

Approaches to Weighting Data Components and Addressing Concerns about Density-Dependent Catchability for Lake Erie Percid Assessments.

February 23, 2006 DRAFT

White paper prepared by Jim Bence (Quantitative Fisheries Center, Michigan State University) in response to request by Andy Cook on behalf of Lake Erie Yellow Perch Task Group (YPTG).

Use this material with due care as it represents the view of the author without benefit of extensive discussion of the issue with the YPTG or peer-review

Abstract

I discuss the current approach to setting weighting factors (lambdas) for different types of information that contribute to the objective function when estimating parameters for Lake Erie percid stock assessments, and a suggested modification based on information indicating that catchability may be density dependent. This white paper was prepared in response to concerns that density dependent changes in catchability could invalidate the current approach for setting lambdas. I suspect changes in catchability are a real issue because the catchability and selectivity blocks are relatively long (all years since 1990 are in one block), but want to stress that density dependent catchability is only a big problem if the changes in catchability are not captured by changes in catchability among time blocks. My understanding of the current approach, based on previous examinations of code and recent discussions with Andy Cook, is that lambdas are set for “effort deviations” prior to the assessment and then relative values of the lambdas within other data types (harvest-at-age or survey indices-at-age) are determined by iteratively changing lambdas until the residual variability is consistent with that assumed by the lambdas. The “best” component within a type is given a lambda of 1.0, the same as the highest effort lambda. I argue that the approach I think is being used to set the relative lambda values for effort components is problematic, even if the underlying assessment assumptions about the relationship between fishing mortality and fishing effort are correct. The iterative approach is reasonable, but setting the best component of each type to 1.0 makes strong assumptions about the relative quality of effort, harvest, and survey information. I argue that a more involved iterative procedure that takes into account prior knowledge about data variability should be considered. I also argue that if there is evidence of density dependent catchability that is not captured by catchability time blocks, then down-weighting effort components is a reasonable response in the context of the current assessment model. However, the presence of such density dependence is not the only consideration for setting effort lambdas, and as an (perhaps longer term) alternative I recommend considering changes to the assessment model so that it can better allow for gradual changes in catchability. I provide some suggestions on how to estimate an overall beta for the power relationship between fishery catch rates and abundance indices, keeping in mind the use of catchability blocks in the assessments.

Discussion of the current approach

Estimation context and the scope of these comments

The primary topic addressed in these comments is how “lambda” values, weighting different components of the objective function used in Lake Erie percid assessment, should be determined. I specifically address questions of how the lambdas should be altered when there is evidence that fishery catchability is density dependent. As a basis for what follows, this subsection describes the estimation approach used in the percid assessments and why the current method used to specify lambda values is an issue.

The stock assessment approach used for Lake Erie percid follows a standard catch-at-age modeling approach adapted from the CAGEAN program developed by Deriso, Quinn, and colleagues. This approach estimates the assessment parameters by minimizing a sum of squares objective function using AD Model Builder software. For the purposes of statistical inference, this has been equated with a likelihood based approach, with assumed lognormal distributions for data. In this context two different variants of the objective functions have been used. The first follows a “concentrated likelihood” approach and the “negative concentrated likelihood” (with additive constants dropped) minimized by AD Model builder is:

$$- conc \log L = \frac{k}{2} \ln(RSS)$$

and the second is a full “negative log likelihood” (again with additive constants dropped):

$$- LogL = k \ln(\sigma) + \frac{1}{2\sigma^2} RSS$$

These two forms are based on identical assumptions and lead to identical point estimates and asymptotic standard errors for the parameters. In both cases RSS represents the weighted residual sum of squares over components making up the objective function:

$$RSS = \sum_{components} \lambda_i RSS_i$$

and k is the total number of squared terms summed over all components. In the full likelihood approach σ^2 or some transformation of it is estimated as a parameter, and represents the variance associated with components having a lambda (λ) of 1.0. The actual code contains algebraic variants of these because some additive constants are sometimes included. The lambda values are constants that are not estimated as parameters as the model is fit.

