Review of the Harvest Control Rule for orange roughy fisheries in New Zealand

P.L. Cordue

11 August 2019

ISL client report to DWG

Table of Contents

Executive summary

In 2014, stock assessments were performed for four New Zealand orange roughy stocks. Later that year, an orange roughy harvest control rule (HCR) was developed as part of a management strategy evaluation (MSE). It is a generic control rule that was developed to be applicable to any of the New Zealand orange roughy stocks which have fisheries based primarily on mature fish. The HCR is used for management purposes for the three orange roughy stocks that have Marine Stewardship Council (MSC) certification (East & South Chatham Rise (ESCR), Northwest Chatham Rise (NWCR), and ORH7A including Westpac Bank).

In 2017, a fifth orange roughy stock (Puysegur) was assessed. The Puysegur fishery has not been put forward for MSC certification and it is not managed using the HCR. However, it does provide an additional set of stock assessment results which include an estimate of natural mortality (*M*). Also, the other four stock assessments have been updated since 2014. These updates provide revised estimates of *M* and stock-recruitment steepness (*h*) which are relevant to the review of the HCR.

In the 2014 MSE, the performance of the HCR was evaluated over a matrix of *M* and *h* values. Together with posterior distributions for *M* and *h* from the 2014 stock assessments, this allowed Bayesian estimation of various performance indicators for the HCR (e.g., the probability of mid-season spawning biomass being below the limit reference point (LRP)). With a fifth stock assessed, and with updated assessments for the four stocks assessed in 2014, revised posterior distributions have been estimated for *M* and *h* to enable updating of the performance indicators for the HCR.

The LRP and target biomass range for orange roughy were also outputs of the MSE and dependent upon the posterior distributions of *M* and *h.* Therefore, the LRP and target biomass range are also reviewed given the revised distributions for *M* and *h.*

For ESCR the likelihood profile suggests that there is very little information on *M* in the data. Most of the apparent information for *M* is coming from the very high correlation with virgin biomass (B_0) and the uncertainty associated with B_0. However, for all of the other stocks the input age frequencies appear to provide some information on *M.*

Mid-East Coast (MEC) stands out from the other stocks with the lowest and most precise estimate of *M.* The fixed value $M = 0.045$ used in the stock assessment is well outside the 95% CI for MEC (0.027–0.036) and it should probably not be used in the base MEC stock assessment model.

For the other stocks, the assumed stock assessment value of $M = 0.045$ is within the 95% CIs and there is very little difference in total likelihood at $M = 0.045$ and the MPD estimate. For these stocks the continued use of $M = 0.045$ in the stock assessment models is justified.

The updated average posterior distribution for *M* was almost identical to the distribution used in the 2014 MSE. The updated posterior distribution for *h* has shifted lower and now includes a higher proportion of low values for which there is an associated increased risk. However, for the base operating model the HCR has moved from being extremely safe with almost zero risk to being very safe with low risk. It is only for the Ricker operating model that any risk for the HCR exceeds 10% (being 11% for depletion risk). For the operating model with 20% positive bias in the estimators of stock status and beginning of year current biomass, both depletion risk and LRP risk are below 10% for the HCR. For all models, the LRP risk (which is more serious than depletion risk) is less than 10%.

Despite the lower updated estimates of h, the HCR is still very safe according to the estimates of the performance indicators, depletion risk, and LRP risk. The HCR, LRP, and target biomass range should therefore not be altered.

Introduction

In 2014, four New Zealand orange roughy stocks were assessed (Cordue 2014a). Subsequently, an orange roughy harvest control rule (HCR) was developed as part of a management strategy evaluation (MSE) (Cordue 2014b). It is a generic control rule that was developed to be applicable to any of the New Zealand orange roughy stocks which have fisheries based primarily on mature fish. The HCR is used for management purposes for the three orange roughy stocks that have Marine Stewardship Council (MSC) certification (East & south Chatham Rise (ESCR), Northwest Chatham Rise (NWCR), and ORH7A including Westpac Bank). In 2014, these stocks were assessed to be at or above the lower bound of the target biomass range (30-50 % B_0) developed in the MSE (although for ESCR the point estimate was 29.6% B_0). The fourth stock, Mid East coast (MEC), was assessed to be some way below 30% B_0 and was not put forward for certification.

In 2017, a fifth orange roughy stock (Puysegur) was assessed (Cordue 2019a). The Puysegur fishery has not been put forward for MSC certification and it is not managed using the HCR. However, it does provide an additional set of stock assessment results which include an estimate of natural mortality (M). There have also been updates of the stock assessments for ESCR & NWCR (Dunn & Doonan 2018, Cordue 2018), ORH7A (Cordue 2019b), and MEC (Cordue 2017). These updates provide revised estimates of *M* and stock-recruitment steepness (h) which are relevant to the review of the HCR.

In the 2014 MSE, the performance of the HCR was evaluated over a matrix of *M* and *h* values. Together with posterior distributions for *M* and *h* from the 2014 stock assessments, this allowed Bayesian estimation of various performance indicators for the HCR (e.g., the probability of mid-season spawning biomass being below the LRP (20% B_0)). With the passage of time, a fifth stock assessed, and with updated assessments for the four stocks assessed in 2014, it is timely to review the HCR to see if it still has acceptable performance given the updated estimates of *M* and *h.*

The LRP and target biomass range for orange roughy were also an output of the MSE and dependent upon the posterior distributions of *M* and *h.* Therefore, the LRP and target biomass have also been reviewed given the revised estimates of *M* and *h.*

Methods

There were three steps necessary to review the acceptability of the HCR given the revised estimates of *M* and *h.* The first step was obtaining the posterior distributions of *M* and *h* for each of the updated stock assessments. Some of the required runs (i.e., estimating *M*, or *M* and *h*) had not been done for some of the stocks. Given the revised posterior distributions, B_{MSY} and the LRP were re-estimated to check if the LRP or the target biomass range needed to be revised. Finally, the performance indicators for the HCR were re-estimated for the base operating model and the most important sensitivities to check that there was still adequate long-term performance.

Estimation of *M* and h

The same methods used in the 2014 stock assessments were used in the updated assessments and for Puysegur (Cordue 2014a, 2017, 2018, 2019a, b, Dunn & Doonan 2018). For this review, existing base models were modified to estimate *M* and/or *h* as required. The stock assessment models were all implemented in CASAL (Bull et al 2012) and used Bayesian estimation to obtain marginal posterior distributions for the desired parameters (*M* and/or *h*). For each run, three MCMC chains were produced, starting from random jumps away from the mode of the posterior distribution (MPD).

One in every one thousand samples were stored and the first one thousand stored samples were discarded as a “burn-in”. The remaining stored samples from all three chains were combined and used to produce the estimates (e.g., median and 95% credibility interval (CI)). Chains were generally each of length 15 million, so the estimates typically used 42 000 samples. The three chains were checked for consistency in the estimates (very similar medians and distributions). Chain diagnostic were generally excellent. The exception was for the estimation of *h* where the histograms were very similar but the medians were quite variable (see Appendix A). However, these chains were considered acceptable based on the running medians when the three chains were combined in parallel (see Appendix A).

In addition to the base models (which had informed priors for *M* and/or *h*), the MCMC estimation of *M* was also done using uniform priors. The MPD estimates of *M* and/or *h* were also calculated using the informed priors and the uniform priors.

The MCMC estimates presented in this report and the best source document(s) for each estimate are given below (BH = Beverton Holt):

Stocks	**Parameters**	**Prior**	**Assessed to fishing year**	**Sources**
ESCR, NWCR	*M*	Normal	2013-14	Cordue 2014a, b
		Normal	2017-18	Cordue 2018
		Uniform	2017-18	This report
Puysegur	*M*	Normal	2016-17	Cordue 2019a
		Uniform	2016-17	This report
MEC	*M*	Uniform	2017-18	This report
	M, *h* (BH)	Normal, beta	2013-14	Cordue 2014a, b
		Normal, beta	2017-18	Cordue 2017, this report
	M, *h* (Ricker)	Normal, lognormal	2013-14	Cordue 2014a, b
		Normal, lognormal	2017-18	This report
ORH7A	*M*	Uniform	2018-19	This report
	M, *h* (BH)	Normal, beta	2013-14	Cordue 2014a, b
		Normal, beta	2018-19	Cordue 2019b
	M, *h* (Ricker)	Normal, lognormal	2013-14	Cordue 2014a, b
		Normal, lognormal	2018-19	This report

The normal prior for *M* was N(mean=0.045, CV=15%) as used by Cordue (2014a, b) which was based on the estimate from Doonan (1994). The informed priors for *h* (beta, lognormal) were as used by Cordue (2014b).

A single average updated posterior distribution for *M* was formed by giving equal weight to each of the most recent assessment estimates over the five stocks (using the posteriors from the runs with the normal priors). The same approach was taken for *h* but there are only two relevant stocks (MEC and ORH7A). The parameter *h* is only estimated for MEC and ORH7A as they are the only stocks for which cohorts have been observed from when spawning biomass was estimated to be below 20% B_0 (the current LRP). In the 2014 MSE the same approach to forming a single posterior was taken but at that time there were only four stocks used for *M* (ESCR, NWCR, ORH7A, and MEC) and only a single stock for *h* (i.e., MEC).

As in the 2014 MSE, 5000 independent random samples of *M* and *h* were taken from the average posterior distributions for use in calculating Bayesian estimates of the LRP, B_{MSY}, and the performance indicators of the HCR.

Estimation of B_{MSY} and the LRP

In the 2014 MSE, B_{MSY} and the LRP (the maximum of 20% B_0 and 50% B_{MSY}) were determined over a grid of *M* and *h* values for the base operating model (Beverton Holt stock-recruitment relationship) and an alternative model with a Ricker stock-recruitment relationship (see Tables 1 and 2 for B_{MSY}).

Table 1: B_{MSY} (%B_0) as a function of *h* and *M* when a Beverton-Holt stock-recruitment relationship was assumed for the base operating model (Cordue 2014b). "–" denotes that B_{MSY} was not defined (i.e., the yield curve did not have a maximum).

	Natural mortality (*M*)						
Steepness (*h*)	**0.02**	**0.025**	**0.03**	**0.035**	**0.045**	**0.05**	**0.06**
0.25	45	45	45	45	45	45	45
0.30	45	44	44	44	43	43	43
0.35	42	42	41	41	41	41	40
0.40	40	40	39	39	39	38	38
0.50	35	34	34	34	34	34	33
0.60	31	30	30	30	29	29	29
0.75	25	24	24	23	23	22	22
0.90	17	17	16	16	15	14	13
1.00	6	4	3	–	–	–	–

Table 2: B_{MSY} (%B_0) as a function of *h* and *M* when a Ricker stock-recruitment relationship was assumed for the base model (Cordue 2014b).

	Natural mortality (*M*)						
Steepness (*h*)	**0.02**	**0.025**	**0.03**	**0.035**	**0.045**	**0.05**	**0.06**
0.25	48	48	48	48	48	48	48
0.30	48	47	47	47	47	47	47
0.35	46	46	46	46	46	46	46
0.40	45	45	45	45	44	44	44
0.50	43	43	43	42	42	42	42
0.60	42	42	41	41	41	41	40
0.75	40	40	39	39	39	39	38
0.90	38	38	38	38	37	37	36
1.00	37	37	37	37	36	36	35
1.20	36	36	36	35	34	34	34

In 2014, the Bayesian estimates of B_{MSY} and the LRP were calculated by weighting the implicit functions given by the *M*, *h* grids using the average posterior distributions of *M* and *h* (Cordue 2014b). The same approach was taken to update the estimates using the updated posteriors for *M* and *h.* For example, a sample from the posterior for B_{MSY} was calculated by running the 5000 M, *h* pairs through the implicit function (with interpolation done by cubic splines in R).

Estimation of the performance indicators for the HCR

In 2014, four performance indicators were estimated for the HCR: mean annual mid-season mature biomass; mean annual yield; the probability of the mid-season mature biomass being above the LRP (denoted LRP_p); and the probability of the mid-season mature biomass being above the lower bound of the biomass target range (denoted LB_p).

Bayesian estimation was used to account for the uncertainty in *h* and *M.* This was achieved in the same fashion as for B_{MSY} and the LRP, using interpolation via cubic splines over the grid of calculated values for fixed *h* and *M.*

The Bayesian posteriors of LRP_p and LB_p were used to derive two summary measures:

- **LRP risk**: the probability that the HCR will allow mid-season mature biomass to be below the LRP more than 5% of the time
- **depletion risk**: the probability that the HCR will allow mid-season mature biomass to be below the lower bound of the biomass target range, LB, more than 30% of the time

The probabilities are the proportion of *h*, *M* pairs where the HCR allows the poor performance with respect to the LRP or LB (e.g., "more than 5% of the time" means the proportion of years, over the long term, is greater than 5%).

