

PHENOTYPIC EVOLUTION STUDIED BY LAYERED STOCHASTIC DIFFERENTIAL EQUATIONS

BY TROND REITAN^{*}, TORE SCHWEDER[†] AND JORIJTJE HENDERIKS^{‡*}

Time series of cell size evolution in unicellular marine algae (division Haptophyta; *Coccolithus* lineage), covering 57 million years, are studied by a system of linear stochastic differential equations of hierarchical structure. The data consists of size measurements of fossilized calcite platelets (coccoliths) that cover the living cell, found in deep-sea sediment cores from six sites in the world oceans and dated to irregularly points in time. To accommodate biological theory of populations tracking their fitness optima, and to allow potentially interpretable correlations in time and space, the model framework allows for an upper layer of partially observed site-specific population means, a layer of site-specific theoretical fitness optima and a bottom layer representing environmental and ecological processes. While the modeled process has many components, it is Gaussian and analytically tractable. A total of 710 model specifications within this framework are compared and inference is drawn with respect to model structure, evolutionary speed and the effect of global temperature.

1. Introduction.

Biological populations evolve with respect to the distribution of organism size and other phenotypic traits by differential fitness. How a phenotypic character like body size evolves over time and what environmental factors influence phenotypic change, are fundamental questions of biology and paleontology. Some key issues include: Is the average size fluctuating around a fixed or a changing fitness optimum [3]? How fast does the population mean track the optimum? What is the inertia of the fitness optimum, and is it tracking a deeper process representing environmental conditions? If there is a deeper process, what qualities does it have? Are population-specific optima related, for example by influence from an unobservable underlying global climatic process? Alternatively, are the phenotypic means at dif-

^{*}CEES, Dept. of Biology, University of Oslo, P. O. Box 1053 Blindern, N-0316 Oslo, Norway. E-mail: trondr@bio.uio.no

[†]Dept. of Economics, University of Oslo, P. O. Box 1095 Blindern, N-0316 Oslo, Norway. E-mail: tore.schweder@econ.uio.no

[‡]Dept. of Earth Sciences, Uppsala University, Villavägen 16, SE-75 236 Uppsala, Sweden. E-mail: jorijntje.henderiks@geo.uu.se

AMS 2000 subject classifications: Time series, Latent processes

Keywords and phrases: Causal model, Coccolith, Fossil data, Latent processes, Time series, Ornstein-Uhlenbeck process

ferent sites correlated directly? Is the optimum responding to global temperature?

Several recent studies have highlighted a covariance between the body size evolution of marine organisms and global temperature. The long-term cooling trend over the past 65 million years appears to be matched by a macroevolutionary size decrease in marine phytoplankton [4] while marine zooplankton and benthos show size increases [17, 14]. With the exception of [14], which only used random walks as a time series model, these studies did however not take time dependency into account when testing this hypothesis. Without consideration of the internal dynamics of two trended time series, the null hypothesis of no interaction between the series may be wrongfully rejected. Here, we study the evolution of body size in marine algae by means of layered stochastic differential equation (SDE) models.

In unicellular algae, cell size and its geometry largely determine the transport rates of dissolved components (e.g. CO_2/O_2 , nutrients) into and out of the cell, which is fundamental to photosynthesis and growth rates. We study the evolution in average log-phenotype of calcifying microalgae belonging to the *Coccolithus* Schwartz 1894 lineage. The *Coccolithus* lineage (division Haptophyta) has extant species in today's oceans and a well-documented fossil record dating back to the early Paleocene [8]. We study coccolith size, which is measured by the largest diameter in the elliptically shaped calcite platelets (coccoliths) that cover the unicellular organism. Coccolith size is used as a proxy for cell size [10]. A total of 205 deep-sea sediment samples, taken from 6 different sites in the Atlantic, Indian and Pacific oceans, offer a final dataset of 19,899 individual size measurements distributed over the last 57 million years (My) (Fig. 1). To current standards of palaeontology and evolutionary biology these data are unusually extensive.

The population of algae behind sampled coccoliths from one site is assumed to track its fitness optimum by natural selection. Fitness (expected number of reproducing offspring) is distributed over the individuals of the population, with the optimum varying over time according to physical and ecological conditions. The fitness optima of different populations may be spatially correlated, due to common environmental conditions, in addition to be temporally correlated within a population. What affects fitness with respect to coccolith size is unknown. Global mean ocean surface temperature might be a contributor, and a measured temperature indicator series [22] is tested in the model framework as potentially driving the fitness optimum processes along with other underlying but unmeasured processes.

In this study, we will work on the logarithmic scale, as untransformed coccolith sizes cannot be normally distributed. Mean logarithmic coccolith size for each population is seen as a stochastic process pulled toward an optimal state, which itself is seen as a stochastic process subject to pull from additional underlying processes. These underlying processes might be population specific and possibly correlated across populations. The model framework thus allows for three hierarchical layers

of processes, each layer having one process for each geographical location. In the building of our framework, the three layers are the population means, population fitness optima, and underlying environmental processes affecting fitness such as an observed global temperature indicator series and also unmeasured environmental variables.

As the data are irregularly distributed in time, a stochastic time series framework that can handle continuous time is called for. Linear SDE models constitute a parsimonious framework for such modeling. Vector processes governed by linear SDEs, as defined in section 3, are suitable for hierarchical models with variables ordered by the flow of causality as determined by the coefficients in the equations, see [20]. Linear vectorial SDEs are analytically tractable, having an explicit likelihood.

A collection of 710 variants of models within this framework were fitted to the sample means at the six locations. These model variants are compared and properties of the evolutionary process are discussed by way of commonalities in models which are singled out by various statistical criteria. We use both Bayesian and classic methods.

Lande [15] was the first to suggest that the population mean of a phenotypic trait, e.g. logarithmic coccolith size, might evolve like an Ornstein-Uhlenbeck (OU) process. This was taken further by Estes and Arnold [3]. The Lande model [15] for the evolution of the mean phenotype, such as logarithmic size, is governed by the one-dimensional linear SDE,

$$(1) \quad dX(t) = -\alpha(X(t) - \mu_0)dt + \sigma dW(t),$$

where $dW(t)$ is white noise and t is time. This is an OU process when $\alpha > 0$. The mean phenotype X is pulled toward the level μ_0 with a force proportional to its displacement $X(t) - \mu_0$ both when having been pushed by randomness above or below its long-term level μ_0 . This level is the fitness optimum in [15]. We will use the term pull for the parameter α and diffusion for σ . Phenotypic evolution is also modeled by OU processes in [6], though in a phylogenetic rather than time series perspective. So do Hansen and co-authors [7], where the level process $\mu(t)$ is defined to be an underlying random walk. A random walk model, which can be seen as a limiting case of Eq. (1) with $\alpha = 0$ has been studied by [12, 13]. Random walks have been in use for a long time, as a proposed null hypothesis in evolutionary models, see for instance [16]. SDE models have been around for years, and have been used in biology (see chap. 8 of [1]) and physics (see [18] chap. 2).

