MARINE AND ATMOSPHERIC RESEARCH
www.csiro.au

Initial exploration of options for inclusion of the close-kin data into the SBT operating model

Rich Hillary, Ann Preece, Campbell Davies, Mark Bravington, J. Paige Eveson, Marinelle Basson

CCSBT-ESC/1208/21

Prepared for the CCSBT Extended Scientific Committee for the 17th Meeting of the Scientific Committee 27-31 August 2012, Tokyo, Japan

CSIRO Marine and Atmospheric Research
Wealth From Oceans National Research Flagship
Castray Esplanade, Batter Point, Hobart TAS 7000, Australia
Telephone : +61 3 6232 5222
Fax : +61 3 6232 5000

Copyright and disclaimer

Important disclaimer

Contents

Abstract

The main analyses for the close-kin project have now been completed. This paper details an initial exploration of options for including the close-kin data into the SBT operating model. While the close-kin project produced its own estimation method (using reproductive data, Indonesian catch-at-age and length, and the close-kin data) we cannot directly include these estimates within the SBT OM. By maintaining the core ideas behind the independent estimator we demonstrate how we can transform and incorporate the close-kin data and a more realistic definition for the effective spawning population into the SBT operating model. With these potential refinements to the definition of the spawning population and the inclusion of the close-kin data in this form, the OM results indicate higher levels of adult abundance and survival probabilities. An in-depth exploration of the true statistical information content in the data in relation to key grid parameters suggests we need to rethink both the most recent grid and how we weight grid elements in future. In summary, the close-kin data are generally well fitted by the OM and do suggest a more optimistic level of spawning population depletion but more work is required to effectively handle the age/length structured nature of the reproductive dynamics of the population as well as the issue of sexual dimorphism.

1 Introduction

The main analyses for the close-kin project have now been completed - for an in depth review of the genetic sample sizes, quality control and matches (parent-offspring pairs or POPs) please see [1]. The project also produced a stand-alone adult-focussed “assessment” of recent adult SBT abundance. There are several structural and data-related issues as to why we cannot simply import the close-kin project abundance estimates directly into the current SBT operating model. This paper begins the process of formally including the close-kin data into the integrated SBT operating model.

To ensure that changes in the results and fits of the OM can be attributed to the appropriate suggested revision to the OM, we have made each potential change sequentially, involving four steps. The first involves a refinement of the definition of SSB or effective spawning population size to be conceptually consistent with the close-kin ”mini-assessment”; Step 2 involves including the close-kin data with the previous grid; Step 3 exploring alternative grid structures and weighting scenarios, given the fitting summaries arising from the revised definition of spawning population and the inclusion of the close-kin data; and Step 4 a full summary of the operating model outputs and data fits. In the final section we identify some of the issues that will require further consideration and work in light of this initial exploration of the implications of the close-kin results.

2 Data & methods

2.1 Form of the close-kin data

There is a lot of information in the close-kin data but for the task of using them in the OM then it is easier to consider simply the following:

- **Juveniles**: there are $i \in \Theta$ samples with a year of capture y_i, length-at-capture l_i, and either a distribution of ages conditional on length $p(a_i|l_i)$ (via the growth relationship) or direct age a_i in the case of a POP. Cohort/year of birth then being $y_i - a_i$.

- **Adults**: there are $j \in \Xi$ samples with a year of capture y_j, length-at-capture l_j, and either a distribution of ages conditional on length age $p(a_j|l_j)$ or direct age a_j in the case of a POP.

2.2 Age, length, reproductive capacity and sexual dimorphism

In reality, the probability of adult j being a parent of juvenile i will vary according to both adult abundance and the effective reproductive contribution of the sampled adult in the birth year of the paired juvenile sample: for equal reproductive success higher/lower levels of one age class would make them more/less likely to be parents; for the same abundance levels more/less effective spawners are more/less likely to be parents.

From [2] there is clear evidence that suggests both age and length-related factors influence what defines an effective spawner (onset of maturity, fecundity-at-length and so on). In the stand-alone close-kin estimator the population dynamics are age-based, but the key factors specifying effective spawning contribution are length-based and so both age and length are jointly (and elegantly) handled within the estimation framework. Also, males and females are handled separately as they appear to have different growth dynamics [1].

Within the current OM, the effective spawning population is defined as the biomass of fish age 10+. While this is a simple compromise that has served its purpose for MP development and evaluation, it is not consistent with the detailed empirical reproductive data [2], the direct information on effective reproductive contribution from the close-kin data and is not consistent with the estimation framework used in [1]. Nor does it fit with the close-kin genetic matches data as the measured age of 4 of the POPs (just less than 10%) was either 8 or 9 years old. Given this, a more appropriate definition would be that the effective reproductive capacity is a product of the time spent on the spawning grounds multiplied by the daily "output" measure (fecundity in females).

Within the close-kin estimator the assumption is made that residence time is effectively the same thing as selectivity within the Indonesian fleet, and the catch composition is used to estimate the residence time/selectivity relationship, as well as aiding the estimation of survival probabilities. Irrespective of whether or not this assumption about residence time on the spawning ground and Indonesian selectivity holds, we are not in a position to simply replicate this assumption within the current OM structure for the following reasons:

Firstly, Indonesian selectivity is age-based and is permitted some freedom to change from year to year. Secondly, the estimated selectivity at the older ages actually begins to decrease even with the reasonably informative (and elevated) estimates of M_{30}. Interpreting this as an effective decrease in the presence of the oldest fish on the spawning grounds (residence time) raises many questions, but the general interpretation is that these oldest fish are simply not as available to the long-lines set for the actual target species: the tropical tunas. So, while there must logically be a strong relationship between Indonesian selectivity and residence time on the spawning grounds, the current OM structure and time-varying selectivity-at-age estimates do not seem to support the hypothesis that this relationship is identical.

To test the structural feasibility of using the mean (over years) Indonesian selectivity as a proxy for residence time a simple example was programmed (no close-kin data included, just the change to the spawning population definition). Even with very accurate estimates of selectivity (i.e. very close to the estimated selectivity from the previous OM) it was not possible to get the model to run given the phasing process within the estimation routine, and the problems that a dynamic (at run time) definition of the spawning population has when interacting with recruitment deviates and so on. While this approach was feasible within the stand-alone close-kin estimator, the more complex form of the OM suggests that we would need to make significant changes to the estimation process and even then with no guarantee that we could end up model definition that was able to be estimated robustly (which is very important when it comes to generating the grid). At least at this stage, it therefore makes sense to define a static definition of the effective spawning population ogive from outside the model, and with the OM structure we can explore options around alternatives quite easily.

Sexual dimorphism is another process that is explored in the stand-alone close-kin estimator but is not currently included within the SBT OM. At least from the Indonesian measurements of age and length, it is apparent that males grow longer than females, though there is little significant support for sexual dimorphism in the other growth parameters [1]. At this stage, we are not in a position to directly include sexually disaggregated data within the OM but a first attempt to include some aspects of these differences was made when constructing the effective relative reproductive contribution-at-age. The relationship between female length and daily batch fecundity was $\propto l^{2.47}$; for males no data was available so the assumption was made that it was $\propto l$. There is also an apparent length-specific relationship between the relative time-frames of the spawning event and resting period [1, 2] that is the last factor required in defining per-capita effective reproductive output by length (and sex).

In this initial exploration, the maturity-at-length relationship (modelled as a logistic function) is used to define residence time, not the Indonesian selectivity. This maturity-at-length is then multiplied by the batch fecundity and spawning-resting ratios (and normalised) to give an effective spawning population ogive-at-length for each sex: $\phi_{s,l}$. Growth relationships (von Bertalanffy with individual L_∞) are then fitted to the male and female adult size-at-age data to obtain the distribution of length-at-age $p(l|a,s)$. The age-based counterparts of the $\phi_{s,l}$ are then calculated as follows:

$$\phi_{s,a} = \int \phi_{s,l} p(l|a,s) \mathrm{d}l. \tag{2.1}$$

In an age-based sexually dimorphic model we would simply use the relationships detailed in Eq. 2.1 but for the sexually aggregated OM the assumption was made of a simple 50/50 sex ratio (not too far off given the Indonesian data) and so the per-capita effective reproductive potential-at-age, ϕ_a, is just the renormalised average of the male and female versions. We are now in a position to define the (sexually aggregated) parental probability, π_{ij}:

For a given juvenile i, then the probability that an adult j (which can be compared to juvenile i) involves the total numbers in the adult population at the time of the juvenile's birth, $c_i = y_i - a_i$. At the time of capture the adult had length l_j and some associated distribution of possible ages, $p(a_j \mid l_j)$, which can easily be translated into a distribution of possible ages at the time of the juvenile's birth, $p(a_j - a_i \mid l_j)$. If $N_{y_i - a_i, a}$ are the OM numbers-at-age in the population at the time of the juvenile's birth then the parental probability is defined as:

$$\pi_{ij} = \frac{2\phi_{\mathbb{E}(a_j - a_i | l_j)}}{\sum\limits_a N_{y_i - a_i, a} \phi_a}, \tag{2.2}$$

and $\mathbb{E}(a_j - a_i \mid l_j)$ is the expected age of the adult at the time of the juveniles birth (given it was of length l_j when caught). The factor of 2 is there as we can detect both male and female parents.

