Fisheries New Zealand
Tini a Tangaroa

Review of the recruitment dynamic used in age-based stock assessment models

New Zealand Fisheries Assessment Report 2021/43

C. Marsh,
N. Sibanda,
A. Dunn,
I. Doonan

ISSN 1179-5352 (online)
ISBN 978-1-99-100960-9 (online)

August 2021

New Zealand Government

Requests for further copies should be directed to:

Publications Logistics Officer
Ministry for Primary Industries
PO Box 2526
WELLINGTON 6140

Email: brand@mpi.govt.nz
Telephone: 0800 00 83 33
Facsimile: 04-894 0300

This publication is also available on the Ministry for Primary Industries websites at:
http://www.mpi.govt.nz/news-and-resources/publications
http://fs.fish.govt.nz go to Document library/Research reports

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Marsh, C.; Sibanda, N.; Dunn, A.; Doonan, I. (2021). Review of the recruitment dynamic used in age-based stock assessment models.

***New Zealand Fisheries Assessment Report 2021/43.* 19 p.**

This report addresses how year class parameters are described and estimated in age-based stock assessments within New Zealand and overseas. We have taken a range of assumptions commonly made with respect to year class strength parameters, including whether it is dealt with in log-space versus natural space, the choice of prior, and whether to include a penalty and/or constraining parameters to have some mean value.

This study applied alternative descriptions, combinations of priors, and penalties of year class strength parameters and applied these to a range of stock assessments based on New Zealand species. This work attempted to review and clarify the assumptions and reasoning behind each parameterisation and address whether there were any consequences of choosing one assumption over another.

Results suggest that some assessments were quite sensitive to the recruitment variance prior value, and that more sensitivities should be conducted on this parameter than is currently done according to the New Zealand Fisheries Assessment Reports reviewed during this project. We explored the inclusion of hierarchical models with the year class parameters and suggest that this method be explored further. If hierarchical models are not a practical option, then we suggest assessment scientists add a constraint to the stock recruitment residuals, either via penalties available within the AD Model Builder software (sum of squares error straying from an expectation), or standardise as is done by CASAL. We found that two of three assessments (orange roughy and ling) investigated had varied estimates of the initial state (B_0) when constraints on year class strength parameters were not included. When constraints were included for these assessments the influence of *a priori* recruitment variability was negligible. For the remaining assessment (southern blue whiting), estimates of the initial state (B_0) and current biomass were sensitive to a priori recruitment variability assumptions across all parameterisations investigated. These results show that for some assessments, constraints are necessary for key management parameter identifiability. We hypothesise that this is due to some models being close to (if not already) over-parameterised.

1. INTRODUCTION

The objective of this research was to review and investigate the assumptions surrounding model parameterisations and prior choices for year class strengths in statistical catch-at-age models. We focused on tactical age-based stock assessment models, which are used to inform management advice for fish stocks. There are a range of generalised stock assessment packages available for stock assessment scientists to model fish stock dynamics. We considered a range of generalised age-based stock assessment packages and found two general formulations describing the error structure associated with the recruitment dynamic. The generalised modelling packages reviewed are listed with references in Table 1.

Table 1: A table of generalised age structured packages reviewed in this project.

Software	Reference
CASAL/Casal2	Bull et al. (2012)
Stock Synthesis	Methot & Wetzel (2013)
MULTIFAN-CL	Fournier et al. (1998)
ASAP	Legault & Restrepo (1999)

We focused on the parameterisations that assumed mean unbiased lognormal stock recruitment residuals. We did not argue the case for that assumption here, rather the argument is outlined elsewhere; for example, by Haddon (2011) and Hilborn & Walters (1992). These sources argue lognormal stock recruitment residuals should be the default assumption in the absence of information and, although some stocks do have pre-recruit surveys (e.g., hoki, Ballara & O'Driscoll 2017), this is not the norm.

We collated and tested varying assumptions for three assessment models which were based on the following stocks: ORH 3B, SBW 6I, and LIN 3&4, to investigate any influence on model fits, estimation, and derived quantities from the assumptions.

This report follows on from Francis & Fu (2015), and a re-occurring discussion in technical deepwater fisheries working groups in New Zealand (A. Dunn, pers. comm.) around the choice of a year class strength prior and prior parameter values. This project investigated the performance of alternative parameterisations, explored current status quo assumptions, and considered implementing hierarchical models for the recruitment dynamic.

2. METHODS

A common formulation of the recruitment dynamic in stock assessment models in fisheries is:

$$R_y = R_0 \times SR\left(SSB_{y-lag}, \theta\right) \times YCS_y$$

Where R_y is the number of recruits that enter the stock in year *y*, $SR\left(SSB_{y-lag}, \theta\right)$ is a term capturing the reproductive capacity of the population, in this case denoted as a simple stock recruitment relationship of Spawning Stock Biomass (*SSB*) from the spawning year $(y - lag)$ and other parameters θ, such as steepness (*h*) and a measure of virgin biomass (B_0). YCS_y is the year class strength multiplier applied in year *y*, also called a stock recruitment residual (Haddon 2011). The recruitment dynamic reflects, in a highly simplified form, the effective fecundity of the spawning stock, and survivorship of pre-recruits, to determine recruitment to the population. This is a general equation, but by no means covers the full range of possibilities. There is a range of hypotheses for describing recruitment variability, accompanied by attempted modelling. We point readers to Anderson (1988) for a review of recruitment hypotheses and corresponding studies. The focus of the present study was on understanding alternative parameterisations and estimation techniques with regard to the YCS_y parameter. For all models here, we assume that the $SR\left(SSB_{y-lag}, \theta\right)$ follows the Beverton-Holt stock recruitment relationship.

There are two common ways of parameterising the lognormal prior distribution of year class strength (YCS_y) when fitting stock assessment models in fisheries. Let YCS_y represent the YCS for year y. The two parameterisations used are:

1. $YCS_y \sim LN(\mu, \sigma_R^2)$, with $\mu = -\frac{1}{2}\sigma_R^2$ chosen so that $E\left[YCS_y\right] = 1$,
2. $YCS_y = e^{\varepsilon_y - \frac{1}{2}\sigma_R^2}$, where $\varepsilon_y \sim N(0, \sigma_R^2)$

Option 1 is the only approach available in the CASAL package, whereas option 2 was the favoured parameterisation in all other stock assessment packages considered in this study. There were variations on option 2, such as Stock Synthesis having an extra ramping function following (Methot & Taylor 2011) and ASAP excluding the $-\frac{1}{2}\sigma_R^2$ term all together. We found that, from a Bayesian perspective, in option 2 if the you apply a normal prior on the ε_y then you need to take into account the change of variable and add a Jacobian. If not, then the resulting YCS_y year class strength parameters will not be *a priori* lognormal and will differ by the Jacobian $\left(\left|\frac{1}{e^{\varepsilon_y - \frac{1}{2}\sigma_R^2}}\right|\right)$ on the negative log likelihood (see Appendix A for the proof of this). It is unclear from reading user manuals that this in fact is applied across packages that formulate the YCS with this parameterisation. We ran models with and without this Jacobian to see if it made any material difference to model fits and quantities for assessments chosen in this investigation.

