Discussion paper

The Ensemble Kalman Filter in Bioeconomics

BY
Sturla F. Kvamsdal AND **Leif K. Sandal**

The Ensemble Kalman Filter in Bioeconomics

Sturla Furunes Kvamsdal, Leif Kristoffer Sandal

Norwegian School of Economics, Helleveien 30

N-5045 Bergen, Norway

sturla.kvamsdal@nhh.no

March 15, 2012

Abstract

We demonstrate the power of the Ensemble Kalman Filter in specifying ecosystem models ideal for bioeconomic analysis. Bioeconomic analysis requires models to be relatively simple, but models must still capture the nature and dynamics of the system. With the Ensemble Kalman Filter, we are able to capture complex dynamics in multispecies models with models simple enough to allow further bioeconomic analysis. While bioeconomic analysis has had limited influence on management decisions, the advent of new methods and the need for high dimensional ecosystem based management models may make bioeconomic research more relevant in the future. The filter is applied to a ecosystem model of the commercially most important species in the Barents Sea. The simpler, aggregated stochastic biomass models capture the complex dynamics of the pelagic stocks on the level needed for making decisions on for example allowable catches.

Keywords: Ensemble Kalman Filter, Ecosystem Management, Bioeconomics, Aggregated Biomass Models.

1 Introduction

Whilst traditional fisheries management has had limited success (Ludwig et al. 1993, Worm et al. 2006), interest in and need for ecosystem-based management of fisheries increases (Holland et al. 2010, Kaufman et al. 2004, May et al. 1979). Economists has spent considerable time and effort on studying efficiency and optimality of fisheries management and more generally renewable resource management models, but the so-called bioeconomics literature has had little impact on real-world fisheries management (Squires 2009). Perhaps the main reason for the lack of impact are the over-simplified biological models typically used. While simple models enhance tractability, the models simply cannot capture the observed dynamics of fish stocks. When it comes to ecosystem-based management, it is obvious that the staple, single-species model in bioeconomics has limited, if any, interest. As such, much of the work in population dynamics, which has had a much larger impact on policy (Wilen 2000), has also focused on single-species models. Thus, the management of most fisheries today is based upon single-species concepts. A case in point is the central position of the maximum sustainable yield concept in the Johannesburg Declaration on Sustainable Development (United Nations 2002). Maximum sustainable yield is a staple single-species concept which leads astray in an ecosystem setting (see Kaufman et al. 2004, p. 694, and references therein, see also Ludwig et al. 1993, p. 17, and May et al. 1979, p. 267). While population dynamics has been the main scientific influence on management decisions, one may ask whether the sole influence is warranted. We subscribe to the criticism raised by Hannesson (2007, p. 699), that 'age-structured models introduce idiosyncratic elements of uncertainty' through unknown parameters, and believe that the much more tractable aggregated biomass models are more relevant 'when they can be reconciled with reality.' Tractability becomes ever more important when the dimensionality of the problem increases. The aim of our present efforts is exactly to demonstrate how aggregated biomass models can be reconciled with the reality of marine ecosystems.

We use the ensemble Kalman filter (Evensen 1994, Burgers et al. 1998) to fit a marine ecosystem model to data. The ensemble Kalman filter is a data assimilation method much used in meteorology and oceanography; sciences which deal with large, high-dimensional, and chaotic systems. Evensen (2003) reviews both

theoretical developments and applications of the ensemble Kalman filter and related methods; Evensen (2009) covers more recent developments. The method can be seen as an extension of the classical Kalman filter to a large class of nonlinear models. The fundamental idea is to use a Markov Chain Monte Carlo approach to solve the Fokker-Planck (or Kolmogorov's) equation which governs the time evolution of the model. The model is written as a stochastic differential equation, and both the model and observations are assumed to contain noise. Importantly, the method facilitates simultaneous estimation of poorly known parameters (Evensen 2009, p. 101). With the ensemble Kalman filter, relatively simple models can capture much of the complexity observed in marine ecosystems. We review the theory of the ensemble Kalman filter and apply it to a three-species model of the Barents Sea ecosystem.

Variants of the Kalman filter has been applied to bioeconomic models earlier. Berck and Johns (1985) considered the extended Kalman filter to fit an aggregated biomass model to the Pacific Halibut fishery. The extended Kalman filter extends the classical filter (Kalman 1960) to nonlinear models by local linearizing of the model. The extended filter has limited applicability as it depends on closure assumptions and costly computations of the Jacobian (Burgers et al. 1998, p. 1719). While the analysis in Berck and Johns (1985) never saw the light of day in a peer-review publication, it was both original and inventive, among other things estimating the value of information about the Pacific Halibut stock (less than perfect information about the 1973 stock was estimated to 2% of the present value of the fishery, p. 18). They also underlined that their approach to estimating the stock level was 'radically different' than the methods used by the management body (the International Pacific Halibut Commission). In their words, the Kalman filter 'sacrifice biological detail while gaining statistical method' (Berck and Johns 1985, p. 19). The extended Kalman filter was later suggested for renewable resource management by Walters (1986). Marine scientists have on some occasions applied the extended Kalman filter to their models, see for example Schnute (1994), Peterman et al. (2000, 2003), and references therein.

More recently, several different data assimilation methods, usually so-called variational adjoint methods, have been suggested to fit aggregated biomass models to data (see Ussif et al. 2003, and references therein). On the other hand, Grønnevik and Evensen (2001) applied different ensemble-based data assimila-

tion techniques to age-structured fish stock assessment models; among them, the ensemble Kalman filter. An advantage of the ensemble Kalman filter when compared to variational adjoint methods is that it does not rely on direct optimization, and all observations are not processed simultaneously. Instead, variable and parameter estimates are updated sequentially according to the filtering procedure. The ensemble Kalman filter also facilitates so-called flow-dependent noise attribution; flow-dependent (or rather, state-dependent) noise processes, it turns out, are fundamental in capturing the dynamics of marine ecosystems.

If, as in Ussif et al. (2003), there is a known or easily identified functional relationship between biological variables and the exploitation strategy, the filter can also estimate economic parameters (the exploitation rate). Similarly, the filter applies to a number of related problems, not only in bioeconomics, but in economics more generally.

The ensemble Kalman filter fits, in an efficient manner, nonlinear, aggregated biomass, ecosystem models to data. It also estimate the model error, which can be translated into uncertainty in model predictions. Combined with developments in high-dimensional, stochastic optimization, we believe the filter can make bioeconomic analysis relevant for real-world fisheries management decisions. The main criticism, over-simplified biological models, looses much of its force when the explanatory power of the fitted biomass models matches, and even competes with, that of age-structured models. The potential of the ensemble Kalman filter reaches further. It has the ability to process large amounts of data in high-dimensional systems with large numbers of poorly known parameters (see Evensen 2003, and references therein) and it should be of interest to any researcher working with large and volatile systems; from macroeconomics to population dynamics and beyond.

2 Theory

Our theoretical presentation of the ensemble Kalman filter is inspired by Evensen (2003, 2009). We depart from the continuous time state-space model:

$$\mathrm{d}x = f(x)\mathrm{d}t + \sigma\mathrm{d}B \tag{1}$$
$$d = M(x) + v \tag{2}$$

3

An incremental change dx in the state variable (or vector) x is the sum of the drift term $f(x)dt$ and the stochastic diffusion term σdB. When x is an aggregated biomass vector, $f(x)$ is the multi-dimensional growth function ($f : \mathbb{R}^n \to \mathbb{R}^n$). The stochastic increments dB are independent, identical, and normal distributed with mean zero and variance dt. The measurement functional $M(x)$ relates the state vector to the observations d. When the state vector is directly observed, the measurement functional is the identity operator. v is a normal distributed error term with mean zero and covariance R. Equation (1) is called the state equation; equation (2) is called the measurement or observation equation.

