Two alternative approaches to the meta-analysis of environmental variability and auto-correlation in reproductive rates of baleen whales

JOHN R. BRANDON[1] AND TOSHIHIDE KITAKADO[2]

Contact e-mail: jbrandon@gmail.com

ABSTRACT

Two approaches are presented for the analysis of select data sets in the estimation of environmental variability and auto-correlation in reproductive rates of baleen whales. Both approaches recognize that for a given stock, the average calving interval is the reciprocal of the average proportion calving, and thus allow for the incorporation of both data types in a unified estimation framework for parameters of interest. For an unknown stock of baleen whale, the extent of environmental variability is (depending on the approach) estimated to be 0.347 and 0.396 (standard deviations in log-space), and the estimates of the auto-correlation parameter are 0.614 and 0.288. The estimates of the hyper-parameters from this meta-analysis framework can be used in simulations to inform the lower end of the range for MSYR values. In general, the resulting parameter estimates appear to be mostly consistent with SC/63/RMP20, which employed a different modelling framework and included more available data on reproductive variability. Therefore, these results may provide some confidence in the robustness of available estimates, given different data sets and modelling assumptions.

INTRODUCTION

An ongoing aspect of the RMP review has been the reconsideration of the plausible range used for maximum sustainable yield rate (MSYR) used for testing the *Catch Limit Algorithm* of the RMP. This range is currently 1% to 7% (expressed in terms of the mature component of the population). Information on observed population growth rates at low population sizes has been used to inform this range (e.g. Best, 1993). However, it has been recently pointed out that ignoring environmental variability and/or temporal autocorrelation in vital rates, and only using observed population growth rates could lead to a bias in the estimated lower end of the range of plausible values for MSYR (Cooke, 2007). In order to provide a basis for estimating the expected level of inter-annual variability and auto-correlation, an intersessional workshop was recently held (IWC, 2011) during which available time series for annual reproductive success in baleen whales were made available. To that end, the goal of this study is to incorporate suitable available data in a meta-analysis framework and provide estimates for the parameters necessary to re-evaluate the lower end of the expected range of MSYR values for an unknown stock of baleen whales.

METHODS

Data

Following the 3rd intersessional workshop on the review of MSYR for baleen whales (IWC, 2011), we determined that several data sets were relatively sparse, i.e. they did not

1. *Greeneridge Sciences, San Francisco, USA*
2. *Tokyo University of Marine Science and Technology, Tokyo, Japan*

have many data points in the time series (e.g., BCB bowheads), or had many years with observations of zero calving females (e.g., Gulf of CA blue whales). Hence, only a subset of the available data was used when fitting the models in these analyses.

Five stocks were included here. Four of these stocks have annual estimates available for both calving interval and proportion calving: 1) Gulf of Maine (GOM) humpback, 2) Gulf of St. Lawrence (GSL) humpback, 3) North Atlantic (NA) right, and; 4) South East Alaska (SE AK) humpback whales. Eastern North Pacific (ENP) gray whales were also included, but this stock only has annual estimates of calving proportions available. The calving interval and proportion calving data are shown in Figures 1 and 2.

General modeling approaches: 'TK' and 'JB'

Two approaches were used to fit the data in these analyses. These are denoted here as the 'TK' and 'JB' approaches. Both approaches are based (at least for certain scenarios) on hierarchical models across stocks and are similar to each other in many respects. They do differ from previous efforts (e.g. Cooke, 2011; Brandon and Kitakado, 2011) in that the estimated parameters for individual stocks are not treated independently when both interval and proportion data are available for the same stock. The key assumption underlying this difference is the recognition here that the average calving interval across years is the inverse of the average proportion calving. This allows the interval and proportion data for the same stock to be analyzed simultaneously by estimating a common parameter for each stock (as opposed to two different parameters for each stock, and likewise two different parameters across stocks for each type of data).

The TK and JB approaches were implemented using the random effects (RE) module of AD Model Builder software and WinBUGS, respectively. The TK approach provided estimates based on maximum likelihood values, and the JB approach provided estimates as probability distributions based on a Bayesian framework.

There are three noteworthy differences between the TK and JB approaches:

1) The TK approach assumed that the annual process error residuals were identical for a given stock, i.e. after taking into account the necessary transformation, the deviations from expected calving proportions and intervals were the same each year. The JB approach only assumed that the process error residuals were *i.i.d.* for each stock, but allowed for the annual residuals to be different between the proportions and intervals for each stock.

2) Each approach assumed a different hierarchical structure for the auto-correlation parameter ρ. Further details on the model structure are provided below.

3) For those interval data that did not have an empirical estimate of the sampling variance (i.e., GSL humpback interval data), the TK approach assumed that the sampling variance was equal to the maximum reported variance, while the JB approach assumed those missing variances were equal to the average variance from years with available estimates.

Proportion Data

Both the TK and JB approach can be considered to have modelled the proportion data in the same way.

Let ψ_{it} be the observed proportion of mature females for stock *i* calving in year *t*, such that:

$$\psi_{it} \mid p_{it} \sim Po(N_{it}\, p_{it}) \tag{1}$$

Where:

p_{it}	is the proportion calving for stock *i* in year *t*, and;
N_{it}	is the number of mature females for stock *i* in year *t*.

The annual proportion calving was assumed to be subject to environmental variability, such that:

$$\text{logit}(p_{it}) = \alpha_i + \varepsilon_{it} \tag{2}$$

Where:

α_i	is the expected proportion calving across years for stock *i*, and;
ε_{it}	is the deviation from that expectation for stock *i* in year *t* due to environmental variability, where the environmental deviations are assumed to be correlated through time, such that:

$$\varepsilon_{it} = \rho_i^{r_{it}} \varepsilon_{it-1} + \sqrt{1-\rho_i^{2r_{it}}}\, \eta_{it} \tag{3}$$

Where:

ρ_i	is the degree of auto-correlation for stock *i*;
r_{it}	is the number of years since the last observation for stock *i* in year *t* (e.g., for the first year r_{it} =0, if successive observations r_{it} =1, if a gap of one year between observations then r_{it} =2, etc.)[3];
η_{it}	is $\sim N(0,\sigma_i^2)$, and;
σ_i	is the standard deviation of the environmental process error for stock *i*.