For this report, the estimates of the performance indicators and the two summary measures were updated using the revised posteriors for *M* and *h.* This was done for the base operating model, the alternative Ricker model, and the most extreme sensitivity which assumed a 20% positive bias (within the operating model) for the estimators of stock status and beginning of year vulnerable biomass.

Results

Estimates of *M* and *h*

For the four stocks assessed in 2014 there are updated assessments with updated estimates of *M* (Table 3). In addition, there has also been an assessment of Puysegur orange roughy with an associated estimate of *M* (Table 3). The updated estimates of *M* are very similar to the previous estimates except for ESCR which has seen a reduction in the median estimate from 0.037 to 0.034 (Table 3). MEC has the lowest and most precise estimate of *M* and stands out as being different from the other four stocks in this regard (Figure 1).

There are some consistent patterns across the stocks for the *M* estimates that use a normal or uniform prior. The effect of the prior is to pull the estimates (that use the uniform prior) upwards towards the mean of the normal prior (0.045) for both the MPD estimates and also the MCMC estimates (Table 3). There is no clear pattern across stocks for the relationship between the MPD and MCMC estimates. For NWCR and Puysegur the MPD estimates are close to the corresponding median MCMC estimates. However, for the other three stocks the median MCMC estimates are much lower than the corresponding MPD estimates (Table 3).

The updated average posterior for *M* is almost unchanged from the previous average posterior for *M* (Table 4, Figure 2).

Table 3: MCMC and MPD estimates of *M* (multiplied by 100 to avoid leading zeroes). The median and 95% CI are given for the MCMC estimates. "Previous" refers to the estimates from 2014 where a normal prior on *M* was used. "Update" refers to the most recent estimates (which used the normal prior on M). "Uniform" denotes that a uniform prior was used for *M* in the updated assessment. There was no previous assessment for Puysegur.

	NWCR		ESCR		MEC		ORH7A		Puysegur	
	MCMC Med.	95% CI	MCMC Med.	95% CI	MCMC Med.	95% CI	MCMC Med.	95% CI	MCMC Med.	95% CI
Previous	4.1	3.3-5.1	3.7	2.7-4.8	3.2	2.8-3.7	3.8	3.1-4.7	–	–
Update	4.0	3.1-5.0	3.4	2.5-4.6	3.2	2.7-3.6	3.7	3.0-4.5	4.0	3.1-5.0
Uniform	3.6	2.7-4.9	2.8	2.1-3.8	3.0	2.6-3.5	3.3	2.6-4.3	3.7	2.8-4.9
	MPD		**MPD**		**MPD**		**MPD**		**MPD**	
Normal	4.1		4.4		3.9		4.2		4.1	
Uniform	3.4		4.0		3.5		4.0		3.9	

Table 4: The previous and updated MCMC estimates for *M* (median and 95% CI from the average posterior distributions for all stocks assessed).

	Med.	95% CI
Previous	0.037	0.028-0.049
Update	0.036	0.027-0.048

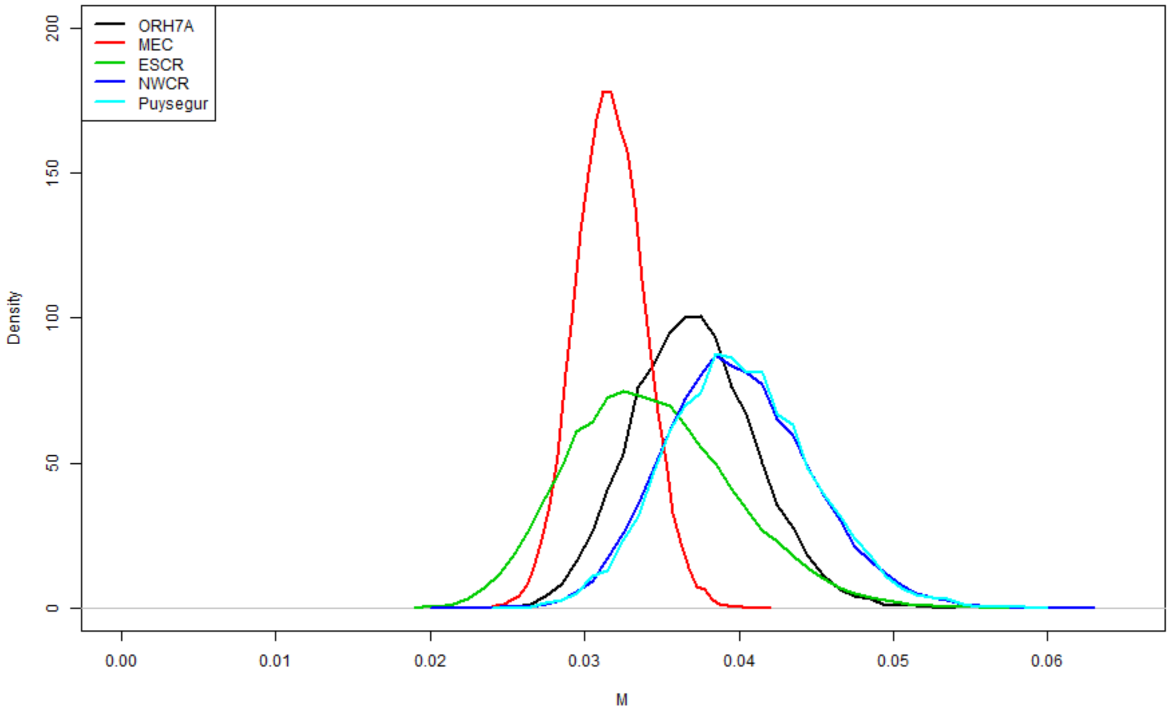

Figure 1: The most recent posterior distributions for *M* for the five orange roughy stocks assessed.

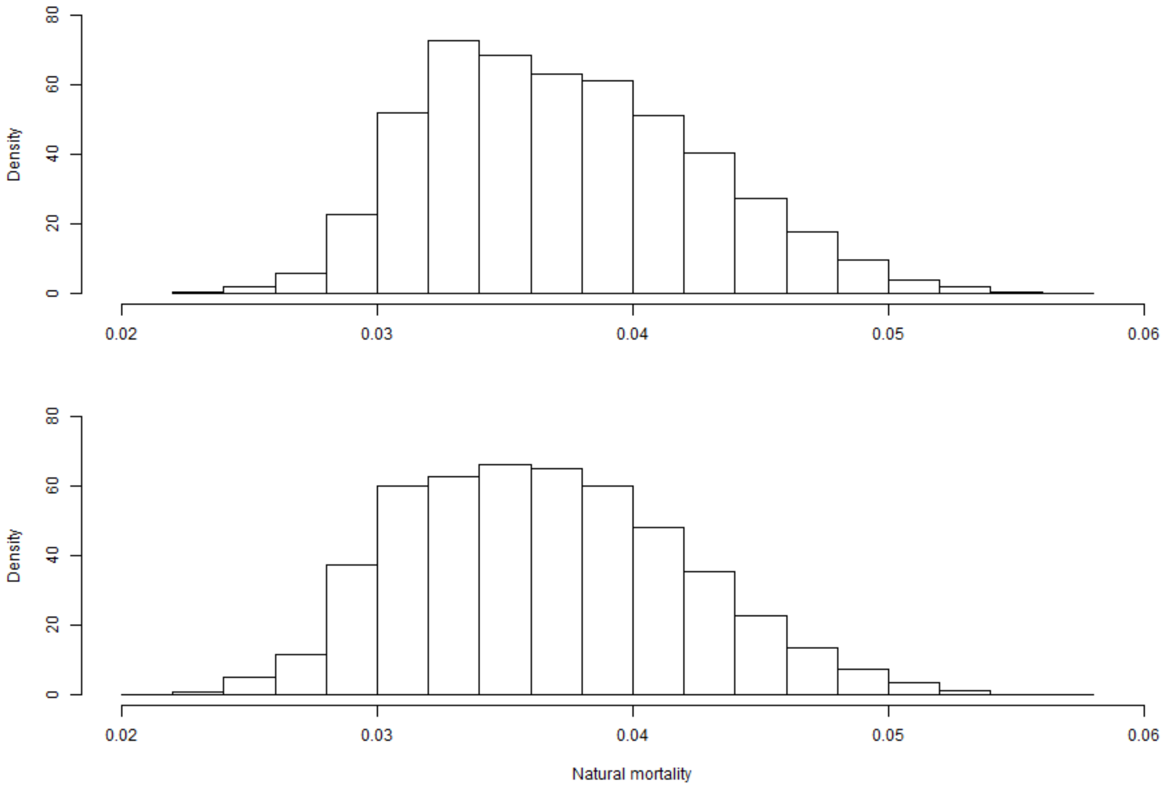

Figure 2: The average posterior distributions for *M* from 2014 (top) and the most recent assessments (bottom).

In 2014, MEC was the only stock for which cohorts spawned at low stock size had been observed. Therefore, the posterior distribution of *h* for MEC was used in the MSE. Since 2014, there has been an update of the MEC assessment and also an assessment for ORH7A where a number of cohorts spawned at low stock size have now been observed.

The updated Beverton Holt and Ricker estimates of *h* for MEC are lower than the previous estimates (Table 5). The estimates for ORH7A are similar to the previous estimates for MEC and higher than the current MEC estimates (Table 5, Figures 3 & 5). The MPD estimates when an informed prior is used are larger than the corresponding MCMC estimates and very much larger than the MPD estimates with a uniform prior (Table 5).

The updated estimates of *h* fall within the 95% CI estimates made in 2014, noting that all of the 95% CI (2014 and update) have a wide spread (Tables 5 and 6), which provides some insight into the uncertainty in the estimation of *h.*

The updated average posterior distributions for *h* are shifted to the left compared to the previous posterior distributions (which were just from MEC) (Table 6, Figures 4 & 6).

Table 5: MCMC and MPD estimates of *h* for a Beverton Holt and a Ricker stock recruitment-relationship. The median and 95% CI are given for the MCMC estimates. "Previous" refers to the estimates from 2014. "Update" refers to the most recent estimates. "Informed" denotes that a beta (Beverton Holt) or a lognormal (Ricker) prior was used for *h* (Cordue 2014b). "Uniform" denotes that a uniform prior was used for *h* in the updated assessment.

	Beverton Holt				Ricker			
	MEC		ORH7A		MEC		ORH7A	
	MCMC Med.	95% CI	MCMC Med.	95% CI	MCMC Med.	95% CI	MCMC Med.	95% CI
Previous	0.68	0.39-0.93	–	–	0.53	0.28-0.99	–	–
Update	0.53	0.26-0.88	0.61	0.31-0.91	0.40	0.23-0.78	0.56	0.27-1.2
	MPD		**MPD**		**MPD**		**MPD**	
Informed	0.77		0.73		0.74		0.62	
Uniform	0.59		0.33		0.58		0.32	

Table 6: The previous and updated MCMC estimates for *h* for a Beverton Holt and Ricker stock-recruitment relationship (median and 95% CI from the average posterior distributions).

	Beverton Holt		Ricker	
	Med.	95% CI	Med.	95% CI
Previous	0.68	0.39-0.93	0.53	0.28-0.99
Update	0.57	0.27-0.90	0.47	0.24-1.07

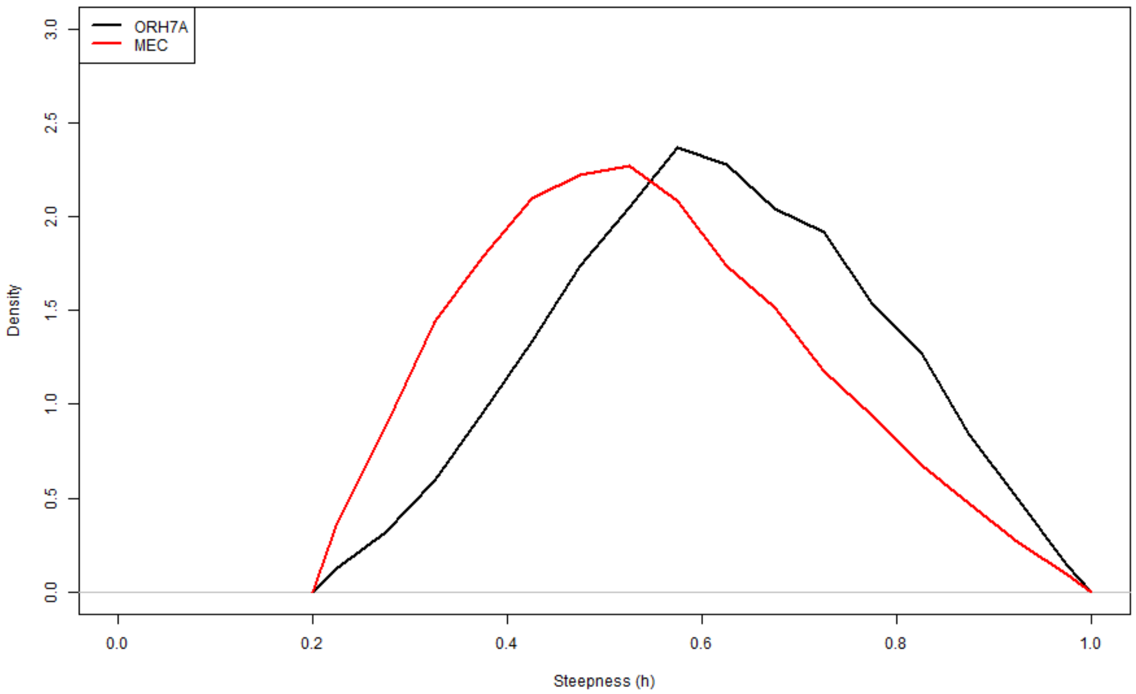

Figure 3: The most recent posterior distributions for Beverton Holt *h* for the two orange roughy stocks for which estimates were made.