The individual RSS_i represent the residual sum of squares on a log-scale for each data component or the “effort deviations”. For example, the RSS_i for the commercial harvest data could be represented by:

$RSS_{com\ catch} = \sum_y \sum_a (\ln C_{ya}^{com} - \ln C_{ya}^{com})^2$ (see equation 3) and the RSS_i for the commercial

effort deviations could be represented by $RSS_{com\ E} = \sum_y \varepsilon_y^{com2}$ (equation 1). The lambda

values play a critical role in the objective function because they determine how closely the final model estimates needs to match data or how closely estimated fishing mortality rates need to track variations in observed fishing effort. Considered in the context of likelihood theory, the lambdas are statements about the relative magnitude of the variance associated with each component of the objective function (they are inversely proportional to the variances).

Technically speaking the assessment approach being used is not maximum likelihood estimation because some of the “parameters” (the effort deviations) are treated as random from an underlying distribution. Some recent literature describes the basic estimation approach used in these assessments as having a Bayesian underpinning and the resulting estimates as being “Highest Posterior Density” (HPD) estimates (Schnute 1994), although the approach is also described in the context of penalized likelihood estimation (Quinn and Deriso 1999). From the Bayesian perspective the sums of squares contributed by harvest-at-age and survey indices at age data are considered part of the likelihood, whereas the sums of squares associated with effort deviations are considered part of the priors for the effort deviations (associated with effort data and catchability) and other parameters are implicitly assumed to have “flat” priors possibly within specified bounds. If one were to move to a full Bayesian approach rather than just estimating point estimates it would be important to use the “full likelihood” rather than “concentrated likelihood” version of the objective function.

In any case, because the RSS associated with effort deviations could become zero (if all effort deviations were zero), their variance, and thus the lambdas for them relative to those for the data, cannot be estimated simply by minimizing the objective function (Schnute 1994, Quinn and Deriso 1999). In principle, with a lambda set for the effort deviations and one data component it could be possible to estimate other lambdas (this is same as estimating variances). This would not be possible if the model is so flexible that it could exactly match any of the other data components, and in practice it can be difficult to estimate these other lambdas by minimizing an objective function. NRC (1998) argued that research on how to specify weighting factors should be a priority, and Quinn and Deriso (1999, page 336) indicated that setting of appropriate weighting factors was an essential part of catch-at-age analysis. There is, however, no consensus on a single best way to set such weighting factors.

Current approach to a priori setting of effort lambdas

The current approach that seems to be used in Lake Erie yellow perch assessments (and which was used in past western Lake Erie walleye assessments I have seen) appears to set the relative lambda values for the different fishery effort values prior to any formal parameter estimation using the assessment model. At least for yellow perch this appears to be done by calculating the sample variances for the sets of

$\ln(E_y^G)$ and $\ln(C_y^G)$, calculating the ratios $\frac{Var(\ln(C_y^G))}{Var(\ln(E_y^G))}$, setting λ_G to 1.0 for the fishery

component (G) with the largest ratio, and setting the other effort lambda equal to its ratio

relative to the largest ratio. Here C_y^G and E_y^G are observed annual total harvest and fishery effort for fishery component G (recreational or commercial), and variances are based on the sample of values over years. Based on discussion with Andy Cook, the basis for this appears to be text on page 336 of Quinn and Deriso (1999) and supporting material in the original CAGEAN paper of Deriso et al. (1985). I had earlier communicated by e-mail with Kevin Kale about this topic with regard to the western basin walleye assessment and he indicated that the ratio approach was instituted based on input from Terry Quinn during workshops (June 24, 2004 email -- I am assuming from the late 1980s workshops).

In my opinion the approach described above for setting the effort lambdas is not consistent with the discussions by Quinn and Deriso (1999) and Deriso et al. (1985). In the ratio $\frac{\sigma_C^2}{\sigma_E^2}$ given by Quinn and Deriso, σ_E^2 describes how closely fishing mortality should track fishing effort, and σ_C^2 describes how closely observed harvest at age is to the actual harvest at age. Temporal variability in annual totals of observed harvest and effort cannot be expected to be proportional to these. Furthermore, the discussion by Quinn and Deriso is in reference to setting the lambda for effort relative to the lambda for harvest rather than for setting the effort lambdas for different fisheries.