The ensemble Kalman filter is a sequential filter method and works as follows. The model is integrated forward in time until measurements become available. Measurements are used to update the model. The updated model is then further integrated until the next measurement time. In the theoretical literature, the update step is called the analysis, thus the notation x^a for the updated state vector. The forward integrated model (the forecast) is denoted x^f. C^f is the covariance of the model forecast; C^a is the covariance of the model analysis.

The ensemble Kalman filter uses, as the name suggests, an ensemble of model states; a cloud of points in the state-space, to represent the probability density function at any given time. With a so-called Markov Chain Monte Carlo method (meaning that the model can be formulated as a Markov Chain and that a large number of simulated solutions are considered, see Evensen 2009), each ensemble member is integrated forward in time according to (1). Errors are simulated. The integrated ensemble represents a forecast of the probability density and the only approximation is the limited number of ensemble members (Evensen 2009, p. 47). The Markov Chain Monte Carlo method is the backbone of the ensemble Kalman filter and is equivalent to solving the Fokker-Planck equation for the time evolution of the probability density; see Evensen (2003, p. 348) for further details.

When measurements are available, each ensemble member is updated as a linear weighting between the forecast and the measurements:

$$x^a = x^f + K(d - Mx^f) \tag{3}$$

The weight K is called the Kalman gain. Assuming M is the identity operator, we see that with $K = 0$, no weight is put on the observation d; with $K = 1$, no

weight is put on the forecast x^f. The Kalman gain is given by:

$$K = C^f M^T (M C^f M^T + R)^{-1} \tag{4}$$

where we assume that M is a linear operator (a matrix); M^T denotes its transpose. It is crucial that observations are treated as uncertain ($R > 0$), and in the ensemble Kalman filter, the observation probability density is represented by an ensemble; observations are perturbed (Burgers et al. 1998, pp. 1720-1721). It is convenient to let the number of ensemble members in the state-space ensemble, denoted X, and in the observation ensemble, denoted D be equal.

In the standard Kalman filter, both the forecast and analysis covariance (C^f and C^a) are in principle unknown; they are defined in terms of the unknown true state (see Evensen 2003, p. 347). In the ensemble Kalman filter, they are defined in terms of the ensemble mean ($\mathbf{E}$ denotes the mean or expected value):

$$C_e^f = \mathbf{E}\left[(X^f - \mathbf{E}[X^f])(X^f - \mathbf{E}[X^f])^T\right] \tag{5}$$

$$C_e^a = \mathbf{E}\left[(X^a - \mathbf{E}[X^a])(X^a - \mathbf{E}[X^a])^T\right] \tag{6}$$

That is, covariances are represented by the ensemble moments which carry the subscript e. The observation covariance is also represented by the ensemble moment:

$$R_e = \mathbf{E}\left[(D - d)(D - d)^T\right] \tag{7}$$

The observation ensemble is defined such that it has the true (given) observation as its mean: $\mathbf{E}[D] = d$. The ensemble Kalman gain is defined as

$$K_e = C_e^f M^T (M C_e^f M^T + R_e)^{-1} \tag{8}$$

We assume that the ensemble is of sufficient size, such that $M C_e^f M^T$ and R_e are nonsingular; see Evensen (2003, p. 349). The analysis step (3) is carried out element by element:

$$X_j^a = X_j^f + K_e(D_j - M X_j^f) \tag{9}$$

where the subscript denote ensemble member j. It can be shown that by updating the ensemble with the perturbed observations D, the updated ensemble X^a has the correct error statistics (Evensen 2003, p. 349). The analysis covariance can be written as

$$C_e^a = (\mathbf{I} - K_e M) C_e^f \tag{10}$$

which is equivalent to the standard Kalman filter expression. $\mathbf{I}$ denotes the identity operator. Please see Evensen (2003) for derivations and further discussion.

The filter can estimate parameters simply by adding them to the state-space, and thus, the state vector. Parameters are treated as model states, and state variables and parameters are estimated simultaneously. As Evensen (2009, pp. 95-97) points out, the approach represents an improvement to more traditional approaches which ignore model error. With parameters in the state space, it is straightforward to adapt involved operators to make them compatible with the extended state vector. For example, in a system with m observations and n states, the measurement operator is $M : \mathbb{R}^n \to \mathbb{R}^m$. With p parameters to be estimated in the system, the measurement operator must be $M : \mathbb{R}^{n+p} \to \mathbb{R}^m$. The Kalman gain must be $K : \mathbb{R}^m \to \mathbb{R}^{n+p}$. Note that the estimate is the ensemble mean. Interpreting the ensemble as representative of the probability density function, the filter parameter estimate is conceptually an alternative to a maximum likelihood estimate; the maximum likelihood estimate is at the mode of the probability density. Hansen and Penland (2007) suggest an alternative interpretation of obtained parameter values, but the approach has only been demonstrated appropriate for high-frequency observations and with a near-perfect model where parameters are unknown constant. The alternative interpretation of parameter estimates could potentially improve model forecasts, but that is auxiliary to our present focus.

The initial ensemble should reflect belief about the initial state of the system (Evensen 2003, p. 350). The filter can be initialized by specifying means and standard deviations which characterize the initial ensemble. In the case of unknown parameters, initialization can be challenging. Our experience is that with large enough standard deviations, such that the initial ensemble cover all eventualities, and enough ensemble members, it is possible to find reasonable traits of the initial ensemble. Often, there is theory and earlier results to rely on. The size of the ensemble is another concern, mostly computational in nature. The ensemble should be large enough to sufficiently represent the probability density function. What is sufficient depends on the dimensionality of the problem, but also on the nature of the data; data with little variation could perhaps require more ensemble members. What is a possible ensemble size depends on computational limitations. An ensemble size of one hundred is not uncommon in the oceanographic and meteorological sciences, sciences where the application

6

of the ensemble Kalman filter to high-dimensional systems with large numbers of unknowns has been successful (Evensen 2009).

To sum up, the ensemble Kalman filter can be interpreted as a statistical Monte Carlo method where the ensemble evolves in state-space with the mean as the best estimate and the spreading of the ensemble as the error variance (Burgers et al. 1998, p. 1720). In many problems, the sequential processing of observations proves to be a better approach than the simultaneous processing which is typical in variational methods (Evensen 2009, p. 101).

3   The Model

The Barents Sea is one of the most productive ocean areas in the world, and is subject to extensive research (Gjøsæter et al. 2009, Huse et al. 2004, O'Brien et al. 2004, see also further references therein). The commercially most important stocks are cod and capelin; cod is highly valued as human food and capelin is an important part of the cod diet. Capelin is also caught for fish meal and oil production. Juvenile herring enters the Barents Sea when large year-classes arise in the Norwegian Sea. Herring has an important influence on the ecosystem; it is preyed on by cod while it preys on capelin larvae. We limit our model to these three fish stocks, for two main reasons. First, we are able to capture the dynamics of the cod stock to a high degree, and the cod fishery is the most important fishery in the region and of our main interest. Second, if the model is to be relevant for bioeconomic analysis, we have to limit the complexity and dimensionality of the model. We have in mind the type of analysis carried out in Sandal and Steinshamn (2010); see also Kugarajh et al. (2006). To limit complexity, we use simple growth functions and interaction terms common in traditional bioeconomic analysis. While dimensionality is based upon technical limitations, we find comfort in the view promoted by Holling and Meffe (1996, p. 333), that the driving forces of an ecosystem are confined to a relatively small subset of variables and relationships. While it is highly improbable that our choice of variables and relationships contain all driving forces of the Barents Sea ecosystem, we observe that our model captures much of the variation detected in stock assessments.