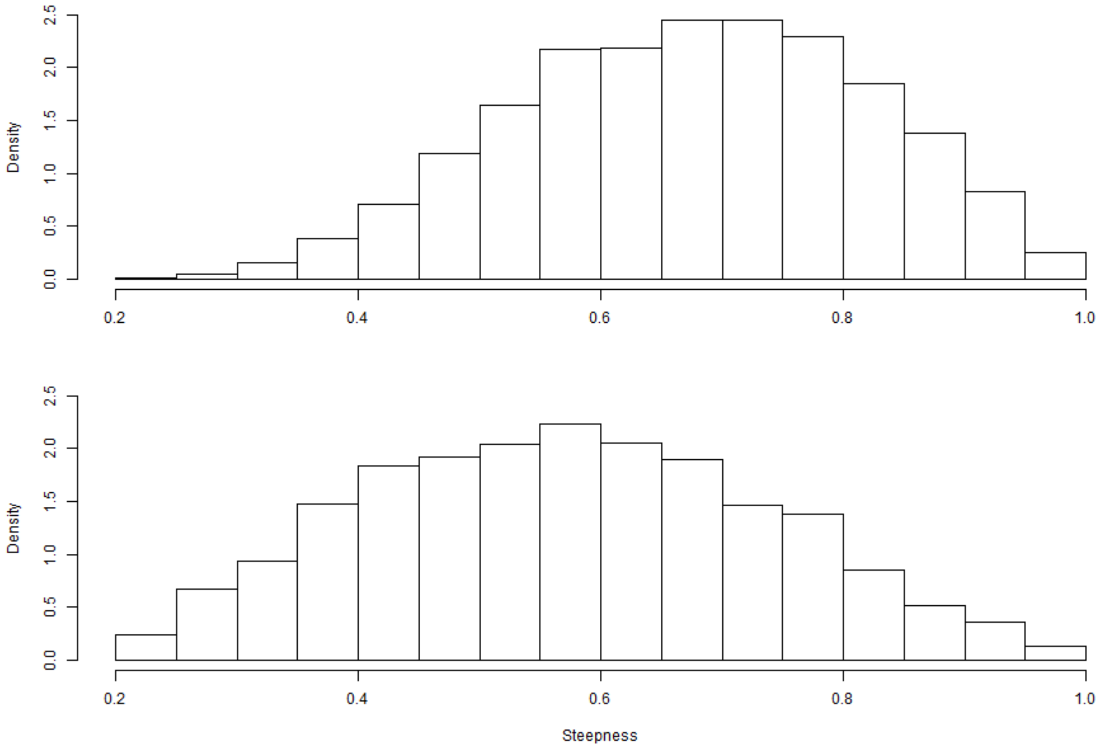

Figure 4: The average posterior distributions for Beverton Holt *h* from 2014 (top) and the most recent assessments (bottom).

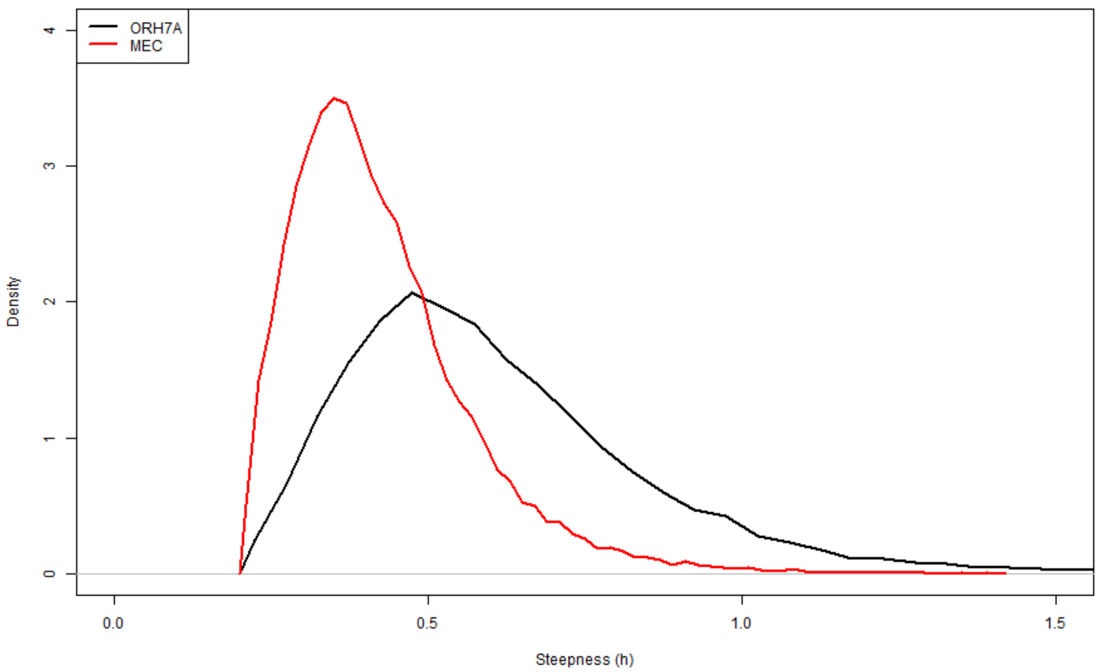

Figure 5: The most recent posterior distributions for Ricker *h* for the two orange roughy stocks for which estimates were made.

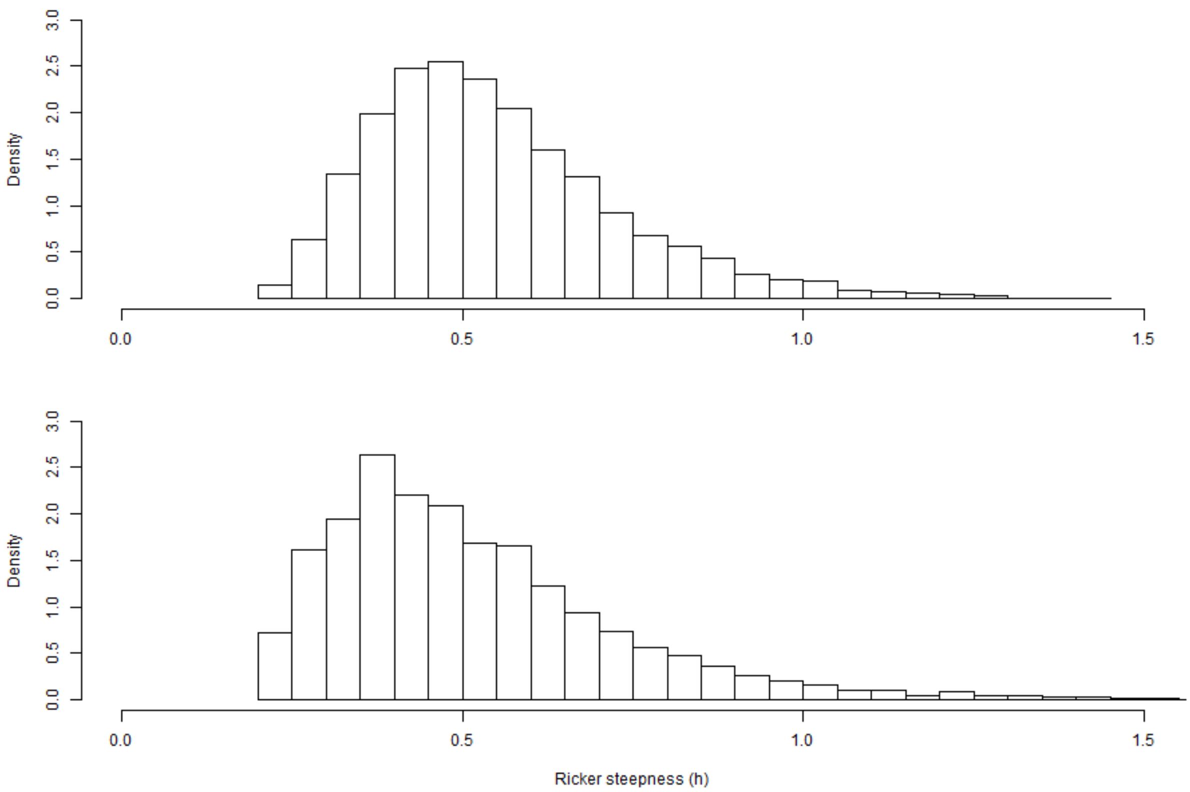

Figure 6: The average posterior distributions for Ricker *h* from 2014 (top) and the most recent assessments (bottom).

Estimates of B_{MSY} and the LRP

There was a substantial shift to the left in the average posterior distribution of Beverton Holt *h* and consequently the estimate of B_{MSY} has increased from 26% B_0 to 31% B_0 (for a Beverton Holt stock-recruitment relationship) (Table 7). The left shift in the updated Ricker *h* posterior was less marked and there is only a small change in the estimate of B_{MSY} for a Ricker stock-recruitment relationship (Table 7).

The updated estimates of the LRP are very similar to the previous estimates with the 95% CIs not much above 20% B_0 (Table 8).

Table 7: The previous and updated Bayesian estimates for B_{MSY} (%B_0) for a Beverton Holt and Ricker stock-recruitment relationship.

	Beverton Holt		Ricker	
	Med.	**95% CI**	**Med.**	**95% CI**
Previous	26	12-39	42	37-47
Update	31	16-45	43	36-48

Table 8: The previous and updated Bayesian estimates for the LRP (%B_0) for a Beverton Holt and Ricker stock-recruitment relationship.

	Beverton Holt		Ricker	
	Med.	**95% CI**	**Med.**	**95% CI**
Previous	20	20-20	21	20-24
Update	20	20-22	22	20-24

Estimates of the performance indicators for the HCR

The lower estimates of *h* are reflected in wider 95% CIs for the performance indicators calculated for the HCR (Table 9). The median estimates were unchanged except for mean yield which was reduced by 0.2% B_0 in each case (Table 9). The estimates of LRP and depletion risk have increased for the base model but are still no more than 5% (Table 10). However, for the Ricker model and the 20% positive bias model these risks are now at about 10% (Table 10).

Table 9: Previous and updated Bayesian estimates of the four performance indicators for the HCR for the base model (Beverton Holt), the Ricker model, and the model with a 20% positive bias for current stock status and beginning of year vulnerable biomass.

Base	Mean B (%B_0)		Mean yield (%B_0)		P(B > 20% B_0)(%)		P(B > 30% B_0)(%)	
	Median	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**
Previous	42	41–43	1.4	0.8–2.1	100	100–100	97	96–98
Update	42	29–42	1.2	0.3–2.0	100	94–100	97	43–98

Ricker	Mean B (%B_0)		Mean yield (%B_0)		P(B > 20% B_0)(%)		P(B > 30% B_0)(%)	
	Median	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**
Previous	42	32–50	1.4	0.4–2.9	100	97–100	96	61–100
Update	42	21–51	1.2	0.1–3.0	100	66–100	96	0–100

Bias	Mean B (%B_0)		Mean yield (%B_0)		P(B > 20% B_0)(%)		P(B > 30% B_0)(%)	
	Median	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**	**Median**	**95% CI**
Previous	35	34–37	1.5	0.8–2.3	100	100–100	78	76–87
Update	35	24–36	1.3	0.3–2.1	100	80–100	78	11–85

Table 10: Previous and updated Bayesian estimates of LRP and depletion risk for the HCR for the base model (Beverton Holt), the Ricker model, and the model with a 20% positive bias for current stock status and beginning of year vulnerable biomass.

		LRP risk (%)	**Depletion risk (%)**
Base	Previous	0	0
	Updated	3	5
Ricker	Previous	2	3
	Updated	8	11
Bias	Previous	0	1
	Updated	4	9

The wider 95% CIs for the performance indicators and the increased LRP and depletion risk are a consequence of the updated average posterior distributions for *h*. These distributions (Beverton Holt and Ricker) contain a greater proportion of low values for *h* compared to the previous distributions. It is the low values of *h* which create the LRP and depletion risk (e.g., Figure 7).

