Estimation of time-varying selectivity in stock assessments using state-space models

Anders Nielsen[a,*], Casper W. Berg[a]

[a]*National Institute of Aquatic Resources, Technical University of Denmark, Charlottenlund Castle, 2920 Charlottenlund, Denmark.*

Abstract

Time-varying selectivity is one of the main challenges in single species age-based assessment models. In classical deterministic VPA-type models the fishing mortality rates are unfiltered representations of the observed catches. As a consequence the selectivity becomes time-varying, but this representation is too fluctuating, because it includes the observation noise. In parametric statistical catch at age models a common assumption is that the selectivity is constant in all years, although time-varying selectivity can be introduced by splitting the data period in blocks with different selectivity, or by using smoothing splines and penalized time-deviances. However, these methods require subjective choices w.r.t. the degree of time-varying allowed. A simple state-space assessment model is presented as an alternative, which among other benefits offers an objective way of estimating time-varying selectivity pattern. The fishing mortality rates are considered (possibly correlated) stochastic processes, and the corresponding process variances are estimated within the model. The model is applied to North Sea cod and it is verified from simulations that time-varying selectivity can be estimated.

Keywords: Stock assessment, Selectivity, State-space models, Catch-at-age analysis

---

1. Introduction

In stock assessment models the selectivity describes the relationship between the age/size composition in the population and the one found in the fisheries or a survey. In age based models the selectivity $S_{a,y}$ is defined as

$$S_{a,y} = \frac{F_{a,y}}{\sum_a F_{a,y}},$$

where $F_{a,y}$ is the instantaneous fishing mortality rate for age group a in year y. Often the selectivity is scaled to have a maximum of one in each year.

This relationship is determined by a combination of underlying processes such as gear selectivity and the spatial-temporal distribution of the fishery and of the species under consideration. These underlying processes may change over time resulting in time-varying selectivity patterns, which, when not accounted for, may lead to biased estimates of key management quantities (Sampson and Scott, 2012). The degree to which selectivity should be allowed to evolve over time in an assessment model involves a bias-variance trade-off: a constant pattern will likely introduce some bias given that the true processes are unlikely to be constant, while allowing for a very fluctuating selectivity over time introduces excessive variance, since the observation noise is not filtered out.

Traditionally one of the two extremes have been chosen. In virtual population analysis (VPA) there are no parametric assumptions or smoothness constraints on the temporal developments in fishing mortalities-at-age, but rather they are allowed to fluctuate freely in order to match the observed catches exactly. While this allows the selectivity

---

*corresponding author

Email address: an@aqua.dtu.dk (Anders Nielsen)

Postprint of publication DOI 10.1016/j.fishres.2014.01.014

to evolve over time, the assumption of having catches without error and the lacking ability to quantify uncertainties in VPA are problematic. In more modern variants of VPAs, such as XSA (Shepherd, 1999), a smoothing parameter named 'shrinkage' is introduced (ICES, 2009b). Shrinkage is a user supplied CV, which restricts how much the fishing mortalities are allowed to change in a single year. In contrast, statistical catch-at-age (SCAA) models treat catches as observations with noise, but often a separability assumption is made about the fishing mortality-at-age ($F_{a,y} = F_a F_y$), which implies that the selectivity is constant over time. To further reduce the number of parameters in the model a parametric relationship between F and age a is often assumed, typically a sigmoid function. However, Sampson and Scott (2012) identified four broad types of selectivity from a meta-analysis: increasing, asymptotic, domed, and having a saddle, as well as sometimes radical changes in shape over time, so it is important to always examine the validity of such assumptions. In data poor situations such assumptions might however be necessary (Methot and Wetzel, 2013).

A simple way of introducing time-varying selectivity is to select blocks of time with constant selectivity, but unless some discrete event caused abrupt changes in selectivity, it may not be straightforward to select the blocks. Other examples include estimating deviations around a mean, random walks deviations, or Gaussian trends (Methot and Wetzel, 2013; Taylor and Methot, 2013). A smooth selectivity surface over both the time and age dimension was used by Butterworth et al. (2003). Smoothness was achieved by adding a curvature penalty to the likelihood function, which prevents irregular shifts in either dimension, but on the other hand allows gradual changes in selectivity as well as in total fishing mortality. This is appealing and relieves the modeler from having to choose between parametric forms and/or time blocks. However, when using the penalized likelihood approach the modeler still has to specify the degree of smoothing rather than estimate it from data. A similar method where the degree of smoothness has to be specified apriori was used by Aarts and Poos (2009).

This paper presents a state-space modelling approach to estimating time-varying selectivity, which is smooth in both the age and time-dimensions as in Butterworth et al. (2003), but where the degree of smoothing is estimated from data within the model. The state-space modelling approach to age-based assessments was introduced by Gudmundsson (1994), but has recently received increased attention (Brinch et al., 2011; Gudmundsson and Gunnlaugsson, 2012; Berg et al., 2013) due to advancements in sophisticated software packages that can estimate efficiently in such models (Fournier et al., 2012). A main feature of state-space models is that it is possible to separate observation noise from process noise, and the degree of temporal development in fishing mortalities-at-age is therefore naturally estimated from data as a process variance parameter within this framework. The potential importance of estimating the degree of smoothness rather than specifying it in an ad-hoc manner is illustrated in figure 1.

A smoothness penalty on the fishing mortalities was used in the XSA-based assessment of Eastern Baltic cod in 2009, but results were highly sensitive to the choice of penalty (ICES, 2009a). Such a smoothness penalty is naturally represented as a process variance parameter in a state-space model. A state-space assessment model (SAM), where this variance parameter was estimated from data using maximum likelihood, gave results lying in between a new candidate and previously chosen value of the penalty (see Figure 1).

In the state-space model introduced in this paper the fishing mortality rates by age are considered stochastic processes, and the amount of correlation between these processes determines how much the selectivity can evolve over time. Four simple correlation structures for these processes are compared using the North Sea Cod stock as an example, and it is demonstrated that including time-varying selectivity greatly improves the model in terms of AIC.