A very simple pseudo-sexual alternative would be the following: assuming little sexual dimorphism in juvenile mortality and that the factors influencing adult mortality are age and not length-based then *maybe* the male and female numbers-at-age are just $N_{y,a,s} = 0.5 \times N_{y,a}$ for $s = m, f$. Given we have the sex of each adult in the sample, s_j, then it's simple to reform the parental probability as follows:

$$\pi_{ij} = \frac{\phi_{s_j,\mathbb{E}(a_j - a_i | l_j, s_j)}}{\sum\limits_a N_{y_i - a_i, a, s_j} \phi_{s_j,a}}, \tag{2.3}$$

and the factor of 2 disappears as we are specifically talking about the probability of detecting either a mother or father and not both (although given the 50/50 split assumed in this case the factor of two reappears and the only sexually-related difference is in the $\phi_{s,a}$ and obviously the distribution of male and female samples and POPs in the data).

2.3 Calculating the likelihood

If we define $j \in \Xi_i$ as subset of adult samples that can actually be sensibly compared to the juvenile i then the total likelihood would be a product of Bernoulli distributions:

$$\prod_{i \in \Theta} \prod_{j \in \Xi_i} \pi_{ij}^{k_{ij}} \left(1 - \pi_{ij}\right)^{1 - k_{ij}}, \tag{2.4}$$

where k_{ij} is equal to 1 if the $\{i, j\}$ pair is a POP and zero if not. As seen in the previous subsection, the probability of j being the parent of i is π_{ij} is complicated but reverting to the simple version and with no sexual disaggregation $\pi_{ij} \equiv 2/N$ where N is the absolute number of adults. In that case, given we have a total of M fair comparisons between juveniles and adults, the total likelihood would just be binomial with with $n = M$ and probability $2/N$ so the expected number of POPs would be $2M/N$ and we recover the cartoon version of the estimator: $\widehat{N} = 2M/\#POPs$. The implicit assumption behind the Bernoulli (and binomial) distributions for the POPs is that the adult samples are independent. In future, if we wish to explore the impact of non-independence then a beta-Bernoulli/beta-binomial model (a univariate version of the Dirichlet-multinomial model used to model over-dispersion in the tagging data) would be a logical option.

In the most naive form (or if age/length was a continuous variable) we would have around 38,000,000 individual Bernoulli distributions to compute. That would take a very long time but thankfully we don't have to do all of that work - the comparisons group into a much smaller number of key "types":

- Juvenile birth year/cohort, c
- Adult capture year, y
- Adult capture age , a (so age in birth year is $\max(0, a - [y - c])$).

For each one of the juvenile-adult comparisons that fall into these types the parental probability, $\pi_{c,y,a}$, is the same. So, if there are $\mathcal{M}_{c,y,a}$ fair comparisons in this group and $\mathcal{P}_{c,y,a}$ detected POPs, then the total likelihood in Eq. 2.4 reduces to a product of binomials:

$$\ell\left(\mathcal{P}_{\bullet,\bullet,\bullet} | \cdots\right) \propto \prod_c \prod_y \prod_a \pi_{c,y,a}^{\mathcal{P}_{c,y,a}} \left(1 - \pi_{c,y,a}\right)^{\mathcal{M}_{c,y,a} - \mathcal{P}_{c,y,a}}. \tag{2.5}$$

We shall see more detail in the data summaries later on but basically there are 6 cohorts sampled (2002-2007), over 5 adult capture years (2006-2010) and over 26 ages (5-30 which is a plus group) so that's 5*6*26=780 types. Still big, but a lot less than 38,000,000 and no bigger really than some of the multinomial likelihoods for the catch composition data. Also, there are some instances with zero comparisons (and POPs obviously) so the *real* number is 663 types.

2.4 Data summaries

The age-based sexually aggregated close-kin OM data is an `R` data frame which has been converted into a style usable with `ADMB` (see new *sbtdata2011.dat*). Table 2.1 details the number of comparisons by juvenile birth year and adult capture year, with Table 2.2 detailing the same information but for confirmed POPs.

	Year				
	2006	2007	2008	2009	2010
Cohort					
2002	30	203	213	199	168
2003	275	1826	1912	1788	1504
2004	361	2323	2381	2290	1971
2005	300	1985	1943	1815	1529
2006	0	1855	1943	1815	1632
2007	0	0	2070	1938	1632

Table 2.1: Number of comparisons (times 10^{-3}) for each cohort and capture year.

	Year				
	2006	2007	2008	2009	2010
Cohort					
2002	0	0	0	0	0
2003	0	5	1	2	0
2004	0	2	0	0	3
2005	1	4	5	4	1
2006	0	4	3	2	0
2007	0	0	3	4	1

Table 2.2: Number of POPs detected for each cohort and capture year.

Extreme caution has been taken to make sure the translation from the original length and sex-based data structures into age based sexually aggregated forms maintained the same underlying unchangeable information (such as those in Tables 2.1 and 2.2). The age structure and lack of full sexual disaggregation is different but the core information content and idea behind the likelihood structure is the same.

3 Results

The results section is split into several subsections, as there are a number of changes to the OM (in terms of definitions and new data) and issues arising with the grid. By taking each one in turn we can understand how things are changing, and hopefully why.

3.1 Effect of alternative spawning population definition

The first change (from `sbtmod25` to `sbtmod26`) replaces the old definition for the effective spawning population (biomass of everything aged 10+) with the new definition: $S = \sum_a \phi_a N_{y,a}$, where ϕ_a (effective spawning potential was as defined in subsection 2.2 and can be seen in Figure 5.1). In Figure 5.2 we plot the level summary in the grid samples and it is clearly very similar to that seen in the previous OM conditioning used to test the candidate MPs [3] and with a similar, but slightly more optimistic level of depletion: median (and 80%CI) of 0.06 (0.04-0.08). Figure 5.3 plots the more familiar likelihood profile plots for steepness (but with both the total negative log-likeliehood *and* the objective function plotted side by side) and Figure 5.4 plots the key penalties that make up the rest of the objective function when added to the negative log-likelihood of all the data. This idea was first explored in [4] as, contrary to some of the language used in various documents, we do not use likelihood weighting (for h, M_0 and M_{10}) we actually use objective function weighting. This will become a familiar idea throughout this paper but this is important as we need to know how much information is coming from the actual data, and how much is coming from the subjective penalties and parameter correlation.

As first observed in [4] there is a high degree of influence coming from the penalties when it comes to which values of steepness are sampled within the grid. Figure 5.3 shows quite clearly the difference between the objective function and the total negative log-likelihood: the data suggest higher levels of steepness but the actual objective function has the clear peak around the 0.73 level, with a slight skew towards higher values. From Figure 5.4 we can see that quite a lot of information is actually coming from mostly the recruitment deviation penalty, but also the selectivity change penalty. The recruitment penalty acts very strongly on higher levels of steepness, until decreasing down to around the same level for the lowest three steepness values. Presumably, this is because for the higher steepness values, there is less of a reduction in mean recruitment as the SSB declines, so the recruitment deviates have to be lower (i.e. further away from zero) to account for this reduction in mean recruitment given the higher steepness. To summarise, the actual sampling distribution of steepness in this grid is a clear mix of data driven (via the likelihood) and subjective (via the penalties whose parameters are fixed by the user) effects.

