

The use of recruitment time-series structure and environmental information in medium-term stock projections

C.L. Needle, C.M. O'Brien, C.D. Darby, and M.T. Smith

Abstract

Foremost amongst the competing aims for fisheries management is the maintenance of fish stocks for the foreseeable future. To this end, it is imperative that managers are equipped with indicators of the expected level and variability of future population levels over the medium-term (typically, a five- to ten-year) time-scale, and that these indicators are sufficiently reliable, realistic and pertinent to the regulatory framework in which the managers operate. We extend the simulation approach currently employed within ICES stock assessment working groups by including characterisations of the time-series structure of residuals to fitted stock-recruitment models. The effect of the imposition of different hypothesised future environmental regimes is also investigated. North Sea cod (*Gadus morhua* L.) is presented as a germane and timely case study. We use the probability of the spawning stock biomass falling below the precautionary level of biomass, B_{pa} , as a diagnostic statistic to monitor projected performance. The relevance and utility of these new models for fisheries management is discussed, together with potential implications.

C.L. Needle: FRS Marine Laboratory, P.O. Box 101, Victoria Road, Aberdeen AB11 9DB, Scotland [tel: +44 (0) 1224 295456, fax: +44 (0) 1224 295511, e-mail: needlec@marlab.ac.uk]. C.M. O'Brien: CEFAS Lowestoft Laboratory, Pakefield Road, Lowestoft, Suffolk NR33 0HT, UK [tel: +44 (0) 1502 524256, fax: +44 (0) 1502 513865, e-mail: c.m.obrien@cefas.co.uk]. C.D. Darby: CEFAS Lowestoft Laboratory, Pakefield Road, Lowestoft, Suffolk NR33 0HT, UK [tel: +44 (0) 1502 524329, fax: +44 (0) 1502 513865, e-mail: c.d.darby@cefas.co.uk]. M.T. Smith: CEFAS Lowestoft Laboratory, Pakefield Road, Lowestoft, Suffolk NR33 0HT, UK [tel: +44 (0) 1502 524242, fax: +44 (0) 1502 513865, e-mail: m.t.smith@cefas.co.uk].

© 2000 Crown Copyright.

1. Introduction

Many strategic management decisions for fisheries within ICES' remit are currently made on the basis of medium-term projections of age-structured population dynamics. Such projections generally take the form of stochastic simulations of stock performance over a future 10-year time period, and are used to determine appropriate levels of fishing mortality that should minimise the risk of stock collapse. The exploitation pressure on most commercially-important species within EU waters is such that few fish survive for longer than five years, and projections are therefore largely driven by the assumed recruitment model and the initial position of the stock on the stock-recruitment curve. The current practice is to fit a parametric curve to the stock-recruitment scatterplot arising from historical assessment, take random samples from the time-series of residuals to this curve, and use these samples to determine recruitment values

in projections. The purpose of this paper is to highlight the principal statistical deficiency in this approach, correct it and compare the new results with the old for the North Sea cod population. The feasibility of incorporating an index of environmental effects into the projection process is also investigated.

2. Medium-term projection methodologies

2.1. The bootstrapped-residual approach

The most widely used method for simulating medium-term projections in the ICES assessment framework is the WGMTERM software and its descendants (Reeves and Cook, 1994). This is a relatively simple program designed to explore the likely response of an assessed population to fixed rates of fishing and natural mortality, given initial population-at-age estimates and an assumed parametric stock-recruitment model. Each WGMTERM analysis consists of a number of simulation runs (up to 1000), each projecting the development of the stock over a number of consecutive years (typically, up to 10). The vector of initial population sizes at the start of each run can be drawn from the distribution of possible sizes determined by the standard deviations of historical population estimates, although this option was not used in the following analysis as the aim was to compare the effects of different recruitment models, without the complication of added variability. The underlying model governing stock size is a standard Baranov (1918) age-structured model, in which exponential population decay is determined by the fixed rates of fishing mortality $F_{i,Y}$ and natural mortality $M_{i,Y}$. Thus the population at age i in year Y is given by

$$N_{i,Y} = N_{i-1,Y-1} e^{-Z_{i-1,Y-1}},$$

where $Z_{i,Y} = F_{i,Y} + M_{i,Y}$ measures the total mortality due to fishing and natural causes. A deterministic value of recruitment for each year in each simulation run is given by the fitted stock-recruitment model. To this is added a residual drawn at random (or *bootstrapped*) from the time-series of residuals to the fitted stock-recruitment curve. The principal output metrics are percentiles of the projected spawning-stock biomass (SSB, denoted by S_Y), where

$$S_Y = \sum_{i=1}^A N_{i,Y} W_{i,Y} Mat_{i,Y}.$$

Here $W_{i,Y}$ are weights-at-age, $Mat_{i,Y}$ are proportions mature-at-age, and A is the maximum age allowed for in the population.