2. Methods

The state-space assessment model contains two parts. The first part describes the process of underlying unobserved states α, which are the log-transformed stock sizes and fishing mortalities, $\alpha = (\log N_1, \ldots, \log N_A, \log F_1, \ldots, \log F_{A^*})'$. The oldest age groups may be grouped together with respect to F, indicated by using subscript (A^*) rather than the maximum age group or "plus group" in the model (A). The transition equation (below) describes the distribution of the next year's state from a given state in the current year.

$$\alpha_y = T(\alpha_{y-1}) + \eta_y$$

The transition function T is where the stock equation and assumptions about stock–recruitment enters the model. For the stock sizes this becomes:

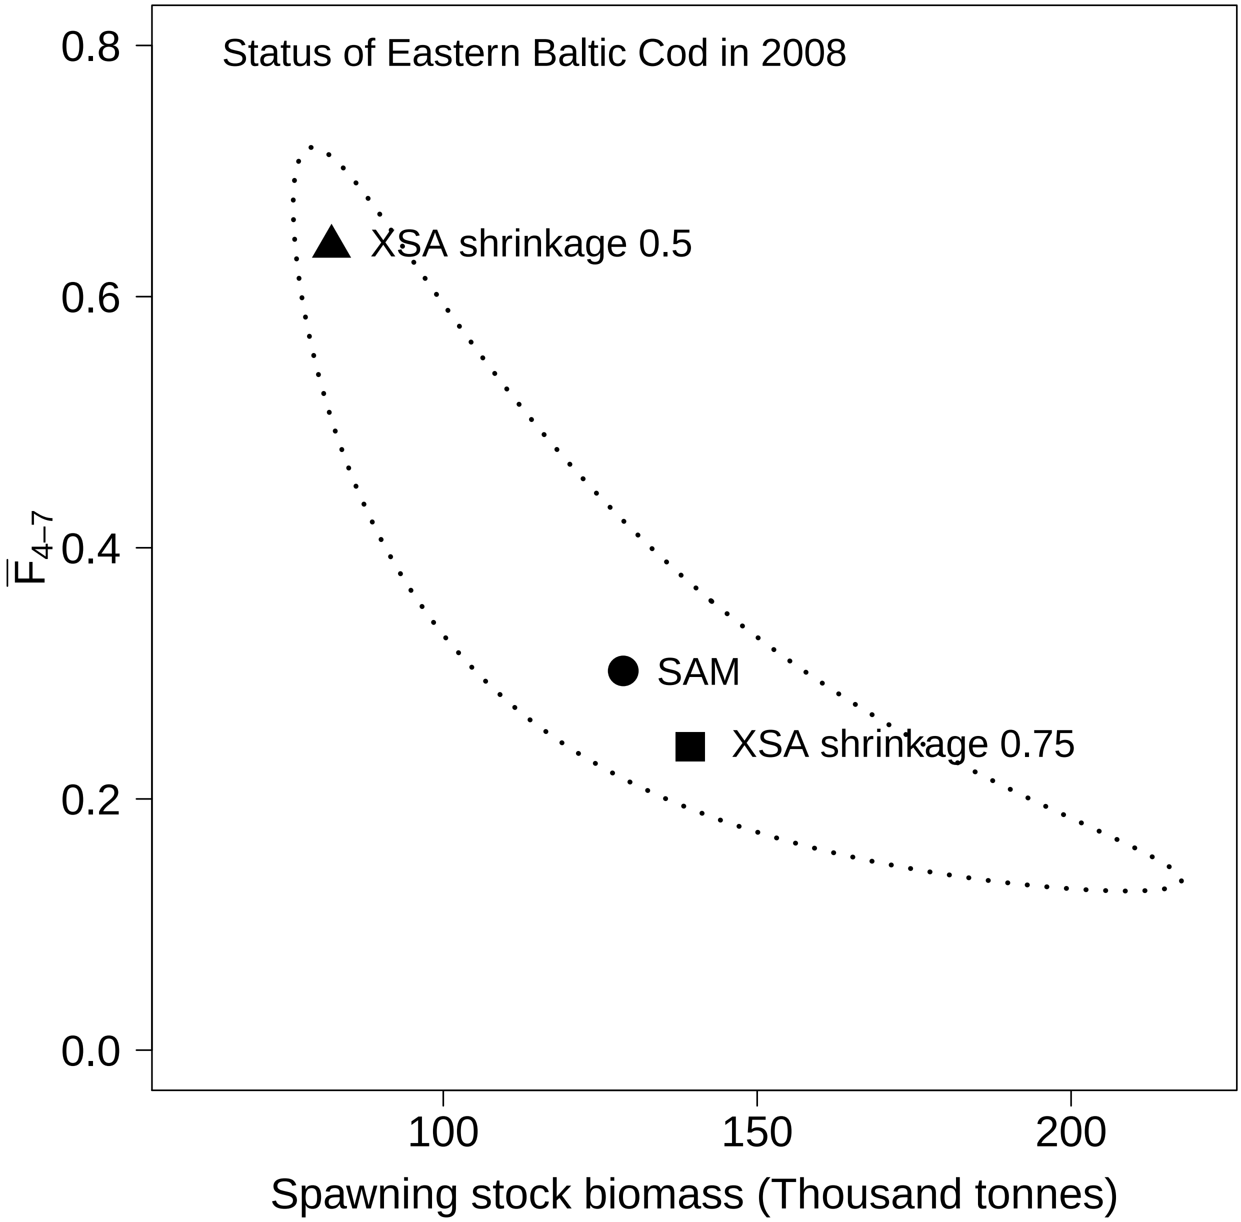

Figure 1: Effect of ad-hoc tuning of smoothness in the time-direction for *F* on the final assessment estimates using an XSA-type model versus an alternative state-space model (SAM) with no ad-hoc tuning (Recreated from ICES (2009a)). Dotted line indicates a 95% confidence region from SAM.

