UNIVERSITY OF
Miami

July 22, 1999

Dr. William W. Fox, Jr.
Director, Office of Science and Technology
National Marine Fisheries Service
1315 East-West Highway
Silver Spring, MD 20910

Dear Dr. Fox:

The Center of Independent Experts has completed another review. For this review, Dr. Mark Bravington (U.K.) participated in a SARC Panel meeting in Woods Hole. A product of his review will be reflected in the SARC documents that are being finalized by the NEFSC. Dr. Bravington also provided us with a more detailed set of comments on the SARC process and on the CIE involvement in that process (enclosed). In addition, he provided very detailed suggestions for improving a difficult scallop assessment. He already forwarded the latter to NEFSC scientists, but I enclose it here for completeness. As usual, I assume that your Office will forward the documents to the appropriate places in NMFS.

Sincerely,

Robert K Cowen

Robert K. Cowen
Professor and Maytag Chair of Ichthyology

c: Steering Committee Members
V. Restrepo

encl.: Review, scallop comments

Rosenstiel School of Marine and Atmospheric Science
Division of Marine Biology and Fisheries
4600 Rickenbacker Causeway
Miami, Florida 33149-1098
305-361-4182

Fax 361-4600

Email: mbf@rsmas.miami.edu

Report on SARC-29 Participation
Mark Bravington, CEFAS
July 21, 1999

Overall, my impression is that the external-assessor scheme is a very good idea. I hope my contributions were useful. From a personal viewpoint, I enjoyed the work and found the exercise very instructive. I certainly hope CEFAS (incl. me) will participate again; I think there are great benefits for both parties.

I was in general very impressed with the quality of the assessments, the background understanding of the stocks and fisheries, the hard work that NEFSC personnel undertook behind the scenes during the week, and the amount of thought that had gone into dealing with difficult and really non-standard issues. It seems to me that the NMFS process of assessing species as & when necessary leads to higher-quality assessment than we generally see in European ICES work, where most species never get assessed at all and, for the remainder, the burden of work entailed in annual assessment tends to lead towards a mindless application of the rulebook.

In a few cases, however, I got the impression that the SAWs may have "kicked upstairs" some difficult issues for the SARC to deal with, while in fact the SAW itself had a pretty good idea what to do and was rather better-placed to get it done. The recent recruitments for witch flounder are a case in point. The SARC struggled with it for quite a while and then Ralph Mayo bailed us out. The witch flounder assessors had already spotted the problem, and in fact knew exactly how to fix it! This had its funny side, and I am not criticizing the assessment team, who were very on-the-ball. However, the SAW/SARC process might benefit from a bit of tinkering.

The demands on NEFSC's time are surely very high already and I certainly wouldn't want to suggest more work. But the time interval between SAW and SARC seems very short to me; if the gap were extended, it might be possible to run a few more analyses before SARC. This would be easier and safer than trying to do all the stuff at 3am during SARC week. The SARC itself is ill-equipped to deal with problems of detail, as we saw very clearly when the issue of scallop growth curves was suddenly brought up at 2pm on Friday. A longer SAW-SARC gap might also allow for clearer reports. I had to reread the scallop assessment document several more times than I would have liked, in order to figure it out.

A few other comments:

1) The nominal time of 7 days isn't really enough to do a thorough job, especially if one is acting as SARC leader and if one is to give extra comments after the meeting. In practice, with travel, preparation and "postparation", I have spent more than the 8 days we are getting covered and could have done with even more time-- not so much for scallops, but to give more thought to the other species. At least 10 days would be more realistic. Having said this, no real problem arose for SARC 29 because of the way the

finances worked out; the extra time was compensated by a higher daily rate. It may be that this compensatory mechanism is exactly what's supposed to happen, but I thought you should be told. Ultimately, it all depends on how much work is wanted from an external assessor.

2) More prior information on the "management advice" side would be very helpful, especially if a non-NMFS assessor is asked to act as SARC leader and therefore to have input to the advisory report. Commenting on the science behind the assessment is of course a very different matter to teasing out the implications for management. It wasn't always clear to me what the management framework was, and what the SARC's options were when constructing the reports. For instance, what would happen if an assessment were rejected outright, and no alternative presented? How far can the SARC go in doing on-the-spot calculations? And so on.

