Refining the November 2019 sardine west component recruitment estimate

D. S. Butterworth and A. Ross-Gillespie[1]

Summary

The estimate of the west component of sardine recruitment from the 2020 recruitment survey has relatively high variance, and it is conventional before basing management advice on such a single estimate alone to reduce that extent of uncertainty by combining the result in some way with the information provided by immediately preceding values of recruitment and spawning stock size. Here that further information is taken to be provided by the sardine assessment estimates of November recruitment for the years 2014 to 2018, and three different approaches are used to combine this information with the 2020 survey estimate. These three approaches are crossed with three variants of a regression method to relate the 2020 survey result to the corresponding recruitment strength as at November 2019. In broad terms, the results from these combinations suggest that the 2019 November sardine recruitment estimate of about 9 billion provided by the mean for the previous five years' recruitment estimates from the assessment to be updated to lie in the range of 11 to 13 when the new information provided by the 2020 survey result is incorporated.

Background

The 2020 recruitment survey provides a survey estimate for the west component of sardine recruitment, which can be transformed into an estimate expressed in terms of a value and associated standard error for the preceding November, which corresponds to the time for which outputs from the assessment model are generated.

The regression approach of equation (4) of FISHERIES/2020/JUL/SWG-PEL/57 provides a basis for such a transformation, with that calibration being conducted in log-space and providing a standard deviation estimate *sig*. This is taken forward below for three implementations of that approach:

a) Calibration based on all historical data, with sig based on 2005+ data only
b) Calibration and *sig* based on 2005+ data only
c) Calibration based on all historical data, with sig based on 2010+ data only

The immediate interest in this estimate from the 2020 survey is, however, related especially to the extent to which it updates inferences based on assumptions concerning this recruitment before the survey took place. At that time, these were based on the estimates of the five preceding November recruitments (2014 – 2018) from the assessment model. These were considered as equally likely, and treated effectively as a discrete prior for the November 2019 recruitment.

The situation is not unlike that for conventional stock assessments, where the most recent recruitment is estimated with poor precision because of limited information available to the assessment algorithm to inform on recruitment strength for that year in comparison to earlier years. In such situations, it is not uncommon to use other information, for example from a stock-recruitment relationship obtained from the assessment, to "improve" the estimate provided by the survey alone.

[1]Marine Resource Assessment and Management Group, Department of Mathematics and Applied Mathematics, University of Cape Town, Rondebosch, 7701

This document provides results for three such approaches, comparing them through tabulations and the use of probability density function plots.

Method 1: Bayesian update of a discrete prior

In this approach the five equally likely (equally weighted) estimates from the assessment for November 2014 to November 2018 are treated as a discrete prior (i.e. as a prior consisting of five delta functions of equal weight of 0.2).

This prior can updated to include the 2020 survey result to provide a discrete posterior (five unequally weighted delta functions) using the Bayesian methodology described in FISHERIES/2020/JUL/SWG-PEL/60.

Comparing these distributions visually when in this discrete form is not too helpful for the reader. Accordingly, we have done so by instead approximating them by lognormal distributions with the same means and variances as the corresponding five (equally or unequally weighted) values for November 2014 to November 2018.

Note nevertheless that such an approximation does not exactly reflect all the features of the distribution (probability density function) concerned. Thus, for example, this approximate plot of the discrete posterior will indicate non-zero probabilities for values of recruitment outside the range of those “prior” values for 2014 to 2018; however, the true posterior is restricted to the five discrete values contained in that range.

Method 2: Bayesian update of a continuous prior

This last point leads to potential “discomfort” about the use of a discrete formulation for the prior based on those five earlier recruitments, which can be considered to reflect realisations from a stock-recruitment relationship, about which the distribution is customarily taken to be lognormal. The discomfort arises because even though the point estimate from the survey can lie outside the range of the five discrete “input” values for the prior, the posterior can give no weight to values outside that range; hence, for example, in the current situation, irrespective of how large or precise the estimate from the 2020 recruitment survey might have been, the approach above could not provide a posterior (mean, median, mode or even 99.9%ile) higher than the largest of the five input recruitment values)

To address this concern, an alternative approach is to first construct an actual lognormal prior in the same manner as the log-normal approximations were constructed for Method 1, and then to update this prior using the likelihood for the 2020 recruitment survey estimate to provide a posterior. The methodology in identical to that in FISHERIES/2020/JUL/SWG-PEL/60, with the Bayesian integration being implemented using a step-size of 0.001 recruitment unit (i.e. 1 million fish) – this coupled to a range from 0.001 to 40 was found to provide more than sufficient accuracy given the very low value of either to prior of the likelihood outside this range.