Interval Data

The TK approach for the interval data is described below. This is also the same as the JB approach, except where otherwise noted (further details on the treatment of the process error residuals for the JB approach are provided under ‘The JB approach to the process error residuals for interval data.’).

Let y_{it} be the observed calving interval for stock *i* in year *t*, such that:

---

[3] Note that due to the different treatment of the process error residuals between the TK and JB approaches, that the values for r_{it} are also different for each approach. For example, the TK approach will on average have lower values for this input, because assuming identical process error residuals between data sets effectively reduces the average gap between available observations.

$$y_{it} \sim N(\phi_{it}, \hat{v}_{it}) \tag{4}$$

Where:

ϕ_{it} is the expected calving interval (equal to the inverse of the proportion calving, i.e. $\phi_{it} = 1/p_{it}$ for the TK approach, see below for the JB approach) for stock *i* in year *t*, and;

$\hat{v}_{it}$ is the known variance of the sampling error associated with the calving interval estimate.

The JB approach to the process error residuals for interval data

The calving interval each year was related to the expected calving proportion for each stock, such that:

$$\phi_{it} = 1/\text{antilogit}\left(\alpha_i + \gamma_{it}\right) \tag{5}$$

Where:

γ_{it} is the deviation from the expected calving interval for stock *i* in year *t* due to environmental variability, where the environmental deviations about the expected interval are assumed to be correlated through time, such that:

$$\gamma_{it} = \rho_i^{g_{it}} \gamma_{i,t-1} + \sqrt{1-\rho_i^{2g_{it}}}\, \xi_{it} \tag{6}$$

Where:

g_{it} is the number of years since the last annual calving interval observation for stock *i* in year *t* (see the definition of r_{it} for the proportion data);

ξ_{it} is $\sim N(0, \sigma_i^2)$, hence the annual process error residuals between proportions and intervals for each stock are not identical in the JB approach, but they are *i.i.d.*

Scenarios

TK Approach

Under the TK approach, five scenarios were examined. These involved different hierarchical structures for the parameters of interest, σ (the extent of process error due to environmental variability) and ρ (the auto-correlation in the process error residuals due to environmental variability). The scenarios were: 1) Common σ and common ρ over stocks; 2) Common σ with independent (but not necessarily *i.i.d.*) ρ_i for each stock; 3)

Independent σ_i for each stock and a common ρ across stocks; 4) Independent σ_i and independent ρ_i for each stock, and; 5) Hierarchical structure for σ_i and ρ_i across stocks, i.e. the values for each parameter were allowed to be different for each stock, but σ_i and ρ_i were assumed to be *i.i.d.* and each parameter shared common estimated hyper-parameters.

For this last scenario, σ_i for each stock was assumed to be drawn from an underlying log-normal distribution with a common variance, i.e. $\ln(\sigma_i) \sim N(0, \sigma_\sigma^2)$, where $\ln(\sigma_\sigma) \sim N(\theta_\sigma, \tau_\sigma^2)$. Likewise, the ρ_i were assumed to be drawn from an underlying distribution with a common variance across stocks, such that: $\text{logit}(\delta_i) \sim N(0, \sigma_\rho^2)$, where $\ln(\sigma_\rho) \sim N(\theta_\rho, \tau_\rho^2)$ and $\rho_i = 2 * \delta_i - 1$.

The TK approach estimated the parameters of interest (including $\theta_\sigma, \tau_\sigma, \theta_\rho$ and τ_ρ for the hierarchical scenario) based on a maximum likelihood framework by analytically integrating over the process error residuals, and compared the ability of the various scenarios to fit the data using AIC.

JB Approach

As stated above, this approach differed in its treatment of the hierarchical structure for ρ_i and also in the assumptions made regarding the process error residuals between the interval and proportion data.

The hierarchical structure for ρ_i was implemented by first transforming a beta random variable, such that:

$$\rho_i = 2\beta_i - 1 \tag{7}$$

Where:

$$\beta_i \sim Beta(\text{shape1}, \text{shape2}) \tag{8}$$

The hyper-parameters, ‘shape1’ and ‘shape2’, for the auto-correlation parameter were assumed to be log-normally distributed (keeping the resulting values positive as required by the assumption of the beta-distribution):

$$\begin{aligned} \ln\left(\text{shape1}\right) &\sim N(0, \sigma_\beta) \\ \ln\left(\text{shape2}\right) &\sim N(0, \sigma_\beta) \end{aligned} \tag{9}$$

Two values for σ_β were explored, corresponding with a ‘strict’ hyper-prior (σ_β = 0.20) and a ‘relaxed’ hyper-prior (σ_β = 1) on the shape parameters.

The JB approach to the process error residuals for interval data allowed the vectors of annual process error residuals to differ between interval and proportion data for the same stock (Eqns. 3 and 6). These random effects were however assumed to be *i.i.d.*, such that:

$$\eta_{it} \sim N(0,\sigma_i^2)$$
$$\xi_{it} \sim N(0,\sigma_i^2) \tag{10}$$

Where:

$$\ln(\sigma_i) \sim N(\theta_\sigma,\sigma_\sigma^2)$$
$$\theta_\sigma \sim N(0,10000) \tag{11}$$
$$\sigma_\sigma^{-2} \sim gamma(0.1,0.1)$$

Finally, in this approach, a hierarchical structure was assumed for the expected reproductive rates between species of baleen whales:

$$\alpha_i \sim N(\theta_\alpha,\sigma_\alpha^2)$$
$$\theta_\alpha \sim N(0,10000) \tag{12}$$
$$\sigma_\alpha^{-2} \sim gamma(0.1,0.1)$$

RESULTS

TK Approach

The results of the TK approach are presented in Table 1. The scenario (#3) with the lowest AIC value, estimated independent σ_i for each stock (with no hierarchical structure) and a single ρ common to all stocks. The average value for σ across stocks was estimated to equal 0.19, with the common ρ estimated to be 0.6242.