In a sense, this isn't the external assessor's problem: constructing the reports is largely the chairman's job, the management implications have to be followed up by NMFS, and the impact is ultimately felt by the fishermen and consumers. But one cannot do meaningful assessments without regard to the management context, and there are issues such as the reference point/open-closed mortalities for scallops where one simply has to say something! So, it would be very useful to get a short piece of paper stating what the SARC's options are, and what would happen post-SARC in each case.

3) When we got to the stage of writing up management advice, several people suggested comments for the report that may well be true, but seemed to come out of nowhere. The comments were presumably borne of many years experience on the stocks and fisheries, but without supporting documentation it's difficult for an external assessor to justify putting them into the reports. This is a shame-- the comments may well deserve to be heard, and the SARC has a role to play in supporting and broadcasting them. But this can't really be done unless a more explicit motivation is available. For example, if someone wants to point out that there is grotesque overcapacity for fishing species X, then a short working paper with a few calculations would do fine. If the assessments were NMFS-only, this might not be necessary. But in the interests of transparency at least, it seems desirable to get it down on paper.

4) Electronic versions of old SARC and SAW documents would be very nice and would save on Fedex delays/costs.

Suggestions on Sea Scallop Modelling
By Mark Bravington
July 23, 1999

[NOTE: These specific comments and suggestions pertain to draft document B1 from SARC-29 entitled "Stock Assessment of Sea Scallop" by the NEFSC Stock Assessment Team]

======================

Most important is to get the data reconciled. By all means, be Bayesian thereafter if you need to be, for example by incorporating a prior on dredge efficiciency. But if the various data sources are incompatible with the underlying dynamics model, then Bayesianism won't help; indeed, it may induce a false sense of security that one has somehow "balanced" the competing signals in the data.

By "reconciled" I mean that the only apparent way to match up the empirical Z-estimates from $n_{t+1}/(n_t + r_t)$ with the catch numbers and the absolute abundance estimates from swept area, is to have unrealistically high dredge efficiencies in the model. In particular, the implied efficiencies are much greater than the values suggested by the depletion experiments and the general literature.

It is difficult to believe that the experimental efficiency estimates are way off, so the resolution must lie elsewhere. Now, roughly speaking $\hat{z} = c_t / n_t$ so $n_t = c_t / \hat{z}$; the n_t are way too low (cf. swept area) so either c_t is too low or $\hat{z}$ is too high. I could see 3 possible explanations for the catch-in-numbers series being too low a measure of c_t:

i. There is significant non-yield fishing mortality, as suggested by published studies and the George's Bank depletion experiment results.

ii. Commercial size distribution data is/was unrepresentative of catches, and is/was biased in favour of large scallops. This would imply that estimates of scallop numbers, based on landings in weight divided by mean weight, are/were too low.

iii. Total landings in weight have been systematically underreported.

The other possibility is that the $\hat{z}$ are too high. Of course, M might be in error, but since it's low anyway (M=0.1), I don't think that the problems could be explained by greatly reducing M. It is, however, possible to get the wrong z-estimate if selectivity is hugely in error. Neglecting M, we have

$$z \approx c_t / n_t \approx (n_t + r_t - n_{t+1}) / n_t = (1 - n_{t+1} / n_t) + r_t / n_{t+1}$$

and for a population in rough equilibrium, $z \approx r_t / n_t$. So if the r_t we are using is too high relative to n, we will get estimates of z that are also too high. This is another reason to investigate the selectivity. I guess problems with the growth curve could also be responsible, e.g. if the lower limit for pre-recruits is too low, or the upper limit overlaps heavily into recruits proper. [While I remember, it really was very confusing to hear about animals being "unavailable" at a size which clearly does get taken by the fishery. Indeed I think that some of the "full recruits" were apparently unavailable. There are surely conventions for this sort of thing which would be easier to understand. This whole growth/recruitment issue needs to get sorted out, but it is too detailed for me to get into here].

One or more of these possibilities (three catch problems, selectivity, or growth) is the culprit. This needs to get straightened out before worrying about the model itself.

I present detailed comments on the existing model below. Of course, if the model structure ends up being changed drastically, then these comments may become redundant. Some of the suggestions are admittedly a bit sketchy, but there doesn't seem much point in going into excessive detail at this stage.

