

Environmental regimes and density-dependence: a Bayesian modeling approach for identifying recruitment regimes.

Stephan B. Munch, Marc Mangel, Athanasios Kottas

Department of Applied Mathematics and Statistics

Center for Stock Assessment Research

Baskin School of Engineering

University of California

1156 High Street

Santa Cruz, Ca. 95064

Abstract

Decadal scale oscillations in the environment result in substantial alterations to population dynamics as evidenced by time series of abundance and recruitment. Regime shifts may occur on very short time scales and are often undetected for several years. Consequently, tools that allow the estimation of regime-specific population dynamic parameters may be of great value. Using a hidden Markov model to describe the unobserved regime state, we develop methods to infer regime-specific parameters for a commonly used model of density dependent recruitment in addition to identifying the unobserved regime state. We apply the method to recruitment data for Japanese sardine and Baltic cod.

Keywords: regime shift, hidden Markov model, density dependence, Japanese sardine, Baltic cod.

Introduction

Fluctuations in climate and ocean conditions result in substantial changes in marine ecosystem structure and dynamics and organisms at all trophic levels may be affected by variations in a variety of environmental variables. Since the responses of organisms to environmental change have substantial ecological and economic impacts, understanding the relationship between climatic variables and population dynamics is of fundamental importance in conservation and management.

Environmental variation may affect species distributions and population dynamics. For instance, shifts in species ranges (Parmesan 1996) have been documented, though the degree to which ranges shift is species specific. For example, a comparison of intertidal invertebrate communities in California between 1933 and 1993 revealed that northern species decreased, southern species increased, and cosmopolitan species remained the same (Barry et al. 1995). Consequently, environmental variability may lead to shifts in community structure (Holbrook et al. 1997) and species interactions (Cattadori et al. 2005).

Environmental variation may also affect a number of vital rates. Kjesbu et al. (1998) found a tight correlation between fecundity and both temperature and prey availability. In steelhead (*Oncorhynchus mykiss*) juvenile survival varies with environmental state (Welch et al. 2000). Moreover, rates of somatic growth are strongly influenced by temperature and prey availability (Elliot 1994). Combined, these effects lead to environmentally induced fluctuations in density dependence (Jacobson and MacCall 1995, Sugimoto et al. 2001).

These observations have prompted many attempts (on the order of hundreds) to incorporate environmental variables into density dependence models for juvenile survival (e.g. Cushing 1982, Drinkwater and Myers 1987). Typically, the approach used is to multiply a standard model by some additional function of the environment (e.g. Madenjian et al. 2005, Majormaki 2004, Kuikka et al. 1999) although Koster et al. (2001) have attempted to build up a model from a series of stage-specific regressions relating survival to some environmental variables. Very few of such environmental correlations, however, are upheld in subsequent analyses (reviewed in Myers 1998).

Recent efforts indicate that aggregate variables such as the NAO index and PDO are better predictors of ecological processes than more mechanistically interpretable variables such as temperature. Hallett et al. (2004) suggest that this result is driven by the complexity of the interactions between weather and ecology; because the specific environmental mechanisms that are most important in any year may vary, aggregate variables such as the NAO do a better job in coarse-grained analyses.

Over the past two decades, long-term environmental fluctuations or "regimes" have become apparent. Regimes are quasi-stable states with residence times on the order of decades. However, regime shifts occur relatively rapidly, typically within the course of a year or two. Many aspects of the ocean environment change in the course of a regime shift – changes in atmospheric circulations, surface layer salinity and temperature, mixed layer depth, and mean sea level. As a consequence, regime shifts may have profound effects on population dynamics. Production indices for sablefish, sardine, English sole, Pacific cod, North Sea plaice, and several salmonids all show pronounced decadal scale fluctuations (McFarlane et al. 2000, MacCall 2002, Kell and Bromley

2004). These fluctuations, when compounded with shifts in species distributions, lead to decadal scale fluctuations in landings in many species including saury, sardine, filefish, sandfish, pollock (Kang et al. 2000, Wooster and Zhang 2004).

Several studies have also demonstrated decadal scale variations in recruitment time series (e.g. McFarlane et al. 2000, Daskalov 2003) that correspond with shifts in physical variables. Hare and Mantua (2000) review a large number of physical and biological time series and suggest that regime shifts may be more easily recognizable in recruitment indices than in the environmental variables themselves. For example, although the physical data do not unambiguously point to a regime shift in 1986, one is clearly visible in recruitment patterns of Pacific groundfishes. This observation led Tian et al. (2004) to suggest the use of saury as a bioindicator of regime shifts.

In light of the many species affected by regime shifts, it will be valuable to be able to incorporate regime effects into models of density dependent juvenile survival. One possibility would be to incorporate a set of environmental variables or aggregate indices into standard models. However, since regime shifts are characterized by relatively substantial changes in many biological and abiotic factors, regime-specific models might be more appropriate. Indeed, Wada and Jacobson (1998) find that the intercepts of models fit to data during each regime were significantly different. Thus a "recruitment-regime" approach seems warranted. However, since it is frequently unclear whether a regime shift has occurred until several years have passed, an alternative method is needed.