3.1 The State-Space Model

To begin with notation, the biomass of the three stocks are the state variables; cod is denoted x_1, capelin is denoted x_2, and herring is denoted x_3. Both cod and capelin are harvested in the Barents Sea; h_1 and h_2 denote fishing mortality on cod and capelin. Herring is not harvested in the Barents Sea, but it flows in from the Norwegian Sea. We denote the herring inflow by i_3. Finally, we denote parameters c_i. The dynamic model for the system is written:

$$\mathrm{d}x_1 = \left(c_1 x_1(1 - x_1/c_2) + c_3 c_7 x_1 x_2 + c_4 c_{12} x_1 x_3 - h_1\right)\mathrm{d}t + \sigma_1(x)\mathrm{d}B_1 \quad (11)$$

$$\mathrm{d}x_2 = \left(c_5 x_2^2(1 - x_2/c_6) - c_7 x_1 x_2 - c_8 x_2 x_3 - h_2\right)\mathrm{d}t + \sigma_2(x)\mathrm{d}B_2 \quad (12)$$

$$\mathrm{d}x_3 = \left(c_9 x_3^2(1 - x_3/c_{10}) + c_{11} c_8 x_2 x_3 - c_{12} x_1 x_3 + c_{13} i_3\right)\mathrm{d}t + \sigma_3(x)\mathrm{d}B_3 \quad (13)$$

The stochastic increments $\mathrm{d}B_i$ are independent, with mean zero and variance $\mathrm{d}t$. Correlations in the noise processes are reflected in the scaling term $\sigma_i(x)$. We use two principal models of the scaling term; white noise $(\sigma_i(x) = \sigma_i)$ and geometric, white noise $(\sigma_i(x) = \sigma_i \cdot x)$.

The first terms (the c_1-, c_5-, and c_9-terms) in each model equation is equivalent to the logistic (cod) or modified logistic (capelin and herring) growth function, and the parameters c_1, c_2, c_5, c_6, c_9, and c_{10} are interpreted accordingly. (The idea of carrying capacity; the standard interpretation of the second parameter in the logistic and modified logistic, becomes unclear in an ecosystem setting. The capacity of the ecosystem to harbor any one specie depends on the state of the entire system. Hence, intrinsic, single species notions such as carrying capacity must be treated with caution in all multispecies approaches.)

Parameters c_7, c_8, and c_{12} are the basic species interaction parameters. c_7 measures the intensity of the capelin-cod interaction; c_8 the capelin-herring interaction, and c_{12} the herring-cod interaction. Parameters c_3, c_4, and c_{11} adjusts the mirror interactions (cod-capelin, cod-herring, and herring-capelin), as biomass is not conserved in the interactions. All parameters are lognormal distributed, and are thus always positive. (Theoretically, they are modeled as $c_i = \exp(\alpha_i)$, where each α_i is a stochastic constant which has white noise.) The signs of the interaction terms thus becomes significant and a matter of modeling; cod benefits from interacting both with capelin and herring (positive sign on interaction terms), and herring benefit from interacting with capelin. Mirror interaction terms have opposite signs; both species cannot benefit from the interaction. (The interactions are predator-prey interactions, where cod is

the top predator, while caplin is preyed upon by both cod and herring. Symbiotic, mutually beneficial interactions between the three species are not likely and not possible in our model. Neither are competitive, mutually destructive interactions.) The multiplicative interaction model is a crude Lotka-Volterra form of predator-prey interaction (May et al. 1979, p. 268); it is also referred to as mass-action (Tschirhart 2009).

The final parameter c_{13} measures the influence of the inflow of herring on the herring stock growth. Most of the time, the amount of herring biomass which enters the Barents Sea is relatively small. After a few years, however, the herring has grown substantially. Thus, we lag the inflow variable two years and multiply it with the scaling parameter c_{13}. The idea is that three year old (and older) herring makes out most of the herring biomass in the Barents Sea, and the biomass influx two years earlier better explains the change in the herring stock. (After three or four years in the Barents Sea, the juvenile herring returns to its main habitat in the Norwegian Sea to mature and eventually spawn.)

The adjustment parameters c_3, c_4, and c_{11} can be interpreted as biomass conversion rates between species. Presumably, regularities exists between the cod-capelin interaction term and the capelin-cod interaction term, and between all other mirror interaction terms. While known or assumed interaction relationships would be helpful in reducing the number of parameters in the system, biologists are sceptical when it comes to the stability of the relationships (Tjelmeland 2011). Thus, we refrain from prescribing fixed relationships. (The alternative would be, if one assumed that, say, 90% of the biomass was lost between two trophic levels, to specify the corresponding mirror interaction parameter $c_i = 0.1$.)

We treat the stock level estimates as direct measurements of the state variables, and the measurement operator is thus the identity operator. Note that parameters are added to the state vector as described above. The measurement operator must thus be adjusted to be compatible with the state vector. The operator is adjusted by adding zeros; parameters are unobserved. The observation equation becomes

$$d = Mx + v \qquad (14)$$

where

$$x = \begin{bmatrix} x_i \\ c_j \end{bmatrix}, \quad i = [1, \ldots, 3], \ j = [1, \ldots, 13], \quad \text{and} \qquad M = [\mathbf{I} \ \mathbf{0}]$$

where $\mathbf{I}$ is a three by three identity matrix and $\mathbf{0}$ is a three by thirteen zero matrix. d is a three-element vector of observations, and v is the error term which is normal, independent, and identically distributed with mean zero and known or assumed variance $\mathbf{R}$.

3.2 Data

The fish stocks in the Barents Sea cannot be observed directly. However, the Institute of Marine Research in Bergen (in collaboration with the Knipovich Polar Research Institute of Marine Fisheries and Oceanography in Murmansk) carries out extensive, yearly ecosystem surveys. Based upon these surveys, they provide yearly estimates of the stock levels of all the important species in the Barents Sea. The stock estimates are published by the International Council for the Exploration of the Sea (ICES), and most of our data are collected from the ICES online database. Uncertainty in stock assessments are, however, usually not reported, and we are left to speculation.

We have stock estimates, catch data and herring inflow estimates from 1950 to 2007. However, the ICES database does not contain data on capelin prior to 1972. For the period, we collected catch data from Røttingen and Tjelmeland (2008, see Figure 2). Capelin stock estimates were collected from Marshall et al. (2000, see Figure 1, p. 2435). The early capelin stock estimates are more uncertain than later estimates, and we assume a 50% increased observation uncertainty on the capelin stock data prior to 1972.

3.3 The Initial Ensemble

The initial ensemble is drawn randomly from a normal distribution. For the three state variables, we use the first observations as the mean of the initial ensemble and 30% of the first observation as standard deviation. For the parameter variables, we use a combination of theory, common sense, and existing estimates to characterize the initial ensemble. As parameters enter the model equations as $c_i = \exp(\alpha_i)$, the parameter variable ensembles are defined in terms of the α_i's, which may be called shadow parameters. Means and variances for the shadow parameter variable ensembles are listed in Table 1. The table also lists the implied parameter mean $\exp(\bar{\alpha}_i)$. Since it is intuitively much easier to relate to the parameters c_i rather than the shadow parameters α_i, we refer to

the parameters in the discussions that follow.

