



Inverse methods for consistent quantification of seafloor anoxia using uranium isotope data from marine sediments

Michael A. Kipp*, François L.H. Tissot

The Isotoparium, Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA, 91125, USA



ARTICLE INFO

Article history:

Received 2 January 2021

Received in revised form 30 August 2021

Accepted 6 October 2021

Available online xxx

Editor: L. Coogan

Keywords:

paleoredox
isotope mass balance
Bayesian inference
MCMC
carbonate

ABSTRACT

Uranium isotopes ($\delta^{238}\text{U}$) have quickly become one of the most widely-used redox proxies in paleoceanographic studies. The quantitative power of the $\delta^{238}\text{U}$ proxy derives from the long marine residence time of uranium and the dominance of reduction in fractionating uranium isotopes during removal from seawater. The seawater $\delta^{238}\text{U}$ value is therefore sensitive to the size of the anoxic sink, and by extension, the area of the seafloor overlain by anoxic waters. Leveraging the ability of carbonates to record and retain the seawater $\delta^{238}\text{U}$ value, and the ubiquity of carbonate sediments in the geologic record, numerous studies have quantified seafloor anoxia across ocean anoxic events, mass extinctions, and global climatic changes. In most cases, forward models of marine uranium isotope mass balance have been used, illustrating potential histories of seafloor anoxia during these events.

Here we show that there are multiple ways in which such forward modeling can lead to spurious inferences of anoxia, including (i) the poor sensitivity of the $\delta^{238}\text{U}$ proxy when fractional anoxia is high, and (ii) the inherent bias in generating illustrative forward model outputs in stratigraphic sections with expected anoxic intervals. We thus explore inverse modeling approaches to constrain the most likely history of seafloor anoxia for a given $\delta^{238}\text{U}$ dataset, and ultimately develop a framework for doing so using Bayesian inference via Markov Chain Monte Carlo simulation. We show that this approach can recover simulated trends, and further reconstruct marine anoxia for eight published $\delta^{238}\text{U}$ datasets. We find that some previous interpretations of anoxic seafloor extent were inaccurate, either because steady state was improperly assumed, or because the illustrative forward models used were poor fits to the data. In order to overcome these issues in future work with the $\delta^{238}\text{U}$ redox proxy, we have made this model publicly available, and also offer suggestions for the judicious use of forward models. By building on this framework, the future quantification of marine anoxia during transient environmental perturbations can be performed consistently, thereby facilitating robust comparison of anoxic extent between events.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

The oxygenation of the ocean dramatically restructured biogeochemical cycles and paved the way for the emergence of animal life on Earth (Nursall, 1959). Since the rise of animals, intervals of extensive marine anoxia have resulted in biogeochemical perturbations and occasionally mass extinctions (e.g., Wignall and Twitchett, 1996). Quantifying the appearance of oxygen in Earth's early history and its disappearance in more recent events is therefore a critical task when studying the role of redox in the evolution of life on our planet. Furthermore, understanding past marine redox fluctuations allows us to reconstruct perturbations to biogeochemical

cycles, which holds implications for understanding future climate change.

In order to reconstruct oxygen levels in ancient seawater, many "paleo-redox proxies" are employed, which are typically elemental or isotopic characteristics of ancient marine sediments with some demonstrated sensitivity to redox conditions in the modern ocean. These proxies range in sensitivity from local to global and qualitative to quantitative. Quantitative, global tracers of ancient marine redox conditions are of great interest, since – in the best case – they enable conclusions to be drawn about the global ocean when sampling a single locality. Using elements that have long marine residence times and thus are well-mixed in the ocean (e.g., Mo, Tl, U), a few global tracers have been developed, each with particular strengths and limitations. Among these is the uranium "stable" isotope ratio ($^{238}\text{U}/^{235}\text{U}$, expressed in delta notation as $\delta^{238}\text{U}$), which in the two decades since natural $\delta^{238}\text{U}$ variability was first analyti-

* Corresponding author.

E-mail address: mkipp@caltech.edu (M.A. Kipp).

cally resolved (Stirling et al., 2007; Weyer et al., 2008) has become one of the most widely-used proxies for global assessments of marine anoxia in deep time (e.g., Tissot and Dauphas, 2015; Lau et al., 2019; Zhang et al., 2020b).

In Earth's surface environment, uranium (U) exists in two main oxidation states: soluble U^{6+} that behaves conservatively in the modern ocean (i.e., [U] varies linearly with salinity, Ku et al., 1977; Owens et al., 2011), and insoluble U^{4+} . Because the mean oceanic residence time of U ($\tau \sim 400$ kyr; Ku et al., 1977; Dunk et al., 2002) is much longer than the global ocean mixing time (1–2 kyr; Broecker and Peng, 1982), the salinity-normalized seawater composition is homogeneous with regard to both U concentration ($[U]_{sw} = 3.24 \pm 0.07$ ng/g,¹ for a salinity of 35 g/L, Chen et al., 1986) and isotopic ratio ($\delta^{238}U_{sw} = -0.39 \pm 0.02\%$; Tissot and Dauphas, 2015). As U inputs to the ocean are dominated by continental weathering, with an isotopic composition identical to that of the continental crust ($-0.30 \pm 0.04\%$; Tissot and Dauphas, 2015; Andersen et al., 2016), changes in $\delta^{238}U_{sw}$ are typically thought to be controlled by the isotopic fractionation associated with U removal into different oceanic sinks.

In particular, U removal via reductive immobilization (i.e., from U^{6+} in dissolved uranyl-carbonate complexes to U^{4+} in uraninite, non-crystalline U^{4+} phases, or organic matter complexes; see e.g., reviews in Lau et al., 2019; Zhang et al., 2020b) in anoxic and/or euxinic (anoxic + sulfidic) settings results in preferential incorporation of ^{238}U in sediments. The magnitude of isotopic fractionation observed in these settings is typically $\sim +0.6\%$ relative to seawater (Andersen et al., 2014; Holmden et al., 2015; Rolison et al., 2017). This is roughly half of the estimated intrinsic fractionation factor of $\sim +1.3\%$ predicted by nuclear field shift calculations (Bigeleisen, 1996) and observed in redox reaction experiments (Fujii et al., 1989; Nomura et al., 1996; Brown et al., 2018), and is thought (e.g., Andersen et al., 2014; Lau et al., 2020; Zhang et al., 2020b) to reflect the typical twofold reduction in expressed fractionation in diffusion-limited settings such as sediment porewaters (e.g., Clark and Johnson, 2008). Although the uniformity of this expressed fractionation factor across environmental gradients has recently been questioned (e.g., Cole et al., 2020; Lau et al., 2020; Zhang et al., 2020b), including suggestions that anoxic, but non-euxinic settings feature smaller expressed isotope fractionation, by considering here this high-end value of $+0.6\%$ we are effectively considering the abundance of environments such as those where this isotopic effect was observed (e.g., Black Sea, Saanich Inlet). One can add complexity to the mass balance by considering intermediate (e.g., “ferruginous”, “hypoxic”) sinks with intermediate scavenging rates and isotopic fractionations. However, to best compare our results to previous work, we here utilize a simplified mass balance in which all other (“non-anoxic”) U sinks in aggregate are presumed to impart a negligible isotopic fractionation (Lau et al., 2016; Zhang et al., 2020b). In this framework, the waxing and waning of the anoxic sink dictates the $\delta^{238}U_{sw}$ value: when the anoxic sink is larger, more ^{238}U is scavenged from seawater, causing $\delta^{238}U_{sw}$ to become more negative (Fig. 1, Section 2.1).

Given the sensitivity of $\delta^{238}U_{sw}$ to the size of the anoxic sink, any archive that records $\delta^{238}U_{sw}$ in the past can be used to quantify, through isotope mass balance, fluctuations in seafloor anoxia. While black shales (Asael et al., 2013; Kendall et al., 2015; Brüske et al., 2020) and ferromanganese crusts (Goto et al., 2014; Wang et al., 2016) have been investigated, both archives are isotopically fractionated relative to seawater, and most studies have instead targeted carbonates (reviewed in Lau et al., 2019; Zhang et al., 2020b). Indeed, early work showed that modern primary carbonate

precipitates have $\delta^{238}U$ values identical to that of modern seawater (Stirling et al., 2007; Weyer et al., 2008; Romaniello et al., 2013; Andersen et al., 2014; Tissot and Dauphas, 2015). Further work confirmed that abiotic carbonate precipitation leads to minor isotopic fractionation ($\sim 0.1\%$; Chen et al., 2017), and that many biological carbonate precipitates have even smaller offsets relative to $\delta^{238}U_{sw}$ (Chen et al., 2018b; Tissot et al., 2018). While diagenetic modification is always a concern in paleo-redox studies (Romaniello et al., 2013; Chen et al., 2018a; Tissot et al., 2018), and will be further discussed in Section 4.2, to a first-order carbonates hold great potential as archives of $\delta^{238}U_{sw}$ on geological timescales.

If one assumes that carbonate sediments (or any other geological archives for that matter) indeed record $\delta^{238}U_{sw}$ on geological timescales, the remaining question is how to best relate these $\delta^{238}U$ data to the amount of U sequestered in anoxic sediments and, ultimately, to the most likely extent of seafloor anoxia. To date, all studies have taken one of two approaches: (i) the marine U isotope mass balance is assumed to be at steady state, in which case each $\delta^{238}U$ value (or the mean of the entire dataset) can be directly equated to the extent of seafloor anoxia (e.g., Bartlett et al., 2018; Zhang et al., 2018), or (ii) dynamic models are used to simulate transient perturbations to the marine U cycle (e.g., Jost et al., 2017; Lau et al., 2017; Clarkson et al., 2018, 2021;). In the latter case, studies have so far used only illustrative forward model outputs to assess marine anoxia using $\delta^{238}U$ datasets.

As will be discussed below, these approaches can result in inaccurate assessments of marine anoxia because of the dynamic nature of $\delta^{238}U$ trends on up to Myr timescales. Here we show that these shortcomings can be readily addressed using an inverse modeling framework, as it quantifies the fit of forward model runs to the data, therefore allowing (i) determination of the best fit for a particular dataset, and thus, (ii) consistent comparison of trends between datasets. We first walk through a few possible ways to handle this inverse modeling, before describing an accessible, robust method for assessing $\delta^{238}U$ and anoxia trends using a Markov Chain Monte Carlo (MCMC) approach. We demonstrate the utility of this method using published datasets, finding that our inverse model can in all cases describe the trajectory of the data, but in some cases implies seafloor anoxia trajectories that are different from those inferred using steady state assumptions or illustrative forward model runs. Overall, we find that this MCMC approach (available at www.github.com/m-kipp/d238U-inverse-model) will be a useful way for future applications of the $\delta^{238}U$ proxy to rigorously quantify seafloor anoxia and robustly compare trends between different datasets in order to determine the magnitude of ocean anoxia during events in Earth history.

2. Methods

2.1. Uranium isotope mass balance

Fig. 1a shows a simplified marine U budget used in many paleo-redox studies (e.g., Brennecke et al., 2011; Clarkson et al., 2018; Lau et al., 2016; Zhang et al., 2020a). It considers rivers as the sole U input to the ocean (J_{riv}), and two sinks: anoxic sediments (J_{anox}) and all other sediments (J_{other}). Although more complex formulations can be used to account for more than two sinks, we follow most previous work in using this simplifying assumption to derive first-order redox constraints. In this framework, the change in the U inventory of the ocean (N_{sw}) over time can be written as:

$$\frac{d(N_{sw})}{dt} = J_{riv} - J_{anox} - J_{other} \quad (1)$$

Accounting for the isotopic composition of riverine inputs ($\delta^{238}U_{riv}$) and the isotopic fractionation associated with U burial in anoxic

¹ Here and throughout the manuscript, errors are reported as 2σ unless otherwise stated.

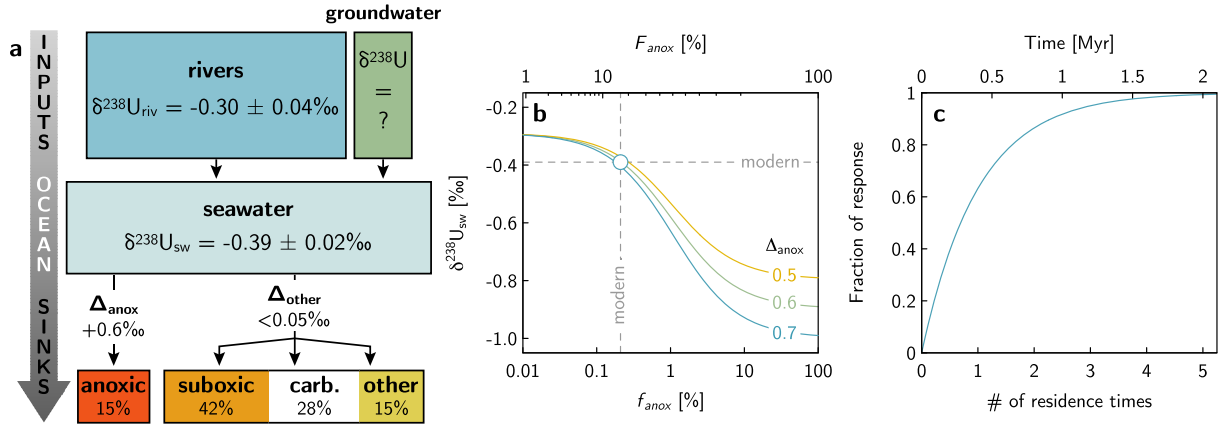


Fig. 1. (a) Simplified U isotope mass balance. Based on data and frameworks presented in Dunk et al. (2002), Andersen et al. (2014, 2016), Tissot and Dauphas (2015), and Lau et al. (2016). See Section 2.1 for details. (b) Steady state $\delta^{238}\text{U}_{\text{sw}}$ as a function of ocean anoxia (expressed as the proportion of U burial in anoxic sediments, F_{anox} , top x-axis, or seafloor area covered by anoxic sediments, f_{anox} , bottom x-axis). At steady state, there is little difference in $\delta^{238}\text{U}_{\text{sw}}$ between oceans with ~10% versus ~100% of the seafloor anoxic. Adapted from Lau et al. (2017). (c) Fraction of oceanic response to a step-wise perturbation as a function of time.