Our aim is to develop a class of models that is capable of describing the data reasonably well, that allows biological interpretations and that is statistically feasible. These models might also be useful elsewhere, especially for time series with irregular temporal resolution.

In section 3 we will first introduce the linear SDE framework tailored for our application. We will then briefly discuss issues of causality and hierarchy. Model selection and inference is described in section 4. These methods are then applied first to artificial data in section 5 and then to the coccolith data in section 6. Some concluding remarks are offered in section 7.

2. The coccolith data.

Microfossils (coccoliths) were measured in a total of 205 sediment samples obtained from Deep Sea Drilling Project (DSDP) and Ocean Drilling Program (ODP) deep-sea sediment cores, taken at six sites in the Atlantic, Indian and Pacific oceans, altogether spanning about 57 My, see Fig. 1.

In each sample, 1 to 400 individual coccoliths were measured on slides using polarized light microscopy [9, 10], resulting in a mean sample size of $19,899/205=97$. The data was transformed from original size in μm to the logarithm of that, before means and variances were calculated. The age of each sample is estimated from biostratigraphic data calibrated to the geological time scale of [2].

While the coccoliths within a single sample may have been formed as much as one thousand years apart, we consider them simultaneously formed on our geological time scale. The sampling process is assumed random, from the historical population through deposition fossilization, drilling and extraction from the drill cores.

The stationary distribution of mean log size is normal in our model. Similarly any lack of normality in the samples should give little cause for concern when modeling phenotypic means by sample means, due to the central limit theorem. In the tradition following Lande [15], we study the evolution of the population mean of the phenotype in question, which in our case is the logarithmic body size. Population variance and other population processes and parameters are of potential interest, but for our study these are nuisance parameters. To keep the complexity of our model within bounds dictated by the data, the population variances are estimated outside the SDE framework and smoothed by a simple GAM analysis so as to reduce the variance uncertainty for samples with a low number of measurements.

We assume the sample means to be normally distributed with known variances, see Eq. (2) in the next section. The 19,899 separate measurements are thus summarized by the 205 sample means. The assumption of normality and known variances is not strictly true, but not far off, and it allows the likelihood to be computed.

To diagnostically check the assumption that the sample means of log coccolith size are normally distributed, skewness and kurtosis in the distributions of the sample means were estimated by bootstrapping for the 192 samples with more than 4 observations, and tested. The p-value plots (see [19]) show only small depar-

tures from normality, with some sample means having a bit of positive skewness, but with no trace of kurtosis. See the supplementary information for more on the normality of the data, the treatment of the measurement variances and other data issues.

3. Layered linear SDE models.

3.1. *The modeling framework.* If we concentrate on a single population, the idea is to connect measurements from different times together using a continuous time series model. Thus the irregular time intervals between observations will not pose a problem.

As described in the introduction, we aim at describing the phenotypic mean as a process responding to another continuous time process which is hidden, namely the size at optimal fitness, which may again respond to environmental changes, partially or entirely unmeasured. Thus the equation describing the dynamics in phenotypic mean (layer one) will not only contain the phenotypic mean, but also a hidden underlying process representing optimal fitness (layer two). Layers are in general processes defined so that one layer can respond to the current state in another layer, which is then said to be below it. The layer may then in its turn affect the layer above it, unless it is the topmost layer.

In our framework, layer two can again be affected by a third layer, understood as the unmeasured environmental conditions. The second layer might also be forced by an external global temperature indicator time series on this layer, as measured by [22].

In an SDE, the change in continuous time series from one time point to another time point infinitesimally further along, is modeled by a transformation of the previous state plus some normally distributed noise. When this transformation is linear and the observational noise is assumed normal and not conditioned on the state, we call this a (vectorial) linear SDE. When one has a linear SDE system and the state at the starting point of the process is normally distributed, then both the marginal distribution and the conditional distribution of the state of a later time point is also normally distributed. When also the sampling errors are normal, a likelihood can be calculated. These conditions provides the modeling framework for this study.

A sample mean Y_t for time t is thus considered a noisy representative of the state of the topmost layer $X_1(t)$, namely the phenotypic population mean at the same time point. So

$$(2) \quad Y_t \sim N(X_1(t), s_t^2/n_t),$$

where s_t^2 is the sample variance and n_t is the sample size. The sampling errors $Y_t - X_1(t)$ are assumed independent. In the language of hidden Markov models, Eq. (2) will be the observation equation while the system equation will be described in the next subsections. Since the state process is a normally distributed linear hidden Markov chain and the observations are normal and independent given the state process, the Kalman filter can be applied for calculating the likelihood, see the supplementary material.

As there are six geographic sites, instead of a single process per layer, we have six components in each layer. To allow for possible instantaneous commonalities in a layer, there might be instantaneous correlations across sites in the layer. We allow only correlations between processes belonging to the same layer.

Collected over the 6 sites, the state of the system is situated in a 18-dimensional vector space and evolves according to a linear hierarchical system of SDEs.

3.2. Linear SDE Basics and examples. The OU process, described in Eq. (1) is a simple linear SDE process. It is stationary when $\alpha > 0$. The stationary distribution is normal with expectation $E(X(t)) = \mu_0$ and standard deviation $\eta \equiv sd(X(t)) = \sigma/\sqrt{2\alpha}$. This can be verified by Eq. (5) later in the text. The stationary correlation is $corr(X(t_1), X(t_2)) = e^{-\alpha(t_2-t_1)} = e^{-(t_2-t_1)/\Delta t}$, $t_1 < t_2$, where $\Delta t \equiv 1/\alpha$ gives a characteristic correlation time for the process, the time for the correlation to drop to 1/e. Such alternate parametrizations, using for instance Δt and η instead of the pull α and the diffusion σ , can make the interpretation of results easier, as well as help in the elicitation of Bayesian priors. Thus with three parameters, one can parsimoniously describe a continuous function of time that is a stationary process.