$$
\begin{aligned}
\log N_{1,y} &= \log\Big(R\big(w_{1,y-1}p_{1,y-1}N_{1,y-1} + \cdots + \\
&\qquad w_{A,y-1}p_{A,y-1}N_{A,y-1}\big)\Big) + \eta_{1,y}\ , \\
\log N_{a,y} &= \log N_{a-1,y-1} - F_{a-1,y-1} - M_{a-1,y-1} + \\
&\quad \eta_{a,y}\ ,\ 2 \le a < A\ , \\
\log N_{A,y} &= \log\Big(N_{A-1,y-1}e^{-F_{A-1,y-1}-M_{A-1,y-1}} + \\
&\qquad N_{A,y-1}e^{-F_{A,y-1}-M_{A,y-1}}\Big) + \eta_{A,y}.
\end{aligned}
$$

Here $M_{a,y}$ is the age specific natural mortality parameter, $w_{a,y}$ is weight in stock, and $p_{a,y}$ is proportion mature, all of which are assumed known. $F_{a,y}$ is the fishing mortality. The function R describes the relationship between stock and recruitment (for North Sea Cod a Beverton-Holt curve is assumed). The parameters of the chosen stock–recruitment function are estimated within the model. All process error terms for the logarithm of the stock sizes ($\log N$) are assumed independent normal distributed.

The logarithm of the fishing mortalities are assumed to follow random walks. Setting up the random walk on the logarithm of F ensures that F itself will be positive. The random walks are allowed to be correlated to mimic the to some degree parallel time series often observed for fishing mortalities in the different age groups. Define $F_y = (F_{1,y}, F_{2,y}, \ldots, F_{A^*,y})'$, then it is assumed that F_y follows a random walk with multivariate normal increments, i.e.

$$
\log F_y = \log F_{y-1} + \xi_y \quad , \qquad \text{where} \qquad \xi_y \sim \mathbf{N}(0, \Sigma)
$$

where Σ is the covariance matrix. Σ is defined via the standard deviation for the individual processes $\sigma_a = \sqrt{\Sigma_{a,a}}$ and the correlation coefficient ρ. Four different models will be explored:

A) The correlation $\rho = 1$ for all combination of ages ($a \neq \tilde{a}$), such that $\Sigma_{a,\tilde{a}} = \sigma_a\sigma_{\tilde{a}}$, and hence the individual $\log F$ processes will develop parallel in time. In case of equal process variances (the diagonal of Σ) this is equivalent to an assumption of constant selectivity in time (having completely parallel increments on $\log F$ will result in constant F proportions).

B) The correlation $\rho = 0$ for all combination of ages ($a \neq \tilde{a}$), such that $\Sigma_{a,\tilde{a}} = 0$, hence the individual $\log F$ processes will develop independently in time.

C) The correlation is set to be a common, but estimated, constant for all combination of ages ($a \neq \tilde{a}$), such that $\Sigma_{a,\tilde{a}} = \rho\sigma_a\sigma_{\tilde{a}}$. Then the individual $\log F$ processes will develop correlated in time, and the correlation will be common for any pair of age specific processes. This correlation structure is commonly named compound symmetry.

D) The correlation is a simple function of the age difference. The function chosen is $\Sigma_{a,\tilde{a}} = \rho^{|a-\tilde{a}|}\sigma_a\sigma_{\tilde{a}}$, then the individual $\log F$ processes will develop correlated in time, and correlation is set to reflect the intuition that neighboring age classes should have more similar fishing mortalities. This correlation structure is commonly named AR(1).

The second part of the state-space assessment model describes the distribution of the observations x given the underlying states α. Here x consists of the log-transformed catches and survey indices.

The combined observation equation is:

$$
x_y = O(\alpha_y) + \varepsilon_y\ ,
$$

The observation function O consists of the familiar catch equations for fleets and surveys, and ε_y of independent measurement noise with separate variance parameters for certain age groups, catches, and survey indices. For the logarithm of the survey catches, a separate variance parameter is used for the youngest age group and a common parameter for all older age groups. For the logarithm of the total catches, a separate variance parameter is used for each of the two youngest age groups, and a common parameter for all older age groups. An expanded view of the observation equation gives:

$$\log C_{a,y} = \log\left(\frac{F_{a,y}}{Z_{a,y}}(1 - e^{-Z_{a,y}})N_{a,y}\right) + \varepsilon_{a,y}^{(\circ)},$$
$$\log I_{a,y}^{(s)} = \log\left(Q_a^{(s)} e^{-Z_{a,y}\frac{day^{(s)}}{365}} N_{a,y}\right) + \varepsilon_{a,y}^{(s)}.$$

Here, Z is the total mortality rate $Z_{a,y} = M_{a,y} + F_{a,y}$, $day^{(s)}$ is the number of days into the year where the survey s is conducted, and $Q_a^{(s)}$ are model parameters describing catchabilities. Finally $\varepsilon_{a,y}^{(\circ)} \sim N(0, \sigma_{\circ,a}^2)$ and $\varepsilon_{a,y}^{(s)} \sim N(0, \sigma_{s,a}^2)$ are all assumed independent and normal distributed.

The likelihood function for this is set up by first defining the joint likelihood of both random effects (here collected in the α_y states), and the observations (here collected in the x_y vectors). The joint likelihood is:

$$L(\theta, \alpha, x) = \prod_{y=2}^{Y}\{\phi(\alpha_y - T(\alpha_{y-1}), \Sigma_\eta)\}\prod_{y=1}^{Y}\{\phi(x_y - O(\alpha_y), \Sigma_\varepsilon)\}$$

Here θ is a vector of the 15 or 16 model parameters. A full list of the model parameters in θ can be seen in Table 1. Since the random effects α are not observed, inference should be carried out using the marginal likelihood:

$$L_M(\theta, x) = \int L(\theta, \alpha, x) d\alpha$$

This integral is difficult to calculate directly, so the Laplace approximation is used. The Laplace approximation is fast and accurate in many applications (Skaug and Fournier, 2006). Model selection is carried out using AIC since model C and D are not nested models, so the likelihood ratio test cannot be used here.

To verify the model, 100 data sets are simulated using the state estimates (N and F) from the North Sea cod case as the underlying truth. Catch-at-age and survey indices are generated by adding observation noise according to the estimated standard deviations, and full estimation is carried out for each of these data sets. This strengthens confidence in model implementation, its identifiability, that confidence bounds are reasonable, and that the procedure is robust. Parameter estimation is carried out using AD Model Builder (Fournier et al., 2012) and the uncertainty estimation follows the method of Skaug and Fournier (2006). The entire source code for all the models is available online at www.stockassessment.org (assessment names using "SELPAP" as prefix).