The scenario (#5: “Random”,”Random”) which has a hierarchical structure, and is therefore more informative with respect to the goals of this project, estimated the expected value of environmental process error for an unknown stock (mean_sigma in Table 3) to equal 0.3472. The standard error of this estimate in log-space (sd_logsigma in Table 3) was equal to 0.0718. The expected value for ρ for an unknown stock was estimated to be (mean_rho in Table 3) 0.6143, with a standard deviation in logit-space (sd_logitrho) of 3.013.

The estimated value (0.347) for the extent of process error for an unknown stock was similar to that (0.396) estimated by the JB approach. Both of these approaches estimated that the expected value of auto-correlation for an unknown stock was positive, but the hierarchical scenario for the TK approach estimated a higher value (0.6143) than the JB approach (median = 0.29).

JB Approach

The scenario with a ‘relaxed prior’ on ρ caused numerical difficulties and continued crashing of JB’s WinBUGS code. Hence, no results are available for alternative hyper-prior values for ρ under this scenario. However, the ‘strict prior’ structure did lead to a uniform prior on ρ (Fig. 3, left panel), and this is perhaps more desirable than the resulting prior on ρ for the relaxed prior scenario (Fig. 3, right panel).

MCMC chains were run for 400,000 iterations, saving every 300th iteration, and discarding the first 10,000 samples as burn-in. This resulted in 1300 samples from the posterior. Diagnostics for the chains were generally positive and indicated that convergence had been achieved.

However, the chains specific to ENP gray whales suggest that running more iterations in future analyses may be warranted. For example, even after thinning the chains by sampling only every 300th iteration, there was still a relatively high level of auto-correlation between the samples for ENP gray whales (Fig. 4). The scenario which excluded ENP gray whales from the analysis provided a sensitivity test, and showed that the parameter estimates were largely indifferent to the inclusion of this data-set. The greatest difference in estimates between each data scenario was for θ_{σ} (0.396 vs. 0.295; with and without ENP gray whale data)(Table 2).

The reason for this difference can be seen in the plots of posterior densities for the stock-specific estimates of σ_i (Fig. 5). The calving proportion data for ENP gray whales are relatively variable, and hence excluding this data set resulted in a lower estimate of average variability across stocks. This pattern is also consistent with the estimates of stock-specific variability from the TK approach (Table 1; see Sigma for ENPg).

The stock-specific estimates for ρ_i are shown in Figure 6. While ENP gray whales were estimated to have had the highest inter-annual variability in reproductive success, they are also estimated to have the lowest level of auto-correlation. This may be due to in part to two factors: 1) it seems likely that the auto-correlation and inter-annual variability parameters are correlated and somewhat confounded with respect to estimation (hence higher values for inter-annual variance may lead to lower estimates of auto-correlation, all else being equal), and; 2) ENP gray whales were the only stock in the data sets fitted in these analyses which were not represented by calving interval data. Observed calving intervals are likely buffered to some extent with respect to inter-annual environmental variability, and can be thought of as a moving-average of reproductive success. Whereas, the annual proportion calving might be expected to be more highly variable. Indeed this pattern is evident in the data sets adopted in these analyses for which both types of observations are available (Figs. 1 and 2).

It is interesting to note that all of the estimates of stock specific auto-correlations were positive, with essentially zero posterior probability for values less than zero (Fig. 6). This is in contrast with the TK approach, where several stocks were estimated to have negative auto-correlation parameters (Table 1). The explanation for this is not immediately obvious. It may have to do with the difference between the estimation frameworks, or perhaps the different prior structures on the auto-correlation parameters (e.g. normal in logit space vs. a beta distribution).

The MCMC chains for the posteriors of the hyper-parameters for the JB approach all appeared to have converged. Auto-correlation between samples was essentially non-existent (Fig. 7) and the running quantiles all appeared to have stabilized (Fig. 8). The posteriors for the hyper-paramters related to the extent of inter-annual variability in reproductive rates are shown in Figure 9 for the JB approach. There was very low posterior probability of values of θ_σ (exponentiated) greater than 1.0 (Fig. 9, left panel). This is generally consistent with the results from SC/63/RMP20.

The shape parameters for the hyper-priors on the auto-correlation coefficient were not greatly updated by the data (Fig. 10), but the effect was large enough that the posterior for the auto-correlation parameter for an unknown stock ρ_0 was shifted away from the implicit uniform prior towards more positive values (Fig. 11, left panel). The median of the posterior for ρ_0 was estimated to be 0.288, although there is substantial probability of values lower and (especially) higher than this. The posterior for the inter-annual variation in reproductive success for an unknown stock σ_0 (Fig. 11, right panel) had a median that corresponded with that for the posterior of its hyper-parameter θ_σ (Fig. 9, Table 2), but the resulting posterior for σ_0 was skewed to the right with a longer tail.

DISCUSSION

In summary, given the posterior distributions for the hyper-parameters of interest, it will be possible to generate stochastic population trajectories and investigate the effects of inter-annual variability and auto-correlation in reproductive rates for the MSYR of an unobserved stock, based on the results of this meta-analysis (e.g., Punt, 2011).

However, several considerations should be kept in mind. Firstly, these analyses did not utilize all of the available data. We had selected several of the most informative data-sets from those available, based largely on the criteria of time-series length. It would be interesting to re-run these analyses incorporating all of the available data. It is not clear if this would alter the parameter estimates however, because the data-sets used in these analyses are likely the most informative available.

For example, excluding the ENP gray whale data from this analysis did not appear to change the results to a great degree, and the other excluded data-sets are perhaps even less informative. There are also some inherent uncertainties in certain available data-sets (e.g. interpreting the values for SA right whales, which are themselves based on somewhat complicated modelling assumptions, D.Butterworth *pers. comm.*) that led us to exclude them from the analyses.

