

Mechanistic site-based emulation of a global ocean biogeochemical model (MEDUSA 1.0) for parametric analysis and calibration: an application of the Marine Model Optimization Testbed (MarMOT 1.1)

J. C. P. Hemmings^{1,2}, P. G. Challenor^{3,1}, and A. Yool¹

¹National Oceanography Centre, Southampton, SO14 3ZH, UK

²Wessex Environmental Associates, Salisbury, UK

³College of Engineering, Mathematics and Physical Sciences, University of Exeter, Exeter, EX4 4QF, UK

Correspondence to: J. C. P. Hemmings (jhemmings@wessexenv.mail1.co.uk)

Abstract

Biogeochemical ocean circulation models used to investigate the role of plankton ecosystems in global change rely on adjustable parameters to capture the dominant biogeochemical dynamics of a complex biological system. In principle, optimal parameter values can be estimated by fitting models to observational data, including satellite ocean colour products such as chlorophyll that achieve good spatial and temporal coverage of the surface ocean. However, comprehensive parametric analyses require large ensemble experiments that are computationally infeasible with global 3-D simulations. Site-based simulations provide an efficient alternative but can only be used to make reliable inferences about global model performance if robust quantitative descriptions of their relationships with the corresponding 3-D simulations can be established.

The feasibility of establishing such a relationship is investigated for an intermediate complexity biogeochemistry model (MEDUSA) coupled with a widely-used global ocean model (NEMO). A site-based mechanistic emulator is constructed for surface chlorophyll output from this target model as a function of model parameters. The emulator comprises an array of 1-D simulators and a statistical quantification of the uncertainty in their predictions. The unknown parameter-dependent biogeochemical environment, in terms of initial tracer concentrations and lateral flux information required by the simulators, is a significant source of uncertainty. It is approximated by a mean environment derived from a small ensemble of 3-D simulations representing variability of the target model behaviour over the parameter space of interest. The performance of two alternative uncertainty quantification schemes is examined: a direct method based on comparisons between simulator output and a sample of known target model “truths” and an indirect method that is only partially reliant on knowledge of target model output.

In general, chlorophyll records at a representative array of oceanic sites are well reproduced. The use of lateral flux information reduces the 1-D simulator error considerably, consistent with a major influence of advection at some sites. Emulator robustness is assessed by comparing actual error distributions with those predicted. With the direct uncertainty quantification scheme, the emulator is reasonably robust over all sites. The indirect uncertainty quantification scheme is less reliable at some sites but scope for improving its performance is identified. The results

demonstrate the strong potential of the emulation approach to improve the effectiveness of site-based methods. This represents important progress towards establishing a robust site-based capability that will allow comprehensive parametric analyses to be achieved for improving global models and quantifying uncertainty in their predictions.

5 1 Introduction

A need for better understanding of the role marine biota will play in influencing the nature and rate of global change in response to human activities has led to the inclusion of process-based models of ocean biogeochemistry in ocean circulation models (Sarmiento et al., 1993) and more recently in models of the whole Earth system (Séférian et al., 2013). They are designed to capture the dominant responses of complex ecosystems to variability in the physical environment. The biogeochemistry models vary in complexity from simple models in which the biota are represented by single phytoplankton and zooplankton types (e.g. Six and Maier-Reimer, 1996; Palmer and Totterdell, 2001) to more complex functional type models in which a much larger range of different planktonic groups are represented (e.g Moore et al., 2004; Gregg et al., 2003; Le Quéré, 2005; Aumont and Bopp, 2006).

The process-based models are often referred to as mechanistic, as distinct from statistical or data-based models. Yet they are also semi-empirical, incorporating adjustable parameters. Such parameters are important in process-based models of complex systems where incomplete knowledge and practical limits on the degree of complexity that can be resolved make it impossible to design a model that represents all relevant mechanisms. Predictions given by each model are thus affected by structural uncertainty, associated with the model's design, and parametric uncertainty, associated with its chosen parameter values. The equivalent parameters in nature are typically highly variable in space and time and among different organisms present in any assemblage, making the optimal values particularly elusive. Effective use of ocean observations to constrain model parameters and reduce parametric uncertainty is necessary to improve the predictive skill of particular models and to gain a better understanding of inadequacies in model design.

Any rigorous exploration of a biogeochemical model's parameter space is computationally intensive, requiring many thousands of simulations. This has generally dictated the use of fast site-based experiments for parametric analyses, following the pioneering work of Fasham and Evans (1995) and Matear (1995). Parameters are optimized to fit observations at individual sites (e.g. Losa et al., 2004; Fasham et al., 2006; Friedrichs et al., 2006, 2007; Dowd, 2011; Kidston et al., 2011; Fiechter et al., 2013; Prieß et al., 2013a; Ward et al., 2013) or at multiple sites simultaneously (Hurtt and Armstrong, 1999; Schartau and Oschlies, 2003; Hemmings et al., 2004; Friedrichs et al., 2007; Kane et al., 2011; Xiao and Friedrichs, 2014). In these experiments, the biogeochemistry model is integrated in a 1-D or 0-D framework representing a single water column at each site, and a local approximation of the physical environment is used as forcing data to drive the simulation.

In the site-based study of Dowd (2011), a sequential data assimilation method with a stochastic configuration of a biogeochemistry model was used to estimate the models' static parameters in combination with its time varying state (i.e. its prognostic variables). Sequential methods use a series of analysis cycles in which analysis steps combine observations with model forecasts, taking into account the uncertainties in each. The forecast for each step is initialized from the previous analysis. Dowd (2011) estimated new joint probability distributions for state and parameters at each observation time on the basis of the new observations and a previous analysis. However, in most cases variational inverse methods are used, the aim being to constrain the parameters of the deterministic free-running model. Parameter values are varied with the objective of minimizing or maximizing some function of the model-data differences. The solution is then the best fit to the complete observational data set that satisfies the model equations exactly (ignoring error introduced by time discretization in the numerical solver). An exception is made in the inverse approach of Losa et al. (2004) where the model equations are used as a weak constraint and both parameters and state are estimated. This allows for sources of simulation uncertainty that are not associated with the adjustable parameters, such as structural error or error in the forcing data.

Sequential data assimilation approaches are particularly useful in short-term forecasting, where the forecast is highly dependent on the initial state and state estimation is the primary

goal. However, for long-term future projections that must rely on free-running models, the estimation of model parameters is paramount. Methods that preserve the integrity of the model dynamics are inherently better suited to this problem but simulation error impacting on the state variables cannot be ignored and a more rigorous treatment of simulation uncertainty is needed before the potential of these methods can be fully realized (Hemmings and Challenor, 2012).

In this study, we focus specifically on simulation uncertainty introduced by the use of 1-D simulations to approximate 3-D model behaviour. The uncertainty is primarily associated with differences in the representation of the physical environment and differences in the horizontal fluxes and initial values of biogeochemical properties. Despite this uncertainty, site-based calibrations have been shown to improve the predictive skill of 3-D models (Oschlies and Schartau, 2005; Kane et al., 2011; McDonald et al., 2012). However, the relationship between 1-D and 3-D simulations is not well understood in quantitative terms. Parameter vectors that are optimal in one context are unlikely to be optimal in the other, inevitably compromising the utility of established parameter estimation methods.

The lack of information about biogeochemical fluxes associated with horizontal advection and diffusion is an obvious source of uncertainty. Some consideration has been given to this problem. Losa et al. (2004) introduced their weak constraint approach primarily to allow for the neglect of horizontal transport. Fasham et al. (2006) parametrized diffusive fluxes based on the analysis of a passive tracer release associated with an iron fertilization experiment, while Friedrichs et al. (2007) included an advective flux divergence term for nutrients based on 3-D model output. Fasham et al. (1999) took a different approach, optimizing parameters in a Lagrangian framework to fit data from a survey of the North Atlantic spring bloom. The survey followed the track of a drogued buoy to minimize the impact of horizontal advection on the biogeochemical system under study. More typically though, horizontal fluxes are ignored in site-based calibration studies.

In a relatively small number of studies, parameters have been optimized for the biogeochemistry model within its host 3-D circulation model. This is practical for limited time and space domains: Garcia-Gorriz et al. (2003) and Huret et al. (2007) estimated parameters for regional models by assimilating satellite-derived chlorophyll data over periods of order 1 month. Doron

et al. (2013) assimilated these data at a single point in time into an eddy-permitting model of the North Atlantic using an adapted Kalman filter analysis with a perturbed parameter ensemble simulation. The ensemble simulation was similarly of 1 month duration. Fan and Lv (2009) estimated spatially varying parameters for the global domain but with an assimilation window limited to 5 days. In contrast, Tjiputra et al. (2007) performed much longer global experiments, assimilating seasonal maps of surface chlorophyll and nitrate into a global model of the annual cycle, but relied on a coarse resolution model (3.5° horizontal resolution) and, in common with a number of other studies, only optimized locally in parameter space.

The type of compromises imposed on parametric analyses of 3-D biogeochemical models by limited computer resources are generic to many different fields in which computer models are used. This problem has motivated the development of statistical emulation techniques that allow more comprehensive investigations of parameter space to be achieved. A good introduction is given by O'Hagan (2006). An emulator provides a prediction of a chosen model output, or a metric used in its assessment, for any setting of the parameter values, together with a measure of uncertainty in that prediction. A relatively small ensemble of model runs is required to provide training data for emulator construction, although this is still a significant overhead for 3-D models.

Statistical emulation techniques have been applied to the estimation of marine biogeochemical model parameters in regional studies. Leeds et al. (2013) used emulators for computational efficiency in a Bayesian hierarchical framework that linked spatially distributed 1-D simulations. In other work, emulators were constructed for relatively expensive 3-D simulations to allow the required coverage of parameter space to be achieved: Hooten et al. (2011) used 50 ensemble members to represent a 7-dimensional parameter space, while Mattern et al. (2012) used a similar ensemble size in a 2 parameter study.

Although, to the authors' knowledge, the application of statistical emulators to ocean biogeochemistry has so far been limited to regional studies, they are starting to be used at the global scale for parametric analyses of other Earth system model components, including the coupled ocean-atmosphere system (Williamson et al., 2013) and atmospheric aerosol concentrations (Lee et al., 2012). These studies involved the use of perturbed parameter ensemble

simulations with global 3-D models. Williamson et al. (2013) investigated a 30-dimensional parameter space, benefitting from a very large ensemble generated using climateprediction.net, a distributed computing project in which personal computers are volunteered by members of the public. Lee et al. (2012) used a much smaller ensemble (80 members) to investigate parametric uncertainty over an 8-dimensional parameter space. The ensemble size was computationally practical owing to the coarse resolution of the model and the limited duration of the runs (4 months).

The application of statistical emulators to global ocean biogeochemical models would make investigation of the models' predictive potential more tractable. However, achieving sufficiently large training ensembles for periods that fully capture the seasonal variability at an appropriate spatial resolution will be challenging. Mesoscale and sub-mesoscale dynamics are known to have a strong impact on biogeochemical processes in the upper ocean (Lévy, 2008), yet global simulations that resolve the ocean mesoscale require considerable computing resources, severely limiting ensemble size.

Given the potential for improving the representation of biogeochemical cycles by increasing model resolution, avoidance of unnecessary trade-offs between resolution and ensemble size is desirable. Improving 1-D modelling capabilities is a potential solution. The goal would be to produce a set of site-based simulators that could serve as an efficient and reliable surrogate model for emulating arbitrary 3-D model outputs with quantified uncertainty. The number of sites could be adapted according to the required ensemble size and the resources available. Like a statistical emulator, the system would provide a prediction of model output and a measure of uncertainty in that prediction. We refer to the proposed system as a mechanistic emulator to distinguish it from statistical site-based emulators (Leeds et al., 2013) that treat the target model as a black box. For some parametric analyses, a mechanistic emulator of this type would be sufficient. Where more comprehensive analyses are required it would be used to bridge the gap between the 3-D target model and one or more statistical emulators of model outputs or metrics.

Here we introduce an experimental mechanistic site-based emulator and use it to explore the feasibility of establishing a robust relationship between 1-D and 3-D simulations. The emulator

predicts annual cycles of surface chlorophyll output produced by a target model of the global ocean. The aim is to provide a way of exploiting satellite chlorophyll or related ocean colour products for making reliable inferences about the target model performance for arbitrary trial parameter vectors, without having to run the corresponding 3-D simulations.

5 Section 2 describes the components of the mechanistic emulator and the method for its construction and Sect. 3 gives the experimental method used to evaluate its performance. The results are presented in Sect. 4. In Sect. 5 the findings are discussed with regard to the potential of the emulation scheme as an enabling tool for improved parametric analyses of global models, using satellite ocean colour data in combination with in situ observations. A summary of the work is
 10 given in Sect. 6.

2 The mechanistic emulator

The site-based emulator combines a surrogate model with a probabilistic prediction of its error with respect to the 3-D target model. The surrogate model takes the form of an array of 1-D
 15 simulators. Variation of the predicted error distribution of surface chlorophyll output from the surrogate model over its time and space domain is fully described. The intention is to establish a form of traceability between the surrogate model and the target model that allows robust inferences about target model skill to be made from analyses of surrogate model output.

Inferences about model performance are often made on the basis of a cost function, summarizing the misfit of a simulation to observational data. The cost function typically takes the form
 20

$$J(\mathbf{y}_P) = (\mathbf{y}_P - \mathbf{y}_O)^T \mathbf{R}^{-1} (\mathbf{y}_P - \mathbf{y}_O) \quad (1)$$

where \mathbf{y}_O is a vector of n observations, \mathbf{y}_P is the corresponding vector of predicted values and \mathbf{R}^{-1} is the inverse of the $n \times n$ error covariance matrix (Stow et al., 2009). The superscript T is the transpose operator. The error covariance matrix describes the predicted error structure of the
 25 model output. It weights the contributions of individual model-data misfits according to their significance, taking into account prior expectations of uncertainty.

It is commonly assumed that the individual misfits are independent. The off-diagonal elements of \mathbf{R} are then zero and the cost function can be written

$$J(\mathbf{y}_P) = \frac{1}{n} \sum_{i=1}^n \frac{(P_i - O_i)^2}{\sigma_{ii}^2} \quad (2)$$

where P_i and O_i are the elements of \mathbf{y}_P and \mathbf{y}_O respectively and σ_{ii}^2 represents the diagonal elements of \mathbf{R} .

If both observation and simulation error are relevant in an analysis, the error variance σ_{ii}^2 is the predicted variance of the combined error from both sources. When using a surrogate model, the simulation error includes the surrogate model error with respect to the target model. It may also include error from other sources such as target model input data or structural error, depending on the objective of the analysis. Hemmings and Challenor (2012) discuss cost function design for different analyses in more detail.

Predicted surrogate model error statistics can be used in a cost function to make the function more informative about the likely misfit between the target model and the observations. They do this by increasing the weight given to model-data misfit where the surrogate model error is expected to be small and decreasing the weight elsewhere. The cost function can then be used to evaluate the goodness-of-fit of the target model simulation to the observations, given the surrogate model output.

In the experimental emulator presented here, the statistical prediction of the error with respect to the target model is restricted to its mean and variance at individual data points. If the emulator were used in a cost function-based analysis, the predicted error variance would contribute directly to σ_{ii}^2 and the predicted mean error would be used to give bias-corrected values for P_i . Estimation of the mean and variance is a first step towards a more complete uncertainty quantification that would include the error covariance structure required to fully specify \mathbf{R} .

The target model in the present study is NEMO-MEDUSA, combining the MEDUSA 1.0 biogeochemistry model (Model for Ecosystem Dynamics, carbon Utilisation, Sequestration and Acidification) described by Yool et al. (2011) with the NEMO ocean model (Nucleus for European Modelling of the Ocean; Madec, 2008).

2.1 The biogeochemical simulator

The 1-D simulator incorporates a representation of the biogeochemistry that is identical to that in the target model. MEDUSA is an intermediate complexity model, representing the plankton ecosystem by 11 compartments in the form of biogeochemical tracers. These include 6 nitrogen pools for two phytoplankton groups (diatoms and non-diatoms), two zooplankton groups (micro- and meso-zooplankton), slow-sinking detritus and dissolved inorganic nitrogen. The remaining compartments represent two additional dissolved nutrients required by the phytoplankton (silicon and iron), the chlorophyll concentrations associated with the two phytoplankton types and the silicon concentration associated with the diatoms. The effect of fast-sinking detritus is represented by instantaneous vertical redistribution of material in the water column.

1-D integrations of MEDUSA are performed in a 3-D context where physical and biogeochemical information from the target model provide environmental input data for the site-based simulations. The physical environment required by the 1-D simulator is independent of the biogeochemical model parameters. However, the biogeochemical environment is parameter-dependent making its representation in a site-based parametric analysis less straightforward. The 1-D simulator for MEDUSA is configured using the Marine Model Optimization Testbed facility described by Hemmings and Challenor (2012). The testbed software, MarMOT 1.1, is open source and freely available as detailed in Appendix A.

