Linearized Carbonate Chemistry for Efficient Modeling of Ocean-Based Carbon Dioxide Removal

Nora Loose[1,2,3], Matthew Long[1], Scott Bachman[1], Pierre Damien[4], Scott Eilerman[1], Alicia Karspeck[1], Jeroen Molemaker[4], Dafydd Stephenson[1], Michael Tyka[5], Abigale Wyatt[1]

[1][C]Worthy LLC, Boulder, CO, USA
[2]Nansen Environmental and Remote Sensing Center, Bergen, Norway
[3]Bjerknes Centre for Climate Research, Bergen, Norway
[4]Department of Atmospheric and Oceanic Sciences, University of California, Los Angeles, CA, USA
[5]Applied Research, Google Inc., Seattle, USA

Key Points:

- A linearized tracer method replaces full biogeochemistry in ocean CDR simulations, carrying many interventions in a single model run
- CDR efficiency from explicit simulations is reproduced within a few percent, and within 1% for plausible intervention amplitudes
- The method achieves speedups of 14–28× over explicit simulations, enabling large ensembles and uncertainty quantification

Corresponding author: Nora Loose, nora.loose@nersc.no

Abstract

Ocean-based carbon dioxide removal (CDR) relies on ocean models for measurement, reporting, and verification (MRV), because induced CO_2 uptake spans basin scales and decades. Explicit calculations with coupled physics–biogeochemistry models are computationally expensive, because biogeochemical models carry many tracers and each CDR intervention must be simulated separately. We introduce a CDR tracer method that represents each direct ocean removal (DOR) intervention by one passive tracer and each ocean alkalinity enhancement (OAE) intervention by two. Air–sea gas exchange is linearized using precomputed carbonate sensitivities, so that many interventions can be simulated in a single integration without a biogeochemical model. We evaluate the method against 24 explicit calculations with a 12 km ROMS configuration coupled to the Marine Biogeochemistry Library (MARBL), and reproduce all 1,380 explicit 1° CESM/MARBL calculations of existing DOR and OAE efficiency atlases in two CESM integrations. The CDR tracer method is 13.5–28× faster than explicit calculations. Its relative errors in CDR efficiency, predominantly positive owing to the linearization, peak shortly after intervention and then decay to less than 1% within a few years. In ROMS, peak errors reach 2–8% for extreme interventions but stay below 1% for plausible ones; in CESM, they average 3.5±2%. Larger local errors, including in pH estimates, arise only where extreme interventions accumulate in semi-enclosed passages. The computational savings enable large ensembles across deployment scenarios and ocean models. Combined with the option to use carbonate sensitivities from different biogeochemical models or observation-based products, these ensembles support uncertainty quantification for MRV.

Plain Language Summary

One way to remove carbon dioxide from the atmosphere is to adjust seawater chemistry so that the ocean absorbs more CO_2, by adding alkaline substances or extracting CO_2 directly from seawater. Because this extra uptake is spread over large ocean areas and many years, it cannot be measured directly and has to be estimated with computer models of the ocean. Standard models of ocean biology and chemistry are computationally expensive. We develop a simpler method that tracks only the chemical changes caused by each intervention and calculates how quickly the ocean takes up CO_2. Many interventions can be simulated in a single model run. We test the method against more than 1,400 detailed simulations in two different ocean models. It reproduces the amount of CO_2 taken up to within a few percent, and often within one percent, while being 14 to 28 times cheaper to run. The errors are small compared to other uncertainties, such as differences between ocean models, and the saved computing power can be used to better understand these uncertainties.

1 Introduction

The Paris Agreement aims to limit global warming to well below 2°C and to pursue efforts to limit it to 1.5°C. Warming is now expected to exceed 1.5°C, however, and returning below this level would require net-negative carbon dioxide (CO_2) emissions (Reisinger et al., 2025; United Nations Environment Programme, 2026). This has motivated interest in carbon dioxide removal (CDR) strategies, which aim to actively remove CO_2 from the atmosphere (IPCC, 2022). While CDR has been explored in terrestrial and industrial contexts for decades, ocean-based CDR has only recently gained more attention as a pathway to enhance oceanic uptake of atmospheric CO_2 (Smith et al., 2024).

At their core, all ocean-based CDR interventions seek to reduce the partial pressure of carbon dioxide (pCO_2) in surface waters, thereby creating a deficit relative to the atmosphere that draws additional carbon dioxide into the ocean. Different strategies achieve this pCO_2 deficit in distinct ways. Direct ocean removal (DOR) lowers surface water pCO_2

by directly extracting CO_2 from seawater; the extracted CO_2 can then be stored, for example in geological reservoirs, or used otherwise (National Academies of Sciences, Engineering, and Medicine, 2022). Ocean alkalinity enhancement (OAE) decreases surface water pCO_2 by shifting the carbonate equilibrium, allowing seawater to hold more CO_2 at a given atmospheric pCO_2 (Renforth & Henderson, 2017; Oschlies et al., 2023). In all cases, the resulting deficit in surface water pCO_2 does not immediately draw down atmospheric CO_2. Instead, atmospheric CO_2 is absorbed as surface waters re-equilibrate with the atmosphere over timescales of years to decades, a process that is controlled by air-sea gas exchange and surface water residence times (He & Tyka, 2023; Zhou et al., 2024).

The vast spatial and temporal scales of the re-equilibration process make quantifying ocean-based CDR a formidable challenge (Fennel, 2026). Because a pCO_2-deficient plume can spread across entire ocean basins while air-sea gas exchange occurs, the intervention's footprint extends beyond what can be captured by direct real-world sampling. Furthermore, the plume dilutes rapidly, suppressing the intervention signal in surface ocean pCO_2 to below 1 µatm (He & Tyka, 2023), at or below the detection threshold of current observational systems (Wanninkhof et al., 2013). These observational hurdles are compounded by the fact that the counterfactual ocean state, i.e., the hypothetical state the ocean would have maintained without intervention, is inherently unobservable. This makes it effectively impossible to directly attribute observed changes in pCO_2 to a specific intervention, particularly when multiple deployments occur simultaneously. For all these reasons, credible measurement, reporting, and verification (MRV) of ocean-based CDR cannot rest on observations alone and will require ocean models as a central component (Fennel, 2026).

To date, modeling for ocean-based CDR has typically relied on coupled physical–biogeochemical simulations. This approach is computationally demanding for three main reasons: (i) the requirement for high-resolution model simulations spanning basin-wide spatial and decadal temporal scales; (ii) the high computational overhead of biogeochemical (BGC) models; and (iii) the requirement for distinct simulations for every intervention scenario. This final factor becomes particularly dominant when exploring various deployment locations, seasons, or start years to account for seasonal and interannual variability in air-sea gas exchange and ocean dynamics.

To address the computational challenge, we introduce a simplified, but more computationally efficient approach: the *CDR tracer* framework. Implemented as passive tracers within a physical ocean model, each CDR tracer pair directly represents the perturbations in alkalinity and DIC that are induced by a specific ocean-based CDR intervention. For the gas exchange calculation, the framework bypasses full chemistry solvers by using a linearized carbonate chemistry formulation with pre-computed carbonate sensitivities. This approximation is expected to hold because, at regional to global scales, the carbonate system responds nearly linearly to perturbations of the magnitude that is relevant to CDR interventions. We evaluate the validity of this linearity assumption in this study within our modeling frameworks and demonstrate that it results in relatively small errors. The CDR tracer approach circumvents two of the three costly aspects identified above. First, the biogeochemical model is needed only once, in a control simulation from which the carbonate sensitivities are precomputed, rather than in every CDR simulation (factor (ii)). Second, multiple interventions can be simulated within a single model integration (factor (iii)), each represented by its own set of CDR tracers.

In Section 2, we derive the CDR tracer method from a linearization of the carbonate chemistry. We then describe its implementation in two ocean models: a 12 km configuration of the Regional Ocean Modeling System (ROMS) for the Pacific Ocean, and a 1° global configuration of the Community Earth System Model (CESM). For comparison, we also introduce explicit calculations, in which the full biogeochemistry is simulated by coupling both physical models to the Marine Biogeochemistry Library (MARBL).

In Section 3, we examine how nonlinearity of the carbonate system is expected to affect the CDR tracer estimate, and characterize how far the interventions displace the carbonate sensitivities from their control values. In Section 4, we verify the CDR tracer method against the explicit calculations (ROMS/MARBL and CESM/MARBL). Section 5 quantifies the computational savings, and Section 6 assesses the robustness of the method to the choice of carbonate sensitivities. We discuss the implications, limitations, and possible extensions of the method in Section 7.

2 Methods

We first derive the CDR tracer method from a linearization of the air–sea CO_2 flux (Sections 2.1 and 2.2). We then describe the ROMS and CESM configurations in which the method is implemented (Section 2.3) and the truth and CDR tracer experiments used to evaluate it (Sections 2.4 and 2.5). Finally, we define CDR efficiency, the primary metric by which the two are compared (Section 2.6).

2.1 Linearization of carbonate chemistry

Ocean models parameterize the net air–sea CO_2 flux, F_{CO_2}, as proportional to the difference in partial pressure between the atmosphere, pCO_2^{atm}, and the surface ocean, pCO_2 (e.g., Wanninkhof, 2014):

$$F_{CO_2} = k\, K_0\, (pCO_2^{atm} - pCO_2), \tag{1}$$

where k denotes the gas transfer velocity and K_0 the solubility of CO_2 in seawater. Here, we define $F_{CO_2} > 0$ as a net flux from the atmosphere into the ocean, i.e., oceanic CO_2 uptake. The parameterization of k varies between models, for example in its dependence on wind speed, sea state, or sea-ice cover. The derivation below relies only on the bulk form of Eq. (1) and holds for any parameterization of k that does not depend on the carbonate system variables.

The evolution of dissolved inorganic carbon (DIC) and alkalinity (ALK) is governed by advection, diffusion, and biogeochemical source-sink processes. In addition, DIC is modified by air–sea CO_2 exchange. Denoting concentrations by square brackets, the tracer transport equations take the form

$$\partial_t[\mathrm{DIC}] = -\nabla \cdot (\mathbf{u}\,[\mathrm{DIC}]) + \nabla \cdot (\kappa\, \nabla[\mathrm{DIC}]) + S^{\mathrm{DIC}}_{\mathrm{bgc}} + \partial_z F_{CO_2}, \tag{2}$$

$$\partial_t[\mathrm{ALK}] = -\nabla \cdot (\mathbf{u}\,[\mathrm{ALK}]) + \nabla \cdot (\kappa\, \nabla[\mathrm{ALK}]) + S^{\mathrm{ALK}}_{\mathrm{bgc}}, \tag{3}$$

where $\mathbf{u}$ is the three-dimensional velocity field, κ the eddy diffusivity tensor, and z the vertical coordinate (positive upward, with $z = 0$ at the sea surface). The terms $S^{\mathrm{DIC}}_{\mathrm{bgc}}$ and $S^{\mathrm{ALK}}_{\mathrm{bgc}}$ represent all net biogeochemical sources and sinks of DIC and alkalinity other than air–sea exchange; their exact form depends on the biogeochemical model. For DIC, they include, for example, photosynthetic uptake, respiration, remineralization, and calcium carbonate production and dissolution. For alkalinity, they are dominated by calcium carbonate cycling, with smaller contributions from, for example, nitrogen cycling and sedimentary processes. $\partial_z F_{CO_2}$ represents the convergence of the air–sea flux in the surface layer, which is equivalent to imposing the surface boundary condition $\kappa_z\, \partial_z[\mathrm{DIC}]|_{z=0} = F_{CO_2}$, where κ_z is the vertical eddy diffusivity. Since F_{CO_2} is defined as positive into the ocean (Eq. (1)), this term acts as a source of surface DIC when atmospheric pCO_2 exceeds its surface oceanic counterpart.

The remainder of this section considers the effects of ocean-based CDR interventions, which act by reducing the oceanic pCO_2. Examples include DOR and OAE. To isolate the impact of a given intervention, we compare a perturbed state (superscript "pert") to a counterfactual or control state (superscript "ctrl") and define the difference operator δ for any variable ψ as

$$\delta\psi = \psi^{\mathrm{pert}} - \psi^{\mathrm{ctrl}}. \tag{4}$$

In the following, we assume a non-responsive atmosphere, in which $p\mathrm{CO}_2^{\mathrm{atm}}$ is unaffected by the intervention. The flux anomaly δF_{CO_2} is then derived from Eq. (1) as

$$\delta F_{\mathrm{CO}_2} = -\,kK_0\,\delta p\mathrm{CO}_2 = -\,k\,\delta[\mathrm{CO}_2] \tag{5}$$

$$\approx -\,k\left(\left.\frac{\partial[\mathrm{CO}_2]}{\partial[\mathrm{DIC}]}\right|_{[\mathrm{ALK}]}\delta[\mathrm{DIC}] + \left.\frac{\partial[\mathrm{CO}_2]}{\partial[\mathrm{ALK}]}\right|_{[\mathrm{DIC}]}\delta[\mathrm{ALK}]\right). \tag{6}$$

The second equality in Eq. (5) follows from Henry's Law, $[\mathrm{CO}_2] = K_0\,p\mathrm{CO}_2$, which relates the oceanic partial pressure to the concentration of dissolved CO_2. The approximation in Eq. (6) is a first-order Taylor expansion treating DIC and ALK as the independent variables of the carbonate system, with the subscripts denoting the variable held constant.

We introduce the carbonate sensitivities

$$\beta = \left.\frac{\partial[\mathrm{DIC}]}{\partial[\mathrm{CO}_2]}\right|_{[\mathrm{ALK}]} \qquad \text{and} \qquad \eta = \left.\frac{\partial[\mathrm{DIC}]}{\partial[\mathrm{ALK}]}\right|_{[\mathrm{CO}_2]} \tag{7}$$

which quantify the sensitivity of DIC to changes in dissolved CO_2 at constant alkalinity, and to changes in alkalinity at constant CO_2, respectively. From β and η, the sensitivity of alkalinity to dissolved CO_2 at constant DIC follows as

$$\left.\frac{\partial[\mathrm{ALK}]}{\partial[\mathrm{CO}_2]}\right|_{[\mathrm{DIC}]} = -\frac{\beta}{\eta}. \tag{8}$$

Analytical expressions for β and η and the derivation of Eq. (8) are given in Appendix A. Substituting Eqs. (7) and (8) into Eq. (6), the flux anomaly δF_{CO_2} takes the compact form

$$\delta F_{\mathrm{CO}_2} \approx -\frac{k}{\beta}\big(\delta[\mathrm{DIC}] - \eta\,\delta[\mathrm{ALK}]\big). \tag{9}$$

We now subtract the control state (Eqs. (2)–(3)) from its perturbed counterpart and substitute the linearized gas exchange term (Eq. (9)), which yields closed equations for the anomalies $\delta[\mathrm{DIC}]$ and $\delta[\mathrm{ALK}]$ of the following form:

$$\partial_t\delta[\mathrm{DIC}] = -\nabla\cdot\big(\mathbf{u}\,\delta[\mathrm{DIC}]\big) + \nabla\cdot\big(\kappa\nabla\delta[\mathrm{DIC}]\big) + \partial_z\Big[-\frac{k}{\beta}\big(\delta[\mathrm{DIC}] - \eta\,\delta[\mathrm{ALK}]\big)\Big] + S_{\delta\mathrm{DIC}}, \tag{10}$$

$$\partial_t\delta[\mathrm{ALK}] = -\nabla\cdot\big(\mathbf{u}\,\delta[\mathrm{ALK}]\big) + \nabla\cdot\big(\kappa\nabla\delta[\mathrm{ALK}]\big) + S_{\delta\mathrm{ALK}}. \tag{11}$$

The source/sink terms $S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$ represent external inputs associated with the ocean-based CDR intervention. The biogeochemical terms $S_{\mathrm{bgc}}^{\mathrm{DIC}}$ and $S_{\mathrm{bgc}}^{\mathrm{ALK}}$ approximately cancel because biogeochemical processes are assumed to be largely unaffected by ocean-based CDR-induced changes in DIC and alkalinity. Some model configurations include higher-order feedbacks such as pH-sensitive ammonia (NH_3) outgassing, sedimentary denitrification, and calcification rates dependent on DIC. We expect these feedbacks to be small at large scales, where CDR-induced perturbations are diffuse. Since the interventions in this study are imposed over large areas, we focus on this regime and discuss the method's applicability to point-source releases in Section 7.

2.2 CDR tracer formulation

Motivated by Eqs. (10) and (11), we introduce the CDR tracers $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$, which track the perturbations in DIC and alkalinity due to the ocean-based CDR intervention. These are passive tracers in arbitrary units per unit volume, i.e., with units of $[\mathrm{m}^{-3}]$, governed by

$$\partial_t c_{\delta\mathrm{DIC}} = -\nabla\cdot(\mathbf{u}\,c_{\delta\mathrm{DIC}}) + \nabla\cdot(\kappa\nabla c_{\delta\mathrm{DIC}}) - \frac{k}{z_0\beta}\,(c_{\delta\mathrm{DIC}} - \eta\,c_{\delta\mathrm{ALK}}) + S_{\delta\mathrm{DIC}}, \tag{12}$$

$$\partial_t c_{\delta\mathrm{ALK}} = -\nabla\cdot(\mathbf{u}\,c_{\delta\mathrm{ALK}}) + \nabla\cdot(\kappa\nabla c_{\delta\mathrm{ALK}}) + S_{\delta\mathrm{ALK}}. \tag{13}$$

The first two terms on the right-hand side represent advection and diffusion, with notation as in Section 2.1. The third term in Eq. (12) acts only in the uppermost model layer, of thickness z_0, and represents the divergence of the intervention-induced change in air–sea CO_2 flux, $\delta\tilde{F}_{CO_2}/z_0$, where

$$\delta\tilde{F}_{CO_2} = -\frac{k}{\beta}\left(c_{\delta\mathrm{DIC}} - \eta\, c_{\delta\mathrm{ALK}}\right), \tag{14}$$

and $\delta\tilde{F}_{CO_2} > 0$ corresponds to additional CO_2 uptake by the ocean. This term acts as a restoring force on the DIC perturbation, consistent with the physical mechanism of ocean-based CDR: an increase in $c_{\delta\mathrm{DIC}}$ reduces the oceanic CO_2 uptake and thereby damps the DIC perturbation, while an increase in $c_{\delta\mathrm{ALK}}$ enhances uptake and amplifies the DIC perturbation. The gas transfer velocity k is computed online exactly as in the host model's air–sea flux parameterization (Section 2.3.1), while the carbonate sensitivities β and η are pre-diagnosed from a control run and provided as spatiotemporally varying external input fields (Figure 1).

The terms $S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$ represent the ocean-based CDR intervention. The form of these source/sink terms is determined by the type of intervention. For DOR, only DIC is modified ($S_{\delta\mathrm{DIC}} \neq 0$, $S_{\delta\mathrm{ALK}} = 0$), and the system reduces to Eq. (12) alone. For OAE, only alkalinity is modified ($S_{\delta\mathrm{DIC}} = 0$, $S_{\delta\mathrm{ALK}} \neq 0$), but both tracer equations are required, because the alkalinity perturbation drives a DIC response through air–sea exchange. Mixed interventions, such as enhanced rock weathering (ERW, e.g., Kanzaki et al., 2023) or partially pre-equilibrated OAE, can be represented by prescribing nonzero sources for both tracers ($S_{\delta\mathrm{DIC}} \neq 0$, $S_{\delta\mathrm{ALK}} \neq 0$) and, like OAE, require the coupled tracer pair.

$S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$ have units of $[\,\mathrm{m^{-3}\,s^{-1}}\,]$ and may be applied at arbitrary locations and depths with an arbitrary spatiotemporal structure. In many applications, the intervention is specified as a surface flux F (units of $[\,\mathrm{m^{-2}\,s^{-1}}\,]$), in which case

$$S = \frac{F}{z_0}. \tag{15}$$

For example, a DOR intervention removing x $\mathrm{mol\,m^{-3}\,s^{-1}}$ of DIC is implemented by setting $S_{\delta\mathrm{DIC}} = -x$ $\mathrm{m^{-3}\,s^{-1}}$, while a surface flux removal of x $\mathrm{mol\,m^{-2}\,s^{-1}}$ corresponds to $F = -x$ $\mathrm{m^{-2}\,s^{-1}}$. Owing to the linearity of the tracer equations, the absolute magnitudes of the source/sink terms are arbitrary: scaling them by a constant factor (for instance, specifying the intervention in mmol rather than mol) scales the tracer response by the same factor, which can be divided out (confirmed by Figure S24). For mixed interventions, only the ratio $S_{\delta\mathrm{DIC}}/S_{\delta\mathrm{ALK}}$ matters.

Figure 1 illustrates the workflow for the CDR tracer method. First, a control integration is performed with the physical model coupled to a full biogeochemical model (Figure 1, left). From this control run, the carbonate sensitivities β and η are obtained as time-varying fields, either diagnosed online or computed offline from output of the carbonate system state variables (temperature, salinity, DIC, alkalinity, silicate, and phosphate). Because β and η are used only in the air–sea gas exchange term of the subsequent CDR tracer run, which acts on the uppermost model layer (Eq. (14)), only their surface fields need to be obtained. In this study, we compute daily β and η fields offline using `pyCO2SYS` (Humphreys et al., 2020) (Appendix A). Because `pyCO2SYS` and MARBL's online carbonate chemistry solver may differ, for example in their equilibrium constants, the offline β and η may deviate slightly from those that would have been diagnosed online. Section 6 assesses how the temporal resolution of β and η affects the results, and Text S2 discusses how β and η should be sampled and averaged in time.

In the second step, the CDR tracer simulation is performed using only the physical model, with the daily β and η fields obtained in the first step prescribed as time-varying input fields (Figure 1, right). The CDR tracer run uses the same physical forc-

ing as the control run (e.g., atmospheric, lateral boundary, and river forcing), together with the prescribed intervention source terms $S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$, and is started either from the same initial conditions or from any restart of the control run. The control run needs to cover the full duration of the CDR tracer run so that β and η are available throughout. Because the CDR tracers are passive by construction, and the biogeochemical tracers of the coupled model are passive as well (i.e., neither feeds back onto the physics), the physical trajectory of the CDR tracer run is, in principle, identical to that of the control run. In practice, the physical trajectories can diverge owing to chaotic sensitivity to machine-level differences, unless the hardware and software environments are tightly controlled. This is discussed in the context of our experiments in Section 4.2.

Multiple CDR interventions can be simulated simultaneously within a single CDR tracer run, where each DOR intervention requires a single tracer $c_{\delta\mathrm{DIC}}$, while each OAE or mixed intervention requires a tracer pair $(c_{\delta\mathrm{DIC}}, c_{\delta\mathrm{ALK}})$. Each intervention has its own prescribed source terms, whereas the physical forcing and the fields β and η are shared by all CDR tracers. Because the physical state is computed only once for all interventions, this can result in significant computational savings. Examples of interventions that can be combined in a single CDR tracer run include:

- interventions applied in different geographic regions,
- interventions at different times (seasonal, interannual, or otherwise),
- different intervention types, such as OAE and DOR,
- or any combination of the above.

From the model output of $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$, the full perturbed state can be reconstructed as

$$[\mathrm{DIC}]^{\mathrm{pert}} = [\mathrm{DIC}]^{\mathrm{ctrl}} + c_{\delta\mathrm{DIC}},$$
$$[\mathrm{ALK}]^{\mathrm{pert}} = [\mathrm{ALK}]^{\mathrm{ctrl}} + c_{\delta\mathrm{ALK}},$$

where $[\mathrm{DIC}]^{\mathrm{ctrl}}, [\mathrm{ALK}]^{\mathrm{ctrl}}$ are saved from the control integration. The remaining carbonate system variables, such as pH, can then be estimated offline from $[\mathrm{DIC}]^{\mathrm{pert}}$ and $[\mathrm{ALK}]^{\mathrm{pert}}$ using `pyCO2SYS` or any other carbonate chemistry solver. Because the carbonate system is nonlinear, these variables, including pH, differ from their time averages when computed from time-averaged $[\mathrm{DIC}]^{\mathrm{pert}}$ and $[\mathrm{ALK}]^{\mathrm{pert}}$. Instantaneous (snapshot) output, as used here for ROMS (Section 2.3.2), avoids this bias but must be frequent enough to resolve the variability of interest.

2.3 Numerical Models

We evaluate the CDR tracer method against explicit coupled physics–biogeochemistry simulations, hereafter *truth* experiments. We use two complementary model configurations: a 12 km mesoscale-resolving regional Pacific ROMS setup (Section 2.3.2) and a 1° non-eddying global CESM setup (Section 2.3.3), both coupled to MARBL (Section 2.3.1). Each configuration supports three kinds of runs:

1. a *control* simulation, with the physical model coupled to MARBL, from which we diagnose the carbonate sensitivities β and η (Figure 2);
2. *truth* experiments, with the physical model also coupled to MARBL, run separately for each intervention; and
3. *CDR tracer* experiments, in which many interventions are carried together as passive tracers in a single run of the physical model alone, with β and η prescribed from the control simulation.

This section describes the biogeochemical model shared by both configurations, the ROMS and CESM model configurations, and the model developments required specif-

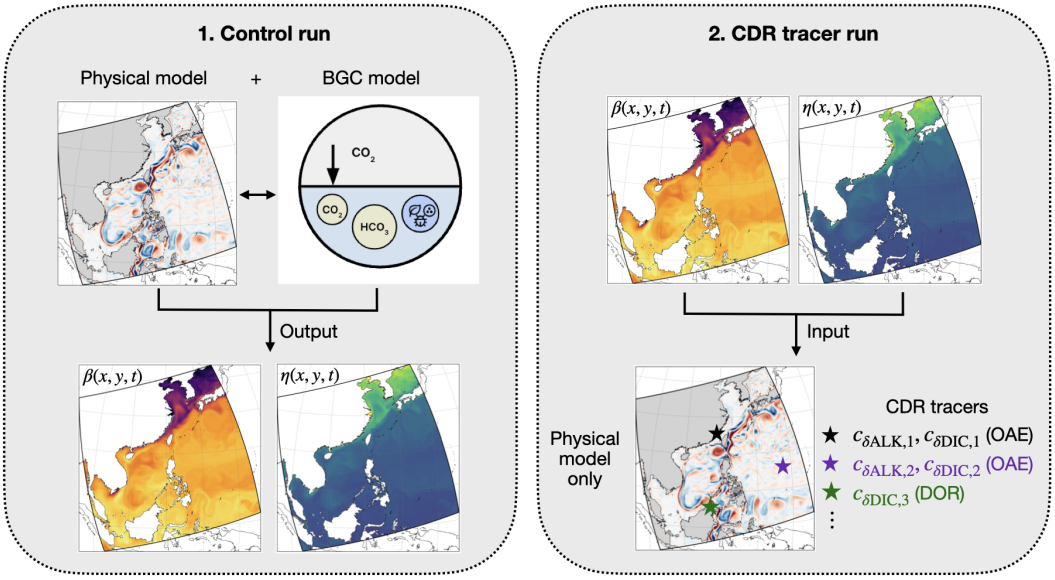


Figure 1. The CDR tracer method workflow. Left: In the control run, the physical model is coupled to the full biogeochemical model; carbonate sensitivities β and η are diagnosed online during the run or offline from its output. Right: The CDR tracer simulation uses only the physical model, with β and η prescribed as input fields. Multiple interventions (different regions, times, or types) can be run as passive tracers within a single CDR tracer run.

ically for the truth and CDR tracer experiments. The design of these experiments (intervention locations, amplitudes, and timing) is given in Sections 2.4 and 2.5.

2.3.1 MARBL

Both configurations use the Marine Biogeochemistry Library (MARBL, Long et al., 2021) for ocean biogeochemistry. Within MARBL, air–sea CO_2 exchange follows Eq. (1), with the gas transfer velocity parameterized after Wanninkhof (2014) as

$$k = \alpha\,(1 - f_{\mathrm{ice}})\,u_{10}^2\,\sqrt{\frac{660}{S_{\mathrm{CO_2}}}}, \tag{16}$$

where $\alpha = 6.97\times10^{-7}\,\mathrm{s\,m^{-1}}$ is a scaling coefficient, f_{ice} is the fractional sea-ice cover, u_{10} is the wind speed at 10 m height, and $S_{\mathrm{CO_2}}$ is the dimensionless Schmidt number for CO_2, computed as a fourth-order polynomial of the sea surface temperature (SST, in °C):

$$S_{\mathrm{CO_2}} = a + b\,\mathrm{SST} + c\,\mathrm{SST}^2 + d\,\mathrm{SST}^3 + e\,\mathrm{SST}^4, \tag{17}$$

with empirical coefficients $a = 2116.8$, $b = -136.25$, $c = 4.7353$, $d = -0.092307$, and $e = 0.0007555$.

Our MARBL configuration includes two higher-order feedbacks that the CDR tracer method is unable to represent. First, pH-sensitive NH_3 outgassing (Paulot et al., 2015): OAE- or DOR-driven increases in surface pH shift the NH_3/NH_4^+ equilibrium toward the volatile form, enhancing outgassing and slightly reducing surface alkalinity. Second, sedimentary denitrification increases alkalinity in the deepest model layer at a rate depending on the particulate organic carbon (POC) flux to the seafloor and bottom-water O_2 and NO_3^- concentrations. A higher-order feedback could arise if the intervention altered

surface carbonate chemistry sufficiently to affect phytoplankton growth and hence POC export, but the resulting deep alkalinity signal would then need to be transported back to the surface to influence air–sea CO_2 exchange. We therefore expect the sedimentary denitrification feedback to be negligible on the timescales considered here.

2.3.2 ROMS

We use the Regional Ocean Modeling System (ROMS, Shchepetkin & McWilliams, 2005) in a configuration covering most of the Pacific Ocean (Figs. 2(a),(c)), with a horizontal resolution of approximately 12 km and 100 topography-following vertical layers, based on the setup of Damien et al. (2026). Whereas Damien et al. (2026) couple ROMS to the Biogeochemical Elemental Cycling (BEC) model, we couple it to MARBL (Section 2.3.1), and refer to this configuration as ROMS/MARBL. Atmospheric forcing is provided at hourly frequency by ERA5 (Hersbach et al., 2020), and open boundary conditions are prescribed from the GLORYS global ocean reanalysis (Lellouche et al., 2021). Biogeochemical boundary conditions combine observational datasets and model solutions (Garcia et al., 2019; Lauvset et al., 2016; Huang et al., 2022; Yang et al., 2020; Yeager et al., 2022), and atmospheric pCO_2 is taken from Landschützer et al. (2016). The configuration has no prognostic sea-ice component; instead, grid cells are treated as ice-covered where the sea surface temperature falls below $-1.8\,^{\circ}\mathrm{C}$. This simplification has no bearing on the experiments presented here, whose intervention sites lie well away from any sea ice. The model is spun up for 78 years under cyclical forcing. All experiments presented here are initialized from this spin-up on January 1, 2000, and span the two-year period 2000–2001, a length limited by computational cost. Full equilibration of the CDR signal can take a decade or longer; a longer-term evaluation is provided by the CESM experiments (Section 2.3.3), which span 15 years.

For the truth experiments, OAE and DOR interventions are implemented as a prescribed surface flux of alkalinity or DIC, read from a forcing file and applied to the corresponding prognostic MARBL tracer over the uppermost model layer. For the CDR tracer experiments, we implement Eqs. (12) and (13) in the ROMS passive-tracer module. Advection and diffusion of $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$ reuse the existing tracer machinery, so only the surface terms are added: the prescribed intervention flux and the air-sea gas exchange. The intervention flux is taken directly from the corresponding truth experiment (the DIC-removal flux for DOR and the alkalinity-addition flux for OAE). It corresponds to the source terms $S_{\delta\mathrm{DIC}}$, $S_{\delta\mathrm{ALK}}$ via Eq. (15), and is applied as a surface flux, distributed over the uppermost layer. The air-sea gas exchange term in Eq. (12) is formed at each time step from the current tracer values, using the prescribed carbonate sensitivities β and η together with a gas-transfer velocity k. The sensitivities β and η are time-varying, supplied as external forcing fields at a user-specified frequency. The velocity k is computed online from sea surface temperature and winds (Eqs. (16), (17)) and set to zero in ice-covered cells, suppressing gas exchange under sea ice. Each $c_{\delta\mathrm{DIC}}$ tracer is paired with its $c_{\delta\mathrm{ALK}}$ partner so that the gas exchange term is evaluated consistently. For DOR the partner is absent, and the term reduces to a damping of $c_{\delta\mathrm{DIC}}$ alone.

2.3.3 CESM

We use the global forced ocean–ice configuration of the Community Earth System Model (CESM) version 2 (Yeager et al., 2022). The ocean component is the Parallel Ocean Program version 2 (POP2), run at a nominal 1° horizontal resolution with 60 vertical levels. POP2 is coupled to MARBL for ocean biogeochemistry (Section 2.3.1) and to the Los Alamos Sea Ice Model version 5 (CICE5) for sea ice, which provides the sea-ice fraction f_{ice} in Eq. (16). We refer to this configuration as CESM/MARBL. Atmospheric forcing is provided by the Japanese 55-year Reanalysis (JRA55, Kobayashi et al., 2015), with transient atmospheric CO_2 forcing. The model is integrated from 1850 to near-present, repeating the 1958–2018 JRA55 forcing cycle. Our analysis focuses on the period Jan-

uary 1999 to December 2013. This configuration, and the truth experiments described below, follow Zhou et al. (2024, 2026).

For the CESM CDR tracer implementation we use the same formulation as for ROMS. Because POP2 organizes each passive-tracer package as a self-contained module wired into the surrounding CESM and Common Infrastructure for Modeling the Earth (CIME) configuration and build layers, implementing the CDR tracer framework in CESM required substantially more development than in ROMS: a dedicated POP2 module with its own initialization, restart handling, diagnostics, and namelist interface, registered through the CIME case-configuration system. Each CDR tracer has its own surface forcing (a prescribed intervention flux, as in ROMS) and, for OAE or mixed interventions, is paired with a $c_{\delta\mathrm{ALK}}$ tracer. The carbonate sensitivities β and η and the per-intervention surface forcing are supplied as time-varying fields through CESM's coupler stream-data framework, the same forcing pathway used by MARBL. Routing the inputs this way ensures that the CDR tracer forcing is identical to the perturbations applied in the truth experiments. The air-sea gas exchange term is assembled as in ROMS, with the gas-transfer velocity k computed online from sea surface temperature, winds, and the prognostic sea-ice fraction (Eqs. (16), (17)).

One aspect of the CESM implementation merits note. In the CESM/MARBL truth experiments (Section 2.4.2), the alkalinity or DIC intervention flux is applied at the ocean surface regardless of sea-ice cover: even under complete sea-ice coverage, the prescribed δALK or δDIC flux is added to the surface layer. Air-sea CO_2 gas exchange, however, is always multiplied by $(1-f_{\mathrm{ice}})$, where f_{ice} is the prognostic sea-ice fraction (Eqs. (16), (1)). In other words, under full ice cover the perturbation accumulates in the surface ocean without any compensating outgassing until the ice retreats. To reproduce this behavior faithfully in the CDR tracer implementation, we adopt the same convention: the intervention source terms $S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$ are applied with no dependence on sea-ice fraction, while the air-sea gas exchange term $\delta\tilde{F}_{\mathrm{CO_2}}$ (Eqs. (16), (14)) uses multiplication by $(1 - f_{\mathrm{ice}})$.

Repeating the 1,380 CESM/MARBL truth experiments of Zhou et al. (2024, 2026) with the CDR tracer method required substantial workflow infrastructure beyond the model source code itself, owing to the large number of interventions carried simultaneously. Although the atmospheric forcing, the gas transfer velocity k, and the carbonate sensitivities β and η are shared across all interventions, each intervention requires its own prescribed intervention flux, and each $c_{\delta\mathrm{DIC}}$ tracer must be paired with the correct $c_{\delta\mathrm{ALK}}$ partner. We therefore developed a Python orchestration layer that configures the simulation programmatically, writes namelists, and assembles the per-intervention flux inputs. This tooling is released as open source alongside the CESM CDR tracer implementation (Appendix B).

2.4 Truth Experiments

The truth experiments comprise two sets: regional ROMS/MARBL experiments at four Pacific locations, and the global CESM/MARBL efficiency atlases of Zhou et al. (2024, 2026). The atlases map OAE and DOR efficiency by applying an intervention separately in each of 690 non-overlapping ocean polygons, which together span the global ocean south of 75°N. In total, this yields 1,404 truth experiments, all summarized in Table 1.

2.4.1 ROMS/MARBL experiments

We select four intervention locations in the Pacific: Japan (J), Vancouver Island (V), Baja California (B), and Ecuador (E), as shown in Figure 2(a),(c). These locations sample a wide range of carbonate sensitivity values and dynamical regimes (Figures 2,

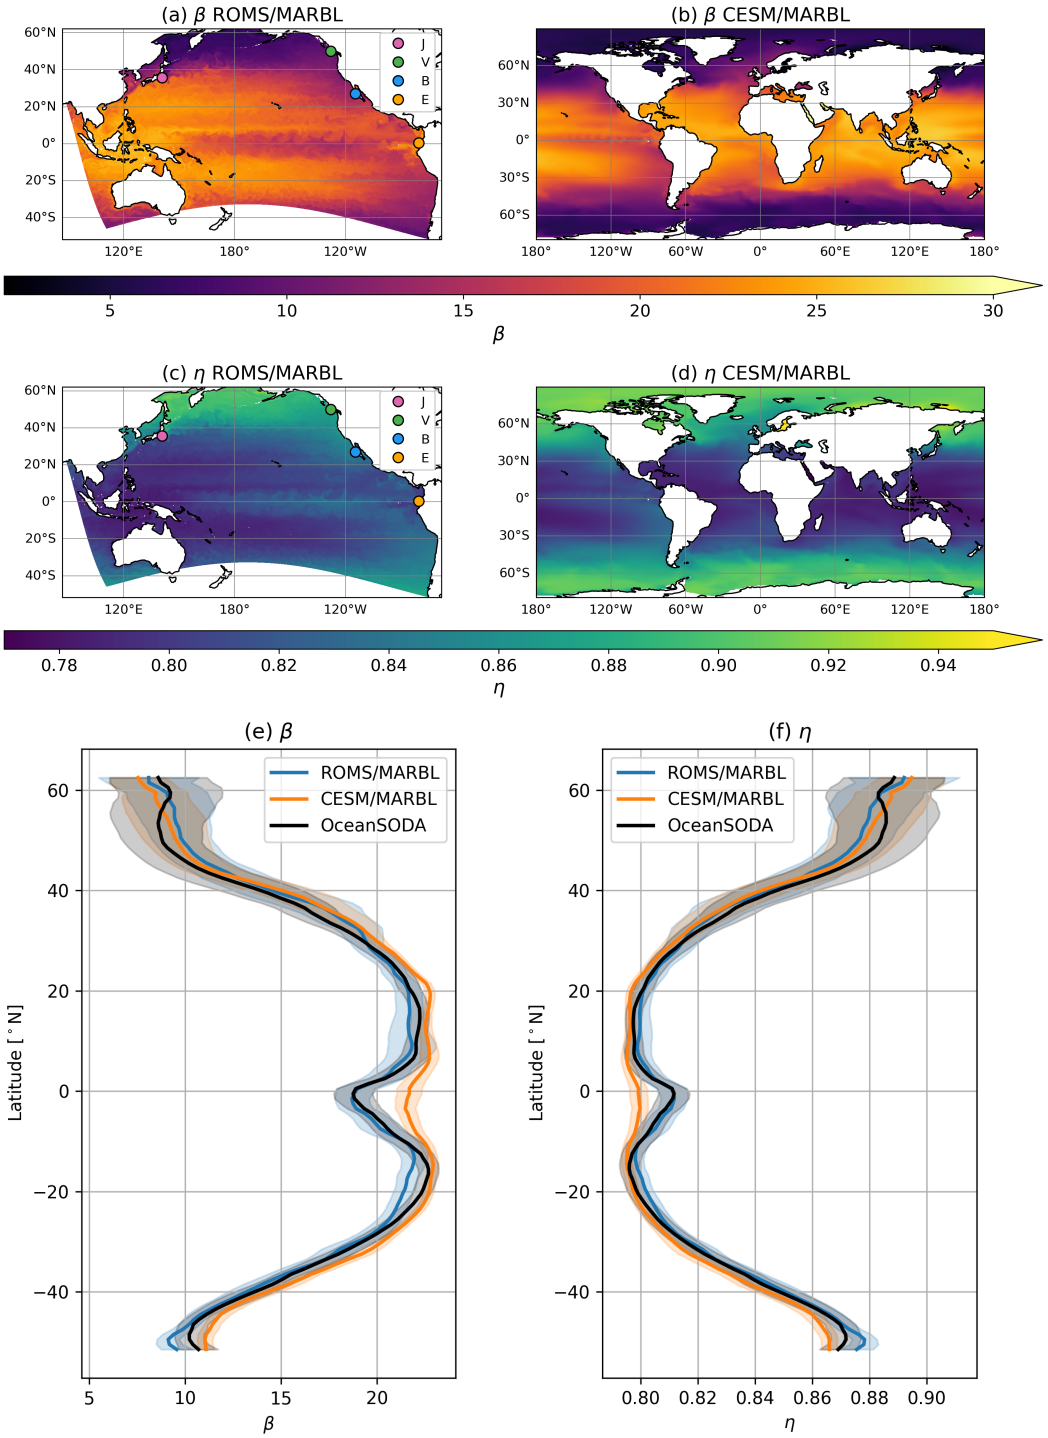


Figure 2. Carbonate sensitivities β and η from the 12 km Pacific ROMS/MARBL and 1° global CESM/MARBL control simulations. (a–d) Spatial maps of β (a,b) and η (c,d) for ROMS/MARBL (a,c) and CESM/MARBL (b,d), shown for January 1, 2000. These time-varying fields serve as input to the respective CDR tracer simulations. The four CDR intervention locations used in the ROMS experiments are marked in (a,c): J = Japan, V = Vancouver Island, B = Baja California, E = Ecuador. (e,f) Zonally averaged β (e) and η (f) as a function of latitude, comparing ROMS/MARBL and CESM/MARBL with the observation-based OceanSODA-ETHZ product (Gregor & Gruber, 2021). All three are averaged over the ROMS Pacific domain and over the two-year period 2000–2001, sampling the same region and time period. Lines show the zonal mean; shading shows temporal variability, given as the per-cell standard deviation over 2000–2001 averaged within each 1° latitude band. The strong anticorrelation between β and η arises because both are controlled by the ratio of carbonate to bicarbonate ions (Text S3).

Table 1. Truth experiments used to evaluate the CDR tracer approach. Experiment names encode location, amplitude, and intervention type (⊕ = OAE, ⊖ = DOR). For ROMS/MARBL, locations are abbreviated as single letters and amplitudes given in Gmol; for CESM/MARBL, names encode polygon numbers. This yields 1,404 experiments in total (24 ROMS/MARBL, 1,380 CESM/MARBL).

Experiment	**Location**	**Amplitude**	**Time**	**Type**	**Model**
J10⊕	Japan	10 Gmol	Jan 2000	OAE	ROMS/MARBL
J100⊕	Japan	100 Gmol	Jan 2000	OAE	ROMS/MARBL
J1000⊕	Japan	1 Tmol	Jan 2000	OAE	ROMS/MARBL
J10⊖	Japan	10 Gmol	Jan 2000	DOR	ROMS/MARBL
J100⊖	Japan	100 Gmol	Jan 2000	DOR	ROMS/MARBL
J1000⊖	Japan	1 Tmol	Jan 2000	DOR	ROMS/MARBL
V10⊕	Vancouver Island	10 Gmol	Jan 2000	OAE	ROMS/MARBL
V100⊕	Vancouver Island	100 Gmol	Jan 2000	OAE	ROMS/MARBL
V1000⊕	Vancouver Island	1 Tmol	Jan 2000	OAE	ROMS/MARBL
V10⊖	Vancouver Island	10 Gmol	Jan 2000	DOR	ROMS/MARBL
V100⊖	Vancouver Island	100 Gmol	Jan 2000	DOR	ROMS/MARBL
V1000⊖	Vancouver Island	1 Tmol	Jan 2000	DOR	ROMS/MARBL
B10⊕	Baja California	10 Gmol	Jan 2000	OAE	ROMS/MARBL
B100⊕	Baja California	100 Gmol	Jan 2000	OAE	ROMS/MARBL
B1000⊕	Baja California	1 Tmol	Jan 2000	OAE	ROMS/MARBL
B10⊖	Baja California	10 Gmol	Jan 2000	DOR	ROMS/MARBL
B100⊖	Baja California	100 Gmol	Jan 2000	DOR	ROMS/MARBL
B1000⊖	Baja California	1 Tmol	Jan 2000	DOR	ROMS/MARBL
E10⊕	Ecuador	10 Gmol	Jan 2000	OAE	ROMS/MARBL
E100⊕	Ecuador	100 Gmol	Jan 2000	OAE	ROMS/MARBL
E1000⊕	Ecuador	1 Tmol	Jan 2000	OAE	ROMS/MARBL
E10⊖	Ecuador	10 Gmol	Jan 2000	DOR	ROMS/MARBL
E100⊖	Ecuador	100 Gmol	Jan 2000	DOR	ROMS/MARBL
E1000⊖	Ecuador	1 Tmol	Jan 2000	DOR	ROMS/MARBL
P000⊕	Polygon 000	10 mol m^{-2} yr^{-1}	Jan 1999	OAE	CESM/MARBL
P001⊕	Polygon 001	10 mol m^{-2} yr^{-1}	Jan 1999	OAE	CESM/MARBL
⋮	⋮	⋮	⋮	⋮	⋮
P689⊕	Polygon 689	10 mol m^{-2} yr^{-1}	Jan 1999	OAE	CESM/MARBL
P000⊖	Polygon 000	10 mol m^{-2} yr^{-1}	Jan 1999	DOR	CESM/MARBL
P001⊖	Polygon 001	10 mol m^{-2} yr^{-1}	Jan 1999	DOR	CESM/MARBL
⋮	⋮	⋮	⋮	⋮	⋮
P689⊖	Polygon 689	10 mol m^{-2} yr^{-1}	Jan 1999	DOR	CESM/MARBL

S3, S4, and S5): a subpolar site (V, 50°N) with low β and high η; a subtropical site (B, 27°N) with intermediate β and η; an eastern-boundary equatorial site (E, 0°) with high β and low η; and a western boundary current region (J, 36°N), at the subtropical-subpolar boundary, with high temporal variability in β and η.

At each location, the surface intervention flux is applied over the first month of the simulation (January 2000) as a horizontal Gaussian with a nominal standard deviation of 300 km, masked to ocean points and rescaled to preserve the total integrated magnitude (Figure S8). We run three intervention amplitudes and both intervention types (OAE and DOR), giving six truth experiments per location and 24 in total (Table 1). Experiments are named by location, amplitude, and type, with $\oplus$ denoting OAE and $\ominus$ denoting DOR (e.g., J100$\oplus$ for a 100 Gmol OAE experiment at Japan). The amplitudes are chosen so that the total alkalinity added or DIC removed equals 10 Gmol, 100 Gmol, and 1 Tmol, spanning three orders of magnitude. This range is deliberate: larger interventions are increasingly expected to violate the linearity assumption underlying the CDR tracer method, and testing across three orders of magnitude reveals whether that assumption holds and at what amplitude it begins to break down. Applied over one month, the three amplitudes correspond to delivery rates of roughly 4×10^3, 4×10^4, and $4\times10^5\,\mathrm{mol\,s^{-1}}$. For comparison, a single ship-based release could deliver 250–1,200 $\mathrm{mol\,s^{-1}}$, and present-day coastal outfall operations deliver about $10\,\mathrm{mol\,s^{-1}}$ (Bach et al., 2026). Even the small-amplitude experiments therefore represent a large deployment by current standards, equivalent to several ships discharging continuously for a month, and the mid-amplitude experiments an ambitious regional deployment. The largest-amplitude experiments are deliberately extreme: annualized, they amount to about 4% of the $\sim$335 $\mathrm{Tmol\,yr^{-1}}$ in the optimistic global gigatonne-scale scenario of Bach et al. (2026), concentrated at a single site, and serve as an upper bound for testing the linearity assumption. The Gaussian source is designed to implicitly parameterize the early near-field mixing of the released material, which cannot be resolved at the 12 km model resolution. All experiments therefore represent the intervention after near-field mixing has taken place, and the resulting perturbations should be compared with far-field rather than near-field alkalinity ranges (Bach et al., 2026).

Because the Gaussian kernel is defined based on a single domain-averaged grid spacing rather than the local grid spacing, the realized standard deviation departs slightly from the nominal 300 km at each location. Combined with differences in the ocean footprint remaining after land masking and renormalization, the four locations experience different flux amplitudes (Figures S8, S9, S10, S11). For the largest-amplitude experiment, peak fluxes reach 44.6, 27.2, 24.8, and 24.6 $\mathrm{mol\,m^{-2}\,yr^{-1}}$ at Vancouver Island, Ecuador, Japan, and Baja California, respectively. These are 2.5–4.5 times the flux of $10\,\mathrm{mol\,m^{-2}\,yr^{-1}}$ that Zhou et al. (2024) chose in order to keep their surface pH increase generally below 0.1. The mid- and small-amplitude peak fluxes, in contrast, remain below this reference flux.

2.4.2 CESM/MARBL experiments

As our global CESM/MARBL truth experiments, we use the OAE and DOR atlasses of Zhou et al. (2024, 2026), in which the global ocean is divided into 690 polygons. For each polygon, these authors ran one OAE and one DOR integration for 15 years. In each, alkalinity was added (OAE) or DIC was removed (DOR) at the ocean surface within a single polygon at a constant rate of $10\,\mathrm{mol\,m^{-2}\,yr^{-1}}$ over the first month. We use only the experiments with interventions starting in January 1999 (the atlasses also include April, July, and October 1999 start dates), giving 1,380 experiments in total (Table 1).

Across the 1,380 truth experiments, polygon areas vary considerably. As a result, the total alkalinity addition or DIC removal ranges from approximately 100 Gmol to 1 Tmol, with median values of about 200 Gmol for exclusive economic zone (EEZ) polygons and

Table 2. ROMS CDR tracer experiments. Each prescribes the carbonate sensitivities β and η, sharing the same spatiotemporal form, indicated by s: $s(x, y, t)$ a time-varying field, $s(x, y)$ a time-constant field, and s a single constant in space and time. `REF` uses the daily, spatially varying ROMS/MARBL carbonate sensitivities; the remaining experiments degrade them, either by temporally and spatially averaging the ROMS/MARBL sensitivities (`MON`, `MEAN`, `CONST`) or by substituting carbonate sensitivities from CESM/MARBL or OceanSODA-ETHZ, regridded to the ROMS grid.

Experiment	Source	Form	Temporal resolution
`REF`	ROMS/MARBL	$s(x, y, t)$	daily
`MON`	ROMS/MARBL	$s(x, y, t)$	monthly
`MEAN`	ROMS/MARBL	$s(x, y)$	time-mean
`CONST`	ROMS/MARBL	s	time-mean
`CESM`	CESM/MARBL	$s(x, y, t)$	daily
`SODA-mon`	OceanSODA-ETHZ	$s(x, y, t)$	monthly
`SODA-clim`	OceanSODA-ETHZ	$s(x, y, t)$	monthly climatology

600 Gmol for open-ocean polygons. For context, these amplitudes fall between the second and third ROMS/MARBL experiment amplitudes (suffixes 100 and 1000 in Table 1). The atlas experiments differ from the ROMS/MARBL truth experiments in a complementary way: here the flux is spatially uniform across the polygon but the total amount varies with polygon area, whereas the ROMS/MARBL experiments fix the total amount and allow the flux to vary spatially via a Gaussian distribution. In both cases, the intervention is distributed over hundreds of kilometers and thus represents the perturbation after near-field mixing has taken place. Together, the two sets of experiments span a wider range of intervention scenarios and provide a more comprehensive test of the CDR tracer method.

2.5 CDR Tracer Experiments

The CDR tracer experiments mirror the truth experiments but use only the physical model, with the carbonate sensitivities β and η prescribed as input fields. For both ROMS and CESM, we prescribe β and η at the daily resolution of the corresponding control simulation; Figure 2 shows the fields for January 1, 2000. We refer to these setups, in which β and η are prescribed at daily resolution, as the CDR tracer *reference configurations*. In the reference configurations, β and η are derived from the coupled physics–biogeochemistry model used for the truth experiments and are resolved in time, so the dominant approximation relative to the truth is the linearization underlying the CDR tracer method (see Text S2 for minor additional approximations arising from the temporal sampling of β and η).

For CESM, we restrict ourselves to this reference configuration and do not perform additional experiments. For ROMS, we run additional experiments in which β and η are degraded relative to the reference configuration (Table 2). These additional experiments probe how robust the method is to coarser or less self-consistent sensitivity inputs, a question with direct practical importance: diagnosing β and η at full resolution requires a dedicated control run with active biogeochemistry, which is costly to spin up and to store. If degraded sensitivities prove adequate, the control output can be stored less often. More importantly, if external sensitivities work, the biogeochemical control run can be skipped altogether, avoiding its costly spin-up.

We degrade the carbonate sensitivities along two axes. The first coarsens the ROMS/MARBL sensitivities in time: from daily (`REF`) to monthly means (`MON`), to a single time-mean field (`MEAN`), to a single domain-wide constant (`CONST`), where the latter is obtained by additionally averaging over space. The `MON` and `MEAN` experiments address how frequently the control output must be stored, while `CONST` removes all spatial and temporal structure and serves as a worst-case bound on how much accuracy can be lost. In all experiments, we calculate β and η at the native daily and spatial resolution and subsequently average these fields, rather than calculating them from temporally and/or spatially averaged state variables. Because β and η are nonlinear functions of the carbonate state, the two orderings differ. However, the resulting differences are small compared to the degradation effects studied here, although regionally significant differences do occur (Figures S6 and S7; Text S2).

