An Initial Response to Legault: "The ability of two age composition error distributions to estimate selectivity and spawning stock biomass in simulated stock assessments"

Doug S. Butterworth and Rebecca A. Rademeyer

Marine Resource Assessment and Management Group (MARAM)
Department of Mathematics and Applied Mathematics

University of Cape Town

Rondebosch 7701, South Africa

July 2014

Summary

The basis of inferences drawn from the simulation studies of Legault (2013), which indicated positive bias in estimates spawning stock biomass by certain SCAA models, together with their tendency to estimate stronger domes in selectivity than actually present, is examined. Legault's results are indicated as likely to be an artefact of inappropriate specifications for his simulation study, which led to the generation of some datasets with unrealistically low proportions of older fish. For the more realistic of his datasets, median bias in the estimation of spawning biomass is smaller, and becomes virtually zero if a constraint on selectivity imposed by the estimation procedure which Legault applied is removed. Accordingly it seems that Legault's concerns, including about certain SCAA models having a tendency to estimate stronger domes in selectivity than are actually present, are likely unfounded.

Introduction

Legault (2013) reports the results of simulation studies where Statistical Catch-at-Age (SCAA) assessment approaches are applied to abundance index and age-composition data of the type which are typically input to such assessment models. He generates observed age composition data from the population simulated using both multinomial and "adjusted lognormal" (Punt and Kennedy 1997) distributional forms, and considers situations both with and without a dome in selectivity. The adjusted lognormal form has been used, for example, by the authors of this paper in assessments of the Gulf of Maine cod population (e.g. Butterworth and Rademeyer 2011).

Legault's (2013) conclusions include the following:

- The adjusted lognormal error distribution form produces positively biased estimates of spawning stock biomass (SSB), which are almost always higher than the true underlying SSB.
- Use of either the adjusted lognormal error or the multinomial distribution forms in the assessment model estimator had tendencies to estimate domes where they did not exist, or stronger domes than actually existed, with the former model consistently estimating stronger domes than the latter. Given the fact that the stronger the dome, the larger the proportion "cryptic" biomass of older fishes that is estimated, the consequences of this are (*inter alia*) positively biased estimates of SSB.
- Utility of the adjusted lognormal model in stock assessments is greatly diminished because of the biased nature of the estimator, high variability of estimates, and inability to properly identify the correct pattern of selectivity.

- The estimation of domed selectivity in SCAA models should be undertaken with caution because of consistent tendencies to estimate stronger domes than in the true underlying data.

This paper duplicates Legault's simulation methodology to examine the reasons underlying his results more closely, and thereby the validity of the inferences which he draws from them.

The paper reports the initial stage of analyses being conducted to support an article to be submitted in response to Legault (2013). It is being tabled in this manner because the matter has relevance to, and has been raised during discussions on the assessment of Atlantic menhaden currently underway, and the deadline for submissions for consideration in this assessment process occurs before the full response can be completed and submitted to the journal that originally published the article.

Methodology

This paper uses identical methodology to that set out in Legault (2013) to generate simulated data[1]. The SCAA assessment model estimation procedure (as used, for example, in Butterworth and Rademeyer 2011) is then applied to various data sets generated by this process to examine Legault's (2013) conclusions further.

The specific scenarios from Legault (2013) which are examined are his runs AEFe and AEDe. Both of these generate catch-at-age observations under the adjusted lognormal distribution assumption with a starting numbers-at-age vector determined by "equilibrium" assumptions, and they estimate selectivity values for every age freely, except that having fixed the selectivity at age 5 to be 1, the selectivities at other ages are constrained not to exceed 1 (consistent with the true values for the underlying scenarios). The difference between the two scenarios is that for ages 5 and above the selectivity is **f**lat for AE**F**e, but **d**omed for AE**D**e.

Results for both the AEFe and the AEDe scenarios as reported in figure 3 of Legault (2013) support his conclusions as listed in the bulleted points in the Introduction section above.