3.2 Close-kin data on previous grid

The next phase was to include the close-kin data and likelihood within the OM and run the OM for the same grid specifications. For some reason, the runs which combined the lowest steepness (0.55) with the highest level of M_0 (0.45) would not converge, so a slightly reformed but similar grid (`baseCKmk3sqrt`) was run where the M_0 range was changed from 0.3, 0.35, 0.4, 0.45 to 0.3, 0.33, 0.37, 0.4. All possible grid permutations converged seemingly removing the low steepness/high M_0 issue. From Figure 5.5 what we observe is a significant switch in preference for generally lower levels of steepness and M_{10} and higher levels of M_0. Concentrating on steepness to begin with:

Figure 5.6 shows quite clearly that this shift in the sampled levels of steepness is not really being driven by new information in the data but by in fact an absence of information in the data. The close-kin data generally prefer lower levels of steepness but this influence is by no means strong; it is also the same trend for both the CPUE and the tagging (which is a reversal of what was the case for the previous grid without the close-kin data). When one actually looks at the total negative log-likelihood it is practically flat, showing little if any preference for a given steepness level. The objective function, however, is quite clearly in favour of lower steepness values and from Figure 5.7 this is clearly driven mostly by the preference for lower steepness coming from the recruitment penalty (as before) and the same consistent trend for higher steepness coming from the selectivity change penalty (as before).

While we have primarily concentrated on the steepness when performing these likelihood profile calculations we can focus on the other key objective function-weighted parameters within the grid: M_0 and M_{10}. Figures 5.8 and 5.9 show the likelihood and penalty profiles for M_0 with each steepness grid value separated out in colour. As one would perhaps expect, the tagging data has by far the strongest level of information on M_0 and it clearly shows preference for higher values. There is very little strong information elsewhere in the data or the penalties so the distribution of samples in the grid seems to genuinely be driven by data. Figures 5.10 and 5.11 show the likelihood and penalty profile plots for M_{10}. While the Indonesian age data has some preference for higher values of M_{10}, the close-kin data shows a strong lack of preference for the highest value (0.16) and a general preference for lower values, as does the tagging data. There is little apparent influence from the penalties in this case.

3.3 Likelihood vs. objective function weighting

While the language used in various documents frequently refers to likelihood weighting in relation to steepness, M_0 and M_{10} we actually use objective function weighting and so admit the influence of the various penalties used within the OM. With the addition of the close-kin data, the general data-driven preference for higher levels of steepness has disappeared (see Figure 5.6). This is because the close-kin data are essentially saying two things: that current adult abundance (and B_0) and survival probabilities are higher than estimated within the previous OM conditioning. This increased adult abundance no longer requires the higher levels of steepness preferred by the abundance data (CPUE and tagging and the survey) and so the preference in fact shifts towards generally lower values of steepness, although the strength of this preference is significantly diminished. This results in little actual information within the data as a whole on steepness and the penalties dominate, pushing it to lower values; actual likelihood weighting would result in something looking like a flat prior.

While the information for M_0 looks fairly convincing and is coming overwhelmingly from the tagging data, it is not so clear when looking at M_{10}. From the likelihood profiles there is a general tendency towards lower values, and a strong tendency for the close-kin data to be incompatible with the highest value of 0.16. There is a subtle but apparent inconsistency between the likelihood profiles and the grid sampling distribution for M_{10}: from Figure 5.10 there is little information to separate the lower levels but in the level plot in Figure 5.5 there is an apparent reduction in preference for lower levels. This is coming from negative correlation between steepness and M_{10}: the higher steepness values are no longer sampled and so the accompanying lower levels of M_{10} which tend to appear with higher steepness levels disappear also. This is a secondary penalty effect, but is still noticeable.

It would seem that some thought needs to be given to exactly how we might reweight grid elements, if not by their priors. The current approach uses the objective function which includes data-driven and subjective/secondary influence via the penalties and this seems to be a strongly driving influence on primarily the steepness but also on M_{10} to a degree. While there is some subjectivity to the selectivity penalties the recruitment penalty is completely driven by the value of $\sigma_R = 0.6$, which is neither estimated (indeed it couldn't be without complex random effect modelling for recruitment) nor tuned in any *ad hoc* way.

3.4 Alternative grid options

A final grid configuration was explored in light of the above results. The M_0 grid was restored back to 0.3, 0.35, 0.4, 0.45. Given the issue with lowest steepness and highest M_0 the $h = 0.55$ option was dropped from the grid leaving four steepness values: 0.64, 0.73, 0.82 and 0.9. The M_{10} grid was also altered to get rid of the highest value and explore lower levels: 0.05, 0.075, 0.1, 0.125. In this case objective function weighting was used *only* for M_0 and M_{10}, with the steepness values assigned a simple uniform prior. It is important to note that this is not a recommended grid combination. It is provided purely to demonstrate the impact of the close-kin data for an example reformulation of the grid and weightings and, in doing so, provide a basis for discussion and further consideration of how these issues may be approached.

3.5 Summary and fits

Figure 5.12 details the levels of each parameter sampled within the revised grid (`baseCKmk5sqrt`). Remember steepness is now sampled from its (uniform) prior; for M_0, the higher levels are sampled but we do now have much better coverage of that sampling distribution than in the previous grid (`baseCKmk3sqrt`). For M_0 there is also some influence apparently coming from the penalties - in particular the selectivity change penalty. As before, the tagging data is the dominant source of information, resulting in a strong preference for the two highest values but with little to choose between them (see Figure 5.13). However, the objective function shows a marginally lower preference for the highest level of M_0 relative to the second highest and this is being driven by the selectivity change penalty (see Figure 5.14). Curiously for M_{10}, while the preference for the higher levels is low as before, there is a gradually increasing preference for the lowest values (down to 0.05). From Figure 5.15 this is not apparently coming from the data (the negative log-likelihood), which seems to show weak but noticeable preference for the two inner values of 0.075 and 0.1), but from the selectivity change penalties (see Figure 5.16). So, while a much weaker effect, we are still getting penalty influence within the grid sampling procedure for both M_0 and M_{10}.

Fits to the CPUE and aerial survey abundance indices are very similar to the previous OM conditioning, as are the quasi-posterior predictive analyses defined in [5], indicating that these data are being both fitted to and explained as they were prior to the inclusion of the close-kin data. To try and minimise the number of graphics in this paper we do not show the detailed fits to the catch composition and tagging data - they are available but, as with the abundance indices, the fits are very very similar to last time with no obvious issues given the inclusion of the close-kin data and changes within the grid and sampling regimes, used here to explore the impacts of including this new data set.

A fairly complex problem arises when trying to find a way to display the fits to the close-kin data. At the base level we have 780 binomial events we are trying to explain but clearly detailing the fits to each of these will yield little processable information. A simple approach was taken whereby we attempt to summarise the number of observed and predicted POPs for each birth year covered in the data (2002-2007) - this would result in a much more compact yet reasonably informative view of how well we are fitting to these data.

Each of the 780 binomial “events” (grouped remember by birth year, adult capture year and age) corresponds to a number of fair comparisons made, $\mathcal{M}_{c,y,a}$, a number of POPs detected, $\mathcal{P}_{c,y,a}$, and the probability of finding a POP in such a case, $\pi_{c,y,a}$. By using the normal approximation to the binomial distribution, the expected number of POPs over all adult years and ages of capture, but for each sampled cohort, can be expressed as the following:

$$\mu_c = \sum_y \sum_a \mathcal{M}_{c,y,a}\pi_{c,y,a}, \tag{3.1}$$

and the variance would then be given by

$$\sigma_c^2 = \sum_y \sum_a \mathcal{M}_{c,y,a}\pi_{c,y,a}(1-\pi_{c,y,a}). \tag{3.2}$$

It is a little more complex than this, however, as we have such a mean and variance in the expected POPs-by-cohort for each grid element $k = 1,\ldots,G$. To obtain a grid-averaged estimate of this mean and variance in the predicted number of POPs we use inverse-variance weighting so that:

$$\mu_c = \sum_{k=1}^{G} \mu_{c,k}\omega_{c,k}, \tag{3.3}$$

where

$$\omega_{c,k} = \frac{\sigma_{c,k}^{-2}}{\sum\limits_{l=1}^{G}\sigma_{c,l}^{-2}}, \tag{3.4}$$

so that $\sum_k \omega_{c,k} = 1$ and

$$\sigma_c^2 = \frac{G}{\sum\limits_{k=1}^{G}\sigma_{c,k}^{-2}}. \tag{3.5}$$

Although this is an approximation, it does give us a relatively simple way to summarise the fits to the close-kin data. Figure 5.17 shows the observed and predicted number of POPs per cohort sampled for the `baseCKmk5sqrt` example grid. Overall the fit is fairly good: only once (for the 2004 cohort) does the observed number of POPs sit outside the 80% CI (and sits within the 95%CI); for the zero observed POPs in 2002 we get an average prediction of less than 1 POP; and there is no clear bias trend of consistent over or under estimation. The OM cannot explain the inferred fairly large changes in adult population size, particularly from 2004 to 2005, and this may be to some degree due to sexual and length-based variation in the POP and comparison sample structure that our comparatively crude sexually aggregated and age-based model is missing. One should keep in mind the relatively low number of actual observed POPs per cohort when considering how well we are (or should be) fitting to these data, and how well we fit the other abundance data (CPUE, aerial survey, tagging). In summary, the OM (and this particular example grid and resampling regime) can fit the close-kin data fairly well.

