COMBINATION OF RECRUIT INDIXES BY WEIGHTED AVERAGES USING RCRTINX2: USER'S GUIDE

J G Shepherd,
Ministry of Agriculture, Fisheries & Food,
Lowestoft
Suffolk NR33 OHT

A) Method and Data file

The method is explained by Shepherd (1987). Data must be supplied on a file, whose format is illustrated in Table 1. The title may be up to 80 characters long. The number of "surveys" and number of yearclasses must be specified, together with the number of the column in which the VPA estimates for each yearclass are specified (usually column number 2 as shown is the most convenient). The date for each yearclass must be inserted in column 1. Data may be in free format. Missing values are indicated by -11. Data should usually be provided in their usual form (the program can accept logged rather than arithmetic quantities, but this is not recommended). The program accepts an identifying code (of up to six characters) for each "survey". These must be appended, one per line, left-justified, at the end of the data.

B) Questions and answers

The program is executed by typing RCRTINX2, and answering the questions. Upper or lower case characters may be used, and only one character is necessary for yes/no responses (more will be ignored). All responses must be terminated by a carriage return (or Enter, or whatever). The responses to the questions are mostly obvious, but the following additional notes may be useful.

1. Data file name: any combination of up to 14 characters (no special extension type is assumed). The full filename is required (including drive:/ path if not in default path).
2. Regression type: C (for calibration) is recommended (following the recommendations of the Methods WG). P (for predictive) and F (for functional) are also legal responses but not recommended.

3. Tapered time weighting: this permits the progressive down-weighting of old data, which is usually a good idea. The shape of the taper may be specified. Typing 0 will give uniform weight up to a sharp (knife-edge) cut-off, and may be used to exclude old data but weight all recent data equally (there is thus no need to delete data from the file to achieve this). 1 gives a linear taper, which seems sensible to me. 3 gives the Cleveland tri-cubic taper which is popular in some circles. 2 is included for completeness (these tapers are all members of the same family) and gives a bisquare weighting which will make users of robust regression feel at home. In general the choice of taper among options 1, 2 or 3 doesn't seem to make much difference.

   Answering 'no' to this question will cause all data on the file to be used and weighted equally: this, and use of uniform (option 0) weighting, can lead to substantially different results.

   If one chooses to use a taper, the range (in years) must be specified. This need not correspond to the actual length of the data series, but should correspond to the time over which data is thought to become useless. Data up to half the taper length old gets substantially more weight than that in the older half - so choosing 20 years as suggested means that one is concentrating on the last 10 years data, which seems reasonable.

4. Prior Weights This permits the results of any particular "survey" to be downweighted when the results are combined. This is not usually necessary or desirable, but may be used to exclude the results of a particular survey completely (by choosing zero as the weight), or to allow for indices which are suspect or very provisional. Note that this applies only to the predictions made using an index: it does not affect the weight they are given as data in the regressions, nor (therefore) the actual estimates from any individual survey.

   Do not use this facility unless you must, because it provides considerable scope for confusion and subjective meddling. If you do select it, you are requested to supply the prior weights for each survey (in the order they occur in the data). These weights should normally lie in the range 0 to 1.

5. Shrinkage towards the Mean This enables one to take account of the historic distribution of yearclass strengths, and inject an element of caution about extreme predictions (see report on Methods WG 1987). If one selects this option (as recommended) the mean recruitment is included as an extra estimate, weighted in accordance with the standard deviation about that mean. If there are any good "surveys" this will have very little effect. If all the indications are imprecise, however, it will lead to more conservative estimates. In particular, any surveys giving estimates whose standard error is larger then the s.d. about the mean will get very little weight, which seems sensible!

   If this option is rejected, a pure calibration estimate is used (which is perfectly OK given any decent data).