Results and Discussion

Figure 1 shows results from simulations which duplicate those for AEFe and AEDe which are reported in figure 3 of Legault (2013). For both scenarios the results are (again) consistent with the inferences drawn by Legault in his conclusions: the estimated spawning biomass is positively biased, can be substantially above the true biomass, and nearly always exceeds that true biomass.

Figure 2 shows the range of spawning biomass and fishing mortality (*F*) values that result from Legault's simulations for the flat (AEFe) selectivity scenario. The random walk model assumed to generate *F* values $F_{t+1} = F_t e^{0.1\varepsilon}$ where ε is drawn from $N(0,1)$ is median unbiased, but because of its lognormality can lead to values of *F* approaching a high of 2 towards the end of the 30 year period simulated. Given that Legault's model has a maximum age of 10+ (a plus group), the age with the smallest representation (on average) in the catch is 9. While the true value of the proportion of age 9 in the catch has a median of about 1%, though this can climb as high as 12%, at the other extreme it can drop to well below 0.1%.

[1] Legault's formula for the estimate of the catch-at-age variance parameter σ of the SCAA model likelihood on the fourth page of his paper is in error. The 1/*n* term should be inside rather than outside the square root sign. Given however that the computations of this paper, with this equation corrected, do duplicate comparable results to those reported by Legault (2013), it seems that this is simply a typographical error.

The adjusted lognormal model assumes the variance of catch-at-age proportions to be inversely proportional to those proportions, so that if the true proportions become very low the associated variance will become very high. Figure 3 shows the consequences of this, in comparing the true proportions of age 9 in the final year of each of the 200 simulations conducted to the values generated for these proportions in the simulations. Note that quite high “observed” values can be generated, even though true proportions are low – in one case the “observed” proportion of age 9 in the catch is (virtually) 100%.

Such instances are hardly realistic. Were they to occur in practice, a routine run of a standard SCAA model would scarcely be contemplated. Indeed given a catch proportions-at-age matrix which includes very low values, it is standard practice to group ages in the application so as to avoid instances of observed proportions below some threshold, typically taken to be about 1 or 2%, to reduce estimation variance.

To examine the consequences of the inclusion of these unrealistic cases in Legault’s simulations, two of the 200 simulations have been selected: Sim A and Sim B. For Sim A the true proportion of 9 year old fish in the catch is typically in the 2-3% range, and seldom falls (much) below 1% during the 30 year simulation period (see Figure 4a for the flat and 4b for the domed selectivity cases). In contrast, for Sim B, this true proportion can drop very low, to less than 0.02%. These Figures also compare the standard deviations of the lognormal errors used to generate the observed from these true proportions. For Sim A this is typically about 0.4; however for Sim B it can reach 6 and even higher towards the end of the period simulated.

For both Sim A and Sim B, 200 datasets were generated keeping the population dynamics parameters (the starting numbers-at-age, the *F* values, and so forth) the same (though they do differ slightly between the Flat and the Domed scenarios because of different selectivities impact the catches-at-age differently). What changes between these datasets is the observation errors generated for the abundance index and catch proportions-at-age.

Figure 5 shows the results of applying the SCAA model to the data sets corresponding to Sim A and to Sim B, for both the flat and the domed selectivity scenarios. For the realistic case (Sim A) the variation of spawning biomass estimates is not large, and these estimates show a relatively small median biased in the 19-20% range for flat selectivity scenario, but are virtually median unbiased for the domed scenario. In contrast, for the unrealistic case (Sim B), all spawning biomass estimates are larger than the true values by some hundreds of percent.

The presence of a median bias for Sim A for the flat but not for the domed selectivity scenario may be surprising. It turns out that it is a consequence of a constraint imposed on the SCAA estimator in this case, specifically that (with selectivity at age 5 fixed at 1) selectivities for ages above 5 were constrained not to exceed 1. Removing this constraint (Sim Adj – see Figure 6) results in estimates of spawning biomass for the flat selectivity scenario that are also all virtually median unbiased. (Results for the domed selectivity scenario are scarcely changed for Sim Adj because the constraint on the selectivities estimated hardly came into play in that case.)