A common constraint in stock assessment models is for mean year class strengths $\overline{YCS} = 1$, over a specified period. In CASAL, this constraint is applied using statistical methods such as priors and penalties to encourage model parameters to fall within some *a priori* space, or through standardisation. There are two standardisation techniques available in CASAL, the Francis method and Haist method. Francis & Fu (2015) found issues with the Francis method and so from here on standardisation refers to the Haist method. The Haist standardisation takes a vector of 'free' estimable year class parameters denoted as Y_y and transforms them to the year class multipliers denoted as YCS_y via the following equation:

$$YCS_y = \begin{cases} \dfrac{Y_y}{\bar{Y}_{y \in S}}, & \text{for } y \in S \\ Y_y & , \text{for } y \notin S \end{cases}$$

where S is the set of years over which the standardisation is applied. Two concerns were expressed by the Francis & Fu (2015) paper about this approach; these were parameter ambiguity and unclear prior definition on the resulting YCS_y. We considered the theory behind this technique to find out what the priors of YCS_y if priors are applied to the 'free' estimable parameters Y_y, as currently implemented in CASAL.

Other generalised packages investigated, which are based on the AD Model Builder software (ADMB, Fournier et al. 2012), use a penalty to enforce the constraint $\sum \varepsilon_y = 0$. The constraint on recruitment deviations did not appear in any manual or usage case that we found. From examining the ADMB source code, the constraint ($\sum \varepsilon_y = 0$) seems to be enforced by the contribution to the negative log likelihood given below:

$$-ll = 10000\ \bar{\varepsilon}^2$$

Where $\bar{\varepsilon}$ is the mean value for the vector of recruitment deviations[1].

Three stock assessment models for three stocks were simplified (similar assumptions were used but the test models were not identical to the assessments models) and used to explore the consequence of the assumptions laid out above. These stocks were chosen due to the varying life histories and quantity of data. The stocks were southern blue whiting (SBW) (*Micromesistius australis*) on the Campbell Island rise (Dunn & Hanchet 2015), orange roughy (ORH) (*Hoplostethus atlanticus*) on the Northwest Chatham Rise (Cordue 2014), and ling (LIN) (*Genypterus blacodes*) on the Chatham Rise (McGregor 2015).

For each assessment model, we conducted runs with the YCS parameterisation without standardisation, with standardisation, and with recruitment deviations. For each of these parameterisations we considered a range of sigma (σ_R) values for the prior choice to see how robust these assessments were to these assumptions.

The choice of σ_R, 0.2, 0.5, 1, 1.5, and nearly uniform (Cordue 2014) were somewhat arbitrary but were broadly based on summaries of worldwide recruitment models (Myers et al.1995), which gave lower quartile = 0.48, median = 0.67, upper quartile = 1, and maximum = 3.11 values. Each model was estimated using both the Maximum *a Posterior* Density (MPD) and Markov chain Monte Carlo (MCMC) estimates to see if these assumptions influenced fits and model quantities, across a range of assessments. The above permutations are visualised in Figure 1.

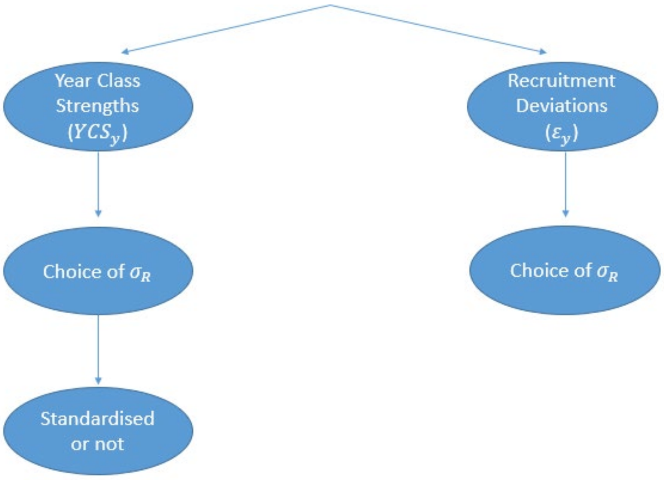

Figure 1: A diagram showing the different parameterisations run for the stock recruitment residuals.

All model parameterisations were fitted to observed data collected from surveys and age sampling. Fitting to real data meant that we were searching for major differences or deviations among assumptions, rather than attempting to find the 'best' parameterisation because, of course, this is unknown.

The last set of model configurations run for the three assessment models were Bayesian hierarchal models (Congdon 2010), which used hyper-priors to estimate hyper-parameters which are usually fixed and can cause much debate in stock assessment working groups (A. Dunn, Fisheries New Zealand, pers. comm.). The current relationship between priors and year class strength parameters can be best described by a Directed Acyclic Graph (DAG) (Figure 2).

[1] Found in source code at https://github.com/admb-project/admb/blob/master/src/nh99/model5.cpp

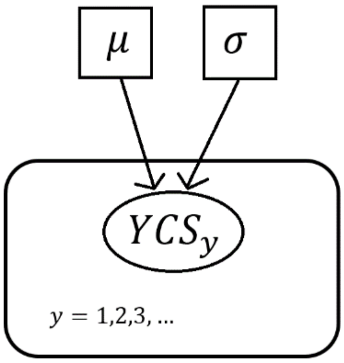

Figure 2: A directed acyclic graph demonstrating the relationship between priors and estimated parameters for a current Bayesian model, where squares indicate fixed constants and circles represent estimated parameters.

With Bayesian hierarchical models you can free up the traditionally constant prior parameters (Figure 3).

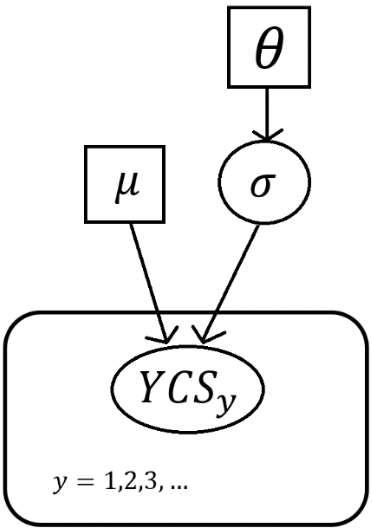

Figure 2: A directed acyclic graph demonstrating the relationship between priors and estimated parameters for a Bayesian hierarchical model. Squares indicate fixed constants and circles represent estimated parameters.

In Figure 3, θ represents hyper-parameters that define the hyper-prior for σ. The variance parameter (σ) can be estimated with the mean fixed, due to the constraint that $E[YCS_y] = 1$, $\mu = -\frac{\sigma^2}{2}$. The benefit of the Bayesian hierarchical model is that instead of conducting the usual discrete sensitivity model runs with alternative fixed prior values, uncertainty (which is usually ignored) can be incorporated in model outputs via parameter estimation of σ_R. Bayesian hierarchical models were applied for the three assessment models, with three hyper-priors on σ_R as expressed below:

$$YCS_y|\sigma_R^2, \theta \sim LN(\mu, \sigma_R^2)$$

$$\sigma_R^2|\ \theta \sim dist(\theta)$$

$$dist(\theta)\begin{cases} uniform(0.1,10) \\ N\ (\mu_{mpd}, 0.4) \\ N\ (\mu_{mpd}, 0.2) \end{cases}$$

For all models run using MCMC, convergence was visually checked by trace plots on likelihood components and important productivity parameters. If convergence was deemed satisfactory, model outputs were displayed and reported.

3. RESULTS

We start by presenting theoretical results of two scenarios and discuss the implications. Assuming a normal prior on recruitment deviations did not imply the resulting year class strength multipliers followed a lognormal distribution. They differed by a Jacobian from the change in variable (see Appendix A for the derivation of this result). The implication is that the models did not match the theory assumed. As pointed out earlier, there has been theoretical justification for recruitment being lognormal, but many models do not mathematically represent this. Figure 4 shows the contribution to the negative log-likelihood when the Jacobian is excluded and included.