I now support the above statements regarding the interpretation of the ratio of variances in detail. The relationship between fully selected fishing mortality (for simplicity I have dropped the fishery component subscript and modifications for catchability blocks) and observed fishing effort currently assumed in the Erie assessments is:

$$F_y = qE_y \exp(\varepsilon_y) \quad (1)$$

Where q is catchability, F_y is fishing mortality in year y and $\exp(\varepsilon_y)$ is a multiplicative error such that ε_y is $N(0, \sigma_E^2)$. These ε_y are the “effort deviations”. Equation 1 is the same as equation 8.65 in Quinn and Deriso (1999). It is worth noting that the above can be rewritten as:

$$\begin{aligned} F_y &= q_y E_y \\ q &\sim LN(\mu_q, \sigma_E^2) \end{aligned} \quad (2)$$

This makes it explicit that the error term can involve both measurement of effort and/or variation in how fishing mortality relates to actual fishing effort (σ_E^2 is the sum of the variances from these two sources). If we think of catchability describing the relationship between actual fishing mortality and observed fishing effort, this latter formulation says that $\ln(q)$ is assumed to vary following a normal distribution.

Now turning to harvest (again only working with one fishery component and dropping the associated subscript), the Erie assessment models assume that actual and observed (denoted by prime) harvest is given by

$$C_{ya} = N_{ya} [1 - \exp(-Z_{ya})] \frac{F_{ya}}{Z_{ya}} \quad (3)$$

$$C'_{ya} = C_{ya} \exp(\delta_{ya})$$

Where $\exp(\delta)$ is the multiplicative measurement error and δ is assumed to be $N(0, \sigma_C^2)$.

This is the same model described for catch by Quinn and Deriso (1999) pages 334-335, and they define σ_E^2 and σ_C^2 the same as I do here.

Thus it seems clear to me that when Quinn and Deriso (page 336) say “the term λ_E may also be thought of as the ratio of log-harvest observations to log-effort

observations $\frac{\sigma_C^2}{\sigma_E^2} \dots$ ” they are referring to the variances σ_C^2 and σ_E^2 in the above

equations, not the sample variances of log annual observed harvest and effort from a time-series. This is a very important distinction because different processes that can influence harvest and effort will influence these different definitions of variance in qualitatively different ways. For example, with all else equal using lambdas based on

$\frac{\sigma_C^2}{\sigma_E^2}$ as defined in equations 1-3 would lead to a lower value for a fishery component

when that fishery component has catchability that is varying from year to year (because σ_E^2 is larger), whereas basing the lambda on the ratio of variances in observed log total harvest and effort (the current Erie method) would generally lead to the lambda for the fishery with varying catchability to be higher because such variation would cause harvest to vary more for a given amount of variation in effort. In this case, which is not a pathological example, using the ratio of the temporal variability in total observed harvest and effort could produce an exactly opposite response to what should be done when catchability is fluctuating in a way consistent with the current assessment model (equation 1).

An additional issue is that even if the true ratios of $\frac{\sigma_C^2}{\sigma_E^2}$ as defined by equations 1-

3 for each fishery component were known, these provide information on the relative lambdas for the fishery effort and harvest for that component, not the relative lambdas among components for effort. For example, if σ_E^2 (recreational) < σ_E^2 (commercial), then recreational effort should get a higher lambda than commercial effort and information that commercial harvest is measured much more accurately than recreational harvest, so that σ_C^2 (commercial) << σ_C^2 (recreational), should not cause us to use a lower lambda for recreational effort.

I recognize that the current approach allows for an operational procedure for setting effort lambdas whereas the interpretation of $\frac{\sigma_C^2}{\sigma_E^2}$ I provide here does not (since I

have not described how these variances can be estimated from data). Nevertheless the critical point is that the current operational procedure could be leading you substantially

astray even if the basic assessment model's assumptions about how fishing mortality and effort relate to one another were correct.

