

Early-warning signals (potentially) reduce uncertainty in forecasted timing of critical shifts

DEREK KARSSENBERG[†] AND MARC F. P. BIERKENS

Department of Physical Geography, Faculty of Geosciences, Utrecht University, Heidelberglaan 2, P.O. Box 80115, 3508 TC, Utrecht, the Netherlands

Citation: Karssenber, D., and M. F. P. Bierkens. 2012. Early-warning signals (potentially) reduce uncertainty in forecasted timing of critical shifts. *Ecosphere* 3(2):15. <http://dx.doi.org/10.1890/ES11-00293.1>

Abstract. Despite the identification of early-warning signals precluding ecosystems regime shifts, limited evidence exists that they can be used to forecast the actual timing of a critical shift. Here, we propose a probabilistic Bayesian approach to forecast the timing of a shift by combining uncertain prior information about the ecosystem dynamics (parameters and drivers) and sampled spatial and temporal correlation and variance of ecosystem states, which are well known early-warning signals. For an ecosystem of logistically growing vegetation under linear increase in grazing pressure, we show that the use of sampled early-warning signals results in lower prediction uncertainty in forecasted timing of shifts compared to forecasts made with sampled mean state variables. In addition, we show that uncertainty in ecosystem parameters decreases well ahead of a shift. An important conclusion of our study is that the use of early-warning signals in forecasting of shifts is promising, provided that a large number of samples are collected ($n \approx 10^4$ in our study). This explains the limited success of finding early warning signals from field studies of real world ecosystems.

Key words: critical transition; data assimilation; early-warning indicator; forecast; logistic growth; particle filter; regime shift; sampling; uncertainty; vegetation.

Received 19 October 2011; revised 12 December 2011; accepted 16 December 2011; final version received 17 January 2012; **published** 14 February 2012. Corresponding Editor: D. P. C. Peters.

Copyright: © 2012 Karssenber and Bierkens. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits restricted use, distribution, and reproduction in any medium, provided the original author and sources are credited.

[†] **E-mail:** d.karssenber@uu.nl

INTRODUCTION

Non-linear systems may undergo abrupt transitions as a result of a gradual change in system drivers. Such regime shifts, or critical transitions, have been described for a large range of spatio-temporal systems, including ecosystems (Scheffer et al. 2009). Regime shifts are considered undesirable because they cause large changes that may be irreversible. For designing adaptation strategies to avoid shifts, the ability to forecast the probability and timing of a shift is essential (Andersen et al. 2009, Biggs et al. 2009, Contamin and Ellison 2009).

However, it is hard to detect upcoming regime

shifts because mean values of observable system state variables show little change before a transition occurs (Scheffer et al. 2009). This problem has sparked research focused on finding alternative properties of a system that show a more marked change before a transition is coming. It turns out that such early-warning signals exist and include various higher order statistics of state variables (Kleinen et al. 2003, Brock and Carpenter 2006, Kéfi et al. 2007, Chisholm and Filotas 2009, Dakos et al. 2009, Fernández and Fort 2009, Takimoto 2009, Drake and Griffen 2010, Carpenter et al. 2011).

Despite the identification of early-warning signals, limited evidence exists that they can be

used to actually forecast the timing of a shift (Biggs et al. 2009, Contamin and Ellison 2009). This is because the use of early-warning signals in forecasting is hampered by a number of issues. One is that changes in higher order statistics sufficiently well ahead of a shift are often considered too small to be detected by sampling schemes (Biggs et al. 2009). Also, absolute values of observed early-warning signals are considered to have limited value in estimating the proximity of a shift (Van Nes and Scheffer 2007, Biggs et al. 2009, Kéfi et al. 2010). This is because changes in early-warning signals are considered too gradual to provide information on the distance to the shift or absolute values depend on ecosystem characteristics that are generally unknown. In addition, the uncertainty of sampled early-warning signals may hamper their use to forecast shifts. Higher order statistics calculated from samples have a high uncertainty compared to first order statistics (de Gruijter et al. 2006). This may seriously reduce their information content for forecasting shifts even when early-warning signals show considerable changes ahead of a shift. It might even be the case that uncertainty in early-warning signals calculated from observations overshadows in many cases observational error, making them virtually useless in detecting shifts in ecosystems. Surprisingly, when analyzing the value of higher order statistics as predictors of regime shifts, this has not been considered thus far.

These issues raise the following questions: Is it possible to quantify the timing of a shift using early-warning signals calculated from a *sampled* ecosystem? And, what is the associated uncertainty in this forecasted timing, relative to the case when (notably small) trends in mean state variables are used as indicators instead of early-warning signals? To examine these questions we use a probabilistic Particle Filter (Van Leeuwen 2009, Karssenberg et al. 2010) framework that in principle allows forecasting the timing of a shift, including the associated uncertainty in the timing. We encode prior, but limited, ecosystem knowledge as a dynamical ecosystem model of logistic vegetation growth (Noy-Meir 1975, Dakos et al. 2009, Fernández and Fort 2009, Guttal and Jayaprakash 2009) where uncertainty in all parameters and drivers is defined by probability distributions. These uncertainties lead to a high

forecast uncertainty of the timing of a shift if no additional observational data were available. To reduce this uncertainty, however, we take samples of a key ecosystem state variable (biomass) at given moments in time, taking into account sampling error caused by limited sampling. As our aim is a proof of concept, we use a synthetic (model generated) world from which we sample. Early-warning signals calculated from these samples are combined with the ecosystem model using a numerical solution of Bayes' theorem.

METHODS

Sequential Bayesian filtering

When forecasting the timing of a regime shift in a particular study area, we can rely on two sources of information: the prior information consisting of a dynamical ecosystem model representing our understanding of the ecosystem and the observational data providing site-specific information on how the ecosystem changes over time. In sequential Bayesian Filtering, these two sources of information are combined using Bayes' theorem. Here we provide a short explanation of the statistical framework, for a full account the reader is referred to Appendix A.

The prior knowledge is described by a forward model and its parameter values. In cases when the ecological process behind the regime shift is partly unknown, one would need to represent the prior knowledge by a set of possible model structures (Clark et al. 2011). As existing statistical methods (e.g., Beven 2002, Ajami et al. 2007, Bulygina and Gupta 2010) dealing with a prior defined by multiple model structures do not fully address simultaneous identification of the model structure and parameters of spatial forward models in a Bayesian framework, it is assumed here that the mechanism behind the regime shift is known. So, the model structure of the forward model is kept fixed. The parameters, drivers and state variables in the model, however, are considered to be uncertain, and the full state of the model is given by a joint probability density of all variables (state variables, parameters, drivers), denoted by $p(\mathbf{x}_t)$, with a time index t , $t = 1, 2, \dots, 3500$. The system is monitored by collecting biomass samples at observation times t . Observation error (e.g., instrument error), and errors due to limited sampling are inherent in the

calculation of early-warning signals \mathbf{y}_t , resulting in a joint probability density of all early-warning signals, denoted as $p(\mathbf{y}_t)$. Below, \mathbf{y}_t is referred to as “observations”.

While monitoring, probabilistic forecasts of the full state $p(\mathbf{x}_t)$ are made by running the ecosystem model forward in time. At observation times, the full state is adjusted by assimilating observations into the model using Bayes' Theorem:

$$p(\mathbf{x}_t|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t)}{p(\mathbf{y}_t)} \quad (1)$$

with $p(\mathbf{x}_t|\mathbf{y}_t)$ the posterior model state, i.e., the joint probability density function of the full model state \mathbf{x}_t given the observations \mathbf{y}_t , and $p(\mathbf{y}_t|\mathbf{x}_t)$ the joint probability density function of the observations \mathbf{y}_t given the model state \mathbf{x}_t . We use the Particle Filter (Van Leeuwen 2009, Karssenberg et al. 2010), which is a Monte Carlo technique, to numerically solve Eq. 1. Unlike ensemble Kalman filters (Evensen 2003), the Particle Filter can deal with non-linear, non-Gaussian models, it does not violate conservation laws, and it provides a straightforward estimation of the joint posterior probability distributions of all state variables and parameters, also for non-Gaussian (e.g., multi-modal) distributions (Doucet et al. 2001). As it is a Monte Carlo technique it may require a large number of model calls, but this is also the case for other Monte Carlo techniques such as the Metropolis-Hastings algorithm (Van Leeuwen 2009). The Particle Filter samples model trajectories (called realizations or particles) $\mathbf{x}_t^{(n)}$, $n = 1, 2, \dots, N$, from the initial joint probability density $p(\mathbf{x}_0)$. Each realization is integrated forward in time up to the first observation time. At the observation time, Eq. 1 is numerically solved by calculating the posterior probability $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$ of each realization. Under the assumption of a normal distribution of the error in the observations, this probability takes the following form (Van Leeuwen 2009):

$$p(\mathbf{x}_t^{(i)}|\mathbf{y}_t) \sim \exp\left(-\frac{1}{2}[\mathbf{y}_t - \mathbf{H}\mathbf{x}_t^{(i)}]^T \mathbf{R}_t^{-1} [\mathbf{y}_t - \mathbf{H}\mathbf{x}_t^{(i)}]\right) \quad (2)$$

for $i = 1, 2, \dots, N$.

In Eq. 2, \mathbf{H} is the measurement operator transforming the model state to the observation, which is here one because no transformation is

required. The matrix \mathbf{R}_t is the covariance matrix of the measurement error at t . It represents all sources of uncertainty in the observations \mathbf{y}_t , including correlation (covariances) between the observation errors. Next, a new set of N realizations is created that consists of exact copies of a subset of realizations. This is done with Residual Resampling (Liu and Chen 1998, Weerts and El Serafy 2006), which copies each realization a number of times proportional to its posterior probability $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$ (Eq. 2). After resampling, the new set of realizations is integrated to the next observation time step using the system model, and the same algorithm is applied. An intuitive explanation of the procedure is that it removes at each observation time those model realizations that do not match the observations available for that time step, while the remaining model realizations are duplicated. At the last observation time, all observations have been incorporated into the ecosystem model resulting in the posterior probability distributions of the full model state and the forecasted timing of the ecosystem shift.