When the level term μ_0 in Eq. (1) is replaced by a process $X_2(t)$ of the OU type, one gets a coupled linear SDE:

$$(3) \quad \begin{aligned} dX_1(t) &= -\alpha_1(X_1(t) - X_2(t))dt + \sigma_1dW_1(t) \\ dX_2(t) &= -\alpha_2(X_2(t) - \mu_0)dt + \sigma_2dW_2(t), \end{aligned}$$

where X_1 is the topmost, partially measured process layer. It is driven by X_2 , which is a hidden process. Since X_2 is unaffected by X_1 , it is an OU process on its own. X_1 is however dynamically affected by X_2 and have a different time correlation function from that of an OU process, as we will show later. This causal structure is described by $X_2 \rightarrow X_1$. When this causality takes the form found in Eq. (3), X_1 is said to be tracking X_2 .

The joint process of X_1 and X_2 stacked in a vector is a special case of the following vectorial linear SDE

$$(4) \quad dX(t) = (m(t) + AX(t))dt + \Sigma dW(t),$$

where the state vector $X(t)$ and $m(t)$ are p -dimensional, A is called the pull matrix and is of dimension $p \times p$. $m(t)$ is regarded a deterministic process. Σ is a $p \times q$ dimensional diffusion matrix and W is a q dimensional Wiener process. The structure of direct causal relations between components of X is determined by the non-zero elements in the pull matrix A . If $A_{i,j} \neq 0$, then component X_j directly affects X_i , and there is an arrow from X_j to X_i in the causality graph. The notion of causality used here is equivalent to that of Granger [5] in Markov process models, see [20]. Since the contributions in Eq. (4) are normal and linear, the distribution will be normal and thus characterized by the expectation and covariance:

$$(5) \quad \begin{aligned} EX(t) &= V^{-1} e^{\Lambda(t-t_0)} V X(t_0) + V^{-1} \int_{t_0}^t e^{\Lambda(t-u)} V m(u) du \\ cov(X(v), X(t)) &= V^{-1} \left[\int_{t_0}^v e^{\Lambda(v-u)} V \Sigma(u) \Sigma(u)' V' e^{\Lambda(t-u)} du \right] (V^{-1})', \end{aligned}$$

for $v \leq t$, where V is the left eigenvector matrix and Λ is the diagonal matrix of eigenvalues of A , so that $A = V^{-1} \Lambda V$ conditioned on the existence of such a representation. A layered causal structure is achieved by specifying which of the elements in A that are non-zero and by restricting the diffusion matrix to a blocked structure. See the supplementary material and [20] for more on this.

Since the covariance structure of $X(t)$ governed by Eq. (4) is available in a closed form in Eq. (5), the covariance of the topmost component of the process of Eq. (3) can be found to be:

$$(6) \quad cov(X_1(0), X_1(t)) = \frac{\sigma_1^2}{2\alpha_1} e^{-\alpha_1 t} + \frac{\sigma_2^2 \alpha_1^2}{\alpha_1^2 - \alpha_2^2} \left(\frac{1}{2\alpha_2} e^{-\alpha_2 t} - \frac{1}{2\alpha_1} e^{-\alpha_1 t} \right),$$

provided $\alpha_1 \neq \alpha_2$ are both positive. The covariance structure here is different from the one-layered OU process described in Eq. (1). This indicates that data generated from the top layer of a two-layered process will be detectably different from that of the one-layered case.

The model we initially made for the coccolith size data had three layers, phenotypic mean, optimum and climate. We assumed the stochastic contributions in the second layer to be identical over the sites, while the stochastic contributions in the two other layers were uncorrelated over the sites. We have $k = 6$ processes indexed by j representing each of the sites in each of the 3 layers, denoted by $X_{1,j}$ in the top layer, $X_{2,j}$ in the middle layer, and $X_{3,j}$ at the bottom. The flow of causality is

$$(7) \quad X_{3,j} \rightarrow X_{2,j} \rightarrow X_{1,j}.$$

So, for a given site $j \in \{1, \dots, 6\}$, we have the following SDE system:

$$\begin{aligned}
 dX_{1,j}(t) &= -\alpha_1(X_{1,j}(t) - X_{2,j}(t))dt + \sigma_1 dW_{1,j}(t) \\
 dX_{2,j}(t) &= -\alpha_2(X_{2,j}(t) - X_{3,j}(t) - \beta T(t))dt + \sigma_2 dW_{2,j}(t) \\
 dX_{3,j}(t) &= -\alpha_3(X_{3,j}(t) - \mu_0)dt + \sigma_3 dW_{3,j}(t),
 \end{aligned}
 \tag{8}$$

where the W processes are independent Wiener processes. The regressor term $\beta T(t)$ defines how the second layer responds to an exogenous process $T(t)$, which in our application will be a global temperature indicator series, as measured by [22]. Thus $m(t)$ in Eq. (4) will have a contribution $\alpha_3 \mu_0$ on the third layer and a contribution $\alpha_2 \beta T(t)$ on the second. The second layer thus tracks a linear combination of the third layer and the external time series $T(t)$. While there are 18 processes in this particular model, it has only 8 parameters, namely $\theta = (\alpha_1, \alpha_2, \alpha_3, \sigma_1, \sigma_2, \sigma_3, \beta, \mu_0)$. The processes are regional (though inter-regional instantaneous correlation is introduced in layer two) but the parameters are global and thus so is the nature of the dynamics. The covariance function on the top (phenotypic) layer for this model is shown in the supplementary.

3.3. Pull identifiability. It is shown in the supplementary material that in any multi-layered process having only one pull parameter per layer, the pulls can be re-shuffled so that if there are l layers, then $\alpha_l < \alpha_{l-1} < \dots < \alpha_2 < \alpha_1$. It thus initially seemed natural to restrict the inference outcomes so that this is the case, which will be called the pull identification restriction. However, in the case of multiple sites, requiring for example that $\alpha_3 < \alpha_{2,s} < \alpha_1$ for all sites s , will exclude some covariance structures. When one analyze data using the model framework without imposing the pull identification restriction, one do however need to keep in mind that a reshuffling of the pull parameters is possible and may in some cases even be necessary. The alternative is either to impose the restriction only on one specific site or to impose them on all and so risk overlooking valid solutions. As we wished to avoid singling out a specific site and did not want to risk being overtly restrictive, we did not enforce the pull identification restriction. The identification issue should however be kept in mind when interpreting the results. In the supplementary material we present results also under the pull the identification restrictions.