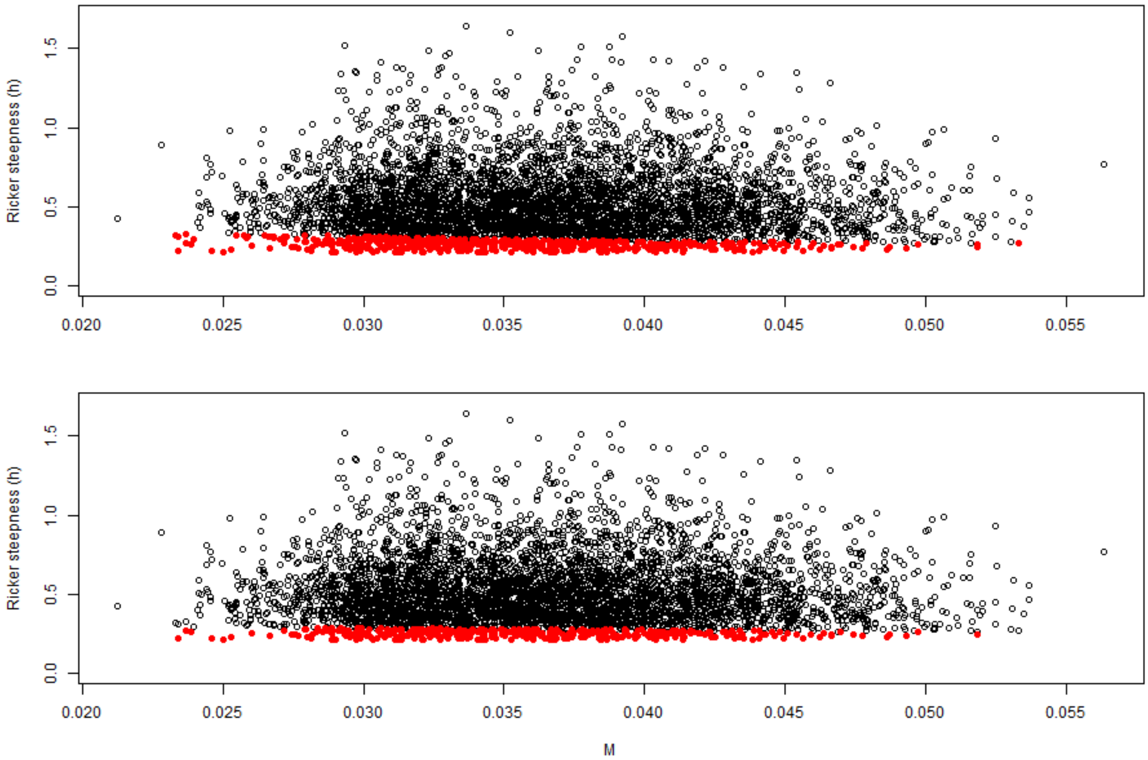

Figure 7: The samples from the updated distributions of *M* (x axis) and Ricker *h* (y axis) with red filled circles indicating that the given *M*-*h* pair contributes to depletion risk (top) or LRP risk (bottom).

Discussion and conclusion

Estimation of *M* and *h*

The updated estimates of *M* for each stock showed little change except for a downward revision for the ESCR stock. In the average posterior distribution, this was compensated for by the introduction of the Puysegur estimate which is the second highest amongst the five stocks (Table 3). The end result was little difference in the updated average posterior distribution for *M* compared to the 2014 distribution (Table 4).

The estimate of *M* for the MEC stock stands out from the others because of the low median (0.032) and the tight distribution (95% CI: 0.027-0.036; Figure 1). MEC is the only stock for which the 95% CI excludes the value of 0.045 which has been used as a fixed value for *M* in the stock assessments (see Table 3). The likelihood profile for *M* in the MEC model shows that the MPD estimate (0.039) is driven by the prior on M, the age frequency data and the YCS prior (Figure 8). This was the same as in the 2014 assessment (Cordue 2014a). For values of *M* which greatly reduce the proportion of old cohorts in the predicted age frequencies the penalty from the YCS prior is elevated because the model attempts to compensate for the lack of old cohorts by introducing unusually high YCS for those cohorts (see Cordue 2014a). The relatively low MPD estimate of *M* is a simple consequence of lots of old fish in the age frequency data (there is a substantial plus group at 120 years in the early age frequencies).

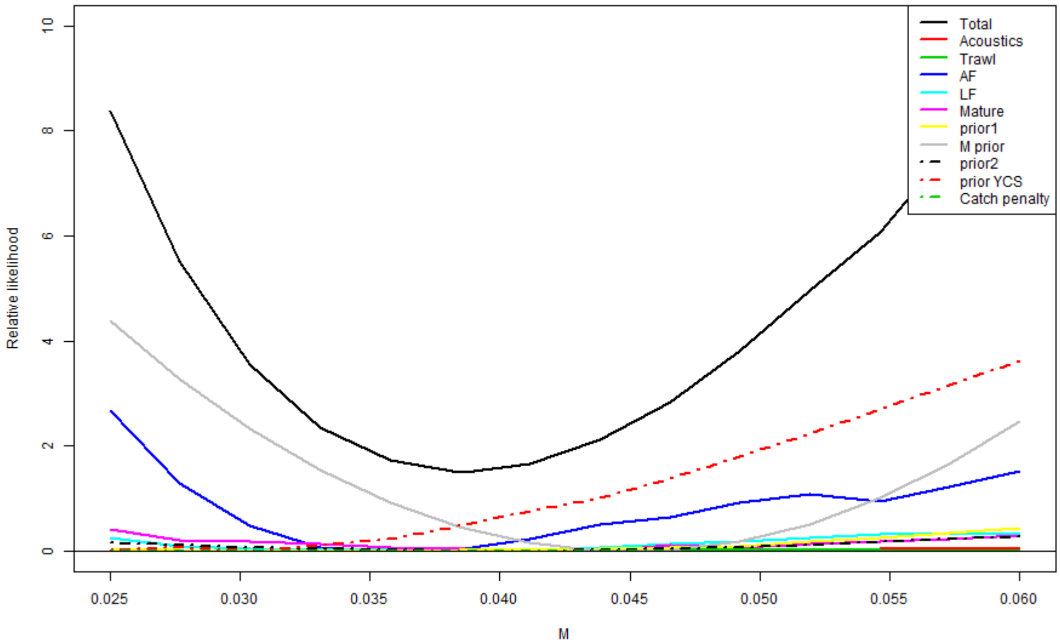

Figure 8: MEC likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (prior 1, prior 2) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

For the ESCR stock, the MPD estimate of *M* is 0.044 which is very close to the mean of the prior (0.045; Table 3) which suggests that there is little information in the data on *M.* This is confirmed by the likelihood profile for *M* which shows the only effect of note coming from the prior on *M* (Figure 9).

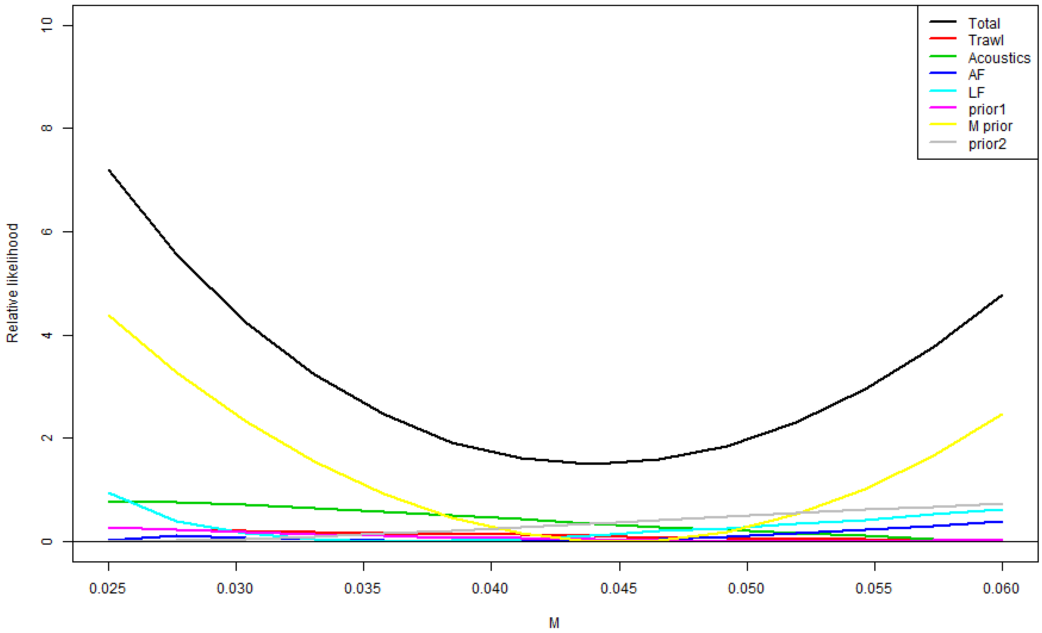

Figure 9: ESCR likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (prior 1, prior 2) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

The MCMC point estimate of *M* for ESCR (0.037) is much lower than the MPD estimate (0.044) despite the apparent lack of information in the data. When a uniform prior is used there is an even larger difference between the MPD estimate (0.040) and the MCMC median (0.028). The ESCR age frequencies do not have lots of old fish that would justify such low values of *M* (the earliest age frequency available for the Chatham Rise was used to estimate *M* outside the model at 0.045 – see Doonan 1994). It appears that the low MCMC estimate of *M* is being driven by structural aspects of the model, the data, and the estimator, rather than information about *M.*

It has previously been observed that estimates of B_0 and *M* can be strongly negatively correlated. This is true for the ESCR stock assessment model with the correlation coefficient (estimated from the first chain after burn in) equal to –0.88 (Figure 10). It is of note that the correlation coefficient for estimates of *M* and annual spawning stock biomass steadily increases from year to year (Figure 10). It is not a coincidence that the correlation coefficient is close to zero near the middle of the period where acoustic biomass estimates are used in the model (Figure 10). The samples from the joint posterior distribution are such that the spawning biomass trajectory is fitting the acoustic biomass estimates with everything from high values of B_0 and low values of *M* through to low values of B_0 and high values of *M* (Figure 11). The point is that the information in the data with regard to B_0 is not sufficient to eliminate “high” values of B_0 which necessarily imply “low” values of *M.*

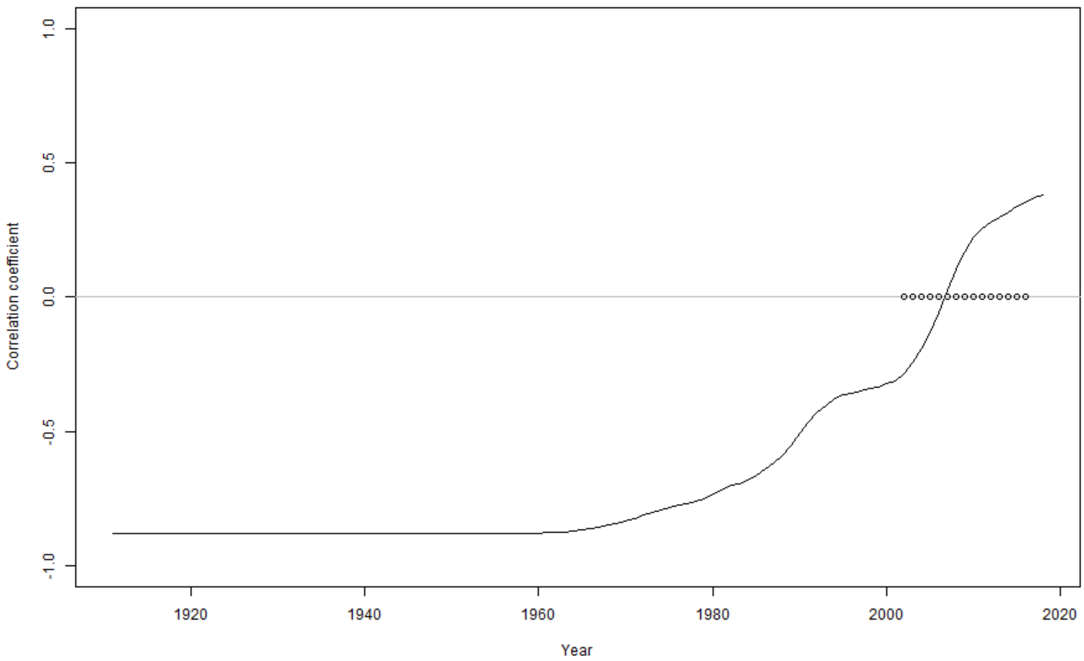

Figure 10: ESCR: the correlation coefficient between estimates of *M* and spawning stock biomass in each year. The years in which acoustic biomass estimates are fitted within the model are plotted as points on the zero line.