Note that the plots provided for the posteriors in this case are exact, and not approximations as for the discrete approach above.

Method 3: Inverse variance weighting

A final approach to combine the prior information and the survey result is to use inverse variance weighting (as common in stock assessments), appropriately conducted in log space given the assumption that the distribution of residuals about the stock recruitment relationship is assumed to be lognormal (with the relationship itself taken to be a constant with a value equal to the geometric mean of the 2014-2018 recruitments, i.e. equivalent to the arithmetic mean of their logarithms). Further details on this inverse variance weighting approach have been included in the Table 1 caption.

This provides a mean and standard deviation for the combined estimate, which can be plotted in the form of a lognormal distribution of equivalent mean and variance in the same way as above for the Bayesian approach with the discrete prior.

Results

Table 1 lists the means and standard deviations of the prior and posterior distributions (or the approximations/equivalences of these distributions in the case of Methods 1 and 3) in normal and log space for the three methods and the three regressions. Figures 1a-c plot the prior and posterior distributions (similarly approximated where indicated) for the three methods. Figure 2 compares these distributions across the three methods for regression (a).

Overall impressions

Ultimately at a broad level, the different methods as well as the different assumption for the regression used to relate survey and model recruitment estimates make little difference to the key results.

The mean of the most recent five November recruitments used before the 2020 recruit survey took place was 9. The updates to take account of the result from that survey lie in the range of 11 to 13 (see Table 1).

Table 1: The means and standard deviations (in parenthesis) of the prior and posterior distributions (or the approximations/equivalences of these distributions for Methods 1 and 3) are listed in the Table below. The calibration factor (ln*k*) and sigma (*sig*) for each regression have been included.

Method 1:
The mean and standard deviation of the "prior" in normal space are the mean ($\bar{R}_5$) and standard deviation (s_5) of the last five recruitment values. The mean (μ) and standard deviation (σ) of the prior in log space are calculated so that $\bar{R}_5 = \exp\,(\mu + \frac{\sigma^2}{2})$ and $s_5 = \left(e^{\sigma^2} - 1\right)e^{2\mu+\sigma^2}$. The estimates reported for the "posterior" distributions are similar, except that $\bar{R}_5$ and s_5 are replaced by the likelihood-weighted mean and standard deviation for the five values.

Method 2:
The prior distribution is constructed in the same manner as the approximation for the "prior" in Method 1. The posterior distribution is calculated by multiplying the prior by the corresponding likelihood (vector), and normalising so that the sum over the elements of the probability vector is one (equivalent to saying that the area under the curve is one). The mean of the posterior in normal space is calculated as $\bar{R} = \sum R_i P_i$ where P_i is the value of the posterior at recruitment value R_i, and the variance is calculated as $s^2 = \sum (R_i - R)^2 P_i$. The mean and standard deviation of the posterior in log space is calculated as for the normal space, but R_i is replaced by $\ln R_i$.

Method 3:
The "prior" distribution is a log-normal distribution with mean (μ_a) and standard deviation (σ_a) equal to the mean and standard deviation of the logs of the last five recruitment values. The mean and standard deviation in normal space are calculated as $\bar{R}_5 = \exp\,(\mu + \frac{\sigma^2}{2})$ and $s_5 = \left(e^{\sigma^2} - 1\right)e^{2\mu+\sigma^2}$. Let μ_b be the 2020 estimate of recruitment from the regression and σ_b the standard deviation of the residuals of the regression (i.e. sig above). Define $w_a = (1/\sigma_a^2)/\left(1/\sigma_a^2 + 1/\sigma_b^2\right)$ and $w_b = \left(1/\sigma_b^2\right)/\left(1/\sigma_a^2 + 1/\sigma_b^2\right)$. Then $\bar{R}' = w_a\mu_a + w_b\mu_b$ and $s' = \sqrt{w_a^2\sigma_a^2 + w_b^2\sigma_b^2}$ are the mean and standard deviation of the "posterior" distribution on log space. The same equations as for the prior distribution have been used to calculate the mean and variance of the "posterior" in normal space.