Here, we propose a new method for fitting regime specific models for density dependent juvenile survival and simultaneously estimating the probability that a regime

shift has occurred. We assume that there are two environmental regimes characterized by distinct survival dynamics. This assumption corresponds well with the observed impacts of regimes on recruitment (Wada and Jacobson 1998, Hare and Mantua 2000). We develop a two-component mixture model where the components are associated with the underlying regimes; i.e. the mean of each mixture component is defined by a regime specific model for density dependent juvenile survival. Temporal aspects are captured using a two-state Markov chain for the mixing distribution. We employ a Bayesian modeling formulation, which allows incorporation of prior information regarding both the regime shifts and the regime specific density dependence models, and yields full and exact inference. Our approach is alternatively referred to as a Markov-dependent mixture model, Markov switching regression model, or hidden Markov model (HMM). (However, note that HMMs do not typically include the regression component.) HMMs have been successfully applied to problems in several fields, including speech recognition (e.g., Juang and Rabiner 1991), bioinformatics (e.g., Liu et al. 1999), econometrics (e.g., Billio et al. 1999), hydrology (e.g., Lu and Berliner 1999), and image processing (Romberg et al. 2001). Scott (2002) provides a review on Bayesian methods for HMMs. To our knowledge, HMMs have not been applied in population dynamics.

The Methods section introduces the statistical modeling approach. The methodology is tested with simulated data and subsequently applied to two real data sets, one for Japanese sardine and one for Baltic cod.

Methods

If the regime states for each year were known, the problem would be straightforward; we could divide the data into regime-specific subsets and fit separate models to each. However, as noted above, the regime state may shift over a short period of time and may not be adequately characterized for several years following a shift. As a consequence, we require a method that allows us to make inferences regarding the regime state as well as the regime-specific model parameters. To do so, we characterize the regime state as a hidden variable whose dynamics are determined by a first order Markov process. As discussed in the Introduction, we work with two environmental regimes, although in principle our methodology can be extended to any specified number of regimes. Specifically, the regime state in year t (r_t) takes on one of two values (here, 1 or 2) with some probability that depends on the regime state in year $t-1$. That is,

$$1) \quad P(r_t = i | r_{t-1} = j) = q_{ij} \quad i, j = 1 \text{ or } 2$$

where the q_j 's are the probabilities of transitions into regime i from regime j . In a Bayesian context, these transition probabilities themselves require a prior probability model. We use a pair of independent Beta distributions for the probability of remaining in each regime, specifically, $q_{jj} \sim B(\pi_j, \rho_j)$, $j = 1, 2$. In everything that follows we used $\pi_j = 40$ and $\rho_j = 4.44$ which establishes a prior expectation of remaining in each regime of 0.9 with a standard deviation of 0.045, in keeping with the fairly extensive time series on regime shifts.

For each regime-specific model of juvenile survival, we use a generalized contest competition model referred to as the Shepherd (1982) function. Specifically, survival from birth to some fixed point in time, hereafter referred to as recruitment (R) is given by

$$2) \quad R_j = \frac{a_j N_0}{1 + (b_j N_0)^{c_j}}$$

where N_0 is the initial number of individuals in a cohort and the subscript indicates parameters for recruitment under regime j . We assume that errors around the recruitment relationship are multiplicative and log normal. Consequently, we switch to the transformed variables, $y = \ln(R)$, $x = \ln(N_0)$. Since each of the three parameters of the Shepherd model must be positive, we use a slightly different parameterization to facilitate model fitting. Our log-transformed, re-parameterized recruitment model is given by

$$3) \quad F(x | \mathbf{a}_j, \mathbf{b}_j, \mathbf{k}_j) = \mathbf{a}_j + x - \ln[1 + \exp\{(\mathbf{b}_j + x)\exp(\mathbf{k}_j)\}]$$

where $a = \exp(\alpha)$, $b^c = \exp(\beta)$, and $c = \exp(\kappa)$. For each of α , β , and κ we define independent, diffuse normal priors with means and variances indicated by μ_α , σ_α^2 , etc. For both regimes we set $\mu = 4, -4, 0$ for α , β , and κ , respectively, and $\sigma^2 = 10$ for each of these parameters. The error variances, V_j , are given independent inverse Gamma priors (Gelman et al. 1995) with the same parameters (γ_V , η_V). We set $\gamma_V = 11$, $\eta_V = 1$ which asserts a prior expectation of 0.1 for V_j with a standard deviation of 0.03. Although this prior is fairly informative, we found that the results did not change if less specific priors

were used (i.e. doubling the S.D. for V_j doesn't change the results). To simplify notation, we gather the regime-specific model parameters into a vector $\mathbf{q} = \{\mathbf{q}_1, \mathbf{q}_2\}$ with $\mathbf{q}_j = \{\alpha_j, \beta_j, \kappa_j\}$, the regime specific error variances into $V = \{V_1, V_2\}$ and the transition probabilities into a vector Q . The fully specified Bayesian model is

$$\begin{aligned}
 4) \quad & y_t | x_t, r_t, \mathbf{q}_{r_t}, V_{r_t} \stackrel{ind.}{\sim} N(y_t | F(x_t | \mathbf{q}_{r_t}), V_{r_t}), \quad t = 1, \dots, T \\
 & \{r_1 \dots r_T\} | Q \sim \prod_{t=2}^T P(r_t | r_{t-1}, Q) \\
 & Q, \mathbf{q}, V \sim p(Q) p(\mathbf{q}) p(V)
 \end{aligned}$$

Here $N(y|\boldsymbol{\mu}V)$ denotes a normal density for y with mean $\boldsymbol{\mu} = F(x|\mathbf{q})$ (as defined in Eq 3) and variance V . $P(r_t|r_{t-1}, Q)$ is given in Eq 1, and the priors for Q , \mathbf{q} , and V are defined above. T is the final year in which recruitment was observed. Since the regime designation is arbitrary, we set $r_1 = 1$ for model identifiability. Sampling from the posterior for all of the parameters and imputed regime states ($\mathbf{r} = \{r_1, \dots, r_T\}$), i.e. $p(Q, \mathbf{q}, V, \mathbf{r} | \text{data})$, was accomplished with standard techniques for HMMs (see, e.g., Scott 2002 and references therein). Details of the sampling algorithm are described in Appendix A.