(Δ_{anox}) and non-anoxic (Δ_{other}) environments, changes in $\delta^{238}\text{U}_{\text{sw}}$ are then described (as in Lau et al., 2016) by:

$$\frac{d(\delta^{238}\text{U}_{\text{sw}})}{dt} = (J_{\text{riv}} * (\delta^{238}\text{U}_{\text{riv}} - \delta^{238}\text{U}_{\text{sw}}) - J_{\text{anox}} * \Delta_{\text{anox}} - J_{\text{other}} * \Delta_{\text{other}}) / N_{\text{sw}} \quad (2)$$

Quantification of oceanic anoxia can then be done in two ways: (i) as the proportion of U sequestered in the anoxic sink [$F_{\text{anox}} = J_{\text{anox}} / (J_{\text{anox}} + J_{\text{other}})$], or (ii) as the fraction of the seafloor that is anoxic (f_{anox}). We note that there is some inconsistency in the literature in the use of the term f_{anox} . We recommend that the notations adopted here – which follow most previous work – be systematically used in future studies to avoid confusion. Uranium removal fluxes can be related to seafloor area as:

$$J_{\text{anox}} = N_{\text{sw}} * K_{\text{anox}} * f_{\text{anox}} \quad (3)$$

$$J_{\text{other}} = N_{\text{sw}} * K_{\text{other}} * (1 - f_{\text{anox}}) \quad (4)$$

where K_{anox} and K_{other} are rate constants that describe the efficiency of U burial in anoxic and non-anoxic sediments, respectively (as in Lau et al., 2016). Combining Equations (2), (3) and (4) allows $\delta^{238}\text{U}_{\text{sw}}$ to be cast as a function of f_{anox} .

For the sake of clarity and completeness, we walk through a determination of the modern extent of seafloor anoxia using the latest understanding of marine $\delta^{238}\text{U}$ systematics. The first step is determining the modern U flux into anoxic and non-anoxic sinks. Assuming the modern ocean is at steady state, Equation (2) becomes:

$$J_{\text{riv}} * (\delta^{238}\text{U}_{\text{riv}} - \delta^{238}\text{U}_{\text{sw}}) = J_{\text{anox}} * \Delta_{\text{anox}} + J_{\text{other}} * \Delta_{\text{other}} \quad (5)$$

Using published estimates (Appendix A) for J_{riv} , $\delta^{238}\text{U}_{\text{riv}}$, $\delta^{238}\text{U}_{\text{sw}}$, Δ_{anox} , and Δ_{other} , with Monte Carlo propagation of uncertainty, we obtain an absolute U flux into anoxic sediments (J_{anox}) of $0.0063^{+0.0035}_{-0.0028}$ Gmol U yr⁻¹, corresponding to $F_{\text{anox}} = 15^{+8}_{-7}\%$. By mass balance, J_{other} must account for the other ~85% of U burial.

These mass fluxes can be equated to the extent of seafloor anoxia (Eq. (3), (4)). Here, one can either (i) assume that the modern f_{anox} value is known (i.e., using estimates from prior studies, e.g., Veeh, 1967; Bertine and Turekian, 1973), and then calculate the necessary rate constants, or (ii) constrain the rate constants with observations of U concentrations in modern sediments and overlying water, and then calculate a modern f_{anox} value. We opt for the latter approach, as it allows us to confirm that our oceanic

U budget indeed produces a reasonable, independent estimate of seafloor anoxia in the present day.

Areal U scavenging rates were determined in modern anoxic settings by Dunk et al. (2002). We can incorporate these constraints by writing the anoxic U burial flux as:

$$J_{\text{anox}} = R_{\text{anox}} * A_{\text{anox}} = R_{\text{anox}} * A_{\text{ocean}} * f_{\text{anox}} \quad (6)$$

where R_{anox} is the areal scavenging rate of U in modern anoxic sediments ($9.2 \mu\text{mol m}^{-2} \text{yr}^{-1}$, range 4.6 to 13.8; Dunk et al., 2002) and A_{anox} is the total area of anoxic seafloor in the modern ocean, which is the product of f_{anox} and the total area of the global seafloor, A_{ocean} ($3.6 \times 10^{14} \text{ m}^2$; Turekian, 1969). It follows from Equations (3) and (6) that

$$K_{\text{anox}} = \frac{R_{\text{anox}} * A_{\text{ocean}}}{N_{\text{sw}}} \quad (7)$$

Solving using the modern U inventory (N_{sw}) of $19,000 \pm 1,200$ Gmol U (Dunk et al., 2002), with propagation of uncertainties on R_{anox} and N_{sw} , gives a K_{anox} value of $1.74^{+0.68}_{-0.63} \times 10^{-4} \text{ yr}^{-1}$. This equates (Eq. (3)) to a modern f_{anox} value of $0.19^{+0.11}_{-0.05}\%$, similar to an earlier estimate of $0.21 \pm 0.09\%$ (Tissot and Dauphas, 2015) and in agreement with previous determinations based on U and Mo elemental mass balances (0.3%, Veeh, 1967; 0.23%, Bertine and Turekian, 1973). By mass balance, the corresponding K_{other} value is $1.88^{+0.22}_{-0.17} \times 10^{-6} \text{ yr}^{-1}$.

2.2. Forward modeling

With the above equations in place, we can explore the relationship between anoxia (f_{anox}) and $\delta^{238}\text{U}_{\text{sw}}$. Let us begin with a steady state system. In this case, there is no change in amount or isotopic composition of U in seawater with time (LHS in Eq. (1) and (2) equal to zero), and using the parameters in Appendix A, one can calculate $\delta^{238}\text{U}_{\text{sw}}$ as a function of f_{anox} (Fig. 1b). Because each f_{anox} value corresponds to a single $\delta^{238}\text{U}_{\text{sw}}$ value (i.e., a bijective function), this framework allows a straightforward back calculation of marine anoxia: locate the position on the plot matching the measured $\delta^{238}\text{U}_{\text{sw}}$ value and retrieve the corresponding f_{anox} value.

This approach is only as valid as the steady state assumption. One way to assess the validity of this assumption is to follow Holland (1978) and calculate the response time of the marine inventory of a given element to a single step-wise perturbation (see Appendix B). Doing so reveals that the system only comes within

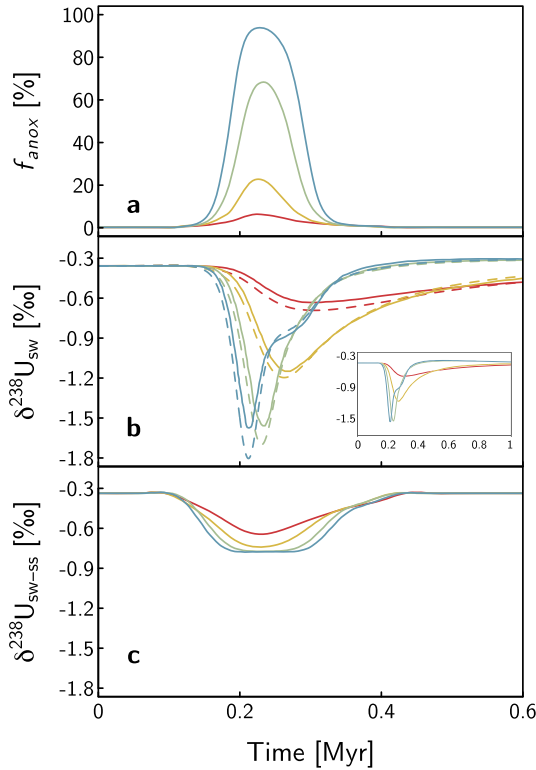


Fig. 2. Forward model runs under different f_{anox} trajectories (a). Solid lines in (b) denote outputs using model equations in Section 2.1; dashed lines denote outputs of Zhang et al. (2020b). Inset in (b) shows recovery toward steady state value over 1 Myr. Solid lines in (c) denote $\delta^{238}\text{U}_{sw}$ value if the system were in steady state at the current f_{anox} value in (a). Note that $\delta^{238}\text{U}_{sw}$ can exceed the lowest possible steady state value during large, rapid perturbations.

1% of the new steady state after 5 residence times (Fig. 1c). Given its modern marine residence time of ~ 400 kyr (Ku et al., 1977; Dunk et al., 2002), this means ~ 2 Myr must pass for the marine U inventory to reach a new steady state following a step perturbation. This implies that for datasets with < 2 Myr between $\delta^{238}\text{U}$ datapoints (as is the case in most studies using the $\delta^{238}\text{U}$ proxy), a dynamic model must be used to accurately recover f_{anox} trends.

When the steady state assumption is invalid and a dynamic model is needed, a simple yet effective approach is to parameterize a secular trend in f_{anox} and determine whether the resulting $\delta^{238}\text{U}_{sw}$ model trend resembles the data (e.g., Clarkson et al., 2018). For instance, and to demonstrate that our formulation of U isotope mass balance agrees with prior work, we modeled four scenarios (from Zhang et al. (2020b), see their Fig. 9g), each representing a single pulse of anoxia of varying magnitude (Fig. 2a). The resulting $\delta^{238}\text{U}_{sw}$ trends (Fig. 2b, solid lines) match closely those of Zhang et al. (2020b) (Fig. 2b, dashed lines), indicating that our model agrees with previous formulations (the residual discrepancy is due to the fact that Zhang et al. modify J_{riv} as a function of $p\text{CO}_2$, which changes as a function of f_{anox}). Importantly, this exercise also demonstrates the inappropriateness of the steady state assumption: $\delta^{238}\text{U}_{sw}$ values calculated assuming steady state at each time point (Fig. 2c) are very different from those in the dynamic model (Fig. 2b) because the system has not had time to fully respond to the forcing (and because during rapid, large perturbations, $\delta^{238}\text{U}_{sw}$ can even exceed the minimum steady state value as the ocean gets quickly depleted of uranium).

One could in theory “hand-tune” this dynamic forward model to re-create $\delta^{238}\text{U}_{sw}$ trends in published datasets. This process would, however, not only be extremely time consuming, but also hindered by significant sources of inaccuracy. First, $\delta^{238}\text{U}_{sw}$ responds very little to changes in $f_{anox} > 10\%$ and $< 0.1\%$ (Figs. 1b, 2).

This means that even in dynamic models, there is little basis for a forward modeler to discern between a transient anoxic pulse of $\sim 60\%$ versus $\sim 90\%$ (the two would give rise to quite similar $\delta^{238}\text{U}_{sw}$ trajectories; Fig. 2b). Second, many $\delta^{238}\text{U}$ datasets are generated using samples that have already been studied for the identification of ocean anoxia. This means that there are known intervals that are thought to have been anoxic; coupled with the insensitivity at high seafloor anoxia noted above, this could lead to bias in inferring moderate vs. severe anoxia when data are ambiguous. Lastly, when using only illustrative forward model outputs, it is unclear whether data points that fall off the simulated $\delta^{238}\text{U}_{sw}$ trajectory can be explained via global U isotope mass balance. This is important because in some cases, rapid, positive $\delta^{238}\text{U}$ excursions cannot be achieved by redox-influenced U isotope mass balance alone (e.g., Clarkson et al., 2018); in these cases, some of the signal must be attributed to diagenesis, but forward models alone leave this distinction ambiguous. For all of these reasons, here we utilize Bayesian inverse analysis – which has demonstrated utility in paleoclimate and isotope geochemistry studies at a range of spatial and temporal scales (e.g., Tierney et al., 2019; Bowen et al., 2020; Krissansen-Totton et al., 2021) – to optimize model fits, propagate parameter uncertainties, and distinguish between competing effects on isotopic data.

2.3. Inverse modeling

2.3.1. A Markov Chain Monte Carlo approach

Here, the term “inverse” modeling simply means that a quantitative framework is used to assess the fit of dynamic forward model outputs to a $\delta^{238}\text{U}$ dataset. Doing so allows optimization of the fit to a $\delta^{238}\text{U}$ dataset and recovery of the corresponding f_{anox} trend. There are multiple ways to approach this problem, which we walk through here to find an efficient and effective solution for the $\delta^{238}\text{U}$ proxy.

First is the question of how to describe the temporal evolution of f_{anox} . There are two approaches: *numerical* and *analytical*. In the *numerical* approach, one simply needs to assign an f_{anox} value at each time point in a model run. This approach is flexible in that one need not describe a functional form for the f_{anox} history; instead one can simply “draw” a curve that looks appropriate. While this enables quick “hand-tuning” to fit a dataset, the numerical approach is computationally costly when trying to optimize the fit, because the possible permutations of f_{anox} histories quickly become very large (see Appendix C). This burden can be reduced by decreasing the temporal resolution in parts of the dataset that appear invariant; however, doing so could lead to unintentional bias influencing data interpretation.