The problem with this standard approach is that there is likely to be some degree of structure in the time-series of recruitment residuals, be it autoregressive or moving-average or both, and simple random sampling of residuals does not account for this. This paper describes new modifications to medium-term methodology to accommodate this issue.

2.2. Time-series data

Data on spawning-stock biomass S and recruitment R for North Sea cod were taken from the relevant stock summary in ICES (2000), ensuring that each value of R was paired with the correct value of S (so that $S_{i-a} \mapsto R_i$ in year i , where a is age at recruitment). Stock-recruitment pairs post-dating the 1996 year-class were excluded from the analysis: the VPA-derived recruitment estimate for 1997 is likely to have been poorly defined due to lack of convergence in the VPA method, whilst those for 1998 and 1999 were RCT3 estimates (Shepherd, 1997) which are thought to be uncertain.

2.3. Stock-recruitment modelling

Scatterplots of R and S are shown in Figure 2.3. Parameters for stock-recruitment models, namely those due to Ricker (1954), Beverton and Holt (1957), and Shepherd (1982), were estimated using non-linear least-squares regression with lognormal errors. In addition, parameters for the Sella-Lorda or gamma model $R = \alpha S^\gamma e^{-\beta S}$ were fitted to ascertain the presence of statistically-significant depensation: that is, decreasing recruits-per-spawner with decreasing stock size (Sella and Lorda, 1980; Iles, 1994; Quinn II and Deriso, 1999). Although the estimated Sella-Lorda depensation parameter $\hat{\gamma} = 1.973 > 1.0$, indicating depensation at low stock sizes, the curve could not be assumed to be significantly depensatory as an F -test comparing the fit of the Sella-Lorda and Ricker curves did not show the former to be statistically better.

The use of an assumed lognormal error distribution leads to a lognormal bias of $\exp(-\sigma^2/2)$ in the *expected* value of the fitted curve, where $\sigma^2 = \text{Var}[\ln R]$ (McCullagh and Nelder, 1983). Hence in analyses where it is the expectation or *mean* of the curve that is of interest, the estimates of R obtained from the curve must be inflated by a bias-correction term of $\exp(\sigma^2/2)$. However, stochastic projection of recruitment produces percentiles of the distribution of projected values: this does not involve the estimation of an expected value and therefore does not call for bias-correction.

Parameter estimates and a goodness-of-fit statistic, R^2 adjusted for the number of parameters, were noted for each model fit (Table 2.3). While such diagnostics are useful in selecting between similar models (for example, assessing whether the additional parameter in the Shepherd model is justified on statistical grounds), equal attention must be paid to the biological plausibility of the chosen model. Cod are thought to be cannibalistic on their young, leading to stock-dependent density effects which are encapsulated by the Ricker and Shepherd models but not by the Beverton-Holt model. A further consideration is that the projection performance of the chosen model should be plausible. In the case of cod, projections will start near the lower limit of the historically-observed range of S and it would be expected that the stock-recruitment trajectory would tend towards the origin at current high levels of fishing mortality. For many stocks, a formulation such as the Beverton-Holt model will have a steep slope at the origin and consequently high R until close to the origin, and may thus not simulate the expected decline in S as well as over-compensatory curves such as the Shepherd model which will often be nearly linear in this region. In the event all the models fitted here (except the depensatory Sella-Lorda model) look very similar at low S , but this needn't always be the case. The whole question

of the systematic selection of recruitment models on the basis of desirable properties they may have for population projections and subsequent management decisions in one that must be pursued further.