1) A general gripe first: the notation is generally a bit strange, and sidesteps the conventions. It is useful to use capital letters for random variables, small letters for realized values of those variables, etc. I certainly would have understood the model much more quickly if the conventions had been followed: the underlying model is actually pretty simple, but the SAW report does not convey this impression. (And NB that simple=good here). One particularly confusing notation is $\hat{n}_{ito}$ for the mean projected population size; hats (circumflexes) normally denote parameter estimates, which $\hat{n}_{ito}$ isn't. It would have been clearer to write

$$\mathbf{E}[N_{t+1,o} \mid n_{t,o}]$$

or similar, to spell out exactly what's going on in the state space. Here $N_{t+1,0}$ is a random variable which depends on the (already fixed) population size $n_{t,0}$ in the previous year.

2) I do not think that pre-specifying the relative variances of different contributions to the likelihood is a good idea. It is certainly very un-Bayesian in spirit, and, what's more, I don't think it is actually necessary in this instance. It is true that classical statisticians cannot do regressions of y on x when the measurement error variances of both x & y are completely unknown. Bayesians would normally get around this by specifying a joint prior on the variances. But in this instance, I think you can do away with the whole problem if you neglect the process errors in the evolution of n_{t+1}. I _think_ this translates to setting $\eta_{i,t,o/c}$ to zero in equations E5,E6 and E7, and setting $\sigma_4^2 = \sigma_5^2 = \sigma_6^2$ in F4,F5,F6.

If you do away with these process errors, and if you also allowed all the $N_{1,t}$ to be free (estimable) parameters-- bear with me-- you would pretty much have a straightforward separable situation. Neglecting M, you'd have

$$\mathbf{E}[U_{2t}] = q_2 N_{2,t} = q_2(N_{1,t-1} + N_{1,t-2} + N_{1,t-3} + \ldots + N_{1,1} - C_t - C_{t-1} - \ldots - C_1)$$
$$\mathbf{E}[U_{1t}] = q_1 N_{1t}$$

(note I have abused my suggested notational conventions here, sorry). q_1 and q_2 are efficiencies for partial & full recruits respectively. This looks like 2T observations for about T+1 parameters. I don't think the parameters are aliased, so this could be estimated straightaway by maximum likelihood methods without recourse to priors on q or on other things. As a rough justification, consider the following. Write

$$\Delta U_{2t} = U_{2t} - U_{2,t-1} = q_2 N_{1t} - q_2 C_t + \varepsilon_t = e^{Q_t} - q_2 C_t + \varepsilon_t \approx 1 + Q_t - q_2 C_t + \varepsilon_t$$
$$\log U_{1t} = \log q_1 + \log N_{1t} + \eta_t = \log q_1 - \log q_2 + Q_t + \eta_t = q^* + Q_t + \eta_t$$

where $q^* = \log q_1 - \log q_2$ and $Q_t = \log(q_2 N_{1t})$. We can then define

$$U_{3t} = \Delta U_{2t} - 1 - \log U_{1t} = -q_2 C_t - q^* + \varepsilon_t - \eta_t$$

from which $-q^*$ and $-q_2$ can be estimated as the intercept and slope, respectively, of a regression of U_{3t} on C_t. Once q^* and q_2 are available, q_1 can be found, then the variances of ε and η can be estimated, and so N_{1t} can be reconstructed from a weighted average of U_{1t}/q_1 and $\Delta U_{2t}/q_2 - C_t$. (Note I am not suggesting anyone actually *uses* this regression; it is based on a very crude linearization of the exponential term. I am just trying to establish that the parameters are estimable without any need to fix variances.)

Even if this no-process-error model can be bettered by a more complicated form, it might be instructive to do the fitting, as the model is so simple. In particular, I think this model would have drawn attention to the discrepancies with much less work.

The next step in complexity would be to make the $N_{1,t}$ into random variables, with an unknown (estimable) mean and variance. In classical statistics, this would be called a "mixed model" or "random-effects" model; there is no problem with estimability. (However, because of the separable-type structure, I don't think you can fit this one with an off-the-shelf command, though I may be wrong here). Bayesians would devise some prior on the mean & variance and use a hierarchical modelling framework. But either way, it's an easy modelling framework to use, as once the $N_{1,t}$ are set, the evolution of that cohort is deterministic. This is a great advantage as it vastly cuts down the number of parameters that need to be considered and AVOIDS the need to fix variances in advance. This random-effects-style model would probably be my ideal level of complexity.