The second axis replaces the ROMS/MARBL sensitivities with those from independent sources, testing whether a dedicated biogeochemical control simulation can be avoided altogether. We use β and η derived from the CESM/MARBL control (`CESM`) and from the observation-based OceanSODA-ETHZ product (Gregor & Gruber, 2021). The latter is used both at its native monthly resolution (`SODA-mon`) and as a 30-year monthly climatology (`SODA-clim`). OceanSODA-ETHZ reconstructs the global surface carbonate system at monthly, 1° resolution from surface pCO_2 and alkalinity observations using a machine-learning mapping approach. We compute β and η from its reconstructed alkalinity and DIC, together with temperature and salinity, using `pyCO2SYS` with constant nutrient concentrations, since OceanSODA-ETHZ does not provide nutrient fields. The time means and temporal standard deviations of β and η from all three sources are shown in Figures S3 and S4, and zonal means of each month of the 30-year OceanSODA-ETHZ climatology in Figure S5. All externally derived sensitivities are regridded to the ROMS grid.

The sensitivities from the external sources differ from the ROMS/MARBL sensitivities in instructive ways. The 1° CESM/MARBL and OceanSODA-ETHZ do not resolve mesoscale dynamics, whereas the 12 km ROMS/MARBL fields show prominent mesoscale structure (Figure 2). Beyond this resolution difference, CESM/MARBL also exhibits large-scale biases in β and η relative to ROMS/MARBL and OceanSODA-ETHZ, particularly at low latitudes (Figures 2, S3, S4). The robustness experiments in Table 2 quantify how these differences propagate to the reconstructed CDR efficiency, presented in Section 6.

2.6 CDR efficiency

The primary metric for MRV of ocean-based CDR is the cumulative oceanic CO_2 uptake implied by the intervention. The CDR efficiency, denoted $\mathcal{E}$ for the truth and $\tilde{\mathcal{E}}$ for the CDR tracer estimate, normalizes this uptake by the total amount of alkalinity added (OAE) or DIC removed (DOR). CDR efficiency is the primary focus of the verification in Section 4. The CDR efficiency can equivalently be computed from either the DIC perturbation (an inventory) or the air-sea CO_2 flux implied by the intervention, since the two are related by the integrated carbon budget and differ only in the diagnostics they require.

Using the DIC perturbation field, the efficiencies for OAE and DOR are

$$\mathcal{E}^{\mathrm{OAE}}(t) = \frac{\int \delta\mathrm{DIC}\,\mathrm{d}V}{\int_{t_0}^{t}\int S_{\delta\mathrm{ALK}}\,\mathrm{d}V\,\mathrm{d}t'}, \qquad \tilde{\mathcal{E}}^{\mathrm{OAE}}(t) = \frac{\int c_{\delta\mathrm{DIC}}\,\mathrm{d}V}{\int_{t_0}^{t}\int S_{\delta\mathrm{ALK}}\,\mathrm{d}V\,\mathrm{d}t'}, \tag{18}$$

$$\mathcal{E}^{\mathrm{DOR}}(t) = 1 - \frac{\int \delta\mathrm{DIC}\,\mathrm{d}V}{\int_{t_0}^{t}\int S_{\delta\mathrm{DIC}}\,\mathrm{d}V\,\mathrm{d}t'}, \qquad \tilde{\mathcal{E}}^{\mathrm{DOR}}(t) = 1 - \frac{\int c_{\delta\mathrm{DIC}}\,\mathrm{d}V}{\int_{t_0}^{t}\int S_{\delta\mathrm{DIC}}\,\mathrm{d}V\,\mathrm{d}t'}. \tag{19}$$

The denominator is the time-integrated source, i.e. the total alkalinity added or DIC removed up to time t. We use the sign convention $S_{\delta\mathrm{ALK}} > 0$ for OAE and $S_{\delta\mathrm{DIC}} < 0$ for DOR. The numerator is the diagnosed ocean carbon response, evaluated from the true perturbation δDIC or its tracer estimate $c_{\delta\mathrm{DIC}}$. For OAE, this is the increase in ocean DIC resulting from the CO_2 drawn into the ocean. For DOR, the numerator is instead the DIC deficit remaining in the ocean. Immediately after removal it equals the amount removed, giving a ratio of 1 and $\mathcal{E} = 0$. As the deficit is refilled, new CO_2 is drawn from the atmosphere into the ocean; this atmospheric drawdown, not the carbon already extracted and stored elsewhere, is the actual climate benefit of DOR. As it accumulates, the ratio falls below 1. Thus 1 minus the ratio, the realized removal, rises toward its equilibrium value, approaching $\mathcal{E} = 1$ if the deficit remains in contact with the atmosphere until fully equilibrated. The formulation in Eqs. (18) and (19) requires instantaneous snapshots of δDIC (truth experiments) and $c_{\delta\mathrm{DIC}}$ (CDR tracer experiments), both available from our ROMS simulations. We therefore use this formulation for the CDR tracer validation in ROMS (Section 4.1).

Equivalently, the CDR efficiency can be expressed from the cumulative air-sea CO_2 flux perturbation,

$$\mathcal{E}^{\mathrm{OAE}}(t) = \frac{\int_{t_0}^{t}\int \delta F_{\mathrm{CO_2}}\,\mathrm{d}A\,\mathrm{d}t'}{\int_{t_0}^{t}\int S_{\delta\mathrm{ALK}}\,\mathrm{d}V\,\mathrm{d}t'}, \qquad \tilde{\mathcal{E}}^{\mathrm{OAE}}(t) = \frac{\int_{t_0}^{t}\int \delta \tilde{F}_{\mathrm{CO_2}}\,\mathrm{d}A\,\mathrm{d}t'}{\int_{t_0}^{t}\int S_{\delta\mathrm{ALK}}\,\mathrm{d}V\,\mathrm{d}t'}, \tag{20}$$

$$\mathcal{E}^{\mathrm{DOR}}(t) = -\frac{\int_{t_0}^{t}\int \delta F_{\mathrm{CO_2}}\,\mathrm{d}A\,\mathrm{d}t'}{\int_{t_0}^{t}\int S_{\delta\mathrm{DIC}}\,\mathrm{d}V\,\mathrm{d}t'}, \qquad \tilde{\mathcal{E}}^{\mathrm{DOR}}(t) = -\frac{\int_{t_0}^{t}\int \delta \tilde{F}_{\mathrm{CO_2}}\,\mathrm{d}A\,\mathrm{d}t'}{\int_{t_0}^{t}\int S_{\delta\mathrm{DIC}}\,\mathrm{d}V\,\mathrm{d}t'}. \tag{21}$$

where $\delta F_{\mathrm{CO_2}}$ is the full air-sea flux perturbation in the truth experiments (Eq. (5)) and $\delta \tilde{F}_{\mathrm{CO_2}}$ its linearized CDR tracer counterpart (Eq. (14)). Note that the DOR efficiencies defined in Eq. (21) are positive because $S_{\delta\mathrm{DIC}}$ is negative by convention. Unlike the inventory form, this requires the cumulative flux and should be evaluated from time-averaged output. We use it for the CDR tracer validation in CESM (Section 4.2), where we have monthly mean flux diagnostics available.

3 Carbonate System Nonlinearity

The carbonate system is nonlinear: the response of dissolved CO_2 to a change in DIC or alkalinity depends on the background carbonate state. The CDR tracer method sidesteps this by linearizing about the control state, prescribing β and η from the control run so that they are unaffected by the intervention. The true sensitivities, however, respond to the carbonate state shift that the intervention itself produces. Figure 3(a,b) shows β and η as functions of ALK and DIC. β increases with ALK and decreases with DIC, while η decreases with ALK and increases with DIC. For DOR, DIC is removed directly from the ocean, resulting in negative δDIC at fixed alkalinity. For OAE, alkalinity is added to the ocean, resulting in positive δALK; this shifts the carbonate equilibrium, drawing down atmospheric CO_2 and producing a positive δDIC response. Both interventions move the carbonate state in the direction of increasing β and decreasing η (white arrows), so the control values used by the CDR tracer method are biased low in β and high in η.

The sign of the resulting error in the CDR tracer estimate can be anticipated from a box model. The box model reduces the full problem to a single well-mixed surface box of depth h with no horizontal transport and admits closed-form solutions (Text S1). Figure 3(c–f) shows the resulting CO_2 flux anomaly and CDR efficiency for illustrative parameter values ($\beta^{\mathrm{ctrl}} = 15$, $\beta^{\mathrm{pert}} = 24$, $\eta^{\mathrm{ctrl}} = 0.85$, $\eta^{\mathrm{pert}} = 0.82$; remaining parameters in Text S1). Because β^{ctrl} is smaller than the true β, the CO_2 flux anomaly estimated by the CDR tracer rises faster and peaks higher than the truth (Figure 3(c)). After the intervention stops at time T, it also decays faster. The instantaneous flux error

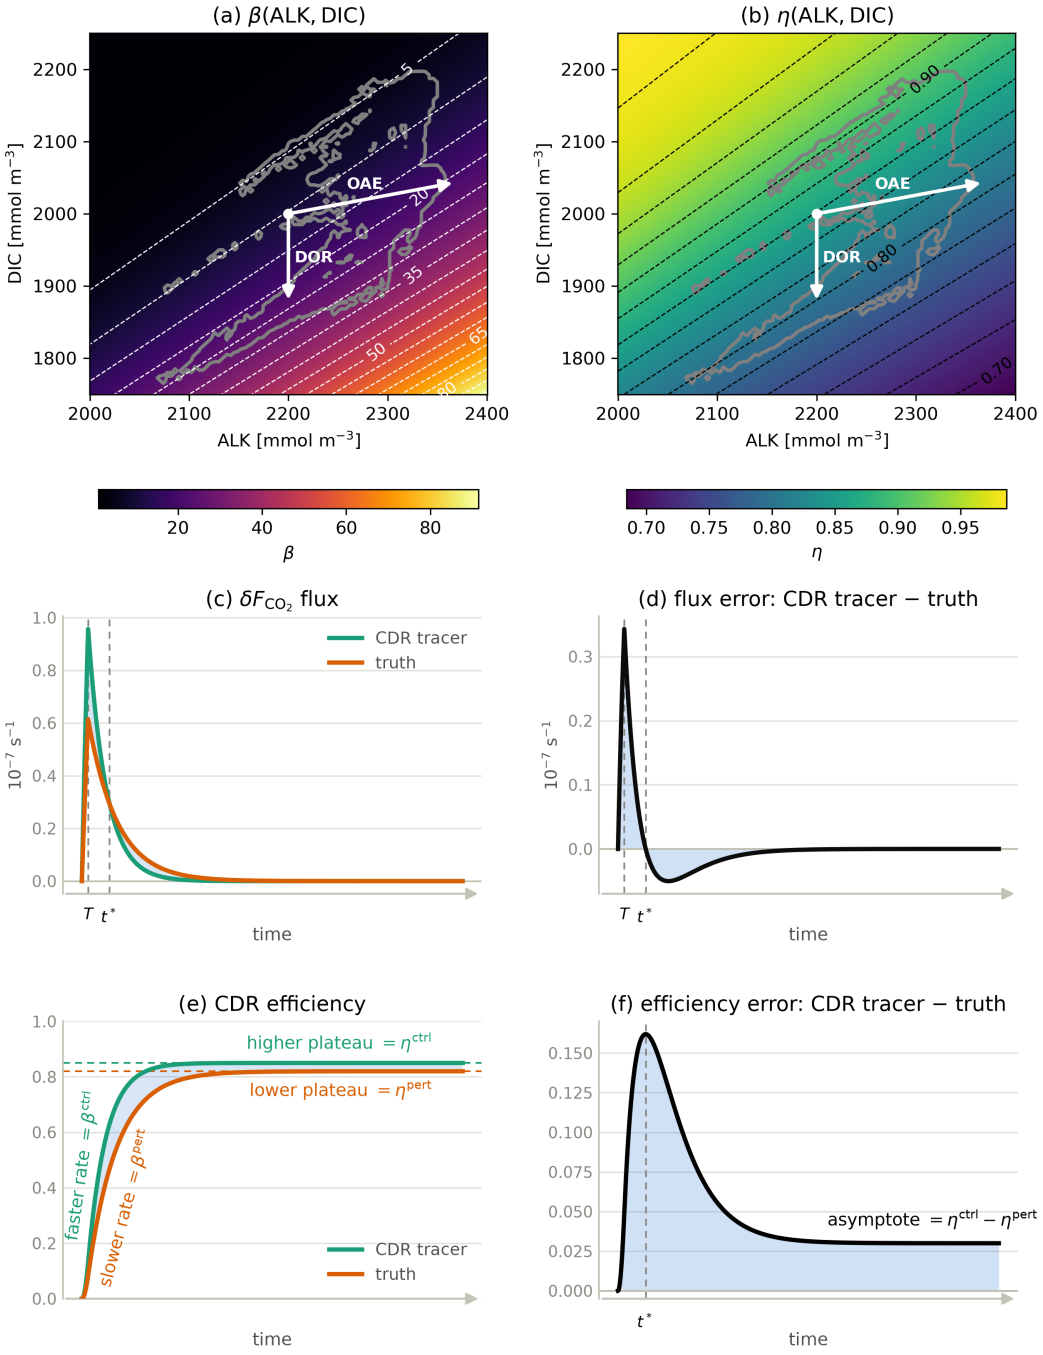


Figure 3. (a,b) Carbonate sensitivities β (a) and η (b) as functions of ALK and DIC, computed at fixed temperature, salinity, phosphate, and silicate set to long-term domain-mean values of the ROMS/MARBL Pacific control run ($T_0 = 21.9\,^\circ$C, $S_0 = 34.5$ psu, $(\mathrm{PO}_4^{3-})_0 = 0.5\,\mathrm{mmol\,m^{-3}}$, $(\mathrm{SiO}_3^{2-})_0 = 3.2\,\mathrm{mmol\,m^{-3}}$). Contour intervals (thin dashed contours) are 5 for β and 0.02 for η. Overlaid are the outlines of the ALK–DIC mean state space occupied by the Pacific ROMS/MARBL control run (gray contours). The white arrows show schematically the direction of the carbonate state shift induced by OAE and DOR; amplitudes are exaggerated for visibility. (c–f) Air-sea CO_2 flux anomaly and CDR efficiency computed with a box model (Text S1), for the CDR tracer method and truth. The intervention is applied over $0 \le t \le T$ and parameter values are given in Text S1. (c) Flux anomaly δF_{CO_2} caused by the intervention. Fluxes are normalized to a unit total intervention (Text S1), so the y-axis carries units of time^{-1} rather than mass per time. (d) Flux error, which changes sign at the crossover time t^*. (e) CDR efficiency, i.e., the time-integrated flux in (c). (f) Efficiency error, positive throughout, maximal at t^*, and converging to $\eta^{\mathrm{ctrl}} - \eta^{\mathrm{pert}}$.

is therefore not sign-definite and changes sign exactly once, at a crossover time t^* (Figure 3(d)). Integrating the flux in time, the estimated efficiency curve rises faster and converges to a higher plateau than the truth (Figure 3(e)), since in the box model the equilibration rate scales as $1/\beta$ and η sets the plateau of the efficiency curve. The efficiency error is positive at all times, maximal at t^*, and converges to $\eta^{\mathrm{ctrl}} - \eta^{\mathrm{pert}}$ (Figure 3(f)). In other words, the CDR tracer is expected to overestimate CDR efficiency at all times. To probe the robustness of this conclusion, we extend the box model numerically to time-varying $\beta(t)$, $\eta(t)$, and mixed layer depth $h(t)$ in Text S1 and Figure S1, partially accounting for the temporal evolution of the carbonate state and mixed layer as the intervention plume spreads. We find the same error structure in all cases (Figure S2), suggesting that the systematic overestimation of efficiency by the CDR tracer method is a robust feature that may extend to full-fledged numerical ocean models. We test this expectation against the full ROMS/MARBL and CESM/MARBL experiments in Section 4.

The magnitude of the error in the CDR tracer flux and efficiency estimates depends on how far the intervention displaces β and η from their control values. This displacement is a direct measure of the nonlinearity of the problem. In the remainder of this section we characterize the shift in surface ALK and DIC induced by our ROMS/MARBL intervention experiments, and the resulting changes in β and η, using only the control and largest-amplitude intervention truth runs. This provides the physical context needed to interpret the CDR tracer verification results in Section 4.1.

Two physical mechanisms dominate how the intervention signal spreads and dilutes after release: horizontal advection and vertical mixing. Figure 4 investigates both. Panels (a)–(d) show the 90% footprint of δALK six months after the OAE intervention (July 1, 2000), overlaid on the time-mean surface current speed and direction. The Japan plume has spread more than 30° to the east, carried by the vigorous Kuroshio current (panel (a)), and the Ecuador plume has spread a similar distance to the west, following the North and South Equatorial currents (panel (d)). In contrast, the Vancouver Island and Baja California plumes have barely moved beyond their intervention footprints, owing to the weak currents in those regions (panels (b),(c)). The footprint of δDIC in both the OAE and DOR experiments closely follows those of δALK (Figure S12).

Panels (e)–(g) show the vertical profiles of the response inventory, integrated horizontally over the full model domain, six months after the intervention start. Focusing on the δALK profiles in the OAE experiments (panel (e)), the four locations show markedly different vertical structures. The Japan δALK anomaly exhibits a subsurface peak at approximately 50 m and extends to around 300 m, reflecting seasonal re-stratification: the winter mixed layer in the Kuroshio region reaches up to 200–300 m (Figure S13(a)), but by July it has shoaled to less than 10 m (Figure S14(a)), leaving a portion of the anomaly stranded below the summer mixed layer. Off Ecuador, by contrast, the shallow mixed layer year-round (Figures S13, S14, S15) confines the δALK signal near the surface, producing the most strongly surface-intensified profile. Baja California shows an approximately uniform signal across the upper 70 m before declining, while the Vancouver Island δALK is also surface-intensified.

The δDIC profiles in the OAE experiments (panel (f)) are more surface-intensified than the δALK profiles (panel (e)), since atmospheric CO_2 uptake driven by the alkalinity perturbation occurs at the surface. However, vertical mixing still redistributes the δDIC signal with depth, and similar mixing-depth signatures as seen in panel (e) emerge in the δDIC profiles later in the year (Figure S16). Vancouver Island shows the largest δDIC mass among the OAE experiments, reflecting its high efficiency: it draws down the most CO_2 per unit alkalinity added, as discussed in Section 4.1. The δDIC profiles in the DOR experiments (panel (g)) closely mirror the δALK profiles of the OAE experiments (panel (e)) in shape but with opposite sign, since the initial DIC deficit is spread by the same physical transport and mixing processes as the alkalinity perturbation. Near the surface, however, the profiles start to diverge from the δALK profiles: CO_2 draw-

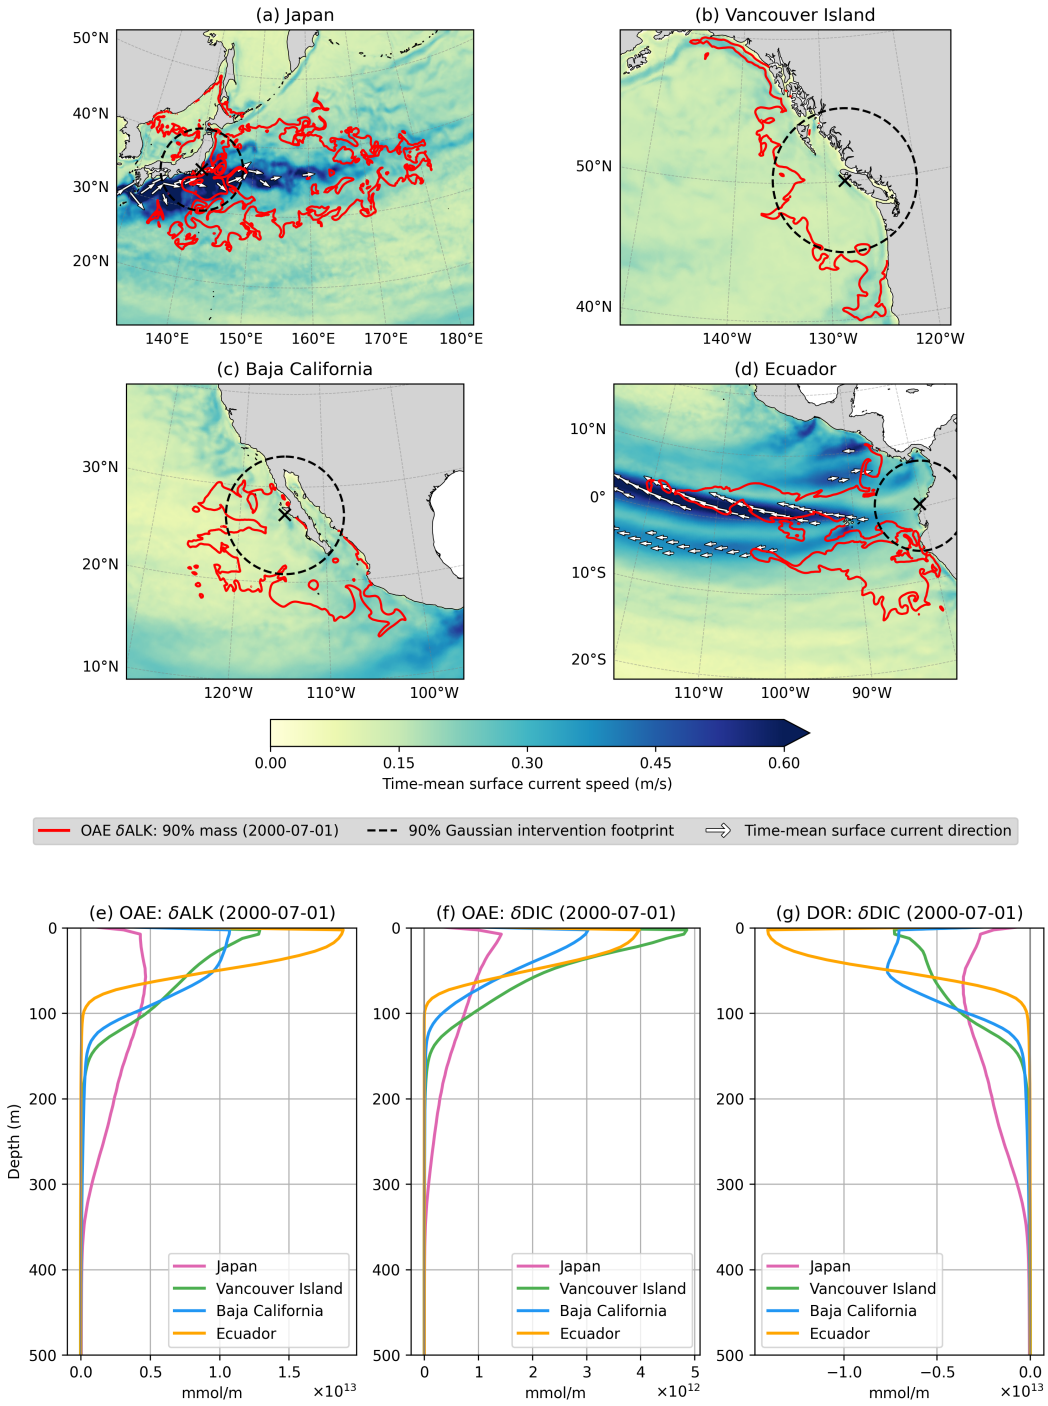


Figure 4. (a–d) Time-mean surface current speed (shading) and direction (white arrows) around each intervention location: Japan (a), Vancouver Island (b), Baja California (c), and Ecuador (d). Arrows are shown only where the time-mean surface current speed exceeds 0.3 m/s and are subsampled to every 15th grid point for clarity. Red contours show the 90% footprint of the depth-integrated δALK from the largest-amplitude OAE experiment on July 1, 2000, where δ denotes intervention minus control (Eq. (4)). The black crosses and dashed circles mark the center and 90% Gaussian footprint of the respective interventions. (e–g) Vertical profiles of the response inventory density (mmol m^{-1}), integrated horizontally over the full model domain, for the largest-amplitude experiments on July 1, 2000, six months after the intervention start. Panels show δALK for OAE (e), δDIC for OAE (f), and δDIC for DOR (g).

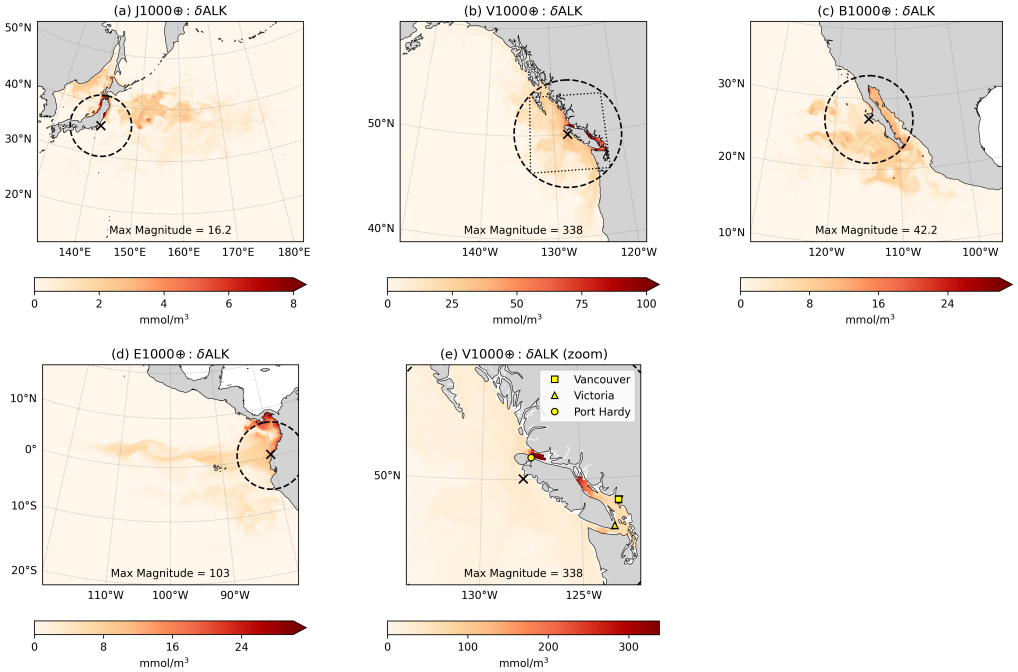


Figure 5. Response of surface alkalinity, δALK, in the largest-amplitude OAE truth experiments for the Japan (a), Vancouver Island (b,e), Baja California (c), and Ecuador (d) intervention locations. δ denotes intervention minus control (Eq. (4)). All panels show uppermost-layer concentrations (mean layer thickness 2.5 m) on July 1, 2000, six months after the intervention start. Black crosses and dashed circles mark the intervention center and 90% Gaussian footprint. Colorbars are saturated in all panels except (e) to keep the plumes visible, with ranges indicating typical concentration scales. Panel (e) shows a zoomed-in view of the region outlined by the dotted box in panel (b), with an unsaturated colorbar to highlight extreme local concentrations around Vancouver Island. The domain-wide absolute maximum of each response field is annotated at the bottom of each panel.

down gradually refills the DIC deficit from above, modifying the surface signal relative to panel (e).

Together, horizontal advection and vertical mixing control the surface concentration of the intervention signal: horizontal advection dilutes it by spreading the perturbation over a larger area, while vertical mixing spreads it away from the surface. The resulting surface δALK fields in the largest-amplitude OAE experiments are shown in Figure 5. At Japan, vigorous Kuroshio advection and deep winter mixing followed by summer restratification produce the lowest surface concentrations; the plumes are stretched eastward along the Kuroshio over a large area (Figure 5(a)). At Vancouver Island, weak currents and the absence of deep mixing keep the δALK signal compact and surface-intensified near the injection site, resulting in the highest surface concentrations of the four locations (Figure 5(b)), including extreme local values (Figure 5(e)) discussed below. The Baja California signal mixes down but does not spread horizontally, while the Ecuador signal spreads westward along the equatorial currents but remains near the surface; both show intermediate concentrations (Figure 5(c,d)). The same concentration ranking holds for δDIC in both the OAE and DOR experiments (Figures S17–S20): highest at Vancouver Island, intermediate at Baja California and Ecuador, and lowest at Japan. The overall difference in typical surface concentrations between Japan and Vancouver Island

is roughly one order of magnitude. Most of this contrast reflects the dynamical differences discussed above, but approximately a factor of 2 arises from differences in intervention flux amplitude: the peak surface flux at Vancouver Island is approximately twice as large as at Japan, owing to differences in the realized Gaussian width and ocean footprint after land masking (Section 2.4.1, Figures S8–S11).

Figure 5(e) zooms into the Vancouver Island region with an unsaturated colorbar, revealing extreme local δALK values. These extremes occur in three isolated spots: the Salish Sea northwest of Vancouver, the Strait of Juan de Fuca southwest of Victoria, and near Port Hardy at the northern tip of Vancouver Island, with a domain maximum of $338\,\mathrm{mmol\,m^{-3}}$. The signal accumulates in these narrow, shallow passages (depths of 10–15 m, with the Strait of Juan de Fuca represented by a single grid cell) where currents are too weak to export it efficiently. The same locations show extreme values in δDIC for both the OAE and DOR experiments (Figure S18). The domain maxima at the other locations (16, 42, and $103\,\mathrm{mmol\,m^{-3}}$ for Japan, Baja California, and Ecuador, respectively; Figure 5(a,c,d)) are similarly dominated by isolated values near the injection sites, though of smaller amplitude. Elevated surface concentrations of δALK and δDIC, whether extreme or more broadly distributed across the plume, shift the carbonate state away from the control and drive deviations in β and η, as we quantify next.

Figure 6 shows the relative difference in β and η between the control run (β_{REF}, η_{REF}) and the perturbed truth runs of the largest-amplitude OAE experiments, six months after the intervention start (see also Figures S21–S23). The spatial pattern follows the spread of the δALK and δDIC plumes (Figures 5, S17–S20): β and η are only modified where the intervention signal has reached. The signs are consistent with Figure 3(a,b): the control run underestimates β and overestimates η relative to the perturbed run, i.e., $\beta_{\mathrm{REF}} < \beta_{\mathrm{pert}}$ and $\eta_{\mathrm{REF}} > \eta_{\mathrm{pert}}$, which appear as negative (blue) and positive (red) relative differences in Figure 6, respectively. Across all four locations, the relative change in η is approximately an order of magnitude smaller than that in β, as the colorbar ranges and annotated maxima show. This identifies β as the dominant source of nonlinearity error in the CDR tracer method, and suggests that the long-term efficiency asymptote is less affected by nonlinearity than the short-term equilibration rate.

The magnitude of the relative differences in β and η follows the surface concentration ranking established above: largest at Vancouver Island, smallest at Japan, and intermediate at Baja California and Ecuador. The extreme local values of δALK and δDIC discussed above are reflected in correspondingly large local deviations in β and η, most prominently in the Salish Sea at Vancouver Island. For β, typical plume values are around 1% at Japan (domain maximum 6.6%) and around 15% at Vancouver Island (domain maximum 72.6% in the Salish Sea); Baja California and Ecuador fall in between, with typical plume values of 5–10% and domain maxima of 14.7% and 22.5%, respectively. How these deviations in β and η propagate into CDR tracer estimates of CO_2 flux anomaly, CDR efficiency, and pH is assessed in Section 4.

4 Verification of CDR tracer method

We now test the CDR tracer method against the coupled truth simulations, first in ROMS and then in CESM.

4.1 ROMS

We evaluate the accuracy of the CDR tracer method for the ROMS/MARBL experiments across the four release locations and three intervention amplitudes, using the reference carbonate sensitivities at daily resolution (`REF`, Table 2). Because the CDR tracer equations are linear, the resulting efficiency estimate is independent of the intervention amplitude (Section 2.2). Figure S24 confirms this, with relative differences across the three

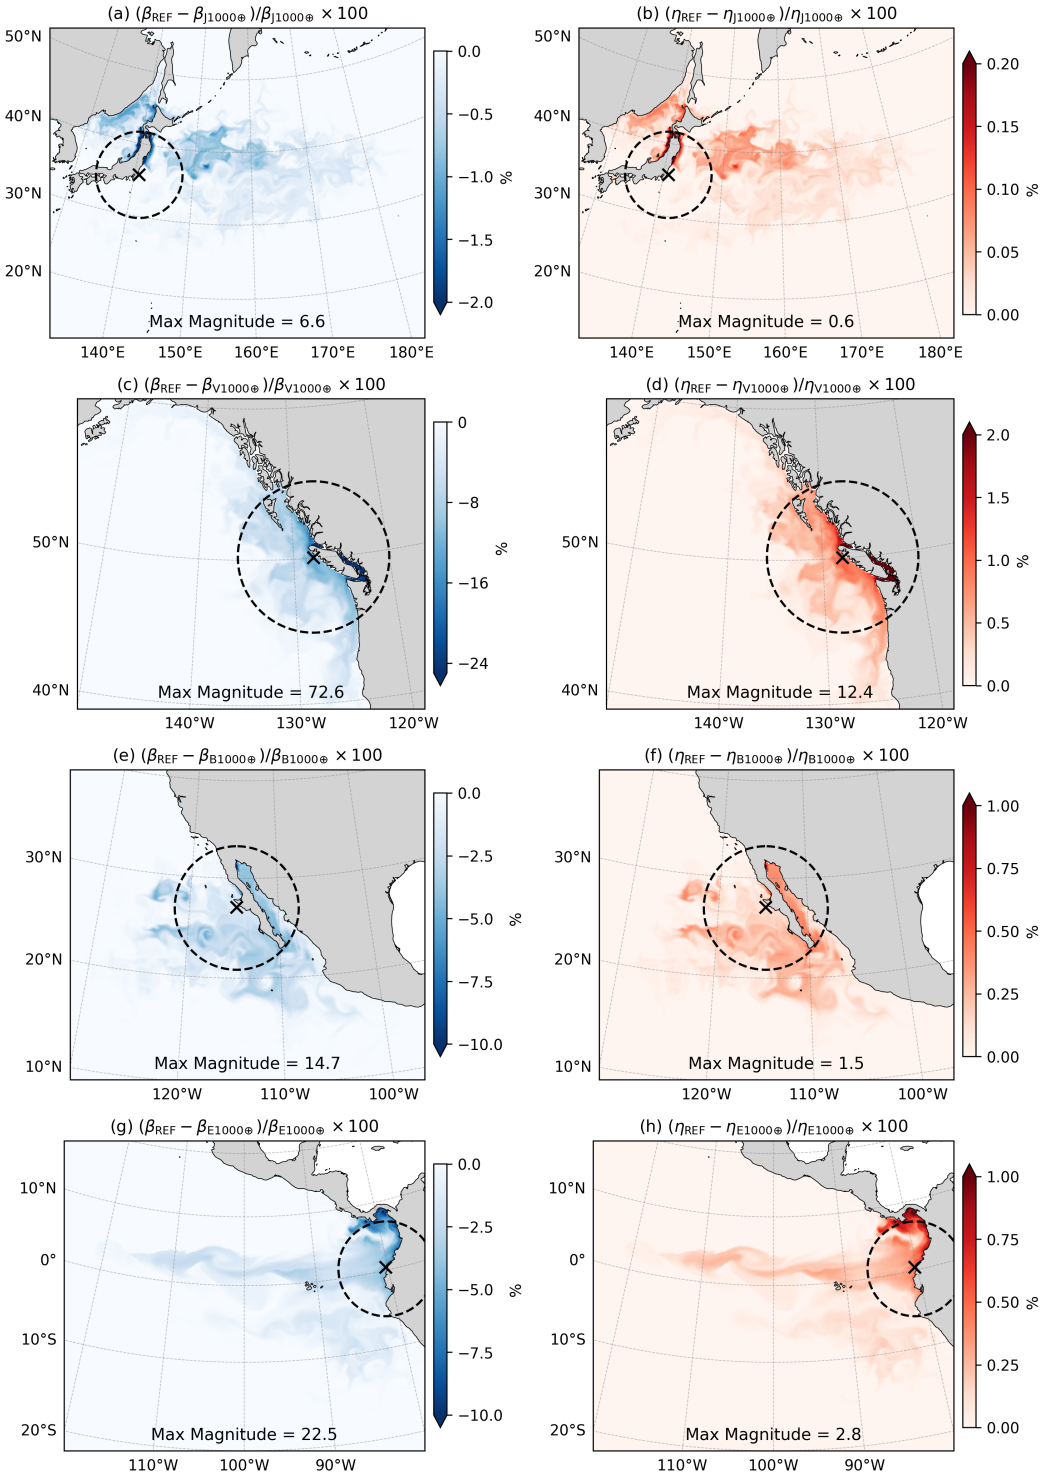


Figure 6. Relative difference in carbonate sensitivities β (left column) and η (right column) between the control run and the largest-amplitude OAE truth experiments, for the Japan (a,b), Vancouver Island (c,d), Baja California (e,f), and Ecuador (g,h) experiments. All panels are shown for July 1, 2000, six months after the start of the intervention. Colorbars are saturated to highlight the large-scale pattern; the domain-wide maximum magnitude is annotated at the bottom of each panel. The black cross marks the intervention center and the dashed circle the 90% Gaussian intervention footprint.

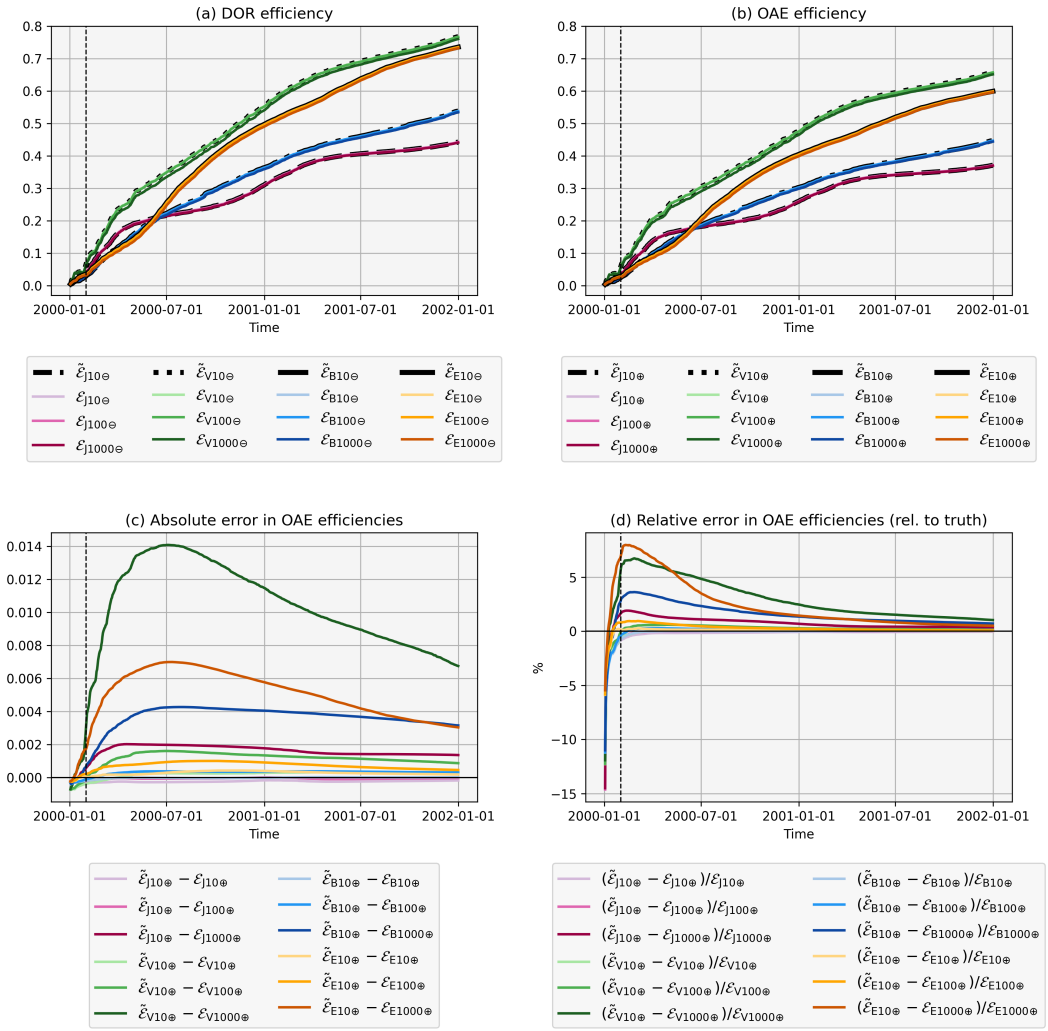


Figure 7. (a) DOR efficiency curves for all 12 ROMS/MARBL truth experiments, colored by release location (pink: Japan, green: Vancouver Island, blue: Baja California, orange: Ecuador) with darker shades indicating larger release amplitudes. Colored lines show the true efficiency $\mathcal{E}^{\mathrm{DOR}}$, thick black lines show the CDR tracer estimate $\tilde{\mathcal{E}}^{\mathrm{DOR}}$ for each location (Eq. (19)). (b) As (a), but for OAE efficiencies $\mathcal{E}^{\mathrm{OAE}}$ and $\tilde{\mathcal{E}}^{\mathrm{OAE}}$ (Eq. (18)). (c) Absolute error of the CDR tracer estimate of OAE efficiency, $\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}}$. (d) As (c), but relative error, $(\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}})/\mathcal{E}^{\mathrm{OAE}} \times 100$ in %. In all panels, the vertical dashed line marks the end of the intervention (February 1, 2000).

amplitudes of $\mathcal{O}(10^{-6})$% for DOR and $\mathcal{O}(10^{-8})$% for OAE. We therefore compare all truth experiments at a given location against a single CDR tracer estimate, obtained from the smallest-amplitude intervention ($\tilde{\mathcal{E}}_{X10}$); any other choice would yield identical results.

Figure 7 shows CDR efficiency curves for all 24 ROMS/MARBL truth experiments alongside their CDR tracer estimates, with DOR in panel (a) and OAE in panel (b). We first consider the truth experiments (colored lines). The three amplitude curves within each location are nearly indistinguishable, confirming that the truth experiments behave close to linearly across the three orders of magnitude tested. Vancouver Island and Ecuador show higher efficiency than Baja California and Japan, consistent with the analysis in Section 3: stronger vertical mixing at the latter two locations carries the alkalinity anomaly or DIC deficit away from the surface, reducing contact with the atmosphere and lowering efficiency. The DOR and OAE efficiency curves are qualitatively very similar in shape, consistent with Zhou et al. (2026). The OAE curves (panel b) reach lower values than the DOR curves (panel a), because the theoretical maximum efficiency for OAE is η rather than 1. The efficiency curves are still rising at the end of the two-year period, as full equilibration can take a decade or longer. A longer-term evaluation over 15 years is provided by the CESM experiments (Section 4.2).

The CDR tracer estimates (thick black dashed lines) are nearly indistinguishable from the truth curves for both DOR and OAE (Figure 7(a,b)). The lower row of Figure 7 quantifies the error. Since the DOR and OAE relative error curves are very similar (Figure S25), we focus on the OAE errors. Panel (c) shows the absolute error $\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}}$, which increases with intervention amplitude, as expected: larger perturbations displace the carbonate state further from the control, increasing the nonlinearity error. The CDR tracer method systematically overestimates CDR efficiency, as predicted by the box model in Section 3, except during the early phase of the intervention (left of the dashed vertical line), which we discuss below. The absolute errors peak between 3.5 to 7 months after the intervention start, with the largest errors at Vancouver Island, which also has the largest absolute efficiency values.

Panel (d) shows the relative error $(\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}})/\mathcal{E}^{\mathrm{OAE}}$ and reveals a distinct two-phase behavior. In the first few days of the intervention the relative error is large and negative, reaching down to -15%. This reflects division by near-zero efficiency values at early times; the corresponding absolute error remains below 0.001 in magnitude for all experiments throughout this initial phase (panel (c)). After this initial phase, the error turns positive, consistent with the box model prediction of systematic overestimation. For the largest-amplitude OAE experiments, it peaks in early to late February, one to four weeks after the end of the intervention, with maximum values of 1.9, 3.6, 6.7, and 8.0% at Japan, Baja California, Vancouver Island, and Ecuador, respectively. The β deviation ranking established in Section 3 (smallest at Japan, largest at Vancouver Island, and intermediate at Baja California and Ecuador) explains much of this ordering, but the spatial extent of the plume over which the deviation is integrated also matters. Vancouver Island has the largest local β deviation, while Ecuador's is intermediate but integrated over a much larger area; the two contributions are of comparable magnitude, resulting in similar peak errors of 6.7 and 8.0%. More generally, less efficient interventions tend to be estimated more accurately by the CDR tracer method, because both CDR efficiency and relative CDR tracer error decrease with increased vertical mixing. Japan and Baja California, where vertical mixing is stronger, indeed show not only the smallest efficiencies but also the smallest relative errors. Because the largest-amplitude experiments are deliberately extreme (Section 2.4.1), the errors discussed above represent an upper bound on the nonlinearity error; for the mid-amplitude experiments, the relative error remains below 1% at all locations and times after the initial phase.

The two-phase error pattern seen in Figure 7 reflects two distinct contributions: nonlinearity of the biogeochemical system (second phase), and inexactness in the implementation of the CDR tracer linearization (first phase). This is analogous to the distinc-

tion in adjoint modeling between errors due to model nonlinearity and errors from inexact adjoint code (see Appendix A in Loose et al., 2020). Some of this inexactness arises from discretization: the CDR tracer equations are derived by linearizing the continuous equations (in Section 2.1) and then discretized independently of the MARBL code, so their discrete form does not exactly match the linearization of the discrete truth model. In our implementation, this manifests, for example, in the order of operations between the MARBL biogeochemistry (used by the truth experiments) and the passive tracer module (used by the CDR tracer experiments) within a single time step: the CDR tracer air–sea gas exchange term is evaluated using tracer concentrations at the beginning of the time step, whereas the MARBL DIC perturbation has already been partially updated by advection before air–sea gas exchange acts on it. Further inexactness arises from differences between the carbonate chemistry solver used offline to compute β and η (`pyCO2SYS`) and the one used online by MARBL, and from the daily temporal sampling of β and η (Text S2).

To separate the nonlinearity of the biogeochemical system from implementation inexactness, Figure S26 shows the efficiency differences between truth experiments at different amplitudes (small minus medium, and small minus large). Because these differences involve only truth experiments and no CDR tracer output, they isolate the nonlinearity of the biogeochemical system independent of the CDR tracer method. Comparing Figure S26(f) to Figure 7(d), the initial negative phase in panel (d) is absent in Figure S26(f): differences between small- and large-amplitude truth experiments start at zero and grow positive from the outset (Figure S26(c–f)), confirming that the initial negative phase in Figure 7(d) is an implementation artifact rather than a signature of system nonlinearity. Once the relative error in Figure 7(d) has also turned positive, it closely mirrors the amplitude-dependent differences in Figure S26(e,f), indicating that the dominant source of error at this stage is system nonlinearity rather than implementation inexactness. As expected, the implementation inexactness error is small and only apparent in the first few days of panel (d), where division by near-zero efficiency values inflates its apparent magnitude.

The CDR efficiency is a quantity integrated over space and time. Figure 8 provides a more stringent field-level assessment by comparing the CDR tracer estimate $c_{\delta\mathrm{DIC}}$ directly against the truth δDIC as spatial snapshots, without the smoothing effect of space-time integration. The true δDIC fields are shown for the largest-amplitude Vancouver Island and Japan OAE experiments (panels (b,c)), and the CDR tracer error $c_{\delta\mathrm{DIC}} - \delta\mathrm{DIC}$ in panels (e,f). Consistent with the systematic overestimation of CDR efficiency seen in Figure 7, the error is positive throughout the plume, with no evidence of compensating errors that would be reflected by a mixed positive-negative error pattern.

The magnitude of the error differs markedly between the two locations. At Japan, the error is approximately two orders of magnitude smaller than the signal (typical plume values of $\sim$3 mmol m^{-3} against an error of $\sim$0.03 mmol m^{-3}). At Vancouver Island, the error is roughly one order of magnitude smaller than the signal over most of the plume, but reaches values of the same order of magnitude in the narrow coastal passages near Port Hardy, Victoria and Vancouver (domain maxima 86.6 vs. 65.2 mmol m^{-3}), where we identified extreme deviations in β in Section 3 (Figures 5, 6). In these passages, the δALK signal is represented accurately by the CDR tracer (panels (a,d); domain maxima 338 vs. 0.14 mmol m^{-3}, three orders of magnitude smaller), confirming that the passive tracer transport is correct. The large δDIC error arises instead from the nonlinearity: the extreme local accumulation of alkalinity produces a drastic, unrepresented change in β, causing the CDR tracer to overestimate CO_2 drawdown in these areas by approximately a factor of two. Despite these local extremes, their impact on the domain-integrated CDR efficiency error remains limited (Figure 7), as the narrow passages represent only a small fraction of the total area covered by the intervention plume. However, the local misrepresentation of δDIC does affect the local pH reconstruction, as we show next.

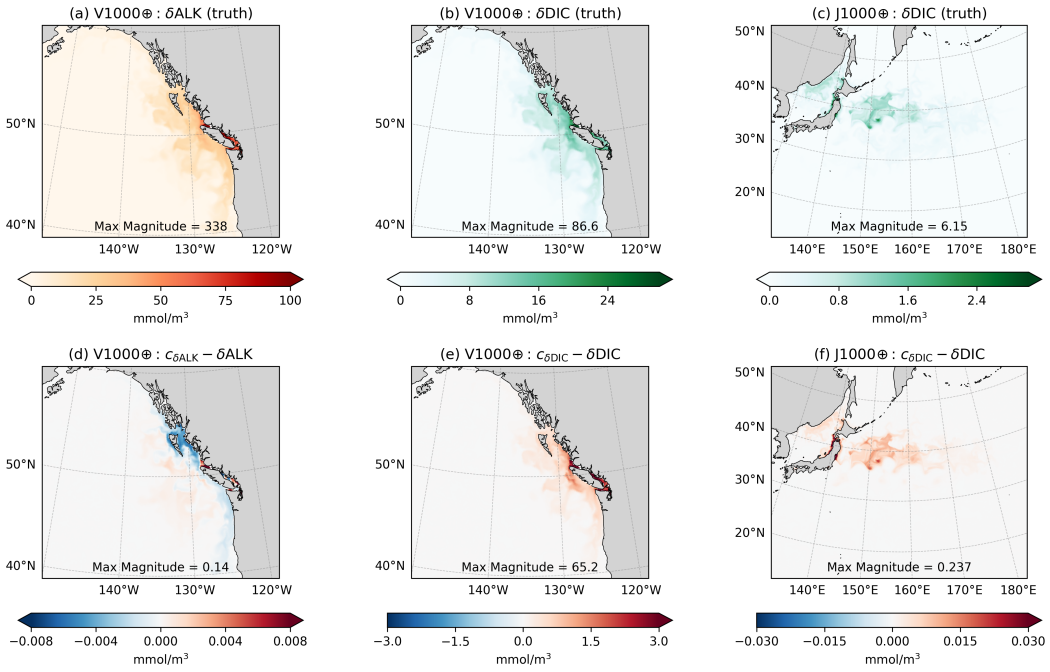


Figure 8. Surface response fields and CDR tracer errors for the largest-amplitude OAE experiments at Vancouver Island (a,b,d,e) and Japan (c,f), on July 1, 2000, six months after the intervention start. Top row: true δALK (a) and δDIC (b) for V1000$\oplus$, and true δDIC (c) for J1000$\oplus$. Bottom row: CDR tracer error ($c_{\delta\mathrm{ALK}} - \delta$ALK and $c_{\delta\mathrm{DIC}} - \delta$DIC) for the corresponding fields in the top row. Absolute rather than relative errors are shown, as near-zero values of δALK and δDIC in the truth fields would render the relative error meaningless over much of the domain. All colorbars are saturated; the domain-wide absolute maximum is annotated in each panel.

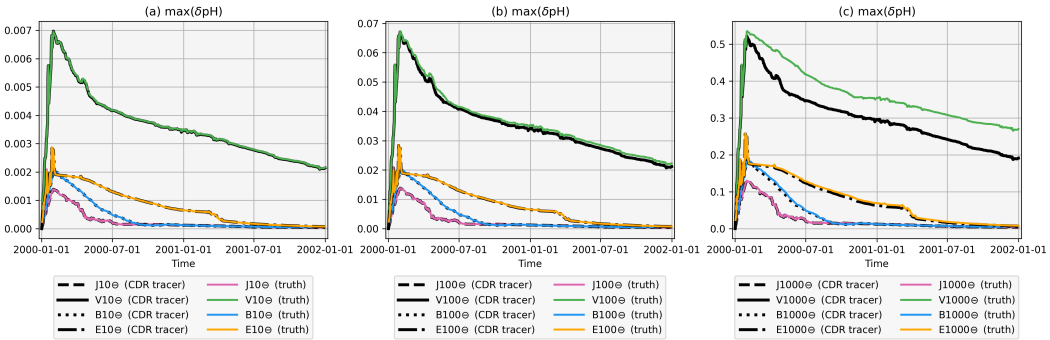


Figure 9. Maximum surface pH change max(δpH), where the maximum is taken over all surface grid points, for the low- (a), mid- (b), and high-amplitude (c) DOR experiments. Colored lines show the truth, with colors indicating release location (pink: Japan, green: Vancouver Island, blue: Baja California, orange: Ecuador). CDR tracer estimates are shown as thick black lines (dashed: Japan, solid: Vancouver Island, dotted: Baja California, dash-dot: Ecuador).

