UPDATE FOR THE BAYESIAN ESTIMATION OF SHRIMP FLEET BYCATCH

Scott Nichols
NMFS Pascagoula

Introduction

I developed methods for applying Markov Chain Monte Carlo (MCMC) and Bayesian techniques to estimation of red snapper bycatch by the offshore shrimp fishery in the Gulf of Mexico in a companion paper (Nichols 2004). That paper included data only up through 1998, and did not consider the impacts of uncertainty in statistics of shrimping effort in the bycatch estimation. Here, I update the bycatch estimates with new observer and research vessel data from 1998-2003, incorporate the new estimates of uncertainty for shrimping effort, and include variable "nets per vessel" estimates based on the NMFS Vessel Operating Units Files (VOUF). The new information is used in a modification of the recommended model (Model 02) from Nichols (2004), to produce estimates of red snapper bycatch for the offshore shrimp fleet for 1972-2003.

Methods

CPUE Data

The new observer data files added here came from the same 02/13/04 request to the Galveston lab described by Foster (2004), and contain data collected from 1998 through December, 2003, with the bulk of the data from 2002. (There were also 5 spurious records dated 1980 or 2004). All data used in the last update of red snapper bycatch estimation (which used data through the major BRD evaluation project of 1998) were also used here. These older data were also the data used in Nichols (2004), and are not described further here. (Scott-Denton (2004) has provided a thorough overview of all observer activity since the start of the Regional Research Program in the early 1990s.) The data newly added for this analysis come from multiple projects, titled BRD (B), Characterization (C), Modified Characterization (M), Rock Shrimp (X), Red Snapper (R), and Effort-Shrimp (E), in the Galveston database. There were additional projects in the Galveston database that were not used here (codes Y, T, N, Z, G, and S). Some of these projects were South Atlantic only. Others involved investigations of experimental BRDS or TEDs, and probably varied in how much emphasis was actually on commercial fishing vs gear experimentation. These experimental projects were presumed not to be generally representative of commercial fishing with widely used gear, and thus were not included here. Sampling protocols for all but the E set were similar to the protocols in general use since the early 1990s (see Scott-Denton 2004 for more details). Most of the data from these projects consisted of BRD / noBRD paired towing, with data in separate 'experimental (BRD)' and 'control (noBRD)' files. A fraction of the tows in the BRD file had indications that the BRD was disabled, and a fraction of the tows in the noBRD file were coded for active BRDS. These 'crossed' data were not used; some are believed to have been tuning trials. Other than this exclusion, all observations coded to indicate a successful tow, coded to indicate that data had been 'archived' (i.e. checked and verified), and coded with suitable time and location information to assign the observation to a stratum in the analysis were included here. Project E used a different approach, which required some adjustments for processing the data. (My thanks to B Gallaway for explaining the nature of the data for the E project.) The goal was to work up the catch in all nets, but field necessity usually required the observers to work with catches of multiple nets that had been dumped together. Indications of which nets were combined appeared only in the Comments fields, and the comments were sometimes ambiguous if some but not all nets were successful. I dealt with this by totaling catches over all records for a 'station' in the data. To put the data on the same basis as the other observer projects, I then divided the catch rate by the number of nets towed (not 'nets worked up,' which is not directly coded). A fraction of the stations included records with 'missing' codes for red snapper catches; I did not use data from these stations. A small fraction of the E data were coded as "no BRD;" these were also not included here. Some BRD data for E were recorded in the "experimental" file, and some in the "control" file. For the E data, I ignored the source file, and accepted all remaining observations coded for having BRDs, coded as 'archived,' and coded with time and location data suitable for assigning to a stratum in the analysis.

There seems to be a greater diversity of extent of catch work-ups in this most recent data compared to pre-1998 data. This presents no problem to this red snapper analysis, as red snapper were worked up preferentially in these projects. However, it looks like some work remains before using these data sets to update bycatch estimates for other species, to be certain that any particular species would have been recorded had it occurred at any particular station.

Research vessel data (consistent with past estimations, limited to the Oregon II samples with 40 ft trawls) were updated through the Summer SEAMAP survey, which ended in July 2003.