Initial means for the logistic parameters for cod, c_1 and c_2, were set to 1 and $5 \cdot 10^6$. The values were inspired by earlier estimates (Kugarajh et al. 2006, p. 75, Table 3) and, when it comes to the capacity parameter c_2, the range of observed stock levels. For the pelagic stocks capelin and herring, we used the modified logistic growth function; see equations (12) and (13). The growth parameter in the modified logistic operates on a different scale than the growth parameter in the pure logistic, and initial means for c_5 and c_6 were set at $10 \cdot 10^{-6}$ and $5 \cdot 10^{-6}$. It is not uncommon with a significantly larger intrinsic growth rate of prey species in many predator-prey relationships (May et al. 1979, p. 269). As with cod, the initial means of the capacity parameters c_6 and c_{10} were inspired by the range of observed levels and set to $2 \cdot 10^7$ and $1 \cdot 10^7$.

The initial ensemble for the basic interaction parameters c_7, c_8, and c_{12} were treated more rigorously. The term $c_7 x_1 x_2$ in (12) reflects the loss of capelin biomass from the interaction with cod. Gjøsæter et al. (2009, see Figure 5, p. 45) estimated the amount of capelin consumed by the Barents Sea cod for the years 1984-2006 from stomach content data. The consumption varies over time, as does the cod and capelin stock levels. To get a reasonable initial measure of c_7, we regressed the total consumption of capelin on the term $x_1 x_2$. Notably, Gjøsæter et al. (2009) provided us with consumption estimates for 1984-2007 (that is, one more year of data than what they based their original analysis upon). The estimated coefficient was $3.46 \cdot 10^{-10}$ (standard error $5.1 \cdot 10^{-11}$, R^2_{adj} 0.63). Similar data for the capelin-herring interaction are not available. Herring is however thought to have a smaller predation rate on capelin than cod; we set the implied mean for c_8 at 10% of the implied mean of c_7. For the herring-cod interaction parameter c_{12}, data are available: Gjøsæter et al. (2009) estimated the amount of herring consumed by the Barents Sea cod. Regressing the herring consumption on the term $x_1 x_3$ yielded a coefficient of $2.49 \cdot 10^{-11}$ (standard error $3.76 \cdot 10^{-12}$, R^2_{adj} 0.61). As with c_7, we set the mean of the initial shadow parameter (α_{12}) ensemble to correspond to the estimated coefficient.

The mirror interaction parameters c_3, c_4, and c_{11} cannot be larger than one as it is assumed that some biomass is lost in the interactions. The biomass loss assumption is not explicitly enforced, but initial implied ensemble means for the three mirror interaction parameters were set to 0.25 for c_3 and 0.1 for c_4 and c_{11}. Typically, one assumes that 90% of the biomass is lost between

11

trophic levels, but cod spends less energy catching capelin and thus we specified a higher mirror interaction parameter for the cod-capelin interaction.

3.4 Practical Details

Some care must be taken when working with stochastic differential equations. We have formulated the model in continuous time, but it is necessary to discretize the equations for the numerical analysis, and in particular to produce the forecast. We us Itô algebra and end up with the following forecast equation:

$$x^{t+1} = x^t + f(x^t)\Delta t + \sqrt{\Delta t} \cdot \sigma(x^t) \cdot w^t \tag{15}$$

where the superscript is a time index, Δt is the discrete time increment, and w^t is a simulated, normal distributed error with zero mean and unit variance. The term $\sqrt{\Delta t}$ conserves the properties of the stochastic process. The term $\sigma(x^t)$ scales the noise process and retains the covariance structure. With a simple, white noise model, $\sigma(x^t)$ is the unique, upper-triangular Cholesky matrix of C_e^a, see equation (6). With a geometric, white noise model, $\sigma(x^t)$ is the Cholesky matrix of C_e^a multiplied with x^t/x^a. (Note that the Cholesky matrix of C_e^a can be written as $\Sigma_e^a x^a$, where Σ_e^a is an upper triangular matrix of coefficients.) The time unit is one year (the same as the observation frequency), and $\Delta t = 1/12$.

We have catch or landings data entering our equations as so-called control variables. We have ample reasons to believe that registered landings are not perfect observations of fishing mortality because of discarding at sea and illegal landings. Thus, we treat the landings data as uncertain and represent them with an uniformly distributed ensemble. The actual observation serves as the lower limit as the official landings certainly is a conservative estimate of fishing mortality, while the upper limit is set 20% higher. In the herring equation (13), landings do not enter. Instead, we have inflow data. The inflow data are estimates based upon virtual population models for the herring stock in the Norwegian Sea, which is coupled with an ocean circulation model which predicts the drift of eggs and larvae into the Barents Sea. While the inflow estimates probably are quite uncertain, we have no reason to believe they are neither upward or downward biased. Thus, we represent them with an ensemble which is normal distributed, with mean at the reported inflow and a 5% standard deviation. (Alternatively, it is possible to not use an ensemble for control variables. The implicit assumption is then that the controls are perfectly observed.)

Table 1: Initial characterization of the initial parameter ensemble. Column 1 (Parameter) identifies the parameters. Column 2 (Ensemble Mean) reports the mean of the ensemble on the underlying parameter α_i. Column 3 (Ensemble Std. Dev.) reports the standard deviation of the α_i ensemble. Column 4 (Implied Mean) reports the initial mean implied on the parameter c_i by the assumed mean on the initial α_i ensemble. The table is grouped with regard to which equation the parameters belong to.

Parameter	Ensemble Mean ($\bar{\alpha}_i$)	Ensemble Std. Dev.	Implied Mean ($\exp(\bar{\alpha}_i)$)
c_1	0.0	1.0	1.0
c_2	15.4249	0.2	$5.0 \cdot 10^6$
c_3	-1.3863	0.5	0.25
c_4	-2.3026	0.5	0.10
c_5	-11.5129	1.0	$1.0 \cdot 10^{-5}$
c_6	16.8112	0.3	$2.0 \cdot 10^7$
c_7	-21.7846	2.0	$3.46 \cdot 10^{-10}$
c_8	-24.0872	1.5	$3.46 \cdot 10^{-11}$
c_9	-12.2061	1.0	$5.0 \cdot 10^{-6}$
c_{10}	16.1181	0.3	$1.0 \cdot 10^7$
c_{11}	-2.3026	0.5	0.1
c_{12}	-24.4162	1.0	$2.49 \cdot 10^{-11}$
c_{13}	2.3026	1.0	10.0

Stock observations are also estimates derived from virtual population models and are uncertain. It is crucial that observations on state variables are represented with an ensemble (Burgers et al. 1998). The stock observation ensemble is normal distributed, with the observation at the mean and a standard deviation of 30%. (Because the capelin stock estimates prior to 1972 are more uncertain, the standard deviation in the capelin observation ensemble is increased with 50%.)

Finally, we use an ensemble size of 200, which currently is at the limit of what our computer system can handle. In comparison, ensemble sizes of 100 or less is not uncommon in problems of larger dimensions than ours (see Evensen 2009, for examples).

4 Results

Figure 1 shows the estimated stock levels (solid lines) for all three species cod (top panel: State 1), capelin (middle panel: State 2), and herring (bottom panel: State 3). The figure also shows observed stock levels (x-marks) which almost always lie within the 95 % confidence intervals (shaded areas). Most observations also lie within the 50 % confidence intervals (dark shaded areas); model uncertainty could be overstated. The general impression is that the model captures much of the dynamics in the ecosystem. In particular, the rather violent dynamics in the pelagic stocks capelin and herring is to a large degree picked up.