As a final check, we verify the induced change in surface pH, the chemical impact most relevant to environmental thresholds for ocean-based CDR. Unlike the quantities above, pH is not carried by the CDR tracer framework but reconstructed offline: we compute $\mathrm{pH}^{\mathrm{pert}}$ from $[\mathrm{DIC}]^{\mathrm{ctrl}}+c_{\delta\mathrm{DIC}}$ and $[\mathrm{ALK}]^{\mathrm{ctrl}}+c_{\delta\mathrm{ALK}}$ using `pyCO2SYS` (Section 2.2). The truth $\mathrm{pH}^{\mathrm{pert}}$ is diagnosed online in the MARBL truth experiments. In both cases the anomaly is $\delta\mathrm{pH} = \mathrm{pH}^{\mathrm{pert}} - \mathrm{pH}^{\mathrm{ctrl}}$, where $\mathrm{pH}^{\mathrm{ctrl}}$ is diagnosed online in the control run. Because the CDR tracer estimate of pH is computed from the reconstructed carbonate state via `pyCO2SYS`, the CDR tracer estimate of δpH does not scale linearly with intervention amplitude (Figure 9): the nonlinearity is introduced by the carbonate chemistry solver itself, operating on a state shifted by $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$.

We assess the accuracy of the CDR tracer method through the domain-wide maximum of δpH over all surface grid points, which captures the peak local impact (Figure 9). For the low- and mid-amplitude experiments, the CDR tracer estimate reproduces this quantity accurately across all locations, with a slight underestimation at Vancouver Island of up to 0.002 (while the signal is in the 0.05 range), occurring around four to five months after the intervention start. For the high-amplitude experiments, the CDR tracer remains accurate at Japan, Baja California, and Ecuador, but shows a significant error at Vancouver Island (Figure 9(c)). The domain-wide maximum δpH at Vancouver Island is about 0.52 and located in the narrow coastal passages near Port Hardy, Victoria, and Vancouver, where extreme signal accumulation was identified in Section 3. The CDR tracer reproduces the initial δpH maximum of 0.52 correctly, but then overestimates the rate of decay, resulting in an underestimation of the domain-wide maximum δpH by up to 0.1 throughout the two-year period.

The high-amplitude experiments produce domain-wide maximum δpH values of 0.12–0.52 across the four locations (Figure 9(c)). As anticipated from the peak fluxes of the interventions (Section 2.4.1), these values exceed the threshold of 0.1 commonly adopted in OAE modeling studies (Zhou et al., 2024; He & Tyka, 2023), which is comparable to the decline in surface-ocean pH since preindustrial times (Doney et al., 2009). The CDR tracer error in δpH at Vancouver Island thus arises only for an intervention intensity beyond what would likely be permitted. Nevertheless, it reflects a genuine limitation of the method in regions of extreme signal accumulation, such as semi-enclosed coastal passages, where local perturbations can remain large even for more moderate interventions.

4.2 CESM

We now turn to the verification of the CDR tracer method in CESM/MARBL, using truth experiments from Zhou et al. (2024, 2026) at 690 intervention locations covering the global ocean, each run for 15 years. This gives 1,380 truth experiments in total (690 DOR and 690 OAE).

We begin with a selection of 12 polygons (Figure 10(a)), drawn from the four broad ocean basins defined by Zhou et al. (2024): the North Atlantic (polygons 000–002), the North Pacific (150–152), the Southern Hemisphere ocean between 0 and 60°S (350–352), and the Southern Ocean south of 60°S (651–653). Figure 10(b,c) show the DOR and OAE truth efficiency curves alongside their CDR tracer estimates over 15 years. The CDR tracer estimates agree closely with the true curves across all 12 polygons and both intervention types. Figure 10(d,e) shows the relative efficiency error. Consistent with the ROMS results and the box model prediction, the CDR tracer systematically overestimates efficiency in most cases. However, the Labrador Sea polygon (000) shows a negative error for DOR throughout, and several other polygons transition from positive to negative error at longer timescales for both OAE and DOR. The Weddell Sea polygon (653) stands out as a persistent positive outlier, with a relative error that remains above 2% throughout the 15-year period. We return to the question of why the CDR tracer underestimates or persistently overestimates efficiency in these cases when we examine the

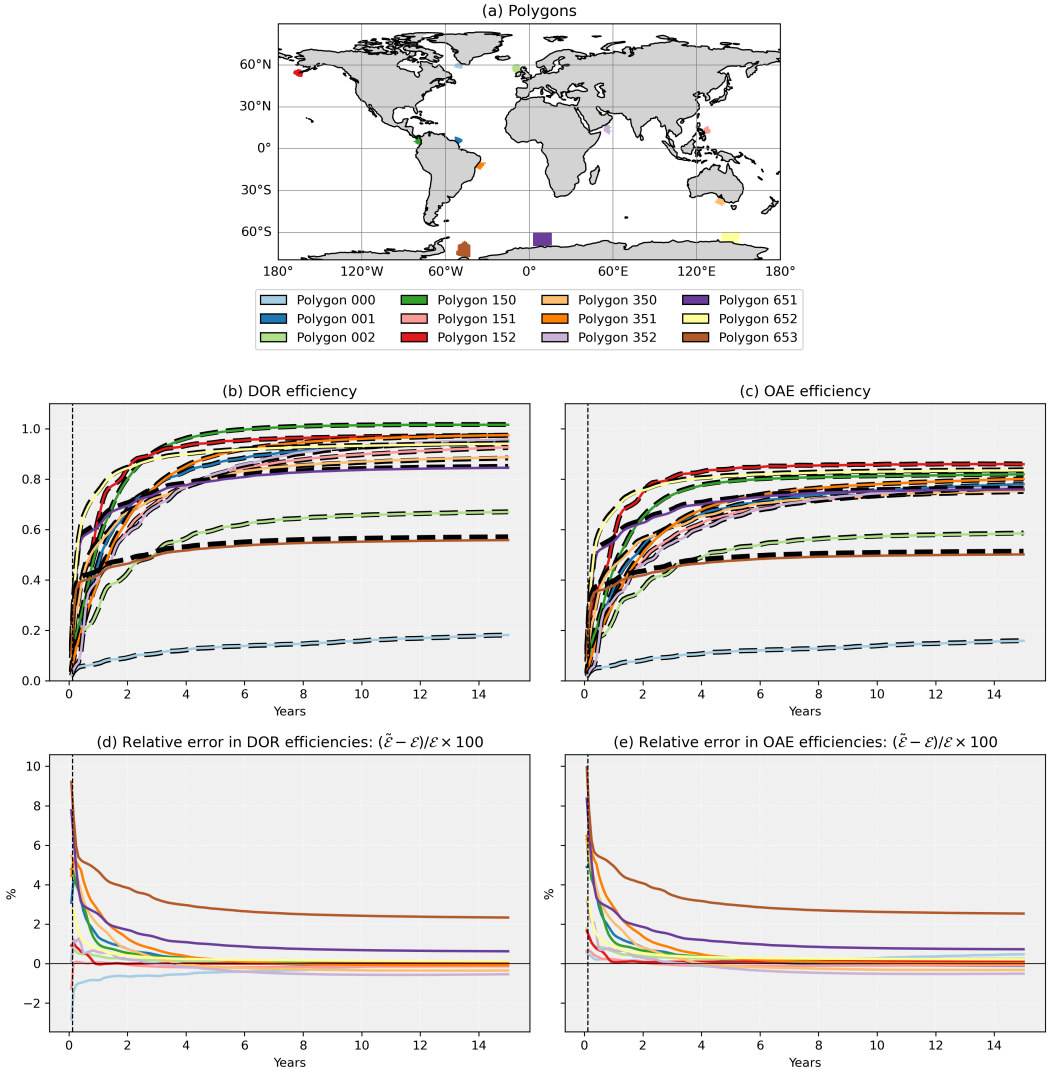


Figure 10. CDR tracer verification for the CESM/MARBL experiments. (a) 12 selected intervention polygons; colors match the efficiency curves in panels (b–e). (b,c) DOR and OAE efficiency curves over 15 years: colored solid lines show the truth $\mathcal{E}^{\mathrm{DOR}}$ (b) and $\mathcal{E}^{\mathrm{OAE}}$ (c); thick black dashed lines show the CDR tracer estimates $\tilde{\mathcal{E}}^{\mathrm{DOR}}$ (b) and $\tilde{\mathcal{E}}^{\mathrm{OAE}}$ (c) (Eqs. (21) and (20)). (d,e) Relative error $(\tilde{\mathcal{E}}-\mathcal{E})/\mathcal{E}\times 100$ (%) for DOR (d) and OAE (e).

full 690-polygon ensemble below. Subtle differences between the DOR and OAE error curves are also apparent. For example, the Labrador Sea polygon (000) shows a negative DOR error but a positive OAE error. The origin of these DOR/OAE differences is not further investigated here.

Figure 11(a) shows the full ensemble of relative OAE error curves across all 690 polygons; Figure S27 shows the DOR equivalent. The general behavior mirrors that of Figure 10: the relative error is largest at the start and decays rapidly, from an ensemble mean of approximately 3.5% at the intervention start to near zero by year 4, where it remains, with a standard deviation of approximately 1%. The absolute error, in contrast, peaks a few months after the intervention (Figures S28, S29).

The 12 polygons with the largest divergence are highlighted in color, and their locations are shown in panel (c). Contrary to what the analysis in Sections 3 and 4.1 would suggest, these are not the polygons with the highest CDR efficiency (panel (b)), where the largest linearization errors would be expected. Instead, they fall into three groups. First, in the Arabian Sea, the CDR tracer begins to underestimate CDR efficiency after a few months; polygon 437 shows the largest error, reaching approximately -9% around year 9 and remaining near -8% at year 15. Second, in the subtropical North Atlantic, the CDR tracer likewise begins to underestimate CDR efficiency after a few months, and the underestimation persists throughout the simulation. Third, on the Antarctic shelf in the Weddell Sea, the CDR tracer persistently overestimates CDR efficiency, as noted above for polygon 653. We investigate these three groups below.

In all three cases, the dominant source of error appears to be differences in the physical model state between the truth and CDR tracer runs, in particular in mixed layer depth, rather than the nonlinearity of the carbonate chemistry, as we argue below. The truth and CDR tracer runs share the same configuration, forcing, initial conditions, and intervention flux, but the CESM/MARBL truth experiments were performed independently and at least two years before the CDR tracer runs. The physical trajectories of the two runs are identical only in an identical software environment (Section 2.2), which was not the case here, contributing an additional source of error.

Because the physical state was not saved in the truth runs, we use the δALK profiles as an indirect diagnostic of physical differences between the two runs: alkalinity is conserved apart from two biogeochemical feedbacks that we discuss below, so the δALK profiles of the two runs should agree if their physical trajectories are identical. Figures S32–S35 show that they do not. Where the CDR tracer method underestimates efficiency (downward-pointing triangles in Figure 11(c), e.g., in the Arabian Sea, the subtropical North Atlantic, and polygon 164 off Baja California), the δALK signal in the CDR tracer run is depleted at the surface and enhanced at depth relative to the truth (Figures S32–S34). This is consistent with a deeper mixed layer, which removes the signal from the surface faster and reduces air–sea gas exchange. In the Weddell Sea, where the CDR tracer method overestimates efficiency, the opposite pattern indicates a shallower mixed layer in the CDR tracer run (Figure S35). Because linearization alone is expected to produce a systematic overestimation (Section 3), downward-pointing triangles in Figure 11(c) generally mark polygons where such physical differences dominate the error. Upward-pointing triangles, in contrast, are consistent with either error source, so the two error sources cannot be distinguished from the sign of the error alone.

In the Weddell Sea, differences in sea ice may contribute as well. These polygons are partially ice-covered during the intervention (Figure S37), and because the gas transfer velocity scales with $(1 - f_{\mathrm{ice}})$ (Eq. (16)), small differences in sea-ice concentration translate directly into differences in CO_2 uptake. Their effect is amplified because the extremely shallow summer mixed layers on the Antarctic shelf confine the intervention signal to a thin surface layer (Figure S36), and because uptake occurs almost entirely within the first two months (Figure S38; Figures 10(b,c) and 11(b)). Any discrepancy

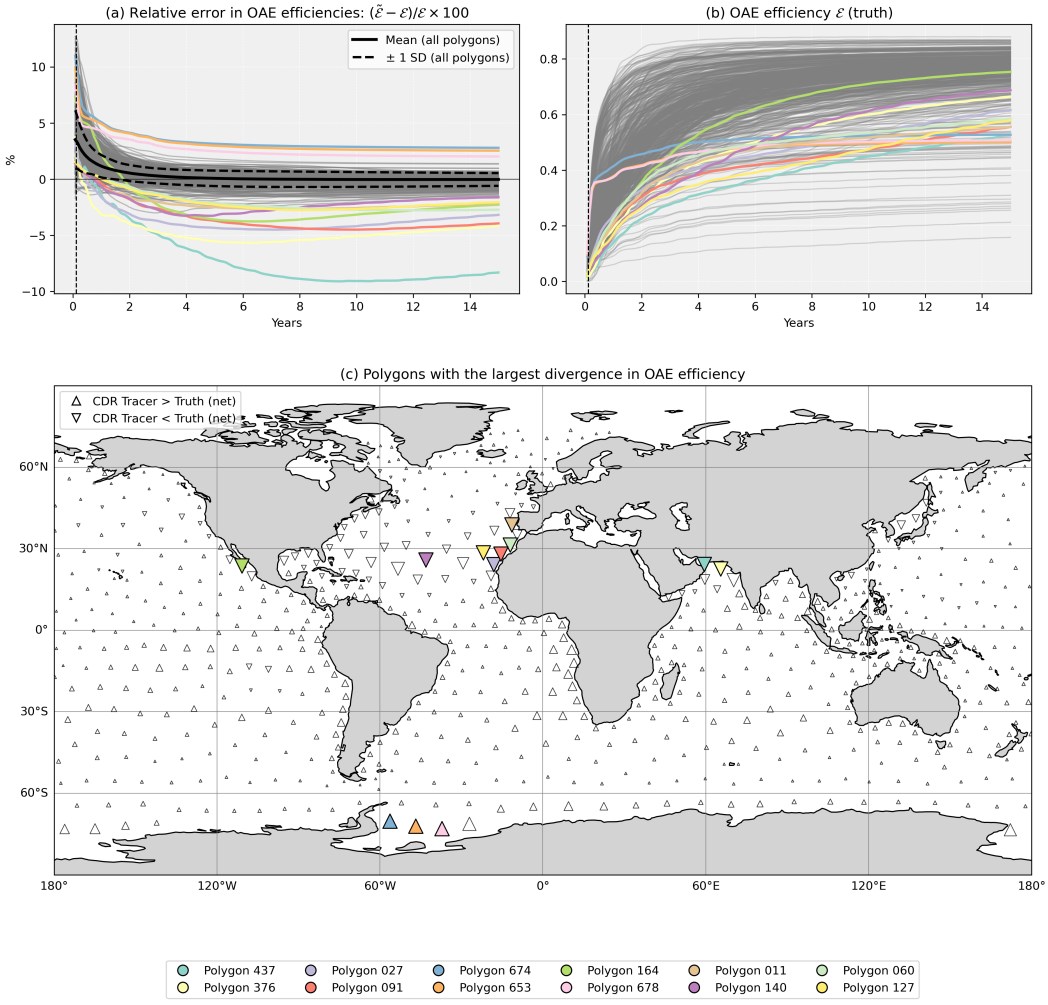


Figure 11. CDR tracer error in OAE efficiency across all 690 CESM/MARBL polygons. (a) Relative error $(\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}})/\mathcal{E}^{\mathrm{OAE}} \times 100$ (%) for all 690 polygons (gray). The 12 polygons with the largest divergence, defined as $\sum_t |R(t)|$ with $R(t) = (\tilde{\mathcal{E}}^{\mathrm{OAE}}(t) - \mathcal{E}^{\mathrm{OAE}}(t))/\mathcal{E}^{\mathrm{OAE}}(t)$, are highlighted in color, with polygon numbers given in the legend below panel (c). The solid and dashed black lines show the mean and ±1 standard deviation across all 690 polygons. (b) True OAE efficiency $\mathcal{E}$ for all 690 polygons (gray), with the same 12 polygons highlighted in color. (c) Location of all 690 polygons. The 12 colored triangles correspond to the highlighted polygons in (a,b). Triangle size is proportional to the divergence $\sum_t |R(t)|$ for all 690 polygons. Upward-pointing triangles indicate polygons where the CDR tracer method net overestimates efficiency ($\sum_t R(t) > 0$) and downward-pointing triangles where it net underestimates ($\sum_t R(t) < 0$). Figure S30 shows the same map using the sign of $\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}}$ at the final time step as the criterion; the two maps are very similar.

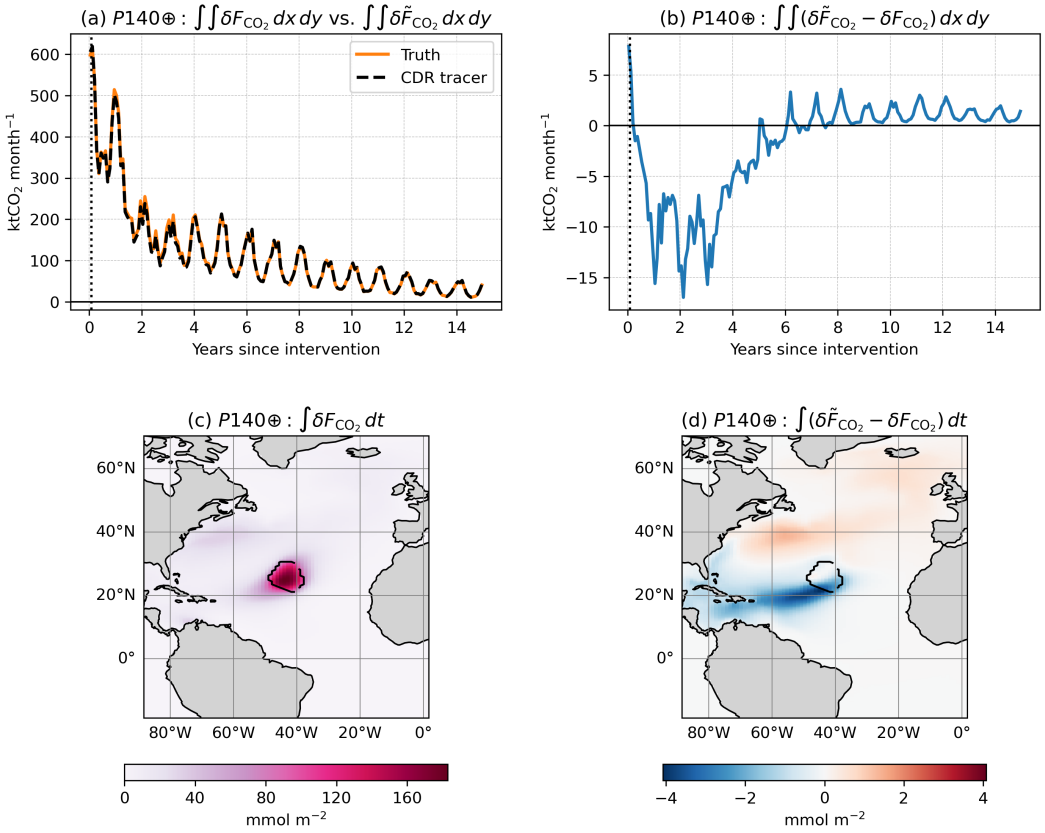


Figure 12. Air-sea CO_2 flux anomaly for the OAE experiment at polygon 140 (P140⊕): comparison between the truth and the CDR tracer estimate. (a) Spatially integrated flux anomaly $\iint \delta F_{CO_2}\,dx\,dy$ (truth, orange) and $\iint \delta \tilde{F}_{CO_2}\,dx\,dy$ (CDR tracer, black dashed) as a function of time, converted to ktCO$_2$ month^{-1}. The dotted vertical line marks the end of the intervention. (b) Difference between the two curves in (a), $\iint (\delta \tilde{F}_{CO_2} - \delta F_{CO_2})\,dx\,dy$. (c) Time-integrated true flux anomaly $\int \delta F_{CO_2}\,dt$, showing where CO_2 uptake occurs spatially over the 15-year period. The black contour marks the polygon boundary. (d) Difference in time-integrated flux anomaly, $\int (\delta \tilde{F}_{CO_2} - \delta F_{CO_2})\,dt$. Negative values (blue) indicate where the CDR tracer underestimates CO_2 uptake relative to the truth.

during this brief window therefore has an irreversible effect on the total uptake, explaining the persistent error throughout the simulation (Figures 10(d,e) and 11(a)).

Finally, the two biogeochemical feedbacks that the CDR tracer method does not represent are unlikely to explain the observed discrepancies in CDR efficiency and δALK. pH-sensitive NH_3 outgassing reduces surface alkalinity in the truth runs and would therefore produce a surface δALK deficit in the truth relative to the CDR tracer run, which is the opposite of what we observe for the underestimating polygons. For the overestimating polygons, the δALK profiles show a vertical redistribution of the signal rather than a surface-confined depletion, which is more consistent with physical differences than with NH_3 outgassing. Sedimentary denitrification could affect air–sea exchange only if the signal reached the seafloor and returned to the surface, which is unlikely within the 15-year simulations.

Table 3. Computational cost of truth experiments versus CDR tracer experiments on the Perlmutter supercomputer at NERSC. Core-hours are reported per simulation year (sim year). Speedup is the ratio of total core-hours for the truth experiments to total core-hours for the CDR tracer experiment.

Configuration	**Core-hours per sim year and run**	**# Runs**	**Total core-hours per sim year**	**Speedup**
ROMS/MARBL truth (32 BGC tracers)	27,684	12	332,210	
ROMS CDR tracer, OAE (2×12 CDR tracers)	18,699	1	18,699	17.8×
CESM/MARBL truth (32 BGC tracers)	446	690	307,940	
CESM CDR tracer, DOR (1×690 CDR tracers)	10,995	1	10,995	28.0×
CESM CDR tracer, OAE (2×690 CDR tracers)	22,890	1	22,890	13.5×

So far in this section, CDR efficiency has been the primary metric of comparison between truth and CDR tracer runs. CDR efficiency is a quantity integrated over both space and time, and a good match could in principle arise from cancelling errors. As a final check, we therefore examine whether the CO_2 flux anomaly predicted by the CDR tracer method occurs in the right locations and at the right times.

Figure 12 shows this comparison for polygon 140 in the subtropical North Atlantic, one of the polygons with the largest divergence in Figure 11 (see Figures S38–40 for the same analysis for selected other polygons with large errors). Panel (a) shows that the CDR tracer method closely reproduces the spatially integrated instantaneous flux anomaly as a function of time. Panel (b) shows the error, which initially follows the two-phase structure predicted by the box model (Figure 3(d)): the CDR tracer initially overestimates the flux anomaly, transitions to underestimation at a crossover time t^* (here occurring one to two months after the end of the intervention), and then underestimates the flux as the alkalinity plume travels anti-cyclonically around the subtropical gyre (panel (d)). After approximately six years, the CDR tracer method switches back to overestimating the flux (panel (b)), with the overestimation concentrated along 40°N and further north into the North Atlantic subpolar gyre (panel (d)). This late-time overestimation is not predicted by the box model and, consistent with the analysis above, we attribute it to subtle differences in mixed layer depth between the truth and CDR tracer runs.

5 Computational savings

The CDR tracer method yields substantial computational savings relative to running individual truth experiments with an explicit BGC model for each intervention. These savings stem from two sources. First, the physical ocean state is computed only once: all interventions are carried simultaneously as passive tracers within a single physical integration, rather than in separate simulations. Second, the full BGC model (32 MARBL tracers plus internal source-sink computations) is replaced by one or two passive tracers per intervention, with only the linearized air–sea gas exchange retained.

Table 3 summarizes the computational cost on the Perlmutter supercomputer at the National Energy Research Scientific Computing Center (NERSC). We report core-hours per simulation year, where a core-hour is defined as one physical CPU core occupied for one hour of wall-clock time. On Perlmutter, each CPU compute node provides 128 physical cores (two AMD EPYC 7763 sockets of 64 cores each).

For ROMS, running 12 OAE truth experiments with MARBL required 332,210 core-hours per simulation year in total, while a single CDR tracer integration carrying 24 CDR tracers (12 pairs of $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$) required only 18,699 core-hours per simulation year, a speedup of 17.8×. Because the CDR tracer equations are linear, a single tracer pair per location would in principle have sufficed to cover all three amplitudes (Section 4.1); we carried one pair per truth experiment to verify the linearity of the implementation (Figure S24). We do not report timings for the ROMS DOR CDR tracer experiments: our implementation carried an additional 12 conservative tracers to track the cumulative DIC sink applied by the intervention, which doubled the tracer count and resulted in a compute time similar to that of the ROMS OAE CDR tracer integration. Future implementations should instead use a dedicated diagnostic for this purpose, as done in the CESM implementation.

For CESM, 690 truth experiments (OAE or DOR) with MARBL required 307,940 core-hours per simulation year. A single CDR tracer integration for DOR, carrying 690 $c_{\delta\mathrm{DIC}}$ tracers, required 10,995 core-hours per simulation year, a speedup of 28.0×. The corresponding OAE integration, carrying 1,380 tracers (690 pairs of $c_{\delta\mathrm{DIC}}$ and $c_{\delta\mathrm{ALK}}$), required 22,890 core-hours per simulation year, a speedup of 13.5×. We performed separate CDR tracer integrations for OAE and DOR for bookkeeping convenience; all tracers could in principle have been carried in a single integration. The number of passive tracers that can be carried simultaneously is ultimately limited by available memory, since each tracer requires storing its full three-dimensional concentration field, setting a practical upper bound on the number of interventions per integration. The experiments presented here ran on 7 compute nodes and memory was not a limiting factor.

Consistent with the fact that tracer advection is generally one of the dominant costs of ocean model integrations, the reduction in tracer count appears to be the primary source of savings. The ratio of BGC tracers to CDR tracers per intervention provides a rough guide: the CESM DOR speedup of 28.0× is broadly consistent with a ratio of 32:1, and the CESM and ROMS OAE speedups of 13.5× and 17.8× with a ratio of 32:2 (i.e., a factor of 16). This correspondence should not be interpreted as a linear relationship between tracer count and cost. Adding tracers increases not only computation but also memory access and MPI communication, whose costs depend on how well a given configuration is scaled and load-balanced. The speedup therefore also depends on the configuration, including the relative cost of physics, BGC, and tracer transport. Given that ROMS and CESM differ in model, resolution, domain, number of tracers, and computational setup (e.g., number of nodes), it is not surprising that their OAE speedups differ; both nevertheless remain close to the tracer-count ratio of 16. A detailed, configuration-specific scaling analysis would be required to determine at what point adding CDR tracers to a single integration becomes less efficient than distributing them across several integrations.

6 Robustness to the Choice of β and η

Figure 13 shows OAE efficiency curves and their relative errors for the CDR tracer experiments using degraded or externally derived carbonate sensitivities (Table 2, Section 2.5). The DOR results are very similar to the OAE results and are shown in Figure S41; absolute error curves for both intervention types are shown in Figures S42, S43.

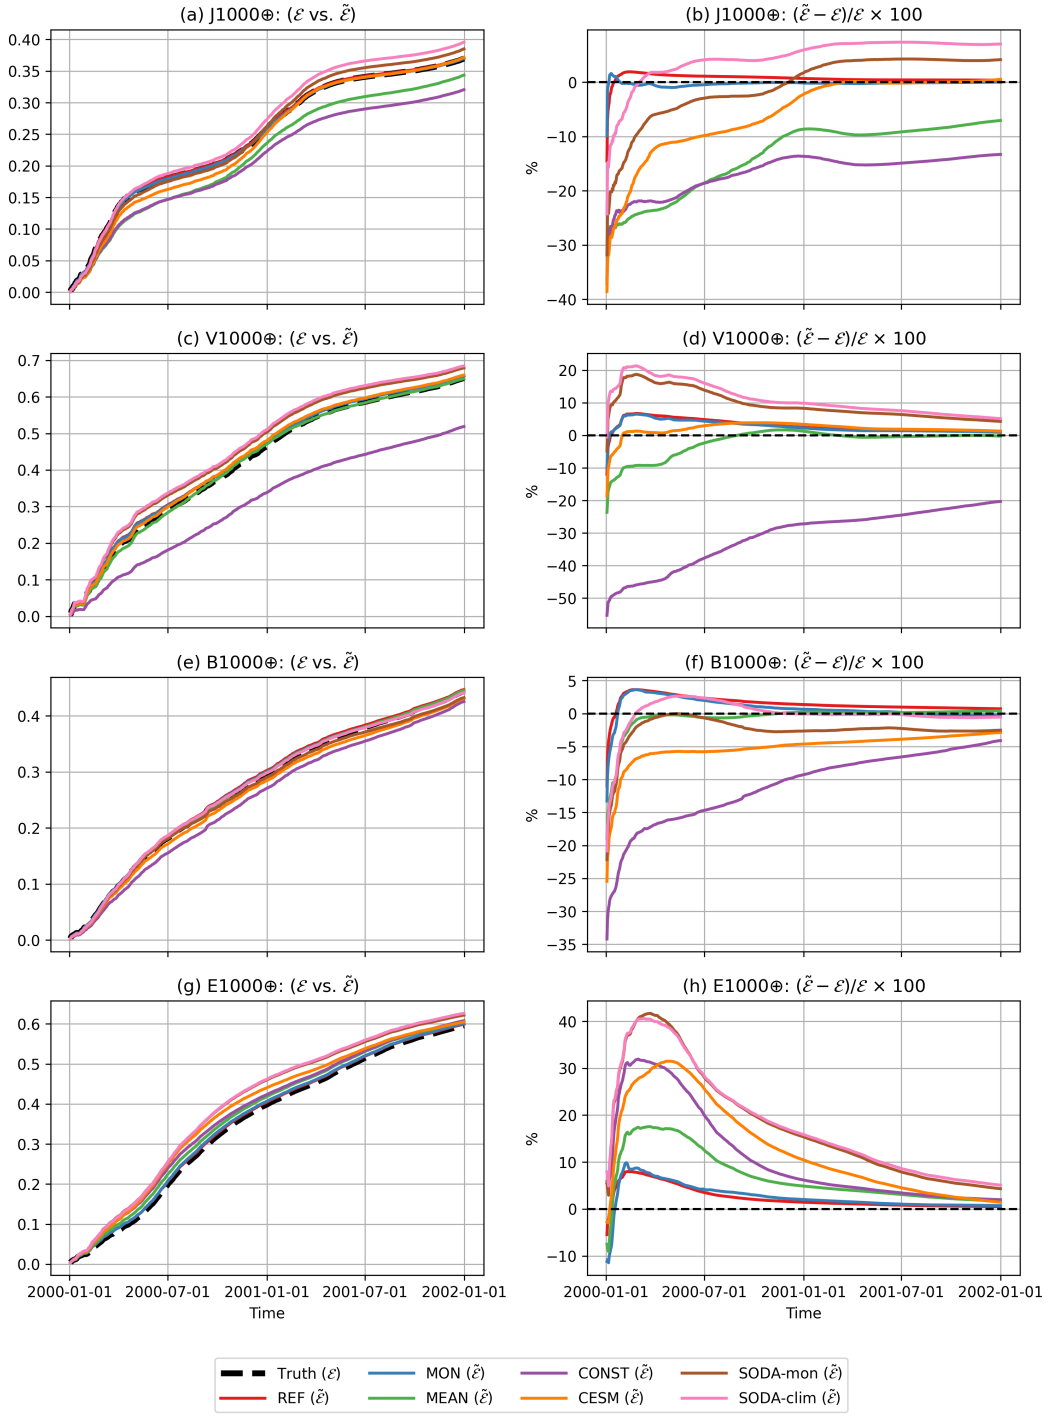


Figure 13. OAE efficiency curves (Eq. (18)) for the largest-amplitude experiments at each of the four release locations: Japan (a,b), Vancouver Island (c,d), Baja California (e,f), and Ecuador (g,h). Left column (a,c,e,g): truth $\mathcal{E}^{\mathrm{OAE}}$ (dashed black) and CDR tracer estimates $\tilde{\mathcal{E}}^{\mathrm{OAE}}$ (solid colored lines) for the seven carbonate sensitivity configurations listed in Table 2. Right column (b,d,f,h): corresponding relative error $(\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}})/\mathcal{E}^{\mathrm{OAE}} \times 100$, expressed as a percentage. The dashed black line marks zero error.

The monthly-mean configuration (`MON`) performs as well as the daily reference (`REF`) across all four locations, with the two relative error curves almost indistinguishable. The similarity reflects the fact that monthly averaging preserves almost the entire spatial variability of the daily carbonate sensitivity fields (Figures S44–S54), with only modest degradation near Ecuador where sub-monthly variability in β and η driven by tropical instability waves is smoothed out by monthly averaging (Figures S53, S54). This manifests in the efficiency curves: Ecuador is the only location where the daily configuration outperforms the monthly one, reducing the peak relative error from approximately 10% for `MON` to 8% for `REF` (Figure 13(h)). At Japan, Vancouver Island, and Baja California, `MON` marginally outperforms `REF` (Figure 13(b,d,f)), likely due to a cancellation of errors. For example, near Baja California `MON` shows a slight positive bias in β (Figure S52) and a slight negative bias in η (not shown) relative to `REF`. Both act to reduce the CDR tracer's systematic overestimation of efficiency and partially offset it.

Further degrading to the time-mean (`MEAN`) leads to a serious deterioration at all locations except Baja California, where a similar cancellation of errors keeps the error near zero after approximately four months after intervention start (Figure 13(f)). The largest degradation occurs at Japan and Ecuador (Figure 13(b,h)): time-averaging wipes out the mesoscale eddy activity along the Kuroshio imprinted on β and η (Figure S47, S48), and eliminates the variability of the equatorial current system driven by tropical instability waves and ENSO, neither of which is captured by a time-mean (Figures S53, S54).

The single domain-wide constants (`CONST`, $\beta = 18.7$, $\eta = 0.82$) were intended as a worst-case scenario. Their performance depends on how close the constants are to the local values at each intervention site. Near Ecuador, the constants happen to be close to the local ranges, so `CONST` outperforms the external products `CESM` and `SODA` at that location (Figure 13(h)). Near Vancouver Island, where the local β and η typically lie in the ranges 8–14 and 0.84–0.89, respectively, the prescribed constants fall far outside the local ranges and the errors are large (Figure 13(d)).

Using β and η from CESM (`CESM`) or OceanSODA-ETHZ (`SODA-mon`, `SODA-clim`) leads to substantially larger errors than the self-consistent carbonate sensitivities derived from the ROMS/MARBL control run. For example, six months after the intervention, `REF` and `MON` give relative errors of 0-5% across all locations, while `CESM` gives errors of 5-24% and `SODA-clim` errors of 2-26%, with the largest errors at Ecuador and Vancouver Island for both external products. The `SODA` products outperform `CESM` at some locations but not others. `SODA-mon` and `SODA-clim` give similar results except at Japan, where strong interannual variability in the Kuroshio drives significant differences between the two (Figure 13(b)).

To summarize, of the degradations tested, only temporal averaging to monthly resolution (`MON`) leaves the CDR efficiency essentially unchanged relative to the daily reference, because it preserves the spatial structure of β and η. Text S2 discusses the approximations involved in such temporal averaging and gives practical recommendations for preparing β and η at daily and monthly frequency. Further temporal or spatial degradation (`MEAN`, `CONST`) substantially increases the error, as does using β and η from an independent model or observation-based product, except in isolated cases that appear to reflect cancellations of errors rather than genuine skill. We note that all results in this section use ROMS/MARBL as the "truth". In real MRV applications, the choice of physical model and BGC model is itself a source of uncertainty, and the spread across configurations in Figure 13 may provide a practical way to map out some of this uncertainty. This idea is further discussed in Section 7.

7 Discussion

We discuss the efficiency and accuracy of the CDR tracer method, the nature of its linearization error, its generality and limitations, and its relation to other approaches, before closing with an outlook on possible extensions and applications.

7.1 Efficiency and accuracy

The CDR tracer method introduced here offers a computationally efficient alternative to explicit coupled physics–biogeochemistry simulations for ocean-based CDR modeling. By replacing the full BGC model with one (DOR) or two (OAE) passive tracers per intervention and a linearized air–sea gas exchange term, the method achieves speedups of 13.5–28× over the equivalent ensemble of MARBL-coupled truth experiments (Table 3). The dominant source of savings is the reduction in tracer count: the CESM DOR speedup of 28× closely matches the 32:1 ratio of MARBL tracers to CDR tracers, and the OAE speedup of 13.5× the 32:2 ratio. Implementation requires only a source term for the intervention and a linearized gas exchange term (Eqs. (12)–(13)), which slot into the existing passive-tracer machinery of any ocean model. The additional code is correspondingly minimal for ROMS (157 lines added, 29 removed), though more substantial for CESM (1,695 added, 66 removed) due to POP2's modular architecture, which required a dedicated passive-tracer package with its own initialization, restart handling, and namelist interface. Both implementations are released as open source (Appendix B). The ROMS implementation is further distributed as CDR LiTE (CDR Linear Tracer Estimation) within [C]Worthy's C-Star platform (C-Star contributors, 2025), which additionally provides the workflow for diagnosing β and η online from a control run and prescribing them as input fields in the CDR tracer run. Practical recommendations for preparing β and η are given in Text S2.

Beyond being computationally efficient and straightforward to implement, the CDR tracer method is accurate: the linearization of the carbonate chemistry introduces only small errors in CDR efficiency, provided that the carbonate sensitivities β and η are prescribed at daily or monthly resolution (Section 6). We consider the relative error *after* the first few days of the intervention, during which it is artificially inflated by division by near-zero efficiency values. In the ROMS experiments, the relative error in OAE efficiency then remains below 1% for the mid-amplitude experiments, which correspond to an ambitious but plausible regional deployment (Section 2.4.1). Even for the deliberately extreme largest-amplitude experiments, it peaks at only 1.9–8.0% one to four weeks after the end of the intervention and decays to below 1% at all four locations by the end of the two-year integrations. In CESM, where the intervention flux of Zhou et al. (2024) was chosen to keep the surface pH increase generally below 0.1, the relative error averaged over all 690 polygons is 3.5±2% at the start of the intervention and decays to 0± 1% by year 4. Only two polygons, both in the Arabian Sea, exceed 5% in magnitude at later times. These late-time errors likely reflect differences in the physical model state between the truth and CDR tracer runs, in particular in mixed layer depth, which are unrelated to the CDR tracer method itself (Section 4.2). More generally, because the CESM truth and CDR tracer runs do not share identical physical trajectories, the CESM errors combine the linearization error with the effect of these physical differences, which we cannot separate.

The errors reported above are small compared with other sources of uncertainty in modeled CDR efficiency. Within the same CESM/MARBL configuration, Yankovsky et al. (2025) find that OAE efficiency varies by up to 50% between seasons, and that interannual variability across 16 January releases produces standard deviations in efficiency of 0.01–0.06. In contrast, the absolute error of the CDR tracer method in CESM peaks at 0.005 ± 0.004 (mean ±1 standard deviation across all 690 polygons) and decays to 0.000±0.004 thereafter (Figure S28). The few polygons with larger errors, of up to 0.04,

likely owe these to differences in mixed layer depth rather than to the linearization (Section 4.2). Notably, Yankovsky et al. (2025) consider this level of error acceptable for MRV: their impulse response approach, which neglects interannual variability, reconstructs multi-year deployments to within 7%, and to within 1% in most regions. The CDR tracer method achieves smaller errors without assuming time invariance, since each intervention is simulated explicitly under the actual time-varying circulation. Differences between models are larger still: comparing CESM/MARBL with ECCO-Darwin, Tyka et al. (2026) find that OAE efficiency typically differs by 10–20% during the first seven years after release, and by up to 50% in extreme cases. They attribute these differences mainly to vertical transport and wind speeds, with carbonate chemistry playing a smaller role. The uncertainty in modeled CDR efficiency thus appears to be dominated by ocean physics rather than by the carbonate chemistry that the CDR tracer method linearizes.

7.2 Nature of the linearization error

Although small, the linearization error of the CDR tracer method is not random: it has a systematic sign and a magnitude that depends predictably on the intervention and its setting. The CDR tracer method overestimates CDR efficiency, as predicted by the box model (Section 3) and confirmed by the ROMS experiments (Section 4.1). This bias follows directly from the linearization: because the CDR tracer uses the control values of β and η, it misses the increase in β and decrease in η that the intervention itself induces, and therefore overestimates both the rate of CO_2 uptake and its long-term plateau. The bias is, however, so small that it is easily masked by differences in the physical model state between the truth and CDR tracer runs, which arise when the two are not performed in identical hardware and software environments (Section 2.2) and are unrelated to the CDR tracer method itself. In CESM, such differences dominate the error at later times and cause the CDR tracer to underestimate efficiency in some regions (Section 4.2). We observed the same effect in ROMS (not shown): when the truth and CDR tracer runs used executables built in different software environments, the resulting differences were comparable to the linearization error of the mid- and small-amplitude experiments.

The magnitude of the linearization error is controlled primarily by vertical mixing. Strong vertical mixing carries the intervention signal away from the surface, which both reduces the CDR efficiency and lowers the surface concentrations of δALK and δDIC. Lower surface concentrations, in turn, cause smaller deviations of β and η from their control values, explaining why less efficient interventions tend to have smaller relative errors (Section 4.1). Horizontal spreading also lowers surface concentrations, but its effect on the error is partly offset because the smaller deviations in β and η then act over a larger area. The error is dominated by the deviation in β, which is about an order of magnitude larger than that in η (Section 3). Since β governs the equilibration rate for both intervention types, DOR is not an easier problem for the CDR tracer method than OAE, consistent with their nearly identical relative error curves. For OAE, the long-term efficiency plateau additionally depends on η, but its smaller deviations make this a minor contribution. To second order, the error for both DOR and OAE may also depend on the background carbonate state: β varies more steeply with ALK and DIC at higher values of β (Figure 3(a)), so a given perturbation would produce a larger deviation in β, and hence a larger error, in regions of high control β. We have not quantified this effect here, and expect it to be secondary to that of vertical mixing.

7.3 Generality and limitations

The CDR tracer method is general: it applies to any model in which air–sea CO_2 exchange follows the bulk form of Eq. (1) with a gas transfer velocity k that does not depend on the carbonate system variables (Section 2.1), and makes no assumptions about model resolution or intervention type. Transport of the CDR tracers is handled entirely by the host model's passive-tracer machinery, so the method inherits whatever physics

the model resolves. We have demonstrated this in two configurations that differ in physical model, resolution, and region: a mesoscale-resolving 12 km regional ROMS setup and a non-eddying 1° global CESM setup. Regarding intervention type, we have tested OAE and DOR, but the formulation equally accommodates mixed interventions such as ERW, which add both alkalinity and DIC, by prescribing both source terms $S_{\delta\mathrm{DIC}}$ and $S_{\delta\mathrm{ALK}}$ (Section 2.2). In principle, any intervention that creates a surface DIC deficit can be represented in this way. For biotic approaches such as ocean fertilization or macroalgae cultivation, however, the induced deficit itself depends on biological processes and the durability of organic carbon export, which the CDR tracer method does not model. More generally, the method tracks the fate of a prescribed perturbation, whereas quantifying that perturbation, for example by accounting for incomplete feedstock dissolution or secondary precipitation, remains a method-specific task outside its scope.

The CDR tracer method also has several limitations. First, although it reproduces CDR efficiency accurately, the linearization can cause substantial local errors where the intervention signal accumulates, such as in the narrow coastal passages near Vancouver Island (Section 4.1). There, the unrepresented local change in β causes the CDR tracer to overestimate δDIC by up to a factor of two and to underestimate the maximum surface pH increase, which peaks at 0.52, by up to 0.1 during its subsequent decay. We found such errors only for the deliberately extreme largest-amplitude experiments, but they can arise in any region with long residence times, such as semi-enclosed basins or shallow coastal passages. Similar errors may occur near concentrated point-source releases, which our experiments, with interventions imposed over large areas, do not represent. In practice, deployments producing such extreme local accumulation would likely not be permitted, since regulatory limits on pH perturbations constrain the rate at which alkalinity can be added at a given site (Bach et al., 2026), although accumulation could still arise from multiple deployments within the same basin. The CDR tracer run itself provides a way to detect such regions: because δALK is approximately conserved, its passive-tracer counterpart $c_{\delta\mathrm{ALK}}$ reproduces it accurately even where $c_{\delta\mathrm{DIC}}$ does not reproduce δDIC (Section 4.1). A simple safeguard is therefore to define a threshold for $c_{\delta\mathrm{ALK}}$. Where this threshold is locally exceeded, the pH reconstruction should not be trusted, while integrated quantities such as CDR efficiency may remain reliable if the affected area is small.

Second, by omitting the biogeochemical model, the CDR tracer method provides no information on biological responses to the intervention, and hence on potential environmental harms or co-benefits beyond changes in the carbonate system, such as pH and saturation state, which can be reconstructed offline. In this respect, however, little is lost relative to the coupled model, since biogeochemical models such as MARBL do not currently represent the organism- and ecosystem-level responses to alkalinity perturbations that environmental assessments require. More relevant for CDR quantification is that the method cannot capture biogeochemical feedbacks on the carbonate system. In our MARBL configuration, these feedbacks, pH-sensitive NH_3 outgassing and sedimentary denitrification, were not a dominant source of error (Section 4). Other biogeochemical models or configurations may include additional feedbacks, for example calcification or secondary $CaCO_3$ precipitation that depend on the saturation state. We expect these to be small at CDR-relevant amplitudes. Near point-source releases, however, perturbations are concentrated in a near field that is not resolved by the models used here, which may enhance these feedbacks. Where such conditions are likely, for example near point-source releases or in particle-rich coastal settings, the error of the CDR tracer method should be reassessed against the corresponding coupled model.

Third, like the truth experiments and many modeling studies of ocean-based CDR (e.g., Zhou et al., 2024; Yankovsky et al., 2025; Tyka et al., 2026), the CDR tracer method assumes a non-responsive atmosphere, in which atmospheric pCO_2 is unaffected by the intervention. It therefore quantifies the additional oceanic CO_2 uptake, but not the subsequent repartitioning of carbon among the ocean, atmosphere, and land reservoirs, which

partially offsets the removal on longer timescales (Oschlies, 2009; Tyka, 2025). Accounting for this repartitioning requires coupling to an interactive atmosphere and land carbon cycle, or a posteriori corrections based on efficiency metrics that account for it (Tyka, 2025). Because the explicit coupled simulations share this assumption, it does not affect the comparison presented here.

7.4 Relation to other approaches

Simplified tracer approaches have been used previously to model ocean-based CDR. Ho et al. (2026) represent the DIC deficit induced by OAE with a single passive tracer that outgasses at a constant gas transfer velocity divided by a fixed value of β. Compared with Eq. (12), this formulation makes three simplifications: k and β are held constant, and the coupling term $\eta\, c_{\delta\mathrm{ALK}}$ is omitted, so that the OAE intervention is effectively treated as a DOR intervention. For their application, a six-day simulation within a single estuary, constant values of k and β are a reasonable approximation. Over larger domains and longer timescales, however, all three simplifications matter. The gas transfer velocity k varies strongly with wind speed, sea-ice cover, and temperature, and these variations shape where and when the induced CO_2 uptake occurs. β likewise varies substantially in space and time, and prescribing a single constant value can lead to large errors in CDR efficiency (Section 6). Finally, the online coupling to $c_{\delta\mathrm{ALK}}$ through a spatially and temporally varying η is essential for reproducing OAE efficiency curves. Offline alternatives, such as scaling DOR efficiency curves by area-weighted values of η, do not fully recover the OAE response (Zhou et al., 2026), whereas the coupled formulation used here does. The linearized equilibration underlying the CDR tracer method has also been used diagnostically. Tyka et al. (2026) express the rate of change of the DIC deficit, $\eta\,\delta\mathrm{ALK} - \delta\mathrm{DIC}$, through the same factor k/β as in Eq. (12), and reconstruct the equilibration rate offline from coupled model output to attribute inter-model differences to individual parameters. The CDR tracer method applies this relation online as a prognostic equation, replacing the coupled biogeochemical model rather than diagnosing it.

A different route to computational savings has been suggested: retaining the carbonate chemistry but switching off the biological model. Tyka et al. (2026) show that switching off 29 of the 31 biogeochemical tracers in ECCO-Darwin has a negligible effect on OAE efficiency. Since both approaches reduce the number of advected tracers to about two per intervention, their computational savings may be broadly comparable. The CDR tracer method nevertheless offers three practical advantages. First, it carries all interventions within a single integration, whereas the alternative requires a separate integration per intervention unless multiple copies of the carbonate system are implemented (MARBL, for instance, provides two copies). Second, it requires no online carbonate chemistry solver, which reduces both computational cost and implementation effort. Third, it avoids drift in the background carbonate state. Both approaches require a spin-up with full biogeochemistry to establish a realistic background carbonate state, in particular the vertical gradients of ALK and DIC maintained by the biological pump. When biology is switched off, these gradients are no longer maintained and the background state drifts. This drift remained small over the 15-year simulations of Tyka et al. (2026), in which all interventions started shortly after biology was switched off. Because the drift accumulates over time, however, it may matter more for interventions that start later in an integration or extend over several years. The CDR tracer method avoids this drift by construction: the background carbonate state, including its biologically maintained structure, enters only through β and η, which are diagnosed from the fully coupled control run and therefore retain the effect of biology throughout the CDR tracer integration.

7.5 Outlook

The linearization of the CDR tracer method could be relaxed in two ways. In the first, the control surface concentrations of DIC, ALK, and nutrients, rather than β and η, would be prescribed as input fields, while temperature and salinity are taken from the physical model, which is identical to the control run. At each time step, the carbonate system would then be solved for both the control and the perturbed surface state, $[\mathrm{DIC}]^{\mathrm{ctrl}} + c_{\delta\mathrm{DIC}}$ and $[\mathrm{ALK}]^{\mathrm{ctrl}} + c_{\delta\mathrm{ALK}}$. The flux anomaly would follow directly from the resulting difference in $p\mathrm{CO}_2$ (Eq. (5)), without linearization. This is essentially a generalization of MARBL's dual carbonate chemistry to an arbitrary number of perturbed states, with the control state prescribed rather than simulated. Since each intervention still requires only one or two advected tracers and a carbonate system solve in the surface layer only, most of the computational savings would be retained. The second, cheaper option is to retain the second-order terms of the Taylor expansion in Eq. (6), with the required derivatives of β and η prescribed as additional input fields. Both extensions would reduce or remove the systematic overestimation of CDR efficiency and the local errors in regions of signal accumulation. The tracer equations would then become nonlinear, however, so each intervention amplitude would have to be simulated separately, and responses to simultaneous interventions could no longer be superposed. The resulting efficiency curves would also violate the linearity on which the impulse response approach of Yankovsky et al. (2025) relies, so they could no longer be combined by convolution. Given the small linearization error at CDR-relevant amplitudes, we consider the linear formulation preferable for most applications, with a nonlinear extension reserved for cases of strong local signal accumulation.

The CDR tracer method still requires a biogeochemical control run to diagnose β and η, and a natural question is whether this run could be avoided altogether. Simply substituting β and η from an independent model or an observation-based product leads to substantially larger deviations from the coupled model (Section 6). Avoiding the control run would therefore require estimating the surface carbonate state of the host model without simulating its biogeochemistry. Surface ALK is strongly constrained by salinity and can be estimated from sea surface salinity and temperature using empirical relationships (Lee et al., 2006; Carter et al., 2021). Surface DIC is more difficult: existing empirical routines estimate it with larger errors when only temperature and salinity are available, and they do not resolve its seasonal variability in the surface ocean, where air-sea exchange decouples DIC from predictors such as temperature and salinity (Carter et al., 2021). Machine-learned emulators trained on coupled simulations may offer a way forward. Such approaches would remove the need for a costly biogeochemical spin-up.

While β and η from external sources are thus no substitute for the self-consistent fields, the spread they produce can be useful in its own right. The carbonate sensitivities β and η compactly summarize the information about the background carbonate state that is relevant for air–sea CO_2 exchange following an intervention, whether this state comes from a biogeochemical model or an observation-based product. Running the CDR tracer method with β and η from several sources, such as the self-consistent control run, an independent model, and an observation-based product (Section 6), therefore isolates the effect of the carbonate system on CDR efficiency while keeping the physical transport fixed. Whereas we treated deviations from the self-consistent configuration as errors in this study, in MRV applications, where no single model can be regarded as the truth, the spread across these configurations (Figure 13) could instead serve as an inexpensive measure of the uncertainty associated with the carbonate system. However, β and η also carry an imprint of the physical model, since they depend on temperature, salinity, and the carbonate state, all of which are shaped by the circulation. In the 12 km ROMS configuration, for example, β and η clearly reflect the mesoscale dynamics (Figures 2, S21–S23, S44-S54). This separation is therefore not perfectly clean, although our results show that monthly averages, which smooth sub-monthly variability, perform as

well as daily fields. How to cleanly separate the uncertainty associated with the carbonate system from that associated with the physical model therefore remains an open question, but the CDR tracer method, by allowing the two to be varied independently, offers a promising step in this direction.

Finally, the computational savings of the CDR tracer method open up applications that are currently prohibitively expensive with explicit coupled physics–biogeochemistry simulations. These include large ensembles spanning deployment locations, seasons, and years, and efficiency atlases at higher resolution than has been feasible so far. Because the largest uncertainties in modeled CDR efficiency arise from ocean physics, in particular vertical mixing (Fennel, 2026; Tyka et al., 2026), the savings are perhaps best invested in exploring this physical uncertainty, by running different ocean models, or the same model with different mixing schemes, each with the CDR tracer method in place of full biogeochemistry. Apart from small linearization errors, the CDR tracer method reproduces explicit simulations at a fraction of the cost, making it a practical substitute for them in the quantification of ocean-based CDR.