		Regression (a)		Regression (b)		Regression (c)	
		Prior	Posterior	Prior	Posterior	Prior	Posterior
		ln*k* = -1.101	*sig* = 0.635	lnk = -0.913	*sig* = 0.635	ln*k* = -1.101	*sig* = 0.729
Normal space	Method 1	8.861 (4.139)	11.499 (2.731)	8.861 (4.139)	11.132 (2.861)	8.861 (4.139)	11.121 (2.911)
	Method 2	8.861 (4.139)	11.772 (4.411)	8.861 (4.139)	11.068 (4.155)	8.861 (4.139)	11.202 (4.384)
	Method 3	9.240 (5.408)	13.087 (5.640)	9.240 (5.408)	12.087 (5.209)	9.240 (5.408)	12.405 (5.666)
Log space	Method 1	2.083 (0.444)	2.415 (0.234)	2.083 (0.444)	2.378 (0.253)	2.083 (0.444)	2.376 (0.257)
	Method 2	2.083 (0.444)	2.400 (0.364)	2.083 (0.444)	2.338 (0.364)	2.083 (0.444)	2.344 (0.379)
	Method 3	2.076 (0.543)	2.486 (0.413)	2.076 (0.543)	2.407 (0.413)	2.076 (0.543)	2.423 (0.435)

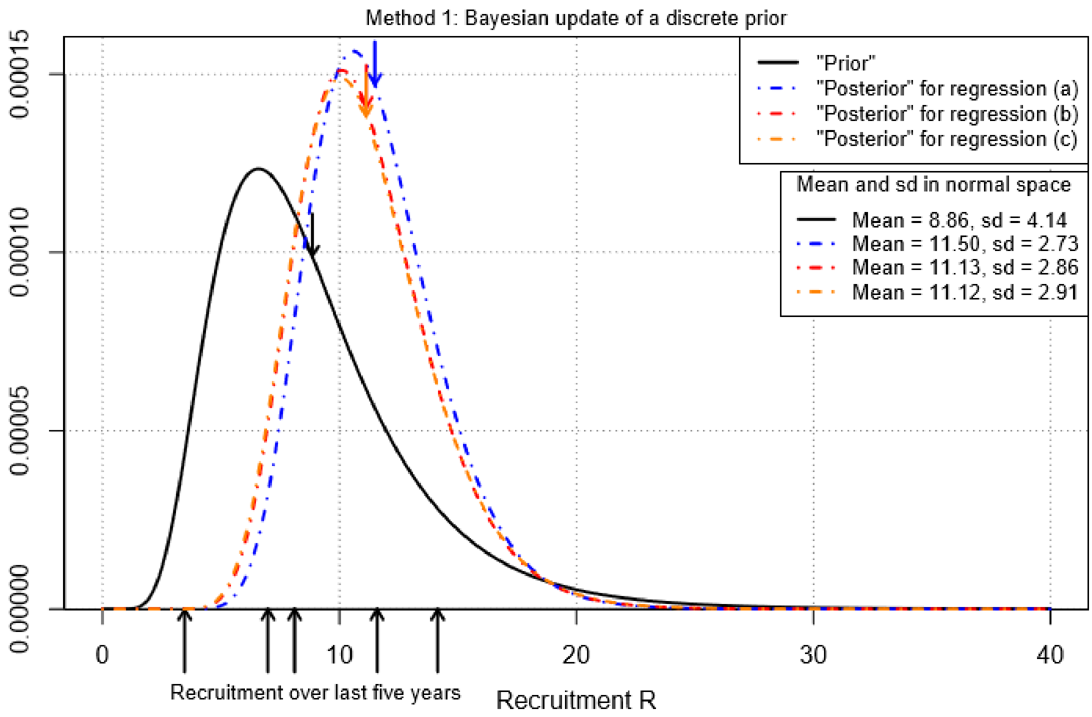

Figure 1a: Probability density functions for **Method 1**. Note that the curves have been normalised so that the area under each curve is one.