Figure 5.18 shows the usual spawning population and recruitment time-series for `baseCKmk5sqrt`. For this OM configuration, current median (and 90%CI) for depletion is 0.08 (0.06-0.11) and one can clearly see both the more precise estimates of the recent spawning population relative to the previous OM (see Figure 14 in [5]) and marginal recent upturn in the most recent estimates of spawning population abundance, opposed to the very gradual decline in the 2000s observed in the previous OM.

4 Discussion

4.1 What are the close-kin data saying?

On the basis of the independent estimate and these initial attempts to incorporate the close-kin data in the OM, these data clearly increase both the recent estimates (and precision) of spawning population size and survival probabilities and, at least for the grid and resampling specification detailed herein, suggest that current spawning population depletion is not as low as for the previous OM conditioning. Clearly the approach provided here, and potential alternatives, need further consideration and review to provide a robust foundation for including this important direct data source on spawner abundance in the SBT OM.

There are a number of as yet unresolved uncertainties, outside of any genetic-specific issues, that are likely to affect the close-kin data. From the parametric side, we have not accounted for uncertainty in certain aspects of what defines the effective spawning contribution ogive. From a more individually focussed perspective, there may be additional variation in the annual reproductive success among the sampled adults due to factors such as local variation in environmental conditions at the time of spawning or individually better/worse-than-average spawners. The former issue is something that can most likely be quantified and therefore might be best handled within an augmented grid configuration. The latter issue, relating to over-dispersion within the close-kin data beyond that assumed in the likelihood, is more difficult:

We do not have a large amount of POPs, unlike the hundreds of tag recaptures, so it is very unlikely that we can do a similar thing to what is done with the tagging data and estimate the level of over-dispersion in the close-kin data by simply looking at expected and observed residual variation. Within any reference OM containing the close-kin data we should not simply fix the over-dispersion at some agreed value as this amounts to subjective and rather arbitrary data weighting. Given the issue of over-dispersion within these data, like it was with the tag mixing hypotheses, will be a largely qualitative issue it seems more sensible to look to deal with this issue within the suite of robustness tests, until such time as additional information is available to move to a more quantitative treatment.

4.2 Information content and resampling the grid

The addition of the close-kin data have seemingly “removed” the information in the *data* on steepness (which has a knock-on effect to M_{10} to some degree). The abundance data (CPUE, aerial survey, tagging) tended to support higher levels of steepness in the previous OM conditioning. Given the close-kin driven increase in recent (and historical) adult abundance this preference has disappeared (and become a very weak preference for lower steepness levels in the abundance data). Across all the data sets, there was no apparent clear information on steepness at all, with all the apparent information coming from primarily the recruitment deviation penalty. Even with objective function weighting turned off for steepness there is still some contribution from the penalties (in particular the selectivity change penalty) toward the sampling distribution of both M_0 and M_{10}. Unlike for the steepness, there is clear apparent information in the data for these parameters, but still one can see how that information trades off against the penalties to form the objective function which ultimately dictates the grid sampling distribution.

The penalties are subjective - especially the recruitment deviation penalty - so one must then consider whether objective function weighting is the best option when actively estimating the grid sampling distribution and not simply using the priors. Likelihood weighting is one obvious alternative option but for steepness (and perhaps for M_{10} also) we may need to reassess the options around the use of prior weighting given the lack of information apparent within the data.

4.3 Future work required for integration of close-kin data to the OM

This paper presents a preliminary exploration of how we might include the close-kin data within the SBT operating model and what the information content within these data might be. There are a number of clear avenues of work that are required to further explore the integration of the close-kin data into the OM:

- Treatment of both age and length: all the key close-kin data are transformed to age-based data prior to inclusion. Care was taken to account for both individual and sexual variation in length-at-age, but given the likelihood that both age and length are key aspects of the reproductive process [2] a more integrated approach, such as that taken in [1], is recommended.
- Sexual dimorphism: while some aspects were accounted for, and the close-kin data are generally fairly well fitted, we must explore options for the inclusion and the importance of the sexually specific nature of the close-kin data. Some simple initial options have been highlighted in this paper but see [1] for a more formal treatment.
- Effective spawning population: the current historical assumption is clearly incompatible with what we observed in the close-kin data relating to the age of adult POPs. In the stand-alone close-kin work [1] Indonesian selectivity-at-length and residence time were considered the same thing, and were used with other key fecundity data [2] to define the effective (male and female) reproductive capacity-at-length. Given the OM structure and what we know about M, there is some uncertainty around whether Indonesian selectivity (at age or length) could be so interpreted in relation to residence time on the spawning grounds. Using maturity-at-length information, as well as the batch fecundity and resting/spawning information from [1, 2], an initial update to the “biomass of 10+ animals” was calculated. This option is just one of a potential range of options (including something linked to residence time as per [1]) and will have an undoubted influence on the results, given it is so strongly associated with what we define the spawning population to be.

5 Acknowledgements

This work was funded by AFMA and CSIRO’s Wealth from Oceans National Research Flagship.

References

[1] Bravington, M. V., Grewe, P. G., and Davies, C. R. (2012) Report of the close-kin project. *CCSBT-ESC/1208/**.

[2] Davis, T., Farley J., Bravington, M. V., and Andamari, R. (2003) Size at first maturity and recruitment into egg production of southern bluefin tuna. *FRDC 1999/106*.

[3] Anonymous. (2011) Report of the Sixteenth meeting of the Extended Scientific Committee of the Commission for the Conservation of Southern Bluefin Tuna.Bali, August 2011.

[4] Sakai, O., Kurota, H., Takahashi, N., and Butterworth, D. S. (2011) Conditioning of the SBT

operating model to inform projection specifications. *CCSBT-ESC/1107/33*.

[5] Hillary, R. M., Preece, A., Davies, C. R., Barnes, B. and Begg, G. (2011) Reconditioning of the southern bluefin tuna operating model: exploratory data analysis, fitting performance, and current stock status. *CCSBT-ESC/1107/11*.

Figures

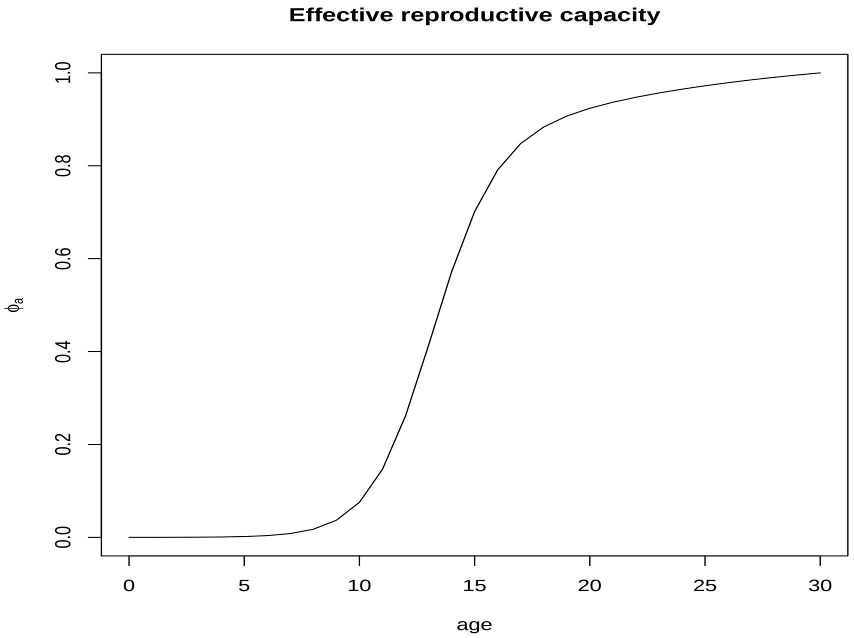

Figure 5.1: Effective relative reproductive capacity-at-age: ϕ_a.

