bayesian hierarchical model: a Dirac/point hyperprior is defined for the variance hyperparameters $\sigma_\bullet^2$, which then form the priors for the $\epsilon_y^\bullet$ parameters. This contrasts with the strict view of a random effects model which utilises expectation/maximisation to estimate all the key parameters: expectation where the joint penalised likelihood of the $\mu_\bullet$ and $\epsilon_y^\bullet$ is integrated over the $\epsilon_y^\bullet$ and maximisation where this marginal likelihood is then maximised for the $\mu_\bullet$.

2.2 Performance on historical data

Figure 1 shows the marginal posterior summary for the parameters μ_R and μ_g and the parameters have posterior mean (and SD) of -1.93 (0.067) and -0.3 (0.037), respectively, with fairly strong negative correlation between these two parameters (-0.52) as one would expect if recruitment makes a significant contribution to the exploitable biomass. The estimated trends in recruit biomass, adult biomass and biomass growth can be seen in Figure 2. For the relative recruitment biomass estimates we observe a sharp decline around 1998 (as seen in 1997 in the aerial survey) hitting the lowest level in 2000. From 2001 to 2004 the estimates are driven by the prior and penalty terms given the absence of data in the aerial survey with the levels of recruitment in 2005 to 2008 staying around the low level but with an upturn in 2009. In the years where there are data to estimate the recruitment trend the CVs ranged from 0.133 to 0.144. For the relative adult biomass estimates we must first point out that we assume that $B_{1994} = I^B_{1994}/q^B$ and that it is known without error (there are no data to estimate it and we assume a relative abundance model anyway). As one would expect the trend follows that in this particular CPUE series (including the gradual decline from 2002-2007 and the sudden upturn in 2008 continuing into 2009). The CVs in the estimates (excluding 1994) range from 0.122 to 0.159 with a sustained increase in uncertainty in the middle of the range given the uncertain recruitment dynamics. For the biomass growth estimates they oscillate below 1 until approaching 1 from 1999-2001 when they show a marked decline as they alone can explain the biomass decline seen in 2002-2007 as recruitment has already dropped to the lower level by 1998. Clearly the sudden increase in 2008 in the biomass cannot be explained by recruitment and so the biomass growth parameter in 2007 increases to a value just above 1 in this year. The continued increase in biomass (given the CPUE) from 2008 to 2009 seems to be attributable to biomass growth also - the estimate of biomass growth in 2008 is just less than 1 and the recruitment estimate from 2008 is the same as 2007 and not above the average. The estimate of biomass growth in 2009 is driven by both the prior and the penalties and should not be viewed with close scrutiny.

With regards to the relative importance of recruitment versus biomass growth in terms of the adult biomass looking at the year-averaged values of the ratio $g_y B_y / R_y$, we observed a median (and 95% credible interval) of 3.27 (2.34-4.54). This suggests that, at least over the data range 1993-2009, the net growth of the biomass from the previous year is the dominant contributor to the biomass in the following year, not incoming recruitment. The model is a relative abundance model but we can derive the trends in relative harvest rate on the juvenile and adult sections of the population by taking the catch of the surface fishery and dividing it by R_y and by taking the catch of all the other fisheries and dividing it by B_y. A summary plot of the mean-standardised juvenile and adult relative harvest rates can be found in Figure 3. In terms of the juveniles 1993-1996 are well below the average, increasing in 1997 and 1998 as the estimates of juvenile biomass from the aerial survey decrease and increasing to well above average levels in 1999 and 2000, staying slightly below average until 2006 and 2007, then dipping below average levels in 2008. For the adult harvest rates, they stay close to but slightly below average levels from 1994 to 2004, but the exploitation rate begins consistently increasing from 2002 to 2007 driven by the apparent decrease in the biomass over this period, and are still well above average levels even with the rise in biomass in 2008.

In terms of fits to the data Figure 4 shows a summary of the estimators performance in this regard. For the aerial survey data they are generally fitted quite well but the extremes in these data (the apparently higher variance earlier on) are not fitted so well, presumably given the assumed value of σ_R. For the CPUE data they are also fitted quite well but the model cannot fit the more extreme changes observed in the CPUE - the data never sit outside the 95% credible interval but the median fitted CPUE is much smoother than the observed data. This again is due to the natural constraints placed upon both the recruitment and biomass growth effects via σ_R and σ_g, respectively.