Although there is sure to be some process error in reality, it seems to me OK to ignore it in this instance. Variable recruitment and measurement errors are probably dominant, particularly so when fishing mortality is high. Also, given the discrepancies that the model has highlighted, one doesn't want to give the model too much leeway in "bending" the population dynamics to fit the assumptions and catch series, especially from a diagnostic viewpoint.

I know lots of other people have used fixed variance ratios in modelling. It does make computation simpler, but it's a pretty dubious practice and I would rather use a more sophisticated statistical approach or (as suggested here) a slightly simplified model. In simple situations, e.g. with one process and one observation error, one can check to see whether conclusions are sensitive to the variance ratio assumed, although it's not obvious how far to extend the tests (what ratios are "reasonable"?). However, when there are 11 variances (document table F), sensitivity testing is no longer really feasible, or at any rate becomes a huge task which requires as much thought as a more satisfying reconstruction of the model anyway.

3) Personally, I would have tried non-autocorrelated recruitments before going to the more complicated autocorrelated case. My guess is, there won't be much of a modelling gain by specifically allowing for autocorrelation. Even if no autocorrelation is built in, the data will still induce it if the signal is strong enough. The pattern of estimated recruitments (maximum posterior marginal densities) might give one a sense of how important autocorrelation really is.

4) I would tend to structure the model as inside-recruitment and outside-recruitment throughout the entire time series, but to only consider composite survey indices and removals before the closing of the areas. This gets around the problem of having no priors on N_{1to} etc. after the closure. Otherwise, the longer the closures go on, the more of an inconsistency there would be with the pre-closure model. If you don't worry too much about autoregression, there should be no problem here. Incidentally, is there any correlation between partial recruitment in the open and closed areas (both before and after closure?).

5) It is definitely not OK to use the proportion variables as well as the raw indices in the model. Each term in the eventual "likelihood" must be viewed as conditional distribution given the data used in the preceding terms, and once the raw survey data have been incorporated there is no variability left to explain or contribute to the likelihood in the proportions. Ignoring this amounts to using the data twice. Naturally, this will tend to stabilise things-- if you replicated all the data in the model 50X over, you would appear to get much more precise results. But you are not allowed to.

6) It is also not OK to set priors on the parameters based on *ad hoc* regressions using the same data that goes into the likelihoods (or other priors). Again, this amounts to doubling up the data. I would rather see "flat" priors, priors based on other experiments, or perhaps priors based on "regressions" using just 1 data point. This problem applies (I think) to q_a and γ.

7) MPLE is untrustworthy for use in parameter-rich situations, such as here where you have all the N_{ito} floating about. It makes much more sense to stick with marginal posterior distributions from MCMC. If your model is too complicated to get MCMC working, try a simpler model. Another reasonable thing would be to

use a hybrid between marginal and profile methods: perhaps the thing to do is to integrate the posterior distribution over the N_{ito} and look at a likelihood profile for the remaining *structural* parameters of the model. I haven't thought about exactly how to do this.

8) I don't have experience of applying MCMC. I am told, though, that it's a bit of a black art to try and get convergence within human lifetimes. Important things to do including keeping the number of parameters low, and trying to make them as "orthogonal" as possible. (A good reason to get rid of as much process error as one can.)

CPUE data

Commercial CPUE is another fertile source of disagreements between models and data. As with the efficiency issue, if commercial and survey data are giving different signals, it is *not* appropriate to just throw them both into the model and rely on Bayes rule to sort it all out. [This is not sniping at Bayesians. It is also not OK to trust to some ad hoc weighting algorithm to do it for you, which is widely done in ICES.] The real issue here is *model uncertainty*-- we aren't fully aware of the mechanisms that give rise to the data, and there are different possibilities-- and this is often not addressed in statistical practice, be it Bayesian or frequentist.

My general feeling these days, is to include commercial CPUE only if it doesn't change the answers much (compared to the survey-only case). Of course, CPUE can potentially improve the precision of estimates because it is based on much larger samples, but IMHO it should basically be viewed as *guilty until proven innocent*. Scallops are a notoriously patchy resource, where you get local concentrations resulting from a good settlement, and fishermen moving around to hit these hot-spots. This is a classic case where CPUE can go wrong as a measure of total abundance. Figure 16 in the SAW report-- which I can't find referred to in the text-- seems to show a trend in CPUE that is different to the trend in survey indices for the Mid-Atlantic at least. And then there are other peculiarities-- such as vessels shifting where they fish in response to the closing of an area, or the changing pattern of days fished to days at sea. For all these reasons, there is a clear basis to be worried about the utility of CPUE data.

