Strategic model reduction by analysing model sloppiness: a case study in coral calcification

Sarah A. Vollert[a,b,c,*], Christopher Drovandi[a,b,c], Gloria M. Monsalve-Bravo[d,e,f] &
Matthew P. Adams[a,b,c,d]

[a]Centre for Data Science, Queensland University of Technology, Brisbane, QLD 4001, Australia

[b]ARC Centre of Excellence for Mathematical and Statistical Frontiers, Queensland University of Technology, Brisbane, QLD 4001, Australia

[c]School of Mathematical Sciences, Queensland University of Technology, Brisbane, QLD 4001, Australia

[d]School of Chemical Engineering, The University of Queensland, St Lucia, QLD 4072, Australia

[e]School of Earth and Environmental Sciences, The University of Queensland, St Lucia, QLD 4072, Australia

[f]Centre for Biodiversity and Conservation Science, The University of Queensland, St Lucia, QLD 4072, Australia

*Corresponding author. E-mail: sarah.vollert@connect.qut.edu.au

Highlights

- We propose a new model reduction framework based on a sloppy analysis.

- Model reduction is guided by identifying mechanisms which weakly inform predictions.

- We successfully reduce a complex coral growth model, preserving predictive accuracy.

- A Bayesian sloppy analysis is advantageous when the likelihood surface is complex.

Abstract

It can be difficult to identify ways to reduce the complexity of large models whilst maintaining predictive power, particularly where there are hidden parameter interdependencies. Here, we demonstrate that the analysis of model sloppiness can be a new invaluable tool for strategically simplifying complex models. Such an analysis identifies parameter combinations which strongly and/or weakly inform model behaviours, yet the approach has not previously been used to inform model reduction. Using a case study on a coral calcification model calibrated to experimental data, we show how the analysis of model sloppiness can strategically inform model simplifications which maintain predictive power. Additionally, when comparing various approaches to analysing sloppiness, we find that Bayesian methods can be advantageous when unambiguous identification of the best-fit model parameters is a challenge for standard optimisation procedures.

Key words: *sensitivity analysis, Bayesian inference, Sequential Monte Carlo, maximum likelihood estimation, parameter interdependence, model reduction*

1 Introduction

Mathematical models are used to understand complex biological and ecological systems, and to make predictions about system behaviours even in extreme conditions (Getz et al., 2018; Jeong et al., 2018). These models ideally include as much of the expected system dynamics as possible in order to gain mechanistic exploratory power (Snowden et al., 2017). However, this constructionist approach can lead to large, complex models, which can become problematic for various reasons. Complexity can introduce practical issues – such as difficulty implementing, calibrating, solving, and interpreting models (Hong et al., 2017) – as well as resulting in over-fitting, poor predictive performance, uncertainty in estimated parameter values and potentially becoming unscientific by being harder to disprove. Models should aim to balance being simple enough to capture general trends without over-fitting, but complex enough to capture the key features of a dataset – this is known as the bias-variance tradeoff (Geman et al., 1992).

Development of strategies to reduce a model's complexity without loss of explanatory power is a research area of increasing interest (Hjelkrem et al., 2017; Jeong et al., 2018; Snowden et al., 2017; Transtrum and Qiu, 2014). However, it can be challenging to identify appropriate model reductions for complex biological or ecological models, because there are many possible ways to reduce a model (Cox et al., 2006; Jeong et al., 2018). More so, parameter interdependencies can make it difficult to determine the informativeness of individual model components on the outputs (Gibbons et al., 2010; Transtrum and Qiu, 2014). Hence, systematic model reduction methods can be helpful for maintaining predictive power when simplifying a model.

There are a variety of existing methods in the literature which propose ways of simplifying an underlying model. For example, projection-based methods aim to reduce the degrees of freedom of a model (Schilders et al., 2008). Systematically removing parameters from models is challenging (Transtrum and Qiu, 2014), so mechanistic-focused methods commonly fix parameter values or state variables (Cox et al., 2006; Crout et al., 2009; Elevitch and Johnson, 2020; Hjelkrem et al., 2017; Lawrie and Hearne, 2007) or lump them together (Huang et al., 2005; Liao and Lightfoot, 1988; Pepiot et al., 2019; Snowden et al., 2017) to make models more efficient and easier to calibrate. Sensitivity analysis methods are typically used to identify the importance of variations in parameter values on model outputs (Mara et al., 2017; Saltelli et al., 2004), and can inform model reductions in a factor-fixing setting (Cox et al., 2006; Hjelkrem et al., 2017; Hsieh et al., 2018; Van Werkhoven et al., 2009). Sensitivity analyses used for model reduction typically focus on the sensitivity of model outputs to changes in parameter values. However, we may instead be interested in which parameters are constrained by an available dataset. Additionally, the effects of changing model parameters on outputs often depend greatly on the assumed values of other parameters; hence we need to consider the effect of parameters in combination on model behaviours, rather than individuals (Brown et al., 2004; Brown and Sethna, 2003; Gutenkunst et al., 2007; Monsalve-Bravo et al., 2022; Transtrum et al., 2011; Transtrum and Qiu, 2014).

This paper presents a new approach to strategic model reduction based on an analysis of model sloppiness. This type of sensitivity analysis captures the sensitivity of model outputs informed by a dataset; it looks at model-data fit sensitivities, rather than just model sensitivities. The analysis of model sloppiness draws on dimension reduction techniques to identify hidden parameter interdependencies (Transtrum et al., 2011, 2015). Such parameter interdependencies, or parameter combinations, are identified by analysing the curvature of the surface which describes how the model-data fit depends on the model parameters (Brown et al., 2004; Brown and Sethna, 2003; Monsalve-Bravo et al., 2022). Consequently, parameter combinations which strongly (or weakly) influence model predictions can be revealed by identifying directions in parameter space

which most (or least) influence model outputs. The analysis accounts for individual parameters acting together or against each other (compensating effects) and identifies sensitive and insensitive parameter combinations for the model-data fit (Brown et al., 2004; Brown and Sethna, 2003; Gutenkunst et al., 2007; Monsalve-Bravo et al., 2022; Transtrum et al., 2011; Transtrum and Qiu, 2014). Hence, an analysis of model sloppiness could be used to identify and remove model mechanisms that only weakly impact predictions (or are weakly informed by the data), whilst accounting for hidden compensatory effects between individual model parameters.

The concept of model sloppiness in model reduction methods has previously been explored by Transtrum and Qiu (2014), who proposed the manifold boundary approximation method (MBAM). The MBAM uses an information theory-based approach whereby a parameter-independent geometric interpretation of the model is used to systematically reduce the effective degrees of freedom (Transtrum and Qiu, 2014). However, alternative methods have been proposed for dimensionality reduction; for example, the active subspace method (Constantine et al., 2016) similarly captures model-data fit sensitivities and can be used for individual parameter rankings through activity scores (Constantine and Diaz, 2017). More recently, Elevitch and Johnson (2020) used a spectral analysis of the Hessian matrix, akin to the non-Bayesian analysis of model sloppiness (Brown et al., 2004), to quantitatively rank parameter importance and thus determine which parameters should be estimated or fixed, rather than simplifying the model structure.

In contrast to these methods, our approach is based solely on the analysis of a model's sloppiness. Our method uses this analysis to identify insensitive groups of parameters which represent processes or mechanisms within a model, whilst accounting for the interdependencies between parameters. Thus, if model outputs are insensitive to changes in model parameters associated with a certain mechanism, such a mechanism is identified to have little effect on the overall model predictions (model-data fit). Using this sensitivity analysis, insensitive mechanisms are thus removed from the model (rather than being fixed or lumped) to produce a conceptually simpler model (Transtrum and Qiu, 2014) which maintains its predictive capability. As a result, the model remains expressed in terms of the parameters of interest and preserves mechanistic interpretability. Additionally, the method we propose can take advantage of Bayesian sensitivity matrices described recently elsewhere (Monsalve-Bravo et al., 2022).

In this work, we showcase the potential for model sloppiness to inform strategic model reductions using a case study on a complex physiological model predicting coral calcification rates. We also use this case study to demonstrate that both Bayesian and non-Bayesian approaches to analysing model sloppiness (Brown and Sethna, 2003; Monsalve-Bravo et al., 2022) may be suitable for model reduction, although Bayesian approaches may be advantageous when the best-fit parameter values cannot be easily identified using standard search procedures (e.g. where the likelihood surface does not have a well defined peak).

2 Methods

2.1 Model calibration

Within complex models, there are often parameters which cannot be measured directly and must instead be estimated through a model-data calibration process. To analyse the model sloppiness, these unknown parameter values must first be estimated using either a classical (frequentist) calibration process to obtain a single-point estimate, or a Bayesian model-data calibration to obtain probabilistic distributions for parameters.

Single point estimates of unknown parameters can be obtained through maximum likelihood estimation (MLE), which is a common frequentist approach used for model calibration (Jackson et al., 2000). More specifically, unknown model parameters $\boldsymbol{\theta}$ are estimated using observed data $\mathbf{y}$ by maximising an appropriately chosen likelihood function $f(\mathbf{y}|\boldsymbol{\theta})$. Under the common assumptions that measurement errors follow a Gaussian distribution of constant standard deviation, and that observations are independent, the likelihood that a set of data $\mathbf{y}$ is observed with a mean given by the model $\mathbf{y}_{\mathrm{model}}(\boldsymbol{\theta})$ and a standard deviation of σ can be defined as

$$f(\mathbf{y}|\boldsymbol{\theta}) = \prod_{i=1}^{n_d} f(y_i|\boldsymbol{\theta}) = \prod_{i=1}^{n_d} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(y_{\mathrm{model},i}(\boldsymbol{\theta}) - y_i)^2}{2\sigma^2}\right), \tag{1}$$

where $\mathbf{y} = \{y_1, y_2, \ldots, y_{n_d}\}$ is a dataset of n_d independent observations, y_i is the ith observation in the dataset, $f(\mathbf{y}|\boldsymbol{\theta})$ is the likelihood function, and $y_{\mathrm{model},i}(\boldsymbol{\theta})$ is the ith model prediction of the data given the conditions of observation y_i and parameters $\boldsymbol{\theta}$. Formally, the MLE ($\boldsymbol{\theta}_{\mathrm{MLE}}$) is obtained by maximising the likelihood function, such that

$$\boldsymbol{\theta}_{\mathrm{MLE}} = \mathrm{argmax}_{\boldsymbol{\theta}} f(\mathbf{y}|\boldsymbol{\theta}). \tag{2}$$

Hence, $\boldsymbol{\theta}_{\mathrm{MLE}}$ represents the "best-fit" parameter values under the assumed statistical model structure and observed dataset. This process is represented as the left (orange) branch in Figure 1.
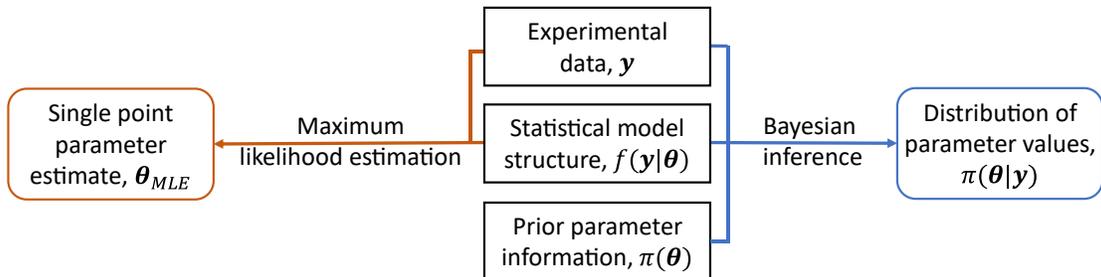

Figure 1: A conceptual figure, highlighting the difference between a model calibrated via MLE (left) and Bayesian inference (right). MLE combines experimental data with a model structure to produce a frequentist model. Bayesian inference also incorporates prior information about the model parameters to produce a probabilistic model that encapsulates parameter uncertainty.

However, this frequentist approach does not account for prior beliefs about the parameter values (Uusitalo et al., 2015), and also ignores parameter values $\boldsymbol{\theta} \neq \boldsymbol{\theta}_{\mathrm{MLE}}$ which may also plausibly represent the data. To analyse the effect of the parameters on the behaviour of the system, the full range of values that a parameter could potentially take should be considered (Jakeman et al., 2006). As an alternative to MLE, Bayesian inference can be used for model-data calibration to obtain a probability distribution for $\boldsymbol{\theta}$ that accounts for prior parameter information, an assumed model structure and the data (Girolami, 2008). Such a distribution is called the posterior distribution as it represents a probability distribution for $\boldsymbol{\theta}$ *after* considering the data.

Using any known information about the parameters, a "prior" distribution $\pi(\boldsymbol{\theta})$ is placed on $\boldsymbol{\theta}$. To obtain the posterior distribution, the prior distribution is multiplied by the likelihood function via Bayes' Theorem,

$$\pi(\boldsymbol{\theta}|\mathbf{y}) \propto f(\mathbf{y}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}), \tag{3}$$

where $\pi(\boldsymbol{\theta}|\mathbf{y})$ is the posterior distribution of the parameters, $f(\mathbf{y}|\boldsymbol{\theta})$ is the likelihood function and $\pi(\boldsymbol{\theta})$ is the prior distribution. The posterior is a distribution of potential parameter values

informed by the data and prior information such that it quantifies parametric uncertainty (Giro-lami, 2008). In this paper, a sequential Monte Carlo (SMC) algorithm (Chopin, 2002; Del Moral et al., 2006) was used to sample from the posterior distribution, providing a representative sample of plausible parameter values, given the available prior information and dataset (see Drovandi and Pettitt (2011) or Jeremiah et al. (2012) for more information on SMC sampling). Bayesian inference for model calibration is depicted as the right (blue) branch of Figure 1.

2.2   Analysing model sloppiness

The analysis of model sloppiness is a type of sensitivity analysis which considers the model sensitivity to all parameters informed by a dataset (Monsalve-Bravo et al., 2022). As with global sensitivity analysis approaches (Geris and Gomez-Cabrero, 2016; Marino et al., 2008; Saltelli et al., 1993), the analysis of model sloppiness accounts for the model-data fit sensitivities of parameter *combinations* across all parameters. Thus, this approach can mathematically characterise the parameter combinations that the model-data fit is most sensitive to (Brown and Sethna, 2003; Transtrum et al., 2015).

A sloppy model refers to a model where most of the model behaviour is captured through a few tightly constrained (stiff) parameter combinations (stiff eigenparameters) (Brown and Sethna, 2003), which are highly influential on model predictions of the data (Transtrum et al., 2015), but remains insensitive to many loosely constrained (sloppy) parameter combinations (sloppy eigenparameters). Stiff and sloppy parameter combinations are identified through the eigendecomposition of a sensitivity matrix (Gutenkunst et al., 2007). While there are various approaches to the construction of the sensitivity matrix – some of these are summarised in Section 2.3 – we focus here on how to obtain stiff/sloppy eigenparameters once this matrix is successfully computed.

Each eigenvector $\mathbf{v}_j$ of the sensitivity matrix indicates a key direction in parameter space that characterises the sensitivity of the model-data fit to changes in multiple parameters simultaneously. We can express each key direction in parameter space as a specified combination of parameter values – known as an eigenparameter. Changing an eigenparameter's value is equivalent to moving the entire set of (original) model parameters along the direction of the eigenvector associated with that eigenparameter. As eigenvectors of the sensitivity matrix are by definition mutually orthogonal, if model parameters are logarithmically-transformed prior to calculation of the sensitivity matrix (which is a common practice when analysing sloppiness), each eigenparameter $\hat{\theta}_j$ of this matrix can be expressed as a linear combination of natural logarithms of model parameters, following Brown et al. (2004),

$$\log \hat{\theta}_j = v_{j,1} \log \theta_1 + v_{j,2} \log \theta_2 + \cdots + v_{j,n_p} \log \theta_{n_p},$$
$$\therefore \hat{\theta}_j = \theta_1^{v_{j,1}} \theta_2^{v_{j,2}} \cdots \theta_{n_p}^{v_{j,n_p}}, \tag{4}$$

where $\mathbf{v}_j = [v_{j,1}, v_{j,2}, ..., v_{j,n_p}]$ is the jth eigenvector of the sensitivity matrix, n_p is the number of model parameters, and θ_i is the ith parameter in the model. It should be noted that expressing eigenparameters $\hat{\theta}_j$ as a product of all model parameters to the power of the different exponents as shown in Equation (4) is only possible if the parameters are scaled by their logarithm when estimating the sensitivity matrix. Additionally, standard renormalisations can be applied such that the exponents $v_{j,i}$ in Equation (4) are rescaled to be between -1 and 1.

In Equation (4), each eigenparameter $\hat{\theta}_j$ has a corresponding eigenvalue λ_j which attributes a magnitude to the direction of the eigenparameter. The largest eigenvalue (λ_1) corresponds to the stiffest (most sensitive) eigenparameter, and the smallest eigenvalue (λ_{n_p}) corresponds to the

sloppiest (least sensitive) eigenparameter (Transtrum et al., 2015). Therefore, comparing the eigenvalues of all eigenparameters indicates the relative impact that each parameter combination $\hat{\theta}_j$ has on the model-data fit. In our implementation of this approach, we scale all eigenvalues λ_j by dividing them by the largest eigenvalue λ_1 to indicate the relative importance of the eigenparameters. Thus, $0 < \lambda_j/\lambda_1 \leq 1$ for all eigenparameters $\lambda_j, j = 1, \ldots, n_p$.

The sloppy analysis results in a list of n_p weighted parameter combinations $\hat{\theta}_j$ ranked by their influence on the model-data fit via the scaled eigenvalues λ_j/λ_1. This process is summarised in Algorithm 1 (see also Monsalve-Bravo et al. (2022) for an overview of model sloppiness).

Algorithm 1: Process used for analysing the model sloppiness

1. Perform model calibration to obtain estimate(s) of parameters $\boldsymbol{\theta}$.

2. Calculate sensitivity matrix S with respect to the log-parameters, $\phi_i = \log \theta_i$.

3. Find all eigenvalues λ_j and eigenvectors $\boldsymbol{v}_j, j = 1, \ldots, n_p$ of the sensitivity matrix S.

4. Order the eigenvectors by their influence using their associated eigenvalues λ_j, such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{n_p}$.

5. Calculate the rescaled eigenvalues λ_j/λ_1, for all $j = 1, \ldots, n_p$.

6. Renormalise each eigenvector $\boldsymbol{v}_j \leftarrow \boldsymbol{v}_j/\max(|\boldsymbol{v}_j|)$.

7. Report the eigenparameters using Equation (4) and the renormalised eigenvectors obtained from Step 6, ordered in importance by the magnitude of the corresponding eigenvalues.