Prior information: the ecosystem model and its drivers

The prior information of the system is assumed to be a spatial model of logistically growing vegetation (Noy-Meir 1975). This system shows a critical transition towards lower biomass when a grazing threshold is exceeded. We selected this system because higher order statistics as early-warning signals have been described elsewhere (Dakos et al. 2009, Fernández and Fort 2009, Guttal and Jayaprakash 2009). Biomass is modelled in response to a linear increase in grazing over time. Biomass $B_{x,y,t}$ at cell coordinates (x, y) , with $x = 1, 2, \dots, 60$, $y = 1, 2, \dots, 60$ is:

$$\begin{aligned} \frac{dB_{x,y,t}}{dt} = & rB_{x,y,t} \left(1 - \frac{B_{x,y,t}}{k}\right) - (c_0 + c_{inc}t) \frac{B_{x,y,t}^2}{B_{x,y,t}^2 + 1} \\ & + dD_{x,y,t} + e_{x,y,t}, \end{aligned} \quad (3)$$

with

$$\begin{aligned} D_{x,y,t} = & B_{x,y-1,t} + B_{x,y+1,t} + B_{x+1,y,t} + B_{x-1,y,t} \\ & - 4B_{x,y,t}, \end{aligned} \quad (4)$$

using $B_{x,y,t} = k$ at $t = 0$. We consider situations with limited prior information, thus all drivers and parameters are taken uncertain, represented by spatially uniform random variables with

a uniform probability distribution. In most cases, prior ranges of parameter values can be estimated from literature, calibration against historic data, or direct observation in the field. These prior ranges of values can be used to define the prior probability distributions of the parameters. Here we assume prior information is available that can narrow the prior distribution as follows: in Eq. 3, r is the growth rate taken uniform in $[0.6, 1.0]$, k is the carrying capacity with a prior in $[9, 11]$, c_0 is the initial grazing pressure with a prior in $[0.75, 1.25]$, and c_{inc} is the expected future increase in grazing with a prior in $[0.0005, 0.0007]$. Eq. 4 represents lateral dispersal of biomass modelled as a diffusion process with the dispersion rate parameter d , with a prior in $[0.24, 0.36]$. The prior distributions are chosen somewhat arbitrarily, as interest lies here in a proof of concept. Results not shown here indicate that increasing the interval of the priors implies smaller information content, and thus slightly increases the forecast uncertainty and the range of the posterior distributions of the parameters for all scenarios described in the paper. The relative performance of scenarios remains approximately the same. In Eq 3, $e_{x,y,t}$ is a random variable that is independent for each location and time step with $e_{x,y,t} \sim N(0, 0.64)$, representing model structural error. For details on the numerical solution of Eqs. 3–4 and closed boundary conditions, see Appendix B. In all calculations, the timing of the shift is the time t at which the average biomass of the grid cells falls below 2.0.

Observational data:

sampled early-warning signals

As real world observations are not used, we generate a synthetic world by creating one realization of the model (Eqs. 3 and 4), using deterministic model input values corresponding to the median values of the prior probability distributions, i.e., $r = 0.8$, $k = 10$, $c_0 = 1.0$, $c_{inc} = 0.0006$, $d = 0.3$. Let $V_{x,y,t}$ be the time series of synthetic world biomass created in this way.

To mimic actual sampling, we add observation errors to retrieve the random variable biomass $C_{x,y,t}$ calculated as $C_{x,y,t} = V_{x,y,t} + S + R_{x,y,t}$ with, S the systematic error, a random variable with $S \sim N(0, \sigma_s^2)$; $R_{x,y,t}$ the random error, a random

variable that is spatially independent with $R_{x,y,t} \sim N(0, \sigma_r^2)$. Note that S is systematic and thus represents an offset (due to instrument error) in the observed value that is constant and spatially uniform. $R_{x,y,t}$ is random error and thus its offset is different for each time step and measurement location. Observed early-warning signals are retrieved by sampling $C_{x,y,t}$. We evaluate the use of spatial and temporal correlation and variance, given as semivariance values, because these are known early-warning signals for our ecosystem (Dakos et al. 2009, Fernández and Fort 2009, Guttal and Jayaprakash 2009). To compare spatial and temporal semivariance as early-warning signals, a sampling scheme is applied that can be used to calculate both types of early-warning signals, using the same number of samples. To retrieve the value of these early-warning signals at a particular observation time t , biomass is sampled for each time step t in a period of 50 time steps preceding t . Each of these 50 time steps, 40 samples are taken on a regular grid with a distance of 5 grid cells between samples. In addition, 10 samples are taken at a distance of one grid cell in the x or y direction away from randomly selected sampling locations on the regular grid, in order to calculate short-range spatial variation (Appendix C). This scheme is feasible and efficient for calculation of spatial semivariance (de Gruijter et al. 2006). Let this set of $50 \cdot 50 = 2500$ sampling locations in the temporal and spatial domains at a particular observation time t be the domain $D(t)$.

Because combining the modelled ecosystem state and sampled early-warning signals requires the covariance matrix \mathbf{R}_t (Eq. 2) of the error in the sampled early-warning signals, we perform a Monte Carlo simulation to calculate \mathbf{R}_t . A number of $M = 500$ realizations $C_{x,y,t}^{(m)}$ with $m = 1, 2, \dots, M$, are drawn from $C_{x,y,t}$ by randomly offsetting the regular sampling grid for each realization (see Appendix C). By this procedure \mathbf{R}_t includes observation error represented by the systematic error and random error incorporated in $C_{x,y,t}$ and sampling error (de Gruijter et al. 2006), i.e., the error caused by finite sampling from the reality. For each realization m , a number of early-warning signals are calculated, which are semivariance values at different separation intervals in the spatial or temporal direction. At a given observation time

t , the semivariance $\gamma_t^{(m)}(\bar{h})$ of sample realization m is calculated within the domain $D(t)$ as (Goovaerts 1997):

$$\gamma_t^{(m)}(\bar{h}) = \frac{1}{2N} \sum_{i=1}^{L(t)} \left(C^{(m)}(\mathbf{s}_i) - C^{(m)}(\mathbf{s}_i + h) \right)^2 - \sigma_r^2, \\ \forall(\mathbf{s}_i, \mathbf{s}_i + h) : h \in [h_{\min}, h_{\max}], \mathbf{s}_i \in D(t). \quad (5)$$

In Eq. 5, $\gamma_t^{(m)}(\bar{h})$ is the semivariance, calculated for time step t ; \mathbf{s}_i is a location (x, y, t) . The semivariance is calculated from biomass at L pairs of locations that have a separation distance between h_{\min} and h_{\max} , and represent semivariance in either the spatial or the temporal direction, with \bar{h} the average separation distance of the L pairs, also referred to as the lag.

The calculation of semivariance at observation time steps allows to assimilate these early-warning signals into the ecosystem model as described in the previous sections, which provides a means to evaluate their use in forecasting the timing of a shift. To compare the use of these early-warning signals with the alternative approach where an observed mean state variable, instead of higher order statistics such as semivariance, is used, we also calculate mean biomass at observation time steps. To make sure these mean values are calculated using the same number of observational data, which is required to make an unbiased comparison with the use of early-warning signals, the mean biomass $M_t^{(m)}$ is calculated in the same domain $D(t)$ which was also used in Eq. 5 for the calculation of semivariance values. In the Bayesian filter used to assimilate the observations into the ecosystem model, the median of $\gamma_t^{(m)}(\bar{h})$ at different separation distances in space and time and/or the median of $M_t^{(m)}$ are taken as the observation vector \mathbf{y}_t (Eq. 2), while \mathbf{R}_t is calculated using the $m = 1, 2, \dots, M$ Monte Carlo samples.

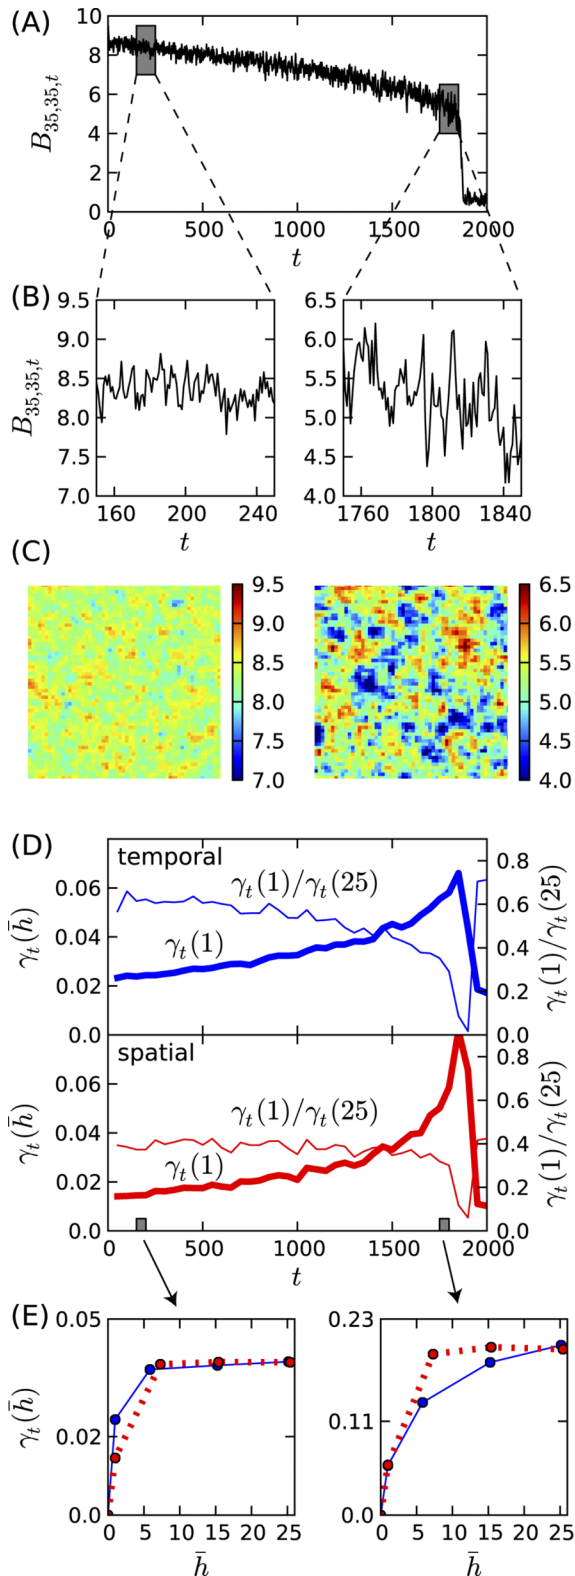
RESULTS

Relations between patterns and timing of transition

To explore the relation between early-warning signals and timing of a shift, we first describe the results of a single realization of the ecosystem model (Fig. 1). This realization was also used as synthetic data in the sequential Bayesian filtering