When initially testing the simulation study, we used the pull identification restriction. The simulation analysis suggested an error in our treatment of the combination of prior distribution and the identification restriction, which resulted in Bayesian model selection support for over-complex models. For one-layered models the prior distribution for the pull parameter remains undisturbed. However, when imposing the pull identification restriction on our prior distribution for multi-layered models, the marginal distribution of each pull parameters became narrower.

Thus if the data suggests a pull parameter within this narrower interval, this is better predicted by the multi-layered model which is then unfairly supported in a BML comparison. We thus corrected our prior distribution so that the marginal distribution for at least one common pull parameter remains the same for models having different number of layers. See the supplementary for more on this issue and for the simulation results after correcting the prior distribution.

4. Likelihood-based inference methodologies.

4.1. *Single model inference.* A process governed by Eq. (4) is Gaussian with known mean and covariance structure. The measurement errors then add extra variance (equal to the sample variance divided by the sample size), which is also regarded as known. The likelihood of observations of the process affected by independent normal measurement errors is thus available in closed form. The likelihood function is however often multimodal since it usually is a non linear restriction of the covariance matrix of the observations, see e.g. [21]. Our experience is also that the coccolith data lead to multimodal likelihood functions for the models we consider. For maximum likelihood (ML) estimation a shotgun approach with a hill-climbing algorithm from some 50 starting points widely distributed in parameter space seemed to work reasonably well, as long as the number of parameters were kept low. Stable and efficient ML results were however only gotten when the hill-climbing algorithm was initialized using posterior Bayesian Markov chain Monte Carlo (MCMC) samples.

In a Bayesian setting using MCMC sampling, the multimodality issue seems to have been dealt with and stable, reproducible results were produced.

Due to the nonlinear parameters and the relatively low number of measurements, the Hessian at the likelihood maximum cannot be expected to provide reliable measures of estimation uncertainty. Bootstrapping might work, but is computationally expensive. Bayesian MCMC samples seem however to work and do yield approximate credibility intervals and scatter plots showing the dependency in the inference on different parameter combinations.

For the application to the fossil coccolith data, to be discussed next, we thus used Bayesian methods both for model choice and for parameter estimation, though we also used Akaike's Information Criterion (AIC) in the model selection and included ML estimates in the final presentation of models.

Further description of the Kalman filter, numerical ML optimization, Bayesian MCMC techniques for single model inference, numerical problems and computational efficiency are provided in the supplementary material.

4.2. *Model variants.* Our model frame for the coccolith data is an 18 dimensional version of Eq. (4) with 3 hierarchical layers with 6 components each. The diffusion matrix Σ is block diagonal accordingly, with Σ_i as the block for layer i . We also consider similar models having one or two hierarchical layers. For each layer i , we consider the following variants:

1) Regionality: It could be that one or several parameters (expectation, layer-specific pull or diffusion) are regional, i.e. different for different sites. Regionality is however only allowed in one of the three layers in a model variant. For more on this, see the supplementary material.

2) Determinism: A layer will be deterministic if it receives no stochastic contributions, i.e. if $\Sigma_i = 0$. Note that the pull and characteristic time may still be finite, so that the layer performs a deterministic filtering of the underlying stochastic layer.

3) Random walk: This is only possible for the lowest layer (interpreted as the climatic layer in this application). With a random walk layer, the process is non-stationary. Approximate random walk is achieved by setting the pull very low. Ideally, the pull should be set to zero, but for practical reasons, we chose to implement $\alpha_i = 0.001$, which means $\Delta t_i = 1Gy$.

4) Regional correlation: The components in a layer might be instantaneously correlated. We allow only one correlation coefficient, $\Sigma_{i,j_1,j_2} = \rho_i \sigma_{i,j_1} \sigma_{i,j_2}$ for any two sites $j_1 \neq j_2$. With $\rho_i = 1$, the instantaneous stochasticity (associated with $dW(t)$ in Eq. (4)) in the layer collapses to one dimension, as for the second layer in Eq. (8).

By crossing model variants and by varying restrictions on the parameters, we end up with 710 different models within our model framework. Without restrictions to only one layer with respect to regionality, the number of model variants to consider would much exceed 710, which already is a big number. Methods for exploring properties of the evolution of *Coccolithus* by confronting these models with the data in a Bayesian way is discussed next.

4.3. *Model and property comparison.* In order to calculate Bayesian model probabilities, one needs the marginal data likelihood,

$$(9) \quad f(D|M) \equiv \int f(D|\theta, M) \pi(\theta|M) d\theta,$$

where D is the data, M is the model, θ is the vector of parameters and π is the Bayesian prior distribution of the parameters. Since this is conditioned on the model, we will call it the Bayesian model likelihood (BML). From this and the prior model probability $Pr(M) = 1/710$ the posterior model probability $Pr(M|D)$ is calculated by Bayes' theorem. Since the prior model probabilities are assumed

equal, $Pr(M|D)$ is proportional to the marginal likelihood, $f(D|M)$. Comparing models by their posterior probabilities, one should note that the posterior probabilities cannot be interpreted in absolute terms. The problem here is that there might be many models that are very similar to each other, and the posterior probability of this type of model will be diluted by the model multiplicity.

The Bayes factor of two competing models is defined by how much the relative probability for the two models change from the prior to the posterior case, $B_{1,2} = (Pr(M_1|D)/Pr(M_2|D))/(Pr(M_1)/Pr(M_2)) = f(D|M_1)/f(D|M_2)$. Thus it is equal to the marginal likelihood ratio. The Bayes factor describes the amount of evidence in the data for model 1 versus model 2.

In Eq. (9) the parameter θ is the same for all model specifications. These models differ however with respect to restrictions imposed on the parameter, and the prior distribution is modified accordingly. We use a very wide prior distribution, see below. Due to the complexity of the prior and nonlinearities in the model, the marginal data likelihood is not analytically available despite the likelihood function being Gaussian. The numerical method we use is an importance sampler described in the supplementary material.

Since we look at many models with different properties, some such properties might be evaluated by expressing the different categories of a given property and identifying the models within each category. Looking at such a multitude of models means that many combinations of properties are tested at once. A model may then sometimes get support simply by statistical fluctuations. It will therefore be beneficial to look for a few general model properties, one at a time, and combine those, when looking for the best model. In order to assess what the data indicate concerning different properties, each category can be given the same property-specific prior probability, which we will call a weight for the application of studying model properties. Each model within that property can be given the same prior weight, so as to better illustrate in which direction the evidence pulls and how much evidence there is.