Since it may be of interest to apply this method when it is uncertain whether two regimes exist, we want to compare the two regime model to a model with only one regime. Several approaches exist for formal Bayesian model comparison (e.g., Bernardo and Smith 2000). Here we utilize a posterior predictive space criterion suggested by Gelfand and Ghosh (1998). This criterion favors the model (m) which minimizes

$$5) \quad D(m) = \sum_{t=1}^T V^{(m)}(y_{t,new} | x_t, data) + \sum_{t=1}^T (y_t - E^{(m)}(y_{t,new} | x_t, data))^2$$

where data indicates the observed $\{y_t, x_t; t=1, \dots, T\}$. $E^{(m)}$ and $V^{(m)}$ denote the posterior mean and variance in the predicted recruitment in year t ($y_{t,new}$) under model m given x_t . The first term can be viewed as a penalty term for model complexity; too simple or too complex models will yield large values. The second term is a sum of squared deviations between the posterior mean recruitment in year t under model m and the observed recruitment. Thus $D(m)$ represents a balance between predictive uncertainty and fidelity to past observations. We estimate $E^{(m)}$ and $V^{(m)}$ for each t using samples from the posterior predictive distribution, $p^{(m)}(y_{t,new} | x_t, data)$.

There are several additional quantities beside the fit of the model that are of interest. These include the posterior predictive distributions (forecasts) for subsequent recruitment and the conditional probability of the regime state given a new recruitment observation. First, we might want to know what the long run probability of recruitment is given the initial cohort size (i.e. $p(y|x)$). This marginal posterior can be obtained from the stationary distribution for the regime states and the regime specific recruitment probabilities conditional on the parameters, i.e.

$$6) \quad p(y | x, Q, \mathbf{q}, V) = \pi_1 N(y | F(x | \mathbf{q}_1), V_1) + (1 - \pi_1) N(y | F(x | \mathbf{q}_2), V_2)$$

where π_1 is the long-run or stationary probability of regime 1, which is calculated as $\pi_1 = (1 - q_{22}) / [2 - (q_{11} + q_{22})]$, for each of the sampled values of Q . The long run recruitment

probability $p(y/x)$ is calculated by averaging over these conditional regime specific recruitment probabilities obtained at each step in the sampling algorithm.

The second quantity of interest is the posterior forecast of the next year's recruitment (y_{T+1}) given a specific initial cohort size (x_{T+1}). This is distinct from the long run probability of recruitment in that we are explicitly taking into account the dependence of next year's regime state on the regime state in year T. This posterior forecast can be obtained as described above by averaging over the posterior

7)

$$p(y_{T+1} | x_{T+1}, Q, \mathbf{q}, V, r_T) = \mathbf{p}^* N(y_{T+1} | F(x_{T+1} | \mathbf{q}_1), V_1) + (1 - \mathbf{p}^*) N(y_{T+1} | F(x_{T+1} | \mathbf{q}_2), V_2)$$

where $\pi^* = p(r_{T+1}=1 | r_T, Q)$ as given in Eq 1.

Finally, we may be interested in assessing the likelihood that a regime shift has occurred given some new recruitment observations. Here we focus on the probability that the regime in year T+1 is 1 given the observed initial cohort size and subsequent recruitment in year T+1. This posterior predictive probability is given by $p(r_{T+1} | x_{T+1}, y_{T+1}, data) = p(r_{T+1}, y_{T+1} | x_{T+1}, data) / p(y_{T+1} | x_{T+1}, data)$. The denominator is calculated by averaging over posterior samples from Eq 7 as described above. The numerator is estimated by averaging over the posterior samples the conditional probability,

$$8) \quad p(r_{T+1}, y_{T+1} | x_{T+1}, Q, \mathbf{q}, V, r_T) = p(r_{T+1} | r_T, Q) N(y_{T+1} | F(x_{T+1} | \mathbf{q}_1), V_1)$$

Testing the method

We tested the method on simulated data generated from mixtures of two Shepherd models with a common error variance. The ability of our method to resolve the regime specific models should be sensitive to the distance between the two models relative to the residual variance. However, how to define the distance between the nonlinear regime specific models is not clear. To address this, we simulated three data sets in which the regime specific models were held fixed and the error variance around each was 0.001, 0.01, and 0.1. This approach generated data sets with recruitment CVs of 23%, 30%, and 44% respectively. The remaining parameters for the underlying models were $\mathbf{q}_1 = \{1.2, 0.01, 1\}$, $\mathbf{q}_2 = \{1, 0.02, 0.8\}$, $Q = \{0.75, 0.6\}$. For these data sets we also tabulated the frequency with which the regimes were correctly classified. Note that our purpose in presenting these simulations is to demonstrate the efficacy of the approach when the true answer is known, not to conduct an exhaustive analysis of the sensitivity of our method. Results of rigorous sensitivity analyses will be published in a more general context in a subsequent paper.