procedure that will be described in the second part of the results. The linear increase in grazing pressure over time results in a small reduction in biomass until the abrupt shift towards a much lower biomass occurs at $t = 1866$ (Fig. 1A). On top of this trend, biomass varies randomly in response to the small random addition or removal of biomass $e_{x,y,t}$ at each time step (Eq. 3). This random variation in biomass increases towards the shift, both when biomass is observed at a single location over time (Fig. 1B) and when biomass is observed at a single time step over the spatial dimension (Fig. 1C). Also, the patterns in biomass change; periods with high or low biomass become longer (Fig. 1B) and patches become larger towards the shift (Fig. 1C). These trends are captured by the early-warning signals used. The increase in variation in biomass towards the shift results in an increase in semivariance at all separation distances, both for the spatial and the temporal semivariance (Fig. 1D, E). The increase in patch size (in space and time), results in an increase in correlation length as shown by the variograms in Fig. 1E, and a decrease in the ratio of the semivariance at separation distance one over semivariance at separation distance 25, indicative for a change in correlation length (Fig. 1D). The results confirm the results found by (Dakos et al. 2009, Fernández and Fort 2009, Guttal and Jayaprakash 2009), who studied the same system, albeit with a different input error model, and without transient inputs.

Although semivariance values change well ahead of the transition as shown above, measurements of semivariance can only be used for forecasting the timing of a shift if semivariance observed at a particular time step is related to the timing of the shift. To explore this relation, we run the ecosystem model in Monte Carlo mode (Liu and Chen 1998) resulting in alternative realizations of the model. These realizations show a non-linear statistical dependence between the timing of the shift and the semivariance observed at a particular time step before the shift (cf. Appendix D). This is the case for each of the individual separation distances of the semivariance (spatial and temporal), and for the ratio of semivariance at separation distance 1 and 25. The statistical dependence is already strong at time step 500 and tends to increase towards the shift.



This exploratory data analysis shows that semi-variances are potential early-warning signals that can be used to forecast the timing of a shift. However, their predictive power does not only depend on the correlation with timing of the shift, but also on the uncertainty in the measured semivariance (due to observation and sampling error). This uncertainty is represented, and thus taken into account, by the covariance matrix \mathbf{R}_t (Eq. 2) that is used in the Bayesian forecast.

Sequential Bayesian forecasting

In order to mimic an ecosystem that is almost continuously monitored, we sample the synthetic world with a time interval of 300, resulting in observation times $t=200, 500, \dots, 1400$. For each observation time, 2500 samples of biomass are collected according to the sampling scheme described in the Methods section. We assume that the standard deviations of the systematic observation error σ_s and the random observation error σ_r are both equal to 0.1. At each observation time, the samples are used to calculate the spatial and the temporal semivariance at four separation intervals ($\bar{h} \approx 1, 6, 15$, and 25 grid cells), the mean biomass, and the covariance matrix \mathbf{R}_t (Eq. 2). Associated with an increase in semivariance values, the error variances of the semivariance values increase towards the shift (for details, see Appendix E). The error variances of mean biomass stay constant, because the mean biomass does not change much towards the shift. Errors are approximately Gaussian distributed, which is required for Eq. 2.

To evaluate the value of different early-warning signals in forecasting shifts, scenarios of different

Fig. 1. Synthetic data set. (A) Time series of biomass at a randomly chosen location (x, y) ; (B) idem, zoomed to three periods; (C) maps of biomass in periods shown in (B), i.e., at $t=200$ and $t=1800$; (D), time series of semivariance at separation distance 1 and 25, upper panel temporal semivariance, lower panel spatial semivariance; (E) semivariance values at four separation distances \bar{h} calculated at $t=200$ and $t=1800$; solid blue line, temporal semivariance; dotted red line, spatial semivariance.

combinations of observed early-warning signals are assimilated into the system model. These are (1) “temporal semivariance”, using observed temporal semivariance at 4 separation distances, (2) “spatial semivariance”, idem, using spatial semivariances, (3) “temporal and spatial semivariance”, using spatial and temporal semivariances, i.e., 8 semivariance values, (4) “mean”, using observed mean biomass, (5) “all information”, using observed temporal, spatial semivariance values and the mean biomass values, (6) “no observations”, without assimilation of measurements into the model, i.e., a free Monte Carlo run. The “no observations” scenario refers to a forecast that only relies on prior estimates of the model parameters and the system driver (grazing pressure).

Fig. 2 shows summary statistics calculated from individual realizations (Appendix F) of the system model. In the “no observations” scenario, the large uncertainty defined by the prior probability distributions of the system model and its parameters results in a large uncertainty in forecasted biomass and timing of the shift (Fig. 2A). This uncertainty is strongly reduced when observed temporal semivariance is assimilated into the model, even when the forecast is made at a moment well in advance of the shift, for instance at $t = 500$ (Fig. 2B). Particularly at the first observation time, at $t = 200$, the Bayesian update strongly reduces the uncertainty in biomass. As a result, the uncertainty in the timing of the forecasted shifts is reduced as well. The assimilation of observed semivariance into the system model results in a larger reduction of uncertainty compared to assimilation of the observed mean biomass. For the forecast made at $t = 500$ (Fig. 2, left panels), the temporal semivariance scenario results in a standard deviation in the timing of the shift of 204, while assimilating mean biomass results in a standard deviation of 365. Results for the spatial semivariance scenario (not shown) are comparable to those for temporal semivariance. The combination of semivariances and mean biomass further reduces uncertainty, with a standard deviation of 114 (Fig. 2E, left panel). The forecasts made at $t = 1400$ (Fig. 2, right panels) show smaller uncertainty compared to $t = 500$ forecasts due to the larger number of data assimilated (5·2500 samples versus 2·2500 samples for the forecast at $t = 500$) and a shorter lead-time.

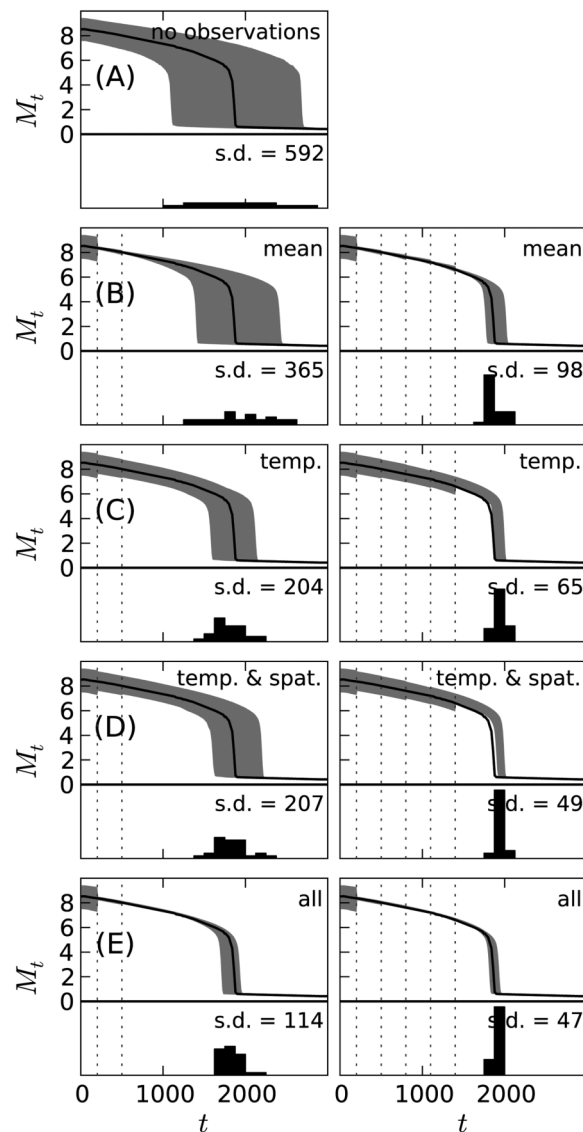


Fig. 2. Forecasted biomass, probability distributions and summary statistics. Top of each panel: Grey area, 80% confidence intervals of forecasted mean biomass M_t ; solid line, mean biomass of synthetic data set; vertical dashed lines, observation times. Bottom of each panel: probability density and standard deviation of forecasted timing of critical shift. (A) Forecast at $t = 0$, i.e., open Monte Carlo run without assimilation of observations into the model. (B–E, left panels) Forecast at $t = 500$, two observation times; (right panels) forecast at $t = 1400$, five observation times; for a description of the scenarios in each row, see main text. Probability densities are calculated with $N = 1000$.