From a Bayesian and MCMC perspective, one final analysis is to check the predictive power of the posterior model - how well does the model not just fit but "explain" the data - to be satisfied that the model is at least able to adequately deal with the historically observed data. An established way to do this is to perform a *posterior predictive* analysis [4]: data are simulated from the likelihood (given the posterior sample) and positive discrepancy statistics, ξ, denoting in some way the "closeness" of the simulated and real data to the model prediction are calculated. Bayesian p-values [5] can then be calculated as the probability that the simulated data are more "extreme" than the real data: $p\left(\xi^{\text{model}} > \xi^{\text{data}}\right)$. Bayesian p-values around 0.5 suggest good performance, in that the model is explaining (not just fitting) the data well - values above and below 0.5 can be indicative of the presence of unaccounted for process error and over-fitting, respectively. For the aerial survey data the p-value was 0.35 and for the CPUE data 0.49 so the aerial survey predictions were less extreme than the data, but for the CPUE they were pretty good. On the whole this suggest both parts of the probability model are performing well on the historical data. It is often useful to plot the data and model-predicted discrepancy statistics (in this case we used the absolute median deviation) and Figure 5 shows a summary of these - both form a fairly circular cloud around the $y = x$ line as we would like with no obvious strange visual patterns. On the whole this suggests both the aerial survey and CPUE parts of the probability model are performing well on the historical data.

2.3 Including catch data in the model

From the 3$^{\text{rd}}$ OMMP meeting members of the group suggested to look at the inclusion of the catch data within the model framework. To explore the potential for this a slight modification of the original model in Eq. 1 was envisaged:

$$B_{y+1} = R_y + g_y B_y - C_y, \tag{5}$$

where C_y is the catch biomass across all fisheries. To deal with initial conditions the following equilibrium assumption was made:

$$B_1 = \frac{\exp(\mu_R) - \tilde{C}}{1 - \exp(\mu_g)}, \tag{6}$$

where $\tilde{C}$ are the "equilibrium" catch levels set at around 14,000t given the pre-1994 catches.

Running the model over the same time-period (1994-2009) using an augmented MCMC Gibbs sampler (now estimating q^B) was not successful. The relative abundance model already had a reasonable amount of freedom and, when widening this freedom with the introduction of the catches and the wish to estimate absolute abundance, we asked more information of the data than it possesses. There was information in the data as to minimum values of μ_R - the population has not undergone huge declines at the current catch levels over the data range. However, the posterior for μ_R at higher levels was extremely flat, as there is no information in the data to tell the model about maximum values of μ_R, not to mention the very strong negative correlation between μ_R and q^B. This should not be surprising - imagine if one tried to estimate even K only from a surplus production model fitted to these data which cover only 1994 to 2009.

One could, in theory, extend the model back in time to use the strong historical decline in the long-line CPUE to better estimate μ_R and q^B but this has three major issues attached to it:

1. The lack of aerial survey data for effectively 25 years: we either restrict the estimation of recruitment effects to the periods for which we have the survey or let the prior and penalty terms deal with the lack of information.
2. The number of extra parameters: going back to 1969 and out to 2040 would yield immensely more parameters, and with it potential estimation problems and time penalties in terms of MP evaluation.
3. We cannot use informative priors: unlike in [1] where we used life-history information to constrain the r parameter from the Pella-Tomlinson models we cannot do such a thing for μ_R. It is likely that population recovery will result in much clearer information on this parameter in the future but up until this occurs we have no way to sensibly constrain the system.

Adding the catch data would perhaps add something more real to the model, and also give us harvest rates that would be useful in an MP setting, but it would involve realistically having to drastically expand the time window of the model - perhaps too far. The model was envisaged as a biological and statistical filter for the aerial survey and CPUE data. Since we have established that the model can more than adequately statistically deal with the data trends and is very stable in terms of “blind” estimation in the initial MP testing work we suggest that we stay with the relative abundance form of the model for the MP work.