Effort Data

Nance (2004) has summarized the shrimping effort data collection and analysis, including summary information from the several previous reviews of the program. One new addition to the effort database is estimation of variance, calculated as if interviews were a truly random sample of fleet activity. The variance estimates are based on a simple ratio estimator for ‘Effort per Unit Catch (EPUC).’ EPUC enters the effort estimate as a product with landings, which are treated as if known exactly. For these bycatch estimates, the Galveston lab supplied effort estimates based on the bycatch strata as unit cells, rather than the original shrimp statistical zone level of resolution (see Nance 2004). This lower resolution effectively eliminates criticism of area assignments, and reduces the ‘holes’ of landings not covered by interviews. This comes at a cost, as pointed out by Nance (2004), if the interviews do not cover the full range of shrimping activity in time and space for each cell in a representative fashion. Galveston was able to supply a file with mean and standard error estimates for effort for 1981-2002. I make bycatch estimates back to 1972, so I used previously existing estimates of effort (based on shrimp statistical zone resolution, summed to bycatch strata resolution), and assumed an arbitrary 6% CV for effort in each bycatch cell for 1972-1980. For 2003, I set all effort point estimates at 90% of their 2002 values, and raised the standard errors by 20%, to serve as an approximation until formal estimates for 2003 are available. Figure 1 shows a comparison of the annual point estimates between calculations based on the shrimp statistical zone resolution (“Old”) and the bycatch cell resolution (“New”). Figure 2 shows the annual totals for effort, and their 95% confidence intervals. (CVs on the annual totals are typically just under 2%. The bycatch estimates themselves are done cell by cell; with CVs on effort typically running 5-15%. Each cell has its own mean and variance estimate for effort.)

VOUF Data

Most observer data are collected on a ‘per net’ basis, with usually only one net worked up (two for BRD / noBRD comparisons). Shrimping effort data are on a ‘per vessel’ basis, thus estimates of ‘nets per vessel’ are needed to link observer CPUE and shrimping effort statistics. Over time, there has been a major trend from two nets to four in the offshore fishery. The only sources of nets per vessel information available throughout the time series come from the NMFS Vessel Operating Units Files. By convention, previous ‘base cases’ of bycatch estimates used the simple assumption of 2 nets per vessel, mainly because the VOUF files were often not up to date, and because there is some ambiguity in the VOUF data as described below. However, VOUF-derived adjustments to the bycatch estimates were presented to the review panels of 1997 as sensitivity cases. The nets per vessel issue has probably been the largest single source of bias in past estimates. In the current analysis, the uncertainty inherent in the nets per vessel value can be carried forward, so this is a good time to start incorporating the nets per vessel information into the standard bycatch estimates.

Unfortunately, nets per vessel has not been available on a trip by trip basis, and it is known that a large number of vessels will change the number of nets fished for different conditions. VOUF data are updated on an annual basis, using a convention that makes the data ambiguous. For an individual vessel, each port agents record the <u>maximum</u> number of nets observed or reported <u>in that port.</u> Thus, a single vessel may have multiple records, with perhaps 2 recorded in one port, and 4 in another. The data files also include many smaller vessels that probably fish mostly or even exclusively inshore, not really relevant to estimation of snapper bycatch. I dealt with these ambiguities by choosing four summary statistics for the VOUF files, and modeled uncertainty from them.

The four statistics are all means of the nets per vessel entries for subsets of each year's VOUF data: 1) all entries, 2) minimum nets per vessel recorded for all vessels, 3) all vessels>45 feet, 4) minimum nets per vessel recorded for all vessels>45 feet. Files for 1972-2002 were queried. Figure 3 shows the results.

Obviously, one would expect each of the four choices to contain some bias with respect to actual offshore fleet activity. I would expect the 'min >45 ft' to be the least biased, but I have no special knowledge to back that up. So to model a central tendency for nets per vessel, and an uncertainty around it, I took the simple mean of the four statistics each year, and calculated the standard deviation among them. That distribution is described by Figure 4 as a 95% confidence interval, and each iteration of the MCMC makes an independent draw from each (annual) distribution for every seasonal / spatial cell for that year. (For 2003, the mean was assumed to be 3.1, and the 2002 precision was cut in half.) The choices made to model nets per vessel are necessarily arbitrary, but the results give a plausible (and probably conservative) increasing trend for nets per vessel in the offshore fishery, with a broad uncertainty that shows we do not know the 'true' mean for fishing in any cell.

Modifications to Model 02

The BUGS Model 02 from Nichols (2004) was modified to add the features needed for this analysis. There are now 3 dataset parameters instead of 2, one for research vessel data, one for observer data without BRDs, and one for observer data with BRDs. Normal distributions were added to model the uncertainty around the effort and nets per vessel statistics. These distributions are not 'narrowed' by the observer data in any way -- the posterior distributions for CPUEs are calculated just an in Nichols (2004). It is the functions of those CPUE posteriors that now include random draws from the normal distributions for effort and nets per vessel, rather than the constant values used in Nichols (2004). Because BRD requirements were phased in during 1998, actual bycatch estimates use the BRD predictions in cells requiring BRDs, and the noBRD predictions in cells not requiring BRDs. That is, each spatial/temporal cell is either a BRD cell or a noBRD cell, with no attempt to subdivide a cell to allow for different requirements in different spatial or temporal areas within cells, and no attempt to incorporate 'degree of compliance' as a factor. Annual totals are presented here for simplicity, but trimester totals are available, just as in past assessments. The BUGS code for the modified Model 02 is attached as an Appendix.