The MEDUSA state variables are the biogeochemical tracer concentrations at each model grid point. The evolution equation for the concentration c_{ik} of the i th biogeochemical tracer at depth level k in the 1-D simulator is

$$\frac{dc_{ik}}{dt} = -(w_p + w_i) \frac{\partial c_i}{\partial z} + \frac{\partial}{\partial z} \left(K_\rho \frac{\partial c_i}{\partial z} \right) + SMS_{ik}(\mathbf{C}, \mathbf{F}) + p_{ik}(\mathbf{C}_k, p_{jk}^*). \quad (3)$$

The first two terms represent the tendencies (i.e. rates of change) due to vertical flux divergence. w_p is the vertical velocity of the water, w_i is the active vertical velocity of the biological material relative to the water and K_ρ is the turbulent diffusion coefficient. SMS_{ik} is the source-minus-sink term from the MEDUSA plankton model. It is a function of the state vector \mathbf{C} and

a forcing vector \mathbf{F} comprising temperature, downwelling solar radiation at the sea surface and input of soluble iron from atmospheric dust deposition. SMS_{ik} is depth-dependent because the light available for phytoplankton photosynthesis and the nutrient sources from the remineralization of fast-sinking detritus depend on tracer concentrations at $k - 1$ shallower levels. w_i is assigned a constant sinking rate for the detritus tracer, corresponding to the MEDUSA sinking rate parameter for slow-sinking detritus. It is zero for all other tracers. Values for w_p , K_ρ and \mathbf{F} are provided by the physical environment from the target model.

The final term in Eq. (3) is a perturbation term used to represent the effect of horizontal flux divergence. The divergence tendency for the i th tracer p_{ik} depends on the local state \mathbf{C}_k (a vector containing the subset of tracer concentrations at depth level k) and an applied perturbation p_{jk}^* . Tracer-specific perturbations are applied to tracers representing dissolved nutrients and the nitrogen content of the plankton. These are referred to as primary tracers. The phytoplankton chlorophyll and silicon tracers (secondary tracers) are affected indirectly, following the perturbations to the corresponding nitrogen tracers in such a way as to preserve the phytoplankton chlorophyll : nitrogen and silicon : nitrogen ratios. For a primary tracer, $j = i$. For a secondary tracer, j indexes the relevant primary tracer.

The input data set required to define the biogeochemical environment comprises the initial state and the applied perturbations controlling the tracers' horizontal flux divergence tendencies. This is the biogeochemical environment vector

$$\mathbf{B} = \{\mathbf{C}(t_0), \mathbf{P}^*\}. \quad (4)$$

$\mathbf{C}(t_0)$ is the initial state vector containing the concentrations of the 11 tracers at each depth level on the model grid at time t_0 and the vector \mathbf{P}^* contains applied perturbations at each depth level for the 8 primary tracers at 5 day period mid-points for $t > t_0$. Perturbations represent the effect of lateral advection inferred from an analysis of local currents and upstream property gradients in the 3-D model output. The effect of horizontal diffusion is ignored.

The advective tendencies of individual tracers are dependent on their upstream gradients and often tend to co-vary with their local concentrations. It is important to give some attention to preserving such relationships that are prevalent in the 3-D simulation as far as possible. A particular example of a prevailing relationship occurs when tracer concentrations are low. If we have

a negative advective tendency it should increase towards zero as the concentration approaches zero, otherwise the concentration will become negative. In the 3-D simulation, this happens naturally because the upstream gradient driving it tends towards zero (assuming the upstream concentration cannot be negative). In the 1-D simulation, adaptation of tendencies to the local concentration is necessary to counter any inconsistencies between the two. This concentration dependency is introduced by using applied perturbations that represent rates of change of transformed tracers. The choice of transformation determines the form of the dependency and is an important consideration in simulator design.

Analysis of 3-D simulations indicate that the concentration dependency of horizontal gradients varies temporally and spatially and between different tracers. Use of the square root transformation protects against the evolution of negative concentrations and was found by Hemmings and Challenor (2012) to be a reasonable compromise between using untransformed and log-transformed concentrations. A square root transformation was therefore chosen for all primary tracers at all sites so that a perturbation p^* specifies the rate of change of \sqrt{c} , where c is the tracer concentration. The implied concentration tendency is then

$$p = 2\sqrt{c}p^*. \quad (5)$$

For secondary tracers the tendency is

$$p_i = \frac{c_i}{c_j} p_j \quad (6)$$

where i is the secondary tracer index and j indexes the associated primary tracer. The applied perturbation diagnosed from 3-D model output is

$$p^* = -\mathbf{u}_h \cdot \nabla_h \sqrt{c} \quad (7)$$

where the subscript h denotes vectors in the horizontal plane and \mathbf{u}_h is the current velocity.

Differences between the simulator output and that of the target model arise due to the combined effects of a number of sources of simulation error. Specifically these are approximation error in the physical environment variables due to temporal averaging of the 3-D target model

data on which they are based, error in the advective flux divergence tendencies, error introduced by ignoring horizontal diffusion and differences in solver numerics. Any differences between the initial state $C(t_o)$ and the target model state at time t_o will contribute an additional source of error.

5 2.2 The uninformed simulator and biogeochemical environment model

In a calibration exercise or other parametric analysis, the 1-D simulator is used to learn about the likely behavior of 3-D target model simulations that have not been performed. For an arbitrary trial parameter vector \mathbf{x}_o , the parameter-specific biogeochemical environment $\mathbf{B}(\mathbf{x}_o)$ is typically unknown. Instead we use an environment vector derived from a statistical model. The corresponding 1-D simulator is referred to as the *uninformed simulator* indicating that it is not informed by parameter-specific environment data. Our surrogate model consists of an array of uninformed simulators at different sites, spanning a range of oceanic conditions.

The statistical model used to define the biogeochemical environment for the uninformed simulator is constructed with reference to a small ensemble of 3-D simulations, designed to be representative of the infinite set of 3-D simulations covering a parameter space of interest χ . If we denote an output value from the simulator with biogeochemical environment vector \mathbf{B} and parameter vector \mathbf{x} by $g(\mathbf{B}, \mathbf{x})$ and the corresponding output from the target model by $f(\mathbf{x})$, then for parameter vector \mathbf{x}_o

$$f(\mathbf{x}_o) = g(\overline{\mathbf{B}}, \mathbf{x}_o) + \epsilon_1 \quad (8)$$

where $\overline{\mathbf{B}}$ is an estimate of the expected environment $E[\mathbf{B}(\mathbf{x})] : \mathbf{x} \in \chi$ and ϵ_1 is a stochastic residual. This is the *uninformed simulator residual* and its negated value is the *uninformed simulator error*. The simulator output may have biases so the residual ϵ_1 is not assumed to have zero mean.

The environment model consists of a model for $E[\mathbf{B}(\mathbf{x})]$, referred to as the *mean environment model*, and a *stochastic environment generator* that is used in quantifying the uncertainty of the simulator output. The environment model assumes multi-variate Gaussian probability

distributions for a vector $\mathbf{S}(t_o)$ that specifies the initial state and for the applied advective flux perturbation vector \mathbf{P}^* . \mathbf{S} is an alternative description of the state \mathbf{C} . It comprises elements \sqrt{c} for each primary tracer concentration c in \mathbf{C} and composition ratios c_i/c_j for each secondary tracer concentration c_i in \mathbf{C} . c_j is the concentration of the associated primary tracer at the same depth level. An estimate of $E[\mathbf{B}(\mathbf{x})]$ is given by the ensemble means of $\mathbf{S}(t_o)$ and \mathbf{P}^* from the 3-D ensemble.

2.3 The uninformed emulator

If an array of 1-D simulators is to be used to make robust inferences about the target model, it must be combined with uncertainty estimates for its predictions of target model output in the form of predicted error statistics. The combination of the uninformed simulator array with its predicted error statistics is referred to here as the *uninformed emulator*. This is the complete mechanistic emulator for the target model.

Two different methods are used in this study for quantifying uncertainty in the uninformed simulator output: a *direct method* and an *indirect method*. In the direct method, statistics for ϵ_1 are estimated by comparing simulator and target model output for matching parameter vectors, using the target model output available from our small 3-D ensemble. In the indirect method, the uncertainty introduced by using the mean environment vector $\bar{\mathbf{B}}$ in place of the unknown environment vector $\mathbf{B}(\mathbf{x}_o)$ is treated separately from that due to other simulator error sources. It is quantified by an uncertainty analysis, using the stochastic environment generator to create multiple realizations of the unknown environment. Uncertainty from other sources is estimated by applying the direct method to $g[\mathbf{B}(\mathbf{x}_o), \mathbf{x}_o]$, referred to as the *informed simulator*. The indirect method is more complicated to apply than the direct method but is less dependent on the small target model ensemble. This means that the indirect method could be more robust than the direct method in situations where the ensemble poorly represents the variability of target model solutions over the parameter space χ .

2.3.1 Direct method for uncertainty quantification

In the direct method, values of ϵ_1 for the variable of interest at each point in space and time are determined from matching pairs of uninformed simulator and target model output values using Eq. (8). Statistics for ϵ_1 are then estimated from this sample. A conceptual overview of the data flow in the emulator construction and evaluation process is given in Fig. 1.

The processing is divided into a *construction phase* and an *application phase*. In a practical application, the construction phase is intended for single execution, whereas the application phase must be executed for each trial parameter vector. The procedure for assessment of the uninformed emulator against a known truth is shown as an extension to the application phase.

Error statistics must be determined using target model data that are independent from those used in the simulation. This means that, in the construction phase, target model ensemble members used to determine ϵ_1 for the simulator output must be different from those used to construct the mean environment model for the simulator input. Furthermore, any target model ensemble member used to assess the uninformed emulator performance must be different from any ensemble member used in the construction phase.

2.3.2 Indirect method for uncertainty quantification

The indirect method requires an explicit quantification of the uncertainty associated with use of the mean environment vector $\bar{\mathbf{B}}$ in lieu of unavailable parameter-specific environment information. Reliance on $\bar{\mathbf{B}}$ introduces a parameter-dependent source of environment-induced error into the simulation. The resulting contribution to simulation error is referred to as the *parametric environment error*. To define it, we consider a perfect simulator $g_{\text{T}}(\cdot, \cdot)$, such that

$$f(\mathbf{x}_o) = g_{\text{T}}[\mathbf{B}_{\text{T}}(\mathbf{x}_o), \mathbf{x}_o] \quad (9)$$

where \mathbf{B}_{T} is the complete and accurate description of the local biogeochemical environment in the 3-D simulation, including advective and diffusive flux perturbations. The simulator is perfect in the sense that it exactly reproduces the results of the 3-D simulation. Introducing parametric uncertainty in the biogeochemical environment and representing the environment

by its expectation then gives

$$f(\mathbf{x}_o) = g_T\{\mathbb{E}[\mathbf{B}_T(\mathbf{x})], \mathbf{x}_o\} + \epsilon_B : \mathbf{x} \in \chi. \quad (10)$$

where ϵ_B is a stochastic residual, possibly with a non-zero mean. This is the negated parametric environment error or *parametric environment residual*.

5 It is important to note that many different designs are possible for a perfect simulator satisfying Eq. (9), having different formulations for concentration dependency in the flux divergence tendencies. Variants of the applied perturbation \mathbf{P}^* will give different results for the simulator term in Eq. (10), where the environment is not consistent with the simulation state, and therefore different residuals. The parametric environment error is therefore not just a property of the target model but depends also on the simulator design.

10 Combining Eqs. (8) and (10), the residual for the target model output with respect to the uninformed simulator output can be expressed as

$$\epsilon_1 = \epsilon_S + \epsilon_B \quad (11)$$

where ϵ_S is a stochastic residual given by

$$15 \quad \epsilon_S = g_T\{\mathbb{E}[\mathbf{B}_T(\mathbf{x})], \mathbf{x}_o\} - g(\overline{\mathbf{B}}, \mathbf{x}_o). \quad (12)$$

ϵ_S is the departure of the hypothetical output of the perfect simulator with the true mean environment from the output of the uninformed simulator. The first term describes a perfect *mean environment simulation*, while the second term described its approximation by the simulator. In this context, we can refer to the uninformed simulator as a mean environment simulator. We refer to ϵ_S as the *mean environment simulation residual*. Mean environment simulation error (the negated residual) is caused by basic simulation errors that are not associated with parametric uncertainty in the environment.

20 It is not possible to evaluate the perfect simulator term in Eq. (12) and directly determine values for ϵ_S . However, we can get a handle on the impact of basic simulation errors from

analysing the informed simulator. The relationship between the target model output for \mathbf{x}_o and that of the corresponding informed simulator is given by

$$f(\mathbf{x}_o) = g[\mathbf{B}(\mathbf{x}_o), \mathbf{x}_o] + \epsilon_2 \quad (13)$$

where $\mathbf{B}(\mathbf{x}_o)$ is the environment data derived from 3-D simulation output for \mathbf{x}_o and ϵ_2 is a stochastic residual, possibly having non-zero mean, referred to as the *informed simulator residual*. Its negated value is the *informed simulator error*.

The residuals ϵ_2 and ϵ_S are closely related, in that the input $\mathbf{B}(\mathbf{x}_o)$ in the informed simulator is intended to approximate the true parameter-specific environment in the same way that $\bar{\mathbf{B}}$ in the uninformed simulator (or mean environment simulator) is intended to approximate the perfect simulator input $E[\mathbf{B}_T(\mathbf{x})]$. Both residuals are affected by basic simulation errors. The difference is that the environment in Eq. (12) is not specific to the parameter vector \mathbf{x}_o .

The uninformed simulator is one of a set of generic simulators, in which the constraint that the input environment is intended to represent the parameter-specific environment does not apply. In generic simulators, inconsistencies between the environment and the simulation state are likely to be greater than in the informed simulator. The mean environment simulation residual ϵ_S may therefore be more sensitive to the concentration-dependency formulation than the informed simulator residual ϵ_2 . Nevertheless, to model ϵ_S we make the pragmatic assumption that it is identically distributed to ϵ_2 . Statistics for ϵ_2 are determined by direct comparison of informed simulator output with true output records from the target model.

The model for the parametric environment residual ϵ_B is derived from a parametric uncertainty analysis, following Hemmings and Challenor (2012). The environment corresponding to the trial parameter vector is unknown so we examine the distribution of the residual over many possible environments, aiming to achieve adequate coverage of the environment space that maps to the parameter space of interest. The method involves running a 1-D ensemble simulation based on a sample of environment realizations. These are generated using the mean environment model and stochastic environment generator introduced in Sect. 2.2.

The environment generator uses independent statistical models for generating the initial state and the input flux perturbations. For each of these two data sets, separate multi-variate Gaussian

models are constructed using Empirical Orthogonal Functions that capture the dominant modes of variability in the target model ensemble output at each site. The statistical models for the initial state preserve spatial covariances (in the vertical) and covariances between the biogeochemical properties, as characterized by the first 5 EOFs of the sample anomalies, anomalies being determined with respect to the ensemble means. The statistical models for the advective flux perturbations preserve temporal and spatial covariances and covariances between the 8 primary tracers, again as characterized by the first 5 EOFs of the anomalies.

To derive the statistical model for a simulator's initial state from a target model ensemble of size n , an $n \times m$ matrix \mathbf{Y}_{3d} is constructed containing the n available instances of the initial state, as defined by the alternative state vector \mathbf{S} . (m is the number of elements in \mathbf{S} .) If $y_{\cdot j}$ is the mean and s_j^2 the variance of the j th column of \mathbf{Y}_{3d} , then the matrix \mathbf{Z}_{3d} with elements

$$z_{ij} = \frac{y_{ij} - y_{\cdot j}}{s_j} \quad (14)$$

is the normalized form of \mathbf{Y}_{3d} for which each column has zero mean and unit variance.

The environment generator uses the eigenvalues and eigenvectors obtained from the spectral decomposition of the correlation matrix for \mathbf{Z}_{3d} :

$$\mathbf{\Sigma} = \mathbf{Z}_{3d}^T \mathbf{Z}_{3d} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T. \quad (15)$$

$\mathbf{\Lambda}$ is a diagonal matrix with diagonal elements $\lambda_1 \geq \lambda_2 \dots \geq \lambda_m$ containing the eigenvalues of $\mathbf{\Sigma}$. Columns of \mathbf{V} are the corresponding eigenvectors.

A data set containing N realizations of the alternative state vector is generated by

$$\mathbf{Z}_{1d} = \mathbf{Q}_1 \mathbf{\Lambda}_p^{\frac{1}{2}} \mathbf{V}_p^T + \mathbf{Q}_2 \sqrt{1 - \text{tr}(\mathbf{\Lambda})} \quad (16)$$

where the subscript p is used to indicate the first p rows and columns of $\mathbf{\Lambda}$ and columns of \mathbf{V} . (Here $p = 5$.) \mathbf{Q}_1 is an $n \times p$ matrix of random values and \mathbf{Q}_2 is a column vector of random values. The random variates are independent and normally distributed with zero mean and unit variance. \mathbf{Z}_{1d} is back-transformed (re-arranging Eq. 14) to obtain an $N \times m$ matrix containing

N realizations of the state vector $\mathbf{S}(t_o)$ for the 1-D environment ensemble. The same analysis is applied to the n available instances of the advective flux perturbation vectors from the 3-D ensemble to generate N realizations of the \mathbf{P}^* vector.