I think another aspect of the text on page 336 of Quinn and Deriso (1999) may have not been entirely clear. It appears to be interpreted as indicating that down weighting from lambda of 1.0 is something you do when the catchability model of equation 1 has been violated (i.e. fishing mortality not directly proportional up to a random error to fishery effort). In my opinion the text below equations 8.65 and 8.66 on page 336 of Quinn and Deriso (1999) are simply describing the implications of equation

1 (equation 8.65) as $\frac{\sigma_C^2}{\sigma_E^2}$ (and hence λ_E) is varied between 0 and infinity. For values

close to zero, the effort variance is near infinity and fishing mortality can essentially take any value with no penalty. This is close to ignoring the effort data and allowing wild fluctuations in catchability (see my equation 3). As values approach infinity, the effort variance approaches zero and this is close to dropping the error term from equation 8.65, so that by my equation 3 catchability becomes strictly constant. The main point here is that the lambdas could and should take a range of possible values (including possible cases where lambda for the best effort component is greater than the lambda for the best harvest and survey components) even if equation 1 were an entirely accurate description of reality.

Current iterative approach to setting lambdas for harvest and survey data

The basic idea of iteratively adjusting lambda values to be consistent with the residual variation seems quite reasonable to me. The basic idea underlying this approach has been suggested as a way to obtain "effective sample sizes" in catch-at-age models when using a multinomial error structure for harvest-at-age proportions rather than the CAGEAN formulation (McCallister and Iannelli 1997) and is an approach similar to this has been used in 1836 treaty water assessments of lake whitefish and lake trout (e.g., Ebener et al. 2005). Francis et al. (2003) compared residual variability to that assumed a priori in assessments as a way to evaluate the amount of temporal variability in catchability and whether prior variance estimates used in assessments were too large or small. However, as implemented I have two main concerns regarding the iterative approach used for Lake Erie percids. First, all the iterative calculations are conditional on fixed effort lambdas, the basis for which is somewhat questionable (see above). Second, the iterative procedure fixes the lambda values for the best harvest and survey data to 1.0, which effectively assumes that the variances associated with these data are the same as σ_E^2 for the effort component given a weight of 1.0. A priori there is no reason to suppose this is true, and it is not clear to me that the YPTG (or WWG) intended to make this assumption. Note that here the variance for the harvest data is the same σ_C^2 described above, and the survey variances (σ_I^2) comes from the following assumed submodel for survey data:

$$\begin{aligned}
I_{ya} &= s_a^s q_s N_{ya} \\
I'_{ya} &= I_{ya} \exp(\xi_{ya}) \\
\xi &\sim N(0, \sigma_I^2)
\end{aligned}
\tag{4}$$

where I_{ya} is the survey index of abundance (cpue) for a given year and age, and the other terms include the survey age-specific selectivity, the survey catchability, abundance (for the year and age) and a measurement error. Given the above model, including the error structure, σ_I^2 describes the measurement error variance in the survey indices of abundance at age, and it is assumed that a very intensive survey would have a variance near zero, which would yield indices equal to I_{ya} . It is widely acknowledged that in reality other sources of variation act on surveys, in particular their catchability (q_s) could vary. However, explicitly addressing this would require a different error structure in the assessment (see “Considerations for Future Assessment Modifications”, point 5).

The idea of setting effort lambdas based on the coefficient in power model and other possible responses to changing catchability

The current assessment model assumes that fully selected fishing mortality for a given fishery component is directly proportional to fishing effort for that component up to a multiplicative error (equation 1) or as indicated by equation 2, that catchability varies randomly with independent (over time) variations about a common mean. An important point to note is that this assumption applies only within catchability blocks, and the assessment model already allows for changes in average catchability among the blocks. The extent to which the basic catchability assumption is false undermines the utility of using the effort data from a given fishery component in the current assessment model. It is not an unreasonable response to evidence that the underlying model (equation 1 modified for catchability blocks) is incorrect, to down-weight the associated effort data. In other applications I have seen that down-weighting data when associated processes are incorrect can improve assessment results (e.g., Radomski et al. 2005). An extreme version of this is to simply stop using fishery effort data whenever survey data are available, because the underlying fishing mortality – effort assumptions in an assessment model are always open to challenge (NRC 1998).