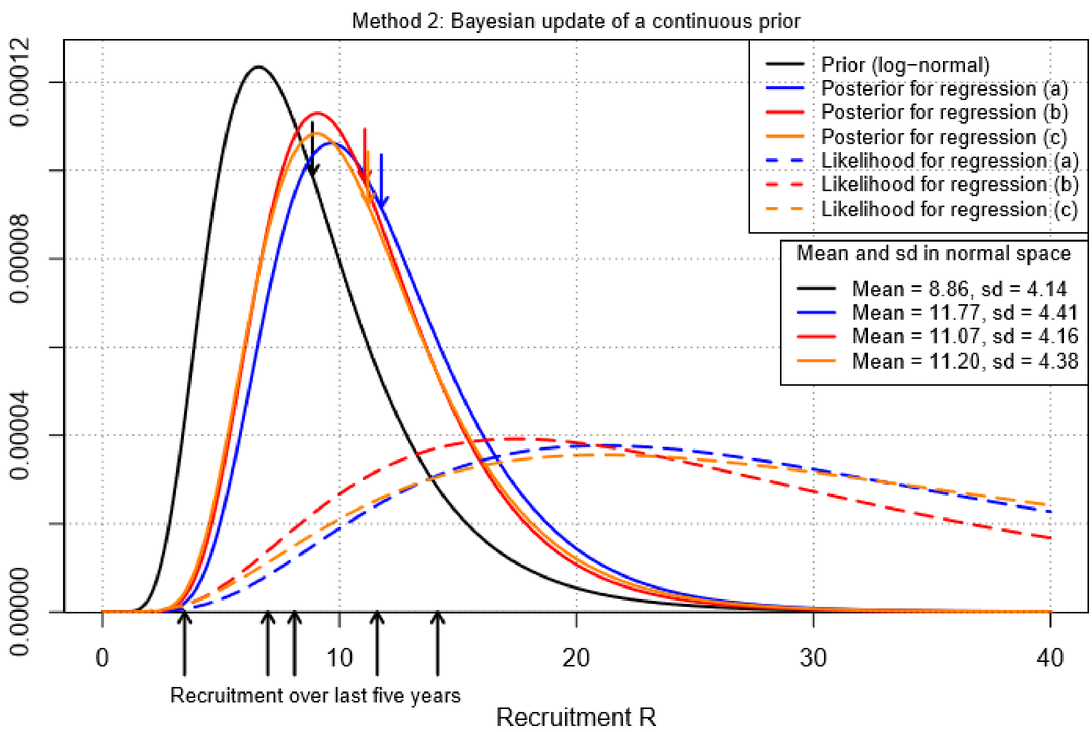

Figure 1b: Probability density functions for **Method 2**, showing the prior and posterior distributions, as well as the likelihood vector arising from the regression. The priors in Figure 1a and 1b are the same.