Figure 7 illustrates what might be a further problem with Legault’s (2013) simulations. Values for observed catch proportions-at-age are generated according to the lognormal distribution assumed for the adjusted lognormal model. However, as proportions-at-age must sum to 1 for any one year, the values so generated have been renormalized to 1. The question thus arises of whether this renormalisation preserves the lognormality intended. While Fig. 7 does not constitute a sufficient examination of this question to provide a definitive answer, it is suggestive that this might be a problem for Sim B situations where the expected proportions of age 9 fish can be very small. Developing procedures to correct for such a bias in generating data could prove a challenging task. However such an attempt would likely be moot, as the arguments above suggest that the associated scenario would not arise in practical applications of the assessment method concerned.

Concluding Remarks

The outcomes illustrated above indicate that Legault's results and associated inferences are likely an artefact of an inappropriately specified simulation study. The bad estimation performance he reports seemingly follows from some of the datasets which his procedure generates having unrealistically low proportions of older fish, which in turn lead to the generation of datasets for proportions-at-age to which SCAA models of the type considered would never be applied were they to occur in practice.

For the more realistic of his datasets, variability of estimates of spawning biomass is much reduced and median bias is smaller. Furthermore median bias becomes virtually zero if the constraint on selectivities imposed by the estimation procedure which he applied is removed.

Accordingly it seems that Legault's concerns about certain SCAA models having a tendency to estimate stronger domes in selectivity than are present (which is linked to positive bias in abundance estimates) are likely unfounded (at least in relation to the reasons he examined). Caution should however be exercised if estimable selectivities are constrained to be less than the value of 1 specified for some reference age, and are found to lie on that constraint boundary – in such cases the constraint should be reconsidered with a view to avoiding possible positive bias in abundance estimation.

Furthermore, in designing simulation studies, care needs to be taken that the datasets generated are at least typical of the situations regarding which they are being used to draw inferences. In the case of a specific stock (Legault 2013 focused particularly on the Gulf of Maine cod stock), the practice of "conditioning (the simulations) on the data" should be followed. This practice was first developed in the Scientific Committee of the International Whaling Committee for conducting MSEs, and was later adapted and detailed by the ICES Working Group on Methods of Fish Stock Assessments (WGMG) for the purpose of simulation testing of assessment methods (ICES 2012).

References

Butterworth DS and Rademeyer RA. 2011. Applications of statistical catch-at-age assessment methodology to Gulf of Maine cod. Document submitted to the 17-21 October, 2011 workshop on the assessment of Gulf of Maine cod, Falmouth. 31pp.

ICES. 2012. Working Group on Methods of Fish Stock Assessments (WGMG), 8–12 October 2012. Lisbon, Portugal. ICES CM 2012/SSGSUE:09. 249pp.

Legault CM. 2013. The ability of two age composition error distributions to estimate selectivity and spawning stock biomass in simulated stock assessments. Fish. Res., http://dx.doi.org/10.1016/j.fishres.2013.12.007 .

Punt AE and Kennedy RB. 1997. Population modeling of Tasmanian rock lobster, *Jasus edwardsii*, resources. Mar. Freshw. Res. 48, 967–980.