It may be that there is somewhat of an asymmetry with "low" values of B_0 being eliminated by the known catch history (there had to be enough biomass to allow the catches to be taken), but "high" values not eliminated because there is relatively little scale information in the available biomass indices (the primary information on scale comes from the acoustic *q* priors in combination with the value of the acoustic estimates). As a consequence, "high" values of *M* associated with "low" virgin biomass are excluded but "low" values of *M* associated with "high" virgin biomass are retained.

The MCMC estimates of *M* are being driven to some extent by a lack of information on B_0, as seen in the other stocks as well, but to a much lesser extent as reflected by lower correlation coefficients between *M* and B_0 (Table 11). The two biggest differences between the MPD and MCMC point estimates are seen for the two stocks with the largest correlation coefficients (Table 11).

Table 11: The correlation coefficient between *M* and B_0 for each stock. Also shown is the difference between the MPD and MCMC estimates of *M* (MPD estimate – MCMC estimate) for each stock.

Stock	Correlation coefficient	Difference between MPD and MCMC estimates
NWCR	–0.49	0.001
ESCR	–0.88	0.010
MEC	–0.57	0.007
ORH7A	–0.33	0.005
Puysegur	–0.37	0.001

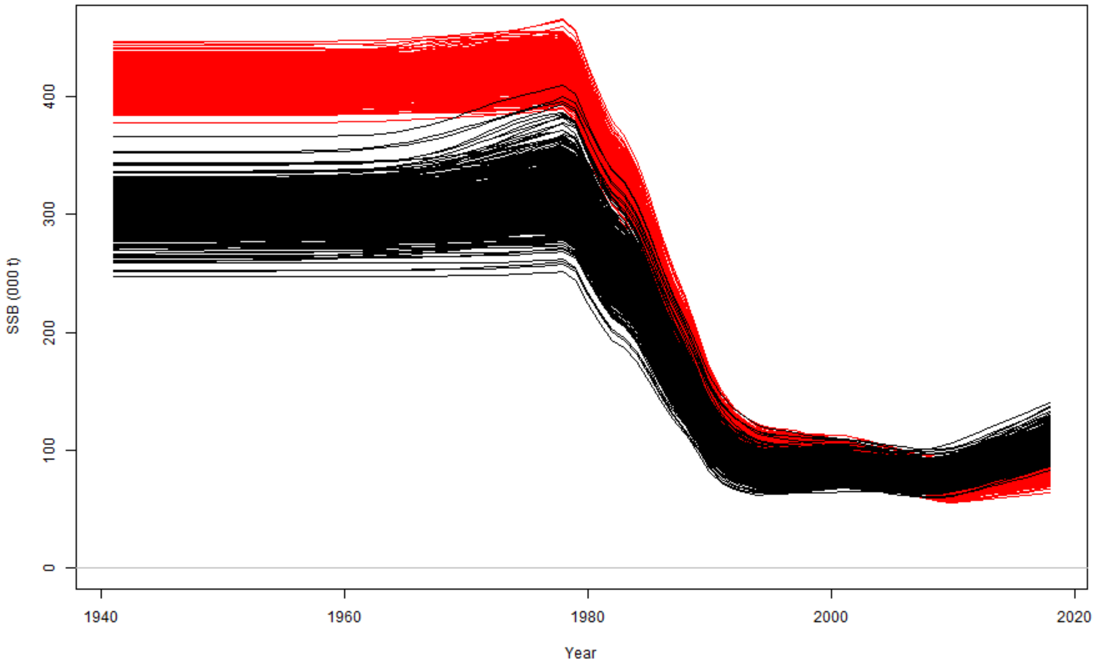

Figure 11: ESCR: spawning stock biomass trajectories for MCMC samples with *M* ≤ 0.026 (red lines) and *M* ≥ 0.045 (black lines).

The full set of likelihood profiles for *M* are given in Appendix B. As already noted, for ESCR the likelihood profile suggests that there is very little information on *M* in the data. Most of the apparent information for *M* for this stock is coming from the very high correlation with B_0 and the uncertainty associated with B_0. However, for all of the other stocks the age frequencies are a driver of the MPD estimate of *M* (Appendix B) and for these stocks there appears to be some information on *M* in the data.

As already noted, MEC stands out from the other stocks with the lowest and most precise estimate of *M.* Indeed, 0.045 is well outside the 95% CI for MEC (0.027–0.036) and it should probably not be used in the base MEC stock assessment model for this reason. MEC was alone in having a substantial plus group at age 120 years in its early age frequencies.

For the other stocks, the assumed stock assessment value of *M* = 0.045 is within the 95% CIs and there is very little difference in total likelihood at *M* = 0.045 and the MPD estimate (Appendix B). For these stocks the continued use of *M* = 0.045 in the stock assessment models is justified. Very much so for ESCR as it would be an error to estimate *M* in the model given the very high correlation between *M* and B_0 (which appears to lead to a spuriously low MCMC estimate of M).

The estimation of *h* is also problematic. For example, the likelihood profile for MEC Beverton Holt *h* has no strong signals from the data with the prior on *h* providing the largest contrast in negative log likelihood (Figure 11). From the data there is a signal from the age frequencies suggesting high *h* and a contrasting signal from the length frequencies suggesting low *h* (Figure 11). The MPD estimate (0.77) is much higher than the median MCMC estimate (0.53). Similarly, for ORH7A, there is little signal in the data (Appendix B, Figure B7). As for M, it appears that there may be some structural features of the model, the data, and the estimator that is driving the estimates away from the prior despite an apparent lack of information on *h.*

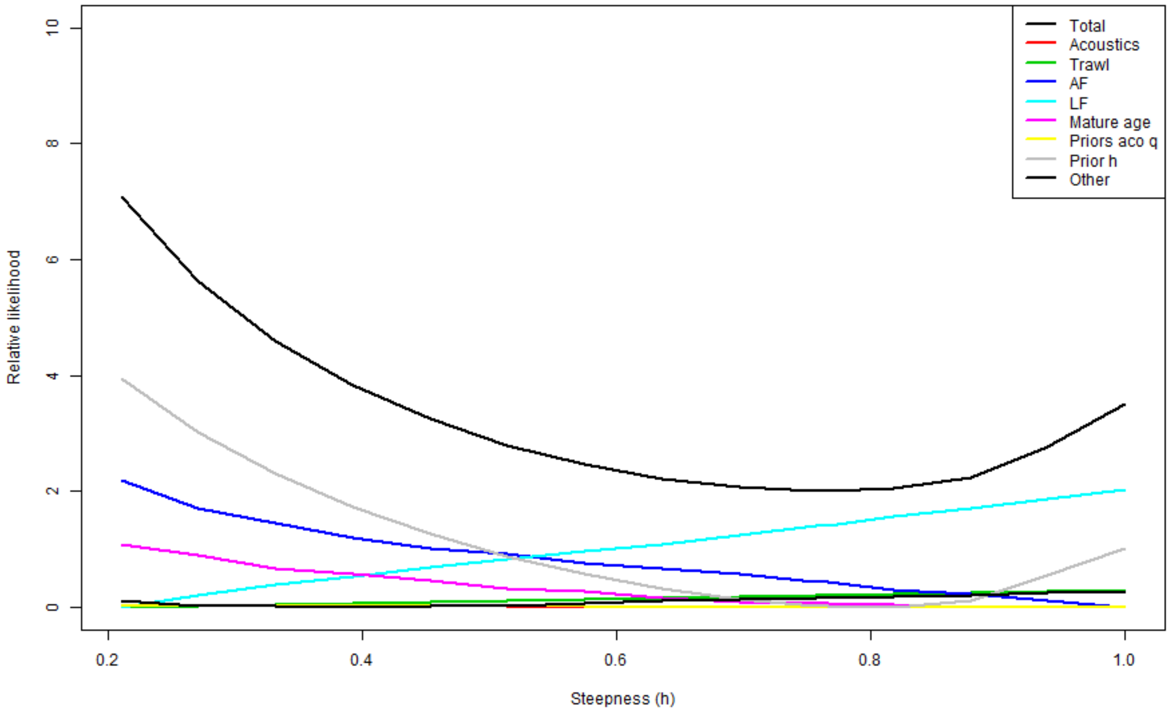

Figure 12: MEC likelihood profile for Beverton Holt *h.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *h.* The total is offset from the x-axis by an arbitrary value.

The estimation of *M* and *h* is problematic but it is necessary in order to estimate reference points and the performance of the HCR. MCMC estimation, with informed priors, is the best way to ensure that a plausible distribution of values are obtained over which to estimate and test.

Safety of the HCR

The updated average posterior distribution for *M* was almost identical to the distribution used in the 2014 Management Strategy Evaluation (Cordue 2014b). The updated posterior distribution for *h* has shifted lower and now includes a higher proportion of low values for which there is increased risk. However, for the base operating model the HCR has moved from being extremely safe with almost zero risk to being very safe with low risk. It is only for the Ricker operating model (a sensitivity) that any risk for the HCR exceeds 10% (being 11% for depletion risk). For the operating model with 20% positive bias in the estimators of stock status and beginning of year current biomass, both depletion risk and LRP risk are below 10% for the HCR. For all models, the LRP risk (which is more serious than depletion risk) is less than 10%.

Despite the lower updated estimates of h, the HCR is still very safe according to the estimates of the performance indicators, depletion risk, and LRP risk.

Acknowledgements

This work was funded by Deepwater Group Ltd, a non-profit organisation that represents quota owners of the deep water fisheries in New Zealand, including those holding almost all of the orange roughy quota.

Thanks to Geoff Tingley for reviewing and editing an earlier version of this report.

References

Bull, B; Francis, R.I.C.C; Dunn, A.; Gilbert, D.J.; Bian, R.; Fu, D. (2012). CASAL (C++ algorithmic stock assessment laboratory): CASAL User Manual v2.30-2012/03/21. *NIWA Technical Report* 135. 280 p.

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

Cordue, P.L. 2014b. A management strategy evaluation for orange roughy. ISL Client Report for Deepwater Group Ltd. 42 p.

Cordue, P.L. 2017. A stock assessment update to the end of the 2017–18 fishing year for the Mid-East Coast orange roughy stock. ISL Client Report for Deepwater Group Ltd. 45 p.

Cordue, P.L. 2018. A brief update of the ORH3B ESCR and NWCR stock assessments to the end of the 2016–17 and 2017–18 fishing years with application of the Harvest Control Rule in both years. ISL Client Report for Deepwater Group Ltd. 59 p.

Cordue, P.L. 2019a. A 2017 stock assessment of ORH 3B Puysegur. *New Zealand Fisheries Assessment Report 2019/20*. 35 p.

Cordue, P.L. 2019b. A 2019 stock assessment of ORH 7A. Draft New Zealand Fisheries Assessment Report 2019/xx. 38 p.

Doonan I.J. 1994. Life history parameters of orange roughy: estimates for 1994. New Zealand Fisheries Assessment Research Document 1994/19. 13 p.

Dunn, *M*.R.; Doonan, I.J. 2018. Assessment of the Chatham Rise orange roughy stocks for 2017. *New Zealand Fisheries Assessment Report 2018/59.* 60 p.

Appendix A: Chain diagnostics for MCMC estimation of *h*

For the estimation of natural mortality (*M*) for each assessment model, the three chains produced nearly identical median estimates in every case. No chain diagnostics are shown for *M.* However, for stock-recruitment steepness (*h*), which was estimated for MEC and ORH7A, the three chains had variable estimates of *h.* The marginal posterior distributions for each of the three chains for each stock and each type of stock-recruitment relationship are shown below (Figures A1, A3, A5, A7). The medians are not "nearly identical" although the three chains are clearly showing a similar picture. A running median and 95% CI for the three chains (combined in parallel) shows, in each case, that there is very little variation in the median and 95% CI over about the last 2000 stored samples (Figures A2, A4, A6, A8). This demonstrates that if the chains were extended that the estimates are very unlikely to change and therefore the chains are of adequate length.

