Research Trawl and Shrimp Bycatch Results Relevant to Red Grouper

Scott Nichols
NMFS Pascagoula

SUMMARY

All of the NMFS and SEAMAP standardized surveys range westward from Mobile Bay, well outside the zone of highest red grouper abundance. The standardized time series contained only one record of red grouper, caught in the Fall Groundfish Survey in 1972. Any expansion of the population into the northern Gulf has thus not been picked up in the surveys.

Red grouper have appeared in the sporadic trawl survey work done by NMFS off Florida. Catch rates were low, on the order of 0.03 fish per hour overall. Red grouper are also reported in Florida stations in the shrimp bycatch observer data base, although they do not appear in all years with Florida coverage. Catch rates in the Florida commercial shrimp trawls on the order of 0.007 fish per hour. Nevertheless, because an appreciable number of observations have been accumulated in Florida when years are combined (407 research vessel, 666 for non-BRD shrimping, and 843 for shrimping with BRDs), an attempt was made to estimate shrimp fleet bycatch for red grouper.

A new modification is introduced here to the Bayesian analysis procedure for shrimp bycatch estimation (previously used in SEDAR7-DW-3, 54; SEDAR9-DW-26; & SEDAR9-AW-3), that should greatly reduce the sensitivity to the year effects' priors seen in the SEDAR9 papers. Lack of an extended time series off Florida prevents meaningful tracking year to year variation, but the average annual bycatch take appears well estimated by the Bayesian procedure, with a plausible confidence interval: point estimate of 8,400 fish, and 95% CI of 3,000-24,000.

INTRODUCTION

The standardized trawl surveys (in the notation of SEDAR7-DW-1: summer surveys SS, ES & TC, 1981-2004; fall surveys FS, FF, & FG, 1972-2004) have produced only one red grouper, in the Fall Groundfish survey of 1972. No further analysis is possible, but the lack of occurrence in recent years may be a bit of a surprise, given that some believe red grouper may be expanding into the northern gulf.

Our data base contains records from scattered research vessel forays into Florida (mostly NW Florida) in 1973, 1980, 1983, and 1986. There was a more thorough Florida shelf survey done in 1978, and reported on by Darcy & Gutherz 1984. Because of Hurricane Katrina, I did not have a copy of their paper readily available, but the catch rates for red grouper in the data records from that cruise showed an average catch rate of 0.053 fish per hour outside 10 fm, with no catch inside 10 fm (from 228 and 73 stations in the depth bands, respectively). These rates seem quite low for a species that supports as much directed catch as it does. Most likely, habitat choice and / or behavior characteristics keep red grouper only minimally vulnerable to capture in trawl gear.

Despite low catch rates, there are enough shrimp bycatch observer records from Florida that development of a catch estimate seemed worth investigating. In SEDAR9-AW-3, I had come up with a very *ad hoc* method for correcting some estimation problems noted in SEDAR9-DW-26. The AW-3 method involved repeated runs with varying priors on year effects, and noting maxima or minima in several diagnostic statistics. Well after the SEDAR9 AW, I realized there would a simpler and more defensible way of accomplishing the same thing. This paper gives me an opportunity to try the revised method on a new, low catch rate species, introduce the method change into the SEDAR process, and document the low bycatch rate for red grouper with an appropriate confidence interval.

METHODS

With the exception of the new modification, the bycatch Bayesian estimation procedure used here closely follows the documentation for Model 2 in SEDAR7-DW-3 and 54. In summary, CPUE is modeled as the sum of a series of (log) main effects (year, season, alongshore area, depth, and data set), with an additional 'local' effect to model perturbations from the main effects predictions. This 'local' term is uniquely Bayesian, in a sense serving as a fixed effect for cells in which data are plentiful, as a random effect for cells without data, and as something between fixed and random in cells having some but not a lot of data. Variation within each cell is modeled by a negative binomial. Estimation is done by Markov Chain Monte Carlo (MCMC) using the freely available BUGS software.