We also applied the method to two real data sets. The first is the sardine data previously analyzed by Wada and Jacobson (1998). We used their egg production estimates and index of age 1 recruits from 1977-1995 (reported in Wada and Jacobson 1998, Table 1). In this case, given the previously established regime shifts in Pacific species, we expect to be able to identify distinct regime-specific differences in survival to age 1. The second data set is for Baltic cod from 1970 to 1993. Baltic cod have experienced a substantial decline in abundance and recruitment through the 1990's, which is thought to have resulted from a combination of high fishing pressure and

environmental effects (Mollman and Koster 1999). The data for our analysis are for Baltic areas 25-32 as downloaded from R. Myers Stock Recruitment Database (<http://www.mscs.dal.ca/~myers/welcome.html>). In this case, recruitment is defined in terms of the numbers of individuals surviving to age 2 and the biomass of mature individuals is used as a proxy for the initial numbers in the cohort.

Results

Under relatively low variability, the method recovers the underlying mixture of density dependence models well (Figure 1). As variability increases, the fit to the data remains good but the deviation between the fitted model and generating model increases substantially. In the low variance case, regimes for all years were correctly identified. In the intermediate variance case, 10% of regimes were misclassified and the degree of confidence in regime classifications was lower. In the high variance case, the method converged on a regime configuration with 30% of regimes misclassified. For the low, medium, and high variance cases, the model selection criterion was 1.6, 1.5, and 3.7 for the single regime model and 0.2, 0.4, and 1.0 for the two regime model. Thus in each case, we would choose the two regime model.

The marginal posterior density for the sardine recruitment series successfully identifies two, clearly separated recruitment regimes (Figure 2A). The inferred regime states agree fairly well with the known environmental regime states (Figure 2B) and the model selection criterion strongly supports the existence of distinct recruitment regimes ($D = 5.8$ for 2 regimes and $D = 40.1$ for a single Shepherd function). There were, however, a number of years in which the regime state differed from expectation based on

prior analyses of environmental variables (Tian et al. 2004, Hare and Mantua 2000), notably 1979, 1982, and 1992. The fit to the Baltic cod data was also quite good resulting in clear separation between recruitment regimes (Figure 3A). Again, the model selection criterion favors two regimes, though less strongly than in the sardine example ($D = 3.1$ for 2 regimes and 7.1 for a single regime model). Interestingly, the inferred regimes suggest a return to the high recruitment relationship for the years 1991-1993 (Figure 3B).

In both of these data sets, the posterior mean error variance was considerably smaller than expected based on previous studies in which single density dependence models were fit. For the cod data we found posterior mean (± 1 S.D) error variances of 0.03 (± 0.01) and 0.04 (± 0.01) and for the sardine data we found posterior mean error variances of 0.05 (± 0.01) and 0.04 (± 0.01). In each case these are an order of magnitude lower than that commonly reported for similar data. Thus, if there are multiple recruitment regimes, estimates of error variance based on a single model will be severely biased upwards.

Discussion

Most efforts to incorporate environmental effects into density dependence models do so by multiplying a standard model with an exponential term containing some linear combination or generalized additive model of environmental variables. Since many of the relevant environmental drivers are highly correlated on long time scales, incorporating them all through a regression approach may produce unreliable fits due to multicollinearity. The regression approach essentially modifies the slope parameter with

the intent of capturing environmentally induced changes in reproductive output or density independent mortality. However, this approach does not actually model density independent factors but confounds density dependent and density independent effects since both are represented in the denominator of Eq 2. Moreover, given that regime changes are characterized by substantial changes in the prey and predator fields to which cohorts are exposed prior to recruitment, it is certainly plausible that the nature of density dependence changes as well. Our method allows both the density dependent and density independent terms to vary under multiple environmental regimes. Moreover, the hidden Markov model on which our approach is based allows inferences to be made regarding the unobserved regime state.

Based on previous analyses (Wada and Jacobson 1999), we had a strong expectation that the data for the Japanese sardine would show a clear regime signal. Although this was the case, there were several odd years that did not fit well into the decadal regime paradigm. Curiously, 1979, 1982, and 1992 exhibited marked deviations from the norm in the Winter North Pacific index reported in Tian et al. (2004). However, visual inspection of the results suggests that had we used a model with an inflection point, 1979 and 1982 would have been assigned to the expected regime, suggesting that Allee effects may be important. Most interestingly, the posterior mean (± 2 SD) for the shape parameter (c) for the high recruitment regime was 0.9 (± 0.3) and for the low recruitment regime was 2.6 (± 0.4) indicating that the nature of density dependence is substantially different in the two regimes.

The Baltic cod data also clearly indicate the presence of two distinct recruitment regimes. Given the well documented difficulties in several cod fisheries (Myers et al.