In contrast, in the *analytical* approach, one can optimize the fit by tuning only a few parameters. Here, we consider the temporal evolution of f_{anox} as a Fourier series:

$$\frac{d(f_{anox})}{dt} = \sum_{i=1}^m a_i * \sin(y_i * t) + b_i * \cos(z_i * t) \quad (8)$$

where a and b set the amplitude, and y and z set the period of oscillations. To describe complex f_{anox} trajectories, several sets of sine and cosine terms can be included; in practice, we find $m \leq 10$ to be sufficient in most cases. In this framework, we only need to optimize a_i , b_i , y_i and z_i to arrive at the best fit to a $\delta^{238}\text{U}$ dataset – a more computationally tractable problem – which we consider below.

The next question is that of the “cost” function used to describe the fit of the model ($\delta^{238}\text{U}_{mod}$) to the data ($\delta^{238}\text{U}_{obs}$). Here we use the negative log-likelihood (NLL) to quantify model fit (MacKay, 2003), which is calculated as:

$$\text{NLL} = \sum_{i=1}^n \frac{(\delta^{238}\text{U}_{\text{obs},i} - \delta^{238}\text{U}_{\text{mod},i})^2}{2\sigma_i^2} + \ln(\sqrt{2\pi}\sigma_i) \quad (9)$$

where n is the number of data points and σ is the uncertainty on each point. The best fit model has the lowest NLL. We note that other inverse modelers sometimes use different cost functions, e.g., the Bayesian Information Criterion, which expands on the log-likelihood by penalizing models with more parameters; for simplicity we opt to utilize the NLL and leave the dimensionality of the model (m value) up to the user.

There are two schools of thought about how to optimize model fit (i.e., minimize NLL): *Frequentist* and *Bayesian*. The *Frequentist* approach is Maximum Likelihood Estimation (MLE). There are MLE functions available in scientific programming languages that often use derivative-based solvers to minimize the cost function. This is computationally efficient, but prone to getting stuck in local minima. While some approaches can circumvent this issue (e.g., “basin-hopping” algorithms), the problem becomes increasingly severe as the number of free parameters increases. In practice, this makes MLE a poor approach for optimizing the fit to $\delta^{238}\text{U}$ datasets, as we need many free parameters (large m values in Eq. (8)) to simulate all plausible trends.

The *Bayesian* approach to optimizing model fit is Maximum A Posteriori estimation (MAP). The chief difference between MAP and MLE is that in MAP each parameter can be assigned a “prior” distribution. This is essentially saying that we have some knowledge about a parameter’s probability distribution when beginning the assessment. MLE can be considered a special case of MAP with uninformative “uniform” priors.

Importantly, MLE and MAP provide “point” estimates of the parameters that yield the best model fit. While uncertainties can be calculated for these point estimates, we are left without information about the probability distribution for each parameter. A widely-used approach for obtaining these probability distributions is *Bayesian inference*. Bayesian inference is often used to expand on MAP and provide *posterior probability distributions* for each model parameter. This is particularly useful for $\delta^{238}\text{U}$ datasets, where we want to constrain a robust confidence interval for f_{anox} . A common computational approach for Bayesian inference is the Markov Chain Monte Carlo (MCMC) method. The basic idea of MCMC is that by running many model iterations (10^5 - 10^6) while varying parameter values within a prescribed range, the model cost can be systematically minimized, approaching the “best fit” model. If the MCMC routine is run long enough that it can converge on the best solution, the posterior distribution of each parameter will be proportional to its true probability distribution (meaning we will have numerically approached a solution without analytically solving a very computationally-costly problem). Current MCMC formulations typically utilize some version of the “Metropolis-Hastings algorithm” (Metropolis et al., 1953; Hastings, 1970), which provides an efficient method of preferentially “accepting” model runs that achieve a better fit, allowing the algorithm to converge on the optimal solution. Here we implement an adaptive MCMC routine (Haario et al., 2001) using the FME package in R (Soetaert and Petzoldt, 2010).

Our MCMC workflow is shown in Fig. 3. The model begins by calculating the NLL (Eq. (9)) using values for each free parameter (f_{anox} , a_i , b_i , y_i , z_i) that yield the modern steady state, i.e., $\delta^{238}\text{U}_{\text{sw}} = -0.39\%$, $f_{\text{anox}} = 0.2\%$, and $\frac{d(f_{\text{anox}})}{dt} = 0$ (Fig. 3a). The algorithm then takes a “step” (Fig. 3b) by randomly selecting a new value for each of the free parameters from within a prescribed range (p_{min} to p_{max}). The forward model is then run again and the NLL is calculated. If the fit is better than in the previous step ($\text{NLL}_{\text{step},i} < \text{NLL}_{\text{step},i-1}$), the new parameter values replace the old ones; if the fit is worse, the new parameter values are rejected with probability P , where $P = 1 - \exp(\text{NLL}_{\text{step},i} - \text{NLL}_{\text{step},i-1})$. This

process is repeated for many steps ($\sim 10^5$ - 10^6 , here n_{iter}), and the evolution of each parameter during these n_{iter} steps describes the path of what we call a “walker” (Fig. 3c). The first $n_{\text{burn-in}}$ steps of this “random walk” through parameter space are discarded, which allows the walker to “forget” where it started, meaning it is no longer biased by our choice of starting values. The rest of the walk then becomes the posterior distribution for a given parameter.

A single walker will eventually converge on the target distribution for each parameter, but as we want to propagate uncertainty on terms typically held constant (e.g., Δ_{anox} , $\delta^{238}\text{U}_{\text{riv}}$) in the mass balance, we deploy many walkers ($\sim 10^2$, here n_{walker}) that sample from the uncertainty ranges for these “constants” (Appendix A) and concatenate their posterior parameter distributions (PPDs) at the end (i.e., the concatenated PPD for each parameter will contain $n_{\text{walker}} * (n_{\text{iter}} - n_{\text{burn-in}})$ values). Armed with these PPDs, a Monte Carlo simulation is done whereby the concatenated PPDs are randomly sampled many times ($\sim 10^3$, here n_{sens}) while running the forward model, thus yielding optimized $\delta^{238}\text{U}_{\text{sw}}$ and f_{anox} trends (Fig. 3d), as well as confidence intervals (here outputs are shown as the median in solid lines and the 16th to 84th percentile in shaded regions, as distributions are non-Gaussian; Fig. 3e).

While this approach works in theory, there are, in practice, many issues to consider in order to ensure that an MCMC routine is converging on a solution. This includes tuning the number of steps taken by each walker (n_{iter}), the number of terms in the f_{anox} expression (m), the range to be explored for each parameter (p_{min} , p_{max}), the size of random steps through that parameter space (p_{step}), and whether the size of those steps is updated as the walkers progress (p_{update}). As datasets differ in the number of data points, density of points, and magnitude of trends, the MCMC routine must be tuned for each individual case, with convergence demonstrated by approaching a minimized NLL value, among other criteria. A thorough discussion of convergence tests can be found in Appendix C.

2.3.2. Model calibration and sensitivity tests

To demonstrate that our MCMC approach can accurately recover trends using a “known” test case, we generated a synthetic dataset by selecting 40 time points from a modeled trend in Fig. 2b, with higher density around the perturbation than before and after (simulating typical data density in $\delta^{238}\text{U}$ studies). The modeled $\delta^{238}\text{U}_{\text{sw}}$ value at each time point became the “measured” value, to which we assigned a conservative analytical uncertainty ($\pm 0.06\%$, 1σ). We then fed this dataset to the MCMC routine and compared the recovered f_{anox} trend to the trend that was used to force the model (Fig. 2a). Doing so revealed that the MCMC approach accurately reconstructed the f_{anox} trend (Fig. 4a).

Input parameter uncertainties: In addition to converging on a best fit trend under certain model assumptions (e.g., parameter values for modern mass balance), we aimed to allow our model to incorporate the uncertainty on those assumptions. To do so, we randomly sampled from the published or herein calculated parameter ranges for Δ_{anox} , Δ_{other} , $\delta^{238}\text{U}_{\text{riv}}$, K_{anox} , and K_{other} (values in Appendix A) when executing the MCMC routine. In this case, the model still found the correct f_{anox} and N_{sw} trends, with only slightly greater uncertainty (Fig. 4b). Importantly, this demonstrates that our model outputs are not unduly reliant on assumptions about modern U isotope mass balance.

Analytical noise: To further demonstrate the robustness of the MCMC approach, we made the recovery more difficult by introducing random noise (e.g., due to analytical inaccuracy) to the data points by adding a random number from a normal distribution with mean = 0% and $1\sigma = 0.06\%$. In this case, the MCMC routine still found the correct trend (Fig. 4c), even while propagating the uncertainty on mass balance parameters, demonstrating that

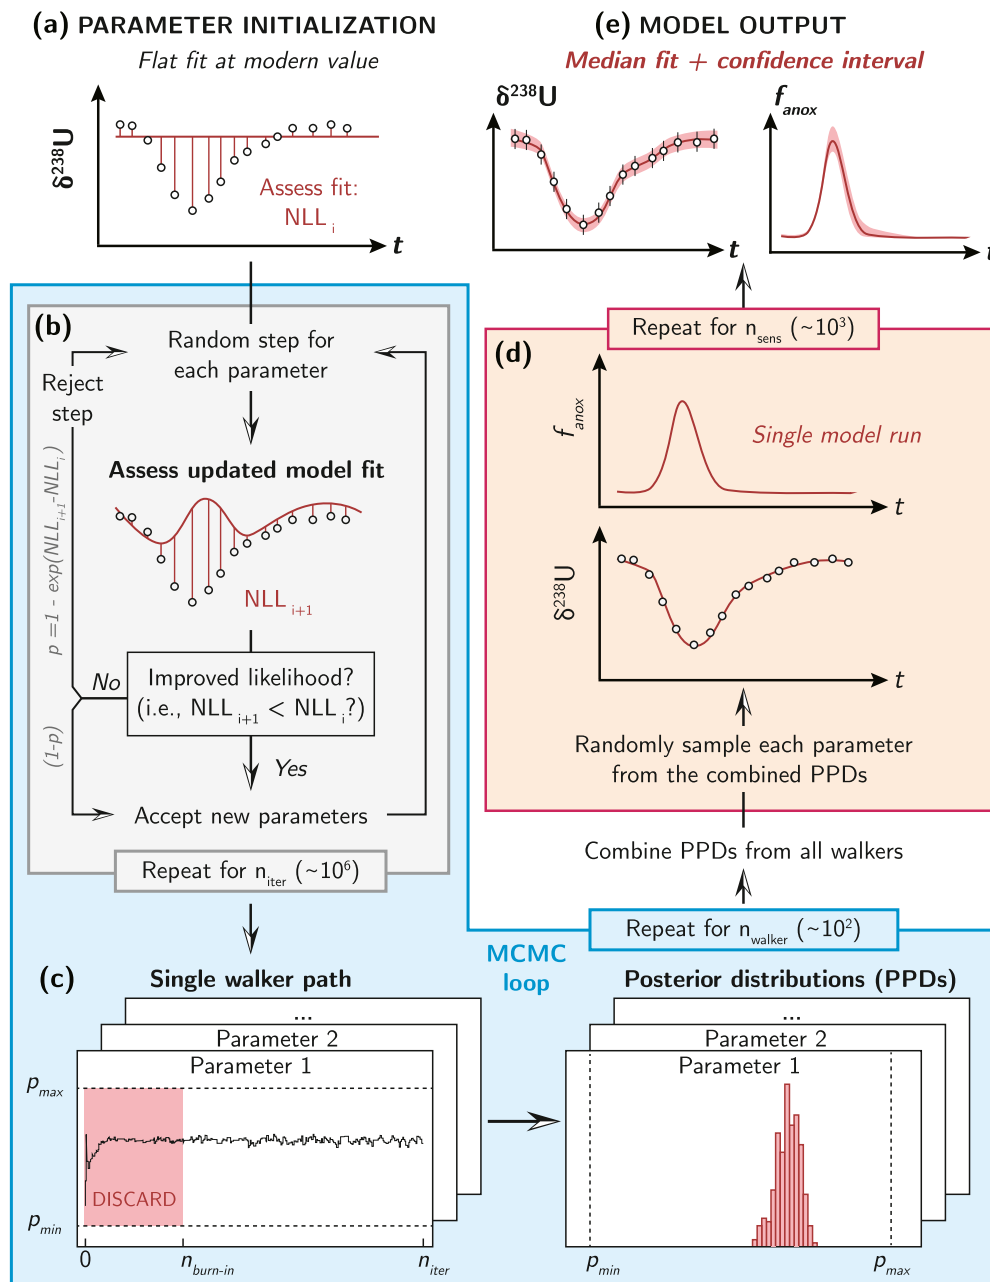


Fig. 3. Schematic representation of Markov Chain Monte Carlo routine employed in this study.

the model can fit trends through noisy datasets (we also experimented with more complicated test cases; see **Appendix C**).

Timescales: We next aimed to demonstrate the timescales on which this inverse model can be robustly used. The ocean mixing time (1-2 kyr) sets a lower limit here, since the assumption of a well-mixed ocean is invalid on timeframes less than several ocean mixing times. We conservatively explored a 100 kyr lower limit here (i.e., data points every ~ 2 kyr) and explored the ability of the model to recover a simple trend at increasingly longer timescales. We began with the f_{anox} trend shown in **Fig. 4** and applied it to shorter and longer time intervals. In doing so, we see that the inverse model properly identifies trends in all cases (**Fig. 5**). At longer timescales (**Fig. 5d**), while the inverse model can fit the trend, the steady state calculation becomes appropriate – and is much faster than running an MCMC routine. We therefore conclude that this inverse model is best applied to datasets that range from $\sim 10^{5-7}$

yr in duration, which in fact encompasses most work published to date using the $\delta^{238}\text{U}$ proxy.