Appendix A Computation of Carbonate Sensitivities

The carbonate sensitivities β and η defined in Eq. (7) can be expressed in terms of the speciation of the carbonate system. The sensitivity η is the reciprocal of the isocapnic quotient introduced by Humphreys et al. (2018). In this work, η is computed numerically using `pyCO2SYS` (Humphreys et al., 2020), taking as inputs physical and biogeochemical model output: temperature, salinity, DIC, ALK, silicate, and phosphate. Total borate, which `pyCO2SYS` includes in the alkalinity, is estimated internally from salinity. The sensitivity β then follows from the analytical expression of Tyka et al. (2026):

$$\beta = \frac{[\mathrm{DIC}]}{[\mathrm{CO_2}]} - \frac{[\mathrm{HCO_3^-}] + 2[\mathrm{CO_3^{2-}}]}{[\mathrm{CO_2}]}\,\eta, \tag{A1}$$

where $[\mathrm{CO_2}]$, $[\mathrm{HCO_3^-}]$, and $[\mathrm{CO_3^{2-}}]$ are the dissolved $\mathrm{CO_2}$, bicarbonate, and carbonate ion concentrations returned by `pyCO2SYS`.

The carbonate equilibrium relates DIC, $\mathrm{CO_2}$, and ALK, so that any one can be expressed as a function of the other two. In Section 2.1, the flux anomaly is linearized by expanding $[\mathrm{CO_2}] = g([\mathrm{DIC}], [\mathrm{ALK}])$ to first order (Eq. (6)). To replace $\partial[\mathrm{CO_2}]/\partial[\mathrm{ALK}]$ appearing in that expansion, it is convenient to take the equivalent view $[\mathrm{DIC}] = h([\mathrm{CO_2}], [\mathrm{ALK}])$. Taking the total differential gives

$$d[\mathrm{DIC}] = \underbrace{\left.\frac{\partial h}{\partial[\mathrm{CO_2}]}\right|_{[\mathrm{ALK}]}}_{\beta} d[\mathrm{CO_2}] + \underbrace{\left.\frac{\partial h}{\partial[\mathrm{ALK}]}\right|_{[\mathrm{CO_2}]}}_{\eta} d[\mathrm{ALK}] \tag{A2}$$

with β and η defined in Eq. 7. Setting $d[\mathrm{DIC}] = 0$ and solving for $d[\mathrm{ALK}]/d[\mathrm{CO_2}]$ yields

$$\left.\frac{\partial[\mathrm{ALK}]}{\partial[\mathrm{CO_2}]}\right|_{[\mathrm{DIC}]} = -\frac{\beta}{\eta},$$

which is Eq. (8).

Note that applying the chain rule directly to $\partial[\mathrm{ALK}]/\partial[\mathrm{CO_2}]$ by expanding via DIC would yield $+\beta/\eta$, with the wrong sign. The error arises because the held-constant variable differs in each factor ($[\mathrm{CO_2}]$ in η^{-1} versus $[\mathrm{ALK}]$ in β), so the standard chain rule does not apply. The total-differential approach above correctly enforces the carbonate equilibrium constraint.

Appendix B Open Research Section

All code needed to reproduce this study is available in three repositories, which document the model source code, experimental designs, and workflows. The ROMS and CESM implementations of the CDR tracer method are provided in `https://github.com/NoraLoose/CDR-Tracer-ROMS` (Loose, 2026c) and `https://github.com/NoraLoose/CDR-Tracer-CESM` (Loose, 2026a), respectively. The former also contains the setup of the ROMS/MARBL truth and ROMS CDR tracer experiments, and the latter the workflow for the global CESM CDR tracer experiments. The notebooks that produce all figures are provided in `https://github.com/NoraLoose/CDR-Tracer-Paper` (Loose, 2026b).

Conflict of Interest declaration

The authors declare there are no conflicts of interest for this manuscript.

Acknowledgments

This work was funded by the Advanced Research Projects Agency–Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001838; the Chan Zuckerberg Initiative; the Grantham Foundation for the Protection of the Environment; and the Ocean Resilience and Climate Alliance (ORCA). This research used resources of the National Energy Research Scientific Computing Center, a Department of Energy User Facility, using award ERCAP0035028. The views and opinions of the authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

Bach, L. T., Tyka, M. D., Wang, B., & Fennel, K. (2026). Lethal by Design? Resolving Differences Between Experimental and Real-World Alkalinity Perturbations in Ocean Alkalinity Enhancement. *Journal of Geophysical Research: Oceans*, *131*(8), e2025JC023598. Retrieved 2026-08-14, from `https://onlinelibrary.wiley.com/doi/abs/10.1029/2025JC023598` (_eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2025JC023598) doi: 10.1029/2025JC023598

C-Star contributors. (2025, December). *C-Star*. [C]Worthy. Retrieved from `https://github.com/CWorthy-ocean/C-Star` (original-date: 2024-08-08T23:01:29Z)

Carter, B. R., Bittig, H. C., Fassbender, A. J., Sharp, J. D., Takeshita, Y., Xu, Y.-Y., ... Barbero, L. (2021, December). New and updated global empirical seawater property estimation routines. *Limnology and Oceanography: Methods*, *19*(12), 785–809. Retrieved 2026-09-25, from `https://aslopubs.onlinelibrary.wiley.com/doi/10.1002/lom3.10461` doi: 10.1002/lom3.10461

Damien, P., et al. (2026). Development of a mesoscale- and tide-resolving ocean–biogeochemistry model for the pacific: Tidal enhancement of the biological carbon pump. *Journal of Advances in Modeling Earth Systems*. doi: 10.1029/2025MS005458

Doney, S. C., Fabry, V. J., Feely, R. A., & Kleypas, J. A. (2009). Ocean Acidification: The Other CO2 Problem. *Annual Review of Marine Science*, *1*, 169–192. Retrieved 2026-09-24, from `https://www.annualreviews.org/content/journals/10.1146/annurev.marine.010908.163834` doi: 10.1146/annurev.marine.010908.163834

Fennel, K. (2026). The Verification Challenge of Marine Carbon Dioxide Removal. *Annual Review of Marine Science*, *18*, 141–164. Retrieved from

https://www.annualreviews.org/content/journals/10.1146/annurev-marine-032123-025717 doi: 10.1146/annurev-marine-032123-025717

Garcia, H. E., Boyer, T. P., Baranova, O. K., Locarnini, R. A., Mishonov, A. V., Grodsky, A., ... Zweng, M. M. (2019). World ocean atlas 2018: Product documentation [Computer software manual]. NOAA/NCEI.

Gregor, L., & Gruber, N. (2021, March). OceanSODA-ETHZ: a global gridded data set of the surface ocean carbonate system for seasonal to decadal studies of ocean acidification. *Earth System Science Data*, *13*(2), 777–808. doi: 10.5194/essd-13-777-2021

He, J., & Tyka, M. D. (2023, January). Limits and CO_2 equilibration of near-coast alkalinity enhancement. *Biogeosciences*, *20*(1), 27–43. Retrieved 2025-05-16, from https://bg.copernicus.org/articles/20/27/2023/ doi: 10.5194/bg-20-27-2023

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., ... Thépaut, J.-N. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, *146*(730), 1999-2049. Retrieved from https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803 doi: 10.1002/qj.3803

Ho, M., Molemaker, J., Damien, P., Long, M. C., Bianchi, D., & McWilliams, J. C. (2026, January). Ocean alkalinity enhancement in an estuary. *Frontiers in Climate*, *7*. Retrieved 2026-09-25, from https://www.frontiersin.org/journals/climate/articles/10.3389/fclim.2025.1665329/full doi: 10.3389/fclim.2025.1665329

Huang, Y., Tagliabue, A., & Cassar, N. (2022, May). Data-Driven Modeling of Dissolved Iron in the Global Ocean. *Frontiers in Marine Science*, *9*. Retrieved from https://www.frontiersin.org/journals/marine-science/articles/10.3389/fmars.2022.837183/full (Publisher: Frontiers) doi: 10.3389/fmars.2022.837183

Humphreys, M. P., Daniels, C. J., Wolf-Gladrow, D. A., Tyrrell, T., & Achterberg, E. P. (2018). On the influence of marine biogeochemical processes over co2 exchange between the atmosphere and ocean. *Marine Chemistry*, *199*, 1-11. Retrieved from https://www.sciencedirect.com/science/article/pii/S0304420317302062 doi: https://doi.org/10.1016/j.marchem.2017.12.006

Humphreys, M. P., Gregor, L., Pierrot, D., van Heuven, S. M. A. C., Lewis, E. R., & Wallace, D. W. R. (2020). *Pyco2sys: marine carbonate system calculations in python.* Zenodo. Retrieved from https://zenodo.org/record/3744275 doi: 10.5281/ZENODO.3744275

IPCC. (2022). Summary for policymakers. In P. R. Shukla, J. Skea, R. Slade, et al. (Eds.), *Climate change 2022: Mitigation of climate change. contribution of working group iii to the sixth assessment report of the intergovernmental panel on climate change.* Cambridge, UK and New York, NY, USA: Cambridge University Press. doi: 10.1017/9781009157926.001

Kanzaki, Y., Planavsky, N. J., & Reinhard, C. T. (2023). New estimates of the storage permanence and ocean co-benefits of enhanced rock weathering. *PNAS Nexus*, *2*(4), pgad059. doi: 10.1093/pnasnexus/pgad059

Kobayashi, S., Ota, Y., Harada, Y., Ebita, A., Moriya, M., Onoda, H., ... Takahashi, K. (2015). The JRA-55 Reanalysis: General Specifications and Basic Characteristics. *Journal of the Meteorological Society of Japan. Ser. II*, *93*(1), 5–48. doi: 10.2151/jmsj.2015-001

Landschützer, P., Gruber, N., & Bakker, D. C. E. (2016). Decadal variations and trends of the global ocean carbon sink. *Global Biogeochemical Cycles*, *30*(10), 1396–1417. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.1002/2015GB005359 doi: 10.1002/2015GB005359

Lauvset, S. K., Key, R. M., Olsen, A., Heuven, S. v., Velo, A., Lin, X., ... Watelet, S. (2016, August). A new global interior ocean mapped climatology: the 1°

× 1° GLODAP version 2. *Earth System Science Data*, *8*(2), 325–340. Retrieved from https://www.earth-syst-sci-data.net/8/325/2016/ doi: 10.5194/essd-8-325-2016

Lee, K., Tong, L. T., Millero, F. J., Sabine, C. L., Dickson, A. G., Goyet, C., ... Key, R. M. (2006, October). Global relationships of total alkalinity with salinity and temperature in surface waters of the world's oceans. *Geophysical Research Letters*, *33*(19). Retrieved 2026-09-25, from https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2006GL027207 doi: 10.1029/2006GL027207

Lellouche, E., Jean-Michel Greiner, Bourdallé-Badie, R., Garric, G., Melet, A., Drévillon, M., Bricaud, C., ... Le Traon, P.-Y. (2021, July). The Copernicus Global 1/12° Oceanic and Sea Ice GLORYS12 Reanalysis. *Frontiers in Earth Science*, *9*. Retrieved from https://www.frontiersin.org/journals/earth-science/articles/10.3389/feart.2021.698876/full (Publisher: Frontiers) doi: 10.3389/feart.2021.698876

Long, M. C., Moore, J. K., Lindsay, K., Levy, M., Doney, S. C., Luo, J. Y., ... Sylvester, Z. T. (2021). Simulations with the marine biogeochemistry library (marbl). *Journal of Advances in Modeling Earth Systems*, *13*(12), e2021MS002647. Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002647 (e2021MS002647 2021MS002647) doi: https://doi.org/10.1029/2021MS002647

Loose, N. (2026a). *NoraLoose/CDR-Tracer-CESM: Version used in Loose et al., JAMES submission.* Zenodo. Retrieved from https://doi.org/10.5281/zenodo.23688722 ([Software]) doi: 10.5281/zenodo.23688722

Loose, N. (2026b). *NoraLoose/CDR-Tracer-Paper: Version used in Loose et al., JAMES submission.* Zenodo. Retrieved from https://doi.org/10.5281/zenodo.23688747 ([Software]) doi: 10.5281/zenodo.23688747

Loose, N. (2026c). *NoraLoose/CDR-Tracer-ROMS: Version used in Loose et al., JAMES submission.* Zenodo. Retrieved from https://doi.org/10.5281/zenodo.23688730 ([Software]) doi: 10.5281/zenodo.23688730

Loose, N., Heimbach, P., Pillar, H. R., & Nisancioglu, K. H. (2020). Quantifying Dynamical Proxy Potential Through Shared Adjustment Physics in the North Atlantic. *Journal of Geophysical Research: Oceans*, *125*(9), e2020JC016112. Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020JC016112 (_eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2020JC016112) doi: 10.1029/2020JC016112

National Academies of Sciences, Engineering, and Medicine. (2022). *A research strategy for ocean-based carbon dioxide removal and sequestration.* Washington, DC: The National Academies Press. doi: 10.17226/26278

Oschlies, A. (2009, August). Impact of atmospheric and terrestrial CO_2 feedbacks on fertilization-induced marine carbon uptake. *Biogeosciences*, *6*(8), 1603–1613. Retrieved 2026-09-25, from https://bg.copernicus.org/articles/6/1603/2009/ doi: 10.5194/bg-6-1603-2009

Oschlies, A., et al. (Eds.). (2023). *Guide to best practices in ocean alkalinity enhancement research* (Vol. 2-oae2023). Copernicus Publications. Retrieved from https://doi.org/10.5194/sp-2-oae2023 doi: 10.5194/sp-2-oae2023

Paulot, F., Jacob, D. J., Johnson, M. T., Bell, T. G., Baker, A. R., Keene, W. C., ... Stock, C. A. (2015). Global oceanic emission of ammonia: Constraints from seawater and atmospheric observations. *Global Biogeochemical Cycles*, *29*(8), 1165–1178. Retrieved 2026-09-09, from https://onlinelibrary.wiley.com/doi/abs/10.1002/2015GB005106 doi: 10.1002/2015GB005106

Reisinger, A., Fuglestvedt, J. S., Pirani, A., Geden, O., Jones, C. D., Maharaj, S., ... Seneviratne, S. I. (2025). Overshoot: A conceptual review of exceeding

and returning to global warming of 1.5°c. *Annual Review of Environment and Resources*, *50*, 185–217. doi: 10.1146/annurev-environ-111523-102029

Renforth, P., & Henderson, G. (2017). Assessing ocean alkalinity for carbon sequestration. *Reviews of Geophysics*, *55*(3), 636–674. Retrieved 2026-08-17, from https://onlinelibrary.wiley.com/doi/abs/10.1002/2016RG000533 (_eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2016RG000533) doi: 10.1002/2016RG000533

Shchepetkin, A. F., & McWilliams, J. C. (2005, January). The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modelling*, *9*(4), 347–404. Retrieved from https://www.sciencedirect.com/science/article/pii/S1463500304000484 doi: 10.1016/j.ocemod.2004.08.002

Smith, S., Geden, O., Gidden, M., Lamb, W., Nemet, G., Minx, J., ... Edwards, M. (2024, May). *The State of Carbon Dioxide Removal – 2nd Edition* (Tech. Rep.). State of Carbon Dioxide Removal. Retrieved from https://osf.io/f85qj

Tyka, M. D. (2025, January). Efficiency metrics for ocean alkalinity enhancements under responsive and prescribed atmospheric pCO_2 conditions. *Biogeosciences*, *22*(1), 341–353. Retrieved 2025-05-16, from https://bg.copernicus.org/articles/22/341/2025/ doi: 10.5194/bg-22-341-2025

Tyka, M. D., Zhou, M., Yankovsky, E., & Carroll, D. (2026, July). Substantial inter-model variation in OAE efficiency between the CESM2/MARBL and ECCO-Darwin ocean biogeochemistry models. *Biogeosciences*, *23*(14), 4943–4966. Retrieved 2026-09-25, from https://bg.copernicus.org/articles/23/4943/2026/ doi: 10.5194/bg-23-4943-2026

United Nations Environment Programme. (2026). *Limiting overshoot: Navigating exceedance of 1.5°c and pathways towards return* (Tech. Rep.). Nairobi: UNEP. Retrieved from https://www.unep.org/resources/limiting-overshoot-navigating-exceedance

Wanninkhof, R. (2014). Relationship between wind speed and gas exchange over the ocean revisited. *Limnology and Oceanography: Methods*, *12*(6), 351–362. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.4319/lom.2014.12.351 (_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.4319/lom.2014.12.351) doi: 10.4319/lom.2014.12.351

Wanninkhof, R., Park, G.-H., Takahashi, T., Feely, R. A., Bullister, J. L., & Doney, S. C. (2013, April). Changes in deep-water CO2 concentrations over the last several decades determined from discrete pCO2measurements. *Deep Sea Research Part I: Oceanographic Research Papers*, *74*, 48–63. Retrieved 2026-08-17, from https://www.sciencedirect.com/science/article/pii/S0967063713000034 doi: 10.1016/j.dsr.2012.12.005

Yang, S., Chang, B. X., Warner, M. J., Weber, T. S., Bourbonnais, A. M., Santoro, A. E., ... Bianchi, D. (2020, June). Global reconstruction reduces the uncertainty of oceanic nitrous oxide emissions and reveals a vigorous seasonal cycle. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(22), 11954–11960. doi: 10.1073/pnas.1921914117

Yankovsky, E., Zhou, M., Tyka, M., Bachman, S., Ho, D. T., Karspeck, A., & Long, M. C. (2025, October). Impulse response functions as a framework for quantifying ocean-based carbon dioxide removal. *Biogeosciences*, *22*(20), 5723–5739. Retrieved 2025-12-08, from https://bg.copernicus.org/articles/22/5723/2025/ doi: 10.5194/bg-22-5723-2025

Yeager, S. G., Rosenbloom, N., Glanville, A. A., Wu, X., Simpson, I., Li, H., ... King, T. (2022, August). The Seasonal-to-Multiyear Large Ensemble (SMYLE) prediction system using the Community Earth System Model version 2. *Geo-*

scientific Model Development, *15*(16), 6451–6493. Retrieved 2025-01-10, from https://gmd.copernicus.org/articles/15/6451/2022/ (Publisher: Copernicus GmbH) doi: 10.5194/gmd-15-6451-2022

Zhou, M., Timmermans, M.-L., Yankovsky, E., Tyka, M. D., & Long, M. C. (2026). Global intercomparison of ocean-based geochemical carbon dioxide removal efficiencies. *Nature Climate Change*. (submitted, Paper NCOMMS-26-017842B)

Zhou, M., Tyka, M. D., Ho, D. T., Yankovsky, E., Bachman, S., Nicholas, T., ... Long, M. C. (2024, November). Mapping the global variation in the efficiency of ocean alkalinity enhancement for carbon dioxide removal. *Nature Climate Change*, 1–7. Retrieved 2024-11-08, from https://www.nature.com/articles/s41558-024-02179-9 doi: 10.1038/s41558-024-02179-9

Linearized Carbonate Chemistry for Efficient Modeling of Ocean-Based Carbon Dioxide Removal

Nora Loose[1,2,3], Matthew Long[1], Scott Bachman[1], Pierre Damien[4], Scott Eilerman[1], Alicia Karspeck[1], Jeroen Molemaker[4], Dafydd Stephenson[1], Michael Tyka[5], Abigale Wyatt[1]

[1][C]Worthy LLC, Boulder, CO, USA
[2]Nansen Environmental and Remote Sensing Center, Bergen, Norway
[3]Bjerknes Centre for Climate Research, Bergen, Norway
[4]Department of Atmospheric and Oceanic Sciences, University of California, Los Angeles, CA, USA
[5]Applied Research, Google Inc., Seattle, USA

SUPPORTING INFORMATION

1 Text S1: Box model for CDR tracer error

The following box model is a simplified version of the CDR tracer framework derived in Sections 2.1–2.2, reduced to a single well-mixed surface box with no horizontal transport. It provides analytical insight into the error structure seen in the flux anomaly and efficiency curves between the CDR tracer estimate and the truth in the presence of nonlinearity.

We consider a surface box of effective depth h, gas transfer velocity k, and intervention source rate S applied over $0 \le t \le T$. For OAE, the alkalinity anomaly is conservative and the DIC anomaly relaxes toward equilibrium with the atmosphere, consistent with Eqs. (10)–(11) (and Eqs. (12)–(13)) but in the absence of advection and diffusion. Denoting the alkalinity anomaly as A and the DIC anomaly as D, the governing equations are

$$\frac{\mathrm{d}A}{\mathrm{d}t} = \frac{S}{h} \quad (t \le T), \tag{S1}$$

$$\frac{\mathrm{d}D}{\mathrm{d}t} = a(\eta A - D), \tag{S2}$$

where $\eta A - D$ is the disequilibrium between alkalinity-driven DIC uptake and the current DIC anomaly, and the relaxation rate

$$a = \frac{k/\beta}{h} \tag{S3}$$

sets the timescale of air-sea CO_2 equilibration (faster for smaller β). The air-sea CO_2 flux anomaly is $F = h\,a\,(\eta A - D)$, corresponding to $\delta\tilde{F}_{CO_2}$ in Eq. (14). The DOR case is analogous, with the source S applied to the DIC equation rather than the alkalinity equation, and $\eta = 1$: the DIC deficit plays the role of the disequilibrium.

We solve for the disequilibrium $G \equiv \eta A - D$, which is proportional to the flux ($F = h\,a\,G$):

$$\frac{\mathrm{d}G}{\mathrm{d}t} = \eta\frac{\mathrm{d}A}{\mathrm{d}t} - \frac{\mathrm{d}D}{\mathrm{d}t} = \eta\frac{\mathrm{d}A}{\mathrm{d}t} - aG. \tag{S4}$$

Corresponding author: Nora Loose, nora.loose@nersc.no

This is a linear ODE with constant coefficients. With forcing on ($\mathrm{d}A/\mathrm{d}t = S/h$, $G(0) = 0$), the solution is $G = \frac{\eta S}{ah}(1 - e^{-at})$; with forcing off ($\mathrm{d}A/\mathrm{d}t = 0$), G decays exponentially from its value at $t = T$. The flux anomaly is therefore

$$F(t) = \begin{cases} \eta\, S\left(1 - e^{-at}\right) & t \le T, \\ \eta\, S\left(1 - e^{-aT}\right) e^{-a(t-T)} & t > T. \end{cases} \tag{S5}$$

Normalizing the total intervention to unity ($S\,T = 1$), the CDR efficiency $C(t) = \int_0^t F\,\mathrm{d}t'$ has the closed form

$$C(t) = \begin{cases} \dfrac{\eta}{T}\left(t - \dfrac{1 - e^{-at}}{a}\right) & t \le T, \\ C(T) + \dfrac{\eta}{T}\left(1 - e^{-aT}\right)\dfrac{1 - e^{-a(t-T)}}{a} & t > T, \end{cases} \tag{S6}$$

and plateaus at $C(\infty) = \eta$.

The CDR tracer method uses control-run sensitivities β^{ctrl} and η^{ctrl}, while the truth uses perturbed sensitivities β^{pert} and η^{pert}. OAE increases β and decreases η, so we have $\beta^{\mathrm{pert}} > \beta^{\mathrm{ctrl}}$ and $\eta^{\mathrm{pert}} < \eta^{\mathrm{ctrl}}$. The flux anomaly error and efficiency error are obtained by evaluating Eqs. (S5) and (S6) with $(\beta^{\mathrm{ctrl}}, \eta^{\mathrm{ctrl}})$ for the CDR tracer estimate and $(\beta^{\mathrm{pert}}, \eta^{\mathrm{pert}})$ for the truth, and taking the difference.

Figure 3(c–f) in the main text illustrates the flux anomaly and CDR efficiency, as well as their estimates by the CDR tracer method for the following parameter values: $\beta^{\mathrm{ctrl}} = 15$, $\beta^{\mathrm{pert}} = 24$, $\eta^{\mathrm{ctrl}} = 0.85$, $\eta^{\mathrm{pert}} = 0.82$, $h = 20\,\mathrm{m}$, $k = 4 \times 10^{-3}\,\mathrm{cm\,s^{-1}}$, $T = 30\,\mathrm{days}$. The CDR tracer estimate of the flux anomaly (panel (c), green) rises faster and reaches a higher peak than the truth (orange), because the lower β^{ctrl} implies a faster relaxation rate a^{ctrl}. After the forcing stops at $t = T$, both fluxes decay exponentially, but the CDR tracer flux decays faster, so the flux error (panel (d)) is initially positive and then crosses zero at a crossover time t^* (dashed vertical line). Equating the $t > T$ expression in Eq. (S5) for the CDR tracer and truth anomaly fluxes and solving gives

$$t^* = T + \frac{1}{a^{\mathrm{ctrl}} - a^{\mathrm{pert}}} \ln\left(\frac{\eta^{\mathrm{ctrl}}\left(1 - e^{-a^{\mathrm{ctrl}}T}\right)}{\eta^{\mathrm{pert}}\left(1 - e^{-a^{\mathrm{pert}}T}\right)}\right), \tag{S7}$$

which for the parameter values above evaluates to $t^* = 133\,\mathrm{days}$.

Integrating the flux, the CDR tracer efficiency curve (panel (e), green) rises faster and converges to a higher plateau η^{ctrl} than the truth (orange, plateau η^{pert}). The efficiency error (panel (f)) shows an initial overestimation that decays toward a persistent positive asymptote of $\eta^{\mathrm{ctrl}} - \eta^{\mathrm{pert}}$, reflecting the difference in the theoretical maximum efficiency between the control and perturbed carbonate states. This asymptotic offset is the dominant long-term error of the CDR tracer method in the presence of nonlinearity.

The analytic solution above assumes constant coefficients. To assess whether the characteristic error structure discussed above holds under more realistic, time-varying parameters, we extend the box model by integrating the governing equations (S1), (S2) numerically using an adaptive Runge–Kutta scheme (RK45; a fourth-order method with fifth-order error control for adaptive step sizing), allowing $\beta(t)$, $\eta(t)$, and $h(t)$ to vary in time. The mixed layer depth enters only through the relaxation rate $a(t) = k/(\beta(t)\,h(t))$. The conditions $\beta^{\mathrm{ctrl}}(t) < \beta^{\mathrm{pert}}(t)$ and $\eta^{\mathrm{ctrl}}(t) > \eta^{\mathrm{pert}}(t)$ are maintained at every instant. The intervention lasts $T = 30\,\mathrm{days}$ and the simulation runs for five years, long enough for the efficiency to plateau.

We consider four scenarios (Figure S1):

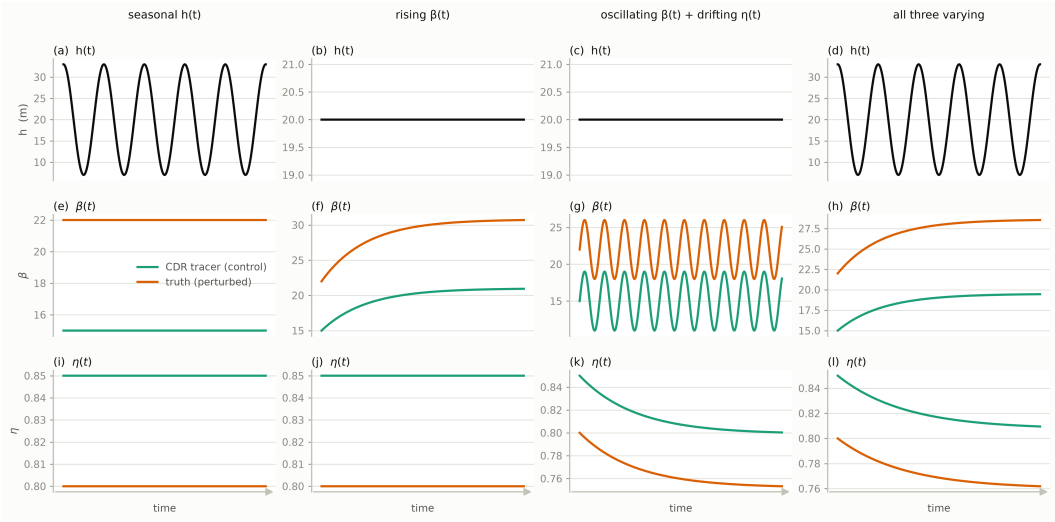


Figure S1: Time-varying parameters used in the four box model scenarios (1–4), shown over the five-year simulation period. Rows show the mixed layer depth $h(t)$ (top), $\beta(t)$ (middle), and $\eta(t)$ (bottom). Columns correspond to scenarios (1) seasonal h, (2) rising β, (3) oscillating β and drifting η, and (4) all three varying simultaneously. Green lines show the control-run values used by the CDR tracer; orange lines show the perturbed values used by the truth. The condition $\beta^{\mathrm{ctrl}}(t) < \beta^{\mathrm{pert}}(t)$ and $\eta^{\mathrm{ctrl}}(t) > \eta^{\mathrm{pert}}(t)$ is maintained at every instant in all scenarios.

(1) Seasonal $h(t)$ (deep in winter, shallow in summer, repeating annually); β and η constant with a fixed gap between control and truth.
(2) Constant h; $\beta(t)$ rising asymptotically to $1.4\,\beta_0$ with the same fractional gap between control and truth; η constant.
(3) Constant h; $\beta(t)$ oscillating sub-annually (180-day period) with a constant offset gap; $\eta(t)$ drifting downward asymptotically at the same fractional rate for both control and truth.
(4) All three varying simultaneously: seasonal $h(t)$, rising $\beta(t)$, and drifting $\eta(t)$.

Figure S2 shows the resulting flux anomaly and CDR efficiency curves for all four scenarios. In all four scenarios the CDR tracer overestimates the flux anomaly during the intervention and underestimates it after the crossover time t^*. The efficiency error remains positive throughout and converges to a positive asymptote in all cases. The error structure seen in the constant-coefficient case (Figure 3(c–f) in the main text) is therefore robust to time-varying parameters, suggesting that the box model conclusions carry over to more realistic settings.

Text S2. Temporal sampling and averaging of β and η

The CDR tracer run requires β and η as time-varying input fields, which are prescribed at a finite frequency (here daily) and interpolated in time by the model. This text addresses three questions: which inputs the CDR tracer run ideally requires, how well common ways of computing β and η approximate these ideal inputs, and how the inputs should be interpolated in time. We then describe the choices made in this study and conclude with practical recommendations.

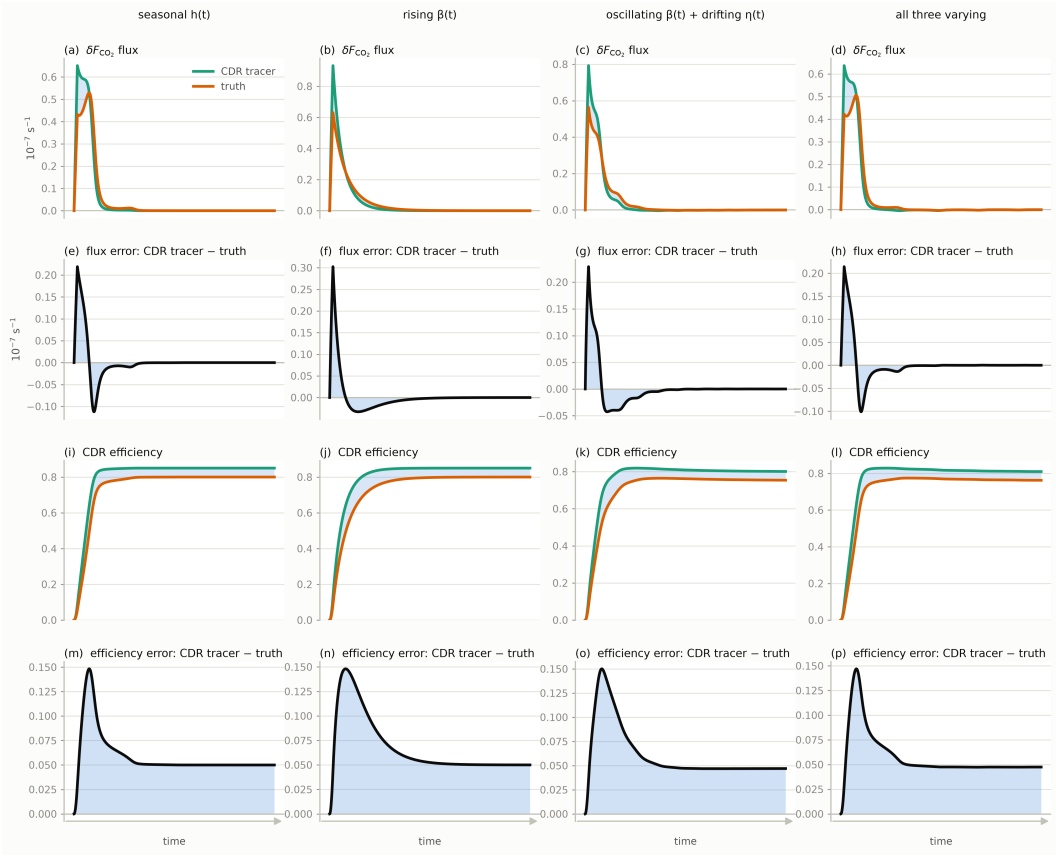


Figure S2: Box model results for the four time-varying scenarios (1–4), analogous to Figure 3(c–f) in the main text. Rows show the CO_2 flux anomaly δF_{CO_2} (first row), flux anomaly error (second row), CDR efficiency (third row), and efficiency error (fourth row). Green lines show the CDR tracer estimate (using control-run carbonate sensitivities); orange lines show the truth (using perturbed carbonate sensitivities). The parameter time series for each scenario are shown in Figure S1. In all four scenarios the efficiency error remains positive throughout and converges to a positive asymptote, consistent with the analytic result for constant coefficients (Figure 3(f) in the main text).

To derive the ideal inputs, consider an input interval of length Δt, and denote interval means by $\langle\cdot\rangle$ and deviations from them by primes. The gas exchange term in the CDR tracer equation (Eq. (9) of the main text) is $G = -(k/\beta)\, D$, with $D = c_{\delta\mathrm{DIC}} - \eta\, c_{\delta\mathrm{ALK}}$, so that

$$\langle G\rangle = -\Big\langle \frac{k}{\beta}\Big\rangle \langle D\rangle - \Big\langle \Big(\frac{k}{\beta}\Big)' D' \Big\rangle. \tag{S8}$$

The second term vanishes if, within the interval, D varies little or its fluctuations are uncorrelated with those of k/β. For daily intervals, this holds well: k varies with the winds, but β and D, which are set by the water mass and the tracer distribution, change little within a day. For longer intervals, β and D can covary, for example when a front or eddy advects a different water mass past a grid cell, which introduces an additional error of second order. Neglecting the second term, we obtain

$$\langle G\rangle \approx -\Big\langle \frac{k}{\beta}\Big\rangle \langle D\rangle. \tag{S9}$$

The gas transfer velocity k is computed online at every time step, whereas β enters as an input field β_{in}. If β_{in} is constant within the interval (i.e., applied piecewise-constant; time interpolation is discussed below), the CDR tracer run produces $\langle k\rangle/\beta_{\mathrm{in}}$, which matches $\langle k/\beta\rangle$ for

$$\beta_{\mathrm{in}} = \frac{\langle k\rangle}{\langle k/\beta\rangle}. \tag{S10}$$

Analogously, matching the term $\langle k\eta/\beta\rangle$ requires

$$\eta_{\mathrm{in}} = \frac{\langle k\eta/\beta\rangle}{\langle k/\beta\rangle}. \tag{S11}$$

The ideal β_{in} is therefore a k-weighted harmonic mean of β, because β appears in the denominator of the flux, and the ideal η_{in} is a k/β-weighted arithmetic mean of η. Writing $\beta = \langle\beta\rangle(1+b)$ and $k = \langle k\rangle(1+\kappa)$ with $\langle b\rangle = \langle\kappa\rangle = 0$, and expanding to second order, gives

$$\beta_{\mathrm{in}} \approx \langle\beta\rangle \left(1 - \langle b^2\rangle + \langle\kappa b\rangle\right). \tag{S12}$$

The arithmetic mean $\langle\beta\rangle$ thus deviates from the ideal input at second order, through the relative variance of β within the interval and its covariance with k.

In practice, β and η are rarely computed in this ideal form. Instead, interval-mean fields are typically obtained in one of two ways: by averaging fields computed from the instantaneous carbonate system state, either online at every time step or offline from high-frequency snapshots, or by computing them from time-averaged state variables $\mathbf{X}$. The first approach yields $\langle\beta\rangle$, which differs from the ideal input as quantified by Eq. (S12). The second yields

$$\beta(\langle\mathbf{X}\rangle) \approx \langle\beta\rangle - \frac{1}{2}\sum_{i,j} \frac{\partial^2\beta}{\partial X_i \partial X_j}\, \mathrm{Cov}(X_i, X_j), \tag{S13}$$

which deviates further from the ideal input because of the nonlinearity of the carbonate system. Both approaches thus differ from the ideal input at second order in the variability within the averaging interval, but the first is preferable because it avoids the additional nonlinearity error. For daily intervals, all of these differences are small.

So far, we have assumed that the input fields are applied piecewise-constant within each interval. Models, however, typically interpolate input fields in time, and whether this is appropriate depends on whether the fields represent snapshots or interval means. Snapshots represent β and η at discrete times and should be interpolated linearly in time: holding them constant until the next snapshot introduces errors of first order in Δt, including a lag of $\Delta t/2$, whereas linear interpolation reduces the error to second order. Interval means, by contrast, represent β and η over the entire interval and are most consistently applied piecewise-constant, as assumed above. If interval means x_i are instead

assigned to interval midpoints and interpolated linearly, the mean of the applied field over interval i becomes $(x_{i-1}+6x_i+x_{i+1})/8 \approx x_i+(\Delta t^2/8)\,\ddot{x}$, so interpolation applies an additional smoothing that preserves the long-term mean but redistributes the signal between neighboring intervals. For a sinusoidal signal with angular frequency ω, interval averaging reduces the amplitude by a factor $\sin(\omega\Delta t/2)/(\omega\Delta t/2)$, and subsequent interpolation by a further factor $(6+2\cos\omega\Delta t)/8$. For the seasonal cycle and monthly intervals, these factors are 0.989 and 0.967; for daily intervals, both are indistinguishable from one. Snapshots, in turn, are appropriate only at high frequency, because variability between snapshots aliases into the input.

We now relate the choices made in this study to these considerations. For ROMS, β and η are computed offline from daily snapshots of the carbonate system state and provided to the CDR tracer run as snapshots, which the model interpolates linearly in time, as recommended above. Because these fields are computed from the instantaneous state, they avoid the nonlinearity error of Eq. (S13). The remaining error arises from variability between the daily snapshots that linear interpolation does not capture (Eqs. (S8) and (S12)), which is small. The coarser fields used in the `MON` and `MEAN` experiments are obtained by averaging these daily β and η snapshot fields (Section 2.5 of the main text), and thus follow the first approach above over longer intervals, for which both the deviation from the ideal input (Eq. (S12)) and the neglected covariance between β and D (Eq. (S8)) grow. In addition, the monthly means in the `MON` experiment are interpolated linearly in time, which further smooths them. The `MON` and `MEAN` experiments therefore measure the combined effect of these approximations (Section 6 of the main text). For CESM, β and η are computed from daily-averaged state variables, following the second approach, which introduces the additional nonlinearity error of Eq. (S13). The daily means are then interpolated linearly, which adds the smoothing described above. Both effects are negligible at daily resolution. Figures S6 and S7 provide an upper bound on the nonlinearity error by comparing both approaches for monthly and time-mean averaging: even at these longer intervals, the differences are small compared to the effects of degrading the sensitivities, although regionally significant differences occur. Consistent with this, the CESM CDR tracer runs agree well with the explicit calculations (Section 4.2 of the main text).

Based on these considerations and results, we recommend the following (summarized in Table 1). At daily frequency, β and η should be computed from the instantaneous carbonate system state. They can then either be provided as daily snapshots and interpolated linearly in time, as in our ROMS reference configuration, or be diagnosed online, averaged over each day, and applied piecewise-constant. Computing β and η from daily-averaged state variables, as in our CESM configuration, is an acceptable shortcut: its nonlinearity error is negligible at this frequency, and so is the choice of interpolation. At monthly frequency, β and η should likewise be computed from the instantaneous state, at daily or higher frequency, and only then averaged over each month, as in our `MON` experiment. They should not be computed from monthly-mean state variables, which incurs the larger nonlinearity error bounded by Figures S6 and S7, nor taken from monthly snapshots, which alias the variability within the month. The resulting monthly means are best applied piecewise-constant over each month, because linear interpolation, as used in our `MON` experiment, further damps their variability; the `MON` results thus represent a slightly conservative estimate of the accuracy achievable with monthly inputs. If β and η are diagnosed online, the ideal inputs of Eqs. (S10) and (S11) can in principle be formed directly by accumulating $\langle k\rangle$, $\langle k/\beta\rangle$, and $\langle k\eta/\beta\rangle$ in the control run, since k is identical in the control and CDR tracer runs; we have not tested this approach here.

Text S3. Approximate expressions for β and η

The carbonate sensitivities β and η are strongly anticorrelated (Figure 2 of the main text). To explain this, we derive approximate expressions for both. They are defined as

Table 1: Approaches for preparing the carbonate sensitivities β and η as input to the CDR tracer run (Text S2).

Input to CDR tracer run	Interpolation	Approximation errors	Used here	Recommended
Daily				
Snapshots of β, η from instantaneous state	Linear	Between-snapshot variability (Eqs. (S8), (S12))	ROMS	Yes
Daily means of β, η diagnosed online	Piecewise-constant	Within-day variability (Eqs. (S8), (S12))	–	Yes
β, η from daily-mean state	Either	As above, plus nonlinearity (Eq. (S13))	CESM	Acceptable
Monthly				
Monthly means of β, η from instantaneous state	Piecewise-constant	Within-month variability (Eqs. (S8), (S12))	–	Yes
Monthly means of β, η from instantaneous state	Linear	As above, plus interpolation smoothing	`MON` (ROMS)	Acceptable
Monthly means, k-weighted (Eqs. (S10), (S11))	Piecewise-constant	Covariance of β and D (Eq. (S8))	–	Best in theory (untested)
β, η from monthly-mean state	–	Plus nonlinearity (Eq. (S13); Figures S6, S7)	–	No
Monthly snapshots	–	Aliasing	–	No

(Eq. (7) of the main text)

$$\beta = \left.\frac{\partial[\mathrm{DIC}]}{\partial[\mathrm{CO_2}]}\right|_{[\mathrm{ALK}]}, \qquad \eta = \left.\frac{\partial[\mathrm{DIC}]}{\partial[\mathrm{ALK}]}\right|_{[\mathrm{CO_2}]}, \tag{S14}$$

where

$$[\mathrm{DIC}] = [\mathrm{CO_2}] + [\mathrm{HCO_3^-}] + [\mathrm{CO_3^{2-}}], \tag{S15}$$

$$[\mathrm{ALK}] = [\mathrm{HCO_3^-}] + 2[\mathrm{CO_3^{2-}}] + [\mathrm{B(OH)_4^-}] + [\mathrm{OH^-}] - [\mathrm{H^+}] + \ldots, \tag{S16}$$

and the dots denote minor contributions from phosphate, silicate, and other acid–base systems. Without approximation, β can be written as (Eq. (A1) of the main text)

$$\beta = \frac{[\mathrm{DIC}]}{[\mathrm{CO_2}]} - \frac{[\mathrm{HCO_3^-}] + 2[\mathrm{CO_3^{2-}}]}{[\mathrm{CO_2}]}\,\eta, \tag{S17}$$

where η must be computed numerically from the full alkalinity expression (Appendix A of the main text).

To obtain analytical insight into how β and η depend on the state of the carbonate system, we approximate alkalinity by its carbonate contribution,

$$[\mathrm{ALK}] \approx [\mathrm{HCO_3^-}] + 2[\mathrm{CO_3^{2-}}], \tag{S18}$$

i.e., we neglect borate, water, and all minor contributions. For brevity, we write $C = [\mathrm{CO_2}]$, $B = [\mathrm{HCO_3^-}]$, $A = [\mathrm{CO_3^{2-}}]$, and $h = [\mathrm{H^+}]$, so that $\mathrm{DIC} = C + B + A$ and $\mathrm{ALK} \approx B + 2A$. The carbonate equilibria are

$$B = \frac{K_1 C}{h}, \qquad A = \frac{K_2 B}{h} = \frac{K_1 K_2 C}{h^2}, \tag{S19}$$

where K_1 and K_2 are the first and second dissociation constants of carbonic acid. At fixed temperature and salinity, the logarithmic differentials of the equilibria are

$$\frac{dB}{B} = \frac{dC}{C} - \frac{dh}{h}, \qquad \frac{dA}{A} = \frac{dC}{C} - 2\,\frac{dh}{h}. \tag{S20}$$

We first derive η. At fixed $\mathrm{CO_2}$ ($dC = 0$), $dB = -B\,dh/h$ and $dA = -2A\,dh/h$, so that

$$d\mathrm{DIC} = dB + dA = -(B + 2A)\,\frac{dh}{h}, \qquad d\mathrm{ALK} = dB + 2\,dA = -(B + 4A)\,\frac{dh}{h}, \tag{S21}$$

and hence

$$\eta \approx \frac{B + 2A}{B + 4A}. \tag{S22}$$

For β, we hold alkalinity fixed:

$$d\mathrm{ALK} = (B + 2A)\,\frac{dC}{C} - (B + 4A)\,\frac{dh}{h} = 0 \quad \Rightarrow \quad \frac{dh}{h} = \eta\,\frac{dC}{C}. \tag{S23}$$

Substituting into $d\mathrm{DIC} = dC + dB + dA$ gives

$$\beta \approx 1 + \frac{(B + A) - (B + 2A)\,\eta}{C} = \frac{\mathrm{DIC}}{C} - \frac{B + 2A}{C}\,\eta, \tag{S24}$$

which recovers the exact expression above, now with η given in closed form. Inserting η and using $(B + A)(B + 4A) - (B + 2A)^2 = AB$ yields

$$\beta \approx 1 + \frac{AB}{C\,(B + 4A)}. \tag{S25}$$

Both sensitivities depend on the ratio of carbonate to bicarbonate ions, $r = A/B$:

$$\eta \approx \frac{1+2r}{1+4r}, \qquad \beta \approx 1 + \frac{B}{C}\,\frac{r}{1+4r}. \tag{S26}$$

Since $d\eta/dr = -2/(1+4r)^2 < 0$ and $d[r/(1+4r)]/dr = 1/(1+4r)^2 > 0$, η decreases and β increases with r. Eliminating r via $r = (1-\eta)/(4\eta-2)$ gives

$$\beta \approx 1 + \frac{B}{C}\,\frac{1-\eta}{2}, \tag{S27}$$

which makes the anticorrelation between β and η explicit. The prefactor $B/C = K_1/h$ is itself larger in warm, high-pH waters, reinforcing the anticorrelation.

To assess the accuracy of these approximations, we compare them with the full solution of `pyCO2SYS` for typical surface seawater with ALK = 2300 μmol kg^{-1}, $S = 35$, and DIC = 2050, 2100, and 2150 μmol kg^{-1} at $T = 28$, 15, and 2°C, respectively. The approximate η overestimates the full value by 2–3%, whereas the approximate β underestimates the full value by 25–27%. This discrepancy arises mainly from neglecting borate: when the comparison is repeated with total borate set to zero, the approximate β agrees with the full value to within 4%, and the approximate η to within 0.3%. The approximations therefore capture the qualitative behavior of β and η, but are not used for any computations in this study. The β and η fields used in this study are computed from the full carbonate system, including borate, silicate, and phosphate contributions to alkalinity (Appendix A of the main text).

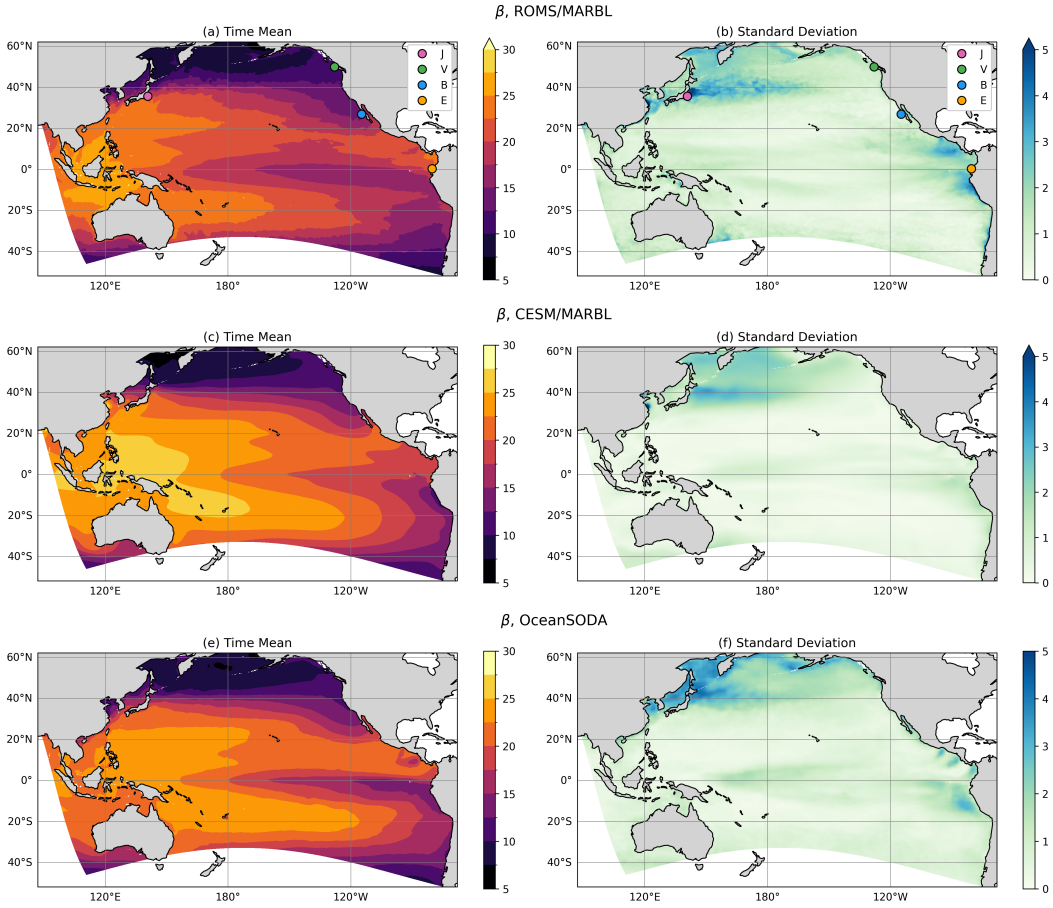


Figure S3: Time-mean (left column) and standard deviation (right column) of the carbonate sensitivity β for three model configurations and observational products: ROMS/MARBL (a,b), CESM/MARBL (c,d), and OceanSODA-ETHZ (e,f), all evaluated over the two-year period 2000-2001 and regridded to the ROMS Pacific domain. Statistics are computed from daily output for ROMS/MARBL and CESM/MARBL, and from monthly output for OceanSODA-ETHZ.

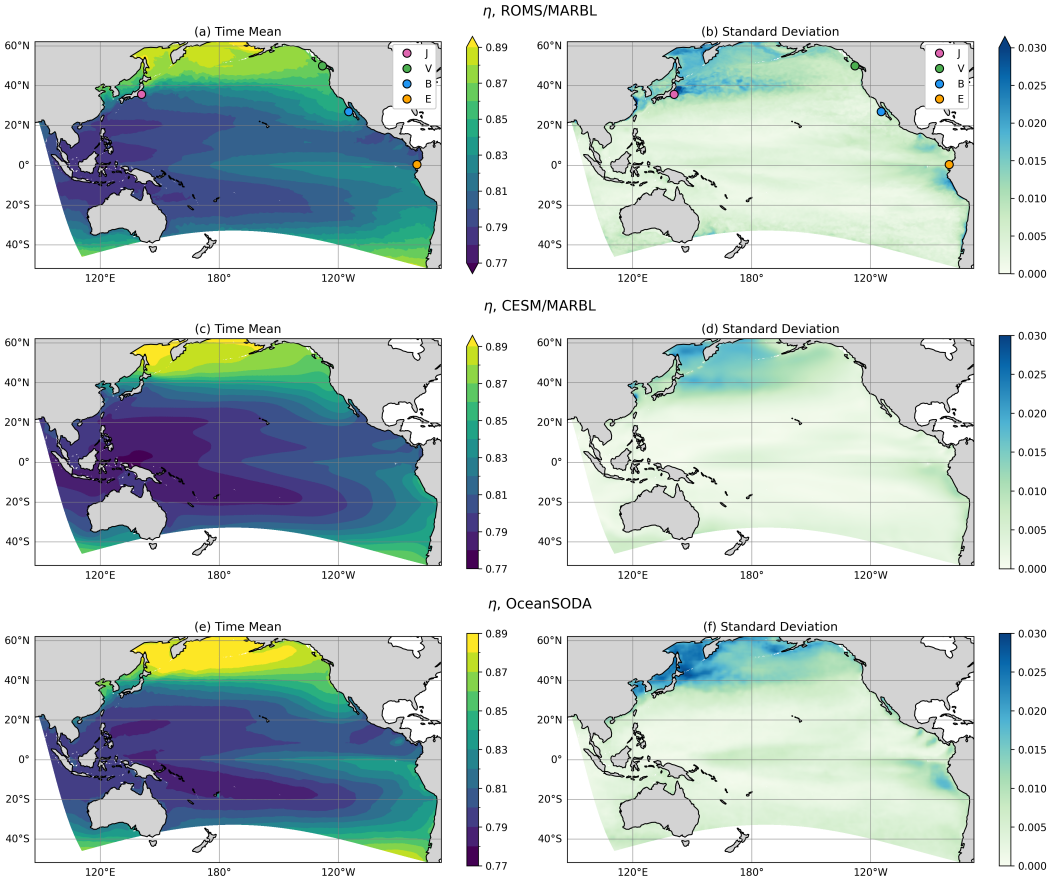


Figure S4: Same as Figure S3, but for carbonate sensitivity η.