If each property was equally probable and each model within each category also was equally probable, the weights would be Bayesian probabilities. We are however not assuming that this is reasonable. Instead, we simply use the formalism of Bayesian probability theory, with probabilities re-labeled as weights, in order to explore how reasonable a given property category is. Note that the ratio of one weight to another will form the traditional Bayes factor for a property under the assumption of equal model probabilities within each property. If different categorizations are made, this will result in different prior weights for each model, so it seems more appropriate to call both the input and the results for weights rather than probabilities.

A categorization allows for an evaluation of the evidence for a given category of

a property by weighting it with the number of models it contains. This is done by evaluating the posterior weights

$$W(C | D) = \frac{\frac{1}{n_C} \sum_{M \in C} f(D | M)}{\sum_{C'} \frac{1}{n_{C'}} \sum_{M' \in C'} f(D | M')}$$

based on the prior weights $\frac{1}{n_C}$, where C is a category, n_C is the number of models within the category, and C' runs over the categories of the property. Thus this method of comparison compensates for the difference in number of models with different properties.

5. Simulations.

In order to test the model comparison, we made artificial datasets based on models with one, two and three layers. The models used for generating the datasets were the best one-, two- and three-layered model according to BML applied on the original dataset, while using the ML estimates for the choice of the parameter set. See the supplementary material for a description of these three models. Twenty datasets were then simulated for each of these three cases, with the same amount of data situated at the same measurement times and having the same measurement uncertainty as the real data, but where the state processes themselves were simulated using these models.

The same three models were then used for analyzing each of these three types of datasets and see how often different statistical model selection methods, AIC, BIC and BML, resulted in the right number of inferred layers. The results are shown in Table 1. BIC seems to perform best when a one-layered model has produced the data, but underperforms for data produced by a three-layered model. In total, it performs slightly worse than AIC. Fisher's exact test for contingency tables suggest that there is no significant difference in performance between AIC and BML for data produced by one- and three-layered model, but that BML underperforms in relation to AIC for data produced by a two-layered model.

It should however be noted that these observations only hold for our particular choice of one-, two- and three-layered models and is conditioned on the parameter values behind the simulated data. We used the maximum likelihood estimates. Since Bayesian methods allow for property inference, inclusion of prior knowledge and measures of parameter uncertainty, we will use both AIC and BML in the analysis. It should also be noted that according to the results, it is more likely than not that any of these three methods will identify the correct number of layers, no matter which of the three models produced the data.

6. Results and discussion for the *Coccolithus* data.

We ran the model selection exercise without imposing the pull identification restrictions (Section 3.3). The model variants selected by BML and AIC as the best violated these restrictions, and were found substantially better than models satisfying the restrictions. As a consequence, the analysis shown here are without pull identification restrictions. Results under these restrictions are shown in the supplementary material.

6.1. *Multi-layered model selection.* Going through all 710 model variants, we opted to do the model selection both by looking at BML and AIC and by doing Bayesian property analysis.

To our knowledge, little information is available on which to base the prior distribution, except broad ideas about the time scales involved in biological processes and general knowledge about the size and variation of size in single-celled organisms. We therefore construct our prior by applying independent normal distributions to each type of parameter (or rather a re-parametrized version so that the new parameters can be allowed to take values over the entire real line). See the supplementary information for more concerning the prior distribution, which was wide but informative.

The best model according to BML (and second best according to AIC), a three-layered model, is shown in Table 2. It had a Bayes factor $B = 215$ compared to the best one-layered model, while the AIC difference between the best multi-layered model according to AIC and the best one-layered model according to AIC was $\Delta AIC = -11.2$. The best model according to AIC was a variant of the best model according to BML only with an extra inter-regional correlation found in the first layer. Since the top layer turned out to have very fast-moving dynamics, any contributions due to inter-regional correlations would quickly vanish. Therefore, the structural difference between the best AIC and best BML model made very little practical difference. Two fairly different statistical model selection methods thus gave approximately the same model, having overall the same structure. There was regional pull without diffusion in the second layer (deterministic filtering of the lowest layer, with different filtering for different sites). Inter-regional correlations were inferred to be on the third layer.

According to Table 2, the second layer contains both sites with greater and smaller pulls than that of the third layer, which is a result not possible if the pull identification restriction is imposed. Note that while layers can be switched in a multiple site setting as well as for single set, there is no way to switch them so as to make pulls progressively lower for all sites.

Compared to our initial model, there are a couple of changes. The inter-regional

correlation is moved to the lowest layer and is no longer perfect. The pull parameters on the second layer are regional rather than global and there is no diffusion on that layer. Thus this model has five more parameters than the initial model.

6.2. *Influence of external data series (global climate).* After finding the best model according to BML, we looked at expansions that allowed for an external data series to influence the second layer through a regression term, β . The 95% credibility interval for β encompassed zero, indicating cell size to be rather unaffected by the global temperature indicator. The Bayes factor between the model described and the same model without the temperature indicator series regression, $B = f(D | M_{no\ temp})/f(D | M_{with\ temp})$, is about 662. The Bayes factor for the regression is sensitive to the choice of prior for β , so that a wider 95% credibility interval $(-10, 10)$, results in the Bayes factor larger than 6000. Thus the strength of the evidence against temperature dependence is sensitive to the prior, though for both the initial and the alternative choice of prior distribution for β , the Bayes factor indicates no evidence for temperature dependence. With a 95% credibility interval encompassing zero and having an upper limit of about 0.012 and with a variation in the temperature indicator series of about 4.7, the impact, if any, of the temperature indicator variations on the second layer would with 95% probability be less than $\exp(4.7 \times 0.012) \approx 6\%$. Thus the global temperature indicator influence on the coccolith size will be disregarded for the rest of the analysis.

6.3. *Inference of model properties.* Inference on the process states themselves can also be performed in the same framework, using the Kalman smoother method. An example of this is shown in Figure 2, where the three layers belonging to a single site in our dataset has been studied using sampled parameter sets (taken from the posterior MCMC samples) on the model considered best. Figure 2a indicates that the top layer is well-adapted to the observations, while the processes of the layers below have expectation values that are progressively smoother versions of the layer above. Figure 2b shows how the uncertainty of the topmost process can be associated with a fixed maximum span far outside the observations but becomes progressively smaller the closer in time one is to an observation. However, the credibility interval does not shrink to zero, because of observational noise.