Fig. 3. Effect of observation error and number of observations on standard deviation of the forecasted timing of the shift. Standard runs with $\sigma_s = \sigma_r = 0.1$; runs with increased systematic error $\sigma_s = 0.3$; runs with increased random error $\sigma_r = 0.3$; runs with fewer observations (17 samples per time step instead of 50) and standard observation error $\sigma_s = \sigma_r = 0.1$; runs with fewer samples (17 samples per time step instead of 50) and increased random error $\sigma_r = 0.3$. Left, forecast at $t = 500$, two observation times; right, forecast at $t = 1400$, five observation times. See main text for a description of the scenarios shown. Standard deviations are calculated with $N = 1000$. Full distributions of forecasted timing of the shift are shown in Fig. 2 (standard runs) and Appendix H.

Importantly, in addition to a reduced uncertainty in the forecasted timing of a shift, the assimilation of measured semivariates provides improved knowledge about the ecosystem itself, and its drivers. This is shown by the posterior probability densities of parameters and inputs, i.e., the probability densities found by assimilation of past-observed semivariance into the system model. The posterior probability densities are considerably narrower compared to the prior densities (for details, see Appendix G).

Effect of observation error and number of observations on forecast uncertainty

To evaluate the effect of observation error on forecast uncertainty, we run our scenarios using two different settings for the observation error. One is an increased systematic error, i.e., $\sigma_s = 0.3$ instead of 0.1. As the systematic error introduces the same bias to all biomass measurements, it does not affect the uncertainty in observed semivariance. However, the error variance of the observed mean biomass does increase, with

approximately 800% (for details, see Appendix E). As a result, the uncertainty in forecasted timing of the shift increases for the scenario that assimilates only mean biomass, compared to the original scenario (Fig. 3, see Appendix H for detailed results). The relative increase is largest when the forecast is made at $t = 1400$. The increased systematic error also increases the uncertainty of the forecasted timing for the “all information” scenario, because this also uses mean biomass. However, the effect is less pronounced because the all “information scenario” also uses observed semivariance.

Unlike increasing systematic error, increasing the random error mainly affects the uncertainty of observed semivariance. With a random error σ_r of 0.3 (instead of 0.1), the variance of the uncertainty in observed semivariates increases with approximately 1000% while the variance of observed mean biomass only increases with approximately 10% (Appendix E). As a result, an increase in random error gives a marked increase in forecast uncertainty for the scenarios

that assimilate semivariance into the model as shown in Fig. 3 (for details, see Appendix H). The scenario that uses only mean biomass is hardly affected by the increased random error.

To study the effect of the number of samples, we run the original scenarios with a number of 17 sampling locations per time step instead of 50, which was used in the original scenarios. The adjusted scheme was created by randomly selecting 14 samples at grid locations of the original scheme and 3 samples at random positions directly next to grid locations. The scheme results in $17 \cdot 50 = 850$ sampling locations at a particular observation time t . The observation error was kept the same as used in the standard scenario ($\sigma_s = \sigma_r = 0.1$). The effect of this scenario (Fig. 3, fewer samples run) on the forecast uncertainty is comparable to the scenario with larger random error.

The forecast uncertainty of the scenario that uses observed semivariance is further reduced when a smaller number of samples (17 sampling locations) and increased random error ($\sigma_r = 0.3$) is used (Fig. 3). With this sampling scheme and observation error, the advantage of using observed semivariance instead of mean observed biomass disappears, and the scenario that uses observed mean biomass has a forecast uncertainty that is considerably smaller than the scenario that uses observed semivariance.

DISCUSSION

Although the literature proposes a wide range of higher order statistics of system variables that show a marked increase well before an upcoming regime shift, the value of these early-warning signals for making quantitative inferences on the timing of a shift and its associated uncertainty remains largely untested. The limited number of studies (Biggs et al. 2009, Contamin and Ellison 2009) that evaluate early-warning signals for forecasting used deterministic ecosystem models with arbitrarily defined thresholds for early-warning indicators signaling a shift, without a theoretical framework to include observation and sampling errors of early-warning signals. Also, a number of current studies states that quantitative values of early-warning signals alone have limited value in diagnosing an impending shift, let alone the calculation of the timing of the shift

(Van Nes and Scheffer 2007, Dakos et al. 2009, Kéfi et al. 2010).

We have shown that the calculation of the absolute timing of the shift (and its uncertainty) is possible using observed early-warning signals. This is the case even when the ecosystem model is largely undefined in its parameters and drivers, as was the case in our runs. The results show that our assumed uncertainty in spatial dispersion of biomass (expressed by the probability distribution of the ecosystem model parameter d) does not hamper the forecast of the timing of the shift. This is notable, as spatial dispersion directly affects semivariance at a particular distance to a shift (Dakos et al. 2009, Fernández and Fort 2009), and this uncertainty could overshadow the trend in semivariance towards a shift, which is not the case here.

The use of spatial or temporal variance in biomass calculated from samples results in a forecast uncertainty of the timing of the shift that is low compared to the case when mean biomass is used. Thus, spatial and temporal variance can indeed be a better leading indicator of forthcoming shifts compared to measured mean state variables, i.e., mean biomass. The best results are found when both semivariance and mean biomass, also referred to as a composite indicator (Contamin and Ellison 2009), are used.

However, a large sample size with small observation error is required. The runs with increased random sampling error and/or a reduced number of samples show that the forecast uncertainty of the timing of the shift strongly depends on the random sampling error and the number of samples used. In general, the estimation of higher order statistics (e.g., semivariance) from samples is more difficult, i.e., results in larger uncertainty, compared to the estimation of first order statistics (de Gruijter et al. 2006). This is confirmed by our calculations on biomass (Appendix E). Thus, the lower uncertainty in forecasted timing of the shift obtained when using the semivariance compared to using the mean, disappears when random observation errors are large and sampling densities are too low. Changing systematic sampling error has the opposite effect of changing random sampling error: the uncertainty in forecasted timing of the shift shows a marked increase when mean biomass is assimilated, while the uncertainty for

the scenarios that use semivariance is not affected. So, the use of early-warning signals to forecast shifts is promising provided that random sampling errors are small, and large data sets are assimilated.

It is difficult to provide quantitative guidelines for the required sampling density to make early-warning signals useful in practice. This is because the effect of early-warning signals on forecast uncertainty depends on many factors, including the type of ecosystem studied, the magnitude of random disturbances of the system, the prior distribution of parameters and inputs including spatial variation in these values due to landscape heterogeneity, the type of early-warning signal used, the magnitude of random and systematic sampling error, the number of samples, and the sampling scheme. A number of these factors were fixed and taken spatially uniform in our study. We do however expect that the problem of high uncertainty in observed semivariance can be extrapolated to other situations, and our study provides a strong indication that dense sampling networks with low measurement error are required to make early-warning signals useful in forecasting shifts. With a smaller sample size, the use of mean state variables in forecasting the shift may give better results in terms of forecast uncertainty.

Our results show the relevance of studying early-warning signals in the context of sampling error. An indicator with a marked increase well ahead of an upcoming shift may still be relatively useless for the purpose of diagnosing an upcoming shift or forecasting its timing when observed values of the indicator have a high uncertainty. Unfortunately this is often the case because ecological sampling networks mostly consist of a limited number of samples. This is also important in the light of studies that test the hypothesis that early-warning signals exist in real world ecosystems. As these studies rely often on relative small sampling networks with inherent large observation errors, relationships between early-warning signals and distance to a shift may be overshadowed by the relatively high error in observed early-warning signals.

ACKNOWLEDGMENTS

We thank two anonymous reviewers for their

constructive comments.

LITERATURE CITED

- Ajami, N. K., Q. Duan, and S. Sorooshian. 2007. An integrated hydrologic Bayesian multimodel combination framework: Confronting input, parameter, and model structural uncertainty in hydrologic prediction. *Water Resources Research* 43:W01403. doi: 01410.01029/02005WR004745.
- Andersen, T., J. Carstensen, E. Hernández-García, and C. M. Duarte. 2009. Ecological thresholds and regime shifts: approaches to identification. *Trends in Ecology and Evolution* 24:49–57.
- Beven, K. 2002. Towards an alternative blueprint for a physically based digitally simulated hydrologic response modelling system. *Hydrological Processes* 16:189–206.
- Biggs, R., S. R. Carpenter, and W. A. Brock. 2009. Turning back from the brink: Detecting an impending regime shift in time to avert it. *Proceedings of the National Academy of Sciences of the United States of America* 106:826–831.
- Brock, W. A., and S. R. Carpenter. 2006. Variance as a leading indicator of regime shift in ecosystem services. *Ecology & Society* 11:9.
- Bulygina, N., and H. Gupta. 2010. How Bayesian data assimilation can be used to estimate the mathematical structure of a model. *Stochastic Environmental Research and Risk Assessment* 24:925–937.
- Carpenter, S. R., J. J. Cole, M. L. Pace, R. Batt, W. A. Brock, T. Cline, J. Coloso, J. R. Hodgson, J. F. Kitchell, D. A. Seekell, L. Smith, and B. Weidel. 2011. Early warnings of regime shifts: A whole-ecosystem experiment. *Science* 332:1079–1082.
- Chisholm, R. A., and E. Filotas. 2009. Critical slowing down as an indicator of transitions in two-species models. *Journal of Theoretical Biology* 257:142–149.
- Clark, M. P., D. Kavetski, and F. Fenicia. 2011. Pursuing the method of multiple working hypotheses for hydrological modeling. *Water Resources Research* 47:W09301. doi:09310.01029/02010WR009827.
- Contamin, R., and A. M. Ellison. 2009. Indicators of regime shifts in ecological systems: What do we need to know and when do we need to know it? *Ecological Applications* 19:799–816.
- Dakos, V., E. H. van Nes, R. Donangelo, H. Fort, and M. Scheffer. 2009. Spatial correlation as leading indicator of catastrophic shifts. *Theoretical Ecology* 3:163–175.
- de Groot, J. J., D. J. Brus, M. F. P. Bierkens, and M. Knotters. 2006. *Sampling for natural resource monitoring*. Springer, Berlin, Germany.
- Doucet, A., N. de Freitas, and N. Gordon. 2001. *Sequential Monte Carlo methods in practice*. Springer, New York, New York, USA.