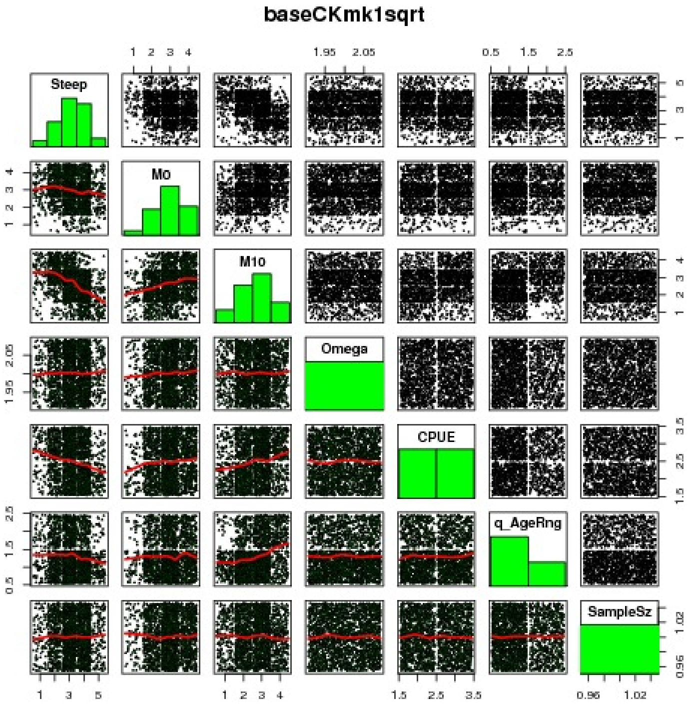

Figure 5.2: Level plot for old grid specifications with new spawning population definition (`baseCKmk1sqrt`)

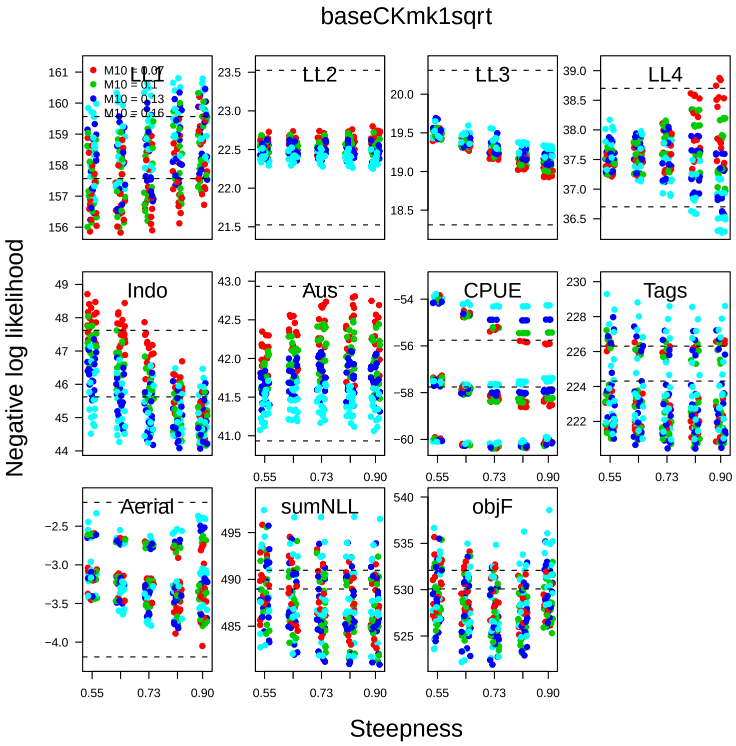

Figure 5.3: Likelihood profile for steepness (split for various values of M_{10}) for grid `baseCKmk1sqrt`.

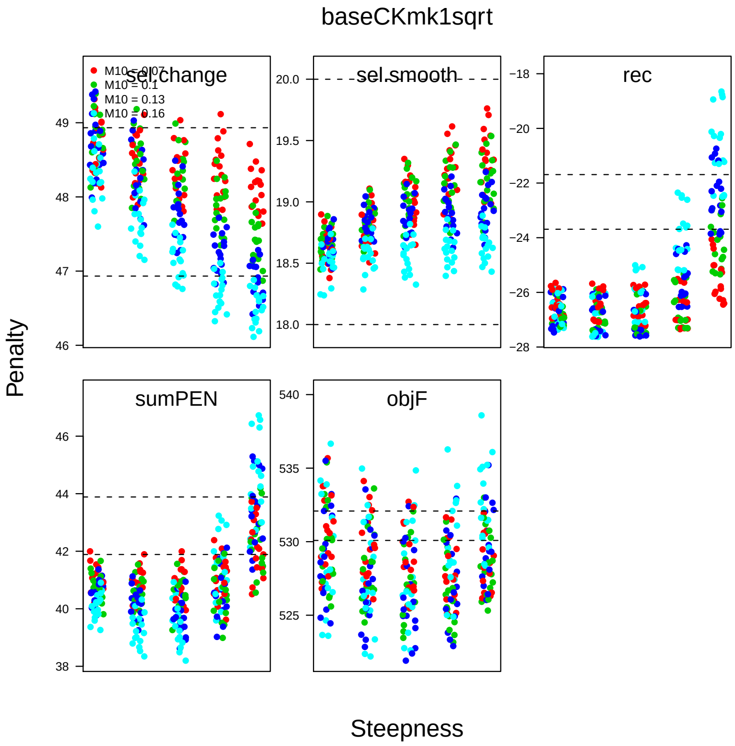

Figure 5.4: Penalty profile for steepness (split for various values of M_{10}) for grid `baseCKmk1sqrt`.

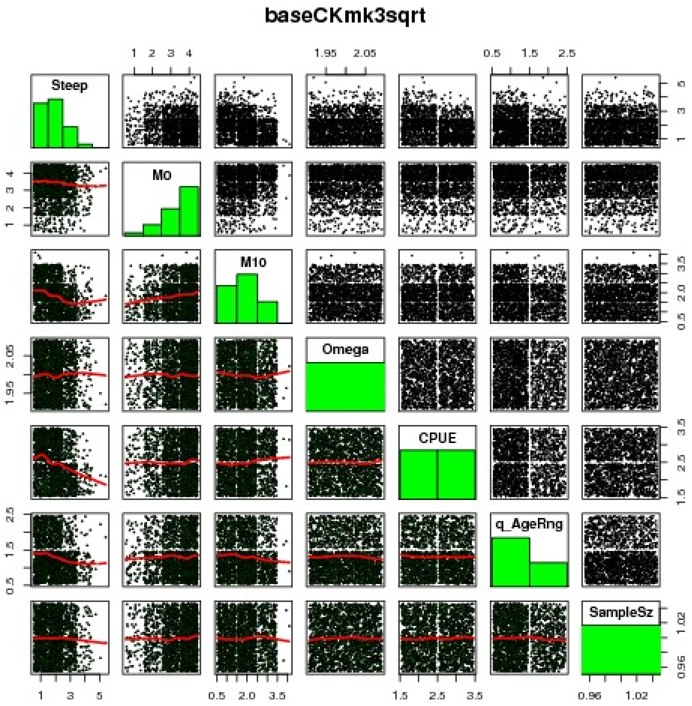

Figure 5.5: Level plot for modified M_0 specifications with new spawning population definition (`baseCKmk3sqrt`).

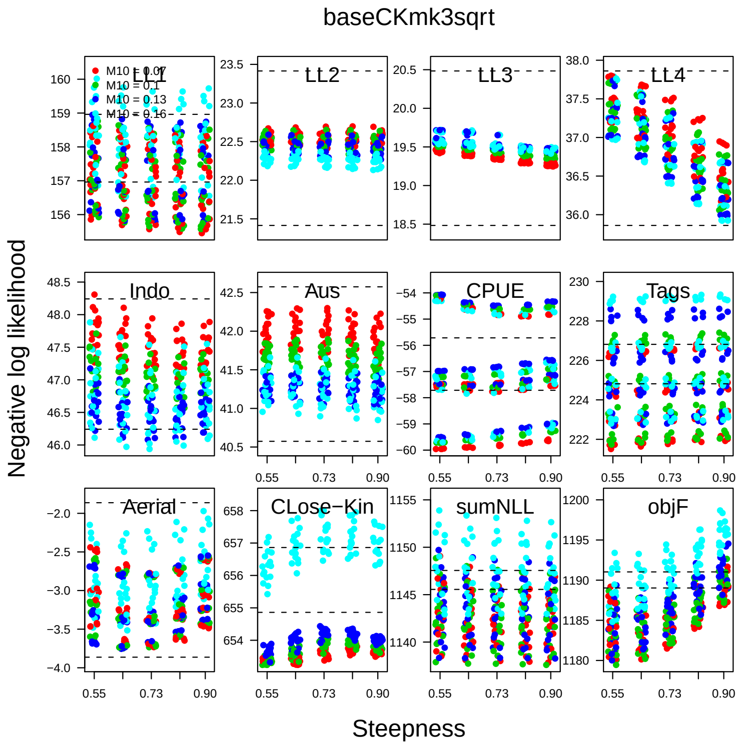

Figure 5.6: Likelihood profile for steepness (split for various values of M_{10}) for grid `baseCKmk3sqrt`.

