RAPID COMMUNICATION / COMMUNICATION RAPIDE

Folly and fantasy in the analysis of spatial catch rate data

Carl Walters

Abstract: Spatial catch per effort data can provide useful indices of population trends provided that they are averaged so as to correct for effects of changes in the distribution of fishing activity. Simple, nonspatial ratio estimates should not be used in such analyses. The averaging for any time period must necessarily make some assumptions about what catch rates would have been in spatial strata that had not yet, or were no longer, being fished. Ignoring the unfished strata (averaging only over the areas that were fished) amounts to assuming that they behaved the same as the fished strata and can lead to severe hyperdepletion in abundance indices for fisheries that developed progressively over large regions.

Résumé : Les données spatiales de captures par unité d'effort fournissent des indices utiles des tendances démographiques à la condition qu'elles soient transformées en moyennes de manière à tenir compte des effets des changements dans la répartition des activités de pêche. Des estimations simples et non spatiales de rapports ne doivent pas être utilisés dans de telles analyses. Le calcul des moyennes pour chaque période de temps doit nécessairement établir des présuppositions sur ce qu'auraient été les taux de capture dans les strates spatiales dans lesquelles on n'a pas encore pêché ou dans lesquelles la pêche a cessé. Si on ne tient pas compte des strates où il ne se fait pas de pêche (en calculant les moyennes seulement à partir des zones de pêche), on assume que toutes les strates réagissent de la même façon, qu'il y ait ou non de la pêche; cela peut mener à une sévère hyperdéplétion dans les indices d'abondance pour les pêches commerciales qui ont évolué progressivement sur de grands espaces.

[Traduit par la Rédaction]

Introduction

Information on trends in relative abundance is critical for most long-term stock assessments, and we typically must use commercial catch rate (catch per unit effort, CPUE) data to estimate such trends. In many cases, detailed logbook records allow spatial analysis of catch rate data, typically by spatial grid cells viewed as sampling strata, so as to hopefully remove the obvious biases that would be expected from nonrandom search behavior by fishers.

Unfortunately, two very serious mistakes have commonly been made in the analysis of spatial catch rate data. The first involves the incorrect use of ratio (catch/effort) estimators, and the second involves failure to recognize hidden assumptions about abundance trends in spatial cells that were not fished at all early or late in fishery development. Both of these mistakes are liable to produce exaggerated trend indices, i.e., "hyperdepletion" (Hilborn and Walters 1992), the appearance that stock size has declined much more than it actually has. The second of these is by far the most serious. This paper reviews both and offers suggestions for dealing with the second.

Folly: incorrect use of ratio estimators

Because catch per effort is obviously a ratio of catch to effort, it has been usual to use ratio estimators for it, i.e., to sum up catches over some units of time and space and then divide the sum by total effort over these units. This calculation should be done only for small enough units (spatial cells, time periods) to have had effectively random sampling within each unit. However, in some cases, e.g., Myers and Worm (2003), catches and efforts have been summed over spatial cells (strata) to produce a single, overall ratio estimate of catch per effort. This method in effect places much more "weight" on data from those cells that were heavily fished than on cells for which catches and efforts were low, whether or not the heavily fished cells were representative in any way of relative abundance in lightly fished or unfished cells. Texts on statistical sampling theory (e.g., Cochran 1963, pp. 167–172) present better estimators for ratios in stratified random sampling, and those estimators are of the form $\bar{c} = \Sigma_h w_h r_h$, where h is the cell or sampling stratum, w_h is the physical stratum size divided by the physical sampling universe size, and r_h is the ratio estimator for cell h. That is, stratified sampling estimators for ratios consist of weighted averages of ratios for individual cells or strata, where the

Received 28 October 2003. Accepted 26 November 2003. Published on the NRC Research Press Web site at http://cjfas.nrc.ca on 8 December 2003.
J17795

C. Walters. Fisheries Centre, The University of British Columbia, Vancouver, BC V6T1Z4, Canada (e-mail: c.walters@fisheries.ubc.ca).

doi: 10.1139/F03-152

weights do not depend in any way on the number of observations made in the individual strata.

In fisheries where effort was initially concentrated in a few cells then later spread widely, as in the Japanese pelagic longline fishery examined by Myers and Worm (2003; see maps at http://fish.dal.ca/~myers/Myers_Supplementary2.pdf), catch rates have typically dropped dramatically in the initial fishing areas before fishers were willing to spread their efforts more widely (for obvious risk-management reasons). The catch-rate declines in such areas are not in any way representative of overall (over all cells) stock change, and to pretend that the overall ratio estimator of catch rate was in fact representative of overall stock change amounts to the folly of pretending that the whole system was declining at the same rate as it did in the initial fishing area.