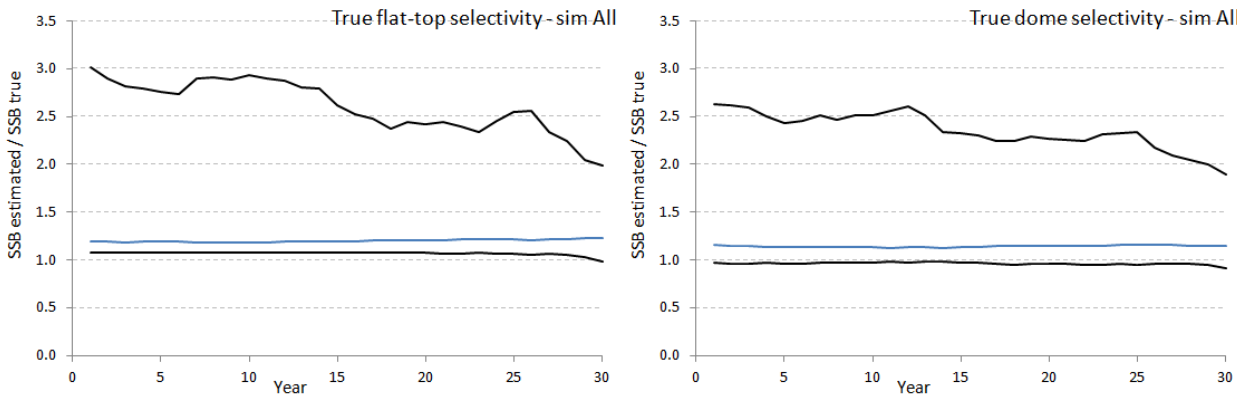

Figure 1: Medians (blue line) and 10th and 90th percentiles (black lines) envelopes of the ratios of estimated over true SSB for the true flat-top selectivity and true dome selectivity scenario simulations, as conducted by Legault (2013), plotted against simulation year.

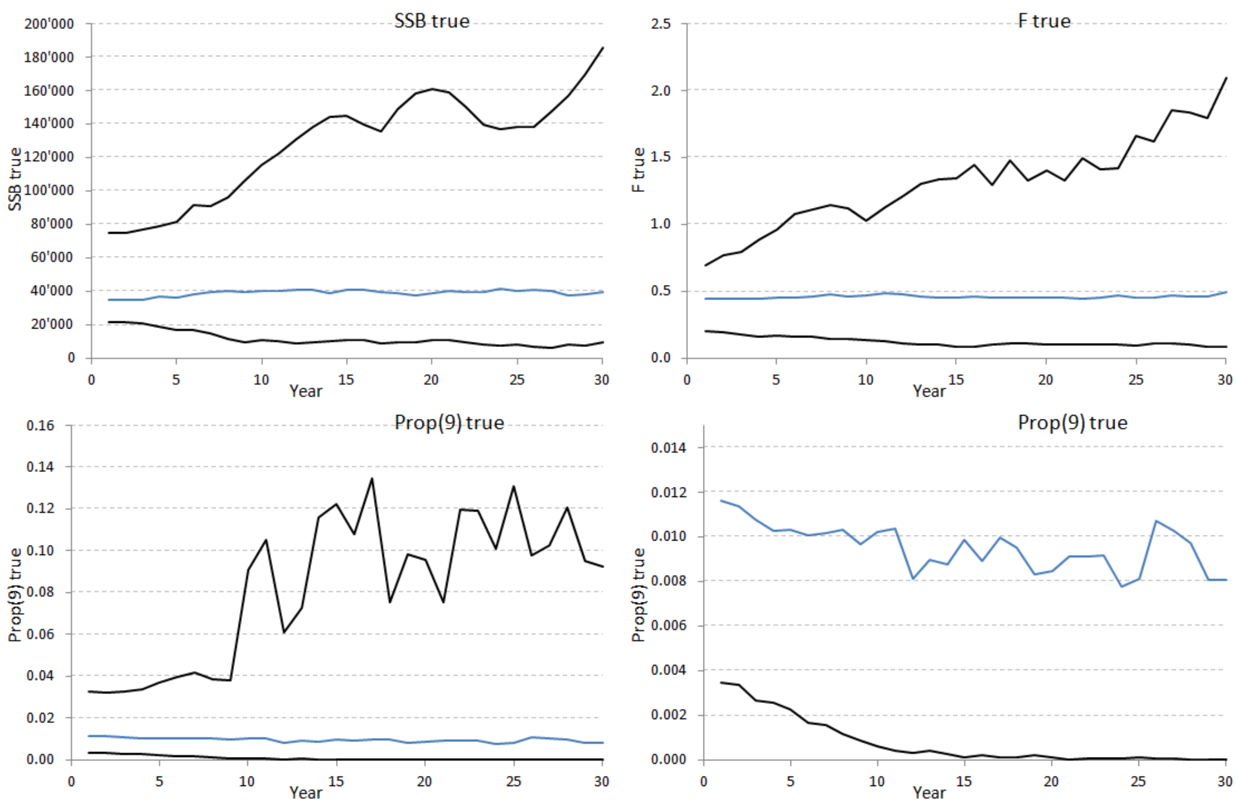

Figure 2: Medians (blue line) and minimum and maximum (black lines) plots against simulation year of true SSB, true fishing mortality and true proportion of fish of age 9 in the catch (the bottom right plot repeats this with a different scale) for the true flat selectivity scenario.