3. Case Study

In this section the method will be applied to 49 years of North Sea cod data (1963-2011). The commercial catch-at-age data, survey indices, natural mortalities, proportion mature, and weight-at-age used in the stock assessments are taken from ICES (2012), as is the number of age groups (1 to 7+). This case was chosen because there have been substantial changes in the fleet and the mesh-sizes used (ICES, 2012, pg. 1230), although it is not obvious apriori how the sum of these changes will affect selectivity. The natural mortalities (M) used are assumed known and equal to estimated values from a stochastic multispecies model, which accounts for predation (ICES, 2011), so these are not constant over time. The estimated selectivity pattern is conditioned on M (He et al., 2011; Punt et al., 2013), however since it is generally not possible to estimate both M and F without extra knowledge or data this assumption is needed, but the dependence on M should be kept in mind when interpreting the results.

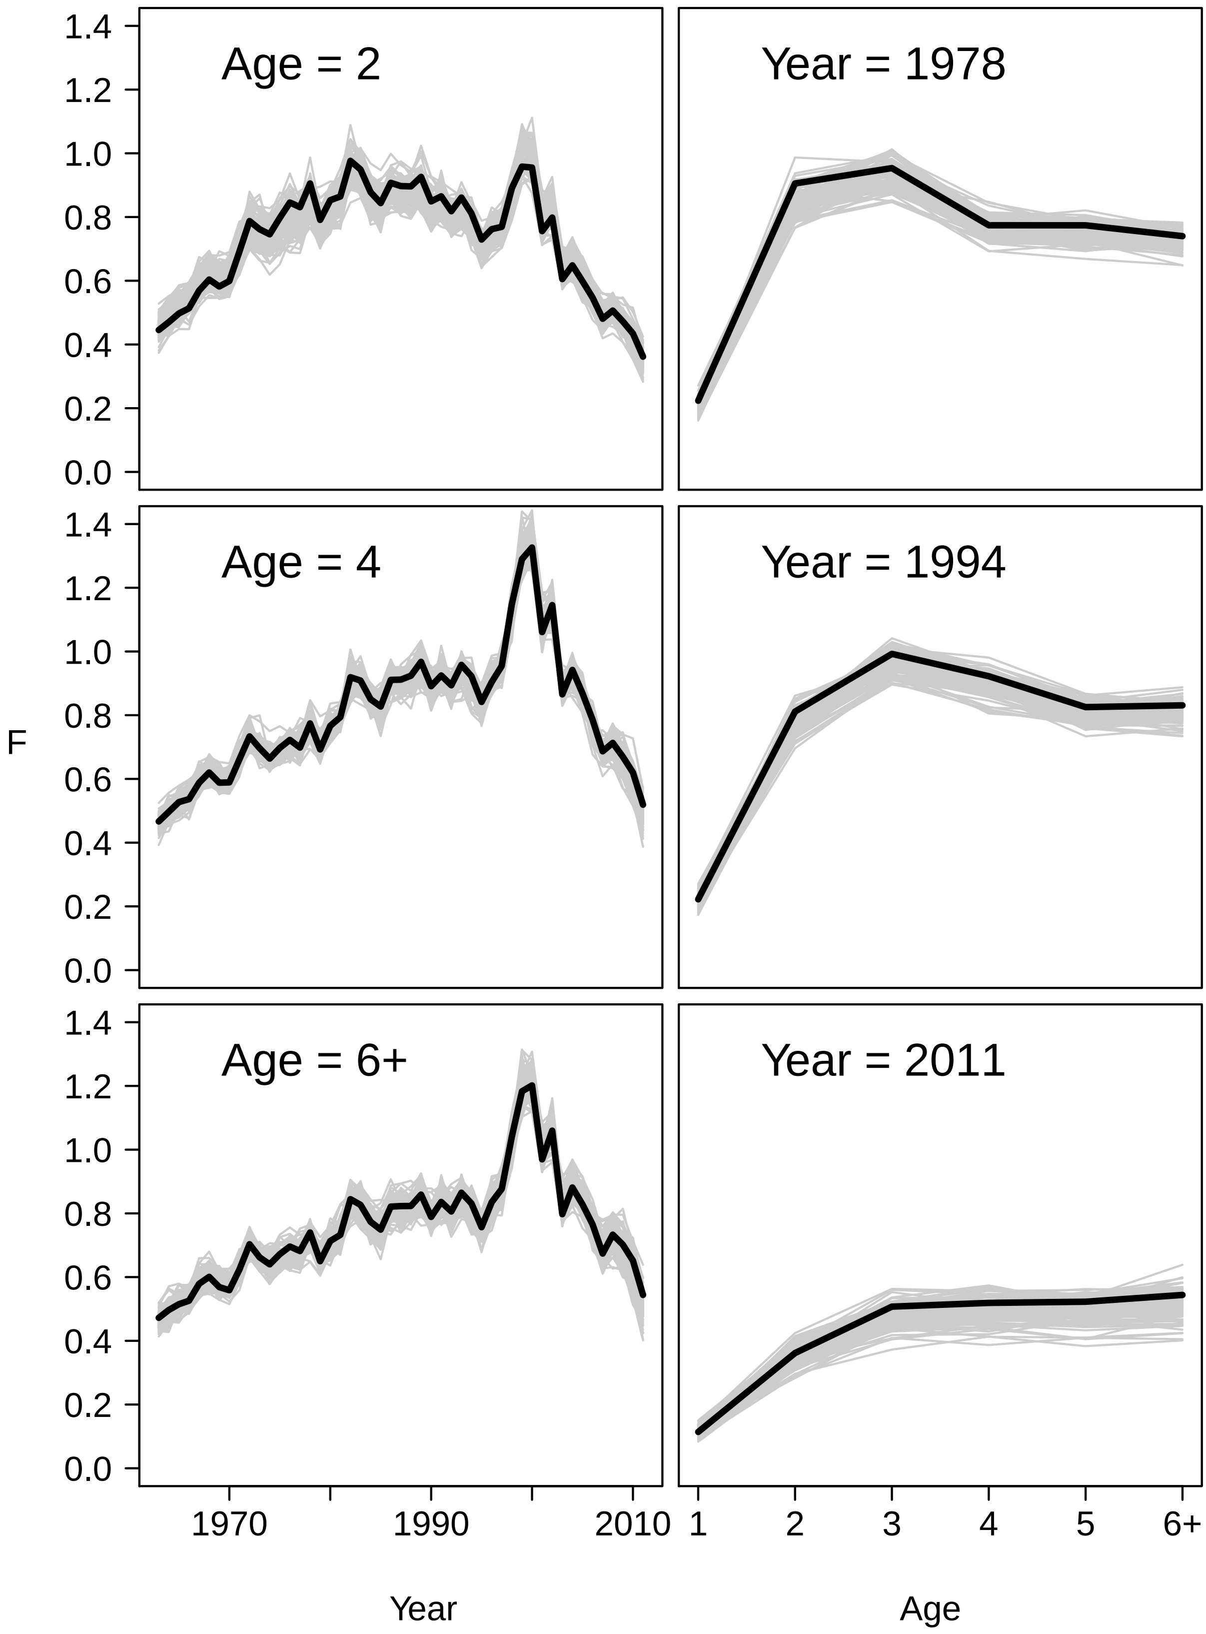

Figure 2: North Sea cod: Fishing mortalities (F) over time for three selected ages (left panel) and by age for three selected years (right panel) using model C. Truth as estimated from actual data (bold black) and re-estimated from conditional simulations (thin gray).

Model parameter	Description
$\sigma_F, \sigma_R, \sigma_S$	Process standard deviations for fishing mortality, stock–recruitment, and survival. All specified on logarithmic scale.
$q_{a=1}, q_{a=2}, q_{a=3}, q_{a=4}, q_{a=5}$	Survey catchability parameters.
a_{sr}, b_{sr}	Beverton–Holt stock–recruitment parameters
$\sigma_{C_{a=1}}, \sigma_{C_{a=2}}, \sigma_{C_{a\geq 3}}, \sigma_{I_{a=1}}, \sigma_{I_{a\geq 2}}$	Observation standard deviations for different ages of catches and surveys. All specified on logarithmic scale.
ρ	Correlation parameter. Present in model configuration (C and D).

Table 1: List on model parameters estimated by the model

model	$-\log L_M$	no. par	AIC
A	151.72	15	333.44
B	171.02	15	372.04
C	138.70	16	309.40
D	130.81	16	293.63

Table 2: North Sea cod: Model selection criteria for models A–D.

3.1. Results

In the following, focus will be on the differences between models A–D, i.e. ($F_{a,y}$ and $S_{a,y}$), key management outputs ($\bar{F}$ and SSB), and predictive power (AIC) for the North Sea cod case, rather on the specifics of the particular case.