In general, one advantage of both of these approaches is the simultaneous analysis of stock-specific data for calving interval and proportion data. This allows for a single parameter estimate for those quantities of interest, as opposed to resulting in a parameter estimate based on proportion data and another based on interval data. There are several details in these approaches which may be improved in future analyses (e.g. assumptions regarding the relationships between process error residuals for different data sets), but in general this modeling framework should serve as a viable alternative for estimating the extent of environmental variability and auto-correlation in reproductive rates for baleen whales. Further, the results between the two approaches explored here are largely consistent with those from Cooke (2011), who utilized a different modeling framework and more of the available data. Therefore, these results may provide some confidence in the robustness of the estimates given different data sets and modeling assumptions.

ACKNOWLEDGEMENTS

Funding for this research was provided by the IWC. Andre Punt is thanked for his input on priors, providing JB with guidance on his WinBUGS code and generally providing helpful advice which improved these analyses. We also thank the scientists that compiled these data and made them available: P. Best, R. Brownell Jr., J. Calambokidis, C. Gabriele, W. Koski, S. Kraus, J. Laake, W. Perryman, C. Ramp, J. Robbins, R. Sears, J. Straley, V. Rowntree.

References

Best, P. B. 1993. Increase rates in severely depleted stocks of baleen whales. ICES J. Mar. Sci. 50: 169-186.

Brandon, J.R. and Kitakado, T. 2010. Preliminary estimates of calving variability for nine stocks of mysticetes. Annex C. Report of the 3rd intersessional workshop on the review of MSYR for baleen whales, Seattle, WA, April 2010. J. Cetacean Res. Manage. 12 (Suppl.): 401-11.

Cooke J.G. 2007. The influence of environmental variability on baleen whale sustainable yield curves. Paper SC/N07/MSYR1 presented to the MSYR workshop, November 2007.

Cooke, J.G. 2011. Bayesian meta-analysis of inter-annual variance and serial correlation in calving rates of baleen whales. Paper SC/63/RMP20, presented to the Scientific Committee of the IWC, June 2011, Tromso, Norway.

IWC. 2011. Report of the 3rd intersessional workshop on the review of MSYR for baleen whales, April 2010. J. Cetacean Res. Manage. 12 (Suppl.): 401-11.

Punt, A.E. 2010. A Revised Bayesian Meta-Analysis for Estimating a Posterior Distribution for the rate of increase for an “Unknown” Stock. SC/62/A10/MSYR2. 9pp. Presented to the 3rd intersessional IWC workshop on the review of MSYR for baleen whales, Seattle, USA, April 2010.

Table 1.

Results for the TK approach. The scenario with the lowest AIC value is highlighted. The scenario in the bottom row (“Random”,”Random”) most closely corresponds with the JB approach.

#Observations	GMh	GSLh	NAr	SEAKh	ENPg
Interval	22	11	25	23	0
Proportions	27	25	29	25	16

Sigma	Rho	loglike	AIC	Sigma					Rho					mean_ sigma	sd_ logsigma	mean_ rho	sd_ logitrho
				GOMh	GSLh	NAr	SEAKh	ENPg	GOMh	GSLh	NAr	SEAKh	ENPg				
Common	Common	-42861.9	85727.8	0.355					0.353								
Common	Different	-42840.1	85692.2	0.524					-0.990	-1.000	0.991	0.992	0.344				
Different	Common	-42839.8	85691.6	0.000	0.000	0.074	0.147	0.720	0.624								
Different	Different	-42838.0	85696.0	0.000	0.375	0.095	0.245	0.612	0.999	-1.000	0.780	0.962	0.411				
Random	Random	-42859.8	85727.6	0.313	0.308	0.305	0.316	0.390	-0.211	0.115	0.833	0.813	0.228	0.347	0.072	0.614	3.013

Table 2.

Results for the JB approach. Medians [lower, upper 95th percentiles] are shown for the posterior distributions of the hyper-parameters. The values for θ_σ have been exponentiated and are in standard space, while those for σ_σ are untransformed and hence represent the standard deviation of the process error variability in log-space. Hence the values of those parameters are directly comparable with the mean_sigma and sd_logsigma estimates in Table 1. The relaxed hyper-prior for ρ scenario crashed JB’s WinBUGS code after repeated attempts, and hence results are not available here.

Dataset	Hyper-prior for rho	$\exp(\theta_\sigma)$	σ_σ	shape1	shape2
with ENP gray	strict prior on rho	0.396 [0.141, 0.816]	0.576 [0.24, 2.13]	1.127 [0.745, 1.7]	0.7674 [0.543, 1.114]
with ENP gray	relaxed prior on rho	Crashed	Crashed	Crashed	Crashed
without ENP gray	strict prior on rho	0.295 [0.075, 0.790]	0.526 [0.20, 2.78]	1.096 [0.763, 1.625]	0.7892 [0.528, 1.156]