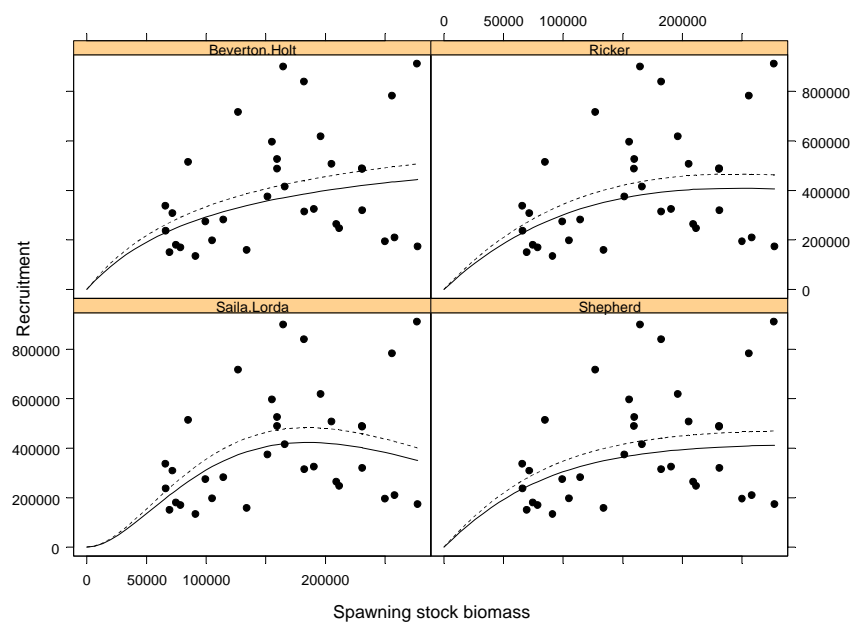
Given that projection performance factors were inconclusive for North Sea cod, the decision on which model to use was based on analysis of the R^2 diagnostic (Table 2.3) and the fact that cannibalism is not modelled by the Beverton-Holt model. Of the two models in which cannibalism could be encapsulated, the Ricker had by far the better adjusted R^2 value and was selected for that reason.

For analyses using the standard ICES Working Group bootstrapped-residual approach, the Ricker formulation used in projections was $R = r_i \alpha S e^{-\beta S}$, where $r_i = R_i / \hat{R}_i$ is a randomly-selected residual ratio from the Ricker model fit to historical stock-recruit data. When this method was modified by the addition of ARMA time-series modelling the equation used became $R = \alpha S e^{-\beta S + x}$, where x is a random variable giving rise to stochastic fluctuations in projections (see §2.5).

Table 2.3. Goodness-of-fit diagnostics for parametric model fits to 1963–1996 North Sea cod stock-recruitment data. The measure given is an R^2 statistic adjusted for the number of parameters. The selected model is highlighted in bold text.

	Ricker	Beverton-Holt	Shepherd	Saila-Lorda
Adjusted R^2	15.20%	13.13%	11.35%	14.93%

Figure 2.3. Beverton-Holt, Ricker, Saila-Lorda and Shepherd models fitted to stock-recruitment data for North Sea cod, 1963–1996. The solid lines give the least-squares fits on a lognormal scale, the dotted lines show the fits corrected for the lognormal bias.



2.4. Time-series model fitting

Time-series models for this study were fitted to $x_i = \ln(R/\hat{R})$, the logarithm of the ratio of observed to fitted recruitments, based on the Ricker stock-recruitment model (see §2.3). The reasons for this were two-fold: to reduce the significance of outliers, and to incorporate the possibility of lower R at low S in projections. Model fitting was carried out using the S-PLUS statistical package (MathSoft Inc., 1999), following and expanding upon methodology outlined in Box and Jenkins (1976) and Venables and Ripley (1999).

ARMA time-series modelling is an attempt to describe the behaviour of a data series in terms of a combination of autoregressive (AR) and moving-average (MA) effects. In order to conform to requirements for stationarity, it may also be necessary in general to *difference* the series. For the analyses described here, however, this was not the case: because the series in question are residuals from a fitted parametric model, they tend to fluctuate around a stationary mean of zero without any requirement for further intervention. If a mean value is not being fitted, an ARMA(p, q) model fitted to a series x_i is given by

$$x_i - \Phi_1 x_{i-1} - \Phi_2 x_{i-2} - \dots - \Phi_p x_{i-p} = a_i - \Theta_1 a_{i-1} - \Theta_2 a_{i-2} - \dots - \Theta_q a_{i-q},$$

where p and q are the order of the AR and MA components of the model respectively, Φ_i and Θ_i are AR and MA parameters to be fitted, and $a_i \sim N(0, \sigma_{\text{ARMA}}^2)$ are independent identically-distributed random variates known as *innovations*.