Figure 2 illustrates two important points: 1) that there have been significant changes in fishing selectivity over time for this stock and 2) that the model is identifiable, i.e. we can estimate these changes from simulated data. The estimated selectivity patterns using models A–D are illustrated in Figure 3. Where model A has constant selectivity over time, model B displays the most wiggly selectivity pattern, and model C and D lie somewhere in between. Notice that it is possible to obtain smooth yet flexible selectivity patterns using only one extra parameter in the model, ρ. Figure 4 shows that significantly different assessment results are obtained using model A versus model B, whereas model C and D produce rather similar results. The profile likelihoods for the ρ-parameters in models C and D are shown in Figure 5. These illustrate how models A and B are both special cases of models C and D ($\rho_A = 1$ and $\rho_B = 0$). Here the constant selectivity of model A was better than the independent random walks in model B, but both were significantly improved by estimating $\hat{\rho}_C$ or $\hat{\rho}_D$. In this case $\hat{\rho}_D$ seems more appropriate. The normalized catch residuals (Supplemental Figures 1–4) exhibit pronounced autocorrelation for model A as opposed to models B–D, which are similar. Some year effects are present in the survey residuals for all models, which means that further refinement of the model should be pursued, but this is out of the scope of this paper. The likelihoods, number of model parameters, and the AIC values for all the models are also shown in Table 2 and variance parameter estimates in Supplemental Table 1, and we must conclude that model D must be preferred here since this has the lowest AIC value.

4. Discussion

A state-space model for stock assessment has been presented, which by using only a single extra parameter can capture a continuum of degrees of smoothness in the temporal development in selectivity – from constant selectivity to completely independent developments in F by age. This is achieved by assuming that the F vector follows a multivariate random walk with a simple covariance structure described by only one parameter (ρ). While the ρ parameter describes the degree to which the F-processes follow parallel trajectories in time, the variance parameter σ_F describes the general temporal variability in fishing pressure and is considered constant over ages and time in this presentation. This may easily be extended, such that σ_F (or ρ) also depends on time, which could for instance be relevant if multiple fleets are fishing on the stock, and their relative effort change over time.