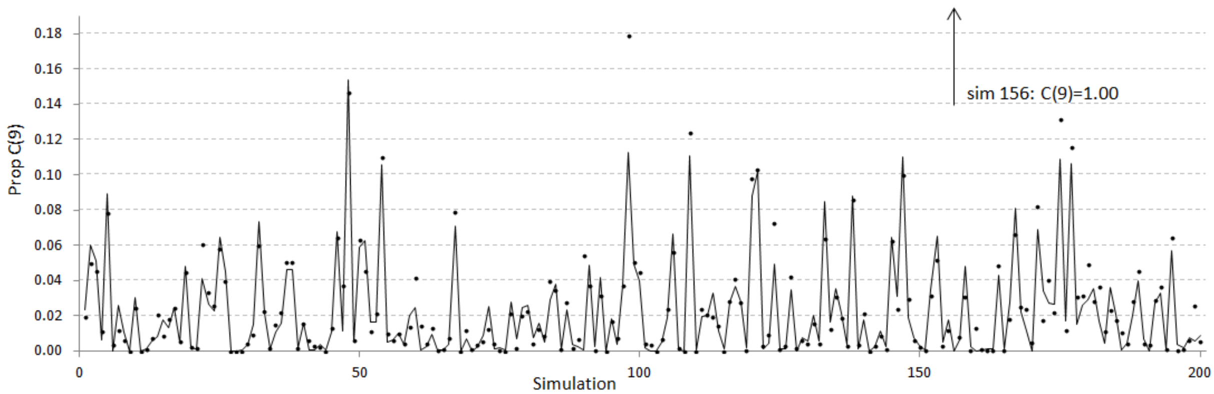

Figure 3: True (black line) and observed (black dots) proportions of fish of age 9 in the catch in the final year (30) for each simulation, for the true flat selectivity scenario.