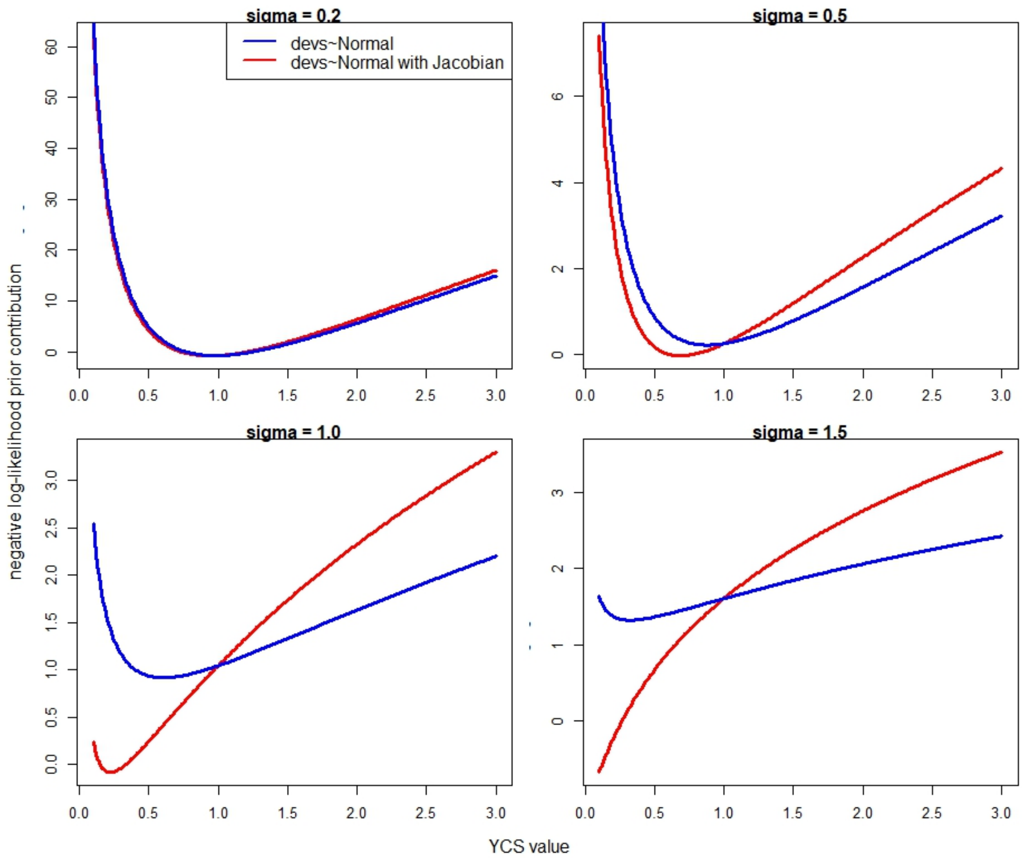

Figure 3: Comparison of negative log-likelihood prior contributions for a given YCS value, when the Jacobian is included (red) and excluded (blue).

We examined the implications of including the Jacobian, or not, and found that when the penalty for recruitment deviations to be *a priori* close to one was applied, as is commonly done in ADMB models, it had no effect. With no penalty constraint, it did make a difference (Figure 5); this is something to consider if using this parameterisation.

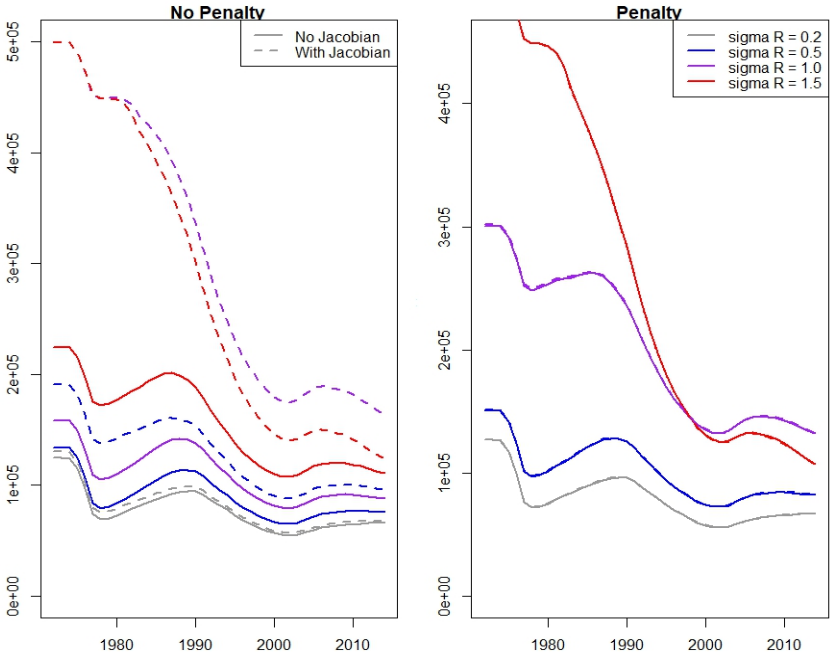

Figure 4: The implications of adding penalties to MPD estimates of *SSB* (t) and whether or not the Jacobian has been included in the objective function (results shown are for the ling based assessment).

The effect of standardisation was also investigated, and it was found that this also affected the underlying distribution (see Appendix B for more information). Standardisation transformed each free year class parameter (Y_y), which was initially assumed independent from a lognormal distribution, to the resulting year class multiplier variables (YCS_y) which are the ratio of two random variables, i.e., $\left(YCS_y = {}^{Y_y}/_{\bar{Y}}\right)$. This changes the interpretation of the resulting prior definition on the final YCS_y which is unlikely to be lognormal (Appendix B).

As well as considering the theory, we applied a combination of these parameterisations (see Figure 1) on the three assessments to see if, in practice, these assumptions changed model fits and outputs in both MPD and MCMC estimation. The following figures show spawning biomass trajectories for each stock and scenario.

The MPD results (Figures 6–8) varied among species and parameterisations. For the ORH-like assessments there were similar patterns across the three parameterisations. Standardised year class strengths were consistent regardless of the choice of σ_R. For non-standardised runs, larger σ_R gave a larger B_0. For the LIN-like assessment, standardised year class strengths were consistent regardless of choice of σ_R. For non-standardised runs, larger σ_R again gave a larger B_0. For the SBW-like assessment, the beginning and end of the *SSB* trajectories were highly variable by choice of σ_R across all methods. Trajectories with large recruitment variances that had extremely large biomass would, in a normal assessment context, be deemed to have not converged; we have added them here for completeness.

Our initial hypothesis was that MPD estimation would be more influenced by these assumptions, because we were viewing a single point on the likelihood surface and considering the full posterior distribution would show these assumptions to be less important. Overall, there was no marked difference between MPD results and median MCMC biomass trajectories. Our results suggested that some assessments were quite sensitive to the recruitment variance prior value, and constraints helped with the estimation of B_0 for some assessments.