2.3 Sensitivity matrices

The key quantity required for analysing sloppiness of a model fitted to data is the sensitivity matrix S. This is a square symmetric matrix of size $n_p \times n_p$, in which n_p is the number of estimated model parameters, excluding parameters that represent measurement error (e.g. standard deviation σ in Equation (1)) as the latter can yield eigenparameters that are trivial or difficult to interpret (Monsalve-Bravo et al., 2022). There are various approaches to calculate the sensitivity matrix (see Monsalve-Bravo et al. (2022) for an overview). Although based on similar dimension reduction techniques (e.g., posterior covariance, likelihood informed subspace, or the active subspace methods), each sensitivity matrix considers different sources of information for the model parameters and their interdependencies. For example, the chosen sensitivity matrix may acknowledge or aim to exclude prior beliefs about the model parameters when identifying key parameter interdependencies (e.g. the posterior covariance method and likelihood informed subspace methods, respectively).

In this paper three sensitivity matrices (Table 1) were used to explore model reduction. Firstly, the Hessian matrix of the log-likelihood is used to capture the model-data fit around one point in parameter space based on the likelihood surface (Brown and Sethna, 2003). Secondly, the posterior covariance is a variance-based method which looks at model-data fit sensitivities over the full posterior parameter-space (Brown and Sethna, 2003; Monsalve-Bravo et al., 2022). Thirdly, the likelihood-informed subspace (LIS) method captures the model-data sensitivities of the posterior sample relative to the prior distribution (Cui et al., 2014), so it can be used in comparison with the posterior covariance method to identify the informativeness of the prior distribution

(Monsalve-Bravo et al., 2022). We chose to use these three sensitivity matrices because the interpretation of the analysis of sloppiness for these matrices can provide uniquely different information, as documented recently for multiple simulation problems (Monsalve-Bravo et al., 2022). However, there are other sensitivity matrices that could also be used, such as the Levenberg-Marquardt Hessian (Brown and Sethna, 2003; Gutenkunst et al., 2007), the matrix arising from the active subspace method (Constantine et al., 2016), or a likelihood-free approximation of LIS (Beskos et al., 2018).

Sensitivity matrices are often calculated using logarithmically-scaled parameter values (Brown et al., 2004). This rescaling enforces positivity constraints on the model parameters, helps prevent scaling issues between parameters with units possessing different orders of magnitude, and allows each eigenparameter to be expressed as a product rather than as a sum, as in Equation (4) (Monsalve-Bravo et al., 2022). Here, we denote the logarithmically-rescaled parameters as $\phi_i = \log\theta_i$, for all $i = 1, \dots, n_p$ parameters (see Step 2 of Algorithm 1).

Sensitivity Matrix	Bayesian	Formula	Features	Reference
Hessian evaluated at the MLE (S_H)	No	$H_{i,j} = \dfrac{\partial^2 \log f}{\partial \phi_i \partial \phi_j}\big|_{\phi=\hat{\phi}_{\mathrm{MLE}}},$ where H_{ij} is the (i,j)th element of Hessian matrix (S_H) evaluated at $\hat{\phi}_{\mathrm{MLE}}$, and f is the likelihood function.	Looks at curvature of model-data fit surface at the MLE.	Brown and Sethna (2003)
Posterior covariance method (S_P)	Yes	$S_P = \Sigma^{-1},$ where Σ is the empirical covariance matrix of $\{\phi_m\}_{m=1}^M$, the set of M posterior samples.	Uses the posterior sample to estimate model-data sensitivities over a range of parameter values.	Brown and Sethna (2003); Monsalve-Bravo et al. (2022)
Likelihood informed subspace method (S_L)	Yes	$S_L \approx \dfrac{1}{M} \sum_{m=1}^{M} L^T H(\phi_m) L,$ where $\{\phi_m\}_{m=1}^M$ is the set of M posterior samples, $H(\phi_m)$ is the Hessian matrix evaluated at ϕ_m, and L is the Cholesky factor of the covariance matrix Ω of the prior distribution $\pi(\phi)$, such that $\Omega = LL^T$.	Estimates model-data fit sensitivities by comparing where the data is most informative relative to the prior distribution. Can identify how informative the prior distribution is on the posterior distribution, when compared to the posterior covariance method.	Cui et al. (2014); Monsalve-Bravo et al. (2022)

Table 1: Three methods of constructing the sensitivity matrix for an analysis of model sloppiness. Each sensitivity matrix captures different features of the model-data fit. For the Bayesian sensitivity matrices, it is assumed here that the posterior distribution is approximated by a sufficiently large number M of equally weighted posterior samples.

2.4 Model reduction

Model reduction aims to simplify the model whilst minimising the loss of predictive power of the model (Jeong et al., 2018; Snowden et al., 2017). Given that model sloppiness can identify sensitive parameter combinations, we propose that model reduction can be informed by considering the removal of mechanisms that contribute negligibly to these sensitive parameter combinations. This approach is similar, although not equivalent, to factor-fixing of sets of parameters from a variance-based sensitivity analysis (Saltelli et al., 2004).

Previous work has shown that analysing model sloppiness can reduce the number of model parameters via iteratively removing the sloppiest eigenparameter and simultaneously adapting the model with limiting approximations (Transtrum and Qiu, 2014). Here, we instead investigate the possibility that, if one or more parameters which together characterise an entire process or mechanism within the model have little contribution to the stiffest eigenparameters, this process or mechanism may contribute very little to the model's ability to predict the data that was used for its calibration. Hence, this analysis can identify which *mechanisms* the model-data fit is insensitive to – similar to variance-based sensitivity analyses, such as Sobol's indices (Sobol, 1993), which can analyse sets of parameters or model structures (Mara et al., 2017). These identified mechanisms can be potentially removed from the model with little effect on the model's predictive power.

We suggest that a general quantitative method for selecting the most insensitive model mechanism to be removed would not be appropriate for all models. Instead, the analysis of model sloppiness is used as a source of information for guiding model reductions, paired with an understanding of the model and data being considered. We also recommend quantitatively comparing the original and reduced models using tools such as model evidence, Bayes factors, the Akaike information criterion or the Bayesian information criterion (Kass and Raftery, 1995; Tredennick et al., 2021). The full process of model reduction informed by the analysis of model sloppiness that we propose and investigate here is depicted in Figure 2.
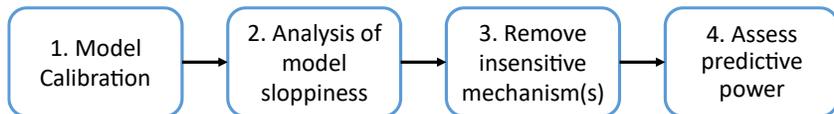
Figure 2: A conceptual diagram of the model reduction process informed by the analysis of model sloppiness. Firstly, the model is calibrated (Section 2.1) using the Bayesian and/or non-Bayesian methods depicted in Figure 1. The second step is an analysis of model sloppiness (Section 2.2) to identify model mechanisms which weakly inform model predictions. Thirdly, the insensitive mechanisms identified in the analysis of model sloppiness are removed from the model and the simplified models are calibrated to the data. Lastly, the predictive power of the reduced model(s) can be assessed, e.g. via goodness-of-fit and/or model selection metrics, to determine the best model for the application.

3 Case study

3.1 Coral calcification model

To demonstrate using the analysis of model sloppiness as a method for strategic model reduction, this method was applied to a data-calibrated model predicting coral calcification rates.

Calcification rates are a common metric for measuring the health of coral reefs (Erez et al., 2011). The process of coral calcification involves the coral laying down its skeleton, resulting in the spatial extension of coral reef structures (Andersson and Gledhill, 2013). This process is vital for the entire ecosystem, because it forms a habitat for a diverse range of marine life as well as providing a stuctural framework for barrier reefs (Hoegh-Guldberg et al., 2017).

The model and data used for this case study are reported by Galli and Solidoro (2018); this is the most recent and comprehensive model for coral reef calcification rates. The model prediction of calcification rates are obtained from the steady state solution of a system of nonlinear ordinary differential equations (ODEs), which together simulate the chemical composition of relevant molecules and ions throughout various physiological compartments of a coral polyp (Figure 3). There are eight mechanisms within the model which represent different chemical processes and reactions through coacting and counteracting flux terms (Table 2), hence, some of these mechanisms could potentially be removed. Further description of the model is provided in Supplementary Material, Section S.1. Galli and Solidoro (2018) calibrated the 20 unknown parameters of the model (see Supplementary Material Section S.2 for details) to an experimental dataset containing 16 data-points obtained from Rodolfo-Metalpa et al. (2010). Each data point consists of paired measurements of calcification rates and the environmental conditions under which this rate was measured (data shown in Table 1 of Galli and Solidoro, 2018).

Mechanism	Description	Parameters
Net calcification	The net effect of calcification and dissolution chemical reactions which produce and remove coral skeleton respectively	N/A
Gross photosynthesis	A chemical reaction caused by zooxanthellae algae yielding carbon reduction in the coelenteron	N/A
Respiration	A chemical reaction caused by zooxanthellae algae yielding carbon gain in the coelenteron	N/A
Seawater-coelenteron diffusion	Diffusion for all chemical species between seawater and coelenteron	s
Coelenteron-ECM paracellular diffusion	Paracellular diffusion for all chemical species between the coelenteron and ECM	k_{pp}
Coelenteron-ECM transcellular diffusion	Transcellular diffusion of carbon dioxide between the coelenteron and ECM	k_{CO2}
Ca-ATPase pump	The active transport of calcium ions through the aboral tissue, driven by ATP	α, β, v_{H_c}, $E0_c$, k_{1f_c}, k_{2f_c}, k_{3f_c}, k_{1b_c}, k_{2b_c}, k_{3b_c}
BAT pump	The active transport of bicarbonate anions through the aboral tissue	$E0_b$, k_{1f_b}, k_{2f_b}, k_{3f_b}, k_{1b_b}, k_{2b_b}, k_{3b_b}

Table 2: The eight model mechanisms and their associated parameters in the model proposed by Galli and Solidoro (2018). Note that the first four mechanisms cannot be removed because they are the rate of interest we wish to output (net calcification), data inputs with no associated parameters (gross photosynthesis and respiration), and a mechanism critical for the model's original purpose of simulating coral responses to ocean acidification (seawater-coelenteron diffusion).
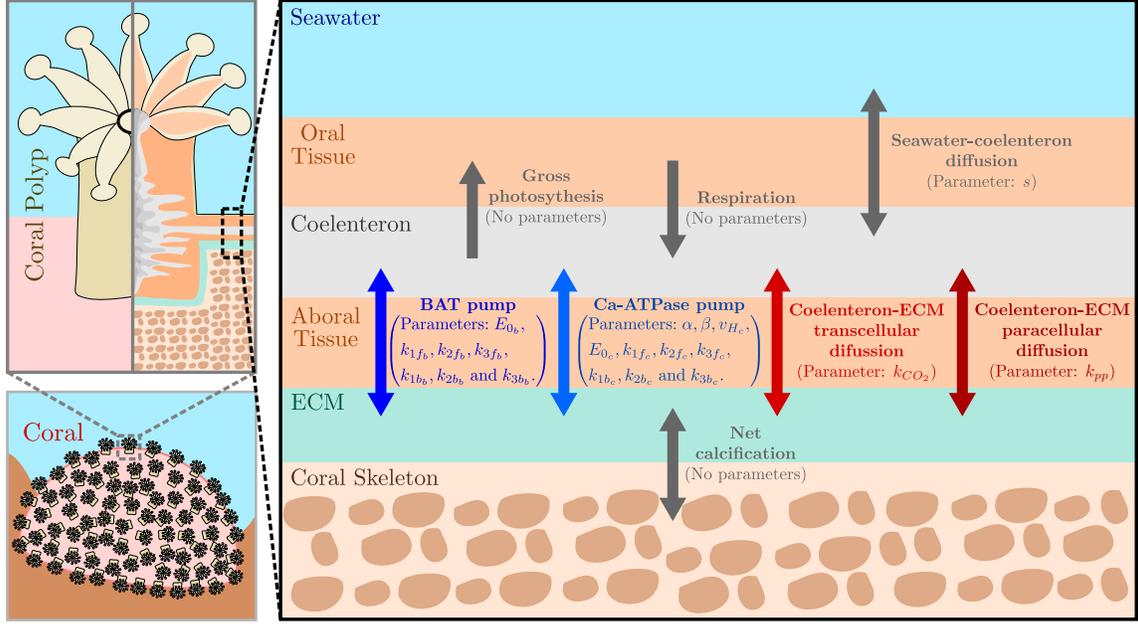
Figure 3: Conceptual model of the chemical processes and reactions within a coral polyp, proposed by Galli and Solidoro (2018). Each of the arrows in this diagram represents a mechanism which affects the chemical species concentrations in the model. This model predicts the calcification reaction rate in the extracellular calcifying medium (ECM) based on environmental conditions.

3.2 Model calibration of the coral calcification model

The coral calcification model was calibrated to the data (see Table 1 of Galli and Solidoro, 2018) using both MLE and Bayesian inference methods. The likelihood function was defined based on Gaussian errors as in Equation (1), where the data was assumed to be normally distributed with a mean calcification rate according to the model proposed by Galli and Solidoro (2018), and a constant and unknown standard deviation of σ.

In cases where the model-data fit surface is complex, local optima can misguide optimisation algorithms such that they do not converge on the global MLE (Transtrum et al., 2011). In the present case study, the likelihood function was high-dimensional due to its dependence on 20 model parameters and had no clearly defined peak, hence the global MLE was difficult to obtain. Instead, 100 local MLEs were obtained and the five with the highest likelihood values were used in further analysis. These local maxima were identified using a gradient-based non-linear function minimisation tool (Matlab's *fmincon* function, described in MathWorks, 2021) using various initial search locations. Figure 4 shows 25 of the identified local maxima with the highest likelihood values (black vertical lines), which are spread across many parameter values, indicating the complexity of the high-dimensional likelihood surface. These distinct parameter estimates each have similarly high likelihoods and this goodness-of-fit is reflected by their predictions of calcification rates when compared to the observations (black asterisks in Figure S1 of Supplementary Material Section S.3). Only one point-estimate (the MLE) is needed for the analysis of model sloppiness if the likelihood surface has a well-defined peak. However, when the likelihood surface is multi-modal and/or flat in certain directions (as is common in sloppy models) it can be difficult to identify the global MLE. Thus, to analyse the reproducibility

of the results, we evaluated the Hessian matrix at five different sets of parameter values – those that yielded the highest values of the likelihood function (orange vertical lines in Figure 4, and orange asterisks in Figure S1 of Supplementary Material Section S.3) algorithm.

To calibrate the coral calcification model a second time, using Bayesian inference instead, we specified uniform prior distributions (grey shaded regions in Figure 4) as most of the parameters for this model were only known to be strictly positive. The posterior distribution was approximated using an adaptive SMC algorithm (Drovandi and Pettitt, 2011) and this algorithm was run multiple times independently to test the reproducibility of both the posterior sample and the later analyses of model sloppiness (we used five independent posterior samples). The estimated marginal distributions were visually indistinguishable for the independently produced posterior samples, indicating that the posterior sample was highly reproducible for a posterior sample size of 5000. Further details of the model calibration procedure performed using Bayesian inference are provided in Supplementary Material, Section S.4. The estimated posterior marginal densities obtained for the model parameters (Figure 4) revealed large uncertainty in many of the model parameters, with only two parameters strongly informed by the data (k_{pp} and β). This result is expected given the limited size of the dataset available for model-data calibration.
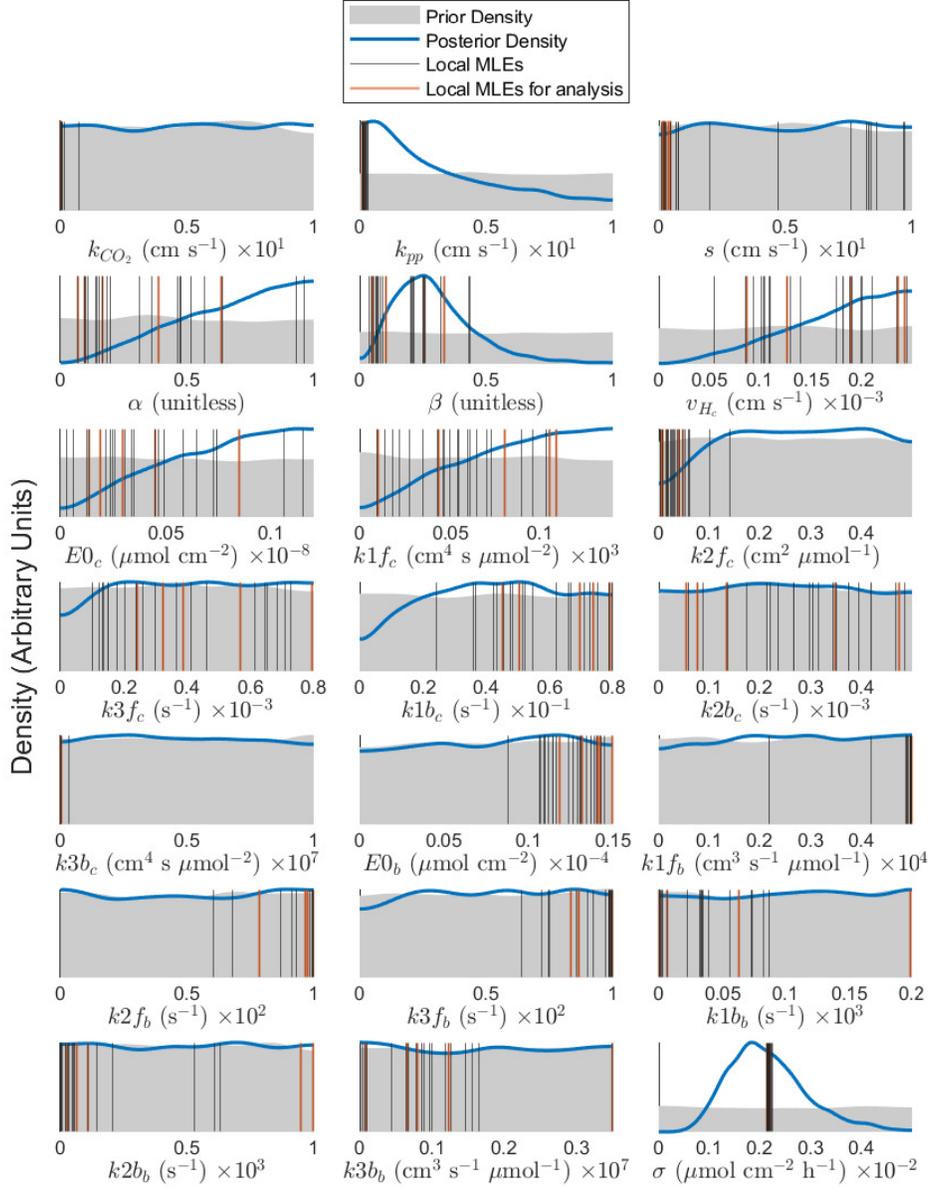
Figure 4: Parameter estimates obtained through both Bayesian inference and MLE model-data calibration methods. The prior and posterior distributions obtained via Bayesian inference are indicated as the grey areas and thick blue lines respectively, showing the change to the distribution as informed by the data. The vertical lines indicate different local likelihood maxima with similarly high likelihoods, with those in orange possessing the highest likelihood function values and thus kept for further analysis. Notice that the local likelihood maxima are distinct and cover a vast range of the parameter space, indicating the complexity of the high-dimensional likelihood surface.