1997), it may be worthwhile to consider the implications of decadal variation in recruitment regimes (MacCall 2002). Although many correlations between recruitment and the environment have been demonstrated, these are rarely used in management (Myers 1998). There are two good reasons for this. The first is that environmental correlations are notoriously unreliable – the vast majority have failed to stand the test of time (Myers 1998). Second, and perhaps more importantly, in order to use them to make predictions of future recruitment success, we must be able to forecast each of the relevant environmental variables. Under the coarse grained approach used here, the information necessary to forecast future regime states is estimated as part of the model. This can be used to forecast future recruitments conditional on the most recent regime estimates (Figure 4). These recruitment forecasts may be far different than the long run probability of recruitment (compare color in Figure 4 to Figures 2 and 3). Moreover, as adumbrated by Tian et al. (2004) new observations of recruitment may be used to estimate the probability of a regime shift (Figure 5).

There are several extensions of this method worth pursuing. The first would be to include environmental variables in the model either as a means of constraining the regime estimates or directly as covariates. Additionally, we might reason that multiple populations or species living in the same geographic region must experience the same time series of regimes, even if they respond to the actual environment in different ways. Thus incorporating multiple populations into the model may improve our ability to determine the unobserved regime state. One major extension which we are currently investigating involves replacing the parametric Shepherd model with a nonparametric Bayesian approach (e.g. Munch et al. 2005). This extension would allow us to approach

regime specific dynamics without the prior assumptions that a fixed parametric form is appropriate when the available biological data do not support such assertions.

We close by noting that there is one other important difference between the regime-specific recruitment approach and the use of environmental correlations; we expect environmental regimes to persist long enough to merit regime-specific management practices (MacCall 2002). For instance, the expected recruitment for sardines under each regime differ by a factor of about 7 indicating that very different management practices are required under each regime for sustainability.

References

- Anon. 1993. Report of the Working Group on the Assessment of Demersal Stocks in the Baltic. ICES Doc. C.M. 1993/Assess: 16.
- Barry, J.P., C.H. Baxter, R.D. Sagarin, S.E. Gilman. 1995. Climate-related, long-term faunal changes in a California rocky intertidal community. *Science* 267: 672-675.
- Bernardo, J.M. and A.F.M. Smith. 2000. *Bayesian Theory*. Wiley, Chichester.
- Billio, M., Monfort, A., and Robert, C.P. (1999). Bayesian estimation of switching ARMA models. *Journal of Econometrics*, 93, 229-255.
- Cattadori, I.M., D.T. Haydon, and P.J. Hudson. 2005. Parasites and climate synchronize red grouse populations. *Nature* 433: 737-741
- Cushing, D.H. 1982. *Climate and Fisheries*. Academic Press. London
- Daskalov, G. 2003. Long-term changes in fish abundance and environmental indices in the Black Sea. *Mar.Ecol.Prog.Ser.* 255:259-270.
- Drinkwater, K.F. and R.A.Myers. 1987. Testing predictions of marine fish and shellfish landings from environmental variables. *Can. J. Fish. Aquat. Sci.* 44:1568-1573.
- Elliot, J.M. 1994. *Quantitative Ecology and the Brown Trout*. Oxford University Press, Oxford.
- Gamerman, D. 1997. *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. CRC Press, London.
- Gelfand, A.E. and S.J. Ghosh. 1998. Model Choice: A minimum posterior predictive loss approach. *Biometrika* 85:1-11.

- Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin. 1995. Bayesian Data Analysis. Chapman and Hall/CRC, New York.
- Hallett, T.B., T. Coulson, J.G. Pilkington, T.H. Clutton-Brock, J.M. Pemberton, and B.T. Grenfell. 2004. Why large-scale climate indices seem to predict ecological processes better than local weather. *Nature* 430: 71-75.
- Hare, S.R. and N.J. Mantua. 2000. Empirical evidence for North Pacific regime shifts in 1977 and 1989. *Prog. Oceanogr* 47: 103-145.
- Holbrook, S.J., R.J. Schmitt, and J.S. Stephens Jr. 1997. Changes in an assemblage of temperate reef fish associated with a climate shift. *Ecol. Appl.* 7:1299-1310.
- Jacobson, L.D. and A.D. MacCall. 1995. Stock-recruitment models for Pacific sardine (*Sardinops-sagax*). *Can.J.Fish.Aquat.Sci* 52: 566-577.
- Juang, B.H., and Rabiner, L.R. (1991). Hidden Markov models for speech recognition. *Technometrics*, 33, 251-272.
- Kang, S., S.Kim, and S-W. Bae. 2000. Changes in ecosystem components induced by climate variability off the eastern coast of the Korean Peninsula during 1960-1990. *Prog. Oceanogr.* 47:205-222.
- Kell, L.T. and P.J. Bromley. 2004. Implications for current management advice for North Sea plaice (*Pleuronectes platessa*): Part II. Increased biological realism in recruitment, growth, density-dependent sexual maturation and the impact of sexual dimorphism and fishery discards. *J. Sea. Res.* 51: 301-312.