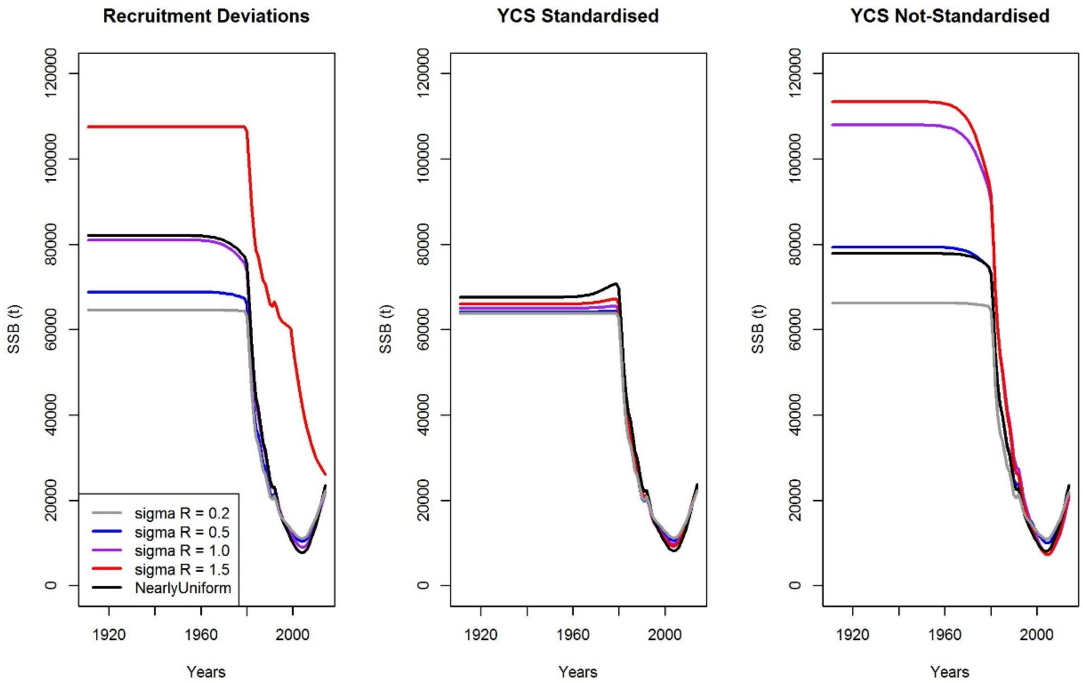

Figure 5: ***SSB* trajectories for the ORH-like assessment after an MPD estimation for each scenario.**

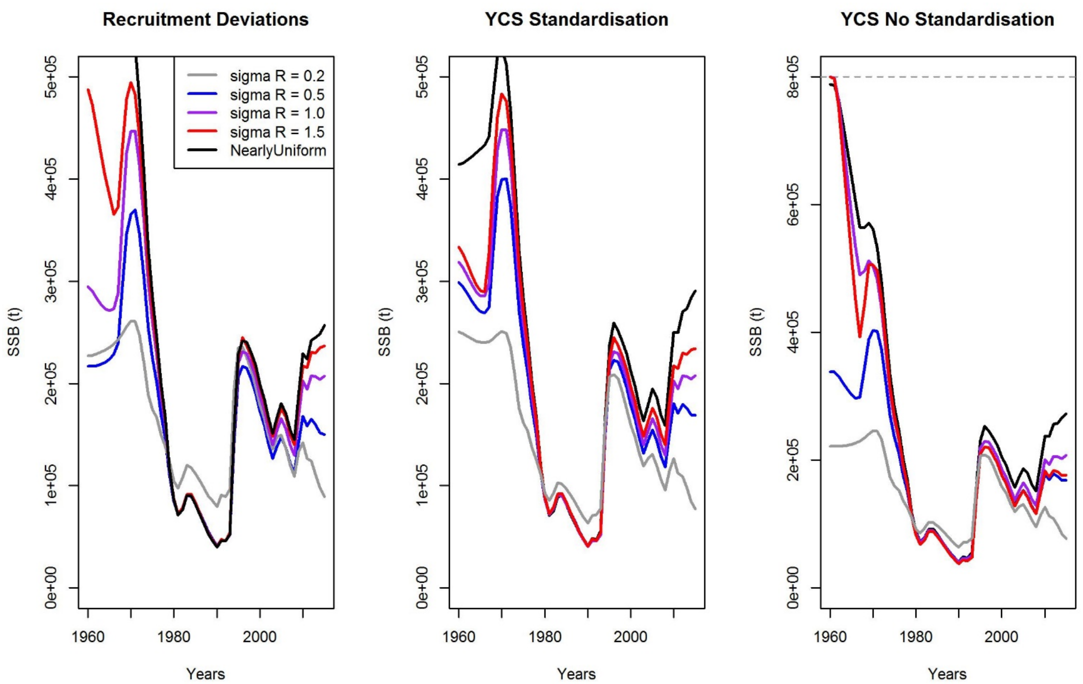

Figure 6: ***SSB* trajectories for the SBW-like assessment after an MPD estimation for each scenario.**

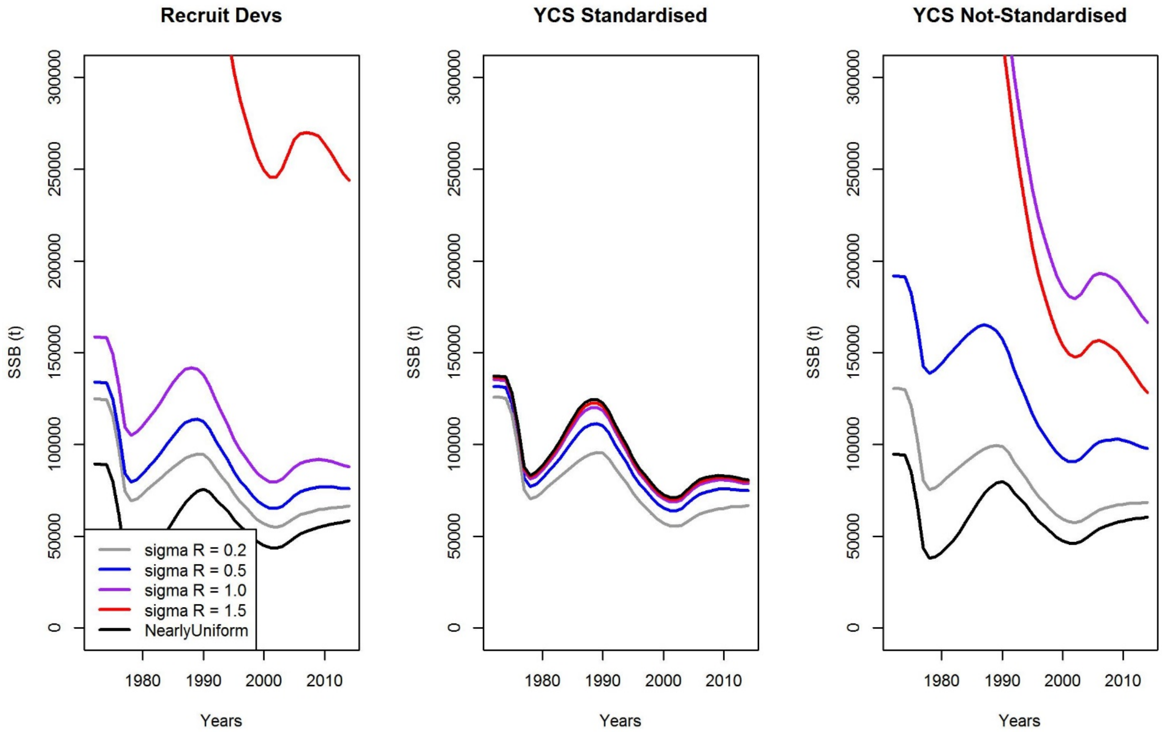

Figure 7: ***SSB* trajectories for the LIN-like assessment after an MPD estimation for each scenario.**

The results from the MCMC runs (Figures 9–11) suggested that standardisation was more robust to assumptions of σ_R relative to the other parameterisations. We also saw that the models broke down in MCMC performance (did not converge) for non-standardised methods when σ_R increased, suggesting not constraining this vector of parameters could lead to some parameters, such initial state parameters (B_0 and R_0), becoming unidentified.