I would suggest re-fitting the models without CPUE data at all in the first instance, then seeing whether CPUE data are giving similar trends to the surveys (I mean similar relative magnitudes of changes, not just similar direction), and only then would I try incorporating CPUE into the model.

Yield-per-recruit calculations under different closure regimes

During the meeting I glibly suggested that this could be done easily. I was wrong. A typical closure regime will at any one time have patches in it that have been open for differing numbers of years. To calculate YPR, one needs to know how fishing effort will distribute itself among these patches, which isn't obvious. We can make a good stab at figuring out the effort distribution, but the inevitable downside is considerable extra complexity in figuring out the F-pattern before YPR can be calculated.

The overall structure that I would use would be a series of trials with varying percentage-of-area-closed, varying rotation periods, and varying levels of overall fishing effort (#hauls/day). The no-closure scheme should obviously be included. There are different ways to do the rotation as well. For instance, one might make a complete switch of closed areas at the end of every n-year term. Another might be a rolling closure in which the total area is divided into n+m bits of which n are closed at any one time, and each of the closures operates on a staggered n-year schedule-- so that every year 100(1-1/(n+m))% of the closed area is in the same place as last year. I would imagine that rolling closures would be vastly preferable to the industry, as they would ensure stability and cut down the "gold rush" effect that is presumably now happening in the newly-reopened area. I have concentrated on this case below, but similar calculations could be done for other rotation schemes.

This is a figure to illustrate the above rolling-closure regime. Here, 40% of the total seabed area is closed at any one time (the shaded bits), and the rotation period is 4 years. The white slices are currently open. Numbers illustrate where each closed slice lies within its closure period, so the darkest slice is currently in its fourth year of closure. Next year, the darkest slice will be opened, one of the white slices will be closed, and the whole picture will rotate one segment anticlockwise. Clearly, the actual spatial arrangement is irrelevant, since scallops have the good grace to not move much. It's not even necessary for each segment to be simply connected (e.g. you could tile the entire seabed with circles like this, and have the same rotation pattern within each circle).

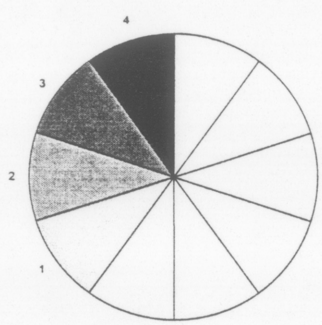

Before getting into details, I point out that the issue of non-yield-fishing-mortality will be critical and definitely needs to be incorporated into any calculations, though I haven't done so below. If no good estimate is available, then one will again have to use a range of possible figures. It may be that the *relative* performance of different closure regimes is unaffected by the level of NYFM. It would also be worth seeing how sensitive the conclusions are to different values of M.

To find the expected yield under such a closure policy, we need to know how fishing effort will distribute itself across newly-opened/less-newly-opened slices. This is tricky. To figure it out, I assume that the fishermen are omniscient, rational and not acting co-operatively. I get the impression that this is a sophisticated and mobile fishery that can hit the best slices until they are just as depleted as everywhere else.

In a sense, the strategy is obvious: start by only fishing in the newest slice until revenue-per-unit-effort in that slice drops to the same level as in the next-newest slice, after which effort is distributed across both the newest and next-newest which are both fished down till RPUE hits the next-next-newest, and so on. Logically, this switchover must happen; suppose it didn't, and they just fished down the new slice for the whole of its first year. The biomass left at the *end* of that year would be the same as the biomass at the *beginning* of the year in the next-newest slice (since they did the same thing last year on that slice). Thanks to growth and recruitment, the biomass in the next-newest slice would increase during the year while it is being left alone, and at some point its increasing trajectory would cross the decreasing trajectory in the newest slice. At that point, the fishermen would make more money if they switched to the next-newest slice.

We also need to figure out how will the effort be split once two slices are being exploited together. We can assume that the proportion of fishing mortality in each of the two slices will vary over time so as to keep RPUE the same in the face of the differing intrinsic rates of growth of RPUE in the two slices-- otherwise everyone would shift to the slice in which the growth rate was higher, until things evened out again.