- Drake, J. M., and B. D. Griffen. 2010. Early warning signals of extinction in deteriorating environments. *Nature* 467:456–459.
- Evensen, G. 2003. The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dynamics* 53:343–367.
- Fernández, A., and H. Fort. 2009. Catastrophic phase transitions and early warnings in a spatial ecological model. *Journal of Statistical Mechanics: Theory and Experiment* 2009:P09014.
- Goovaerts, P. 1997. Sampling for natural resources evaluation. Oxford University Press, Oxford, UK.
- Guttal, V., and C. Jayaprakash. 2009. Spatial variance and spatial skewness: Leading indicators of regime shifts in spatial ecological systems. *Theoretical Ecology* 2:3–12.
- Karssenberg, D., O. Schmitz, P. Salamon, K. De Jong, and M. F. P. Bierkens. 2010. A software framework for construction of process-based stochastic spatio-temporal models and data assimilation. *Environmental Modelling & Software* 25:489–502.
- Kéfi, S., M. Rietkerk, C. L. Alados, Y. Pueyo, V. P. Papanastasis, A. ElAich, and P. C. De Ruiter. 2007. Spatial vegetation patterns and imminent desertification in Mediterranean arid ecosystems. *Nature* 449:213–217.
- Kéfi, S., M. Rietkerk, M. Roy, A. Franc, P. C. De Ruiter, and M. Pascual. 2010. Robust scaling in ecosystems and the meltdown of patch size distributions before extinction. *Ecology Letters* 14:29–35.
- Kleinen, T., H. Held, and G. Petschel-Held. 2003. The potential role of spectral properties in detecting thresholds in the Earth system: Application to the thermohaline circulation. *Ocean Dynamics* 53:53–63.
- Kreyszig, E. 1999. Advanced engineering mathematics. Eighth edition. Wiley, New York, New York, USA.
- Liu, J., and M. West. 2001. Combined parameter and state estimation in simulation-based filtering. Pages 197–223 in A. Doucet, N. de Freitas, and N. Gordon, editors. Sequential Monte Carlo methods in practice. Springer, New York, New York, USA.
- Liu, J. S., and R. Chen. 1998. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* 93:1032–1044.
- Noy-Meir, I. 1975. Stability of grazing systems: an application of predator-prey graphs. *Journal of Ecology* 63:459–481.
- Scheffer, M., J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara. 2009. Early-warning signals for critical transitions. *Nature* 461:53–59.
- Simon, D. 2006. Optimal state estimation: Kalman, H infinity, and nonlinear approaches. Wiley-Interscience, Hoboken, New Jersey, USA.
- Takimoto, G. 2009. Early warning signals of demographic regime shifts in invading populations. *Population Ecology* 51:419–426.
- Van Leeuwen, P. J. 2009. Particle filtering in geophysical systems. *Monthly Weather Review* 137:4089–4114.
- Van Nes, E. H., and M. Scheffer. 2007. Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *American Naturalist* 169:738–747.
- Weerts, A. H., and G. Y. H. El Serafy. 2006. Particle filtering and ensemble Kalman filtering for state updating with hydrological conceptual rainfall-runoff models. *Water Resources Research* 42:W09403. doi:09410.01029/02005WR004093.

SUPPLEMENTAL MATERIAL

APPENDIX A

This appendix provides a concise description of the concepts and theory behind the Particle Filter, which is the sequential Bayesian data assimilation scheme used in this study. For details and reviews of particle filter approaches and applications, we refer to van Leeuwen (2009), Weerts and El Serafy (2006), Simon (2006), and Liu and West (2001). We used the software framework for data assimilation developed by Karssenberg et al. (2010).

Introduction

When forecasting the timing of a regime shift in a particular study area, we can rely on two sources of information: prior information and observational data. The prior information in-

cludes our understanding of how the ecosystem works, in particular how it changes over time leading to a possible regime shift. For quantitative predictions, this knowledge needs to be encoded in a dynamical (temporal) ecosystem model, for instance a vegetation growth model, as is used in our study. In addition, the prior information includes estimates of the magnitude and direction of change in drivers that push the ecosystem towards a regime shift. These drivers are represented as external inputs to the ecosystem model. An example of a driver would be grazing pressure, as is used here. When calculating the forecast of the timing of a regime shift, initial estimates, so called priors, are made of model parameters and drivers. Values of model parameters are estimated from literature or by calibration against historic data. Future changes

in drivers of the ecosystem are for instance estimated by extrapolating historic trends.

As these prior estimates rely on assumptions or uncertain data, they need adjustment using observational data of ecosystem state variables collected in the study area, in the period of time before a potential regime shift. Here, we use time series of observations of biomass. This is our second source of information for forecasting the timing of a regime shift. The temporal sampling interval and spatial resolution of the observations need to be sufficiently small to derive temporal changes in values of early-warning signals. Observations that can be used are for instance field samples of vegetation taken at a regular time interval, or time series of biomass derived from daily remote sensing images.

As both prior information (model parameters and drivers) and observational data collected in the study area are uncertain, they need to be described in probabilistic terms, i.e., using probability distributions for prior model parameters and observations. This allows weighting both sources of information when making the forecast of the timing of the shift. Here, prior information and observations are combined by a sequential Bayesian filter that numerically solves Bayes' Theorem. This results in updated values, referred to as posteriors, for model parameters and drivers, and a forecast of the timing of the shift and its uncertainty, calculated using all available information and observations. The approach is to calculate this forecast using different sets of observational data, i.e., either using early-warning signals (i.e., higher order statistics) calculated from samples of ecosystem state variables or using mean values of ecosystem state variables as observational data. If the forecast uncertainty is indeed lowest when early-warning signals are used, our study confirms the assumed value of early-warning signals to predict regime shifts.

Sequential Bayesian filtering

Sequential Bayesian filtering combines two sources of information: prior knowledge of the system and observational data of the state of the system. These sources of information are combined by using a numerical solution of the Bayes' theorem. The prior knowledge of the system is given by a forward stochastic model:

$$\mathbf{x}_t = \mathbf{f}(\mathbf{x}_{t-1}) \quad (\text{A.1})$$

In Eq. A.1, \mathbf{x}_t is the full state of the system, described by a joint probability density of all variables (state, parameters, drivers), denoted by $p(\mathbf{x}_t)$, with the time index $t, t = 0, 1, \dots, T$, and \mathbf{f} is the system transition function that represents real-world processes. Eq. A.1 of this appendix is represented by Eqs. 2–3 in the main text of the article. Observational data are collected by monitoring the real state of the system. This is done by sampling one or more state variables of the system. Observations \mathbf{y}_t are taken at observation times t_m , with t_m a subset of t .

These sources of information are combined as follows, as illustrated in Fig. A1A. At $t=0$, which is generally a moment just before the first observation time, the system model is initialized with prior probability density functions of state variables, parameters, and drivers. The system is integrated forward in time with the system model (Eq. A.1) up to the first observation time. At this time step, the observational data are used to adjust the full state of the model. From this adjusted state, the system is propagated forward in time with the system model, up to the next observation time where a similar adjustment of the full state is made. This procedure is followed up to the last time step T .

The adjustment of the full state of the model at an observation time is done by solving Bayes' theorem:

$$p(\mathbf{x}_t|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t)}{p(\mathbf{y}_t)}. \quad (\text{A.2})$$

In Eq. A.2, $p(\mathbf{x}_t)$ is the joint probability density of the system variables at t calculated by propagating the system model from the previous observation time up to the current observation time. As it refers to the density at the observation time before assimilating observations collected at the observation time, it is referred to as the prior probability density at the observation time. The term $p(\mathbf{y}_t|\mathbf{x}_t)$ is the joint probability density of the observations at t given the model state, also referred to as the likelihood of the observations given the model. The term $p(\mathbf{y}_t)$ is the joint probability density function of the observations at t . The term $p(\mathbf{x}_t|\mathbf{y}_t)$ is the adjusted joint probability density of the model. As it refers to the density at the observation time after assim-

ilating observations collected at the observation time, it is referred to as the posterior probability density.

The Particle Filter solves this sequential Bayesian scheme with Monte Carlo simulation (Fig. A1B). At $t=0$, a number of N independent model realizations (also referred to as particles) $\mathbf{x}_t^{(n)}$, $n = 1, 2, \dots, N$, are sampled from the initial joint probability density $p(\mathbf{x}_0)$ of the model. The joint probability density of the model is represented by these realizations as:

$$p(\mathbf{x}_t) = \sum_{n=1}^N \delta(\mathbf{x}_t - \mathbf{x}_t^{(n)}) \quad (\text{A.3})$$

with $\delta(\cdot)$ the Dirac delta function. Each realization is integrated forward in time up to the first observation time. At the observation time, Bayes' theorem is evaluated for each realization:

$$p(\mathbf{x}_t^{(i)}|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t^{(i)})p(\mathbf{x}_t^{(i)})}{p(\mathbf{y}_t)},$$

for each $i = 1, 2, \dots, N$. (A.4)

Since

$$p(\mathbf{y}_t) = \sum_{j=1}^N p(\mathbf{y}_t|\mathbf{x}_t^{(j)})p(\mathbf{x}_t^{(j)}) \quad (\text{A.5})$$

we can rewrite Eq. A.4 as:

$$p(\mathbf{x}_t^{(i)}|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t^{(i)})p(\mathbf{x}_t^{(i)})}{\sum_{j=1}^N p(\mathbf{y}_t|\mathbf{x}_t^{(j)})p(\mathbf{x}_t^{(j)})},$$

each $i = 1, 2, \dots, N$. (A.6)

As shown hereafter, because realizations are re-sampled after each observation time, the prior probability $p(\mathbf{x}_t^{(i)}) = 1/N$, for all i , $i = 1, 2, \dots, N$. Therefore Eq. A.6 reduces to:

$$p(\mathbf{x}_t^{(i)}|\mathbf{y}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t^{(i)})}{\sum_{j=1}^N p(\mathbf{y}_t|\mathbf{x}_t^{(j)})} \text{ for each realization}$$

$i = 1, 2, \dots, N$. (A.7)

Under the assumption of a Gaussian distributed observation error in the observations, the likelihood of the observations given the model takes the form:

$$p(\mathbf{y}_t|\mathbf{x}_t^{(i)}) \sim \exp\left(-\frac{1}{2}[\mathbf{y}_t - \mathbf{H}_t(\mathbf{x}_t^{(i)})]^T \mathbf{R}_t^{-1} [\mathbf{y}_t - \mathbf{H}_t(\mathbf{x}_t^{(i)})]\right).$$

(A.8)

By combining Eqs. A.7 and A.8, $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$, i.e., the posterior probability of a realization given the observations, can be calculated for each realization. In the literature, $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$ is often referred to as the weight of a particle or realization.