3.3   Analysing sloppiness of the coral calcification model

After the model-data calibration, we used the analysis of model sloppiness to unravel strong parameter interdependencies. This analysis was conducted for each of the three sensitivity

matrices listed in Table 1: the Hessian approach evaluated at each of the five local MLEs, the posterior covariance method, and the LIS approach, with both of the latter evaluated from the posterior distribution samples obtained from the SMC algorithm.

Figure 5 shows the size of the eigenvalues relative to the largest, after eigendecomposition of each sensitivity matrix. Here, the MLE Hessian approach (Table 1, first row) leads to an inconsistent decay in eigenvalue spectra for the five different local MLEs used to evaluate the sensitivity matrix. This result is unsurprising, as the parameter values of each local MLE are in distinctly different locations of the likelihood surface, so the sensitivity of the model-data fit changes in these different local parameter spaces considered. In contrast, the posterior covariance method (Table 1, second row) produces a consistent decay in eigenvalue size across independent posterior samples. The consistent decay is likely because the full range of parameter values from the posterior are considered when analysing the model-data sensitivity, achieving a global analysis of the posterior surface. Finally, the decline in relative eigenvalue size is much more rapid for the LIS framework (Table 1, third row, and Figure 5, blue triangles), when compared to the posterior covariance approach. Here, the relative eigenvalue size for the fifth LIS eigenparameter is comparable to the twentieth (i.e. smallest eigenvalue) from the posterior covariance method.
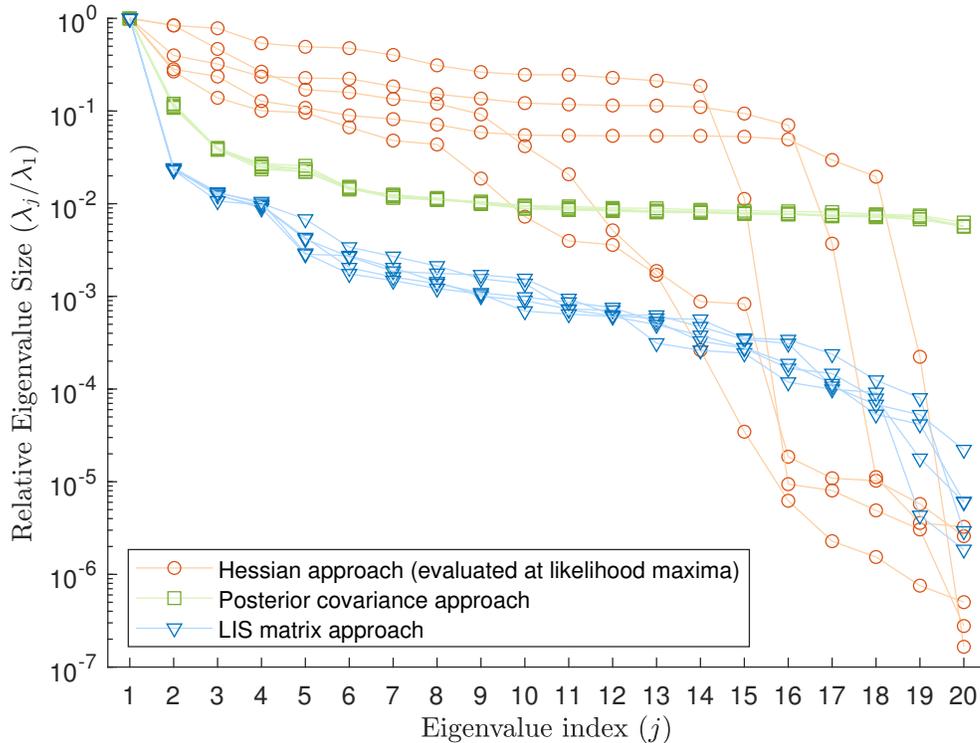
Figure 5: Relative eigenvalue sizes (λ_j/λ_1) compared to the largest eigenvalue (λ_1) for each of the approaches to analysing model sloppiness. Each of the MLE Hessian analyses of model sloppiness were based on parameter estimates (orange vertical lines in Figure 4) from different local maxima. Both the posterior covariance method and LIS sensitivity matrix analyses were conducted on five independent sets of posterior samples obtained using Bayesian inference. Notice that the posterior covariance method leads to extremely similar eigenvalue spectra for the five independent sets of posterior samples.

3.3.1 Hessian matrix evaluated at local likelihood maxima

First, the non-Bayesian approach to analysis of sloppiness (i.e. using the Hessian matrix) was considered. As the Hessian matrix could not be analytically computed for the model in this case study, the finite differences method (Beers, 2007) was used to numerically approximate the Hessian matrix. When using finite differences, the likelihood function is evaluated very close to the MLE, using a step-size $\Delta\theta_i = \delta \times \theta_i, i = 1, \ldots, n_p$, which is as small as possible to minimize truncation error. However, a step-size too small will result in numerical issues (round-off errors). For this application, we used $\delta = 10^{-2}$, as larger step-sizes yielded the same results, but smaller step-sizes yielded inconsistent results which we attributed to numerical errors. The Hessian matrix was evaluated at five different parameter sets – those that yielded the highest values of the likelihood function (orange vertical lines in Figure 4).

After eigendecomposition of the Hessian matrix, evaluated at each parameter set, we used Equation (4) to identify parameter interdependencies (eigenparameters). Figure 6 shows the eigenparameters that correspond to the four largest eigenvalues ($\lambda_1, \lambda_2, \lambda_3$ and λ_4 for the orange circles in Figure 5), for each of the five likelihood maxima considered. In each row of the matrix depicted in Figure 6, the relative contribution of a given (ith) parameter to the (jth) eigenparameter (parameter combination) is indicated by the listed value. This value mathematically represents the eigenvector element value $v_{j,i}$ and can be interpreted as the magnitude of the ith parameter's exponent in the expression (Equation (4)) for the jth eigenparameter. Standard renormalisations during eigendecomposition ensure that $-1 \leq v_{j,i} \leq 1, \ \forall i, j$, such that exponents that are close to 1 or -1 indicate strong contribution of the parameter to the eigenparameter, and exponents close to 0 indicate negligible contribution of the parameter to the eigenparameter. For example, in the first row of Figure 6, the parameter k_{pp} has an exponent of -0.9, so strongly contributes to the first (stiffest) eigenparameter. In contrast, the parameter k_{CO2} in the first row contributes negligibly to the stiffest eigenparameter.

The parameters that contribute to the stiffest eigenparameter are largely associated with the coelenteron-ECM paracellular diffusion mechanism (represented as k_{pp}) and the Ca-ATPase pump mechanism (Figure 6). In contrast, parameters associated with the coelenteron-ECM transcellular diffusion mechanism (represented as k_{CO2}), seawater-coelenteron diffusion mechanism (represented as s) and the BAT pump mechanism all contribute to a lesser extent to the stiffest eigenparameter. While this result is generally true when considering different point estimates (local MLEs), the relative contribution of each individual parameter to the eigenparameters depends on the local MLE used to evaluate the Hessian matrix (Figure 6). As each of the likelihood maxima had distinct parameter values, this indicates that the sensitivity of the model-data fit to parameter combinations depends on the local parameter space considered. This motivates the need for a sensitivity matrix which captures the model-data fit sensitivity across a range of potential parameter values, such as the posterior covariance method.

14

Eigen-parameter	Local maximum likelihood estimate	λ_j/λ_1	3 x Diffusion Mechanisms			Ca-ATPase Pump Mechanism										BAT Pump Mechanism						
			k_{CO_2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1f_c}	k_{2f_c}	k_{3f_c}	k_{1b_c}	k_{2b_c}	k_{3b_c}	E_{0_b}	k_{1f_b}	k_{2f_b}	k_{3f_b}	k_{1b_b}	k_{2b_b}	k_{3b_b}
1	1	1.0000	0.0	-0.9	0.2	0.1	0.8	1.0	0.8	0.9	0.9	0.2	-0.9	-0.2	0.0	0.1	0.0	0.1	0.1	0.0	0.0	0.0
	2	1.0000	0.0	1.0	-0.1	-0.4	-0.4	-0.5	-0.7	-0.8	-0.8	-0.3	0.7	0.3	0.0	-0.3	0.0	-0.2	-0.2	0.0	0.0	0.0
	3	1.0000	0.0	0.5	-0.1	-0.1	-0.9	-0.2	-0.7	-1.0	-1.0	0.2	1.0	-0.2	0.0	-0.3	0.0	-0.3	-0.3	0.0	0.0	0.0
	4	1.0000	0.0	0.0	-0.1	0.0	-0.3	-0.1	-0.1	-0.1	-0.4	1.0	0.7	-1.0	0.0	0.4	0.0	0.0	0.1	0.0	0.0	0.0
	5	1.0000	0.0	1.0	-0.1	-0.1	-0.5	-0.2	-0.5	-0.5	-0.7	-0.6	0.9	0.6	0.0	-0.4	0.0	-0.2	-0.3	0.0	0.0	0.0
2	1	0.2674	0.0	1.0	0.4	-0.2	0.5	-0.5	0.5	0.2	0.2	0.6	-0.2	-0.6	0.0	-0.9	-0.1	-0.6	-0.6	0.0	0.0	0.0
	2	0.2825	0.0	-0.2	0.0	0.0	0.0	1.0	-0.3	-0.4	-0.4	0.2	0.3	-0.2	0.0	0.6	0.0	0.2	0.1	0.0	0.0	0.0
	3	0.8395	0.0	-0.2	-0.7	0.0	-0.3	0.2	-0.1	-0.2	-0.2	-0.9	0.2	0.9	-0.2	0.4	0.1	1.0	1.0	0.0	0.0	0.0
	4	0.8390	0.0	1.0	0.0	-0.2	0.0	-0.5	-0.1	-0.1	-0.6	-0.3	0.7	0.3	0.0	-0.7	0.0	-0.3	-0.3	0.0	0.0	0.0
	5	0.3980	0.0	0.1	0.4	0.1	-0.7	0.5	0.1	0.1	-0.6	0.9	1.0	-0.9	0.0	0.7	0.0	0.6	0.6	0.0	0.0	0.0
3	1	0.1390	0.0	-0.2	-0.2	0.1	-0.5	0.2	1.0	-0.4	-0.4	0.5	0.4	-0.5	0.0	0.7	0.0	-0.3	-0.3	0.0	0.0	0.0
	2	0.2357	0.0	0.3	0.1	0.2	0.3	-0.2	0.0	-0.1	-0.1	1.0	0.0	-0.9	0.0	-0.3	0.0	-0.2	-0.5	0.0	0.0	0.0
	3	0.4677	0.0	1.0	-0.2	-0.1	0.3	-0.8	0.0	0.1	0.1	-0.2	-0.1	0.2	0.0	-0.5	0.0	0.1	0.1	0.0	0.0	0.0
	4	0.7815	0.0	-0.3	0.0	0.1	-0.6	1.0	0.2	0.2	-0.4	-0.3	0.6	0.3	0.0	0.0	0.0	0.2	0.2	0.0	0.0	0.0
	5	0.3219	0.0	1.0	0.4	-0.2	0.8	-0.8	-0.4	-0.4	0.1	0.8	-0.4	-0.8	0.0	-0.2	0.0	0.1	0.0	0.0	0.0	0.0
4	1	0.1007	0.0	0.8	0.3	0.1	-0.1	0.1	-0.1	0.1	0.1	0.7	-0.1	-0.7	0.0	0.4	0.0	1.0	1.0	0.0	0.0	0.0
	2	0.1280	0.0	-0.1	-0.3	1.0	-0.8	0.5	0.0	0.0	0.0	-0.1	0.0	0.2	0.0	-0.5	0.0	-0.4	-0.4	0.0	0.0	0.0
	3	0.2657	0.0	0.4	0.2	0.0	0.0	-0.1	0.0	0.0	0.0	-0.4	0.0	0.4	0.3	1.0	-0.3	-0.4	-0.5	0.0	0.0	0.0
	4	0.5384	0.0	0.1	-0.1	0.0	1.0	0.8	-0.6	-0.6	-0.3	0.1	-0.1	-0.1	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0
	5	0.2351	0.0	0.2	0.0	0.0	-0.5	0.3	-0.4	-0.4	1.0	0.0	-0.1	0.0	0.0	0.2	0.0	0.1	0.1	0.0	0.0	0.0

Figure 6: Eigenvector element values for eigenparameters identified using the MLE Hessian approach to an analysis of sloppiness. These eigenparameters correspond to the four largest eigenvalues, and so are ordered from highest relative importance to lowest. For each eigenparameter, the results of five high likelihood parameter estimates are compared to indicate the consistency of results. Here we only consider four of twenty eigenparameters to show the inconsistency between results based on different local MLEs in the four most sensitive parameter combinations. For each eigenparameter, the values were normalised by the leading eigenvector value, such that they are rescaled to be between -1 and 1 inclusive. Here the colour darkens as the absolute values of the eigenvector values increase from 0 to 1, such that the model-data fit is more sensitive to darker eigenvector values. Notice that each of the parameters has been grouped based on its mechanistic function in the model (Figure 3). Additionally, the relative size of each eigenvalue when compared to the leading eigenvalue for each sample has been included in the column λ_j/λ_1.

3.3.2 Posterior covariance method

As an alternative to the MLE Hessian approach, the posterior covariance method considers a sample of the posterior distribution to obtain the key eigendirections (parameter interdependencies). For the analysis performed on one representative set of posterior samples (Figure 7), the two stiffest eigenparameters indicate that the coelenteron-ECM paracellular diffusion (represented as k_{pp}), and Ca-ATPase pump mechanisms strongly inform model predictions, in agreement with the results obtained for the Hessian matrix. This suggests that within this dataset, the calcification rate was strongly informed by the balance of these mechanisms within the coral polyp.

Unlike the Hessian-based results, the parameters associated with the BAT pump, seawater-coelenteron diffusion mechanism (represented as s) and the coelenteron-ECM transcellular diffusion (represented as k_{CO2}) mechanisms contribute little to the model behaviour within the seven stiffest eigenparameters. Given the decay in relative eigenvalue size beyond the seventh eigenparameter (more than two orders of magnitude, see Figure 5), the analysis suggests that these latter mechanisms may not be necessary to maintain a good model-data fit for calcification

rate predictions.

In our case study, the posterior covariance method yielded substantially more consistent results across different sets of posterior samples (Supplementary Material Section S.5) when compared to the MLE Hessian approach across different local MLEs (Figure 6). When using a posterior covariance method, the analysis yielded consistent eigenvector values for the stiffest nine eigenparameters in Figure 7 across independently generated sets of posterior samples (Figure S2 of Supplementary Material Section S.5). Here, differences amongst eigenparameters having small eigenvalues are expected as the model-data fit is weakly sensitive to sloppy eigendirections (Brown and Sethna, 2003).

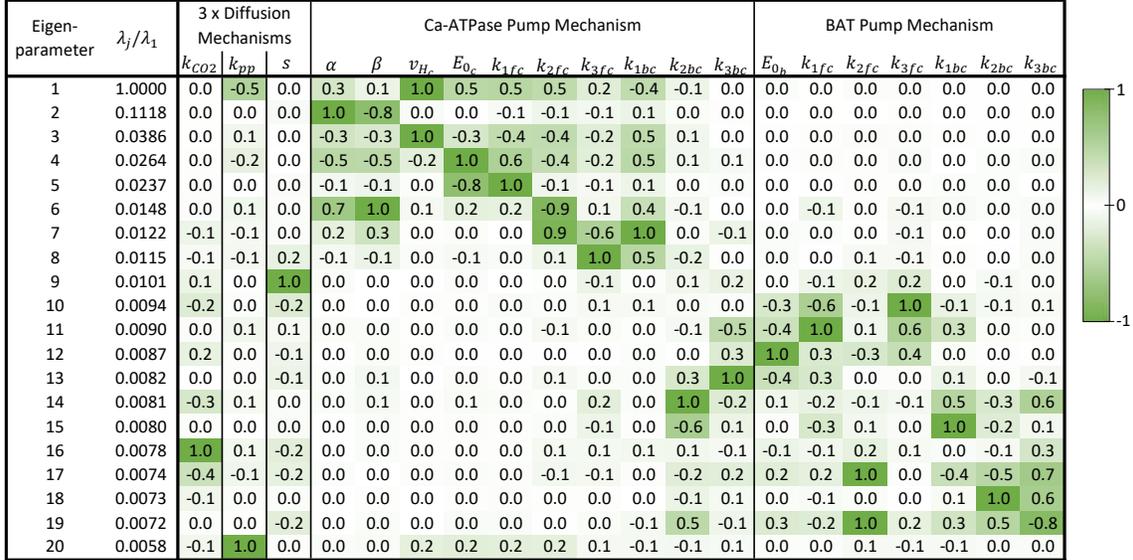| Eigen-parameter	λ_j/λ_1	3 x Diffusion Mechanisms			Ca-ATPase Pump Mechanism										BAT Pump Mechanism						
		k_{CO2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1fc}	k_{2fc}	k_{3fc}	k_{1bc}	k_{2bc}	k_{3bc}	E_{0_h}	k_{1fc}	k_{2fc}	k_{3fc}	k_{1bc}	k_{2bc}	k_{3bc}
1	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.5	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	0.1118	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	0.0386	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	0.0264	0.0	-0.2	0.0	-0.5	-0.5	-0.2	1.0	0.6	-0.4	-0.2	0.5	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
5	0.0237	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.8	1.0	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
6	0.0148	0.0	0.1	0.0	0.7	1.0	0.1	0.2	0.2	-0.9	0.1	0.4	-0.1	0.0	0.0	-0.1	0.0	-0.1	0.0	0.0	0.0
7	0.0122	-0.1	-0.1	0.0	0.2	0.3	0.0	0.0	0.0	0.9	-0.6	1.0	0.0	-0.1	0.0	0.0	0.0	-0.1	0.0	0.0	0.0
8	0.0115	-0.1	-0.1	0.2	-0.1	-0.1	0.0	-0.1	0.0	0.1	1.0	0.5	-0.2	0.0	0.0	0.0	0.1	-0.1	0.0	0.0	0.0
9	0.0101	0.1	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.1	0.2	0.0	-0.1	0.2	0.2	0.0	-0.1	0.0
10	0.0094	-0.2	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	0.0	-0.3	-0.6	-0.1	1.0	-0.1	-0.1	0.1
11	0.0090	0.0	0.1	0.1	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.0	-0.1	-0.5	-0.4	1.0	0.1	0.6	0.3	0.0	0.0
12	0.0087	0.2	0.0	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	1.0	0.3	-0.3	0.4	0.0	0.0	0.0
13	0.0082	0.0	0.0	-0.1	0.0	0.1	0.0	0.0	0.0	0.1	0.0	0.0	0.3	1.0	-0.4	0.3	0.0	0.0	0.1	0.0	-0.1
14	0.0081	-0.3	0.1	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.2	0.0	1.0	-0.2	0.1	-0.2	-0.1	-0.1	0.5	-0.3	0.6
15	0.0080	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	-0.6	0.1	0.0	-0.3	0.1	0.0	1.0	-0.2	0.1
16	0.0078	1.0	0.1	-0.2	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1	-0.1	-0.1	-0.1	0.2	0.1	0.0	-0.1	0.3
17	0.0074	-0.4	-0.1	-0.2	0.0	0.0	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.2	0.2	0.2	0.2	1.0	0.0	-0.4	-0.5	0.7
18	0.0073	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.1	0.0	-0.1	0.0	0.0	0.1	1.0	0.6
19	0.0072	0.0	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.5	-0.1	0.3	-0.2	1.0	0.2	0.3	0.5	-0.8
20	0.0058	-0.1	1.0	0.0	0.0	0.0	0.2	0.2	0.2	0.2	0.1	-0.1	-0.1	0.1	0.0	0.0	0.1	-0.1	-0.1	0.0	0.0

Figure 7: Eigenvector values for eigenparameters identified using a posterior covariance analysis of sloppiness, based on one set of posterior samples. Here we show the full analysis of model sloppiness, i.e. all 20 eigenparameters. (See Supplementary Material Section S.5 for confirmation of the consistency across different sets of posterior samples.) See the caption of Figure 6 for further interpretation of this figure.