Looking at property inference, as described in section 4.3, we get results which are shown in Table 3. The data give little posterior weight to a process having only one layer, while higher weights are given to two or three layers. Also, there is an OU process rather than a random walk in the lowest layer (thus stationarity) and there is evidence for inter-regional correlations. Three layers seem to be slightly preferred over two layers and regional pull parameters are found in the middle layer. There seems to be support for the existence of at least one deterministic layer. These results support models having the same overall structure as the top

model described in Table 2. All of the top five models under both AIC and BML selection, described in the supplementary, also seem to have the same properties.

Bayesian model comparison can be hampered by sensitivity to the prior probabilities. We therefore did a new analysis with another prior distribution deemed reasonable, namely one where the characteristic times were given a 95% credibility bands spanning from 1y to 300My rather than 1ky to 1Gy and where the distribution of the diffusion parameters were widened. The top two models according to BML remained the same and the parameter estimates remained essentially the same also. The property weights were subtly shifted, but not enough to change any previous conclusions. These results suggest to us that the Bayesian inference was robust. See the supplementary material for details concerning various sensitivity tests.

The structure of the inter-regional correlations can be investigated by categorizing according to in which layer the inter-regional correlations exist. The posterior weights (Table 4) indicate support for inter-regional correlations in the lowest layer and possibly also on the top layer. According to the best model in Table 2, the top layer is found to be very fast moving (having a characteristic time less than 10ky which is less than the smallest time intervals in the data). With the rather coarse time resolution in our data, the issue of inter-regional correlations in the upper layer must therefore remain unresolved.

Although our study is explorative, we are confident that *Coccolithus* has phenotypic evolution with motion in at least two layers. In addition, the posterior distributions for the pull parameter(s) in the second layer indicate that the fitness optima are not evolving as random walks, which was also the case for the optimum in [7], though there the optimum was modeled to track a random walk in the lower layer, while in our study the data suggests a stationary lower layer.

7. Summary.

A modeling framework for systems of related processes evolving in continuous time has been constructed. This framework has been applied to fossil data of size variability in marine unicellular algae spanning nearly 60 million years. A simulation study showed that for the amount of data in this application, the number of layers could be inferred.

There is basic consensus among the models considered best using different criteria and in the property analysis as to the following model properties:

(1) There is more than one layer of stochastic processes at play. (2) There is dependency between what happens in two different regions. (3) There are some regional differences in the nature of the dynamics, i.e. there are some parameters that are site-specific. (4) The regional dependency does not stem primarily from the top

layer, but from something further down the causal chain. (5) The hidden process seems to exhibit slow variation resulting in long autocorrelation in cell size.

This suggests that both global and local processes influence the mean phenotype, through dynamic site-specific fitness optima that respond to an underlying process that is partly global. However, as already mentioned, we find no support for a forcing by a global temperature indicator series on *Coccolithus* size, in contrast to reports on other biotic groups [17, 4]. The regional differences in estimated model parameters might reflect contrasts in local environmental conditions, differences in morphotype (possibly (sub)species) composition, or a combination of both. These interpretations and additional sensitivity tests (e.g. how age models affect the model outcome) will be further explored in another publication [11].

Thus it seems that the framework can be used for studying and reaching tentative conclusions about the driving forces behind a phenotypic time series. It should also be possible to use this framework in other settings both within and outside paleontology and evolutionary biology. Time series with irregular temporal resolution, due to missing data, breaks in the observational scheme, or for other reasons, are not uncommon. For such data our framework provides an alternative to methods based on auto regression with regular temporal resolution.

The class of models described by linear SDEs is wide, but computationally feasible. They allow causality to be modeled and studied, see [20], and they accommodate latent hierarchical structures. Our model could for instance be expanded by including more layers, allowing for time series that are phylogenetically related, allowing for external processes being included as exogenous forcings in different layers or by using the framework also for modeling these time-series, thus making them an integral part of the analysis. Furthermore, geographic information could be incorporated by letting the correlation between sites depend on the distance between the sites, for example in relation to ocean circulation.

Acknowledgments. We would like to thank Thomas Hansen for helpful comments. We also thank the reviewers and the editor for their extensive advice. Funding for this research was provided by CEES at the University of Oslo, the Research Council of Norway (RCN project 197823), and by the Royal Swedish Academy of Sciences through funding by the Knut and Alice Wallenberg Foundation (KAW 2009.0287). This project used DSDP/ODP samples provided by the Integrated Ocean Drilling Program (IODP). IODP is sponsored by the U.S. National Science Foundation (NSF) and participating countries under the management of Joint Oceanographic Institutions (JOI), Inc.

Source code. The source codes for our analysis programs can be found on the web page <http://folk.uio.no/trondr/layered>.

References.