The modification introduced here is quite simple. Instead of independent priors for each year effect, an overall (log) mean term is added and given its own normal prior. Year effects are redefined to have a mean of zero (log scale) over all years. Table 1 shows a short segment of the BUGS code to illustrate the change (full code is still very similar to that tabled in SEDAR7-DW-54), and to show the numerical values used. There was an unexpectedly big cost in running time from this modification. A single standard run now takes about 70 hours, *vs* about 20 hours the old way. This increase precluded much experimentation for sensitivity, etc. I have not yet been able to complete runs for the SEDAR9 species. However, trials with red snapper and total bycatch showed that the dependence of average bycatch level on the year effects' priors was just about eliminated. The analyses returned values very similar to those obtained on completion of the ad hoc procedure of SEDAR9-AW-3. Therefore, I now recommend this modification as the standard procedure for shrimp fleet bycatch estimation.

Red grouper is not included among the 'short list' of species worked up under the Evaluation Protocol (list available from NMFS Galveston Lab), and thus analysis is limited to the stations worked up under the Characterization protocol (all species recorded individually). The appropriate data come from the Galveston base, largely from the subsets coded C and X, and from the 'historical' (1972-82) base in Pascagoula. The data used here are from the same April 25, 2005 update from the Galveston base used in SEDAR9. For species prevalent in the western Gulf, research vessel hauls by the Oregon II provide much of the basis for following interannual variation. For red grouper, the Oregon II data set west of Florida provides merely a large number of observations of zero. Nevertheless, to keep the structure of the previous analyses, I've included all western Gulf observations in this analysis.

As the SEDAR9 experience showed, it is wise to compare the Bayesian analysis output with approximate values based on very simple analyses. I considered two approximations, similar to the ones used in SEDAR9-DW-26. The average red grouper CPUE for all commercial (non-BRD) observations was 0.00161 (from a SAS Proc Tabulate run). The median of the annual nets per vessel was 2.52, and the median effort gulfwide was 4.6 million hours. The product of these factors implies an annual bycatch of approximately 19k fish. Isolating Florida data alone, the average CPUE over all years was 0.00692, effort averaged 0.594 million hours, and I've kept the same nets per vessel figure of 2.52. The product of these Florida-specific factors implies ~10k fish per year. If the Bayesian analysis gives a median well away from these values, and particularly, if the confidence interval on the Bayesian estimate does not include these approximate values, we may still accept the Bayesian estimate, but would want to investigate possible reasons for the difference.

RESULTS

The distribution of the median annual bycatch (median over years) from each iteration of the BUGS run was chosen as the best statistic to describe an 'average' bycatch across years. The distribution is summarized by BUGS-generated quantiles in Table 2, and by a graph of the distribution in Figure 1. Similar summaries for that statistic on a log scale appear as Table 3 and Figure 2. The medians of these statistics are good choices for summarizing central tendency, and the 2.5% and 97.5% quantiles provide 95% confidence intervals. Stock assessment models usually need a parameterized summary of these results. Even on a log scale, the resulting distributions are never quite normal, but the lognormal is usually a reasonable approximation. I customarily use the median on the log scale as the mean of a paramterized distribution, and calculate the standard deviation as the difference between the 97.5 and 2.5 quantiles divided by 3.92, resulting in a (log)normal distribution with mean -4.781 and variance 0.2922, which is the form usually used to pass the information to the stock assessment models. Calculating the variance of a normal distribution matching the MCMC confidence interval effectively writes off any asymmetry or elongated tails in the log distribution as unrealistic or not meaningful. In this particular case, little is written off: the parameters calculated are very similar to the mean and standard deviation columns to the left in Table 3.

Note that the confidence intervals on the average bycatch easily contain both approximate values produced by multiplying average catch rates by average efforts.