3.3.3 Likelihood informed subspace sensitivity matrix

Lastly, the analysis of model sloppiness using a LIS method was considered. One advantage of using both a posterior covariance and LIS approach to analysing sloppiness is that the influence of the prior distribution on the posterior sample can be identified (Monsalve-Bravo et al., 2022). More specifically, if there are substantial differences in the directions of stiff eigenparameters between the posterior covariance and LIS methods, this indicates that the prior is having a large impact on the stiff eigenparameter directions found using posterior covariance, and therefore the prior itself is strongly informing the posterior. For calculation of the LIS matrix, finite differences was also used to approximate the Hessian, with the step-size $\delta = 10^{-2}$.

Looking at the two stiffest eigenparameters (Figure 8), the LIS analysis of sloppiness gives a similar result to that of the posterior covariance method – the relative magnitude of the elements within the two eigenvectors are very similar between methods. However, beyond the second eigenparameter, the results of the two approaches begin to differ and the LIS results be-

16

come inconsistent across independent posterior samples (Supplementary Material Section S.6). Here we note that the relative decay in eigenvalue size is much more rapid for an analysis of sloppiness using the LIS approach (blue triangles in Figure 5), when compared to the posterior covariance method (green squares in Figure 5). For example, the relative eigenvalue size of the tenth posterior covariance eigenvalue is of similar magnitude to that of the third LIS eigenvalue (Figure 5). Additionally, we previously saw that the eigenvector values became inconsistent when $\lambda_j/\lambda_1 \sim \mathcal{O}(10^{-2})$ using the posterior covariance method (Supplementary Material Section S.5, tenth eigenparameter), so it is not unreasonable to expect inconsistency between samples following the second eigenparameter using the LIS method, when this order of magnitude approximation is similar. Such inconsistencies are expected when eigenvalues are very small because the model-data fit is less sensitive to the corresponding sloppy eigenparameters, and these sloppy eigenparameters are therefore difficult to uniquely identify. Hence, the results of the analysis of model sloppiness used here appears to be consistent across both the LIS and posterior covariance approaches.

A similar result between the posterior covariance and LIS methods for the stiffest eigenparameters indicates that the prior distribution used for the application of Bayesian inference does not substantially influence on the shape of the posterior distribution. Hence, in our case we concluded that the prior used was weakly informative (as intended) and had little influence on the posterior distribution.

Eigen-parameter	Sensitivity Matrix	λ_j/λ_1	3 x Diffusion Mechanisms			Ca-ATPase Pump Mechanism										BAT Pump Mechanism						
			k_{CO_2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1f_c}	k_{2f_c}	k_{3f_c}	k_{1b_c}	k_{2b_c}	k_{3b_c}	E_{0_b}	k_{1f_b}	k_{2f_b}	k_{3f_b}	k_{1b_b}	k_{2b_b}	k_{3b_b}
1	S_P	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.5	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
2	S_P	0.1118	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	0.0243	-0.1	0.1	0.0	-0.9	1.0	-0.1	-0.1	0.1	0.2	0.1	-0.2	-0.2	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	S_P	0.0386	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	0.0133	-0.1	0.0	0.0	0.4	-0.1	-0.7	-0.2	0.0	0.3	1.0	-0.3	-1.0	-0.3	0.5	0.1	0.3	0.3	0.0	0.0	0.0

Figure 8: Comparison of eigenvector values of the three stiffest eigenparameters of the analysis of model sloppiness from LIS (S_L) and posterior covariance (S_P) sensitivity matrices (blue and green respectively). Here we only consider three of twenty eigenparameters to compare the consistency between approaches for the three most sensitive parameter combinations. See the caption of Figure 6 for further interpretation of this figure.

3.4 Model reduction for the coral calcification model

The analysis of model sloppiness revealed the parameter combinations which most and least inform model predictions for the coral calcification case study – a useful tool for identifying potential model simplifications that have little impact on model predictions for the given dataset. Each analysis of model sloppiness (Section 3.3) indicated that the coelenteron-ECM paracellular diffusion (characterised by parameter k_{pp}) and the Ca-ATPase pump mechanisms contribute substantially to the stiffest eigenparameter (Figures 6–8), suggesting that these two mechanisms (characterised by ten parameters, see Table 2) have a strong influence on model behaviours. However, the analysis indicates that model predictions are relatively insensitive to the seawater-coelenteron diffusion mechanism (characterised by parameter s), the BAT pump mechanism (characterised by seven parameters, see Table 2), and the coelenteron-ECM transcellular diffusion mechanism (characterised by parameter k_{CO2}) because they did not contribute to the stiff eigenparameters (see the first eight parameter combinations in Figure 7). Hence, the results suggest that each of these three mechanisms, or a combination of them could be removed.

17

However, the seawater-coelenteron diffusion mechanism is an integral part of the model, as it is the primary mechanism by which ocean acidification affects coral calcification rates in the model. Since this model's original purpose in Galli and Solidoro (2018) was to identify the way that environmental factors affect calcification rates (including but not limited to ocean acidification), this result could suggest that the experimental dataset did not sufficiently capture the effects of ocean acidification on calcification rates. Regardless of this issue, removing the diffusion between the coral polyp and external seawater seems nonsensical as it would yield a model whereby there is no interaction between the local carbon chemistry of seawater and the coral host animal (except indirectly through effects on net photosynthesis which is unlikely; see, e.g. Kroeker et al. 2013).

It was biologically unreasonable to remove the seawater-coelenteron diffusion mechanism, so the potential simplified models could include one without the BAT pump, one without the coelenteron-ECM transcellular diffusion, and one without both mechanisms (see Figure 3 for a conceptual depiction of these mechanisms). All three of these reduced models were investigated by removing one or both of the selected mechanisms from the ordinary differential equations. The reduced models were recalibrated via Bayesian inference using the same dataset (see Table 1 of Galli and Solidoro, 2018) and likelihood function (Equation 1), and for the remaining parameters the same independent uniform distributions were used as the prior distribution (Figure 4 and Table S1). Removing both the BAT pump and coelenteron-ECM transcellular diffusion mechanisms from the model reduces the number of parameters from 20 to 12.

The reduced models suggested by the analysis of model sloppiness yielded very similar predictions of the coral calcification rate data when compared to the original model, despite having up to eight parameters removed. Visually, the goodness-of-fit between the model and data was similar for the original and our proposed reduced model without the BAT pump and coelenteron-ECM transcellular diffusion mechanisms (Figure 9) – and a similar result is observed when only one of the two mechanisms were removed (Supplementary Material Section S.7). The estimated model evidence quantitatively suggests that the original model and its three proposed reductions are similarly supported by the data, and the estimated Bayes factors (Kass and Raftery, 1995) do not indicate a strong preference between models (Supplementary Material Section S.8).

Additionally, removing the two insensitive mechanisms (BAT pump and coelenteron-ECM transcellular diffusion) from the model did not lead to clear differences in the estimated marginal posterior densities for each parameter, or to the analysis of model sloppiness between the original model and the reduced model (Supplementary Material Section S.9). Whilst there is limited sacrifice in predictive ability, there was a significant gain in computation time – the reduced model required less than half the computation time of the original model for calibration (13.2 hours for the original model and 6.0 hours for the reduced model, using a high-performance workstation with 12 cores).

For comparison, we also removed a mechanism from the model that was represented within the stiffest eigenparameter and therefore very sensitive to the model-data fit. Removing this sensitive mechanism resulted in a much worse model-data fit, both visually (Supplementary Material Section S.7) and quantitatively (Supplementary Material Section S.8). For this case study, these results indicate that analysing the model sloppiness is an appropriate way to inform model reductions for maintaining a good fit between the model and calibration dataset.
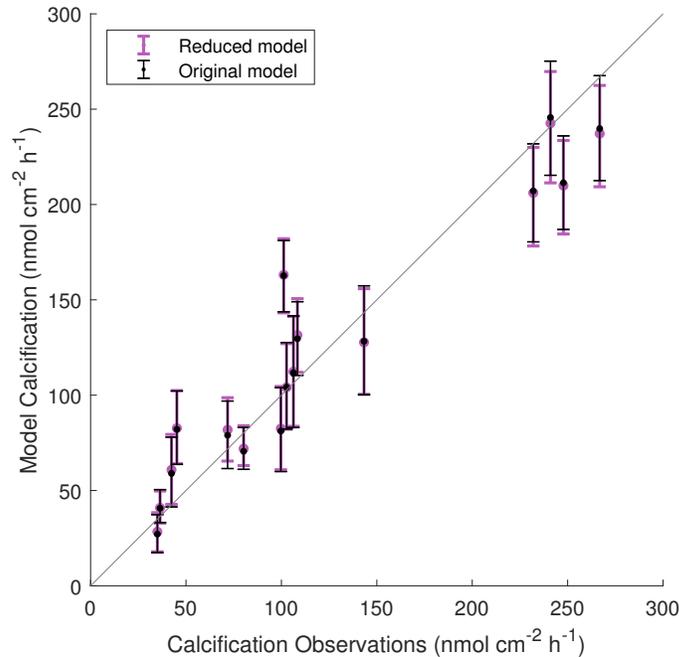
Figure 9: Posterior predictive distribution showing observed calcification compared to modelled calcification for both the original and reduced model. The reduced model is one without both the BAT pump mechanism and coelenteron-ECM transcellular diffusion. The dots indicate the median prediction value, and error bars indicate the 95% central credible interval for calcification predictions from the respective posterior samples of each model.

4 Discussion and Conclusion

In this work, we have proposed and demonstrated a new method for simplifying models based on the analysis of model sloppiness. This analysis can identify the informativeness of parameter combinations on model behaviours by analysing the topology of the surface which describes the fit of the model to data in parameter space (Monsalve-Bravo et al., 2022). As such, it can identify insensitive model mechanisms whose parameters contribute little to the model's ability to fit the available data, whilst accounting for parameter interdependencies. We showed that identifying and removing such insensitive mechanisms can be used for model reduction whilst minimising the loss of predictive capability.

4.1 Coral calcification case study

Our model reduction method, informed by the analysis of model sloppiness, was demonstrated on a case study on a model of coral calcification reaction rates (Galli and Solidoro, 2018). In this case study, the data was not sufficient to provide narrowly constrained estimates for most of the model parameters, but it was also not immediately clear how the model could be simplified. To address this issue, the proposed analysis identified two mechanisms within the model which only weakly informed predictions and were sensible to remove; removing these processes reduced the number of model parameters from 20 to 12. A comparison of the goodness-of-fit for both the

19

original and simplified models indicated that similar predictions were produced by both (Figure 9). It should be noted here that the simplified model has, to our best knowledge, not been considered in the literature for modelling coral calcification, so the outcome from analysis of sloppiness yields a new model for practitioners to consider.

For the present case study, the analysis of model sloppiness may also be useful for better understanding the physiology of the coral polyp. The analysis indicated that the model-data fit was highly sensitive primarily (i.e. first eigenparameter) to the Ca-ATPase pump (through parameters $v_{H_c}, E0_c, k_{1f_c}, k_{2f_c}$ and k_{1b_c}) relative to the paracellular diffusion of chemical species (through parameter k_{pp}) between the coral polyp compartments that the pump mechanism connects – the coelenteron and ECM (see Figure 3). This may be expected since it is the balance of the active and passive fluxes between the ECM and coelenteron at steady state that determines the aragonite saturation state of the ECM, which itself directly controls the calcification rate. Secondarily (i.e. second eigenparameter), the model-data fit was sensitive to the ATP availability from gross photosynthesis and respiration (through parameters α and β). In the model this process determines the amount of ATP available to fuel the Ca-ATPase pump responsible for the active flux.

From our results, we also gain insights from the mechanisms which are removed from the model. Given that the analysis of model sloppiness is closely related to the concept of parameter identifiability (Bellman and Åström, 1970; Browning et al., 2020; Chis et al., 2016; Raue et al., 2009), it can inform the structural and/or practical identifiability of parameters. In the explored coral case study, the removed mechanisms had similar functions to other mechanisms in the model, such that these processes are partially structurally unidentifiable. The BAT pump mechanism is an active transport for bicarbonate anions (HCO_3^-), and so the mechanism alters the dissolved inorganic carbon (DIC) levels between the coelenteron and ECM in a similar way to the Ca-ATPase pump. Additionally, the model includes two passive transport mechanisms between the coelenteron and ECM (see Figure 3): a paracellular pathway for all chemical species and a transcellular pathway for carbon dioxide, which have very similar functions. Hence, the analysis of model sloppiness here indicates that certain components of the model are not necessary for a good model-data fit because their functions are partially or fully replicated through other mechanisms. So, when considering the aggregate behaviour of these processes within the coral polyp, including these "redundant" physiological processes may not lead to better predictions because of structural and/or practical identifiability issues, and instead may cause poor parameter estimation (Raue et al., 2009).

Importantly, any physiological conclusions drawn from the results of this analysis are in relation to this dataset alone. The dataset used for this analysis is small, given that the number of data points is of the same order of magnitude as the number of model parameters. Additionally, any misspecification or inaccurate assumptions in this model are carried through to the reduced model. That being said, all models are unavoidably approximate representations of the real world. Hence, any model mechanisms excluded through model reduction cannot be considered unimportant for coral reef physiology, as the analysis indicates that model *predictions* of calcification rates for this dataset are negligibly affected by exclusion of these mechanisms. We also acknowledge that these two criticisms are common to data-driven methods which derive simplifications from an originally more complicated model (Crout et al., 2009).

The explored coral calcification case study demonstrated the potential for an analysis of model sloppiness to inform model reductions. However, this method could be used generally for strategically proposing simpler models. The method is restricted to parametric models where an appropriate sensitivity matrix can be defined, and is best suited to models where mechanisms can be easily removed (e.g. process-based models). However, future work could examine how this

model reduction technique performs on various other types of models. For instance, Monsalve-Bravo et al. (2022) describes how an analysis of model sloppiness could be used to identify critical parameter combinations in a stochastic setting, and this idea could potentially be explored for application to model reduction.

4.2 Sensitivity matrix selection

In this paper, we compared the results of an analysis of model sloppiness using three different approaches to constructing the sensitivity matrix (Table 1). For the explored case study, all approaches – the Hessian evaluated at MLE, the posterior covariance method and the LIS method – agreed on which mechanisms strongly inform the model-data fit (Figures 6–8), so the model reduction informed by each approach was equivalent in this case. However, this conclusion will not always hold for other model-data fitting problems (e.g., see Monsalve-Bravo et al. 2022). Just as each approach to capturing the model-data sensitivities differ, so might the model reductions informed by this analysis. The results from our case study lead to two key general findings regarding selection of a sensitivity matrix.

Firstly, this case study demonstrated that the non-Bayesian sensitivity matrix may have limited utility when the likelihood surface is rugged. In the coral calcification case study, the likelihood surface had no well-defined peak in parameter space, so single parameter estimates of this peak based on MLE could not reproducibly encapsulate the model-data fit sensitivity. Results based on the local sensitivity of the likelihood surface may not capture important features of the model-data fit in such circumstances, and should be interpreted with caution. For complex models, the global likelihood surface including the full range of parameter values represented by the posterior distribution should instead be considered (e.g. using the posterior covariance approach).

Though the Hessian sensitivity matrix may become problematic for some complex models, that does not mean that the MLE Hessian approach cannot be used to gain useful information for model reduction. Rather, a non-Bayesian method of analysing sloppiness is easier to implement and far more computationally efficient. Thus, a non-Bayesian method of analysing sloppiness – such as the Hessian matrix evaluated at MLE, or the Levenberg-Marquardt approximation (Marquardt, 1963) of the Hessian (used for computationally intensive models) – could provide a simpler, faster method of suggesting model reductions, in place of the Bayesian counterparts. However, where the likelihood surface is expected to be complicated, sufficient consideration should be given to the choice of optimisation algorithm used to identify the best-fit parameter values.

Secondly, the case study demonstrated the benefits of using the posterior covariance and LIS methods together. In the explored coral calcification model, both the posterior covariance and LIS results for analysing sloppiness were similar, indicating that the prior distribution was not having a substantial influence on the model-data calibration process (i.e. on the posterior distribution). Hence, this case study demonstrates the value of using multiple methods for analysing sloppiness, as together these methods provide richer information than each method by itself (Monsalve-Bravo et al., 2022).

However, because in our case study the prior distribution was weakly informative on the posterior, it is difficult to state what the outcomes of informing model reductions via a LIS analysis would show if informative priors were instead used. Although we leave this exploration for future work, we hypothesise that model reduction informed by LIS could yield reduced models based purely on retaining mechanisms for which the data is highly informative relative to prior beliefs

for each mechanism. At the very least, the LIS method is a useful check to identify the influence of the data on the model calibration process within a Bayesian framework.

While we only considered three sensitivity matrices, there are various approaches in the literature that could be used within this model reduction framework. Methods such as the Levenberg-Marquardt Hessian, a likelihood-free approximation of LIS, or the active subspace method each capture the model-data fit differently to the methods considered here and may provide advantages for different applications – such as computationally intensive or stochastic models (Monsalve-Bravo et al., 2022).