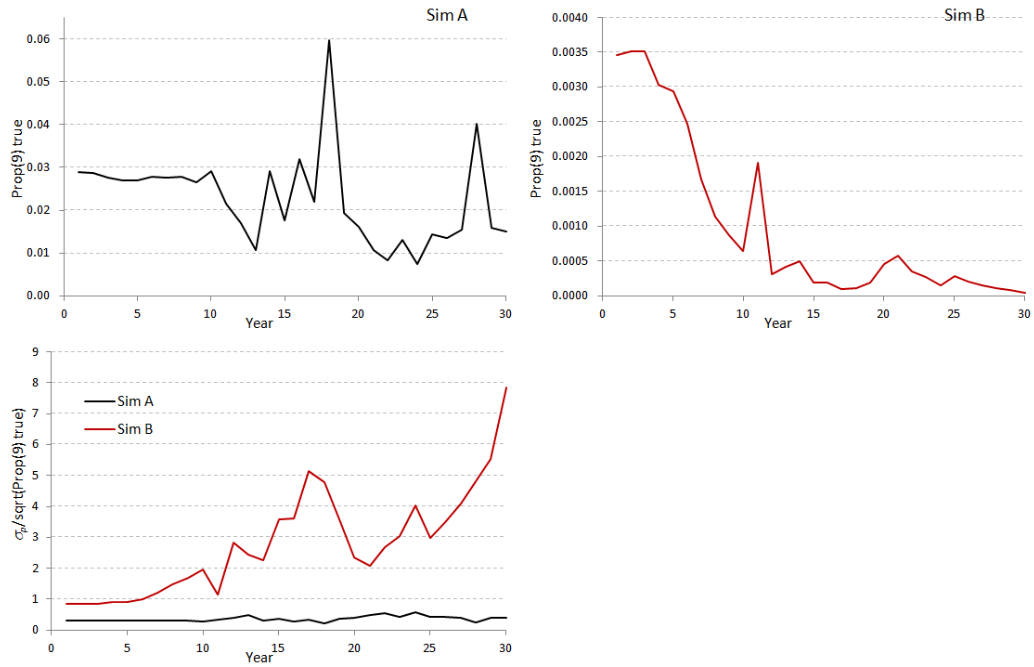

Figure 4a: Time series of true proportions of fish of age 9 in the catch for selected simulations Sim A and Sim B, and corresponding standard deviation used for generating observations for the true **flat selectivity scenario**.

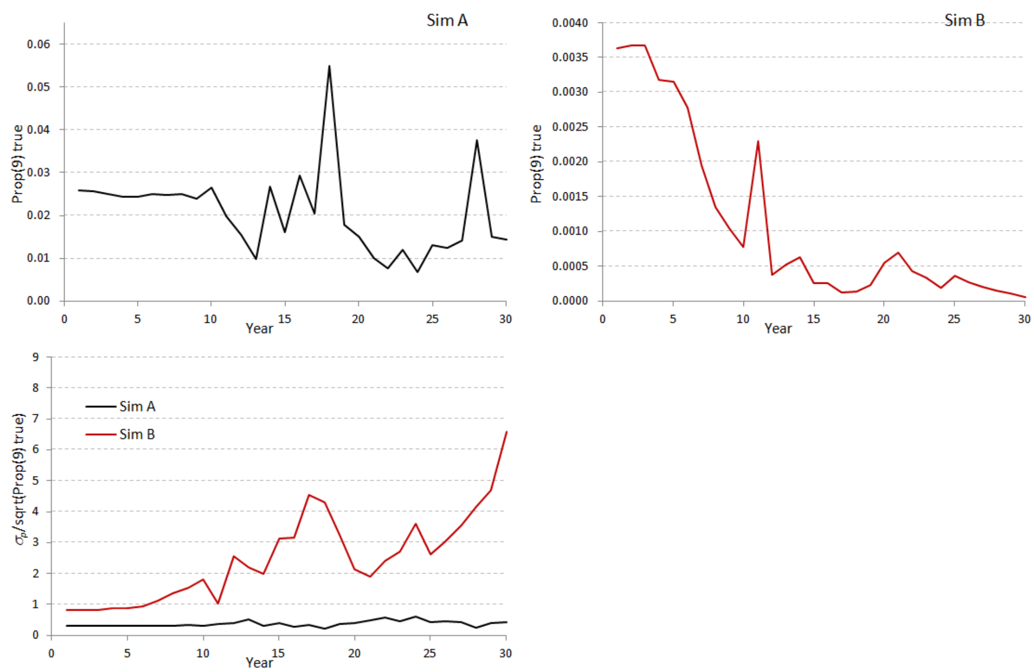

Figure 4b: Time series of true proportions of fish of age 9 in the catch for selected simulations Sim A and Sim B, and corresponding standard deviation used for generating observations for the true **dome-selectivity scenario**..