3. Application to published datasets

To demonstrate the model's ability to reconstruct f_{anox} trends in various contexts and to conduct a robust comparison of marine anoxia across events in Earth's history, we applied our inverse model to eight published carbonate $\delta^{238}\text{U}$ datasets in time intervals ranging from the Ordovician to the Eocene (**Fig. 6**). These datasets all capture dynamic perturbations to the marine U cycle caused by expansions of marine anoxia, including mass extinction events (**Fig. 6a-f**), Cretaceous Ocean Anoxic Event 2 (**Fig. 6g**) and the hyperthermal event at the Paleocene-Eocene Thermal Maximum (**Fig. 6h**). The MCMC parameters required for efficient convergence were different for each case, highlighting the importance of conducting an individual assessment of convergence (see **Ap-**

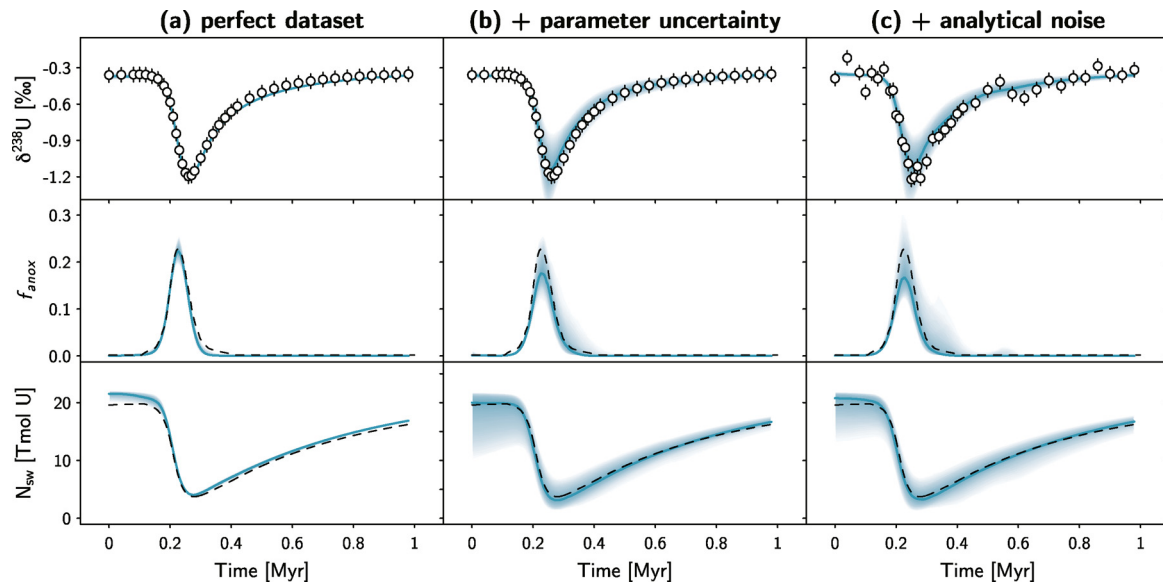


Fig. 4. MCMC accuracy tests. Three scenarios are depicted: (a) dataset exactly equivalent to forward model output, (b) same dataset with Monte Carlo sampling of uncertainties on Δ_{anox} , $\delta^{238}\text{U}_{\text{riv}}$, K_{anox} and K_{other} , (c) Monte Carlo error propagation plus random analytical noise introduced to dataset. Dashed lines in middle and lower panels denote forward model trajectory used to generate datapoints. In this and subsequent plots, the solid blue lines denote median MCMC outputs, and the 16th to 84th percentile window is shown in shading with increasing opacity toward the median. In all cases, the MCMC routine accurately recovers the correct f_{anox} trend and seawater U inventory (N_{sw}), even when parameter uncertainties are taken into account and analytical noise is introduced. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)




Fig. 5. Age range tests. The same magnitude of perturbation was spread over increasingly longer time intervals: (a) 100 kyr, (b) 1 Myr, (c) 10 Myr, (d) 100 Myr. Dashed lines denote model forcing used to generate “synthetic” data points shown in white in top panels. Inverse model recoveries of the $\delta^{238}\text{U}$, f_{anox} and N_{sw} trajectories are shown as the solid blue lines and shading. Red lines show calculated f_{anox} trajectory if steady state is assumed for each data point. The MCMC routine successfully recovers trends at timescales ranging from 10^5 – 10^8 yr.

pendix C). For the calculations in Fig. 6, we have taken authors’ assessments of diagenetic effects on $\delta^{238}\text{U}$ records at face value, meaning that we accepted their preferred values for $\delta^{238}\text{U}_{\text{sw}}$ as inferred from $\delta^{238}\text{U}_{\text{carb}}$ data. In Section 4.2 we will critically evaluate this assumption.

Previous inferences of f_{anox} vary in their resemblance to those recovered in the MCMC routine (e.g., Fig. 6a vs. 6e). This is not simply due to differences in modeling approaches used in previous studies; in both Fig. 6a (Bartlett et al., 2018) and Fig. 6e (Zhang et al., 2018), the authors used a steady state framework to calculate f_{anox} . Instead, this reflects (i) the poor ability of the $\delta^{238}\text{U}$ proxy to distinguish between high f_{anox} values at steady state (Fig. 1b), and (ii) the inappropriateness of the steady state assumption on

sub-Myr timescales. The latter is particularly evident in Fig. 6e, where the steady state assumption leads to gross overestimation of seafloor anoxia. In contrast, datasets across longer timespans (a few Myr) that were evaluated using a steady state inference of f_{anox} have better agreement between steady state calculations and the MCMC inversion (e.g., Fig. 6b).

Dynamic forward model outputs used to fit published datasets also varied in their match to the MCMC inversions. Some dynamic models match the MCMC inference quite well (Figs. 6d, f), while others differ substantially (Fig. 6g). In the latter case (Clarkson et al., 2018), the MCMC inversion provided a much better fit (lower NLL) to the dataset than the illustrative forward model. The previous inference of stronger anoxia may therefore have been biased

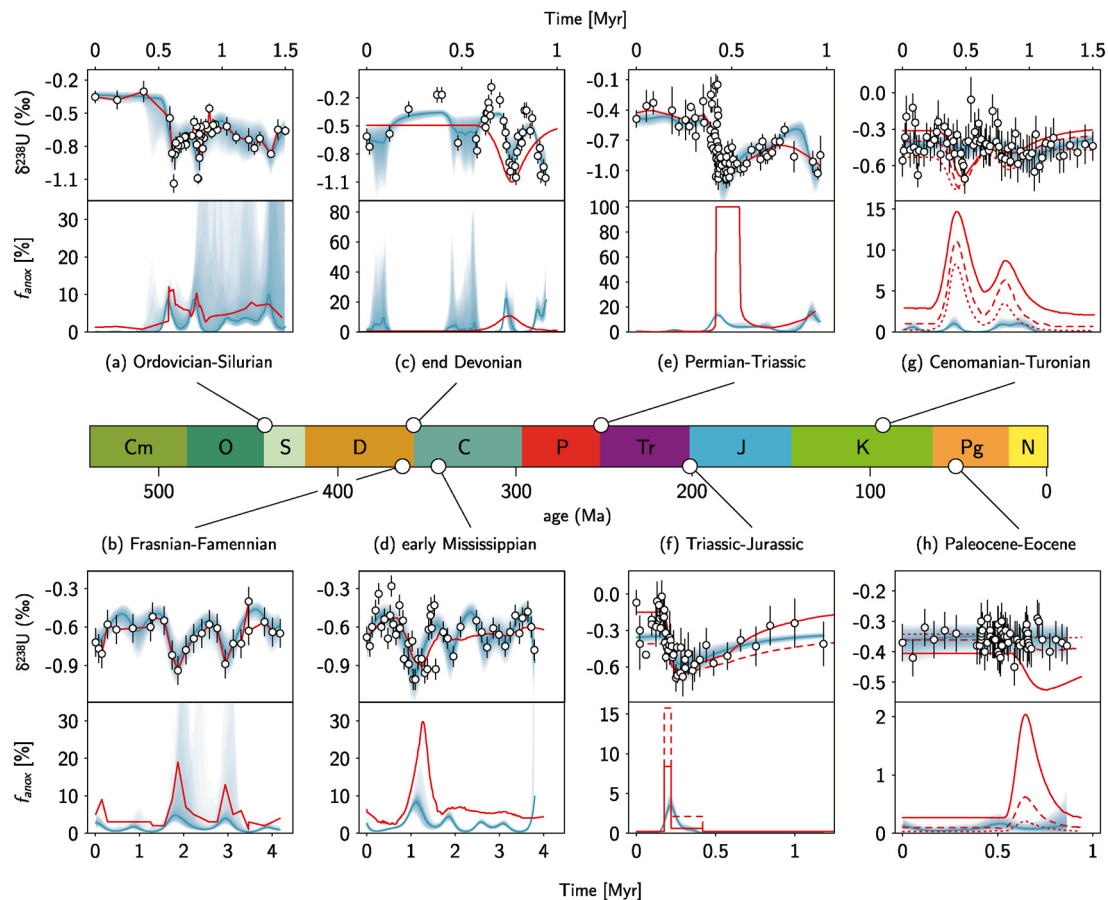


Fig. 6. Comparison of inverse model results and published forward models describing $\delta^{238}\text{U}$ and f_{anox} trends. Datasets: (a) Ordovician-Silurian boundary (Bartlett et al., 2018), (b) Frasnian-Famennian boundary (Song et al., 2017), (c) Hangenberg biotic crisis in latest Devonian (Zhang et al., 2020a), (d) Tournasian stage in the early Mississippian (Cheng et al., 2020), (e) Permian-Triassic boundary (Zhang et al., 2018 and references therein), (f) Triassic-Jurassic boundary (Jost et al., 2017), (g) Ocean Anoxic Event 2 at the Cenomanian-Turonian boundary in the middle Cretaceous (Clarkson et al., 2018), (h) Paleocene-Eocene Thermal Maximum (Clarkson et al., 2021). Red lines show $\delta^{238}\text{U}$ and f_{anox} trends inferred by the authors of the cited studies (dashed and dotted lines are used to denote additional published model scenarios where applicable).

by the prior knowledge of an anoxic interval in the stratigraphy.

We see here that the inverse model can successfully fit published datasets, allowing rigorous assessment of trends within and between events. Furthermore, the uncertainty on model parameters (*i.e.*, Δ_{anox} , Δ_{other} , $\delta^{238}\text{U}_{\text{riv}}$, K_{anox} , K_{other}) has been accounted for in these calculations. This all speaks to the robustness of the inversion technique in a range of situations. However, the MCMC approach has its own limitations. For instance, datasets with variable data density pose problems for the MCMC routine, with highly uncertain trajectories in large temporal gaps (Fig. 6c). The optimal solution to this problem is to obtain a uniformly dense dataset across a timeseries. In the event of unavoidable temporal heterogeneity, binning data into larger time intervals can alleviate some of the problem for the MCMC routine, but at the expense of smoothing out some of the variability in the record. On a case-by-case basis, an inverse modeler can decide the most appropriate way to handle the data.

4. Remaining limitations of the $\delta^{238}\text{U}$ proxy

Besides the finer details of tuning the MCMC approach to fit a dataset, there remain two foundational ways in which even the conservative assumptions of our modeling may still lead to inaccuracies in our reconstruction of f_{anox} . We discuss the magnitude of these effects below.

4.1. Extrapolation of rate constants

Propagating the uncertainty of parameter values (Appendix A) in our MCMC inversions allows us to account for imperfection in the modern estimation of (or natural variability in) fractionation factors, river input fluxes and composition, and rate constants. However, in the case of U burial rate constants, the uncertainty derived from modern oceanographic studies may not be representative of past environments that were quite different from today. As discussed in Section 2.1, we use rate constants that were derived from areal constraints on U burial in modern anoxic sediments (Dunk et al., 2002). Although these values are likely representative of U burial in modern anoxic and non-anoxic sediments, they may not be representative of ancient oceans that were strongly anoxic. This is because anoxic sediments today are found overwhelmingly in productive continental margin settings with high organic matter and/or sulfide burial rates, whereas oxygen-depleted pelagic settings (which would prevail in a fully-anoxic ocean) are likely to have much lower rates of organic carbon and sulfide burial, resulting in less pronounced redox-sensitive trace metal (including U) scavenging per surface area – and perhaps even muted isotopic fractionation (Cole et al., 2020).

Reinhard et al. (2013) noted this problem when modeling the effect of expanding ocean anoxia on chromium and molybdenum burial. To avoid extrapolating from continental margin-derived rate constants to pelagic settings, they scaled their rate constant for anoxic metal burial using an algorithm that coupled predicted or-




Fig. 7. Effect of lower K_{anox} value. Panel (a) shows the end-Permian dataset of Zhang et al. (2018) fitted with the base model (K, blue), a two-fold reduction in K_{anox} (K/2, green), five-fold reduction in K_{anox} (K/5, gold), and ten-fold reduction in K_{anox} (K/10, red). While all scenarios fit the dataset well (panel a) and imply similar (b) N_{sw} and (c) F_{anox} trajectories, panel (d) shows that lower K_{anox} values lead to inferences of greater f_{anox} (solid lines denote median output in all panels). The true K_{anox} value is likely within an order of magnitude of that used in the base model, since lower values imply an unrealistic extent of marine anoxia in the early Triassic.

ganic carbon export profiles to a modern bathymetric configuration. While each of those steps is subject to uncertainty, it is likely that such an approach provides a better first-order approximation of redox sensitive trace element burial in anoxic pelagic settings than simply using the margin-derived value.