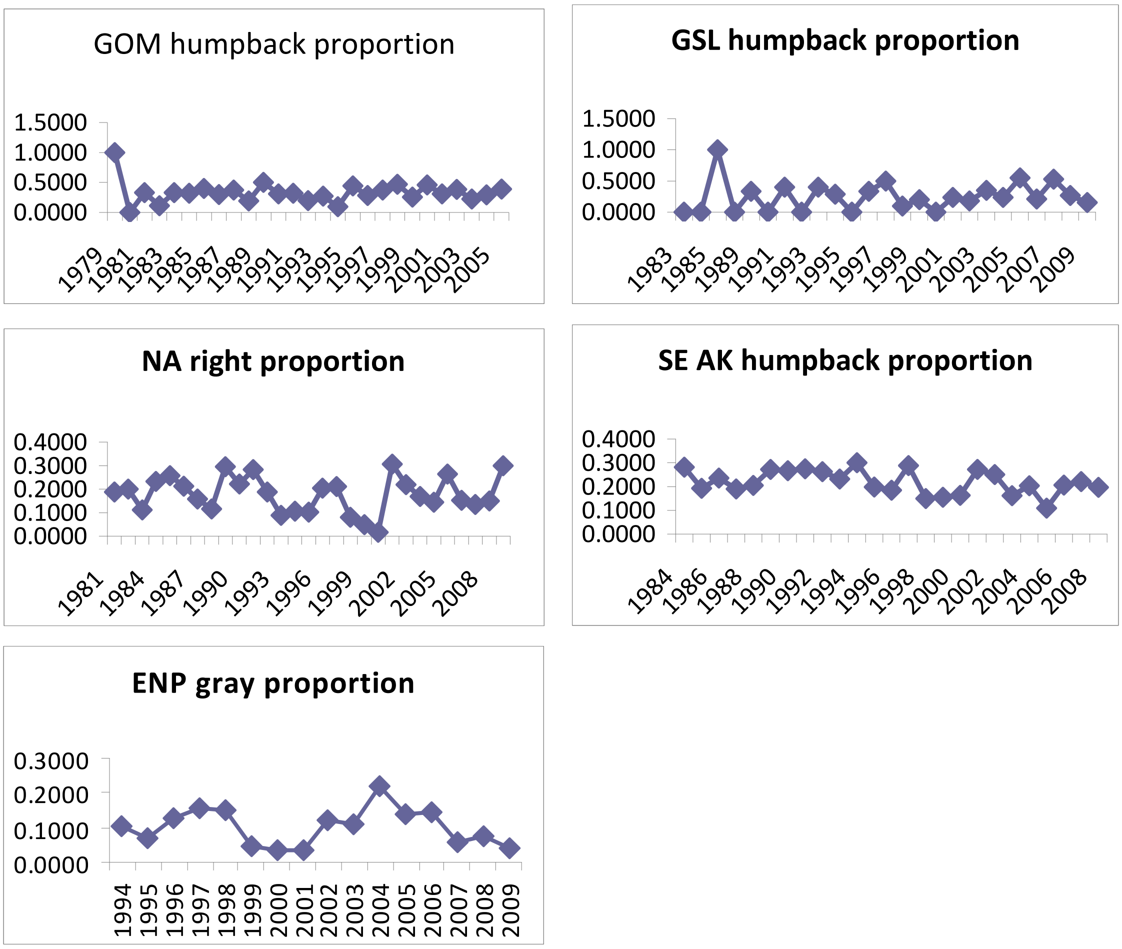

Figure 1. Annual proportion of mature females calving are shown.

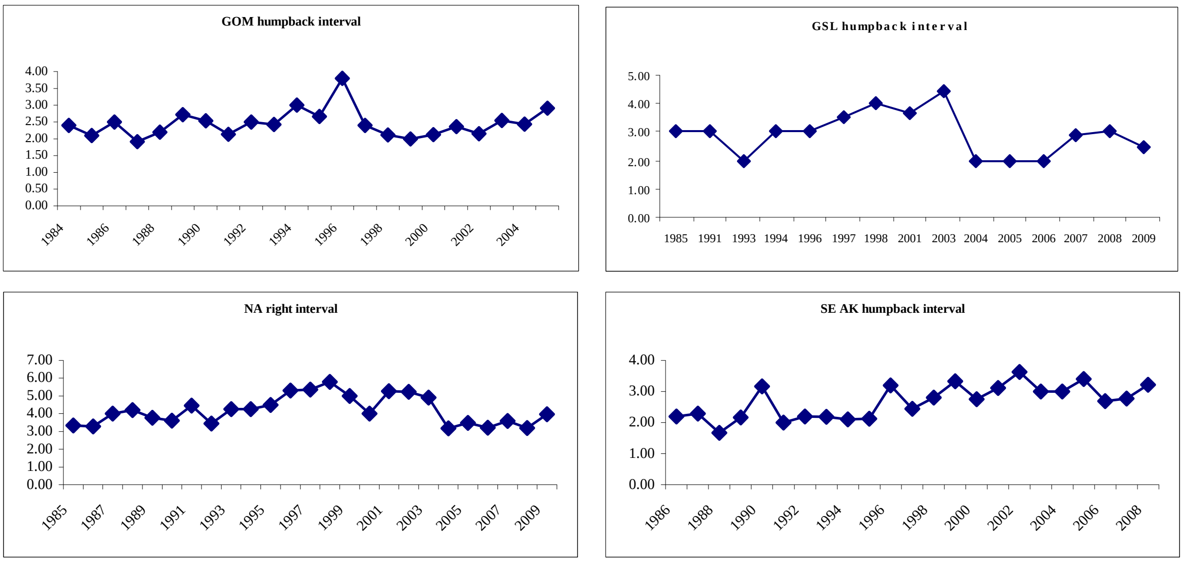

Figure 2. Annual calving intervals are shown.

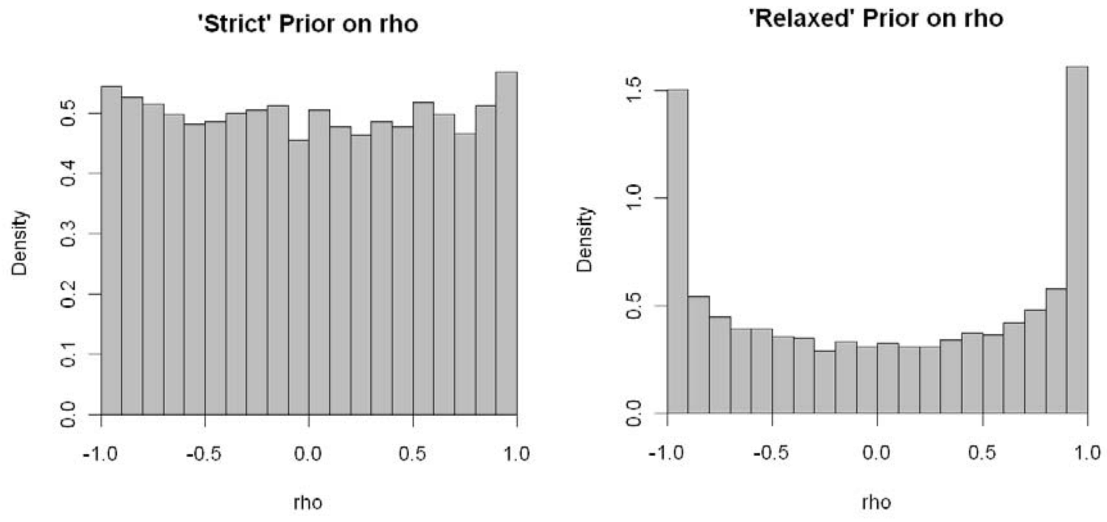

Figure 3. Ten thousand samples from the prior distribution of ρ, with the standard deviation in log-space for the hyper-priors (shape1 and shape2) of the Beta distribution equal to 0.20 (left panel) vs. 1.0 (right panel).