The posterior probabilities of the realizations can in principle be used to calculate the posterior full state of the model, and the same procedure could be applied for each observation time. However, this will almost always lead to posterior probabilities of realizations that differ more and more through time. A small number of realizations will receive high posterior probabilities while a large number will have a posterior probability close to zero. As a result, the probability density of the full state is not properly represented anymore. This is called particle collapse or filter degeneracy. A number of methods have been developed to avoid filter degeneracy. We selected a so-called resampling technique because of its straightforward implementation and proven effectiveness. After the calculation of the posterior probabilities at observation time, the realizations are resampled in order to retrieve a new set of realizations with equal posterior probabilities. This set of realizations is propagated using the system model to the next observation time (Fig. A1B). A number of different approaches exist to do the resampling step. Each approach represents the intuitive idea that the number of copies of a realization should be approximately proportional to $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$. In the most basic approach, a cumulative distribution function is constructed from the posterior probabilities $p(\mathbf{x}_t^{(i)}|\mathbf{y}_t)$. From this distribution, N samples are randomly drawn with replacement using a uniform distribution between zero and one. If a limited number of particles are used, this approach has the disadvantage that sampling noise is introduced. This is why we selected Residual Resampling (Liu and Chen 1998, Weerts and El Serafy 2006, Van Leeuwen 2009), which copies samples in two steps. In the first step, a sample is copied a number of times equal to $k_t^{(i)} = \text{round}(p(\mathbf{x}_t^{(i)}|\mathbf{y}_t) \cdot N)$, where $\text{round}(x)$ is a




Fig. A1. (A) Sequential data assimilation for a complete model run. Line, median; grey area, confidence interval; vertical dashed lines, observation times. (B) Solution of Bayes' theorem at an observation time. Curves, realizations; dots, probabilities of realizations represented by size. (B) is modified after van Leeuwen (2009).

Table A1. Number of unique realizations at the end of a model run.

	$t = 500$				$t = 1400$			
Scenario	S	ISE	IRE	FS	S	ISE	IRE	FS
No observations	1000	1000	1000	1000	1000	1000	1000	1000
Mean	139	362	149	149	35	189	42	44
Temporal semivariance	106	105	232	193	34	33	98	83
Spatial semivariance	100	107	230	203	31	31	102	87
Temporal and spatial semivariance	87	87	219	185	21	21	81	67
All information	19	50	54	49	6	9	12	11

Notes: Rows give scenarios of early warning signals that are assimilated, as described in the main text. Columns give the scenarios; S, standard observation error ($\sigma_s = \sigma_r = 0.1$); ISE, increased systematic error $\sigma_s = 0.3$; IRE, increased random error $\sigma_r = 0.3$; FS, fewer samples (17 samples per time step instead of 50) with standard observation error ($\sigma_s = \sigma_r = 0.1$). Number of unique realizations are given for the forecast made at $t = 500$ and $t = 1400$.

function that rounds to the nearest integer toward zero. In the second step, residual probabilities $r_t^{(i)}$ are calculated:

$$r_t^{(i)} = \frac{p(\mathbf{x}_t^{(i)} | \mathbf{y}_t) \cdot N - k_t^{(i)}}{N - \sum_{j=1}^N k_t^{(j)}}. \quad (\text{A.9})$$

A number of $N - \sum_{j=1}^N k_t^{(j)}$ additional copies of samples are made using the residual probabilities $r_t^{(i)}$. This is done in a similar way as in the basic resampling technique by uniform sampling from a cumulative distribution function constructed now from the residual weights $r_t^{(i)}$. The samples copied in step one and two result again in N samples representing the posterior probability density of the full model state.

A potential well-known problem of the particle filter is that the number of unique realizations (particles) that is propagated through time becomes too small to provide a precise representation of the probability density of the system.

Fig. F1 shows that a large number of unique realizations is retained in the system, and for all scenarios the number of unique realizations is sufficient to approximate probability distributions (Table A1).

APPENDIX B

Details regarding the ecosystem model

Eq. 3 is solved for each time step $t = 1, 2, \dots, 3500$, using a time-explicit scheme with 10 integration steps per time step. We compared this scheme with the implicit Heun's method (Kreyszig 1999) and differences were negligible at the input ranges used. The model boundaries in space were closed to avoid spatial correlation across periodic boundaries that would be interpreted as variation over large distances in the spatial correlation analysis. A number of model runs with periodic boundary conditions showed negligible differences in the results.

APPENDIX C




Fig. C1. Schematic of sampling scheme for sampling the synthetic reality. A number of 40 samples is taken on a regular grid with 10 samples at a distance of one grid cell away from regular grid locations. For calculation of \mathbf{R} in the Monte Carlo scheme, we randomly positioned the sampling grid on the modelling area for each realization. This was done by drawing random offset distances O_x and O_y from a uniform distribution such that the sample grid always falls within the model area. The schematic shows one of these realizations.

APPENDIX D

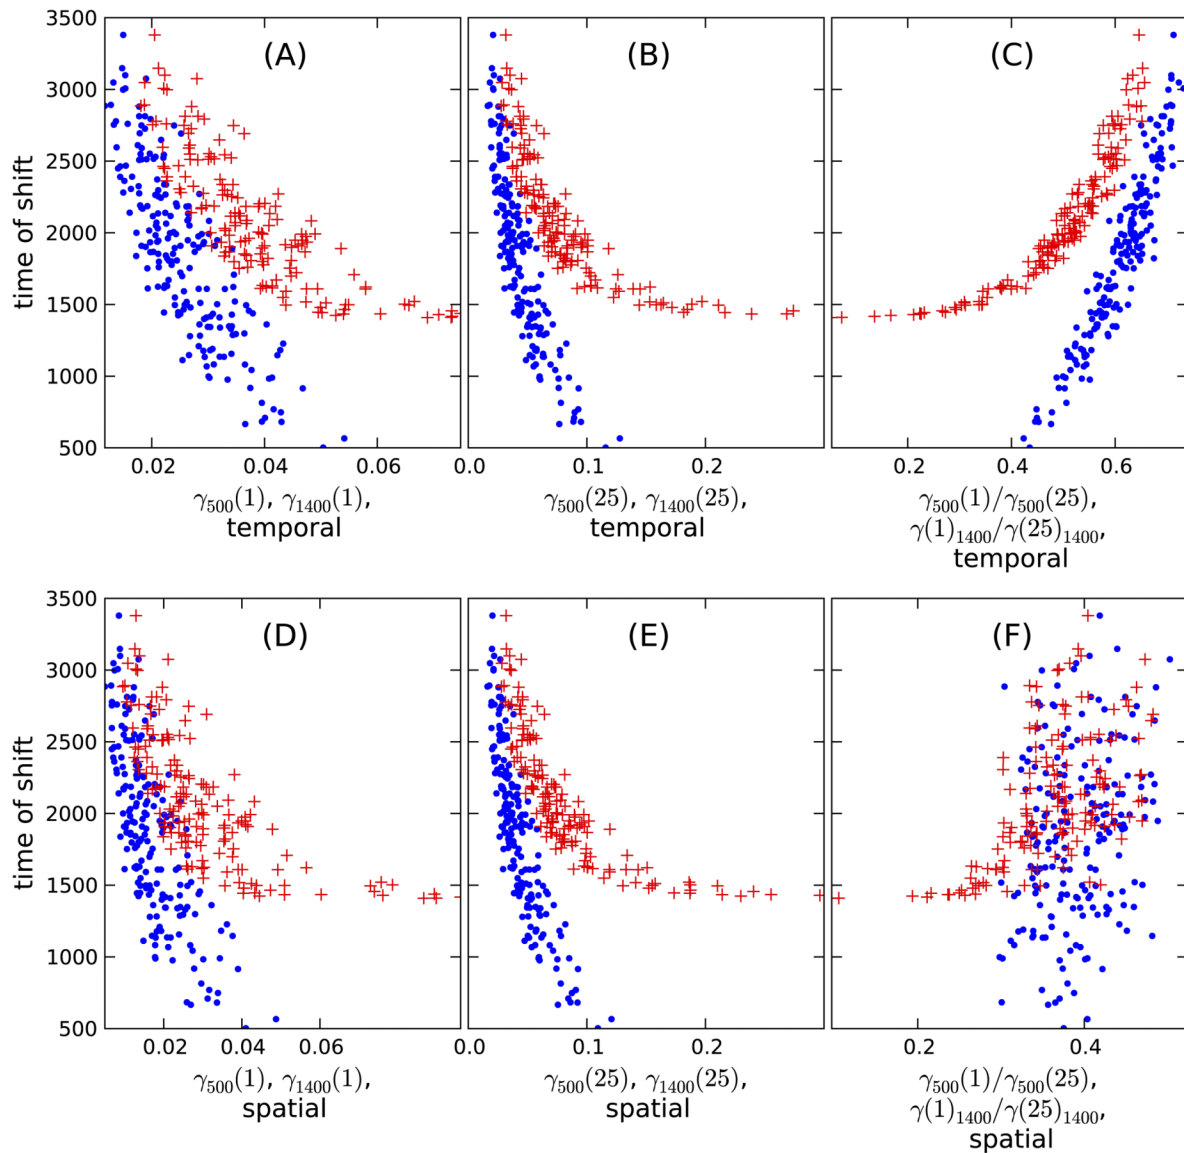


Fig. D1. Relation between time of critical shift and early-warning signals. $N = 200$ realizations of the ecosystem model created with prior probability density distributions given in the main text. Dots give values at $t = 500$ (e.g., $\gamma_{500}(1)$ in panel A); crosses at $t = 1400$. (A) Temporal semivariance at separation distance 1; (B) idem, separation distance 25; (C) ratio of temporal semivariance at separation distance 1 and 25; (D–F), same as (A–C), but spatial semivariance.