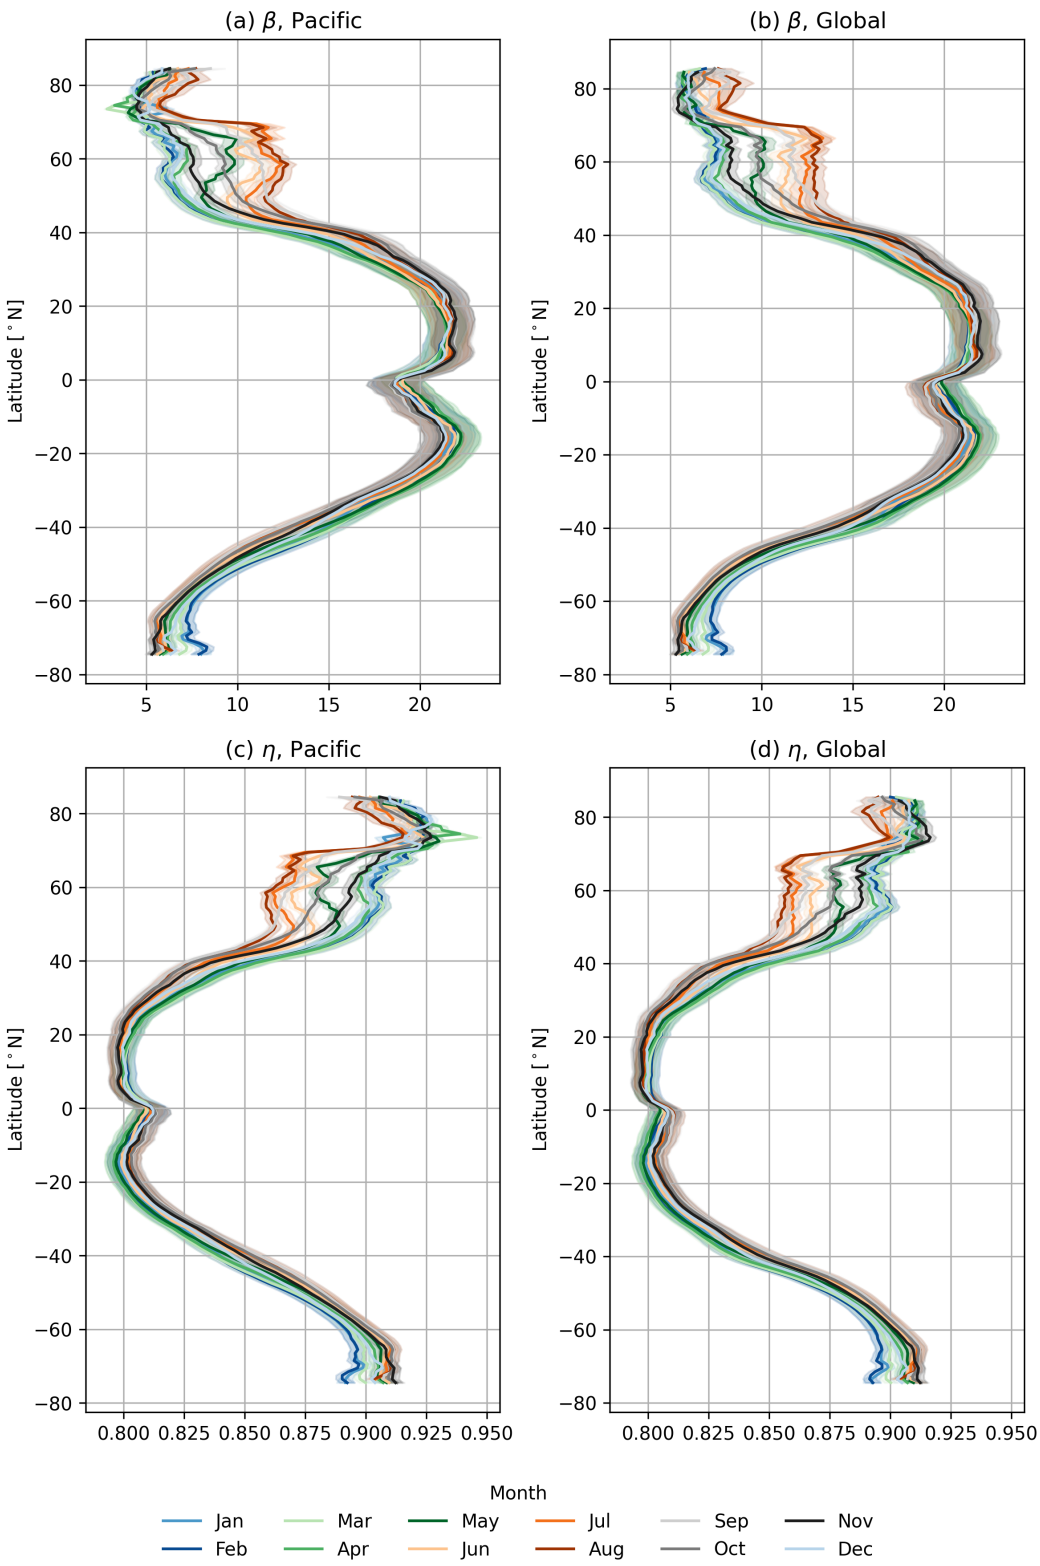


Figure S5: Latitudinal dependence of β (a,b) and η (c,d) from OceanSODA-ETHZ, shown as zonal means over the ROMS Pacific domain (a,c) and the global ocean (b,d). Lines show the climatological monthly mean; shading indicates interannual variability (± 1 standard deviation). All statistics are computed over 1993-2022.

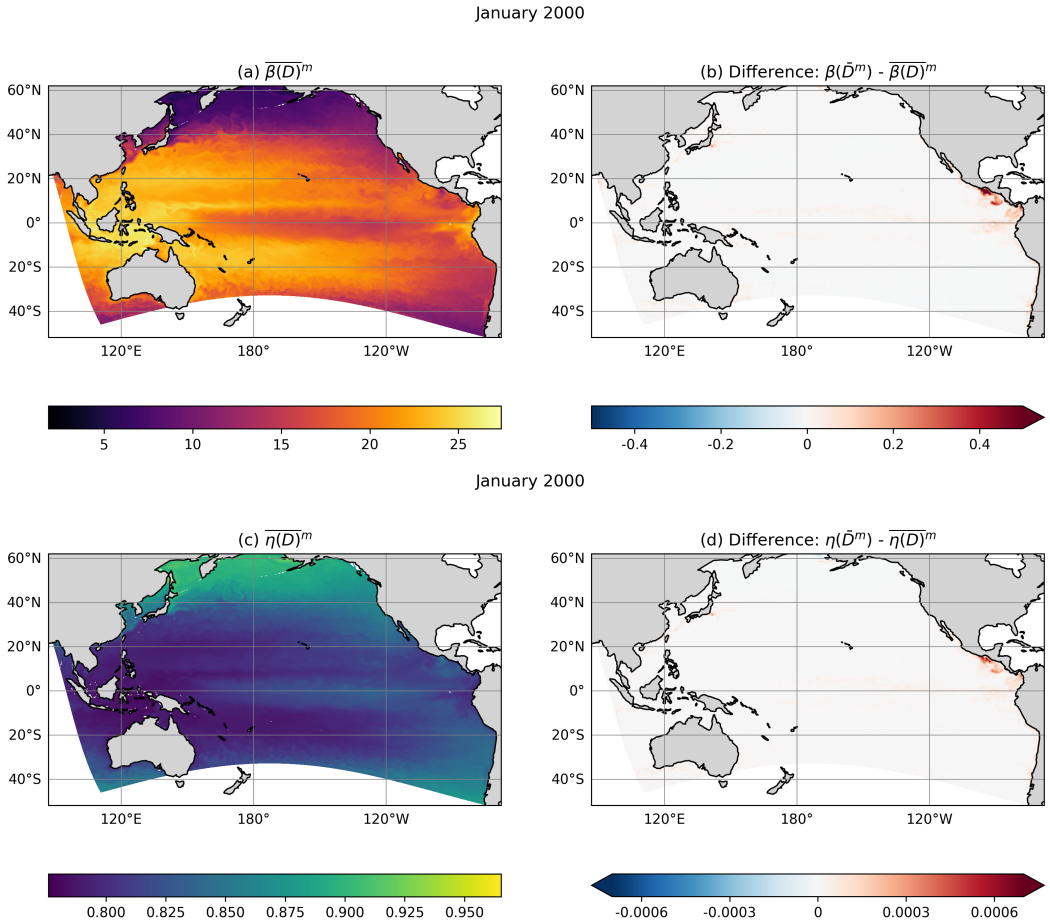


Figure S6: All panels show January 2000. (a), (c) Carbonate sensitivities β and η computed from daily state variables and subsequently averaged to monthly resolution, i.e., $\overline{\beta(\mathbf{X})}^m$ and $\overline{\eta(\mathbf{X})}^m$, where $\mathbf{X}$ denotes the daily state variables and $\overline{(\cdot)}^m$ denotes a monthly mean. (b), (d) Differences arising from reversing the order of operations, i.e., computing the sensitivities from monthly-averaged state variables: $\beta(\overline{\mathbf{X}}^m) - \overline{\beta(\mathbf{X})}^m$ and $\eta(\overline{\mathbf{X}}^m) - \overline{\eta(\mathbf{X})}^m$.

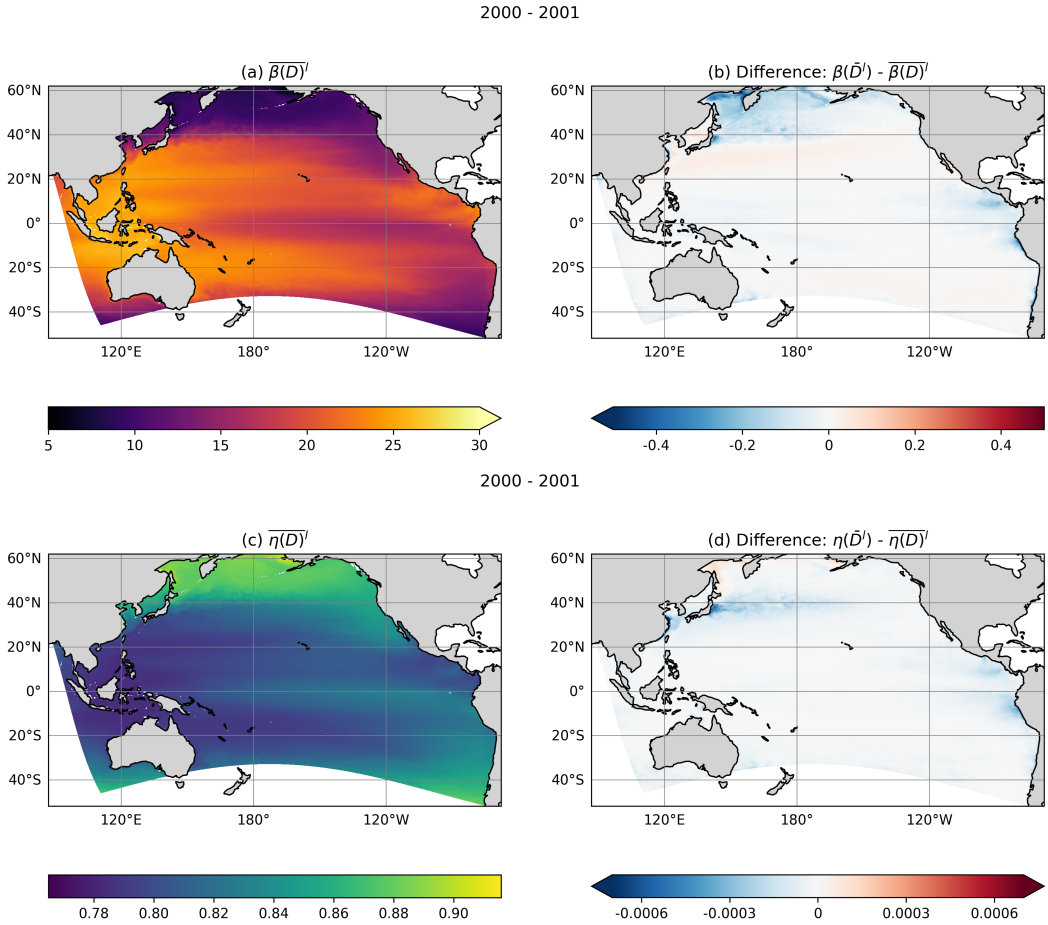


Figure S7: As in Figure S6, but for the two-year mean over 2000–2001 instead of the monthly mean for January 2000, i.e., $\overline{(\cdot)}^m$ is replaced by $\overline{(\cdot)}^l$, the long-term mean over the two-year period. Panels (a) and (c) show the two-year means, and panels (b) and (d) the corresponding differences.

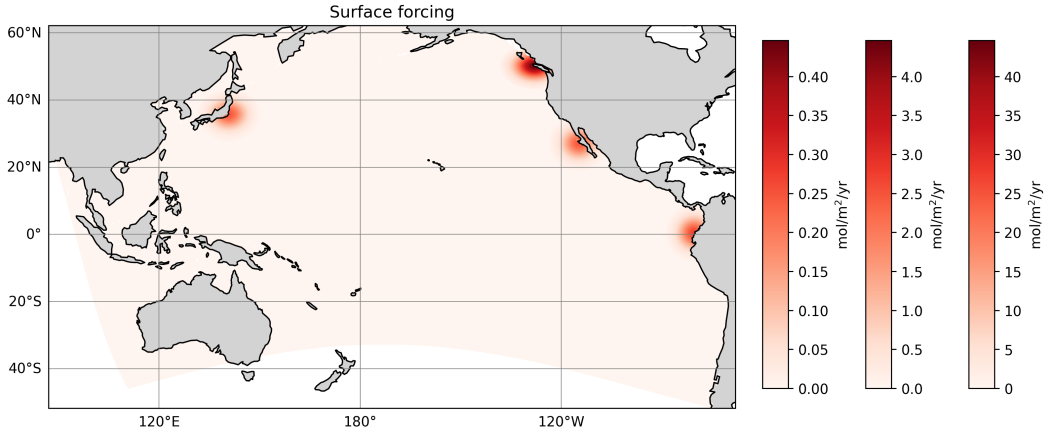


Figure S8: Surface forcing for the ROMS experiments at the four release locations (Table 1). The figure shows the sum of all four surface forcings, which do not overlap spatially. The three colorbars correspond to the three intervention amplitudes (experiment suffixes 10, 100, 1000). DOR and OAE experiments share the same spatial distribution, with the DIC flux in DOR assigned a negative sign. The forcing is applied over January 2000 as a Gaussian with a nominal standard deviation of 300 km, masked to ocean points and renormalized to preserve the total integrated magnitude. For the same total intervention mass, peak surface fluxes differ by up to a factor of 1.8 across locations (e.g. 44.6, 27.2, 24.8, and 24.6 $\mathrm{mol\,m^{-2}\,yr^{-1}}$ at Vancouver Island, Ecuador, Japan, and Baja California for the 1 Tmol experiments), owing to differences in the realized Gaussian width and in the ocean footprint after land masking and renormalization. These effects are further investigated in Figures S9, S10, and S11.

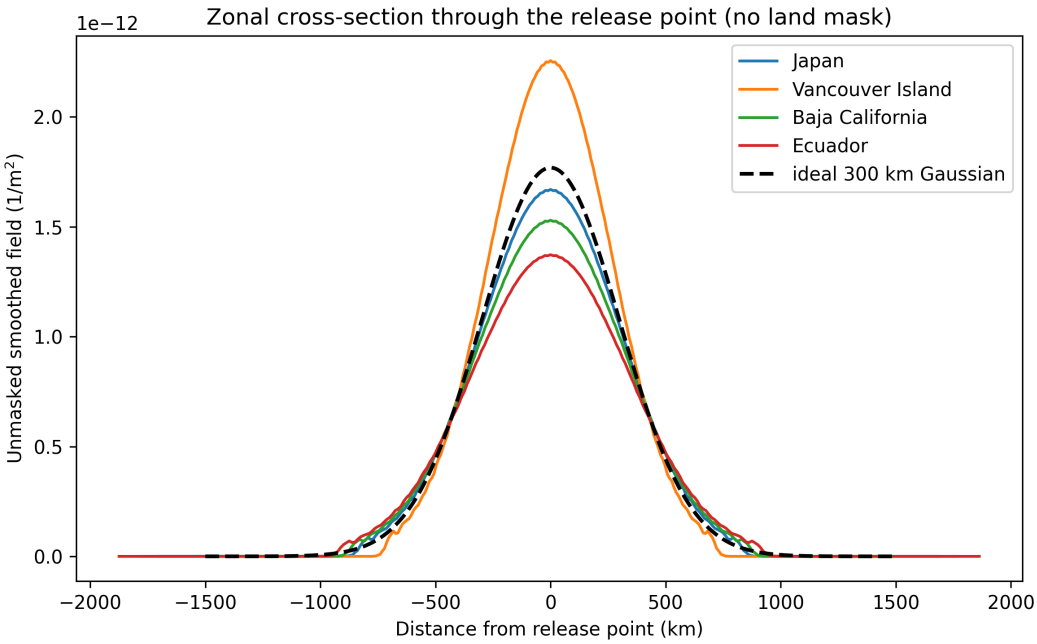


Figure S9: Zonal cross-section through the release point of the unmasked Gaussian forcing field, shown for each of the four release locations alongside an ideal 300 km Gaussian (black dashed). The Gaussian filter is calibrated using a single domain-mean grid spacing, but the ROMS grid is only approximately spatially uniform, so physical cell size varies with location. A fixed filter scale in grid cells therefore corresponds to different physical widths at different locations, explaining the location-dependent peak amplitudes and widths seen here. The land mask has not yet been applied, isolating the effect of the non-uniform grid spacing from the additional effect of land masking and renormalization.

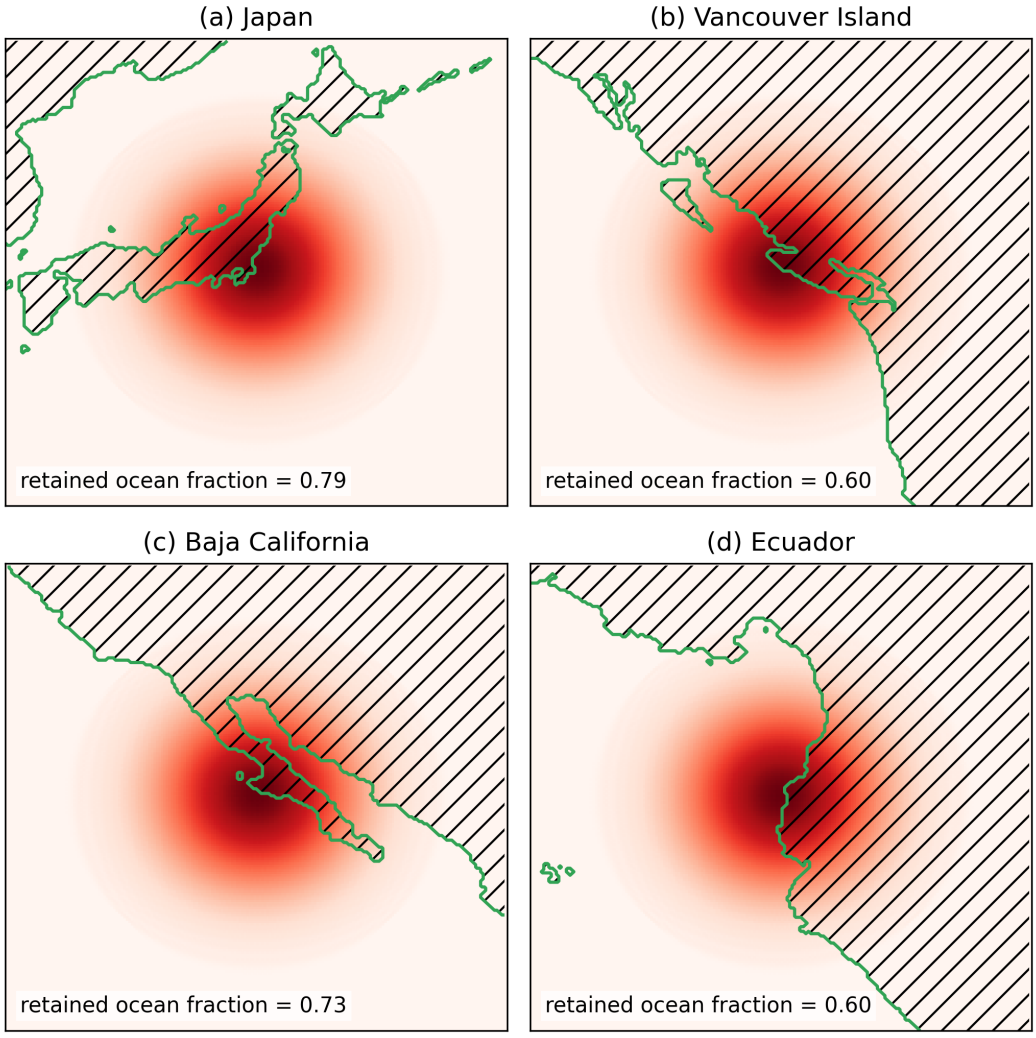


Figure S10: Unmasked Gaussian forcing field (shading) and land mask (hatching) for each of the four release locations: Japan (a), Vancouver Island (b), Baja California (c), and Ecuador (d). The fraction of the total Gaussian mass retained over ocean points after land masking is annotated in each panel (retained ocean fraction). This fraction varies across locations depending on how much of the Gaussian footprint overlaps with land; the masked flux is redistributed uniformly over the remaining ocean points to preserve the total intervention mass.

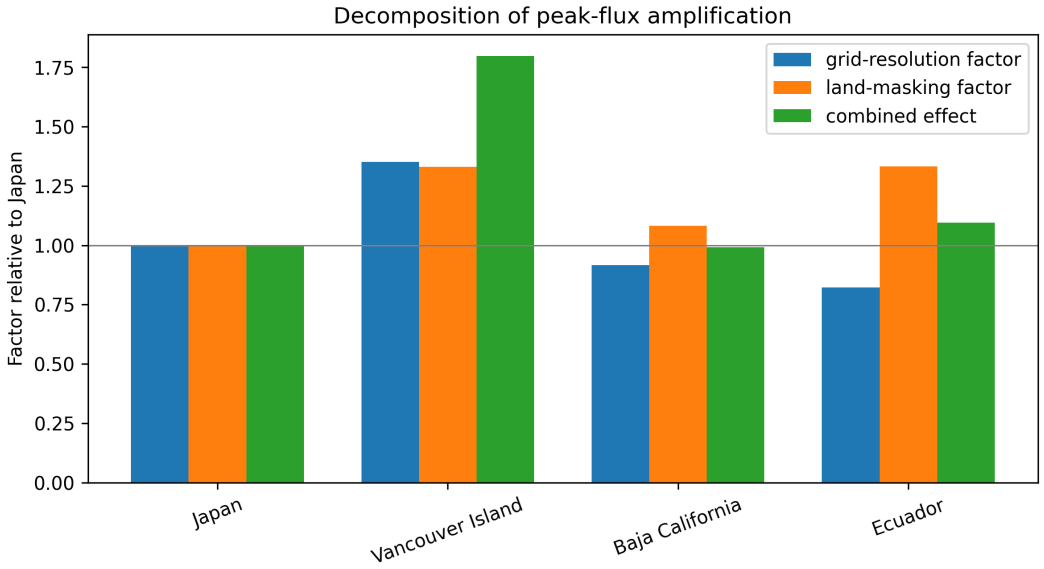


Figure S11: Decomposition of peak surface flux amplification relative to Japan across the four release locations. The grid-resolution factor (blue) reflects the location-dependent physical width of the Gaussian (Figure S9). The land-masking factor (orange) reflects the amplification of the peak flux by $1/f$, where f is the retained ocean fraction after land masking (Figure S10), since the total intervention mass is redistributed over the remaining ocean points. The combined effect (green) is the product of the two factors and reproduces the reported peak fluxes in Figure S8. Vancouver Island has the largest combined amplification ($\sim$1.80$\times$), driven by approximately equal contributions from both factors ($\sim$1.35 each). Ecuador has a modest combined amplification ($\sim$1.10$\times$) despite a grid-resolution factor below 1 (0.82), because the land-masking factor (1.33) more than compensates. Baja California is nearly indistinguishable from Japan ($\sim$0.99$\times$), reflecting a near-cancellation between the two factors (0.92 and 1.08).

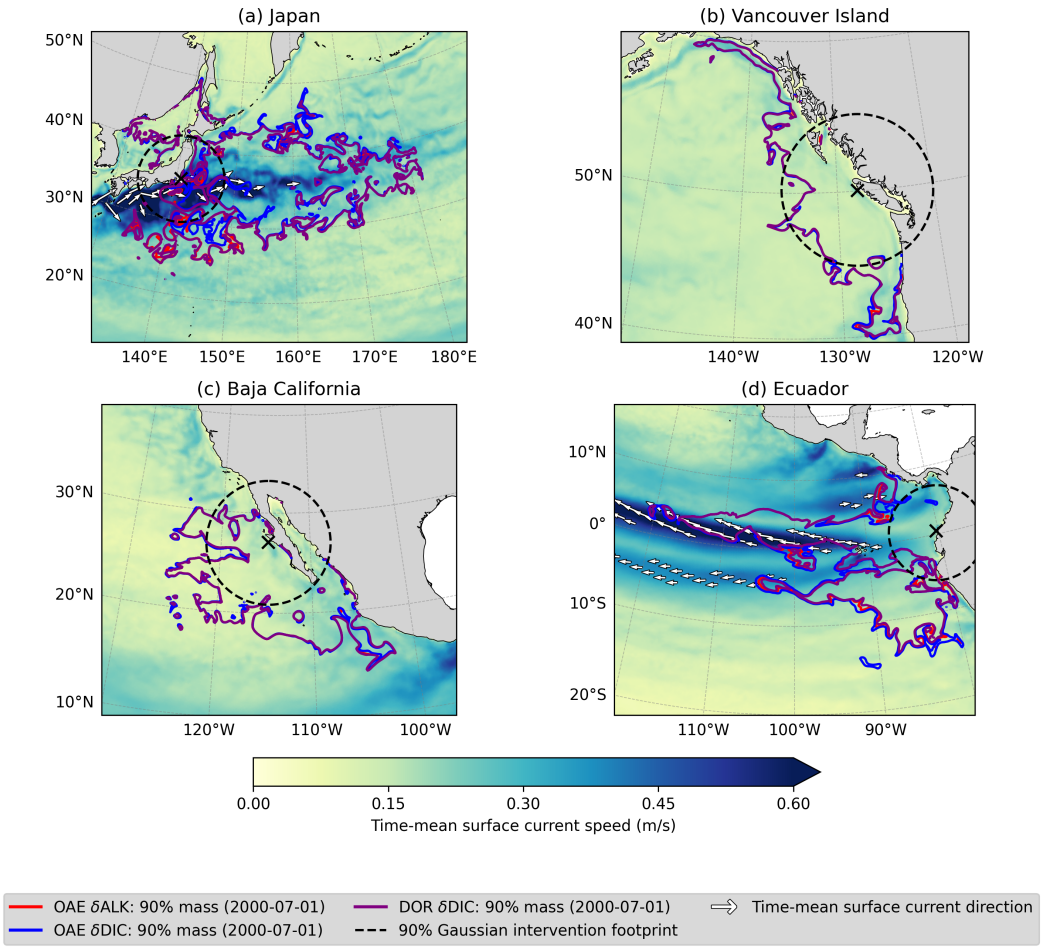


Figure S12: As Figure 4(a)–(d), but showing the 90% mass footprints of all three response fields: δALK from OAE (red), δDIC from OAE (blue), and δDIC from DOR (purple). The three contours are nearly coincident at all four locations, confirming that horizontal advection transports the DIC and alkalinity perturbations together. Minor differences between the contours, most visible at the Japan location, reflect the imprint of mesoscale eddy variability on the spreading of individual tracer patches.

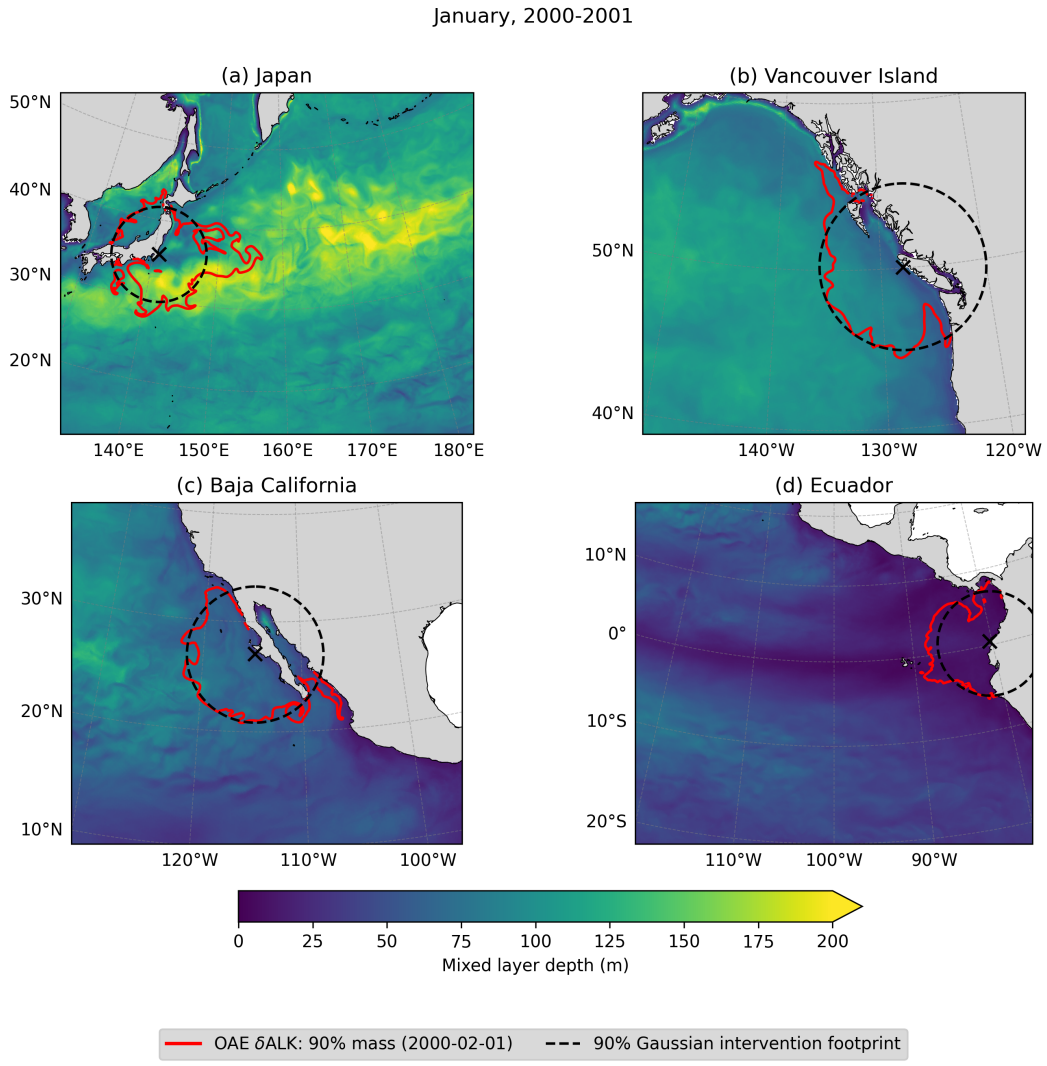


Figure S13: January mean mixed layer depth around each intervention location: Japan (a), Vancouver Island (b), Baja California (c), and Ecuador (d), averaged over 2000–2001. The mixed layer depth is defined by a density-threshold criterion as the depth at which the potential density anomaly first exceeds 0.03 kg m^{-3} relative to the surface value. Red contours show the 90% mass footprint of δALK from the largest-amplitude OAE experiment on February 1, 2000, one month after the intervention start, when the tracer footprint is still close to the intervention site. Black crosses and dashed circles mark the intervention center and 90% Gaussian footprint, respectively. The deep mixed layers in the Japan intervention region (up to 200–300 m) explain the pronounced subsurface penetration of the Japan tracer profiles (Figure 4(e)–(g)). In contrast, Ecuador has a very shallow January mixed layer of less than 25 m, consistent with its strongly surface-intensified vertical profiles.

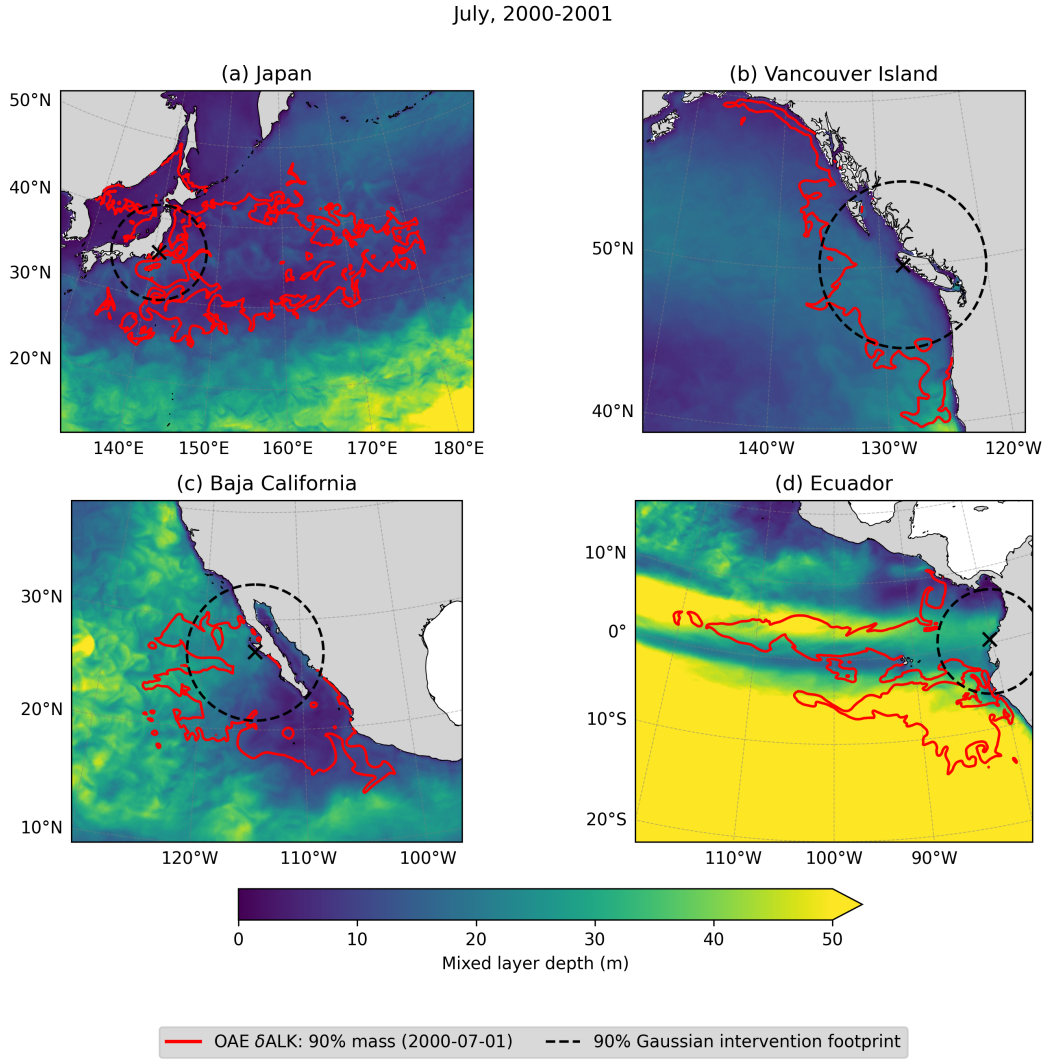


Figure S14: As Figure S13, but for July mean mixed layer depth, with the 90% mass footprint of δALK shown on July 1, 2000, six months after the intervention start. By this time the plume has spread considerably. The Japan mixed layer has shoaled dramatically relative to January, to less than 10 m, consistent with the seasonal restratification that traps tracer below the summer mixed layer (Figure 4(e)–(g)). Although the Ecuador intervention site itself has a shallow mixed layer year-round, the westward-spreading plume encounters mixed layer depths exceeding 50 m, explaining why the Ecuador vertical profile is less surface-intensified than expected from the intervention site alone. Note the different colorbar scale relative to Figure S13.

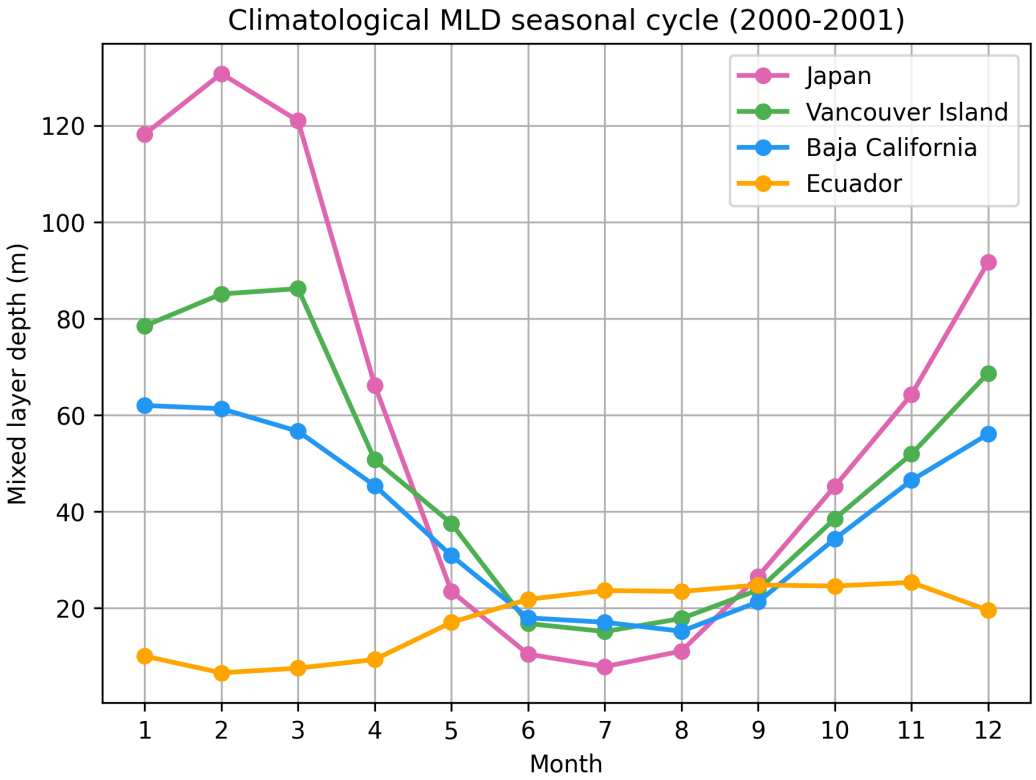


Figure S15: Seasonal cycle of the area-mean mixed layer depth within the 90% Gaussian footprint of each intervention location, averaged over 2000–2001, computed using the same density-threshold criterion as in Figures S13 and S14. Japan exhibits the largest seasonal amplitude, with a winter maximum of approximately 130 m in February and a summer minimum of approximately 10 m in July–August, driving the seasonal restratification that traps tracer below the summer mixed layer (Figure 4(e)–(g)). Vancouver Island shows a similarly pronounced but shallower seasonal cycle, while Baja California has a more moderate amplitude. Ecuador has the weakest seasonal cycle, remaining shallow year-round at the intervention site; however, the westward-spreading plume encounters deeper mixed layers away from the release location (Figure S14(d)).

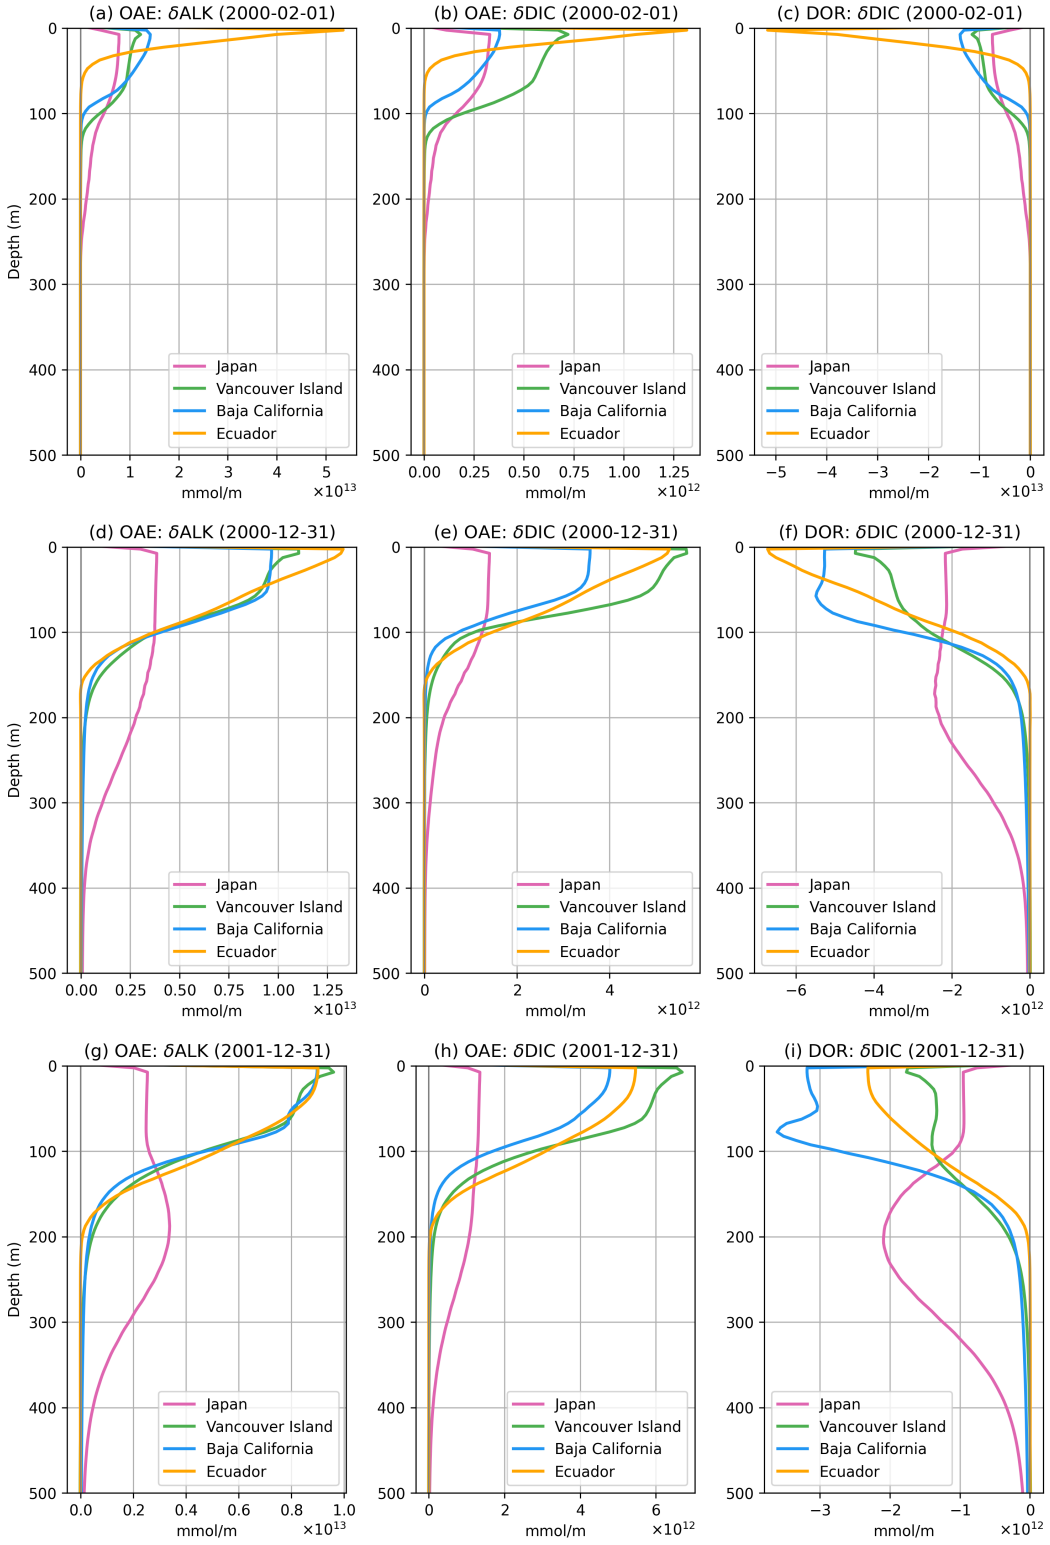


Figure S16: Vertical profiles of the response inventory density as in Figure 4(e)–(g), but shown at three additional times: one month after the intervention start (February 1, 2000; panels a–c), twelve months after the intervention start (December 31, 2000; panels d–f), and two years after the intervention start (December 31, 2001; panels g–i). One month after the intervention, all profiles show a vertically uniform signal down to the mixed layer depth, reflecting homogenization by winter convection. By December 2000, the Japan profile has developed a pronounced subsurface peak through summer restratification and subsequent winter deepening; the peak shifts further downward by December 2001 after a second annual cycle. The δDIC profiles in the OAE experiments (middle column) are more surface-intensified than the left and right columns, reflecting the imprint of air-sea gas exchange.

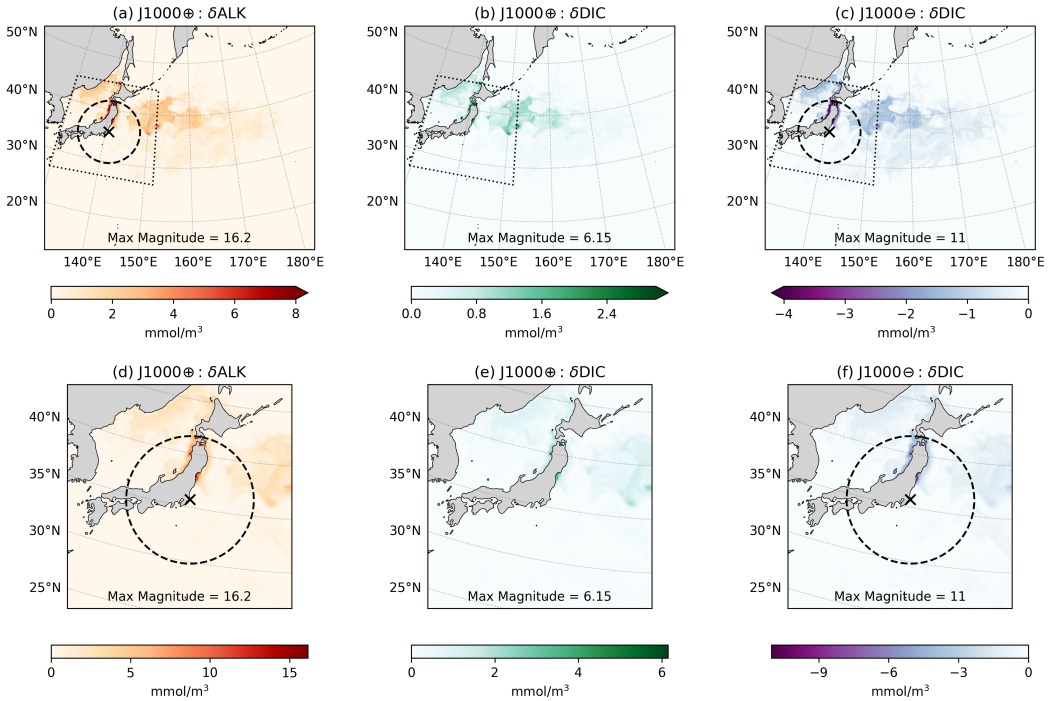


Figure S17: Surface alkalinity and DIC anomalies for the Japan OAE (J1000⊕) and DOR (J1000⊖) experiments on July 1, 2000, six months after the intervention start. Top row (a–c): wide-domain view with saturated colorbars; the dotted box indicates the zoomed region shown in the bottom row (d–f). Bottom row (d–f): zoomed view with unsaturated colorbars. Left column (a,d): δALK for J1000⊕. Middle column (b,e): δDIC for J1000⊕. Right column (c,f): δDIC for J1000⊖. The domain-wide absolute maximum is annotated in each panel. Black crosses and dashed circles mark the intervention center and 90% Gaussian footprint.

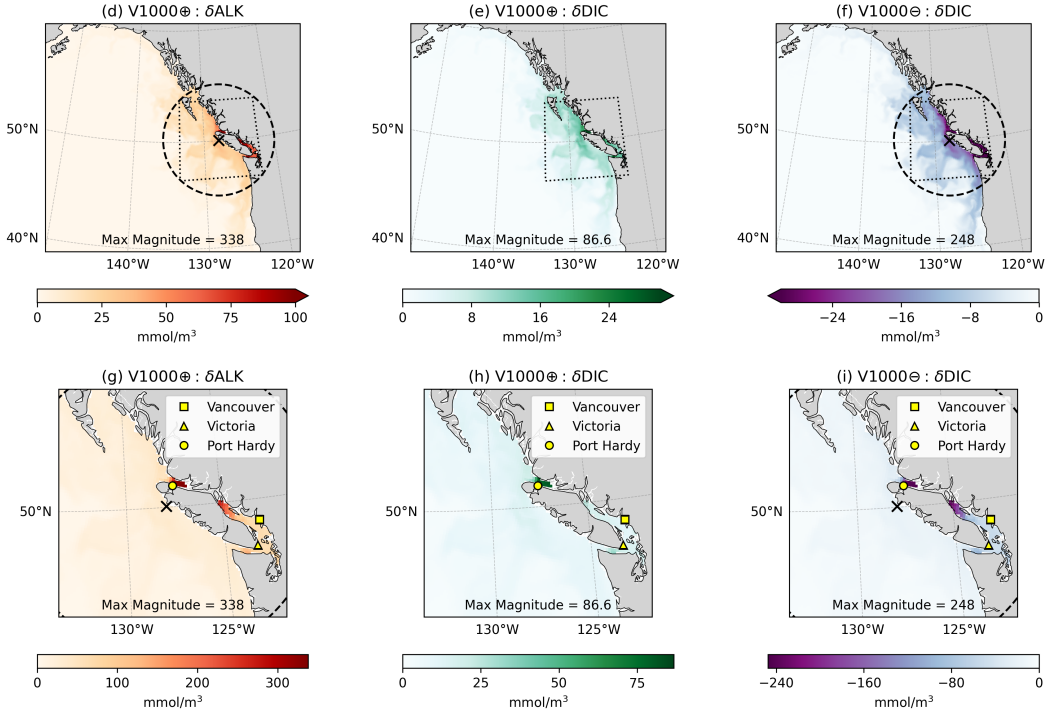


Figure S18: As Figure S17, but for the Vancouver Island location.

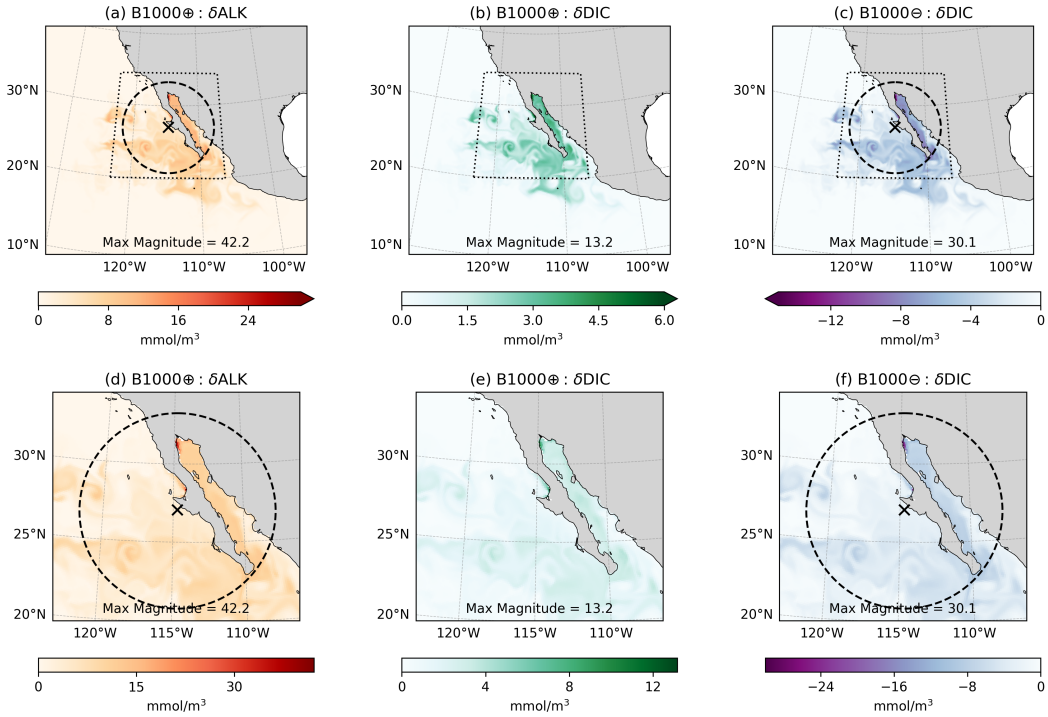


Figure S19: As Figure S17, but for the Baja California location.