The active subspace method (Constantine et al., 2016) is a dimension reduction approach that constructs a sensitivity matrix based on the gradient of the log-likelihood relative to the prior distribution, capturing informativeness of the data relative to the prior on the model-data fit similarly (but not exactly) as the LIS approach does (Zahm et al., 2022). Constantine et al. (2016) define the active/inactive subspace of eigenvectors via eigendecomposition of the sensitivity matrix and identifying the first largest gap between eigenvalues. As noted by Monsalve-Bravo et al. (2022), the active subspace sensitivity matrix could be used to assess model sloppiness and we suggest it can also be used within our model reduction framework. In addition, Constantine and Diaz (2017) proposed activity score metrics to rank individual model parameters based on the eigenvalues and eigenvectors of the active subspace matrix and can be used for model reduction. We leave explorations of active subspace for model reduction and a comparison with model reductions produced by the sensitivity matrices used in this manuscript for future research.

4.3   Complexity versus parsimony in models

The processes, phenomena and systems of the world around us are extremely complex; so should the models we create to represent these ideas be equally as complex? The simple answer is that it depends on the purpose of the model, whether that be accurate predictions of collective behaviour (Dietze et al., 2018), inference for physically meaningful parameters (Adams et al., 2017), or analysis for understanding and/or changing processes within a system (e.g. Galli and Solidoro 2018, Verspagen et al. 2014).

If the aim of the model is accurate prediction, it is the aggregate behaviour of a system that we aim to recreate through models, and in many cases, including additional underlying processes is not beneficial in modelling the collective behaviours (Machta et al., 2013; Transtrum and Qiu, 2014). Including many potential mechanisms within a model may mean that some processes are being fit to noise in the data, leading to poor predictions in new situations (Cox et al., 2006). However, a general model may be too simple for accurate prediction and could distort the importance of processes in the model (Lawrie and Hearne, 2007; Van Nes and Scheffer, 2005). Models should minimise the bias-variance tradeoff, balancing the complexity so that the model is simple enough to predict new data well and complex enough to capture features of the data (Geman et al., 1992).

What if accurate parameter inference for meaningful parameters is the goal? The more complex a model is, the more difficult parameterisation becomes (Van Nes and Scheffer, 2005), and more data is required as a consequence. If unidentifiable parameters are included in the model calibration problem, the values of these parameters become meaningless as they are often interdependent on others (Van Werkhoven et al., 2009). On the other hand, removing important and existing processes from a model could mean that calibrated parameters lose their physical meaning as they compensate for processes missing from the model and become more like aggregate parameters for modelling the collective model dynamics (Elevitch and Johnson, 2020).

If there are specific processes within the model that we wish to understand or change within a system, the relevant process needs to be included for analysis (Hannah et al., 2010). However, even in this case the modeller must still consider whether the parameters retain their physical meaning (due to structural identifiability issues or data limitations) as well as the potential for overfitting caused by the inclusion of the process (e.g. Jakeman and Hornberger, 1993). Additionally, there are other issues that come with complex models, such as difficulty implementing, reproducing, interpreting, validating and communicating the models, as well as computation times and difficulty or costs associated with updating or replacing these models (Van Nes and Scheffer, 2005).

So where is the line between too simple and too complex? There are many arguments both for and against complexity (e.g. Anderson, 1972; Hong et al., 2017; Hunt et al., 2007; Logan, 1994; Wigner, 1990). The desired complexity of a model should be based on both the model's purpose and statistical assessments of quality (Saltelli, 2019). Our goal in the present work is to highlight that, in some circumstances, model reduction could benefit the predictions and parameter estimates of a model. Our model reduction framework provides a principled and intuitive approach for model simplifications to address this goal.

Software and data availability

The code used for this analysis was implemented in MATLAB (R2021b), and is freely available for download on Figshare at https://doi.org/10.6084/m9.figshare.19529626.v2. This code (14.8MB) contains 37 functions as MATLAB code files, which runs the analyses for the coral calcification case study and produces the corresponding plots described within this manuscript. In addition, 7 MATLAB data files are included, which contain the dataset used for analyses (available to access through Galli and Solidoro (2018)), and generated by the analysis (including SMC samples for each model, and for independent runs, local MLE samples, and calculated sensitivity matrices).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

of the manuscript; SAV, CD and MPA contributed to the design of the research; all authors contributed to the coding, analysis of results, review and editing of the manuscript.

References

Adams, M. P., C. J. Collier, S. Uthicke, Y. X. Ow, L. Langlois, and K. R. O'Brien (2017). Model fit versus biological relevance: evaluating photosynthesis-temperature models for three tropical seagrass species. *Scientific Reports 7*(1), 1–12.

Anderson, P. W. (1972). More is different: broken symmetry and the nature of the hierarchical structure of science. *Science 177*(4047), 393–396.

Andersson, A. J. and D. Gledhill (2013). Ocean acidification and coral reefs: Effects on breakdown, dissolution, and net ecosystem calcification. *Annual Review of Marine Science 5*, 321–348.

Beers, K. J. (2007). *Numerical Methods for Chemical Engineering: Applications in MATLAB*. Cambridge: Cambridge University Press.

Bellman, R. and K. J. Åström (1970). On structural identifiability. *Mathematical Biosciences 7*(3-4), 329–339.

Beskos, A., A. Jasra, K. Law, Y. Marzouk, and Y. Zhou (2018). Multilevel sequential Monte Carlo with dimension-independent likelihood-informed proposals. *SIAM/ASA Journal on Uncertainty Quantification 6*(2), 762–786.

Brown, K. S., C. C. Hill, G. A. Calero, C. R. Myers, K. H. Lee, J. P. Sethna, and R. A. Cerione (2004). The statistical mechanics of complex signaling networks: nerve growth factor signaling. *Physical Biology 1*(3), 184.

Brown, K. S. and J. P. Sethna (2003). Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E 68*(2), 021904.

Browning, A. P., D. J. Warne, K. Burrage, R. E. Baker, and M. J. Simpson (2020). Identifiability analysis for stochastic differential equation models in systems biology. *Journal of the Royal Society Interface 17*(173), 20200652.

Chis, O.-T., A. F. Villaverde, J. R. Banga, and E. Balsa-Canto (2016). On the relationship between sloppiness and identifiability. *Mathematical Biosciences 282*, 147–161.

Chopin, N. (2002). A sequential particle filter method for static models. *Biometrika 89*(3), 539–552.

Constantine, P. G. and P. Diaz (2017). Global sensitivity metrics from active subspaces. *Reliability Engineering & System Safety 162*, 1–13.

Constantine, P. G., C. Kent, and T. Bui-Thanh (2016). Accelerating Markov chain Monte Carlo with active subspaces. *SIAM Journal on Scientific Computing 38*(5), A2779–A2805.

Cox, G., J. Gibbons, A. Wood, J. Craigon, S. Ramsden, and N. Crout (2006). Towards the systematic simplification of mechanistic models. *Ecological Modelling 198*(1-2), 240–246.

Crout, N. M., D. Tarsitano, and A. T. Wood (2009). Is my model too complex? evaluating model formulation using model reduction. *Environmental Modelling & Software 24*(1), 1–7.

Cui, T., J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini (2014). Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems 30*(11), 114015.

Del Moral, P., A. Doucet, and A. Jasra (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68*(3), 411–436.

Dietze, M. C., A. Fox, L. M. Beck-Johnson, J. L. Betancourt, M. B. Hooten, C. S. Jarnevich, T. H. Keitt, M. A. Kenney, C. M. Laney, L. G. Larsen, H. W. Loescher, C. K. Lunch, B. C. Pijanowski, J. T. Randerson, E. K. Read, A. T. Tredennick, R. Vargas, K. C. Weathers, and E. P. White (2018). Iterative near-term ecological forecasting: Needs, opportunities, and challenges. *Proceedings of the National Academy of Sciences 115*(7), 1424–1432.

Drovandi, C. C. and A. N. Pettitt (2011). Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics 67*(1), 225–233.

Elevitch, C. R. and C. R. J. Johnson (2020). A procedure for ranking parameter importance for estimation in predictive mechanistic models. *Ecological Modelling 419*, 108948.

Erez, J., S. Reynaud, J. Silverman, K. Schneider, and D. Allemand (2011). Coral calcification under ocean acidification and global change. In *Coral reefs: an ecosystem in transition*, pp. 151–176. Springer.

Galli, G. and C. Solidoro (2018). ATP supply may contribute to light-enhanced calcification in corals more than abiotic mechanisms. *Frontiers in Marine Science 5*, 68.

Geman, S., E. Bienenstock, and R. Doursat (1992). Neural networks and the bias/variance dilemma. *Neural Computation 4*(1), 1–58.

Geris, L. and D. Gomez-Cabrero (2016). An introduction to uncertainty in the development of computational models of biological processes. In *Uncertainty in Biology*, pp. 3–11. Springer.

Getz, W. M., C. R. Marshall, C. J. Carlson, L. Giuggioli, S. J. Ryan, S. S. Romañach, C. Boettiger, S. D. Chamberlain, L. Larsen, P. D'Odorico, et al. (2018). Making ecological models adequate. *Ecology letters 21*(2), 153–166.

Gibbons, J., A. Wood, J. Craigon, S. Ramsden, and N. Crout (2010). Semi-automatic reduction and upscaling of large models: A farm management example. *Ecological Modelling 221*(4), 590–598.

Girolami, M. (2008). Bayesian inference for differential equations. *Theoretical Computer Science 408*, 4 – 16.

Gutenkunst, R. N., J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna (2007). Universally sloppy parameter sensitivities in systems biology models. *PLoS Computational Biology 3*(10), 1871–1878.

Hannah, C., A. Vezina, and M. S. John (2010). The case for marine ecosystem models of intermediate complexity. *Progress in Oceanography 84*(1-2), 121–128.

Hjelkrem, A.-G. R., M. Höglind, M. Van Oijen, J. Schellberg, T. Gaiser, and F. Ewert (2017). Sensitivity analysis and Bayesian calibration for testing robustness of the BASGRA model in different environments. *Ecological Modelling 359*, 80–91.

Hoegh-Guldberg, O., E. S. Poloczanska, W. Skirving, and S. Dove (2017). Coral reef ecosystems under climate change and ocean acidification. *Frontiers in Marine Science 4*, 158.

Hong, E.-M., Y. A. Pachepsky, G. Whelan, and T. Nicholson (2017). Simpler models in environmental studies and predictions. *Critical Reviews in Environmental Science and Technology 47*(18), 1669–1712.

Hsieh, N.-H., B. Reisfeld, F. Y. Bois, and W. A. Chiu (2018). Applying a global sensitivity analysis workflow to improve the computational efficiencies in physiologically-based pharmacokinetic modeling. *Frontiers in Pharmacology 9*, 588.

Huang, H., M. Fairweather, J. Griffiths, A. Tomlin, and R. Brad (2005). A systematic lumping approach for the reduction of comprehensive kinetic models. *Proceedings of the Combustion Institute 30*(1), 1309–1316.

Hunt, R. J., J. Doherty, and M. J. Tonkin (2007). Are models too simple? arguments for increased parameterization. *Groundwater 45*(3), 254–262.

Jackson, L. J., A. S. Trebitz, and K. L. Cottingham (2000). An introduction to the practice of ecological modeling. *BioScience 50*(8), 694–706.

Jakeman, A. and G. Hornberger (1993). How much complexity is warranted in a rainfall-runoff model? *Water Resources Research 29*(8), 2637–2649.

Jakeman, A. J., R. A. Letcher, and J. P. Norton (2006). Ten iterative steps in development and evaluation of environmental models. *Environmental Modelling & Software 21*(5), 602–614.

Jeong, J. E., Q. Zhuang, M. K. Transtrum, E. Zhou, and P. Qiu (2018). Experimental design and model reduction in systems biology. *Quantitative Biology 6*(4), 287–306.

Jeremiah, E., S. A. Sisson, A. Sharma, and L. Marshall (2012). Efficient hydrological model parameter optimization with Sequential Monte Carlo sampling. *Environmental Modelling & Software 38*, 283–295.

Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association 90*(430), 773–795.

Kroeker, K. J., R. L. Kordas, R. Crim, I. E. Hendriks, L. Ramajo, G. S. Singh, C. M. Duarte, and J.-P. Gattuso (2013). Impacts of ocean acidification on marine organisms: quantifying sensitivities and interaction with warming. *Global Change Biology 19*(6), 1884–1896.

Lawrie, J. and J. Hearne (2007). Reducing model complexity via output sensitivity. *Ecological Modelling 207*(2-4), 137–144.

Liao, J. C. and E. N. J. Lightfoot (1988). Lumping analysis of biochemical reaction systems with time scale separation. *Biotechnology and Bioengineering 31*(8), 869–879.

Logan, J. (1994). In defense of big ugly models. *American Entomologist 40*(4), 202–207.

Machta, B. B., R. Chachra, M. K. Transtrum, and J. P. Sethna (2013). Parameter space compression underlies emergent theories and predictive models. *Science 342*(6158), 604–607.

Mara, T. A., B. Belfort, V. Fontaine, and A. Younes (2017). Addressing factors fixing setting from given data: A comparison of different methods. *Environmental Modelling & Software 87*, 29–38.

Marino, S., I. B. Hogue, C. J. Ray, and D. E. Kirschner (2008). A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of theoretical biology 254*(1), 178–196.

Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics 11*(2), 431–441.

MathWorks (2021). fmincon. https://au.mathworks.com/help/optim/ug/fmincon.html.

Monsalve-Bravo, G. M., B. A. J. Lawson, C. Drovandi, K. Burrage, K. S. Brown, C. M. Baker, S. A. Vollert, K. Mengersen, E. McDonald-Madden, and M. P. Adams (2022). Analysis of sloppiness in model simulations: Unveiling parameter uncertainty when mathematical models are fitted to data. *Science Advances 8*(38), eabm5952.

Pepiot, P., L. Cai, and H. Pitsch (2019). Model reduction and lumping procedures. In *Computer Aided Chemical Engineering*, Volume 45, pp. 799–827. Elsevier.

Raue, A., C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer (2009). Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics 25*(15), 1923–1929.

Rodolfo-Metalpa, R., S. Martin, C. Ferrier-Pagès, and J.-P. Gattuso (2010). Response of the temperate coral *Cladocora caespitosa* to mid-and long-term exposure to pco_2 and temperature levels projected for the year 2100 AD. *Biogeosciences 7*, 289–300.

Saltelli, A. (2019). A short comment on statistical versus mathematical modelling. *Nature Communications 10*(1), 1–3.

Saltelli, A., T. Andres, and T. Homma (1993). Sensitivity analysis of model output: an investigation of new techniques. *Computational Statistics & Data Analysis 15*(2), 211–238.

Saltelli, A., S. Tarantola, F. Campolongo, and M. Ratto (2004). *Sensitivity analysis in practice: a guide to assessing scientific models*, Volume 1. Wiley Online Library.

Schilders, W. H., H. A. Van der Vorst, and J. Rommes (2008). *Model order reduction: theory, research aspects and applications*, Volume 13. Springer.

Snowden, T. J., P. H. van der Graaf, and M. J. Tindall (2017). Methods of model reduction for large-scale biological systems: a survey of current methods and trends. *Bulletin of Mathematical Biology 79*(7), 1449–1486.

Sobol, I. M. (1993). Sensitivity analysis for non-linear mathematical models. *Mathematical Modelling and Computational Experiment 1*, 407–414.

Transtrum, M. K., B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna (2015). Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics 143*(1), 010901–1–13.

Transtrum, M. K., B. B. Machta, and J. P. Sethna (2011). Geometry of nonlinear least squares with applications to sloppy models and optimization. *Physical Review E 83*(3), 036701.

Transtrum, M. K. and P. Qiu (2014). Model reduction by manifold boundaries. *Physical Review Letters 113*(9), 098701.

Tredennick, A. T., G. Hooker, S. P. Ellner, and P. B. Adler (2021). A practical guide to selecting models for exploration, inference, and prediction in ecology. *Ecology 102*(6), e03336.

Uusitalo, L., A. Lehikoinen, I. Helle, and K. Myrberg (2015). An overview of methods to evaluate uncertainty of deterministic models in decision support. *Environmental Modelling & Software 63*, 24–31.

Van Nes, E. H. and M. Scheffer (2005). A strategy to improve the contribution of complex simulation models to ecological theory. *Ecological Modelling 185*(2-4), 153–164.

Van Werkhoven, K., T. Wagener, P. Reed, and Y. Tang (2009). Sensitivity-guided reduction of parametric dimensionality for multi-objective calibration of watershed models. *Advances in Water Resources 32*(8), 1154–1169.

Verspagen, J. M., D. B. Van de Waal, J. F. Finke, P. M. Visser, and J. Huisman (2014). Contrasting effects of rising CO_2 on primary production and ecological stoichiometry at different nutrient levels. *Ecology Letters 17*(8), 951–960.

Wigner, E. P. (1990). The unreasonable effectiveness of mathematics in the natural sciences. In *Mathematics and Science*, pp. 291–306. World Scientific.

Zahm, O., T. Cui, K. Law, A. Spantini, and Y. Marzouk (2022). Certified dimension reduction in nonlinear Bayesian inverse problems. *Mathematics of Computation 91*(336), 1789–1835.

<div align="center">

Supplementary Material

Strategic model reduction by analysing model sloppiness: a case study in coral calcification

Sarah A. Vollert[a,b,c,*], Christopher Drovandi[a,b,c], Gloria M. Monsalve-Bravo[d,e,f] &
Matthew P. Adams[a,b,c,d]

[a]Centre for Data Science, Queensland University of Technology, Brisbane, QLD 4001, Australia

[b]ARC Centre of Excellence for Mathematical and Statistical Frontiers, Queensland University of Technology, Brisbane, QLD 4001, Australia

[c]School of Mathematical Sciences, Queensland University of Technology, Brisbane, QLD 4001, Australia

[d]School of Chemical Engineering, The University of Queensland, St Lucia, QLD 4072, Australia

[e]School of Earth and Environmental Sciences, The University of Queensland, St Lucia, QLD 4072, Australia

[f]Centre for Biodiversity and Conservation Science, The University of Queensland, St Lucia, QLD 4072, Australia

[*]Corresponding author. E-mail: sarah.vollert@connect.qut.edu.au

</div>

S.1 Further details of the coral calcification model

The deterministic model proposed by Galli and Solidoro (2018) is used to predict the rate at which calcification occurs based on the steady-state solution of a system of nonlinear ODEs. There are two modelled compartments within the coral polyp, the coelenteron, which is a fluid compartment considered to be the stomach of the polyp, and the extracellular calcifying medium (ECM), the compartment where calcification reactions occur within the coral. Additionally, the model also includes the surrounding seawater, as seawater exchanges chemical species with the coelenteron of the coral polyp.