- Kjesbu, O.S., P.R. Witthames, P. Solemdal, and M.G. Walker. 1998. Temporal variations in the fecundity of Arcto-Norwegian cod (*Gadus morhua*) in response to natural changes in food and temperature. *J.Sea.Res* 40: 303-321.
- Koster, F.W., H.-H. Hinrichsen, M.A.St. John, D. Schnack, B. R. MacKenzie, J. Tomkiewicz, and M. Plikshs. 2001. Developing Baltic cod recruitment models. II. Incorporation of environmental variability and species interaction. *Can. J. Fish. Aquat. Sci.* 58:1534-1556.
- Kuikka, S. M. Hilden, H. Gislason, S. Hansson, H. Sparholt, and O. Varis. 1999. Modeling environmentally driven uncertainties in Baltic cod (*Gadus morhua*) management by Bayesian influence diagrams. *Can. J. fish. Aquat. Sci.* 56: 629-641.
- Liu, J.S., Neuwald, A.F., and Lawrence, C.E. (1999). Markovian structures in biological sequence alignments. *Journal of the American Statistical Association*, 94, 1-15.
- Lu, Z.-Q., and Berliner, M.L. (1999). Markov switching time series models with application to a daily runoff series. *Water Resources Research*, 35, 523-534.
- MacCall, A.D. 2002. Fishery-management and stock-rebuilding prospects under conditions of low-frequency environmental variability and species interactions. *Bull. Mar. Sci.* 70: 613-628.
- Madenjian, C.P., T.O. Hook, E.S. Rutherford, D.M. Mason, T.E. Croley II, E.B. Szalai, and J.R. Bence. 2005. Recruitment variability of alewives in Lake Michigan. *Trans.Am.Fish.Soc.* 134:218-230.

- Majormaki, T.J. 2004. Analysis of the spawning stock-recruitment relationship of vendace (*Coregonus albula*) with evaluation of alternative models, additional variables, biases, and errors. *Ecol.F.W.Fish.* 13:46-60.
- McFarlane, G.A., J.R.King, R.J. Beamish. 2000. Have there been recent changes in climate? Ask the fish. *Prog. Oceanogr.* 47: 147-169.
- Mollmann C.and F.W. Koster 1999. Food consumption by clupeids in the Central Baltic: evidence for top-down control? *ICES J.Mar. Sci.* 56: 100-113
- Myers, R.A. 1998. When do environment-recruitment correlations work? *Reviews in Fish Biology and Fisheries.* 8:285-305.
- Myers, R.A., J.A. Hutchings and N. J. Barrowman, 1997, Why do fish stocks collapse? The example of cod in Eastern Canada, *Ecological Applications* 7:91-106.
- Parmesan, C. 1996. Climate and species' range. *Nature* 382: 765-766
- Romberg, J.K., Choi, H., and Baraniuk, R.G. (2001). Bayesian tree-structured image modeling using wavelet-domain hidden Markov models. *IEEE Transactions on Image Processing*, 10, 1056-1068.
- Scott, S.L. (2002). Bayesian methods for hidden Markov models: Recursive computing in the 21st century. *Journal of the American Statistical Association*, 97, 337-351.
- Shepherd, J.G. 1982. A versatile new stock-recruitment relationship for fisheries and the construction of sustainable yield curves. *J. Cons. Int. Explor. Mer* **40**: 67–75.
- Sugimoto, T. S. Kimua, and K. Tadokoro. 2001. Impact of El Nino events and climate regime shift on living resources in the western North Pacific. *Prog. Oceanorg.* 49: 113-127.

- Tian, Y., Y. Ueno, M. Suda, and T. Akamine. 2004. Decadal variability in the abundance of Pacific saury and its response to climatic/oceanic regime shifts in the northwestern subtropical Pacific during the last half-century. *Journal of Marine Systems*. 52: 235-257.
- Wada, T. and L. Jacobson. 1998. Regimes and stock-recruitment relationships in Japanese sardine (*Sardinops melanostictus*), 1951-1995. *Can. J. Fish. Aquat. Sci.* 55:2455-2463.
- Welch, D.W., B.R. Ward, B.D. Smith, J.P. Eveson 2000. Temporal and spatial responses of British Columbia steelhead (*Oncorhynchus mykiss*) populations to ocean climate shifts. *Fish. Ocean.*9: 17-32.
- Wooster, W.S. and C.I. Zhang. 2004. Regime shifts in the North Pacific: Early indications of the 1976-1977 event. *Prog.Ocean.* 60:183-200.

Figure 1. Results for simulated data. Left column: Black dots are simulated data, white lines indicate the generating model, black lines indicate the posterior mean SR function for each regime. The color in each plot is proportional to the posterior density. Right column: Black bars indicate the posterior probability that r_t is 2 for each year. White dots indicate the true regime, low for regime 1, high for regime 2. Rows A-C indicate results for $V = 0.001, 0.01, \text{ and } 0.1$, respectively.




Figure 2. Results for Pacific sardines. A. Numbers indicate year and recruitment for the observed data with dotted lines connecting consecutive years. The color in each plot is proportional to the long-run posterior density (see Eq 6). B. Black bars indicate the posterior probability of being in the low recruitment regime (regime 2) for each year.