By its structure, the model can return annual estimates. These are shown in Figure 3, but are not recommended for use. Note that many years' estimates are very similar, reflecting that the only data were the large numbers of observations of zero from the western Gulf. Even for the later years, the fluctuations seen are probably not very meaningful, representing only chance inclusions or exclusions of small numbers of fish. If the data were truly random, we might accept some of those fluctuations as real to the extent supported by the confidence intervals; but as the sampling usually was far from random, we should not expect much from the year to year variations for such a rarely-taken species.

Other parameters returned by the model may be of some interest. The 'r' parameter used in generating the negative binomial is essentially pegged at its lowest allowed value (0.03). This was the same pattern seen with king mackerel in SEDAR9-DW-3, so the discussion in that paper is not repeated here. The posterior marginal for the mean parameter is shown in Figure 4. Relative to the prior, the data have relocated the posterior quite effectively. The BUGS history plot for the mean parameter (Fig. 5) shows a rather slow mixing pattern (chains remaining apart for extended periods). This is probably the slowest mixing pattern I've seen among the species examined to date, but there do not appear to be any problems from it. Basically, one would probably not want to use any fewer iterations than I used here. Adding more iterations would probably smooth out Figure 4, but the distribution has been pretty well established by the 40k points used. The area main effects (Fig. 6) are interesting, given that almost all observations in zones 2, 3, and 4 were zero catches. Structurally, the model assumes that there are low abundances throughout the Gulf, translated into very rare occurrences via the negative binomial. Numerically, the model estimates the underlying abundances in the Florida area must be about 135x higher than the abundances in the rest of the Gulf. This actual value is probably not of great interest -- the point really is that the model seems able to handle very low abundance situations without numerical foul-ups.

Sizes were reported for 3 fish in the observer data base; all were about 10 inches. Not exactly strong evidence, but based on other species, and lacking evidence to the contrary, it seems reasonable to presume that trawl-catchable red grouper are largely age zero.

DISCUSSION

Compared to the directed catches, even the upper confidence limit for red grouper bycatch is a small value. I suspect the SEDAR12 DW may thus choose to omit shrimp bycatch from the assessment, but the parameters to include it are available, if desired. However, the real points of this paper are documentation of the method modification; and demonstration, with confidence interval, of the low take of red grouper in the shrimp fishery. I've presented extended discussions of data limitations, modeling choices, performance, and sensitivity in my series of Bayesian bycatch papers (all in the SEDAR literature, and cited in the Summary here) written prior to this paper. Beyond the minor method modification, there is nothing new here, so I won't repeat those general discussions.

Regarding the method modification, the structure now used, mean + main effects, is actually more in line with the structure most commonly used in analysis of variance and general linear models. I did not use this structure originally, because I wanted to keep the years as independent as possible, largely so as not to flatten out estimates of year class strength downstream in stock assessments. However, the SEDAR9 results made it clear that tracking interannual variation in bycatch would not be meaningful for most species with low catch rates in trawls, so the main reason for having independent year-effects priors evaporated. The mean + (centered) year effects structure also allows separation of interannual variation (usually large) from uncertainty about an overall level of catch rate (usually much smaller, if one is willing to incorporate an advance look at average catch rate as part of the prior on mean). Keeping the prior uncertainty on overall level relatively small reduces the risk of the MCMC wandering into numerically unstable areas, and seems worth the 'peeking ahead' compromise to the prior. Those wanting more rigor could easily select a small, random subset of the data, and use that to set the mean of the prior. I prefer to simply take the mean CPUE from a SAS Proc Tabulate run for the entire commercial data set and use that as the mean for the lognormal prior. The nature of the error structure usually moves the median of the posterior on the mean parameter an e-fold or two below the prior value, with a much narrowed spread compared to the prior.