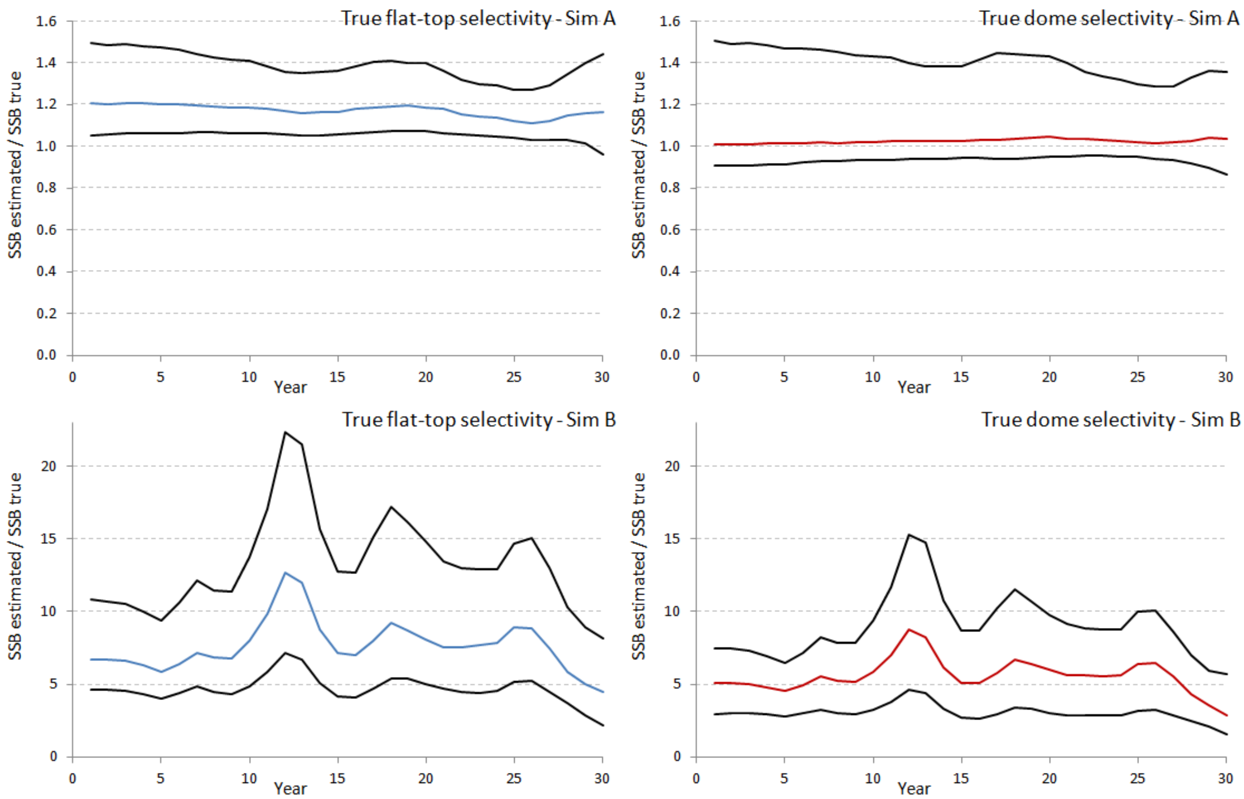

Figure 5: Medians (blue and red lines) and 10th and 90th percentiles (black lines) envelopes for the ratios of estimated over true SSB for the true flat selectivity (blue) and true dome selectivity (red) scenarios for selected simulations Sim A and Sim B.

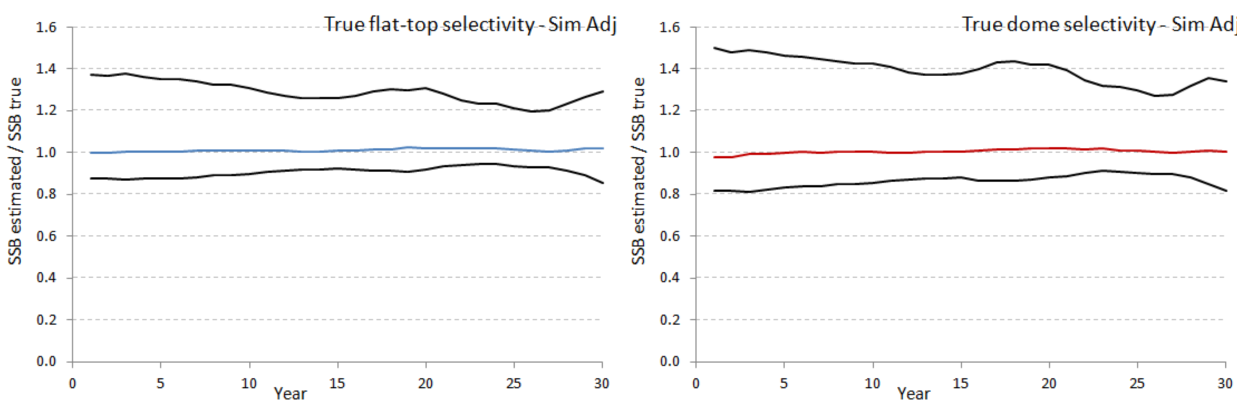

Figure 6: Medians (blue and red lines) and 10th and 90th percentiles (black lines) envelopes for the ratios of estimated over true SSB for the true flat-selectivity (blue) and true dome selectivity (red) for Sim Adj for which constraints of a maximum of 1 for estimates of selectivities above age 5 are removed .