3 Updated BREM harvest control rule

At the 3rd OMMP meeting the group requested some changes to the HCR in the `BREM1` CMP [6]:

- To increase the strength of reaction of the MP to low biomass and recruitment signals - more specifically to make the responses greater than linear in nature.
- To smooth in some fashion the catch time series - to reduce if possible the strength of the early declines and the later increases without removing the ability of the MP to still react to adverse situations.
- To change the part of the HCR which deals with the estimated recruitment trends - instead of a moving (geometric) average of the relative recruitment ratio use a moving (arithmetic) mean of the recruitment compared to that estimated over the period for which we have aerial survey data.

To try and smooth the catch time series it was suggested [6] to impose a "memory" effect on the HCR as follows:

$$TAC_y = \psi_y TAC_{y-1} + (1-\psi_y) TAC_y^{\text{brem}}, \tag{7}$$

where $\psi_y \in [0,1]$ is a memory weighting term and TAC_y^{brem} is the TAC produced by the underlying `BREM` HCR. Similar to the general form defined in [1] we have that

$$TAC_y^{\text{brem}} = C_y^{\text{targ}} \times \Delta_y^R \times \Delta_y^g, \tag{8}$$

where

$$C_y^{\text{targ}} = \begin{cases} \delta \left[\dfrac{B_{y-2}}{B^*}\right]^{1-\varepsilon_b} & \text{for} \quad B_{y-2} \geq B^* \\ \delta \left[\dfrac{B_{y-2}}{B^*}\right]^{1+\varepsilon_b} & \text{for} \quad B_{y-2} < B^* \end{cases} \tag{9}$$

and $\varepsilon_b \in [0,1]$ so represents the degree to which the response to a biomass level above or below the target level B^* is asymmetric. It is hard-wired into the HCR that the response to biomass levels above the target can never be stronger than linear or greater than the response to levels below the target. The reason for such a change is to both accommodate the recommendation of stronger responses to low biomass levels and to aid in smoothing the response of the HCR when the stock has recovered sufficiently to and above the target levels.

The recruitment adjustment Δ_y^R is defined as follows:

$$\Delta_y^R = \begin{cases} \left[\frac{\bar{R}}{\mathcal{R}}\right]^{1-\varepsilon_r} & \text{for} \quad \bar{R} \geq \mathcal{R} \\ \left[\frac{\bar{R}}{\mathcal{R}}\right]^{1+\varepsilon_r} & \text{for} \quad \bar{R} < \mathcal{R} \end{cases} \tag{10}$$

and $\varepsilon_r \in [0,1]$ is the level of asymmetry in response to the current moving (arithmetic) average recruitment levels, $\bar{R}$:

$$\bar{R} = \frac{1}{\tau} \sum_{i=y-\tau-1}^{y-2} R_i, \tag{11}$$

of length τ relative to the average, $\mathcal{R}$, calculated over the years for which the estimates are based on observed data (1994-2009).

The final term is the stock growth term and no asymmetry in action is assumed so

$$\Delta_y^g = \left[\frac{\bar{g}}{\mathcal{G}}\right]^{\gamma}, \tag{12}$$

where

$$\bar{g} = \frac{1}{\tau} \sum_{i=y-\tau-1}^{y-2} g_i, \tag{13}$$

and $\mathcal{G}$ is mean value of g_y over which the estimates are based on real data (1994-2009). The term $\gamma \in [0,1]$ in Eq. 12 is an importance weighting term.

4 Tuning and fixed MP parameters

As before, there is only one tuning parameter: the reference catch level, δ. Instead of alternative CMPs we now have a single harvest control rule as defined in Section 3 but allow for various scenarios based on different settings for some of the key control parameters of the HCR.

Key HCR parameters that will vary across scenarios are:

- ψ_y: the TAC memory weighting which is also permitted to change across years.
- ε_b & ε_r: the asymmetry levels in the response to levels of biomass and recruitment above and below the target/historical mean levels.