If we look closer at the cod stock estimates (Figure 1, top panel), there seems to be a pattern in that when the stock level is high or increasing, the estimate is below the observations. When the stock level is low or decreasing, the estimate is typically above the observations. The pattern could suggest that some of the dynamics is lost. The pattern could also be interpreted as a delay; it takes some time before the system adapts to new information. Some delay is inevitable in multidimensional systems, as the estimate at any given time depends on the forecast (old observations processed through the model equations with adapted parameters), the Kalman gain (covariances of the forecast and the observations; thus not directly depend upon the latest observation), and the latest observation; see equations (8) and (9). When the system is truly stochastic or if the model is imperfect (in our case, both), new observations

14

will always contain some new information which is not reflected in the forecast model. The estimate is a convex combination of the forecast and the observation; again, see equation (9), and the estimate will always lie somewhere between the forecast and the observation.

The dynamics of the pelagic stocks (capelin: Figure 1, middle panel; herring: Figure 1, bottom panel) is also largely captured by the model. For capelin, the confidence intervals are significantly larger prior to 1972, when measurements were more uncertain. The difference is visible for the naked eye when one compares, for example, the intervals in 1970 and 1980; years with comparable stock level measurements. After 1972, the capelin stock estimates, in addition to being more precise, lie closer to the measurements. The increased fit with the measurements follows from the definition of the Kalman gain (8): with less uncertain measurements (a smaller R_e), the gain increases and measurements gains more weight in the estimates (9). For herring, the model fits well to all observations up until the most recent observations. There are a couple of possible explanations for the relative poor fit for late herring observations. One is that fundamental changes in the Barents Sea ecosystem could be under way, and as discussed earlier, the model takes some time to adapt to new information. Another potential explanation is that late measurements are under review; our measurements results from virtual population models, and recent stock estimates are often reviewed when more raw data are fed to the virtual population models and upon further study of both biology and technology.

Figure 1 shows that there is correlation between stock levels and confidence intervals. Level dependent noise is built in both the model; see equations (11 - 13), and the measurement uncertainty. The noise is modeled as $\sigma \cdot x$, where σ is an upper triangular matrix of coefficients. The matrix of coefficients will depend upon time in the same manner as state and parameter estimates; the final estimate is:

$$\sigma = \begin{bmatrix} 0.2595 & 0.0180 & -0.0915 \\ 0 & 0.3080 & 0.0492 \\ 0 & 0 & 0.2682 \end{bmatrix} \tag{16}$$

Relative to the diagonal terms, the off-diagonal terms, which measure covariations in the noise structure, are small, but the covariation between cod and herring (-0.09) is probably too large to be ignored. We also estimated the model with pure white noise and obtained qualitatively similar results. Clearly, a geo-

metric noise model is better suited to model the covariance behavior. Thus, with the white noise model, it would require a post-analysis to obtain the relevant σ-matrix.

Figure 2 shows the parameter estimates (solid lines) with 95 % confidence intervals (shaded areas) for the model in equations (11 - 13). All parameters are stable over most of the time series, and most have a clear tendency of a decreasing spread. In particular, the growth parameters (c_1, c_2, c_5, c_6, c_9, and c_{10}) all show stability and ensemble convergence. The interaction adjustment parameters (c_3, c_4, and c_{11}) are stable, but convergence is not particularly pronounced. One reason could be that it simply is difficult to tease out the interaction adjustment parameters from the data. Another potential reason is that the interaction model we use is not perfect. In that regard, our estimates are the best possible with the given structure. While a more sophisticated and complicated interaction model could improve parameter estimates, it would not necessarily be ideal as basis for adaptive, ecosystem management. Our aim is to strike a balance between enough complexity to capture observed dynamics and enough simplicity to allow dynamic, stochastic optimization. But the evidence on the interaction model is mixed, as the basic interaction parameter estimates are promising. The basic interaction parameters (c_7, c_8, and c_{12}) are less stable than for example the growth parameters, but are still not unstable, and convergence is present. (In Figure 2, the time profile of the basic interaction parameter estimates cannot be detected because of the relatively wide confidence intervals. Figure 3 shows only the time profiles of the central estimates.) The final parameter, c_{13}, basically measures how much the inflow of herring larvae biomass from two years earlier has grown. The estimate is stable around 10 and the confidence interval steadily decreases throughout the time series.

Holling and Meffe (1996, p. 334) highlighted the problem of identifying stable parameters in ecosystem models as a major challenge to establish conceptually sound management regimes. Our approach identifies relatively stable parameters, and the key to the challenge is to allow for model error, as is done with the ensemble Kalman filter in particular, and in all state-space approaches in general.