In the current circumstances I believe that down weighting effort data based on evidence that catchability is density dependent in ways that are not already captured by the catchability blocks is not the best approach (although it might be pragmatic in the short-term). In such circumstances Quinn and Deriso (page 336) suggest replacing the direct proportionality in equation 1 with a nonlinear relationship, and point to their section 1.3.4 where power models are discussed. Although the issue of nonlinear relationships are often discussed, actual implementation of such models inside of assessments is nearly absent from actual stock assessments and published literature. Fournier (1983) describe one attempt to modify equation 1, where q is replaced by q_y , and q_y is assumed to be linearly related to the deviation of the year y stock biomass from the average over the assessment period. In the end he concluded that allowing for density dependent catchability did not fit the data better for his application. More recently, my

former student Mike Wilberg's Ph.D. dissertation addressed the issue of conducting stock assessment in the face of slowly varying catchability, and evaluated alternatives through simulations (Wilberg 2005). Portions of this work are now being revised for publication in CJFAS based on positive reviewer comments (Wilberg and Bence MSS in revision). He assumed that catchability varied in a variety of ways in the simulations, and used a variety of estimation approaches. He found that replacing equation 1 with a power function of abundance (and no error term), as seems to be suggested by Quinn and Deriso, did not reliably produce assessments that converged to valid estimates, unless the true generating model was a power function, and not some other time varying catchability pattern. He also found that using a random walk model generally performed well among those he tried, including when catchability was a power function of abundance. For the random walk model, one simply modifies equation 1 in the assessment model to be:

$$F_y = qE_y \exp(\zeta_y) \quad (5)$$

$$\zeta_y = \zeta_{y-1} + \varepsilon_y$$

One possibility that Mike Wilberg did not look into in his simulations was to simply replace q in equation 1 by q_y , with q_y made a function of abundance (say a power function or a linear function of biomass as in the Fournier 1983 application), which retains random variation as well as invoking a nonlinear relationship for q . With respect to these simulations one important result (not reported in the dissertation or manuscript) was that using a random walk model outperformed simply downweighting the effort data. The simulations also showed that making use of fishery effort data could provide substantial benefits in terms of better stock assessment estimates, but if left as white noise (i.e., following equation 1) when catchability really changed gradually over time the assessment results could be quite poor.

In other work concerned with time-varying selectivity (Radomski et al. 2005) we discovered that moderate length (five year) selectivity blocks could capture important aspects of changing selectivity patterns. Similarly, as I suggested above, it is possible that the current use of catchability blocks is already capturing an important part of any changes in catchability either due to population abundance or other factors. However, given the relatively long blocks that appear to be used in at least the yellow perch assessments as suggested by the "yptg_effort_lambda.xls" template I saw, these time blocks seem unlikely to help much in this regard.

My basic recommendation here is to consider two new alternative approaches for dealing with potentially density dependent catchability or other gradual changes in catchability over time. The first would replace q in equation 1 by q_y , with q_y made a function of abundance. Obviously there are many choices in such functions, although one reasonable one would be to assume a power relationship between catchability and the sum of abundance at ages, with each age weighted by its corresponding selectivity (sometimes referred to as fishable abundance). The second would be to use a random walk model for the effort deviations (i.e., replace equation 1 by equation 5). Note that my recommendation here addresses the concern about modeling catchability when it is varying so that fishing mortality is not proportional to fishing effort as assumed in equation 1. It does not address the original and remaining issue of how one should weight (set lambdas) for the different components of the sum of squares in the assessment

model's objective function after doing this. Both these recommended approaches need to be thought about in light of your current catchability blocks. They can be thought of as alternatives to or additions to the catchability blocks. As alternatives you would drop out the current catchability blocks when using them. But if you wish to model both sharp changes between blocks and more gradual changes within blocks the approaches could be combined. My first thought is that if the current blocks are defined by known system changes that you think caused sharp changes then they should be retained.