Here we use a simple test to explore the sensitivity of our MCMC retrievals to inaccurate rate constants. We decreased the K_{anox} value by factors of 2 to 10 and re-ran the MCMC routine for the Permian-Triassic dataset of Zhang et al. (2018). As expected, while the inferred value of F_{anox} (amount of U sequestered in anoxic sediments) is similar in all cases, using lower K_{anox} values leads to inferences of higher f_{anox} (extent of seafloor anoxia, Fig. 7). Importantly, the extreme case of a 10-fold reduction in K_{anox} is likely a conservative estimate of the deviation from the “true” K_{anox} value, since that scenario invokes f_{anox} values that are unrealistically high (100%) for unrealistically long periods (e.g., early Triassic). We therefore conclude that our adopted K_{anox} value is within roughly an order of magnitude of the value for pelagic settings.

While this constraint on the legitimacy of the margin-based rate constant is helpful, it still leaves us with potentially large uncertainties in our reconstruction of f_{anox} using $\delta^{238}\text{U}$ data (Fig. 7). Future work with the $\delta^{238}\text{U}$ proxy would therefore benefit from a more detailed treatment of potential variability in rates of U scavenging in anoxic settings. In the simplest case, one could modify K_{anox} with a non-dimensional scaling factor to account for differences in rate constants across burial environments (e.g., Chen et

al., 2021). For datasets when a more detailed assessment of ocean bathymetry and export production is possible, an approach similar to that of Reinhard et al. (2013) would be more robust. Accounting for this effect will be critical for the $\delta^{238}\text{U}$ proxy to be truly useful in interrogating redox dynamics in strongly anoxic oceans. In lieu of a clear way around this source of uncertainty at present, we encourage the reporting of F_{anox} along with f_{anox} to facilitate inter-comparison of trends across datasets.

4.2. Identifying and accounting for diagenetic offsets

Beyond uncertainties in input parameters and differences in rates of U scavenging across marginal and pelagic sites, the largest remaining hurdle to the accurate inference of both F_{anox} and f_{anox} using $\delta^{238}\text{U}$ data is the isotopic offset between ancient carbonates and coeval seawater that can arise during diagenetic U addition and/or removal. As noted in Section 1, studies of recent carbonate sediments have documented sizable offsets between $\delta^{238}\text{U}_{\text{carb}}$ and $\delta^{238}\text{U}_{\text{sw}}$ in Bahamian drill cores (typically 0 to +0.6‰; Romaniello et al., 2013; Chen et al., 2018a; Tissot et al., 2018). The generally elevated U content of these sediments relative to primary biological carbonate precipitates suggests that this isotopic effect derives from U addition during diagenesis (Tissot et al., 2018). Unfortunately, no other geochemical proxies closely correlate with observed deviations from $\delta^{238}\text{U}_{\text{sw}}$ (Tissot et al., 2018), making the prospects for precisely correcting diagenetic offsets rather poor. When faced with obvious signs of diagenetic enrichment, many studies have therefore taken the simple approach of applying a constant diagenetic offset to their entire dataset (e.g., Song et al., 2017; Zhang et al., 2018, 2020a). Given that the observed offsets in Bahamian drill core archives <1.5 Myr old vary considerably ($+0.23 \pm 0.15\%$, 1σ ; Tissot et al., 2018), this is an insufficient remedy to addressing diagenesis in ancient $\delta^{238}\text{U}$ datasets.

Here we explore the effect of diagenesis on $\delta^{238}\text{U}$ records in two ways. First, we consider the identification of diagenesis in $\delta^{238}\text{U}$ datasets. While an extensive literature surrounds the petrographic and geochemical study of carbonate diagenesis (e.g., Fantle et al., 2020) – and petrographic screening as part of U isotope work is recommended (e.g., Hood et al., 2016) – here we simply consider the expression of diagenesis on bulk-rock $\delta^{238}\text{U}$ data. As diagenesis tends to enrich carbonate sediments in ^{238}U (by adding either detrital U with a continental $\delta^{238}\text{U}$ signature, or authigenic U that is enriched in ^{238}U during reductive precipitation), we can ask whether certain positive excursions in $\delta^{238}\text{U}$ datasets are diagenetic artifacts or primary redox signals.

We can begin with a simple screening criterion: carbonates with $\delta^{238}\text{U}$ values much greater than $\delta^{238}\text{U}_{\text{riv}}$ ($-0.30 \pm 0.04\%$) have likely been isotopically perturbed during diagenesis. We know this because the “global redox” framework for U isotope mass balance only includes a pathway for depleting seawater of ^{238}U (reductive immobilization); if seafloor anoxia were entirely absent, the framework of Fig. 1 would predict $\delta^{238}\text{U}_{\text{sw}} = \delta^{238}\text{U}_{\text{riv}}$. Estimating $\delta^{238}\text{U}_{\text{riv}}$ in deep time is difficult, and it is possible that different continental configurations or exposures of different lithologies could alter $\delta^{238}\text{U}_{\text{riv}}$ (e.g., Jost et al., 2017). There are also smaller U sinks (e.g. Fe-Mn oxide adsorption) that preferentially remove ^{235}U and could in theory generate a positive $\delta^{238}\text{U}_{\text{sw}}$ excursion, though it is perhaps unlikely that these sinks could generate a globally significant isotope effect. So while uncertain, $\delta^{238}\text{U} \gg -0.3\%$ in ancient carbonates can provide a qualitative, yet strong, hint that diagenesis ought to be investigated further.

A more nuanced way to identify diagenesis in $\delta^{238}\text{U}$ datasets is to consider the possible timescale of a global-redox-driven positive $\delta^{238}\text{U}$ excursion. Positive $\delta^{238}\text{U}_{\text{sw}}$ excursions occur when F_{anox} (and thus f_{anox}) decreases and the ocean re-fills with U of continental (riverine) composition ($\sim -0.3\%$). Because the riverine in-




Fig. 8. Forward model runs showing fastest possible positive $\delta^{238}\text{U}$ recoveries. Grey band denotes immediate return to modern anoxia after 0.1 Myr of expanding anoxic seafloor coverage; dashed line denotes modern $\delta^{238}\text{U}_{\text{sw}}$ value.

flux of U is small relative to the size of the modern U reservoir (**Appendix A**), these positive $\delta^{238}\text{U}_{\text{sw}}$ excursions occur quite slowly. The fastest possible positive redox-driven $\delta^{238}\text{U}_{\text{sw}}$ excursions would occur in the extreme scenario where widespread anoxia is followed by an immediate return to the modern anoxia value. As shown in **Fig. 8**, when the marine U reservoir is large (*i.e.*, $N_{\text{sw}} \approx 100\%$ of modern), it takes hundreds of kyr for $\delta^{238}\text{U}_{\text{sw}}$ to increase by $\sim 0.1\%$. In contrast, when the marine U reservoir is depleted ($N_{\text{sw}} < 10\%$ of modern), $\delta^{238}\text{U}_{\text{sw}}$ can increase by $>1\%$ in ~ 50 kyr.

Based on these observations, it appears that nearly every dataset in **Fig. 6** shows indications of diagenesis, either in the form of $\delta^{238}\text{U} > -0.3\%$, or by containing apparent positive $\delta^{238}\text{U}$ excursions too rapid to be redox-driven. This points to the ubiquity of diagenesis in shaping ancient carbonate archives. Some studies (*e.g.*, Clarkson et al., 2018; Cao et al., 2020) have constrained plausible diagenetic offsets by studying coeval pelagic and platform carbonates (where pelagic settings tend to have smaller/negligible diagenetic effects, as observed in Tissot et al., 2018; Clarkson et al., 2020). Ancient pelagic carbonates do indeed appear to more directly record the $\delta^{238}\text{U}_{\text{sw}}$ value without requiring diagenetic correction (**Fig. 6g, h**, Clarkson et al., 2018, 2021). However, diagenetic “noise” is still apparent in some pelagic data (**Fig. 6g**). For all of these reasons, it is imperative that diagenetic offsets be considered when attempting to quantify seafloor anoxia using ancient carbonate $\delta^{238}\text{U}$ values, particularly if comparing pelagic and platform datasets.

Our second aim is to quantify the possible effect of inaccurate diagenetic corrections on f_{anox} reconstructions. We consider a single dataset (platform carbonates from the Permian-Triassic boundary; Zhang et al., 2018 and references therein) through three diagenetic correction schemes, assuming (i) zero diagenetic offset (**Fig. 9a**), (ii) a uniform diagenetic offset of $+0.3\%$ (**Fig. 9b**; as in Zhang et al., 2018), and (iii) a random diagenetic offset drawn from the distribution observed in Bahamian carbonate sediments (**Fig. 9c**; $+0.23 \pm 0.15$, 1σ ; Tissot et al., 2018). In doing so, we see that these different correction schemes create different histories of f_{anox} , in some cases missing entire anoxic intervals. We also see

that the assumption of zero diagenesis makes it impossible for the model to simulate $\delta^{238}\text{U}$ values $\gg -0.3\%$ (**Fig. 9a**).

Although our inverse modeling approach can provide strong quantitative constraints on oceanic anoxia using the $\delta^{238}\text{U}$ proxy, the above sensitivity tests show that the assessment of diagenetic alteration of primary isotopic signatures is the largest remaining hurdle to the accurate utilization of this proxy. The most conservative approach to these corrections is to randomly sample from the range of diagenetic offsets observed in modern platform carbonate sediments (**Fig. 9c**); pelagic datasets may be best left uncorrected (as in Clarkson et al., 2018, 2021), though it is important to acknowledge the possible role of diagenesis is obscuring subtle trends. Screening individual samples for diagenesis using geochemical and/or petrographic indicators can also improve the precision of f_{anox} reconstructions by filtering out clearly altered samples. In this way, it is data quality rather than quantity that will dictate the precision and accuracy of f_{anox} reconstructions using $\delta^{238}\text{U}$ data moving forward.

5. Concluding remarks

When quantifying the extent of seafloor anoxia using sedimentary $\delta^{238}\text{U}$ data, we found that a Bayesian inverse approach enables both the rigorous interpretation of various datasets and robust comparison of trends between datasets. The largest remaining hurdles to accurately inferring seafloor anoxia with sedimentary $\delta^{238}\text{U}$ data are (i) accurate assessments of rate constants for anoxic U burial in marginal versus pelagic settings, and (ii) accurate corrections for diagenetic alteration of primary isotopic signatures. We follow others in the community in stressing that screening sample sets for preservation of primary signatures is a critical prerequisite to obtain accurate inferences about marine anoxia. If approached with awareness of the potential pitfalls outlined above, we find that a combination of judicious forward modeling and well-tuned inverse modeling can provide a powerful framework for quantitative assessment of anoxic seafloor extent within and among critical intervals in Earth’s history.

To demonstrate this promise, we compiled the posterior probability densities of F_{anox} and f_{anox} at the height of the four most recent anoxic events studied here (**Fig. 10**). Doing so shows that when accurately fitting trends in the data and propagating uncertainty in the isotope mass balance, one can indeed resolve meaningful differences in marine redox conditions across critical events in Earth history. For the four events considered here, we constrain the peak anoxic seafloor extent (f_{anox}) to $0.15^{+0.40}_{-0.14}\%$ at the PETM, $0.6^{+1.1}_{-0.5}\%$ during OAE-2, $3.7^{+5.1}_{-1.8}\%$ across the Triassic-Jurassic boundary, and $10.0^{+14.0}_{-5.6}\%$ across the Permian-Triassic boundary (all reported as median and 5th to 95th percentile). This represents relative extents of anoxia ranging from $\sim 1\times$ the modern value to $\sim 3\times$, $\sim 20\times$ and $\sim 50\times$, reflecting the ability of the $\delta^{238}\text{U}$ proxy to track global redox perturbations across this range of intensities. Two of these estimates (OAE-2, PTB) were also substantially revised compared to prior inferences (**Fig. 6e, g**); this re-analysis thus enables a consistent comparison of anoxic intensity across these different events. With this modeling framework in hand, we are therefore poised to gain a deeper quantitative understanding of these and other anoxic events, as well as their co-evolution with life and climate.

CRediT authorship contribution statement

Michael A. Kipp: Conceptualization, Methodology, Software, Writing – original draft. **François L.H. Tissot:** Conceptualization, Methodology, Writing – review & editing.