In the ensemble Kalman filter literature (see reviews in Evensen 2003, 2009), a traditional tool to evaluate filter and model performance is to inspect plots of root mean squared innovations and root mean squared errors. Innovations
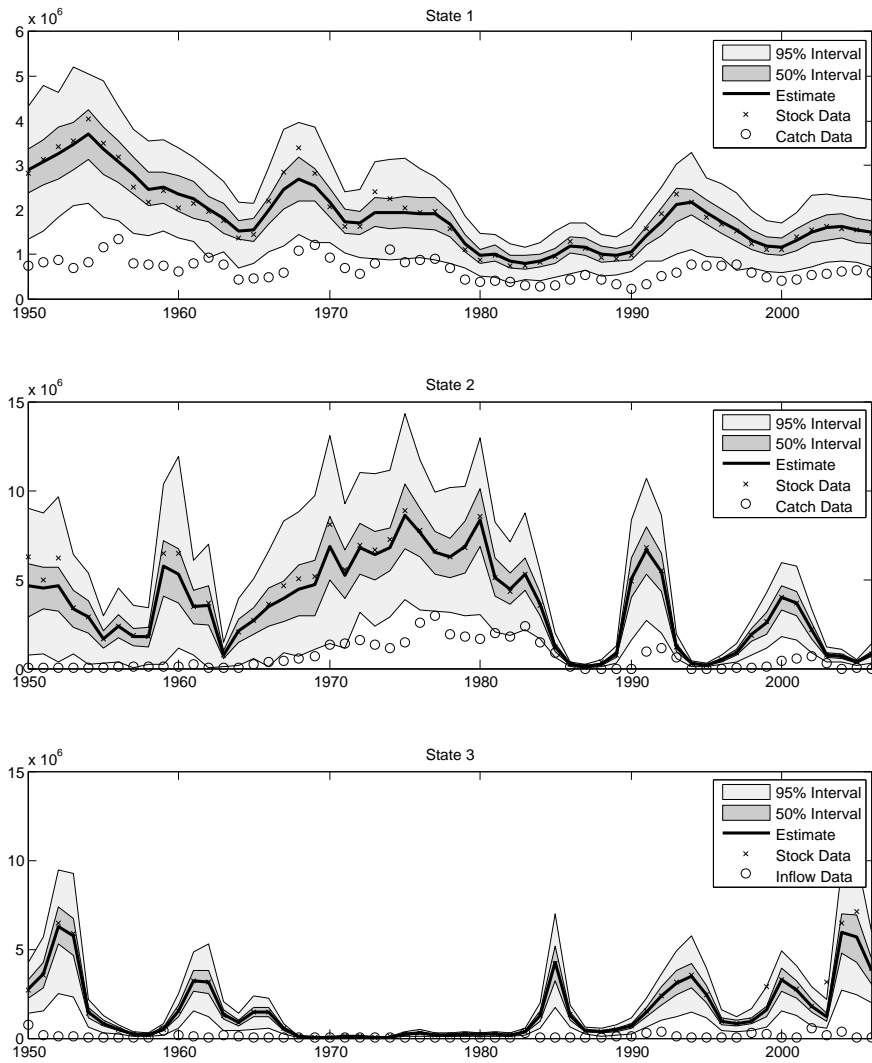
Figure 1: Estimated stock levels (solid line) with 95 % confidence intervals (shaded area), 50 % confidence intervals (dark shaded area), stock observations (x-marks), catch and inflow data (circles).
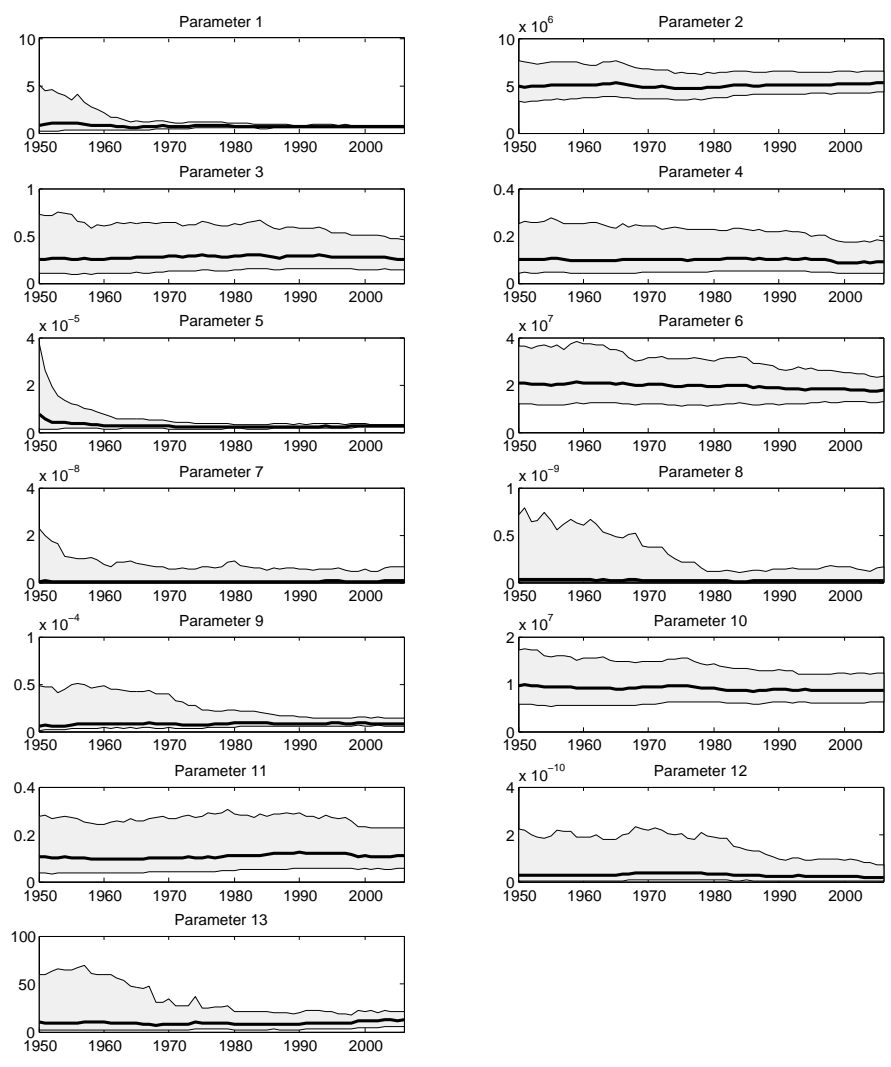
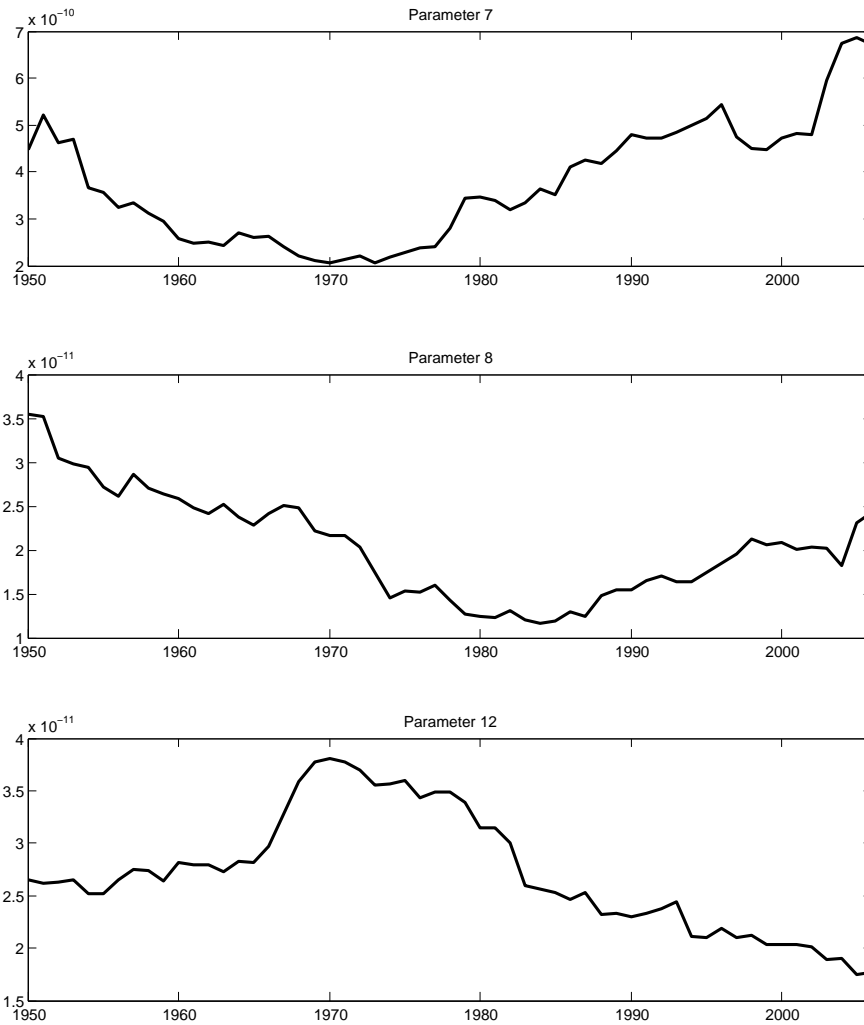Figure 2: Estimated parameter values (solid lines) with 95 % confidence intervals (shaded areas).

Figure 3: Basic interaction parameters estimates. See Figure 2 for confidence intervals.

are the corrections or updates which occur in the analysis step; see equation (9). Errors are also called ensemble anomalies, and are the mean difference of the updated ensemble and the estimate (the ensemble mean). Thus, the root mean squared errors are given by $\mathbf{E}\left[\mathbf{E}[X^a] - X^a\right]$. The absolute ensemble mean innovation should in theory decline over the time series and eventually stabilize. A lack of decline in the innovations can be a general sign of model misspecification, filter divergence, too little measurement error, or other, more esoteric problems (see Anderson and Anderson 1999). Since our parameters are modeled as $c_i = \exp(\alpha_i)$, where α_i is represented by an ensemble, the confidence intervals of the c_i may not give the correct impression of wether an estimate improves over the time series or not (wether the ensemble converges or not). In comparison, errors decline by construction. Too much error decline is not healthy, however. If the root mean squared error goes to zero, the filter exhibits ensemble collapse and filter divergence.