As a result of the reactions occurring within the coral polyp, as well as exchanges of species between the compartments, the ODE model is comprised of flux terms that adjust the chemical species concentrations. Calcification is one such reaction, where calcium ions (Ca^{2+}) and carbonate ions (CO_3^{2-}) precipitate into calcium carbonate ($CaCO_3$) in the ECM. The process of calcification is mathematically described as a flux term within the system of ODEs. Calcification rates depend solely on the concentration of the relevant chemical species (Ca^{2+} and CO_3^{2-}), so prediction of calcification rates requires knowledge of the concentrations of the chemical species in each compartment.

The model also assumes the presence of the enzyme carbonic-anhydrase within the coral host, so that it catalyses the equilibration of the carbonate system (Bertucci et al., 2013). Therefore, it is assumed that the carbonate system is at separate equilibria within both the coelenteron and ECM. This assumption allows the entire carbonate system to be described uniquely by knowledge of any two chemical species of the carbonate system. Hence, the carbonate system is modelled using only dissolved inorganic carbon (DIC) and total alkalinity (TA), and equations describing carbonate species equilibrium relationships are used to calculate the concentration of all other carbonate system species (H^+, OH^-, CO_2, HCO_3^- and CO_3^{2-}).

The various flux terms, which either represent the flow of chemical species between compartments or chemical reactions, are derived from current understanding of coral physiology. Justification

and description of the mathematical forms of the ODEs and these flux terms are provided in Galli and Solidoro (2018); here we summarise the purpose of these fluxes and the parameters requiring estimation that characterise them. There are four types of processes that are assumed to control the carbonate system and calcium ion concentrations within the coral host:

1. **Gross photosynthesis and respiration reactions** – These are the processes of carbon exchange due to the zooxanthellae algae which are in symbiosis with the coral host. Photosynthesis and respiration reactions remove and supply DIC from the coral host respectively, and both reactions produce ATP energy. In the model of Galli and Solidoro (2018), the chemical fluxes due directly to photosynthesis and respiration are treated as model inputs, so do not have corresponding free parameters that require estimation.

2. **Passive transport** – There are three routes for chemical species to passively diffuse between the coral compartments within the polyp and the external seawater. These three passive transport mechanisms are:

   (a) Coelenteron-ECM transcellular diffusion: a transcellular pathway for carbon dioxide between the coelenteron and ECM (characterised by parameter k_{CO2}),

   (b) Coelenteron-ECM paracellular diffusion: a paracellular pathway for all chemical species between the coelenteron and ECM (characterised by parameter k_{pp}), and

   (c) Seawater-coelenteron diffusion: a passive exchange mechanism between the coelenteron and external seawater (characterised by parameter s).

3. **Active transport (Ca-ATPase and BAT pumps)** – The model assumes there are two active pump mechanisms between the coelenteron and ECM compartments of the coral polyp, which each increase the aragonite saturation state of the ECM.

   (a) The Ca-ATPase pump increases the calcium ion concentration in the ECM in exchange for protons being transported to the coelenteron. This pump mechanism is assumed to be driven by ATP energy sourced from photosynthesis and respiration fluxes, so is characterised by 10 parameters requiring estimation (α, β for the ATP energy budget, and the 8 Ca-ATPase parameters indicated in Table S1).

   (b) The BAT pump controls the movement of bicarbonate anions (HCO_3^-) between the coelenteron and ECM, and is characterised by 7 parameters requiring estimation (see Table S1).

4. **Calcification reactions** – The key process rate which is being predicted. The rate of this reaction depends on concentration of calcium and carbonate ions in the ECM; the parameters associated with this reaction do not require estimation in the model of Galli and Solidoro (2018).

The coral compartments and flux terms associated with these four processes are visualised in Figure 3 of the manuscript. These underlying components of the model are combined together into a system of ordinary differential equations (see Equations (23–28) in Galli and Solidoro (2018)). This system models the flow of calcium ions, DIC and TA throughout the coral polyp. To keep the carbonate species in equilibrium at each numerical timestep of the ODE, the pH

and carbonate species concentrations (H^+, OH^-, CO_2, HCO_3^- and CO_3^{2-}) are recalculated based on the current DIC, TA, temperature and salinity using the MATLAB version of CO2SYS (van Heuven et al., 2011). This CO2SYS algorithm is a commonly used software package for calculating carbonate species equilibrium concentrations; full details of this algorithm are provided in Orr et al. (2015). The ODE model in tandem with the CO2SYS algorithm in the coelenteron and ECM compartments are together solved at steady state, so that the concentrations of species and fluxes have stabilised, to gain an estimate for the coral polyp's calcification rate. Such stabilisation is expected to occur over timescales of an hour or less (Al-Horani et al., 2003; Tambutté et al., 1996).

S.2 Coral calcification model parameters

Parameters of the coral calcification model that are estimated in this work via MLE and Bayesian inference are summarised in Table S1, with assumed values indicated where possible.

Mechanism	Parameter	Description	Value Range	Units	Source/Justification of Range
Coelenteron-ECM transcellular diffusion	k_{CO2}	CO_2 permeability constant	0-0.1	cm s^{-1}	Strictly positive, maximum velocity should not exceed passive diffusion of seawater (Zeebe, 2011)
Coelenteron-ECM paracellular diffusion	k_{pp}	Paracellular pathway permeability	0-0.1	cm s^{-1}	Strictly positive, maximum velocity should not exceed passive diffusion of seawater (Yuan-Hui and Gregory, 1974; Zeebe, 2011)
Seawater-coelenteron diffusion	s	Diffusion coefficient	0-0.1	cm s^{-1}	Same as k_{pp}
ATP energy budget (for Ca-ATPase mechanism)	α	Fraction of Pg allocated to calcification	0-1	-	Fraction must be between 0 and 1
	β	Fraction of R allocated to calcification	0-1	-	Fraction must be between 0 and 1
Ca-ATPase mechanism	v_{H_c}	Proportionality constant (Ca-ATPase)	0-250	cm s^{-1}	Strictly positive
	$E0_c$	Ca-ATPase concentration	0-1.2×10^7	μmol cm^{-2}	Strictly positive
	k_{1f_c}	Ca-ATPase rate constant	0-1.4×10^{-4}	cm^4 s μmol^{-2}	Strictly positive
	k_{2f_c}	Ca-ATPase rate constant	0-0.5	cm^2 μmol^{-1}	Strictly positive
	k_{3f_c}	Ca-ATPase rate constant	0-800	s^{-1}	Strictly positive
	k_{1b_c}	Ca-ATPase rate constant	0-8	s^{-1}	Strictly positive
	k_{2b_c}	Ca-ATPase rate constant	0-500	s^{-1}	Strictly positive
	k_{3b_c}	Ca-ATPase rate constant	0-1.0×10^{-7}	cm^4 s μmol^{-2}	Strictly positive
BAT mechanism	$E0_b$	BAT concentration	0-1500	μmol cm^{-2}	Strictly positive
	k_{1f_b}	BAT rate constant	0-5.0×10^{-5}	cm^3 s^{-1} μmol^{-1}	Strictly positive
	k_{2f_b}	BAT rate constant	0-0.01	s^{-1}	Strictly positive
	k_{3f_b}	BAT rate constant	0-0.01	s^{-1}	Strictly positive
	k_{1b_b}	BAT rate constant	0-2.0×10^{-4}	s^{-1}	Strictly positive
	k_{2b_b}	BAT rate constant	0-1.0×10^{-3}	s^{-1}	Strictly positive
	k_{3b_b}	BAT rate constant	0-3.5×10^{-8}	cm^3 μmol^{-1} s^{-1}	Strictly positive

Table S1: Calibrated model parameters as defined by Galli and Solidoro (2018). Value ranges stated for estimated parameters represent the bounds for the uniform prior distributions assumed in this work.

S.3   Goodness-of-fit of local maxima

In this paper we found that the likelihood surface for the fit of the coral calcification model to data did not contain a well-defined peak. As such, many local maxima with similarly high likelihoods could be found by initiating a gradient-based search function at different locations in parameter space. These local maxima have distinct parameter values, yet each reflects a similarly good model-data fit. The goodness-of-fit for each local maxima is shown in Figure S1.
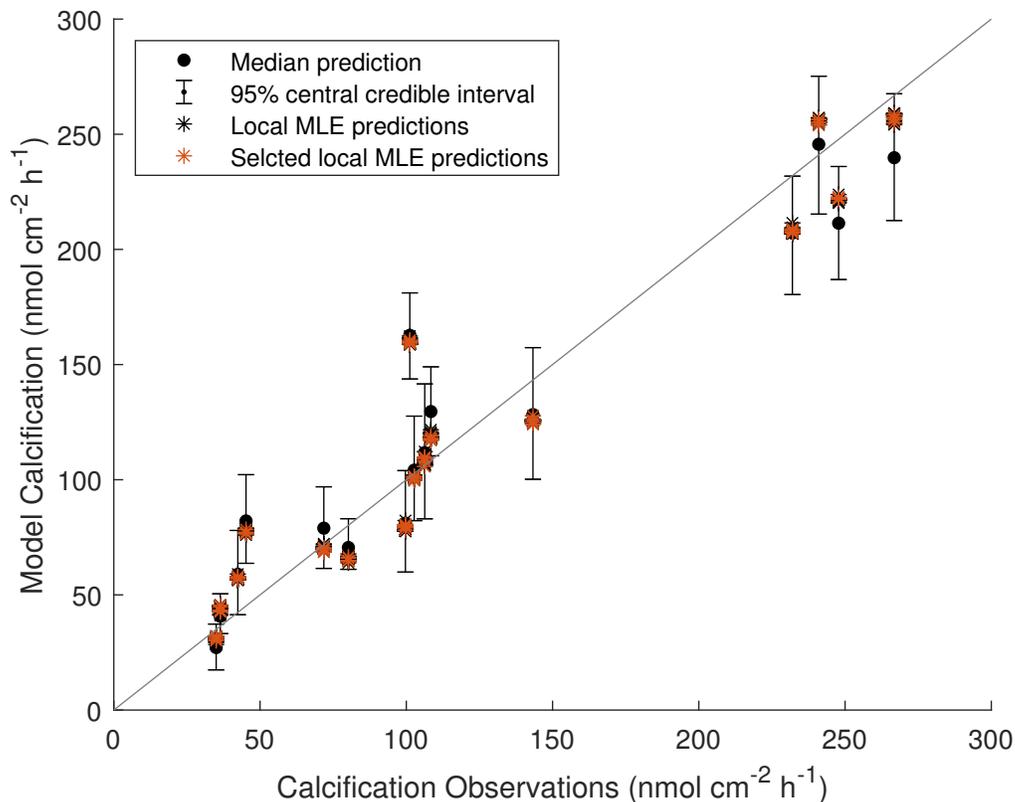
Figure S1: Posterior predictive distribution showing the goodness-of-fit for observed calcification versus model-predicted calcification, from Bayesian inference (black dots and error bars). The local likelihood maxima identified using MLE are shown for 25 parameter estimates (black asterisks) and the five estimates with the highest likelihood for the analysis of model sloppiness are also shown (orange asterisks). Each of the local maxima produce similarly good fits to the data despite having distinctly different parameter values (see corresponding orange and black vertical lines in Figure 4 of the manuscript).

S.4 Further details for calibration of the coral calcification model

S.4.1 Prior distribution

When Bayesian inference was applied to this problem, a prior distribution was specified for each of the unknown parameters,

$$
\begin{aligned}
\boldsymbol{\theta} = (&k_{CO2},\ k_{pp},\ s,\ \alpha,\ \beta,\ v_{H_c},\ E0_c,\ k_{1f_c},\ k_{2f_c},\ k_{3f_c},\ k_{1b_c}, \\
&k_{2b_c},\ k_{3b_c},\ E0_b,\ k_{1f_b},\ k_{2f_b},\ k_{3f_b},\ k_{1b_b},\ k_{2b_b},\ k_{3b_b},\ \sigma),
\end{aligned}
\tag{S1}
$$

where σ (units of μmol cm^{-2} h^{-1}) is the standard deviation characterising measurement noise, that is also estimated via Bayesian inference.

Some of the parameters had clearly definable limits. For example, α and β defined the fractions of ATP from gross photosynthesis and respiration utilised by the Ca-ATPase pump, and therefore must possess values between 0 and 1. We could also sensibly surmise that the speed of passive chemical species movement k_{CO2}, k_{pp} and s has a maximum velocity equivalent to passive diffusion within seawater (Table S1). However, most of the parameters for this model were only known to be strictly positive, yet had no other obvious choice for prior distributions. As there was little information available on the parameters, uniform priors were used for each parameter. There was also no prior knowledge of covariance between parameters, so the joint prior $\pi(\boldsymbol{\theta})$ was chosen to be the product of independent uniform priors for each parameter,

$$
\pi(\boldsymbol{\theta}) = p(k_{CO2}) \times \cdots \times p(\sigma),
\tag{S2}
$$

where $p(k_{CO2}), ..., p(\sigma)$ are the independent uniform prior distributions for the model parameters $k_{CO2}, ..., \sigma$, listed in Equation (S1). The upper and lower bounds of the uniform priors for each parameter were chosen to match the range of values specified in Table S1. Additionally, σ was assumed to have a uniform prior bounded between 0 and 50 μmol cm^{-2} h^{-1}.

During preliminary simulations, some steady state solutions of the ODE model did not converge when calculating the carbonate species equilibrium using the CO2SYS algorithm. To deal with this issue, parameter values that caused divergence of CO2SYS were excluded from the prior distribution. This modified the joint prior to

$$
\pi(\boldsymbol{\theta}) \propto p(k_{CO2}) \times \cdots \times p(\sigma) \times \mathbf{1}_c(\boldsymbol{\theta}),
\tag{S3}
$$

where the indicator function $\mathbf{1}_c$ was defined here to output one if the CO2SYS converges and zero otherwise. This rejection procedure did not substantially alter the shape of the marginal prior distributions (shaded grey regions in Figure 4 of the manuscript). In addition, we found in practice that rejections due to divergence of the CO2SYS algorithm reduced as the SMC algorithm we used for Bayesian inference moved towards areas of high posterior density. Hence, for this coral calcification model, areas of low posterior density are more likely to involve parameter combinations that yield highly unreasonable predictions that subsequently cause divergence of CO2SYS. However, we also cannot rule out the possibility that this issue was due to the numerical procedures we used. Whether this divergence issue was a result of areas of low posterior

density or the numerical procedures, the resulting prior distributions which appear uniform, and the observation that the rate of rejection due to divergence reduced as areas of high posterior density were visited more often, suggested that the approach of rejecting these parameter samples was reasonable.

S.4.2 Posterior simulation

The SMC algorithm used to estimate the posterior distribution was adapted from Drovandi and Pettitt (2011). For this application, an SMC algorithm based on a likelihood annealing approach was used. This approach allowed the ensemble of particles to steadily converge onto the posterior distribution.

Within the algorithm, a transform was placed on each of the parameters, such that any real number could be proposed and converted to a proposal for $\boldsymbol{\theta}$ within the prior bounds. This transform was applied for each parameter, where the transforms for the jth parameter $\theta_j \to \tilde{\theta}_j$ and the equivalent inverse form $\tilde{\theta}_j \to \theta_j$ were defined as

$$\tilde{\theta}_j = \log\left(\frac{\theta_j - l_j}{u_j - \theta_j}\right), \qquad \theta_j = \frac{l_j + u_j e^{\tilde{\theta}_j}}{1 + e^{\tilde{\theta}_j}}, \tag{S4}$$

where l_j and u_j are the lower and upper bounds for parameter θ_j, respectively. This transform then allowed any real number proposed as $\tilde{\theta}_j$ to be converted to a θ_j value within the uniform prior bounds, such that if $-\infty < \tilde{\theta}_j < \infty$, then $l_j \leq \theta_j \leq u_j$.

The final result of the SMC algorithm was then visually examined to ensure that the results were reproducible over multiple independent runs of the algorithm, such that the posterior distributions were similar across independent sets of samples. For this application, the SMC algorithm was run independently five times to produce five independent sets of posterior samples. The estimated marginal distributions and posterior predictive distributions were visually very similar for each independent set of 5000 samples. Hence, the results were found to be highly reproducible for 5000 posterior samples. Each of these five independent samples were then used to ensure reproducibility of the results of an analysis of model sloppiness, using a posterior covariance or LIS sensitivity matrix.