With the additional data and parameters incorporated here, the modified Model 02 is now right at the memory limit of the best machine I have available. For that reason, I have limited the analysis to a 2-chain run of 20k iterations.

Results

Figure 5 shows the BUGS boxplot of total annual bycatch. (By convention in BUGS, the thin line ends at the 2.5% and 97.5% quantiles – a 95% confidence interval, while the thicker bar shows the interquartile range. Figure 6 shows the boxplot of the distribution for the natural log of the annual totals. The distributions are closer to symmetric on the log scale, but there is still some extension of the high-side tail.

The same type of goodness of fit plots used in Nichols (2004) shown in Figure 7 (log scale) and Figure 8 (arithmetic scale).

Discussion

One might be able to see some expansion of the confidence intervals from incorporation of the uncertainties for effort and nets per vessel, and some adjustment of the central tendencies due to the nets per vessel trends if Figures 5 and 6 are compared with the similar figures of Nichols (2004, Fig. 8), but the results are not dramatic. These effects can be discerned in tabled values that generated the plots, but in fact these structural changes probably had less effect on the overall results than the new data additions did. However, the real message remains the overall uncertainty in the bycatch estimates, which is dominated completely by the uncertainty in the bycatch CPUE estimates.

In the past, we were always concerned that effort variance was a major, unestimated component of bycatch uncertainty. Frankly, I was expecting CVs for effort to be around 20-25% for the annual totals, not the 2% actually obtained. On reflection, I had no real basis for a 20-25% expectation. Precision is clearly not a serious issue regarding shrimping effort. There are still concerns in the general community about accuracy, but past reviews have not identified potential biases with certain direction. Differences in results among the partitioning choices considered by Nance (2004) had a higher spreads than confidence intervals calculated under this single choice used here, but even those spreads would be negligible compared to CIs of the bycatch CPUE estimates. The factor related to effort that might be most limiting to the accuracy of the bycatch estimates is probably the mismatch in time resolution between observers and interviews. Time fished is recorded to the minute by observers, at the time of each trawl. Interviews provide summary estimates after the fact, probably to an accuracy of +/- a few hours or more per trip. No direction is necessarily implied by this difference. If the community found the issue worrisome enough, switching to a permit / logbook system could probably improve the situation, but it probably wouldn't be cheap.

Estimating nets per vessel via the VOUF files is clearly not ideal. We know vessels change their configurations as conditions change, but the unusual conventions of the VOUF data complicate interpretation of any estimate. I have consider other approximations beyond using a common mean for all cells in a year (things like assuming 2 nets nearshore, 4 nets offshore), but I could not come up with a good reason to justify these more elaborate assumptions. The distributions for nets per vessel used in this report cover broad ranges, and are dominated by a real uncertainty about what the "true" mean number of nets per vessel is at any time. Real variation around each "true" mean is probably a minor factor by comparision.

I have in the past expressed great reluctance to continue any indirect model approach into the early years of BRD implementation. Imposing successful BRDs would change the temporal and spatial distributions of the bycatch species, making the "main effect" coefficient "wrong" until a substantial history of new data accumulated. Foster's (2004) estimates for current BRD performance shows my concerns to be unfounded, at least at present. The effects of BRDs right now on spatial and temporal distributions must be small, comparable to the effects of fluctuations in effort that we have always ignored in the modeling. Should BRDs become more successful, our models may need more complicated structures, if indirect modeling must be continued. However, I continue to urge development of a stable, comprehensive observer program, as recommended by the 1997 peer reviews. With such a program, indirect models like the GLM or this Bayesian approach may become unnecessary.

References

Foster D. 2004. 1999-2003 north-central and western Gulf of Mexico BRD performance. Report to the Gulf of Mexico Fishery Management Council. Red snapper data SEDAR, April 2004.

Nance J M. 2004. Estimation of effort in the offshore shrimp trawl fishery of the Gulf of Mexico. Report to the Gulf of Mexico Fishery Management Council. Red snapper data SEDAR, April 2004.

Nichols S. 2004. Some Bayesian approaches to estimation of shrimp fleet bycatch. Report to the Gulf of Mexico Fishery Management Council. Red snapper data SEDAR, April 2004.

Scott-Denton E. 2004. Observer coverage of the US Gulf of Mexico and southeastern Atlantic shrimp fishery, February 1992 – December 2003 – methods report to SEDAR. Report to the Gulf of Mexico Fishery Management Council. Red snapper data SEDAR, April 2004.