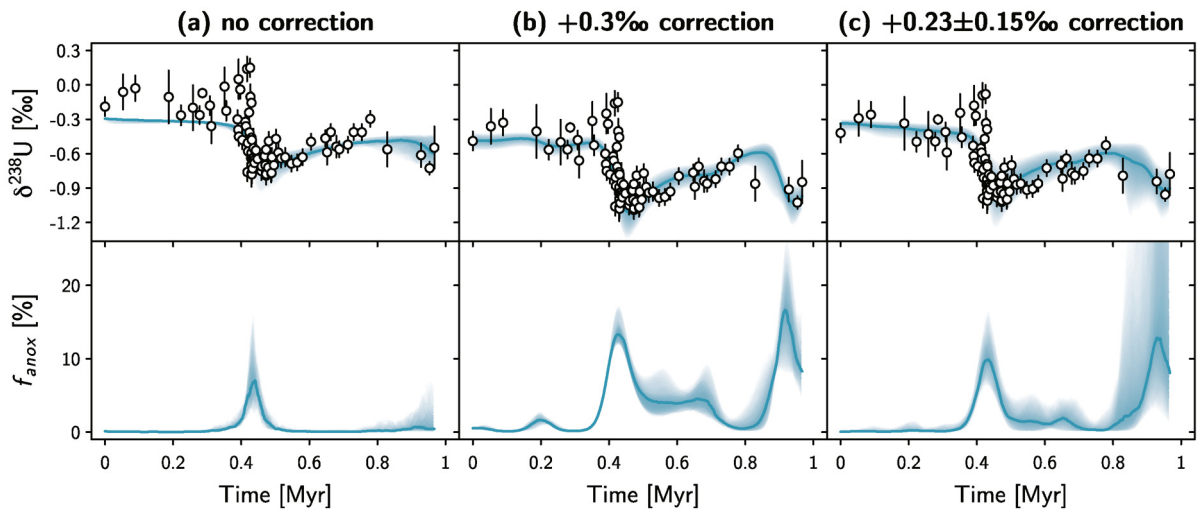


Fig. 9. Effect of diagenesis on f_{anox} reconstructions. Three methods of accounting for diagenetic effects were applied to the same dataset (Zhang et al., 2018): (a) no correction, (b) uniform correction of 0.3‰, (c) random correction drawn from distribution of observed offsets in Bahamas drill core data (0.23 ± 0.15 , 1σ ; Tissot et al., 2018). The different methods of accounting for diagenetic effects on carbonate $\delta^{238}\text{U}$ values have moderate to severe effects on the inferred f_{anox} trajectory.

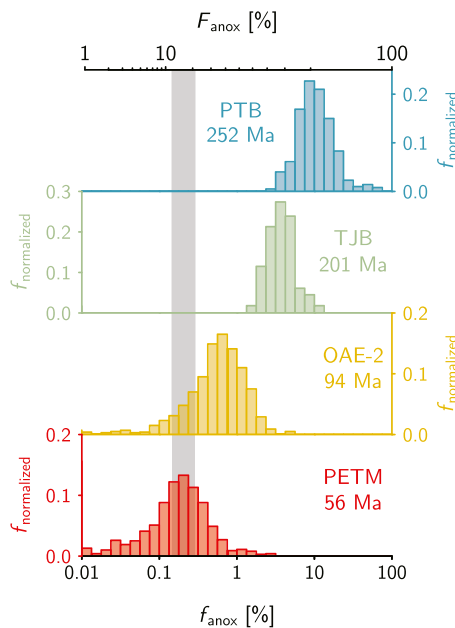


Fig. 10. Comparison of f_{anox} and F_{anox} at the peak of different anoxic events. Distributions represent 1000 forward model runs sampling the posterior parameter distributions; each distribution samples a time slice at the peak of the respective event. Events and their ages: Permian-Triassic boundary (PTB, 252 Ma); Triassic-Jurassic boundary (TJB, 201 Ma); Ocean Anoxic Event 2 (OAE-2, 94 Ma); Paleocene-Eocene Thermal Maximum (PETM, 56 Ma). The PTB results include propagation of uncertainty on the diagenetic correction, as in Fig. 9c; for other datasets no diagenetic correction was applied (following the originally published results).

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

This manuscript benefitted from the insightful comments of Kimberly Lau, Chris Reinhard, an anonymous reviewer, and Editors Lou Derry and Laurence Coogan. We also thank Joshua Krissansen-Totton and Jess Adkins for helpful discussions. Some of the compu-

tations presented here were conducted in the Resnick High Performance Computing Center, a facility supported by the Resnick Sustainability Institute at the California Institute of Technology. This work was supported by a Postdoctoral Fellowship in Geobiology from the Agouron Institute to MAK, and NSF grants EAR-1824002 and MGG-2054892, and start-up funds (provided by Caltech) to FLHT.

Appendix A

Below are the published and herein derived parameter values for the modern U isotope mass balance, along with their associated 2σ uncertainty intervals and references, if applicable. In our dynamic model runs, J_{riv} , $\delta^{238}\text{U}_{riv}$, Δ_{anox} , Δ_{other} , K_{anox} and K_{other} are held constant, while N_{sw} , $\delta^{238}\text{U}_{sw}$, F_{anox} and f_{anox} are calculated iteratively (with an assumption that at $t = 0$, the system is at steady state). As noted in Section 2.3.2, propagation of uncertainty on the constant terms can be achieved by iteratively sampling from a Gaussian distribution with the mean and 2σ as shown below, where each walker in the MCMC routine will pick a different set of values and maintain those values for the entire random walk.

$$N_{sw}: 19,000 \pm 1,200 \text{ Gmol U (Dunk et al., 2002)}$$

$$J_{riv}: 0.042 \pm 0.015 \text{ Gmol U yr}^{-1} \text{ (Dunk et al., 2002)}$$

$$\delta^{238}\text{U}_{riv}: -0.30 \pm 0.04\text{‰} \text{ (Tissot and Dauphas, 2015; Andersen et al., 2016)}$$

$$\delta^{238}\text{U}_{sw} \text{ (modern)}: -0.39 \pm 0.02\text{‰} \text{ (Tissot and Dauphas, 2015; Andersen et al., 2016; Rolison et al., 2017; Chen et al., 2018b)}$$

$$\Delta_{anox}: +0.6 \pm 0.2\text{‰} \text{ (Andersen et al., 2014; Holmden et al., 2015; Rolison et al., 2017)}$$

$$\Delta_{other}: 0.0 \pm 0.05\text{‰} \text{ (Lau et al., 2016; Zhang et al., 2020b)}$$

$$F_{anox} \text{ (modern)}: 15^{+8}_{-7}\% \text{ (Derived in Section 2.1)}$$

$$f_{anox} \text{ (modern)}: 0.19^{+0.11}_{-0.05}\% \text{ (Derived in Section 2.1)}$$

$$K_{anox}: 1.74^{+0.68}_{-0.63} \times 10^{-4} \text{ yr}^{-1} \text{ (Derived in Section 2.1)}$$

$$K_{other}: 1.88^{+0.22}_{-0.17} \times 10^{-6} \text{ yr}^{-1} \text{ (Derived in Section 2.1)}$$

Appendix B

In order to determine the time required for a system to reach a new steady state following a stepwise perturbation, we follow the approach of Holland (1978, p. 6-7). Let us consider a system with a rate of input (F_{in}) that remains at a constant value (a_0) for a time period long enough to reach a new steady state (t_1). The

input rate then suddenly changes to a new, constant level of input (a_1). The input flux of said element can then be described using the equations

$$F_{in} = \begin{cases} a_0, & 0 < t \leq t_1 \\ a_1, & t_1 \leq t \end{cases} \quad (\text{B.1})$$

If the output rate is proportional to the inventory of the element (N), then the change in inventory can be described as

$$\frac{dN}{dt} = \begin{cases} a_0 - kN, & 0 < t \leq t_1 \\ a_1 - kN, & t_1 \leq t \end{cases} \quad (\text{B.2})$$

where k is a rate constant describing the relationship between removal rate and the inventory of the element. Since we specified that t_1 is long enough for steady state to be attained, it follows that at time t_1 ,

$$N(t_1) = N_1 = \frac{a_0}{k} \quad (\text{B.3})$$

and after t_1 ,

$$N(t) = N_{final} - (N_{final} - N_1) \exp(-k(t - t_1)) \quad (\text{B.4})$$

where N_{final} is the inventory of the element at the new steady state, such that

$$N_{final} = \frac{a_1}{k}. \quad (\text{B.5})$$

The rate at which the system reaches the new steady state is therefore governed by k . When the system reaches the new steady state

$$\frac{N_{final}}{a_1} = \frac{1}{k} = \tau \quad (\text{B.6})$$

where τ is the residence time of the element in the reservoir. We can simplify Equation (B.4) by introducing new terms for the time elapsed in the approach toward the new steady state (Δt , where $\Delta t = t - t_1$), the magnitude of the change in inventory achieved when the new steady state is reached (ΔN_{final} , where $\Delta N_{final} = N_{final} - N_1$), and the amount of change in inventory toward the new steady state value that remains at time t (ΔN , where $\Delta N = N_{final} - N$). Substituting these terms into Equation (B.4), and using Equation (B.6) to substitute τ for k , we obtain

$$\frac{\Delta N}{\Delta N_{final}} = 1 - \exp\left(-\frac{\Delta t}{\tau}\right) \quad (\text{B.7})$$

which casts the fraction of the response toward the new steady state ($\Delta N/\Delta N_{final}$) as a function of the number of residence times elapsed ($\Delta t/\tau$). This equation was used to generate Fig. 1c in the main text.

Appendix C. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.epsl.2021.117240>.

References

Andersen, M.B., Romaniello, S., Vance, D., Little, S.H., Herdman, R., Lyons, T.W., 2014. A modern framework for the interpretation of $^{238}\text{U}/^{235}\text{U}$ in studies of ancient ocean redox. *Earth Planet. Sci. Lett.* 400, 184–194.

Andersen, M.B., Vance, D., Morford, J.L., Bura-Nakić, E., Breitenbach, S.F., Och, L., 2016. Closing in on the marine $^{238}\text{U}/^{235}\text{U}$ budget. *Chem. Geol.* 420, 11–22.

Asael, D., Tissot, F.L.H., Reinhard, C.T., Rouxel, O., Dauphas, N., Lyons, T.W., Ponzevera, E., Liorzou, C., Cheron, S., 2013. Coupled molybdenum, iron and uranium stable isotopes as oceanic paleoredox proxies during the Paleoproterozoic Shunga Event. *Chem. Geol.* 362, 193–210.

Bartlett, R., Elrick, M., Wheeley, J.R., Polyak, V., Desrochers, A., Asmerom, Y., 2018. Abrupt global-ocean anoxia during the Late Ordovician–early Silurian detected using uranium isotopes of marine carbonates. *Proc. Natl. Acad. Sci.* 115, 5896–5901.

Bertine, K.K., Turekian, K.K., 1973. Molybdenum in marine deposits. *Geochim. Cosmochim. Acta* 37, 1415–1434.

Bigeleisen, J., 1996. Nuclear size and shape effects in chemical reactions: isotope chemistry of the heavy elements. *J. Am. Chem. Soc.* 118, 3676–3680.

Bowen, G.J., Fischer-Femal, B., Reichart, G.J., Sluijs, A., Lear, C.H., 2020. Joint inversion of proxy system models to reconstruct paleoenvironmental time series from heterogeneous data. *Clim. Past* 16, 65–78.

Brennecke, G.A., Herrmann, A.D., Algeo, T.J., Anbar, A.D., 2011. Rapid expansion of oceanic anoxia immediately before the end-Permian mass extinction. *Proc. Natl. Acad. Sci.* 108, 17631–17634.

Broecker, W.S., Peng, T.-H., 1982. Tracers in the Sea. Lamont-Doherty Geological Observatory, Columbia University.

Brown, S.T., Basu, A., Ding, X., Christensen, J.N., DePaolo, D.J., 2018. Uranium isotope fractionation by abiotic reductive precipitation. *Proc. Natl. Acad. Sci.* 115, 8688–8693.

Brüske, A., Weyer, S., Zhao, M.-Y., Planavsky, N.J., Wegwerth, A., Neubert, N., Dellwig, O., Lau, K.V., Lyons, T.W., 2020. Correlated molybdenum and uranium isotope signatures in modern anoxic sediments: implications for their use as paleoredox proxy. *Geochim. Cosmochim. Acta* 270, 449–474.

Cao, M., Daines, S.J., Lenton, T.M., Cui, H., Algeo, T.J., Dahl, T.W., Shi, W., Chen, Z.-Q., Anbar, A., Zhou, Y.-Q., 2020. Comparison of Ediacaran platform and slope $\delta^{238}\text{U}$ records in South China: implications for global-ocean oxygenation and the origin of the Shuram Excursion. *Geochim. Cosmochim. Acta* 287, 111–124.

Chen, J.H., Edwards, R.L., Wasserburg, G.J., 1986. ^{238}U , ^{234}U , and ^{232}Th in seawater. *Earth Planet. Sci. Lett.* 80, 241–251.

Chen, X., Romaniello, S.J., Anbar, A.D., 2017. Uranium isotope fractionation induced by aqueous speciation: implications for U isotopes in marine CaCO_3 as a paleoredox proxy. *Geochim. Cosmochim. Acta* 215, 162–172.

Chen, X., Romaniello, S.J., Herrmann, A.D., Hardisty, D., Gill, B.C., Anbar, A.D., 2018a. Diagenetic effects on uranium isotope fractionation in carbonate sediments from the Bahamas. *Geochim. Cosmochim. Acta* 237, 294–311.

Chen, X., Romaniello, S.J., Herrmann, A.D., Samankassou, E., Anbar, A.D., 2018b. Biological effects on uranium isotope fractionation ($^{238}\text{U}/^{235}\text{U}$) in primary biogenic carbonates. *Geochim. Cosmochim. Acta* 240, 1–10.

Chen, X., Tissot, F.L.H., Jansen, M.F., Bekker, A., Liu, C.X., Nie, N.X., Halverson, G.P., Veizer, J., Dauphas, N., 2021. The uranium isotopic record of shales and carbonates through geologic time. *Geochim. Cosmochim. Acta* 300, 164–191.