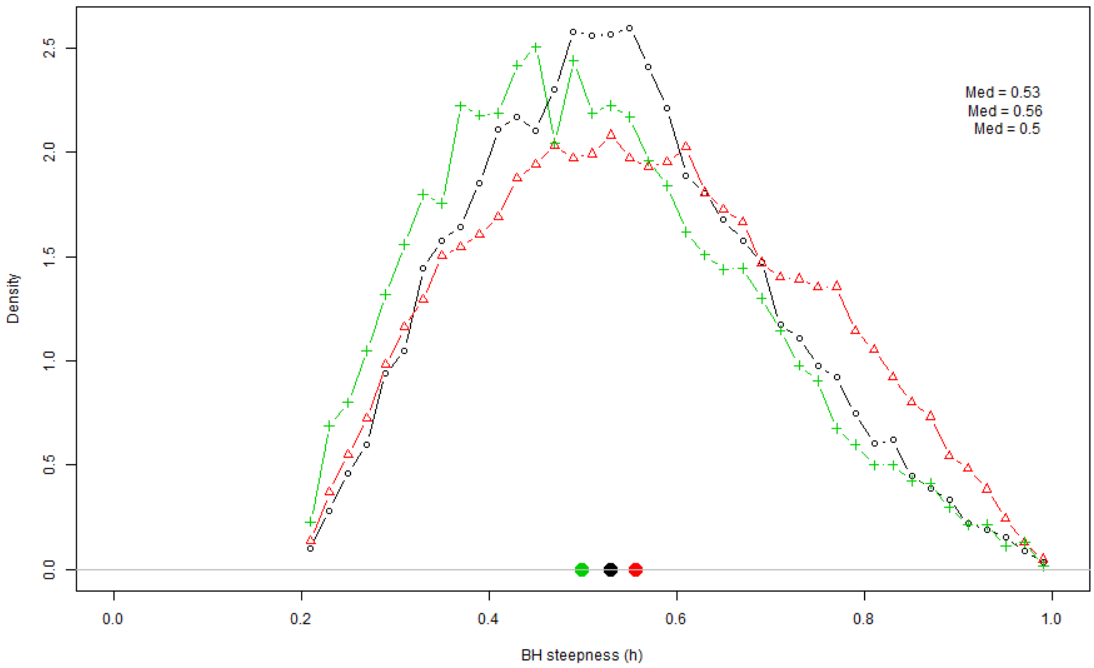

Figure A1: MEC marginal posterior distributions for steepness, *h*, when assuming a Beverton Holt stock-recruitment relationship. The distribution of the stored samples from each of the three chains (after burn-in) are shown together with the median for each chain (filled circles).

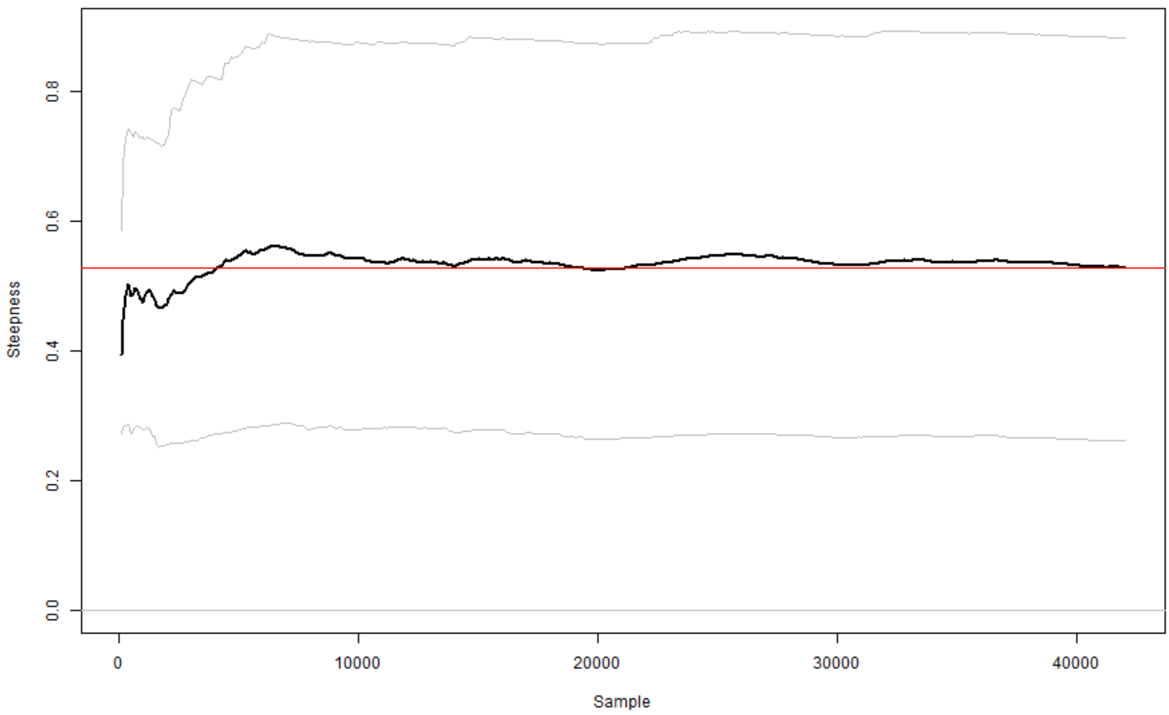

Figure A2: MEC: running median and 95% CI for steepness when assuming a Beverton Holt stock-recruitment relationship (three chains combined in parallel after burn in). A red line is plotted at the terminal median value.

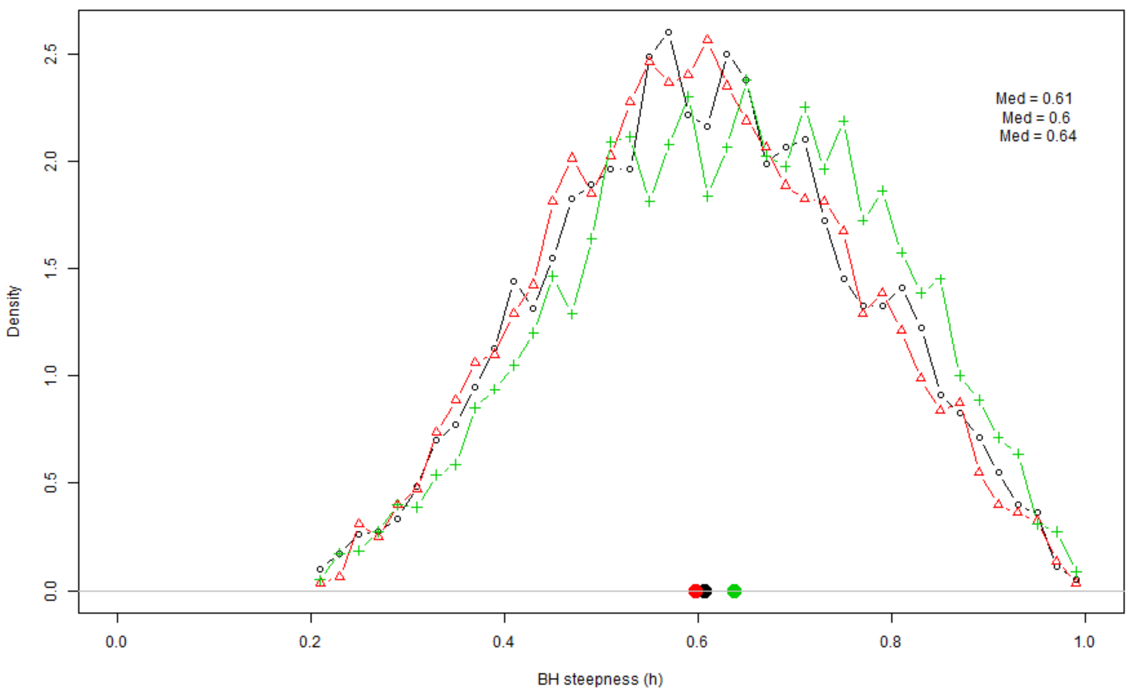

Figure A3: ORH7A marginal posterior distributions for steepness when assuming a Beverton Holt stock-recruitment relationship. The distribution of the stored samples from each of the three chains (after burn-in) are shown together with the median for each chain (filled circles).

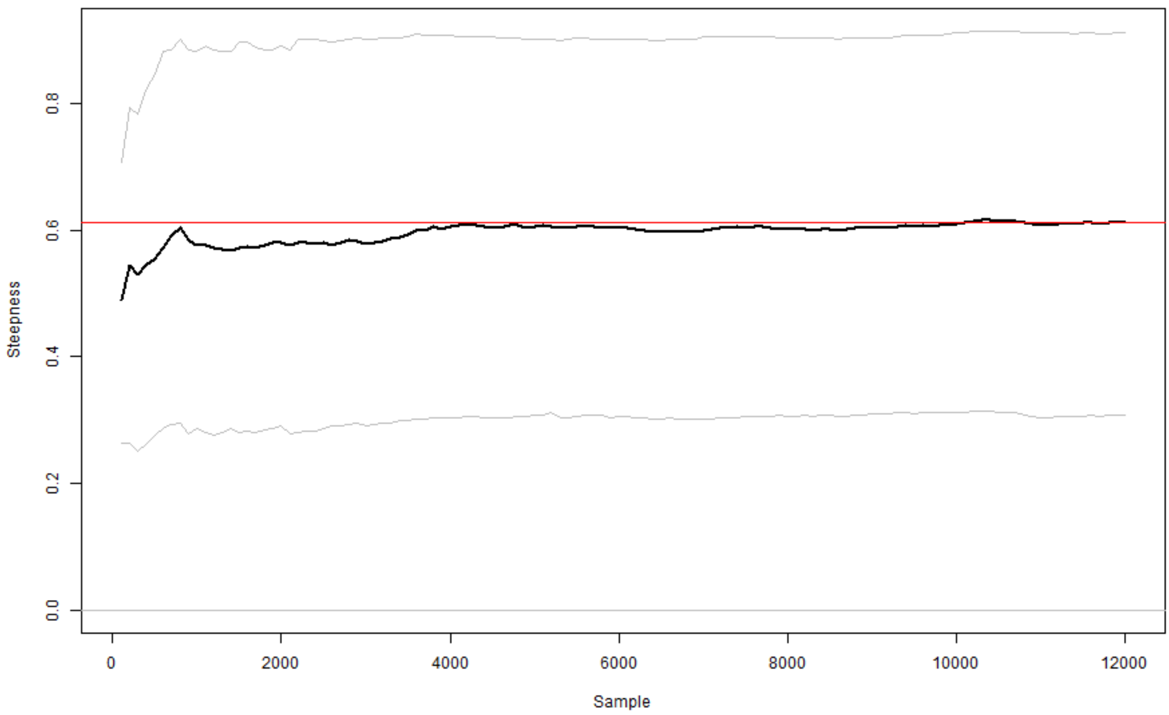

Figure A4: ORH7A: running median and 95% CI for steepness when assuming a Beverton Holt stock-recruitment relationship (three chains combined in parallel after burn in). A red line is plotted at the terminal median value.

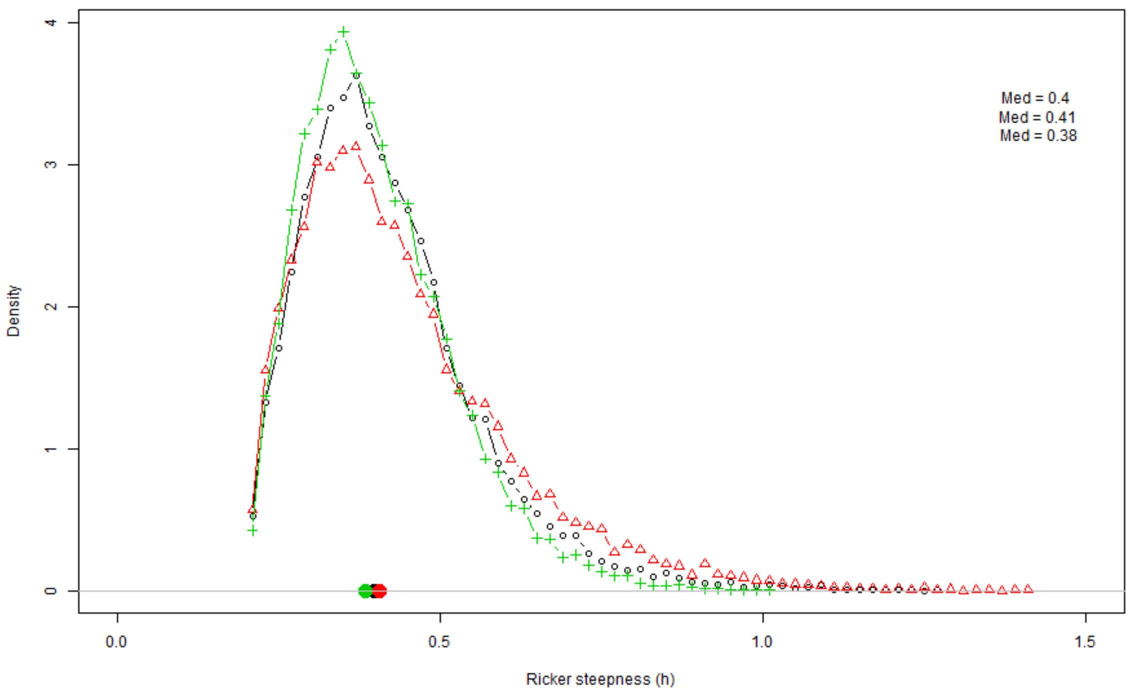

Figure A5: MEC marginal posterior distributions for steepness when assuming a Ricker stock-recruitment relationship. The distribution of the stored samples from each of the three chains (after burn-in) are shown together with the median for each chain (filled circles).