Figure 1. Comparison of effort calculated at the shrimp statistical zone level of resolution, then summed to the bycatch strata (old), vs effort calculated using the bycatch strata as the effort cells (new).

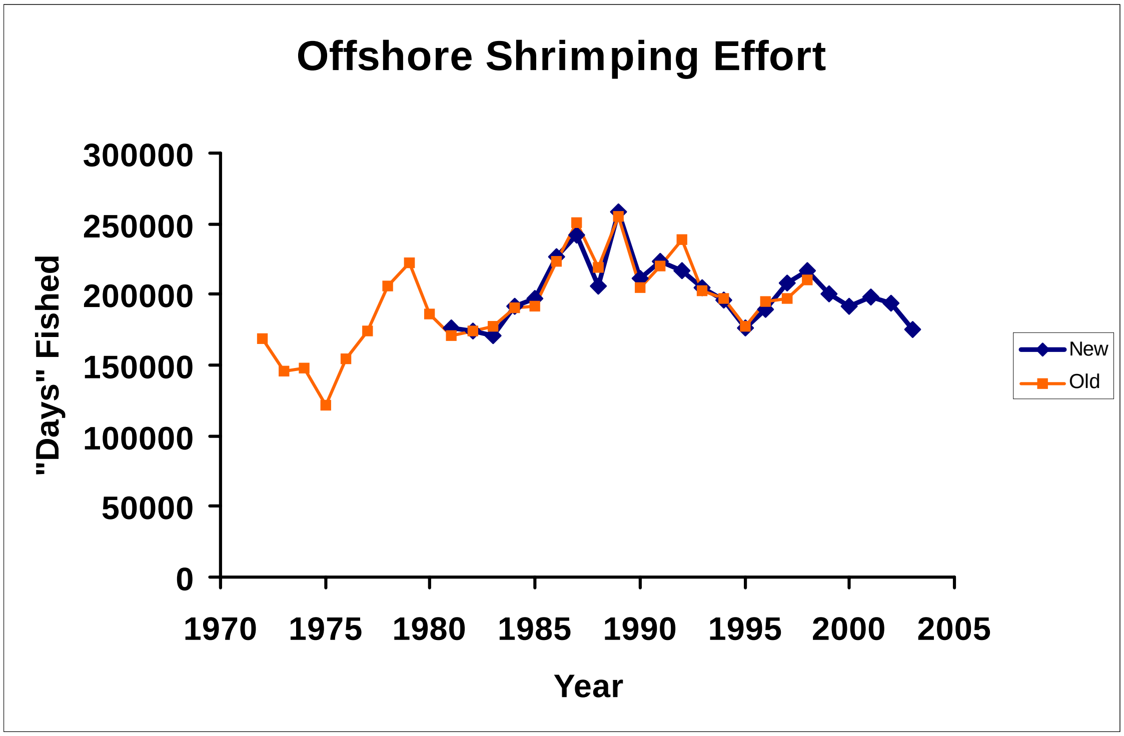

Figure 2. Confidence intervals (95%) for annual totals of shrimping effort, estimated by treating interviews as if they were random samples of fleet activity. Variances for the 1972-1980 were not estimated; instead a 6% CV for each (bycatch) cell was assumed. The higher spread in 1972-80 compared to the 1980s suggest that choice was very conservative.