Figure 4 shows the root mean squared innovations (solid lines) and root mean squared errors (dashed lines) of the α_i ensembles. To make comparisons easier, all quantities are scaled such that 1 is the maximum of each serie . All parameters show ensemble convergence and there are no signs of filter divergence. The growth parameters (parameters 1, 2, 5, 6, 9, and 10) all have errors close to or below half of the maximum at the end of the time series. Parameter 6, the capacity parameter in the capelin equation, has some erratic behavior in the innovations late in the series, but the average statistic has a downward trend. For the interaction parameters, the ensemble convergence is less pronounced and innovations are more erratic with no clear trends. But, at least innovations are stable (no increasing trends) for all parameters. Parameters 11 and 12 demonstrate the potential problem of only inspecting the plots of confidence intervals in Figure 2: While the confidence intervals of c_{12} clearly have stronger convergence, the innovations and errors plots show that the parameters have similar ensemble convergence, and both have erratic innovations. The inflow scaling parameter in the herring equation, parameter 13, shows satisfactorily ensemble convergence with errors converging to less than half of the maximum.

From Figure 4, we can conclude that the distributions of most parameters have not converged to the true distributions; the root mean squared errors (dashed lines) have not stabilized. Thus, the model probably overstates the uncertainty in the system. While overstated uncertainty not necessarily is prob-
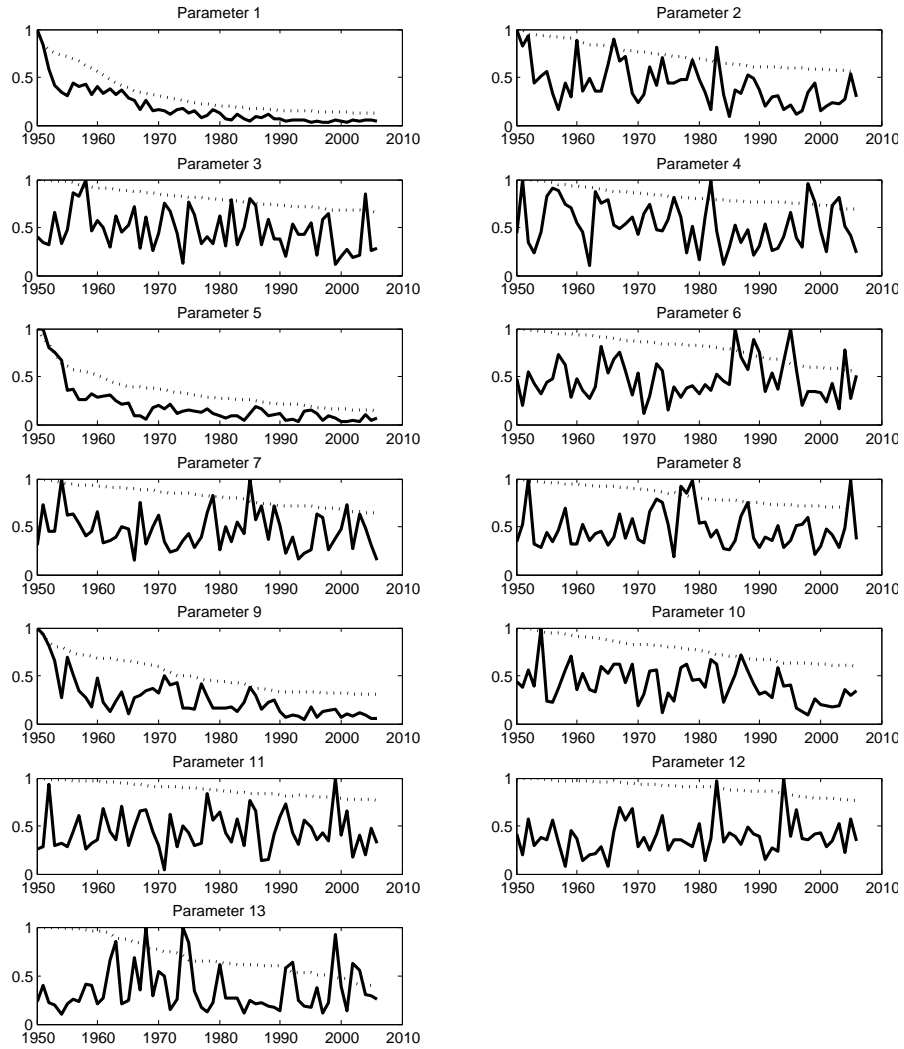
Figure 4: Root mean squared innovations (solid lines; $\mathbf{E}\left[X^a - X^f\right]$, see equation (3)) and root mean squared errors (dashed lines; $\mathbf{E}\left[\mathbf{E}[X^a] - X^a\right]$), for all parameters.

lematic in a management model, it explains why much more than half of the observations in Figure 1 lie within the 50 % confidence intervals.

5   Conclusions

In applying the ensemble Kalman filter, we have demonstrated how relatively simple aggregated biomass models, typical in bioeconomic analysis, can capture much of the dynamics of ecosystems. When compared to earlier efforts (Ussif et al. 2003), our results are far superior. But our model still has potential for improvement. The results overstate uncertainty as only a few of the parameter distributions have converged. More data or a more accurate initial ensemble could remedy the situation. In addition, a number of possible variations of the model exist; model noise as pure white (non-geometric) noise, assume perfect observations of the control variables (catch and inflow), model herring inflow as a state variable, and model herring inflow as pure noise with non-zero mean. The listed variations were all tested and lead to only small changes in the model fit and small variations in parameter estimates. A future research topic is how to treat more general noise terms, for example terms including unknown parameters.

Compared to the standard Kalman filter and the extended Kalman filter, the ensemble Kalman filter has no disadvantages, but a number of advantages. The ensemble Kalman filter is suitable for large-scale problems, it requires no linearisation, it is robust, it extends to asyncronous observations, and solves rank problems which may occur with large numbers of observed variables. Compared to variational adjoint methods, the ensemble Kalman filter is simpler (it requires no adjoint operator) and it has flow-dependent covariance, but the ensemble integration can be computationally costly and it is not iterative (iterative filtering is seen as a solution to strongly nonlinear systems, but is still under investigation). A range of related methods exist; see Evensen (2003). We mention only two; square root filters and the method of ensemble inflation. Square root filters remedy a tendency to underestimate the covariance (a typical problem when working with relatively few ensemble members). Ensemble inflation is a remedy for filter divergence which occurs when the ensemble collapses to a single solution; see Anderson and Anderson (1999) or Hansen and Penland (2007, p. 91) for a brief introduction. Hansen and Penland (2007) also provide

an alternative take on parameter estimation with the ensemble Kalman filter.

If simple models are to capture complex dynamics, they must rely on adaptive parameters. Estimates of parameters in chaotic systems are not likely to be very precise, and management models should be flexible and adaptive (Holling and Meffe 1996, p. 332). It is important that management models take the uncertainty of the dynamics into account. Adaptive management models such as feedback models are already well understood in the bioeconomic literature (Sandal and Steinshamn 1997). The challenge is to solve models of higher dimensionality which must underly ecosystem-based management. We believe the ensemble Kalman filter has an important role to play in both theoretical and operational management research, particularly in light of the recent calls for ecosystem-based management (Commission 2003).