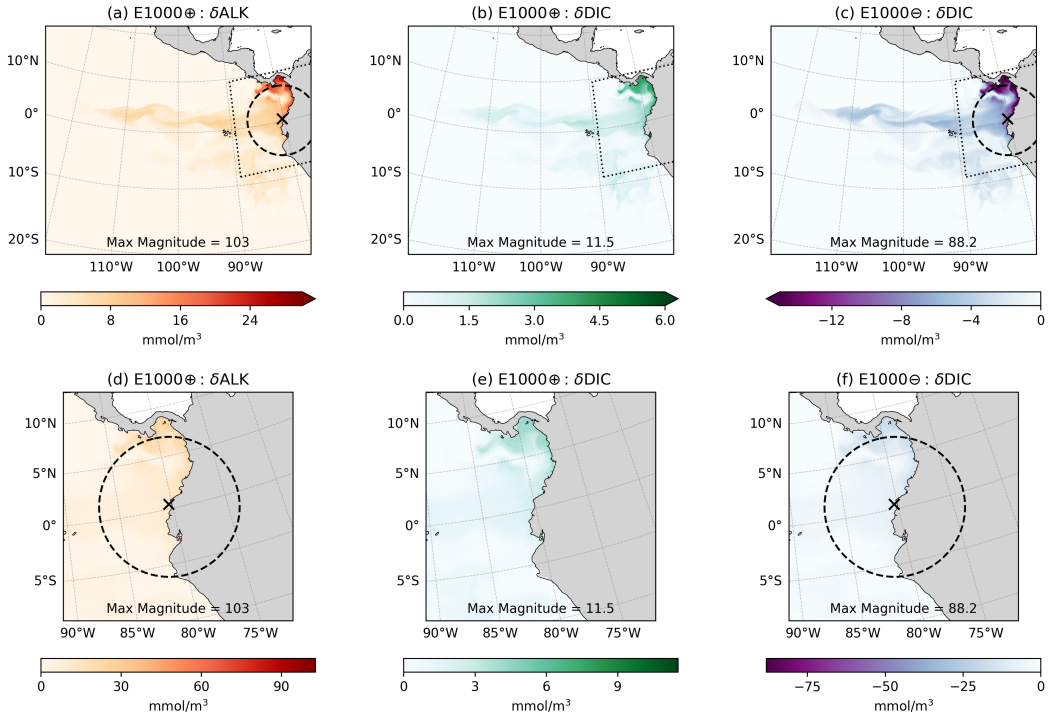


Figure S20: As Figure S17, but for the Ecuador location.

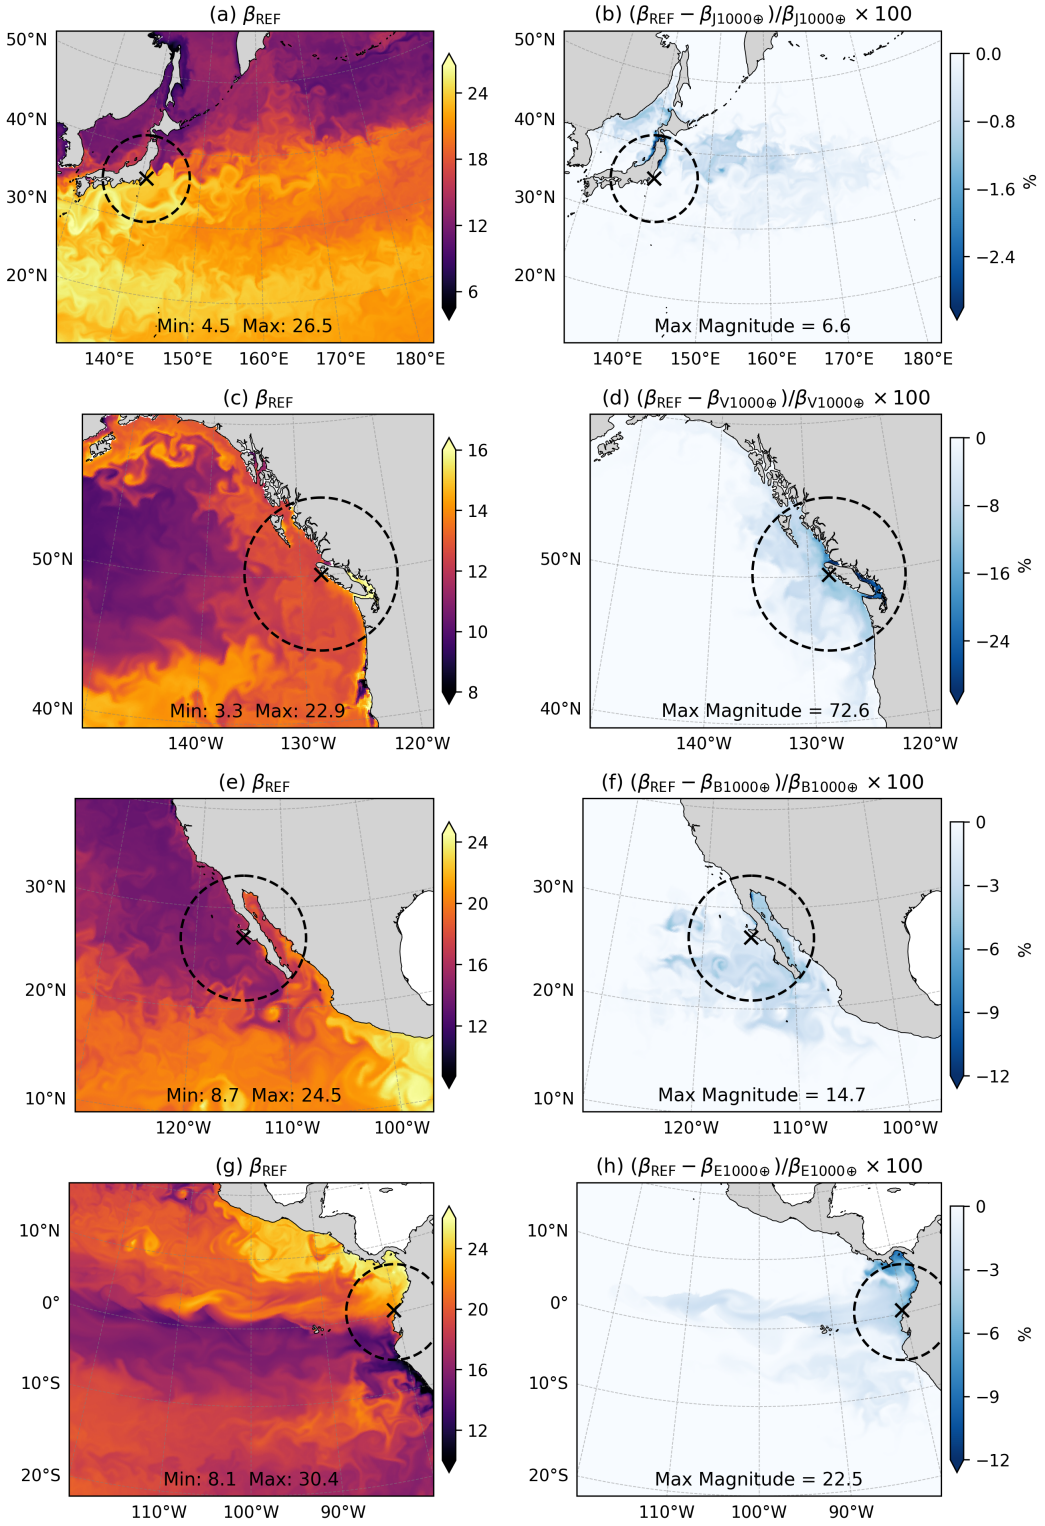


Figure S21: Spatial structure of β_{REF} and its relative difference from the perturbed truth for the largest-amplitude OAE experiment at each intervention location, on July 1, 2000, six months after the intervention start. Rows correspond to Japan (a,b), Vancouver Island (c,d), Baja California (e,f), and Ecuador (g,h). Left column: β_{REF}, shown to provide the background value against which the relative differences in the right column are computed; this is the same global field zoomed into each region, with annotated min and max values restricted to that region. Right column: relative difference $(\beta_{\mathrm{REF}} - \beta_{\mathrm{X1000}\oplus})/\beta_{\mathrm{X1000}\oplus} \times 100$, as in the first column of Figure 6. Black crosses and dashed circles mark the intervention center and 90% Gaussian footprint; the domain-wide absolute maximum magnitude is annotated in each panel.

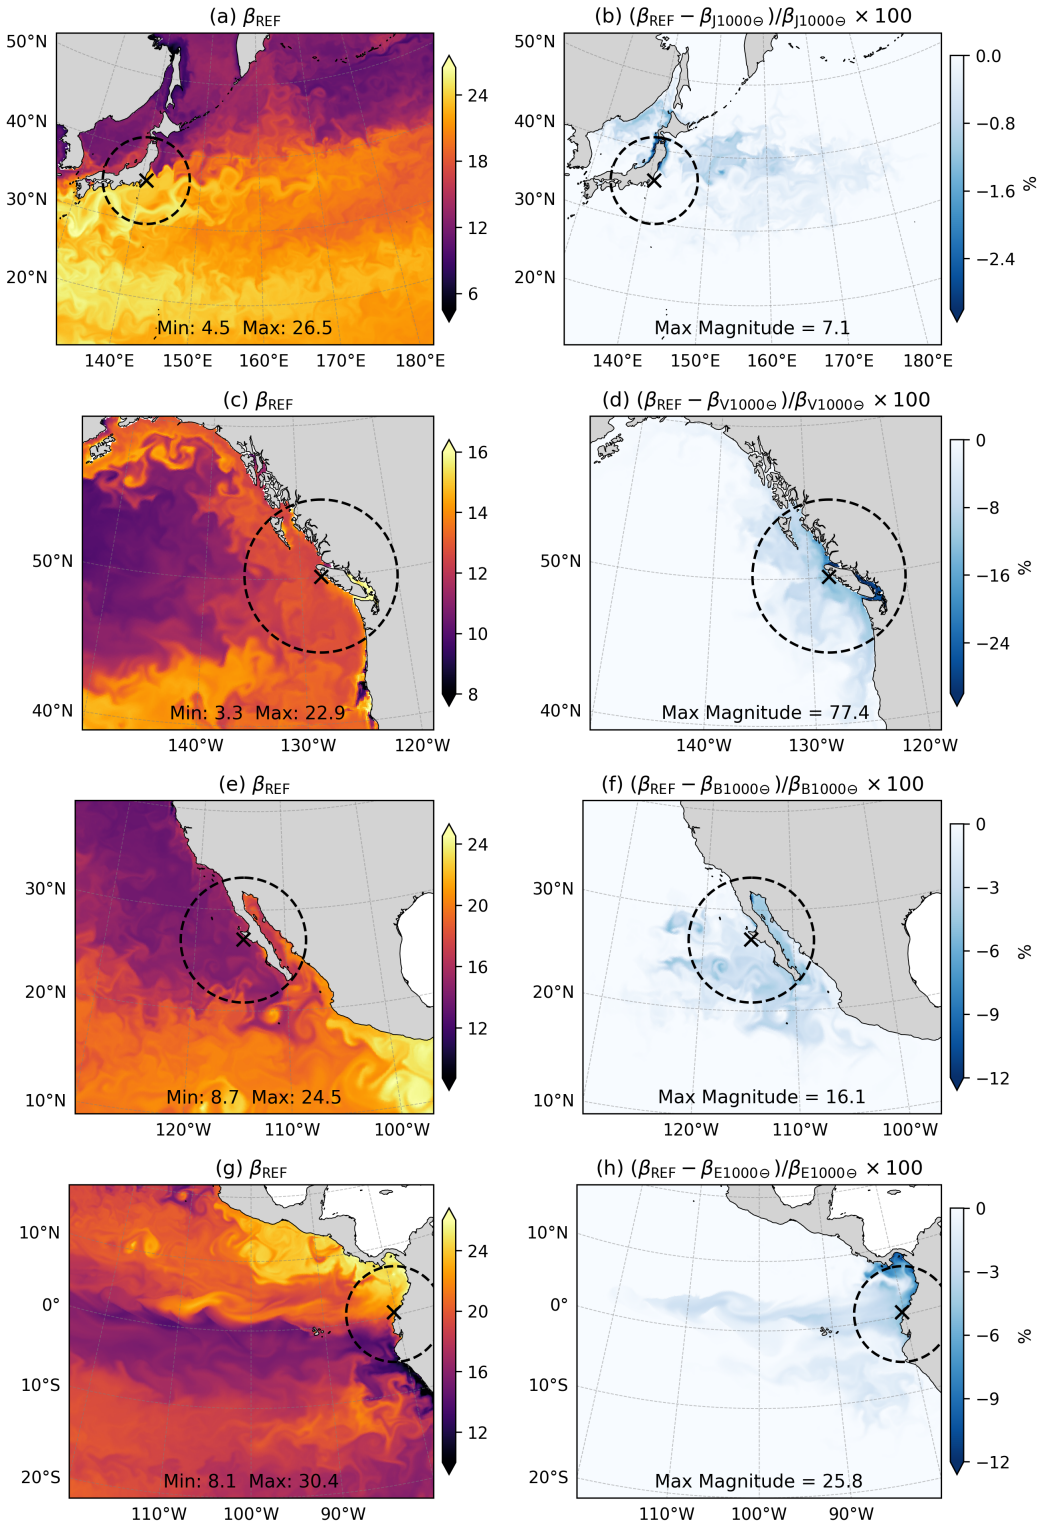


Figure S22: As Figure S21, but for the largest-amplitude DOR truth experiments.

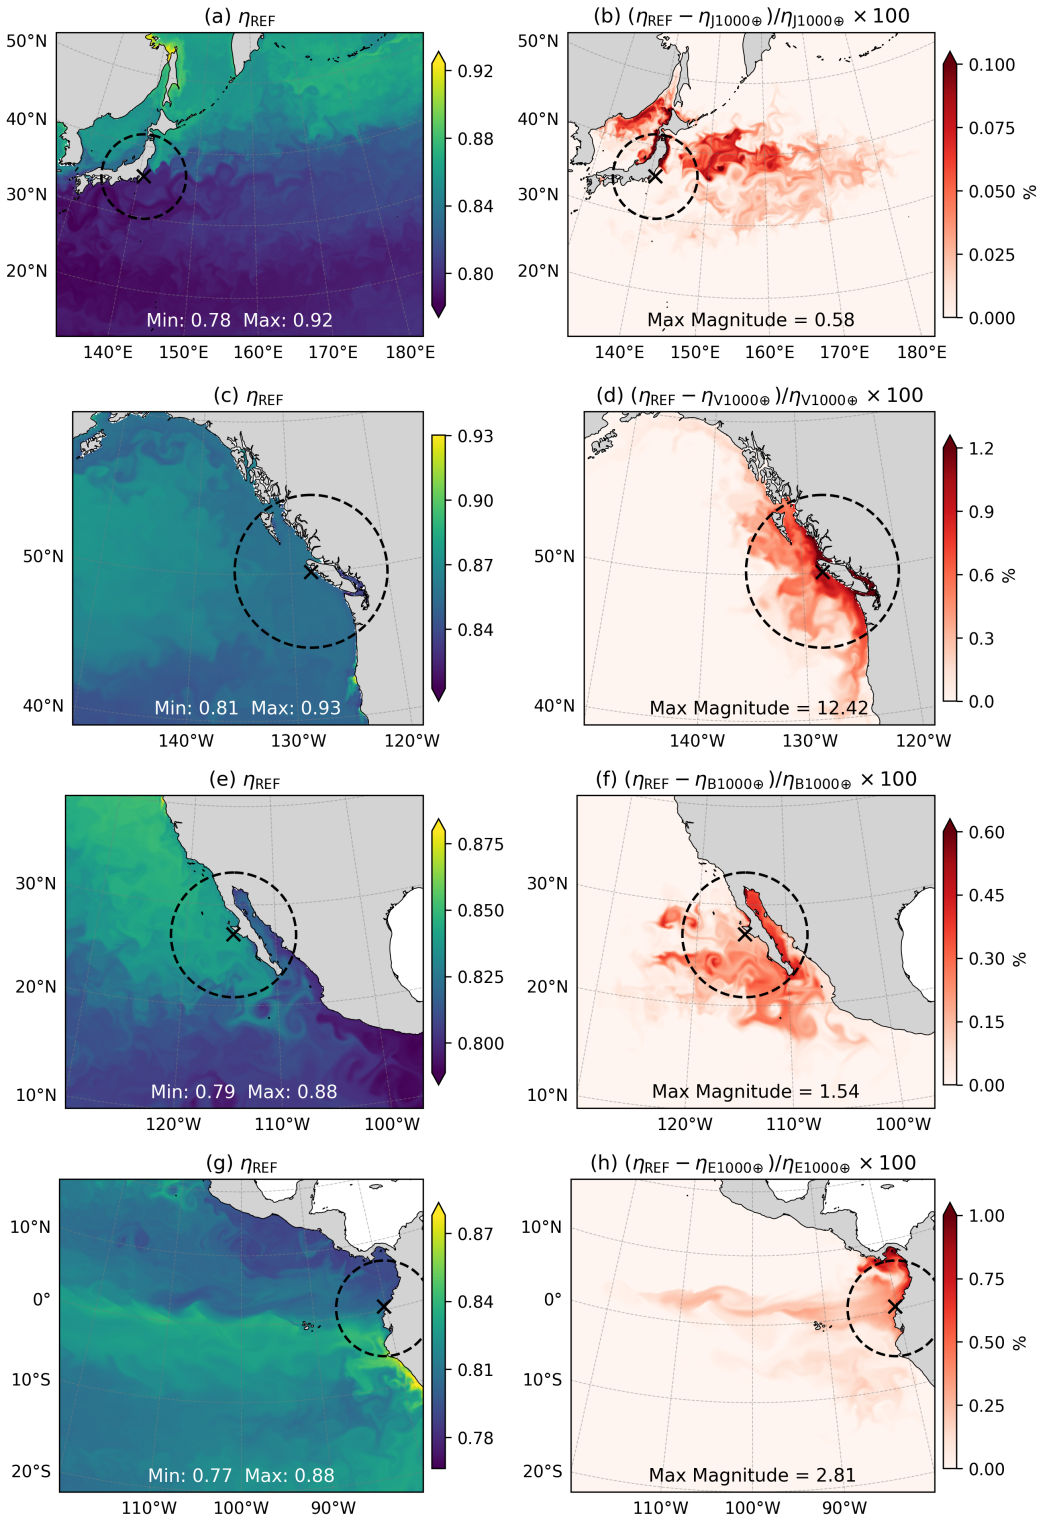


Figure S23: As Figure S21, but for η.

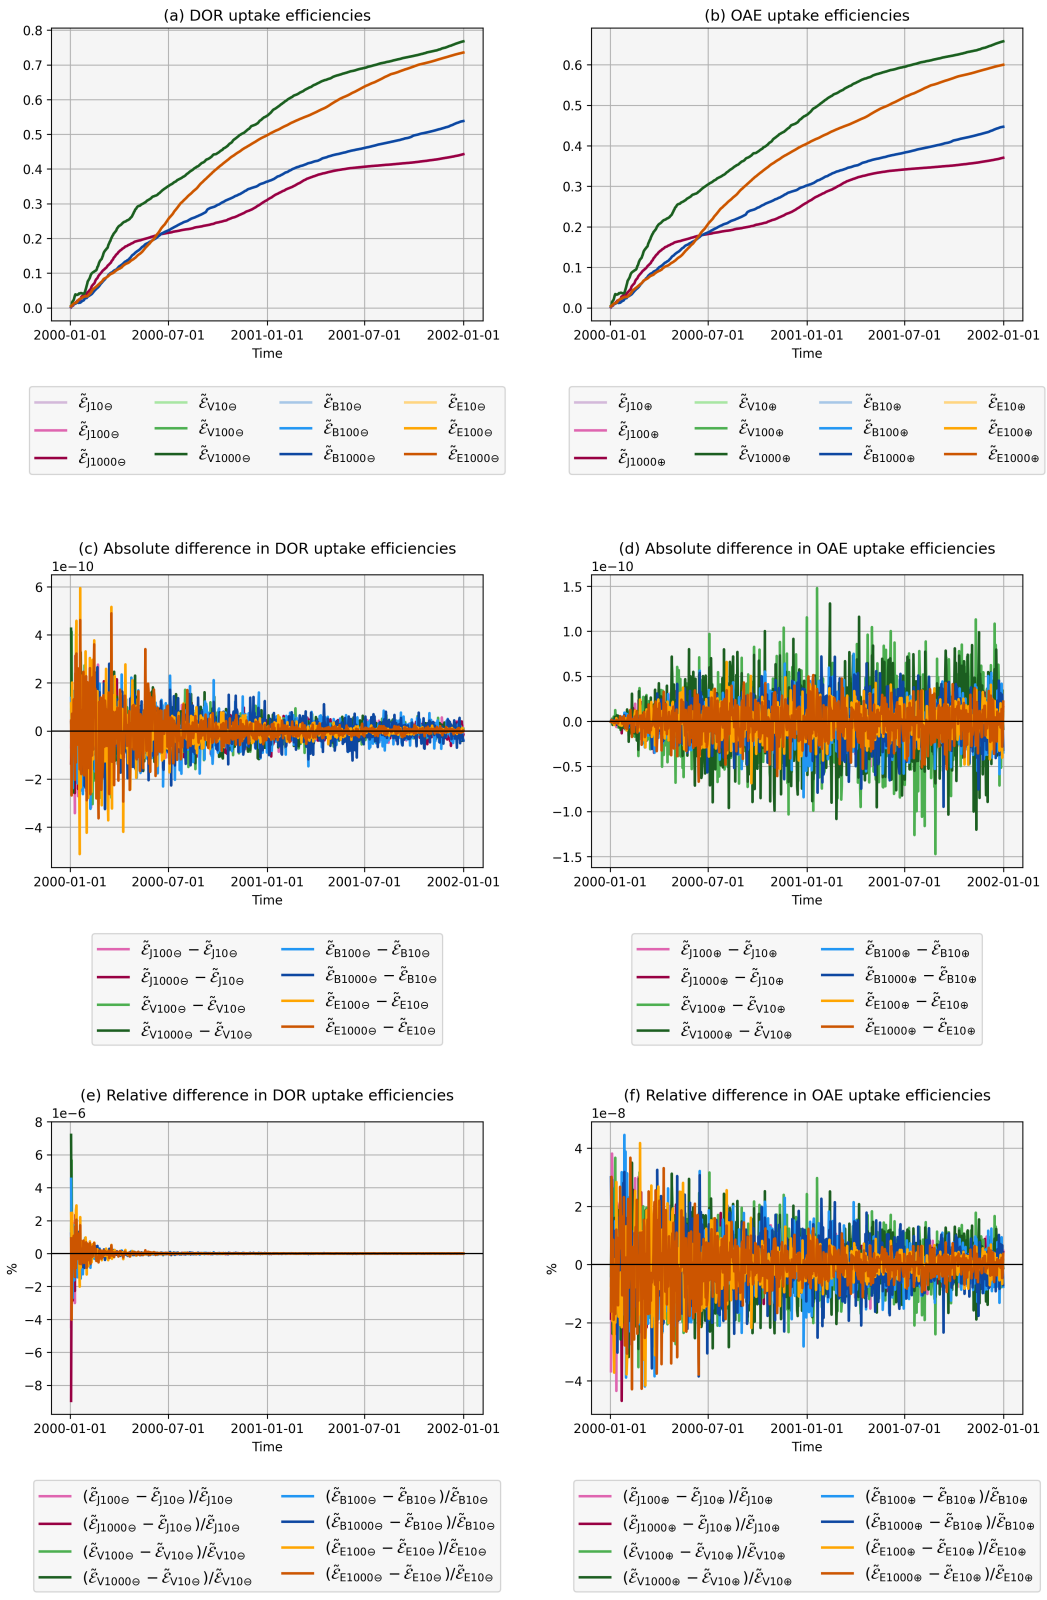


Figure S24: Linearity of the CDR tracer method, confirmed by the ROMS CDR tracer experiments. All experiments use β_{REF} and η_{REF} (Table 2). Left column: DOR experiments; right column: OAE experiments. Colors indicate release location (pink: Japan, green: Vancouver Island, blue: Baja California, orange: Ecuador); darker shades indicate larger release amplitudes. (a,b) CDR tracer efficiency curves $\tilde{\mathcal{E}}$ for DOR (a) and OAE (b), for all 12 ROMS CDR tracer experiments per intervention type. (c,d) Absolute difference between large-/mid- and small-amplitude CDR tracer efficiencies. (e,f) As (c,d) but relative differences (%).

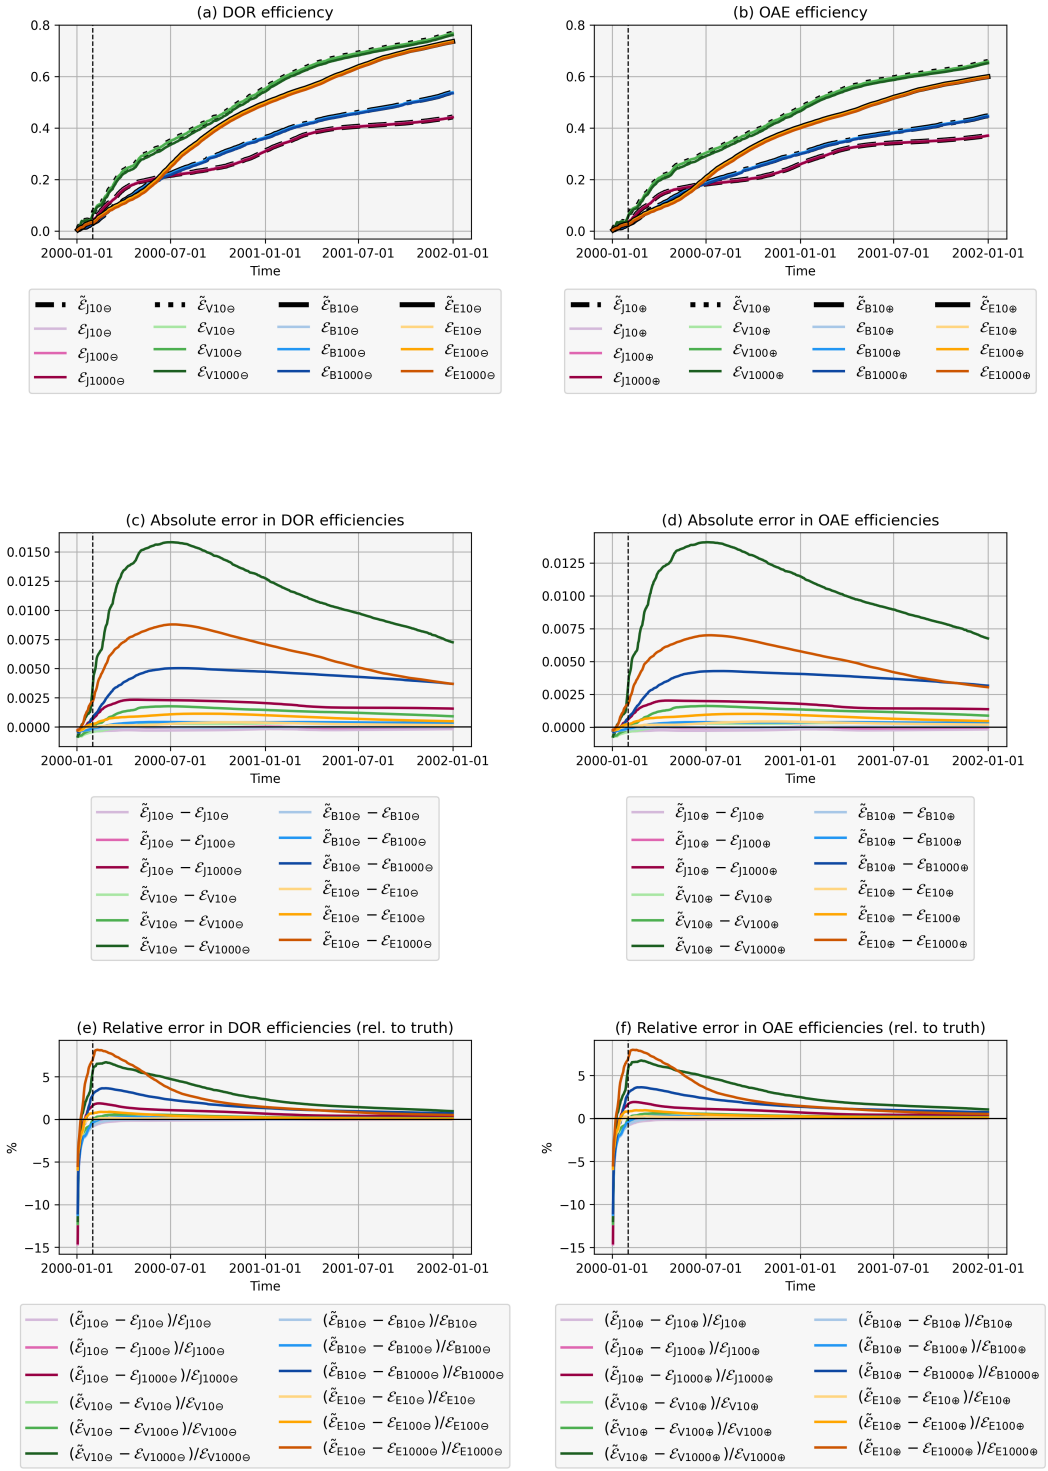


Figure S25: (a) Same as Figure 7(a), showing DOR efficiency curves $\mathcal{E}^{\mathrm{DOR}}$ and $\tilde{\mathcal{E}}^{\mathrm{DOR}}$. (b) Same as Figure 7(b), showing OAE efficiency curves $\mathcal{E}^{\mathrm{OAE}}$ and $\tilde{\mathcal{E}}^{\mathrm{OAE}}$. (c) Absolute error of the CDR tracer estimate of DOR efficiency, $\tilde{\mathcal{E}}^{\mathrm{DOR}} - \mathcal{E}^{\mathrm{DOR}}$. (d) Same as Figure 7(c), showing the absolute error of the CDR tracer estimate of OAE efficiency. (e) Relative error of the CDR tracer estimate of DOR efficiency, $(\tilde{\mathcal{E}}^{\mathrm{DOR}} - \mathcal{E}^{\mathrm{DOR}})/\mathcal{E}^{\mathrm{DOR}} \times 100$, expressed as a percentage. (f) Same as Figure 7(d), showing the relative error of the CDR estimate of OAE efficiency.

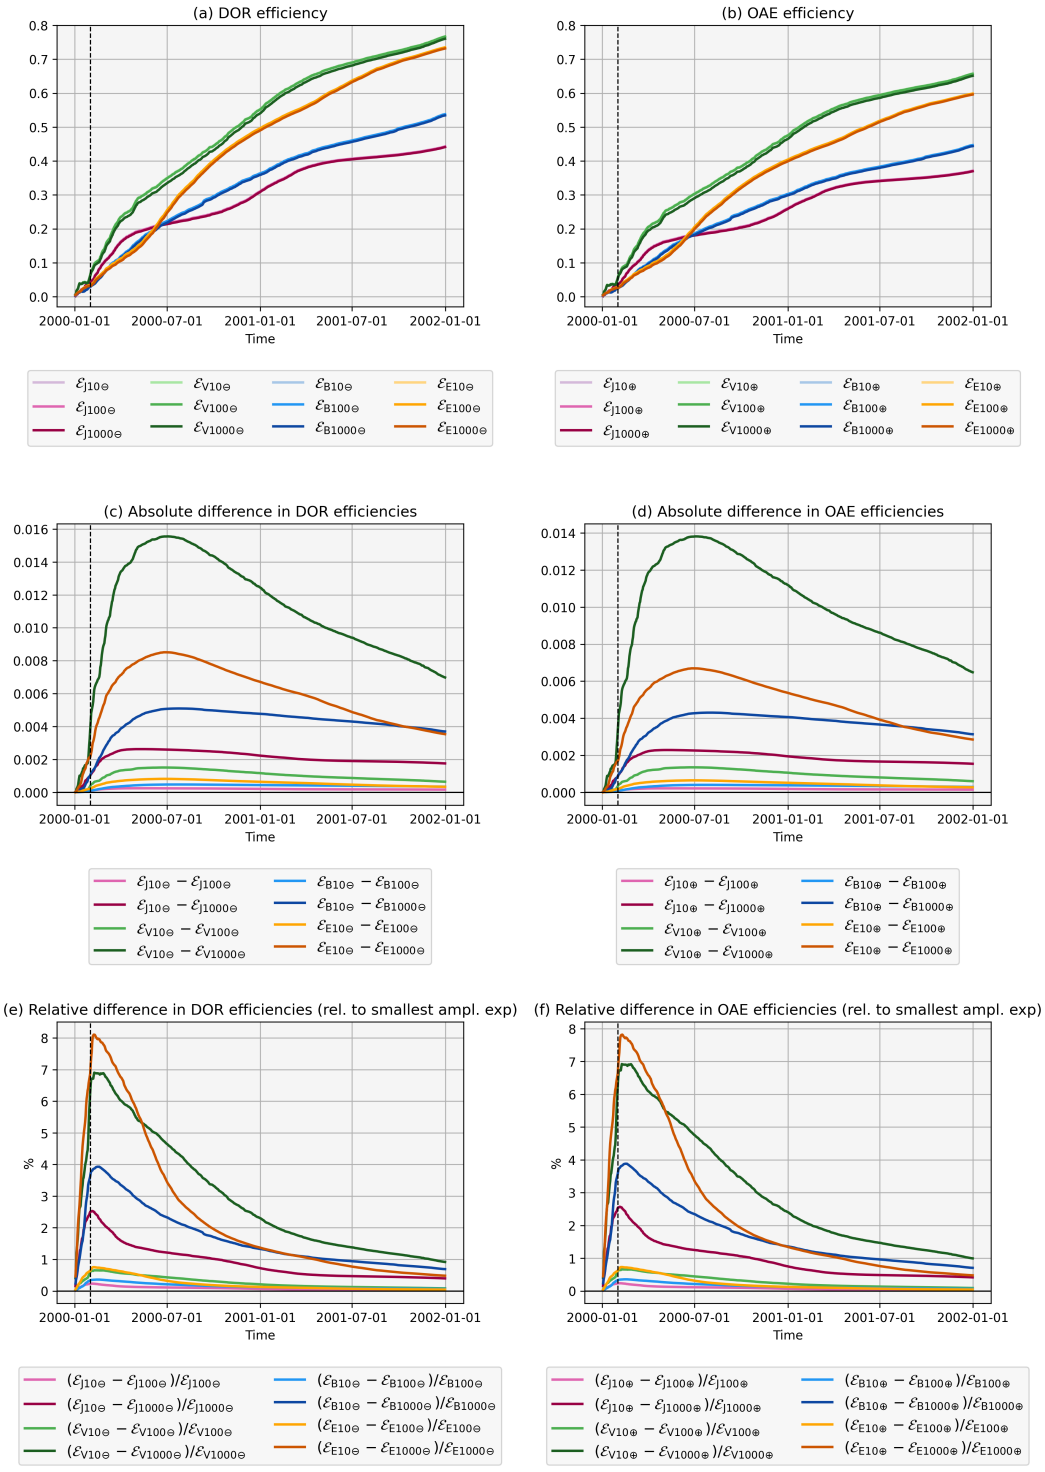


Figure S26: As Figure S25, but isolating system nonlinearity by comparing truth experiments of different amplitudes, rather than comparing CDR tracer estimates to truth experiments as in Figures 7 and S25. (a,b) DOR (a) and OAE (b) efficiency curves from the truth experiments only, as in Figure S25(a,b) but without the CDR tracer estimates. (c,d) Absolute difference between small- and large/mid-amplitude truth experiments, $\mathcal{E}^{\mathrm{DOR}}_{X10\ominus} - \mathcal{E}^{\mathrm{DOR}}_{XA\ominus}$ (c) and $\mathcal{E}^{\mathrm{OAE}}_{X10\oplus} - \mathcal{E}^{\mathrm{OAE}}_{XA\oplus}$ (d), for the intervention locations X $\in$ {J, V, B, E} and the amplitudes A $\in$ {100, 1000}. (e,f) As (c,d), but relative difference (%).

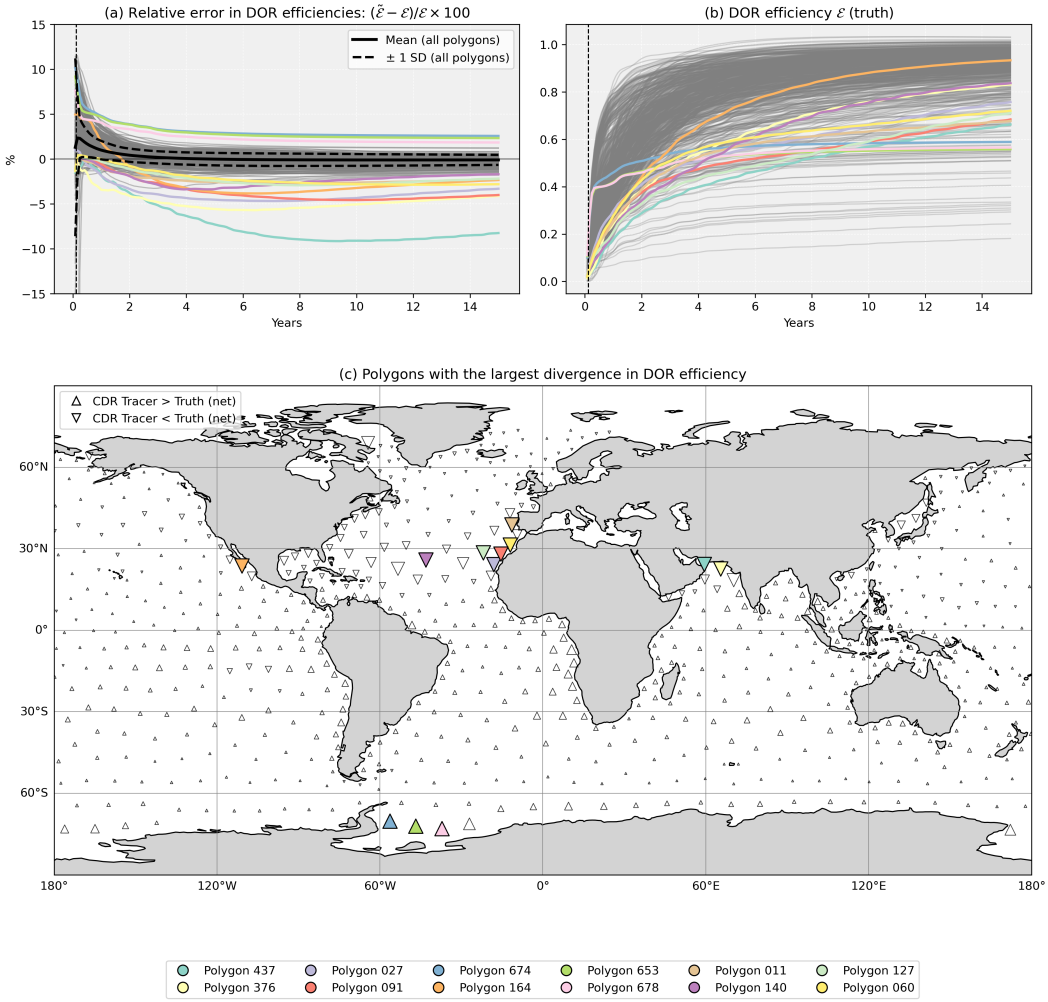


Figure S27: As Figure 11, but for the DOR experiments. The 12 highlighted polygons are selected by the same divergence criterion applied to the relative DOR error curves and slightly differ from those in Figure 11.

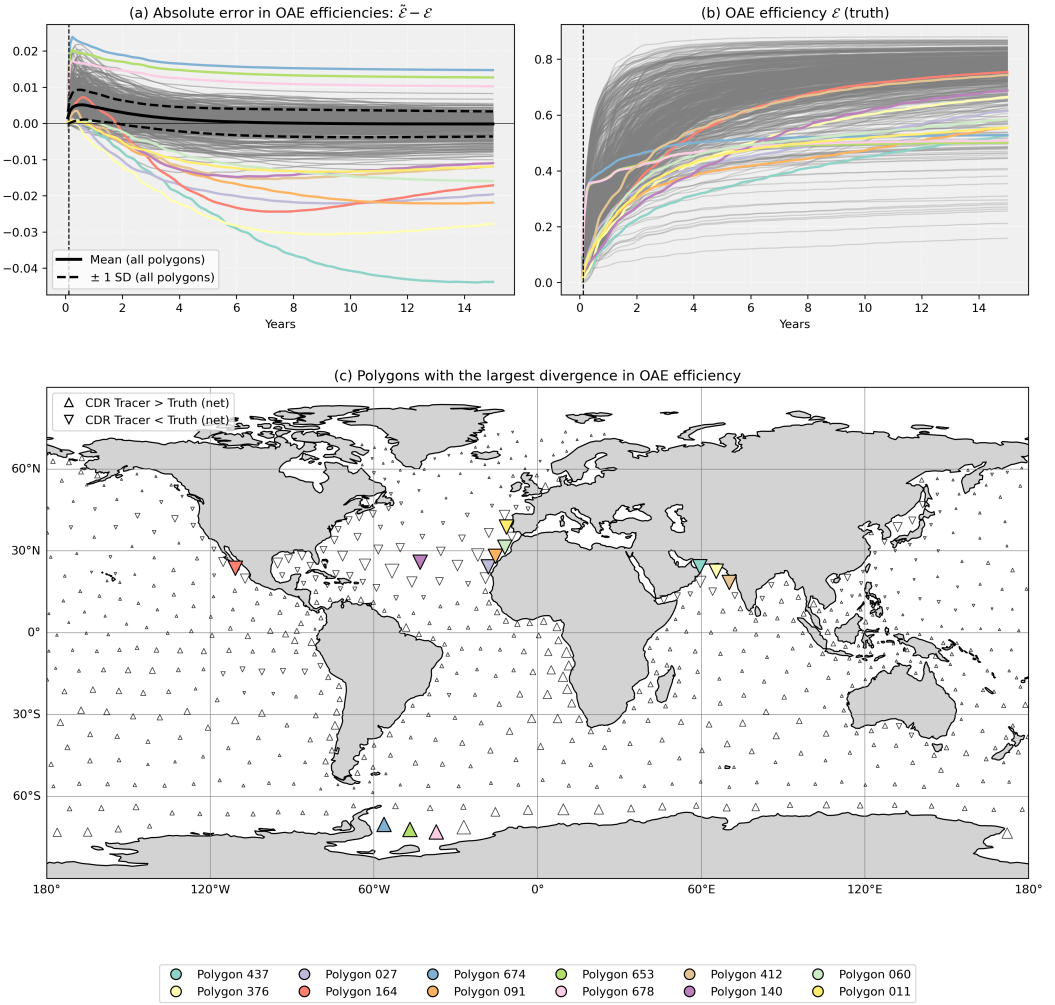


Figure S28: As Figure 11, but showing absolute OAE error curves $\tilde{\mathcal{E}}(t) - \mathcal{E}(t)$ in panel (a) instead of relative errors. The 12 highlighted polygons are selected by the same divergence criterion applied to the absolute OAE error curves ($R(t) = (\tilde{\mathcal{E}}(t) - \mathcal{E}(t))$) and slightly differ from those in Figure 11.

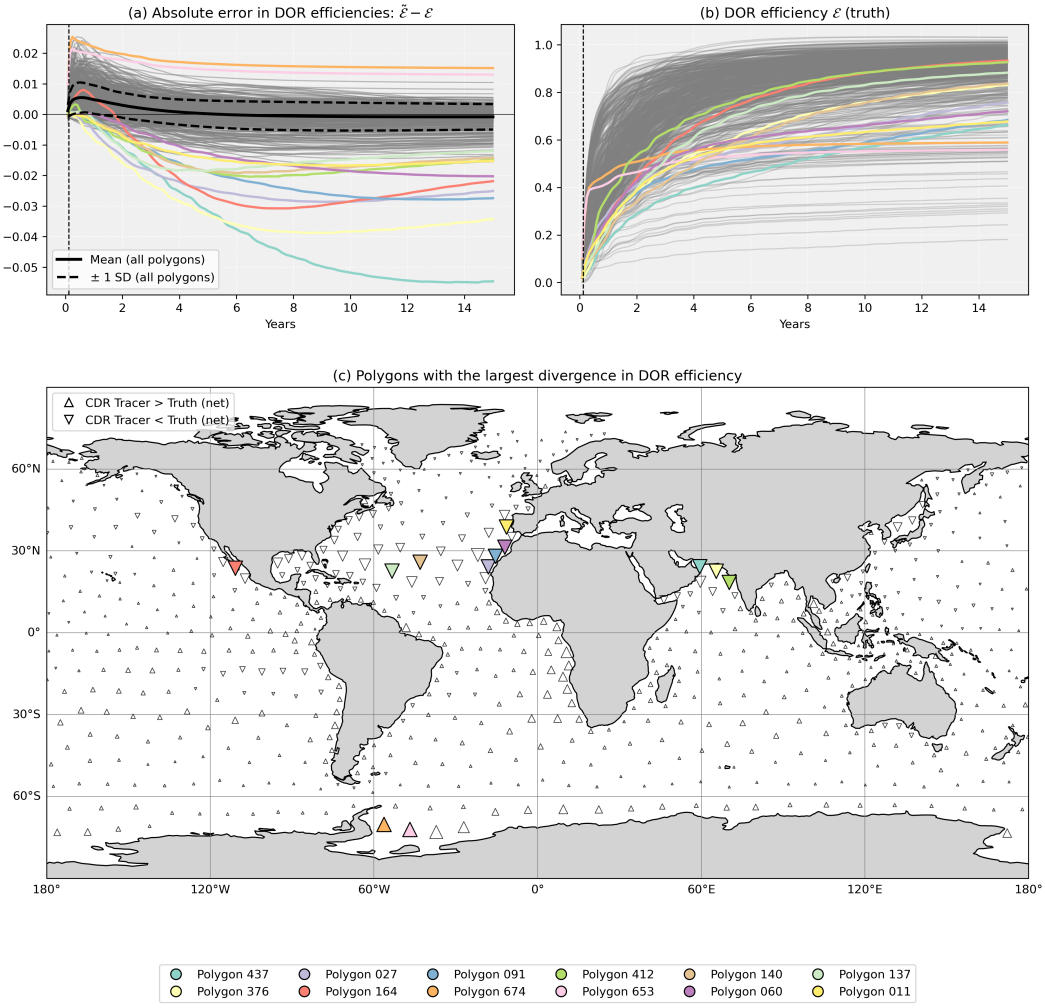


Figure S29: As Figure S28, but for the DOR experiments. The 12 highlighted polygons are selected by the same divergence criterion applied to the absolute DOR error curves.

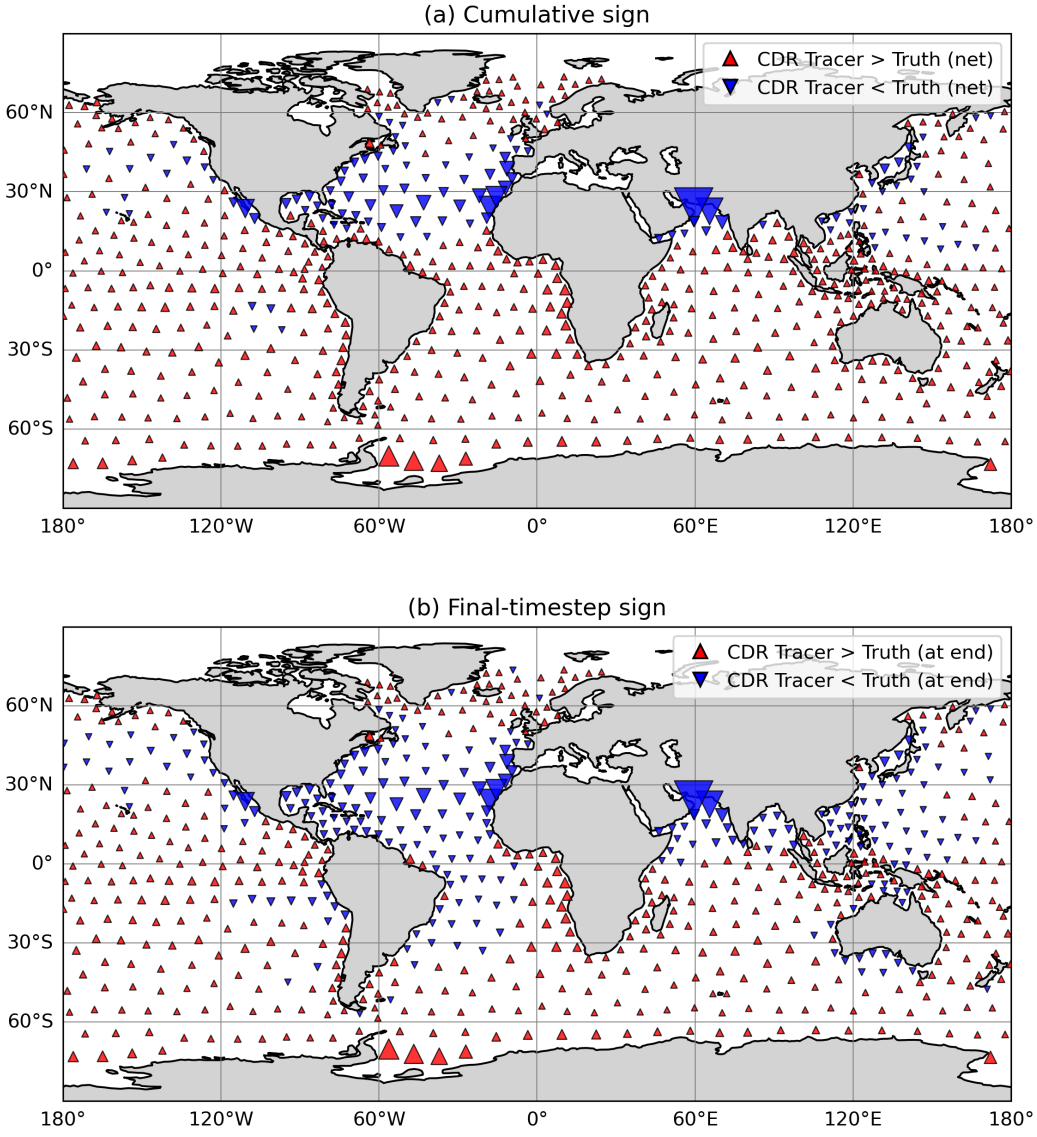


Figure S30: Sign of the relative CDR tracer OAE efficiency error for all 690 polygons, colored red (upward triangle: CDR tracer overestimates) and blue (downward triangle: CDR tracer underestimates). Triangle size is proportional to the divergence $\sum_t |R(t)|$ with $R(t) = (\tilde{\mathcal{E}}(t) - \mathcal{E}(t))/\mathcal{E}(t)$. (a) Sign based on the cumulative criterion $\sum_t R(t)$, as in Figure 11(c). (b) Sign based on the error at the final time step, $\tilde{\mathcal{E}}(t_{\text{end}}) - \mathcal{E}(t_{\text{end}})$. The two maps are very similar, confirming that the cumulative and final-time-step criteria are largely consistent.