S.5 Analysis of model sloppiness using the posterior covariance method

Multiple independent results of the analysis of model sloppiness via a posterior covariance sensitivity matrix were produced to analyse the consistency of results. Five sets of posterior samples were independently produced using the SMC algorithm and each of these were used to analyse the sloppiness of the coral calcification model. Figure S2 shows the eigenparameters for three of the five sets of independent samples, revealing that the results are reproducible for the first nine eigenparameters (the results for the two posterior samples not shown here were similar). After this point, the independent analyses show different model-data fits; however, this can be expected because these parameter combinations commonly become less constrained in a sloppy model.

| Eigen-parameter	Set of posterior samples	λ_j/λ_1	3 x Diffusion Mechanisms			Ca-ATPase Pump Mechanism										BAT Pump Mechanism						
			k_{CO_2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1f_c}	k_{2f_c}	k_{3f_c}	k_{1b_c}	k_{2b_c}	k_{3b_c}	E_{0_b}	k_{1f_b}	k_{2f_b}	k_{3f_b}	k_{1b_b}	k_{2b_b}	k_{3b_b}
1	1	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.5	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.4	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	3	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.4	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	1	0.1118	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2	0.1083	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	3	0.1103	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	0.0	-0.1	-0.1	0.2	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	1	0.0386	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2	0.0381	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.4	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	3	0.0399	0.0	0.1	0.0	-0.4	-0.3	1.0	-0.2	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	1	0.0264	0.0	-0.2	0.0	-0.5	-0.5	-0.2	1.0	0.6	-0.4	-0.2	0.5	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2	0.0235	0.0	-0.2	0.0	-0.5	-0.4	-0.2	0.7	1.0	-0.5	-0.2	0.6	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	3	0.0248	0.0	-0.1	0.0	-0.4	-0.3	-0.2	1.0	0.4	-0.3	-0.2	0.4	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
5	1	0.0237	0.0	-0.1	0.0	-0.1	-0.1	0.0	-0.8	1.0	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2	0.0223	0.0	-0.1	0.0	-0.1	-0.1	-0.1	1.0	-0.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	3	0.0221	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.6	1.0	-0.2	-0.1	0.2	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
6	1	0.0148	0.0	0.1	0.0	0.7	1.0	0.1	0.2	0.2	-0.9	0.1	0.4	-0.1	0.0	0.0	-0.1	0.0	-0.1	0.0	0.0	0.0
	2	0.0146	0.0	0.1	0.2	0.6	1.0	0.1	0.2	0.0	-0.7	0.0	0.4	0.0	-0.1	-0.1	0.0	0.0	0.0	-0.1	0.0	0.0
	3	0.0142	-0.1	0.1	0.1	0.7	1.0	0.1	0.2	0.1	-0.7	-0.2	0.4	0.0	-0.1	0.1	-0.1	-0.1	0.0	0.0	0.0	0.0
7	1	0.0122	-0.1	-0.1	0.0	0.2	0.3	0.0	0.0	0.0	0.9	-0.6	1.0	0.0	-0.1	0.0	0.0	0.0	-0.1	0.0	0.0	0.0
	2	0.0115	0.0	-0.1	-0.2	0.2	0.3	0.0	0.0	0.0	1.0	-0.8	0.9	0.1	0.0	-0.1	0.1	0.0	0.0	0.0	0.0	0.0
	3	0.0126	0.1	-0.1	0.0	0.1	0.2	0.0	-0.1	0.0	1.0	-0.8	0.7	0.1	0.1	0.0	0.0	0.1	-0.1	0.1	0.0	0.0
8	1	0.0115	-0.1	-0.1	0.2	-0.1	-0.1	0.0	-0.1	0.0	0.1	1.0	0.5	-0.2	0.0	0.0	0.0	0.1	-0.1	0.0	0.0	0.0
	2	0.0109	0.0	-0.1	0.1	0.0	-0.1	0.0	0.0	-0.1	0.2	1.0	0.7	-0.1	-0.1	0.0	0.0	0.1	-0.1	0.0	0.0	0.1
	3	0.0113	0.0	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.2	1.0	0.8	-0.1	0.0	0.0	0.0	0.0	0.1	-0.1	-0.1	0.0
9	1	0.0101	0.1	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.1	0.2	0.0	-0.1	0.2	0.2	0.0	-0.1	0.0
	2	0.0103	0.0	0.0	1.0	0.0	-0.1	0.0	0.0	0.0	0.1	-0.1	0.0	-0.1	-0.1	0.0	0.0	-0.1	-0.1	0.0	0.1	0.0
	3	0.0099	0.0	0.0	1.0	0.0	-0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.2	0.1	-0.2	0.0	0.0	-0.1	0.0	
10	1	0.0094	-0.2	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	0.0	-0.3	-0.6	-0.1	1.0	-0.1	-0.1	0.1
	2	0.0089	-0.1	0.0	-0.1	0.0	0.0	0.0	0.0	0.0	0.1	-0.1	0.2	0.0	0.0	0.1	1.0	-0.2	-0.2	-0.3	0.1	-0.2
	3	0.0089	1.0	0.1	-0.1	0.0	0.1	0.0	0.0	0.0	-0.1	0.1	0.0	-0.1	0.4	0.6	0.3	-0.2	0.2	0.3	0.5	-0.1
11	1	0.0090	0.0	0.1	0.1	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.0	-0.1	-0.5	-0.4	1.0	0.1	0.6	0.3	0.0	0.0
	2	0.0087	-0.1	0.0	0.0	0.0	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	-0.2	-0.4	-0.5	0.3	0.1	1.0	0.0	0.2	-0.2
	3	0.0085	0.0	0.0	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.4	-0.3	-0.2	0.0	0.3	-0.9	0.0	1.0	-0.3
12	1	0.0087	0.2	0.0	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	1.0	0.3	-0.3	0.4	0.0	0.0	0.0
	2	0.0085	0.3	-0.1	0.2	0.1	0.1	0.0	0.0	-0.1	0.0	0.1	0.2	0.8		1.0	0.0	-0.7	0.9	0.0	-0.2	-0.4
	3	0.0085	-0.3	-0.1	0.2	0.0	0.0	0.0	-0.1	0.0	0.0	-0.1	0.1	-0.5	0.1	-0.1	0.2	0.7	1.0	-0.1	0.7	0.8
13	1	0.0082	0.0	0.0	-0.1	0.0	0.1	0.0	0.0	0.0	0.1	0.0	0.0	0.3	1.0	-0.4	0.3	0.0	0.0	0.1	0.0	-0.1
	2	0.0083	1.0	0.3	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	-0.1	0.2	-0.8	0.9	0.2	1.0	0.0	0.0	0.0	-0.2
	3	0.0084	0.7	0.1	0.1	0.1	0.2	0.0	0.0	0.0	-0.2	0.0	-0.1	0.1	1.0	-1.0	0.9	1.0	-0.8	0.0	-0.7	0.6
14	1	0.0081	-0.3	0.1	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.2	0.0	1.0	-0.2	0.1	-0.2	-0.1	-0.1	0.5	-0.3	0.6
	2	0.0082	-0.5	-0.1	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.7	-0.1	0.1	1.0	-0.1	0.3	0.0	-0.8
	3	0.0083	-0.3	0.1	-0.2	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.1	-0.8	0.8	1.0	-0.1	-0.6	0.1	-0.1	0.9
15	1	0.0080	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	-0.6	0.1	0.0	-0.3	0.1	0.0	1.0	-0.2	0.1
	2	0.0078	-1.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.5	-0.9	0.7	-0.2	-0.5	-0.1	1.0	0.3	-0.4
	3	0.0080	0.1	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.5	-0.7	1.0	-0.4	0.5	0.2	0.1	-0.6
16	1	0.0078	1.0	0.1	-0.2	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1	-0.1	-0.1	-0.1	0.2	0.1	0.0	-0.1	0.3
	2	0.0077	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	-0.1	0.2		0.2	0.4	0.3	0.2	0.6	-0.3	1.0
	3	0.0078	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.1		-0.3	0.6	0.2	1.0	0.2	-0.5	-0.3	-0.8
17	1	0.0074	-0.4	-0.1	-0.2	0.0	0.0	0.0	0.0	0.0	-0.1	-0.5	0.2	0.2		0.2	0.2	1.0	0.0	-0.4	-0.5	0.7
	2	0.0075	0.5	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.1	-0.1	1.0	0.0	-0.6	0.0	-0.1	0.0	0.4	-0.3	0.0
	3	0.0076	-0.9	-0.3	-0.2	0.0	0.0	0.0	-0.1	-0.1	-0.2	-0.2	-0.1	-0.6	0.7	0.4	0.3	0.2	-0.1	1.0	-0.2	-0.5
18	1	0.0073	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.1		0.0	-0.1	0.0	0.0	0.1	1.0	0.6
	2	0.0071	0.7	0.1	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.5	0.3		-0.2	0.1	-0.2	-0.2	0.9	1.0	-0.2
	3	0.0075	0.6	0.0	0.2	0.0	-0.1	0.0	0.0	-0.1	-0.1	-0.1	-0.4	-0.8		-0.2	-0.4	0.4	0.1	1.0	-0.4	0.2
19	1	0.0072	0.0	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.5	-0.1		0.3	-0.2	1.0	0.2	0.3	0.5	-0.8
	2	0.0068	-0.2	0.2	0.1	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.7	0.2		0.1	-0.2	0.2	0.2	-0.5	1.0	0.5
	3	0.0072	-0.2	0.1	0.0	0.0	0.1	0.0	0.0	0.0	0.1	0.2	-0.1	1.0		0.0	0.0	0.3	0.1	0.5	0.4	0.0
20	1	0.0058	-0.1	1.0	0.0	0.0	0.0	0.2	0.2	0.2	0.2	0.1	-0.1	-0.1	0.1	0.0	0.0	0.1	-0.1	-0.1	0.0	0.0
	2	0.0057	-0.1	1.0	0.0	0.0	0.0	0.2	0.2	0.2	0.2	0.1	-0.1	-0.1	0.1	-0.1	0.0	-0.1	0.1	0.0	-0.1	-0.1
	3	0.0057	-0.2	1.0	0.0	0.1	0.0	0.2	0.2	0.2	0.2	0.1	0.0	-0.2	0.1	0.0	0.0	0.0	0.0	0.1	0.0	-0.1
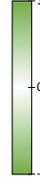
Figure S2: Eigenvector element values for eigenparameters identified using a posterior covariance analysis of sloppiness. These eigenparameters are ordered from highest relative importance to lowest according to the size of the eigenvalues. For each eigenparameter, the results of three independent sets of posterior samples with 5000 particles each are compared to indicate the consistency of results. For each eigenparameter, the values were normalised by the leading eigenvector value, such that they are rescaled to be between -1 and 1 inclusive. Here the colour darkens as the absolute values of the eigenvector values increase from 0 to 1, such that the model-data fit is more sensitive to darker eigenvector values. Additionally, the relative size of each eigenvalue when compared to the leading eigenvalue for each sample has been included in the column λ_j/λ_1.

S.6 Analysis of model sloppiness using the LIS method

Here, we compare multiple independent analyses of model sloppiness using a LIS sensitivity matrix to show the consistency of results. Five sets of posterior samples were independently produced using the SMC algorithm and each of these were used to analyse the sloppiness of the coral calcification model. Figure S3 shows the eigenparameters for each of these five independent sets of samples, revealing that the results are somewhat reproducible for the two stiffest eigenparameters, but not any others. This figure also shows the results of a posterior covariance method approach for comparison, as this reveals that the two approaches have a similar model-data fit for the two stiffest eigenparameters.

Eigen-parameter	Sensitivity Matrix	Sample	$\frac{\lambda_j}{\lambda_1}$	k_{CO_2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1f_c}	k_{2f_c}	k_{3f_c}	k_{1b_c}	k_{2b_c}	k_{3b_c}	E_{0_b}	k_{1f_b}	k_{2f_b}	k_{3f_b}	k_{1b_b}	k_{2b_b}	k_{3b_b}
1	S_P	1	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.5	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
		2	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
		3	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
		4	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
		5	1.0000	0.0	-0.6	0.0	0.3	0.2	1.0	0.5	0.5	0.5	0.2	-0.5	-0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
2	S_P	1	0.1118	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1	0.0243	-0.1	0.1	0.0	-0.9	1.0	-0.1	-0.1	0.1	0.2	0.1	-0.2	-0.2	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
		2	0.0237	0.0	0.1	0.0	-0.9	1.0	-0.1	0.0	0.0	0.2	0.2	-0.2	-0.2	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
		3	0.0239	-0.1	0.1	0.0	-0.9	1.0	-0.2	-0.1	0.0	0.3	0.3	-0.3	-0.3	-0.1	-0.1	0.0	0.0	0.0	0.0	0.0	0.0
		4	0.0229	0.0	0.1	0.0	-0.9	1.0	-0.1	0.0	0.0	0.2	0.2	-0.1	-0.2	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
		5	0.0232	0.0	0.0	0.0	-1.0	1.0	-0.1	-0.1	0.1	0.1	0.2	-0.2	-0.2	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	S_P	1	0.0386	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1	0.0133	-0.1	0.0	0.0	0.4	-0.1	-0.7	-0.2	0.0	0.3	1.0	-0.3	-1.0	-0.3	0.5	0.1	0.3	0.3	0.0	0.0	0.0
		2	0.0125	-0.1	-0.1	0.2	-0.6	0.1	0.8	0.1	0.0	-0.3	-1.0	0.3	1.0	0.3	-0.3	-0.1	-0.1	-0.2	0.0	0.0	0.0
		3	0.0123	0.1	0.4	0.0	0.7	-0.2	-0.4	-0.4	0.0	0.3	1.0	-0.3	-1.0	-0.4	-0.2	0.0	-0.1	0.0	0.0	0.0	0.0
		4	0.0107	0.0	-0.7	0.0	0.0	-0.1	-0.5	-0.2	-0.1	0.0	0.4	0.0	-0.4	-0.2	1.0	0.1	0.5	0.5	0.0	-0.1	0.0
		5	0.0130	0.0	-0.5	-0.1	0.3	-0.1	-0.6	-0.5	-0.1	0.1	0.6	-0.1	-0.6	-0.4	1.0	0.1	0.5	0.4	0.0	0.0	0.0
4	S_P	1	0.0264	0.0	-0.2	0.0	-0.5	-0.5	-0.2	1.0	0.6	-0.4	-0.2	0.5	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1	0.0100	0.2	-0.9	0.0	-0.1	0.1	-0.4	0.0	-0.1	0.0	-0.5	0.0	0.5	0.2	1.0	0.1	0.5	0.4	0.0	0.0	0.0
		2	0.0089	0.0	-0.8	0.0	-0.2	0.0	-0.2	-0.1	-0.1	-0.2	-0.2	0.2	0.2	0.0	1.0	0.0	0.6	0.5	0.0	0.0	0.0
		3	0.0095	0.0	-0.6	0.1	-0.1	0.0	-0.4	-0.2	-0.1	-0.1	0.1	0.0	-0.1	0.0	1.0	0.1	0.4	0.5	0.0	0.0	0.0
		4	0.0097	0.0	-0.6	0.0	-0.5	0.2	0.5	0.0	-0.1	-0.4	-1.0	0.4	1.0	0.4	0.6	0.1	0.2	0.2	0.0	-0.1	0.0
		5	0.0105	0.0	1.0	0.1	0.4	-0.1	0.0	-0.4	0.0	0.4	0.9	-0.4	-0.9	-0.3	-0.9	0.0	-0.4	-0.3	0.0	0.0	0.0
5	S_P	1	0.0237	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.8	1.0	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	S_L	1	0.0068	0.1	-0.3	0.1	-0.8	-0.5	0.8	0.1	-0.1	-0.9	0.5	1.0	-0.5	-0.5	0.5	0.0	0.0	0.2	0.0	0.0	0.0
		2	0.0028	0.1	0.5	-0.5	-0.2	0.0	0.8	-0.3	0.7	-0.9	0.4	1.0	-0.6	-0.1	0.2	0.1	0.2	0.2	0.0	0.0	0.0
		3	0.0042	0.8	-0.1	0.6	-0.7	-0.3	0.4	0.2	0.5	-0.9	0.6	1.0	-0.8	0.6	-0.2	0.1	0.0	0.2	0.0	-0.1	0.0
		4	0.0043	-0.5	-0.4	0.4	-0.8	-0.6	0.7	0.5	-0.4	-0.9	0.6	1.0	-0.8	0.0	-0.5	0.1	0.0	0.0	0.0	0.0	0.0
		5	0.0029	-0.2	-0.3	0.2	-0.2	0.0	0.9	0.1	-0.9	-0.8	0.7	1.0	-0.7	0.7	0.0	0.1	0.0	0.2	0.0	-0.2	0.0

The "3 x Diffusion Mechanisms" header spans the k_{CO_2}, k_{pp}, s columns; the "Ca-ATPase Pump Mechanism" header spans α through k_{3b_c}; the "BAT Pump Mechanism" header spans E_{0_b} through k_{3b_b}.

Figure S3: Comparison of eigenvector values of an analysis of model sloppiness using a LIS (S_L) and posterior covariance (S_P) sensitivity matrix (blue and green respectively). For each LIS matrix eigenparameter, the results of five independent sets of posterior samples with 5000 particles each are compared to indicate the consistency of results. Five eigenvector values are also shown for a posterior covariance method of analysing model sloppiness (green) from a single sample of 5000 particles, since the eigenvector values from this analysis were fairly consistent across independent samples (Figure S2). See the caption of Figure S2 for further interpretation of this figure.

S.7 Goodness-of-fit comparisons for other reduced models

Firstly, two other model reductions informed by the analysis of model sloppiness were tested: one where only the BAT pump mechanism was removed and one where only the coelenteron-ECM transcellular diffusion was removed. The goodness-of-fit of the model without BAT pump mechanism and the model without the coelenteron-ECM transcellular diffusion are shown in Figures S4 and S5 respectively.
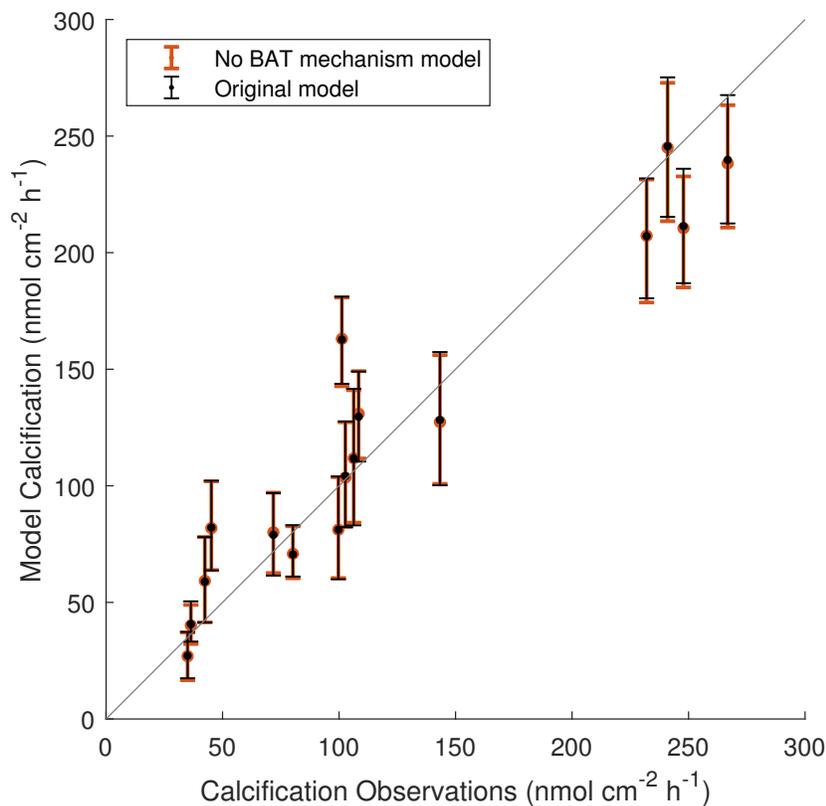

Figure S4: Posterior predictive distribution showing observed calcification compared to modelled calcification for both the original and reduced model. Here the reduced model is one without the BAT pump mechanism. Here, the dots indicate the median prediction value, and error bars indicate the 95% central credible interval for calcification predictions from the respective posterior samples of each model.
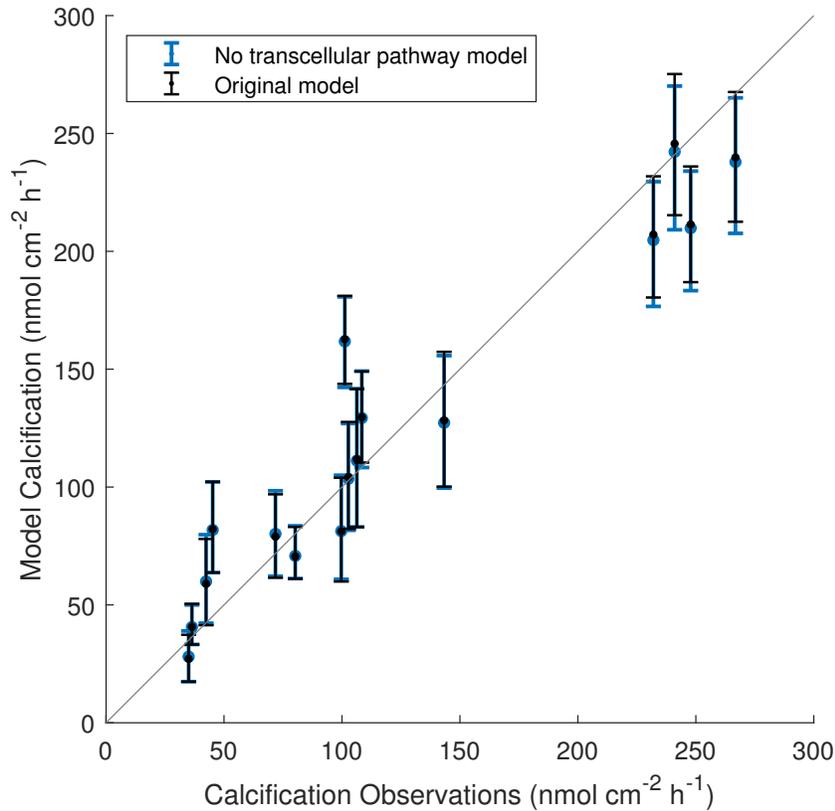
Figure S5: Posterior predictive distribution showing observed calcification compared to modelled calcification for both the original and reduced model. Here the reduced model is one without the coelenteron-ECM transcellular diffusion for CO_2. Here, the dots indicate the median prediction value, and error bars indicate the 95% central credible interval for calcification predictions from the respective posterior samples of each model.