7. Minimum Number of Years for Regression The absolute minimum is 3 (because the method needs to estimate standard errors about the regressions), but a larger number (4 or 5, say) may be chosen if one wishes to be cautious about new datasets entering the analysis. The high standard error of predictions from regressions based on few points is already taken into account, so it is not essential to use this option (see also comments relating to question 9).

8. Exclusion of surveys with SEs greater than that of mean This is not really necessary or effective if one has selected shrinkage towards the mean. Otherwise it does permit one to exclude completely datasets which appear to be useless. In the long run it would be sensible to delete such datasets from the file, although this should not be done casually, because some indication of high or low yearclasses may be useful and useable, even if imprecise. So shrinkage to the mean is a better choice, excluding only really bad datasets.

9. Minimum log S.E. This permits one to insert an estimate of the minimum log S.E. (CV as a fraction) one would be prepared to believe of any "survey" (because of known sampling variability for example). This "under-pinning" of the estimated S.E.s (since it is used to replace any lower estimates when weightings are calculated) prevents (for example) a very high weight being given to a "survey" which just happens to have an excellent correlation with only a few points. The suggested value of 0.2 seems to be a reasonable choice in practise (and better estimates are hard to come by!) The same value is applied to all "surveys".

10. First Yearclass for Prediction The program automatically carries out a retrospective analysis of as much of the data as you wish (which is very instructive the first few times you examine a particular set of data). Once you know where it's at, you will probably wish to examine the predictions only for the most recent yearclasses.

11. Surveys already logged This permits you to inform the program if you have included data of which logarithms have already been taken by mistake (or even deliberately).

C) Results

The printed results include a header which includes a complete listing of the options selected.

Then, for each yearclass in turn, the results of predictions based on each "survey" are given. The (log) index value, slope and intercept of the log-log regression, value of r^2, number of points used and the predicted value of the log yearclass strength are listed. In addition the residual standard error about the regression, the standard error of the prediction, and the weight attached to the estimate in the final weighted mean are given. In these results, a slope near unity, a high value of r^2, and a low value of standard error are indicative of a good prediction (and vice versa). The standard error of the prediction (again in logarithms) is taken as the paramount criterion.

Finally, the weighted average predictions are listed for each yearclass (in both log and re-exponentiated units), together with their internal and external standard errors, and their ratio, plus the log and re-exponentiated VPA estimates, for retrospective comparison.

The internal standard error is what would be expected from the combination of the various estimates whose standard errors are known. The external value is the actual value estimated from the (weighted) scatter of these estimates about their (weighted) mean. If all is well these should be comparable, so values of the ratio less than or around one indicate consistency. However, if the external value is much larger than the internal, this is evidence of statistically significant discrepancies among the estimates. A variance ratio test may be applied to the square of the external/internal ratio, if desired (using the number of "surveys" minus one as the degree of freedom in both cases). The larger of the two values should be taken as the final estimate of the standard error of the final estimate, and thus confidence limits may be constructed. Note that all the standard errors are still log quantities.

NORTH SEA COD AS 1-GROUP : POST 1976 , 0 & 1 GROUP DATA

7 11 2 (No. of surveys, No. of yearclasses, VPA column no.)

1976	709	-11	-11	36.7	6818	-11	28.2	-11
1977	426	1559	-11	12.9	2372	-11	27.2	-11
1978	454	1679	-11	9.9	2265	-11	31.1	-11
1979	800	1856	-11	16.8	5150	163.8	35.5	-11
1980	271	1006	43.2	2.9	1232	46.9	14.1	-11
1981	556	7963	176.8	9.2	3234	83	23.2	53.8
1982	265	254	26.9	3.9	1541	21.8	9	23.3
1983	552	9595	121.5	15.2	6122	121.3	43	63.1
1984	93	45	1.3	.9	419	3.6	.9	6
1985	-11	798	143.6	17.2	3464	111.2	-11	86
1986	-11	96	37	9.6	-11	-11	-11	-11

EGFS0
DGFS0
IYFS1
EGFS1
DGFS1
FRGSH1
SGFS