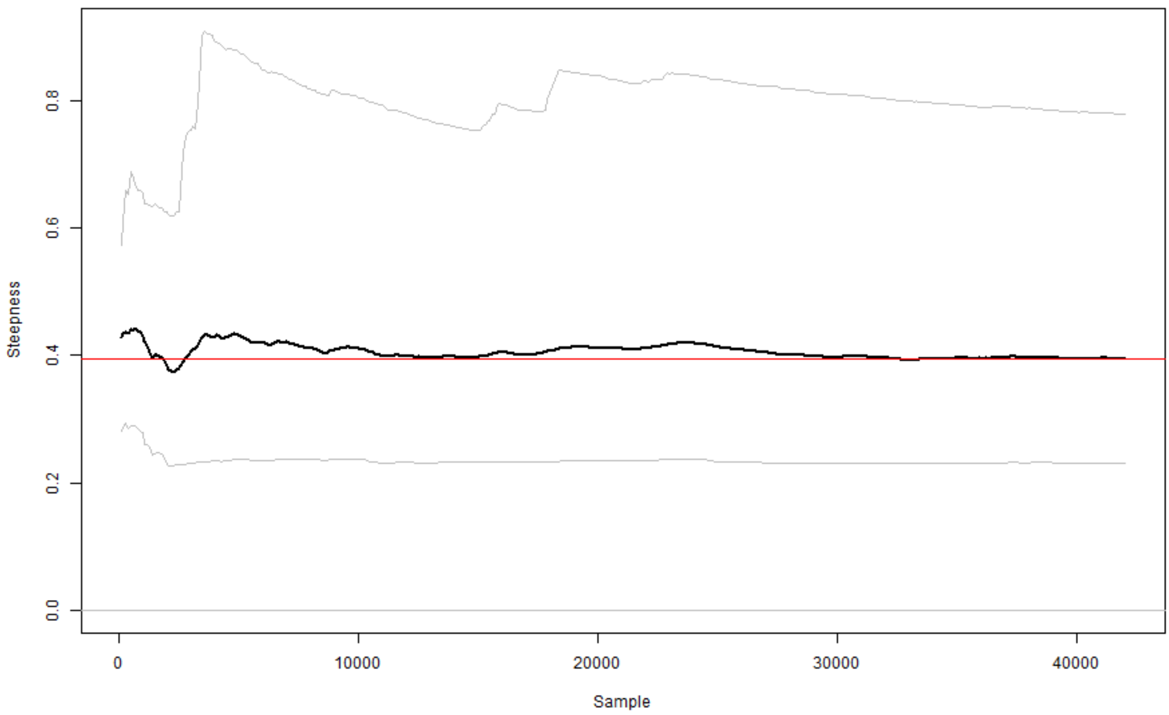

Figure A6: MEC: running median and 95% CI for steepness when assuming a Ricker stock-recruitment relationship (three chains combined in parallel after burn in). A red line is plotted at the terminal median value.

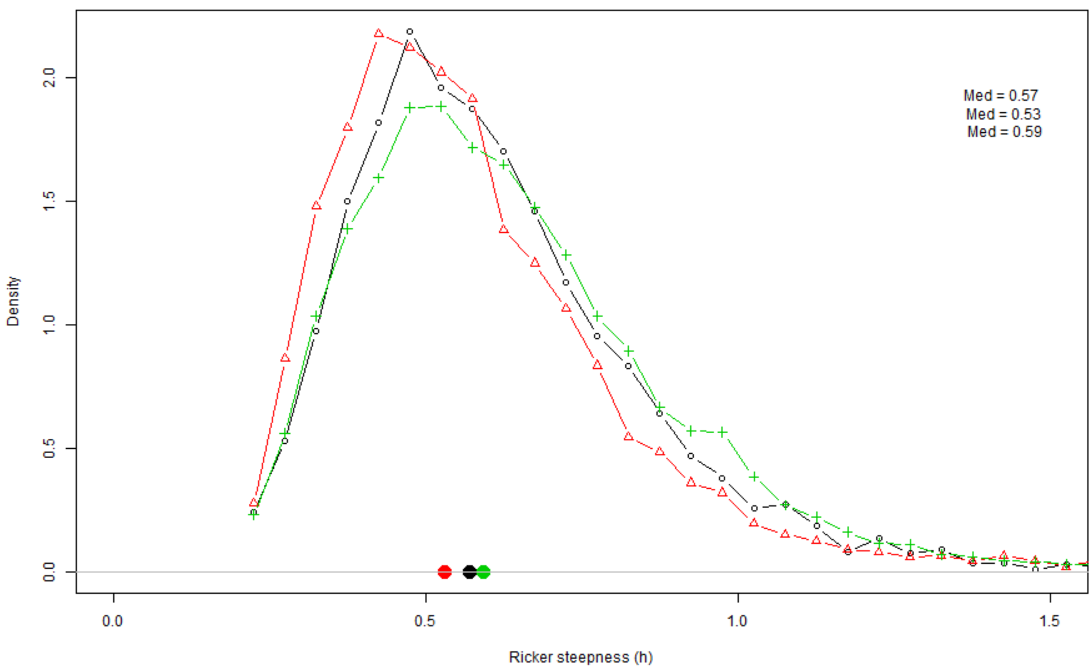

Figure A7: ORH7A marginal posterior distributions for steepness when assuming a Ricker stock-recruitment relationship. The distribution of the stored samples from each of the three chains (after burn-in) are shown together with the median for each chain (filled circles).

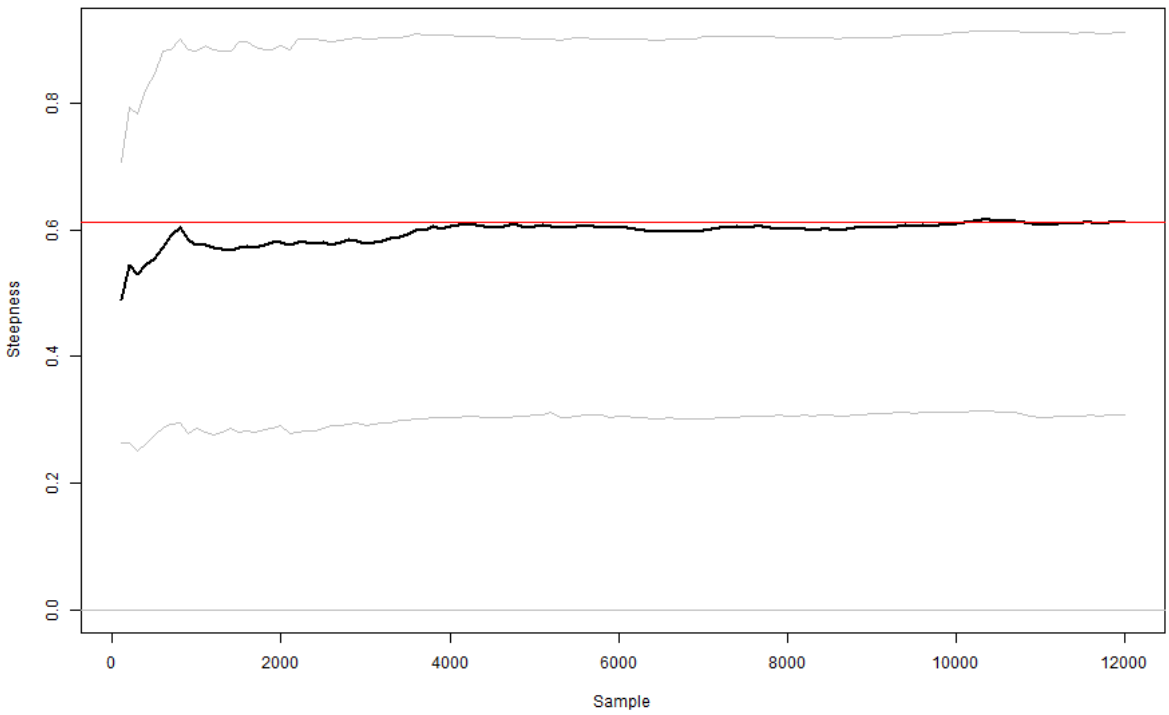

Figure A8: ORH7A: running median and 95% CI for steepness when assuming a Ricker stock-recruitment relationship (three chains combined in parallel after burn in). A red line is plotted at the terminal median value.

Appendix B: Likelihood profiles for *M* and *h*

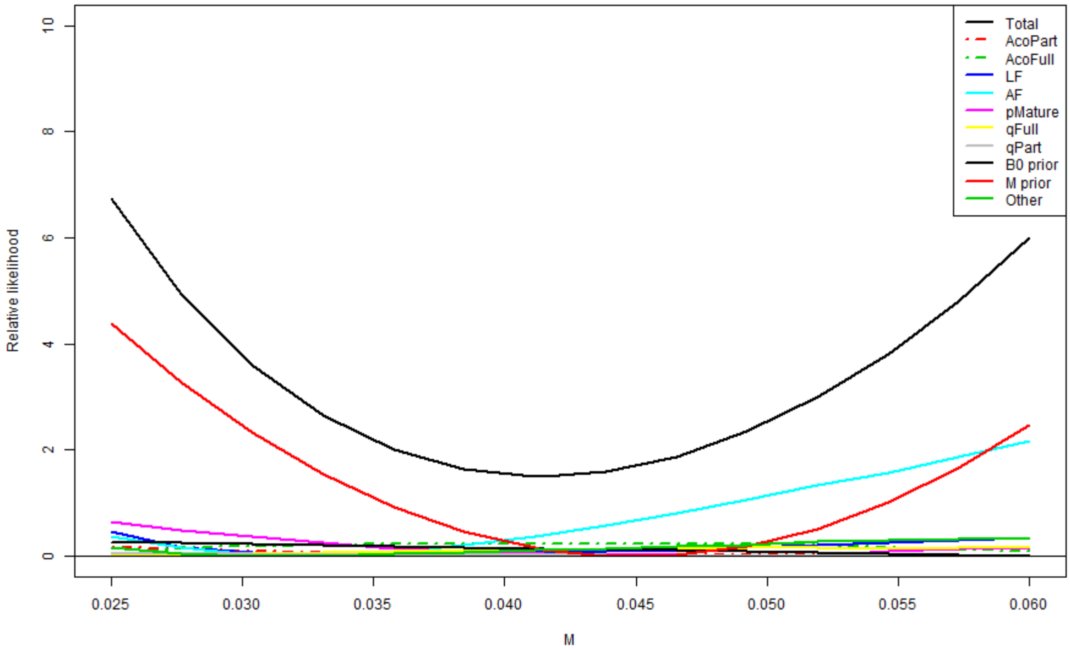

Figure B1: NWCR likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

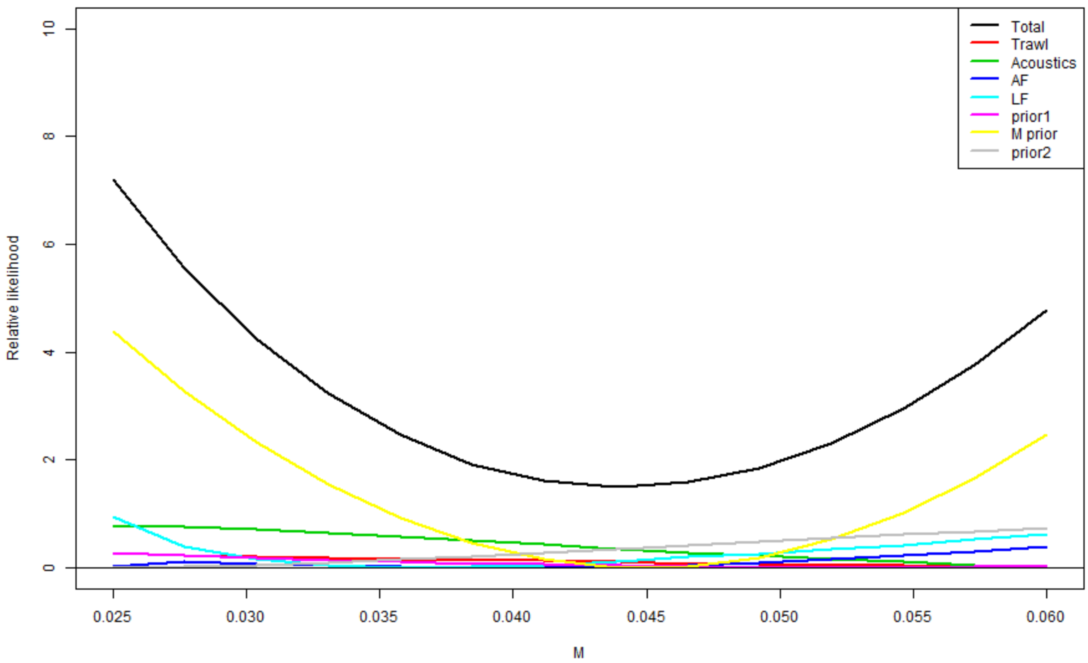

Figure B2: ESCR likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (prior 1, prior 2) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