The customary initial exploratory analysis for ARMA modelling consists of plots of the autocorrelation function (ACF) and partial autocorrelation function (PACF) for the original series x_i . However, given the restricted length of the North Sea cod series, it was not possible to ascertain much evidence of autoregressive or moving-average time-series dynamics from these plots, and there was little noticeable pattern in the ACF and PACF. Because of this, and to ensure that well-fitting models did not get overlooked, we estimated systematically ARMA model fits for all combinations of AR and MA orders from zero up to three. The maximal order of the AR and MA terms was chosen to ensure that only parsimonious models are considered.

Resulting values of AIC (Akaike, 1973), and the corresponding value adjusted for a small sample size (AICc: Hurvich and Tsai, 1989), are shown in Table 2.4, while ranked order plots are given in Figure 2.4.1. For both diagnostics there is a clear step between the two best-fitting models (ARMA(1,0) and ARMA(0,1)) and the rest. However, the distinction between ARMA(1,0) and ARMA(0,1) in terms of AIC or AICc is not convincing, so a series of additional time-series diagnostics were produced. The results for each model are given in Figure 2.4.2 and can be interpreted as follows. The three left-hand plots for each model (standardised residuals, ACF and PACF) should show few outliers and little clear autocorrelation (that is, series of all-negative or all-positive values). The standardised residuals should also follow an approximately normal distribution. Hence, the histogram on the top right should be roughly unimodal and symmetric, and the accompanying quantile (or Q-Q) plot (Wilk and Gnanadesikan, 1968) should follow the fitted diagonal with little deviation. The cumulative periodogram of standardised residuals should also follow as closely as possible the diagonal: any persistent bias suggests a poorly fitting model (Bartlett, 1955). Finally, the sample projections plotted at the bottom right-hand corner should evince a similar level of variability to that seen in the

historical time-series. Unfortunately, in this case the diagnostics for the ARMA(1,0) and ARMA(0,1) models shed no further light on which should be used in projections, and it would seem that either would be equally valid. For the purposes of pedagogy and familiarity, the autoregressive ARMA(1,0) model has been used in this paper. A theoretical comparison of the ARMA(1,0) and ARMA(0,1) models warrants attention but will not be considered further in this paper.

Table 2.4. Values of Akaike information criterion (AIC) resulting from ARMA(p,q) model fits to the 1963–1996 North Sea cod time-series of $\ln(R/\hat{R})$, the logged ratio of observed to fitted recruitment for a Ricker curve. The values of AICe, a small-sample adjustment to AIC, are given in parentheses. The selected model is highlighted in bold text. nc = ARMA fit procedure did not converge.

Autoregressive order p	Moving average order q			
	0	1	2	3
0		48.597 (48.984)	50.578 (51.378)	52.577 (53.956)
1	48.579 (48.966)	50.577 (51.377)	52.577 (53.956)	nc
2	50.577 (51.377)	52.575 (53.954)	nc	nc
3	52.577 (53.956)	54.560 (56.703)	nc	nc

Figure 2.4.1. AIC (open points) and AICe (closed points) diagnostics for ARMA time-series model fits to log residual ratios $\ln(R/\hat{R})$ for North Sea cod, 1963–1996. Models range from ARMA(1,0) (model 1) and ARMA(0,1) (model 2), to ARMA(3,1) (model 10).