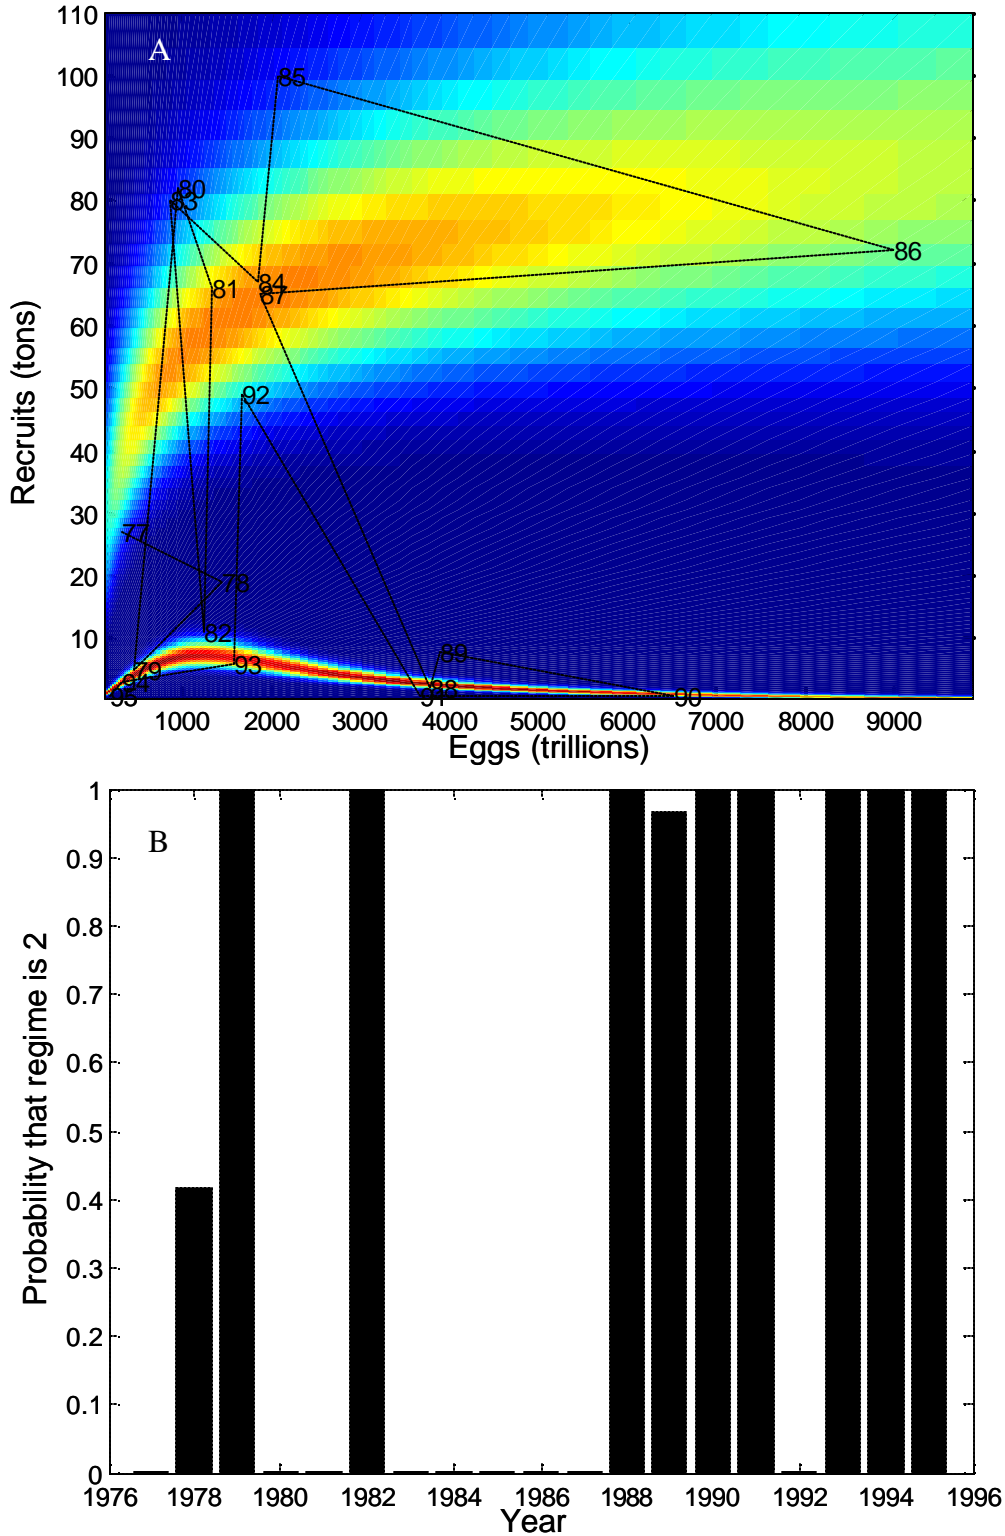


Figure 2. Results for Baltic cod. A. Numbers indicate year and recruitment for the observed data with dotted lines connecting consecutive years. The color in each plot is proportional to the long-run posterior density (see Eq 6). B. Black bars indicate the posterior probability of being in the low recruitment regime (regime 2) for each year.




Figure 4. Posterior forecast density for next year's recruitment as a function of initial population size for the sardine (A) and for the Baltic cod (B), conditional on previous regime states. Recruits are in tons for panel A and numbers of individuals for panel B. Color indicates posterior density with dark blue being zero and warmer color indicating higher probability.


