Reviewer 1:

Thank you for providing additional details about the impulse response function (IRF) framework, particularly the role of η_{max}, which you note is derived based on baseline carbonate chemistry (i.e., the DIC/TA ratio), and how other parameters such as downwelling constrain the system from reaching η_{max}.

I have a few clarification questions and suggestions:

1. Could the authors clarify how η_{max} was determined? I assume it was calculated using GLODAP data. If so, does the seasonally varying IRF (and the interannual ensemble predictions) also reflect variability in η_{max} across seasons or years?

In the introduction, we have highlighted in red the description of how η_{max} was determined. Indeed, it is done using the GLODAP climatology and does not take into account seasonal variations. Please see the Figure 3 caption and the two paragraphs below equation 6. This is for the purpose of assessing linearity and time invariance.

However, the η_{max} is not computed a priori when deriving the IRFs. We added text beneath Equation 8 (the IRF model) to emphasize that this is done as a curve-fitting. η_{max} will indeed vary seasonally and will depend on where the plume goes, thus we do not impose it as a constraint but rather simply fit an analytical IRF curve to the observed curve. Please see the next comment for more explanation. How to conceptually understand " η_{max}" is indeed tricky as the reviewer points out due to the fact that the plume sees many carbonate chemistry states as it propagates out of its initial region and the effective η_{max} will be a function of time and space.

2. Some key parameters in the IRF formulation are currently under-described. For example, in Equation 8, could the authors provide the actual values of τ_1, τ_2, and τ_3? Also, do these time constants vary by subregion, or are the same values applied uniformly across the domain?

We've expanded upon this, please see the text surrounding Equation 8. There is also clarification about η_{max} there.

3. If my understand it right, one of the major advantages of the IRF approach appears to be its ability to quantify whether—and when—OAE might approach η_{max}. To enhance reader understanding, would the authors consider including η_{max} as a reference line or visual marker in one or more of the figures? This could help clarify the framework's predictive purpose and the degree to which local or seasonal factors constrain OAE efficiency relative to the theoretical maximum.

We appreciate this comment, but believe the map of η_{max} shown in Figure 3 accomplishes this goal to first order. Our figures already have a great deal of information with ensembles and seasonal curves, and we prefer to keep the figures as is. Also, it's challenging to identify a single purely chemical η_{max} for a given alkalinity plume as it propagates through various regions around the globe (experiencing differing ALK, DIC, T, and S values), so we believe Figure 3 will suffice here. More information on η_{max} and equilibration timescales are presented in Zhou et al. 2024.

Reviewer 2:

1) General comments:

I was asked to review the responses to Reviewer 1's concerns; therefore, my main suggestions are based on that request.

From a physical perspective, how can we show that the LTI (Linear Time-Invariant) assumption is a reasonable approximation? Another question I have is: how realistic is it to use alkalinity pulses that last for one month and cover several hundred square kilometers? I wonder how closely this setup reflects real-world applications. Moreover, the biological non-linearity is ignored on the basis of complexity. As a result, the IRF framework is supported only by the LTI assumption from the chemical perspective, which raises concerns about its overall suitability for OAE and mCDR approaches in general.

I believe Reviewer 1’s main — that the manuscript “is sufficiently miss-marketed and occasionally overstated that I believe it needs to be rewritten to be shorter, more focused, simpler, and more straightforward in its aims” — has not been fully addressed. While I cannot completely assess Reviewer 1’s reasoning behind the recommendation to shorten the manuscript, I agree that presenting this work as a general proof of concept for MRV applications may be too ambitious at this stage. Therefore, I believe additional effort is needed to make the manuscript more digestible and to the point, in line with Reviewer 1’s detailed comments.

That said, the authors have, in many cases, provided sound reasoning for the statements questioned and/or made appropriate revisions. I, therefore, recommend the acceptance of the manuscript after the remaining points raised by Reviewer 1 have been more fully addressed.

We appreciate the overall positive assessment of our revisions, and address the comments listed below point-by-point.

2) Specific comments:

Here, I focus only on those specific comments from Reviewer 1 that I believe were not fully addressed by the authors.