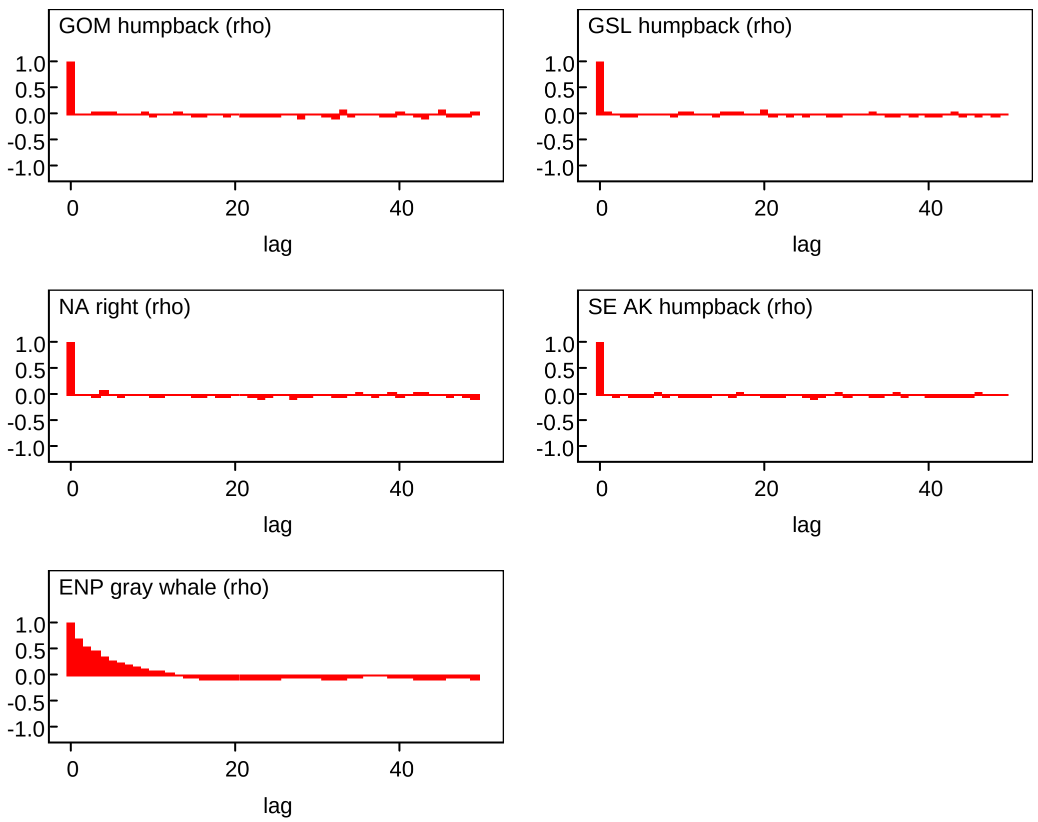

Figure 4. Auto-correlations are shown between MCMC samples for the parameter ρ_i. These plots are indicative of the results for σ_i as well. The diagnostic results for the scenario excluding ENP gray whales were also very similar for the remaining stocks. That is, only ENP gray whales showed some signs of benefiting from running longer chains in future analyses.

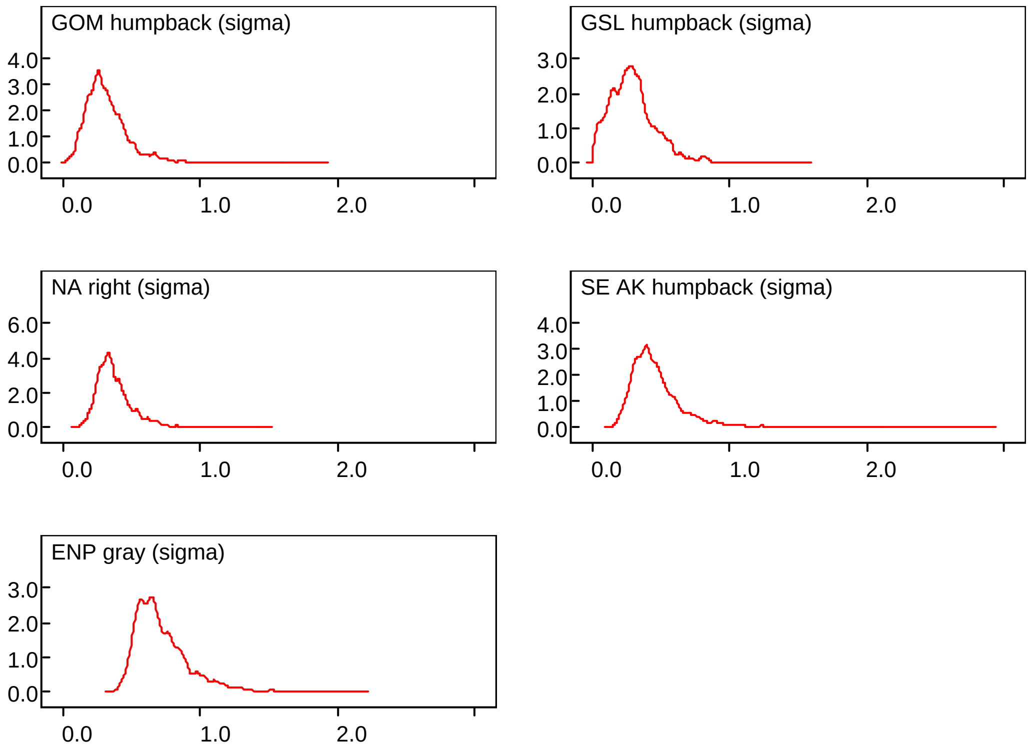

Figure 5. Stock-specific estimates of σ_i are shown for the JB approach.