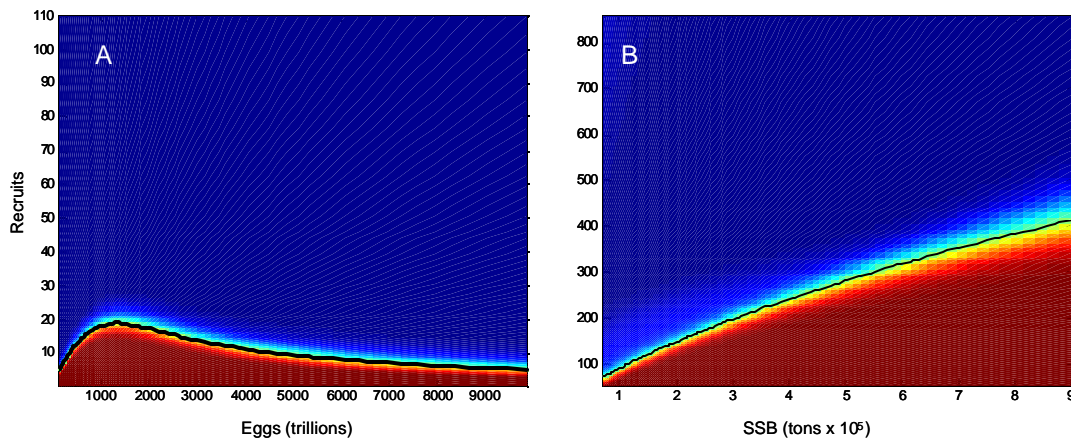


Figure 5. Probability of the next year's regime, as a function of initial cohort size and recruitment, updated based on a new observation of recruitment for the sardine (A) and Baltic cod (B). Recruits are in tons for panel A and numbers of individuals for panel B. Color indicates probability with dark red indicating probability 1 of being in the low recruitment regime and dark blue indicating probability zero of being in the low recruitment regime. The black line indicates the point where this probability is 0.5.



Appendix A. *Sampling from the posterior*

The full conditional for the error variance V_j is an inverse Gamma distribution with parameters

$$A1) \quad \gamma_{V_j} = \gamma_V + T_j/2, \quad T_j = |\{t: r_t = j\}|$$

$$\eta_{V_j} = \eta_V + \frac{1}{2} \sum_{t: r_t = j} [y_t - F(x_t | \mathbf{q}_j)]^2$$

Note that T_j is the total number of years for which $r_t = j$ and the summation is over those years as well. The probability that the regime is in state 1 in year t conditional on the data, parameters, and all other regime states is given by

A2)

$$P(r_t = 1 | data, \mathbf{q}, \mathcal{Q}, V, r_{-t}) = \frac{N(y_t | F(x_t | \mathbf{q}_1), V_1) P(r_t = 1 | r_{t-1}, \mathcal{Q}) P(r_{t+1} | r_t = 1, \mathcal{Q})}{\sum_{j=1}^2 N(y_t | F(x_t | \mathbf{q}_j), V_j) P(r_t = j | r_{t-1}, \mathcal{Q}) P(r_{t+1} | r_t = j, \mathcal{Q})}$$

where r_{-t} indicates the set of regime states with r_t removed. To sample from the full conditionals for the q_{ij} 's, we first define n_{ij} as the number of transitions into state i from state j in the current sample of r . Then, conditional on all other parameters, the q_{ij} 's $\sim B(\pi_j + n_{jj}, \rho_j + n_{ji})$.

Finally, the Metropolis step (Gammerman 1997) for the vector \mathbf{q} consists of drawing new values θ^{new} from a multivariate normal proposal distribution centered on θ with covariance matrix S , and accepting these values with probability u ,

$$A3) \quad u = \frac{p(\mathbf{q}^{new}) \prod_{j=1}^2 \prod_{t:r_t=j} N(y_t | F(x_t | \mathbf{q}_j^{new}), V_j)}{p(\mathbf{q}^{old}) \prod_{j=1}^2 \prod_{t:r_t=j} N(y_t | F(x_t | \mathbf{q}_j^{old}), V_j)}$$

The covariance matrix S was initially $0.05 \cdot I$, where I is a six dimensional identity matrix. S was updated every 5000 iterations based on the observed covariance among sampled elements of θ . Specifically, we used $S = 0.5 \text{Chol}(\Sigma)$ where Σ is the covariance among sampled parameters, and $\text{Chol}(\bullet)$ denotes the Cholesky factorization. The factor 0.5 seemed to allow reasonable rejection rates.

Unfortunately, the model structure described above is unidentifiable. Since the regime designation is arbitrary, for any data set there are two equally probable solutions corresponding to swapping the regime labels. That is, under one solution, some set of observations are assigned to group 1, the remainder to group two and parameters for each group are determined appropriately. The other equally reasonable solution is to swap group assignments for all members of each group and update parameter distributions accordingly. Thus, for sufficiently close regime-specific parameter sets, the algorithm will switch regime assignments for all (or nearly all) observations between each jump resulting in bimodal posteriors for the parameters and equiprobable group assignment.

A possible approach to overcoming this identifiability problem involves placing order restrictions on the parameters for each of the regime-specific models. In the present case, the nonlinearity of the models makes it difficult to rank the parameter vectors *a priori*. However, recognizing that this constraint is functionally equivalent to ignoring the mirror image of the regime assignments, we designed our fitting algorithm to avoid this label switching problem by limiting the step sizes between iterations and by providing the algorithm with good initial estimates obtained in the following way. First, a single Shepherd model was fit to the data using maximum likelihood. Preliminary regime assignments were made based on whether residuals from this fits were positive or negative. Finally, preliminary regime-specific Shepherd models were fit given these preliminary regime assignments. The resulting parameters and regime states were used as our initial values for the MCMC run.

As noted above, to keep the label switching problem to a minimum, relatively small jump sizes were used. This, however, induces a substantial autocorrelation among samples. To overcome this, we sampled 10×10^6 draws from the posterior and used a thinning rate of 0.0004 to construct our final posterior.