Some thoughts on setting lambdas

As discussed above, theory tells us that we must set at least the lambdas for the fishery effort deviations and one other data source. After doing this it may be possible to iteratively adjust the other lambdas until they stop changing, although it also may be necessary to specify more about the relative values of some of the other lambdas. My recommendation is to try only specifying three lambdas (those for the effort deviations and say commercial fishery harvest-at-age) initially, and specify more a priori only if that proves necessary. This will require some exploration of what actually works as in my experience if too many lambdas are being adjusted the iterative procedure sometimes fails to converge. I think prior knowledge needs to be used to come to some judgment about the specified lambdas. Expert knowledge about the fishery should be considered and a judgment made on the likely relative differences in the variances for effort deviations based on how well effort is measured and how much fishery catchability is likely to fluctuate. I recommend initially fixing these (and possibly returning to them later (see below)). For the one other lambda that is fixed (I will assume for now the commercial harvest-at-age), after fixing its value, iteratively adjust the other "non-fixed" lambdas. I would then compare how well the resulting amount of residual variation in commercial harvest-at-age data corresponds to prior estimates of how well the harvest-at-age data are estimated. I would think those familiar with the harvest monitoring and aging procedures could provide rough estimates of sampling precision for harvest-at-age data in terms of CVs. These CVs, unless very large, approximate the logscale standard deviations. Thus your prior information suggests a logscale variance for the harvest data of $CV_{Catch\ com}^2$, which can be compared with $MSE_{Catch\ com} = RSS_{Catch\ com} / n_{Catch\ com}$. This suggests a procedure of fixing the commercial harvest lambda at a range of values, adjusting the lambdas for the other non-fixed lambdas (i.e., the other lambdas except efforts) using an iterative approach, and selecting the commercial harvest lambda that produces a best match between the prior information on variance and the resulting residual estimate. This procedure is similar to one used by Richards et al. (1997). I more or less arbitrarily picked commercial harvest as the one to fix the lambda for at a range of values, but the procedure described above indicates you should choose one for which you have good information on precision.

I suggested fixing relative effort lambdas based on expert opinion, but it may prove possible to also evaluate alternative choices about these lambdas based on residual variation. In particular, the ratio of lambdas for the effort deviations should be inversely proportional to the ratio of the corresponding variances, so the ratio of the MSEs after model fitting could be compared with what is expected based on the lambdas. Possibly this would allow you to repeat the iterative procedure while setting one of the effort

lambdas to several different values, and choose among the results based on the match between residual variation and prior expectations based on the lambdas.

At this point I am pretty deeply into the weeds and my point is not to outline exactly what to do, just the general flavor of an approach. It seems to me that getting further than general ideas may require some sort of working meeting or workshop.

Estimating and making use of a power relationship

I caution here that I am not convinced that estimation of the power relationship should be the only basis for setting the lambdas for effort data. As indicated above there may be other reasons to consider different effort lambdas. However as a first step it might be useful to use evidence of a stronger nonlinear relationship for one fishery component versus another to downweight the relative effort lambda for that fishery component versus the other or even versus other data. This might be a reasonable approach until alternative approaches for incorporating time varying catchability (beyond the current blocking approach) are incorporated into the assessment models.

In this regard, here are some ideas. First, some very careful thought is needed with regard to how beta is estimated, how it is to be used, and the current assessment model structure with respect to both catchability and selectivity blocks. Each distinct combination of selectivity and catchability block in effect defines a time period with a unique age-specific mean catchability in the current assessment models. It seems to me that you should be down-weighting the effort data only if you uncover evidence that there are changes in catchability that are not already captured by these blocks. This may not be a real concern as the example template I examined included survey data starting in 1990 and this represented a single time block. In any case, my first thought is that you should be allowing separate intercepts (but a common slope) in your regression of log fishery catch rate on log survey cpue, at least for each catchability block and age, and perhaps restrict the analysis only to those ages for which selectivity is not expected to change too much among selectivity blocks (separate intercepts for each combination of selectivity and catchability block probably requires estimating too many parameters). In effect my recommendation here is to estimate a common beta by doing an ANCOVA on log fishery catch rate, with log survey index being the covariate (assuming homogeneity of slopes) and the age-catchability block combinations represent groups with different intercepts (adjusted means). In a case where all the survey data come from one time block this just means doing the ANCOVA with age being the grouping variable. If the analysis is restricted to just a few adjacent ages this could probably be done as a standard ANCOVA. However, it may be the case that different ages used in the analysis have substantially different residual variances, and if so this should be allowed for.