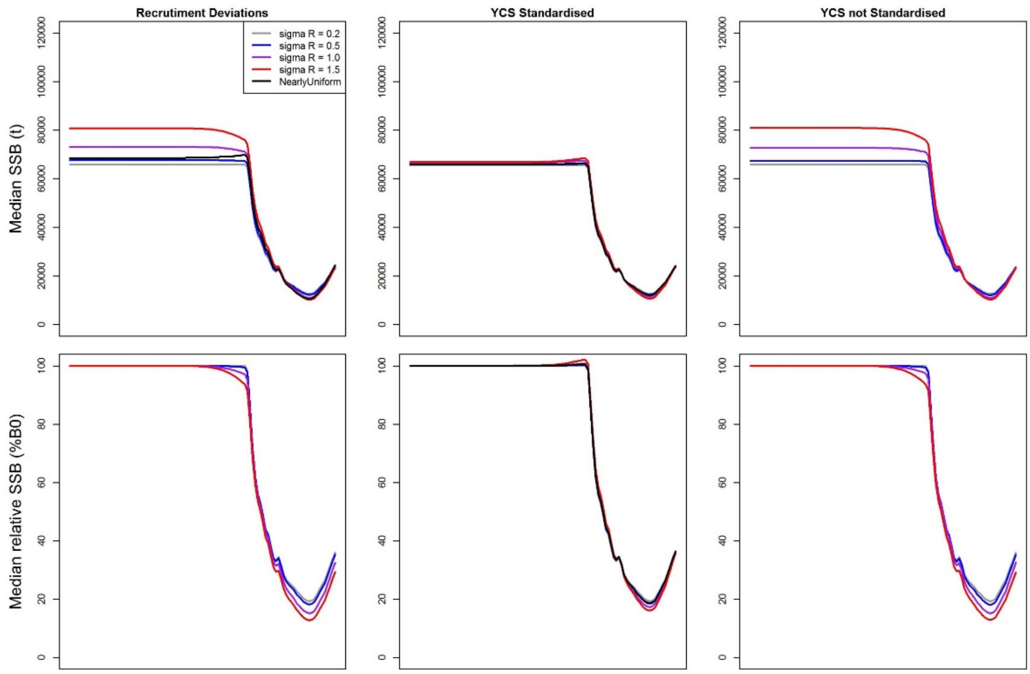

Figure 8: **Absolute and relative median *SSB* trajectories for converged MCMC runs for the ORH-like assessment.**

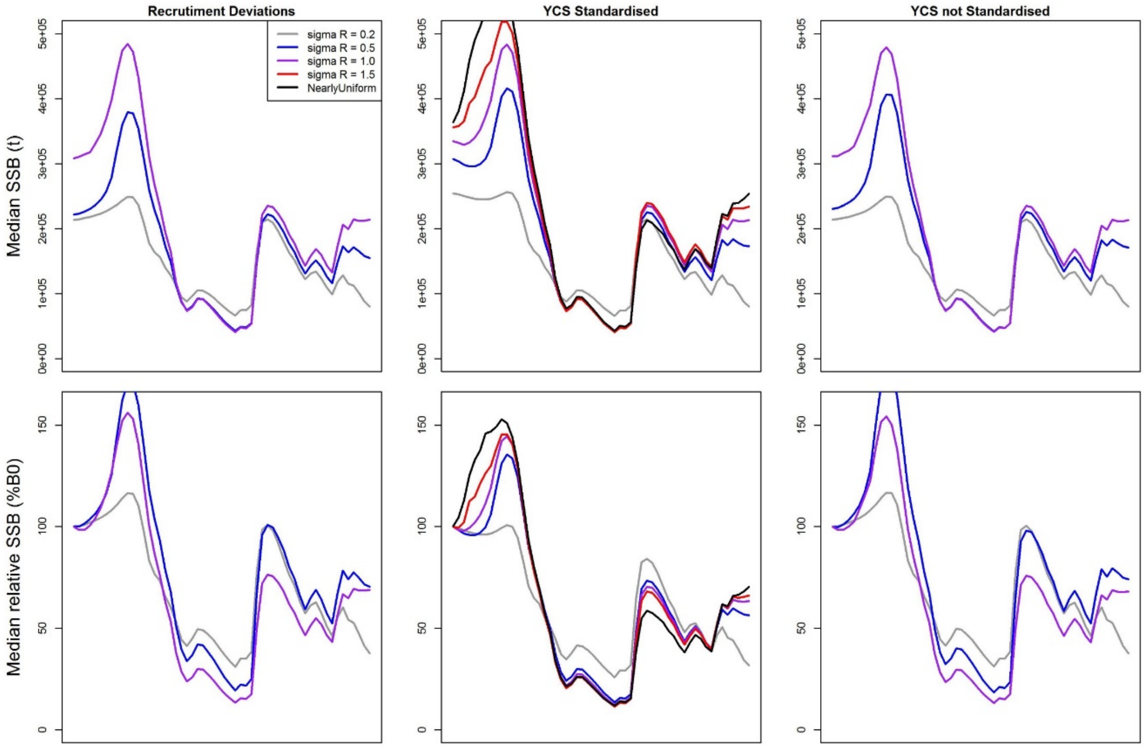

Figure 9: Absolute and relative median *SSB* trajectories for converged MCMC runs for the SBW-like assessment.

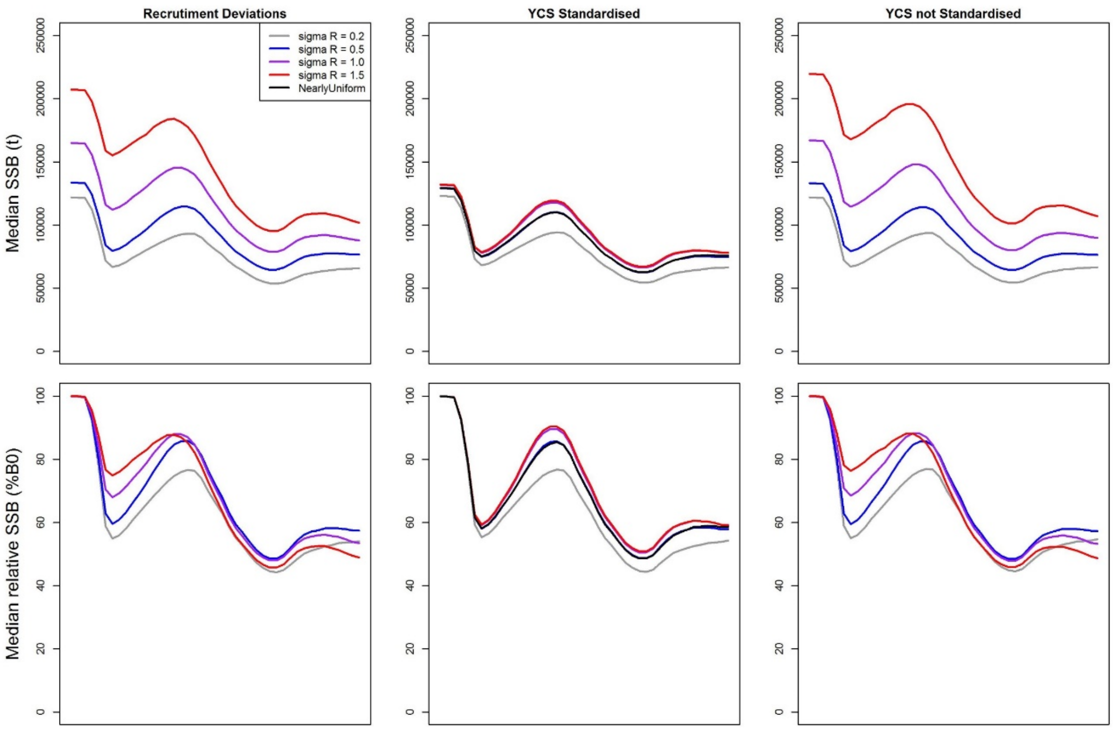

Figure 10: Absolute and relative median *SSB* trajectories for converged MCMC runs for the LIN-like assessment.