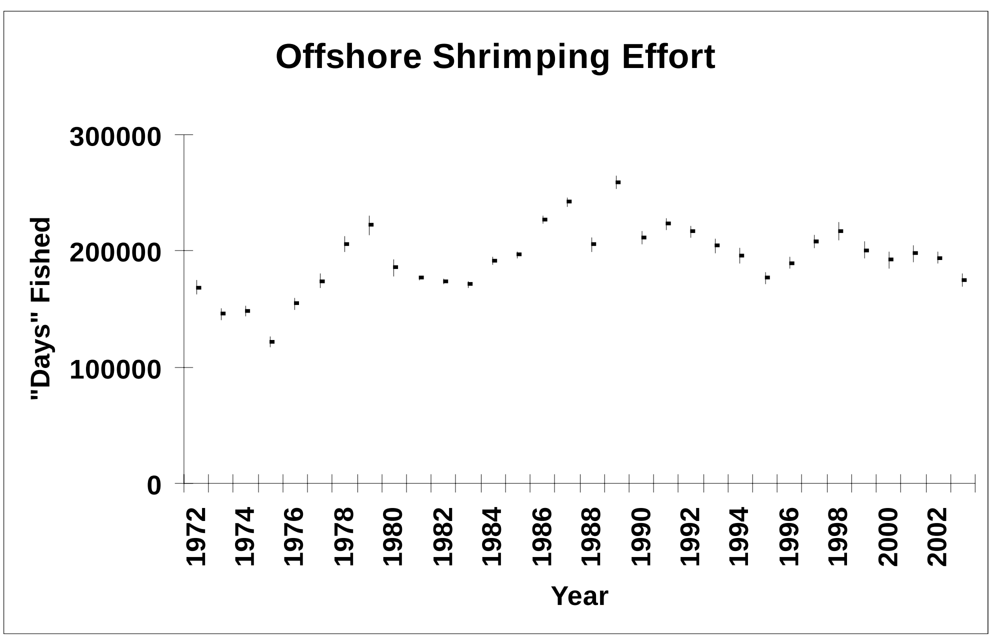

Figure 3. Average nets per vessel for the four statistics calculated from the annual VOUF files. The anomalous points for >45 in 1998 are probably an artifact of a large number of omissions of vessel lengths in that year's files. (No adjustment was made for this anomaly.)

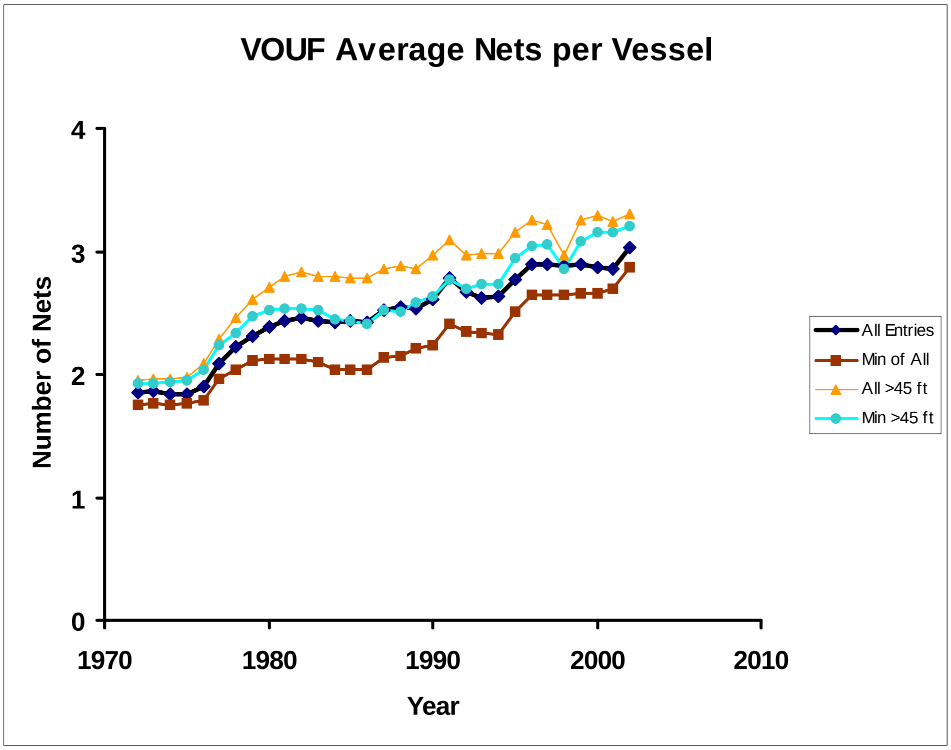

Figure 4. The distribution used to model nets per vessel in the MCMC simulation. The distributions were normal, with the 95% confidence intervals shown.

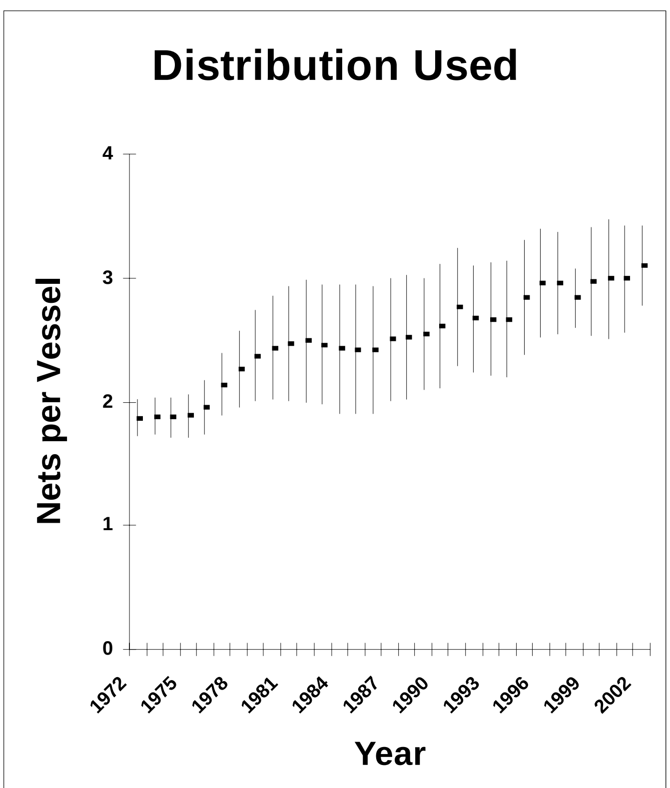

Figure 5. Box plot for estimates of annual totals of red snapper (millions of fish) taken in the offshore shrimp fishery in the Gulf of Mexico. The horizontal axis is years, with 1 being 1972 and 32 being 2003.