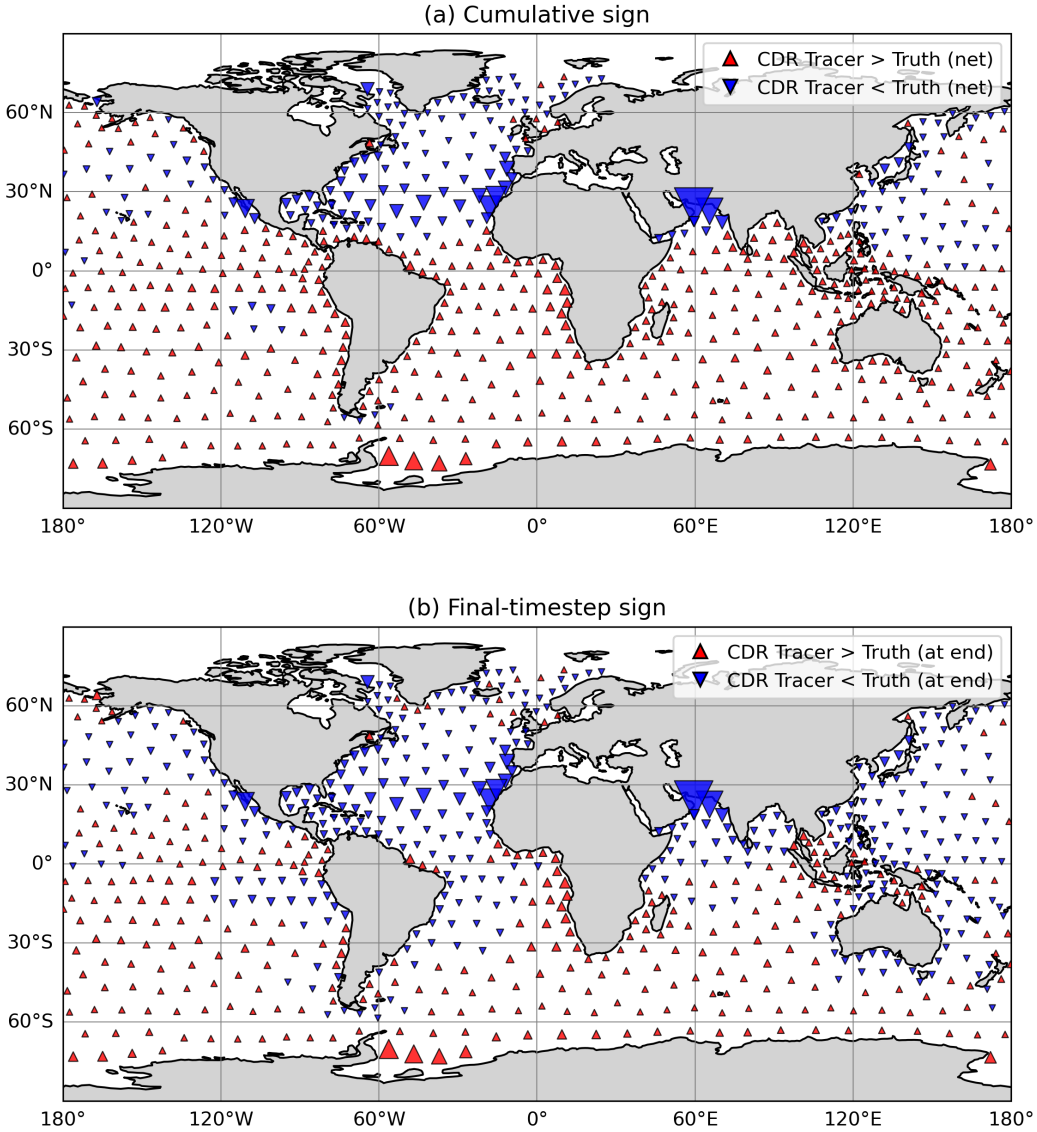


Figure S31: As Figure S30, but for the relative DOR efficiency error.

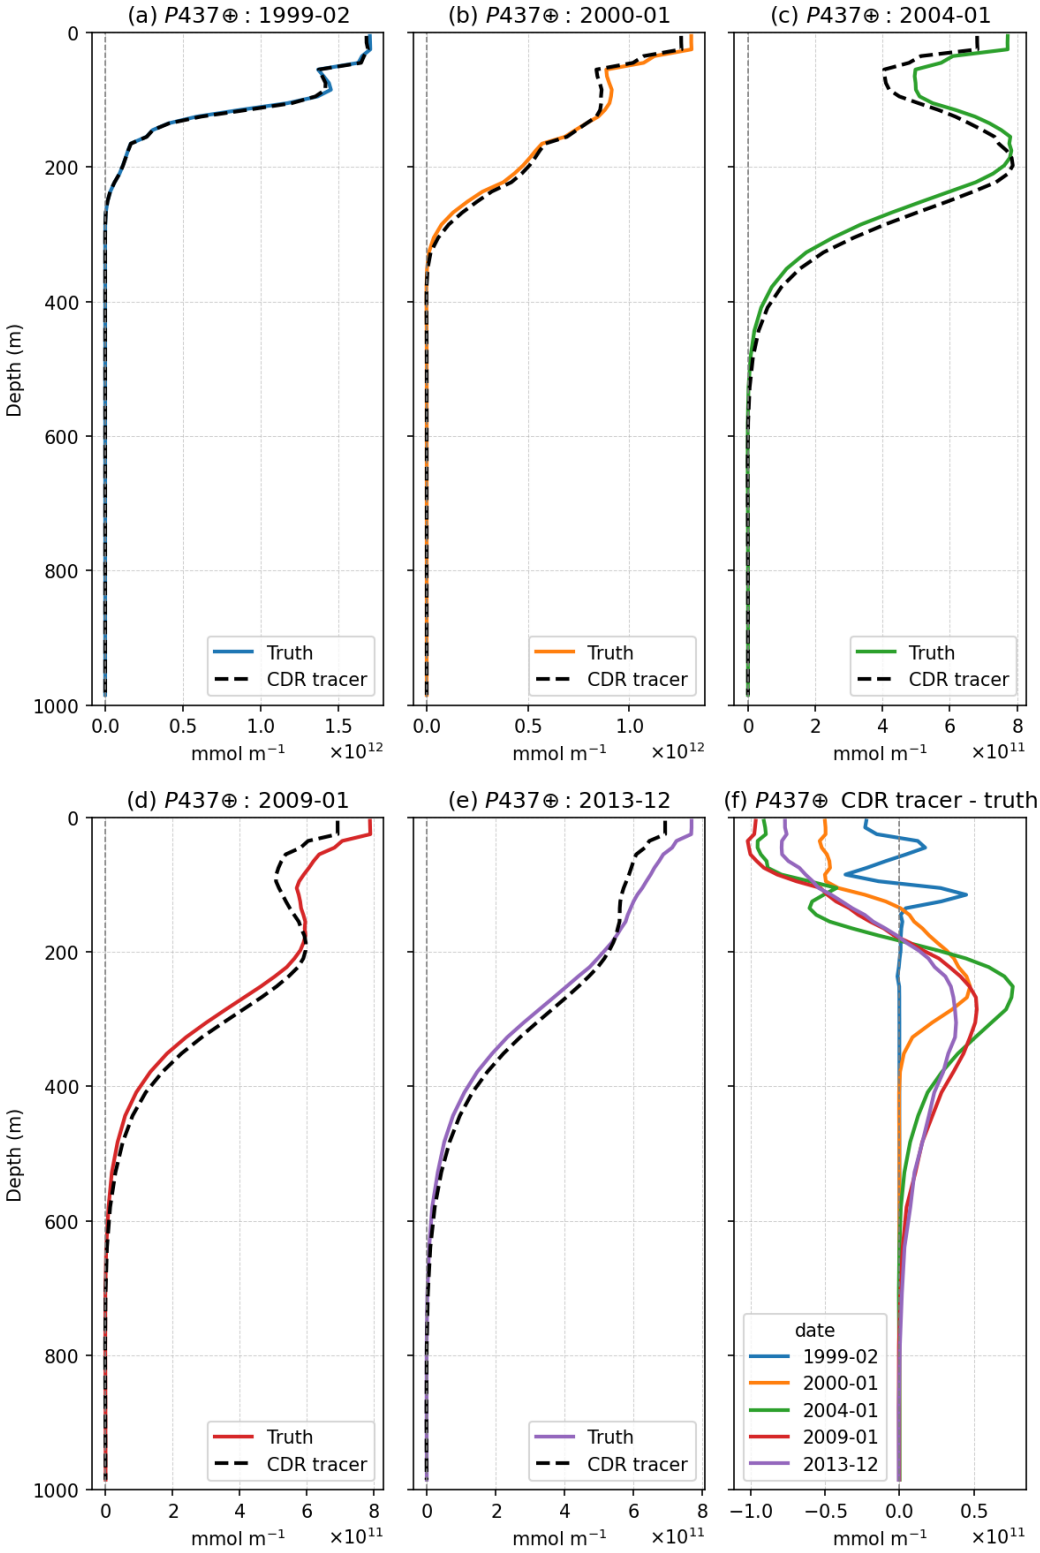


Figure S32: Vertical profiles of the horizontally integrated alkalinity anomaly $\sum_{i,j} \delta\mathrm{ALK}_{ij}\, A_{ij}$ (mmol m^{-1}) for the OAE experiment at the Arabian Sea polygon 437 (P437⊕), at five times spanning the 15-year period. Here, A_{ij} is the horizontal area of grid cell (i, j), and the sum is taken over all grid cells at a given depth level. Panels (a–e) show the truth (colored solid lines) and CDR tracer estimate (black dashed lines) in February 1999, January 2000, January 2004, January 2009, and December 2013, respectively. (f) Difference between CDR tracer estimate and truth (CDR tracer minus truth) for all five times, shown on a common axis.

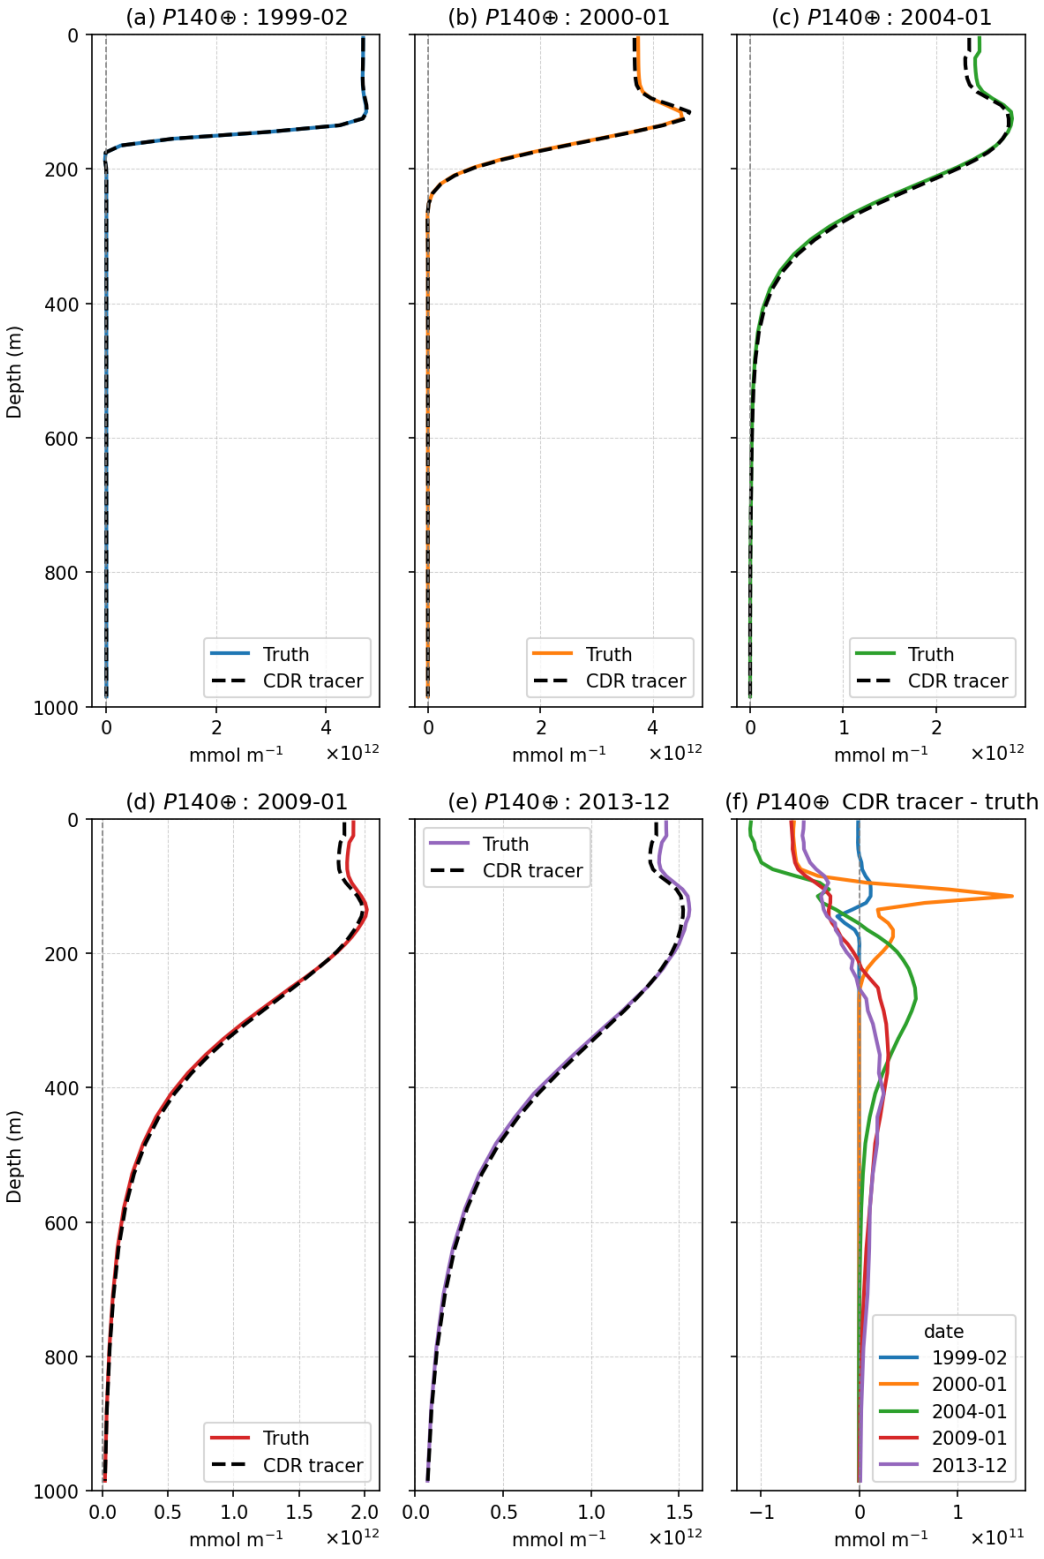


Figure S33: As Figure S32, but for polygon 140 in the subtropical Atlantic.

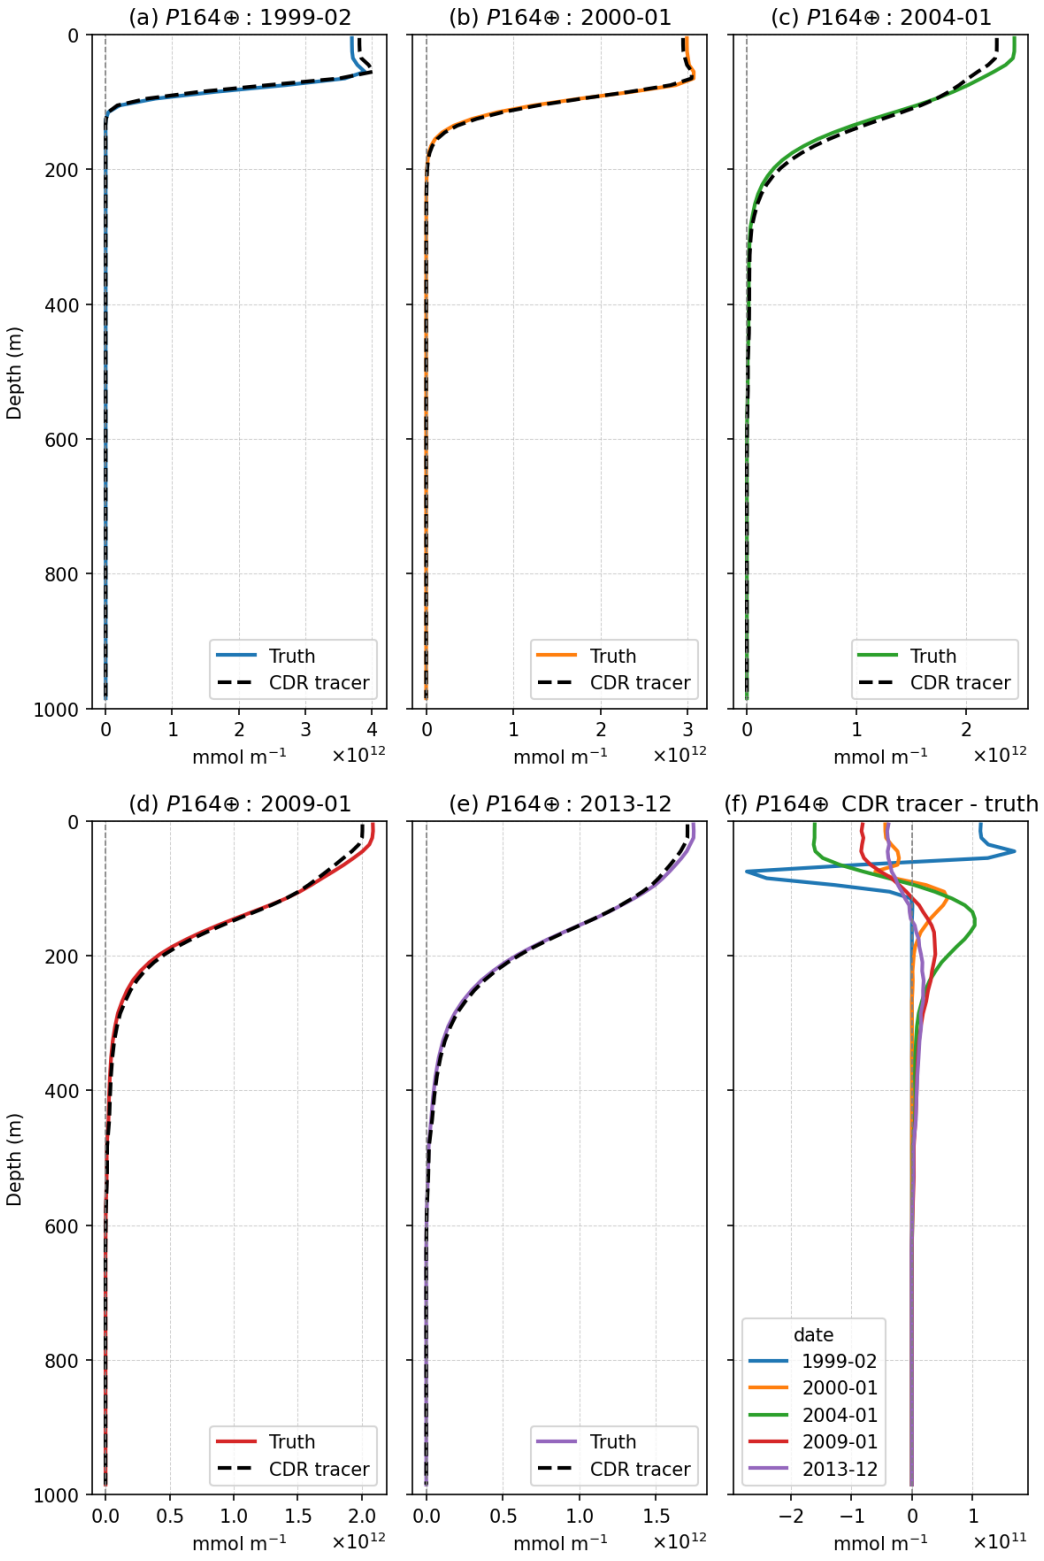


Figure S34: As Figure S32, but for polygon 164 near Baja California.

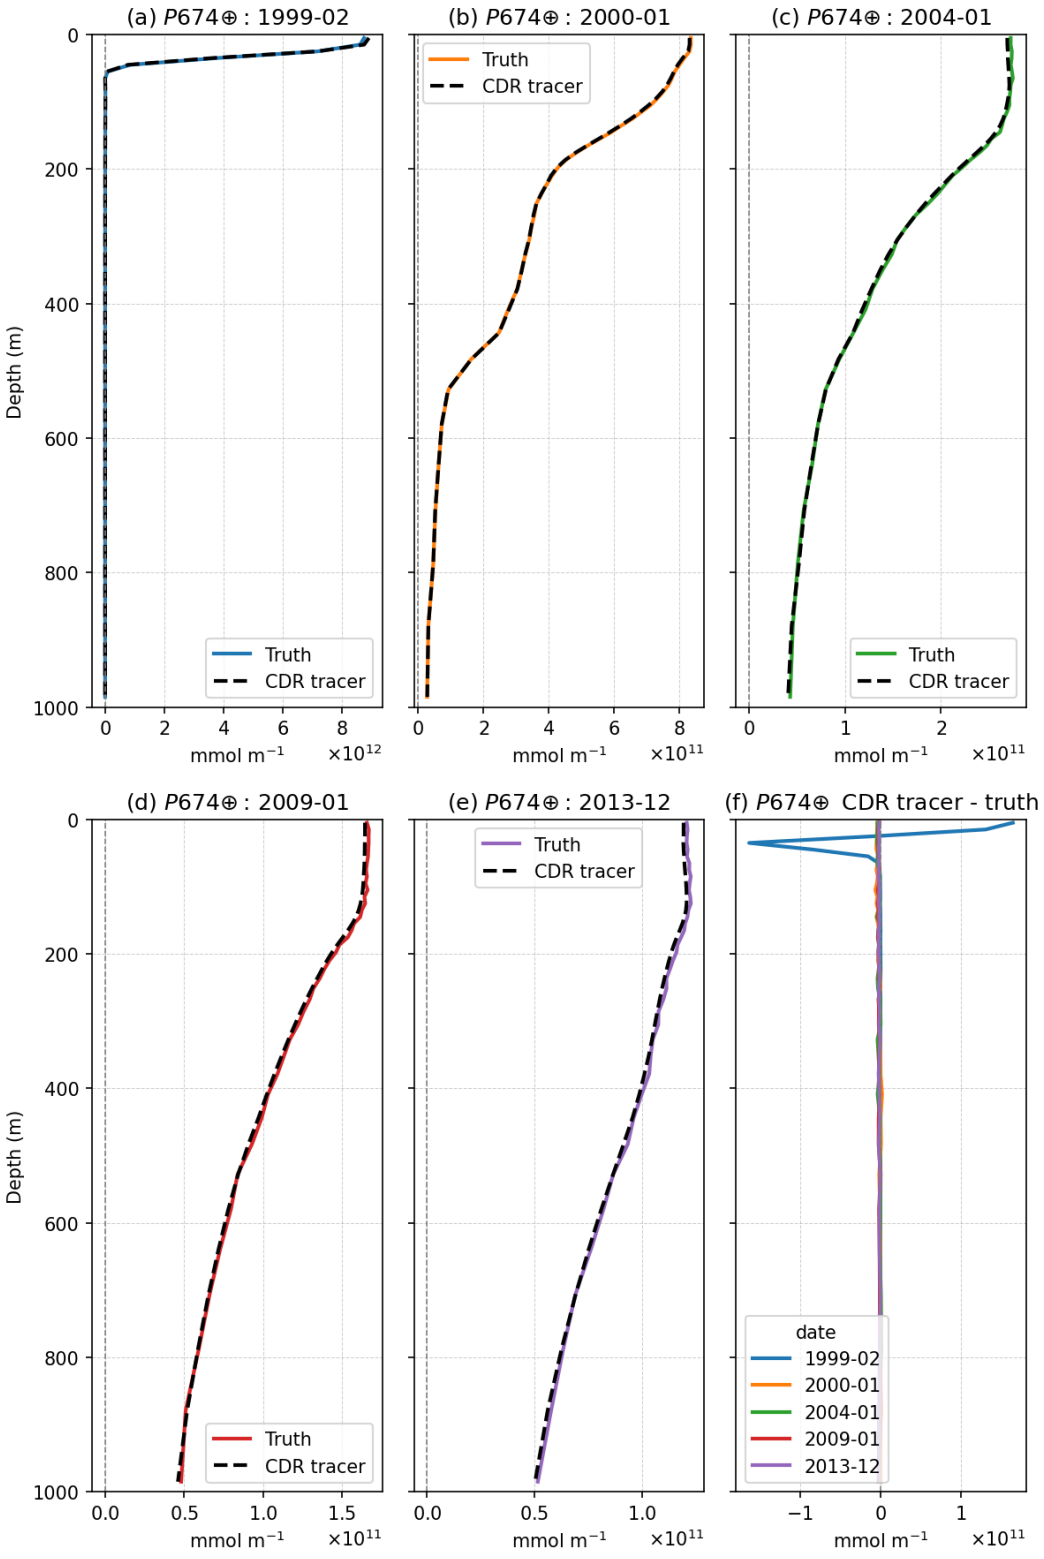


Figure S35: As Figure S32, but for the Weddell Sea polygon 674.

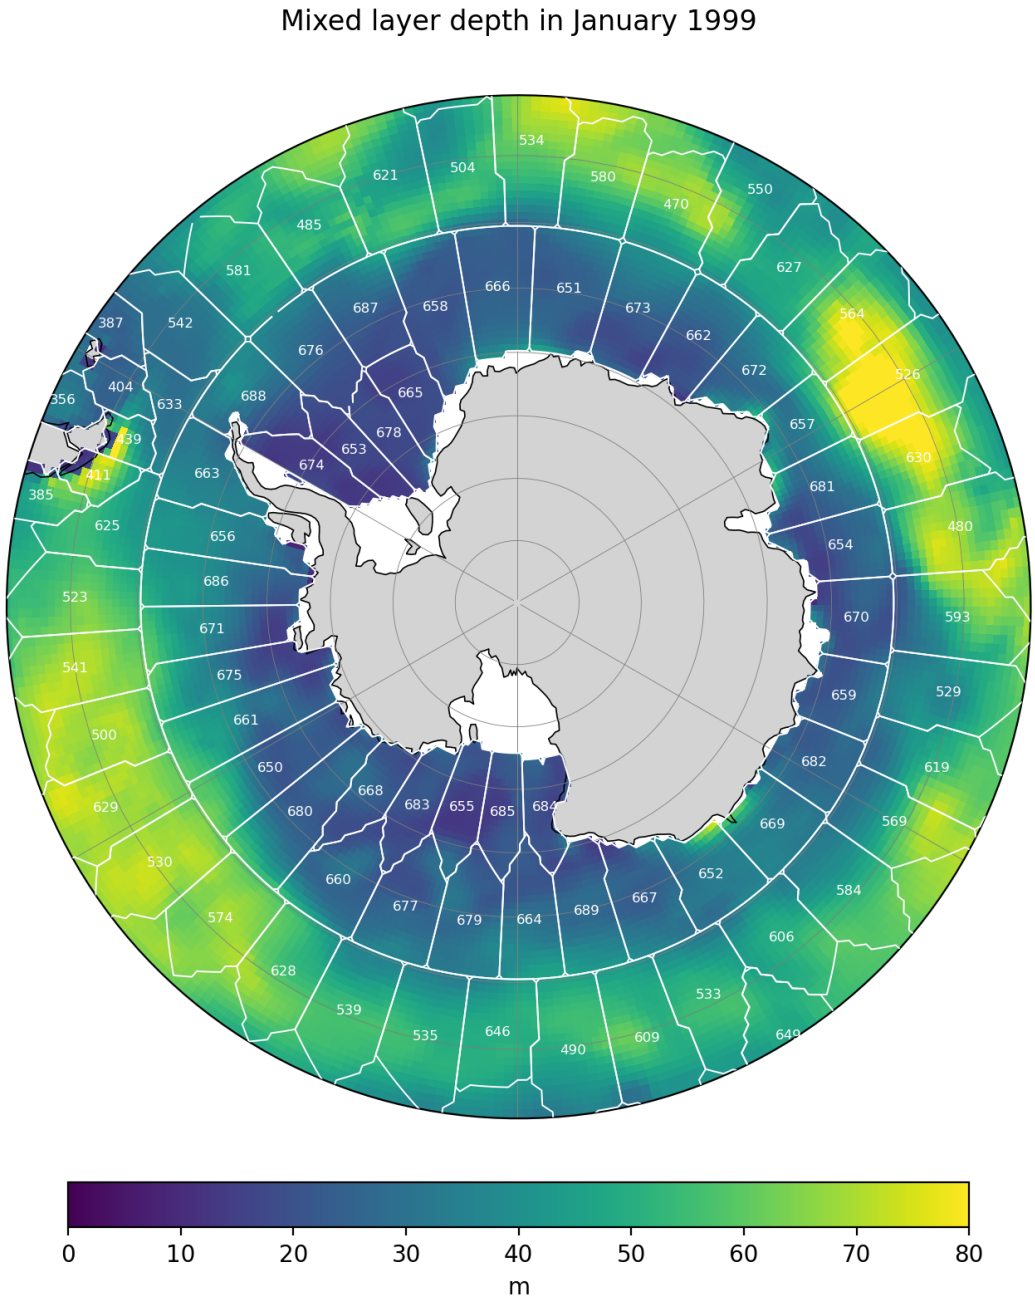


Figure S36: Mixed layer depth in CESM during the intervention month (January 1999), shown in a circum-Antarctic view with polygon outlines and numbers labeled. Mixed layers are shallow throughout the Southern Ocean in austral summer, particularly on the Antarctic shelf.

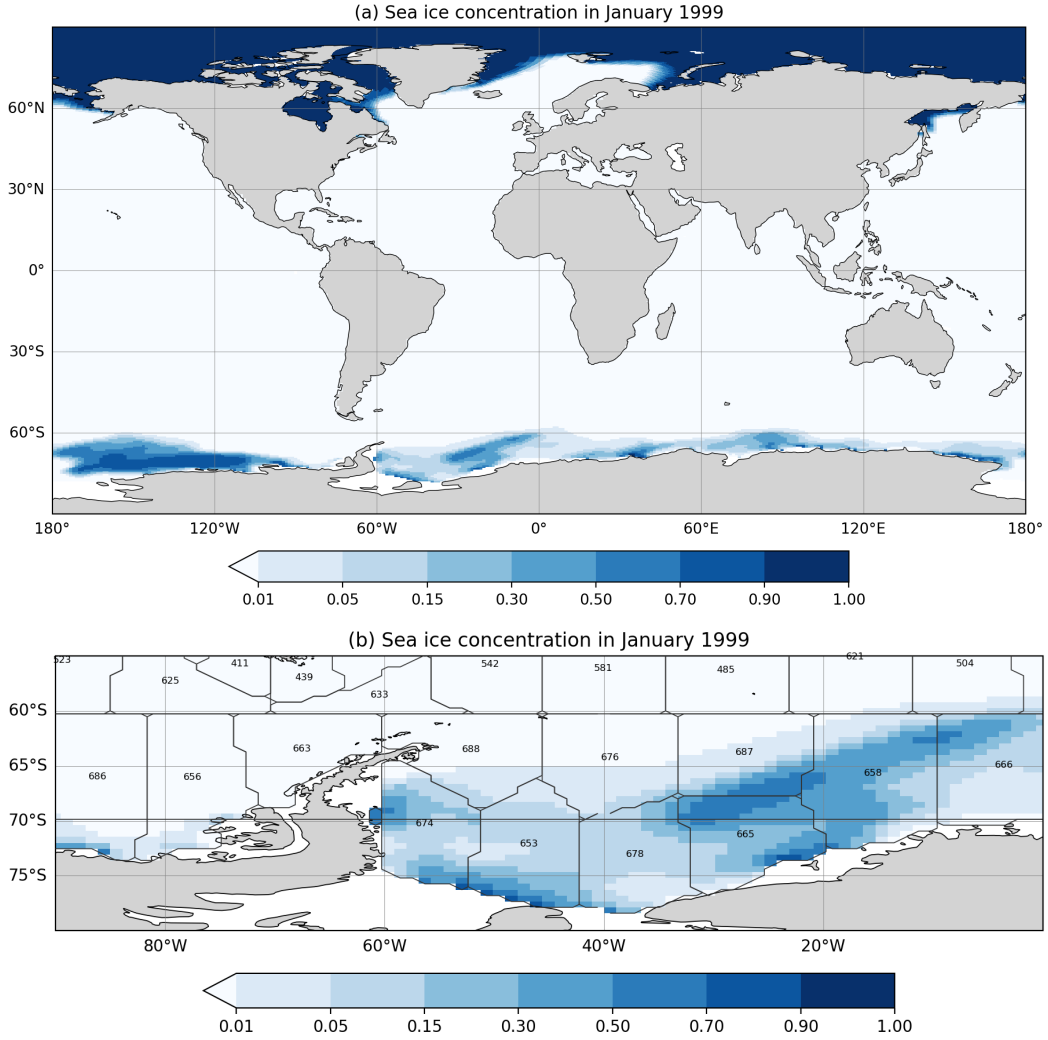


Figure S37: Sea ice concentration in the CESM CDR tracer run during the intervention month (January 1999). (a) Global view. (b) Zoom into the Weddell Sea and surrounding Southern Ocean, with polygon outlines and numbers labeled. The Antarctic shelf polygons (e.g. 653) are subject to partial sea ice cover, which modulates air-sea gas exchange. Combined with the extremely shallow mixed layers in austral summer (Figure S36), small differences in sea ice concentration between the truth and CDR tracer runs can produce disproportionately large errors in CO_2 uptake and CDR efficiency.

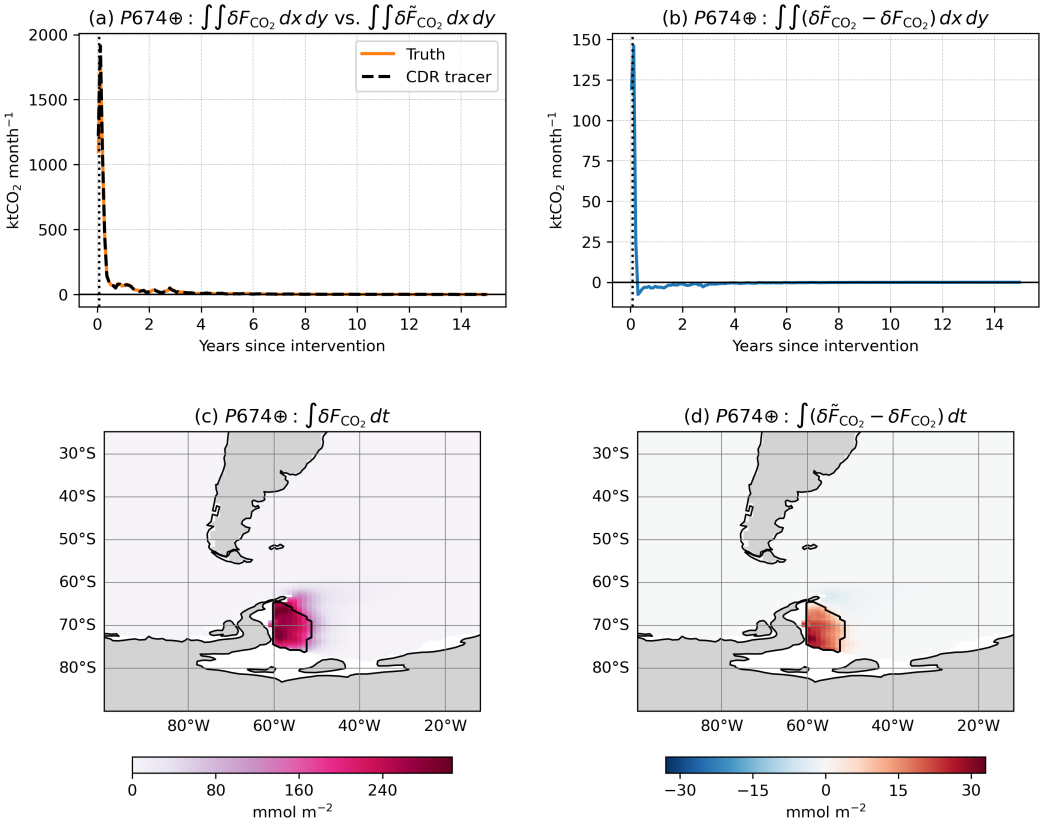


Figure S38: As Figure 12, but for the Weddell Sea polygon 674.

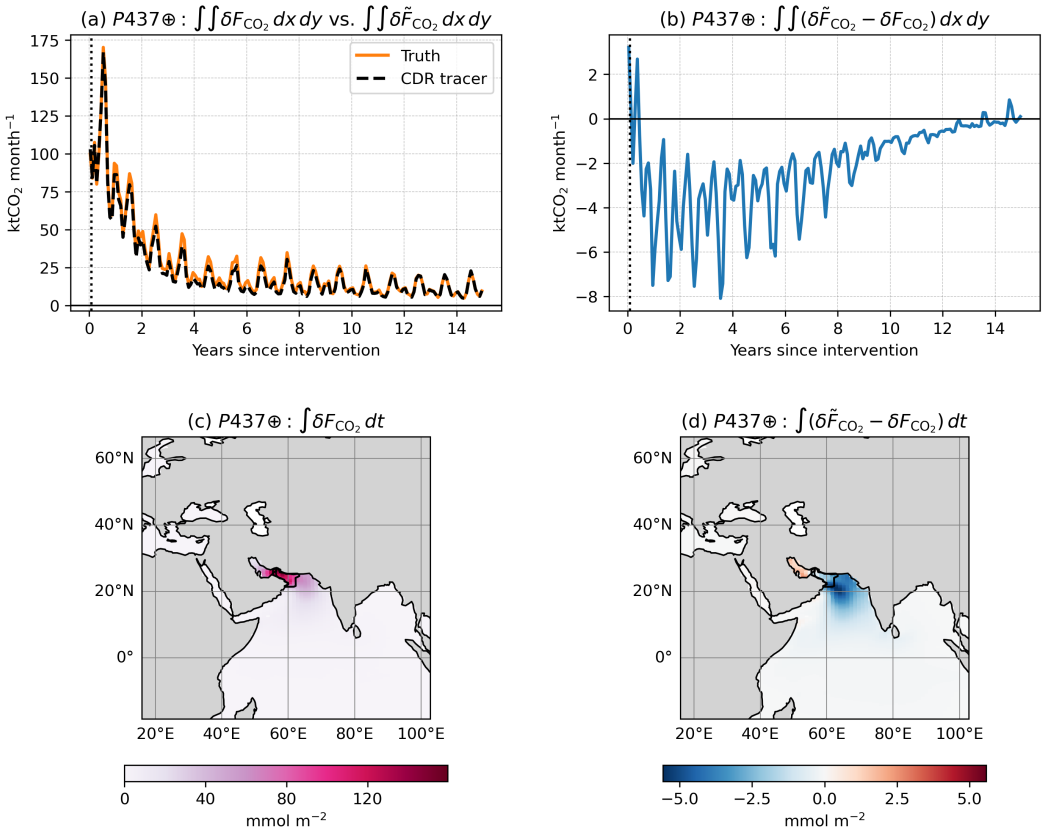


Figure S39: As Figure 12, but for the Arabian Sea polygon 437.

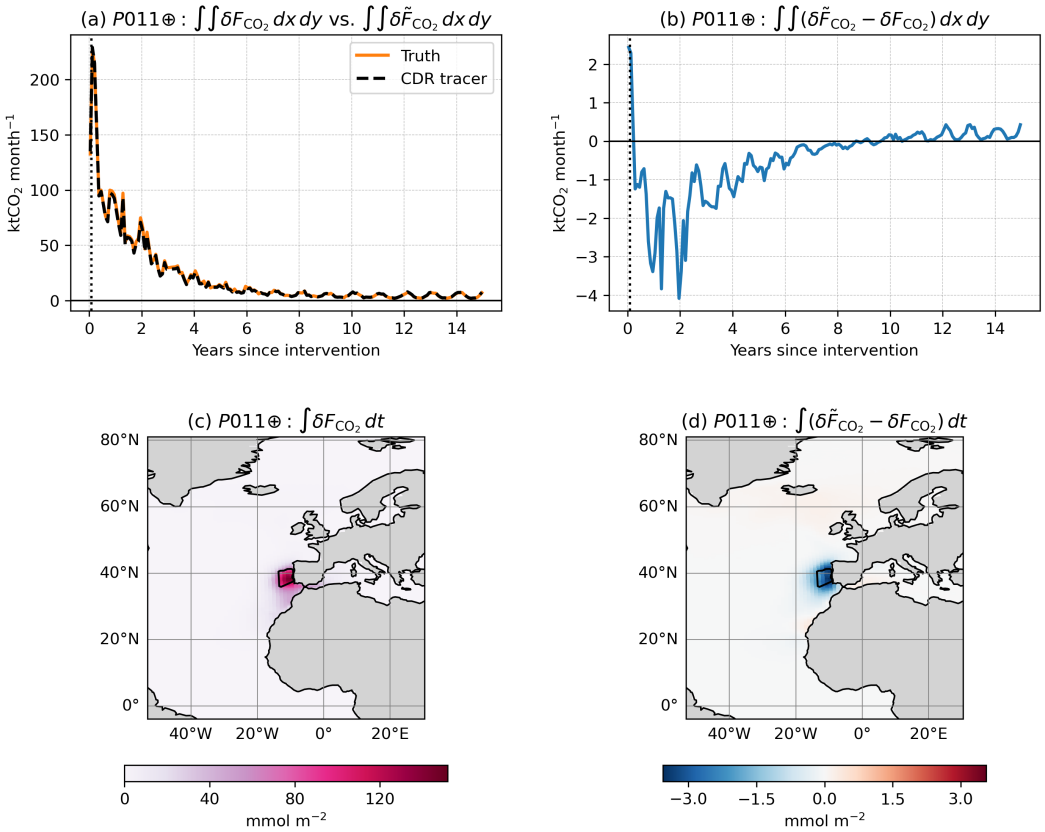


Figure S40: As Figure 12, but for the North Atlantic polygon 11.

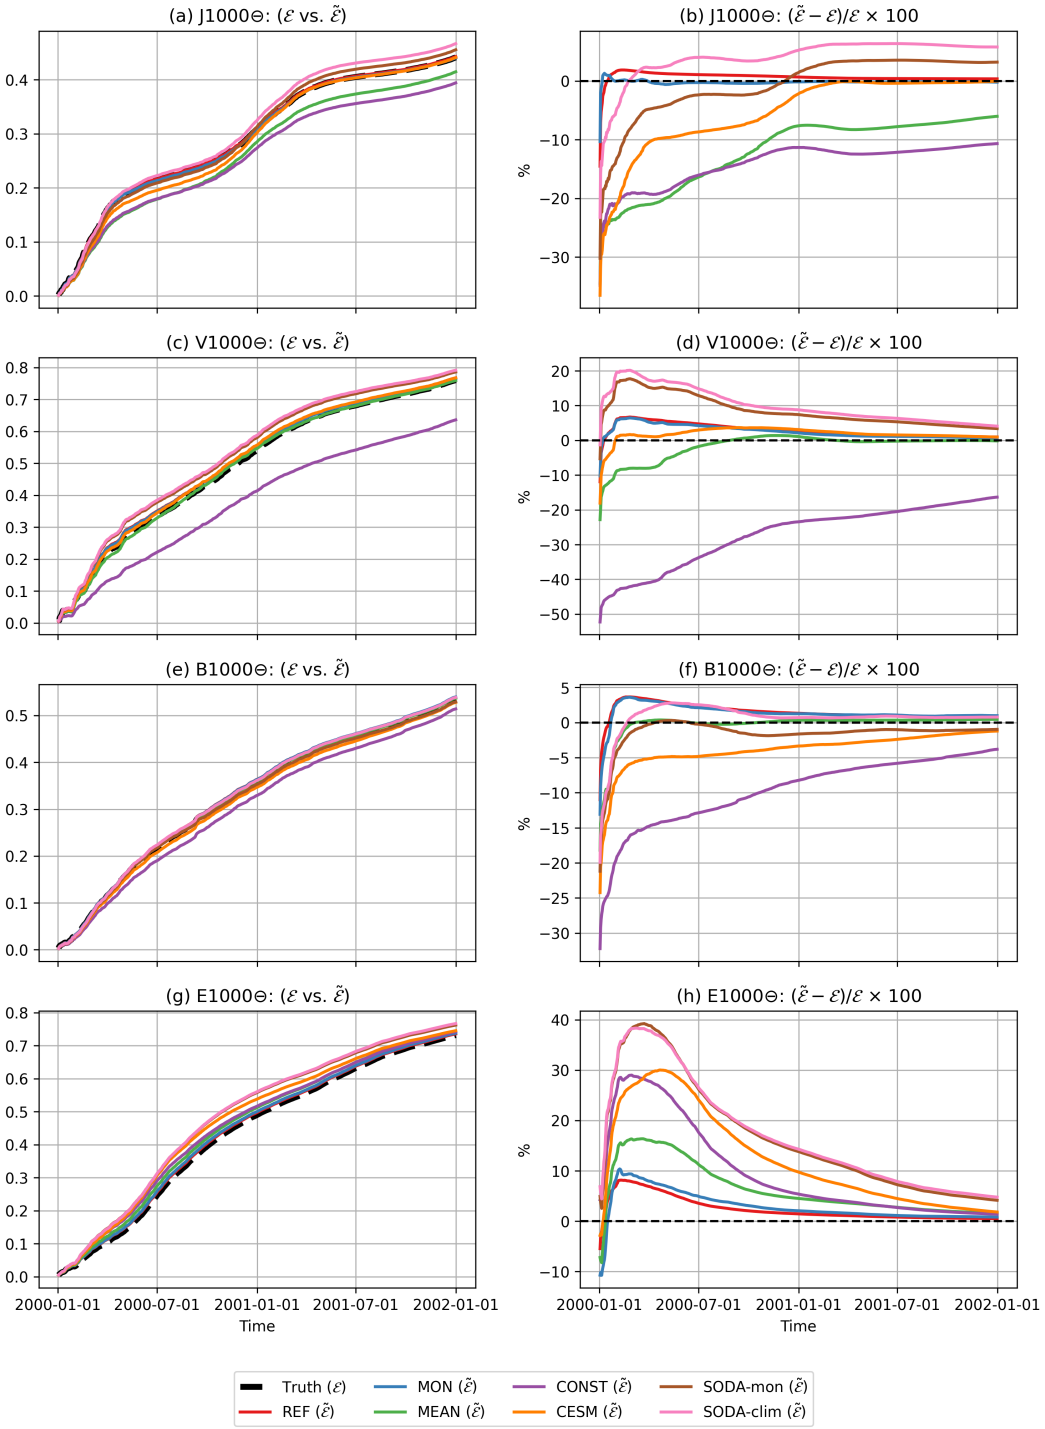


Figure S41: As Figure 13, but for DOR (Eq. 19) rather than OAE.

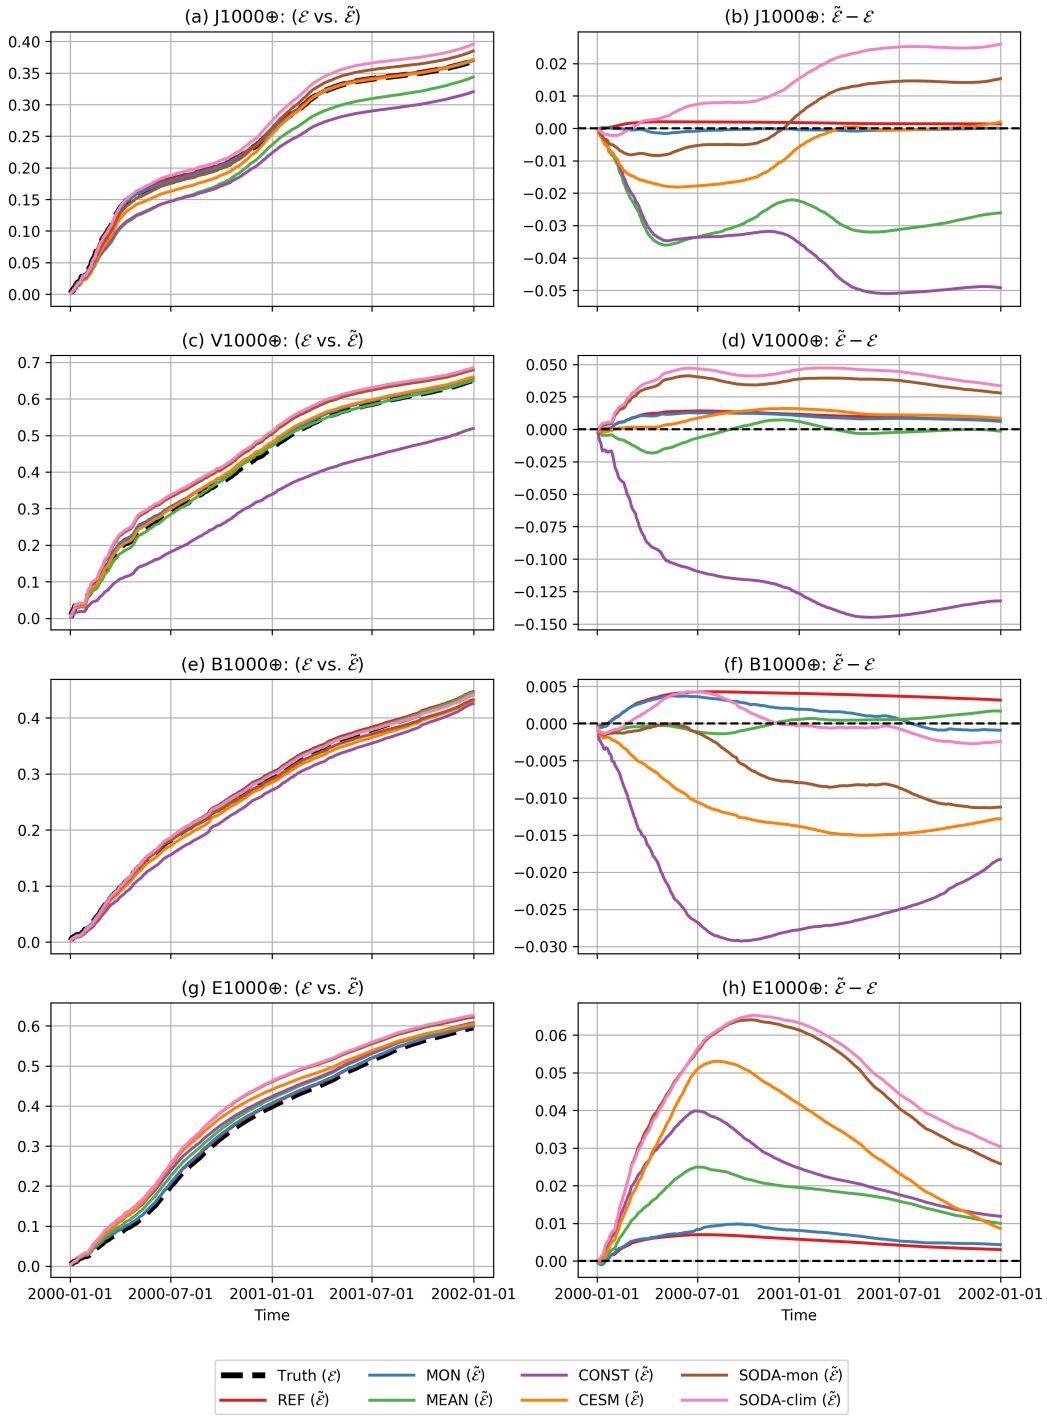


Figure S42: As Figure 13, but showing absolute errors $\tilde{\mathcal{E}}^{\mathrm{OAE}} - \mathcal{E}^{\mathrm{OAE}}$ in the right column rather than relative errors.

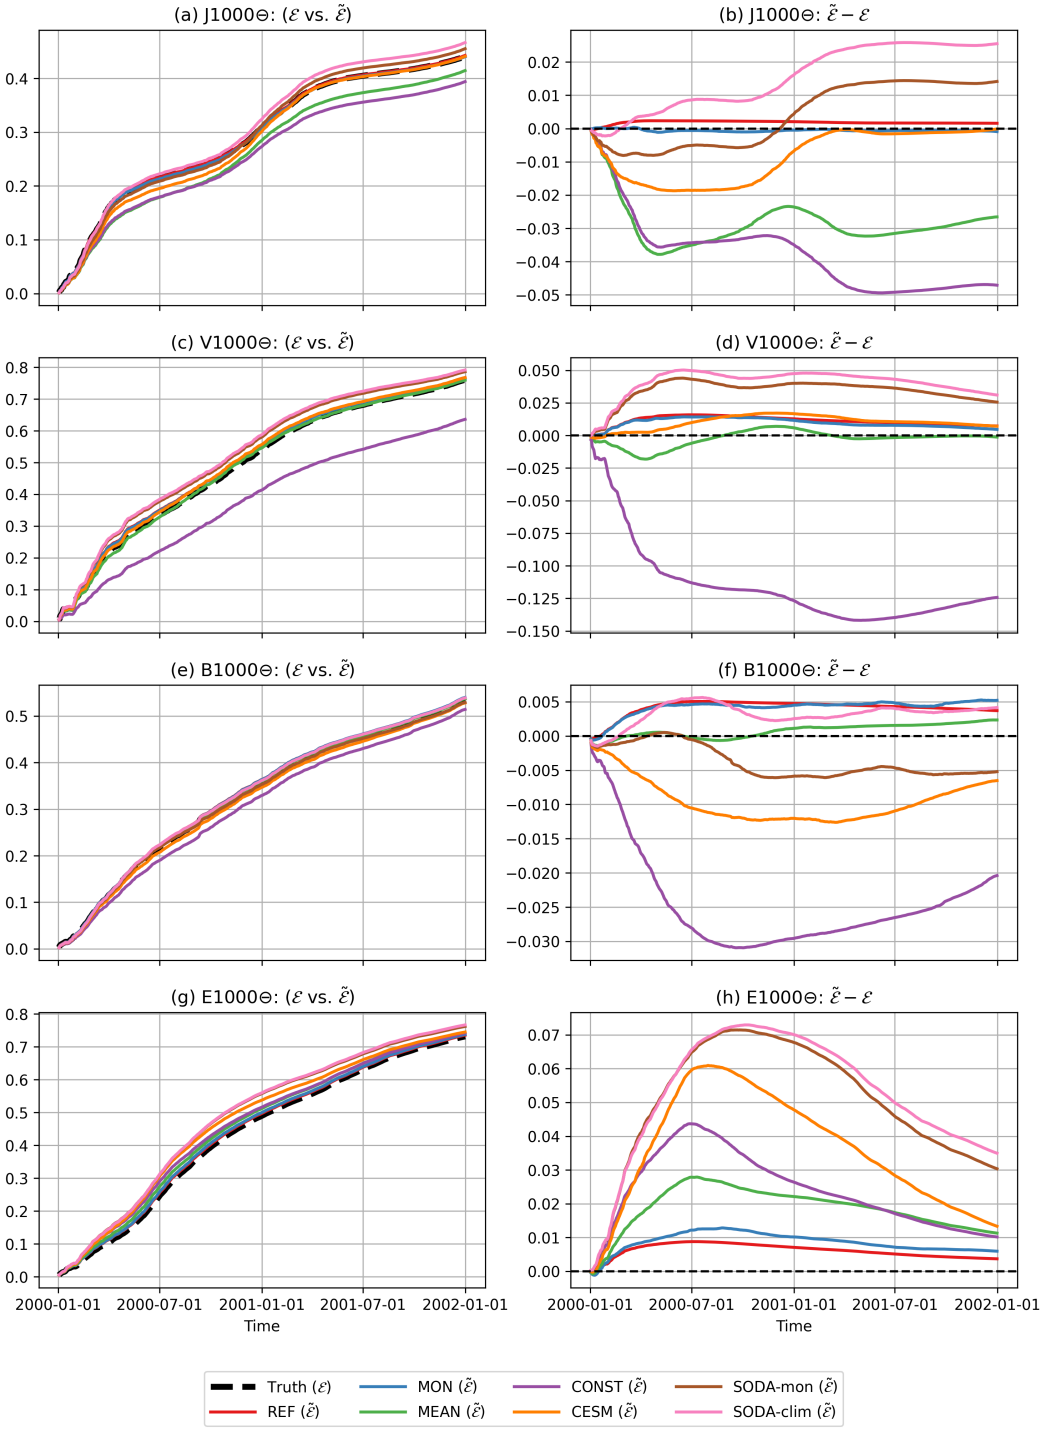


Figure S43: As Figure S42, but showing absolute errors for DOR rather than OAE.