Cheng, K., Elrick, M., Romaniello, S.J., 2020. Early Mississippian ocean anoxia triggered organic carbon burial and late Paleozoic cooling: evidence from uranium isotopes recorded in marine limestone. *Geology* 48, 363–367.

Clark, S.K., Johnson, T.M., 2008. Effective isotopic fractionation factors for solute removal by reactive sediments: a laboratory microcosm and slurry study. *Environ. Sci. Technol.* 42 (21), 7850–7855.

Clarkson, M.O., Stirling, C.H., Jenkyns, H.C., Dickson, A.J., Porcelli, D., Moy, C.M., von Strandmann, P.A.P., Cooke, I.R., Lenton, T.M., 2018. Uranium isotope evidence for two episodes of deoxygenation during Oceanic Anoxic Event 2. *Proc. Natl. Acad. Sci.* 201715278.

Clarkson, M.O., Musing, K., Andersen, M.B., Vance, D., 2020. Examining pelagic carbonate-rich sediments as an archive for authigenic uranium and molybdenum isotopes using reductive cleaning and leaching experiments. *Chem. Geol.* 539, 119412.

Clarkson, M.O., Lenton, T.M., Andersen, M.B., Bagard, M.L., Dickson, A.J., Vance, D., 2021. Upper limits on the extent of seafloor anoxia during the PETM from uranium isotopes. *Nat. Commun.* 12, 399.

Cole, D.B., Planavsky, N.J., Longley, M., Boning, P., Wilkes, D., Wang, X., Swanner, E.D., Wittkop, C., Loydell, D.K., Busigny, V., Knudsen, A.C., Sperling, E.A., 2020. Uranium isotope fractionation in non-sulfidic anoxic settings and the global uranium isotope mass balance. *Glob. Biogeochem. Cycles* 34 (8), e2020GB006649.

Dunk, R.M., Mills, R.A., Jenkins, W.J., 2002. A reevaluation of the oceanic uranium budget for the Holocene. *Chem. Geol.* 190, 45–67.

Fantle, M.S., Barnes, B.D., Lau, K.V., 2020. The role of diagenesis in shaping the geochemistry of the marine carbonate record. *Annu. Rev. Earth Planet. Sci.* 48, 549–583.

Fujii, Y., Nomura, M., Okamoto, M., Onitsuka, H., Kawakami, F., Takeda, K., 1989. An anomalous isotope effect of ^{235}U in U(IV)-U(VI) chemical exchange. *Z. Naturforsch.* 44 (5), 395–398.

Goto, K.T., Anbar, A.D., Gordon, G.W., Romaniello, S.J., Shimoda, G., Takaya, Y., Tokumaru, A., Nozaki, T., Suzuki, K., Machida, S., 2014. Uranium isotope systematics of ferromanganese crusts in the Pacific Ocean: implications for the marine $^{238}\text{U}/^{235}\text{U}$ isotope system. *Geochim. Cosmochim. Acta* 146, 43–58.

Haario, H.H., Saksman, E., Tamminen, J., 2001. An adaptive Metropolis algorithm. *Bernoulli* 7, 223–242.

Hastings, W.K., 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57, 97–109.

Holland, H.D., 1978. The Chemistry of the Atmosphere and Oceans. John Wiley & Sons.

- Holmden, C., Amini, M., Francois, R., 2015. Uranium isotope fractionation in Saanich Inlet: a modern analog study of a paleoredox tracer. *Geochim. Cosmochim. Acta* 153, 202–215.
- Hood, A.v.S., Planavsky, N.J., Wallace, M.W., Wang, X., Bellefroid, E.J., Gueguen, B., Cole, D.B., 2016. Integrated geochemical-petrographic insights from component-selective $\delta^{238}\text{U}$ of Cryogenian marine carbonates. *Geology* 44, 935–938.
- Jost, A.B., Bachan, A., van de Schootbrugge, B., Lau, K.V., Weaver, K.L., Maher, K., Payne, J.L., 2017. Uranium isotope evidence for an expansion of marine anoxia during the end-Triassic extinction. *Geochem. Geophys. Geosyst.* 18, 3093–3108.
- Kendall, B., Komiya, T., Lyons, T.W., Bates, S.M., Gordon, G.W., Romaniello, S.J., Jiang, G., Creaser, R.A., Xiao, S., McFadden, K., Sawaki, Y., Tahata, M., Shu, D., Han, J., Li, Y., Chu, X., Anbar, A.D., 2015. Uranium and molybdenum isotope evidence for an episode of widespread ocean oxygenation during the late Ediacaran Period. *Geochim. Cosmochim. Acta* 156, 173–193.
- Krissansen-Totton, J., Kipp, M.A., Catling, D.C., 2021. Carbon cycle inverse modeling suggests large changes in fractional organic burial are consistent with the carbon isotope record and may have contributed to the rise of oxygen. *Geobiology* 19, 342–363.
- Ku, T.-L., Knauss, K.G., Mathieu, G.G., 1977. Uranium in open ocean: concentration and isotopic composition. *Deep-Sea Res.* 24, 1005–1017.
- Lau, K.V., Maher, K., Altiner, D., Kelley, B.M., Kump, L.R., Lehrmann, D.J., Silva-Tamayo, J.C., Weaver, K.L., Yu, M., Payne, J.L., 2016. Marine anoxia and delayed Earth system recovery after the end-Permian extinction. *Proc. Natl. Acad. Sci.* 113, 2360–2365.
- Lau, K.V., Macdonald, F.A., Maher, K., Payne, J.L., 2017. Uranium isotope evidence for temporary ocean oxygenation in the aftermath of the Sturtian Snowball Earth. *Earth Planet. Sci. Lett.* 458, 282–292.
- Lau, K.V., Romaniello, S.J., Zhang, F., 2019. *The Uranium Isotope Paleoredox Proxy*. Cambridge University Press.
- Lau, K.V., Lyons, T.W., Maher, K., 2020. Uranium reduction and isotopic fractionation in reducing sediments: insights from reactive transport modeling. *Geochim. Cosmochim. Acta* 287, 65–92.
- MacKay, D.J., 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. *J. Chem. Phys.* 21, 1087–1092.
- Nomura, M., Higuchi, N., Fujii, Y., 1996. Mass dependence of uranium isotope effects in the U(IV)-U(VI) exchange reaction. *J. Am. Chem. Soc.* 118 (38), 9127–9130.
- Nursall, J., 1959. Oxygen as a prerequisite to the origin of the Metazoa. *Nature* 183, 1170–1172.
- Owens, S.A., Buesseler, K.O., Sims, K.W.W., 2011. Re-evaluating the ^{238}U -salinity relationship in seawater: implications for the ^{238}U - ^{234}Th disequilibrium method. *Mar. Chem.* 127, 31–39.
- Reinhard, C.T., Planavsky, N.J., Robbins, L.J., Partin, C.A., Gill, B.C., Lalonde, S.V., Bekker, A., Konhauser, K.O., Lyons, T.W., 2013. Proterozoic ocean redox and biogeochemical stasis. *Proc. Natl. Acad. Sci. USA* 110, 5357–5362.
- Rolison, J.M., Stirling, C.H., Middag, R., Rijkenberg, M.J., 2017. Uranium stable isotope fractionation in the Black Sea: modern calibration of the $^{238}\text{U}/^{235}\text{U}$ paleo-redox proxy. *Geochim. Cosmochim. Acta* 203, 69–88.
- Romaniello, S.J., Herrmann, A.D., Anbar, A.D., 2013. Uranium concentrations and $^{238}\text{U}/^{235}\text{U}$ isotope ratios in modern carbonates from the Bahamas: assessing a novel paleoredox proxy. *Chem. Geol.* 362, 305–316.
- Soetaert, K., Petzoldt, T., 2010. Inverse modelling, sensitivity and Monte Carlo analysis in R using package FME. *J. Stat. Softw.* 33, 1–28.
- Song, Huyue, Song, Haijun, Algeo, T.J., 2017. Uranium and carbon isotopes document global-ocean redox-productivity relationships linked to cooling during the Frasnian-Famennian mass extinction. *Geology* 45, 887–890.
- Stirling, C.H., Andersen, M.B., Potter, E.-K., Halliday, A.N., 2007. Low-temperature isotopic fractionation of uranium. *Earth Planet. Sci. Lett.* 264, 208–225.
- Tierney, J.E., Malevich, S.B., Gray, W., Vetter, L., Thirumalai, K., 2019. Bayesian calibration of the Mg/Ca paleothermometer in planktic foraminifera. *Paleoceanogr. Paleoclimatol.* 34, 2005–2030.
- Tissot, F.L., Dauphas, N., 2015. Uranium isotopic compositions of the crust and ocean: age corrections, U budget and global extent of modern anoxia. *Geochim. Cosmochim. Acta* 167, 113–143.
- Tissot, F.L.H., Chen, C., Go, B., Naziemiec, M., Healy, G., Bekker, A., Swart, P.K., Dauphas, N., 2018. Controls of eustasy and diagenesis on the $^{238}\text{U}/^{235}\text{U}$ of carbonates and evolution of the seawater ($^{234}\text{U}/^{238}\text{U}$) during the last 1.4 Myr. *Geochim. Cosmochim. Acta*.
- Turekian, K.K., 1969. The oceans, streams, and atmosphere. *Handb. Geochem.*, 297–323.
- Veeh, H.H., 1967. Deposition of uranium from the ocean. *Earth Planet. Sci. Lett.* 3, 145–150.
- Wang, X., Planavsky, N.J., Reinhard, C.T., Hein, J.R., Johnson, T.M., 2016. A Cenozoic seawater redox record derived from $^{238}\text{U}/^{235}\text{U}$ in ferromanganese crusts. *Am. J. Sci.* 316, 64–83.
- Weyer, S., Anbar, A.D., Gerdes, A., Gordon, G.W., Algeo, T.J., Boyle, E.A., 2008. Natural fractionation of $^{238}\text{U}/^{235}\text{U}$. *Geochim. Cosmochim. Acta* 72, 345–359.
- Wignall, P.B., Twitchett, R.J., 1996. Oceanic anoxia and the end Permian mass extinction. *Science* 272, 1155–1158.
- Zhang, F., Dahl, T.W., Lenton, T.M., Luo, G., Shen, S., Algeo, T.J., Planavsky, N., Liu, J., Cui, Y., Qie, W., 2020a. Extensive marine anoxia associated with the Late Devonian Hangenberg Crisis. *Earth Planet. Sci. Lett.* 533, 115976.
- Zhang, F., Lenton, T.M., del Rey, Á., Romaniello, S.J., Chen, X., Planavsky, N.J., Clarkson, M.O., Dahl, T.W., Lau, K.V., Wang, W., 2020b. Uranium isotopes in marine carbonates as a global ocean paleoredox proxy: a critical review. *Geochim. Cosmochim. Acta* 287, 27–49.
- Zhang, F., Algeo, T.J., Romaniello, S.J., Cui, Y., Zhao, L., Chen, Z.-Q., Anbar, A.D., 2018. Congruent Permian-Triassic $\delta^{238}\text{U}$ records at Panthalassic and Tethyan sites: confirmation of global-oceanic anoxia and validation of the U-isotope paleoredox proxy. *Geology* 46, 327–330.

Inverse methods for consistent quantification of seafloor anoxia using uranium isotope data from marine sediments

Michael A. Kipp and François L. H. Tissot

Assessing convergence of the MCMC routine

1) In order to assess the ability of an MCMC routine to optimize free parameters, one must **first** decide on the number of free parameters to include. Here it can be useful to consider model *collinearity* and parameter *identifiability*. Models with high *collinearity* have many covariant parameters. This poses an issue for the MCMC routine, since changes in certain parameters can be largely compensated by changes in others; the net effect is that these cases make convergence on the “best fit” solution very inefficient (*i.e.*, poor *identifiability* of optimal parameter values).

We used the `collin` function in the FME package in R (Soetaert and Petzoldt, 2010) to assess collinearity in our model. This follows the approach of Brun et al. (2001, their Eq. 13), who defined a collinearity index (γ_K) that describes the ability of other model parameters to compensate changes in a single parameter value. For instance, $\gamma_K = 10$ indicates that if a single parameter value is altered, the model output can be matched to within $1/10 = 10\%$ by compensating changes in other parameters. While no γ_K value specifically precludes convergence, higher values imply more difficult convergence (the threshold for parameter *identifiability* is often found to be in the range of $\gamma_K = 5-20$). As expected, we find that with increasing m values (Eq. 8, main text), the collinearity of our model increases (Fig. S1). It is therefore recommended to start with low m values (fewer free parameters) for a given dataset, assess convergence, and iteratively increase m as possible and as is warranted by the data.




Figure S1. Model collinearity as function of m value. Higher m values result in higher collinearity and poorer identifiability.

2) Having used the *collinearity* assessment above to develop an approach to selecting the number of free parameters, the **second** step in building the MCMC routine is to determine the number of steps (n_{iter}) a walker needs to take to converge on the optimal parameter values. There is a vast literature on this topic (reviewed in [Cowles and Carlin, 1996](#)) that includes several indices for assessing whether a walker (also called a “chain”) has converged upon the stationary (target) distribution. Here we consider perhaps the most widely-used convergence metric: the *potential scale reduction factor* ([Gelman and Rubin, 1992](#)).