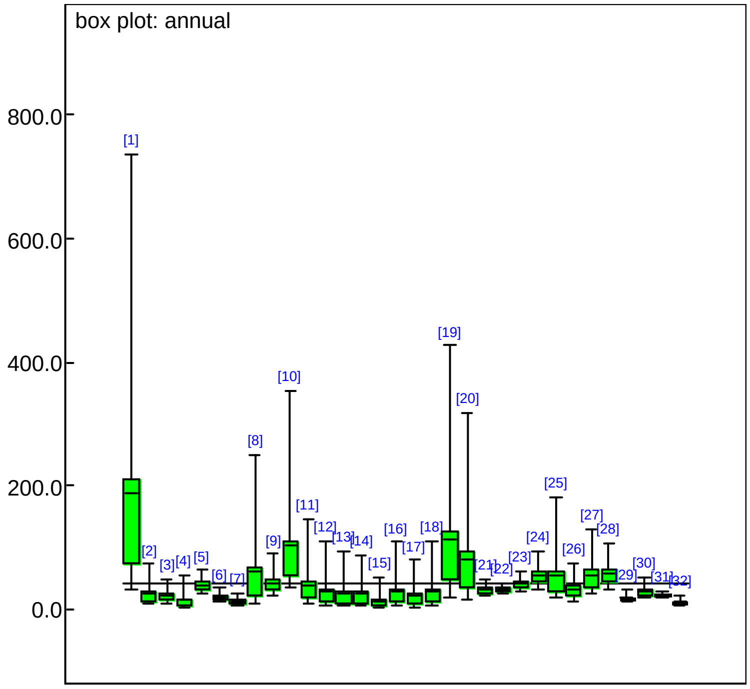

Figure 6. Box plot for red snapper bycatch estimates shown on a log scale.

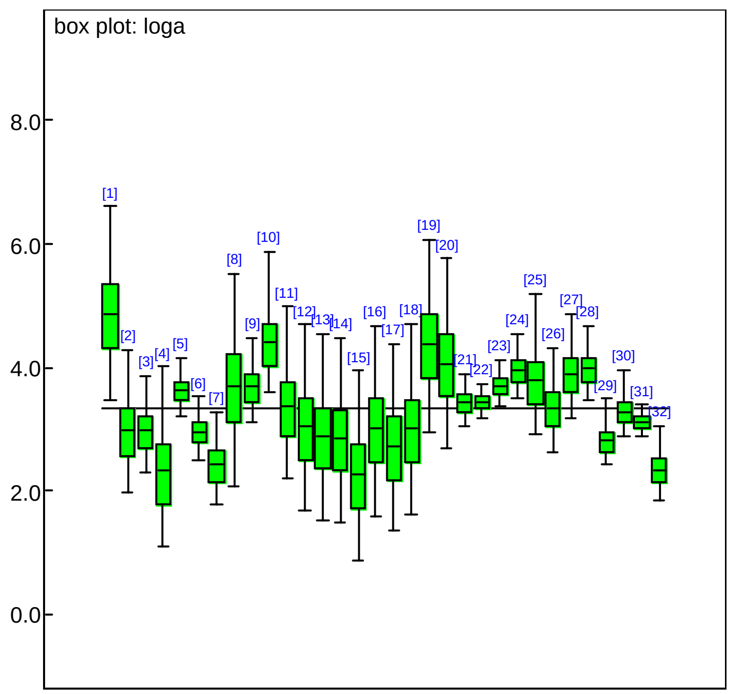

Figure 7. Goodness of fit plot (medians of log CPUE posteriors vs observed log (mean CPUE) in cells with observed means > 0). The points are coded to dataset: 1 is observer, no BRD, 2 is research vessel, and 3 is observer, with BRD.