The model we have estimated is not the final model. Uncertainty is overstated, and the model must be revised upon further observations. New observations are always imminent, and future management regimes must therefore be adaptive, not only in the sense that uncertainty is explicitly modeled, but in the sense that the underlying model, and in turn management decisions, are revised at appropriate times. The management model we envision here is thus not much different than current management regimes in terms of model and decision revisions. In the broader scope of things, we aim to answer calls for 'flexible, adaptive, and experimental' management models (Holling and Meffe 1996, p. 332), who further write that 'effective natural resource management that promotes long-term system viability must be based on an understanding of the key processes that structure and drive ecosystems, and on acceptance of both the natural ranges of ecosystems variation and the constrains of that variation for long-term success and sustainability' (p. 335). We think, when models are simplified and reduced down to the key driving phenomena, that the ensemble Kalman filter can capture variabilities and stabilities of ecosystems and serve tractable management models.

Acknowledgements

References

Anderson, Jeffrey L., Stephen L. Anderson. 1999. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review* **127**(12) 2741–2758.

Berck, Peter, Grace Johns. 1985. Policy consequences of better stock estimates in Pacific Halibut fisheries. *Working Paper, University of California, Berkeley, Department of Agricultural and Resource Economics* **348**.

Burgers, Gerrit, Peter Jan van Leeuwen, Geir Evensen. 1998. Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review* **126**(6) 1719–1724.

Commission, Pew Oceans. 2003. America's living oceans-charting a course for sea change. recommendations for a new ocean policy. Tech. rep., Pew Foundation, Washington, D.C.

Evensen, Geir. 1994. Sequential data assimilation with a nonlinear quasi-geostropic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research* **99**(C5) 10143–10162.

Evensen, Geir. 2003. The Ensemble Kalman Filter: Theoretical formulation and practical implementation. *Ocean Dynamics* **53** 343–367.

Evensen, Geir. 2009. *Data Assimilation: The ensemble Kalman filter*. 2nd ed. Springer-Verlag, Berlin Heidelberg.

Gjøsæter, Harald, Bjarte Bogstad, Sigurd Tjelmeland. 2009. Ecosystem effects of the three capelin stock collapses in the Barents Sea. *Marine Biology Research* **5** 40–53.

Grønnevik, Rune, Geir Evensen. 2001. Application of ensemble-based techniques in fish stock assessment. *Sarsia* **86** 517–526.

Hannesson, Rögnvaldur. 2007. Cheating about the Cod. *Marine Policy* **31** 698–705.

Hansen, James A., Cécile Penland. 2007. On stochastic parameter estimation using data assimilation. *Physica D* **230**(1-2) 88–98.

Holland, Daniel S., James N. Sanchirico, Robert J. Johnston, Deepak Joglekar. 2010. *Economic Anlysis for Ecosystem-Based Managment: Applications to Marine and Costal Environments*. RFF Press, Washington. D.C.

Holling, C. S., Gary K. Meffe. 1996. Command and control and the pathology of natural resource management. *Conservation Biology* **10**(2) 328–337.

Huse, Geir, Geir O. Johansen, Bjarte Bogstad, Harald Gjøsæter. 2004. Studying spatial and thropic interactions between capelin and cod using individual-based modelling. *ICES Journal of Marine Science* **61** 1201–1213.

Kalman, Rudolph Emil. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering* **82**(Series D) 35–45.

Kaufman, Les, Burr Heneman, J. Thomas Barnes, Rod Fujita. 2004. Transition from low to high data richness: An experiment in ecosystem-based fishery management from California. *Bulletin of Marine Science* **74**(3) 693–708.

Kugarajh, Kanaganayagam, Leif K. Sandal, Gerhard Berge. 2006. Implementing a stochastic bioeconomic model for the North-East Artic cod fishery. *Journal of Bioeconomics* **8**(1) 35–53.

Ludwig, Donald, Ray Hilborn, Carl Walters. 1993. Uncertainty, resource exploitation, and conservation: Lessons from history. *Science* **260**(5104) 17, 36.

Marshall, C. Tara, Nathalia A. Yaragina, Bjørn Ådlandsvik, Andrey V. Dolgov. 2000. Reconstructing the stock-recruit relationship for Northeast Artic cod using a bioenergetic index of reproductive potential. *Canadian Journal of Fisheries and Aquatic Sciences* **57**(12) 2433–2442.

May, Robert M., John R. Beddington, Colin W. Clark, Sidney J. Holt, Richard M. Laws. 1979. Management of multispecies fisheries. *Science* **205** 267–277.

O'Brien, Karen, Heather Tompkins, Siri Eriksen, Pål Prestrud. 2004. Climate vulnerability in the Barents Sea ecoregion: A multi-stressor approach. CICERO Report 2004:07, Center for International Climate and Environmental Research - Oslo.

Peterman, Randall M., Brian J. Pyper, Jeff A. Grout. 2000. Comparison of parameter estimation methods for detecting climate-induced changes in productivity of Pacific salmon (*Oncorhynchus* spp.). *Canadian Journal of Fisheries and Aquatic Sciences* **57**(1) 181–191.

Peterman, Randall M., Brian J. Pyper, Brice W. MacGregor. 2003. Use fo the Kalman filter to reconstruct historical trends in productivity of Bristol Bay sockeye salmon (*Oncorhynchus nerka*). *Canadian Journal of Fisheries and Aquatic Sciences* **60**(7) 809–824.

Røttingen, Ingolf, Sigurd Tjelmeland. 2008. A quest for management objectives - case study on the Barents Sea Capelin. *ICES 2008 Annual Science Conference, Halifax, Nova Scotia, Canada*. O:08, ICES CM Documents 2008.

Sandal, Leif K., Stein I. Steinshamn. 1997. A stochastic feedback model for optimal management of renewable resources. *Natural Resource Modeling* **10**(1) 31–52.

Sandal, Leif K., Stein I. Steinshamn. 2010. Rescuing the prey by harvesting the predator: Is it possible? Endre Bjørndal, Mette Bjørndal, Panos M. Pardalos, Mikael Rönnqvist, eds., *Energy, Natural Resources and Environmental Economics*. Energy Systems, Springer-Verlag Berlin Heidelberg, 359–378.

Schnute, Jon T. 1994. A general framework for developing sequential fisheries models. *Canadian Journal of Fisheries and Aquatic Sciences* **51**(8) 1676–1688.

Squires, Dale. 2009. Opportunities in social science research. Richard J. Beamish, Brian J. Rothschild, eds., *The Future of Fisheris Science in North America*, *Fish & Fisheries Series*, vol. 31, chap. 32. Springer, 637–696.

Tjelmeland, Sigurd. 2011. Personal communication.

Tschirhart, John. 2009. Integrated ecological-economic models. *Annual Review of Resource Economics* **1** 381–407.

United Nations, The World Summit on Sustainable Development. 2002. Plan of Implementation of the World Summit on Sustainable Development: Johannesburg Declaration on Sustainable Development. UN Documents A/Conf.199/20, United Nations.

Ussif, Al-Amin M., Leif K. Sandal, Stein I. Steinshamn. 2003. A new approach of fitting biomass dynamics models to data. *Mathematical Biosciences* **182** 67–79.

Walters, Carl. 1986. *Adaptive Management of Renewable Resources*. Macmillan Publishing Company, New York.

Wilen, James E. 2000. Renewable resource economists and policy: What differences have we made? *Journal of Environmental Economics and Management* **39**(3) 306–327.

Worm, Boris, Edward B. Barbier, Nicola Beaumont, J. Emmett Duffy, Carl Folke, Benjamin S. Halpern, Jeremy B. C. Jackson, Heike K. Lotze, Fiorenza Micheli, Stephen R. Palumbi, Enric Sala, Kimberley A. Selkoe, John J. Stachowicz, Reg Watson. 2006. Impacts of biodiversity loss on ocean ecosystem services. *Science* **314**(5800) 787 – 790.