The reason why the modification works is fairly obvious – all data now directly affect the posterior on the mean parameter, and so the results will do a better job of returning an overall average. Two things were given up. 1) Results for years with modest amounts of data will drawn closer to the overall average than under the previous structure. This situation now appears to be more desirable overall. Results for years with less data turned out to be unreliable for low abundance species; higher abundance species are scarcely affected by the method change. 2) The real cost turned out to be the ~3x increase in running time. It is not obvious to me why there was such a large change, and it does limit the amount of work that can be done exploring alternative structures, sensitivity cases, etc.

Why include the western Gulf in this analysis at all? It's a fair question. The answers are that 1) there is a finite, if miniscule, contribution from the west for red grouper, so the analysis should allow the possibility, and 2) among the species we may eventually consider, red grouper are rather extreme in their concentration in one zone – I wanted to learn if the model structure would work in that situation. All indications are the model deals with the distribution properly, and I'm now more confident that it will work for most species with restricted distributions. There may still be problems with some species. The method modification will not help the vermilion snapper situation, which was limited by the applicability of the negative binomial to the extreme patchiness that vermilion exhibit. However, even that problem might be greatly reduced if we can develop more balanced and more random sampling than has proved possible in the past.

CITATIONS

Darcy, G H & Gutherz, E J. 1984. Abundance and density of demersal fishes on the West Florida shelf, January 1978. Bull.. Mar. Sci. 34(1): 81-105.

All other citations are to the SEDAR literature, maintained on the Southeast Fisheries Science Center web page.

TABLES AND FIGURES

Table 1. Short segments of BUGS code comparing the structure used in previous analyses vs the modification introduce in this paper. For an example of the full code, consults SEDAR7-DW-54.

Old version:

```
for (i in 1:33)  {
  yx[i]~dnorm(-6,0.7)
  }
...
l          logy[i,j,k,l,m]<-yx[i]+sx[j]+ax[k]+zx[l]+dx[m]+local[i,j,k,l,m]
```

New version:

```
meen~dnorm(-6,1)
for (i in 1:33)  {
  yraw[i]~dnorm(0,1)
  yx[i]<-yraw[i]-mean(yraw[])
  yef[i]<-meen+yx[i]
  }
...
      logy[i,j,k,l,m]<-meen+yx[i]+sx[j]+ax[k]+zx[l]+dx[m]+local[i,j,k,l,m]
```

Table 2. BUGS results (marginal posterior) for the statistic describing average annual bycatch (arithmetic scale.) In this and all other BUGS runs, these results are based on 2-chain runs of 24k iterations. The first 4000 points in each chain are used for adaptation and burn-in, leaving 40k points to characterize the posterior distribution. (The BUGS 'start' and 'sample' columns with these numbers were omitted, to make the tables fit the page.) Units are millions of fish.

```
Node statistics
node          mean       sd         MC error  2.5%       25.0%      median     75.0%     97.5%
mofam        0.00956    0.005443   2.4E-4    0.002846   0.005884   0.008389   0.01174   0.02368
```

Figure 1. Graph of the posterior marginal for average annual bycatch from BUGS output.

Kernel density

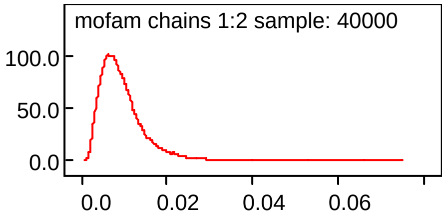

Table 3. BUGS mariginal posterior for the statistic describing average annual bycatch on a log (ln of millions of fish) scale.

```
Node statistics
 node         mean       sd         MC error  2.5%       25.0%      median     75.0%     97.5%
mofamlog     -4.79      0.5298     0.02436   -5.862     -5.135     -4.781     -4.445    -3.743
```

Figure 2. Graph of posterior marginal for log of average annual bycatch (ln of millions of fish).