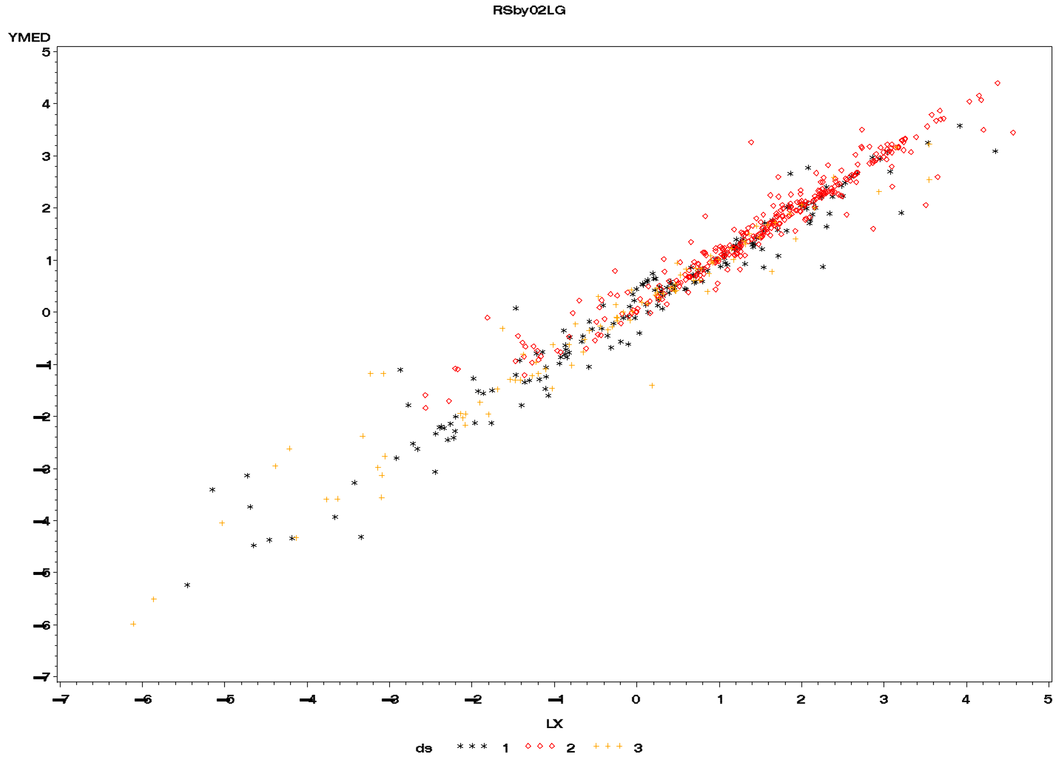

Figure 8. Goodness of fit plot on arithmetic scales. (Medians of CPUE posteriors vs means of observations for cells having data.)