APPENDIX E




Fig. E1. Error variance and covariance of observed temporal semivariance and measured mean biomass. To show the trend over time, the values are plotted for measurement times $t = 50, 100, \dots, 2000$. Note that the two sampling schemes used in the paper use only the values at measurement times $t = 200, 500$ and $t = 200, 500, \dots, 1400$, respectively. In each panel, the three lines represent: S, main (standard) scenario with standard deviation of systematic observation error $\sigma_s = 0.1$ and standard deviation of the random observation error $\sigma_r = 0.1$; ISE, increased systematic observation error scenario with $\sigma_s = 0.3$, $\sigma_r = 0.1$; IRE, increased random observation error scenario with $\sigma_s = 0.1$, $\sigma_r = 0.3$. (A), error variance of temporal semivariance at separation distance 1, (B) error variance of temporal semivariance at separation distance 25, (C) error covariance between temporal semivariance at separation distance 1 and 25, (D) error variance of mean biomass. Error variances and covariances for spatial semivariance, and other separation intervals show comparable trends (not shown). The observation errors of the semivariance values increase towards the shift, which can be explained by an increase in mean semivariance. The error variances of mean biomass, however, stay constant. Errors are approximately Gaussian distributed, which is required for Eq. 2 (main text).

APPENDIX F




Fig. F1. Realizations of the system model. (A) Forecast at $t=0$, i.e., open Monte Carlo run without assimilation of observations into the model. (B) Forecast at $t=500$, i.e., assimilation of observed temporal semivariance (temporal scenario) at observation times $t=200$ and $t=500$. Left, temporal semivariance at separation distance 1; centre, temporal semivariance at separation distance 25; right, mean biomass. $N=1000$. Thick blue line: synthetic data set from which measurements are taken in the data assimilation runs in panel B. Vertical dashed lines: observation times.

APPENDIX G

The posterior probability densities of model parameters and drivers (Figs. G1 and G2) are considerably narrower compared to the prior densities. The ranking in the performance of the different scenarios is comparable to what we found for the uncertainty in the forecasted timing

of the shift, with the sampled semivariance resulting in lower uncertainties than using mean biomass. Parameter and input uncertainty reduces when a larger number of observations is assimilated, as shown by the difference in the results between $t = 500$ and $t = 1400$.

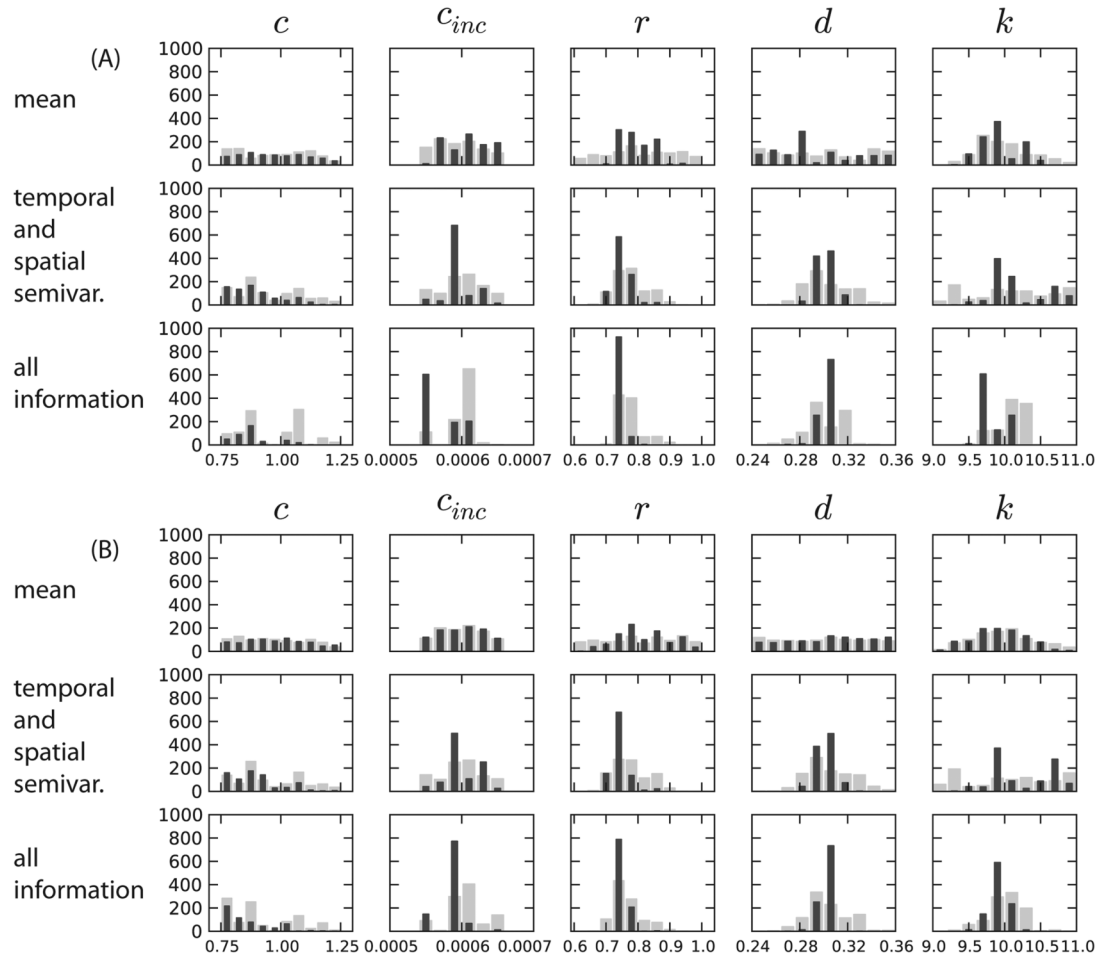


Fig. G1. (A) Runs with standard observation error $\sigma_s = \sigma_r = 0.1$; (B) Runs with increased systematic observation error $\sigma_s = 0.3$. The plots show initial grazing pressure c , expected future increase in grazing pressure c_{inc} , growth rate r , dispersion rate parameter d , and carrying capacity k . Grey bars, posteriors at $t = 500$, two observation times, black bars, posteriors at $t = 1400$, five observation times. Upper and lower limits of the x-axis correspond to bounds of the uniform prior distributions used. See main text for a description of the scenarios in each row. Probability densities are calculated with $N = 1000$.

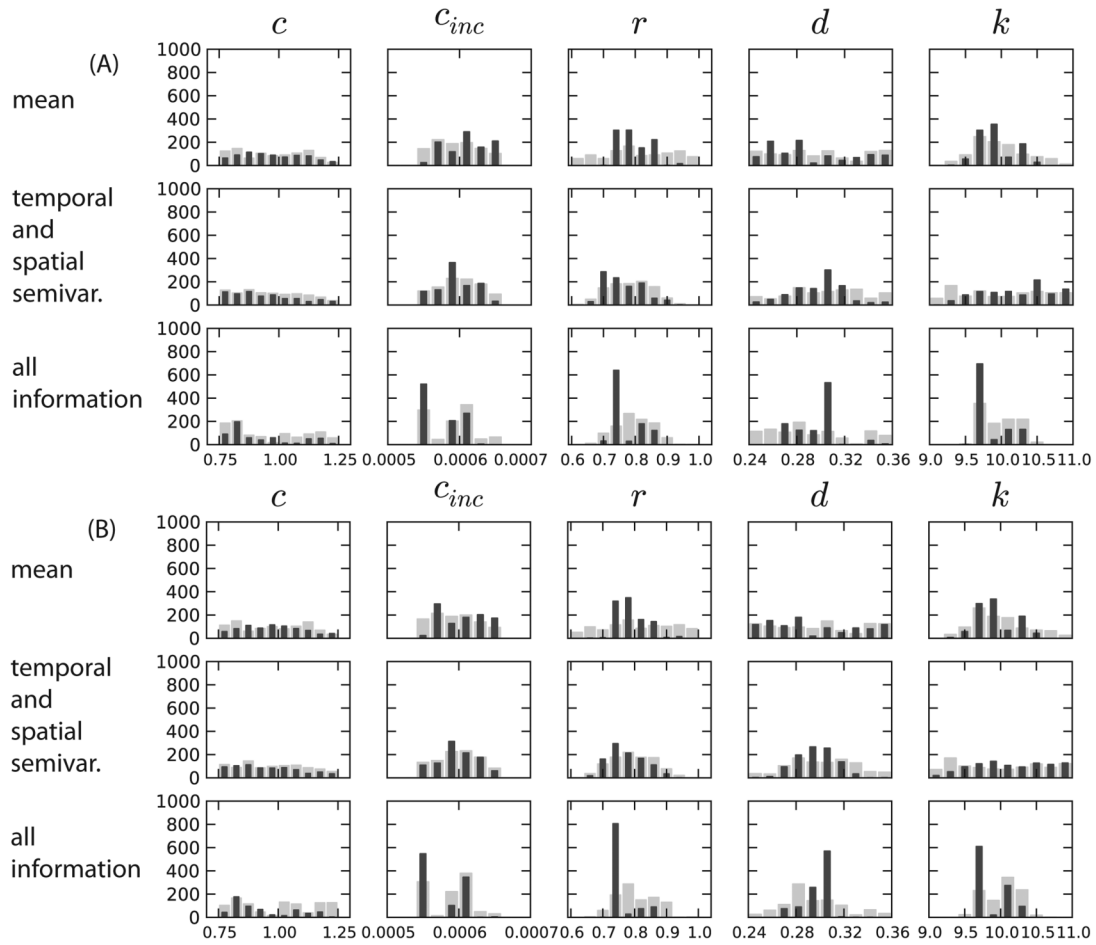


Fig. G2. (A) Runs with increased random observation error $\sigma_r = 0.3$. (B) Runs with fewer samples (17 samples per time step instead of 50) and standard observation error $\sigma_s = \sigma_r = 0.1$. For explanation see Fig. G1.