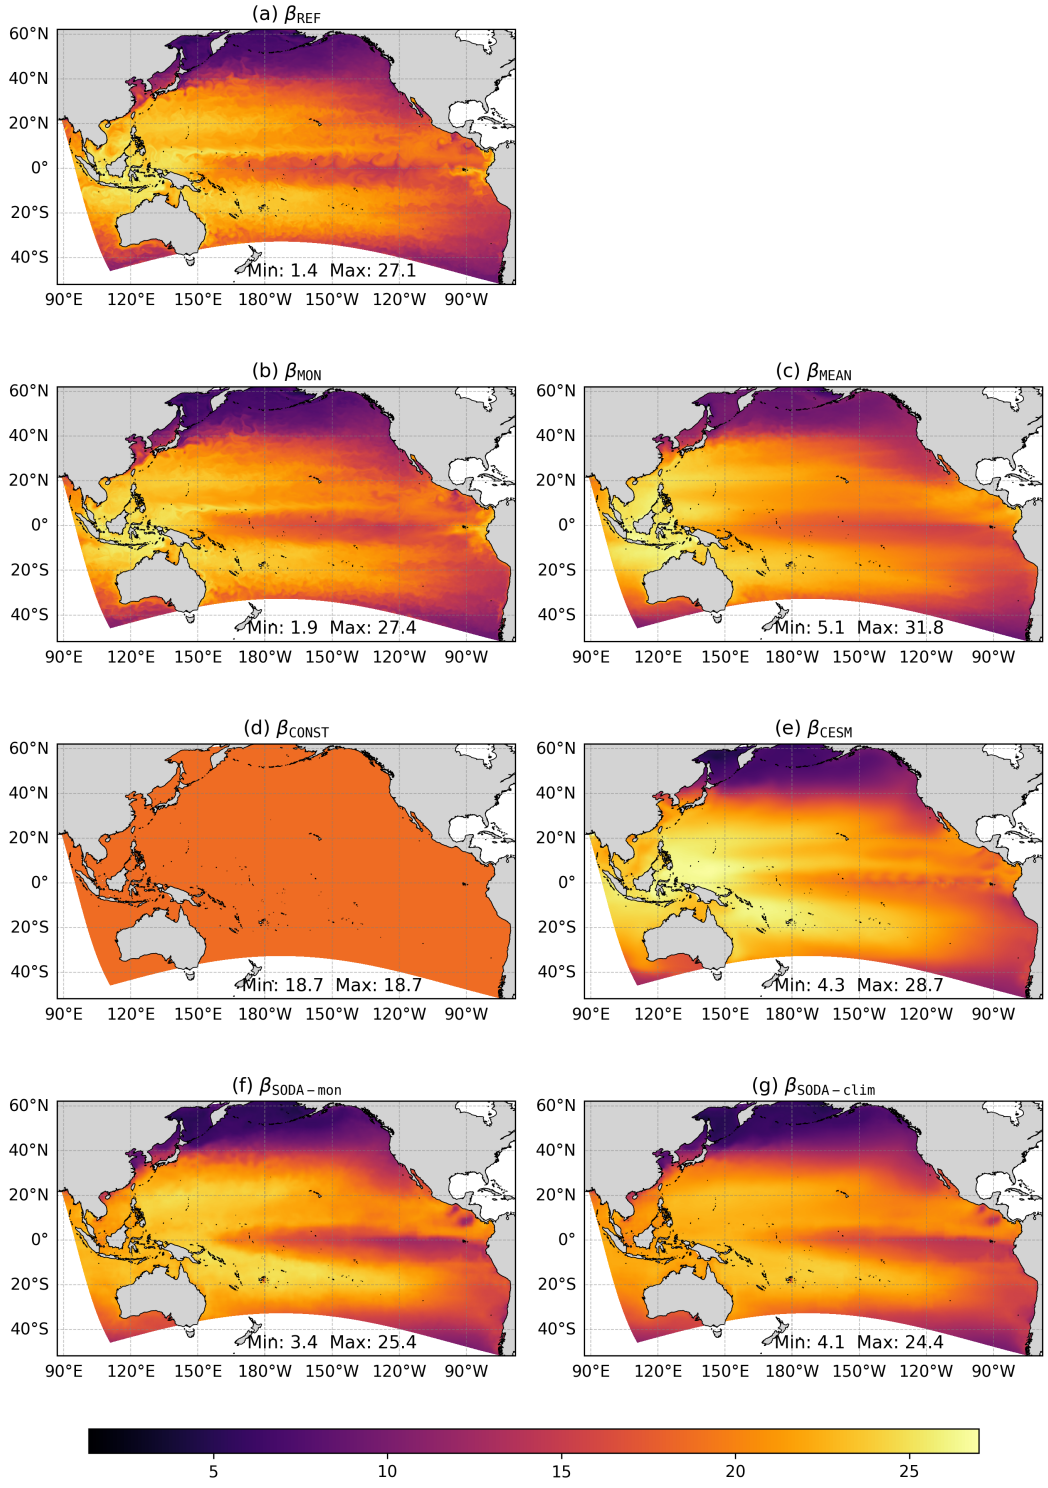


Figure S44: Carbonate sensitivity β for each of the seven CDR tracer configurations (Table 2), shown over the full ROMS Pacific domain on January 1, 2000. β_{REF} (a) is the daily ROMS/MARBL reference field; β_{MON} and β_{MEAN} (b,c) are obtained by temporal averaging to monthly and time-mean resolutions, respectively; β_{CONST} (d) is the single domain-wide constant obtained by additionally averaging over space. β_{CESM} (e), $\beta_{\mathrm{SODA-mon}}$, and $\beta_{\mathrm{SODA-clim}}$ (f,g) are derived from the CESM/MARBL control run and the OceanSODA-ETHZ product, respectively, regridded to the ROMS grid. All panels share the same colorbar. The domain-wide minimum and maximum of each field are annotated in each panel.

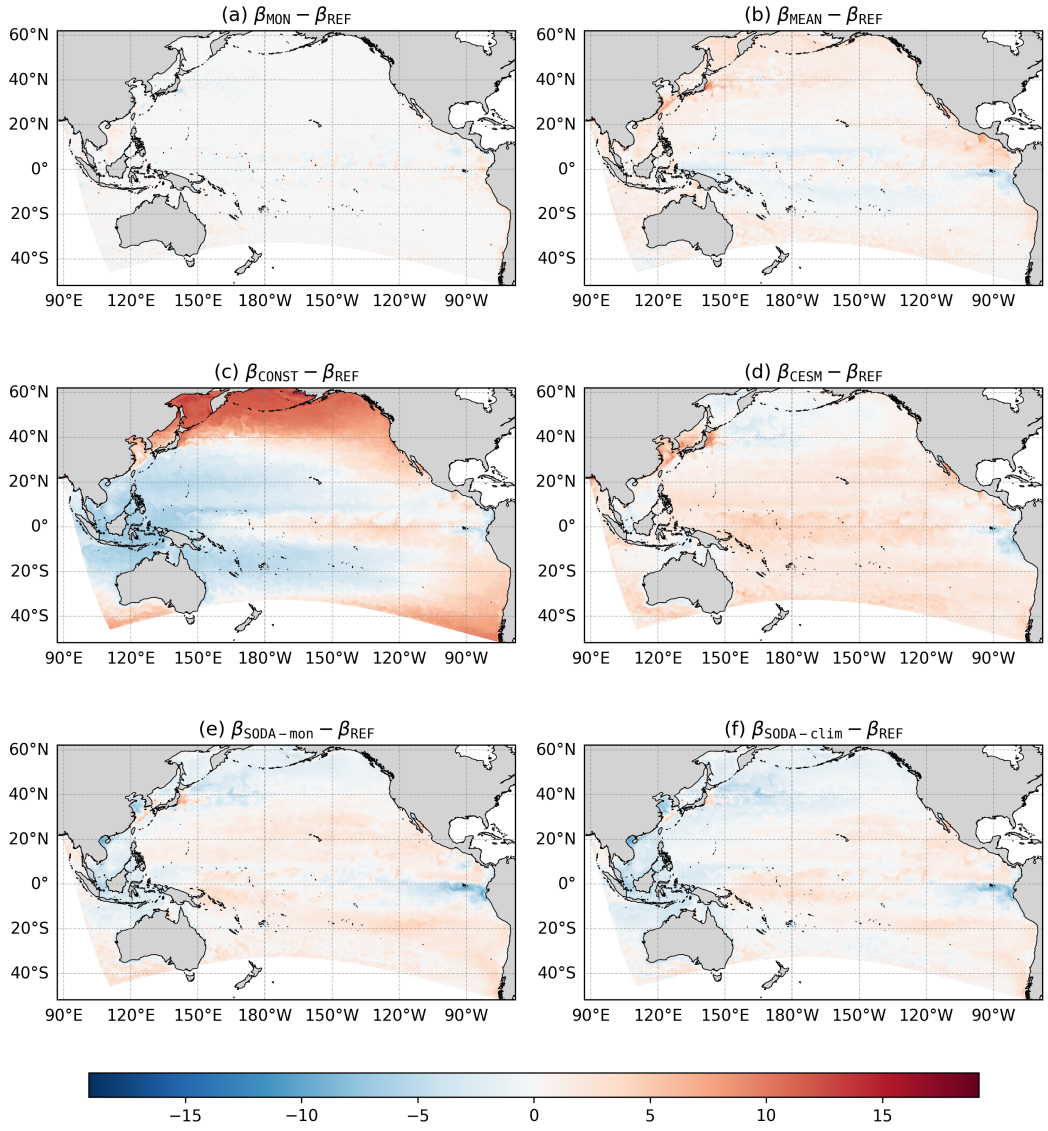


Figure S45: As Figure S44, but showing the absolute difference $\beta_X - \beta_{\mathrm{REF}}$ for each configuration $X \in \{\mathrm{MON}, \mathrm{MEAN}, \mathrm{CONST}, \mathrm{CESM}, \mathrm{SODA\text{-}mon}, \mathrm{SODA\text{-}clim}\}$.

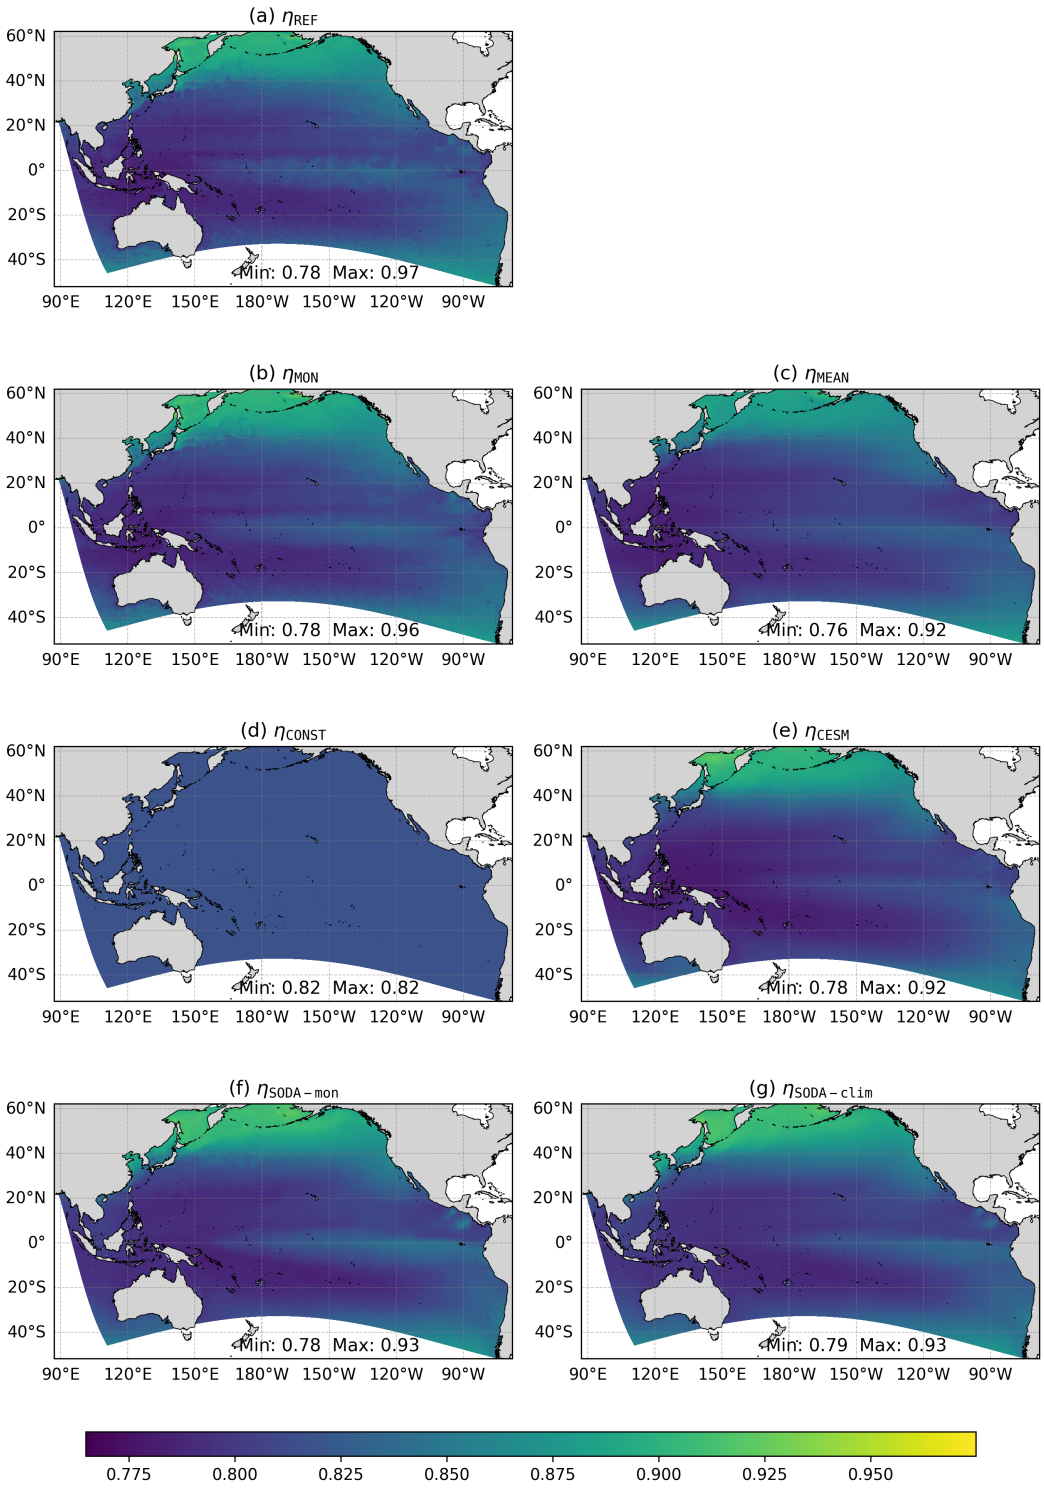


Figure S46: As Figure S44, but for η.

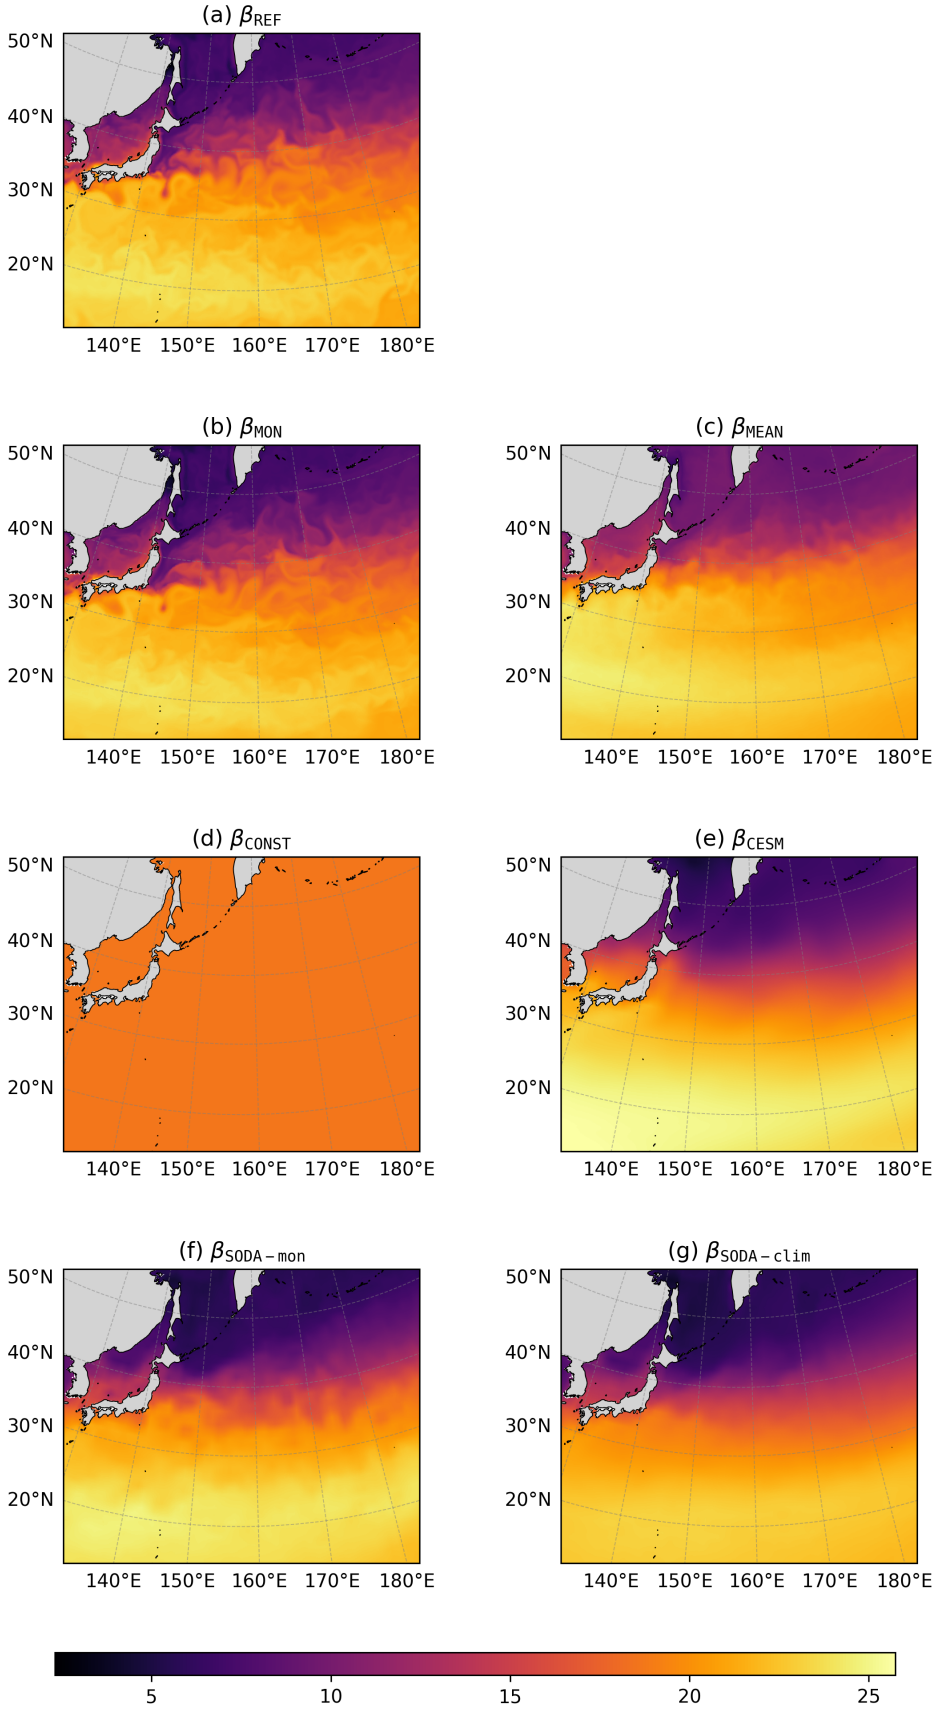


Figure S47: As Figure S44, but zoomed into the Japan intervention region. The colorbar range has been adjusted to the values present in this region.

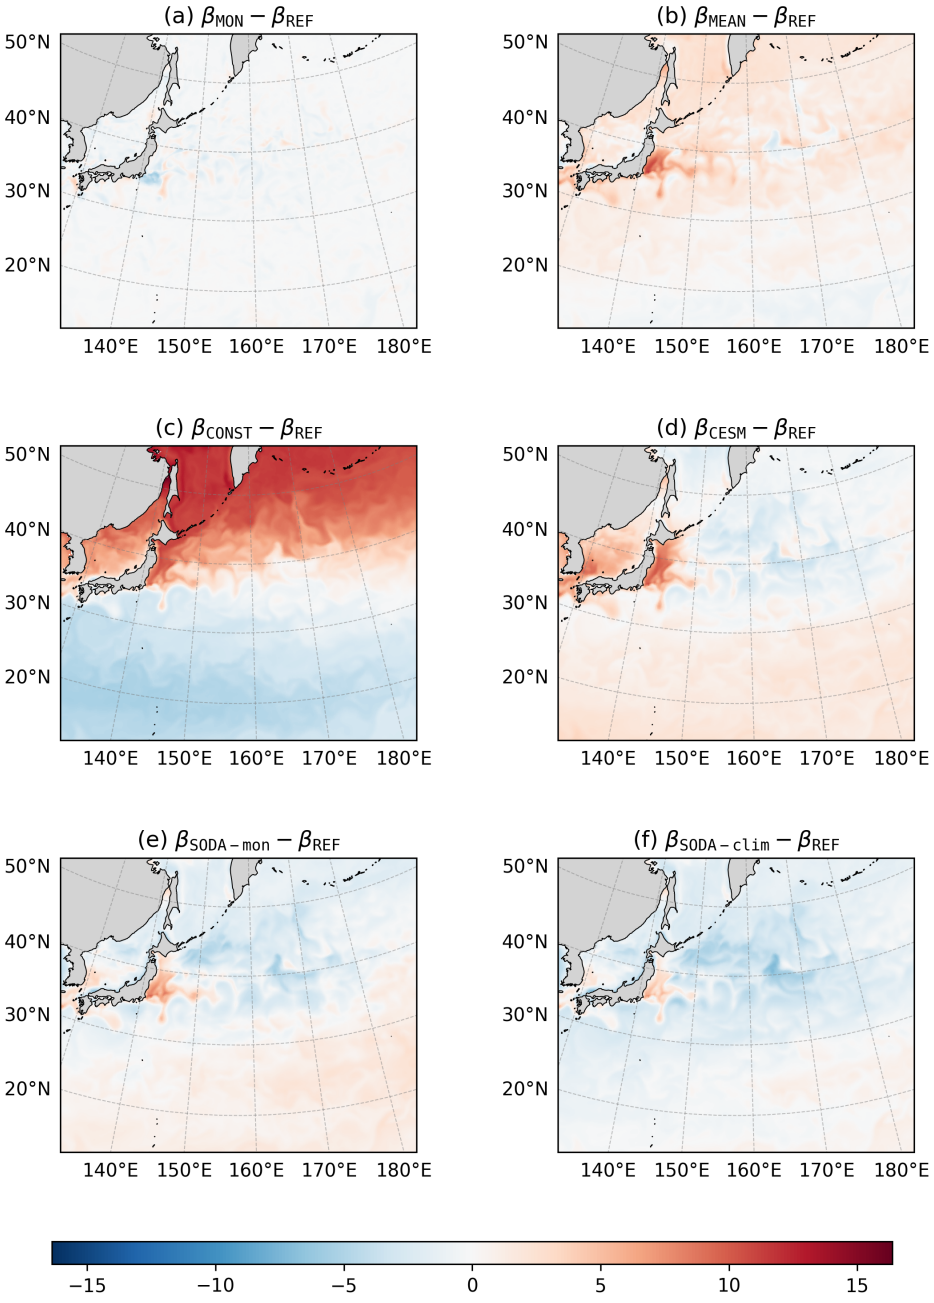


Figure S48: As Figure S45, but zoomed into the Japan intervention region. The colorbar range has been adjusted to the values present in this region.

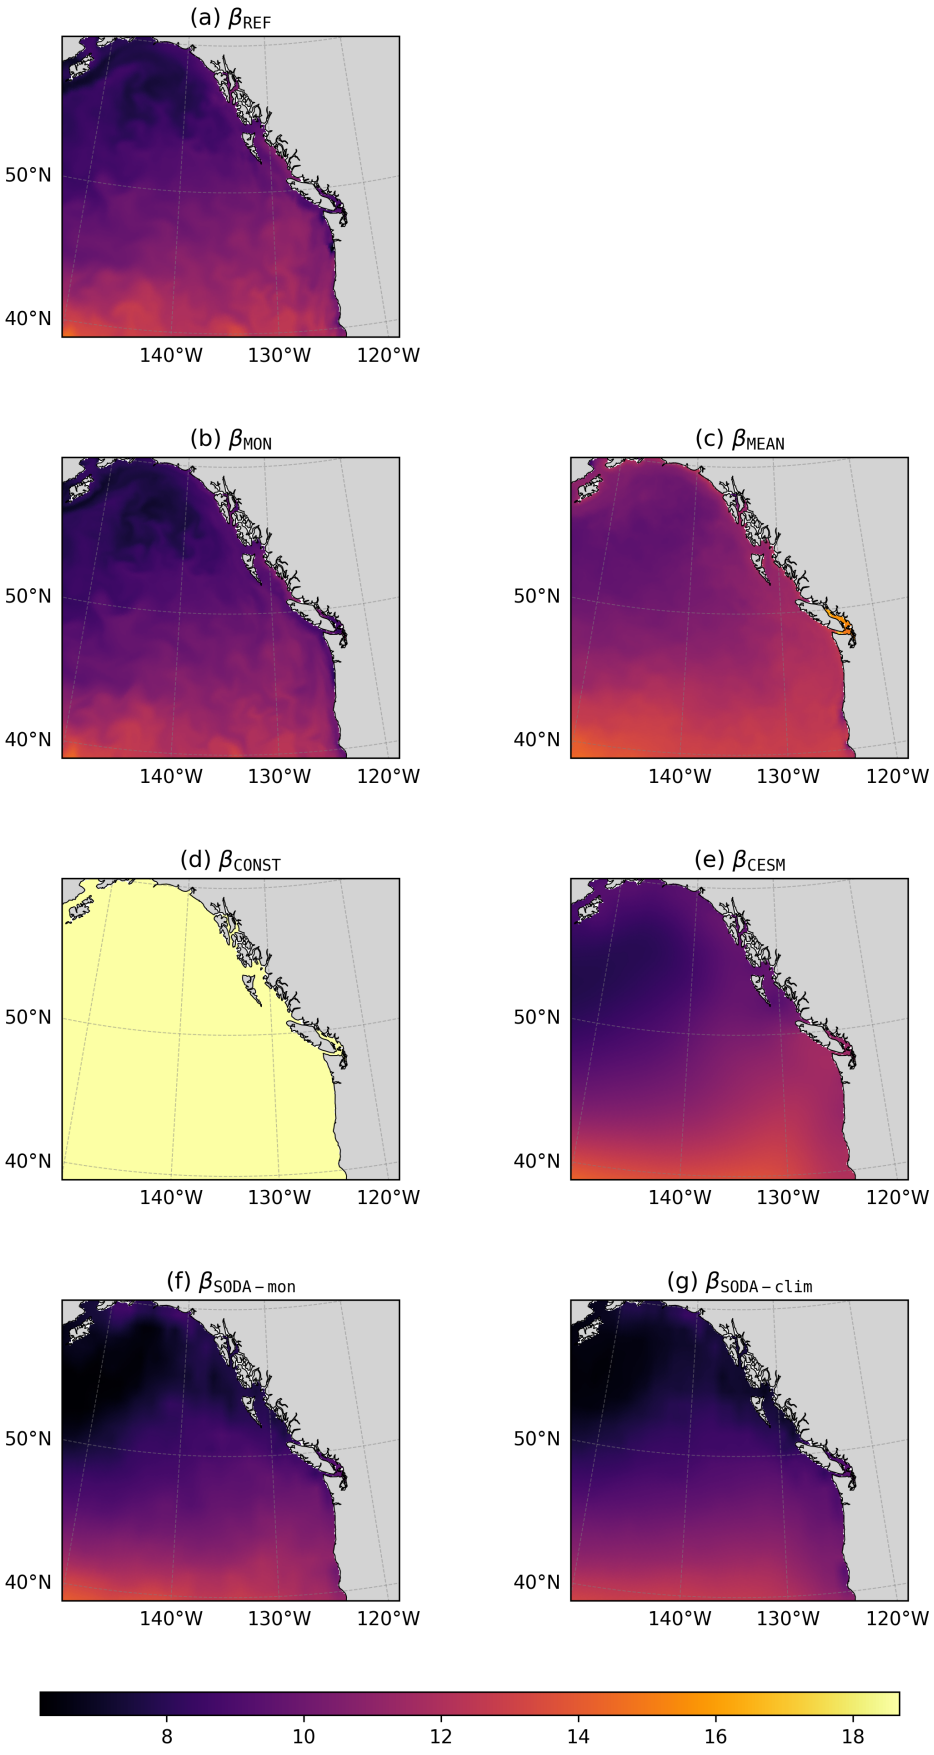


Figure S49: As Figure S44, but zoomed into the Vancouver Island intervention region. The colorbar range has been adjusted to the values present in this region.

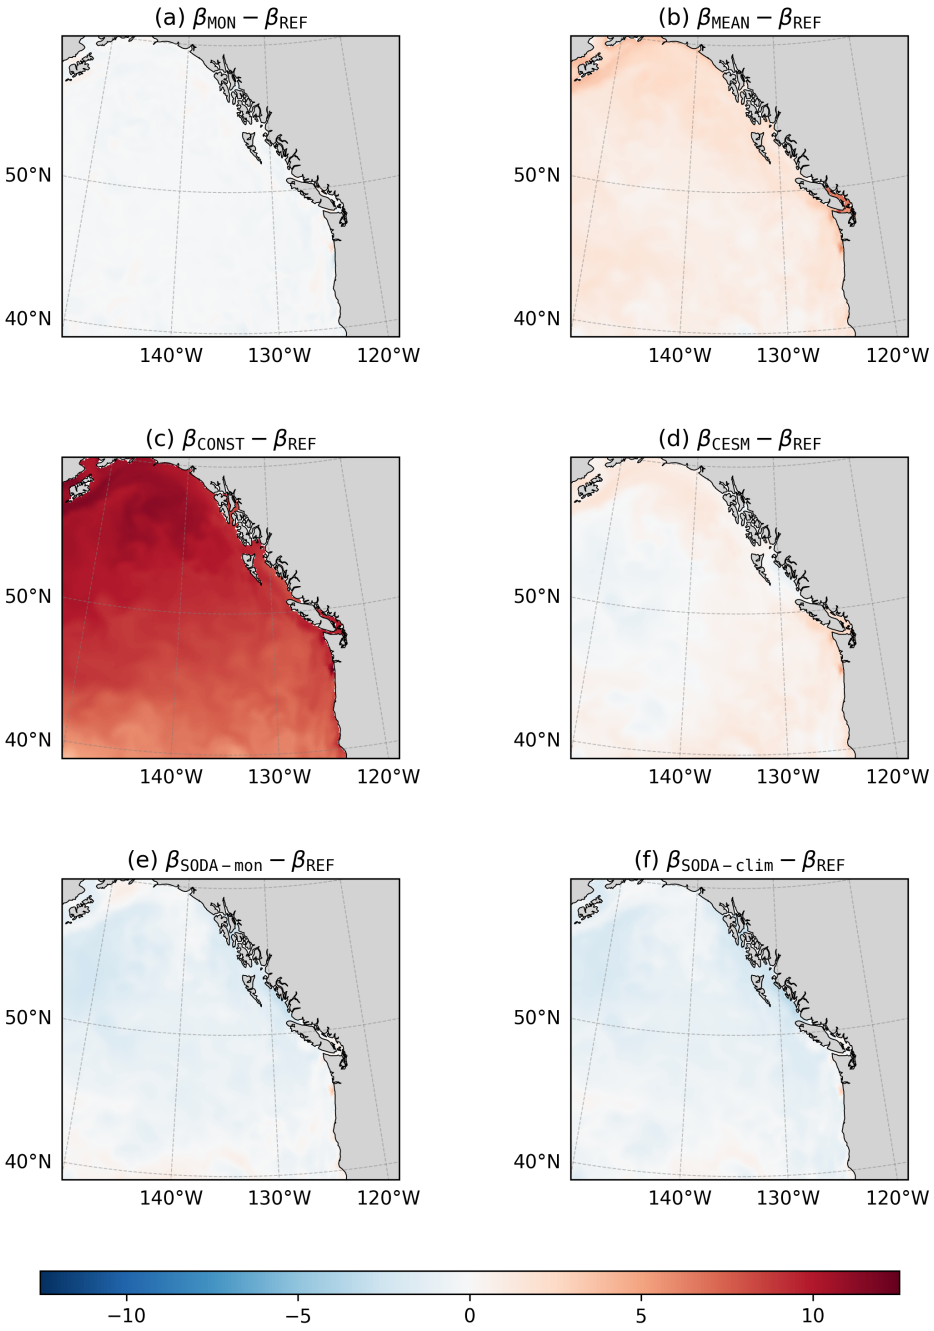


Figure S50: As Figure S45, but zoomed into the Vancouver Island intervention region. The colorbar range has been adjusted to the values present in this region.

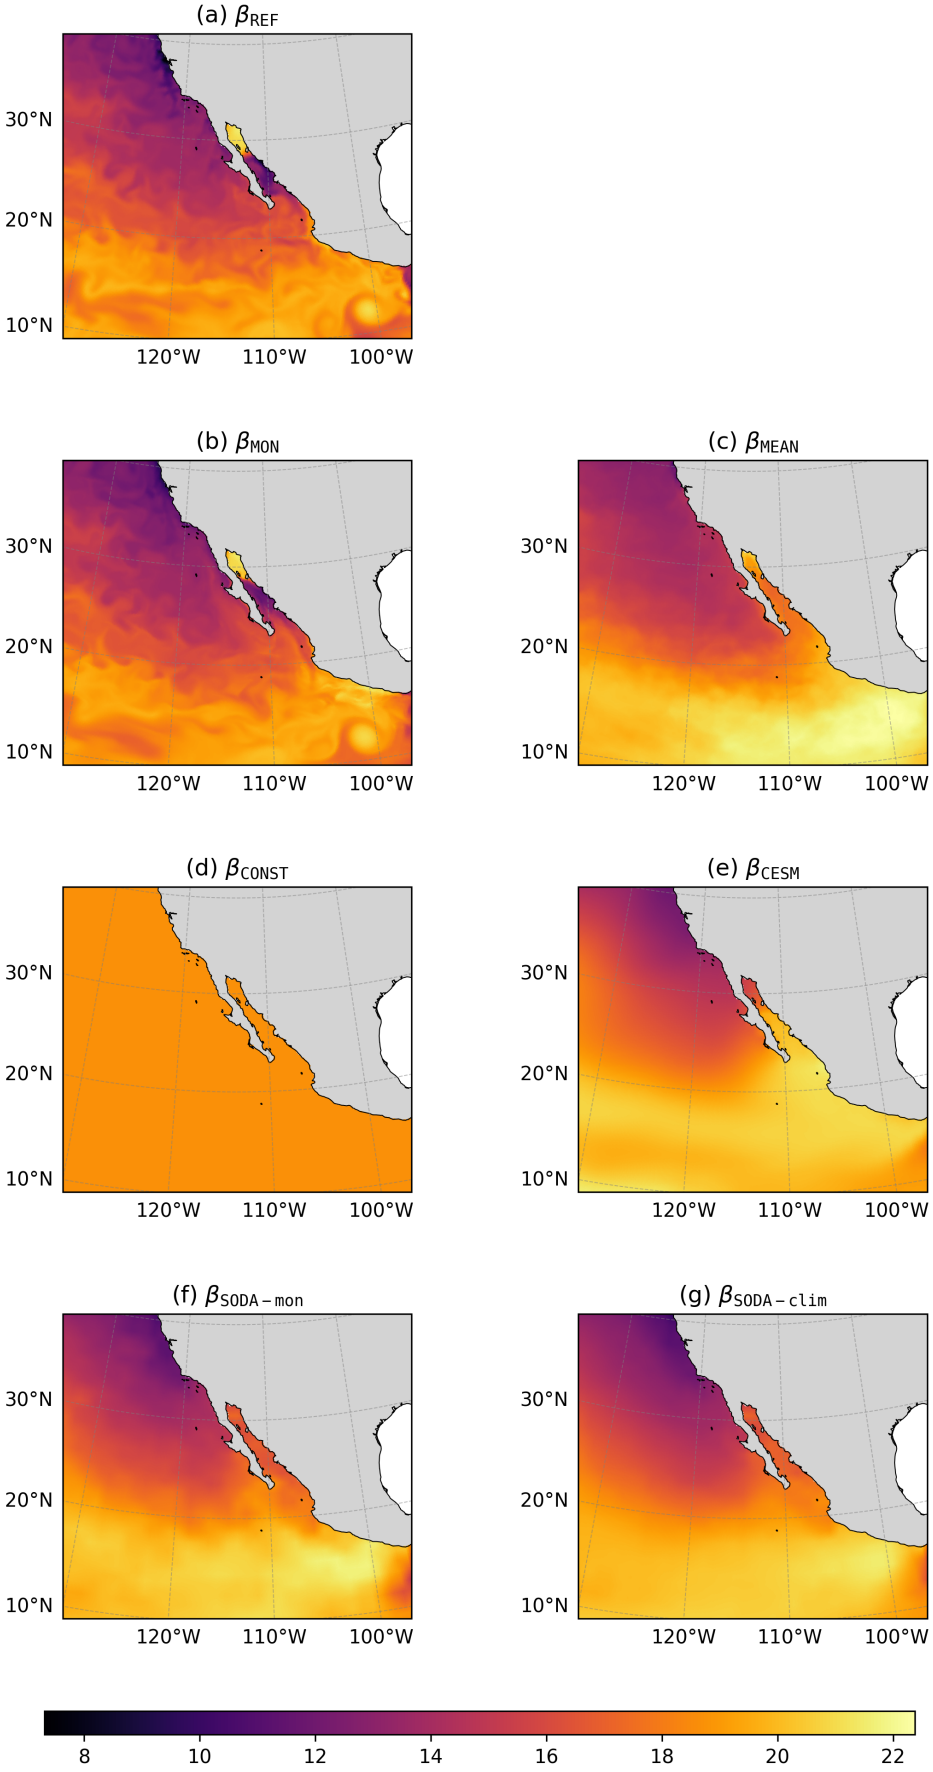


Figure S51: As Figure S44, but zoomed into the Baja California intervention region. The colorbar range has been adjusted to the values present in this region.

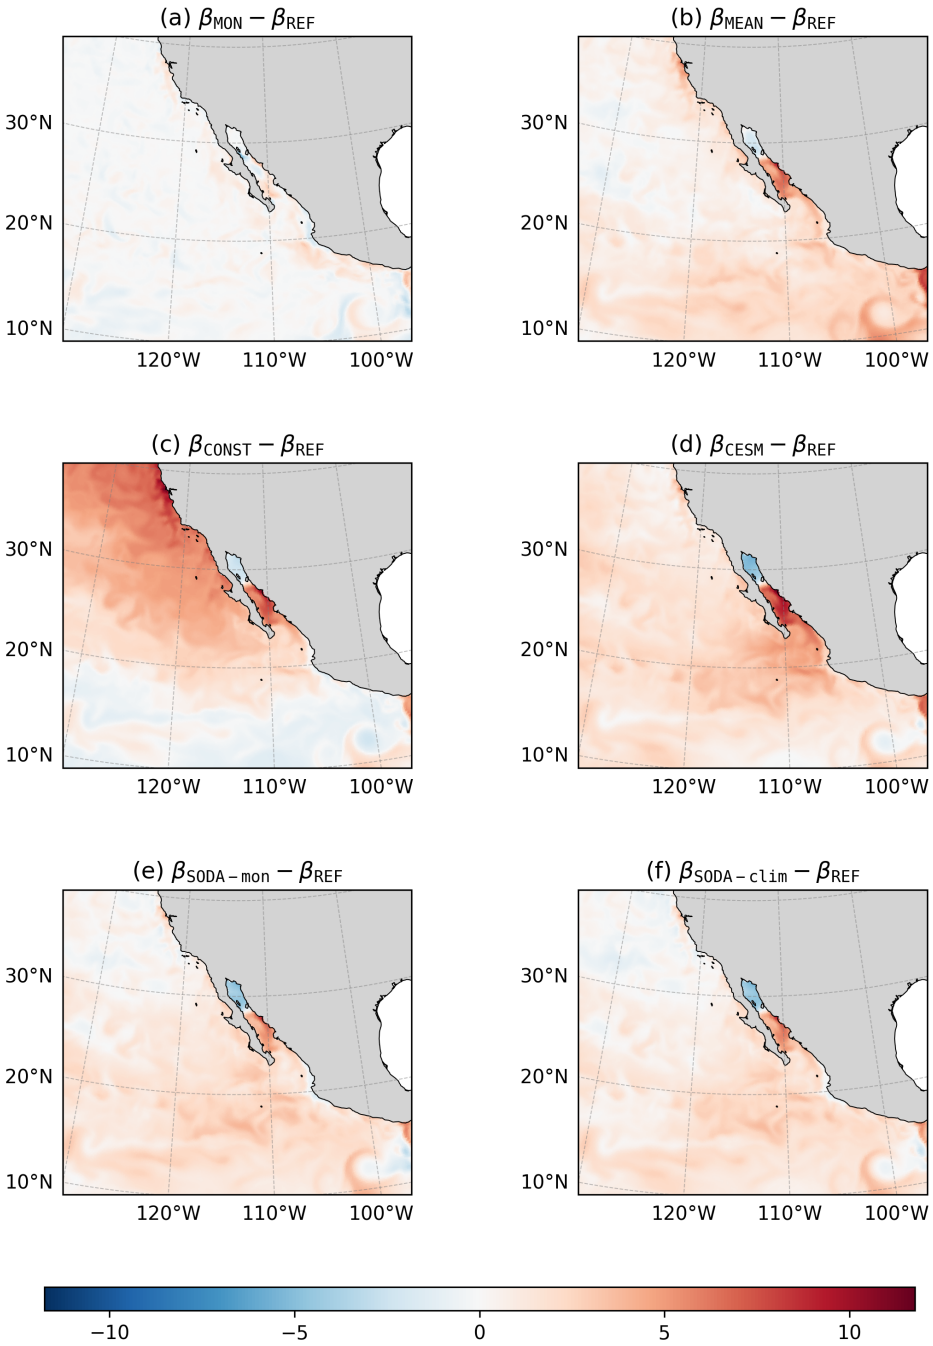


Figure S52: As Figure S45, but zoomed into the Baja California intervention region. The colorbar range has been adjusted to the values present in this region.

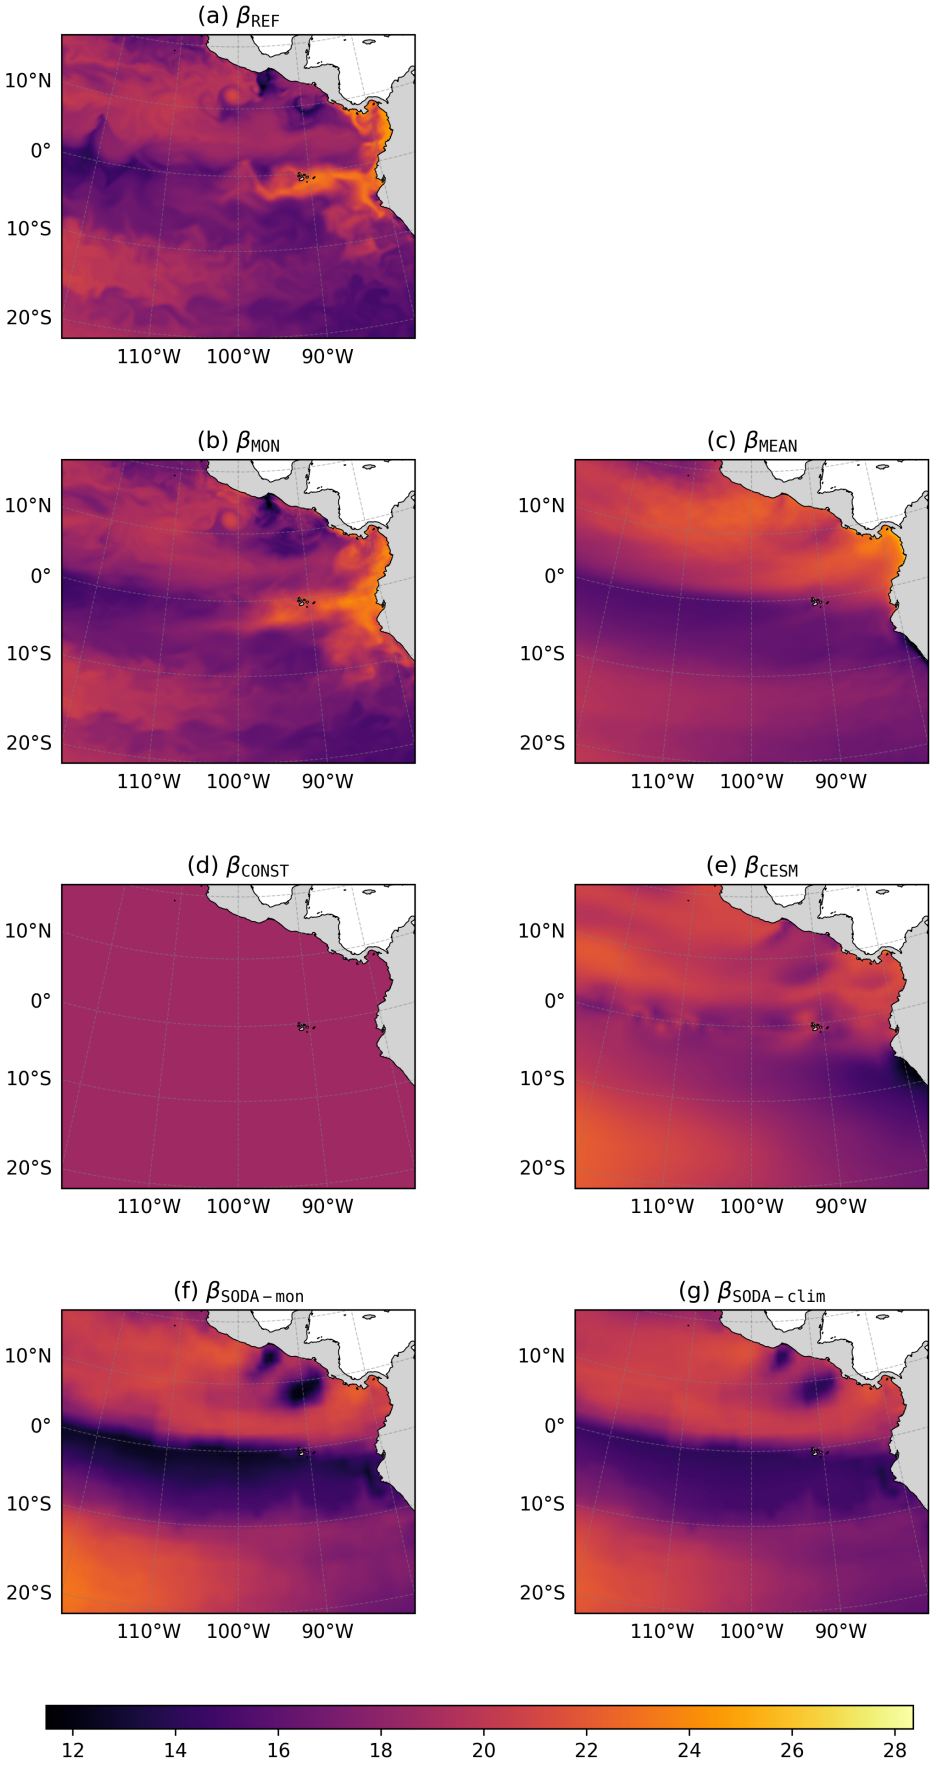


Figure S53: As Figure S44, but zoomed into the Ecuador intervention region. The color-bar range has been adjusted to the values present in this region.

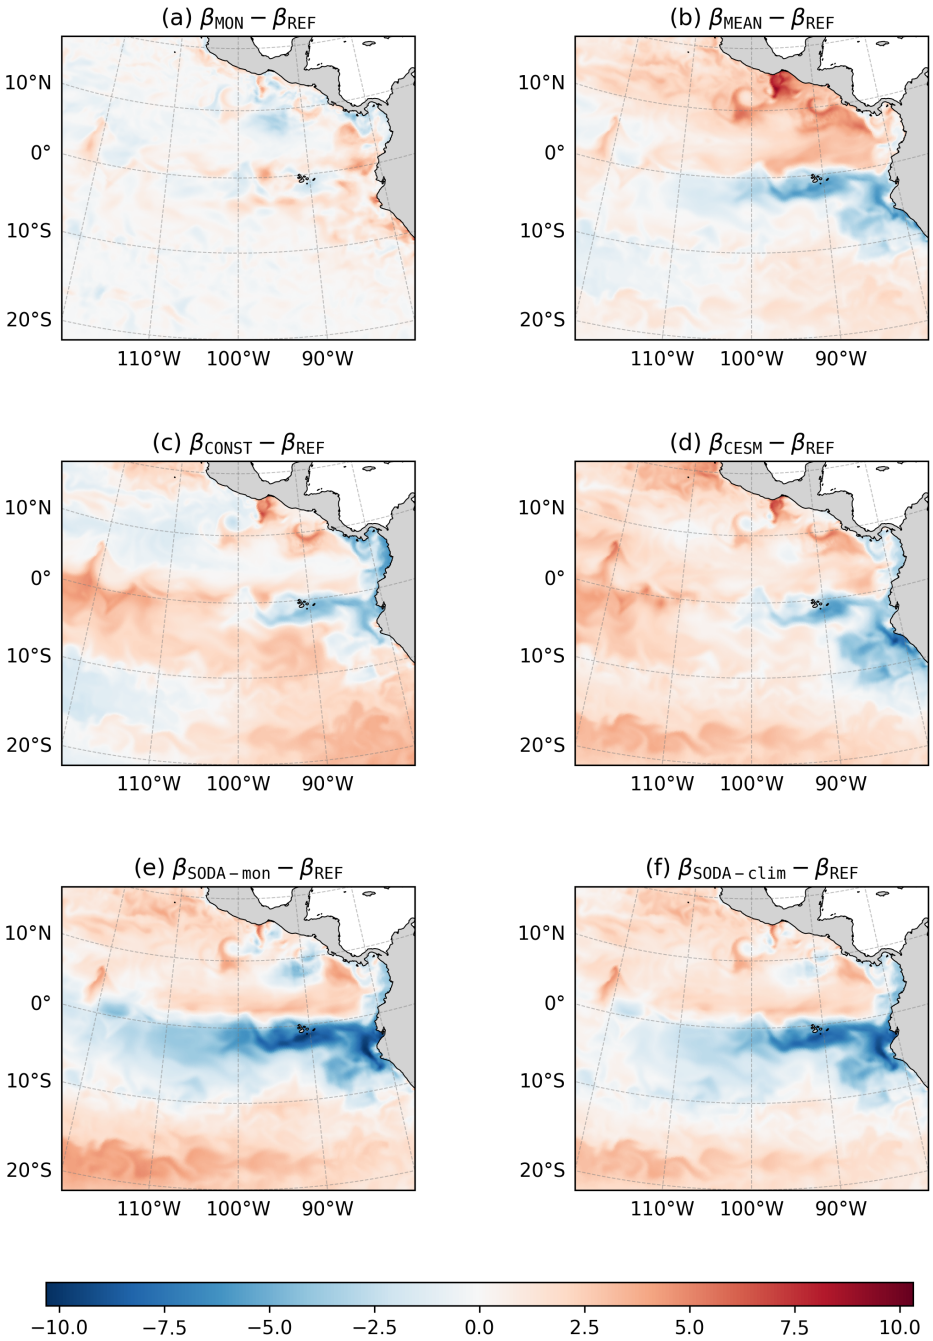


Figure S54: As Figure S45, but zoomed into the Ecuador intervention region. The color-bar range has been adjusted to the values present in this region.