1-5: I understand why the authors chose to focus on OAE; however, I agree with Reviewer 1 that it may be confusing to introduce it so abruptly. The first part of the abstract could be revised to incorporate some of the background information currently presented in the introduction (as the authors mention in their response to this comment). Additionally, including a clarifying sentence at the end of the abstract could help bridge the abstract with the title and reinforce the points raised by the authors in their reply.

We have edited the second sentence to introduce OAE more gently, and the ending to say that the IRF approach is broadly applicable to ocean-based CDR (we discuss the similarity to direct ocean removal later in the manuscript).

10-15: I believe Reviewer 1 was referring to the statements: "We find that the IRF prediction can typically reconstruct the carbon uptake in continuous-release simulations within several percent error. Our simulations elucidate the influences of oceanic variability and deployment duration on carbon uptake efficiency." While this is addressed in the manuscript, I agree that the abstract should make it clear that these findings only refer specifically to the model used by the authors.

Agreed, we've modified the text to include "in our model": "We find that the IRF prediction can typically reconstruct the carbon uptake in continuous-release simulations in our model within several percent error."

39: I agree with Reviewer 1 that the mixing of timescales is problematic. If the main point is that OAE aims to accelerate rock weathering to ultimately enhance CO_2 uptake, then I concur with the reviewer that if the added alkalinity sinks, that objective is not achieved. I understand the authors' intention with the "thermostat" analogy, but I also agree with the reviewer that this could be omitted to maintain focus on the primary goal of OAE—enhancing CO2 uptake at the ocean surface.
We appreciate the reviewer's perspective, but really prefer to keep the thermostat analogy. The paragraph includes justification about alkalinity remaining at the surface. If the reviewers feel strongly about rewriting this paragraph we will, but we believe it adds to the narrative to present OAE as the acceleration of the natural silicate weathering cycle.

106: The authors state that they have modified the text to clarify their point, but they have only added "h" to that sentence. They could strengthen the manuscript by expanding on this point in more detail—similar to how they do in their response to Reviewer 1.
We explain this in the end of the Figure 2 caption, by saying "However, provided we have a sufficiently LTI system, we can compute the convolution of the IRF and the forcing, thus avoiding the need for an additional model integration." This is also explained further in lines 108-109 (i.e. how Figure 2 illustrates the strength of the IRF, which is what Reviewer 1 initially asked about).

Figures 3 and 4: The authors have added new text to the discussion section in response to the comment about how different Figures 3 and 4 would be if atmospheric pCO2 were not fixed. However, this addition only partially addresses Reviewer 1's question. There is no clear answer provided regarding how different the figures would actually be. Instead, the authors mention that including an interactive atmosphere and terrestrial carbon pools "may be important" future considerations, as this could reduce the sensitivity of the biological pump to changes in carbon uptake. I believe that simply acknowledging this known limitation, without offering any estimate or indication of how the final results might change, does not fully address the reviewer's concern. One possible solution would be to add a note directly to Figures 3 and 4, flagging this limitation for readers who may focus on the figures without reading the full discussion. However, I still think the authors should include some numbers to their estimation of the error introduced by not having interactive atmosphere and terrestrial carbon pools.
Figures 3 and 4 view OAE as an idealized chemical process, neglecting the nonlinearities of biosphere, atmosphere, and lithosphere interactions. This is acknowledged, and unfortunately we can only speculate what the net effect of including and fully resolving all of these constituents would be (including numbers would be speculative/ unfounded).
However, we have added text to the Figure 3 caption as suggested, saying: "**Note: these calculations assume an idealized, non-interactive atmosphere**." We have cited work that addresses the influence of an interactive atmosphere and have included caveats several times in the manuscript. For instance, please see lines 185-190: "Although OAE will decrease the

atmospheric pCO2 and thus impact CO2 uptake, for small OAE deployments we can accurately capture the first-order carbon uptake curve by assuming a non-interactive atmosphere, making modeling more affordable (Tyka, 2024)." The Tyka paper does quantify the role of an interactive atmosphere, generally found to be small but the reader may consult that paper for additional information. For the sake of keeping the paper concise (as Reviewer 1 suggested) we've simply cited this paper and briefly discussed the limitations of our non-interactive atmosphere.

Figures 6, 7, and 8 do not include subplot labels, although they are referenced in the text using subplot letters.
Thank you for pointing this out. We identified one instance where this mistake was made (referencing Fig. 7b), and this has been corrected.