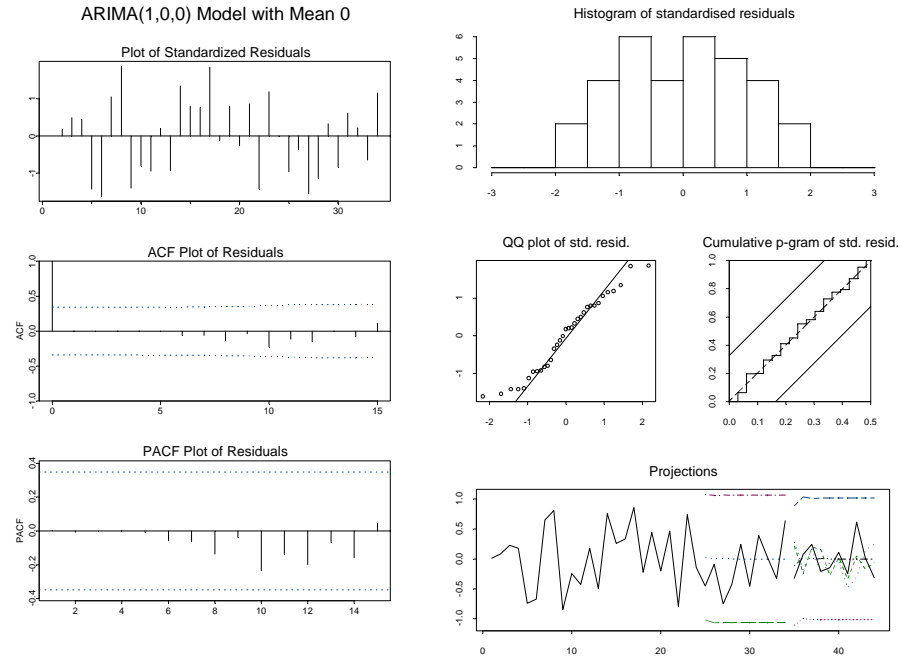
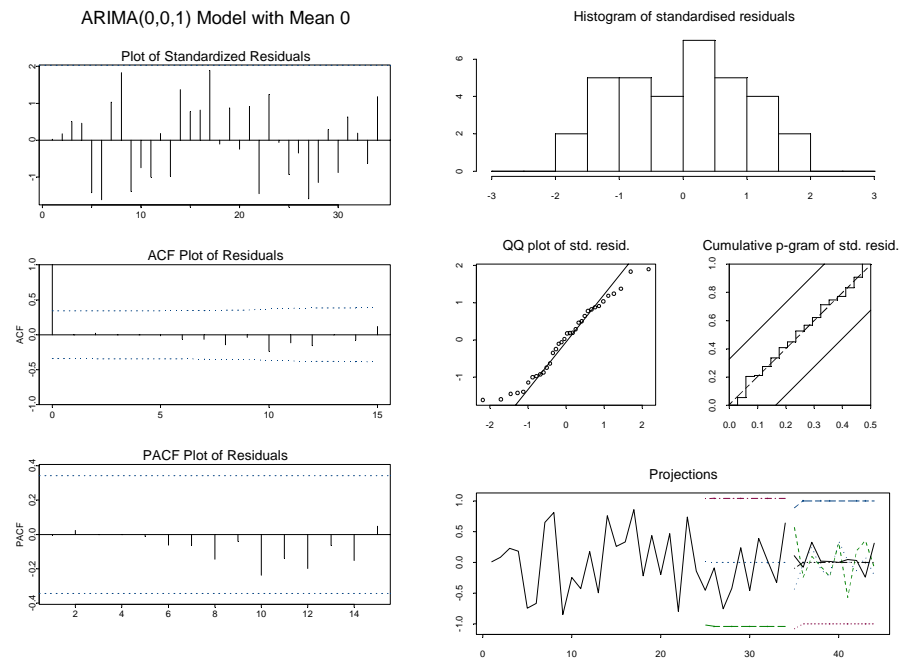


Figure 2.4.2. Time-series diagnostic plots for North Sea cod, 1963–1996. **a.** ARMA(1,0): AIC = 48.5787, $\phi_1 = -0.1842$, $\sigma^2_{\text{ARMA}} = 0.2631$. **b.** ARMA(0,1): AIC = 48.5966, $\theta_1 = 0.1805$, $\sigma^2_{\text{ARMA}} = 0.2632$.

a.



b.