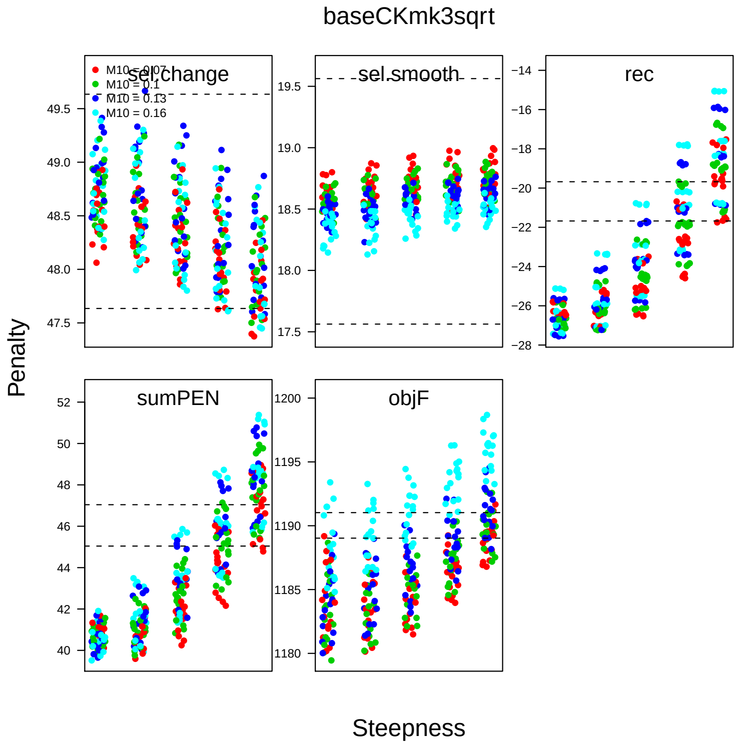

Figure 5.7: Penalty profile for steepness (split for various values of M_{10}) for grid `baseCKmk3sqrt`.

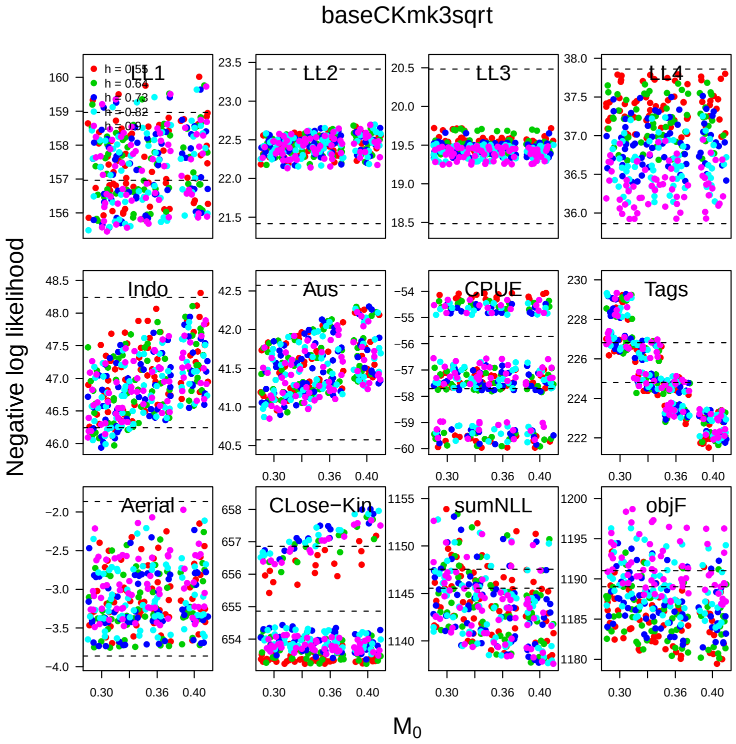

Figure 5.8: Likelihood profile for M_0 (split for various values of steepness) for grid `baseCKmk3sqrt`.

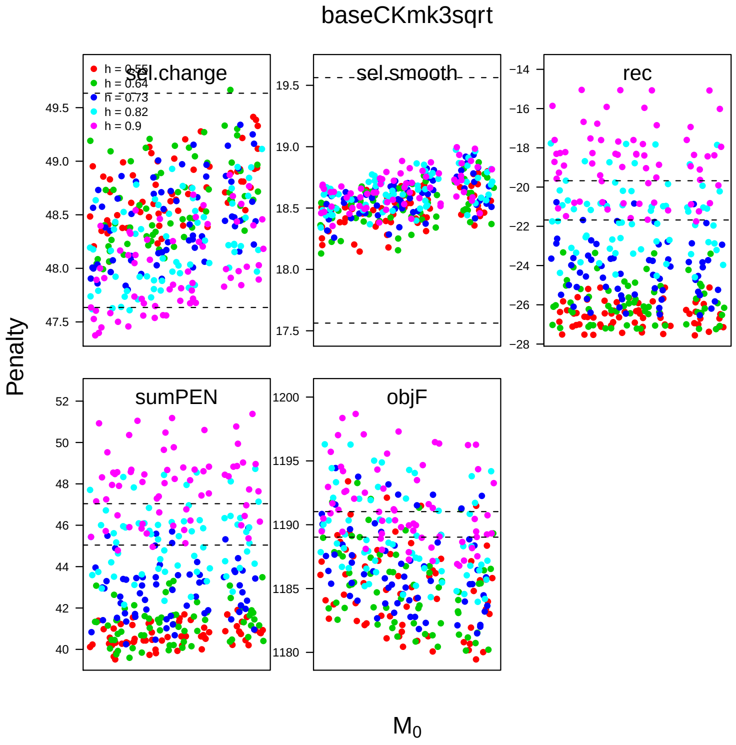

Figure 5.9: Penalty profile for M_0 (split for various values of steepness) for grid `baseCKmk3sqrt`.

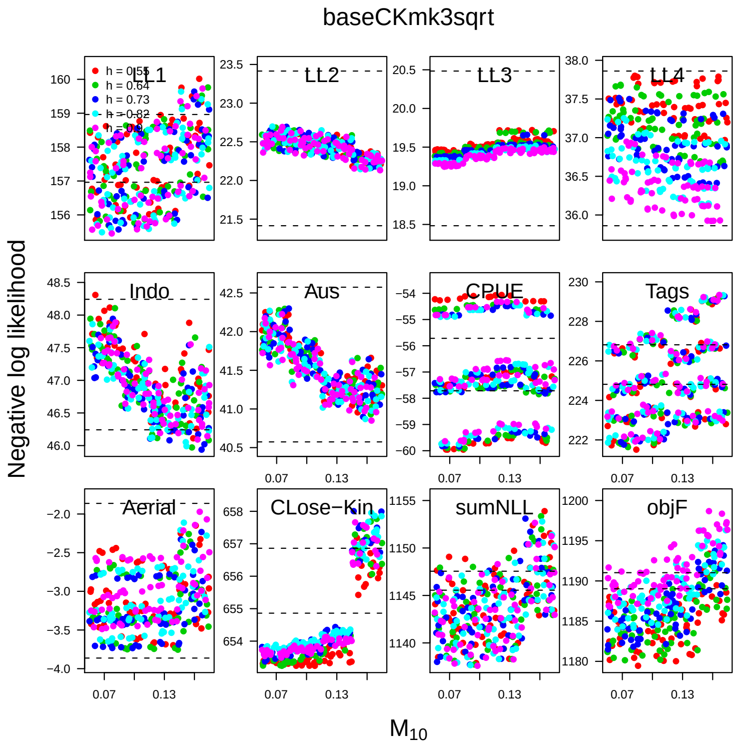

Figure 5.10: Likelihood profile for M_{10} (split for various values of steepness) for grid `baseCKmk3sqrt`.