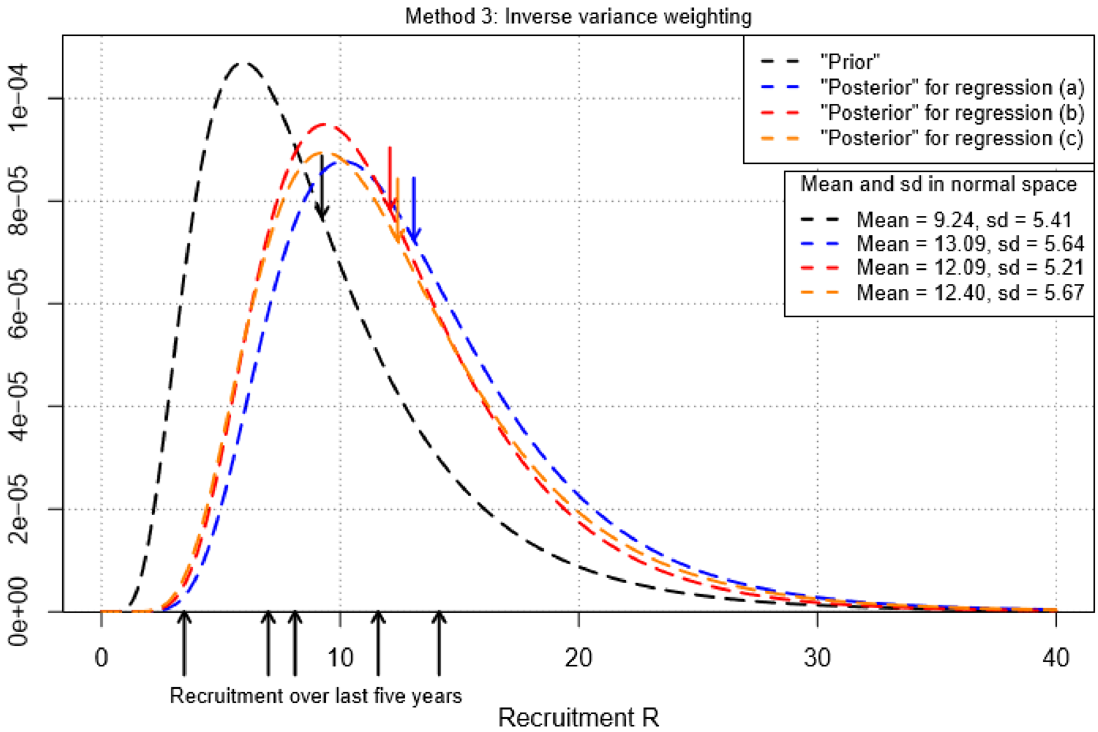

Figure 1c: Probability density functions for **Method 3**. Note that the curves have been normalised so that the area under each curve is one. Note that the prior in Figure 1c differs from that for Figures 1a and 1b – further explanation can be found in the Table 1 caption.

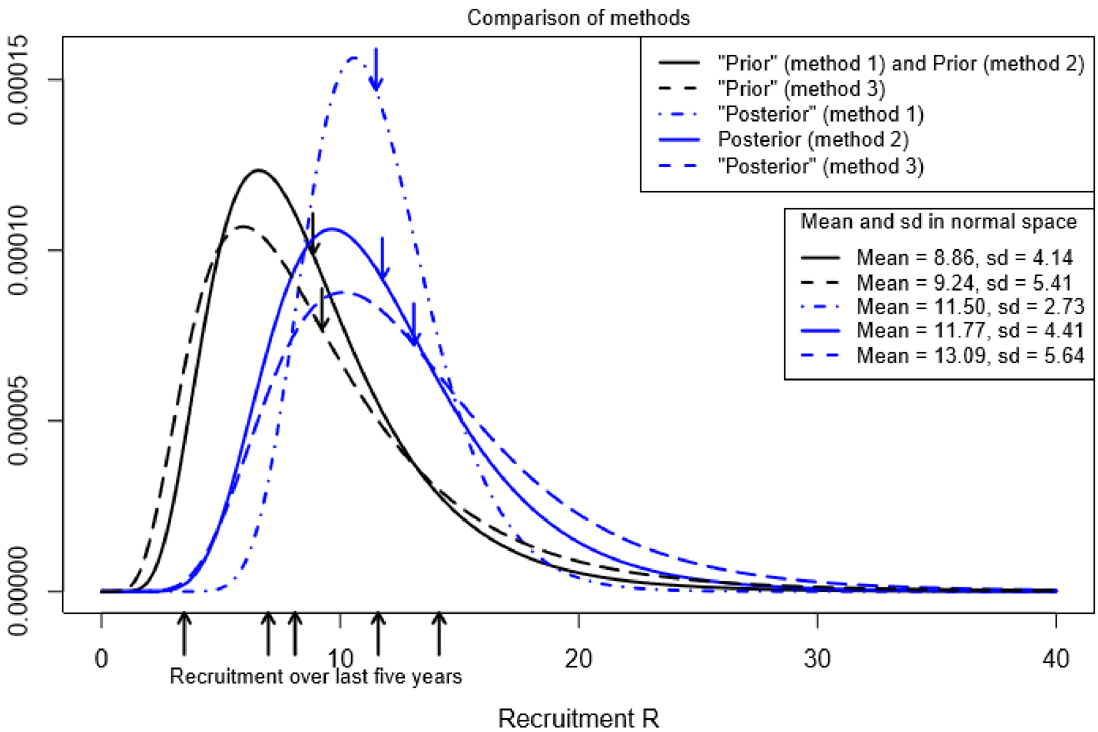

Figure 2: Comparison of the priors/"priors" and posteriors/"posteriors" for the three methods, for regression (a) which estimates the calibration factor from data from all years, but uses the *sig* value from the regression based on data from 2005 onwards only.