- [1] ALLEN, L. J. S (2003). *An Introduction to Stochastic Processes with Applications to Biology*. Pearson Education, Inc., New Jersey.
- [2] CANDE, S. C. and KENT, D. V. (1995). Revised calibration of the geomagnetic polarity timescale for the Late Cretaceous and Cenozoic. *J. Geophys. Research* **100**, 6093–6095.
- [3] ESTES, S. and ARNOLD, S. J. (2007). Resolving the Paradox of Stasis: Models with Stabilizing Selection Explain Evolutionary Divergence on All Timescales. *Am. Naturalist* **169(2)**, 227–244.
- [4] FINKEL, Z. V., KATZ, M. E., WRIGHT, J. D., SCHOFIELD, O. M. E. and FALKOWSK, P. G. (2007). Climatically Driven Macroevolutionary Patterns in the Size of Marine Diatoms over the Cenozoic. *Proc. Nat. Acad. Sci. U.S.A.* **102(25)** 8927–8932.
- [5] GRANGER, C. W. J. (1969). Investigating Causal Relations by Econometric Models and cross-spectral methods. *Econometrica* **37(3)**, 424–438.
- [6] HANSEN, T. F. (1997). Stabilizing Selection and the Comparative Analysis of Adaptation. *Evolution* **51(5)**, 1341–1351.
- [7] HANSEN, T. F., PIENAAR, J. and ORZACK, S. H. (2008). A Comparative Method for studying Adaptation to a Randomly Evolving Environment. *Evolution* **62(8)**, 1965–1977.
- [8] HAQ, B. U. and LOHMANN, G. P. (1976). Early Cenozoic calcareous nanoplankton biogeography of the Atlantic Ocean. *Marine Micropal.* **1**, 119–194.
- [9] HENDERIKS, J. and TÖRNER, A. (2006). Reproducibility of coccolith morphometry: Evaluation of spraying and smear slide preparation techniques. *Marine Micropal.* **58**, 207–218.
- [10] HENDERIKS, J. (2008). Coccolithophore size rules - Reconstructing ancient cell geometry and cellular calcite quota from fossil coccoliths. *Marine Micropal.* **67**, 143–154.
- [11] HENDERIKS, J., REITAN, T., SCHWEDER, T. and HANSEN, T. (2011) Probing phenotypic adaptation in marine algae using stochastic equations. In preparation for *Paleobiology*.
- [12] HUNT, G. (2006). Fitting and Comparing Models of Phyletic Evolution: Random Walks and Beyond. *Paleobio.* **32(4)**, 578–601.
- [13] HUNT, G., BELL, M. A. and TRAVIS, M. P. (2008). Evolution Toward a New Adaptive Optimum: Phenotypic Evolution in a Fossil Stickleback Lineage. *Evolution* **62(3)**, 700–710.
- [14] HUNT, G., WICAKSONO, S. A., BROWNS, J. E and MACLEOD, K. G. (2010). Climate-driven Body-size Trends in the Ostracod Fauna of the Deep Indian Ocean. *Palaeont.* **53** 1255–1268.
- [15] LANDE, R. (1976). Natural Selection and Random Genetic Drift in Phenotypic Evolution. *Evolution* **30**, 314–334.
- [16] RAUP, D. M. (1977). Probabilistic models in evolutionary paleobiology. *Amer. Sci.* **65(1)** 50–57.
- [17] SCHMIDT, D. N., THIERSTEIN, H. R., BOLLMANN, J. and SCHIEBEL, R. (2004). Abiotic forcing of plankton evolution in the Cenozoic. *Science* **303** 207–210.
- [18] SCHUSS, Z. (1980). *Theory and Applications of Stochastic Differential Equations*. John Wiley and Sons, Inc., New York.
- [19] SCHWEDER, T. and SPJØTVOLL, E. (1982) Plots of P-Values to Evaluate Many Tests Simultaneously. *Biometrika*, **69(3)**, 493–502.
- [20] SCHWEDER, T. (2012). Causal sufficiency and Markov completeness. Submitted to *Scandinavian Journal of Statistics*.
- [21] SUNDBERG, R. (2010), Flat and Multimodal Likelihoods and Model Lack of Fit in Curved Exponential Families. *Scandinavian Journal of Statistics*, —.
- [22] ZACHOS, J., PAGANI, M., SLOAN, L., THOMAS, E. and BILLUPS, K. (2001). Trends, Rhythms, and Aberrations in Global Climate 65 Ma to Present. *Science* **292(5517)**, 686–693.

TABLE 1

Number of correct identifications of a model (as given by the number of layers) according to different selection schemes, when a total of 20 data sets are simulated, see text.

| Actual model | AIC | BIC | BML |
|--------------|-----|-----|-----|
| One layer | 16 | 20 | 15 |
| Two layers | 20 | 18 | 12 |
| Three layers | 19 | 15 | 19 |

TABLE 2

Maximum likelihood (ML) estimates, Bayesian posterior median (B.median), 95% prior and posterior credibility intervals for the model considered best according to BML. The units of the characteristic times $\Delta t_{layer,site}$ are specified in the table, where the following units are used: y =year, ky = 10^3 years, My = 10^6 years, Gy = 10^9 years. The fixed median of the state e^{μ_0} , specified in the original measurement scale, is given in units of μm = micrometers. The parameter β is the effect of temperature indicator on the second layer process in units of $\log(size)/C^\circ$. The diffusion parameters σ_{layer} are in units of $\log(size)/My^{1/2}$.

| | Parameters | ML | B.median | prior 95% | posterior 95% |
|--------|--------------------|--------|----------|--------------|----------------|
| | e^{μ_0} | 7.42 | 7.43 | 0.001 - 1000 | 7.15 - 7.69 |
| | β | 0.0006 | 0.0005 | -1 - 1 | -0.002 - 0.012 |
| Upper | Δt_1 | 11ky | 16ky | 1ky - 1Gy | 0.3ky - 80ky |
| layer | σ_1 | 0.48 | 0.43 | 0.02 - 3.5 | 0.22 - 2.8 |
| Middle | $\Delta t_{2,525}$ | 130ky | 120ky | 1ky - 1Gy | 1.0ky - 0.52My |
| layer | $\Delta t_{2,612}$ | 215My | 40My | 1ky - 1Gy | 0.58My - 3.2Gy |
| | $\Delta t_{2,516}$ | 0.4ky | 11ky | 1ky - 1Gy | 90y - 150ky |
| | $\Delta t_{2,752}$ | 1.0My | 1.0My | 1ky - 1Gy | 250ky - 2.9My |
| | $\Delta t_{2,806}$ | 3.7Gy | 88My | 1ky - 1Gy | 2.0My - 10Gy |
| | $\Delta t_{2,982}$ | 0.4ky | 12ky | 1ky - 1Gy | 100y - 140ky |
| Lower | Δt_3 | 1.2My | 1.4My | 1ky - 1Gy | 0.64My - 3.7My |
| layer | σ_3 | 0.17 | 0.16 | 0.02 - 3.5 | 0.11 - 0.24 |
| | ρ_3 | 0.66 | 0.64 | -0.18 - 0.98 | 0.29 - 0.85 |

TABLE 3

Posterior weights for different properties, with the number of models in parenthesis. Note that while regional parameters can be found either in the first, second, third layer in a three-layered model, such a feature is only possible in the first and second layer in a two-layered model and only in the first layer in a one-layered model. Similarly, a three-layered model may be allowed to have no diffusion either in the first or second layer or both, but a two-layered model may only be without diffusion in the first layer and a layer without diffusion is not an option for a one-layered model.

| Property | Options | | | |
|--|-------------------------------|----------------------------------|---------------------------|-----------------------------------|
| | 1 | 2 | 3 | |
| Number of Layers: | 7.5 % (18) | 32.4% (114) | 60.1% (578) | |
| Regionality in: | none or μ_0 0.2% (177) | pull 88.3% (259) | diffusion 11.5% (274) | |
| Regional parameters in: | no layer 0.2% (177) | layer 1 13.1% (205) | layer 2 69.1% (196) | layer 3 17.5% (132) |
| Inter-regional instantaneous correlations: | none 4.9% (50) | intermediate (6D) 85.2% (212) | perfect (1D) 9.9%(448) | |
| No diffusion in | no layer 7.7% (486) | layer 1 1.2% (132) | layer 2 91.1% (72) | both layer 1 and 2 0.007% (20) |
| Random walk in the lowest layer: | no 99.1% (414) | yes 0.9% (296) | | |