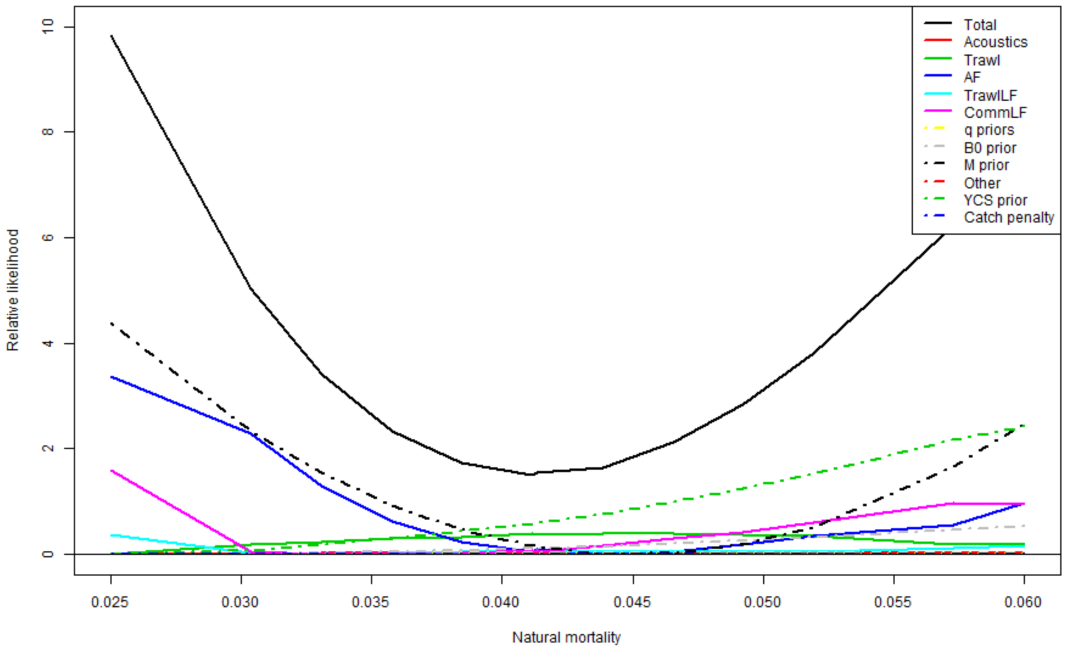

Figure B3: Puysegur likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

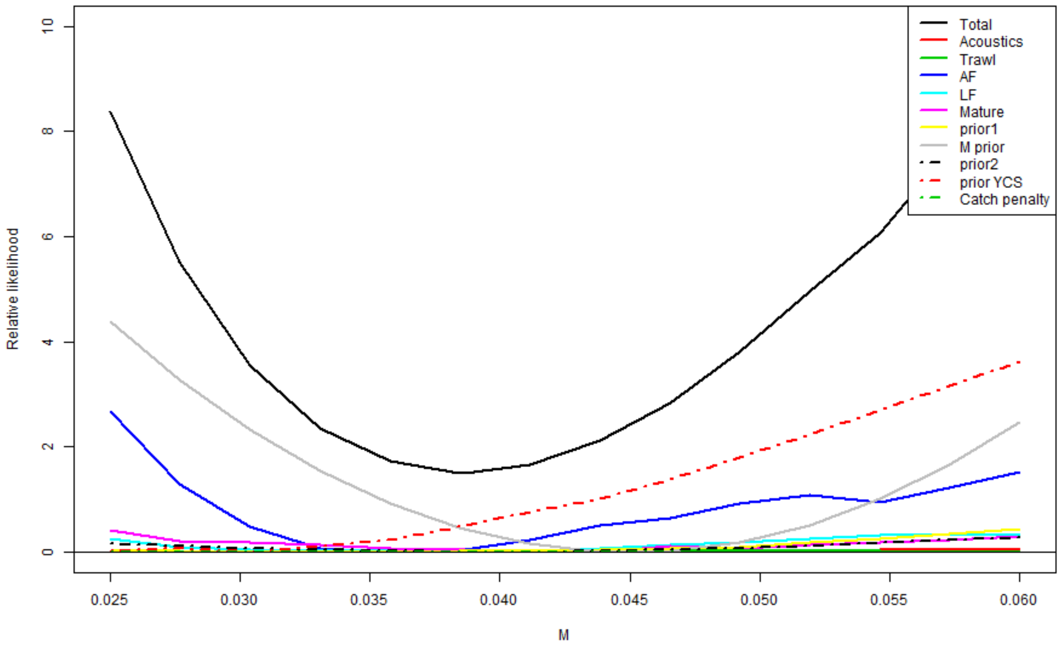

Figure B4: MEC likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (prior 1, prior 2) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

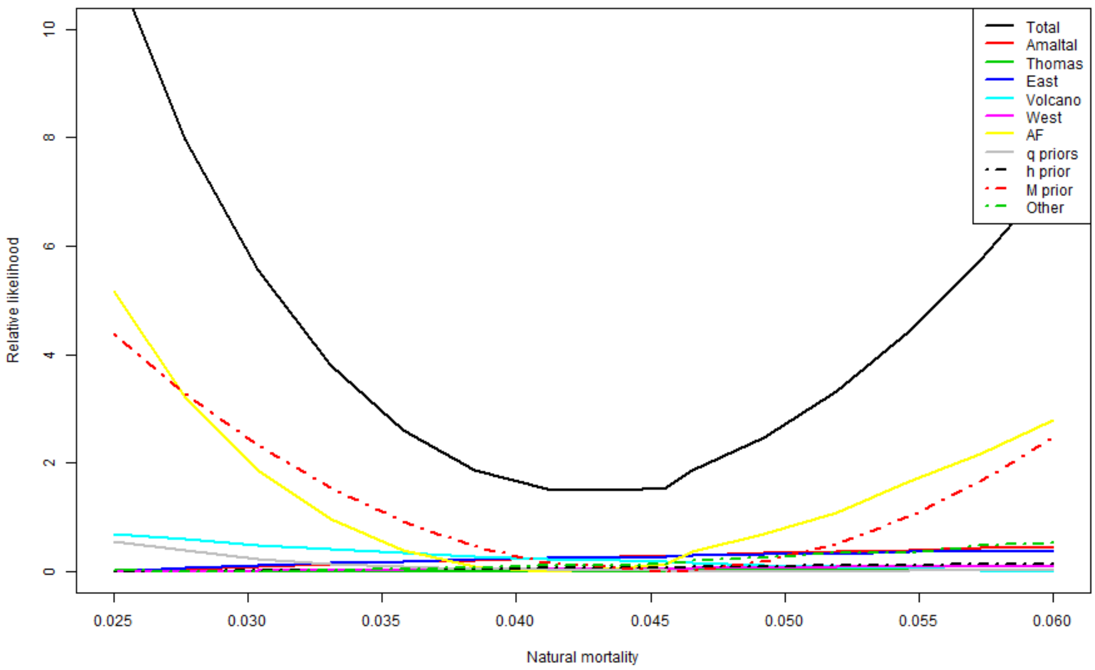

Figure B5: ORH7A likelihood profile for *M.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *M.* The total is offset from the x-axis by an arbitrary value.

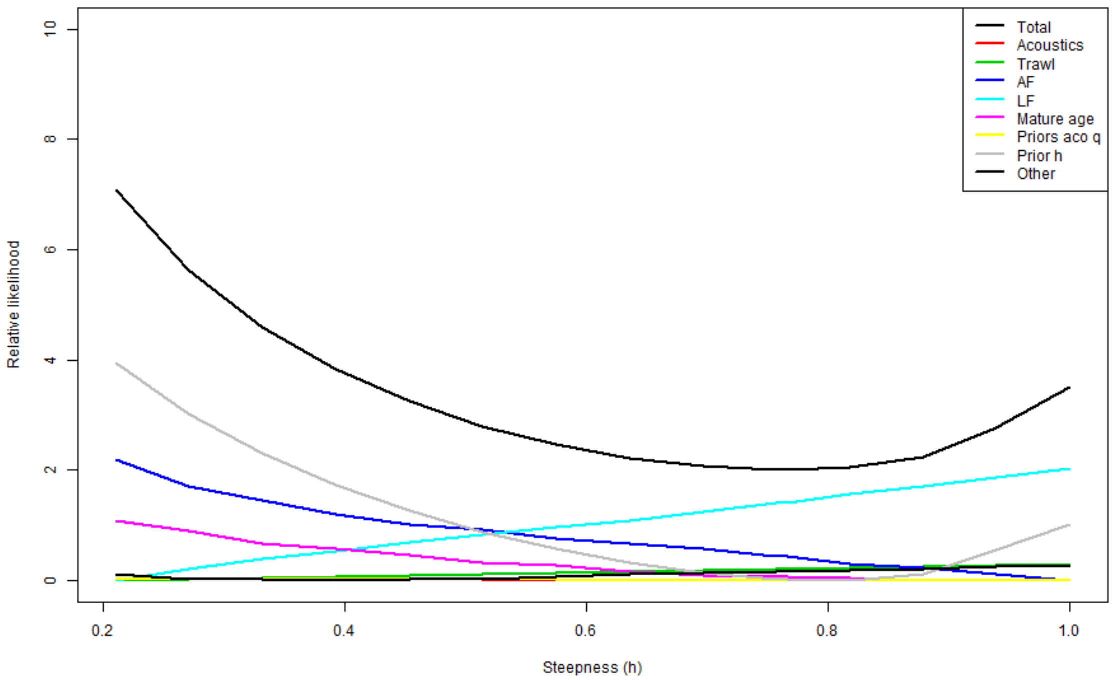

Figure B6: MEC likelihood profile for Beverton Holt *h.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *h.* The total is offset from the x-axis by an arbitrary value.

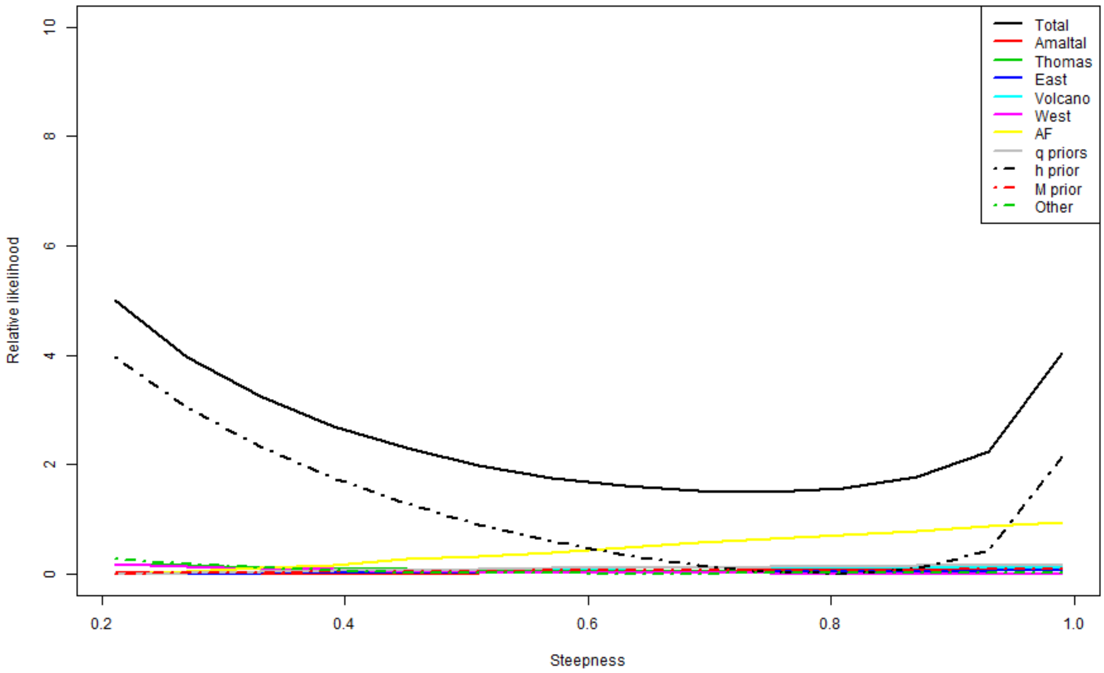

Figure B7: ORH7A likelihood profile for Beverton Holt *h.* The relative negative log likelihood for the total, each type of data, and prior or groups of priors (Other) are shown at fixed values of *h.* The total is offset from the x-axis by an arbitrary value.