Each of the N randomly generated environment realizations is used to provide a separate estimate of the parametric environment residual corresponding to a possible truth. For the i th ensemble member this is

$$\epsilon_{B_i} = g(\mathbf{B}_i, \mathbf{x}_o) - g(\overline{\mathbf{B}}, \mathbf{x}_o) \quad (17)$$

where \mathbf{B}_i is the i th environment realization generated by the environment model. For the true environment, \mathbf{B}_i would be $\mathbf{B}(\mathbf{x}_o)$, as in the informed simulator. The environment residual statistics $\text{var}(\epsilon_B)$ and $E(\epsilon_B)$ are approximated by $\text{var}(\epsilon_{B_i})$ and $E(\epsilon_{B_i}) : i \in \{1, \dots, N\}$. In Eq. (17), we rely on the simulator $g(.,.)$ to provide estimates for the terms $f(\mathbf{x}_o)$ and $g_T(E[\mathbf{B}_T(\mathbf{x})], \mathbf{x}_o)$ in Eq. (10). Thus, the estimated environment residual statistics are to some extent affected by basic simulation errors and will not be strictly independent of the statistics for the mean environment simulation residual ϵ_S .

It should be noted that the residual ϵ_B and its predicted distribution are dependent on the trial parameter vector \mathbf{x}_o . Hemmings and Challenor (2012) demonstrated that the dependency of environment error variance estimates on variations in the simulation trajectory over the parameter space is potentially important in the context of a parametric analysis. For this reason, estimation of the environment residual statistics must be performed for each trial parameter vector in the analysis, so is a significant overhead.

If the underlying distributions of the residuals ϵ_S and ϵ_B are taken to be Gaussian then they are fully described by their means and variances. Statistics for the uninformed simulator residual ϵ_1 are obtained under the assumption that ϵ_S and ϵ_B can be considered only weakly dependent such that

$$E(\epsilon_1) = E(\epsilon_S) + E(\epsilon_B) \quad (18)$$

and

$$\text{var}(\epsilon_1) \approx \text{var}(\epsilon_S) + \text{var}(\epsilon_B). \quad (19)$$

Any indirect dependency between ϵ_S and ϵ_B that might arise from their dependencies on the simulator design are ignored. The uninformed simulator statistics are determined by substituting our estimates for the residual statistics for each error component in Eqs. (18) and (19). In doing so, we also ignore potential dependency arising from the effect of basic simulation errors on $\text{var}(\epsilon_{Bi})$.

A conceptual overview of the data flow for the indirect method is given in Fig. 2. Once again, the processing is divided into a construction phase intended for single execution and an application phase to be applied with each trial parameter vector. The procedure for assessment of the uninformed emulator is included in the application phase.

3 Experimental method

Anticipating the use of satellite ocean colour data for model calibration, an emulator was constructed for the NEMO-MEDUSA surface chlorophyll output at an array of oceanic sites. The surface chlorophyll concentration is the sum of the surface level chlorophyll concentrations for the two phytoplankton types. Data for defining the biogeochemical environment were provided by a 10 member reference ensemble of global 3-D simulations with the NEMO-MEDUSA target model. For emulator assessment, the known “truth” for a given trial parameter vector is defined by chlorophyll output from a target model simulation with that parameter vector.

3.1 1-D experimental framework

To provide a representative range of oceanic conditions for the experiments, 12 sites were selected, located on a meridional transect along 20° W in the North Atlantic at 5° intervals from 5 to 60° N. This spans the sub-tropical gyre and temperate regions further north where large spring blooms are typical, extending into the sub-polar gyre south of Iceland. To the south, it also crosses a high productivity region off the East African coast between the shelf break and the Cape Verde Islands.

Physical forcing data for the 1-D experiments, in the form of vertical velocity w_p , the vertical diffusion coefficient K_ρ and temperature are taken from 5 day mean output common to all of the 3-D NEMO-MEDUSA simulations. 5 day mean time series of downwelling solar radiation at the sea surface and the soluble iron flux from dust deposition are likewise taken from 5 day data common to all reference simulations.

Biogeochemical environment vectors for the 1-D experiments are based on initial state vectors and applied perturbation vectors from 1 or more 3-D simulations. Initial concentrations are taken from NEMO-MEDUSA restart files. Approximate values for the applied perturbation p^* are derived from the target model's 5 day mean current vector and primary tracer concentration fields using Eq. (7).

1-D simulations use the same vertical grid as the 3-D NEMO-MEDUSA simulations. The dynamics of interest are largely confined to the upper ocean where the seasonal signal is most pronounced. A depth threshold of 1000 m was therefore chosen for the simulations, reducing the number of model levels from 63 to 37 with consequent computational savings. Level 36 spans the 1000 m threshold and Level 37 is included purely to act as a sink for detritus entering from above. In the target model, sinking detritus is re-mineralized at the bottom of the water column. In the simulator it is re-mineralized in Level 37 instead and the vertical velocity and diffusion at the bottom of Level 36 are set to zero to prevent any interaction between Level 37 and the water column above. Zeroing the vertical velocity does have the effect of introducing an anomalous divergence in the vertical flow but the effect on the overall simulation is negligible. The upper ocean levels have boundaries at depths 6, 12, 19, 25, 32, 39, 46, 54, 62, 71, 80, 90, 100, 112, 124, 137, 152, 168, 187, 207, 229, 254, 281, 312, 347, 386, 429, 477, 531, 591, 656, 729, 809, 896, 991 and 1093 m.

The schemes used for vertical tracer transport are the same as those used in the target model and are described by Madec (2008). The diffusion scheme is an implicit scheme and the advection scheme is the Monotonic Upstream Scheme for Conservative Laws (Van Leer, 1977; Hourdin and Armengaud, 1999), introduced into NEMO for use in biogeochemical modelling studies by Lévy et al. (2001). A 1 h forward Euler time step is used.

3.2 Model parameter space

Full details of the derivation of the parameter space for the emulation experiments are given in Appendix B. Initially, a 28-dimensional parameter space of interest was defined; 28 parameters of particular relevance to the seasonal plankton dynamics in the upper ocean were selected from a set of 60 potential input parameters in the MarMOT 1.1 implementation of MEDUSA. The parameter bounds were defined according to a set of rules designed to ensure that parameter values within the bounds are biologically plausible with respect to their defined roles.

The set of adjustable input parameters differs from the set of internal model parameters defined by Yool et al. (2011) due to a number of modifications made to facilitate parametric analyses. For example, where pairs of parameters such as rate parameters are used in the model for the two different phytoplankton types, the diatom parameter has been replaced in the input vector by the ratio of the two internal parameters. The input non-diatom parameter then scales both of the internal phytoplankton parameter values without affecting their relationship, while the new input parameter controls the relationship. The zooplankton parameters are treated similarly. The changes allow us to consider the effects of a phytoplankton rate parameter or a zooplankton rate parameter on the system without having to consider the impact of directly changing the relationship between rates for closely related plankton types. It is then easier to interpret parameter effects at a high level of abstraction which facilitates comparison with simpler models where parameters represent rates for more aggregated plankton compartments.

The dimensionality of the initial parameter space was reduced further with reference to a sensitivity analysis, performed at the experimental sites, to identify parameters that are influential with respect to annual primary production and sinking particle flux outputs from the model (see Appendix B). Improving the reliability of these outputs in the target model will be important for understanding and predicting change in the global carbon cycle. 8 model parameters were chosen on the basis of the findings. The corresponding parameter space is defined by Table 1.

One finding of the sensitivity analysis was that the input parameters controlling the relationship between associated internal parameters for different plankton types were less influential than the input parameters exerting control over the different plankton types jointly. None of

the input parameters from the first set were selected. The mapping of input parameters to internal parameters means that varying any of the 5 non-diatom phytoplankton parameters in Table 1 will also change the corresponding internal diatom parameters in proportion. The non-diatom density-independent loss rate and half-saturation concentration for density-dependent loss will additionally affect the corresponding internal parameters for both zooplankton types in proportion and the microzooplankton grazing half-saturation concentration will affect the corresponding internal parameter for mesozooplankton in the same way.

3.3 3-D reference simulations

A 10 member ensemble of 3-D simulations was used to create a reference sample of NEMO-MEDUSA output data that is representative of variability in the target model solution over the defined parameter space. The 10 parameter vectors are distributed in parameter space according to a Latin hypercube design (McKay et al., 1979). For improved coverage, a “maximin” criterion (Johnson et al., 1990) was applied to 1000 randomly generated hypercubes: the hypercube design is selected that maximizes the smallest Euclidean distance between parameter vector pairs in terms of their positions on a parameter space grid with an equal number of intervals in each dimension. Grid intervals are in log units for rate parameters and half-saturation concentrations.

The chosen parameter vectors are given in Table 2. NEMO-MEDUSA integrations were performed for each of the 10 parameter vectors to provide representative output for a 2 year period, beginning in 1997. The second year, 1998, is the first complete year for which satellite ocean colour data from the SeaWiFS sensor are available (although these data are not used in the present study). The integrations, at 1° horizontal resolution, were initialized from the NEMO-MEDUSA simulation of Yool et al. (2011) at the beginning of 1995 and integrated for 4 years with their respective modified parameter sets, thereby allowing a 2 year spin-up period prior to any analysis to attenuate the worst effects of transient behaviour with respect to the seasonal cycle in the upper ocean. A longer spin-up time would normally be envisaged for a practical application, consistent with the intended use of the target model.

The 3-D reference sample is used in two ways. Chlorophyll records are used for evaluating 1-D simulation error, while the initial concentrations and horizontal gradients of the biogeochemical tracers are used to provide parameter-specific environment information for 1-D simulator construction.

5 3.4 Emulator construction and assessment

Performance of the basic 1-D simulator array is evaluated, with respect to surface chlorophyll, for a set of trial parameter vectors for which the true target model output is known. The performance of emulators constructed using the two uncertainty quantification methods is then assessed. Finally, to explore the importance of the lateral flux perturbations, we assess the performance of simulator arrays in which these are omitted. In this context, the behaviour of an alternative array employing informed simulators is examined in addition to that of the uninformed simulator array used in the emulator. Doing this allows us to see the impact of omitting lateral flux perturbations in a scenario where other error sources are minimized. The experimental methods for the assessments are as follows.

15 3.4.1 Simulator assessment

Informed simulator skill is described by error statistics calculated from a set of 10 experiments with the representative parameter vectors defined in Table 2, so that each experiment corresponds to one of the available 3-D reference simulations. In each experiment, the informed simulator is initialized at the start of 1997 and run for 2 years. If the set of representative parameter vectors is denoted by $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{10}\}$ then the trial parameter vector for the i th experiment is \mathbf{x}_i and the environment is defined by the 3-D ensemble member with parameter vector \mathbf{x}_i .

The error statistics describing the skill of the uninformed simulator were determined from 10 similar experiments, covering the same time period. One experiment was performed for each parameter vector in \mathbf{X} but simulator construction was performed on a leave-one-out basis: in the i th experiment, the trial parameter vector is \mathbf{x}_i and the mean environment is derived from the 9 NEMO-MEDUSA ensemble members with $\mathbf{x} \neq \mathbf{x}_i$, $\mathbf{x} \in \mathbf{X}$, leaving the NEMO-MEDUSA

output $f(\mathbf{x}_i)$ as independent data for validation. Thus, each experiment uses a slightly different version of the simulator, constructed by applying the same method to a different 9 member ensemble.

Error statistics are calculated with respect to the log-transformed 5 day mean chlorophyll output. The log transformation applied to the 5 day means acts to stabilize the error variance which otherwise tends to increase with increasing chlorophyll concentration. Its use in the analysis of surface chlorophyll variability is strongly supported by theoretical considerations and empirical data (Campbell, 1995).

3.4.2 Assessment of the full emulator

Validation of the complete uninformed emulator for surface chlorophyll is by analysis of the results from the 10 leave-one-out experiments, taking into account the predicted simulator error statistics to determine the emulator robustness. These uncertainty estimates are, like the simulator itself, required to be independent of parameter-specific environment information. Thus, for the i th experiment, they are derived using the 9 NEMO-MEDUSA ensemble members with $\mathbf{x} \neq \mathbf{x}_i$. The uninformed emulator uncertainty is quantified using the direct and indirect methods.

When the indirect method is used, the 9 NEMO-MEDUSA ensemble members are used to derive statistics for the two component residuals ϵ_S and ϵ_B . In the estimation of the statistics for the mean environment simulation residual ϵ_S (assumed identically distributed to the informed simulator residual ϵ_2), the 3-D ensemble members are required for comparison with the corresponding informed simulators to determine informed simulator error. In the estimation of the statistics for the parametric environment residual ϵ_B , the 3-D ensemble is required for building the environment model used in the parametric uncertainty analysis.

When the direct method is used, the 9 NEMO-MEDUSA ensemble members are used to derive statistics for the uninformed simulator residual ϵ_1 . Each of the 9 corresponding uninformed simulators require independent data for their mean environment input. In the i th experiment, the mean environment for the uninformed simulator with parameter vector \mathbf{x}_j is derived from the 8 NEMO-MEDUSA members with $\mathbf{x} \neq \mathbf{x}_i \cap \mathbf{x} \neq \mathbf{x}_j$. As a result, simulators must be constructed

with 90 different mean environment estimates to calculate the uncertainty estimates for the 10 experiments.

For the uncertainty quantification analyses, Gaussian error distributions in log-transformed chlorophyll are assumed so that the resulting probability density functions for the residuals are fully described by their mean and variance, both of which are allowed to vary in time and between sites. The residuals are defined with respect to log-transformed 5 day mean chlorophyll concentrations. Their predicted distributions are described by their monthly means and variances, interpolated to 5 day intervals. Appendix C gives the estimation method for the residual statistics and the resulting time series.

4 Results

The surface chlorophyll records from the 3-D NEMO-MEDUSA reference ensemble at each of the experimental sites are shown in Fig. 3. This shows the spatial variation in chlorophyll from values a little above 0.001 mg m^{-3} in the oligotrophic gyre at 30° N and 35° N for Parameter Set 6 to seasonal highs associated with the spring bloom in temperate regions ($45\text{--}60^\circ \text{ N}$), approaching 10 mg m^{-3} for a number of the parameter vectors. It also illustrates the variability in the seasonal response of the plankton dynamics which is generally stronger at the more northerly sites.

The variation between records produced by different parameter vectors is large compared with the seasonal variability. At some sites, particularly $5\text{--}10^\circ \text{ N}$ and $25\text{--}35^\circ \text{ N}$, the parameter dependency manifests primarily as a control on the overall chlorophyll concentration level in the surface layer, throughout the annual cycles. These are generally the more oligotrophic sites, where concentrations remain below or very close to 1 mg m^{-3} for all parameter vectors. At other sites, particularly in the north, the different parameters also have a notable influence on the dynamic range and there is some evidence of an impact on the characteristics of the spring bloom.

Some parameter vectors tend to have the same effect on overall surface chlorophyll levels at all sites. For example, Parameter Set 10 gives elevated levels over the whole data set. However,

this is not generally the case. Parameter Set 6, for example, shows a strong tendency to give low chlorophyll concentrations at many of the sites but gives some of the higher concentrations at 55 and 60° N. With this parameter vector, the phytoplankton light-response controlled by α_{Pn} is exceptionally strong and nutrient-limitation is reduced by low half-saturation concentrations $k_{N,Pn}$ and $k_{Fe,Pn}$. As a result, the phytoplankton can achieve very high growth rates. This can cause blooms that lead to long-term nutrient depletion as a consequence of organic material sinking out of the euphotic zone. Subsequent growth is then inhibited. At 4 sites (5, 10, 30 and 35° N), nitrogen depletion during the 2 year spin-up period results in very low chlorophyll concentrations at the start of 1997 which remain relatively low throughout 1997 and 1998.

Parameter Sets 1 and 4 also lead to some interesting site-specific impacts. They are associated with very low winter-time chlorophyll concentrations at the most northerly sites, particularly in 1997, although are associated with some of the highest concentrations throughout 1997 and 1998 at the most southerly sites. These parameter vectors combine low α_{Pn} values with low values for the grazing half-saturation concentration k_{μ} , reducing phytoplankton production at low light levels and making them more susceptible to zooplankton grazing. This makes the phytoplankton less well-suited to over-wintering at the high latitude sites where light availability is very low due to the combination of low surface irradiance and deep winter mixing.

The strong variation between parameter vectors indicates the potential for significant constraints on the parameter values to be realized by the assimilation of satellite chlorophyll data.

4.1 Emulator prediction of target model output

Chlorophyll concentrations given by the uninformed simulator at all sites are compared against the corresponding values from the matching 3-D experiment in Fig. 4. Data are shown for the 1998 annual cycle only so are representative of the simulator performance one year on from its initialization year, during which errors have had time to develop.