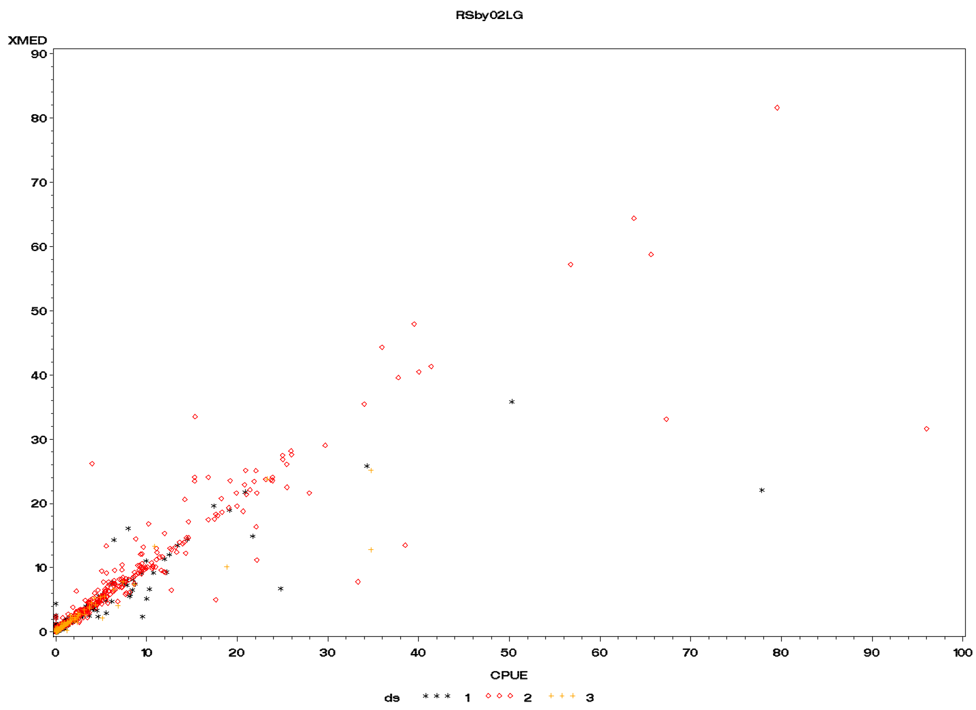

Appendix. Bugs code for the 3-dataset, error-in-effort model used in this paper.

```
model rsbycatch02 {

r~dunif(0.03,5)
tau~dlnorm(0,3.5)
center~dnorm(0,tau)

for (i in 1:32)  {
  yx[i]~dnorm(1,0.7)
  }
for (j in 1:3)  {
  sraw[j]~dnorm(0,1)
  sx[j]<-sraw[j]-mean(sraw[])
  }
for (k in 1:4)  {
  araw[k]~dnorm(0,0.2)
  ax[k]<-araw[k]-mean(araw[])
  }
for (l in 1:2)  {
  zraw[l]~dnorm(0,0.2)
  zx[l]<-zraw[l]-mean(zraw[])
  }
for (m in 1:3)  {
  draw[m]~dnorm(0,1)
  dx[m]<-draw[m]-mean(draw[])
  }
for (i in 1:32)  {
  for (j in 1:3)   {
    for (k in 1:4)  {
      for (l in 1:2)  {
        for (m in 1:3)  {
          local[i,j,k,l,m]~dnorm(0,tau)
          logy[i,j,k,l,m]<-yx[i]+sx[j]+ax[k]+zx[l]+dx[m]+local[i,j,k,l,m]
          y[i,j,k,l,m]<-exp(logy[i,j,k,l,m])
          mu[i,j,k,l,m]<-r/y[i,j,k,l,m]
          }
        }
      }
    }
  }
for (h in 1:40550) {
  lamb[h]~dgamma(r,mu[yr[h],seas[h],ar[h],dp[h],ds[h]])
  lambda[h]<-lamb[h]*hrsfishd[h]
  catch[h]~dpois(lambda[h])
  }
for (i in 1:26)  {
  for (j in 1:3)   {
    for (k in 1:4)  {
      for (l in 1:2)  {
        effort[i,j,k,l]~dnorm(effmean[i,j,k,l],efftau[i,j,k,l])
        npv[i,j,k,l]~dnorm(voufmean[i],vouftau[i])
        take[i,j,k,l]<-y[i,j,k,l,1]*npv[i,j,k,l]*effort[i,j,k,l]
        }
      }
    }
  }
    for (k in 1:4)  {
      for (l in 1:2)  {
        effort[27,1,k,l]~dnorm(effmean[27,1,k,l],efftau[27,1,k,l])
        npv[27,1,k,l]~dnorm(voufmean[27],vouftau[27])
        take[27,1,k,l]<-y[27,1,k,l,1]*npv[27,1,k,l]*effort[27,1,k,l]
        }
      }
      for (l in 1:2)  {
        effort[27,2,1,l]~dnorm(effmean[27,2,1,l],efftau[27,2,1,l])
        npv[27,2,1,l]~dnorm(voufmean[27],vouftau[27])
        take[27,2,1,l]<-y[27,2,1,l,1]*npv[27,2,1,l]*effort[27,2,1,l]
        }
```

```
     for (k in 2:4)  {
      for (l in 1:2)  {
        effort[27,2,k,l]~dnorm(effmean[27,2,k,l],efftau[27,2,k,l])
        npv[27,2,k,l]~dnorm(voufmean[27],vouftau[27])
        take[27,2,k,l]<-y[27,2,k,l,3]*npv[27,2,k,l]*effort[27,2,k,l]
        }
      }
    for (k in 1:4)  {
      for (l in 1:2)  {
        effort[27,3,k,l]~dnorm(effmean[27,3,k,l],efftau[27,3,k,l])
        npv[27,3,k,l]~dnorm(voufmean[27],vouftau[27])
        take[27,3,k,l]<-y[27,3,k,l,3]*npv[27,3,k,l]*effort[27,3,k,l]
        }
      }
for (i in 28:32)  {
  for (j in 1:3)   {
    for (k in 1:4)  {
      for (l in 1:2)  {
        effort[i,j,k,l]~dnorm(effmean[i,j,k,l],efftau[i,j,k,l])
        npv[i,j,k,l]~dnorm(voufmean[i],vouftau[i])
        take[i,j,k,l]<-y[i,j,k,l,3]*npv[i,j,k,l]*effort[i,j,k,l]
        }
      }
    }
  }
for (i in 1:32) {
  annual[i]<-sum(take[i,,,])
  loga[i]<-log(annual[i])
 }
for (i in 1:32) {
  for (j in 1:3)    {
    trimester[i,j]<-sum(take[i,j,,])
    logt[i,j]<-log(trimester[i,j])
   }
  }
}
list(tau=0.1)
list(tau=1.2)
```