Contrasting with the above suggestion, the existing template appears to either do the regression for all years with available data or for a selected catchability and/or selectivity block. (I am not entirely sure what happens when data from both surveys and the fishery are available for more than one block although it seems like the template is set up to do the analysis only for either all the data or a single selected catchability and/or selectivity block.) Two approaches mentioned for combining ages are weighting by R^2 and simply establishing an overall relationship based on “total abundance.” The rationale for weighting by R^2 was that this would in effect weight more highly selected ages more.

It seems to me looking at total catch rate versus total survey cpue in effect does the same thing since more highly selected ages by the fishery will dominate catch rates and more highly selected ages by the survey will dominate survey indices. I do not like the idea of weighting by R^2 . My concern is that relationships having a near zero slope will by definition have a small R^2 and be given little weight – if slope is zero then R^2 has to be zero. Thus a low R^2 might reflect noisy data but it could also reflect the case of serious concern where fishery cpue tells us virtually nothing about abundance because of how catchability is changing. It seems to me this would substantially bias the analysis.

Among the specific choices discussed in materials provided to me: (1) use the estimated beta for the best age (highest R^2), (2) use a beta for pooled ages 2 and older, (3) weight the age-specific results by R^2 , I prefer (2) mainly because of the concern I express above about weighting by R^2 . I think if “best age” is defined a priori (rather than by R^2), for example by being well represented in both fishery and survey, I would think that might be as good or better than using the pooled data if data. The fourth option, of calculating the absolute deviation from 1.0 and weighting this by R^2 , I do not care for. Ultimately it may be necessary to convert the results on density dependent catchability to a metric by which you devalue the effort data when assumptions are broken, and clearly both positive and negative deviations from direct proportionality should lead to some down-weighting. But I think the first task is to come up with a best estimate of beta. I have two specific objections to this fourth option. First, is the same problem with using R^2 in the weighting I describe above. The second is that deviations in beta below and above 1.0 are not equal. Clearly fishery cpue has no direct information on abundance if there is no relationship between fishery cpue and abundance (beta of zero), whereas if beta is 2.0 there is some relationship. Perhaps lambda should fall from 1.0 to zero beta goes from 1.0 to infinity. This could be accomplished by some sort of rescaling or transformation of beta.

Considerations for Future Assessment Modifications

My above notes and related thinking lead to the following thoughts regarding the Erie percid assessments. Some of this goes well beyond the current issue of setting lambdas.

1. In general I advocate a more flexible approach to the assessment model, when the evidence clearly indicates a change is needed. If it is indeed the case that catchability is varying gradually in ways not captured well by the assessment model, then the assessment model should be changed, perhaps by allowing for a random walk in the effort deviations or power model relationship between catchability and abundance. The attention to potential density dependence in catchability is a positive step in the Lake Erie assessment work and puts this assessment group ahead of many others.
2. Distinct from the issue of density dependent catchability, there still remains the issue of deriving best (or at least better) weighting factors. These should not be viewed as synonymous issues. The current approach to setting these weighting factors appears to make strong and possibly unintended assumptions regarding the variances. Even if it proves necessary to specify the lambdas for each set of effort deviations and one for harvest and one for a survey (as is now done), I think the

- stock assessment working groups could probably use their knowledge to establish default values that with better support than the current use of 1.0 for the “best” effort, harvest, and survey components. In the short term it might make sense to use evidence of density dependence as one ground for the relative weighting of effort data or down-weighting it relative to other components.
3. One longer term possibility is to move away from the highest posterior density estimate and to a more fully Bayesian approach. This would then require specification of prior distributions for variance parameters (rather than specific lambda values). I have a current student (Brian Linton) comparing in simulations a version of the iterative adjustment of weightings to the fully Bayesian approach. He also is exploring using a frequentist approach where random quantities like effort deviations are treated as random effects rather than parameters (using the new ADMB-RE package). The latter approach has not worked well in preliminary tests.
 4. After fixing the effort deviation and one other lambda, it might be possible to estimate the other lambdas as formal parameters, rather than using an iterative approach. This would require revision of the objective function being used now so that the variances for each component in the objective function appear explicitly as they should for estimated parameters.
 5. Consideration of lambdas opens up a broader issue of assumed distributions and whether the assumed lognormal distributions are appropriate. It has been argued that often harvest-at-age and survey index-at-age data might better follow an error structure where “errors” for different ages within a year are correlated, and variances on a log scale are not constant for all ages and years. Correlations among ages within a year might arise for a survey if all or a range of ages were similarly influenced by changes in catchability. One alternative is to use a lognormal for total harvest or survey cpue and a multinomial for proportions at age (Fournier and Archibald 1982, Crone and Sampson 1998), and another is to use a multivariate lognormal distribution (Myers and Cadigan 1995).