The correlation between simulator and target model values is good. Pearson's correlation coefficient r for the simulator and target model output is 0.91, indicating that 83 % of the variance in the log-transformed surface chlorophyll from the simulator array is explained by the target model output. There are some notable examples of poor performance though. In particular, the

results for Parameter Set 6 indicate a strong positive bias, with the simulator array overestimating some surface chlorophyll values by an order of magnitude. There are some fairly large negative biases for other parameter sets, notably Parameter Sets 7 and 10 at mid-range concentrations, although these are less systematic. Also, the simulator array poorly reproduces the relatively low variability in chlorophyll associated with Parameter Set 1.

The chlorophyll output from the uninformed emulator includes a bias correction term which depends on the uncertainty quantification method. (This corrects for spatio-temporal biases rather than for parameter-related biases.) When using the direct uncertainty quantification method, the bias-corrected error in log-transformed 5 day mean chlorophyll is

$$d_{Ud} = g(\bar{\mathbf{B}}, \mathbf{x}_o) + \bar{u}_1 - f(\mathbf{x}_o) \quad (20)$$

where \bar{u}_1 is our estimate of $E(\epsilon_1)$. When using the indirect method, the bias correction includes corrections for both the mean environment simulation bias and the bias associated with parametric environment uncertainty. The bias-corrected error is then

$$d_{Ui} = g(\bar{\mathbf{B}}, \mathbf{x}_o) + \bar{u}_S + \bar{u}_B(\mathbf{x}_o) - f(\mathbf{x}_o) \quad (21)$$

where \bar{u}_S and \bar{u}_B are our estimates of $E(\epsilon_S)$ and $E(\epsilon_B)$ respectively and $\bar{\mathbf{B}}$ is our estimate of the mean environment. The estimates \bar{u}_1 and \bar{u}_S were determined without reference to results for Parameter Set 6. These were excluded on the basis of the unrepresentative simulator performance, to avoid excessive influence from a single outlier. Time series of \bar{u}_2 and \bar{u}_S are therefore based on an ensemble size of 8 (or 9, when Parameter Set 6 is the trial parameter vector).

Error statistics for the uninformed emulator results are given in Fig. 5. Results are presented for the basic simulator array with no bias correction ($\bar{u}_1 = \bar{u}_S = \bar{u}_B = 0$) and for the full emulator with bias correction. There are only minor differences between the mean and r.m.s. values for d_{Ud} and d_{Ui} .

Biases are reduced by the emulator's bias correction scheme, irrespective of the method used. Time series of simulator bias before and after correction show that in both cases the bias correction is effective at all sites, with the possible exception of 20° N where d_{Ui} shows the introduction of a negative bias in the summer of 1998 when using the indirect uncertainty quantification

method. In particular, note that the summer 1998 bias at 60° N is largely removed and the correction is particularly effective in removing negative bias at some of the more oligotrophic sites (5° N and 25–30° N) and at 50° N in 1997.

The relatively high r.m.s. errors for early 1997 at most sites are the consequence of transient behaviour associated with error in the initial conditions. This source of error seems to influence the model primarily in the early half of the year, before the local dynamics start to dominate over the environmental influence. The lack of parameter-specific information about the lateral fluxes appears to be a less dominant source of simulation error. Nevertheless, it does contribute strongly to the relatively large 1998 errors at 5° N and at 50° N.

4.2 Robustness of the emulator

The robustness of the uninformed emulator is assessed by comparing the MarMOT-MEDUSA chlorophyll records with the NEMO-MEDUSA results for the matching parameter sets, taking into account the quantified emulator uncertainty in terms of the predicted bias and error variance. The results are presented here in terms of the normalized emulator error, which is the error in the bias-corrected simulator output scaled by the reciprocal of its predicted standard deviation. The scaling factor ensures that the predicted normalized error distribution for both versions of the emulator is Gaussian with zero mean and unit standard deviation at all times and locations.

The normalized uninformed emulator error for each log-transformed 5 day mean surface chlorophyll concentration depends on the uncertainty quantification method. For the direct method, it is given by

$$D_{\text{Ud}} = \frac{g(\bar{\mathbf{B}}, \mathbf{x}_o) + \bar{u}_1 - f(\mathbf{x}_o)}{s_1} \quad (22)$$

where s_1^2 is our estimate for $\text{var}(\epsilon_1)$. For the indirect method, it is

$$D_{\text{Ui}} = \frac{g(\bar{\mathbf{B}}, \mathbf{x}_o) + \bar{u}_S + \bar{u}_B(\mathbf{x}_o) - f(\mathbf{x}_o)}{\sqrt{s_S^2 + s_B^2(\mathbf{x}_o)}} \quad (23)$$

where s_S^2 and s_B^2 are our estimates for $\text{var}(\epsilon_S)$ and $\text{var}(\epsilon_B)$ respectively. s_1^2 and s_S^2 , like the residual mean estimates \bar{u}_1 and \bar{u}_S , were determined without reference to the results for Parameter Set 6, so were likewise based on a sample size of 8 (or 9 when Parameter Set 6 is the trial parameter vector). The denominator in Eq. (22) varies between 0.014 and 0.62 \log_{10} units and that in Eq. (23) varies between 0.015 and 0.50 \log_{10} units (with chlorophyll concentration in mg m^{-3}). Further details of the residuals' statistics and their variation in time and between sites can be found in Appendix C.

The normalized uninformed emulator errors for each experiment are shown in Fig. 6. In Experiment 6 (pertaining to trial Parameter Set 6), the positive errors already noted are extreme, relative to the predicted error variance. This is a consequence of the unusually large simulator errors associated with Parameter Set 6. The atypical behaviour associated with this parameter vector may be truly representative of the model dynamics over a significant region of parameter space. However, such detail is not resolved with our small sample so is not represented in the data used for emulator construction. Large normalized error values in Experiment 6 are therefore unsurprising.

When the indirect uncertainty quantification method is used, Fig. 6 shows that there are also very large extremes associated with the post-initialization phase, particularly at 55°N and 60°N . These high D_{Ui} values occur for experiments with the two parameter vectors that were seen to cause unusually low winter-time chlorophyll concentrations at the start of 1997 in the target simulation (Fig. 3, Parameter Sets 1 and 4). Fortunately, at these sites, the extreme error appears fairly transient, lasting only a few months. At other sites, in particular at 5°N , D_{Ui} remains correlated to some extent with its early 1997 value over the whole 2 year period, suggesting that parametric error in the initial state may be introducing persistent biases. This pattern seems to be a common feature of the more oligotrophic sites, being reflected also at latitudes from 25 to 35°N . At more northerly sites, there is a tendency for persistent biases over long time periods where relatively large errors occur (e.g. at 45°N and 50°N for the indirect method) but this pattern develops later with no obvious connection to initialization error.

Comparing the two uncertainty quantification methods, it is seen that D_{Ui} initially tends to be larger than D_{Ud} at all sites. The post-initialization D_{Ud} values are more consistent with their

5 predicted distribution. In particular, the extreme positive D_{U_i} values seen in early 1997 are not replicated in D_{U_d} . From these observations, it is clear that the indirect method is generally less effective at quantifying initial uncertainty. Furthermore, at the oligotrophic sites where the early 1997 biases tend to persist, there is a general tendency for D_{U_i} to be larger than D_{U_d} over the 2 year period.

10 The normalized error distributions for the uninformed emulators are compared with the predicted distribution in Fig. 7. Results, including 1998 data only, are shown for each site. Experiment 6 is excluded to allow the results for the remaining experiments to be more clearly represented. The emulator with direct uncertainty quantification appears fairly robust with D_{U_d} distributions broadly similar to the predicted distribution at all sites. The worst performance is arguably at 30° N where there are a significant proportion of anomalously low values associated with persistent negative errors in the experiments with Parameter Sets 1 and 4 (Fig. 6). However, D_{U_i} shows a strong tendency to be larger than expected at a number of the sites. In general, these are the sites that have already been associated with persistent error in some of the experiments (5° N, $25\text{--}35^\circ$ N, $45\text{--}50^\circ$ N). A smaller proportion of the D_{U_d} values at 15 and 20° N are rather larger than predicted. These are associated with extreme negative biases occurring in Experiment 9 that persist only for a month or two.

20 Table 3 summarizes the uninformed emulator results in terms of the mean and standard deviation of the normalized errors. Statistics are given for all 10 experiments combined and in brackets for the 9 experiments excluding Experiment 6. The difference between the two sets of results illustrate to some extent the sensitivity of the evaluation statistics to sampling error.

25 When the emulator performance with direct uncertainty quantification is evaluated over all experiments and all sites, the D_{U_d} standard deviation is rather high at 1.41, suggesting that the emulator is a little over-confident. When Experiment 6 is excluded from the evaluation, the standard deviation drops to 1.13. Whether or not this is a more appropriate measure of performance depends on the extent to which the model dynamics with Parameter Set 6 are representative of its behaviour over a significant region of parameter space. The performance with respect to the other parameter vectors is fairly reliable at all sites, with standard deviations from 0.98 to 1.31 and very little sign of post-correction bias shown by D_{U_d} mean values. All

but 2 of the standard deviations are above 1, indicating a slight tendency for the spread of the simulator residuals to be under-estimated. When Experiment 6 results are included in the evaluation data set, this tendency for over-confidence is more evident and there are notable positive biases at a number of sites (D_{Ud} mean greater than 0.3 at 10, 15 and 50° N). These are associated with relatively large D_{Ud} standard deviations (1.55 to 2.07).

The high standard deviation in D_{Ui} of 1.82, is consistent with results already presented that show the emulator with indirect uncertainty quantification has a clear tendency towards over-confidence in its predictions. If Experiment 6 is excluded, the overall standard deviation is less at 1.39, but the performance still leaves some room for improvement. The over-confidence is particularly notable at the highly oligotrophic site at 5° N, with a standard deviation of just over 2 reflecting the persistent parameter-specific biases already noted at this site (Fig. 6). There is also a tendency for the emulator with indirect uncertainty quantification to significantly underestimate chlorophyll concentrations. In particular, the 9 parameter vector sample shows large negative biases of around -0.7 at some of the other oligotrophic sites (20, 25 and 35° N). Fairly large negative biases of -0.30 and -0.38 are also seen at 5 and 15° N respectively. Nevertheless, the performance at a number of the sites is good. The subset of 5 sites at 10° N, 40–45° N and 55–60° N has standard deviations in the range 0.74 to 1.32 with small biases (-0.1 to 0.2).

4.3 The importance of lateral advection

In the majority of site-based calibration studies, the effect of lateral advection is ignored. It is useful then to examine the extent to which the skill of our 1-D simulations is dependent on the explicit representation of the advective flux divergence term. Fig. 8 shows the chlorophyll values given by the uninformed simulator array compared with the matching target model output when the uninformed simulators are run with all lateral flux perturbations removed. Comparison with Fig. 4 shows that the omission of lateral flux perturbations degrades the performance of the simulator array considerably. Pearson's correlation coefficient r for the simulator and target model output drops from 0.91 to 0.75, indicating that just 56% of the variance in the

log-transformed surface chlorophyll from the simulator array is explained by the target model output, compared with 83 % in the standard simulator array with lateral flux perturbations.

The impact of omitting lateral flux perturbations is most clearly seen in the performance of the informed simulator array, where removing the effects of the parametric environment uncertainty minimizes other sources of error. The initial state error for this simulator array is zero and the lateral flux perturbations are parameter-specific. The error for each log-transformed 5 day mean chlorophyll concentration is defined by

$$d_t = g[\mathbf{B}(\mathbf{x}_o), \mathbf{x}_o] - f(\mathbf{x}_o) \quad (24)$$

where $\mathbf{B}(\mathbf{x}_o)$ is the appropriate set of environmental input data, either including or not including lateral flux perturbations. Error statistics for the informed simulator results, with and without perturbations, are given for each site in Fig. 9.

The use of lateral flux perturbations leads to strong reductions in bias and r.m.s. error at most of the low and mid-latitude sites to 40° N, and at 50° N from the summer of 1998 onwards. The improvement is particularly notable at 10° N, 25° N, 35° N and 40° N, where the addition of these perturbations correct a long-term drift very effectively, albeit with slight over-correction of the positive bias at 10° N. Performance is a little more equivocal at 20° N where perturbation of the simulation leads to a relatively large over-correction of a negative bias but the overall r.m.s. error is still reduced.

The perturbed simulator does not perform better everywhere. The main exception is seen at 60° N, where the simulator shows a tendency to over estimate chlorophyll in the summer of 1998. Another exception is an over correction of the positive bias at 50° N in 1997 which leads to a bias of larger magnitude over some parts of the year. These detrimental effects are minor compared with the overall improvement achieved.

It is clear from Fig. 9 that omitting lateral flux perturbations altogether can lead to particularly large biases associated with serious drifts. Biases of magnitude 0.6 log₁₀ units, representing a factor 4 error in surface chlorophyll, are not uncommon. Examination of the uninformed simulator results in Fig. 5 before any bias correction shows that even at the sites where the error is relatively large, the biases are not. The largest biases are of magnitude 0.3 log₁₀ units,

equivalent to a factor of 2. This indicates that a scheme based on average flux perturbations for the parameter space (i.e. the mean environment) can reduce the problem of drift to a large extent, even though the environment information is not parameter specific.

5 Discussion

- 5 In this section, the performance of the experimental mechanistic emulator is first examined and scope for its improvement identified. Practical application of the site-based emulation scheme is then considered and its envisaged role in enabling advances in the parametric analysis and calibration of global biogeochemical models is discussed.

5.1 Mechanistic emulator performance

- 10 Two alternative versions of a mechanistic emulator for surface chlorophyll from global NEMO-MEDUSA simulations have been evaluated. Each of these site-based emulators uses the same set of site-specific 1-D simulators. The two emulators differ in the method they employ to quantify uncertainty in the simulator predictions.

- 15 The site-based emulator with direct uncertainty quantification is able to predict the 1998 chlorophyll record for a given parameter vector to an accuracy broadly consistent with its uncertainty prediction at all sites. It should therefore serve as a reasonably reliable emulator of the target model for parametric analyses. There is a slight tendency to under-estimate the uncertainty, which is likely to be a consequence of the small target model ensemble size used to represent the known truth (8 or 9). This interpretation would be consistent with a parametric uncertainty analysis of a regional 3-D biogeochemical model by Fiechter (2012), spanning
20 a similar parameter space, in which an ensemble size of 10 was found to give significantly low estimates for ensemble spread compared with 25, 50 and 100 member ensembles. In a practical application, the tendency towards over-confidence could be compensated for by a small inflation factor applied to the residual variance estimate. The optimal factor would be the normalized er-

ror variance from the evaluation experiments (i.e. 1.28, based on the standard deviation of 1.13 for the 9 trial parameter vector experiments in Table 3).

The emulator with indirect uncertainty quantification is able to predict the 1998 record to an accuracy consistent with its uncertainty prediction at about half of the experimental sites, so clearly has some potential. However, it shows a tendency to be over confident in its predictions at other sites, particularly at the more oligotrophic sites studied. Its performance therefore requires some improvement before it can be considered generally robust over a wide range of oceanic conditions.

The most notable instances of poor emulator performance occur for parameter vectors associated with the more extreme behaviour in the target model. This raises the question of whether it is really necessary to emulate the target model over such large parameter ranges. Certainly, restricting the parameter space further should help to make our reference sample more representative. In principle, comparisons with observational data at an early stage could be used to identify implausible target model behaviour and suggest ways in which the parameter space might be reduced. However, any such constraints based on the sparse sampling of parameter space achieved by the target model ensemble could greatly increase the risk of excluding promising parameter combinations and should be undertaken with care. Modifications of the parameter space that are consistent with our biological understanding of the parameters are the most easily justified but, acknowledging the high level of abstraction involved in modelling a system of such complexity, we should avoid over-reliance on subjective priors. Increasing the sample size or improving the emulation methods may be preferable.

While the indirect uncertainty quantification method is currently less robust than the direct method, it has the advantage of being less reliant on the small target model ensemble. Simulator uncertainty due to basic simulation error and parametric environment error are quantified separately, the latter being the uncertainty due to substitution of the true parameter-specific environmental input by a mean environment. The quantification method for basic simulation uncertainty relies wholly on the target model ensemble. However, that for parametric environment uncertainty relies on it only for providing environmental data for a 1-D uncertainty analysis. The output uncertainty depends on the way in which these input data interact with the

parameter-specific dynamics in the 1-D simulators and the 1-D ensemble size can be relatively large. As found by Hemmings and Challenor (2012), the output standard deviation can be highly dependent on the trial parameter vector (see Appendix C). This parameter dependency cannot be accounted for by the direct method. For this reason, a refined version of the indirect method could prove to be more robust than the direct method, particularly if basic simulation errors can be reduced so that the uncertainty quantification for this error component becomes less critical.

The presence of very large normalized error values early in 1997 when the indirect uncertainty quantification method is used suggests that the environment model for the initial conditions should be improved, perhaps through the use of different variance-stabilizing transformations in the EOF analysis used to characterize the environmental uncertainty. Tracer-specific transformations should be considered in place of the square root transformations applied to all primary tracers. Another refinement that may improve performance in the post-initialization phase would be to include covariances between the initial state and the advective flux divergences of the transformed tracer concentrations, instead of modelling the two separately. The persistence of biases at some sites over the whole simulation period, in particular those associated with poor emulator performance, suggests that such improvements could improve robustness of the emulation of the 1998 chlorophyll records.