The last sequence of runs implemented hierarchical models. Instead of fixing σ_R we estimated it, assuming hyper-priors on that parameter.

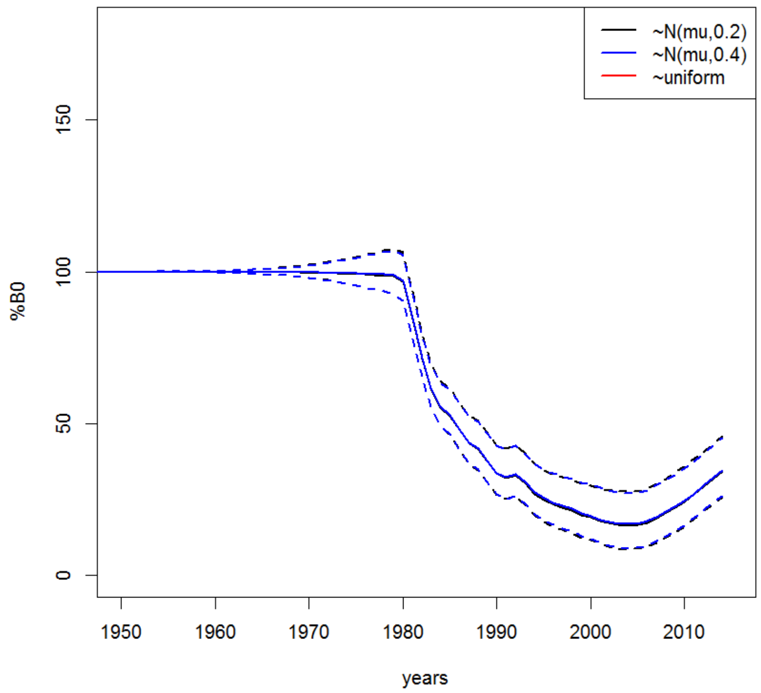

Figure 11: Relative *SSB* trajectories with 95% credible intervals for the ORH-like assessment using different hyper-priors on $\boldsymbol{\sigma_R}$.

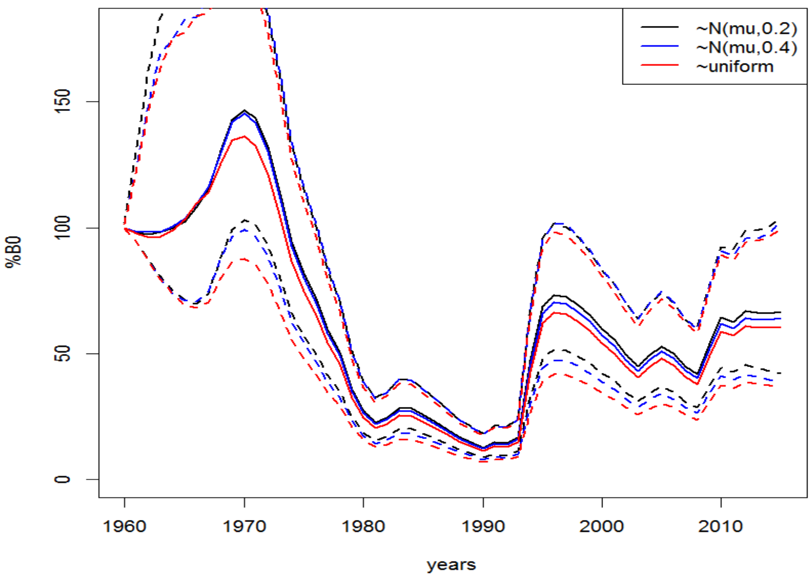

Figure 13: Relative *SSB* trajectories with 95% credible intervals for the SBW-like assessment using different hyper-priors on $\boldsymbol{\sigma_R}$.

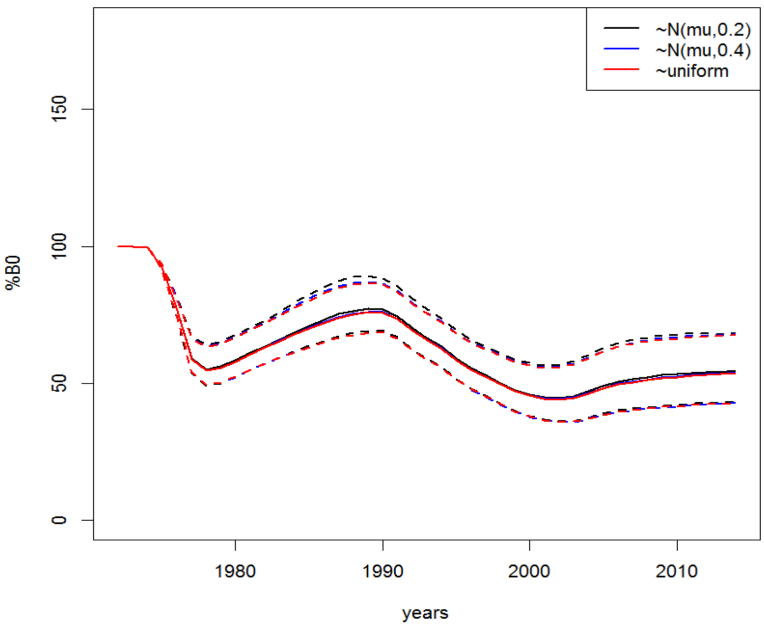

Figure 14: Relative *SSB* trajectories with 95% credible for the LIN-like assessment using different hyper-priors on $\boldsymbol{\sigma_R}$.

From the model outputs shown in Figures 12–14, we see that the outputs were not sensitive to the choice in hyper-priors. Although convergence was not ideal, and more work is needed before this class of models should be implemented in practice using Casal2, the preliminary results did estimate σ_R consistent with *a priori* expectations regarding the amount of recruitment variability, i.e., the LIN-like assessment estimated a lower σ_R whereas for the SBW-like assessment the σ_R was high, which was consistent with current hypotheses.

4. CONCLUSIONS

For the majority of model results for both MPD and MCMC estimation with varying σ_R, we found there were no obvious benefits to parameterisation of year class strengths in log space as recruitment deviations, over the current implementation of year class strengths by Casal2. Standardisation was more robust to assumptions about this parameter, which we considered a positive benefit of the method. We also found that in some model formulations the theory was not consistent with the model setup, but that often this did not change the resulting model outputs in an obvious way. For the SBW-like assessment, it was observed that the choice of σ_R was more influential than the parameterisation choice, so we recommend further investigation for this assessment. We also began to examine hierarchical models, which we suggest should be further tested in assessments.