Gelman and Rubin ([1992, their Eq. 20](#)) defined a potential scale reduction factor (psrf, also called the Gelman-Rubin statistic or Rhat) that can determine whether a walker (or “chain”) has converged by comparing the variance within a walker to that between walkers. This approach therefore requires that at least two walkers be deployed in parallel. If the walkers have converged on the stationary distribution, they will have nearly identical variance, such that psrf is near unity. As in all cases, no single psrf value guarantees convergence, but the authors of this diagnostic typically view psrf values < 1.2 as suggesting convergence ([Kass et al., 1998](#)). We monitored psrf during all MCMC runs. We found that when psrf was too high ($>>1.2$), increasing n_{iter} was typically necessary, along with tuning of the size of proposed “jumps” in parameter values (p_{step}) and the frequency of updates (p_{update}) in proposed jumps using the covariance matrix (after [Haario et al., 1999; 2001](#)). Another statistic worth tracking when trying to get the walkers to converge is the *acceptance fraction*, *i.e.*, the fraction of proposed steps that are accepted as new parameter values. It has been demonstrated that accepting $\sim 20\%$ of runs is most efficient in a wide range of MCMC applications (*e.g.*, [Geyer and Thompson, 1995; Gelman et al., 1996](#)). We therefore viewed acceptance of $< 10\%$ or $> 50\%$ of runs as suggestive of slow convergence, and accordingly updated p_{step} and/or p_{update} to bring the value closer to $\sim 20\%$ as possible and/or necessary. In most cases, a psrf value close to unity co-occurred with an optimized acceptance fraction. Ultimately, we found that psrf < 1.1 could typically be achieved with $n_{iter} = 10^5$ - 10^6 when $m \leq 10$.

3) Having determined that the walkers went long enough to converge on the target distribution, the **third** step in developing the MCMC routine is to decide how many steps to keep from each walker. This is primarily a computational issue; more steps would give a larger sample size and thus more precise MCMC estimate, though with the cost of extra computational resources (*i.e.*, time to execute a longer walk, and memory to manipulate large matrices of posterior samples). A relevant metric for deciding the number of samples to keep is the *integrated autocorrelation time* (IAT). The IAT is a measure of the inefficiency of the MCMC method, describing the number of MCMC samples (or steps in the random walk) required to generate one independent sample. The MCMC method is not perfectly efficient because the samples are in fact not independent. The MCMC-derived error on our estimates therefore scales as the inverse of the IAT, where IAT = 1 would mean that there is no MCMC-derived error (*e.g.*, [Sherlock et al. 2010; their Eq. 6](#)). In practice, it is prudent to let each walker go for > 10 autocorrelation times (*e.g.*, [p. 380 in MacKay, 2003](#)). This gives a relative error of a few percent in the estimates derived from the MCMC approach – often smaller than the uncertainty ranges constrained in the MCMC analysis – meaning that we aren’t inhibited from drawing meaningful quantitative conclusions from our MCMC outputs. Multiple authors describe methods for estimating the IAT (*e.g.*, [Sherlock et al., 2010; Soetaert and Petzoldt, 2010; Foreman-Mackey et al., 2013](#)). Here we use the `IAT` function from the `LaplacesDemon` package in R. We monitored IAT for all MCMC runs, finding that for most datasets, IAT was ≤ 100 steps. This means that

keeping 1000 steps would sufficiently minimize MCMC-derived error (*i.e.*, yielding ~ 10 or more effectively independent samples). The walkers can go for longer and more samples can be retained, but as we ultimately want to conduct many parallel walks to propagate uncertainty in the isotopic mass balance terms (see below), we want our runs to be as concise as possible. So we typically retained the final 1000 steps from each walker once it was demonstrated that the walkers had converged on the stationary distribution. In practice, we did this by setting $n_{burn-in}$ to 1000 less than n_{iter} . This also allows the adaptive MCMC routine of the FME package to update the proposed steps using the covariance matrix up until the point that we retain runs; after the burn-in period, the proposals are not adjusted, meaning that the convergence metrics described above can be applied (whereas they cannot be applied while the proposals are subject to change).

4) Finally, we can consider the number of walkers ($n_{walkers}$) to deploy in parallel. If all we care about is obtaining a large number of samples from the stationary distribution, the most computationally efficient method would be to let a single walker converge on the target distribution and then retain as many samples as wanted in order to minimize MCMC-derived error (p. 381; MacKay, 2003). This means that only once does a walker need to take the time to converge on the target distribution. However, with a single chain it is not possible to assess convergence using a metric like the psrf. Proponents of the psrf thus often deploy a few walkers (3-5) and compare variance between walkers via psrf to demonstrate convergence (Kass et al., 1998). This is only slightly less computationally efficient than the first method, but has the added benefit of a more rigorous convergence assessment. Additionally, if computations for each walker are conducted in parallel, then no extra time is needed. A final end-member is to deploy many parallel chains, running only long enough to converge and then retain a sufficient number of samples, and combining all the samples to comprise the final posterior distribution. One advantage of this method is that samples from different chains will likely be less correlated than those from a single, long chain. In our case, a further advantage is that we can randomly sample from slightly different values in the constant terms of the isotopic mass balance in each of the parallel chains. This enables us to propagate the uncertainty on these terms into our inversions. We therefore took the approach of first tuning the MCMC routine for convergence using a few chains and holding mass balance terms constant. Then having found the conditions necessary for convergence, we deployed ≥ 100 walkers in parallel, each sampling from the uncertainty range surrounding each parameter in **Appendix A**. In doing so, many chains had a harder time converging and required a higher n_{iter} value. In the end, though, this large sample set was compiled to comprise the total posterior distribution. As a final step, the forward model was then run 1000 times while sampling from this posterior distribution in order to generate the time series estimates and confidence intervals in the figures in the main text.

Age uncertainty

In addition to propagating uncertainty on the parameters in the isotopic mass balance, we can account for uncertainty in sample ages. While the studies examined here do not report age uncertainties, we took a conservative approach by assuming a 20% uncertainty (1σ) in the age difference between adjacent samples, such that the law of superposition is not violated. We compared the MCMC inversion results with and without age uncertainty for the Frasnian-Famennian dataset of Song et al. (2017), shown in Fig. 6b in the main text, as this dataset had the largest time intervals between points (in denser datasets, the uncertainty would be negligible).

We found that accounting for this potential age uncertainty had a negligible effect on the inferred history of anoxia (Fig. S2). We therefore infer that our inversion results presented in the main text, which do not account for age uncertainty, are likely robust. We note, however, that future studies can account for this effect when reporting isotopic data from samples with poorly constrained relative ages.

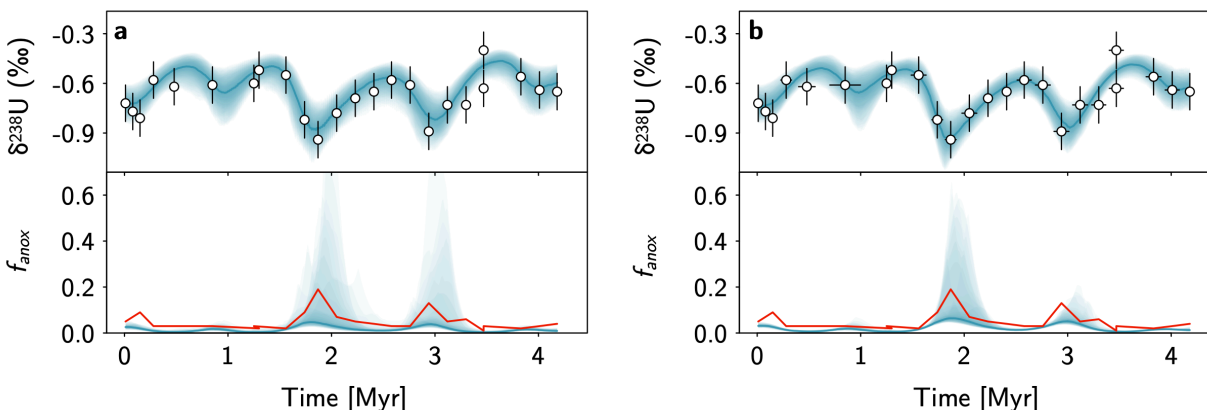


Figure S2. MCMC inversion results for Frasnian-Famennian dataset of Song et al. (2017), conducted (a) without and (b) with age uncertainty. The inferred trajectory of seafloor anoxia is quite similar in the two cases, suggesting that this potential uncertainty does not undermine the reconstructions.

End-member cases

Here we tested the ability of the model to successfully recover a step change in seafloor anoxia. We used three simple cases, where f_{anox} increases from the modern value (0.2%) to 1%, 10% and 80% instantaneously at the midpoint of a 1 Myr time interval. The resulting U isotope trends (data points, **Figs. S3a,c,e**) were successfully fitted by the MCMC routine. However, the recovered f_{anox} trajectory was erroneously high in the highest fractional anoxia scenario (**Fig. S2f**). This is because there is very little difference in the steady-state U isotope composition of seawater at high extents of seafloor anoxia ($\gg 10\%$; **Fig. 2a** in main text). We therefore recommend that interpretations of such datasets (where f_{anox} is persistently in the range of several 10's of percent) be approached cautiously, as minor changes in U isotope ratios can implicate vastly different extents of anoxia.




Figure S3. Step change test. Data points in panels (a,c,e) represent the trend implicated by the anoxia trajectories shown in panels (b,d,f) as the dashed line. The MCMC routine was then used to fit this “known” dataset, with the resulting fits shown as the median (blue lines) and 16th to 84th percentile confidence intervals (blue shading), here conducted without propagation of uncertainty for the constant terms in the mass balance (**Appendix A**), such that the effect of the step change on the inversion could be isolated. The poor sensitivity of the U isotope proxy when f_{anox} is high ($\gg 10\%$) causes the model to struggle to identify the correct trend in the 80% f_{anox} scenario (f) despite fitting the data precisely (e).

Numerical approach to model fit

Last, we demonstrate the efficacy of an alternative approach to finding the most likely f_{anox} trend given a $\delta^{238}\text{U}$ dataset. As noted in **Section 2.3.1** of the main text, one could avoid analytically parameterizing the trajectory of f_{anox} by numerically prescribing an f_{anox} history for a dataset. This is a more agnostic approach to the problem of finding the most likely f_{anox} history, so we explored whether it would be computationally feasible to optimize the fit to the data taking this approach. To do so, we ran 100,000 forward model simulations with the f_{anox} history in each simulation described as a random walk from a random starting value (between 0.01% and 100%) and taking a random step (± 0.25 log units) at each time point (1 kyr). We then calculated the NLL for each model run and retained only the best 0.5% of model runs. The median, along with 16th and 84th percentile ranges, of these best model runs displays a good fit to the data (**Fig. S4**). However, the fit is not quite as good as that obtained in our MCMC routine (where temporal changes in f_{anox} were parameterized as a Fourier series), despite a longer computational time. Tuning parameters in this approach (the size of steps in the random walk, percentage of runs retained) did not appreciably improve the outcome. We therefore opted to proceed with the analytical description of f_{anox} trajectories in our final model runs, while noting that other approaches are capable of minimizing NLL.




Figure S4. Comparison of MCMC-derived time series (blue shading) with time series derived from numerical f_{anox} simulation (red lines). The numerical f_{anox} parameterization approach finds a good fit to the data and similar f_{anox} trajectory (median in solid red lines, 16th and 84th percentile shown in dashed lines) to that recovered by the MCMC routine. However, the fit is not quite as good and took longer to obtain. We therefore opted to describe the f_{anox} history as a Fourier series (Eq. 8, main text) and optimize the values of the coefficients via MCMC, as this more efficiently found the best fit to a dataset.

References

- Brun, R., Reichert, P., Kunsch, H.R. 2001. Practical identifiability analysis of large environmental simulation models. *Water Resources Research*. 37, 1015—1030.
- Cowles, M.K. and Carlin, B.P. 1996. Markov Chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*. 91, 883—904.
- Foreman-Mackey, D., Hogg, D.W., Lang, D., Goodman, J., 2013. Emcee: the MCMC hammer. *PASP*. 125, 306.
- Gelman, A. and Rubin, D.B. 1992. Inference from iterative simulation using multiple sequences. *Statistical Science*. 7, 503—507.
- Gelman, A., Roberts, G.O., Gilks, W.R. 1996. Efficient Metropolis jumping rules. In J.M. Bernardo, J.O. Berger, A.P. Dawid, and A.F.M. Smith (eds), *Bayesian Statistics 5: Proceedings of the Fifth Valencia International Meeting*, pp. 559—607. Oxford University Press, Oxford.
- Geyer, C.J. and Thompson, E.A., 1995. Annealing Markov chain Monte Carlo with applications to ancestral inference. *Journal of the American Statistical Association*. 90, 909—920.
- Haario, H., Saksman, E., Tamminen, J. 1999. Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*. 14, 375—395.
- Haario, H., Saksman, E., Tamminen, J. 2001. An adaptive Metropolis algorithm. *Bernoulli*. 7, 223—242.
- Kass, R.E., Carlin, B.P., Gelman, A., Neal, R.M. 1998. Markov Chain Monte Carlo in practice: A roundtable discussion. *The American Statistician*. 52, 93—100.
- MacKay, D.J., 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.
- Sherlock, C., Fearnhead, P., Roberts, G.O. 2010. The Random Walk Metropolis: Linking theory and practice through a case study. *Statistical Science*. 25, 172—190.
- Soetaert, K., Petzoldt, T., 2010. Inverse modelling, sensitivity and Monte Carlo analysis in R using package FME. *J. Stat. Softw.* 33, 1—28.