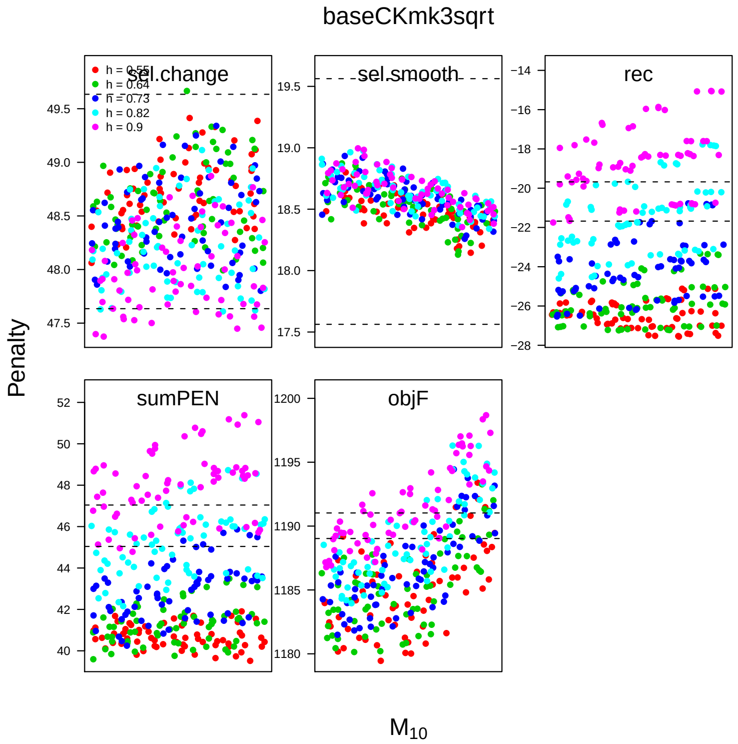

Figure 5.11: Penalty profile for M_{10} (split for various values of steepness) for grid `baseCKmk3sqrt`.

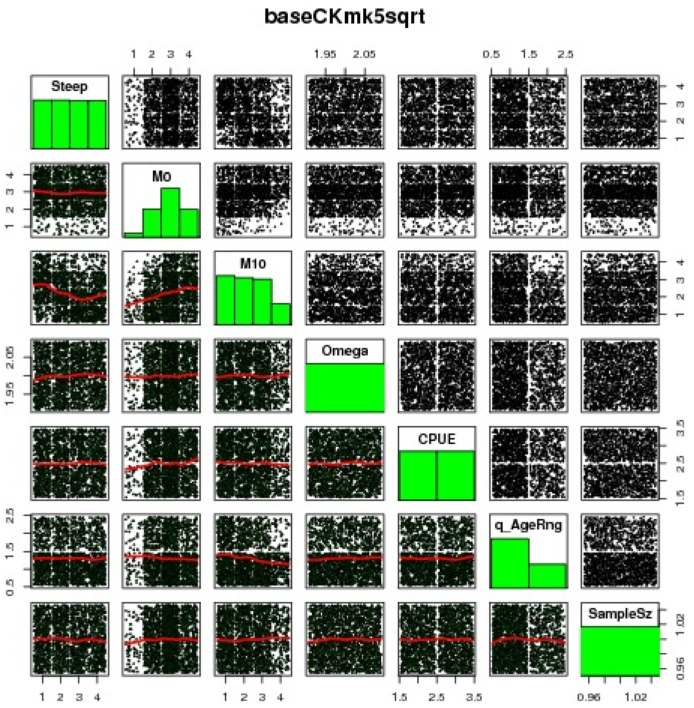

Figure 5.12: Level plot for modified steepness and M_{10} specifications with new spawning population definition (`baseCKmk5sqrt`).

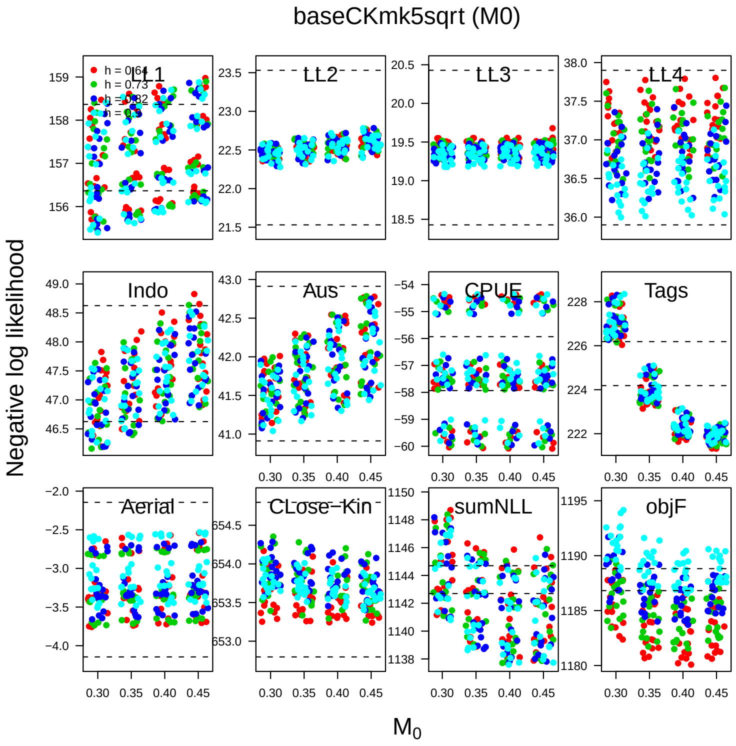

Figure 5.13: Likelihood profile for M_0 (split for various values of steepness) for grid `baseCKmk5sqrt`.

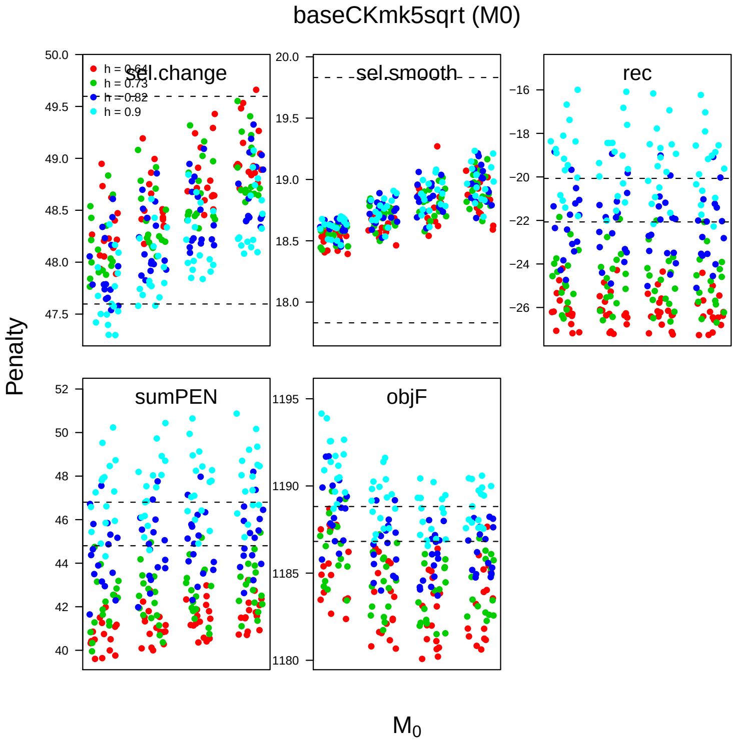

Figure 5.14: Penalty profile for M_0 (split for various values of steepness) for grid `baseCKmk5sqrt`.