Future work and recommendations include: this work showed that sensitivity runs focused on the year class strength parameter should be conducted more commonly than is currently done, or if it is done, commented on in the assessment report. We investigated the use of hierarchical models as the natural next step to estimate the σ_R parameter. We have just scratched the surface on the use of hierarchical models in practice on the assessments currently being implemented. Further work directly related to the Casal2 assessment package should include the incorporation of hierarchical model structure at the MPD level. Currently this is ignored, and this has been found to lead to follow-on effects when using the Metropolis-Hastings MCMC algorithm with the proposed distribution based on the covariance of the MPD, which is currently the only approach implemented by Casal2. For users to confidently use this class of model, this would need to be addressed. Theoretical work that we recommend includes the question: under what conditions do hierarchical models estimate the correct hyper-parameters? One would want to examine this question with respect to data quality, including the sample sizes and ageing error that are common in fish stock assessments. Also, how influential are hyper-parameters? Are we just propagating the uncertainty into another tier of model structure? One could argue that if the choice of hyper-parameters was influential, the problem has only been shifted up a level and not addressed.

We would also like to see continued development of constrained estimation that is statistically consistent with the lognormal assumptions that are commonly made *a priori*.

5. ACKNOWLEDGMENTS

This work was funded by the Fisheries New Zealand project SCM2016-01. Our thanks to Pamela Mace (Fisheries New Zealand) and the Fisheries New Zealand Stock Assessment Methods Working Group for valuable assistance and feedback during the project.

6. REFERENCES

Anderson, J.T. (1988). A review of size dependant survival during pre-recruit stages of fishes in relation to recruitment. *Journal of Northwest Atlantic Fishery Science 8*: 55–65.

Asmussen, S.; Rojas-Nandayapa. L. (2008). Asymptotics of sums of lognormal random variables with gaussian copula. *Statistics & Probability Letters 78(16)*: 2709–2714.

Ballara, S.; O'Driscoll, R. (2017). Catches, size, and age structure of the 2015–16 hoki fishery, and a summary of input data used for the 2017 stock assessment. *New Zealand Fisheries Assessment Report 2017/31*. 119 p.

Bull, B.; Francis, R.I.C.C.; Dunn, A.; McKenzie, A.; Gilbert, D.; Smith, M. H.; Fu, D. (2012). CASAL (C++ algorithmic stock assessment laboratory): CASAL user manual v2.30-2012/03/21. *NIWA Technical Report 135*. 280 p.

Congdon, P.D. (2010). *Applied Bayesian hierarchical methods*. Chapman and Hall. 604 p.

Cordue, P. (2014). The 2014 orange roughy stock assessments. *New Zealand Fisheries Assessment Report 2014/50*. 135 p.

Crow, E. L.; Shimizu, K. (1988). *Lognormal distributions: Theory and applications.* New York: M. Dekker. 387 p.

Dunn, A.; Hanchet, S. (2015). Southern blue whiting (*Micromesistius australis*) stock assessment for the bounty platform up to, and including, the 2011 season. *New Zealand Fisheries Assessment Report 2014/55*. 30 p.

Fournier, D.A.; Hampton, J.; Sibert, J.R. (1998). MULTIFAN-CL: A length-based, age-structured model for fisheries stock assessment, with application to south pacific albacore, *Thunnus alalunga*. *Canadian Journal of Fisheries and Aquatic Sciences 55(9)*: 2105–2116.

Fournier, D.A.; Skaug, H.J.; Ancheta, J.; Ianelli, J.; Magnusson, A.; Maunder, M.N.; Sibert, J. (2012). AD model builder: Using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software 27(2)*: 233–249.

Francis, R.I.C.C.; Fu, D, (2015). Rationale for year-class strength priors. *New Zealand Fisheries Assessment Report 2015/62*. 23 p.

Gao, X.; Xu, H.; Ye, D. (2009). Asymptotic behaviour of tail density for sum of correlated lognormal variables. *International Journal of Mathematics and Mathematical Sciences*. Article ID 630857. https://doi.org/10.1155/2009/630857.

Haddon, M. (2011). *Modelling and quantitative methods in fisheries*. Chapman & Hall/CRC. 465 p.

Hilborn, R.; Walters, C. J. (1992). *Quantitative fisheries stock assessment: Choice, dynamics and uncertainty.* Springer USA. 570 p.

Legault, C.; Restrepo, V. R. (1999). A flexible forward age-structured assessment program. ICCAT working document, SCRS/98/58. 15 p.

McGregor, V. (2015). Stock assessment of ling (*Genypterus blacodes*) on the Chatham Rise (LIN 3&4) for the 2014–15 fishing year. *New Zealand Fisheries Assessment Report 2015/82*. 50 p.

Methot, R.D.; Taylor, I.G. (2011). Adjusting for bias due to variability of estimated recruitments in fishery assessment models. *Canadian Journal of Fisheries and Aquatic Sciences 68(10)*: 1744–1760.

Methot, R.D.; Wetzel, C.R. (2013). Stock synthesis: A biological and statistical framework for fish stock assessment and fishery management. *Fisheries Research 142*: 86–99.

Myers, R.A.; Bridson, J.; Barrowman, N.J. (1995). Summary of worldwide stock and recruitment data. *Canadian Technical Report on Fisheries and Aquatic Sciences No. 2024*: 327 p.

7. APPENDIX A: Distribution of year class strengths when parameterised as recruitment deviations

Considering the distribution of the year class strengths parameter (YCS_y)

$YCS_y = e^{\varepsilon_y - \frac{1}{2}\sigma_R^2}$
Where,
$\varepsilon_y \sim N(0, \sigma_R^2)$

In this case, $YCS_y = e^{\varepsilon_y - \frac{1}{2}\sigma_R^2}$ implies
$\ln(YCS_y) = \varepsilon_y - \frac{1}{2}\sigma_R^2$

and

$\mathrm{E}[\ln(YCS_y)] = E\left[\varepsilon_y - \frac{1}{2}\sigma_R^2\right] = -\frac{1}{2}\sigma_R^2$
Since $\mathrm{E}[\varepsilon_y] = 0$. We also have

$\mathrm{Var}[\ln(YCS_y)] = \mathrm{Var}\left[\varepsilon_y - \frac{1}{2}\sigma_R^2\right] = \sigma_R^2$

Therefore

$\ln(YCS_y) \sim N(-\frac{1}{2}\sigma_R^2, \sigma_R^2)$

Thus, this parameterisation results in the same distribution for YCS_y as if we used the lognormal formulation. However, if the negative log-likelihood in an estimation framework applies to ε_y, adjustments need to be made so that *a priori* YCS_y is assumed to be lognormal. If the likelihood is applied to ε_y the prior contribution would be,

$$f(\varepsilon_y) = \frac{1}{\sigma_R\sqrt{2\pi}} e^{-\frac{1}{2\sigma_R^2}\varepsilon_y^2}$$

And therefore, if we consider the negative log-likelihood,

$-\ln f(\varepsilon_y) = \ln(\sigma_R) + \frac{1}{2}\ln 2\pi + \frac{1}{2\sigma_R^2}\varepsilon_y^2$
If we make $\varepsilon_y = \ln(YCS_y) + \frac{1}{2}\sigma_R^2$ this becomes
$-\ln f(\varepsilon_y) = \ln(\sigma_R) + \frac{1}{2}\ln 2\pi + \frac{1}{2\sigma_R^2}\left(\ln(YCS_y) + \frac{1}{2}\sigma_R^2\right)^2$

If we consider the negative log-likelihood contribution for $YCS_y \sim LN(-\frac{1}{2}\sigma_R^2, \sigma_R^2)$

$$-\ln f(YCS_y) = \ln(YCS_y) + \ln(\sigma_R) + \frac{1}{2}\ln 2\pi + \frac{1}{2\sigma_R^2}\left(\ln(YCS_y) - \left(-\frac{1}{2}\sigma_R^2\right)\right)^2$$

$$= \ln(YCS_y) + \ln(\sigma_R) + \frac{1}{2}ln2\pi + \frac{1}{2\sigma_R^2}\left(\ln(YCS_y) + \frac{1}{2}\sigma_R^2\right)^2$$