A fairly simple way of improving the simulator itself would be to provide physical forcing based on 3-D model output at higher temporal resolution for the experimental sites, as the impacts of important weather events are attenuated in the 5 day mean output. Improvements in the representation of concentration dependency in the simulator's lateral flux divergence tendencies are also likely to be beneficial.

Concentration dependency in the 1-D simulations is controlled by the transformation applied to the tracer concentrations. A promising approach to improving its representation might be to introduce tracer-specific transformations, possibly varying in space and time, based on statistical analyses of 3-D model output. A key consideration will be the need to reduce the potential for positive feedback cases, where concentration errors reinforce error in the advective tendencies. This type of positive feedback can cause the growth of large positive errors, particularly in the dissolved nutrient tracers. It may also lead to excessive nutrient depletion

rates where an initial tendency towards negative bias in nutrient concentrations is increased by reduction in lateral supply. Such errors are likely to have a greater impact on surface chlorophyll at oligotrophic sites, where the phytoplankton dynamics are more sensitive to nutrient concentration. It is at these sites where the emulator with indirect uncertainty quantification appears least robust. However, an investigation of the surface nutrient records output by the simulator (not presented) did not show evidence of severe nutrient depletion that might be expected from positive feedback.

5.2 Application of the emulation scheme

For calibration of global ocean biogeochemical models against ocean colour data, the spatial extent of the simulator array can readily be extended to produce a mechanistic emulator with truly global coverage based on a larger set of representative sites. Similarly, the emulation procedure could be extended to records of the annual cycle from multiple years. Importantly, we expect the method to be applicable to models of much higher resolution than the 1° target model used in the present demonstration, with minimal adaptation. The requirement for a small ensemble of 3-D reference simulations is relatively modest, making useful parametric analyses feasible for eddy-permitting and eddy-resolving global models.

While the emulation scheme has the potential to make considerable reductions in the number of 3-D simulations required in a parametric analysis, it must be recognized that even a single 3-D simulation may be a large overhead if long spin-up periods are required. The 2 year spin-up period employed for producing the reference ensemble in our experiments is sufficient to demonstrate proof-of-concept. However, biogeochemical models and carbon cycle models in particular require long spin-up times, typically thousands of years, to reach equilibrium. This implies that in many practical applications much longer spin-up times would be needed. Fortunately, recent advances in the estimation of steady state annual cycles for global models (Khatiwala, 2007, 2008) promise to alleviate this problem. The efficient Transport Matrix Method of Khatiwala (2007) has recently been exploited in parametric analyses where simulations are evaluated against global nutrient data (Kriest et al., 2010, 2012). It has also been combined with a surrogate-based optimization tech-

Discussion Paper | Discussion Paper | Discussion Paper

nique for practical parameter estimation (Prieß et al., 2013b). Incorporating these new steady state estimation techniques into the target model prior to site-based emulation would be particularly advantageous.

Continued development of the indirect uncertainty quantification method is motivated by its potential in situations where a known truth is unavailable. Such a situation arises if we want to emulate a target model for which we have no model-specific 3-D ensemble but must rely on results for a related model. For example, we might try to emulate a high-resolution model, for which we have perhaps just one simulation, by adapting the method to make use of biogeochemical information from lower resolution ensembles. In this scenario, the statistical environment model could be constructed using the high resolution flow field in combination with upstream gradient and initial state information from the low resolution model. Additional uncertainty in the gradient and state information associated with the change of resolution would be quantified with reference to the equivalent high resolution model fields. The effect of basic simulation errors would, of course, have to be quantified with reference to the single high resolution simulation but this is less likely to be a problem if the basic simulation errors can be made small compared with the parametric environment error.

In applying the emulator with indirect uncertainty quantification to each trial parameter vector, the requirement for a parameter-specific set of 1-D ensemble simulations in the environmental uncertainty analysis imposes a large overhead. The significance of this overhead depends on the experimental set up. For a 1° target model emulated by a global array of simulators at 10° intervals, the computational savings in replacing the 3-D simulation by the emulator array would be fairly limited if an ensemble size of 100 were used as in the present study (being largely those due to the reduced vertical domain and use of pre-calculated physical fields). However, for a 0.25° model with the same array, savings would be considerable. Moreover, it seems likely that the ensemble size could be reduced and investigation of the sensitivity of performance measures to ensemble size would certainly be worthwhile.

In a practical calibration exercise where the uncertainty statistics are required for weighting model-data misfit to account for simulation uncertainty, we should not ignore temporal covariance in simulation error. Although the covariance structure of the error has not been quantified

in this study, the results are indicative of strong temporal correlation over long time scales at some sites. This suggests that it will be important to extend the chosen uncertainty quantification procedure to predict the temporal error covariances for each site-specific simulator. Correlation between sites may also need to be considered, particularly if sites are relatively close together.

5 Although the emphasis of the present study has been on emulating surface chlorophyll, the method can in principle be used to emulate other observable variables associated with the target model. A full set of model outputs are available from the 1-D simulations at each site and simulation uncertainty measures can similarly be predicted for any of these variables, although the robustness of such predictions is as yet untested. Use of in situ observations in conjunction
10 with the satellite ocean colour data will provide valuable additional constraints on parameter values, making this an important extension to the mechanistic emulator capability.

5.3 The role of a site-based mechanistic emulator

Thorough investigation of the large multi-dimensional parameter spaces associated with mechanistic biogeochemistry models like MEDUSA will inevitably place great demands on our computer resources. For most parametric analyses, it is envisaged that the mechanistic emulator
15 would be used in combination with one or more statistical emulators for which it would provide the training data and associated uncertainty estimates. This would facilitate the use of rigorous Bayesian analysis techniques which would otherwise not be computationally feasible. Introducing mechanistic emulation as an intermediate step should greatly decrease the number of
20 expensive 3-D simulations that are needed.

Modern Bayesian calibration methods, following Kennedy and O'Hagan (2001), provide a comprehensive statistical framework for addressing issues of parametric uncertainty as well as uncertainty from other sources. They allow estimation of joint posterior distributions for model parameters and model discrepancy. Model discrepancy, originally referred to as model
25 inadequacy, quantifies error associated with the model design that cannot be corrected by parameter adjustment. Arhonditsis et al. (2008) and Zhang and Arhonditsis (2009) demonstrate the application of Bayesian calibration methods to aquatic biogeochemical modelling in a 1-D framework, indicating the value of these methods for quantifying uncertainty associated with

model predictions. A capability for routine application of these methods to biogeochemistry at the global scale would contribute to more robust probabilistic predictions of global change.

A flexible alternative to full Bayesian calibration is the well-established history matching approach adopted by Williamson et al. (2013) in their coupled ocean–atmosphere model analysis.

5 This relatively simple technique uses perturbed parameter ensembles in combination with an implausibility metric to rule out regions of parameter space. The implausibility function takes into account the relevant uncertainties and can be applied iteratively, introducing additional observational data at each stage, to rule out successive regions. The initial focus can be on simple model outputs that are easy to model statistically over the whole parameter space. Subsequent
10 re-focussing of computational effort on smaller regions of parameter space can then be used to develop statistics for more complex outputs.

In this way, history matching can be used as a precursor to Bayesian calibration or, if the region of parameter space not ruled out by the history matching process is sufficiently small, further calibration may be omitted in favour of an averaged parameter vector. The emphasis
15 on defining a “not-ruled-out-yet” region of parameter space, rather than finding the optimal parameter vector, is well-suited to ecosystem modelling where the “underdetermination problem” highlighted by Ward et al. (2010) is ubiquitous.

It is important to recognize that the site-based experimental framework is designed to investigate relatively short time-scale responses of the biogeochemistry to physical drivers. The efficiency of the method makes the corresponding output relatively easy to model statistically
20 and so is well suited to the early stages of history matching. However, we cannot rule out the possibility of interactions with the ocean circulation that would compromise performance of particular parameter vectors in much longer simulations. Further tests would be needed in 3-D simulations to fully determine suitability.

25 In designing a calibration strategy for ocean biogeochemical models, we can take advantage of the relatively weak coupling between the upper ocean and the interior and the different time-scales associated with upper ocean processes and the sinking and remineralization of material in the deep ocean. Site-based methods are best suited to the optimization of parameters associated with seasonal productivity cycles in the upper ocean, occurring on short time scales compared

with those for the redistribution of plankton by the large scale circulation. Parameters associated primarily with slow deep water processes that interact more strongly with the circulation can be optimized separately in 3-D experiments, without compromising the seasonal dynamics.

There are parallels with an established system used in terrestrial carbon cycle modelling. This is the Carbon Cycle Data Assimilation System (Rayner et al., 2005), which uses a two stage process to calibrate a terrestrial biogeochemistry model. The first step involves optimization of parameters controlling phenology and soil moisture by assimilating satellite data related to vegetation activity. The second step then uses fields from the optimized model as input to a simpler model version, combined with a 3-D atmospheric transport model, for constraining the remaining model parameters to fit atmospheric CO₂ data.

5.4 Site-based process model analysis

As a final point, it should be stressed that we have focused here on enabling parametric analyses for a coupled model system, where the optimal parameter values are conditional on a particular representation of the physical ocean. This is important for applications of biogeochemistry models in specific host model configurations. However, there is also a need to be able to evaluate and improve the fidelity of the biogeochemistry model with respect to the processes it is designed to represent, independently of a particular physical simulation. This is emphasized by parameter optimization experiments of Friedrichs et al. (2006) which show that likely error in the physical forcing data can have a large effect on the biogeochemical simulations, leading to inappropriate posterior parameter values.

Site-based methods can be adapted to allow for such error by including a quantification of uncertainty in the physical environment in the analysis as suggested by Hemmings and Challenor (2012). By doing this, we aim to emulate the output that would be obtained from the biogeochemistry model if it were embedded in a perfect physical simulation. History matching could then be used to rule out areas of parameter space that are inconsistent with a plausible representation of the biogeochemical dynamics. Computing effort would be focused primarily on data-rich sites, including established biogeochemical time series observatories.

A statistical model of the biogeochemical environment would be required for the 1-D simulations at each site. The methods introduced here provide the basis for constructing such a model. However, they would need to be refined to allow for additional uncertainty involved in making inferences about a hypothetical perfect physics ensemble from analysis of a practical 3-D ensemble. The development of a robust method is more likely to be achievable if a good observationally-constrained statistical description of the local flow field can be established. Then, only the upstream tracer gradient and initial state information would need to be inferred from the 3-D model analysis. Furthermore, it should be possible to take initial state information from an observation-based statistical model of the real-world state, say from a climatology. Inferences about the model would then be restricted to its behaviour over relatively short time scales. However, this seems likely to be the most practical approach.

In principle, the site-based capability could be adapted for use in a Lagrangian framework allowing a Eulerian simulator array to be augmented by 1-D simulations following Argo floats or surface drifter trajectories. Physical data from Eulerian observatories and Lagrangian platforms, in combination with satellite Earth observation data could be used in conjunction with 3-D simulations to develop observationally-constrained statistical representations of the physical environment to which the biogeochemistry responds. Bringing these different components of the global observation system together in a robust statistical framework for model calibration and assessment will be an important step in developing a reliable predictive capability for the Earth system that accounts for the role of marine biogeochemistry in global change.

6 Summary and conclusions

A mechanistic site-based emulator for annual cycles of surface chlorophyll output from the global NEMO-MEDUSA model was presented. The emulation scheme introduces two fundamental improvements to our site-based biogeochemical modelling capabilities: an explicit representation of the lateral flux divergences of the model tracers, following Hemmings and Challenor (2012), and a quantification of output uncertainty with respect to the target model.

The emulator relies on an array of 1-D simulators of the target model dynamics. In the absence of parameter-specific 3-D model information about the environment at each site, the simulators use a mean environment provided by a small ensemble of target model simulations. This 3-D ensemble is designed to be representative of variability in the model dynamics over the parameter space of interest. It provides information about the local environment in the form of estimates of the required initial state and lateral flux divergences, together with their uncertainties. The use of lateral flux information reduces simulator error considerably, consistent with a major influence of advection at some sites, and this has been instrumental in achieving a promising level of performance.

Two different versions of the mechanistic emulator have been evaluated. One is constructed using a direct uncertainty quantification method, in which output uncertainty is quantified by comparison with a known truth. The other is constructed using an indirect method, in which output uncertainty is inferred from separate analyses for two contributing factors: the set of basic simulation errors and the parametric environment error. Uncertainty due to basic simulation errors is quantified by applying the direct method to the simulator with a known parameter-specific environment. Parametric environment error is the error in the simulator output when an unknown parameter-specific environment is approximated by the mean environment (an estimate of the expectation of the environment over the parameter space of interest). Uncertainty associated with this error is quantified by 1-D uncertainty analyses.

The analysis for NEMO-MEDUSA indicates that the emulator with direct uncertainty quantification should provide a reasonably robust site-based emulation capability for the surface chlorophyll output from 3-D models. The indirect uncertainty quantification scheme, although more expensive in terms of the number of 1-D simulations required, has the advantage of accounting for the dependency of simulation uncertainty on the trial parameter vector. However, as implemented here, it was found to be less robust. Nevertheless, a number of improvements to the method have been suggested which are expected to improve its reliability. Irrespective of whether this leads to the performance of the indirect method exceeding that of the direct method in terms of robustness, the indirect method provides the basis for a more flexible approach that is less reliant on target model simulations. The potential of both versions of the

emulation scheme to improve the effectiveness of site-based approaches to parametric analysis of ocean biogeochemical models is clear.

Our experimental mechanistic emulator serves as a prototype for an improved site-based capability. This facility would allow robust inferences to be made about the parameter-dependent behaviour of global biogeochemical models on the basis of analyses performed on representative arrays of 1-D simulators. It would thus enable the routine execution of relevant parameter perturbation ensembles with 100s of members. In conjunction with statistical emulators, this would enable comprehensive investigations of large parameter spaces to be performed.

In addition, the new developments in the treatment of lateral advection and quantification of environmental uncertainty for 1-D simulators will be important for performing analyses of biogeochemistry models that are based on their representation of the biogeochemical dynamics, rather than being conditional on a particular representation of the physical circulation. This type of process-based analysis is essential for assessing and improving the fidelity of process representation in biogeochemical models.

Site-based analyses of both coupled and stand-alone biogeochemistry models promise to make important contributions to our ability to constrain model parameters and quantify biogeochemical uncertainty in ocean and Earth system model predictions.

Appendix A: Code availability

MarMOT 1.1 is open source software available under the CeCILL Free Software License Agreement. It is designed for use on UNIX-based systems, including LINUX and Mac OS X. The original code was released on 21st November 2013. The current version, MarMOT 1.1.1, released on 23rd January 2015, is functionally equivalent to the original but includes modifications to address a known portability issue and improve reliability. A tar archive containing the MarMOT 1.1.1 distribution can be downloaded from the National Oceanography Centre's web site at <http://noc.ac.uk/project/marmot> or supplied by the corresponding author on request. The software release includes a set of command line tools for handling MarMOT-compatible data tables. Full documentation and test data are included with the distribution.

The MEDUSA 1.0 code is available as a supplement to Yool et al. (2011). A version of this original code with adaptations for interfacing with the MarMOT testbed is included in the MarMOT 1.1.1 distribution.

Appendix B: Defining the parameter space

5 The first step in parametric analysis of a model, whether for purposes of uncertainty analysis or calibration, is defining the parameter space to be investigated. Our primary interest here is in exploring uncertainty in the seasonal cycle and its impact on annual primary production and the export of material from the euphotic zone. We therefore want to investigate plankton system parameters that have a significant influence on these processes. These are identified by a formal
10 sensitivity analysis involving 28 relevant model parameters varied over ranges consistent with their defined roles in the model.

B1 Initial parameter selection

The MEDUSA 1.0 model as described by Yool et al. (2011) has over 60 parameters. Our focus is on the seasonal cycle in the euphotic zone with the ultimate aim of using satellite-derived
15 chlorophyll data to constrain upper ocean plankton dynamics in the model. On this basis, a number of parameter groups are excluded from the model analysis. These are the parameters of the inorganic iron and carbonate systems and parameters associated with the remineralization of sinking particles that occurs mainly in the ocean interior. Parameters related to stoichiometry are, in general relatively well known compared with many of the other parameters and are
20 also excluded from the analysis. However, this is largely a pragmatic decision to reduce the size of the parameter space; sensitivity to these parameters within their expected ranges should ideally be explored in future studies. The parameters referred to are the carbon : nitrogen and iron : nitrogen ratios for the organic components and the parameters controlling the variable chlorophyll : carbon ratios for the two phytoplankton types and the diatom silicon : nitrogen ratios.
25

The remaining set of parameters used in MEDUSA includes parameters that are conceptually related in such a way as to complicate the interpretation of parametric analyses in which they are varied independently. For example, the two phytoplankton types each have their own set of rate parameters, so adjusting a rate parameter for one phytoplankton type affects the relative rates for each type. There are no individual parameters controlling the overall rates associated with phytoplankton as an aggregated biotic group. To avoid problems of this kind, the input parameter set in the MarMOT 1.1 configuration of MEDUSA has been modified from the parameter set used internally.