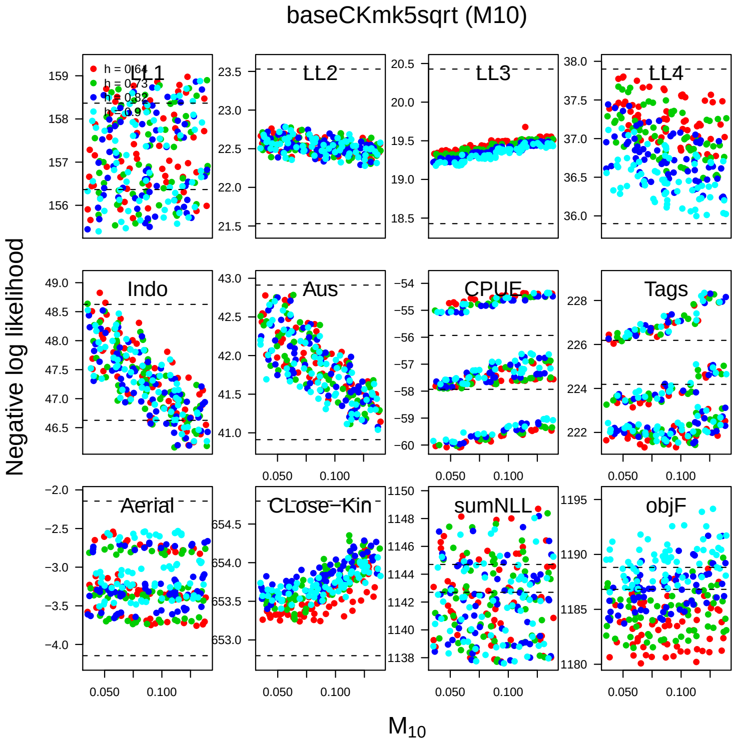

Figure 5.15: Likelihood profile for M_{10} (split for various values of steepness) for grid `baseCKmk5sqrt`.

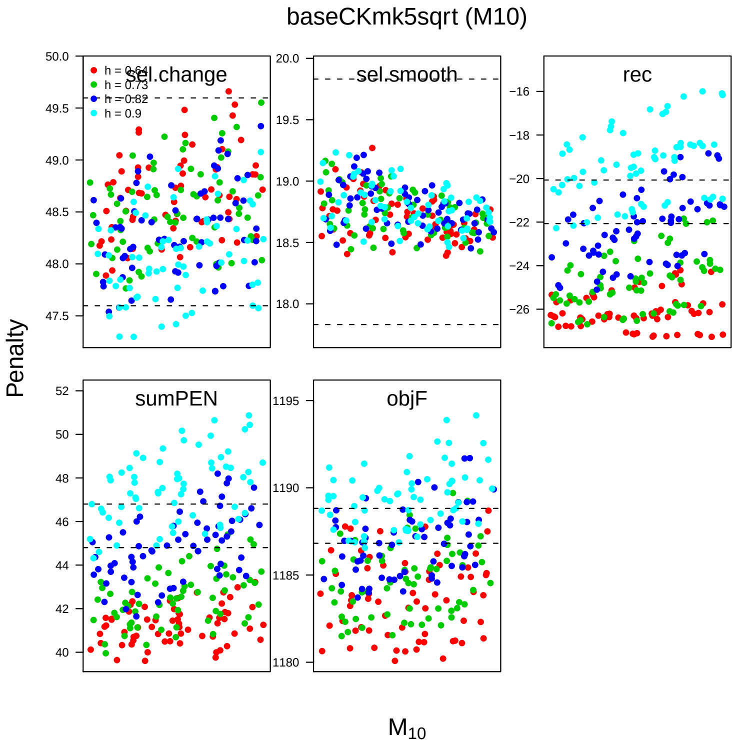

Figure 5.16: Penalty profile for M_{10} (split for various values of steepness) for grid `baseCKmk5sqrt`.

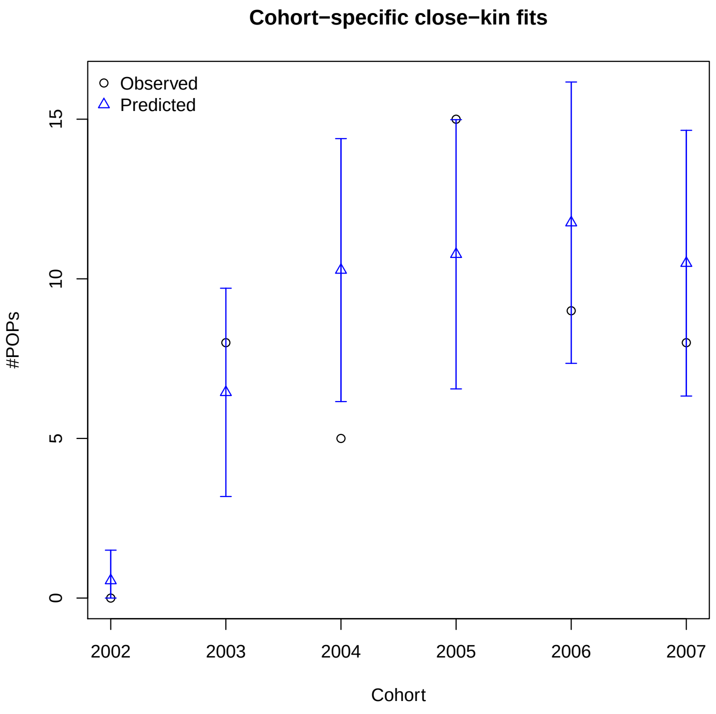

Figure 5.17: Observed (points) and median (blue triangles) and 80%CI (blue whiskers) number of POPs for each sample cohort, using the normal approximation to the binomial distribution, for grid `baseCKmk5sqrt`.

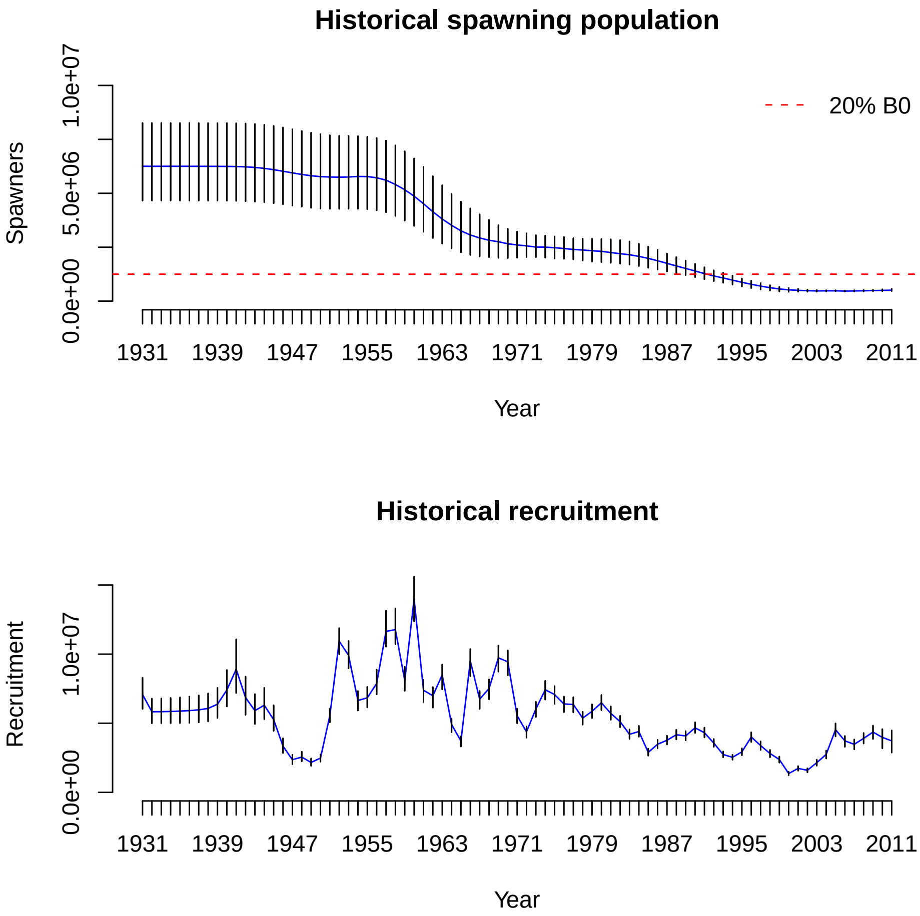
Historical spawning population

Historical recruitment

Figure 5.18: Median (blue line) and 80%CI (black whiskers) estimates of spawning population abundance (top) and recruitment (bottom) for grid `baseCKmk5sqrt`. The median estimate of 20% B_0 is denoted by the dashed red line.

CONTACT US

t 1300 363 400
+61 3 9545 2176
e enquiries@csiro.au
w www.csiro.au

YOUR CSIRO

Australia is founding its future on science and innovation. Its national science agency, CSIRO, is a powerhouse of ideas, technologies and skills for building prosperity, growth, health and sustainability. It serves governments, industries, business and communities across the nation.

FOR FURTHER INFORMATION

CSIRO Marine and Atmospheric Research
Rich Hillary
t +61 3 6232 5452
e Rich.Hillary@csiro.au
w