2.5. Time-series projection methodology

For each simulation run, a vector $\mathbf{A} = [a_i]$ of innovations was created by random draws from a normal distribution with mean 0 and variance given by the variance of the ARMA model fit, so that $a_i \sim N(0, \sigma_{\text{ARMA}}^2)$. North Sea cod recruit at age 1, so the first value of the projected time-series vector $\mathbf{X} = [x_i]$ was given by the logged residual ratio for the final historical assessment year. Subsequent values of \mathbf{X} were generated from this point using the innovations vector. For an autoregressive ARMA(1,0) model with parameter ϕ_1 the i th projection value is

$$x_i = \phi_1 x_{i-1} + a_i,$$

while the equivalent for a moving-average ARMA(0,1) model with parameter θ_1 is

$$x_i = a_i - \theta_1 a_{i-1}.$$

The required projected Ricker recruitment is then

$$R_i = \alpha S_{i-1} e^{-\beta S_{i-1} + x_i},$$

assuming age at recruitment to be 1. Once recruitment is calculated, population dynamics are processed as for the standard bootstrapped-residual approach described in §2.1.

2.6. Inclusion of a temperature index in the cod stock recruitment relationship

The potential effect of changes in climate on fish stocks is a topical issue, and for some stocks strong links between climatic signals and recruitment have been demonstrated (Borja *et al.* 1996, 1998; Planque and Frédou, 1999), although the mechanisms driving these relationships are not well understood. For cod stocks towards the southern margin of their range, such as the North Sea cod, increasing temperature has been shown to be negatively correlated with recruitment (O'Brien *et al.*, 2000). Temperature may be incorporated in the Ricker stock-recruitment model by including the temperature signal and associated parameter as an extra exponential term

$$R = \alpha S e^{-\beta S} e^{\varphi T},$$

where T is the temperature index and φ a parameter relating recruitment to temperature (Hilborn and Walters, 1992). As before, parameters were estimated by least-squares minimisation of the log-recruitment residuals. The temperature regime for the projections was obtained by randomly selecting a temperature for each year from the North Sea temperature signal for the period (1987-1996). The temperature signal used was the mean annual North Sea sea surface temperature for the months February to June, derived from the Comprehensive Ocean Atmosphere Dataset (COADS) and provided by the National Center for Atmospheric Research (NCAR, Boulder, Co.). It must be emphasised that these temperature data were chosen because they correlate reasonably well with cod recruitment, and temperature seems to be an appropriate proxy for some of the underlying (and unknown) biological drivers of recruitment. It is not possible to precisely predict future temperature regimes, and the use of temperature as a predictor of recruitment is therefore limited to the investigation of alternative future scenarios and subsequent management responses.

3. Medium-term projection results

Projections of population dynamics for a standard Working Group bootstrapped-residual model (denoted in output figures by *Working.Group*), the alteration to incorporate a temperature index (*Temperature*), and the ARMA model (*ARMA*) were performed for 10 years and 1000 iterations each, respectively. Stochasticity was introduced through recruitment projections only: all other variables, including starting population values, were fixed in order to clarify the effect of the different methods of resampling recruitment. Plots of percentiles of projected recruitment and SSB are given in Figures 3.1 and 3.2, along with point estimates produced by the 1999 Working Group (ICES 2000) based on final-year VPA results and RCT3 recruitment estimates. The point estimates for recruitment in 1998 and 1999 were the two smallest recruiting year-classes in the entire time-series, so it should not be surprising that these values lie outside the 90% confidence intervals of all of the models. It is more pertinent to note that the point estimate of SSB for 1999 lies very close to the median projected value of every model. For both recruitment and SSB the spread of projections from the ARMA model is very much tighter than from the other two: this is to be expected as time-series variation is explicitly modelled in this case. The median recruitment and SSB levels are slightly higher for the bootstrapped-residual than for the ARMA model: inclusion of a temperature effect in the former slightly reduces the medians, as would be expected given that the temperature data on which the projections were based were taken from a relatively warm period (1987-1996).

Figure 3.3 gives probability profiles for projected SSB in 2001 and 2006 under a range of F -multipliers. These support the conclusions reached from the projection plots, namely: a) ARMA projections are less uncertain than those from bootstrapped-residual models, and b) the incorporation of a temperature signal reduces recruitment, although this is a function of the period from which temperature data were taken. A caveat is provided by the values of \bar{F}_{2-8} giving a 10% probability of SSB in 2006 being less than B_{pa} given in Table 3.1: although the median SSB from the ARMA model is lower than that from the bootstrapped-residual model, the reduced uncertainty of the former mean that the 10th percentile of the SSB distribution is actually higher, leading to a higher estimate for a precautionary F .