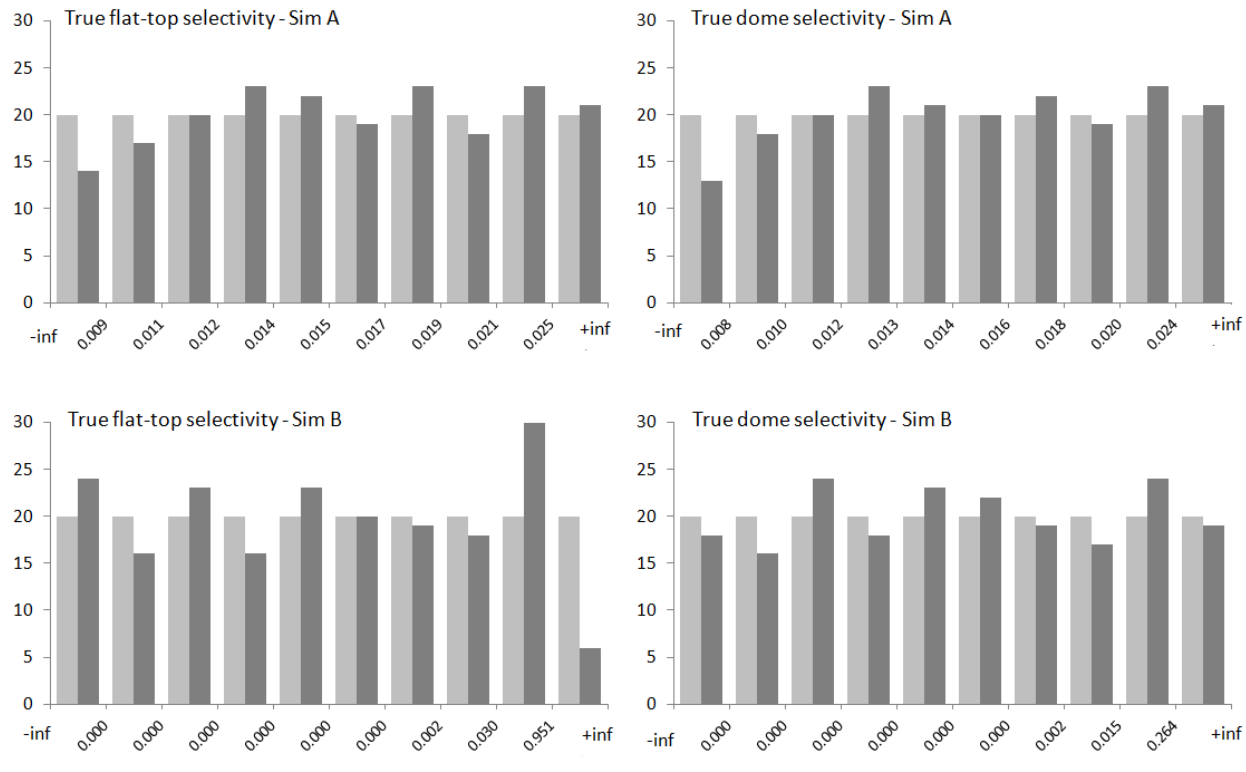

Figure 7: Frequency distribution (dark grey) of observed proportion of fish of age 9 in year 30 in the catch for the selected simulations Sim A and Sim B for the true flat and dome selectivity scenarios. The expected lognormal distribution (light grey) is also shown. Note that the horizontal axis is non-linear, with the bin edges chosen so that each bin would be expected to contain 20 of the 200 simulations generated accordingly to the lognormal distribution form applicable.