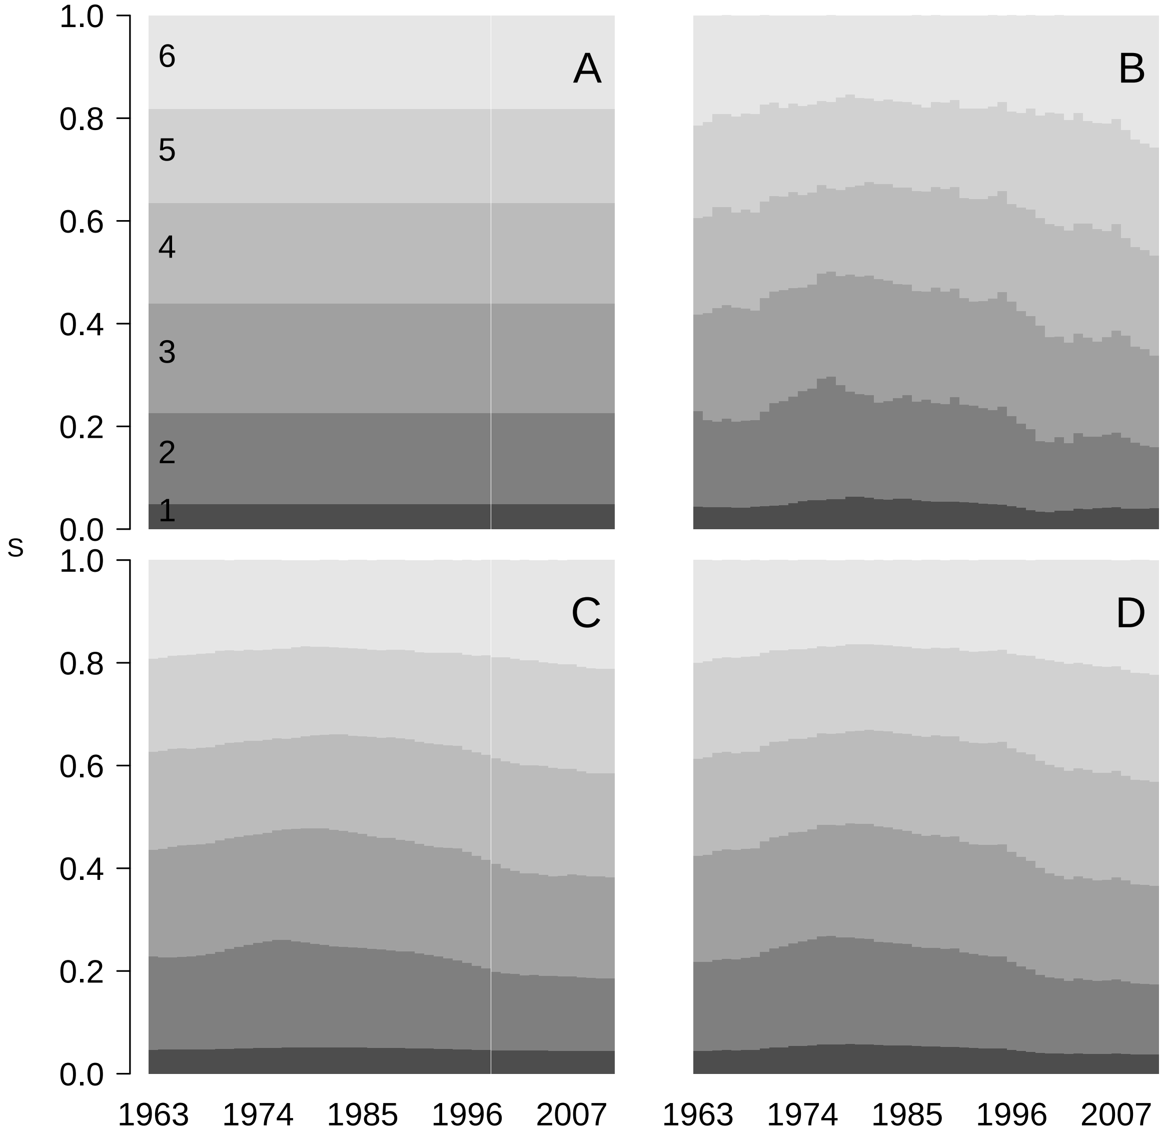

Figure 3: North Sea cod: Proportions of total fishing mortality (i.e. selectivity) by age group over time for each of the examined correlation structures (A–D).