Table 3.1. Values of \bar{F}_{2-8} giving a 10% probability of SSB in 2006 being less than $B_{pa} = 70000$ kg.

Model	\bar{F}_{2-8}
Working Group bootstrapped-residual	0.730
Working Group bootstrapped-residual with temperature covariate	0.685
ARMA	0.770

A final diagnostic tool is presented in Figure 3.4. Here the median recruitment projection for each model has been plotted alongside a deterministic projection: that is, one in which recruitment is based purely on the fitted stock-recruitment curve, with no stochasticity. Deviation between the median and deterministic projections for a given model would suggest that there is structure in the historical data that has not been adequately encapsulated by that model. It is clear that the ARMA model has less systematic bias and is thus preferable on this basis.

Figure 3.1. Projected recruitment for North Sea cod, 1963–1996 (10 years, 1000 simulation runs), under three different projection models. Plotted percentiles are 5% (dotted), 25% (dashed), 50% (solid), 75% (dashed), and 95% (dotted). Recruitment estimated by the Working Group (ICES, 2000) is shown as points.




Figure 3.2. Projected SSB for North Sea cod, 1963–1996 (10 years, 1000 simulation runs), under three different projection models. Plotted percentiles are 5% (dotted), 25% (dashed), 50% (solid), 75% (dashed), and 95% (dotted). SSB estimated by the Working Group (ICES, 2000) is shown as points.




Figure 3.3. Probability profiles for North Sea cod SSB in 2001 and 2006, under three projection models and a range of imposed rates of fishing mortality (F_{bar}). Displayed percentiles are 50% (solid), 5%, 10%, 25% and 95% (all dotted). The reference and limit levels of SSB (respectively, $B_{pa} = 140,000$ kt and $B_{lim} = 70,000$ kt) are shown as solid horizontal lines.





Figure 3.4. Diagnostic plot of deterministic (dotted line) and median (solid line) projections for North Sea cod, under three different projection models. For the *Temperature* model, the deterministic recruitment is derived assuming temperature to be equal to the mean of the last ten years, while the median recruitment is the median of 1000 stochastic runs where temperature was selected randomly from that ten year period. For the other models, the deterministic recruitment is obtained by removing any variability in the projection of recruitment.



4. Discussion

While we await the results of several contemporaneous large-scale biological process studies the need for simple empirical methods of population projection is clear. If these are to be useful we would be well-advised to make them as statistically appropriate as possible, giving due regard to the use to which the methods are to be put. The ARMA model does not pretend to model biology directly, but is an attempt to improve the current empirical approach by accounting for time-series structure in the historical data and producing projections conforming to that same structure. The outputs of SSB distributions from the ARMA model are at a slightly lower level than those from the bootstrapped-residual model, and with less uncertainty: the second feature is due to the fact that the time-series structure in the recruitment residuals is now modelled explicitly rather than being allowed to influence randomly the outputs. Given this, we propose that ARMA time-series models might be contemplated for medium-term projections by assessment working groups when evidence of historical time-series structure exists. Including a temperature signal as a covariate in the stock-recruitment relationship may help to explain more of the variation for some stocks. However, since the mechanisms driving the relationship are poorly understood it may be that temperature is merely a proxy for the true causal environmental factor(s). Further, prediction of climatic variables into the future, particularly for several years, is difficult or impossible so projections are limited to "what if?" scenarios under various broad environmental regimes. Nonetheless given current concerns of global warming there may be value in such exercises.

There are a number of additional analyses to be carried out in the future to extend the relevance of these findings for fisheries management. The statistical properties of the parametric estimation procedure are worthy of further investigation since the estimating equations are reminiscent of those first proposed by Aitken (1935) and might be improved. The starting date for projections could be pushed progressively further back in time, allowing more extensive evaluation of the degree to which projections from different models agree with estimated SSB values. Bootstrapped sensitivity tests could be used to build up a distribution of median projected SSB levels, thus yielding a measure of the robustness of this metric of projection performance. Where ARMA modelling is inappropriate, investigations could be undertaken into alternative schemes, whereby blocks of residuals are resampled, which would retain some of the time-series structure in projections. The selection of models with due regard to plausible performance in projections should be studied further.