We can see that the likelihood contributions differ by $\ln(YCS_y)$, if we put *a priori* contributions on transformed parameters such as ε_y but our prior belief actually is on the resulting multipliers then we need to take account for the change in variable, following the change of variable theorem,

$$g(YCS_y) = f(s(YCS_y))\left|\frac{ds(YCS_y)}{dYCS_y}\right|$$

where $s(YCS_y) = \varepsilon_y(YCS_y)$ is the result of the conversion from YCS_y to ε_y and $\left|\frac{ds(YCS_y)}{dYCS_y}\right|$ is the Jacobian of the transformation. We find $s(YCS_y)$ by expressing ε_y as a function of YCS_y,

$$YCS_y = e^{\varepsilon_y - \frac{1}{2}\sigma_R^2}$$
$$\ln(YCS_y) = \varepsilon_y - \frac{1}{2}\sigma_R^2$$
$$\varepsilon_y = \ln(YCS_y) + \frac{1}{2}\sigma_R^2$$

Therefore:

$$\frac{ds(YCS_y)}{dYCS_y} = \frac{d\varepsilon_y(YCS_y)}{dYCS_y} = \frac{d}{dYCS_y}\left(\ln(YCS_y) + \frac{1}{2}\sigma_R^2\right) = \frac{1}{YCS_y}$$

So if we plug $\varepsilon_y = \ln(YCS_y) + \frac{1}{2}\sigma_R^2$ into $-ln\,f(\varepsilon_y)$ and add the Jacobian above we get the following prior contribution

$$g(YCS_y) = f(\varepsilon_y(YCS_y))\frac{1}{YCS_y}$$

For $\varepsilon_y = \ln(YCS_y) + \frac{1}{2}\sigma_R^2$,

$$= \frac{1}{\sigma_R\sqrt{2\pi}}e^{-\frac{1}{2\sigma_R^2}\left[\ln(YCS_y)+\frac{1}{2}\sigma_R^2\right]^2}\frac{1}{YCS_y}$$
$$= \frac{1}{YCS_y\sigma_R\sqrt{2\pi}}e^{-\frac{1}{2\sigma_R^2}\left[\ln(YCS_y)+\frac{1}{2}\sigma_R^2\right]^2}$$

which is the density function for the lognormal distribution with parameter $\mu - \frac{1}{2}\sigma_R^2$ and σ_R^2 thus will have a resulting lognormal prior on the quantity of interest YCS_y.

8. APPENDIX B: The effect of standardising year class strength parameters.

Taking year class strength values, Y_y, for each time point we can obtain standardised values, YCS_y using the formula:

$$YCS_y = \frac{Y_y}{\bar{Y}}$$

for $y = 1, \dots, T$ (T is the number of time points) and

$$\bar{Y} = \frac{\sum_{y=1}^{T} Y_y}{T}$$

The following calculations examine the resulting assumptions for the parameters YCS_y given that we make *a priori* assumptions on Y_y being lognormal. First, we consider the distribution of $\bar{Y}$, assuming:

$$Y_y \sim LN(-\frac{1}{2}\sigma_R^2, \sigma_R^2)$$

The distribution of $\bar{Y}$ relies on the distribution of a sum of lognormal random variables, $\sum_{y=1}^{T} Y_y$. No closed form probability distribution exists for such a sum, and it requires approximation. Some approximation methods date back over 80 years and most take one of two approaches; either: 1) an approximate probability distribution is derived mathematically, or 2) the sum is assumed to follow a lognormal distribution and the parameters of that distribution are then approximated. We take the later approach and assume:

$$\sum_{y=1}^{T} Y_y \sim LN(\mu_s, \sigma_s^2)$$

where μ_s and σ_s^2 can be approximated through the commonly used Fenton-Wilkinson method to be:

$$\sigma_s^2 = ln\left[\frac{\sum e^{2\mu_y + \sigma_y^2}(e^{\sigma_y^2} - 1)}{\left(\sum e^{\mu_y + \frac{1}{2}\sigma_y^2}\right)^2} + 1\right]$$

$$\mu_s = ln\left[\sum e^{\mu_y + \frac{1}{2}\sigma_y^2}\right] - \frac{1}{2}\sigma_y^2$$

In our case, $\mu_y = -\frac{1}{2}\sigma_R^2$ and $\sigma_y^2 = \sigma_R^2$ for $y = 1, \dots, T$. Substituting these values into the expressions above gives

$$\sigma_s^2 = ln\left[\frac{e^{\sigma_R^2} + T - 1}{T}\right]$$

And

$$\mu_s = -\frac{1}{2}\sigma_R^2$$

Using variable transformation techniques, we can show that

$$\bar{Y} = \frac{\sum_{y=1}^{T} Y_y}{T} \sim LN(\mu_s - lnT, \sigma_s^2)$$

This also follows from the general rule for lognormal random variables that if $X \sim LN(\mu, \sigma^2)$, then $aX \sim LN(\mu + \ln a, \sigma^2)$. In our particular example then,

$$\bar{Y} \sim LN\left(lnT - \frac{1}{2}\sigma_R^2 - lnT, \sigma_s^2\right)$$

Now we consider the distribution of YCS_y the ratio of two lognormal random variables

$$YCS_y = \frac{Y_y}{\bar{Y}}$$

The following property about the ratio of lognormal random variables is described by (Crow & Shimizu 1988)

> if $X_1 \sim LN(\mu_1, \sigma_1^2)$ and $X_2 \sim LN(\mu_2, \sigma_2^2)$ are independent lognormal random variables and $Z = \frac{X_1}{X_2}$, then $Z \sim LN(\mu_1 - \mu_2, \sigma_1^2 + \sigma_2^2)$.

If we assume Y_y is independent of $\bar{Y}$, then we have:

$$YCS_y \sim LN(-\frac{1}{2}\sigma_R^2 - [\mu_s - lnT], \sigma_R^2 + \sigma_s^2)$$

$$YCS_y \sim LN(-\frac{1}{2}\sigma_R^2 - [-\frac{1}{2}\sigma_R^2], \sigma_R^2 + \sigma_s^2)$$
$$YCS_y \sim LN(0, \sigma_R^2 + \sigma_s^2)$$

But, is Y_y independent of $\bar{Y}$? The mean $\bar{Y}$, is a function of all the Y_y values. However, as $T \rightarrow \infty$ the effect of a single Y_y on $\bar{Y}$ becomes smaller and we can therefore say they are approximately independent. However, for smaller T values the independence assumption is unlikely to hold, and therefore the lognormal distribution assumption about YCS_y may not hold.

In summary, the distribution of standardised YCS_yvalues has been shown to be approximately lognormal. However, the following caveats should be noted:

1. The lognormal density used for the sum of lognormal random variables is only an approximation because there is no closed form expression for this distribution. The approximation is reasonably accurate only in the right tail (i.e., for larger sums) (Asmussen & Rojas-Nandayapa 2008). The density function in the neighbourhood of 0 has been investigated and does not resemble a lognormal distribution (Gao et al. 2009).

2. We assumed that the unstandardised year class strengths, Y_y, are each independent of their mean $\bar{Y}$. This assumption is counter-intuitive and will not hold for shorter time series or values of Y_y in the right tail.

Further investigation is required to determine the distribution of the ratio $Y_y/\bar{Y}$ under the assumption that Y_y and $\bar{Y}$ are correlated. This would give a more accurate approximation to the distribution of YCS_y but is unlikely to result in a lognormal distribution.