The 37 input parameters relevant to this study and their relationships to the internal parameters specified in Yool et al. (2011) are shown in Tables 4 to 6. The standard values tabulated are those used in the standard simulation of Yool et al. (2011) or their equivalents. The standard simulation is referred to in the National Oceanography Centre's archive as EXP276 (available on request from A. Yool; axy@noc.ac.uk). There are inconsistencies between values for 3 of the zooplankton density-dependent loss parameters in Table 6 ($f_{\mu 2, Z\mu}$, $f_{kZ\mu}$ and $f_{\mu 2, Zm}$) and values appearing in Yool et al. (2011) since the latter were incorrect. The correct standard simulation values for the microzooplankton maximum loss rate and half saturation concentration are $\mu_{2, Z\mu} = 0.1$ and $k_{Z\mu} = 0.5$ respectively (in units of d^{-1} and mmol N m^{-3}). These match the corresponding standard simulation values for phytoplankton. The correct value for the mesozooplankton maximum loss rate $\mu_{2, Zm}$ is 0.2 d^{-1} .

Pairs of rate or half-saturation concentration parameters for the different phytoplankton or zooplankton types have been replaced by a base value, pertaining to the smaller plankton type (non-diatoms or microzooplankton), and a relative value for the larger type (diatoms or mesozooplankton). This leads to new parameters that are non-dimensional factors. For the diatom growth process these are $f_{\alpha Pd}$, f_{VPd} , $f_{kN, Pd}$ and $f_{kFe, Pd}$. For mesozooplankton growth we have f_{gm} and f_{km} . The new parameters for the diatom loss processes are $f_{\mu 1, Pd}$, $f_{\mu 2, Pd}$, f_{kPd} . For the zooplankton loss processes, the microzooplankton values $f_{\mu 2, Z\mu}$ and $f_{kZ\mu}$ are defined in terms of the non-diatom phytoplankton values and the mesozooplankton values $f_{\mu 2, Zm}$, f_{kZm} are defined in terms of the microzooplankton values. This suite of modifications allow individual parameters, the base values, to be varied without affecting the relationships between closely

associated parameters. The parameter relationships can be controlled independently using the new parameters.

A similar approach is taken for assimilation efficiencies and feeding preference parameters. The carbon assimilation efficiency for zooplankton grazers has been re-expressed in terms of their nitrogen assimilation efficiency by a non-dimensional offset parameter $a_{\beta C}$. The value is the fraction of the maximum possible offset determined by the constraint that assimilation efficiencies must logically be within the range 0–1. Mesozooplankton feeding preferences have been re-expressed in a hierarchical way so that instead of preference factors for each individual food type, there is an overall preference for live food (as opposed to detritus) p_{mLive} and two conditional preferences: a preference for phytoplankton given live food $p_{c, mP}$ and a preference for non-diatoms given phytoplankton $p_{c, mPn}$.

Yool et al. (2011) used identical values for some parameter pairs and groups to avoid introducing arbitrary complexity. The new definition of the input parameter set described here allows the values of associated internal parameters to be kept the same while varying their values via the base parameter. Adding additional complexity over that of the original model is not justified for the present calibration experiments so the relevant non-dimensional factors are fixed at 1 wherever identical parameter values were used by Yool et al. (2011), thereby further reducing dimensionality of the parameter space. By the same argument, $a_{\beta C}$ is fixed at 0.

The standard value for the fast detritus fraction of mesozooplankton losses $D2_{frac}$ is 1, implying that all mesozooplankton losses are treated as fast-sinking detritus. Adjusting this value would cause the losses to be divided between slow and fast sinking detritus adding a small amount of additional complexity to the model processes. Again, we chose to avoid introducing this new complexity and left this parameter fixed.

As a consequence of excluding less relevant parameter groups from the analysis and choosing to avoid the introduction new complexity, an initial parameter space of 28 dimensions was considered in the present study. The remaining parameters are constrained a priori to take their standard values; this constraint effectively becomes part of the model design. Further dimension reduction was performed objectively on the basis of a sensitivity analysis.

B2 Parameter ranges

Acceptable ranges for each of the parameters to be included in the analysis are defined according to a set of rules as follows.

5 Rule 1: for all positive parameters with no inherent upper limit, bounds are symmetric about the prior value on a geometric scale. This applies to rate parameters and half-saturation concentrations, whether expressed in absolute or relative units. Rate parameter bounds are set initially at half and double the prior. A factor of 5 is used for half-saturation concentrations.

10 Rule 2: for fractions, such as efficiencies and feeding preferences, limits are initially set at ± 0.25 . Limits of 0.05 and 0.95 are imposed on the lower and upper bounds respectively and the bounds are adjusted if necessary.

15 Rule 3: the sign of differences between associated internal parameters is preserved. This is done for rates and half-saturation concentrations by imposing 1 as a lower or upper limit for the ranges of the parameters that are expressed as relative values, depending on whether their priors are greater than or less than 1. The relevant bound is adjusted if necessary.

20 Rule 4: if one or other bound is adjusted in applying Rule 3, then symmetry is used to reset the opposite bound. Geometric symmetry is applied to rates and half-saturation concentrations. This rule applies a constraint on the difference between associated parameters that is dependent on their difference in the prior parameter set.

The resulting parameter space is defined by Table 7. Log-transformed values are used for some parameters when dividing up the parameter space for sampling purposes. The dimensions to which this applies are indicated in the table.

B3 Parameter sensitivity analysis

25 Following the initial parameter selection, further reduction in the dimensionality of the parameter space to be explored in the calibration process is based on the potential impact of parameters on annual primary production and the ratio of annual particulate export to annual primary production, referred to as the pe-ratio. (The inorganic fraction of particulate carbon export asso-

ciated with carbonate production is excluded.) The value of the pe-ratio at 207 m is used since this is the greatest depth at which photosynthesis can occur in the model.

Annual mean values for 1998 at 12 sites were determined for 5000 different parameter vectors in the 28 dimensional parameter space. The parameter vectors were distributed in parameter space using a Latin hypercube design (McKay et al., 1979) with a “maximin” criterion (Johnson et al., 1990) applied to 10 randomly generated hypercubes. For generating the design points, distance is defined in terms of positions on a parameter space grid with an equal number of intervals in each dimension. Grid intervals are in log units for rate parameters and half-saturation concentrations. The sensitivity analysis was performed using the 1-D experimental framework described in Sect. 3, with the time step increased to 2 h for efficiency. 1-D simulations were initialized from the standard 3-D simulation of Yool et al. (2011) at the start of 1997, allowing one complete annual cycle for adjustment to the new parameter values and the 1-D context to reduce the impact of transient behaviour. Lateral fluxes were ignored.

The results of an initial sensitivity analysis for all 28 parameters were examined to identify parameters that have a clear impact on the primary production and the pe-ratio. Parameters that individually explained less than 5 % of the variance in both variables at all sites were then automatically excluded. The sensitivities of the two variables to the remaining parameters are summarized in Table 8 in terms of the number of sites out of 12 at which the parameter explains at least 5 % of the variance and the proportion of variance explained given by the squared Pearson correlation coefficient r^2 .

There are 9 parameters that explain more than 5 % of the variance in both model outputs. Of these, k_C has a relatively weak effect on both and is excluded. Of the remaining 3 parameters, V_{Pn} is the only one with any stronger influence than k_C on either output, having some impact on primary production. However, its effect does not appear to be any greater than the least influential of the other parameters to be retained. Given its lack of influence on pe-ratio, it is discarded along with f_{km} and $\mu_{2,Pn}$ leaving an 8-dimensional parameter space for the emulation experiments.

The sensitivity analysis was repeated in the 8 dimensional parameter space, again with a sample size of 5000 parameter vectors. Discarding the other 20 parameters reduced the total vari-

ance in primary production at each site by between 5 and 38 %. The reduction in the pe-ratio variance was generally less, varying from 6 to 19 %. The parametric uncertainty in primary production and pe-ratio associated with the final 8-dimensional parameter space is illustrated by the coefficient of variation (ratio of standard deviation to mean) for the two variables at each site. The coefficient of variation for primary production ranges from 0.29 (at 15° N) to 0.48 (at 55° N). That for the pe-ratio is generally greater, ranging from 0.38 (at 60° N) to 1.06 (at 30° N).

Appendix C: Quantification of simulator uncertainty

Uncertainty for the log-transformed 5 day mean chlorophyll output is quantified in terms of time series of the predicted monthly means and variances of the uninformed simulator residual. In the direct uncertainty quantification method, these statistics are derived from differences between the 5 day uninformed simulator output and the corresponding target model output over all parameter vectors in the Construction Phase ensemble. In the indirect method, they are derived from the sums of the mean and variance estimates for the mean environment simulation residual ϵ_S and the parametric environment residual ϵ_B . The ϵ_S statistics are estimated from differences between the 5 day informed simulator output and the target model output over the parameter vectors in the Construction Phase ensemble. The ϵ_B statistics are estimated from the 5 day output of a parametric uncertainty analysis using 100 ensemble members.

For each residual, the mean and variance of the 5 day probability distributions are estimated from the relevant ensemble-based sample: $u_i, i \in \{1, \dots, n\}$. The unbiased population variance estimator

$$s_u^2 = \frac{\sum_{i=1}^n (u_i - \bar{u})^2}{n - 1} \quad (\text{C1})$$

is used. The 5 day statistics are then used to derive monthly means and variances which are interpolated to give continuous time series $\bar{u}_m(t)$ and $s_m^2(t)$ respectively for uncertainty quantification. The procedure for calculating the time series from the 5 day statistics is as follows.

5 day samples are grouped in pseudo-monthly bins (intervals of 30.42 days) and the monthly mean residual \bar{u}_m is estimated from the k sample means in each bin using the unweighted

average, so

$$\bar{u}_m = \frac{1}{k} \sum_{i=1}^k \bar{u}_i \quad (C2)$$

where \bar{u}_i is the mean of the i th 5 day sample. \bar{u}_m is then linearly interpolated between monthly mid-points to obtain $\bar{u}_m(t)$. Values for early January 1997 and late December 1998 are equated to those at the respective monthly mid-point. $\bar{u}_m(t)$ is the estimate of the expected residual used for bias correction.

The true residual for the trial parameter vector \mathbf{x}_o can be expressed as

$$\psi_o(t, \mathbf{x}_o) = \bar{u}_m(t) + \epsilon_\mu + \epsilon_\psi \quad (C3)$$

where ϵ_μ is the departure of the true residual mean from the estimated residual mean:

$$\epsilon_\mu = \mu(t) - \bar{u}_m(t) \quad (C4)$$

and ϵ_ψ is the departure of the true residual from the true residual mean:

$$\epsilon_\psi = \psi_o(t, \mathbf{x}_o) - \mu(t). \quad (C5)$$

For the purposes of uncertainty quantification, these departures are assumed to be independent Gaussian random variables with zero means and variances $s_\mu^2(t)$ and $s_\psi^2(t)$ respectively, derived from the sample data. Variances s_μ^2 and s_ψ^2 are determined for each pseudo-monthly bin. The monthly variance estimate for the residual is then given by

$$s_m^2 = s_\mu^2 + s_\psi^2. \quad (C6)$$

This is converted to a continuous time series by interpolation and end-point extrapolation, as for the residual means, to obtain $s_m^2(t)$.

For each bin, s_μ^2 is given by the monthly variance of the anomaly between the 5 day sample mean \bar{u} and the expected residual estimate \bar{u}_m at the 5 day interval mid-point. So

$$s_\mu^2 = \frac{\sum_{i=1}^k (a_i - \bar{a})^2}{k - 1} \quad (C7)$$

where

$$a_i = \bar{u}_i - \bar{u}_m(t_i). \quad (C8)$$

s_ψ^2 is given by the pooled estimates of the residual variance

$$s_\psi^2 = \frac{1}{k} \sum_{i=1}^k s_{u,i}^2 \quad (C9)$$

where $s_{u,i}^2$ is the variance estimated from the i th 5 day sample.

Determination of monthly means and variances for the residuals from the 5 day samples is expected to give more robust estimates. However, the increase in effective sample size depends on the extent to which samples are temporally correlated over each pseudo-monthly bin. This is not quantified in the present study.

Time series of uninformed simulator residual statistics given by the direct and indirect uncertainty quantification methods are shown in Fig. 10. (Note that for an arbitrary residual ϵ_X , \bar{u}_m is denoted \bar{u}_X and s_m is denoted s_X .) For both methods, the time series determined for all 10 trial parameter experiments are shown. The statistics for the uninformed simulator residual ϵ_1 predicted by the direct method do not account for dependency of the true residual distributions on the trial parameter vectors. Thus, variation in the time series between experiments is due only to sampling uncertainty. The ϵ_1 statistics predicted by the indirect method do account for this parameter dependency and the variation between experiments is then in part due to the parameter-specific dynamics of the environment ensemble simulation used for the parametric uncertainty analysis.

Time series for the statistics of the component residuals contributing to the uninformed simulator statistics given by the indirect method are shown in Fig. 11. The statistics for the mean environment simulation residual ϵ_S , like the ϵ_1 statistics given by the direct method, differ between experiments only due to sampling uncertainty. They exhibit less variation between experiments than the ϵ_1 statistics, reflecting the lack of dependency of the true distribution of ϵ_S on the trial parameter vector. The statistics for the parametric environment residual ϵ_B , the component residual that explicitly accounts for the trial parameter vector dependency in the uninformed simulator uncertainty, show much greater variation between experiments.

**The Supplement related to this article is available online at
doi:10.5194/gmdd-0-1-2015-supplement.**

Acknowledgements. The authors would like to thank Andrew Coward and the NEMO development team at the National Oceanography Centre, Southampton for their technical support of the global NEMO-MEDUSA modelling activities. Thanks are also due to Markus Schartau and one anonymous referee for their helpful and constructive comments on the original manuscript. The financial support of the Natural Environment Research Council (NERC) is gratefully acknowledged. Support for the present study was provided via the National Centre for Earth Observation's Carbon Cycle research theme (grant reference: EARTH010003). NCEO is a NERC research centre.

References

- Arhonditsis, G. B., Papantou, D., Zhang, W., Perhar, G., Massos, E., and Shi, M.: Bayesian calibration of mechanistic aquatic biogeochemical models and benefits for environmental management, *J. Marine Syst.*, 73, 8–30, 2008.
- Aumont, O. and Bopp, L.: Globalizing results from ocean in situ iron fertilization studies, *Global Biogeochem. Cy.*, 20, GB2017, doi:10.1029/2005GB002591, 2006.
- Campbell, J. W.: The lognormal distribution as a model for bio-optical variability in the sea, *J. Geophys. Res.*, 100, 13237–13254, 1995.

- Doron, M., Brasseur, P., Brankart, J.-M., Losa, S. N., and Melet, A.: Stochastic estimation of biogeochemical parameters from Globcolour ocean colour satellite data in a North Atlantic 3-D ocean coupled physical-biogeochemical model, *J. Marine Syst.*, 117–118, 81–95, 2013.
- 5 Dowd, M.: Estimating parameters for a stochastic dynamic marine ecological system, *Environmetrics*, 22, 501–515, doi:10.1002/env.1083, 2011.
- Fan, W. and Lv, X.: Data assimilation in a simple marine ecosystem model based on spatial biological parameterizations, *Ecol. Model.*, 220, 1997–2008, doi:10.1016/j.ecolmodel.2009.04.050, 2009.
- Fasham, M. J. R. and Evans, G. T.: The use of optimization techniques to model marine ecosystem dynamics at the JGOFS station at 47° N 20° W, *Philos. T. Roy. Soc. B*, 348, 203–209, 1995.
- 10 Fasham, M. J. R., Boyd, P. W., and Savidge, G.: Modeling the relative contributions of autotrophs and heterotrophs to carbon flow at a Lagrangian JGOFS station in the Northeast Atlantic: the importance of DOC, *Limnol. Oceanogr.*, 44, 80–94, 1999.
- Fasham, M. J. R., Flynn, K. J., Pondaven, P., Anderson, T. R., and Boyd, P. W.: Development of a robust marine ecosystem model to predict the role of iron in biogeochemical cycles: a comparison of results for iron-replete and iron-limited areas, and the SOIREE iron-enrichment experiment, *Deep-Sea Res. Pt. I*, 53, 333–366, 2006.
- 15 Fiechter, J.: Assessing marine ecosystem model properties from ensemble calculations, *Ecol. Model.*, 242, 164–179, doi:10.1016/j.ecolmodel.2012.05.016, 2012.
- Fiechter, J., Herbei, R., Leeds, W., Brown, J., Milliff, R., Wikle, C., Moore, A., and Powell, T.: A Bayesian parameter estimation method applied to a marine ecosystem model for the coastal Gulf of Alaska, *Ecol. Model.*, 258, 122–133, 2013.
- 20 Friedrichs, M. A. M., Hood, R. R., and Wiggert, J. D.: Ecosystem model complexity versus physical forcing: quantification of their relative impact with assimilated Arabian Sea data, *Deep-Sea Res. Pt. II*, 53, 576–600, 2006.
- 25 Friedrichs, M. A. M., Dusenberry, J. A., Anderson, L. A., Armstrong, R. A., Chai, F., Christian, J. R., Doney, S. C., Dunne, J., Fujii, M., Hood, R., McGillicuddy Jr., D. J., Moore, K., Schartau, M., Spitz, Y., and Wiggert, J. D.: Assessment of skill and portability in regional marine biogeochemical models: role of multiple planktonic groups, *J. Geophys. Res.*, 112, C08001, doi:10.1029/2006JC003852, 2007.
- 30 Garcia-Gorriz, E., Hoepffner, N., and Ouberdous, M.: Assimilation of SeaWiFS data in a coupled physical-biological model of the Adriatic Sea, *J. Marine Syst.*, 40–41, 233–252, 2003.
- Gregg, W., Ginoux, W. P., Schopf, P. S., and Casey, N. W.: Phytoplankton and iron: validation of a global three-dimensional ocean biogeochemical model, *Deep-Sea Res. Pt. II*, 50, 3143–3169, 2003.