Acknowledgements

This paper was prepared with funding support provided by the Ministry of Agriculture, Fisheries and Food (contract MF0316) to the three authors at CEFAS.

References

- Aitken, A.C. (1935). On least squares and linear combinations of observations. *Proc. R. Soc. Edinburgh*, **55**:42–48.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In *Proceedings of the Second International Symposium on Information Theory*, pp. 267–281. Ed. by B.N. Petrov and F. Csàki. Akadémiai Kiadó, Budapest.
- Baranov, T.I. (1918). On the question of the biological basis of fisheries. *Nauch. issledov. iktiolog. Inst. Izv.*, **1**:81–128.
- Bartlett, M.S. (1955). *Stochastic Processes*. Cambridge: Cambridge University Press.
- Beverton, R.J.H. and Holt, S.J. (1957). *On the Dynamics of Exploited Fish Populations*. Vol. 19 of *Fisheries Investigations, Series 2*. London: Ministry of Agriculture and Fisheries.
- Borja, A., Uriate, A., Valencia, V., Motos, L. and Uriarte, Ad. (1996). Relationship between anchovy (*Engraulis encrasicolus* L.) recruitment and the environment in the Bay of Biscay. *Sci. Mar.*, **60** (Suppl. 2):179–192.
- Borja, A., Uriate, A., Egaña, J., Motos, L. and Valencia, V. (1998). Relationship between anchovy (*Engraulis encrasicolus* L.) recruitment and the environment in the Bay of Biscay. *Fish. Oceanogr.*, **7**:375–380.
- Box, G.E.P. and Jenkins, G.M. (1976). *Time Series Analysis: Forecasting and Control* (2nd edition). San Francisco: Holden-Day.
- Hilborn, R. & Walters, C.J. (1992). *Quantitative Fisheries Stock Assessment: Choice, Dynamics and Uncertainty*. Chapman & Hall, Inc.
- Hurvich, C.M. and Tsai, C.-L. (1989). Regression and time series model selection in small samples. *Biometrika*, **76**:297–307.
- ICES (2000). Report of the 1999 Working Group on the Assessment of Demersal Stocks in the North Sea and Skagerrak. *ICES CM 2000/ACFM*:7.
- Iles, T.C. (1994). A review of stock-recruitment relationships with reference to flatfish populations. *Netherlands Journal of Sea Research*, **32**(3/4): 399–420.
- MathSoft Inc. (1999). *S-PLUS 2000 Programmer's Guide*. Seattle: Data Analysis Products Division, MathSoft.
- McCullagh, P. and Nelder, J.A. (1983). *Generalized linear models*. London: Chapman & Hall.

- O'Brien, C.M., Fox, C.J., Planque, B. and Casey, J. (2000). Climate variability and North Sea cod. *Nature*, **404**:142
- Planque, B. and Frédou, T. (1999). Temperature and the recruitment of Atlantic cod (*Gadus morhua*). *Canadian Journal of Fisheries and Aquatic Science*, 56:2069-2077.
- Quinn II, T.J. and Deriso, R.B. (1999). *Quantitative Fish Dynamics*. Oxford: Oxford University Press.
- Ricker, W.E. (1954). Stock and recruitment. *Journal of the Fisheries Research Board of Canada*, **11**, 559–623.
- Reeves, S.R. and Cook, R. (1994). Demersal assessment programs, September 1994. *Working document to the 1994 meeting of the ICES North Sea Demersal Working Group*.
- Saila, S.B. and Lorda, E. (1980). A critique and generalization of the Ricker stock-recruitment curve. Unpublished manuscript.
- Shepherd, J.G. (1982). A versatile new stock-recruitment relationship for fisheries, and the construction of sustainable yield curves. *Journal du Conseil International pour l'Exploration de la Mer*, **40**:67–75.
- Shepherd, J.G. (1997). Prediction of year-class strength by calibration regression analysis of multiple recruit index series. *ICES Journal of Marine Science*, **54**: 741–752.
- Venables, W.N. and Ripley, B.D. (1997). *Modern Applied Statistics With S-PLUS* (2nd edition). New York: Springer-Verlag, Inc.
- Wilk, M.B. and Gnanadesikan, R. (1968). Probability plotting methods for the analysis of data. *Biometrika*, **55**:1–17.