Fantasy: ignoring unfished cells/strata

The ratio estimation mistake above is simply a failure to apply well-known statistical methods for analysis of stratified sampling data. A much more subtle but ultimately more important mistake has been to apply the correct formulae for stratified sampling (separate ratio estimator for each cell, take average of these weighted by cell sizes), but to do so only for those cells that were actually fished in a given time period. After all, how can we justify trying to make any claim about what was going on in those cells that were not even visited or sampled by fishers?

The construction of an overall trend index for change in a spatially structured stock is the same as constructing a very large table, with a row for each time period (typically year or month) and a column for each spatial cell or stratum that might have been occupied at each time (i.e., each cell for which there was ever any fishing, whether or not the cell was fished in all time periods). For standard theorems from stratified sampling to apply to the calculation of the index for each time period (as an average across a row of the table), we would need to have catch-rate data for every cell, in every period. This almost never happens, especially early in fishery development. Faced with a table that has many missing entries, we have three choices: (*i*) we can refuse to construct any overall stock trend index and instead restrict statistical analysis and inference to only those cells that were fished each year (in which case the index is meaningless as an input for overall stock assessment); (*ii*) we can assume that the mean catch rate over those cells that were fished in any period is somehow representative of what would have been seen in the cells that were not fished, i.e., that cells were chosen for fishing at random or that it would be acceptable to fill all empty table cells in each row with the mean of the cells that were fished in that row; or (*iii*) we can somehow fill the empty table cells with our best estimates of what would have been seen in those cells had they been fished, which almost certainly would not be the mean of the cells that were fished in each period.

Notice that there is an element of fantasy in either of approaches (*ii*) or (*iii*). For fisheries that have moved progressively across large grounds, like pelagic long-liners crossing the Pacific during the 1950s (see Myers and Worm 2003), it would be unwise to use (*ii*) and assert that catch rates obtained in the western Pacific as the fishery first started to spread were representative of what stock abundances were doing in the (unfished) eastern Pacific at that time (as witness high catch rates obtained later, when the fishery did reach further afield). A far more credible estimate of the catch rate that would have been obtained in such unfished cells would be the catch rates that were obtained later when the fishery actually reached those cells, but even these “best estimates” require the fantasy that stock size in each cell was stable over time (at the level first seen by fishing) before fishing arrived at it.

So to analyze the long-term data without making some bad assumption (e.g., that unfished areas were behaving the same as fished ones), we are forced to think very carefully about how best to “fill in” the gaps in space–time tables of catch rates. Just as we cannot analyze even crude catch and effort data without making some dangerous assumptions (e.g., that local catch rate is in fact proportional to local abundance), the table filling has to involve some strong judgmental decisions about what assumptions to use. The important point is that once we recognize that the table has to be filled in somehow, we can at least elect to make reasonable assumptions rather than unwise ones based on the mistaken notion that statistical rigor or scientific objectivity (the main bases for averaging only over the fished cells) can somehow substitute for judgment.

When we deliberately and consciously ask the question “what is our best estimate of the catch rate for any particular space–time table entry for which we do not have direct observations”, several alternatives become apparent (besides the unwise one of just using the average over row elements that were actually fished). First, we can look forward and backward in time (up and down the table column) for catch rate information that may be representative of what would have been observed, essentially performing some sort of trend analysis and interpolation (extrapolation) from the time series that is available for the cell. Second, if there are data for the same time period from spatially “nearby” cells and if we can show that there is spatial correlation structure in catch rates, we can use spatial statistics methods to interpolate (extrapolate) across the spatial grid. Third, we can use data on spatial covariates of distribution (e.g., water temperature) to improve both the spatial and temporal trend estimates, or to at least provide estimates of changing range boundaries (which cells likely had zero abundance). Fourth, we can use estimators based on spatial models of dispersal–migration and local renewal dynamics, essentially fitting the time–space data to a much larger model than would typically be used for overall stock assessment.