Key HCR parameters that are kept fixed are:

- B^*: target relative biomass level (effectively a target CPUE level) of 1.2 as before.
- τ: length of the moving averages for the recruitment and biomass growth parameters set to 5 in all cases.
- γ: importance weighting of the biomass growth adjustment in the HCR set to 1 in all cases.

References

[1] Hillary, R., Eveson, J.P., Basson, M., and Davies, C. 2010. Technical specifications and proof of concept analyses for candidate management procedures for southern bluefin tuna. *CCSBT-OMMP/1006/4*

[2] Hillary, R., Giannini, F., Eveson, J.P., Basson, M., Davies, C. 2010. Results of the performance of the BREM suite of candidate management procedures. *CCSBT-ESC/1009/11*

[3] Trenkel, V. 2008. A two-stage biomass random effects model for stock assessment without catches: What can be estimated using only biomass survey indices? *Can. J. Fish. Aquat. Sci.* **65**(6): 1024–1035

[4] Gelman, A., Carlin, B.J., Stern, H.S. and Rubin, R.B. 1995. Bayesian data analysis. Chapman and Hall, London.

[5] Meng, X. 1994. Posterior predictive p-values Ann. Stat. 22, 1142–1160.

[6] Anonymous. 2010. Report of the Third Operating Model and Management Procedure Technical Meeting. Seattle, USA.

5 Figures

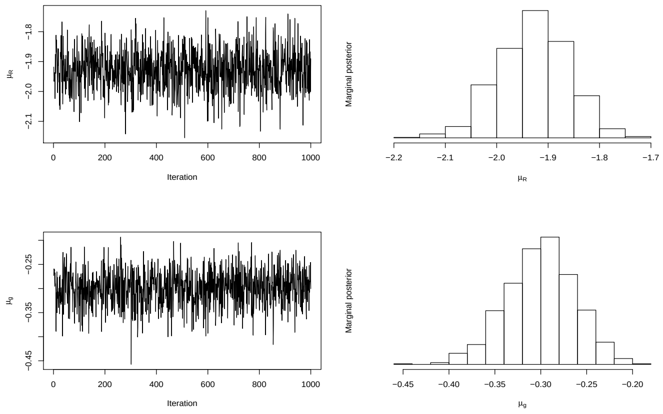

Figure 1: *Trace plots (left) and histograms (right) for the marginal posteriors of μ_R (top) and μ_g (bottom).*

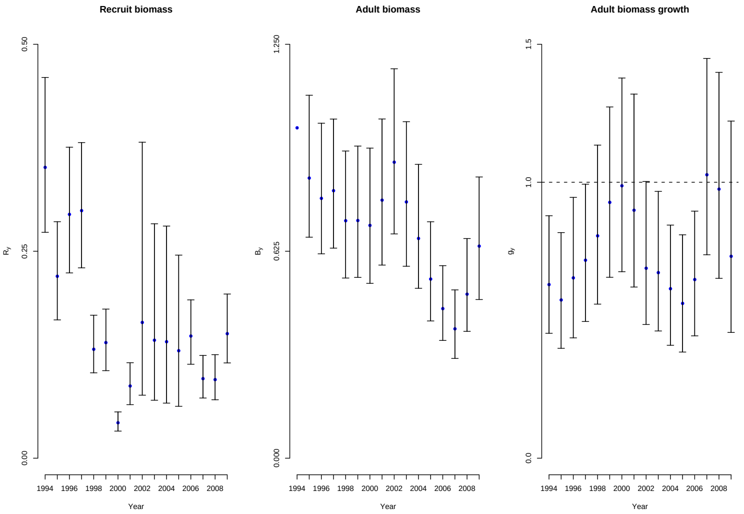

Figure 2: *Summary (median, circle; whiskers, 95% credible interval) of the relative recruitment biomass (left), relative adult biomass (middle) and net biomass growth (right) using the aerial survey and the CPUE data.*