Secondly, we investigated models where mechanisms necessary for a good model-data fit were removed. The analysis of model sloppiness identified that the coelenteron-ECM paracellular diffusion and Ca-ATPase pump mechanisms contributed substantially to the stiffest eigenparameter (Figure 7 of the manuscript). Hence, we sought to analyse models with these mechanisms excluded as a test case for what happens if the sloppy analysis recommendations for which mechanisms should be removed are not followed. Removing the coelenteron-ECM paracellular diffusion results in the exchange of chemical species between the coelenteron and ECM to being controlled entirely though the exchange of carbon dioxide. Consequently, calibrating this model resulted in numerical issues, and so we did not further consider the model without coelenteron-ECM paracellular diffusion. Analysing the remaining unrecommended model reduction – removal of the Ca-ATPase pump mechanism – resulted in a visibly poorer goodness-of-fit in comparison to the original model (Figure S6).
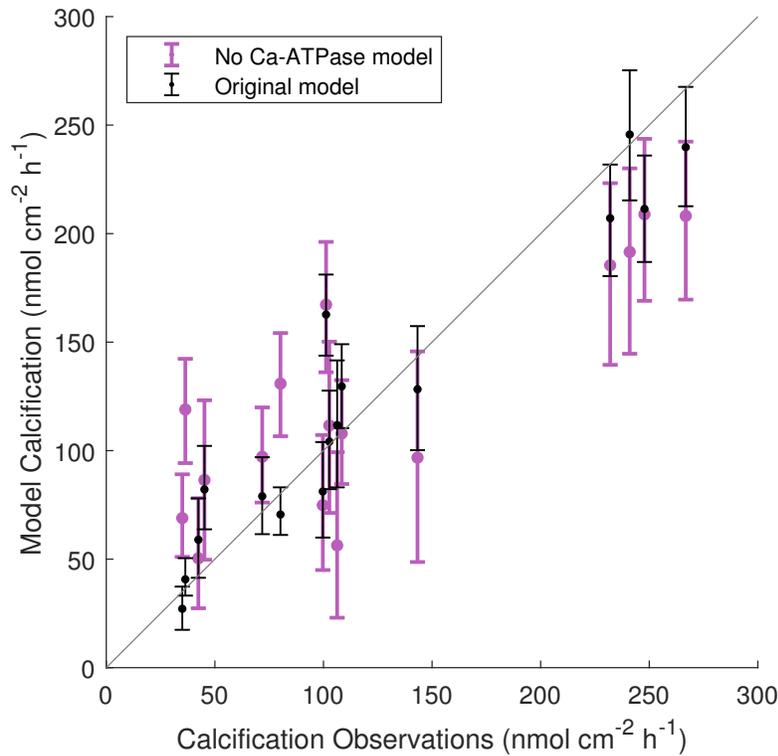
Figure S6: Posterior predictive distribution showing observed calcification compared to modelled calcification for both the original and reduced model. Here the reduced model is one without the Ca-ATPase pump mechanism. Here, the dots indicate the median prediction value, and error bars indicate the 95% central credible interval for calcification predictions from the respective posterior samples of each model.

S.8 Bayes' factors for investigated models

For each of the models considered in this case study, the Bayesian model evidence was estimated for the sample and used to produce approximate Bayes' factors (Kass and Raftery, 1995) for comparing models to the original model. Values close to one identify that the model evidence between the compared models is similar, whereas a value of 10 or more might suggest strong evidence to prefer the original model. The estimated Bayes' factors quantitatively suggest that each of the model reductions suggested by the analysis of model sloppiness are comparable to the original model (Table S2). However, when looking at a model reduction which was not recommended by the analysis of model sloppiness there is strong evidence to prefer the original model for this dataset.

Model mechanism(s) removed	Estimated Bayes' factor
Coelenteron-ECM transcellular diffusion	1.12
BAT pump	1.26
Coelenteron-ECM transcellular diffusion and BAT pump	1.25
Ca-ATPase pump	7.30×10^6

Table S2: The estimated Bayes' factors for each reduced model compared to the original model. Note that the analysis of sloppiness suggests the first three model reductions, however, it suggests that the Ca-ATPase pump is important for a good model-data fit.

S.9   Reduced model compared to original model

In this section, we compare the original coral calcification model proposed by Galli and Solidoro (2018) to the reduced version which excludes both the BAT pump mechanism and the coelenteron-ECM transcellular diffusion (characterised by parameter k_{CO2}). Firstly, Figure S7 demonstrates that the estimated marginal densities of each of the remaining parameters appear to be similar to that of the original model calibration. Secondly, Figure S8 shows the decay in relative eigenvalue size using a posterior covariance analysis of model sloppiness. Here, the reduced model has a faster decay in relative eigenvalue importance. Finally, Figure S9 shows the corresponding eigenparameters for both models. Each of the first six eigenparameters indicate that there is a similar model-data fit for both models.
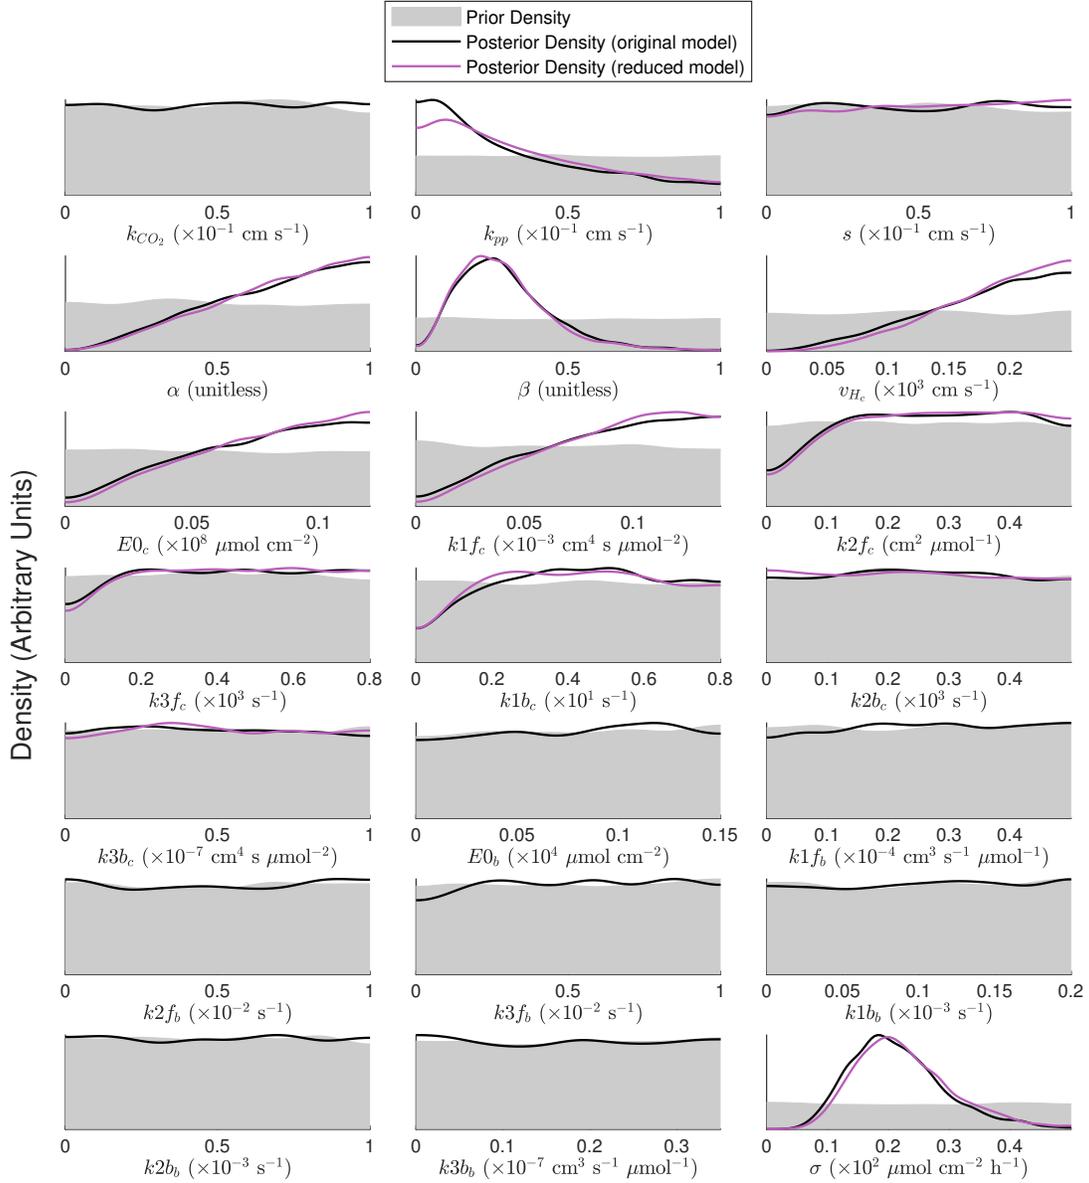
Figure S7: Estimated posterior marginal densities for the model parameters from 5000 posterior samples. This figure shows the uncertainty of model parameters for both the original and reduced model. Parameters k_{CO_2}, $E0_b$, k_{1f_b}, k_{2f_b}, k_{3f_b}, k_{1b_b}, k_{2b_b}, and k_{3b_b} are present in the original model but not in the reduced model, so for these parameters results are only shown for the original model. The bounds of the uniform prior distributions for each parameter are the upper and lower limits on the x-axes.
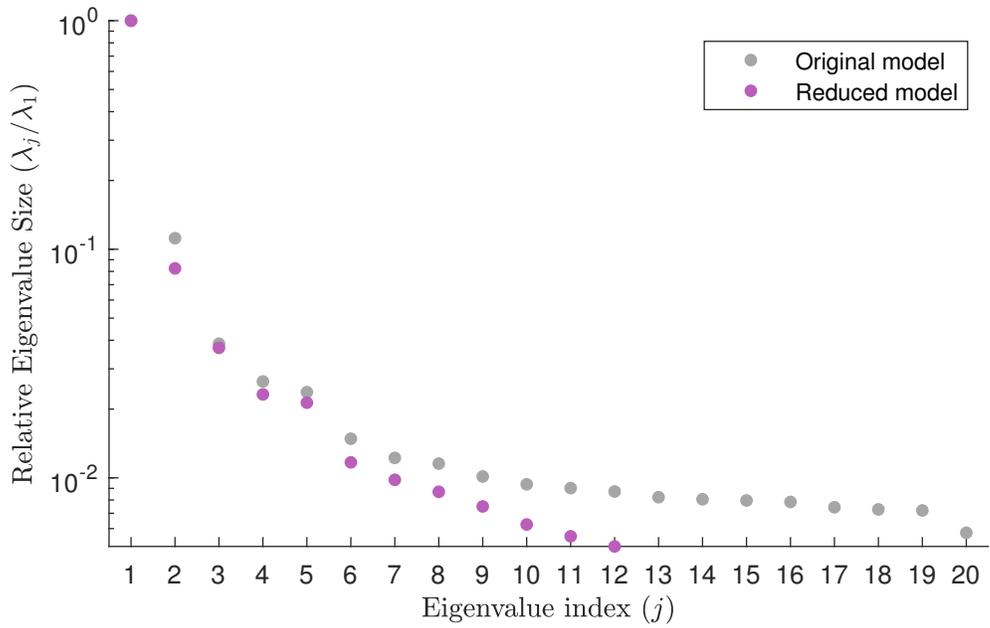
Figure S8: Relative eigenvalue sizes when compared to the leading eigenvalue. This posterior covariance analysis of sloppiness was conducted on 5000 posterior samples for each of the models. For both the original and reduced model, the relative eigenvalue sizes of independent samples produced similar results so only one sample is shown for each model here.

Eigen-parameter	Model	$\frac{\lambda_j}{\lambda_1}$	k_{CO_2}	k_{pp}	s	α	β	v_{H_c}	E_{0_c}	k_{1f_c}	k_{2f_c}	k_{3f_c}	k_{1b_c}	k_{2b_c}	k_{3b_c}	E_{0_b}	k_{1f_b}	k_{2f_b}	k_{3f_b}	k_{1b_b}	k_{2b_b}	k_{3b_b}
1	Original	1.0000	0.0	-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.5	0.2	-0.4	-0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1	Reduced	1.0000		-0.5	0.0	0.3	0.1	1.0	0.5	0.5	0.4	0.2	-0.4	-0.1	0.0							
2	Original	0.1118	0.0	0.0	0.0	1.0	-0.8	0.0	0.0	-0.1	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
2	Reduced	0.0825		0.0	0.0	1.0	-0.8	0.0	0.0	0.0	-0.1	-0.1	0.1	0.0	0.0							
3	Original	0.0386	0.0	0.1	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
3	Reduced	0.0371		0.2	0.0	-0.3	-0.3	1.0	-0.3	-0.4	-0.4	-0.2	0.4	0.1	0.0							
4	Original	0.0264	0.0	-0.2	0.0	-0.5	-0.5	-0.2	1.0	0.6	-0.4	-0.2	0.5	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
4	Reduced	0.0232		-0.1	0.0	-0.2	-0.2	-0.1	1.0	-0.2	-0.2	-0.1	0.3	0.0	0.0							
5	Original	0.0237	0.0	0.0	0.0	-0.1	-0.1	0.0	-0.8	1.0	-0.1	-0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
5	Reduced	0.0213		0.0	0.0	-0.2	-0.2	0.0	-0.1	1.0	-0.3	-0.2	0.3	0.1	0.0							
6	Original	0.0148	0.0	0.1	0.0	0.7	1.0	0.1	0.2	0.2	-0.9	0.1	0.4	-0.1	0.0	0.0	-0.1	0.0	-0.1	0.0	0.0	0.0
6	Reduced	0.0117		0.1	0.0	0.6	1.0	0.1	0.1	0.1	-0.7	-0.2	0.3	0.0	-0.1							
7	Original	0.0122	-0.1	-0.1	0.0	0.2	0.3	0.0	0.0	0.0	0.9	-0.6	1.0	0.0	-0.1	0.0	0.0	0.0	-0.1	0.0	0.0	0.0
7	Reduced	0.0098		0.0	0.0	-0.1	-0.2	0.0	0.0	0.0	-0.7	1.0	-0.4	-0.1	0.0							
8	Original	0.0115	-0.1	-0.1	0.2	-0.1	-0.1	0.0	-0.1	0.0	0.1	1.0	0.5	-0.2	0.0	0.0	0.0	0.1	-0.1	0.0	0.0	0.0
8	Reduced	0.0087		-0.1	0.2	0.1	0.1	0.0	0.0	0.0	0.4	0.7	1.0	-0.1	0.0							
9	Original	0.0101	0.1	0.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	-0.1	0.0	0.1	0.2	0.0	-0.1	0.2	0.2	0.0	-0.1	0.0
9	Reduced	0.0075		0.0	1.0	0.0	-0.1	0.0	0.0	0.0	-0.1	-0.1	-0.1	0.0	0.1							

The "Ca-ATPase Pump Mechanism" spans columns α through k_{3b_c}. The "BAT Pump Mechanism" spans columns E_{0_b} through k_{3b_b}. The "3 x Diffusion Mechanisms" spans columns k_{CO_2}, k_{pp}, s.

Figure S9: Eigenvector values for eigenparameters identified using a posterior covariance analysis of sloppiness, for both the original model (in grey) and reduced model (in purple). These eigenparameters correspond to the nine largest eigenvalues, and so are ordered from highest relative importance to lowest. Results beyond the ninth eigenparameter were inconsistent for the original model (Figure S2), and so are not shown in this comparison. See the caption of Figure S2 for further interpretation of this figure.

Supplementary Material References

Al-Horani, F., S. Al-Moghrabi, and D. De Beer (2003). The mechanism of calcification and its relation to photosynthesis and respiration in the scleractinian coral *Galaxea fascicularis*. *Marine Biology 142*(3), 419–426.

Bertucci, A., A. Moya, S. Tambutté, D. Allemand, C. T. Supuran, and D. Zoccola (2013). Carbonic anhydrases in anthozoan corals—A review. *Bioorganic & Medicinal Chemistry 21*(6), 1437–1450.

Drovandi, C. C. and A. N. Pettitt (2011). Estimation of parameters for macroparasite population evolution using approximate Bayesian computation. *Biometrics 67*(1), 225–233.

Galli, G. and C. Solidoro (2018). ATP supply may contribute to light-enhanced calcification in corals more than abiotic mechanisms. *Frontiers in Marine Science 5*, 68.

Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association 90*(430), 773–795.

Orr, J., J.-M. Epitalon, and J.-P. Gattuso (2015). Comparison of ten packages that compute ocean carbonate chemistry. *Biogeosciences 12*(5), 1483–1510.

Tambutté, É., D. Allemand, E. Mueller, and J. Jaubert (1996). A compartmental approach to the mechanism of calcification in hermatypic corals. *Journal of Experimental Biology 199*(5), 1029–1041.

van Heuven, S., D. Pierrot, J. Rae, E. Lewis, and D. Wallace (2011). MATLAB program developed for CO_2 system calculations. ORNL/CDIAC-105b. Technical report, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tennessee.

Yuan-Hui, L. and S. Gregory (1974). Diffusion of ions in sea water and in deep-sea sediments. *Geochimica et Cosmochimica Acta 38*(5), 703–714.

Zeebe, R. E. (2011). On the molecular diffusion coefficients of dissolved CO_2, HCO_3^-, and CO_3^{2-} and their dependence on isotopic mass. *Geochimica et Cosmochimica Acta 75*(9), 2483–2498.