TABLE 4

Posterior weights of correlation structure properties. Only three-layered models with non-degenerate pull at the bottom layer (333 in total) were examined.

| Inter-regional correlation in | | | | | | | | |
|-------------------------------|-----|-----|-----|-----|------|------|------|-----|
| | N | Y | N | Y | N | Y | N | Y |
| Top layer | | | | | | | | |
| Intermediate layer | N | N | Y | Y | N | N | Y | Y |
| Bottom layer | N | N | N | N | Y | Y | Y | Y |
| Number of models | 15 | 16 | 30 | 32 | 56 | 60 | 60 | 64 |
| Posterior probability (%) | 1.6 | 1.1 | 4.8 | 2.5 | 49.9 | 25.2 | 10.1 | 4.9 |

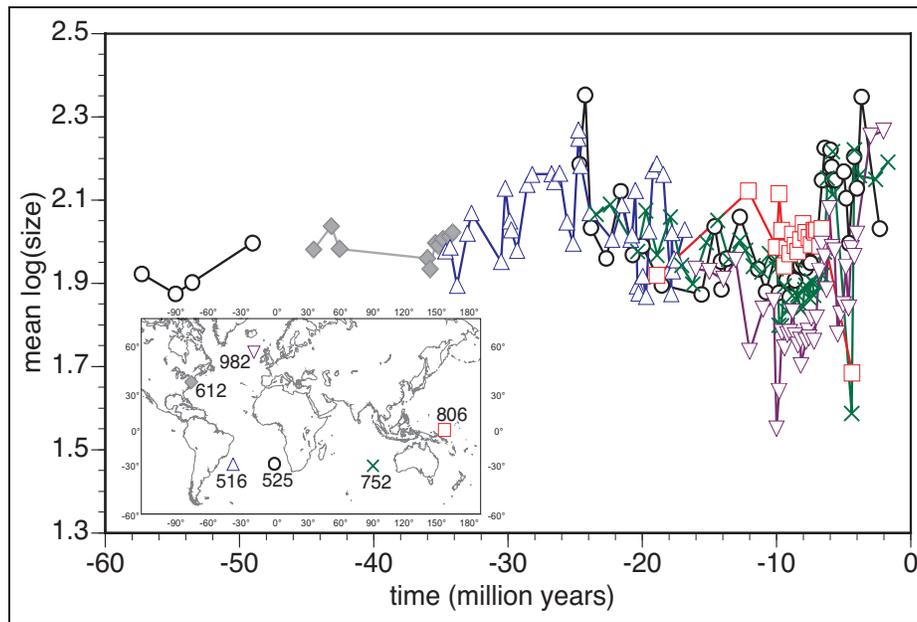
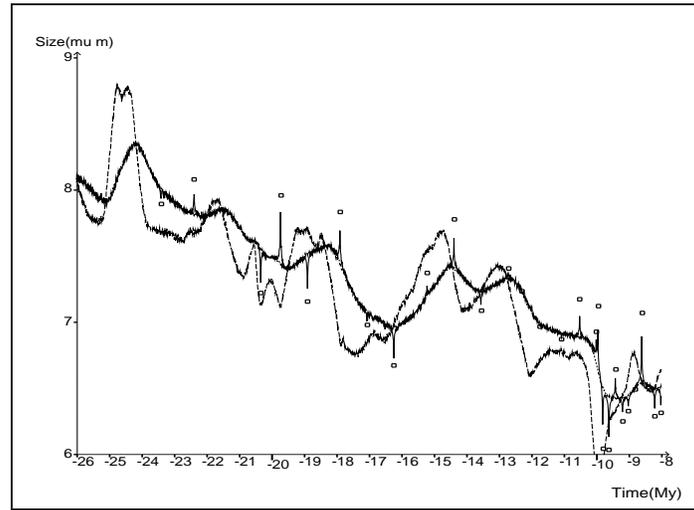
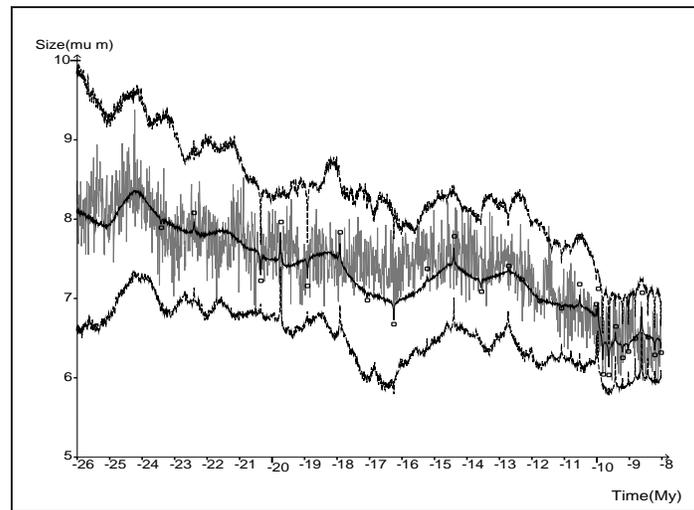


FIG 1. DSDP and ODP site locations (inset map) and site-specific mean logarithmic size values spanning the last 57 My. One symbol is shown for each of the 205 observations of mean log coccolith size in the time series plot.



(a)



(b)

FIG 2. Process inference and measurements for Site 752 running in the time interval -26My to -8My , using the model in Table 2, which is considered best according to BML. The inference takes Bayesian parameter uncertainty into account as well as process uncertainty given parameter values.

(a) Three-layered posterior means. Solid line: upper layer, short dashed line: middle layer, long dashed line: lower layer, circles: measurements. The upper and middle layer may be hard to separate, except for small time intervals close to measurements. Note that there are patterns in the lowest layer before the first measurement belonging to this site. This is because the lowest layer has correlations to other sites, for which there are measurements before this time.

(b) Upper layer posterior mean, uncertainty and one realization. Solid line: mean, dashed lines: limits of a 95% credibility interval, grey wiggly line: realization, circles: measurements. Except for the right side of the graph, where there are many high quality measurements, the credibility interval is wide. Note that while the upper layer realization may seem noisy due to the short-memory stochastic contributions to that layer, it is also influenced by lower layers with longer memory. This creates correlations over large time spans, so that the realization can be over or under the top layer mean for large time spans.