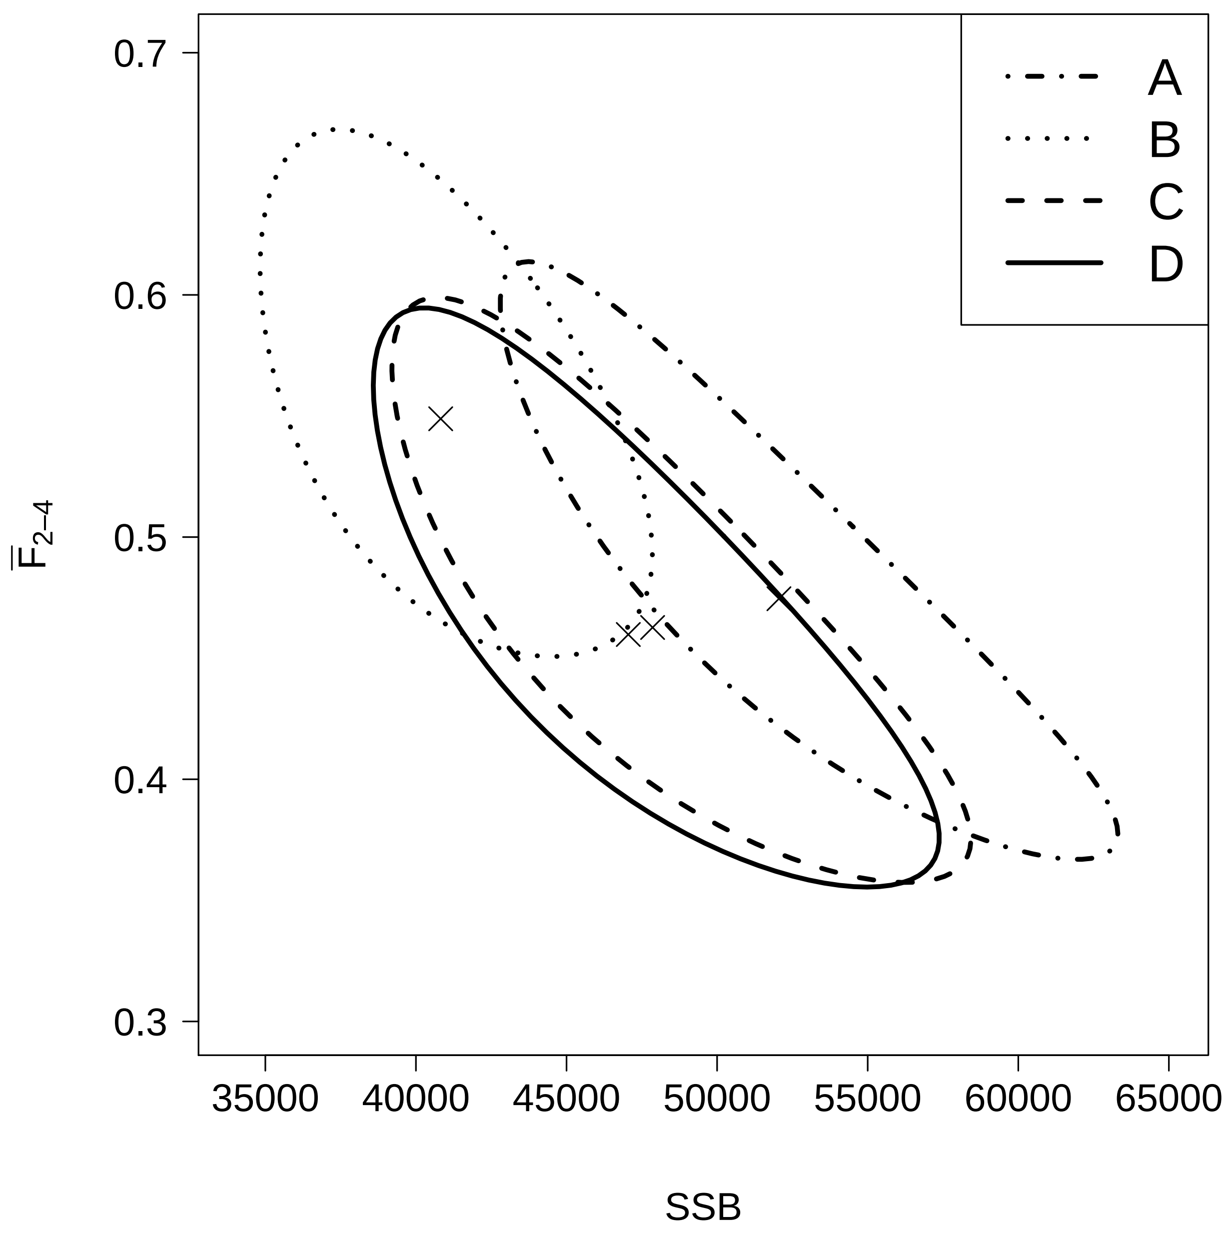

Figure 4: North Sea cod: 95% confidence region for the joint distribution of $\bar{F}$ and SSB in the last data year (2011) for each of the examined correlation structures (A–D).

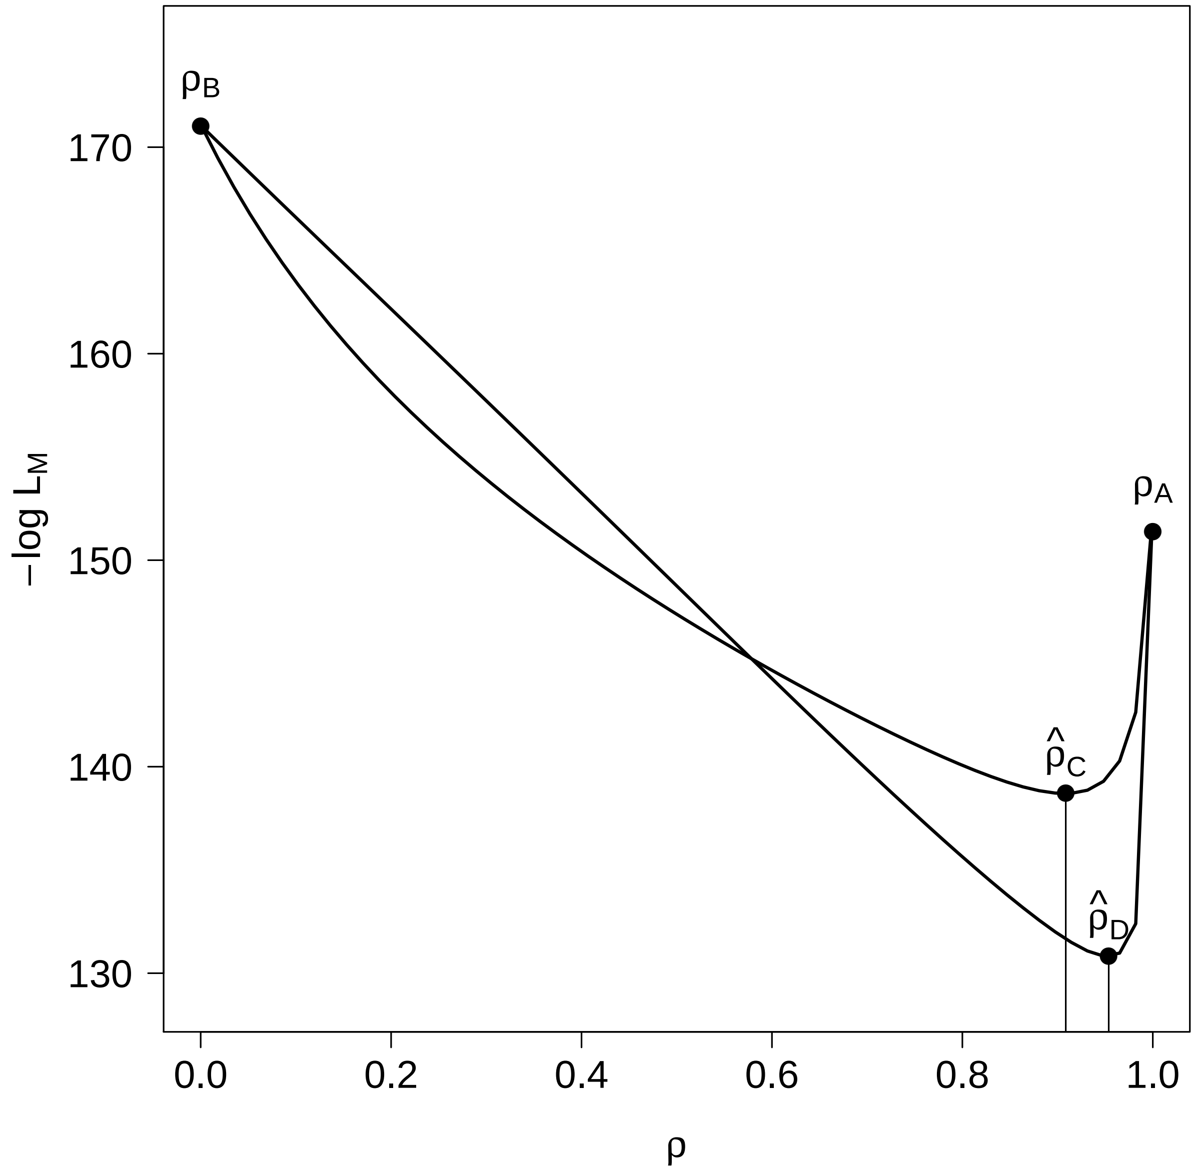

Figure 5: North Sea cod: Profile likelihood for the ρ-parameter for models C and D, $\rho = 1$ corresponds to model A, and $\rho = 0$ corresponds to model B.