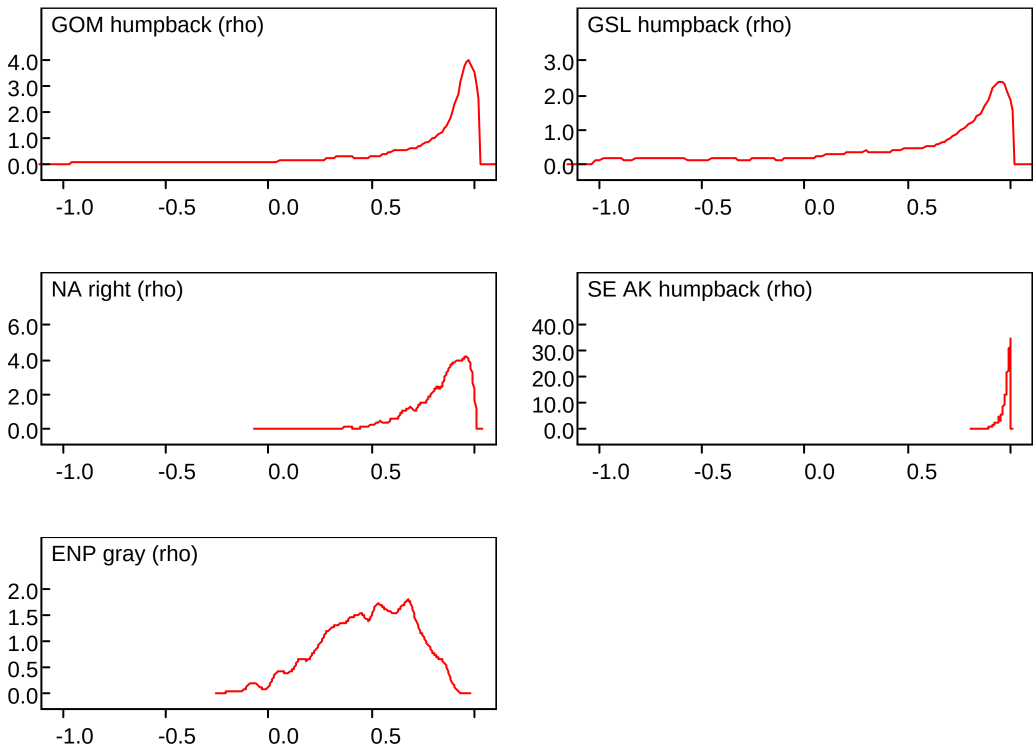

Figure 6. Stock-specific estimates of ρ_i are shown for the JB approach.

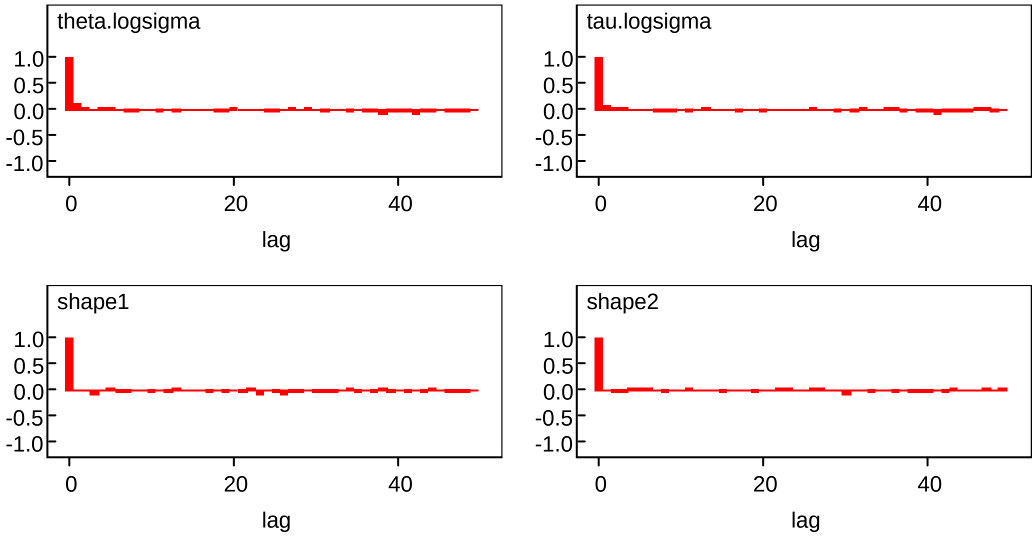

Figure 7. Auto-correlations between MCMC samples for the hyper-parameters of the model. Note that theta.logsigma is in log-space and tau.logsigma is in units of precision (1/variance).

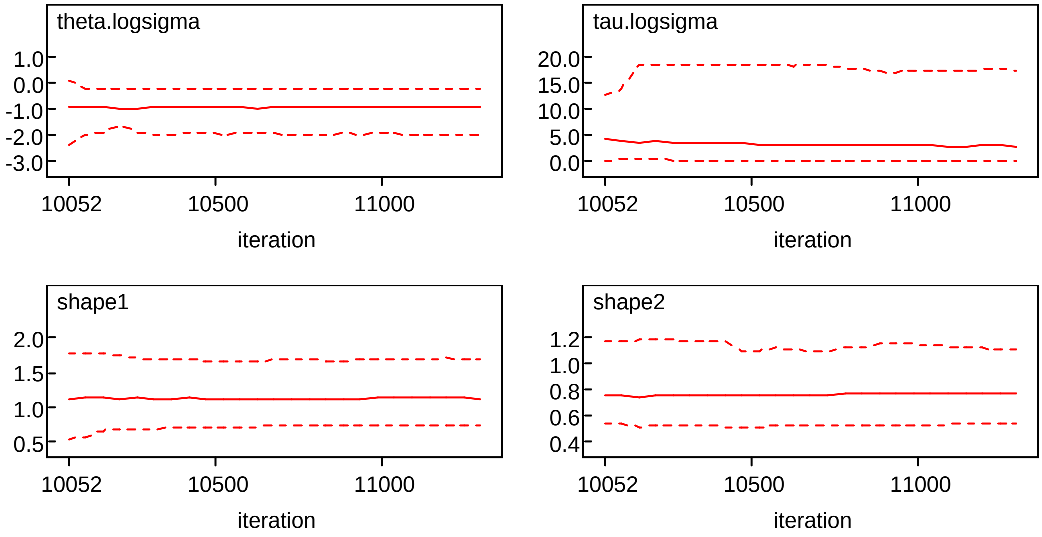

Figure 8. Running quantiles for the hyper-parameters of the model. Note that theta.logsigma is in log-space and tau.logsigma is in units of precision (1/variance).