- Hemmings, J. C. P. and Challenor, P. G.: Addressing the impact of environmental uncertainty in plankton model calibration with a dedicated software system: the Marine Model Optimization Testbed (Mar-MOT 1.1 alpha), *Geosci. Model Dev.*, 5, 471–498, doi:10.5194/gmd-5-471-2012, 2012.
- 5 Hemmings, J. C. P., Srokosz, M. A., Challenor, P., and Fasham, M. J. R.: Split-domain calibration of an ecosystem model using satellite ocean colour data, *J. Marine Syst.*, 50, 141–179, 2004.
- Hooten, M. B., Leeds, W. B., Fiechter, J., and Wikle, C. K.: Assessing first-order emulator inference for physical parameters in nonlinear mechanistic models, *J. Agric. Biol. Envir. S.*, 16, 475–494, doi:10.1007/s13253-011-0073-7, 2011.
- 10 Hourdin, F. and Armengaud, A.: The use of finite-volume methods for atmospheric advection of trace species. Part I: Test of various formulations in a general circulation model, *Mon. Weather Rev.*, 127, 822–837, 1999.
- Huret, M., Gohin, F., Delmas, D., Lunven, M., and Garçon, V.: Use of SeaWiFS data for light availability and parameter estimation of a phytoplankton production model of the Bay of Biscay, *J. Marine Syst.*, 65, 509–531, 2007.
- 15 Hurtt, G. C. and Armstrong, R. A.: A pelagic ecosystem model calibrated with BATS and OWSI data, *Deep-Sea Res. Pt. I*, 46, 27–61, 1999.
- Johnson, M., Moore, L., and Ylvisaker, D.: Minimax and maxmin distance designs, *J. Stat. Plan. Infer.*, 26, 131–148, 1990.
- Kane, A., Moulin, C., Thiria, S., Bopp, L., Berrada, M., Tagliabue, A., Crépon, M., Aumont, O., and Badran, F.: Improving the parameters of a global ocean biogeochemical model via variational assimilation of in situ data at five time series stations, *J. Geophys. Res.*, 116, C06011, doi:10.1029/2009JC006005, 2011.
- 20 Kennedy, M. C. and O’Hagan, A.: Bayesian calibration of computer models, *J. Roy. Stat. Soc. B*, 63, 425–464, 2001.
- Khatiwala, S.: A computational framework for simulation of biogeochemical tracers in the ocean, *Global Biogeochem. Cy.*, 21, GB3001, doi:10.1029/2007GB002923, 2007.
- Khatiwala, S.: Fast spin up of Ocean biogeochemical models using matrix-free Newton–Krylov, *Ocean Model.*, 23, 121–129, doi:10.1016/j.ocemod.2008.05.002, 2008.
- 25 Kidston, M., Matear, R., and Baird, M. E.: Parameter optimisation of a marine ecosystem model at two contrasting stations in the Sub-Antarctic Zone, *Deep-Sea Res. Pt. II*, 58, 2301–2315, doi:10.1016/j.dsr2.2011.05.018, 2011.
- 30

- Kriest, I., Khatiwala, S., and Oschlies, A.: Towards an assessment of simple global marine biogeochemical models of different complexity, *Prog. Oceanogr.*, 86, 337–360, doi:10.1016/j.pocean.2010.05.002, 2010.
- Kriest, I., Oschlies, A., and S. Khatiwala, S.: Sensitivity analysis of simple global marine biogeochemical models, *Global Biogeochem. Cy.*, 26, GB2029, doi:10.1029/2011GB004072, 2012.
- Lee, L. A., Carslaw, K. S., Pringle, K. J., and Mann, G. W.: Mapping the uncertainty in global CCN using emulation, *Atmos. Chem. Phys.*, 12, 9739–9751, doi:10.5194/acp-12-9739-2012, 2012.
- Leeds, W. B., Wikle, C. K., Fiechter, J., Brown, J., and Milliff, R. F.: Modeling 3-D spatio-temporal biogeochemical processes with a forest of 1-D statistical emulators, *Environmetrics*, 24, 1–12, doi:10.1002/env.2187, 2013.
- Le Quéré, C., Harrison, S. P., Prentice, I. C., Buitenhuis, E. T., Aumont, O., Bopp, L., Claustre, H., Cotrim da Cunha, L., Geider, R., Giraud, X., Klaas, C., Kohfeld, K. E., Legendre, L., Manizza, M., Platt, T., Rivkin, R. B., Sathyendranath, S., Uitz, J., Watson, A. J., and Wolf-Gladrow, D.: Ecosystem dynamics based on plankton functional types for global ocean biogeochemistry models, *Glob. Change Biol.*, 11, 2016–2040, doi:10.1111/j.1365-2486.2005.1004.x, 2005.
- Lévy, M.: The modulation of biological production by oceanic mesoscale turbulence, *Lect. Notes Phys.*, 744, 219–261, doi:10.1007/978-3-540-75215-8_9, 2008.
- Lévy, M., Estublier, A., and Madec, G.: Choice of an advection scheme for biogeochemical models, *Geophys. Res. Lett.*, 28, 3725–3728, 2001.
- Losa, S. N., Kivman, G. A., and Ryabchenko, V. A.: Weak constraint parameter estimation for a simple ocean ecosystem model: what can we learn about the model and data?, *J. Marine Syst.*, 45, 1–20, 2004.
- Madec, G.: NEMO Reference Manual, Ocean Dynamic Component: NEMO-OPA, Note du Pole de modélisation, l’Institut Pierre-Simon Laplace, Paris, France, No. 27., ISSN 1288–1619, 2008.
- Matear, R. J.: Parameter optimization and analysis of ecosystem models using simulated annealing: a case study at Station P, *J. Mar. Res.*, 53, 571–607, 1995.
- Mattern, J. P., Fennel, K., and Dowd, M.: Estimating time-dependent parameters for a biological ocean model using an emulator approach, *J. Marine Syst.*, 96–97, 32–47, 2012.
- McDonald, C. P., Bennington, V., Urban, N. R., and McKinley, G.: A 1-D test-bed calibration of a 3-D Lake Superior biogeochemical model, *Ecol. Model.*, 225, 115–126, 2012.
- McKay, M. D., Conover, W. J., and Beckman, R. J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, 21, 239–245, 1979.

- Moore, J. K., Doney, S. C., and Lindsay, K.: Upper ocean ecosystem dynamics and iron cycling in a global three-dimensional model, *Global Biogeochem. Cy.*, 18, GB4028, doi:10.1029/2004GB002220, 2004.
- 5 Oschlies, A. and Schartau, M.: Basin-scale performance of a locally optimized marine ecosystem model, *J. Mar. Res.*, 63, 335–358, 2005.
- O’Hagan, T.: Bayesian analysis of computer code outputs: a tutorial, *Reliab. Eng. Syst. Safe.*, 91, 1290–1300, 2006.
- Palmer, J. R. and Totterdell, I. J.: Production and export in a global ocean ecosystem model, *Deep-Sea Res. Pt. I*, 48, 1169–1198, 2001.
- 10 Prieß, M., Koziel, S., and Slawig, T.: Marine ecosystem model calibration with real data using enhanced surrogate-base optimization, *J. Comput. Sci.*, 4, 423–437, doi:10.1016/j.jocs.2013.04.001, 2013a.
- Prieß, M., Piwonski, J., Koziel, S., Oschlies, A., and Slawig, T.: Accelerated parameter identification in a 3-D marine biogeochemical model using surrogate-based optimization, *Ocean Model.*, 68, 22–36, doi:10.1016/j.ocemod.2013.04.003, 2013b.
- 15 Rayner, P. J., Scholze, M., Knorr, W., Kaminski, T., Giering, R., and Widmann, H.: Two decades of terrestrial carbon fluxes from a carbon cycle data assimilation system (CCDAS), *Global Biogeochem. Cy.*, 19, GB2026, doi:10.1029/2004GB002254, 2005.
- Sarmiento, J. L., Slater, R. D., Fasham, M. J. R., Ducklow, H. W., Toggweiler, J. R., and Evans, G. T.: A seasonal three-dimensional ecosystem model of nitrogen cycling in the North Atlantic Euphotic Zone, *Global Biogeochem. Cy.*, 7, 417–450, doi:10.1029/93GB00375, 1993.
- 20 Schartau, M. and Oschlies, A.: Simultaneous data-based optimization of a 1-D-ecosystem model at three locations in the North Atlantic: Part I – Method and parameter estimates, *J. Mar. Res.*, 61, 765–793, 2003.
- Séférian, R., Bopp, L., Gehlen, M., Orr, J. C., Ethé, C., Cadule, P., Aumont, O., Salas y Méliá, D., Voldoire, A., and Madec, G.: Skill assessment of three earth system models with common marine biogeochemistry, *Clim. Dynam.*, 40, 2549–2573, doi:10.1007/s00382-012-1362-8, 2013.
- Six, K. D. and Maier-Reimer, E.: Effects of plankton dynamics on seasonal carbon fluxes in an ocean general circulation model, *Global Biogeochem. Cy.*, 10, 559–583, 1996.
- Stow, C. A., Jolliff, J., McGillicuddy Jr., D. J., Doney, S. C., Allen, J. I., Friedrichs, M. A. M., Rose, K. A., and Wallhead, P.: Skill assessment for coupled biological/physical models of marine systems, *J. Marine Syst.*, 76, 4–15, 2009.
- 30

Tjiputra, J. F., Polzin, D., and Winguth, A. M. E.: Assimilation of seasonal chlorophyll and nutrient data into an adjoint three-dimensional ocean carbon cycle model: sensitivity analysis and ecosystem parameter optimization, *Global Biogeochem. Cy.*, 21, GB1001, doi:10.1029/2006GB002745, 2007.

5 Van Leer, B.: Towards the ultimate conservative difference scheme IV: a new approach to numerical convection, *J. Comput. Phys.*, 23, 276–299, 1977.

Ward, B. A., Friedrichs, M. A. M, Anderson, T. R., and Oschlies, A.: Parameter optimisation techniques and the problem of underdetermination in marine biogeochemical models, *J. Marine Syst.*, 81, 34–43, 2010.

10 Ward, B. A, Schartau, M., Oschlies, A., Martin, A. P., Follows, M. J., and Anderson, T. R.: When is a biogeochemical model too complex? Objective model reduction and selection for North Atlantic time-series sites, *Prog. Oceanogr.*, 116, 49–65, 2013.

Williamson, D., Goldstein, M., Allison, L., Blaker, A., Challenor, P., Jackson, L., and Yamazaki, K.: History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble, *Clim. Dynam.*, 41, 1703–1729, doi:10.1007/s00382-013-1896-4, 15 2013.

Xiao, Y. and Friedrichs, M.: The assimilation of satellite-derived data into a one-dimensional lower trophic level marine ecosystem model, *J. Geophys. Res.-Oceans*, 119, 2691–2712, doi:10.1002/2013JC009433, 2014.

20 Yool, A., Popova, E. E., and Anderson, T. R.: Medusa-1.0: a new intermediate complexity plankton ecosystem model for the global domain, *Geosci. Model Dev.*, 4, 381–417, doi:10.5194/gmd-4-381-2011, 2011.

Zhang, W. and Arhonditsis, G. B.: A Bayesian hierarchical framework for calibrating aquatic biogeochemical models, *Ecol. Model.*, 220, 2142–2161, doi:10.1016/j.ecolmodel.2009.05.023, 2009.

Table 1. 8-dimensional MEDUSA parameter space for target model emulation.

Parameter	Description and units	Lower bound	Upper bound
α_{Pn}	chlorophyll-specific initial slope of P-I curve for non-diatoms $\text{g C (g Chl)}^{-1} (\text{W m}^{-2})^{-1} \text{d}^{-1}$	7.5	30
$k_{N,Pn}$	N nutrient uptake half-saturation concentration for non-diatoms mmol N m^{-3}	0.1	2.5
$k_{Fe,Pn}$	Fe nutrient uptake half-saturation concentration for non-diatoms mmol Fe m^{-3}	0.000066	0.0017
k_{μ}	microzooplankton grazing half-saturation concentration mmol N m^{-3}	0.16	4
ϕ	zooplankton grazing inefficiency –	0.05	0.45
$\mu_{l,Pn}$	non-diatom phytoplankton density-independent loss rate d^{-1}	0.01	0.04
k_{Pn}	non-diatom phytoplankton half-saturation concentration for density-dependent loss mmol N m^{-3}	0.1	2.5
w_g	detrital sinking rate m d^{-1}	1.5	6

Table 2. Representative sample from 8-dimensional MEDUSA parameter space.

Parameter set	α_{Pn}	$k_{\text{N,Pn}}$	$k_{\text{Fe,Pn}}$	k_{μ}	ϕ	$\mu_{1,\text{Pn}}$	k_{Pn}	w_{g}
1	12.2	1.54	0.00104	0.19	0.27	0.0325	0.31	1.61
2	10.6	1.12	0.00021	0.94	0.39	0.0283	0.22	4.87
3	18.5	2.13	0.00011	0.36	0.23	0.0123	1.54	5.60
4	8.0	0.31	0.00145	0.26	0.15	0.0246	0.81	3.22
5	14.0	0.81	0.00055	0.68	0.35	0.0107	0.43	2.12
6	28.0	0.12	0.00008	1.79	0.11	0.0214	0.12	2.44
7	9.2	0.43	0.00015	3.41	0.19	0.0187	0.16	1.85
8	21.2	0.22	0.00076	1.30	0.07	0.0141	0.59	3.69
9	24.4	0.16	0.00039	0.49	0.43	0.0162	1.12	4.24
10	16.1	0.59	0.00028	2.47	0.31	0.0373	2.13	2.80

Table 3. Uninformed emulator robustness evaluation for all 10 experiments and for the 9 experiments excluding Experiment 6.

Site	Direct UQ Method		Indirect UQ Method	
	D_{U_d} mean	D_{U_d} std. dev.	D_{U_i} mean	D_{U_i} std. dev.
60° N	0.03 (0.08)	1.17 (1.16)	0.15 (0.20)	0.98 (0.89)
55° N	0.01 (−0.03)	1.17 (1.07)	−0.02 (−0.07)	1.10 (0.97)
50° N	0.48 (0.05)	1.88 (0.98)	0.58 (−0.15)	3.06 (1.43)
45° N	0.16 (0.03)	1.19 (0.99)	0.08 (−0.07)	1.48 (1.32)
40° N	−0.19 (−0.07)	1.41 (1.29)	−0.12 (−0.10)	0.98 (0.99)
35° N	0.16 (0.04)	1.06 (1.02)	−0.47 (−0.72)	1.89 (1.77)
30° N	0.07 (−0.04)	1.15 (1.15)	0.36 (0.03)	1.88 (1.53)
25° N	−0.07 (−0.03)	1.08 (1.12)	−0.76 (−0.70)	1.63 (1.69)
20° N	−0.04 (−0.07)	1.33 (1.31)	−0.64 (−0.66)	1.04 (1.01)
15° N	0.47 (−0.01)	2.07 (1.21)	0.23 (−0.38)	2.60 (1.29)
10° N	0.36 (0.01)	1.55 (1.03)	0.19 (−0.10)	1.21 (0.74)
5° N	0.03 (−0.07)	1.23 (1.22)	−0.09 (−0.30)	2.10 (2.01)
ALL	0.12 (−0.01)	1.41 (1.13)	−0.04 (−0.25)	1.82 (1.39)

Table 4. MEDUSA phytoplankton parameters (MarMOT 1.1 configuration).