Kernel density

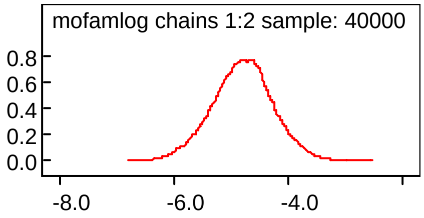

Figure 3. BUGS box plots of individual annual bycatch distributions on a log scale.

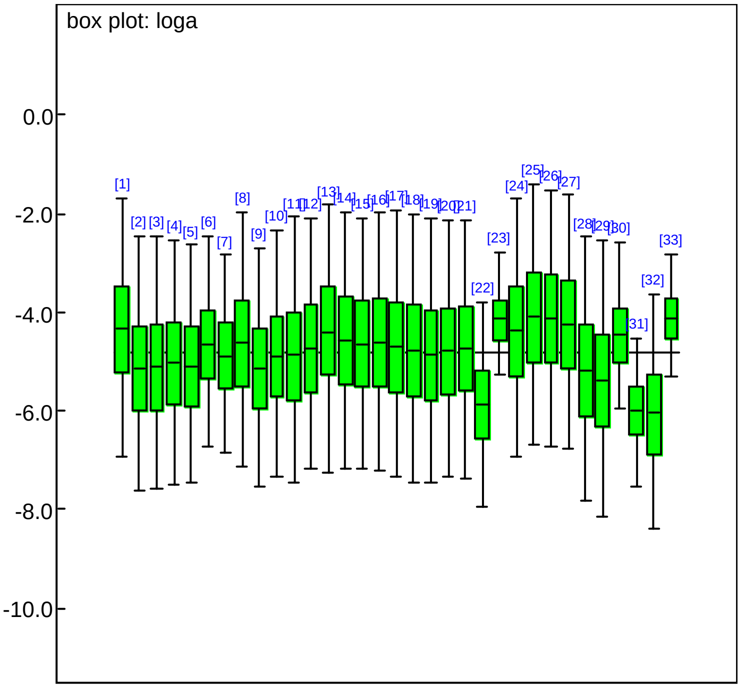

Figure 4. Posterior for the mean parameter. This and all main effects are on a log scale.

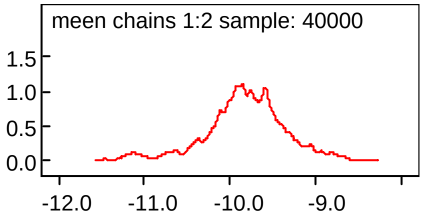

Figure 5. BUGS history plot for the mean parameter.

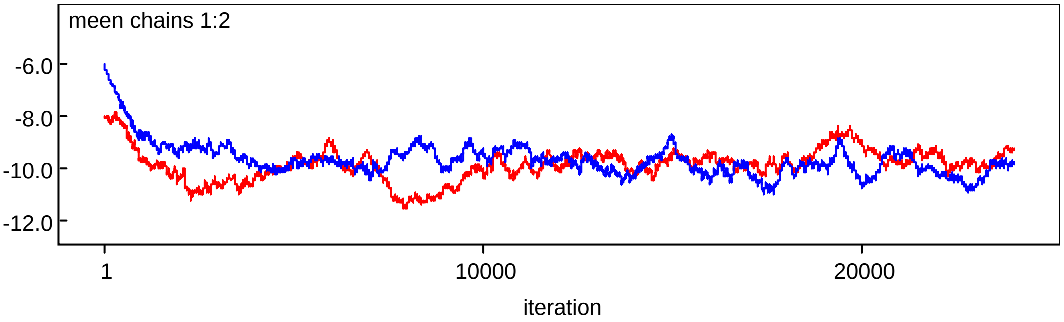

Figure 6. BUGS box plot for the area main effects. (1 is Florida, 2 is 'east of the river', 3 is LA west of the river, and 4 is Texas, using the same stratum boundaries as in all previous Bayesian bycatch papers.)