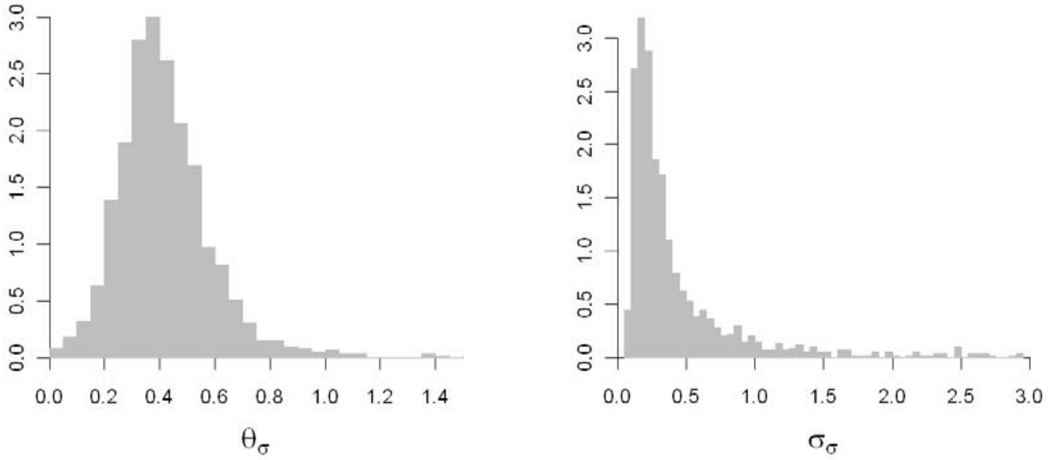

Figure 9. Posterior distributions are shown for the hyper-parameters θ_σ (exponentiated) and σ_σ (in log-space), i.e. the expectation and standard deviation for the extent of environmental process error σ in reproductive rates. These estimates correspond to the scenario with strict hyper-priors for ρ, and are conditioned on the dataset including ENP gray whales.

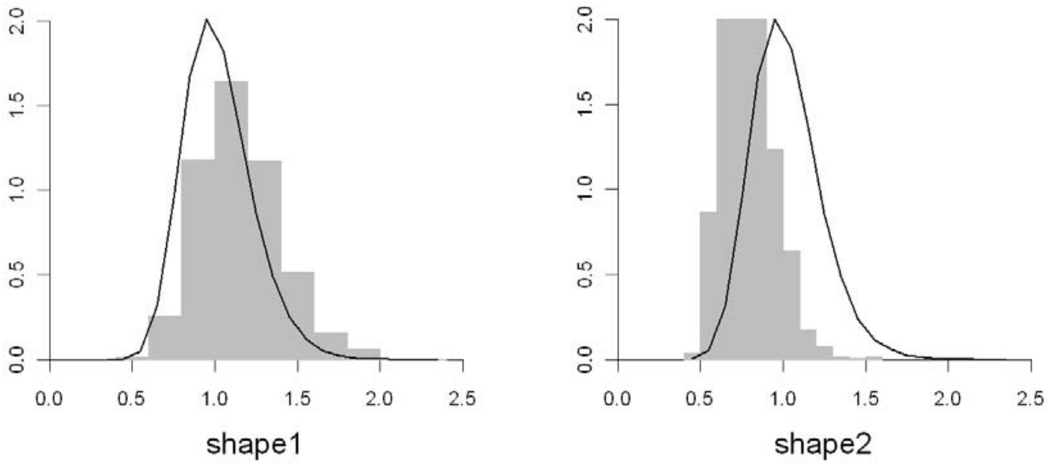

Figure 10. Prior (black lines) and posterior (gray areas) distributions are shown for the 'shape1' and 'shape2' hyper-priors on ρ.

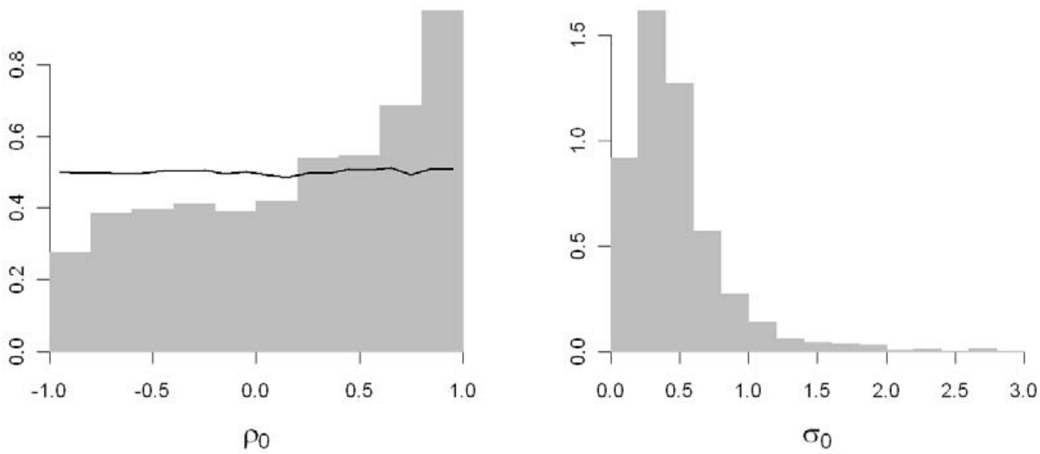

Figure 11. Prior (black line) and Posterior (gray area) distributions are shown for the auto-correlation parameter ρ_0 and extent of process error σ_0 for an unknown stock.