Figure 1 illustrates just how big a difference the catch rate averaging method can make, using the Japanese long-line data for albacore (*Thunnus alalunga*) and yellowfin (*Thunnus albacares*) tuna in the Pacific. Three time trend patterns are shown: (1) the ratio estimator (total catch)/(total effort) reported by Myers and Worm (2003); (2) spatial average catch rate, only over those cells that were actually fished each year; and (3) catch rate averaged over all cells that were ever fished, with two filling conventions. First, the catch rate before first fishing in each cell was set to the average of the catch rates observed for the cell in the first 3 years after fishing reached it (i.e., abundance stable before fishing). Second, for cells that were “abandoned” by the fishery, the

Fig. 1. Catch per unit effort (CPUE) trends for large tuna and billfish (total number of fish per hook) from the 5 × 5 degree cell Japanese long-line database (Myers and Worm 2003), estimated by three alternative methods for (*a*) Atlantic, (*b*) Pacific, and (*c*) Indian oceans. Full spatial (solid line) assigns mean of first three observed catch rates to each cell for years before it was first fished and the last observed catch rate for years after it was last fished. Restricted spatial (▲) is the mean catch rate over only those cells that were actually fished each year. Ratio (■) is simply total catch summed over all cells divided by total effort.

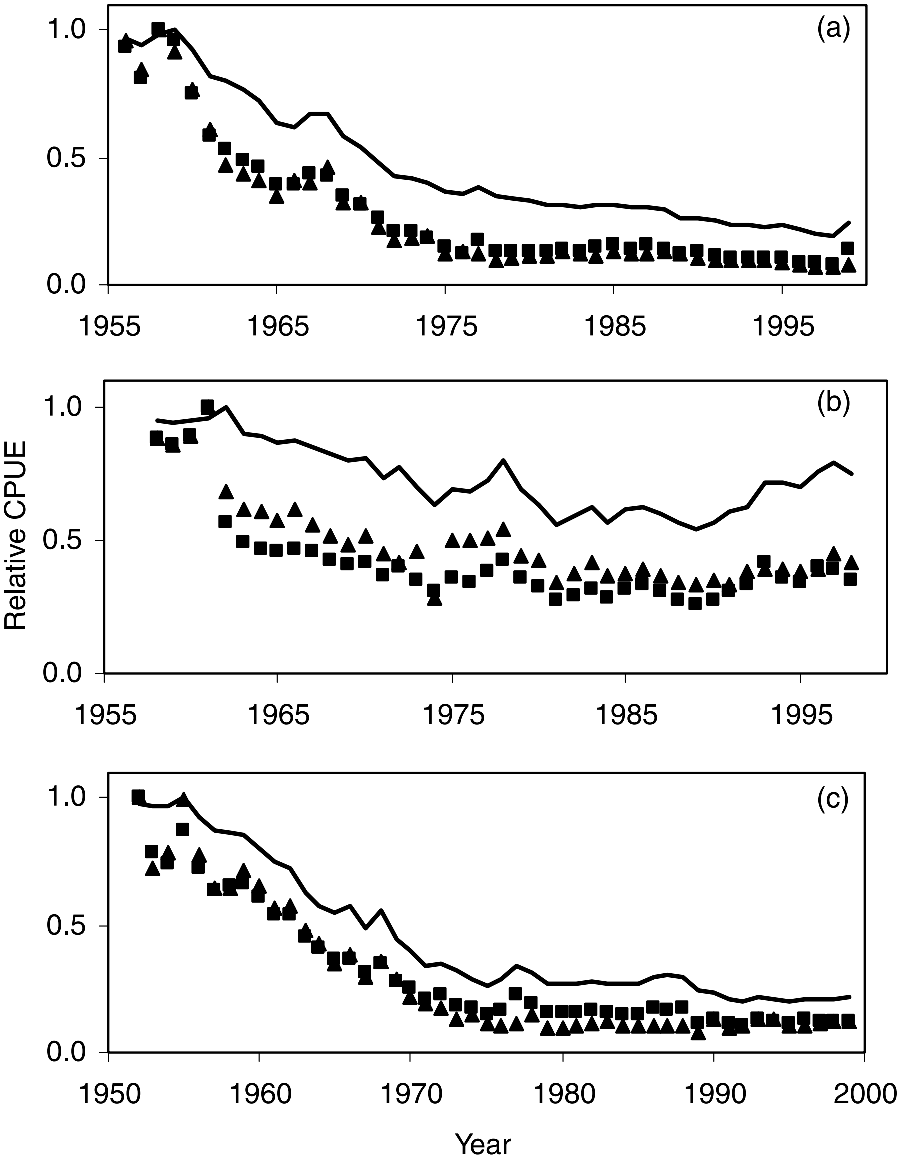

catch rate in later years was set to the rate for the last year of fishing (i.e., abundance assumed not to recover after fishing departed). The difference between these methods is very obvious; although all three indicate stock depletion, the third implies that the stocks were not as severely reduced as the fishery developed.