Symbol	Description and units	Standard value
α_{Pn}	chlorophyll-specific initial slope of P-I curve for non-diatoms $\text{g C (g chl)}^{-1} (\text{W m}^{-2})^{-1} \text{d}^{-1}$	15
$f_{\alpha\text{Pd}} = \frac{\alpha_{\text{Pd}}}{\alpha_{\text{Pn}}}$	chlorophyll-specific initial slope of P-I curve for diatoms relative to that for non-diatoms –	0.75
V_{Pn}	maximum non-diatom growth rate at 0°C d^{-1}	0.53
$f_{V\text{Pd}} = \frac{V_{\text{Pd}}}{V_{\text{Pn}}}$	maximum growth rate at 0°C of diatoms relative to that of non-diatoms –	0.9434
$k_{\text{N,Pn}}$	N nutrient uptake half-saturation concentration for non-diatoms mmol N m^{-3}	0.5
$f_{k_{\text{N,Pd}}} = \frac{k_{\text{N,Pd}}}{k_{\text{N,Pn}}}$	N nutrient uptake half-saturation concentration for diatoms relative to that for non-diatoms –	0.5
k_{Si}	Si nutrient uptake half-saturation concentration for diatoms mmol Si m^{-3}	0.75
$k_{\text{Fe,Pn}}$	Fe nutrient uptake half-saturation concentration for non-diatoms mmol Fe m^{-3}	0.00033
$f_{k_{\text{Fe,Pd}}} = \frac{k_{\text{Fe,Pd}}}{k_{\text{Fe,Pn}}}$	Fe nutrient uptake half-saturation concentration for diatoms relative to that for non-diatoms –	2.03

Table 5. MEDUSA zooplankton parameters (MarMOT 1.1 configuration).

Symbol	Description and units	Standard value
g_{μ}	maximum microzooplankton grazing rate d^{-1}	2
$f_{gm} = \frac{g_m}{g_{\mu}}$	maximum grazing rate of mesozooplankton relative to that of microzooplankton –	0.25
k_{μ}	microzooplankton grazing half-saturation concentration mmol N m^{-3}	0.8
$f_{km} = \frac{k_m}{k_{\mu}}$	grazing half-saturation concentration for mesozooplankton relative to that of microzooplankton –	0.375
ϕ	zooplankton grazing inefficiency –	0.2
β_N	zooplankton N assimilation efficiency –	0.69
$a_{\beta C} = \frac{\beta_C - \beta_N}{\beta_N}, \beta_C \leq \beta_N$ $a_{\beta C} = \frac{\beta_C - \beta_N}{1 - \beta_N}, \beta_C > \beta_N$	offset of zooplankton C assimilation efficiency from that of N as a fraction of maximum offset possible –	0
k_C	zooplankton net C growth efficiency –	0.8
$p_{\mu Pn}$	microzooplankton grazing preference for live food (non-diatom phytoplankton) –	0.75
$p_{mLive} = p_{mPn} + p_{mPd} + p_{mZ\mu}$	mesozooplankton grazing preference for live food (phytoplankton or microzooplankton) –	0.85
$p_{C, mP} = \frac{p_{mPn} + p_{mPd}}{p_{mPn} + p_{mPd} + p_{mZ\mu}}$	mesozooplankton conditional grazing preference for phytoplankton, given live food –	0.5882
$p_{C, mPn} = \frac{p_{mPn}}{p_{mPn} + p_{mPd}}$	mesozooplankton conditional grazing preference for non-diatoms, given phytoplankton –	0.3

Table 6. MEDUSA plankton loss-related parameters (MarMOT 1.1 configuration).

Symbol	Description and units	Standard value
$\mu_{1,Pn}$	non-diatom phytoplankton density-independent loss rate d^{-1}	0.02
$f_{\mu 1,Pd} = \frac{\mu_{1,Pd}}{\mu_{1,Pn}}$	density-independent loss rate of diatoms relative to that of non-diatom phytoplankton –	1
$f_{\mu 1,Z\mu} = \frac{\mu_{1,Z\mu}}{\mu_{1,Pn}}$	density-independent loss rate of microzooplankton relative to that of non-diatom phytoplankton –	1
$f_{\mu 1,Zm} = \frac{\mu_{1,Zm}}{\mu_{1,Z\mu}}$	density-independent loss rate of mesozooplankton relative to that of microzooplankton –	1
$\mu_{2,Pn}$	non-diatom phytoplankton maximum density-dependent loss rate d^{-1}	0.1
k_{Pn}	non-diatom phytoplankton half-saturation concentration for density-dependent loss $mmol\ N\ m^{-3}$	0.5
$f_{\mu 2,Pd} = \frac{\mu_{2,Pd}}{\mu_{2,Pn}}$	maximum density-dependent loss rate of diatoms relative to that of non-diatom phytoplankton –	1
$f_{kPd} = \frac{k_{kPd}}{k_{Pn}}$	density-dependent loss half-saturation concentration of diatoms relative to that of non-diatom phytoplankton –	1
$f_{\mu 2,Z\mu} = \frac{\mu_{2,Z\mu}}{\mu_{2,Pn}}$	maximum density-dependent loss rate of microzooplankton relative to that of non-diatom phytoplankton –	1
$f_{kZ\mu} = \frac{k_{kZ\mu}}{k_{Pn}}$	density-dependent loss half-saturation concentration of microzooplankton relative to that of non-diatom phytoplankton –	1
$f_{\mu 2,Zm} = \frac{\mu_{2,Zm}}{\mu_{2,Z\mu}}$	maximum density-dependent loss rate of mesozooplankton relative to that of microzooplankton –	2
$f_{kZm} = \frac{k_{kZm}}{k_{Z\mu}}$	density-dependent loss half-saturation concentration of mesozooplankton relative to that of microzooplankton –	1.5
$D1_{frac}$	fast detritus fraction of diatom losses –	0.75
$D2_{frac}$	fast detritus fraction of mesozooplankton losses –	1
Diss	diatom frustule dissolution rate d^{-1}	0.006
w_g	detrital sinking rate $m\ d^{-1}$	3

Table 7. MEDUSA parameter space for 28-dimensional sensitivity analysis.

Parameter	Standard value	Lower bound	Upper bound	Transformation
α_{Pn}	15	7.5	30	log
$f_{\alpha\text{Pd}}$	0.75	0.56	1	log
V_{Pn}	0.53	0.27	1.1	log
f_{VPd}	0.9434	0.89	1	log
$k_{\text{N,Pn}}$	0.5	0.1	2.5	log
$f_{\text{kN,Pd}}$	1.5	1	2.3	log
k_{Si}	0.75	0.15	3.8	log
$k_{\text{Fe,Pn}}$	0.00033	0.000066	0.0017	log
$f_{\text{kFe,Pd}}$	2.03	1	4.1	log
g_{μ}	2	1	4	log
f_{gm}	0.25	0.13	0.5	log
k_{μ}	0.8	0.16	4	log
f_{km}	0.375	0.14	1	log
ϕ	0.2	0.05	0.45	
β_{N}	0.69	0.44	0.94	
k_{C}	0.8	0.55	0.95	
$p_{\mu\text{Pn}}$	0.75	0.5	0.95	
p_{mLive}	0.85	0.6	0.95	
$p_{\text{c, mP}}$	0.5882	0.34	0.84	
$p_{\text{c, mPn}}$	0.3	0.05	0.55	
$\mu_{1,\text{Pn}}$	0.02	0.01	0.04	log
$\mu_{2,\text{Pn}}$	0.1	0.05	0.2	log
k_{Pn}	0.5	0.1	2.5	log
$f_{\mu 2,\text{Zm}}$	2	1	4	log
f_{kZm}	1.5	1	2.3	log
DI_{frac}	0.75	0.5	0.95	
Diss	0.006	0.003	0.012	log
w_{g}	3	1.5	6	log

Table 8. Parameter sensitivity of annual mean model output from 28 dimensional analysis, showing parameters that explain 5 % or more of the variance in either variable at 1 or more sites.

Parameter	Primary Production		Particulate Export Ratio at 207 m		Selected ?
	No. of sites with $r^2 \geq 0.05$	Maximum r^2	No. of sites with $r^2 \geq 0.05$	Maximum r^2	
α_{Pn}	8	0.44	4	0.08	yes
V_{Pn}	4	0.15	0	< 0.05	no
$k_{N,Pn}$	5	0.22	2	0.10	yes
$k_{Fe,Pn}$	9	0.34	3	0.08	yes
k_{μ}	5	0.17	2	0.13	yes
f_{km}	0	< 0.05	1	0.05	no
ϕ	4	0.15	7	0.17	yes
k_C	3	0.07	2	0.10	no
$\mu_{1,Pn}$	6	0.11	11	0.10	yes
$\mu_{2,Pn}$	3	0.06	0	< 0.05	no
k_{Pn}	4	0.22	7	0.15	yes
w_g	6	0.17	11	0.38	yes

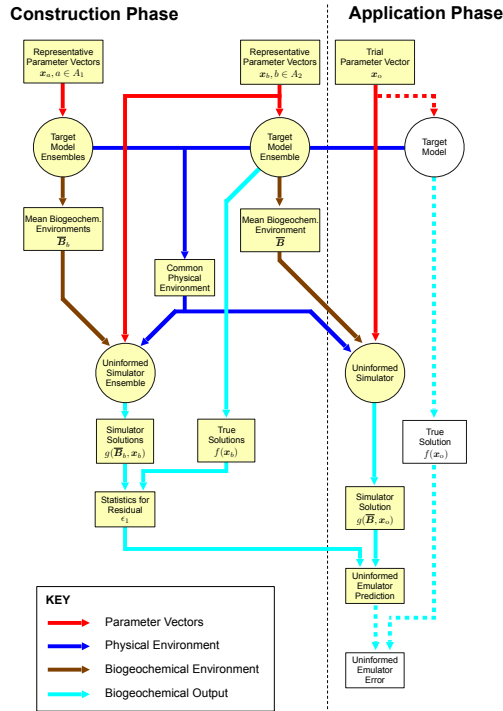


Figure 1. Data flow for emulator construction and application to the prediction of target model output where simulator uncertainty is quantified by the direct method. A_1 and A_2 are arbitrary sets of indices satisfying $A_1 \cap A_2 = \emptyset$. Simulation steps are indicated by circles. The dotted lines and uncoloured boxes indicate data flow for validating emulator performance against a known truth. They are not part of the practical application procedure, where the truth would be unknown.

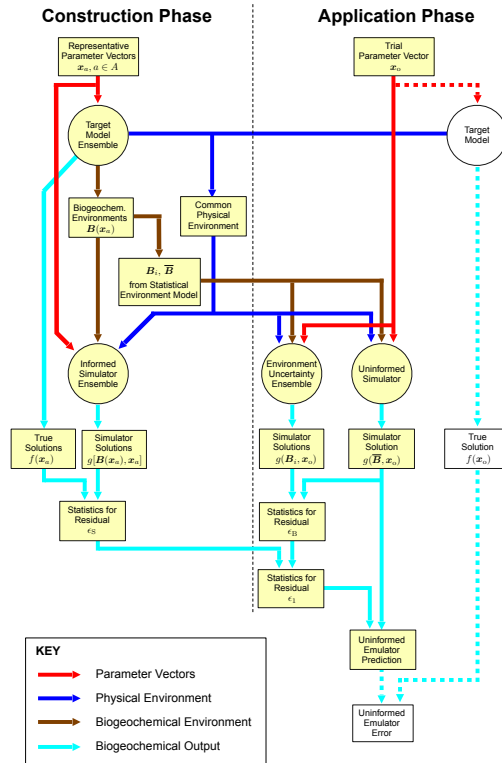


Figure 2. Data flow for emulator construction and application to the prediction of target model output where simulator uncertainty is quantified by the indirect method. A is an arbitrary set of indices. Simulation steps are indicated by circles. The dotted lines and uncoloured boxes indicate data flow for validating emulator performance against a known truth. They are not part of the practical application procedure, where the truth would be unknown.




Figure 3. 5 day mean surface chlorophyll output from 3-D NEMO-MEDUSA simulations for the 10 parameter vectors in Table 2, colour coded by Parameter Set number.




Figure 4. 5 day mean surface chlorophyll output for 1998 at all 12 sites from the uninformed simulator, compared with that from the matching 3-D NEMO-MEDUSA reference simulation. Results are shown for the 10 different parameter vectors in Table 2, colour coded by Parameter Set number.




Figure 5. Uninformed emulator error statistics for $\log_{10}(\text{surface chlorophyll})$ (mg m^{-3}) over 10 experiments, one experiment for each of the parameter vectors in Table 2: **(a)** bias and **(b)** r.m.s. error. The statistics are shown for the simulators without bias correction and for the bias-corrected simulator array, which is the uninformed emulator. The emulator statistics are given for emulator versions constructed using direct and indirect uncertainty quantification methods (i.e. for errors d_{U_d} and d_{U_i}).




Figure 6. Normalized uninformed emulator error for emulator versions constructed using (a) the direct uncertainty quantification method (D_{Ud}) and (b) the indirect uncertainty quantification method (D_{Ui}). Errors are shown for the 10 different parameter vectors in Table 2, colour coded by Parameter Set number. Off scale D_{Ui} values not shown at the beginning of 1997 go up to about 26 at 55° N and about 35 at 60° N.

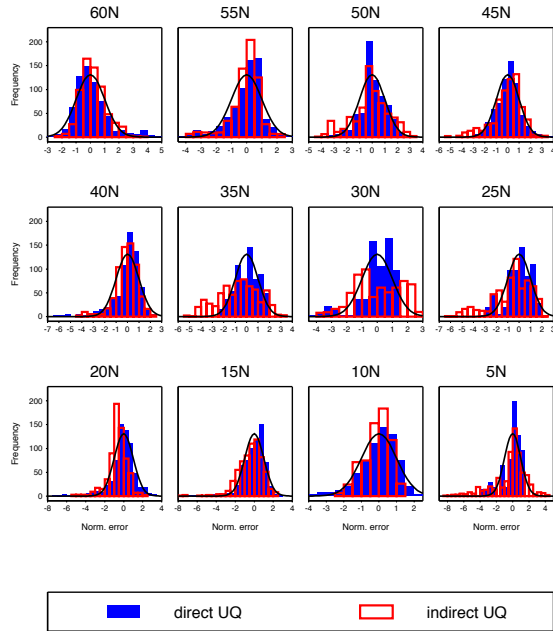


Figure 7. 1998 distributions of the normalized error for the uninformed emulator constructed using the direct and indirect uncertainty quantification methods: D_{Ud} and D_{Ui} . Results for 9 of the 10 parameter vector experiments are combined. Experiment 6, for which large extremes occur, is excluded. The predicted normalized error distribution, over-plotted for reference, is Gaussian with zero mean and unit standard deviation at all times and locations.




Figure 8. 5 day mean surface chlorophyll output for 1998 at all 12 sites from the uninformed simulator with lateral flux perturbations set to zero, compared with that from the matching 3-D NEMO-MEDUSA reference simulation. Results are shown for the 10 different parameter vectors in Table 2, colour coded by Parameter Set number.

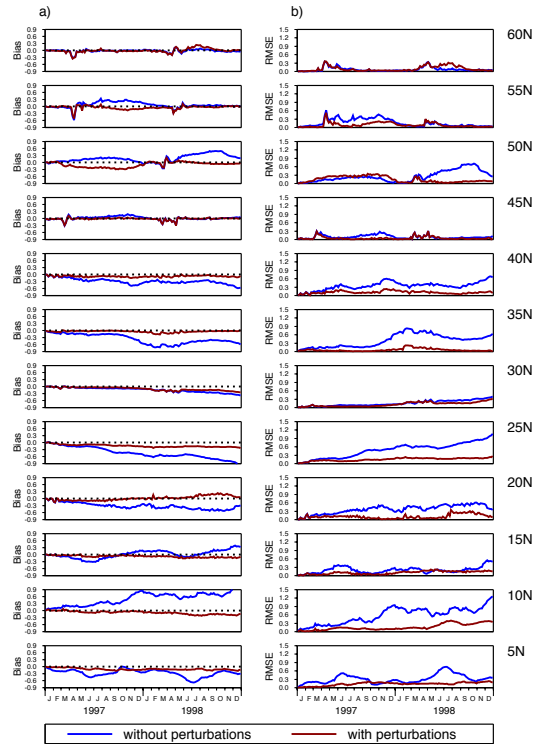


Figure 9. Informed simulator error statistics for $\log_{10}(\text{surface chlorophyll})$ (mg m^{-3}) over 10 experiments, one experiment for each of the parameter vectors in Table 2. **(a)** bias and **(b)** r.m.s. error. The statistics are shown for informed simulators with and without lateral flux perturbations.




Figure 10. Statistics for the uninformed simulator residual ϵ_1 , predicted by the direct and indirect uncertainty quantification methods for all 10 experiments: **(a)** residual means \bar{u}_1 and $\bar{u}_S + \bar{u}_B$; **(b)** residual standard deviations s_1 and $\sqrt{s_S^2 + s_B^2}$. Values are in \log_{10} (chlorophyll) units with chlorophyll in mg m^{-3} .




Figure 11. Predicted statistics for the mean environment simulation residual ϵ_S and the parametric environment residual ϵ_B for all 10 experiments: **(a)** residual means \bar{u}_S and \bar{u}_B ; **(b)** residual standard deviations s_S and s_B . Values are in $\log_{10}(\text{chlorophyll})$ units with chlorophyll in mg m^{-3} .