Acknowledgments

This represents a white paper prepared as part of MSU Quantitative Fisheries Center operations. It benefited from input by Mike Wilberg.

Sources Cited

- Crone, P.R. and Sampson, D.B. 1998. Evaluation of assumed error structure in stock assessment models that use sample estimates of age composition. In: Fishery Stock Assessment Models, Funk, F., Quinn, T.J. II, Heifetz, J., Ianelli, J.N., Powers, J.E., Schweigert, J.F., Sullivan, P.J., and Shang, C.-I. Alaska Sea Grant College Program Report No. AK-SG-98-01, University of Alaska, Fairbanks.
- Deriso, R.B., Quinn, T.J., II, and Neal, P.R. 1985. Catch-age analysis with auxiliary information. *Can. J. Fish. Aquat. Sci.* 42: 815–824.
- Ebener, M. P., J. R. Bence, K. Newman, and P. Schneeberger. 2005. Application of statistical catch-at-age models to assess lake whitefish stocks in the 1836 treaty-ceded

- waters of the upper Great Lakes. Pages 271–309 in L. C. Mohr and T. F. Nalepa, editors. Proceedings of a workshop on the dynamics of lake whitefish (*Coregonus clupeaformis*) and the amphipod *Diporeia* spp. in the Great Lakes. Great Lakes Fishery Commission, Technical Report 66, Ann Arbor, Michigan.
- Fournier, D., and Archibald, C.P. 1982. A general theory for analyzing catch at age data. *Can. J. Fish. Aquat. Sci.* 39: 1195–1207.
- Fournier, D.A. 1983. An analysis of the Hectate Strait Pacific cod fishery using an age-structured model incorporating density-dependent effects. *Can. J. Fish. Aq. Sci.* 40: 1233-1243.
- Francis, R.I.C., Hurst, R.J., and Renwick, J.A. 2003. Quantifying annual variation in catchability for commercial and research fishing. *Fishery Bulletin* 101: 293-304.
- McAllister, M. K., and J. N. Ianelli. 1997. Bayesian stock assessment using catch-at-age data and the sampling-importance resampling algorithm. *Canadian Journal of Fisheries and Aquatic Sciences* 54: 284–300.
- Myers, R.A. and Cadigan, N.G. 1995. Statistical analysis of catch-at-age data with correlated errors. *Can. J. Fish. Aq. Sci.* 52: 1265-1273.
- NRC (National Research Council). 1998. Improving fish stock assessments. National Academy Press, Washington, D.C.
- Quinn, T.J., and Deriso, R.B. 1999. Quantitative fish dynamics. Oxford University Press, New York, NY.
- Radomski, P., J.R. Bence, and T.J. Quinn II. 2005. Comparison of virtual population analysis and statistical kill-at-age analysis for a recreational, kill-dominated fishery. *Canadian Journal of Fisheries and Aquatic Sciences* 62: 436-452.
- Richards, L.J., Schnute, J.T., and N. Olsen. 1997. Visualizing catch-age analysis: a case study. *Can. J. Fish. Aq. Sci.* 54: 1646-1658.
- Schnute., J.T. 1994. A general framework for developing sequential fisheries models. *Can. J. Fish. Aquat. Sci.* 51: 1676-1688.
- Wilberg, M.J. 2005. Improving statistical catch-at-age stock assessments. Ph.D. Dissertation. Michigan State University, East Lansing. 132 pages.
- Wilberg, M.J., and J.R. Bence. Submitted MSS. Performance of time-varying catchability estimators in statistical catch-at-age analysis. *Canadian Journal of Fisheries and Aquatic Sciences*.