It has been typical in the analysis of commercial catch rate data to be very concerned about "hyperstability" (Hilborn and Walters 1992) caused by shoaling behavior and range contraction during stock declines. In terms of the table representation of trend index construction, hyperstability is represented by having a shrinking number of cells (if cells are small enough) occupied by fishers over time. In this case, to use the average catch rate in just the fished cells as an overall index (i.e., to fill the unfished cells with that average catch rate) would obviously create a hyperstable index. To avoid this hyperstability, a simple estimation tactic is to fill the unfished cells outside the fished area for each year with zeros (to represent the notion that fishers would still have been going to those cells if there had still been fish in them).

Discussion

It is apparently not often realized that constructing a long-term abundance trend index from catch-rate data is necessarily equivalent to filling in a large data table with many missing entries. And it may also be surprising to realize that using only the data at hand can be equivalent to (produces same numerical result as) filling in the table with obviously incorrect averages. But for far too long, we have been "putting our heads in the sand", pretending that things will average out in the end and that routine application of statistical techniques will somehow absolve us of responsibility for thinking carefully about what we are assuming in analysis. That in fact is a much more dangerous fantasy than to use various gap-filling methods, and it has to come to an end; examples like Fig. 1, where anyone who has examined overall catch trends would suspect that the stock declines could not have been as bad as indicated by ratio catch rate trends, can seriously undermine the credibility of fisheries assessment scientists. This will in turn make it harder to argue for conservation measures in cases where there really has been severe overfishing. Further, to simply ignore the long-term data entirely would, in most cases, be an invitation to the potentially even worse errors of the shifting baseline syndrome (Pauly 1995).

Once we realize the importance of gap-filling spatial catch rate tables to avoid both hyperstability and hyperdepletion in indices, we open the door to various innovative and relatively inexpensive approaches to design of field survey programs that can make the gap-filling more objective. For example, in cases in which behavior (shoaling, etc.) would lead us to suspect hyperstability, even a very simple, wide-area system for presence–absence sampling may be all that is needed to fill the spatial data gaps. Likewise, in many cases there may be simple, local information to warn us of whether back-filling of tables with prefishing catch rates for spreading fisheries would be misleading because of changes in stock distribution and (or) impacts of fishing on dispersal rates of fish to (from) areas that have not yet been heavily fished.

Even with care in the spatial data analysis, there always remains a serious risk that catch per effort is not in fact proportional to abundance, even at the scale of small spatial cells. Hyperstability at small scales can arise as the result of gear saturation and handling time effects. Hyperdepletion can arise as the result of localized depletion of fine-scale spatial aggregations within cells. Apparent hyperdepletion can be caused by interference competition (Gillis and Peterman 1998; Swain and Wade 2003), exploitation of a pool of more vulnerable fish that exchanges with a pool of less vulnerable individuals (Walters and Bonfil 1999), and perhaps other mechanisms than can cause direct decrease in catch per effort with increasing effort, even when stock size is not being depleted.

References

Cochran, W.G. 1963. Sampling techniques. John Wiley and Sons, New York.

Gillis, D.M., and Peterman, R.M. 1998. Implications of interference among fishing vessels and the ideal free distribution to the interpretation of CPUE. Can. J. Fish. Aquat. Sci. **55**: 37–46.

Hilborn, R., and Walters, C. 1992. Quantitative fisheries stock assessment: choice, dynamics, and uncertainty. Chapman and Hall Inc., New York.

Myers, R.A., and Worm, B. 2003. Rapid world-wide depletion of predatory fish communities. Nature (Lond.), **423**: 280–283.

Pauly, D. 1995. Anecdotes and the shifting baseline syndrome of fisheries. Trends Ecol. Evol. **10**: 430.

Swain, D.P., and Wade, E.J. 2003. Spatial distribution of catch and effort in a fishery for snow crab (*Chionoecetes opilio*): tests of predictions of the ideal free distribution. Can. J. Fish. Aquat. Sci. **60**: 897–909.

Walters, C., and Bonfil, R. 1999. Multispecies spatial assessment models for the British Columbia groundfish trawl fishery. Can. J. Fish. Aquat. Sci. **56**: 601–628.