An attractive feature of this model is that the degree of smoothness is estimated from data using an objective criterion (maximum likelihood) rather than being specified apriori or using more ad-hoc arguments. The North Sea Cod example illustrated that estimating some degree of correlation in the random walks led to substantial improvements in AIC as opposed to assuming either completely correlated or uncorrelated random walks for F. Since AIC is a measure of the one-step prediction error of the model, we can conclude that the forecasting power of the model has been substantially increased by this. The state-space formulation of the model has several additional benefits which include: estimation of observation error which allows for objective weighting of multiple data sources, quantification of uncertainties on all parameters of interest, and finally it provides a fast and flexible framework using very few model parameters compared to traditional SCAA models. As an example, Butterworth et al. (2003) used more than 500 parameters in a SCAA model with time-varying selectivity. The model was specified via penalized likelihood, which provides information and reduces the effective number of parameters. In contrast, the state-space model used a total 16 parameters in the North Sea Cod case, and the degree of smoothness was included as a model parameter to be estimated from data.

The state-space model formulation also enables the use of information criteria such as AIC for model selection as opposed to the penalized likelihood approach, which requires alternative and more computationally demanding techniques such as cross-validation (Maunder and Harley, 2011; Punt et al., 2013).

It should be noted that only correlations between age groups in the *process* equation for the F-random walks were considered, whereas zero correlation was assumed between ages in the *observation* equations, which can be substantial at least in the survey data for highly aggregated species like herring and whiting in the North Sea (Berg et al., 2013). More research should be made to investigate under what circumstances both types of correlation are present, can be estimated, and how they might possibly interact and affect the results.

5. Acknowledgements

The authors wish to thank Uffe H. Thygesen, Peter Lewy, Hans Skaug and an anonymous reviewer for their valuable input to this manuscript.

References

Aarts, G., Poos, J., 2009. Comprehensive discard reconstruction and abundance estimation using flexible selectivity functions. ICES Journal of Marine Science: Journal du Conseil 66, 763–771.

Berg, C.W., Nielsen, A., Kristensen, K., 2013. Evaluation of alternative age-based method for estimating relative abundance from survey data in relation to assessment models. Fisheries Research `http://dx.doi.org/10.1016/j.fishres.2013.10.005`.

Brinch, C.N., Eikeset, A.M., Stenseth, N.C., 2011. Maximum likelihood estimation in nonlinear structured fisheries models using survey and catch-at-age data. Canadian Journal of Fisheries and Aquatic Sciences 68, 1717–1731.

Butterworth, D., Ianelli, J., Hilborn, R., 2003. A statistical model for stock assessment of southern bluefin tuna with temporal changes in selectivity. African Journal of Marine Science 25, 331–361.

Fournier, D.A., Skaug, H.J., Ancheta, J., Sibert, J., Ianelli, J., Magnusson, A., Maunder, M.N., Nielsen, A., 2012. AD Model Builder: Using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. Optimization Methods and Software 27, 233–249.

Gudmundsson, Gunnlaugsson, 2012. Selection and estimation of sequential catch-at-age models. Canadian Journal of Fisheries and Aquatic Sciences 69, 1760.

Gudmundsson, G., 1994. Time series analysis of catch-at-age observations. Journal of the Royal Statistical Society: Series C (Applied Statistics) 43, 117–126.

He, X., Ralston, S., MacCall, A.D., 2011. Interactions of age-dependent mortality and selectivity functions in age-based stock assessment models. Fishery Bulletin 109, 198–216.

ICES, 2009a. Report of the Baltic Fisheries Assessment Working Group (WGBFAS). ICES Document CM 2009/ACOM:07 .

ICES, 2009b. Report of the Working Group on Methods of Fish Stock Assessment (WGMG). ICES CM 2009/RMC:12 .

ICES, 2011. Report of the Working Group on Multispecies Assessment Methods (WGSAM). ICES CM 2011/SSGSUE:10 .

ICES, 2012. Report of the Working Group on the Assessment of Demersal Stocks in the North Sea and Skagerrak (WGNSSK). ICES CM 2012/ACOM:13 .

Maunder, M.N., Harley, S.J., 2011. Using cross validation model selection to determine the shape of nonparametric selectivity curves in fisheries stock assessment models. Fisheries Research 110, 283–288.

Methot, R.D., Wetzel, C.R., 2013. Stock synthesis: A biological and statistical framework for fish stock assessment and fishery management. Fisheries Research 142, 86 – 99.

Punt, A.E., Hurtado-Ferro, F., Whitten, A.R., 2013. Model selection for selectivity in fisheries stock assessments. Fisheries Research .

Sampson, D.B., Scott, R.D., 2012. An exploration of the shapes and stability of populationselection curves. Fish and Fisheries 13, 89–104.

Shepherd, J., 1999. Extended survivors analysis: An improved method for the analysis of catch-at-age data and abundance indices. ICES Journal of Marine Science: Journal du Conseil 56, 584–591.

Skaug, H., Fournier, D., 2006. Automatic approximation of the marginal likelihood in non-gaussian hierarchical models. Comput. Stat. Data An. 51, 699–709.

Taylor, I.G., Methot, R.D., 2013. Hiding or dead? a computationally efficient model of selective fisheries mortality. Fisheries Research 142, 75 – 85.