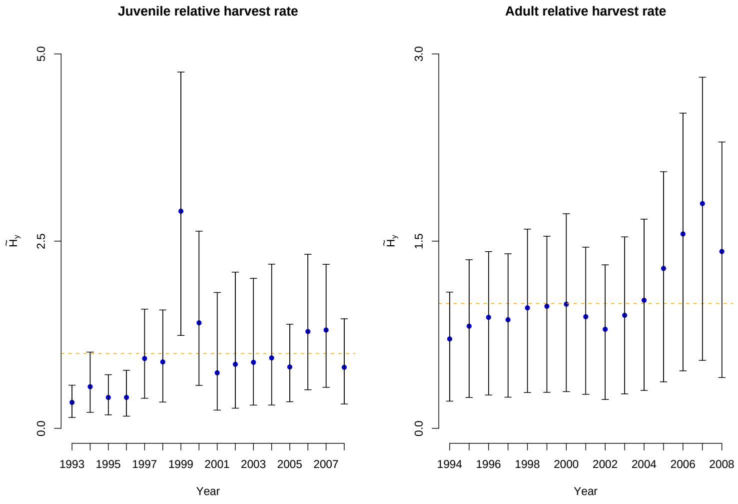

Figure 3: *Summary (median, circle; whiskers, 95% credible interval) of the relative juvenile (left) and adult (right) harvest rates, expressed as catch divided by relative biomass and are mean standardised with the orange dotted line being equal to 1.*

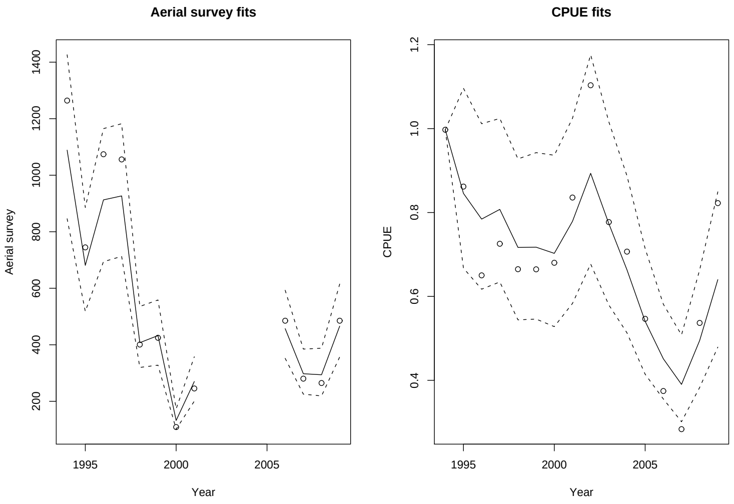

Figure 4: *Fits of the relative abundance biomass dynamic model to the aerial survey data (left) and the commercial CPUE data (right). The points are the data with the full and dashed lines representing the median and 95% credible intervals, respectively.*

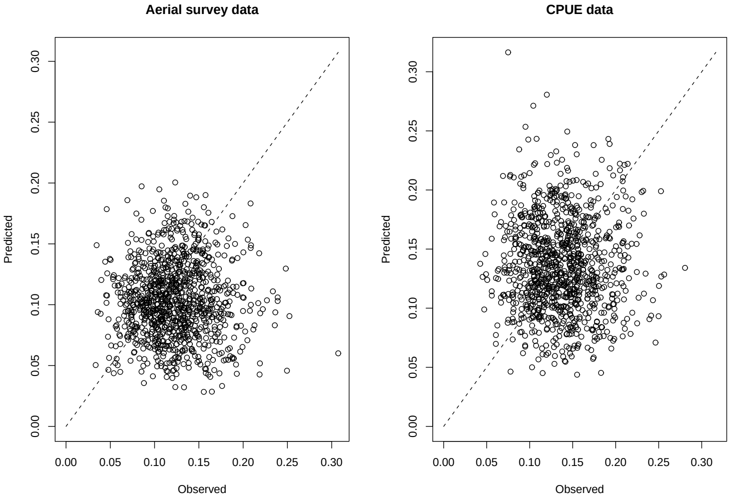

Figure 5: *Observed (x-axis) and predicted (y-axis) discrepancy statistics from the posterior predictive analysis for the aerial survey data (left; p-value of 0.35), and CPUE data (right; p-value of 0.49). The dotted diagonal line is the* $y = x$ *line.*