Here's a mathematical formulation of the way that fishing effort is expected to split. Definitions: RPUE in slice i is R_i; the per-capita intrinsic rate of RPUE growth at time t for animals of age a is $r_a(t)$; the number of such animals in slice i is $n_{ia}(t)$; the selectivity pattern is $q_a(t)$; there are $P(t)$ slices being fished at time t (those with the equal highest RPUE); natural mortality at age is $m_a(t)$; the proportion of fishing effort going into slice i is $\alpha_i(t)$ (with $\sum_i \alpha_i(t) \equiv 1$). Dropping the dependence on t in the notation, we can write

$$
\begin{aligned}
&R_i = \sum_a n_{ia} r_a \\
&\Rightarrow \frac{\mathrm{d}R_i}{\mathrm{d}t} = \sum_a r_a \frac{\mathrm{d}n_{ia}}{\mathrm{d}t} + n_{ia} \frac{\mathrm{d}r_a}{\mathrm{d}t} = \sum_a r_a \dot{n}_{ia} + n_{ia} \dot{r}_{ia} \\
&\dot{n}_{ia} = -z_{ia} n_{ia}, \quad z_{ia} = m_a + \alpha_i q_a F \\
&\Rightarrow \dot{R}_i = \sum_a n_{ia} (\dot{r}_a - r_a (m_a + \alpha_i q_a F))
\end{aligned}
$$

For equality among the $\dot{R}_i$ in the P exploited slices, we need

$$\alpha_i S_{1i} + S_{2i} \equiv \alpha_j S_{1j} + S_{2j}$$

where $S_{1i} = F \sum_a q_a r_a n_{ia}$, $S_{2i} = \sum_a n_{ia}(r_a - m_a r_a)$. Writing out $P(t)$-1 equations of the form

$$\alpha_p S_{1p} - \alpha_{p+1} S_{1,p+1} = S_{2p} - S_{2,p+1}$$

and introducing the linear constraint $\sum_{p=1} \alpha_p = 1$, we have P linear equations in P unknowns which can be used to solve for α. Hopefully all the components of the solution will be positive, otherwise it's impossible to maintain equal rates of change of RPUE. If this happens, fishing will switch back out of the slices with negative α_p. Under reasonable conditions (e.g. sufficiently high F), this shouldn't happen... I think.

[NB this neglects differential size-based mortality within an age class due to non-uniform selectivity; the so-inclined may wish to keep track of sizes as well as ages.]

How does one actually go about finding the times of switching to other slices and the effort distribution among exploited slices? There is no closed-form expression, because the time of the first switch (for example) depends on the end-point RPUEs in the previous year, and therefore on the whole past exploitation history of each slice within the management plan.

I guess one might try the following simulation. Start with zero population in all the white slices, and an equilibrium unexploited population in all the dark slices. Open up the darkest slice, and the fishermen will exploit it all season long. In the next year, the fishermen will start on the new slice but at some point the previous slice will catch it up, at which time we apply the matching-rate-of-change-of-RPUE rule. (Doing this will require weekly or more frequent time steps in the simulation). Continue round and round the circle until-- we hope-- things settle down. All one needs to keep track of is the age composition in each slice.

Once you have determined the pattern of fishing effort on each slice over time, you can proceed to working out yield-per-recruit. You need to track the lifetime yield from a cohort, and you need to do this for each of its $n+m$ possible birthplaces within the rotation cycle. Since you now know how much fishing pressure it will receive at each year of life, the calculation is simple. The overall YPR is just the average of all these possible birthplace-dependent YPRs. You can also calculate average spawning stock biomass, etc.

For more realism, it might be necessary to consider different bottom types and how these affect scallop density and growth. Density-dependent growth might also be significant. There might also be some important aspect of fishermen's behaviour that are neglected here. A further caveat is that I've neglected one aspect of non-uniform selectivity, which is to give differential exploitation on sizes of scallops within an age class (the above formulation does still allow for different selectivities at different ages, of course). In other words, the simple form of these calculations makes the assumption that 4-year-old scallops will have the same length distribution regardless of their exploitation history, though obviously there will be different total numbers.

As NEFSC actually has economists on hand, it would be interesting to incorporate some value-at-size factors rather than just using yield = weight * numbers (partly to determine when to switch slices, but more importantly to compare payoffs from different management strategies). I can foresee potential difficulties with substitution and elasticities among different size categories, that's what economists should sort out. It would be interesting to keep a good track of scallop supply & price by size class this season, now the ex-closed area is ex.