APPENDIX H

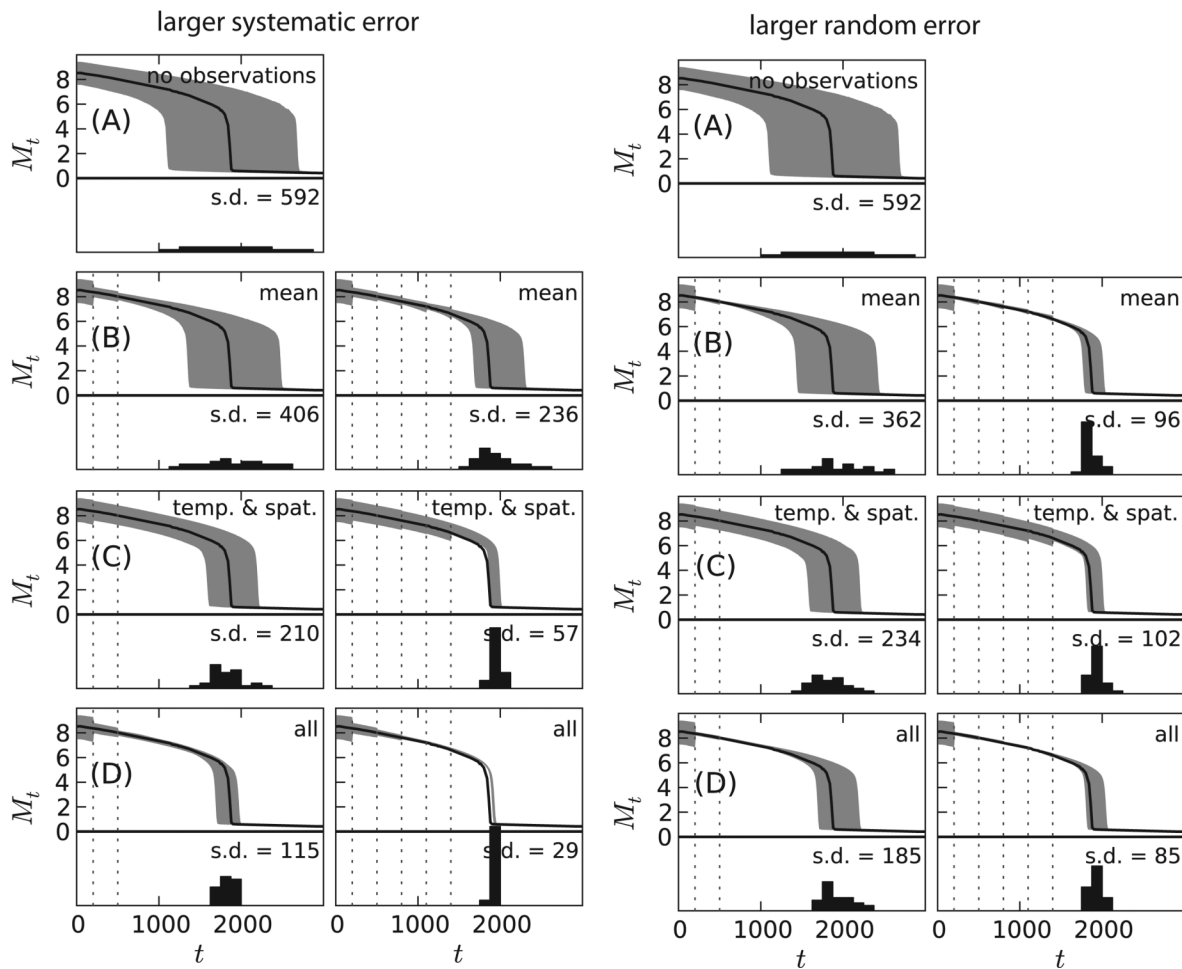


Fig. H1. Forecasted biomass, probability distributions and summary statistics. Larger systematic error, runs with increased systematic error $\sigma_s = 0.3$; larger random error, runs with increased random error $\sigma_r = 0.3$. Top of each panel: Grey area, 80% confidence intervals of forecasted mean biomass M_t ; solid line curve, mean biomass of synthetic data set; vertical dashed lines, measurement times. Bottom of each panel: probability distribution and standard deviation of timing of critical shift. (A) Forecast at $t = 0$, i.e., open loop Monte Carlo run without assimilation of observations into the model; (B–D, left panels) forecast at $t = 500$, two observation times; (right panels) forecast at $t = 1400$, five observation times; (B) “mean”; (C) “temporal and spatial semivariance”; and (D) “all information” scenario. Probability densities are calculated with $N = 1000$.

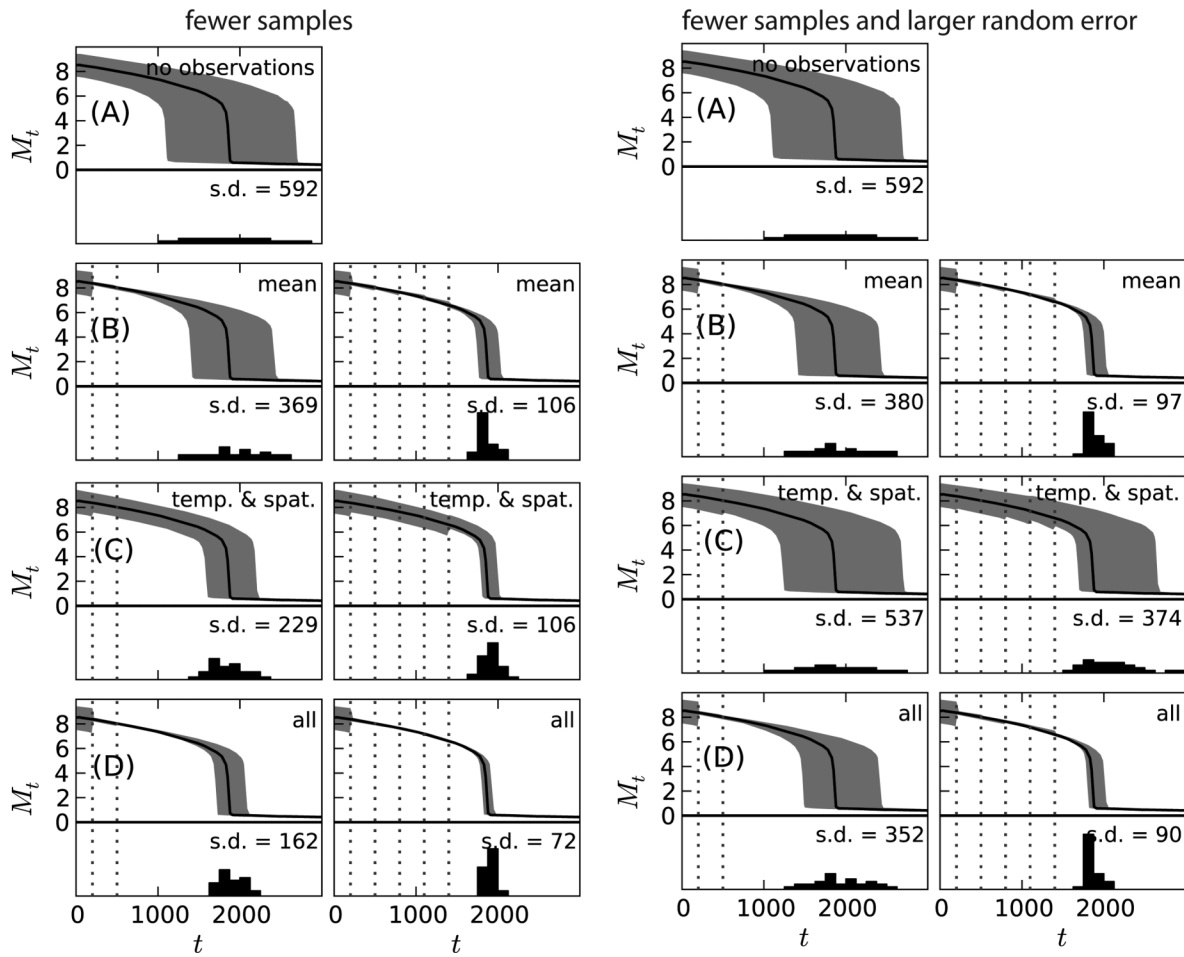


Fig. H2. Forecasted biomass, probability distributions and summary statistics. Fewer samples (17 samples per time step instead of 50) with standard observation error ($\sigma_s = \sigma_r = 0.1$); fewer samples (17 samples per time step instead of 50) with increased random error $\sigma_r = 0.3$. See Fig. H1 for explanation.